arXiv:2611.16295v1 [math.CO] 12 Nov 2026

Polynomial Swap Coboundary Expansion with Applications to LTCs and PCPs

Yotam Dikstein *, Izhar Oppenheim†, Inga Valentiner-Branth ‡

November 12, 2026

Abstract

Swap coboundary expansion, introduced independently by Bafna–Minzer and Dikstein–Dinur [BM23; DD23a], is a local-to-global property that underlies agreement tests in the low-acceptance regime and recent constructions of succinct PCPs. Previous works obtained bounds of $\exp(-O(\sqrt{r}))$ for the r-swap coboundary expansion of Chapman–Lubotzky complexes [DDL24; BLM24]. We establish the first inverse-polynomial bounds for bounded-degree complexes: suitable skeleta of Kac–Moody–Steinberg (KMS) complexes have r-swap coboundary expansion at least $r^{-O(1)}$.

Our theorem supplies a key ingredient in forthcoming work of Ben Yaacov–Dikstein–Liu [BDL26], which obtains local tests for sparse algebraic codes with guarantees even at arbitrarily small acceptance probability. These tests use $\exp(\mathrm{poly}(1/\varepsilon))$ queries and guarantee that any word accepted with probability at least ε agrees with a codeword on an $\varepsilon^{O(1)}$ fraction of coordinates. Separately, substituting our bound into the existing analysis of Bafna–Minzer–Vyas–Yun [Baf+25] improves the soundness dependence of their two-query PCP construction, giving perfect completeness, soundness ε, and proof length $n(\log n)^{\exp(\mathrm{poly}(1/\varepsilon))}$.

Our proof exploits the product structure of the groups arising in KMS complexes. The main new ingredient is a decomposition lemma for Dehn functions, which separates the analysis of group relations into independent problems on much smaller factors, with only polynomial overhead.

1 Introduction

High-dimensional expanders (HDXs) are generalizations of expander graphs into hypergraphs and other high dimensional structures. These objects support global structural guarantees from local constraints, and therefore are extremely useful in constructing locally testable and decodable codes, PCPs and other types of property tests [DLZ25; Dik+26b; Baf+25].

In this work, we study swap coboundary expansion, a high dimensional analogue of expansion that captures topological properties. We prove that families of Kac–Moody–Steinberg (KMS) complexes with arbitrarily many vertices and uniformly bounded degree are swap coboundary expanders with parameter $\beta = r^{-34}$, where r is the face dimension (an r-dimensional face has $r+1$ vertices). This improves the previously known subexponential bounds for similar complexes to an inverse-polynomial bound [DDL24; BLM24; OS25].

Very informally, coboundary expansion is a statement about when a labeling of directed edges of a simplicial complex can be explained by a labeling defined on the vertices (aka coboundaries). Given a labeling of directed edges by group elements, we test whether it comes from a vertex assignment by evaluating it on triangles of the complex: on each triangle, we multiply the values on its three directed edges. If the labeling is induced by a vertex assignment, then this product is always equal to the identity. Coboundary expansion is a robust converse to this statement — if most triangle evaluations yield the identity, then the labeling must be close to a coboundary.

Swap coboundary expansion is the analogous statement in a higher-order setting, where the objects being labeled are higher dimensional sets rather than vertices and edges.

*Tel Aviv University, Tel Aviv, Israel, yotam.dikstein@gmail.com
†Ben-Gurion University of the Negev, Be'er-Sheva, Israel.
‡UC Berkeley, Berkeley, United States, ingavb@berkeley.edu

This local-to-global mechanism is the core tool underlying recent constructions of agreement tests, locally testable codes, and probabilistically checkable proofs [DLZ25; Baf+25]. In these settings, one is given a collection of local functions defined on overlapping parts of a larger domain, and the goal is to "stitch" them into a single global function that simultaneously agrees with many of the local functions. Swap coboundary expansion provides the structural guarantee that makes this stitching possible.

1.1 Definitions and our results

We now formally define coboundary expansion. Let X be a simplicial complex with vertices $X(0)$, edges $X(1)$ and triangles $X(2)$. Let $\vec{X}(1)$ be the set of directed edges (i.e. all ordered pairs uv such that $\{u,v\} \in X(1)$). Let Sym_m be the symmetric group on m elements, and $\beta \in (0,1]$. For a function $g : X(0) \to \mathrm{Sym}_m$ we denote by $\delta g : \vec{X}(1) \to \mathrm{Sym}_m$ the function $\delta g(uv) = g(v)g(u)^{-1}$. This function is called *the coboundary of* g. A function $f : \vec{X}(1) \to \mathrm{Sym}_m$ is called antisymmetric if $f(uv) = f(vu)^{-1}$ for every $uv \in \vec{X}(1)$. We say that X is a (β, m)-coboundary expander if for every antisymmetric function $f : \vec{X}(1) \to \mathrm{Sym}_m$ there exists $g : X(0) \to \mathrm{Sym}_m$ such that

$$\beta \mathop{\mathbb{P}}_{uv \in \vec{X}(1)} \left[f(uv) \neq g(v)g(u)^{-1} \right] \leqslant \mathop{\mathbb{P}}_{uvw \in X(2)} \left[f(uv)f(wu)f(vw) \neq Id \right].^{1} \tag{1.1}$$

In words, if an antisymmetric function fails the triangle identity on a δ-fraction of triangles, then it can be changed on at most a δ/β-fraction of edges to obtain a coboundary. This gives a local-to-global test: sample a random triangle uvw and check whether $f(uv)f(wu)f(vw) = Id$. For swap coboundary expansion, our theorem makes the correction loss $1/\beta$ polynomial in the dimension of the faces; this is the quantitative gain used in the applications below.

As for swap coboundary expansion, let X be a d-dimensional simplicial complex. Let $X(r)$ denote the faces of X of size $r+1$. The r-faces complex, denoted $F^r X$, is a complex whose vertices are $F^r X(0) = X(r)$. Its k-faces are all $t = \{s_0, s_1, \ldots, s_k\}$ such that every pair of $s_i, s_j \in t$ are disjoint and such that the union $\bigcup_{i=0}^{k} s_i \in X$.

Definition 1.1 (Swap coboundary expander)**.** A simplicial complex X is an (r, β, m)-swap coboundary expander if $F^r X$ is a (β, m)-coboundary expander.

Our main theorem is about the KMS complexes, a variant of the complexes appearing in [KO18; GV25; OV25].

Theorem 1.2 (Main)**.** *For every r, there exist $q_0, d_0 = \mathrm{poly}(r)$ such that the following holds. Let $q \geqslant q_0$ be a prime power and $d \geqslant d_0$. Let X be a congruence KMS complex of type A_d over $\mathbb{F}_q$. Then X is an $(r, \Omega(r^{-34}), m)$-swap coboundary expander for every integer $1 \leqslant m < \mathrm{Char}(\mathbb{F}_q)$.*

In particular, if moreover $q > d^2$, there exists a sequence of d-dimensional complexes whose number of vertices tends to infinity and whose vertex degrees are at most $\exp(O(d^2 \log q))$, all of which are simultaneously $\frac{1}{\sqrt{q}-(d-1)}$-local spectral expanders and $(r, \Omega(r^{-34}), m)$-swap coboundary expanders for every integer $1 \leqslant m < \mathrm{Char}(\mathbb{F}_q)$.

1.2 Applications of our theorem

Both our applications rely on the same mechanism: using swap coboundary expansion to convert local consistency into global structure.

Derandomized low-degree testing. The Reed-Muller error correcting code $RM(m, d, q)$ consists of all (truth tables of) degree d polynomials over $\mathbb{F}_q^m$. This code plays an important role in constructions of PCPs and of locally testable codes [KMS02]. In particular, locally testing whether a function is (close to a)

[1] The probabilities in this inequality are with respect to the natural probability distribution on X. For the purpose of this introduction we encourage the readers to think of the uniform probability over directed edges or triangles respectively. See Section 2 for details.

low-degree polynomial by only querying few values from its truth table is a problem that has been extensively studied [RS96; RS97; AS03; MR08; Har+24]. Such a test is called a *low degree test*.

One of the most common tests is the point vs. subspace test. The tester in this test receives as input a function $f : \mathbb{F}_q^m \to \mathbb{F}_q$ which we want to test and determine whether it is a low degree function. In addition, it receives an additional 'witness' F: for every k-dimensional subspace (for some constant parameter k) $W \subseteq \mathbb{F}_q^m$, F has a degree d polynomial $F_W : W \to \mathbb{F}_q$. The tester samples a random $W \subseteq \mathbb{F}_q^m$ and a point $x \in W$ and accepts if $F_W(x) = f(x)$.

Recently, [DLZ25] analyzed a derandomized version of this test, where both f and the witness F are only defined on a small subset of the points S, and a small subset of all possible subspaces. As a result they get a low degree test whose witness F has linear size in f and has the following guarantee. For every function $f : S \to \mathbb{F}_q$, if

$$\exists F \quad \mathop{\mathbb{P}}_{x,W} [(f, F) \text{ passes the test}] \geqslant 1 - \varepsilon,$$

then there exists a codeword $g : S \to \mathbb{F}_q$, whose restriction to every prescribed local domain has low degree, such that

$$\operatorname{dist}(f, g) \leqslant O(\varepsilon).$$

Here the distance is the normalized Hamming distance over S.

A very natural question is what happens in the high error regime, i.e. what happens if

$$\exists F \quad \mathop{\mathbb{P}}_{x,W} [(f, F) \text{ passes the test}] \geqslant \varepsilon$$

where $\varepsilon > 0$ is just a small constant (or even subconstant). The analysis of [DLZ25] does not generalize to the this setup. In particular, if the probability of passing the test is less than $\frac{4}{5}$, the analysis in [DLZ25] does not predict any conclusion). However, in upcoming work, [BDL26] show that a modified version of the code in [DLZ25] still has a non-trivial guarantee in this case: passing the test with probability ε guarantees that there exists a codeword $g : S \to \mathbb{F}_q$ such that

$$\operatorname{dist}(f, g) \leqslant 1 - \operatorname{poly}(\varepsilon).$$

The proof of [BDL26] uses swap coboundary expansion of the KMS complexes underlying the code to stitch local corrections into a single codeword. Our theorem establishes this property for these complexes, for which even an inverse-exponential swap coboundary expansion bound was not previously known.

The polynomial strength of our bound also improves the query complexity of the resulting local tester. This tester reads all the values of f on a sampled local domain and checks the prescribed low-degree constraints there. Our bound yields $\exp(\operatorname{poly}(1/\varepsilon))$ queries. If one instead had an expansion bound of the previously known subexponential strength, the same analysis would yield $\exp(\exp(\operatorname{poly}(1/\varepsilon)))$ queries. Thus our bound gives a testing guarantee at smaller acceptance thresholds for a given query budget.

Succinct PCPs. A probabilistically checkable proof (PCP) allows a randomized polynomial-time verifier to check a proof by reading only a few of its symbols. The number of symbols read is the query complexity, while the total number of symbols in the proof is its length. For satisfiability, perfect completeness means that every satisfiable instance has a proof that is always accepted; soundness error ε means that, for every unsatisfiable instance, every purported proof is accepted with probability at most ε.

A central goal in PCP constructions is to simultaneously achieve:

1. A constant number of queries (ideally 2 queries),
2. Arbitrarily small soundness error $\varepsilon > 0$ (even subconstant error), and
3. Near-linear proof length.

Achieving all three simultaneously was a longstanding challenge. A breakthrough of [Baf+25] achieves this combination for every fixed $\varepsilon > 0$: they construct a two-query PCP with soundness error ε and proof length

$$|w| = O(|x| \log^{C(\varepsilon)} |x|), \tag{1.2}$$

where $|x|$ is the input size and $|w|$ is the proof length.

More recently, [Dik+26a] obtained two-query PCPs with proof length $n(\log n)^{O(1)}$, where the exponent is independent of ε, for soundness down to $1/\operatorname{polylog}(n)$. Their construction uses HDXs of constant dimension, independent of the soundness parameter, and does not require bounding swap coboundary expansion.

Although this newer construction achieves better parameters, our theorem also improves the parameters of the original construction of [Baf+25]. This improvement follows simply by substituting our expansion bound into their existing analysis: it requires no new PCP analysis and does not use the new arguments of [Dik+26a].

To explain this dependence, the core technical ingredient in [Baf+25] is a family of agreement tests analyzed by [DD23a; BM23; DD23c; DDL24; BLM24]. These tests check whether a collection of local functions on overlapping domains is consistent with a single global function. Their quantitative guarantees depend on swap coboundary expansion, and this dependence propagates to the proof-length exponent $C(\varepsilon)$. The previously available expansion bounds give $C(\varepsilon) = \exp(\exp(\operatorname{poly}(1/\varepsilon)))$. Our polynomial bound reduces this to $C(\varepsilon) = \exp(\operatorname{poly}(1/\varepsilon))$, removing one exponential from the dependence of the polylogarithmic overhead on the soundness parameter.

Theorem 1.3 (Improved parameters for the PCP of [Baf+25])**.** *There exists an absolute constant $a > 0$ such that, for every fixed $\varepsilon \in (0, 1/2)$ and all sufficiently large n, the PCP construction of [Baf+25], instantiated using Theorem 1.2, gives a polynomial-time reduction from a 3-SAT instance φ of size n to a two-query PCP with perfect completeness, soundness at most ε, and proof length at most*

$$n(\log n)^{\exp(O(\varepsilon^{-a}))}.$$

The proof alphabet has size at most $\exp\left((1/\varepsilon)^{O(\log(1/\varepsilon))}\right)$.

1.3 Proof overview

To give an overview of the proof of Theorem 1.2, we present an analysis of a family of complexes that are subcomplexes of links of the KMS complex. We dub these 'matrix block complexes'. We will discuss how to bound their coboundary expansion. In the end of the overview, we will explain how a series of reductions (all of which appear in previous works) lift coboundary expansion lower bounds on these complexes to coboundary expansion lower bounds on the KMS complexes.

The key idea is that the groups defining the matrix block complexes possess a natural product structure. Our main new ingredient is a decomposition lemma that allows us to analyze the factors separately, with only a polynomial overhead in r. This reduces the problem of proving coboundary expansion for an r-faces complex to independent problems on smaller groups, whose defining subgroups have only $O(\log r)$ diagonal blocks.

What does this reduction buy us? The block boundaries are randomly positioned. For example, if all boundaries of a subgroup lie in the first fifth of the matrix, then the remaining four fifths form a single large block. The probability of this event is exponentially small in the number of boundaries. With only $O(\log r)$ boundaries, it is therefore inverse polynomial in r, so polynomially many sampled subgroups suffice to find such a subgroup. Applying the same idea to several prescribed regions gives the arrangement of large blocks needed for our direct analysis. These large blocks let us handle whole regions of the matrix together, using an argument in the spirit of [KO21] whose bounds are independent of the matrix size. Combining these bounds using the decomposition lemma yields the desired inverse-polynomial expansion bound.

We begin by introducing the matrix block complexes and recalling the connection between coboundary expansion and Dehn functions for coset complexes. We then explain how the product structure leads to the decomposition described above, and how this reduction exposes the simple geometric structure needed for the analysis.

Finally, we give a disclaimer. To explain our argument clearly, we skip the definition of the KMS complexes in this overview. Instead, we will explain how to bound the swap coboundary expansion for *a vertex link* of the KMS complex. The same decomposition and Dehn-function arguments apply to the full KMS complexes; the additional vanishing-cohomology input is discussed at the end of the overview.

1.3.1 Setup: matrix block complexes and the Dehn function

The matrix block complexes are coset complexes.

Definition 1.4 (Coset complexes)**.** Let Γ be a group and $H_0, H_1, \ldots, H_d \leqslant \Gamma$ be distinct subgroups. The coset complex $X = \mathcal{CC}(\Gamma; H_0, H_1, \ldots, H_d)$ is the complex whose vertices are cosets

$$X(0) = \{AH_i \mid \ 0 \leqslant i \leqslant d,\ A \in \Gamma\}.$$

Its k-faces are

$$X(k) = \left\{\{AH_{j_0}, AH_{j_1}, \ldots, AH_{j_k}\} \,\middle|\, \ 0 \leqslant j_0 < j_1 < \cdots < j_k \leqslant d,\ A \in \Gamma\right\}.$$

Coset complexes have been studied in the context of high dimensional expanders in many works such as [KO18; KO21; OP22; GV25; KOW25; OV25; OS25].

Fix a dimension $d > 0$ and a field $\mathbb{F}_q$. Let $\mathrm{Unip}_{d+1} \subseteq SL_{d+1}(\mathbb{F}_q)$ be the subgroup consisting of upper triangular unipotent matrices. A vertex link of the d-dimensional KMS complex is isomorphic to the coset complex $\mathcal{CC}(\mathrm{Unip}_{d+1}; H_1, H_2, \ldots, H_d)$, where, for $1 \leqslant i \leqslant d$,

$$H_i = \left\{ \begin{pmatrix} A_1 & 0 \\ 0 & A_2 \end{pmatrix} \,\middle|\, \ A_1 \in \mathrm{Unip}_i,\ A_2 \in \mathrm{Unip}_{d+1-i} \right\}.$$

The swap complex of this complex is not a coset complex, but nevertheless it has group-inherited geometry. The vertices of the r-swap complex of this complex can be identified with cosets of the subgroups $H_J = \cap_{i \in J} H_i$, for all subsets $J \in \binom{\{1,\ldots,d\}}{r+1}$. The faces in the r-swap complex are all $\{AH_{J_0}, AH_{J_1}, \ldots, AH_{J_\ell}\}$ for $A \in \mathrm{Unip}_{d+1}$ and *pairwise disjoint* index sets $J_0, J_1, \ldots, J_\ell \in \binom{\{1,\ldots,d\}}{r+1}$. The pairwise disjointness is the reason this is not a coset complex per se. Albeit, if $J_0, J_1, \ldots, J_T$ are pairwise disjoint, the complex induced by the cosets of $H_{J_0}, H_{J_1}, \ldots, H_{J_T}$ is indeed a coset complex. In a nutshell, the matrix block complexes are these subcomplexes. A reduction in [DD23c] shows that, assuming sufficiently strong local spectral expansion, if most collections of pairwise disjoint color sets $J_1, \ldots, J_T$ induce coboundary expanders, then the whole faces complex is a coboundary expander.

Therefore, the main object in this proof overview are actually these subcomplexes, which we call matrix block complexes and denote by $Y^{\mathcal{I}}$. Every subcomplex is indexed by collections $\mathcal{I} = \{J_1, J_2, \ldots, J_T\} \subset \binom{\{1,\ldots,d\}}{r+1}$. That is, $\mathcal{I}$ is a collection of $T = \mathrm{poly}(r)$ subsets of $\{1, \ldots, d\}$, each of size $|J_i| = r + 1$ which are pairwise disjoint.

Let us describe these subgroups explicitly. For $J = \{i_1 < \cdots < i_{r+1}\}$, set $i_0 = 0$ and $i_{r+2} = d + 1$. Then

$$H_J = \left\{ \begin{pmatrix} A_1 & 0 & \ldots & 0 \\ 0 & A_2 & \ldots & 0 \\ & & & \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \ldots & A_{r+2} \end{pmatrix} \,\middle|\, A_\ell \in \mathrm{Unip}_{i_\ell - i_{\ell-1}} \text{ for every } 1 \leqslant \ell \leqslant r+2 \right\}.$$

In other words, we can decompose

$$H_J = \mathrm{Unip}_{i_1} \oplus \mathrm{Unip}_{i_2 - i_1} \oplus \cdots \oplus \mathrm{Unip}_{d+1-i_{r+1}}$$

as diagonal blocks.

The matrix block complexes are $Y^{\mathcal{I}} \coloneqq \mathcal{CC}(\mathrm{Unip}_{d+1}; \{H_J\}_{J \in \mathcal{I}})$. The main technical result of the paper is that most $Y^{\mathcal{I}}$ are coboundary expanders.

Lemma 1.5 (Informal, combining Lemma 3.7 and Proposition 3.8)**.** *For suitable polynomial choices of T and d in r, and under the field-size and coefficient assumptions of Theorem 1.2, the following holds. With probability at least 90% over a uniformly random collection $\mathcal{I}$ of T mutually disjoint $(r+1)$-element subsets of $\{1, \ldots, d\}$, the complex $Y^{\mathcal{I}}$ is an $(r^{-O(1)}, m)$-coboundary expander.*

To prove Lemma 1.5 we use the group theoretic properties of this complex. More specifically, we use the connection between coboundary expansion of a coset complex and its associated Dehn function (or 'word problem'), as observed by [KO21]. Let $X = \mathcal{CC}(\Gamma; H_0, H_1, \ldots, H_d)$ be as above. Let

$$g_1 \cdot g_2 \cdot \cdots \cdot g_\ell = 1_\Gamma$$

where every $g_j \in \bigcup_{i=0}^d H_i$.[2] The area of the product is the minimal number of 'reduction steps' necessary to go from this product to the trivial product 1_Γ, where a reduction step is one of the following

1. Add or remove gg^{-1} for some $g \in \bigcup_{i=0}^d H_i$ somewhere in the sequence.
2. Replace the pair $g_j g_{j+1}$ with g' such that $g_j g_{j+1} = g'$ and all three g_j, g_{j+1}, g' are in the same subgroup H_i for some $i = 0, 1, \ldots, d$.

The Dehn function associated with the complex X,

$$\mathrm{Dehn}_X(t) = \max \left\{ area(g_1 \cdot g_2 \cdot \cdots \cdot g_\ell) \;\middle|\; \ell \leqslant t, \;\; g_1, g_2, \ldots, g_\ell \in \bigcup_{i=0}^d H_i, \;\; g_1 \cdot g_2 \cdot \cdots \cdot g_\ell = 1_\Gamma \right\}.$$

In words, this is the maximal number of reduction steps necessary in order to reduce a word of length at most t to the trivial product (and define $\mathrm{Dehn}_X(t) = \infty$ if no such reduction exists). For a formal definition see Subsection 2.6.

A theorem of [KO21] states that if X has a small diameter and if every short word can be reduced to the identity in a few steps, then the coset complex is a coboundary expander.

Theorem 1.6 ([KO21])**.** *Let X be a d-dimensional coset complex, with diameter k and $\mathrm{Dehn}_X(2k+1) < \infty$. Then X is a $(\mathrm{poly}(\frac{1}{d}, \frac{1}{k}, \frac{1}{\mathrm{Dehn}_X(1+2k)}), m)$-coboundary expander for every $m > 1$.*

Let us denote by $\Delta_X = \mathrm{Dehn}_X(1 + 2\mathrm{diam}(X))$. That is, the smallest integer such that every product of $1 + 2\mathrm{diam}(X)$ elements as above that are equal to the identity, can be reduced to the identity in Δ_X reduction steps. Next we see how the Dehn function interplays with the product structure of the links of $Y^{\mathcal{I}}$.

1.3.2 The product decomposition

Once sufficiently strong local spectral expansion and vanishing first cohomology have been established, the local-to-global machinery of [KKL16; EK16; DD23b] reduces the remaining task to proving coboundary expansion of every vertex link of $Y^{\mathcal{I}}$. We will therefore analyze the Dehn functions of these links. Specifically, we will apply Theorem 1.6 with both the diameter and $\Delta_{Y_v^{\mathcal{I}}}$ bounded by $\mathrm{poly}(r)$. In this overview we focus on the Dehn bound; the diameter bound follows from a simpler argument.

Fix $J_1 \in \mathcal{I}$. The link of a vertex AH_{J_1} is isomorphic to the coset complex with ambient group H_{J_1} and defining subgroups $H_{J_1} \cap H_{J'}$, where J' ranges over $\mathcal{I} \setminus \{J_1\}$. The key structural observation is that the ambient group and all these subgroups decompose along the same diagonal blocks. We use this common product structure to bound the Dehn function of the link.

Indeed, recall that

$$H_{J_1} = \prod_{\ell=1}^{r+2} G_\ell$$

where every $G_\ell \cong \mathrm{Unip}_{j_\ell}$ for some integer j_ℓ. These correspond to the different blocks in the matrices in H_{J_1}. The same holds for $H_{J_1} \cap H_{J'}$ which is equal to $H_{J_1 \cup J'}$. One observes that every block in $H_{J_1} \cap H_{J'}$ sits entirely inside one of the blocks in H_{J_1}. For a pictorial example, see Figure 1.

Therefore, we can decompose $H_{J_1} \cap H_{J'} = \prod_{\ell=1}^{r+2} H_\ell(J')$, where each $H_\ell(J') \leqslant G_\ell$ is itself a direct product of upper triangular unipotent groups.

Let $Y^0 = \mathcal{CC}(H_{J_1}; \{H_{J_1 \cup J'}\}_{J' \in \mathcal{I} \setminus \{J_1\}})$ be the link of any coset AH_{J_1} and let $Y_\ell = \mathcal{CC}(G_\ell; \{H_\ell(J')\}_{J' \in \mathcal{I} \setminus \{J_1\}})$ (these depend on $\mathcal{I}$ even though it doesn't appear in the notation).

[2] Formally, the object we are looking at is a finite sequence of group elements $(g_1, g_2, \ldots, g_\ell)$ whose product, taken in the order above, is equal to the identity.

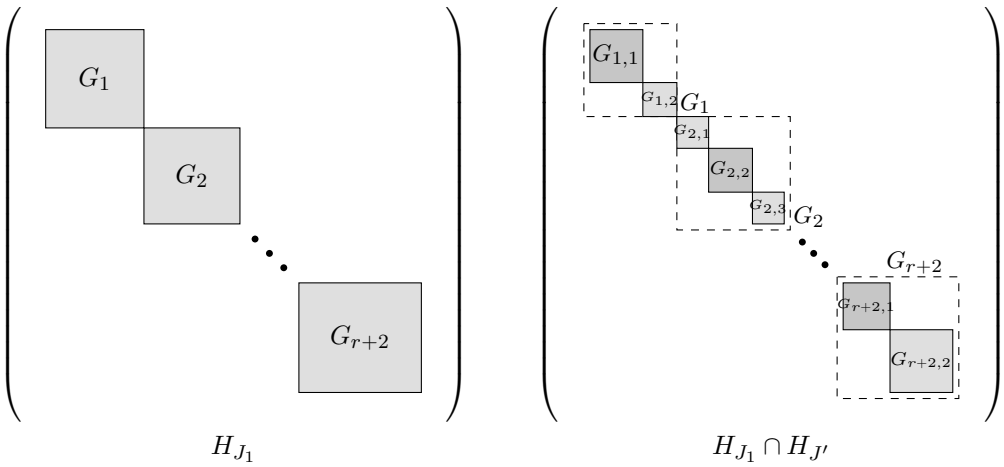


Figure 1: The blocks

The heart of the proof is the following bound on Δ_{Y^0} using the Dehn functions of the individual factors Dehn_{Y_ℓ} on words whose length is $1 + 2\mathrm{diam}(Y^0)$. More explicitly we show that with high probability over the choice of $\mathcal{I}$,

$$\Delta_{Y^0} \leqslant \sum_{\ell=1}^{r+2} \mathrm{Dehn}_{Y_\ell}(1 + 2\mathrm{diam}(Y^0)) + \mathrm{poly}(r). \tag{1.3}$$

Thus, instead of analyzing the Dehn function of the original link directly, it suffices to analyze the Dehn function for much simpler complexes Y_ℓ.

To prove (1.3), we sort a word of length at most $1 + 2\,\mathrm{diam}(Y^0)$ into separate words, one for each factor G_ℓ. Although elements from distinct factors commute, exchanging them must be justified using the allowed reduction steps, which operate within a single defining subgroup. For example, suppose $ghk = Id$ and write $g = g_1g_2g_3$, $h = h_1h_2h_3$, and $k = k_1k_2k_3$, where $g_i, h_i, k_i \in G_i$. Each letter can be split into its components using reductions within the subgroup containing that letter. We then transform $g_1g_2g_3h_1h_2h_3k_1k_2k_3$ into the 'sorted product' $(g_1h_1k_1)(g_2h_2k_2)(g_3h_3k_3)$. Since $ghk = Id$, each component word equals the identity and can be reduced separately using the Dehn bound for Y_i. In general there are $r + 2$ components, but the same argument applies.

We perform the sorting by a bubble-sort algorithm, repeatedly exchanging adjacent elements from different factors. With high probability over $\mathcal{I}$, each exchange can be implemented by splitting the two elements into a bounded number of pieces. Each pair of pieces that we need to exchange lies in a common subgroup $H_{J_1} \cap H_{J'}$, for a suitable $J' \in \mathcal{I} \setminus \{J_1\}$, and still belongs to distinct factors. We can therefore exchange the pieces using reductions within that subgroup, and then recombine them. Since the original word length and the number of components are polynomial in r, the total sorting cost is $\mathrm{poly}(r)$. Each resulting component word has length at most the original word length, so reducing these words separately gives (1.3).

1.3.3 The direct analysis

What do we gain from (1.3)? The gain is that the complexes Y_ℓ are much simpler than Y^0. They have a similar structure to Y^0, being coset complexes defined by block-unipotent matrix subgroups. However, although the number of subgroups remains $T - 1 = \mathrm{poly}(r)$, a random choice of $\mathcal{I}$ ensures with high probability that each subgroup $H_\ell(J')$ has only $O(\log r)$ blocks. More specifically, if $G_\ell \cong \mathrm{Unip}_{d'}$, then with high probability $d' = O(d \log r / r)$, and each J' places only $O(\log r)$ indexes inside this block. A probabilistic argument then shows that, for every ℓ, there are three subgroups $H_\ell(J'), H_\ell(J''), H_\ell(J''')$ whose indexes are confined to the first, last, and middle fifth of the block, respectively. Up to rounding, the first two subgroups contain a bottom block and a top block of size at least $4d'/5$, respectively, while the third contains both a top block and a bottom block of size at least $2d'/5$. This is an advantage because we can now use any relation that lives in these huge blocks to reduce our words.

Our goal is now to reduce any word of length t whose product is the identity using $O(t)$ allowed reduction steps, independently of the matrix size. The huge blocks are useful because one allowed step can perform a calculation on many matrix entries at once: whenever two neighboring letters belong to the same defining subgroup, we may replace them by their product. For example, consider two matrices that differ from the identity only within the same rectangular block above the diagonal, containing arrays A and B. Their product simply has $A + B$ in that block. If the rectangle lies inside a huge diagonal block of one of our subgroups, this multiplication is one allowed reduction step, regardless of how many entries A and B contain.

We exploit this by viewing each matrix as a 5×5 array of blocks of size approximately $d'/5$, and treating entire blocks as entries. The example above illustrates how addition of these entries can be performed in one step. More generally, we show that the arrangement and overlap of the three large-block subgroups provide the relations needed for the block-matrix calculations in our reduction, each at a bounded cost. Thus, enlarging the matrices only enlarges the arrays involved in these calculations; it does not increase the number of reduction steps needed to perform them. Using these relations, we obtain

$$\mathrm{Dehn}_{Y_\ell}(t) \leqslant Ct \tag{1.4}$$

for a universal constant $C > 0$. This argument is in the spirit of [KO21].

We stress that in the more common case where there are many small blocks, even proving a bound on the Dehn function is more challenging (partially from the fact that even the diameter in this case is super constant), but we completely sidestep this issue by using the big-block triple of subgroups.

Therefore, together with a bound on the diameter of Y^0 we get a bound of

$$\Delta_{Y^0} \overset{(1.3),(1.4)}{\leqslant} C(r+2)(1+2\mathrm{diam}(Y^0)) + \mathrm{poly}(r) = \mathrm{poly}(r).$$

1.3.4 Putting everything together

The argument above bounds the coboundary expansion of every vertex link of most subcomplexes $Y^{\mathcal{I}}$. Together with sufficiently strong local spectral expansion and vanishing first cohomology, the local-to-global machinery of [KKL16; EK16; DD23b] implies that these $Y^{\mathcal{I}}$ are themselves coboundary expanders. The reduction in [DD23c] then transfers these bounds to the faces complex of $\mathcal{CC}(\mathrm{Unip}_{d+1}; H_1, \ldots, H_d)$. In the body of the paper, we apply the same decomposition and Dehn-function arguments to the KMS complexes, together with a vanishing-cohomology argument based on [OV25]. This last input accounts for the restriction on the coefficient group in Theorem 1.2.

1.4 Related works

Coboundary expansion was defined by [LM06] and independently by [Gro10] to study topological properties of simplicial complexes. The definition we use in this paper is due to [DM22] (see also [CL23]). Swap coboundary expansion was later defined by [BM23] and [DD23a] independently for purposes of agreement testing. Both works observed that an $\exp(-O(r))$ lower bound on swap coboundary expansion is insufficient for the intended agreement-testing applications. Subsequent works obtained the first subexponential lower bounds on swap coboundary expansion for bounded-degree simplicial complexes, together with the first applications to agreement testing [DD23c; DDL24; BLM24] and PCP constructions [Baf+25]. O'Donnell and Singer [OS25] obtained a subexponential bound for a variant of the [KO18]-coset complexes constructed by [KOW25], following the proof techniques of [BLM24]. Hopkins and Ray [HR26] constructed swap coboundary expanders from projections of spherical buildings, with degree quasipolynomial in the number of vertices and expansion $\exp(-O(\sqrt{r}\log^2 r))$. Our work is the first to achieve polynomial bounds on bounded-degree complexes. As mentioned above, the work [Dik+26a] used a different route for constructing PCPs that does not use swap coboundary expansion. We also mention the work of [BDL26] which uses Theorem 1.2 to analyze locally testable codes.

Bounding coboundary expansion of various complexes has proved difficult. Gromov suggested 'the cone method' [Gro10], which is the main method we rely on in this work. This method bounds coboundary expansion for symmetric complexes by relating expansion to generalized diameter bounds. This method was later applied to various complexes, improved and generalized (see e.g. [LMM16; KO21; KM19; DD23b]). A second method is the local-to-global method which originated in [KKL16] for dimension 1 (which is what is

relevant to this paper) and [EK16] (for all dimensions). This method also underwent several quantitative improvements and was generalized to cover all groups as in our current definition of coboundary expansion [KM18; KM21; KM22; DD23b]. Our proof combines cone-type arguments with structural properties of KMS complexes.

Finally, the local-to-global method discussed above requires a separate proof for vanishing 1-cohomology. The first constructions of swap coboundary expanders appealed to the congruence subgroup property of the symplectic group [DDL24; BLM24]. More elementary vanishing-cohomology arguments were obtained for the Kaufman–Oppenheim construction by [KOW25] and for congruence KMS complexes by [OV25]; the latter applies in odd characteristic p when the coefficient group has no nontrivial elements of order p.

1.5 Open questions

It is not clear what is the optimal coboundary expansion constant β one can get on these KMS complexes. Our current cone construction gives an inverse-polynomial bound, with losses coming from the diameter and the number of word reductions used to contract loops. At present, it is conceivable that the true expansion is bounded below by a constant independent of r, or even tends to 1 as the field size and dimension grow.

Another interesting question is whether one can get a blackbox reduction from coboundary expansion of X to swap coboundary expansion. A black-box reduction achieving $\beta = \exp(-O(r))$ was established [DD23c]. However, all current applications require substantially stronger quantitative bounds.

1.6 Acknowledgements

We thank Irit Dinur for insightful discussions on this result. This work was done in part while the second and third named authors were visiting the Simons Institute for the Theory of Computing. The second named author was partially supported by ISF grant no. 242/24. The third named author was supported by the FWO and the F.R.S.–FNRS under the Excellence of Science (EOS) program (project ID 40007542) for the main part of the project.

1.7 Organization of this paper

Section 2 contains the necessary preliminaries. In Section 3, we state the main technical ingredients and use them to prove the main theorem. These ingredients are proved in Section 4, Section 5, Section 6 and Section 7. Section 8 contains the AI methodology statement. In Appendix A, we prove the PCP theorem by adapting the argument of [Baf+25] to our complex. Finally, in Appendix B, we prove that the cone method is equivalent to bounding the Dehn function in coset complexes. One direction of this equivalence already essentially appeared in [KO21]; we include both directions for completeness and because the equivalence may be of independent interest.

2 Preliminaries

Throughout we use the notation $[m] = \{0, 1, \ldots, m\}$ where $m > 0$ is an integer.

2.1 Simplicial complexes

We start by giving the definition of simplicial complexes and recall some basic properties.

Definition 2.1. A *simplicial complex* X consists of a set of vertices V and a collection of subsets of vertices $X \subseteq \mathcal{P}(V)$ such that

1. $\{v\} \in X$ for all $v \in V$,
2. for all $A \in X, B \subset A$ we have $B \in X$ (i.e. X is closed under taking subsets).

The elements of X are called simplices or faces. Given a simplex $\sigma \in X$ we define its dimension $\dim(\sigma) = |\sigma| - 1$. We will denote $X(k) = \{\sigma \in X \mid \dim(\sigma) = k\}$ for all $k \in \mathbb{N}_0$. We set $X(-1) = \{\emptyset\}$.

As is customary, we will identify the singletons $\{v\}$ with the vertices $v \in V$. In particular, we identify $X(0)$ with V and also write $X(0)$ for the set of vertices.

We call a simplicial complex finite if its vertex set is finite. We call it d-dimensional for some $d \in \mathbb{N}$ if $\max_{\sigma \in X} \dim(\sigma) = d$. Further, X is pure d-dimensional if every face is contained in a d-dimensional one.

The *degree* of a simplex σ in a pure d-dimensional complex X is the number of maximal faces containing σ, i.e. $\deg(\sigma) = |\{\tau \in X(d) \mid \sigma \subseteq \tau\}|$.

The *k-skeleton* of X, for some $1 \leqslant k \leqslant d$, is $X^{\leqslant k} = \bigcup_{i=0}^{k} X(i)$. This is a k-dimensional simplicial complex obtained by forgetting the higher-dimensional structure of X. In particular, the 1-skeleton of X is its underlying graph.

We call a simplicial complex *connected* if its 1-skeleton is connected as a graph.

Definition 2.2. Given a simplicial complex X and a simplex $\sigma \in X$, the *link* of σ in X is defined by

$$X_\sigma = \{\tau \in X \mid \tau \cap \sigma = \emptyset, \tau \cup \sigma \in X\}.$$

Definition 2.3. We call a simplicial complex X to be $d+1$-*partite* if there exists a surjective map $\mathrm{col} : X(0) \to [d] = \{0, \dots, d\}$ such that the induced map $\mathrm{col} : X \to \mathcal{P}([d]), \{v_0, \dots, v_k\} \mapsto \{\mathrm{col}(v_0), \dots, \mathrm{col}(v_k)\}$ satisfies $|\sigma| = |\mathrm{col}(\sigma)|$ for all $\sigma \in X$, i.e. each simplex consists of vertices with different colors. The map col is called the *type* or *color* function.

We use coset complexes to construct nice simplicial complexes from a group and a family of subgroups.

Definition 2.4. Let G be a group and $\mathcal{H} = (H_0, \dots, H_d)$ be a family of subgroups of G. Then the coset complex $\mathcal{CC}(G; \mathcal{H})$ is defined as the simplicial complex

- with vertex set $\bigsqcup_{i=0}^{d} G/H_i$
- such that a set of vertices $\{g_1 H_{i_1}, \dots, g_k H_{i_k}\}$ forms a $(k-1)$-simplex in $\mathcal{CC}(G; \mathcal{H})$ if and only if $\bigcap_{j=1}^{k} g_j H_{i_j} \neq \emptyset$ and all i_j's are distinct.

Remark 2.5. Given a simplex $\{g_1 H_{i_1}, \dots, g_k H_{i_k}\} \in \mathcal{CC}(G; \mathcal{H})$ we define its color to be the set of indices $\{i_1, \dots, i_k\}$. The coset complex $\mathcal{CC}(G; \mathcal{H})$ is a pure d-dimensional, $(d+1)$-partite, simplicial complex. The group G acts by $g.\{g_1 H_{i_1}, \dots, g_k H_{i_k}\} = \{g g_1 H_{i_1}, \dots, g g_k H_{i_k}\}$. This action is color preserving and transitive on the maximal faces. The vertex stabilizers of the vertices $1H_i$ are precisely the subgroups H_i.

Indeed, every d-dimensional, $d+1$-partite simplicial complex with color preserving group action which is transitive on the maximal faces can be described as a coset complex.

A key property of coset complexes is that we have a very good description of their links, as the following well-known lemma shows.

Lemma 2.6 (see e.g. [HS19, Lemma 3.4])**.** *Let G be a group with subgroups $\mathcal{H} = (H_0, \dots, H_d)$. Set $X = \mathcal{CC}(G; \mathcal{H})$ and let $\sigma \in X$ be a simplex of type $T \subset \{0, \dots, d\}$. Then the link of σ in X can again be described as a coset complex:*

$$X_\sigma \cong \mathcal{CC}(H_T; (H_{T \cup \{i\}})_{i \in \{0,\dots,d\} \setminus T})$$

where $H_T = \cap_{j \in T} H_j$.

We equip a finite, pure d-dimensional complex X with a weight function called Garland weights, which we later use to define a norm on the cochains. Let $w_d : X(d) \to (0, 1]$ be a probability distribution over d faces. We can extend w_d to a probability distribution over chains of faces $\{(\sigma_0, \sigma_1, \dots, \sigma_d) \mid \forall i\ \sigma_i \in X(i),\ \forall i < j,\ \sigma_i \subset \sigma_j\}$ by the following probabilistic experiment:

1. Sample $\sigma_d \sim w_d$.
2. Order the vertices of σ_d uniformly at random $(v_0, v_1, \dots, v_d)$.
3. Output $\sigma_i = \{v_0, v_1, \dots, v_i\}$ for all $i = 0, 1, \dots, d$.

The Garland weights $w_i : X(i) \to (0,1]$ are the marginal distributions of this joint distribution, explicitly

$$w_i(\sigma) = \frac{\sum_{\tau \in X(d), \tau \supset \sigma} w_d(\tau)}{\binom{d+1}{i+1}}.$$

When the distribution is not otherwise specified, we use $w_d(\tau) = \frac{1}{|X(d)|}$ the uniform distribution as top-level distribution. In the following two cases we will choose a different top-level distribution by convention. Given a complex X with distributions w_i as above, let $1 \leqslant n \leqslant d$. When we consider the n-skeleton of X, $X^{\leqslant n}$ we equip it with the top-level distribution $w_n^{\leqslant n} : X(n) \to (0,1]$ which is the n-th marginal of the top level distribution of X, i.e. $w_n^{\leqslant n} = w_n$. In particular, we get $w_i^{\leqslant n} = w_i$ for all $i \leqslant n$. Note that the resulting distribution is different then if we would consider $X^{\leqslant n}$ as an individual complex, ignoring its connection to X. In that case, we would choose the top level distribution $w_n^{\leqslant n}$ to be the uniform distribution. The resulting distribution is however comparable to the one we defined respecting the connection to X. A detailed comparison can be found in [Val26, Remark 2.3.15].

When X_{σ_i} is a link of $\sigma_i \in X(i)$ then we equip X_{σ_i} with the conditional distribution

$$w_{d-i-1}(\tau) \propto \begin{cases} w_d(\tau \cup \sigma_i) & \tau \cup \sigma_i \in X(d) \\ 0 & \text{otherwise} \end{cases}$$

where w_d is the top level distribution over X.

2.2 High-dimensional expansion

In this section, we recall the definition of local spectral expansion, coboundary and cosystolic expansion with (potentially) non-abelian coefficients and swap coboundary expansion.

2.2.1 Local spectral expansion

Let X be a finite, pure d-dimensional simplicial complex together with the weight function $w : X \to \mathbb{R}_{\geqslant 0}$ from above. For each $\tau \in X$ of dimension k, w induces a weight function on X_τ, which we denote by w_τ. Let $X_\tau^{\leqslant 1}$ denote the 1-skeleton of X_τ. We consider the random walk on $X_\tau^{\leqslant 1}$ with respect to w_τ and denote the second largest eigenvalue of the random walk matrix by $\lambda_2(X_\tau^{\leqslant 1})$.

Definition 2.7. Let X be a finite, pure d-dimensional simplicial complex with weight function w and let $\varepsilon \in [0,1)$. We call X an ε-(one sided) *local spectral expander* if for every $\tau \in X$ with $-1 \leqslant \dim \tau \leqslant d-2$ we have

$$\lambda_2(X_\tau^{\leqslant 1}) \leqslant \varepsilon.$$

Note that in particular the 1-skeleton of $X = X_\emptyset$ should be an ε-expander.

2.2.2 Coboundary and cosystolic expansion with non-abelian coefficients

For this section, we fix a finite, pure d-dimensional simplicial complex X and an arbitrary group Λ. We denote by $\vec{X}$ the set of ordered simplices of X. I.e. for a simplex $\{v_0, \dots, v_k\} \in X$ we have for every permutation π that $v_{\pi(0)} \dots v_{\pi(k)} \in \vec{X}$. We define the spaces of $-1, 0$ and 1-cochains:

$$\begin{aligned} C^{-1}(X,\Lambda) &= \{\phi : \{\emptyset\} \to \Lambda\}, \\ C^0(X,\Lambda) &= \{\phi : X(0) \to \Lambda\}, \\ C^1(X,\Lambda) &= \{\phi : \vec{X}(1) \to \Lambda \mid \forall uv \in \vec{X}(1), \phi(uv) = (\phi(vu))^{-1}\}. \end{aligned}$$

Furthermore, we define the coboundary maps $d_k : C^k(X,\Lambda) \to C^{k+1}(X,\Lambda)$ as follows:

– $k = -1$: For every $\phi \in C^{-1}(X,\Lambda)$ and every $v \in X(0)$, we define $d_{-1}\phi(v) = \phi(\emptyset)$.

– $k = 0$: For every $\phi \in C^0(X, \Lambda)$ and every $v_0v_1 \in \vec{X}(1)$, we define $d_0\phi(v_0v_1) = \phi(v_1)\phi(v_0)^{-1}$.

– $k = 1$: We also define the map $d_1 : C^1(X, \Lambda) \to \{\phi : \vec{X}(2) \to \Lambda\}$ as follows: for every $\phi \in C^1(X, \Lambda)$ and every $v_0v_1v_2 \in \vec{X}(2)$, we define $d_1\phi(v_0v_1v_2) = \phi(v_0v_1)\phi(v_2v_0)\phi(v_1v_2)$.

Note that $(d_0 \circ d_{-1}) = 1$ and $d_1 \circ d_0 = 1$. Thus we can define cocycles and coboundaries for $k = 0, 1$ by $Z^k(X, \Lambda) = \{\phi \in C^k(X, \Lambda) : d_k\phi \equiv 1\}$, $B^k(X, \Lambda) = \mathrm{Im}(d_{k-1})$ and get $B^k(X, \Lambda) \subseteq Z^k(X, \Lambda)$ (as sets).

We further define an action of $C^{k-1}(X, \Lambda)$ on $Z^k(X, \Lambda)$:

– An element $\psi \in C^{-1}(X, \Lambda)$ acts on $Z^0(X, \Lambda)$ by

$$\psi.\phi(v) = \psi(\emptyset)^{-1}\phi(v), \forall \phi \in Z^0(X, \Lambda), \forall v \in X(0).$$

– An element $\psi \in C^0(X, \Lambda)$ acts on $Z^1(X, \Lambda)$ by

$$\psi.\phi(v_0v_1) = \psi(v_1)\phi(v_0v_1)\psi(v_0)^{-1}, \forall \phi \in Z^1(X, \Lambda), \forall v_0v_1 \in \vec{X}(1).$$

We denote by $[\phi]$ the orbit of $\phi \in Z^k(X, \Lambda)$ under the action of $C^{k-1}(X, \Lambda)$ and define the reduced k-th cohomology to be

$$H^k(X, \Lambda) = \{[\phi] : \phi \in Z^k(X, \Lambda)\}.$$

We will say that the k-th cohomology is trivial if $H^k(X, \Lambda) = \{[\phi \equiv 1]\}$, denoted $H^k(X, \Lambda) = 1$. Note that the cohomology is trivial if and only if $B^k(X, \Lambda) = Z^k(X, \Lambda)$.

We define the support of cochains as follows. For $\phi \in C^0(X, \Lambda)$, we define

$$\mathrm{supp}(\phi) = \{v \in X(0) : \phi(v) \neq 1\}.$$

Also, for $\phi \in C^1(X, \Lambda)$, we define

$$\mathrm{supp}(\phi) = \{\{u, v\} \in X(1) : \phi(uv) \neq 1\},$$

and

$$\mathrm{supp}(d_1\phi) = \{\{v_1, v_2, v_3\} \in X(2) : (d_1\phi)(v_1v_2v_3) \neq 1\}.$$

With this notation, we define the following norms: for $\phi \in C^k(X, \Lambda)$, $k = 0, 1$, we set

$$\|\phi\| = \sum_{\tau \in \mathrm{supp}(\phi)} w(\tau).$$

Also, for $\phi \in C^1(X, \Lambda)$, we define

$$\|d_1\phi\| = \sum_{\tau \in \mathrm{supp}(d_1\phi)} w(\tau).$$

Lastly, for two maps $\phi, \psi \in C^k(X, \Lambda), k = 0, 1$, we define $\mathrm{dist}(\phi, \psi)$ as follows: For $k = 0$, we set

$$\mathrm{dist}(\phi, \psi) = w\left(\{v \in X(0) \mid \phi(v)(\psi(v))^{-1} \neq 1\}\right).$$

For $k = 1$, we define

$$\mathrm{dist}(\phi, \psi) = w\left(\{\{u, v\} \in X(1) \mid \phi(uv)(\psi(uv))^{-1} \neq 1\}\right).$$

Definition 2.8. Using the notation above, we will define for $k = 0, 1$ the *k-cocycle expansion constant* of X with respect to Λ as

$$h^k(X, \Lambda) = \min_{\phi \in C^k(X,\Lambda) \setminus Z^k(X,\Lambda)} \frac{\|d_k\phi\|}{\min_{\psi \in Z^k(X,\Lambda)} \mathrm{dist}(\phi, \psi)}.$$

Definition 2.9. For a constant $\beta > 0$, we say that

– The complex X is a (Λ, β) 1-*coboundary expander*, if $h^0(X, \Lambda) \geqslant \beta$, $h^1(X, \Lambda) \geqslant \beta$ and $H^1(X, \Lambda)$ is trivial.

– The complex X is a (Λ, β) 1-*cocycle expander*, if $h^0(X, \Lambda) \geqslant \beta$, $h^1(X, \Lambda) \geqslant \beta$. If furthermore for every $\phi \in Z^1(X, \Lambda) \setminus B^1(X, \Lambda)$ it holds that $\|\phi\| \geqslant \beta$ then we call X a (Λ, β) 1-*cosystolic expander*.

Observation 2.10. A lower bound on h^0 is usually easy to verify: If the 1-skeleton of X is a connected finite graph, then for any non-trivial group Λ, $h^0(X, \Lambda)$ is exactly the Cheeger constant of the weighted 1-skeleton (where each edge $\{u, v\}$ is weighted by $w(\{u, v\})$).

Proof. Note that the weighted Cheeger constant is

$$h(X^{\leqslant 1}) = \min_{\emptyset \neq A \subsetneq X(0)} \frac{w(E(A, A^c))}{\min(w(A), w(A^c))}$$

where $A^c = X(0) \setminus A$ denotes the complement of A and $E(A, A^c) = \{\{v, w\} \in X(1) \mid v \in A, w \notin A\}$.

To see that $h^0(X, \Lambda) \leqslant h(X^{\leqslant 1})$, fix $\emptyset \neq A \subsetneq X(0)$ minimizing for $h(X^{\leqslant 1})$ and define a cochain $f : X(0) \to \Lambda$ by fixing an arbitrary $1 \neq g \in \Lambda$ and setting $f(v) = \begin{cases} g & \text{if } v \in A \\ 1 & \text{else} \end{cases}$. Then $\operatorname{supp}(d_0 f) = E(A, A^c)$ and hence $\|d_0 f\| = w(E(A, A^c))$. Furthermore, note that since $X^{\leqslant 1}$ is connected $Z^0(X, \Lambda) = B^0(X, \Lambda)$ are all constant functions. Thus

$$\min_{b \in Z^0(X,\Lambda)} \operatorname{dist}(f, b) = \min(\operatorname{dist}(f, [v \mapsto 1]), \operatorname{dist}(f, [v \mapsto g]), \operatorname{dist}(f, [v \mapsto h] : h \neq g, 1) = \min(w(A), w(A^c)).$$

Together we get

$$h(X^{\leqslant 1}) = \frac{w(E(A, A^c))}{\min(w(A), w(A^c))} = \frac{\|d_0 f\|}{\min_{b \in Z^0(X,\Lambda)} \operatorname{dist}(f, b)} \geqslant h^0(X, \Lambda).$$

To see the other inequality, $h^0(X, \Lambda) \leqslant h(X^{\leqslant 1})$, fix an arbitrary cochain $f \in C^0(X, \Lambda) \setminus Z^0(X, \Lambda)$. Then f partitions the set of vertices

$$X(0) = \sqcup_{g \in \Lambda} V_g \text{ with } V_g = f^{-1}(\{g\}).$$

Firstly, observe that for a constant map $b = [v \mapsto g]$ we have $\operatorname{dist}(f, b) = w(\{v \in X(0) \mid f(v) \neq g\}) = w(V_g^c)$. Next, we use the definition of $d_0 f$ and that $h(X^{\leqslant 1}) \min(w(A), w(A^c)) \leqslant w(E(A, A^c))$ for all $\emptyset \neq A \subset X(0)$ to obtain

$$\|d_0 f\| = \frac{1}{2} \sum_{g \in \Lambda} w(E(V_g, V_g^c)) \geqslant \frac{1}{2} \sum_g h(X^{\leqslant 1}) \min(w(V_g), w(V_g^c)).$$

To compare this to the distance of f to the cocycles/coboundaries, we consider two cases. First, assume that $\min(w(V_g), w(V_g^c)) = w(V_g)$ for all g. Thus, in particular $w(V_g) \leqslant \frac{1}{2}$ and hence $\min_{b \in Z^0(X,\Lambda)} \operatorname{dist}(f, b) = \min_g w(V_g^c) \geqslant \frac{1}{2}$. And we have $\sum_{h \in \Lambda} \min(w(V_h), w(V_h^c)) = \sum_{h \in \Lambda} w(V_h) = w(X(0)) = 1$.

On the other hand, assume that there exists a $g \in \Lambda$ such that $w(V_g) > \frac{1}{2}$ (note the there can exist at most one such g). Then

$$\sum_{h \in \Lambda} \min(w(V_h), w(V_h^c)) = w(V_g^c) + \sum_{h \neq g} w(V_h = 2(w(V_g^c)) \geqslant 2 \min_b \operatorname{dist}(f, b).$$

Together, we get that

$$\|d_0 f\| \geqslant \frac{1}{2} \sum_g h(X^{\leqslant 1}) \min(w(V_g), w(V_g^c)) \geqslant \frac{1}{2} h(X^{\leqslant 1}) 2 \min_b \operatorname{dist}(f, b)$$

and hence $h^0(X, \Lambda) \geqslant h(X^{\leqslant 1})$.

□

2.2.3 Swap coboundary expansion

Swap coboundary expansion of a complex X in essence means that the r-faces complex of X is a coboundary expander. This notion was defined independently by [DD23a] and [BM23] in order to analyze property tests used in PCPs.

We start by defining the faces complex and related notions to then state the precise definition of swap coboundary expansion below in Definition 2.12.

Definition 2.11 ([DD23c, Definition 2.11])**.** Let X be a d-dimensional simplicial complex and $1 \leqslant r \leqslant d$. We denote by $F^r X$ the simplicial complex with vertex set $X(r)$ and where a set of vertices $\{\sigma_0, \dots, \sigma_k\}$ forms a simplex if and only if the σ_i are pairwise disjoint and $\bigcup_{i=0}^{k} \sigma_i \in X$. We call $F^r X$ the *r-faces complex* of X.

The dimension of $F^r X$ is $t = \lfloor \frac{d+1}{r+1} \rfloor - 1$. Let $t' = (t+1)(r+1) - 1$. We equip $F^r X$ with the following measure. On top level faces of $F^r X$, the measure is given by sampling a t'-dimensional simplex in X (according to the measure of X and then partitioning it into $t+1$ faces of dimension r uniformly at random. For lower dimensional faces, we consider the induced measure as usual. Explicitly, if w_i denote the weight functions on X we get

$$w_{F^r X, \ell}(\{\sigma_0, \dots, \sigma_\ell\}) = \frac{(\ell+1)!((r+1)!)^{\ell+1}}{((l+1)(r+1))!} w_{(r+1)(\ell+1)-1}(\cup_i \sigma_i).$$

Definition 2.12. Let X be finite, pure d-dimensional simplicial complex, Λ a group and $\beta > 0, 1 \leqslant r \leqslant d$ be constants. We call X a (β, r, Λ)-*swap coboundary expander* if $F^r X$ is a (β, Λ) 1-coboundary expander.

Observation 2.13. For a d-dimensional $(d+1)$-partite complex Y, mapping each face to its set of colors induces a surjective graph homomorphism from the one-skeleton of $F^r Y$ to the Kneser graph $K(d+1, r+1)$, and similarly for links with the corresponding parameters. Here $K(d+1, r+1)$ is the graph whose vertices are the $(r+1)$-element subsets of $\{0, 1, \dots, d\}$, with two vertices adjacent if their sets are disjoint.

Thus, whenever the corresponding Kneser graph is disconnected, so is the link. In particular, if $r+1$ divides $d+1$, the codimension-2 links map surjectively onto $K(2(r+1), r+1)$, which is a disjoint union of edges, and hence these links are disconnected. Even when $d+1$ is not a multiple of $r+1$, a similar spectral argument shows that the codimension-2 links cannot be λ-one-sided spectral expanders for any $\lambda < \frac{1}{4}$.

The faces complex of a coset complex is not again a coset complex. Thus, it is much easier to analyze the colored faces complex of a coset complex, which we will define now, and which will turn out to be again a coset complex.

Definition 2.14. Let X be a finite, pure d-dimensional, $d+1$-partite simplicial complex with color function $\mathrm{col} : X \to \mathcal{P}(\{0, \dots, d\})$. Let $1 \leqslant r, T \leqslant d$. Fix a set of colors $\mathcal{I} = \{I_1, \dots, I_T\}$ such that $I_j \subseteq \{0, \dots, d\}$, $|I_j| = r+1$, and such that the I_j are pairwise disjoint. We denote by $X[I_j] = \{\sigma \in X \mid \mathrm{col}(\sigma) = I_j\}$ the faces with color I_j.

Then we define the *colored faces complex* $F^{\mathcal{I}} X$ of X with respect to $\mathcal{I}$ as simplicial complex

- with vertex set $\bigcup_{j=0}^{T} X[I_j]$,
- such that a set of vertices $\{\sigma_0, \dots, \sigma_\ell\}$ forms a simplex if there exist distinct $j_0, \dots, j_\ell$ such that $\sigma_i \in X[I_{j_i}]$ and $\bigcup_i \sigma_i \in X$.

Note that $F^{\mathcal{I}} X$ is $|T|$-partite and $(T-1)$-dimensional. We equip $F^{\mathcal{I}} X$ with the measure obtained by the measure on $F^r X$ conditioned on sampling a face whose vertices have colors $\mathcal{I}$. Denote the weight on X by w_i, we get for $\sigma = \{\sigma_0, \dots, \sigma_\ell\} \in F^{\mathcal{I}} X$ the weight

$$w_{F^{\mathcal{I}} X, \ell}(\sigma) = \frac{w_{(\ell+1)(r+1)-1}(\cup_i \sigma_i)}{\sum_{\{\tau_0, \dots, \tau_\ell\} \in F^{\mathcal{I}} X(\ell)} w_{(\ell+1)(r+1)-1}(\cup_i \tau_i)}.$$

Fact 2.15. *Let $X = \mathcal{CC}(G; (H_i)_{i \in [d]})$ be a coset complex and $\mathcal{I} = \{I_1, \dots, I_T\}$ a set of disjoint colors as above. Set $H_I = \cap_{i \in I} H_i$ for each $I \in \mathcal{I}$. Then*

$$F^{\mathcal{I}} X \cong \mathcal{CC}(G; (H_I)_{I \in \mathcal{I}}).$$

Definition 2.16. Let $1 \leqslant r \leqslant d \in \mathbb{N}$ and $T \in \mathbb{N}$. We denote by $\binom{[d]}{r+1}$ the set of subsets of size $r+1$ of $[d] = \{0, \dots, d\}$. We also introduce the following notation for the set of all sets of T disjoint colors of size $r+1$ of $[d]$:

$$\binom{[d]}{r+1}^{[T]} = \left\{ \mathcal{I} \subseteq \binom{[d]}{r+1} \mid |\mathcal{I}| = T \text{ and } \forall I, J \in \mathcal{I}, I \neq J : I \cap J = \emptyset \right\}.$$

2.3 KMS complexes

We now describe the main objects of our study, the congruence KMS complexes. We restrict to the congruence KMS complexes of type A_d, hence we work with the Chevalley group of type A_d which is SL_{d+1}. We fix a prime power q and a size $d \geqslant 2$. For a polynomial $f \in \mathbb{F}_q[t]$ and $1 \leqslant i, j \leqslant d+1, i \neq j$ we denote by $e_{i,j}(f)$ the matrix in $\mathrm{SL}_{d+1}(\mathbb{F}_q[t])$ with ones on the diagonal and f at position (i, j). We then define for $i = 1, \dots, d$

$$U_i = \langle e_{i,i+1}(\lambda) \mid \lambda \in \mathbb{F}_q \rangle \leqslant \mathrm{SL}_{d+1}(\mathbb{F}_q[t])$$

and

$$U_0 = \langle e_{d+1,1}(\lambda t) \mid \lambda \in \mathbb{F}_q \rangle \leqslant \mathrm{SL}_{d+1}(\mathbb{F}_q[t]).$$

For each subset $J \subseteq [d], i \in [d]$ we define

$$U_J = \langle U_i \mid i \in J \rangle \leqslant \mathrm{SL}_{d+1}(\mathbb{F}_q[t]), \qquad H_i = U_{[d]\setminus\{i\}}$$

and we call $\mathrm{KMS}(d, q) = U_{[d]}$ the *concrete KMS group in* $\mathrm{SL}_{d+1}(\mathbb{F}_q[t])$.

Remark 2.17. For $d \geqslant 3$ the concrete KMS group $\mathrm{KMS}(d, q)$ is isomorphic to the KMS group of type $\widetilde{A}_d$ over $\mathbb{F}_q$ as defined in [GV25; OV26; OV25].

Example 2.18. For the case $d = 2$ we have

$$U_{\{1,2\}} = \begin{pmatrix} 1 & \mathbb{F}_q & \mathbb{F}_q \\ 0 & 1 & \mathbb{F}_q \\ 0 & 0 & 1 \end{pmatrix}, \qquad U_{\{0,1\}} = \begin{pmatrix} 1 & \mathbb{F}_q & 0 \\ 0 & 1 & 0 \\ \mathbb{F}_q t & \mathbb{F}_q t & 1 \end{pmatrix}, \qquad U_{\{0,2\}} = \begin{pmatrix} 1 & 0 & 0 \\ \mathbb{F}_q t & 1 & \mathbb{F}_q \\ \mathbb{F}_q t & 0 & 1 \end{pmatrix}.$$

Here we use a shortened notation, by which we for example mean

$$\begin{pmatrix} 1 & 0 & 0 \\ \mathbb{F}_q t & 1 & \mathbb{F}_q \\ \mathbb{F}_q t & 0 & 1 \end{pmatrix} = \left\{ \begin{pmatrix} 1 & 0 & 0 \\ \lambda_1 t & 1 & \lambda_2 \\ \lambda_3 t & 0 & 1 \end{pmatrix} \mid \lambda_i \in \mathbb{F}_q \right\}.$$

Then $\mathrm{KMS}(2, q)$ is the group generated by $U_{\{1,2\}}, U_{\{0,1\}}, U_{\{0,2\}}$ in $\mathrm{SL}_3(\mathbb{F}_q[t])$. It consists of all matrices that are unipotent upper triangular when considered modulo t.

Fix an irreducible polynomial $f \in \mathbb{F}_q[t]$ of degree at least two. The quotient map $\mathbb{F}_q[t] \to \mathbb{F}_q[t]/(f)$ induces a map

$$\pi_f : \mathrm{SL}_{d+1}(\mathbb{F}_q[t]) \to \mathrm{SL}_{d+1}(\mathbb{F}_q[t]/(f)).$$

We set $H_i^f = \pi(H_i) = \pi_f(U_{[d]\setminus\{i\}})$ for all $i \in [d]$. The *congruence KMS complex of type* A_d *over* $\mathbb{F}_q[t]/(f)$ is

$$X(A_d, \mathbb{F}_q, f) = \mathcal{CC}(\mathrm{SL}_{d+1}(\mathbb{F}_q[t]/(f)); (H_i^f)_{i\in[d]}).$$

For $J \subset [d]$ we set $H_J = \cap_{i\in J} H_i$. Note that $H_J = U_{[d]\setminus J}$ and $\pi_f(H_J) = \cap_{i\in J} H_i^f$.

Let $v \in X(A_d, \mathbb{F}_q, f)(0)$ be a vertex. Then the link of v in $X = X(A_d, \mathbb{F}_q, f)$ can be described as follows. Let U^+ be the subgroups of unipotent upper triangular matrices in $\mathrm{SL}_{d+1}(\mathbb{F}_q)$ and let $K_i = \langle e_{j,j+1}(\lambda) \mid \lambda \in \mathbb{F}_q, j \in \{1, \dots, d\} \setminus \{i\} \rangle$. Then

$$X_v \cong \mathcal{CC}(U^+; (K_i)_{i\in\{1,\dots,d\}}).$$

This complex is isomorphic to a subcomplex of the spherical building of type A_d called the opposite complex. However, we will not make use of this fact in this work.

2.4 Previous results on KMS complexes

In [GV25] the authors show that KMS complexes are local spectral expanders. More precisely and adapted to our case they show the following.

Theorem 2.19. *Let $q \geqslant 4, d \geqslant 2$ and $f \in \mathbb{F}_q[t]$ be an irreducible polynomial of degree at least 2. Further, assume that $\sqrt{q} \geqslant d$. Then $X(A_d, \mathbb{F}_q, f)$ is a $\frac{1}{\sqrt{q}-(d-1)}$-local spectral expander.*

We use a slightly better local-to-global theorem by [LO25] as a black box to obtain local spectral expansion that is independent of dimension.

Theorem 2.20. *Let $d \geqslant 2$, let $q > 16$ be a prime power, and let $f \in \mathbb{F}_q[t]$ be irreducible with $\deg f \geqslant 2$. Set $X = X(A_d, \mathbb{F}_q, f)$. Then for every face $\tau \in X$, including the empty face, with $|\tau| \leqslant d-1$, we have*

$$\lambda_2(X_\tau^{\leqslant 1}) \leqslant \frac{2}{(d-|\tau|)(\sqrt{q}-2)}.$$

In particular X is a $\frac{2}{\sqrt{q}-2}$-one-sided local spectral expander. Moreover, for any $2 \leqslant n \leqslant d$, then $X^{\leqslant n}$, with the weights inherited from X, is a $\frac{1}{d-n+1}$-two-sided local spectral expander.

Proof. The links of X are connected. A codimension-two link whose two remaining colors are adjacent in the cyclic diagram of type $\widetilde{A}_d$ has second eigenvalue at most $q^{-1/2}$; for nonadjacent colors, the link is complete bipartite and its second eigenvalue is zero [GV25]. Thus the pairwise spectral influence matrix is bounded entrywise by $q^{-1/2}$ times the adjacency matrix of a cycle, and its largest eigenvalue is at most $2/\sqrt{q}$. Applying [LO25, Theorem 1.3] gives the stated bound.

For the two-sided conclusion, recall that the smallest eigenvalue of the weighted graph of a pure D-dimensional complex is at least $-1/D$, since its edge distribution is a mixture of the edge distributions of D-simplices. Here $\dim X_\tau = d - |\tau|$. When $q \geqslant 16$, we have $\frac{2}{\sqrt{q}-1} \leqslant 1$, so every nontrivial eigenvalue of $X_\tau^{\leqslant 1}$ has absolute value at most $1/(d-|\tau|)$. Every face whose link must be checked in $X^{\leqslant n}$ has size at most $n-1$, and its weighted link graph is unchanged by taking the skeleton. This proves the claim. □

[OV25] also proved these complexes are 1-coboundary expanders. We will only use part of the proof method for the vanishing of cohomology later, but state the result for completenes.

Theorem 2.21. *Let $d \geqslant 3$. There exist $q_0, \varepsilon > 0$ such that*

- *for every prime power $q = p^m \geqslant q_0$,*
- *for every group Λ without non-trivial elements of order p (e.g. if Λ is finite and $|\Lambda| < p$),*
- *for every irreducible polynomial $f \in \mathbb{F}_q[t]$ of degree at least 2*

we have

$$h^1(X(A_d, \mathbb{F}_q, f), \Lambda) \geqslant \varepsilon \text{ and } H^1(X(A_d, \mathbb{F}_q, f), \Lambda) = 1.$$

In particular, $X(A_d, \mathbb{F}_q, f)$ is a (Λ, ε) 1-coboundary expander.

2.5 Dehn functions

The notion of Dehn functions will be useful to use below as a tool to bound coboundary expansion in coset complexes. In the section, we cover the basic definitions and some results regarding Dehn function that we will need in the sequel.

Let $G = \langle S | R \rangle$ be a finitely presented group, where S is a symmetric generating set of G and R is a set of relations. Denote $\mathrm{Free}(S)$ to be the free group with a generating set S.

Definition 2.22 (Word length)**.** Any element in $\mathrm{Free}(S)$ can be written as a word in the alphabet S and for a word w in the alphabet S, we denote $|w|_S$ to be the length of the word (i.e., the number of letters).

Without loss of generality, we can assume that $R \subseteq \mathrm{Free}(S)$ are cyclically reduced words. A *relation* in G is a word $w \in \mathrm{Free}(S)$ such that $w = 1$ in G. We note that every relation is in the normal closure of R in $\mathrm{Free}(S)$. In the following, we will deal with different generating sets for the same group G, in that situation (when G is obvious from the context), we will say that w is a *relation in the alphabet* S to specify to what generating set we are referring to.

Given a relation $w \in \mathrm{Free}(S)$, the *area with respect to* R of w, denoted $\mathrm{Area}_R(w)$, is the minimal number $k \in \mathbb{N} \cup \{0\}$ such there are $\rho_1, ..., \rho_k \in R \cup R^{-1}$ and $u_1, ..., u_k \in \mathrm{Free}(S)$ such that

$$w = u_1^{-1}\rho_1 u_1 ... u_k^{-1}\rho_k u_k \text{ in } \mathrm{Free}(S).$$

Definition 2.23 (The Dehn function)**.** The *Dehn function of* G *with respect to* S, R is the function $\mathrm{Dehn}_{S,R} : \mathbb{N} \to \mathbb{N}$ defined as

$$\mathrm{Dehn}_{S,R}(m) = \max\{\mathrm{Area}_R(w) : w = 1 \text{ in } G, |w|_S \leqslant m\}.$$

The idea behind the Dehn function is that it counts how many reduction moves are needed to reduce w to the trivial word. More precisely, following [CM17, Chapter 8], we define the following moves on a word w:

1. Free reduction: remove a sub-word of the form ss^{-1} within the word w, where $s \in S$.
2. Free expansion: Replace $w = uv$ (where u, v are words) with $w' = uss^{-1}v$ or with $w' = uev$.
3. Applying a relation: replace a subword w' in w with a new subword w'' where $w'(w'')^{-1}$ or $w''(w')^{-1}$ is a cyclic permutation of a word in $R \cup R^{-1}$.

Every $w = 1$ in G can be reduced to 1 by applying the above moves, and a different way to define $\mathrm{Dehn}_{S,R}(m)$ is as follows: $\mathrm{Dehn}_{S,R}(m)$ is the minimal number N such that for every w with $|w| \leqslant m$, and $w = 1$, the word w can be reduced to 1 by applying most N relations. It does not take into account the number of free reductions or free expansions that are applied, and we will refer to these relations as *trivial relations*.

Given a presentation as above and two words w, w', we will say that *the word* w' *can be deduced from* w, if one can apply the above moves on w and get w'.

In the sequel, it will be useful to compare Dehn function of the same group with two different presentations. Let G be a group with two different presentations in terms of generators and relations: $G = \langle S|R\rangle$ and $G = \langle S'|R'\rangle$. It is well-known that $\mathrm{Dehn}_{S,R}$ and $\mathrm{Dehn}_{S',R'}$ are the same up to a linear distortion, but we will need to keep track of the constants of this distortion. We define homomorphism $\phi : \mathrm{Free}(S) \to \mathrm{Free}(S')$ and $\phi' : \mathrm{Free}(S') \to \mathrm{Free}(S)$ as follows: Every element $s \in S$ is in particular an element in G and as such can be written as a word in S'. For every element $s \in S$ we choose $w_s \in \mathrm{Free}(S')$ such that $s = w_s$ in G and $|w_s|_{S'}$ is minimal. We define $\phi(s) = w_s$ and extend ϕ to a homomorphism $\phi : \mathrm{Free}(S) \to \mathrm{Free}(S')$. Similarly, for every $s' \in S'$, we choose $w_{s'} \in \mathrm{Free}(S)$ such that $s' = w_{s'}$ in G and $|w_{s'}|_S$ is minimal and define a homomorphism $\phi' : \mathrm{Free}(S') \to \mathrm{Free}(S)$ as above.

Proposition 2.24. *Let* G *be a group with two different presentations in terms of generators and relations:* $G = \langle S|R\rangle$ *and* $G = \langle S'|R'\rangle$*. For choices of* ϕ, ϕ' *as above, we define the following constants:*

- $C_1 = \max_{s\in S} |\phi(s)|_{S'}$.
- *We note that for every* $\rho' \in R'$, $\phi'(\rho')$ *is a relation in the alphabet* S *and define* $C_2 = \max_{\rho'\in R'} \mathrm{Area}_R(\phi'(\rho'))$.
- *We note that for every* $s \in S$, $s(\phi'(\phi(s)))^{-1}$ *is a relation in the alphabet* S *and define* $C_3 = \max_{s\in S} \mathrm{Area}_R(s(\phi'(\phi(s)))^{-1})$.

Then it holds for every $m \in \mathbb{N}$ *that*

$$\mathrm{Dehn}_{S,R}(m) \leqslant C_2\,\mathrm{Dehn}_{S',R'}(C_1 m) + C_3 m.$$

Proof. Fix $m \in \mathbb{N}$ and let w be a relation in the alphabet S with $|w|_S \leqslant m$. We denote $w = s_1...s_j$ where $s_1, ..., s_j \in S$ (not necessarily distinct) where $j \leqslant m$. We note that by applying a free expansion $w = s_1(\phi'(\phi(s_1)))^{-1}(\phi'(\phi(s_1)))s_2...s_j$. Thus, by the definition of C_3, we can deduce $(\phi'(\phi(s_1)))s_2...s_j$ from w by applying at most C_3 relations in R. Continuing in the same manner, we can deduce $(\phi'(\phi(s_1)))...(\phi'(\phi(s_j))) = \phi'(\phi(s_1...s_j))$ by applying at most $C_3 j \leqslant C_3 m$ relations in R (and m free expansions that we do not count). Thus,

$$\mathrm{Area}_R(s_1...s_j) \leqslant C_3 m + \mathrm{Area}_R(\phi'(\phi(s_1...s_j))).$$

We will finish the proof by showing that $\mathrm{Area}_R(\phi'(\phi(s_1...s_j))) \leqslant C_2 \,\mathrm{Dehn}_{S',R'}(C_1 m)$. We note that $w = s_1...s_j$ is a relation in the alphabet S and thus $\phi(s_1...s_j)$ is a relation in the alphabet S' and

$$|\phi(s_1...s_j)|_{S'} \leqslant \sum_{i=1}^{j} |\phi(s_i)|_{S'} \leqslant C_1 j \leqslant C_1 m.$$

By the definition of $\mathrm{Dehn}_{S',R'}$ there is $k \leqslant \mathrm{Dehn}(C_1 m)$, words $u_1, ..., u_k \in \mathrm{Free}(S')$ and relations $\rho'_1, ..., \rho'_k \in R'$ such that

$$\phi(s_1...s_j) = u_1^{-1}\rho'_1 u_1 u_2^{-1}\rho'_2 u_2 ... u_k^{-1}\rho'_k u_k.$$

Therefore

$$\begin{aligned}
&\mathrm{Area}_R(\phi'(\phi(s_1...s_j))) \leqslant \\
&\mathrm{Area}_R(\phi'(u_1^{-1})\phi'(\rho'_1)\phi'(u_1)...\phi'(u_k^{-1})\phi'(\rho'_k)\phi'(u_k) \leqslant \\
&\sum_{i=1^k} \mathrm{Area}_R(\phi'(\rho'_i)) \leqslant C_2 k \leqslant C_2 \,\mathrm{Dehn}(C_1 m),
\end{aligned}$$

as needed. □

2.6 Dehn functions, cone functions and coboundary expansion

Non-abelian cone functions were introduced in [DD23c], our definition is very similar. We follow the exposition in [OV25].

Throughout, X is an n-dimensional simplicial complex. We again denote by $\vec{X}(k)$ the set of ordered k-dimensional simplices of X, and for a simplex $\{v_0, \dots, v_k\} \in X$ we denote by $v_0 \dots v_k \in \vec{X}$ the ordered simplex.

Given two vertices u, v of X, a *path from u to v* is a $(k+1)$-tuple of vertices $(u_0; ...; u_k)$ where $u = u_0$ and $v = u_k$ such that for every $0 \leqslant i \leqslant k-1$, $u_i u_{i+1} \in \vec{X}(1)$. The constant k is called the *length of the path*. Given two paths $P_0 = (u_0; ...; u_k), P_1 = (u_k; ...; u_m)$, we define $P_0 \circ P_1 = (u_0; ...; u_m)$. Also, given a path $P = (u_0; ...; u_k)$, we denote $P^{-1} = (u_k; ...; u_0)$.

A *loop around v* is a path from v to itself. We note that (v) is also considered as a loop of length 0 and we will call it the *trivial loop around v*.

Definition 2.25. Fix v_0 to be a vertex in X. We define the following symmetric relations between two loops P, P' around v_0, see also their visualizations in Figure 2.

- We denote $P \overset{(BT)}{\sim} P'$ if there exist vertices u, v and paths Q_0, Q_1 from v_0 to u such that $P = Q_0 \circ (u; v; u) \circ Q_1^{-1}$ and $P' = Q_0 \circ (u) \circ Q_1^{-1}$ ((BT) stands for "backtracking").

- We denote $P \overset{(TR)}{\sim} P'$ if there is a triangle uvw in X and paths Q_0, Q_1 from v_0 to u and w correspondingly such that $P = Q_0 \circ (u; v; w) \circ Q_1^{-1}$ and $P' = Q_0 \circ (u; w) \circ Q_1^{-1}$ ((TR) stands for "triangle"). We denote by $t_{P,P'} = uvw$ the triangle used for the relation.

Remark 2.26. Note that the relations above are not defined to be transitive. Indeed, in the sequel, it will be important to keep track of how many times the relation $\overset{(TR)}{\sim}$ is used to pass from one loop to another.

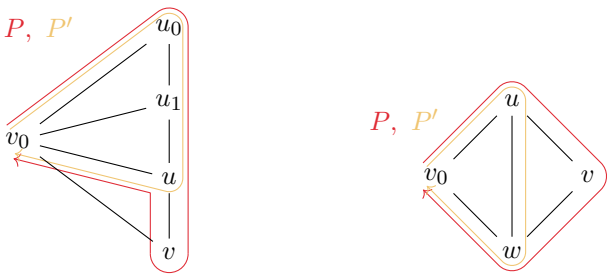


Figure 2: Visualization of $\overset{(BT)}{\sim}$ on the left and $\overset{(TR)}{\sim}$ on the right.

Definition 2.27. A 0-*cone* on X is a pair $\mathrm{Cone}_X = \left(v_0, (P_u)_{u\in X(0)}\right)$ such that

1. $v_0 \in X(0)$.
2. For every $u \in X(0), u \neq v_0$, P_u is a path from v_0 to u. For v_0, P_{v_0} is the trivial loop around v_0.

A 1-*cone* on X is a triple $\mathrm{Cone}_X = \left(v_0, (P_u)_{u\in X(0)}, (T_{uv})_{uv\in \vec{X}(1)}\right)$ such that

1. $\left(v_0, (P_u)_{u\in X(0)}\right)$ is a 0-cone on X.
2. For every $uv \in \vec{X}(1)$, T_{uv} is a sequence of loops around v_0, $T_{uv} = (P_0, P_1, \ldots, P_m)$ such that the following holds:
   (a) $P_0 = P_u \circ (u; v) \circ P_v^{-1}$.
   (b) For every $0 \leqslant i \leqslant m-1$, either $P_i \overset{(TR)}{\sim} P_{i+1}$ or $P_i \overset{(BT)}{\sim} P_{i+1}$.
   (c) P_m is the trivial loop around v_0.

We refer to the fixed vertex v_0 as the apex of the cone. Furthermore, we call $T_{uv} = (P_0, P_1, ..., P_m)$ as above a *contraction* and denote the number of triangle relations in the contraction by

$$|T_{uv}| = |\{0 \leqslant i \leqslant m-1 : P_i \overset{(TR)}{\sim} P_{i+1}\}|.$$

We further define the *radii of the cone* as follows:

1. $\mathrm{Rad}_0(\mathrm{Cone}_X) = \sup_{u\in X(0)} \mathrm{length}(P_u)$.
2. $\mathrm{Rad}_1(\mathrm{Cone}_X) = \sup_{uv\in\vec{X}(1)} |T_{uv}|$.

Definition 2.28. A topological space T is 1-*connected* if it is path-connected and each continuous map from the 1-sphere to T, $f : S^1 \to T$ is nullhomotopic, i.e. homotopic to a constant map that maps all points of S^1 to a single point $t_0 \in T$. We call a simplicial complex 1-connected if its geometric realization is 1-connected.

Remark 2.29. Let P, P' be two loops around a vertex v in X. If $P \overset{(BT)}{\sim} P'$ or $P \overset{(TR)}{\sim} P'$ then the corresponding paths in the geometric realization of X are homotopic. Conversely, any two loops P, P' in $|X|$ along the edges of X around v which are homotopic can be seen as loops in the abstract simplicial complex and we can find a sequence of loops $P = P_0, \ldots, P_m = P'$ such that $P_i \overset{(BT)}{\sim} P_{i+1}$ or $P_i \overset{(TR)}{\sim} P_{i+1}$ for all $0 \leqslant i \leqslant m-1$.

Remark 2.30. Any connected non-empty simplicial complex admits a 0-cone: pick an arbitrary vertex v_0, and by connectivity, we can find a path from any other vertex u to v_0. The converse is also true.

Similarly, any 1-connected (connected and simply connected) simplicial complex X admits a 1-cone function. By connectedness we get a 0-cone. Then each loop $P_u \circ (u; v)P_v^{-1}$ is contractible in the geometric realization of X. By Remark Remark 2.29, we can realize this homotopy equivalence by a sequence of paths connected by the relations $\overset{(BT)}{\sim}$, $\overset{(TR)}{\sim}$.

Lemma 2.31. *Let X be a finite n-dimensional simplicial complex with 1-cone* $\mathrm{Cone}_X = \left(v_0, (P_u)_{u\in X(0)}, (T_{uv})_{uv\in \vec{X}(1)}\right)$. *Then $|X|$ is 1-connected and hence $H^0(X,\Lambda) = 1$ and $H^1(X,\Lambda) = 1$ for any group Λ.*

Proof. Note that X is connected and hence $|X|$ is path connected. Hence we can without loss of generality consider a loop around v_0 in $|X|$. This loop is homotopic to a loop along the edges of X. Let $P = (v_0; u_1; \dots; u_m; v_0)$ denote this loop. Then

$$\begin{aligned} P =& (v_0) \circ (v_0; u_1, \dots, u_m, v_0) \sim_{T_{(u_1,v_0)}} P_{u_1} \circ (u_1; v_0) \circ (v_0; u_1, \dots, u_m, v_0) \\ & \overset{(BT)}{\sim} P_{u_1} \circ (u_1, \dots, u_m, v_0) \overset{(BT)}{\sim} \dots \overset{(BT)}{\sim} P_{u_1}(u_1; u_2) \circ P_{u_2}^{-1} \circ P_{u_2} \circ (u_2, \dots, u_m, v_0) \\ & \sim_{T_{(u_1,u_2)}} P_{u_2} \circ (u_2, \dots, u_m, v_0). \end{aligned}$$

Here we denoted by $\sim_{T_{(u_i,u_{i+1})}}$ the series of relations that are described by $T_{(u_i,u_{i+1})}$.

Further iterating this process yields

$$P \sim P_{u_m} \circ (u_m; v_0) \sim_{T_{(u_m,v_0)}} (v_0).$$

Thus, we have shown that P is contractible and hence $|X|$ is simply connected.

The vanishing of the cohomology groups follows from the fact that cohomology is homotopy invariant. □

A bound on the cone radius gives rise to a bound on the coboundary expansion, see Proposition 3.14 for the detailed statement. Thus, we can intuitively think of coboundary expansion as a quantitative way of (higher) connectedness.

In particular, the existence of a cone function with bounded cone radius implies coboundary expansion.

Next, we recall a method for bounding the cone radius using the Dehn function of a specific presentation. We start by recalling some facts that describe on the group theoretic level when a coset complex is 1-connected.

Lemma 2.32 (see e.g. [HS19, Lemma 3.3])**.** *Let G be a group with subgroups $\mathcal{H} = (H_0, \dots, H_d)$. Then $\mathcal{CC}(G;\mathcal{H})$ is connected (i.e. the underlying graph is connected) if and only if $G = \langle H_0, \dots, H_d\rangle$.*

Definition 2.33. Let G be a group and $H_i = \langle S_i \mid R_i\rangle, i \in I$ be a finite family of subgroups of G given by their presentation. Let $\phi_{ij} : H_i \cap H_j \hookrightarrow H_i$ denote the embedding of $H_i \cap H_j$ into H_i for all $i, j \in I$. The free product of the subgroups H_i amalgamated along their intersections is the group with presentation

$$\langle \bigcup_i S_i \mid \bigcup_i R_i \cup \{\phi_{ij}(h) = \phi_{ji}(h) \mid \text{ for all } i,j \in I, h \in H_i \cap H_j\}\rangle.$$

Intuitively, in the above definition, we consider the free product of the H_i but identify the elements in the intersection of the subgroups with their images in both subgroups.

Proposition 2.34 ([AH93, Theorem 2.4])**.** *Let G be a group with subgroups $\mathcal{H} = (H_0, \dots, H_d)$. Then $\mathcal{CC}(G;\mathcal{H})$ is 1-connected if and only if G is the free product of the subgroups H_i amalgamated along their intersections.*

For most of our work, we consider a specific presentation for a group associated to a coset complex.

Let G be a group with subgroups $H_i, i \in I$, where I is a finite set with $|I| \geqslant 3$. Assume that G is the free product of the subgroups H_i amalgamated along their intersections. For each $i \in I$ let R_i denote the set of all the nontrivial relations in the multiplication table of H_i of length at most 3. Then we obtain a presentation of G by considering all the elements of the H_i as generators and $\bigcup_i R_i$ as the set of relations. (We do not need to consider the relations given by the identifications in the definition of the amalgamated product, as we already consider the H_i as subgroups of G.)

Definition 2.35. Given a 1-connected coset complex $X = \mathcal{CC}(G;(H_i)_{i\in I})$, we denote by Dehn_X the Dehn function of the presentation of G obtained by considering all the elements of the H_i as generators and $\bigcup_i R_i$ as the set of relations as above.

We will see below in Theorem 3.19, following the work of [KO21], that a bound on the Dehn function of a coset complex implies a bound on its cone radius.

3 Proof of the Main Theorem

In this section we prove Theorem 3.1, our main theorem on swap coboundary expansion of congruence KMS complexes of type A_d. The statement below specifies the parameter choices and the coefficient groups for which the result holds.

Theorem 3.1. *Let $r \in \mathbb{N}$, $r \geqslant 10^4$. We fix the following parameters:*

- $d \in \mathbb{N}, d \geqslant r^{55}$,
- *a desired spectral expansion for the final complex* $\frac{r+1}{d-r^{11}-r^{10}-2} \leqslant \gamma_1 < \frac{1}{r^{11}}$,
- $n \in \mathbb{N}, r^{11} + r^{10} - 1 \leqslant n \leqslant d - \frac{r+1}{\gamma_1} - 1$,
- $q_0 \in \mathbb{N}$ *large enough with respect to the other parameters, polynomial in r.*

Then there exists $\beta > 0$ which is of order $\frac{1}{r^{34}}$ such that

- *for all prime powers $q = p^m \geqslant q_0$,*
- *for all groups Λ without non-trivial elements of order p,*
- *for all irreducible polynomials $f \in \mathbb{F}_q[t]$ with $\deg f \geqslant 2$,*

the n-skeleton of the KMS complex $X(A_d, \mathbb{F}_q, f)$ is a (β, r, Λ)-swap coboundary expander and γ_1-local spectral expander.

Corollary 3.2. *In the above setting, the KMS complex $X(A_d, \mathbb{F}_q, f)$ is a (β, r, Λ)-swap coboundary expander, proving Theorem 1.2.*

Note that it the full r-faces complex is in general not a local spectral expander, see Observation 2.13.

Proof. We need to show that $F^r X$ is a (β, Λ) 1-coboundary expander. Now, 1-coboundary expansion is fully determined by the 2-skeleton of the corresponding complex (using our convention about the weights of a skeleton). By Lemma 3.5 below, we have

$$(F^r X)^{\leqslant 2} = F^r(X^{\leqslant 2r+r+2}) = F^r((X^{\leqslant n})^{\leqslant 2r+r+2}) = (F^r(X^{\leqslant n}))^{\leqslant 2}$$

given $2r + r + 2 \leqslant n$ which holds with our choice of constants. □

Roadmap. The proof proceeds in four steps. The first three are standard, the fourth is the main new step of this work. We explain the reductions in this section, stating the main technical ingredients whose proofs appear in later sections.

1. **Reduction to colored faces complexes (Subsection 3.1).** We first reduce swap coboundary expansion to coboundary expansion of colored faces complexes for a sufficiently large fraction of choices of colors. We identify a collection of "nice" choices for which we will prove expansion. The color sampling lemma (Lemma 3.7), proved in Section 4, shows that a random choice is nice with high probability.
2. **Reduction to vertex links (Subsection 3.2).** For each nice choice of colors, we use the local-to-global theorem of [DD23c] to reduce the expansion bound to coboundary expansion of every vertex link. These links have a simpler algebraic description, which allows us to analyze them in the remaining steps.
3. **Reduction to diameter and Dehn bounds (Subsection 3.3).** We describe each vertex link as a coset complex and use the cone method to reduce its coboundary expansion to constructing a cone of bounded radius. In turn, this radius can be bounded using the diameter and the Dehn function of the coset complex, which measures how many applications of the defining relations are needed to reduce a word representing the identity. The precise connection between cones and Dehn functions is proved in Appendix B, extending the result of [KO21].

4. **Decomposition into bins (Subsection 3.3.1).** The group underlying each link decomposes as a direct product of smaller upper-triangular matrix groups, corresponding to intervals of colors called bins. Our main new ingredient is the decomposition lemma, proved in Section 6, which reduces the Dehn bound to separate bounds for these factors, with an additional cost polynomial in r and the word length. For each factor, we obtain a linear Dehn bound whose constant is independent of the matrix size. This uses the matrix-group theorem proved in Section 5; its application to the factors, together with the required diameter bound, is proved in Section 7.

Finally, in Subsection 3.4, we combine these ingredients and track the parameters to complete the proof of Theorem 3.1.

3.1 Step 1: Reduction to colored faces complex

The following result by Dikstein-Dinur reduces the question to showing coboundary expansion for a large set of colored faces complexes, instead of considering the full faces complex $F^r X$. This has the advantage that the colored faces complex of a coset complex can again be described as a coset complex.

Theorem 3.3 ([DD23c, Lemma 8.2])**.** *Let $T, d, r \in \mathbb{N}$, such that $3 \leqslant rT \leqslant (d+1)^{1/5}$. Let $\beta, p, \lambda \in (0,1]$. Let Λ be any group. Let X be a $(d+1)$-partite simplicial complex such that*

1. *$F^r X$ is a γ-local spectral expander for $\gamma < \frac{1}{rT}$, and*
2. $$\mathbb{P}_{\mathcal{I} \in \binom{[d]}{r+1}^{[T]}} \left[h^1(F^{\mathcal{I}} X, \Lambda) \geqslant \beta \wedge H^1(F^{\mathcal{I}} X, \Lambda) = 1 \right] \geqslant p.$$

Then $H^1(F^r X, \Lambda) = 1$ and $h^1(F^r X, \Lambda) \geqslant \Omega(p^2 \beta)$.

This statement is proven in [DD23c, Section 8.3] for the case where $T = \sqrt{r+1}$ and for X which is a quotient of the affine building of type $\tilde{A}$, but no properties of the building quotient were used in the proof other than the ones that appear in the theorem statement.

Proof intuition for Theorem 3.3. The proof of Theorem 3.3 in [DD23c, Section 8.3] is based on extending a coboundary from a colored faces complex to the entire faces complex. Given a 1-cochain ϕ on $Z = F^r X$, the goal is to show that if ϕ violates the cocycle condition on few triangles, then it is close to a coboundary.

The first step in the proof is to find a collection $\mathcal{I}$ of disjoint colors such that $Y = F^{\mathcal{I}} X$ has good coboundary expansion, and such that ϕ violates the cocycle condition on relatively few triangles involving vertices of Y. This is obtained by averaging over choices of colors, using the assumption that a p-fraction of the colored faces complexes are coboundary expanders. The argument also considers larger collections of colors, whose expansion is deduced from that of smaller restrictions.

Having found such a subcomplex Y, its coboundary expansion implies that there is a 0-cochain g on its vertices such that $\phi = d_0 g$ on almost all edges of Y. The next step is to extend g to the remaining vertices. Every neighbor $u \in Y$ of a vertex $v \notin Y$ suggests a value for $g(v)$, determined by $g(u)$ and $\phi(uv)$. The key observation is that two such suggestions agree whenever they come from vertices u, w for which the edge uw and triangle uvw satisfy the relevant coboundary and cocycle conditions. The spectral expansion of the links ensures that most suggestions are consistent, allowing $g(v)$ to be defined by the most common suggestion.

Finally, the values of g are shown to agree with ϕ on most edges whose endpoints both lie outside Y. Indeed, for such an edge vw, a triangle uvw with $u \in Y$ allows one to recover $\phi(vw)$ from the other two edges, provided the cocycle condition holds. The remaining technical part of the proof consists of showing that all these comparisons involve sufficiently many triangles, and bounding the resulting errors using local spectral expansion. This yields the claimed coboundary expansion bound.

To be able to apply Theorem 3.3, we need to show local spectral expansion for the faces complex.

Lemma 3.4. *Let $r, n, d \in \mathbb{N}$ satisfy $2r+1 \leqslant n \leqslant d$, and let $\gamma_1 \in (0,1)$. Suppose that*

$$d \geqslant n + \frac{r+1}{\gamma_1} + 1.$$

For every prime power $q > 16$ and every irreducible $f \in \mathbb{F}_q[t]$ of degree at least two, the complex $F^r(X^{\leqslant n})$, where $X = X(A_d, \mathbb{F}_q, f)$, is a γ_1-local spectral expander.

Proof. By Theorem 2.20, the complex $Y = X^{\leqslant n}$ is a $\frac{1}{d-n+1}$-two-sided local spectral expander. The graph underlying each link of $F^r Y$ is the swap walk between disjoint $(r+1)$-vertex faces in a link of Y. The swap-walk bound of [GLL22, Lemma 3.3] therefore bounds its nontrivial eigenvalues in absolute value by

$$\frac{r+1}{d-n+1} \leqslant \gamma_1.$$

□

Thus, given that we choose the parameters (r, n, d, γ, q_0) as in Theorem 3.1 and setting $X = X(A_d, \mathbb{F}_q, f)$ and Y the n-skeleton of X, we get that $F^r Y$ is a γ_1-local spectral expander. Given our choice of r, n, d the colored faces complexes of X and Y are the same, as the following lemma shows.

Lemma 3.5. *Let X be a $(d+1)$-partite complex, and fix integers $1 \leqslant r, n < d, 1 \leqslant T \leqslant \frac{d+1}{r+1}$. Let $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$ be a set of T disjoint size $r+1$ color sets. Recall that $X^{\leqslant n}$ denotes the n-skeleton of X. Then*

1. $(F^r X)^{\leqslant n} = F^r(X^{\leqslant nr+n+r})$,
2. $\left(F^{\mathcal{I}} X\right)^{\leqslant n} = F^{\mathcal{I}}(X^{\leqslant nr+n+r})$,
3. *if $n \geqslant T(r+1) - 1$ then $F^{\mathcal{I}} X = F^{\mathcal{I}}(X^{\leqslant n})$.*

In all three cases, the equility also respects the weights on the complexes, using the convention introduced in Subsection 2.1 for the weights on skeleta.

Proof. Following the definition of $F^r X$ we have for a set $\{\sigma_0, \ldots, \sigma_i\}$ with $\sigma_j \in X(r)$ pairwise disjoint

$$\{\sigma_0, \ldots, \sigma_i\} \in F^r X(i) \iff \cup_j \sigma_j \in X(ir + r + i)$$

where we use that, since the σ_j are pairwise disjoint, $|\cup_j \sigma_j| = \sum_j |\sigma_j| = (i+1)(r+1)$, and hence if $\cup_j \sigma_j \in X$ then it will be of dimension $(ir + i + r)$.

Now when considering the n-skeleton of $F^r X$ we have

$$\{\sigma_0, \ldots, \sigma_i\} \in (F^r X)^{\leqslant n} \iff \cup_j \sigma_j \in X(ir + r + i) \wedge i \leqslant n \iff \cup_j \sigma_j \in X^{\leqslant nr+r+n}$$

which shows point 1.. The proof for point 2. is analogues.

Point 3. follows directly from point 2. as $\dim F^{\mathcal{I}} X = T - 1$, hence $F^{\mathcal{I}} X = (F^{\mathcal{I}} X)^{\leqslant T-1} = F^{\mathcal{I}}(X^{\leqslant (T-1)r+(T-1)+r}) = F^{\mathcal{I}}(X^{\leqslant T(r+1)-1})$.

Regarding the weights, note that for each $0 \leqslant i \leqslant n$ by our convention, the weight on the i-simplices of $F^r X$ and $(F^r X)^{\leqslant n}$ coincide. On the other hand, we have that $w_X, i = w_{X^{\leqslant nr+n+r}}, i$ for all $0 \leqslant i \leqslant nr + n + r$ and thus they induce the same weight on the r-faces complex. Together we get

$$w_{(F^r x)^{\leqslant n}, i} = w_{F^r X, i} = w_{F^r(X^{\leqslant nr+n+r}), i}$$

for all $0 \leqslant i \leqslant n$ as required. The argument for the other two points is analogous. □

Note that a statement similar to part 3. for $F^r X$ is not very useful, since $\dim F^r X = \lfloor \frac{d+1}{r+1} \rfloor - 1$ and hence we would require $n \geqslant (\lfloor \frac{d+1}{r+1} \rfloor - 1 + 1)(r+1) - 1$, which in the case where $d+1$ is divisible by $r+1$ results in requiring $n \geqslant d$ for which the statement is trivial.

Using the above lemma, we will from now on consider $F^{\mathcal{I}} X$ instead of the colored faces complex of the skeleton. For part 2. of Theorem 3.3 we show coboundary expansion for all colored faces complexes given by "nice" sets of colors. Furthermore, we show that sets of colors are nice with high probability.

The conditions defining a nice collection will describe the matrix blocks that appear in vertex links. Before defining them, let us explain why these blocks correspond to intervals in a linear order, even though the root indices of the KMS complex are cyclic. Let $X = X(A_d, \mathbb{F}_q, f)$, let $s \in X(r)$ represent a vertex of $F^{\mathcal{I}} X$, and

write $J = \mathrm{col}(s)$. Choose a vertex $u \in s$ and write $s = s' \cup \{u\}$. A collection of vertices belongs to $(F^{\mathcal{I}}X)_s$ exactly when its members have colors in $\mathcal{I} \setminus \{J\}$, are pairwise disjoint, and their union together with s is a face of X. Equivalently, their union is a face of X_s. Since $X_s = (X_u)_{s'}$, we obtain

$$(F^{\mathcal{I}}X)_s = F^{\mathcal{I}\setminus\{J\}}(X_s) = F^{\mathcal{I}\setminus\{J\}}\big((X_u)_{s'}\big).$$

Thus we can study the same link by first passing to X_u and then taking the link of s' there. The vertices of the resulting faces complex still represent faces with $r+1$ vertices.

As described in Subsection 2.3, the vertex link X_u is a coset complex of upper triangular unipotent matrices. If u has color j_*, its simple root indices follow the cyclic order $j_*+1, \dots, d, 0, \dots, j_*-1$, which becomes a path after deleting j_*. We relabel this path as $1, \dots, d$. The remaining r colors of s' divide it into $r+1$ intervals, and the root groups in these intervals act on disjoint diagonal blocks. These are the bins used in the definition below.

For the rest of the paper we will need a parameter polynomial in r to denote the number of colors we want to restrict to. We will keep this parameter abstract in statements but can set it to $T = r^{10}$ for concreteness.

Definition 3.6. 1. Let $J \subseteq [d]$ have size $r+1$. Choose $j_0 \in J$ and cyclically relabel all colors so that j_0 becomes 0 and the colors following it become $1, \dots, d$. Write the relabeled set as $J = \{0 = j_0 < j_1 < \cdots < j_r\}$, and set $j_{r+1} = d+1$. We define the *bins of* J, or J*-bins*, by

$$B_\ell = \{j_\ell + 1, \dots, j_{\ell+1} - 1\}, \qquad 0 \leqslant \ell \leqslant r.$$

We write $B = [a, b]$ as a shorthand for $B = \{a, a+1, \dots, b\}$. When considering the bins of a fixed J, we apply the same relabeling to every color set in $\mathcal{I}$. Choosing a different $j_0 \in J$ only cyclically permutes the bins and translates their coordinates, preserving the order within each bin. Thus the conditions below do not depend on this choice. Empty bins contribute trivial factors to the group decomposition and are omitted from the niceness requirements.

2. Let $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$ be a set of T disjoint color sets. We say $\mathcal{I}$ is *nice* if the following holds for every $J \in \mathcal{I}$, in the path order chosen for J.
   (a) Every bin $B = [a, b]$ has size $|b - a| > 100$.
   (b) For every bin $B = [a, b]$ of J there exists a color $I' \in \mathcal{I} \setminus \{J\}$ such that $I' \cap B \subseteq [a, a + \lfloor \frac{b-a}{5} \rfloor]$.
   (c) For every bin $B = [a, b]$ of J there exists a color $I' \in \mathcal{I} \setminus \{J\}$ such that $I' \cap B \subseteq [b - \lfloor \frac{b-a}{5} \rfloor, b]$.
   (d) For every bin $B = [a, b]$ of J there exists a color $I' \in \mathcal{I} \setminus \{J\}$ such that $I' \cap B \subseteq [a + 2\lceil \frac{b-a}{5} \rceil, a + 3\lfloor \frac{b-a}{5} \rfloor]$.
   (e) For every pair of J-bins $B = [a, b], B' = [a', b']$ and every $\ell, \ell' \in \{0, 1, 2, 3\}$ there exists some $I' \in \mathcal{I} \setminus \{J\}$ such that

$$I' \cap [a + \ell \lfloor \frac{b-a}{5} \rfloor, \min\{a + (\ell+2)\lceil \frac{b-a}{5} \rceil + 4, b\}] = \emptyset, \text{ and}$$

$$I' \cap [a' + \ell' \lfloor \frac{b'-a'}{5} \rfloor, \min\{a' + (\ell'+2)\lceil \frac{b'-a'}{5} \rceil + 4, b'\}] = \emptyset.$$

Indeed, sufficiently many sets of colors are nice:

Lemma 3.7. *Let $r \geqslant 10^4, T \geqslant r^{10}$ and $d \geqslant 100(r+1)^2T^2$ then with probability $\geqslant 0.9$ a randomly selected set of T disjoint colors $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$ is nice.*

We prove this lemma in Section 4.

The main next step is now to show that the nice colored faces complexes are coboundary expanders.

Proposition 3.8. *Let $r, T, d \in \mathbb{N}$ such that $T \geqslant 4$ and $(r+1)T \leqslant d+1$. There exist $\beta_1 > 0, q_0 \in \mathbb{N}$ such that for all prime powers $q = p^m \geqslant q_0$, all groups Λ without non-trivial elements of order p, all irreducible polynomials $f \in \mathbb{F}_q[t]$ of degree $\geqslant 2$ and all nice sets of T disjoint colors $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$, we have*

$$h^1(F^{\mathcal{I}}X, \Lambda) \geqslant \beta_1 \text{ and } H^1(F^{\mathcal{I}}X, \Lambda) = 1$$

where $X = X(A_d, \mathbb{F}_q, f)$.

We start by showing the vanishing of the first cohomology group.

Lemma 3.9. *Let $r, d, T \in \mathbb{N}, d, T \geqslant 3$. For all prime powers $q = p^m \geqslant 5$, all groups Λ without non-trivial elements of order p, all irreducible polynomials $f \in \mathbb{F}_q[t]$ of degree $\geqslant 2$ and all nice sets of disjoint colors $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$, we have*

$$H^1(F^{\mathcal{I}}X(A_d, \mathbb{F}_q, f), \Lambda) = 1.$$

Proof. Let $X = X(A_d, \mathbb{F}_q, f)$. Recall that

$$X = \mathcal{CC}(G; (H_i^f)_{i\in[d]})$$

where $G = \mathrm{SL}_{d+1}(\mathbb{F}_q[t]/(f))$ and $H_i^f = \pi_f(U_{[d]\setminus\{i\}})$ with $\pi_f : \mathrm{SL}_{d+1}(\mathbb{F}_q[t]) \to \mathrm{SL}_{d+1}(\mathbb{F}_q[t]/(f))$, $U_i = \langle e_{i,i+1}(\lambda) \mid \lambda \in \mathbb{F}_q\rangle \leqslant \mathrm{SL}_{d+1}(\mathbb{F}_q[t]), U_0 = \langle e_{d+1,1}(\lambda t) \mid \lambda \in \mathbb{F}_q\rangle$ and $U_J = \langle U_i \mid i \in J\rangle$ for al $J \subseteq [d]$. Let $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$ be the set of disjoint colors from above. For $J \in \mathcal{I}$ we set $H_J^f = \bigcap_{i\in J} H_i^f$. Then by Lemma 2.15 we can describe $F^{\mathcal{I}}X$ as

$$F^{\mathcal{I}}X = \mathcal{CC}(G; (H_J)_{J\in\mathcal{I}}).$$

We want to apply basically the same argument as in [OV25], thus we also refer to that work for more details on the terminology.

We can see each such complex (i.e. for each choice of f) as a quotient of the same underlying infinite complex, which we describe now. Let $\mathrm{KMS}(d, q) \leqslant \mathrm{SL}_{d+1}(\mathbb{F}_q[t])$ be the corresponding concrete KMS group, denote $H_i = U_{[d]\setminus\{i\}}$, and consider the infinte complex

$$Y = \mathcal{CC}(\mathrm{KMS}(d, q); (H_i)_{i\in[d]}).$$

We refer to this as fundamental KMS complex. Let φ_f denote the map π_f restricted to $\mathrm{KMS}(q, d)$ and let $N_f = \ker\varphi_f$. In [OV25, Lemma 3.5], it was shown that $X \cong N_f\backslash Y$. Analogously, it holds that $F^{\mathcal{I}}X \cong N_f\backslash F^{\mathcal{I}}Y$. We want to apply the following result.

Theorem 3.10 ([KOW25])**.** *Let X be a connected pure 2-dimensional simplicial complex and N be a group acting on X. Assume $p : X \to N\backslash X$ is rigid, i.e. it is simplicial and for all $\sigma \in X$ we have $|\sigma| = |p(\sigma)|$. For every group Λ, if $H^1(X, \Lambda) = 0$ and $H^1(N, \Lambda) = 0$ then $H^1(N\backslash X, \Lambda) = 0$ (the cohomology of N is taken with the trivial action of N on Λ).*

In [OV25, Cororllary 3.7] it was shown that if Λ has no non-trivial elements of order p (recall $q = p^m$), then $H^1(N_f, \Lambda) = 1$. Thus, it remains to show that $H^1(F^{\mathcal{I}}Y, \Lambda) = 1$ which we will achieve by showing that $F^{\mathcal{I}}Y$ is 1-connected using Proposition 2.34. We know by the work of [DM07; Abr96] (see also [GV25, Proposition 3.8]) that $\mathrm{KMS}(d, q)$ is the free product of the subgroups $U_{\{i,j\}}, \{i, j\} \subseteq [d]$ amalgamated along their intersections, given $d \geqslant 3, q \geqslant 5$. Hence, it remains to show that for each $\{i, j\} \subseteq [d]$ there exists an $L \in \mathcal{I}$ such that $\{i, j\} \subseteq [d] \setminus L$ (and hence $U_{\{i,j\}} \subseteq H_L = \cap_{i\in L} U_{[d]\setminus\{i\}} = U_{[d]\setminus L}$. This clearly holds if $|\mathcal{I}| \geqslant 3$, as all sets in $\mathcal{I}$ are disjoint. ◻

We further reduce the task of bounding the cocycle expansion constant in the next step.

3.2 Step 2: Reduction to the links via the local-to-glocal machinery

To show coboundary expansion of the nice colored faces complexes, we reduce the question to showing coboundary expansion for all vertex links via the local-to-global result by Dikstein-Dinur.

Theorem 3.11 ([DD23b, Theorem 8])**.** *Let $0 \leqslant \lambda < 1, \beta > 0$ be constants and Λ be a group. For a finite pure n-dimensional simplicial complex Y with $n \geqslant 3$, if*

- *for every vertex $v \in Y(0)$, $h^1(Y_v, \Lambda) \geqslant \beta$ and $H^1(Y_v, \Lambda)$ is trivial,*
- *and Y is a λ-local spectral expander,*

then

$$h^1(Y,\Lambda) \geqslant \frac{(1-\lambda)\beta}{24} - e\lambda,$$

where $e \approx 2.71$ *is the Euler constant.*

To apply this theorem, we first of all need to show local spectral expansion for $F^{\mathcal{I}}X$.

Lemma 3.12. *Let* $X = X(A_d, \mathbb{F}_q, f)$*, where* $q > 16$ *and* f *is irreducible of degree at least two. Let* $\mathcal{I}$ *consist of* $T \geqslant 3$ *pairwise disjoint color sets, each of size* $r+1$*. Then* $F^{\mathcal{I}}X$ *is a* $\frac{2(r+1)}{\sqrt{q}-2}$*-local spectral expander.*

Proof. We first bound the bipartite walk between any two specified colors in any link X_τ. Write $D = d - |\tau|$ and $\eta_q = 2/(\sqrt{q}-2)$. If this bipartite walk has a nontrivial singular value s, extending its centered singular vectors by zero to the other colors gives a mean-zero function on X_τ with Rayleigh quotient s/D. By Theorem 2.20, this is at most η_q/D, so $s \leqslant \eta_q$.

Thus every two-color walk in every link has nontrivial singular values at most η_q. By [GLL22, Lemma 3.3], the walk between faces on two disjoint color sets I, J has nontrivial singular values at most $\sqrt{|I||J|}\eta_q$. For two members of $\mathcal{I}$, this is $(r+1)\eta_q = \lambda$.

The same bound holds after fixing any face of $F^{\mathcal{I}}X$, since this amounts to passing to a link of X. In each such link, decompose a function into its mean on each remaining color class and a function with mean zero on every class. The first subspace carries the walk on a complete graph, whose nonconstant eigenvalues are negative. On the second subspace, averaging the bipartite walks between classes gives operator norm at most λ. Hence every link has second eigenvalue at most λ. □

Finally, we need to show that the links are coboundary expanders.

Proposition 3.13. *Let* $r, d \in \mathbb{N}, d \geqslant 3$ *and* $T \in \mathbb{N}, T \geqslant 4$*. There exist* $\beta_2 > 0, q_0 \in \mathbb{N}$ *such that for all prime powers* $q = p^m \geqslant q_0$*, all groups* Λ *without non-trivial elements of order* p*, all irreducible polynomials* $f \in \mathbb{F}_q[t]$ *of degree* $\geqslant 2$ *and all nice sets of* T *disjoint colors* $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$*, we have for all vertices* $v \in F^{\mathcal{I}}X(0)$

$$h^1(F^{\mathcal{I}}X_v, \Lambda) \geqslant \beta_2 \text{ and } H^1(F^{\mathcal{I}}X_v, \Lambda) = 1$$

where $X = X(A_d, \mathbb{F}_q, f)$*.*

In the next step, we further reduce the statement of this proposition using the cone method.

3.3 Step 3: Coboundary expansion of the links via the cone method

The next step is to show that each link $F^{\mathcal{I}}X_v$ admits a 1-cone function and bounding its radius. Then the following provides a lower bound on the cocycle expansion constant.

Proposition 3.14 ([DD23c, Lemma 1.6, Lemma 4.3])**.** *Let* Y *be a pure* d*-dimensional finite simplicial complex,* $d \geqslant 2$*. Assume that there is a 1-cone on* Y *such that* $\mathrm{Rad}_i(\mathrm{Cone}_Y) < \infty, i = 0, 1$ *and that the automorphisms group* $\mathrm{Aut}(Y)$ *acts transitively on* $Y(d)$*. Then for every group* Λ*,*

$$h^1(Y,\Lambda) \geqslant \frac{1}{\binom{d+1}{3}\,\mathrm{Rad}_1(\mathrm{Cone}_Y)}.$$

Note that any coset complex satisfies the symmetry assumption about the action of the automorphism group.

Thus, our goal is to show the following.

Proposition 3.15. *Let* $r, d \in \mathbb{N}, d \geqslant 3$*, and* $T \in \mathbb{N}, T \geqslant 4$*. There exist* $\beta_3 > 0, q_0 \in \mathbb{N}$ *such that for all prime powers* $q = p^m \geqslant q_0$*, all irreducible polynomials* $f \in \mathbb{F}_q[t]$ *of degree* $\geqslant 2$*of degree* $\geqslant 2$ *and all nice sets of* T *disjoint colors* $\mathcal{I} \subseteq \binom{d+1}{r+1}$*, we have for all vertices* $v \in F^{\mathcal{I}}X(0)$ *that* $(F^{\mathcal{I}}X)_v$ *admits a 1-cone function* Cone_v *with* $\mathrm{Rad}_1(\mathrm{Cone}_v) \leqslant \beta_3$ *(again* $X = X(A_d, \mathbb{F}_q, f)$*).*

To show the existence of a cone function, we have a closer look at the structure of links.

Let $X = X(A_d, \mathbb{F}_q, f)$. For every vertex $v \in F^{\mathcal{I}}X(0)$ we can describe $F^{\mathcal{I}}X_v$ explicitly as coset complex. Recall that

$$X = \mathcal{CC}(G; (H_i^f)_{i\in[d]})$$

where $G = \mathrm{SL}_{d+1}(\mathbb{F}_q[t]/(f))$ and $H_i^f = \pi_f(U_{[d]\setminus\{i\}})$ where $\pi_f : \mathrm{SL}_{d+1}(\mathbb{F}_q[t]) \to \mathrm{SL}_{d+1}(\mathbb{F}_q[t]/(f))$, $U_i = \langle e_{i,i+1}(\lambda) \mid \lambda \in \mathbb{F}_q\rangle \leqslant \mathrm{SL}_{d+1}(\mathbb{F}_q[t]), U_0 = \langle e_{d+1,1}(\lambda t) \mid \lambda \in \mathbb{F}_q\rangle$ and $U_J = \langle U_i \mid i \in J\rangle$ for al $J \subseteq [d]$. Let $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$ be the set of disjoint colors from above. For $J \in \mathcal{I}$ we set $H_J = H_J^f = \bigcap_{i\in J} H_i^f$. Then by Lemma 2.15 we can describe $F^{\mathcal{I}}X$ as

$$F^{\mathcal{I}}X = \mathcal{CC}(G; (H_J)_{J\in\mathcal{I}}).$$

Thus, for a vertex $gH_J \in F^{\mathcal{I}}X$, we can describe its link again as coset complex using Lemma 2.6 as follows.

$$F^{\mathcal{I}}X_{gH_J} \cong \mathcal{CC}(H_J; (H_J \cap H_I)_{I\in\mathcal{I}\setminus\{J\}}).$$

We can now use the fact ([GV25, Lemma 4.2, 4.6]) that $H_J = \pi_f(U_{[d]\setminus J}) \cong U_{[d]\setminus J}$ and get

$$(F^{\mathcal{I}}X)_{gH_J} \cong \mathcal{CC}(U_{[d]\setminus J}; (U_{[d]\setminus(J\cup I)})_{I\in\mathcal{I}\setminus\{J\}}).$$

For each fixed J, we now use the upper triangular matrix model obtained by first passing to X_u, with $\mathrm{col}(u) = j_* \in J$, as explained before Definition 3.6. We relabel the color sets and root groups in the path order chosen there, retaining the notation U_i for the corresponding simple root groups. Thus $0 \in J$, and for $1 \leqslant i \leqslant d$ the group U_i consists of the matrices $e_{i,i+1}(\lambda)$, $\lambda \in \mathbb{F}_q$. We use these coordinates whenever J is fixed in the rest of this section and in Section 6 and Section 7.

Corollary 3.16. *Every link $(F^{\mathcal{I}}X)_v$ is a $(|\mathcal{I}\setminus\{J\}| - 1 = T - 2)$-dimensional coset complex. Hence its automorphism group acts transitively on the maximal, i.e. $(T-2)$-dimensional, simplices. Thus, if $\mathrm{Rad}_1(\mathrm{Cone}_v) \leqslant \beta_3$ then $h^1((F^{\mathcal{I}}X)_v, \Lambda) \geqslant \frac{1}{\binom{T-1}{3}\beta_3}$.*

But before showing the bound on the cone radius, we show that the links are 1-connected and hence in particular have vanishing first cohomology group.

Lemma 3.17. *Assume that the size of the underlying field $q \geqslant 5$. Then the complex $\mathcal{CC}(U_{[d]\setminus J}; (U_{[d]\setminus(J\cup I)})_{I\in\mathcal{I}\setminus\{J\}})$ is connected if $|\mathcal{I}| \geqslant 3$ and it is 1-connected if $|\mathcal{I}| \geqslant 4$. Thus in particular $H^1(F^{\mathcal{I}}X_{gH_J}, \Lambda) = 1$ if $|\mathcal{I}| \geqslant 4$.*

Proof. Let $Y = \mathcal{CC}(U_{[d]\setminus J}; (U_{[d]\setminus(J\cup I)})_{I\in\mathcal{I}\setminus\{J\}})$. Then by Proposition 2.32, Y is connected if and only if $U_{[d]\setminus J} = \langle U_{[d]\setminus(J\cup I)} \mid I \in \mathcal{I}\setminus\{J\}\rangle$. The later is equivalent to

$$\forall i \in [d]\setminus J\ \exists L \in \mathcal{I}\setminus\{J\} : i \in [d]\setminus(L\cup J) \iff \bigcup_{L\in\mathcal{I}\setminus\{J\}} [d]\setminus(L\cup J) = [d]\setminus J.$$

Now, we have

$$\bigcup_{L\in\mathcal{I}\setminus\{J\}} [d]\setminus(L\cup J) = \left(\bigcup_{L\in\mathcal{I}\setminus\{J\}} [d]\setminus L\right) \cap [d]\setminus J = \left([d]\setminus \bigcap_{L\in\mathcal{I}\setminus\{J\}} L\right) \cap [d]\setminus J.$$

Hence Y is connected if and only if $\bigcap_{L\in\mathcal{I}\setminus\{J\}} L = \emptyset$ which, since the color sets in $\mathcal{I}$ are distinct, holds if and only if $|\mathcal{I}| \geqslant 3$.

For the 1-connectedness, we want to use Proposition 2.34. We know by the work of [Tit86; Abr96] (see also [GV25, Proposition 3.8]) that $U_{[d]\setminus J}$ is the free product of the subgroups $U_{\{i,j\}}, \{i,j\} \subseteq [d]\setminus J$ amalgamated along their intersections. Hence, it remains to show that for each $\{i,j\} \subseteq [d]\setminus J$ there exists an $L \in \mathcal{I}\setminus\{J\}$ such that $\{i,j\} \subseteq [d]\setminus L$. To show that this holds if $|\mathcal{I}| \geqslant 4$, we fix $\{i,j\} \subseteq [d]\setminus J, i \neq j$. Both i,j can be contained in at most one set from $\mathcal{I}$, call them $i \in L_i, j \in L_j$ respectively. Since $|\mathcal{I}| \geqslant 4$, there will be at least one set in $L \in \mathcal{I}\setminus\{J, L_i, L_j\}$. Then $i,j \notin L$ and hence $\{i,j\} \subseteq [d]\setminus L$ as required. ◻

Thus we know that, given $|\mathcal{I}| \geqslant 4$, every link $F^{\mathcal{I}}X_{gH_J}$ is 1-connected and thus admits a cone function, see Remark Remark 2.30.

Remark 3.18. The statement above holds in a more general setting. Consider the coset complex $X = \mathcal{CC}(G;(H_i)_{i\in[d]})$, set for $J \subseteq [d], H_J = \cap_{j\in J}H_j$ and assume $\langle H_J, H_L\rangle = H_{J\cap L}$ for all $J, L \subseteq [d]$ (this for example holds if $G, H_i, i \in [d]$ form a subgroup geometric system in the sense of [KO23]). Let $\mathcal{I} \subseteq \binom{[d]}{r+1}$ be a set of disjoint colors. Assume that the chamber system generated by $G, H_i, i \in [d]$ is m-simply connected, in the sense of e.g. [GV25, Definition 2.13]. Then $F^{\mathcal{I}}X$ is 1-connected if $|\mathcal{I}| \geqslant m+1$. The general proof is similar to the one stated above. Since we do not need the general result here, we will not discuss the detailed proof here.

It remains to bound the radius of the cone function. We will do that by considering the Dehn function of a specific presentation of $U_{[d]\setminus J}$ in the next step.

3.3.1 Bounding the cone radius using the Dehn function

We bound the radius of the cone function using the Dehn function of a specific presentation using the following.

Theorem 3.19. *Let G be a finite group with subgroups H_i, $i \in I$, where $|I| \geqslant 3$, and assume that the coset complex $Y = \mathcal{CC}(G;(H_i)_{i\in I})$ is 1-connected. Let Dehn_Y be as in Definition 2.35, and write $D = \mathrm{diam}(Y)$. Then Y admits a 1-cone Cone_Y satisfying*

$$\mathrm{Rad}_0(\mathrm{Cone}_Y) \leqslant D, \qquad \mathrm{Rad}_1(\mathrm{Cone}_Y) \leqslant 15\big(\mathrm{Dehn}_Y(2D+1) + 2D + 1\big).$$

The relation between Dehn functions and fillings of coset complexes was previously studied in [KO21, Theorem 6.6] where this theorem was first discovered. We extend their theorem in a rather straightforward manner in Appendix B, where we also prove that the converse direction is also true — a cone can also serve as a bound to the Dehn function.

Thus, in order to bound the cone radius, it suffices to bound the diameter and the Dehn function of the relevant coset complex. We will prove the following theorem.

Note that the polynomial p is increasing in both variables. Thus, we have reduced the original problem to showing the following theorem.

Theorem 3.20. *Let $r, d \in \mathbb{N}, d \geqslant 3$, and $T \in \mathbb{N}, T \geqslant 4$. There exist $\beta_4, \beta_5 > 0, \beta_6, \beta_7 \geqslant 0$ and $q_0 \in \mathbb{N}$ such that for every prime power $q \geqslant q_0$ and every nice set of T disjoint colors $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$ for every $J \in \mathcal{I}$ the complex*

$$Y = \mathcal{CC}(U_{[d]\setminus J};(U_{[d]\setminus(J\cup I)})_{I\in\mathcal{I}\setminus\{J\}})$$

has $\mathrm{diam}(Y) \leqslant \beta_4$ and $\mathrm{Dehn}_Y(m) \leqslant \beta_5 m^2 + \beta_6 m + \beta_7$ for all $m \in \mathbb{N}$.

Thus, it remains to bound the diameter and the Dehn function of Y. For the diameter, we get the following bound, which will be proven in Section 7, see Proposition 7.1.

Lemma 3.21. *It holds that $\mathrm{diam}(Y) \leqslant 6(r+1)+1$.*

We reduce the task of bounding the Dehn function further below.

3.4 Step 4: Decomposing into bins

Let $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$ be a nice set of T disjoint color sets, fix $J \in \mathcal{I}$, and let $B_0, \dots, B_r$ be its bins in the path order of Definition 3.6. In the upper triangular matrix model, a nonempty bin $B = [a, b]$ acts on the coordinate block $\{a, \dots, b+1\}$, so $U_B \cong \mathrm{Unip}_{b-a+2}(\mathbb{F}_q)$. The blocks corresponding to different bins are disjoint, and every simple root indexed by $[d] \setminus J$ belongs to one of the bins. Therefore the group decomposes as the internal direct product

$$U_{[d]\setminus J} = \prod_{\ell=0}^{r} U_{B_\ell}.$$

Let $Y_\ell = \mathcal{CC}(U_{B_\ell};(U_{B_\ell\setminus I})_{I\in\mathcal{I}\setminus\{J\}})$.

The following lemma is a specific case of [Abr96, Lemma 16].

Lemma 3.22.

$$\mathcal{CC}(U_{[d]\setminus J};(U_{[d]\setminus(J\cup\{i\})})_{i\in[d]\setminus J}) = *_{\ell=0}^{r}\mathcal{CC}(U_{B_\ell};(U_{B_\ell\setminus\{i\}})_{i\in B_\ell})$$

With this in mind, we investigate the colored faces complex of a join.

Lemma 3.23. *Let X be a finite, pure d-dimensional, $(d+1)$-partite simplicial complex. Assume that $X = Y_1 * Y_2$ is the join of two complexes Y_1, Y_2, and let $\mathcal{I}$ be a collection of pairwise disjoint nonempty color sets. Write $C_i = \mathrm{col}(Y_i)$ and $\mathcal{I}_i = (J \cap C_i)_{J\in\mathcal{I}}$ for $i = 1, 2$, retaining the original label J on each restricted color set, including empty ones.*

Define the colored join $F^{\mathcal{I}_1}Y_1 \overset{*}{\times} F^{\mathcal{I}_2}Y_2$ *as follows. Its vertices are pairs (σ_1, σ_2) such that $\sigma_i \in Y_i$ and, for some $J \in \mathcal{I}$, $\mathrm{col}(\sigma_i) = J \cap C_i$ for $i = 1, 2$. We assign this vertex color J. Empty faces are allowed in either coordinate. A collection $\{(\sigma_{1,j}, \sigma_{2,j})\}_{j=0}^{k}$ forms a face if its vertices have distinct colors and $\bigcup_{j=0}^{k}\sigma_{i,j} \in Y_i$ for each $i = 1, 2$.*

Then the map $(\sigma_1, \sigma_2) \mapsto \sigma_1 \sqcup \sigma_2$ induces an isomorphism

$$F^{\mathcal{I}_1}Y_1 \overset{*}{\times} F^{\mathcal{I}_2}Y_2 \cong F^{\mathcal{I}}X.$$

Proof. Every face $\sigma \in X = Y_1 * Y_2$ decomposes uniquely as $\sigma = \sigma_1 \sqcup \sigma_2$ with $\sigma_i \in Y_i$. If $\mathrm{col}(\sigma) = J$, then $\mathrm{col}(\sigma_i) = J \cap C_i$. Thus the proposed map is a color-preserving bijection between the vertex sets.

Let $\{(\sigma_{1,j}, \sigma_{2,j})\}_{j=0}^{k}$ be a collection of vertices, and write $\sigma_j = \sigma_{1,j} \sqcup \sigma_{2,j}$. Their images form a face of $F^{\mathcal{I}}X$ precisely when their colors are distinct and $\bigcup_{j=0}^{k}\sigma_j \in X$. By the definition of the join,

$$\bigcup_{j=0}^{k}\sigma_j \in X \iff \bigcup_{j=0}^{k}\sigma_{i,j} \in Y_i \text{ for each } i = 1, 2.$$

These are exactly the conditions defining a face of the 'colored join', proving the isomorphism. □

We can iterate the result and get the following corollary.

Corollary 3.24. *Let $X = Y_1 * \cdots * Y_k$ be a finite, pure d-dimensional, $(d+1)$-partite simplicial complex. Further, let $\mathcal{I}$ be a set of distinct colors. Then*

$$F^{\mathcal{I}}X = \overset{*}{\prod_{1\leqslant i\leqslant k}} F^{\mathcal{I}_i}Y_i.$$

Here $\overset{*}{\prod}$ denotes the straightforward generalization of $\overset{*}{\times}$ to multiple complexes.

Let $X = \mathcal{CC}(U_{[d]\setminus J};(U_{[d]\setminus(J\cup\{i\})})_{i\in[d]\setminus J}), Y = F^{\mathcal{I}}X$. By the above results, we have

$$Y = \overset{*}{\prod_{0\leqslant\ell\leqslant r}} \mathcal{CC}(U_{B_\ell};(U_{B_\ell\setminus I})_{I\in\mathcal{I}\setminus\{J\}}).$$

Denote by $Y_\ell = \mathcal{CC}(U_{B_\ell};(U_{B_\ell\setminus I})_{I\in\mathcal{I}\setminus\{J\}})$ the ℓ-th component of this product.

Lemma 3.25. *Assume that every Y_ℓ has a bound of $\mathrm{Dehn}_{Y_\ell}(t) \leqslant At$ on its Dehn function. Then Y has a bound of $\mathrm{Dehn}_Y(t) \leqslant O((t(r+1))^2) + (r+1)At$ on its Dehn function.*

The idea is that we take a word written with generators from $K_{I_0\cup I_m}$, and use $\mathrm{poly}(r)$ commuting relations in order to write the word as $r+1$ words that live only in the respective K^{B_ℓ}, and then to contract each word separately. We prove this lemma in Section 6.

Thus, as a last step, we need to bound the Dehn function of each component Y_ℓ.

Lemma 3.26. *Let Y_ℓ be as above. Then there exists a universal constant A such that $\mathrm{Dehn}_{Y_\ell}(t) \leqslant At$.*

We prove this lemma in Section 7, see Lemma 7.2.

3.5 Proof of the main theorem

We now put all the pieces together to prove Theorem 3.1. We fix the parameters r, d, γ_1, n as in the statement of the theorem and set $T = r^{10}$. Let $q = p^m$ be a prime power and Λ a group without non-trivial elements of order p. Fix an irreducible polynomial $f \in \mathbb{F}_q[t]$ of degree at least 2. Let $X = X(A_d, \mathbb{F}_q, f) = \mathcal{CC}(G, (H_i^f)_{i\in[d]})$ be the corresponding KMS complex. Let $\mathcal{I} \in \binom{[d]}{r+1}^{[T]}$ be a nice set of disjoint colors. Then $F^{\mathcal{I}}X \cong \mathcal{CC}(G, (H_J^f)_{J\in\mathcal{I}})$ (where $H_J^f = \cap_{j\in J} H_j^f$) and for any vertex $gH_J \in F^{\mathcal{I}}X$ its link has the form

$$(F^{\mathcal{I}}X)_{gH_J} \cong \mathcal{CC}(U_{[d]\setminus J}; (U_{[d]\setminus(J\cup I)})_{I\in\mathcal{I}\setminus\{J\}}).$$

We fix one such link and denote it by Y. We use the path order associated with J in Definition 3.6, and let $B_0, \dots, B_r$ be the resulting J-bins. Then Y decomposes as restricted join

$$Y = \prod_{0\leqslant\ell\leqslant r}^{*} \mathcal{CC}(U_{B_\ell}; (U_{B_\ell\setminus I})_{I\in\mathcal{I}\setminus\{J\}}),$$

and we denote by $Y_\ell = \mathcal{CC}(U_{B_\ell}; (U_{B_\ell\setminus I})_{I\in\mathcal{I}\setminus\{J\}})$ the ℓ-th component of this decomposition. By Lemma 3.26, we know that there exists a universal constant such that $\mathrm{Dehn}_{Y_\ell}(t) \leqslant At$. From Lemma 3.25 we deduce that $\mathrm{Dehn}_Y(t) \leqslant \beta_5 t^2 + \beta_6 t + \beta_7 = O((t(r+1))^2)$ with $\beta_5, \beta_6, \beta_7$ at most quadratic in r and independent of all other parameters. By Lemma 3.21, we know that $\mathrm{diam}(Y) \leqslant 6(r+1)+1 = \beta_4$. Thus we have established Theorem 3.20.

Next, we want to apply Theorem 3.19 to Y. By Lemma 3.17, since $|\mathcal{I}| = T > 4$ and assuming $q \geqslant 5$, we know that Y is simply connected. Thus we get

$$\mathrm{Rad}_0(\mathrm{Cone}_Y) \leqslant \beta_4, \qquad \mathrm{Rad}_1(\mathrm{Cone}_Y) \leqslant 15(\beta_5(2\beta_4+1)^2 + \beta_6(2\beta_4+1) + \beta_7 + 2\beta_4 + 1) =: \beta_3.$$

Since β_4 is of order r, $\beta_5, \beta_6, \beta_7$ are of order r^2, we get that β_3 is of order r^4 and independent of all other parameters. Thus we have established Proposition 3.15 with $q_0 = 5$. Using Corollary 3.16, we get

$$h^1((F^{\mathcal{I}}X)_{gH_J}, \Lambda) \geqslant \frac{1}{\binom{T-1}{3}\beta_3} =: \beta_2.$$

Since $T = r^{10}$, $\binom{T-1}{3}$ is of order $T^3 = r^{30}$. In total, β_2 is of order $\frac{1}{r^{34}}$. Together with Lemma 3.17, we have shown Proposition 3.13.

In the following step, we want to combine coboundary expansion of the links with local spectral expansion to deduce cocycle expansion of $F^{\mathcal{I}}X$ via the local-to-global theorem. By Lemma 3.12 we know that for each γ_2 there exists q_0 such that if $q \geqslant q_0$ then $F^{\mathcal{I}}X$ is a γ_2-local spectral expansion. Given β_2 we choose γ_2 small enough such that $\frac{(1-\gamma_2)\beta_2}{24} - e\gamma_2 =: \beta_1$ is positive. This choice of γ_2 then determines a lower bound for q_0. After fixing q_0, β_1 will grow with the same order as β_2, i.e. order r^{-34}. Together with Lemma 3.9 we have shown Proposition 3.8. Finally, we will apply Theorem 3.3 to deduce coboundary expansion of $F^r X$. By our choice of n and T, we have $F^{\mathcal{I}}X = F^{\mathcal{I}}(X^{\leqslant n})$, in particular $F^{\mathcal{I}}(X^{\leqslant n})$ is a β_1-coboundary expander for any nice set of colors $\mathcal{I}$. Our choice of n ensures that we can apply Lemma 3.4 which gives a lower bound on q_0 and shows that $F^r(X^{\leqslant n})$ is a γ_1-local spectral expander. Since

$$\mathbb{P}_{\mathcal{I}\in\binom{[d]}{r+1}^{[T]}}[\mathcal{I} \text{ is nice}] \geqslant 0.9$$

by Lemma 3.7, we have

$$\mathbb{P}_{\mathcal{I}\in\binom{[d]}{r+1}^{[T]}}\left[h^1(F^I(X^{\leqslant n}), \Lambda) \geqslant \beta_1 \wedge H^1(F^{\mathcal{I}}(X^{\leqslant n}), \Lambda) = 1\right] \geqslant 0.9$$

and thus get $H^1(F^r X, \Lambda) = 1$ and $h^1(F^r X, \Lambda) \geqslant \Omega(r^{-34})$ which finishes the proof.

4 The color sampling lemma

In this subsection we prove Lemma 3.7. We start with the following observation. Cutting the cyclic order at one of the $r+1$ colors of a color set leaves a path of d positions with r selected colors. These divide the remaining $d-r$ coordinates into $r+1$ bins, so the average bin size is $\frac{d-r}{r+1} \approx \frac{d+1}{r+1}$. We first show that, with high probability, no bin is larger than this average by more than a logarithmic factor (and to cover the first item, also that no bin is smaller than 101). More precisely, for most collections $\mathcal{I}$, all bins of all $I \in \mathcal{I}$ have size $\lesssim \frac{\ln T}{r+1}(d+1)$. This is the content of (4.1), up to a union bound over the color sets.

Next, fix a color set defining a small bin B, and a forbidden subset $S \subseteq B$ occupying at most four fifths of the bin, up to rounding. We show that a random color set I' avoids S with probability $\gtrsim (Tr)^{-\alpha}$, where $\alpha = \frac{4(r+1)}{5r} < 0.81$. Since we have $T-1$ other color sets to try, the probability that all of them intersect S is exponentially small in $T(Tr)^{-\alpha}$. For $T \geqslant r^{10}$, this is small enough to take a union bound over all the requirements in the definition of niceness.[3] The same argument applies when S is contained in two bins of this same color set, provided its total size satisfies the same bound; fixing that color set determines both bins before we sample the candidates.

Finally, we apply this observation to the forbidden sets appearing in the definition of niceness. For example, to find a color set I' whose intersection with B lies in its first fifth, we take S to be the remaining four fifths of B. Avoiding S is exactly the required condition. The last-fifth and middle-fifth conditions are handled similarly. For the fourth condition, S is the union of two intervals, each occupying approximately two fifths of its respective bin. Since both bins are small, the total size of S is again at most four fifths of our bin-size cutoff, up to rounding. Thus the observation applies to every requirement, and a union bound shows that they all hold simultaneously with high probability.

Proof of Lemma 3.7. Before proving the lemma, we will prove that all sizes of all bins are small with high probability. Then we will show for each bin, and for every pair of bins, that there exists a color set that avoids the respective intervals appearing in the definition of niceness. This will make use of the fact that the intervals we need to avoid occupy only a constant fraction of the relevant bins. Note that different requirements in the definition may have different color sets witnessing them.

We first sample $I_0, \dots, I_{T-1}$ independently and uniformly from $\binom{[d]}{r+1}$. These sets need not be disjoint. Let D be the event that they are pairwise disjoint. We will estimate the probability that the niceness requirements hold in this independent sampling, and only at the end condition on D.

Indeed, the probability that two independently sampled color sets intersect is at most $(r+1)^2/(d+1)$. To see this, fix the first set. Each of its $r+1$ colors belongs to the second set with probability $(r+1)/(d+1)$, and we take a union bound over these colors. Taking another union bound over all pairs of color sets gives

$$\mathbb{P}\left[D^c\right] \leqslant \binom{T}{2}\frac{(r+1)^2}{d+1} \leqslant \frac{1}{200}.$$

Conditioned on D, all ordered collections of pairwise disjoint color sets are equally likely. Forgetting the ordering gives precisely the distribution in the lemma, since each collection has the same number $T!$ of orderings.

Set $A \coloneqq 100T(r+1)$, $m \coloneqq d(1-A^{-1/r})$, and $\alpha \coloneqq \frac{4(r+1)}{5r}$. We claim that

$$\mathop{\mathbb{P}}_{I_0}\left[\exists \text{ a bin } B \text{ of } I_0 \text{ with } |B| > m\right] \leqslant \frac{1}{100T}. \tag{4.1}$$

This probability is small enough so that we could later on union bound on all T sets I_j.

To prove the claim, choose j_* uniformly from I_0 and condition on its value. Conditional on j_*, the other r colors of I_0 form a uniformly random r-subset of $[d] \setminus \{j_*\}$. Rotate the colors so that $j_* = 0$. The remaining selected colors are then a uniformly random r-subset of the path $\{1, \dots, d\}$, whose $r+1$ gaps are exactly the bins of Definition 3.6. Changing j_* only permutes these bins, so the event in (4.1) is independent of the choice of j_*. There is no additional union bound over possible choices of j_*.

[3] The meticulous reader may object and say that the color sets in $\mathcal{I}$ must be disjoint, so they are not independent: why should having many candidates make it so unlikely that they all intersect S? The point is that, under our lower bound on d, independently sampled color sets are already disjoint with high probability. Thus requiring disjointness changes the failure probability only slightly.

Fix one of the $r+1$ bin positions, and let B be its bin. For an integer $0 \leqslant s \leqslant d-r$, under this conditional distribution we have

$$\mathbb{P}\left[|B| \geqslant s\right] = \frac{\binom{d-s}{r}}{\binom{d}{r}}.$$

The denominator counts all possible choices of the r selected positions. For the numerator, take a choice whose specified bin has at least s unused positions, delete the first s of these positions, and shift all subsequent positions down by s. This produces an r-subset of $\{1, \ldots, d-s\}$. Conversely, inserting s unused positions at the start of the specified bin recovers a unique choice with a bin of size at least s. This bijection proves the equality.

Canceling factorials gives

$$\frac{\binom{d-s}{r}}{\binom{d}{r}} = \prod_{j=0}^{r-1}\left(1 - \frac{s}{d-j}\right) \leqslant \left(1 - \frac{s}{d}\right)^r.$$

Taking $s = \lfloor m \rfloor + 1$, we obtain

$$\mathbb{P}\left[|B| > m\right] \leqslant \left(1 - \frac{m}{d}\right)^r = \frac{1}{A}.$$

If $s > d-r$, the probability is zero, since there are only $d-r$ unselected positions. A union bound over the $r+1$ bins gives $(r+1)/A = 1/(100T)$. This holds for every value of j_*, so averaging over j_* proves (4.1). A further union bound over the T color sets shows that the probability that any bin has size greater than m is at most $1/100$.

We also need to rule out bins of size at most 100. Let E be the event that such a bin occurs in one of the sampled color sets. The two selected colors bounding such a bin have clockwise distance at most 101. View each color set as sampled in a uniformly random order. For two fixed, distinct sampling positions, conditioned on the first color, the second is uniform among the d remaining colors, so it lies in the next 101 positions with probability $101/d$. A union bound over the $r(r+1)$ ordered pairs in each of the T color sets gives

$$\mathbb{P}\left[E\right] \leqslant \frac{101Tr(r+1)}{d} \leqslant \frac{2}{T} < 10^{-5},$$

where we used $d+1 \geqslant 100(r+1)^2T^2$, $T \geqslant r^{10}$, and $r \geqslant 10^4$.

The next observation estimates the probability of finding a color set that avoids a prescribed set S. Later, S will consist of the colors that are forbidden to a witness for one particular niceness requirement. The exponent α accounts for the fact that the bin cutoff came from r random cut points, whereas a candidate color set contains $r+1$ colors.

Observation 4.1.

1. Let $S \subseteq [d]$ be a fixed set of size at most $\frac{4}{5}m + 20$. For a uniformly sampled color set I_1,

$$\mathop{\mathbb{P}}_{I_1}\left[I_1 \cap S = \emptyset\right] \geqslant \frac{1}{2A^\alpha}.$$

2. If t color sets are sampled independently of one another and of S, then the probability that none of them avoids S is at most

$$\left(1 - \frac{1}{2A^\alpha}\right)^t \leqslant \exp\left(-\frac{t}{2A^\alpha}\right).$$

To prove the first item, we sample the $r+1$ colors of I_1 one at a time without repetition. Suppose that the first j sampled colors were all outside S. There are then $d+1-j$ colors left to choose from, of which $d+1-|S|-j$ are outside S. Therefore, conditioned on having avoided S so far, the probability that the next color also avoids S is $(d+1-|S|-j)/(d+1-j)$. Multiplying these conditional probabilities for $j = 0, \ldots, r$ gives

$$\mathop{\mathbb{P}}_{I_1}\left[I_1 \cap S = \emptyset\right] = \prod_{j=0}^{r} \frac{d+1-|S|-j}{d+1-j}.$$

This uses conditional probabilities; the individual sampled colors are not independent.

Put $\theta = \frac{1}{5} + \frac{4}{5}A^{-1/r}$. Since $m = d(1 - A^{-1/r}) \leqslant (d+1)(1 - A^{-1/r})$, the bound on $|S|$ gives

$$d + 1 - |S| \geqslant d + 1 - \frac{4}{5}m - 20 \geqslant (d+1)\theta - 20.$$

Since $j \leqslant r$, each numerator in the product above is at least $(d+1)\theta - (r+20)$, while its denominator is at most $d+1$. This lower bound on the numerator is positive under our assumptions. Hence

$$\mathop{\mathbb{P}}_{I_1}[I_1 \cap S = \emptyset] \geqslant \left(\theta - \frac{r+20}{d+1}\right)^{r+1} \geqslant \theta^{r+1}\left(1 - \frac{5(r+20)}{d+1}\right)^{r+1},$$

where the last inequality uses $\theta \geqslant 1/5$.

By the elementary inequality $(1-x)^{r+1} \geqslant 1 - (r+1)x$ for $0 \leqslant x \leqslant 1$, the last factor is at least

$$1 - \frac{5(r+1)(r+20)}{d+1} \geqslant \frac{1}{2}.$$

Here we used $d + 1 \geqslant 100(r+1)^2T^2 \geqslant 10(r+1)(r+20)$. Finally, the weighted arithmetic–geometric mean inequality gives

$$\theta = \frac{1}{5}\cdot 1 + \frac{4}{5}\cdot A^{-1/r} \geqslant \left(A^{-1/r}\right)^{4/5},$$

and therefore $\theta^{r+1} \geqslant A^{-4(r+1)/(5r)} = A^{-\alpha}$. Combining these inequalities proves the first item of the observation. For the second item, each candidate fails to avoid S with probability at most $1 - 1/(2A^\alpha)$. Independence allows us to multiply these failure probabilities, and the last inequality follows from $1 - x \leqslant e^{-x}$.

We now explain how to apply the observation to each of the four items in the definition of niceness. For each item, we must identify a forbidden set S such that $I' \cap S = \emptyset$ guarantees the required property of I'. We must also check that $|S| \leqslant \frac{4}{5}m + 20$ when the relevant bins have size at most m. Only colors in $[d]$ can belong to I', so any forbidden set below is intersected with $[d]$ if necessary; this can only decrease its size. For each fixed color set I, we use its path order from Definition 3.6 for all color sets and intervals below.

item 1 Fix a bin $B = [a, b]$ of I. We want a color set $I' \neq I$ whose colors in B all lie in the first fifth of the bin. Thus we take $S = (a + \lfloor \frac{b-a}{5} \rfloor, b]$. Avoiding S means precisely that $I' \cap B \subseteq [a, a + \lfloor \frac{b-a}{5} \rfloor]$. The number of forbidden colors is

$$|S| = (b-a) - \left\lfloor \frac{b-a}{5} \right\rfloor \leqslant \frac{4}{5}(b - a + 1) = \frac{4}{5}|B|.$$

Here we used $\lfloor \frac{b-a}{5} \rfloor \geqslant \frac{b-a-4}{5}$. In particular, $|S| \leqslant \frac{4}{5}m$ if $|B| \leqslant m$.

item 2 We want the colors of I' in B to lie in its last fifth. Accordingly, we take $S = [a, b - \lfloor \frac{b-a}{5} \rfloor)$. Avoiding S means precisely that $I' \cap B \subseteq [b - \lfloor \frac{b-a}{5} \rfloor, b]$. This forbidden set has the same size as in item 1, so again $|S| \leqslant \frac{4}{5}m$ whenever $|B| \leqslant m$.

item 3 We want the colors of I' in B to lie in the specified middle interval. We therefore take $S = [a, b] \setminus [a + 2\lceil \frac{b-a}{5} \rceil, a + 3\lfloor \frac{b-a}{5} \rfloor]$. Avoiding S is exactly the required condition. The middle interval contains

$$\max\left\{0, 3\left\lfloor \frac{b-a}{5} \right\rfloor - 2\left\lceil \frac{b-a}{5} \right\rceil + 1\right\}$$

colors, where the maximum with zero accounts for the possibility that the interval is empty. Since $\lceil \frac{b-a}{5} \rceil \leqslant \lfloor \frac{b-a}{5} \rfloor + 1$, this number is at least

$$\left\lfloor \frac{b-a}{5} \right\rfloor - 1 \geqslant \frac{|B|}{5} - 2.$$

Hence the number of colors outside the middle interval is at most $\frac{4}{5}|B| + 2$, and $|S| \leqslant \frac{4}{5}m + 2$ when $|B| \leqslant m$.

item 4 Fix two bins $B = [a, b]$, $B' = [a', b']$ of the same color set I and a pair $(\ell, \ell') \in \{0, 1, 2, 3\}^2$. This time we need the same color set I' to avoid both intervals in the fourth item. Thus the forbidden set is their union:

$$\begin{aligned} S = [a + \ell\lfloor \tfrac{b-a}{5} \rfloor, \min\{a + (\ell+2)\lceil \tfrac{b-a}{5} \rceil + 4, b\}] \\ \cup [a' + \ell'\lfloor \tfrac{b'-a'}{5} \rfloor, \min\{a' + (\ell'+2)\lceil \tfrac{b'-a'}{5} \rceil + 4, b'\}]. \end{aligned}$$

Avoiding this union means avoiding both intervals simultaneously, as required. The first interval has at most

$$(\ell+2)\left\lceil \frac{b-a}{5} \right\rceil - \ell\left\lfloor \frac{b-a}{5} \right\rfloor + 5 \leqslant 2\left\lfloor \frac{b-a}{5} \right\rfloor + \ell + 7 \leqslant \frac{2}{5}|B| + 10$$

colors. The $+5$ in the first expression accounts for the $+4$ in the upper endpoint and for counting both endpoints. The first inequality uses $\lceil \frac{b-a}{5} \rceil \leqslant \lfloor \frac{b-a}{5} \rfloor + 1$, and the second uses $\ell \leqslant 3$. The same argument bounds the second interval by $\frac{2}{5}|B'| + 10$. Consequently,

$$|S| \leqslant \frac{2}{5}|B| + \frac{2}{5}|B'| + 20 \leqslant \frac{4}{5}m + 20$$

whenever both bins have size at most m.

We have now checked that every individual niceness requirement can be satisfied by finding a color set that avoids a set S of the size allowed in the observation. We next bound the probability that such a color set does not exist. First fix which requirement we are considering: the excluded color set I, the position of its bin in items 1–3, or the two bin positions and (ℓ, ℓ') in item 4. At this point the bins, and hence S, are still random.

Reveal the one or two color sets defining these bins. For example, if the two bins in item 4 belong to I_j and $I_{j'}$, we reveal these sets. This determines the bins and S. The other color sets are still independently and uniformly sampled. We look for a witness among them, also excluding I, since the definition requires $I' \neq I$. In total we exclude at most three color sets, so at least $T - 3$ independent candidates remain. Whenever the revealed bins have size at most m, the observation shows that the probability that all these candidates fail is at most

$$\exp\left(-\frac{T-3}{2A^{4/5}}\right).$$

This bound holds for every choice of the revealed color sets with the required bin sizes. Averaging over those choices, the probability that the defining bins have size at most m and this particular requirement has no witness is also at most the displayed bound. Note that we have not conditioned on all bins of all color sets being small. We fixed only the sets defining the requirement, so the candidate witnesses retained their original independent distributions.

It remains to make all requirements hold simultaneously. For items 1–3 there are T choices of I, $r + 2$ bin positions, and three items to check. Thus these items give at most $3T(r + 2)$ requirements. For item 4 there are at most T choices of the excluded set I, at most $T(r + 2)$ choices for each of the two bin positions, and 16 choices of (ℓ, ℓ'). This gives at most $16T^3(r + 2)^2$ requirements. This count allows the two bins to belong to arbitrary color sets, as in the definition.

Let F be the event that some niceness requirement fails in the independent sampling. Either there is a bin larger than m, or there is a failed requirement whose defining bins all have size at most m. The first event has probability at most $1/100$, and we have just bounded the probability of the second event for each individual requirement. Taking a union bound gives

$$\mathbb{P}[F] \leqslant \frac{1}{100} + \left(3T(r+2) + 16T^3(r+2)^2\right)\exp\left(-\frac{T-3}{2A^{4/5}}\right).$$

We finish by checking explicitly that this error bound is small for the stated values of r and T. The number of requirements satisfies

$$3T(r+1) + 16T(r+1)^2 \leqslant 20T(r+1)^2 \leqslant 21Tr^2.$$

Also, $T-1 \geqslant T/2$, $r+1 \leqslant 1.01r$, and $\alpha \leqslant 0.81$, so

$$\frac{T-1}{2A^\alpha} \geqslant \frac{T^{1-\alpha}}{4[100(r+1)]^\alpha} \geqslant \frac{T^{1-\alpha}}{200r^\alpha},$$

where we used $4 \cdot 101^{0.81} < 200$. Consequently,

$$\mathbb{P}[F] \leqslant \frac{1}{100} + 21Tr^2 \exp\left(-\frac{T^{1-\alpha}}{200r^\alpha}\right).$$

To bound the second term, first hold r fixed and increase T. Since $1-\alpha \geqslant 0.19$ and

$$\frac{T^{1-\alpha}}{r^\alpha} \geqslant r^{10-11\alpha} = r^{6/5-44/(5r)} \geqslant r^{11/10},$$

the derivative of the logarithm of this term with respect to T is

$$\frac{1}{T}\left(1 - \frac{(1-\alpha)T^{1-\alpha}}{200r^\alpha}\right) < 0 \qquad \text{for } T \geqslant r^{10},\ r \geqslant 10^4.$$

Thus the term is largest at $T = r^{10}$, where it is at most $21r^{12}\exp(-r^{11/10}/200)$. The logarithmic derivative of this upper bound with respect to r is

$$\frac{1}{r}\left(12 - \frac{11r^{11/10}}{2000}\right) < 0 \qquad \text{for } r \geqslant 10^4.$$

It is therefore enough to evaluate at $r = 10^4$. At this value its logarithm is

$$\ln 21 + 48\ln 10 - \frac{10^{22/5}}{200} < -12,$$

so the term is less than 10^{-5}. We conclude that $\mathbb{P}[F] \leqslant 1/100 + 10^{-5}$.

Finally, we return to the distribution in the lemma by conditioning on pairwise disjointness. We proved at the start that $\mathbb{P}[D] \geqslant 1 - 1/200$. Thus

$$\mathop{\mathbb{P}}_{D}[F \cup E] \leqslant \frac{\mathbb{P}[F] + \mathbb{P}[E]}{\mathbb{P}[D]} \leqslant \frac{1/100 + 10^{-5} + 2/T}{1 - 1/200} < 0.011 < 0.1.$$

On D, the collection is nice whenever neither F nor E occurs. Hence a uniformly sampled collection of T pairwise disjoint color sets is nice with probability at least 0.9, as claimed. □

5 Bound on the Dehn function for $\mathrm{Unip}_n(\mathcal{R})$

In Lemma 3.26 in the proof outline above, we bound the Dehn function of coset complex Y_ℓ based on the relations of the vertex stabilizers under the niceness assumption. We note that Y_ℓ is a coset complex of an upper-triangular matrix group and bounding the Dehn function requires studying the Dehn function of such groups under group presentations related to the niceness assumption. Roughly speaking, the niceness assumption is meant to ensure that the vertex stabilizers in Y_ℓ will have "enough" relations to yield a good bound on the Dehn function. In this section, we study presentations of upper-triangular matrix groups and derive a presentation that on one hand has enough relations to ensure a linear bound on the Dehn function and on the other hand its set of relations in "small" enough to be derived from the niceness assumption.

Throughout let $\mathcal{R}$ be a commutative unital ring in which we denote the unit as 1. For our application below, $\mathcal{R}$ will be a finite field, but the results of the section are general and do not depend on the field or finiteness properties.

We denote $\mathrm{Unip}_n(\mathcal{R})$ to be the group of $n \times n$ upper-triangular matrices with entries in $\mathcal{R}$ and 1's along the main diagonal with the binary operation of usual matrix multiplication. Below, we will abbriviate and refer to such matrices as upper-triangular (without explicitly mentioning the 1's on the main diagonal or the $\mathcal{R}$ entries).

The key result of this section is to find a bound for the Dehn function of $\mathrm{Unip}_n(\mathcal{R})$ for specific presentation that will be useful in the sequel.

5.1 Generalities regarding commutators in groups

Here we cover some general definitions and facts regarding commutators in groups that will be used below.

Let G be a group. For two elements $g, h \in G$, the *commutator of g and h* is the group element defined as $[g,h] = g^{-1}h^{-1}gh$. We say that g and h *commute* if $[g,h] = e$ or equivalently $gh = hg$. The following identities are well-known (and straightforward to prove):

Claim 5.1. Let G be a group and $g, h, g_1, g_2 \in G$ be elements. Then

- $[g,h]^{-1} = [h,g]$
- $[g,e] = [e,g] = e$
- If $[g,h] = e$, then $[g^{-1},h] = [g^{-1},h^{-1}] = [g,h^{-1}] = e$
- If $[g_1,h] = [g_2,h] = e$, then $[[g_1,g_2],h] = e$.

Using the notation $g^h = h^{-1}gh$, the following identity is known as the Hall–Witt identity: For every $g_1, g_2, g_3 \in G$ it always hold that

$$[[g_1,g_2^{-1}],g_3]^{g_2}[[g_2,g_3^{-1}],g_1]^{g_3}[[g_3,g_1^{-1}],g_2]^{g_1} = e.$$

In the sequel, we will need the following implication of this identity:

Proposition 5.2. *Let G be a group and g_1, g_2, g_3 be elements of G. If $[g_3, g_1^{-1}] = e$ and $[[g_1, g_2^{-1}], g_3]$ commutes with g_2 and g_3^{-1}, then*

$$[[g_1,g_2^{-1}],g_3] = [g_1,[g_2,g_3^{-1}]].$$

Proof. We note that the assumption that $[g_3, g_1^{-1}] = e$ implies that $[[g_3, g_1^{-1}], g_2]^{g_1} = e$. Also, by the assumption that $[[g_1, g_2^{-1}], g_3]$ commutes with g_2 it follows that $[[g_1, g_2^{-1}], g_3]^{g_2} = [[g_1, g_2^{-1}], g_3]$. Thus, the Hall–Witt identity is reduced to

$$[[g_1,g_2^{-1}],g_3][[g_2,g_3^{-1}],g_1]^{g_3} = e.$$

Using the assumption that $[[g_1, g_2^{-1}], g_3]$ commutes with g_3^{-1} yields that

$$e = [[g_1,g_2^{-1}],g_3][[g_2,g_3^{-1}],g_1]^{g_3} = [[g_1,g_2^{-1}],g_3]g_3^{-1}[[g_2,g_3^{-1}],g_1]g_3 = g_3^{-1}[[g_1,g_2^{-1}],g_3][[g_2,g_3^{-1}],g_1]g_3,$$

and thus

$$[[g_1,g_2^{-1}],g_3][[g_2,g_3^{-1}],g_1] = e$$

or equivalently

$$[[g_1,g_2^{-1}],g_3] = [[g_2,g_3^{-1}],g_1]^{-1} = [g_1,[g_2,g_3^{-1}]],$$

as needed. □

We will also need the following result

Proposition 5.3 ([BD01, Lemma 1])**.** *For a group G and elements $g_1, g_2, h \in G$. Assume that g_2 commutes with $[g_1, h]$. Then $[g_1,h][g_2,h] = [g_1g_2,h]$.*

5.2 Basic presentation and block matrices

For $1 \leqslant i < j \leqslant n$ and $a \in \mathcal{R}$, we define the elementary matrix $e_{i,j}(a) \in \mathrm{Unip}_n(\mathcal{R})$ to be the matrix with 1's along the main diagonal, a in the (i,j) entry and 0's in all other entries. The following fact is well-known (see for instance [BD01, Theorem 2] where a much smaller presentation is shown to suffice):

Theorem 5.4. *The group $\mathrm{Unip}_n(\mathcal{R})$ can be presented in terms of generators and relations as follows: The set $S_{\mathrm{elm}} = \{e_{i,j}(a) : 1 \leqslant i < j \leqslant n, a \in \mathcal{R}\}$ is a generating set of $\mathrm{Unip}_n(\mathcal{R})$ and $\mathrm{Unip}_n(\mathcal{R}) = \langle S_{\mathrm{elm}} | R_{\mathrm{elm}} \rangle$ where R_{elm} are the following relations:*

1. *For every $1 \leqslant i < j \leqslant n$ and every $a_1, a_2 \in \mathcal{R}$, $e_{i,j}(a_1)e_{i,j}(a_2) = e_{i,j}(a_1 + a_2)$.*

2. *For every* $1 \leqslant i, j, k, l \leqslant n$ *with* $i < j$ *and* $k < l$ *and every* $a_1, a_2 \in \mathcal{R}$,

$$[e_{i,j}(a_1), e_{k,l}(a_2)] = \begin{cases} e_{i,k}(a_1 a_2) & j = k \\ e_{k,j}(-a_1 a_2) & i = l \\ e & \textit{otherwise} \end{cases}.$$

The basic presentation above is well suited for describing $\mathrm{Unip}_n(\mathcal{R})$, but it is less convenient for bounding the Dehn function. The goal of this section is to introduce a block presentation that is adapted to a recursive analysis of the Dehn function.

A key idea is to partition matrices into five large blocks and to regard each off-diagonal block as a single generator carrying an entire matrix rather than a single ring element. This yields a presentation whose generators are naturally indexed by pairs (i, j) of block positions. Explicitly, assuming that $n \geqslant 5$, we denote $k_0 = \lfloor \frac{n}{5} \rfloor$ and $k_1 = n - 4k_0$ (note that $n \geqslant 5$ implies that k_0, k_1 are positive integers). A matrix $A \in \mathrm{Unip}_n(\mathcal{R})$ can always be written in the following block form:

$$A = \begin{pmatrix} A_{1,1} & A_{1,2} & A_{1,3} & A_{1,4} & A_{1,5} \\ 0_{k_0} & A_{2,2} & A_{2,3} & A_{2,4} & A_{2,5} \\ 0_{k_1 \times k_0} & 0_{k_1 \times k_0} & A_{3,3} & A_{3,4} & A_{3,5} \\ 0_{k_0} & 0_{k_0} & 0_{k_0 \times k_1} & A_{4,4} & A_{4,5} \\ 0_{k_0} & 0_{k_0} & 0_{k_0 \times k_1} & 0_{k_0} & A_{5,5} \end{pmatrix}$$

where

- $0_{k_i \times k_j}$ denotes $k_i \times k_j$ 0 matrix and 0_{k_i} denotes the $k_i \times k_i$ 0 matrix
- $A_{1,1}, A_{2,2}, A_{4,4}, A_{5,5} \in \mathrm{Unip}_{k_0}(\mathcal{R})$
- $A_{3,3} \in \mathrm{Unip}_{k_1}(\mathcal{R})$
- For $i \neq 3, j \neq 3$, $A_{i,j} \in M_{k_0}(\mathcal{R})$
- $A_{1,3}, A_{2,3} \in M_{k_0 \times k_1}(\mathcal{R})$
- $A_{3,4}, A_{3,5} \in M_{k_1 \times k_0}(\mathcal{R})$

With this block form in mind, we introduce the following notation:

1. For $1 \leqslant i < j \leqslant 5$, and a matrix A (of the suitable dimensions), we denote $E_{i,j}(A)$ to be the block matrix in the form above, where $A_{1,1}, ..., A_{5,5}$ are identity matrices (of the suitable dimensions), $A_{i',j'} = 0$ for every $(i', j') \neq (i, j)$ and $A_{i,j} = A$. E.g.,

$$E_{1,2}(A) = \begin{pmatrix} I & A & 0_{k_0 \times k_1} & 0_{k_0} & 0_{k_0} \\ 0_{k_0} & I & 0_{k_0 \times k_1} & 0_{k_0} & 0_{k_0} \\ 0_{k_1 \times k_0} & 0_{k_1 \times k_0} & I & 0_{k_1 \times k_0} & 0_{k_1 \times k_0} \\ 0_{k_0} & 0_{k_0} & 0_{k_0 \times k_1} & I & 0_{k_0} \\ 0_{k_0} & 0_{k_0} & 0_{k_0 \times k_1} & 0_{k_0} & I \end{pmatrix},$$

   where $A \in M_{k_0}(\mathcal{R})$.

2. For $1 \leqslant i \leqslant 5$, and an upper triangular matrix A (of the suitable dimensions), we denote $E_{i,i}(A)$ to be the block matrix of the form above, where $A_{i,i} = A$, all the other $A_{i',i'}$'s are identity matrices and the $A_{i,j}$'s are zero matrices. E.g.,

$$E_{1,1}(A) = \begin{pmatrix} A & 0_{k_0} & 0_{k_0 \times k_1} & 0_{k_0} & 0_{k_0} \\ 0_{k_0} & I & 0_{k_0 \times k_1} & 0_{k_0} & 0_{k_0} \\ 0_{k_1 \times k_0} & 0_{k_1 \times k_0} & I & 0_{k_1 \times k_0} & 0_{k_1 \times k_0} \\ 0_{k_0} & 0_{k_0} & 0_{k_0 \times k_1} & I & 0_{k_0} \\ 0_{k_0} & 0_{k_0} & 0_{k_0 \times k_1} & 0_{k_0} & I \end{pmatrix},$$

   where $A \in \mathrm{Unip}_{k_0}(\mathcal{R})$.

Above, we used the imprecise phrase "of suitable dimension", in order to make it precise, we define the following notion: A matrix A will be called (i,j)-*admissible*, if $E_{i,j}(A)$ is well-defined. Explicitly, A is (i,j)-admissible is equivalent to:

$$\begin{cases} A \in M_{k_0}(\mathcal{R}) & i<j, i\neq 3, j\neq 3 \\ A \in M_{k_0\times k_1}(\mathcal{R}) & i<3, j=3 \\ A \in M_{k_1\times k_0}(\mathcal{R}) & 3<j, i=3 \\ A \in \mathrm{Unip}_{k_0}(\mathcal{R}) & i=j, i\neq 3 \\ A \in \mathrm{Unip}_{k_1}(\mathcal{R}) & i=j=3 \end{cases}$$

With this notation, we will define a presentation of $\mathrm{Unip}_n(\mathcal{R})$ and prove a bound for the Dehn function of this presentation.

5.3 Main result - a bound on the Dehn function

We define the generating set of our presentation to be

$$S_{\mathrm{Elm}} = \{E_{i,j}(A) : 1 \leqslant i \leqslant j \leqslant 5, (i,j) \neq (1,5), A \text{ is } (i,j)\text{-admissible}\}.$$

We also define the following set of relations R_{Elm} on the alphabet S_{Elm}:

(Elm1) For every $1 \leqslant i \leqslant j \leqslant 5, (i,j) \neq (1,5)$ and all matrices A_1, A_2 that are (i,j)-admissible,

$$E_{i,j}(A_1)E_{i,j}(A_2) = \begin{cases} E_{i,j}(A_1+A_2) & i<j \\ E_{i,j}(A_1A_2) & i=j \end{cases}.$$

(Elm2) For every $1 \leqslant i,j,k,l \leqslant 5$ such that one of the following holds:

- $1 \leqslant i < j \leqslant 4$ and $1 \leqslant k < l \leqslant 4$
- $2 \leqslant i < j \leqslant 5$ and $2 \leqslant k < l \leqslant 5$
- $\{(i,j),(k,l)\} = \{(1,2),(4,5)\}$

and all matrices A_1, A_2 that are (i,j) and (k,l) admissible correspondingly,

$$[E_{i,j}(A_1), E_{k,l}(A_2)] = \begin{cases} E_{i,k}(A_1A_2) & j=k \\ E_{k,j}(-A_2A_1) & i=l \\ e & \text{otherwise} \end{cases}.$$

(Elm3) For every $1 \leqslant i,j,l \leqslant 5, i \leqslant j$ such that one of the following holds:

- $1 \leqslant i \leqslant j \leqslant 4$ and $1 \leqslant l \leqslant 4$
- $2 \leqslant i \leqslant j \leqslant 5$ and $2 \leqslant l \leqslant 5$
- $1 \leqslant i \leqslant j \leqslant 2$ and $l = 5$
- $4 \leqslant i \leqslant j \leqslant 5$ and $l = 1$

and all matrices A, B that are (i,j) and (l,l) admissible correspondingly,

$$(E_{l,l}(B))^{-1}E_{i,j}(A)E_{l,l}(B) = \begin{cases} E_{i,j}(A) & i\neq l \text{ and } j\neq l \\ E_{i,j}(B^{-1}A) & i=l \text{ and } j\neq l \\ E_{i,j}(AB) & j=l \text{ and } i\neq l \\ E_{l,l}(B^{-1}AB) & i=j=l \end{cases}.$$

We note that currently, we have not yet shown that S_{Elm} is a generating set of $\mathrm{Unip}_n(\mathcal{R})$ or that $\langle S_{\mathrm{Elm}} | \mathcal{R}_{\mathrm{Elm}} \rangle$ is a presentation of $\mathrm{Unip}_n(\mathcal{R})$. Our main result in this section is the following theorem:

Theorem 5.5. *The group* $\mathrm{Unip}_n(\mathcal{R})$ *can be presented as* $\langle S_{\mathrm{Elm}}|R_{\mathrm{Elm}}\rangle$*. Moreover, there is a universal constant* C' *(independent on* n *and* $\mathcal{R}$*) such that the Dehn function with respect to this presentation can be bounded as follows:*

$$\mathrm{Dehn}_{S_{\mathrm{Elm}},R_{\mathrm{Elm}}}(m) \leqslant C'm.$$

The strategy of proving this theorem goes as follows: we start with a different "larger" presentation of $\mathrm{Unip}_n(\mathcal{R})$, which we denote by $\langle S'_{\mathrm{Elm}}|R'_{\mathrm{Elm}}\rangle$. For this larger presentation, we bound the Dehn function linearly by a universal constant. We then show that moving to our "smaller" presentation $\langle S_{\mathrm{Elm}}|R_{\mathrm{Elm}}\rangle$ defined above still has almost the same Dehn function - the only thing that changes is the universal constant in the linear bound which increases.

5.4 Bounding the Dehn function for a larger presentation

We define the following set

$$S'_{\mathrm{Elm}} = \{E_{i,j}(A) : 1 \leqslant i \leqslant j \leqslant 5, A \text{ is } (i,j)\text{-admissible}\}$$

By block matrix multiplication, we derive the following relations which we denote by R'_{Elm}:

(Elm1') For every $1 \leqslant i \leqslant j \leqslant 5$ and all matrices A_1, A_2 that are (i,j)-admissible,

$$E_{i,j}(A_1)E_{i,j}(A_2) = \begin{cases} E_{i,j}(A_1+A_2) & i<j \\ E_{i,j}(A_1A_2) & i=j \end{cases}.$$

(Elm2') For every $1 \leqslant i,j,k,l \leqslant 5$ with $i<j$ and $k<l$ and every all matrices A_1, A_2 that are (i,j) and (k,l) admissible correspondingly,

$$[E_{i,j}(A_1), E_{k,l}(A_2)] = \begin{cases} E_{i,k}(A_1A_2) & j=k \\ E_{k,j}(-A_2A_1) & i=l \\ e & \text{otherwise} \end{cases}.$$

(Elm3') For every $1 \leqslant i,j,l \leqslant 5, i \leqslant j$ and all matrices A, B that are (i,j) and (l,l) admissible correspondingly,

$$(E_{l,l}(B))^{-1}E_{i,j}(A)E_{l,l}(B) = \begin{cases} E_{i,j}(A) & i \neq l \text{ and } j \neq l \\ E_{i,j}(B^{-1}A) & i=l \text{ and } j \neq l \\ E_{i,j}(AB) & j=l \text{ and } i \neq l \\ E_{l,l}(B^{-1}AB & i=j=l \end{cases}.$$

Proposition 5.6. *It holds that* $\langle S'_{\mathrm{Elm}}|R'_{\mathrm{Elm}}\rangle$ *is a presentation of* $Unip_n(\mathcal{R})$*.*

Proof. We note that for every $1 \leqslant i < j \leqslant n$ and every $a \in \mathcal{R}$, $e_{i,j}(a) \in S'_{\mathrm{Elm}}$. We also note that all the relations between the $e_{i,j}(a)$'s specified in Theorem 5.4 appear in R'_{Elm}. Thus By Theorem 5.4, $\langle S'_{\mathrm{Elm}}|R'_{\mathrm{Elm}}\rangle$ is a presentation of $Unip_n(\mathcal{R})$. □

As we explained above, we will bound the Dehn function of $Unip_n(\mathcal{R})$ with respect to the presentation $\langle S'_{\mathrm{Elm}}|R'_{\mathrm{Elm}}\rangle$ by a universal linear bound (this bound will be independent of n and $\mathcal{R}$).

Our method of bounding the Dehn function for the presentation $\langle S'_{\mathrm{Elm}}|R'_{\mathrm{Elm}}\rangle$ is to decompose the group into a finite nested sequence of groups (where the last group is the complete group $\mathrm{Unip}_n(\mathcal{R})$) and then bound the Dehn function of each of these subgroups recursively (thus achieving a bound for the Dehn function of $\mathrm{Unip}_n(\mathcal{R})$).

We define the following order on the set $\{(i,j) : 1 \leqslant i \leqslant j \leqslant 5\}$: For (i,j) and (i',j') we define $(i,j) < (i',j')$ as follows:

$$(i,j) < (i',j') \Leftrightarrow \begin{cases} j'-i' > j-i \\ \text{or} \\ j'-i' = j-i \text{ and } i' > i \end{cases}.$$

We denote $(i,j) \leqslant (i',j')$ if either $(i,j) < (i',j')$ or $(i,j) = (i',j')$.

With this order, we define the following subgroups of $\mathrm{Unip}_n(\mathcal{R})$ for $1 \leqslant i_0 \leqslant j_0 \leqslant 5$:

$$G_{\geqslant(i_0,j_0)} = \langle E_{i,j}(A_{i,j}) : (i_0,j_0) \leqslant (i,j), A_{i,j} \text{ is } (i,j)\text{-admissible}\rangle,$$

$$G_{>(i_0,j_0)} = \langle E_{i,j}(A_{i,j}) : (i_0,j_0) < (i,j), A_{i,j} \text{ is } (i,j)\text{-admissible}\rangle,$$

$$G_{(i_0,j_0)} = \langle E_{i_0,j_0}(A_{i_0,j_0}) : A_{i_0,j_0} \text{ is } (i_0,j_0)\text{-admissible}\rangle.$$

For $1 \leqslant i \leqslant j \leqslant 5$, we further define $S_{\geqslant(i,j)} = S'_{\mathrm{Elm}} \cap G_{\geqslant(i,j)}$ and $R_{\geqslant(i,j)}$ to be the subset of the relations in R'_{Elm} that involve only elements in $S_{\geqslant(i,j)}$. We similarly define $S_{>(i,j)}, S_{(i,j)}, R_{>(i,j)}, R_{(i,j)}$.

Observation 5.7. We note is that the (i_0,j_0)-block in $G_{>(i_0,j_0)}$ is trivial, i.e., it is the 0 matrix when $i<j$ and the identity matrix if $i=j$. For example,

$$G_{>(2,3)} = \left\{ \begin{pmatrix} I & 0 & A_{1,3} & A_{1,4} & A_{1,5} \\ 0 & I & 0 & A_{2,4} & A_{2,5} \\ 0 & 0 & I & A_{3,4} & A_{3,5} \\ 0 & 0 & 0 & I & A_{4,5} \\ 0 & 0 & 0 & 0 & I \end{pmatrix} : \text{the non-zero } A_{i,j} \text{ are } (i,j)\text{-admissible} \right\}$$

and the $(2,3)$-block is indeed trivial.

The following notion will play an important role in the sequel. A group G is called a semidirect product of subgroups H and N, denoted $G = H \ltimes N$, if every element $g \in G$ admits a unique decomposition

$$g = hn,$$

where $h \in H$ and $n \in N$, and multiplication of elements is compatible with this decomposition. Equivalently, N is normal in G, $G = HN$ and $H \cap N = \{e\}$.

One may think of a semidirect product as a way of representing every element of G by two pieces of information: an H-component and an N-component. The main advantage of such a decomposition is that it allows one to separate a problem on G into a problem on H and a problem on N. In particular, if

$$g = hn = e,$$

where $h \in H$ and $n \in N$, then necessarily $h = e$ and $n = e$.

The next proposition shows that every group $G_{\geqslant(i_0,j_0)}$ admits such a decomposition, where $G_{(i_0,j_0)}$ records the (i_0,j_0)-block and $G_{>(i_0,j_0)}$ records all higher-order blocks.

Proposition 5.8. *For every $1 \leqslant i_0 \leqslant j_0 \leqslant 5$, it holds that $G_{\geqslant(i_0,j_0)} = G_{(i_0,j_0)} \ltimes G_{>(i_0,j_0)}$. In particular, for every $g \in G_{\geqslant(i_0,j_0)}$ there are unique A_{i_0,j_0} and $g' \in G_{>(i_0,j_0)}$ such that $g = E_{i_0,j_0}(A_{i_0,j_0})g'$.*

Proof. We note that $G_{(i_0,j_0)}$ normalizes $G_{>(i_0,j_0)}$ either by (Elm2'), if $i_0 < j_0$ or by (Elm3') if $i_0 = j_0$.

By the definition of $G_{\geqslant(i_0,j_0)}$ this group is generated by $G_{>(i_0,j_0)}$ and $G_{(i_0,j_0)}$. It follows that $G_{>(i_0,j_0)}$ is normal in $G_{\geqslant(i_0,j_0)}$ and that $G_{\geqslant(i_0,j_0)} = G_{(i_0,j_0)}G_{>(i_0,j_0)}$.

Last, by Observation 5.7, the groups $G_{(i_0,j_0)}$ and $G_{>(i_0,j_0)}$ intersect trivially, and thus it follows that $G_{\geqslant(i_0,j_0)} = G_{(i_0,j_0)} \ltimes G_{>(i_0,j_0)}$. □

Corollary 5.9. *Every element $g \in \mathrm{Unip}_n(\mathcal{R})$ can be written as a product*

$$g = \prod_{1 \leqslant i \leqslant j \leqslant 5} E_{i,j}(A_{i,j}),$$

where the order of the elements in the product is according to the order $<$ on $\{(i,j) : 1 \leqslant i \leqslant j \leqslant 5\}$ defined above.

Proof. We note that $\mathrm{Unip}_n(\mathcal{R}) = G_{\geqslant(1,1)}$. Fix $g \in G_{\geqslant(1,1)}$. By Proposition 5.8, there are $g' \in G_{>(1,1)} = G_{\geqslant(2,2)}$ and $A_{1,1}$ such that $g = E_{1,1}(A_{1,1})g'$. Continuing in the same manner, using Proposition 5.8 gives the needed result. □

Proposition 5.10. *Let $1 \leqslant i \leqslant j \leqslant 5, (i,j) \neq (1,5)$, $g \in G_{>(i,j)}$ and $s \in S_{(i,j)}$. The word $s^{-1}gs$ can be transformed to a word g' in the alphabet $S_{>(i,j)}$ by applying at most $|g|_{S_{>(i,j)}}$ relations in $R_{\geqslant(i,j)}$ such that*

$$|g'|_{S_{>(i,j)}} \leqslant \begin{cases} |g|_{S_{>(i,j)}} & i = j \text{ or } (i,j) \geqslant (2,4) \\ 2|g|_{S_{>(i,j)}} & i < j \text{ and } (i,j) \leqslant (1,3) \end{cases}.$$

Proof. We denote $s = E_{i,j}(A)$ and $g = s_1...s_k$ where $s_1, ..., s_k \in S_{>(i,j)}$ and $|g|_{S_{>(i,j)}} = k$. By free expansion moves, we have that

$$s^{-1}gs = s^{-1}(s_1s_2...s_k)s = (s^{-1}s_1s)(s^{-1}s_2s)...(s^{-1}s_ks).$$

Assume first that $(i,j) \geqslant (2,4)$. By (Elm2'), for every $s' = E_{i',j'}(A') \in S_{>(i,j)}$, it holds that $[E_{i',j'}(A'), E_{i,j}(A)] = e$ or equivalently,

$$s^{-1}s's = E_{i,j}(-A)E_{i',j'}(A')E_{i,j}(A) = E_{i',j'}(A')[E_{i',j'}(A'), E_{i,j}(A)] = E_{i',j'}(A').$$

Thus, by applying one relation for every $s^{-1}s_is$, we deduce that $s^{-1}gs = g$ and in particular $|s^{-1}gs|_{S_{>(i,j)}} = |g|_{S_{>(i,j)}}$.

Next, we assume that $i = j$. By (Elm3'), for every $s' = E_{i',j'}(A') \in S_{>(i,j)}$ it holds that $s^{-1}s's \in S_{>(i,j)}$. Thus, by applying one relation for every $s^{-1}s_is$, we get a word in the alphabet $S_{>(i,j)}$ of length $|g|_{S_{>(i,j)}}$.

Last, we assume that $i < j$ and $(i,j) \leqslant (1,3)$. In that case, for $s' = E_{i',j'}(A') \in S_{>(i,j)}$ there are three possible cases. If $i' \neq j$, then by (Elm2'), $s^{-1}E_{i',j'}(A')s = E_{i',j'}(A')$. If $i' = j$, then by (Elm2'), $s^{-1}E_{i',j'}(A')s = E_{i',j'}(A')E_{i,j'}(-AA')$. If $j' = i$, then by (Elm2'), $s^{-1}E_{i',j'}(A')s = E_{i',j'}(A')E_{i',j}(A'A)$. Thus, by applying one relation for every $s^{-1}s_is$, we get a word in the alphabet $S_{>(i,j)}$ with a length that is $\leqslant 2|g|_{S_{>(i,j)}}$. ◻

We define $\mathrm{Dehn}^{G_{\geqslant(i,j)}}$ to be the Dehn function of $G_{\geqslant(i,j)}$ with respect to the presentation $\langle S_{\geqslant(i,j)} | R_{\geqslant(i,j)} \rangle$ and similarly define $\mathrm{Dehn}^{G_{>(i,j)}}$, $\mathrm{Dehn}^{G_{(i,j)}}$.

Proposition 5.11. *We have the following recursive bounds on the Dehn functions:*

- $\mathrm{Dehn}^{G_{(1,5)}}(m) \leqslant m$.
- *For every $1 \leqslant i \leqslant j \leqslant 5$, $(i,j) < (1,5)$,*

$$\mathrm{Dehn}^{G_{\geqslant(i,j)}}(m) \leqslant \begin{cases} 2m + \mathrm{Dehn}^{G_{>(i,j)}}(m) & i = j \text{ or } (i,j) \geqslant (2,4) \\ 2m + \mathrm{Dehn}^{G_{>(i,j)}}(2m) & i < j, (i,j) \leqslant (1,3) \end{cases}.$$

Proof. We start by proving a generalized version of the first bound: Let $g \in G_{(i,j)}$, then by applying at most $|g|_{S_{(i,j)}} - 1$ relations in $R_{(i,j)}$, g can be reduced to an element in $S_{(i,j)}$. Indeed, if $|g|_{S_{(i,j)}} = l$, then $g = E_{i,j}(A_1)...E_{(i,j)}(A_l)$. By applying (Elm1') $l-1$ times, g can be reduced to

$$\begin{cases} E_{i,j}(A_1 + ... + A_l) \in S_{(i,j)} & i < j \\ E_{i,j}(A_1...A_l) \in S_{(i,j)} & i = j \end{cases}.$$

As a special case, when $(i,j) = (1,5)$ and w is a relation in $G_{(i,j)}$ with $|w|_{S_{(1,5)}} = m$, it can be reduced by applying $m-1$ relations to an element in $S_{(1,5)}$, i.e., to e. Thus, $\mathrm{Dehn}^{G_{(1,5)}}(m) \leqslant m$.

Assume that $(i,j) \neq (1,5)$ and let w be a relation in $G_{\geqslant(i,j)}$ with $|w|_{S_{\geqslant(i,j)}} = m$. Such w can be written as a product

$$w = g_1h_1...g_lh_lg_{l+1}$$

where $g_1, ..., g_{l+1} \in G_{>(i,j)}$ (where g_1, g_{l+1} may be trivial) and $h_1, ..., h_l \in G_{(i,j)}$ and

$$|g_1|_{S_{>(i,j)}} + |h_1|_{S_{(i,j)}} + ... + |g_{l+1}|_{S_{>(i,j)}} = m.$$

By the previous bound, every h_i can be reduced to $s_i \in S_{(i,j)}$ by applying at most $|h_i|_{S_{(i,j)}} - 1$ relations. Thus, by applying $\leqslant |h_1|_{S_{(i,j)}} + ... + |h_l|_{S_{(i,j)}}$ relations, w can be reduced to

$$w = g_1 s_1 ... g_l s_l g_{l+1}.$$

By a free expansion, we get that $w = g_1 s_1 ... s_{l-1} s_l (s_l^{-1} g_l s_l) g_{l+1}$ and by applying a relation of type (Elm1'), we can reduce $s_{l-1}s_l$ to $s'_{l-1} \in S_{(i,j)}$. Thus, we get that $g_1 s_1 ... g_{l-1} s'_{l-1} (s_l^{-1} g_l s_l) g_{l+1}$. Continuing in the same manner (applying another $l-1$ relations of type (Elm1')), we reduce the word to $s'_1((s'_1)^{-1} g_1 s'_1)((s'_2)^{-1} g_2 s'_2)...((s'_{l-1})^{-1} g_{l-1} s'_{l-1})(s_l^{-1} g_l s_l) g_{l+1}$. By Proposition 5.10, for each $(s'_i)^{-1} g_i s'_i$ we can apply $|g_i|_{S_{>(i,j)}}$ relations to change it into $g'_i \in G_{>(i,j)}$. To conclude, we applied at most

$$|h_1|_{S_{(i,j)}} + ... + |h_l|_{S_{(i,j)}} + l + |g_1|_{S_{>(i,j)}} + ... + |g_{l+1}|_{S_{>(i,j)}} \leqslant 2m$$

relations and changed the relation w into the relation $s'_1(g'_1 ... g'_{l+1})$. We note that $g'_1 ... g'_{l+1} \in G_{>(i,j)}$ and since $s'_1(g'_1 ... g'_{l+1})$ is a relation, it follows from Proposition 5.8, that $s'_1 = e$ and that $w' = g'_1 ... g'_{l+1}$ is a relation in $G_{>(i,j)}$. Therefore we deduce that

$$\mathrm{Area}_{R_{\geqslant(i,j)}}(w) \leqslant 2m + \mathrm{Area}_{R_{>(i,j)}}(w').$$

By Proposition 5.10,

$$|w'|_{S_{>(i,j)}} = |g'_1 ... g'_{l+1}|_{S_{>(i,j)}} \leqslant$$
$$\begin{cases} \sum_{i=1}^{l+1} |g_i|_{S_{>(i,j)}} & i = j \text{ or } (i,j) \geqslant (2,4) \\ \sum_{i=1}^{l+1} 2|g_i|_{S_{>(i,j)}} & i < j \text{ and } (i,j) \leqslant (1,3) \end{cases} \leqslant$$
$$\begin{cases} m & i = j \text{ or } (i,j) \geqslant (2,4) \\ 2m & i < j \text{ and } (i,j) \leqslant (1,3) \end{cases}.$$

The above bounds were for any arbitrary relation w of length $\leqslant m$ and therefore

$$\mathrm{Dehn}^{G_{\geqslant(i,j)}}(m) \leqslant \begin{cases} 2m + \mathrm{Dehn}^{G_{>(i,j)}}(m) & i = j \text{ or } (i,j) \geqslant (2,4) \\ 2m + \mathrm{Dehn}^{G_{>(i,j)}}(2m) & i < j, (i,j) \leqslant (1,3) \end{cases},$$

as needed. □

Corollary 5.12. *The Dehn function of* $\mathrm{Unip}_n(\mathcal{R})$ *with respect to the presentation* $\langle S'_{\mathrm{Elm}} | R'_{\mathrm{Elm}} \rangle$ *can be bounded as follows:*

$$\mathrm{Dehn}_{S'_{\mathrm{Elm}}, R'_{\mathrm{Elm}}}(m) \leqslant 360m.$$

Proof. We note that by definition $G_{\geqslant(1,1)} = \mathrm{Unip}_n(\mathcal{R})$.

We spell out the order on the (i,j): $(1,1) < (2,2) < (3,3) < (4,4) < (5,5) < (1,2) < (2,3) < (3,4) < (4,5) < (1,3) < (2,4) < (3,5) < (1,4) < (2,5) < (1,5)$. Also, observe that for two consecutive $(i,j) < (i',j')$ in the list, we have $G_{>(i,j)} = G_{\geqslant(i',j')}$. Thus, iteratively using Proposition 5.11, we get

$$\begin{aligned}
&\mathrm{Dehn}_{S'_{\mathrm{Elm}}, R'_{\mathrm{Elm}}}(m) = \mathrm{Dehn}^{G_{\geqslant(1,1)}}(m) \\
&\leqslant 2m + \mathrm{Dehn}^{G_{\geqslant(2,2)}}(m) \leqslant \ldots \leqslant 8m + \mathrm{Dehn}^{G_{\geqslant(5,5)}}(m) \leqslant 10m + \mathrm{Dehn}^{G_{\geqslant(1,2)}}(m) \\
&\leqslant 10m + 2m + \mathrm{Dehn}^{G_{\geqslant(2,3)}}(2m) \leqslant 12m + 4m + \mathrm{Dehn}^{G_{\geqslant(3,4)}}(4m) \leqslant 16m + 8m + \mathrm{Dehn}^{G_{\geqslant(4,5)}}(8m) \\
&\leqslant 24m + 16m + \mathrm{Dehn}^{G_{\geqslant(1,3)}}(16m) \leqslant 40m + 32m + \mathrm{Dehn}^{G_{\geqslant(2,4)}}(32m) \leqslant 72m + 64m + \mathrm{Dehn}^{G_{\geqslant(3,5)}}(32m) \\
&\leqslant 136m + 64m + \mathrm{Dehn}^{G_{\geqslant(1,4)}}(32m) \leqslant 200m + 64m + \mathrm{Dehn}^{G_{\geqslant(2,5)}}(32m) \\
&\leqslant 264m + 64m + \mathrm{Dehn}^{G_{\geqslant(1,5)}}(32m) \leqslant 328m + 32m = 360m.
\end{aligned}$$

□

5.5 Proof of our main Dehn bound

Corollary 5.12 provides a bound for the Dehn function of $\mathrm{Unip}_n(\mathcal{R})$ with respect to the presentation $\langle S'_{\mathrm{Elm}}|R'_{\mathrm{Elm}}\rangle$. As explained above, we will now show how this bound yields a bound for the Dehn function for the "smaller" presentation $\langle S_{\mathrm{Elm}}|R_{\mathrm{Elm}}\rangle$ defined in 5.3 above.

We start by showing that S_{Elm} defined in 5.3 is a generating set:

Claim 5.13. The set S_{Elm} generates $\mathrm{Unip}_n(\mathcal{R})$. Moreover, every element $g \in \mathrm{Unip}_n(\mathcal{R})$ can be written as a product

$$g = \left(\prod_{1\leqslant i\leqslant j\leqslant 5,(i,j)\neq(1,5)} E_{i,j}(A_{i,j}) \right) [E_{1,4}(A_{1,5}), E_{4,5}(I)],$$

where the order of the elements in the product is according to the order $<$ on $\{(i,j) : 1 \leqslant i \leqslant j \leqslant 5\}$ defined above and I is the $k_0 \times k_0$ identity matrix.

Proof. We note that for every $(1,5)$-admissible $A_{1,5}$, it holds that $[E_{1,4}(A_{1,5}), E_{4,5}(I)] = E_{1,5}(A_{1,5})$ and use Corollary 5.9. □

After that we can prove our main theorem in the section which gives a linear bound on the Dehn function of $\mathrm{Unip}_n(\mathcal{R})$ with respect to the presentation $\langle S_{\mathrm{Elm}}|R_{\mathrm{Elm}}\rangle$:

Proof of Theorem 5.5. Whenever we write $E_{i,j}(A)$, we mean that A is (i,j)-admissible (without mentioning it explicitly). We also note that for $1 \leqslant i < j \leqslant 5$, if $i \neq 3, j \neq 3$, an (i,j)-admissible matrix is a square matrix and when denoting $E_{i,j}(I)$, we mean that I is the identity matrix.

By Corollary 5.12, it is enough to show that there is a universal constant C'' such that $\mathrm{Dehn}_{S_{\mathrm{Elm}},R_{\mathrm{Elm}}}(m) \leqslant C'' \mathrm{Dehn}_{S'_{\mathrm{Elm}},R'_{\mathrm{Elm}}}(m)$. In order to do so, we will use Proposition 2.24. Since $S_{\mathrm{Elm}} \subseteq S_{\mathrm{Elm}'}$ it follows that $\mathrm{Free}(S_{\mathrm{Elm}}) \subseteq \mathrm{Free}(S_{\mathrm{Elm}'})$ and we define $\phi : \mathrm{Free}(S_{\mathrm{Elm}}) \to \mathrm{Free}(S_{\mathrm{Elm}'})$ via this inclusion. Explicitly, for every $1 \leqslant i \leqslant j \leqslant 5$, $(i,j) \neq (1,5)$ and every A that is (i,j)-admissible, we define $\phi(E_{i,j}(A)) = E_{i,j}(A)$. The map in the other direction $\phi' : \mathrm{Free}(S'_{\mathrm{Elm}}) \to \mathrm{Free}(S_{\mathrm{Elm}})$ is induced from the map

$$\phi'(E_{i,j}(A)) = \begin{cases} E_{i,j}(A) & (i,j) \neq (1,5) \\ [E_{1,4}(A), E_{4,5}(I)] & (i,j) = (1,5) \end{cases}.$$

We will bound the constants of Proposition 2.24 given by this map. Two bounds are very easy: first

$$C_1 = \max_{s\in S_{\mathrm{Elm}}} |\phi(s)|_{S_{\mathrm{Elm}'}} = 1$$

and second

$$C_3 = \max_{s\in S_{\mathrm{Elm}}} \mathrm{Area}_{R_{\mathrm{Elm}}}(s(\phi'(\phi(s)))^{-1}) = \max_{s\in S_{\mathrm{Elm}}} \mathrm{Area}_{R_{\mathrm{Elm}}}(ss^{-1}) = 0.$$

Thus by Proposition 2.24 it is enough to bound $C_2 = \max_{\rho'\in R_{\mathrm{Elm}'}} \mathrm{Area}_{R_{\mathrm{Elm}}}(\phi'(\rho'))$, and if we show that C_2 can be bounded universally, then for $C'' = C_2$ it will hold that $\mathrm{Dehn}_{S_{\mathrm{Elm}},R_{\mathrm{Elm}}}(m) \leqslant C'' \mathrm{Dehn}_{S'_{\mathrm{Elm}},R'_{\mathrm{Elm}}}(m) \leqslant C''360m$.

Next, we will go systematically over all the relations $\rho' \in R_{\mathrm{Elm}'}$ and show that for every such ρ' the $\mathrm{Area}_{R_{\mathrm{Elm}}}(\phi'(\rho'))$ can be bounded independently of $n, \mathcal{R}$. We recall that $\mathrm{Area}_{R_{\mathrm{Elm}}}(\phi'(\rho'))$ is the number of relations in R_{Elm} one needs to to make $\phi'(\rho')$ into the trivial word. This will also show that $\langle S_{\mathrm{Elm}}|R_{\mathrm{Elm}}\rangle$ is indeed a presentation of $\mathrm{Unip}_n(\mathcal{R})$.

We will not keep track of the actual area book-keeping, but only show that every relation $\phi'(\rho')$ can be deduced from the relations in R_{Elm} (and that the number of relations in R_{Elm} needed is independent of $n, \mathcal{R}$).

We note that every $\rho' \in R_{\mathrm{Elm}} \subseteq R_{\mathrm{Elm}'}$ is trivially deduced from R_{Elm}.

Step 1: We will show how to deduce the relations in $R_{\mathrm{Elm}'}$ that involve only $E_{i,j}(A)$'s in which $i < j$ and $(i,j) \neq (1,5)$ (and are not included in R_{Elm}).

***.** $[E_{1,2}(A), E_{3,5}(B)] = e$: By (Elm2), $E_{3,5}(B) = [E_{3,4}(B), E_{4,5}(I)]$ and also by (Elm2), $[E_{1,2}(A), E_{3,4}(B)] = [E_{1,2}(A), E_{4,5}(I)] = e$. It follows from Claim 5.1 that $[E_{1,2}(A), E_{3,5}(B)] = e$.

*. $[E_{1,3}(A), E_{4,5}(B)] = e$: By (Elm2), $E_{1,3}(A) = [E_{1,2}(I), E_{2,3}(A)]$ and also by (Elm2), $[E_{1,2}(I), E_{4,5}(B)] = [E_{2,3}(A), E_{4,5}(B)] = e$. It follows from Claim 5.1 that $[E_{1,3}(A), E_{4,5}(B)] = e$.

*. $[E_{1,3}(A), E_{2,5}(B)] = e$: By (Elm2), $E_{2,5}(B) = [E_{2,4}(I), E_{4,5}(B)]$ and also by (Elm2), $[E_{1,3}(A), E_{2,4}(I)] = e$. We have shown above that $[E_{1,3}(A), E_{4,5}(B)] = e$ and thus it follows from Claim 5.1 that $[E_{1,3}(A), E_{2,5}(B)] = e$.

*. $[E_{1,4}(A), E_{3,5}(B)] = e$: By (Elm2), $E_{1,4}(A) = [E_{1,2}(I), E_{2,4}(A)]$ and also by (Elm2), $[E_{2,4}(A), E_{3,5}(B)] = e$. We have shown above that $[E_{1,2}(I), E_{3,5}(B)] = e$ and thus it follows from Claim 5.1 that $[E_{1,4}(A), E_{3,5}(B)] = e$.

*. $[E_{1,4}(A), E_{2,5}(B)] = e$: We define B_1 to be the $k_0 \times k_1$ matrix given in block form $B_1 = \begin{pmatrix} B & 0_{k_0 \times (k_1 - k_0)} \end{pmatrix}$. We also define B_2 to be the $k_1 \times k_0$ matrix given in block form $B_2 = \begin{pmatrix} I_{k_0 \times k_0} \\ 0_{(k_1 - k_0) \times k_0} \end{pmatrix}$. By (Elm2), $E_{2,5}(B) = [E_{2,3}(B_1), E_{3,5}(B_2)]$ and also by (Elm2), $[E_{1,4}(A), E_{2,3}(B_1)] = e$. We have shown above that $[E_{1,4}(A), E_{3,5}(B_2)] = e$ and thus it follows from Claim 5.1 that $[E_{1,4}(A), E_{2,5}(B)] = e$.

Step 2: We will show how to deduce the relations in $R_{\mathrm{Elm}'}$ of the form $\phi'([E_{1,5}(A), E_{i,j}(B)])$ for every $1 \leqslant i < j \leqslant 5, (i,j) \neq (1,5), (4,5)$ and every A, B, i.e., by our definition of ϕ', we need to show how to deduce the relations of the form $[[E_{1,4}(A), E_{4,5}(I)], E_{i,j}(B)] = e$ for every $1 \leqslant i < j \leqslant 5, (i,j) \neq (1,5), (4,5)$ and every A, B. We will actually show more: we will show how to deduce the relations of the form $[[E_{1,4}(A_1), E_{4,5}(A_2)], E_{i,j}(B)] = e$ for every $1 \leqslant i < j \leqslant 5, (i,j) \neq (1,5), (4,5)$ and every A_1, A_2, B.

*. $[[E_{1,4}(A_1), E_{4,5}(A_2)], E_{i,j}(B)] = e$ for $(i,j) \in \{(1,2), (1,3), (2,3), (2,5), (3,5)$: For every $(i,j) \in \{(1,2), (1,3), (2,3), (2,5), (3,5)$, we have either by (Elm2) or by the previous step that $[E_{1,4}(A_1), E_{i,j}(B)] = [E_{4,5}(A_2), E_{i,j}(B)] = e$ and thus it follows from Claim 5.1 that $[[E_{1,4}(A_1), E_{4,5}(A_2)], E_{i,j}(B)] = e$.

*. $[[E_{1,4}(A_1), E_{4,5}(A_2)], E_{3,4}(B)] = e$: By Hall-Witt identity:

$$\begin{aligned} &[[E_{1,4}(A_1), E_{4,5}(-A_2)], E_{3,4}(B)]^{E_{4,5}(A_2)} \cdot \\ &\cdot [[E_{4,5}(A_2), E_{3,4}(-B)], E_{1,4}(A_1)]^{E_{3,4}(B)} \cdot \\ &\cdot [[E_{3,4}(B), E_{1,4}(-A_1)], E_{4,5}(A_2)]^{E_{1,4}(A_1)} = e. \end{aligned}$$

By (Elm2), $[E_{3,4}(B), E_{1,4}(-A_1)] = e$ and thus $[[E_{3,4}(B), E_{1,4}(-A_1)], E_{4,5}(A_2)]^{E_{1,4}(A_1)} = e$. By (Elm2), $[E_{4,5}(A_2), E_{3,4}(-B)] = E_{3,5}(BA_2)$. By the previous step, $[E_{3,5}(BA_2), E_{1,4}(A_1)] = e$ and thus

$$[[E_{4,5}(A_2), E_{3,4}(-B)], E_{1,4}(A_1)]^{E_{3,4}(B)} = [E_{3,5}(A_2B), E_{1,4}(A_1)]^{E_{3,4}(B)} = e.$$

It follows from the Hall-Witt identity that $[[E_{1,4}(A_1), E_{4,5}(-A_2)], E_{3,4}(B)]^{E_{4,5}(A_2)} = e$ which implies that $[[E_{1,4}(A_1), E_{4,5}(-A_2)], E_{3,4}(B)] = e$. Since this is true for every A_1, A_2, B, we can replace A_2 with $-A_2$ and get $[[E_{1,4}(A_1), E_{4,5}(A_2)], E_{3,4}(B)] = e$ as needed.

*. $[[E_{1,4}(A_1), E_{4,5}(A_2)], E_{i,4}(B)] = e$ for $i = 1, 2$: We define B_1 to be the $k_0 \times k_1$ matrix given in block form $B_1 = \begin{pmatrix} B & 0_{k_0 \times (k_1 - k_0)} \end{pmatrix}$. We also define B_2 to be the $k_1 \times k_0$ matrix given in block form $B_2 = \begin{pmatrix} I_{k_0 \times k_0} \\ 0_{(k_1 - k_0) \times k_0} \end{pmatrix}$. By (Elm2), $[E_{i,3}(B_1), E_{3,4}(B_2)] = E_{i,4}(B)$. We showed above that $[[E_{1,4}(A_1), E_{4,5}(A_2)], E_{i,3}(B_1)] = [[E_{1,4}(A_1), E_{4,5}(A_2)], E_{3,4}(B_2)] = e$ and thus by Claim 5.1 it follows that $[[E_{1,4}(A_1), E_{4,5}(A_2)], E_{i,4}(B)] = e$.

Step 3: We will show how to deduce the relations in $R_{\mathrm{Elm}'}$ of the form $\phi'([E_{1,j}(A), E_{j,5}(B)](E_{1,5}(AB))^{-1})$ for every $2 \leqslant j \leqslant 4$ and every A, B. By the definition of ϕ', we need to show how to deduce $[E_{1,4}(AB), E_{4,5}(I)] = [E_{1,j}(A), E_{j,5}(B)]$ for every $2 \leqslant j \leqslant 4$ and every A, B.

*. $[E_{1,4}(A_1A_2), E_{4,5}(A_3)] = [E_{1,j}(A_1), E_{j,5}(A_2A_3)]$ for every $j = 2, 3$ and every A_1, A_2, A_3: We denote $g_1 = E_{1,j}(A_1), g_2 = E_{j,4}(-A_2)$ and $g_3 = E_{4,5}(A_3)$ and use the Hall-Witt identity:

$$[[g_1, g_2^{-1}], g_3]^{g_2}[[g_2, g_3^{-1}], g_1]^{g_3}[[g_3, g_1^{-1}], g_2]^{g_1} = e.$$

By (Elm2) and our choice of g_1, g_2, g_3 it holds that: $[g_1, g_2^{-1}] = E_{1,4}(A_1A_2)$, $[g_2, g_3^{-1}] = E_{j,5}(A_2A_3)$ and $[g_3, g_1^{-1}] = e$. Thus, the Hall-Witt identity gives:

$$[E_{1,4}(A_1A_2), E_{4,5}(A_3)]^{E_{j,4}(-A_2)}[E_{j,5}(A_2A_3), E_{1,j}(A_1)]^{E_{4,5}(A_3)} = e.$$

We showed above that $[[E_{1,4}(A_1A_2)], E_{4,5}(A_3)]$ commutes with $E_{j,4}(-A_2)$ and thus $[[E_{1,4}(A_1A_2)], E_{4,5}(A_3)]^{E_{j,4}(-A_2)} = [[E_{1,4}(A_1A_2)], E_{4,5}(A_3)]$. Also, by (Elm2), $[E_{j,5}(A_2A_3), E_{4,5}(A_3)] = [E_{1,j}(A_1), E_{4,5}(A_3)] = e$ and thus $[E_{j,5}(A_2A_3), E_{1,j}(A_1)]^{E_{4,5}(A_3)} = [E_{j,5}(A_2A_3), E_{1,j}(A_1)]$. Therefore we get

$$[E_{1,4}(A_1A_2), E_{4,5}(A_3)][E_{j,5}(A_2A_3), E_{1,j}(A_1)] = e$$

or equivalently, $[E_{1,4}(A_1A_2), E_{4,5}(A_3)] = [E_{1,j}(A_1), E_{j,5}(A_2A_3)]$.

*. $[E_{1,4}(AB), E_{4,5}(I)] = [E_{1,j}(A), E_{j,5}(B)]$ for every $j = 2, 3$ and every A, B: We showed that $[E_{1,4}(A_1A_2), E_{4,5}(A_3)] = [E_{1,j}(A_1), E_{j,5}(A_2A_3)]$. Take $A_1 = A, A_2 = B, A_3 = I$.

*. $[E_{1,4}(AB), E_{4,5}(I)] = [E_{1,4}(A), E_{4,5}(B)]$ for every A, B: We showed that $[E_{1,4}(A_1A_2), E_{4,5}(A_3)] = [E_{1,2}(A_1), E_{2,5}(A_2A_3)]$. For $A_1 = I, A_2 = A, A_3 = B$, then $[E_{1,4}(A), E_{4,5}(B)] = [E_{1,2}(I), E_{2,5}(AB)]$. In order to conclude the proof, we use the fact that we already shown that $[E_{1,2}(I), E_{2,5}(AB)] = [E_{1,4}(AB), E_{4,5}(I)]$.

Step 4: We will show how to deduce the relations in $R_{\mathrm{Elm}'}$ that involve only $E_{i,j}$'s with $i < j$. We have two such relations: First, $\phi'([E_{1,5}(A), E_{4,5}(B)])$ or equivalently by the definition of ϕ', $[[E_{1,4}(A), E_{4,5}(I)], E_{4,5}(B)] = e$. Second, $\phi'((E_{1,5}(A)E_{1,5}(B))(E_{1,5}(A+B))^{-1})$ or equivalently by the definition of ϕ', $[E_{1,4}(A), E_{4,5}(I)][E_{1,4}(B), E_{4,5}(I)] = [E_{1,4}(A+B), E_{4,5}(I)]$.

*. $[[E_{1,4}(A), E_{4,5}(I)], E_{4,5}(B)] = e$: We showed above that $[E_{1,4}(A), E_{4,5}(I)] = [E_{1,2}(A), E_{2,5}(I)]$. By (Elm2), $[E_{1,2}(A), E_{4,5}(B)] = [E_{2,5}(I), E_{4,5}(B)] = e$. It follows from Claim 5.1 that $[[E_{1,4}(A), E_{4,5}(I)], E_{4,5}(B)] = e$.

*. $[E_{1,4}(A), E_{4,5}(I)][E_{1,4}(B), E_{4,5}(I)] = [E_{1,4}(A+B), E_{4,5}(I)]$: We showed above that $E_{1,4}(B)$ commutes with $[E_{1,4}(A), E_{4,5}(I)]$ and therefore by Proposition 5.3 and (Elm1),

$$[E_{1,4}(A), E_{4,5}(I)][E_{1,4}(B), E_{4,5}(I)] = [E_{1,4}(A)E_{1,4}(B), E_{4,5}(I)] = [E_{1,4}(A+B), E_{4,5}(I)]$$

as needed.

Step 5: We will show how to deduce the relations in $R_{\mathrm{Elm}'}$ that involve $E_{l,l}$'s and $E_{i,j}$'s with $(i,j) \neq (1,5)$ and do not appear in R_{Elm}.

*. $(E_{1,1}(B))^{-1}E_{i,5}(A)E_{1,1}(B) = E_{i,5}(A)$ for $i = 2, 3$: By (Elm2), $E_{i,5}(A) = [E_{i,4}(A), E_{4,5}(I)]$. By free expansion,

$$(E_{1,1}(B))^{-1}[E_{i,4}(A), E_{4,5}(I)]E_{1,1}(B) = [(E_{1,1}(B))^{-1}E_{i,4}(A)E_{1,1}(B), (E_{1,1}(B))^{-1}E_{4,5}(I)E_{1,1}(B)].$$

By (Elm3), $(E_{1,1}(B))^{-1}E_{i,4}(A)E_{1,1}(B) = E_{i,4}(A)$ and $(E_{1,1}(B))^{-1}E_{4,5}(I)E_{1,1}(B) = E_{4,5}(I)$ and thus

$$[(E_{1,1}(B))^{-1}E_{i,4}(A)E_{1,1}(B), (E_{1,1}(B))^{-1}E_{4,5}(I)E_{1,1}(B)] = [E_{i,4}(A), E_{4,5}(I)]$$

as needed.

***.** $(E_{5,5}(B))^{-1}E_{1,j}(A)E_{5,5}(B) = E_{1,j}(A)$ for $j = 3, 4$: By (Elm2), $E_{1,j}(A) = [E_{1,2}(I), E_{2,j}(A)]$. By free expansion,

$$(E_{5,5}(B))^{-1}[E_{1,2}(I), E_{2,j}(A)]E_{5,5}(B) = [(E_{5,5}(B))^{-1}E_{1,2}(I)E_{5,5}(B), (E_{5,5}(B))^{-1}E_{2,j}(A)E_{5,5}(B)].$$

By (Elm3), $(E_{5,5}(B))^{-1}E_{1,2}(I)E_{5,5}(B) = E_{1,2}(I)$ and $(E_{5,5}(B))^{-1}E_{2,j}(A)E_{5,5}(B) = E_{2,j}(A)$ and thus

$$[(E_{5,5}(B))^{-1}E_{1,2}(I)E_{5,5}(B), (E_{5,5}(B))^{-1}E_{2,j}(A)E_{5,5}(B)] = [E_{1,2}(I)), E_{2,j}(A)]$$

as needed.

Step 6: We will show how to deduce the relations in $R_{\mathrm{Elm}'}$ that involve $\phi'(E_{1,5}(A))$. By our definition of ϕ', we need to show for every A, B that

$$(E_{l,l}(B))^{-1}[E_{1,4}(A), E_{4,5}(I)]E_{l,l}(B) = \begin{cases} [E_{1,4}(A), E_{4,5}(I)] & 2 \leqslant l \leqslant 4 \\ [E_{1,4}(B^{-1}A), E_{4,5}(I)] & l = 1 \\ [E_{1,4}(AB), E_{4,5}(I)] & l = 5 \end{cases}.$$

We note that by free expansion,

$$(E_{l,l}(B))^{-1}[E_{1,4}(A), E_{4,5}(I)]E_{l,l}(B) = [(E_{l,l}(B))^{-1}E_{1,4}(A)E_{l,l}(B), (E_{l,l}(B))^{-1}E_{4,5}(I)E_{l,l}(B)].$$

To finish, we will go case by case:

For $l = 2, 3$, we have by (Elm3), that $(E_{l,l}(B))^{-1}E_{1,4}(A)E_{l,l}(B) = E_{1,4}(A)$ and $(E_{l,l}(B))^{-1}E_{4,5}(I)E_{l,l}(B) = E_{4,5}(I)$ and thus we are done.

For $l = 1$, we have by (Elm3), that $(E_{1,1}(B))^{-1}E_{1,4}(A)E_{1,1}(B) = E_{1,4}(B^{-1}A)$ and $(E_{1,1}(B))^{-1}E_{4,5}(I)E_{1,1}(B) = E_{4,5}(I)$ and thus we get $[E_{1,4}(B^{-1}A), E_{4,5}(I)]$ as needed.

For $l = 5$, we have by (Elm3), that $(E_{5,5}(B))^{-1}E_{1,4}(A)E_{5,5}(B) = E_{1,4}(A)$ and $(E_{5,5}(B))^{-1}E_{4,5}(I)E_{5,5}(B) = E_{4,5}(B)$ and thus we get $[E_{1,4}(A), E_{4,5}(B)]$ and by Step 3, this is equal to $[E_{1,4}(AB), E_{4,5}(I)]$ as needed.

For $l = 4$, we have by (Elm3), that $(E_{4,4}(B))^{-1}E_{1,4}(A)E_{4,4}(B) = E_{1,4}(AB)$ and $(E_{4,4}(B))^{-1}E_{4,5}(I)E_{4,4}(B) = E_{4,5}(B^{-1})$ and thus we get $[E_{1,4}(AB), E_{4,5}(B^{-1})]$ and by Step 3, this is equal to $[E_{1,4}(ABB^{-1}), E_{4,5}(I)] = [E_{1,4}(A), E_{4,5}(I)]$ as needed. □

5.6 Element decomposition in $\mathrm{Unip}_n(\mathcal{R})$

Below, we will also need the following result (which is not about the Dehn function):

Proposition 5.14. *Let $H_1, H_2 < \mathrm{Unip}_n(\mathcal{R})$ be the following groups:*

$$H_1 = \langle E_{i,j}(A) : 1 \leqslant i \leqslant j \leqslant 4, A \text{ is } (i,j)\text{-admissible}\rangle,$$

$$H_2 = \langle E_{i,5}(A) : 2 \leqslant i \leqslant 5, A \text{ is } (i,j)\text{-admissible}\rangle.$$

Then every element $g \in \mathrm{Unip}_n(\mathcal{R})$ can be written as a product of at most 6 elements of $H_1 \cup H_2$.

Proof. We note that by definition, H_1 is composed of all the matrices of the form:

$$\begin{pmatrix} A_{1,1} & A_{1,2} & A_{1,3} & A_{1,4} & 0_{k_0} \\ 0_{k_0} & A_{2,2} & A_{2,3} & A_{2,4} & 0_{k_0} \\ 0_{k_1\times k_0} & 0_{k_1\times k_0} & A_{3,3} & A_{3,4} & 0_{k_1\times k_0} \\ 0_{k_0} & 0_{k_0} & 0_{k_0\times k_1} & A_{4,4} & 0_{k_0} \\ 0_{k_0} & 0_{k_0} & 0_{k_0\times k_1} & 0_{k_0} & I \end{pmatrix}$$

and H_2 is composed of all the matrices of the form:

$$\begin{pmatrix} I & 0_{k_0} & 0_{k_1\times k_0} & 0_{k_0} & 0_{k_0} \\ 0_{k_0} & I & 0_{k_1\times k_0} & 0_{k_0} & A_{2,5} \\ 0_{k_1\times k_0} & 0_{k_1\times k_0} & I & 0_{k_1\times k_0} & A_{3,5} \\ 0_{k_0} & 0_{k_0} & 0_{k_0\times k_1} & I & A_{4,5} \\ 0_{k_0} & 0_{k_0} & 0_{k_0\times k_1} & 0_{k_0} & A_{5,5} \end{pmatrix}.$$

Therefore, for every element $g \in \mathrm{Unip}_n(\mathcal{R})$ there are elements $h_1 \in H_1, h_2 \in H_2$ and a matrix A such that $g = h_1h_2E_{1,5}(A)$. We finish by $E_{1,5} = [E_{1,2}(I), E_{2,5}(A)]$ and noting that $E_{1,2}(I) \in H_1$ and $E_{2,5}(A) \in H_2$. □

6 The direct product lemma

For $B, I \subseteq [d]$, we use the notation

$$H_I^B \coloneqq U_{B \setminus I}.$$

Thus, H_I^B is the subgroup generated by the elementary unipotent matrices indexed by $B \setminus I$. We use this notation throughout this section.

For the rest of this section, let $\mathcal{I}$ be a nice set of indexes and let $J \in \mathcal{I}$. The complex we analyze is

$$Y \coloneqq \mathcal{CC}\left(U_{[d]\setminus J}; \left(H_I^{[d]\setminus J}\right)_{I \in \mathcal{I}\setminus\{J\}}\right).$$

Let $B_0, B_1, \ldots, B_r$ be the J-bins. The block decomposition gives the internal direct product decompositions

$$U_{[d]\setminus J} = \prod_{\ell=0}^{r} U_{B_\ell} \qquad \text{and} \qquad H_I^{[d]\setminus J} = \prod_{\ell=0}^{r} H_I^{B_\ell} \quad \text{for every } I \in \mathcal{I} \setminus \{J\}.$$

For every $0 \leqslant \ell \leqslant r$, denote by

$$Y_\ell \coloneqq \mathcal{CC}\left(U_{B_\ell}; \left(H_I^{B_\ell}\right)_{I \in \mathcal{I}\setminus\{J\}}\right)$$

the coset complex corresponding to the ℓ-th factor of the direct product.

We recall which presentations are used in the Dehn functions below. For every $I \in \mathcal{I} \setminus \{J\}$, let R_I be the multiplication-table relations of $H_I^{[d]\setminus J}$, and let R_I^ℓ be the multiplication-table relations of $H_I^{B_\ell}$. In the notation of Definition 2.35, let

$$\mathrm{Dehn}_Y = \mathrm{Dehn}_{S,R}, \qquad S = \bigcup_{I \in \mathcal{I}\setminus\{J\}} H_I^{[d]\setminus J}, \qquad R = \bigcup_{I \in \mathcal{I}\setminus\{J\}} R_I,$$

whereas

$$\mathrm{Dehn}_{Y_\ell} = \mathrm{Dehn}_{S_\ell,R_\ell}, \qquad S_\ell = \bigcup_{I \in \mathcal{I}\setminus\{J\}} H_I^{B_\ell}, \qquad R_\ell = \bigcup_{I \in \mathcal{I}\setminus\{J\}} R_I^\ell.$$

Thus, each Dehn_{Y_ℓ} is computed using the presentation associated with the component complex Y_ℓ, while Dehn_Y is computed using the presentation associated with Y.

Lemma 6.1. *Assume that there is a constant A such that*

$$\mathrm{Dehn}_{Y_\ell}(s) \leqslant As$$

for every $0 \leqslant \ell \leqslant r$ and every $s \in \mathbb{N}$, where each Dehn function is taken with respect to the coset-complex presentation of Y_ℓ described above. Then, with respect to the coset-complex presentation of Y,

$$\mathrm{Dehn}_Y(t) \leqslant O((t(r+1))^2) + (r+1)At$$

for every $t \in \mathbb{N}$.

We will start with the following more basic claim that we use as a building block. We will say that an upper-triangular matrix g is *supported on block size k*, if

$$g = \begin{pmatrix} I_{k'} & 0 & 0 \\ 0 & A & 0 \\ 0 & 0 & I_{k''} \end{pmatrix},$$

where A is a $k \times k$ upper-triangular matrix. We will call A *the support of g*.

Claim 6.2. Let $0 \leqslant \ell \leqslant r$ and denote $B_\ell = [a, b]$. There is a universal polynomial p (independent of $\ell, r, [a, b]$) such that for every $C \in \mathbb{N}, 20 \geqslant C \geqslant 1$, every $I \in \mathcal{I} \setminus \{J\}$ and every $g \in H_I^{B_\ell}$, there are $g_1, ..., g_s \in H_I^{B_\ell}$ such that $s \leqslant p(C)$, $g = g_1 ... g_s$, and each g_i is supported on block size $\leqslant \frac{b-a}{C}$. Moreover, g can be transformed into $g_1 ... g_s$ by applying at most $p(C)$ relations in R_I^ℓ.

Proof. Fix $0 \leqslant \ell \leqslant r$, denote $B_\ell = [a, b]$. Also fix $I \in \mathcal{I} \setminus \{J\}$ and $g \in H_I^{B_\ell}$.

We first observe that if $g \in H_I^{B_\ell}$ is decomposed into a product $g = g_1 \cdots g_6$ of matrices in U_{B_ℓ}, we can assume that all six matrices belong to $H_I^{B_\ell}$. Indeed, replace each g_i by the matrix obtained by setting to zero all entries outside the diagonal blocks determined by $I \cap B_\ell$. Since the diagonal blocks of a product of upper-triangular matrices are the products of the corresponding diagonal blocks, the resulting matrices still multiply to g. Moreover, this operation yields matrices in $H_I^{B_\ell}$.

We show by induction on $m \geqslant 0$ that g can be transformed into a product of at most 6^m elements of $H_I^{B_\ell}$, each supported on block size at most $\max\{(b-a+2)(9/10)^m, 5\}$, using at most $6^m - 1$ relations in R_I^ℓ.

For $m = 0$, we simply take g itself. Assume that the statement holds for m. Consider one of the resulting matrices h, supported on block size k. If $k \leqslant 5$, we leave it unchanged. Otherwise, by Proposition 5.14, h can be decomposed into at most six matrices supported on block size at most

$$k - \lfloor k/5 \rfloor \leqslant 9k/10.$$

By the observation above, we can take all these matrices to belong to $H_I^{B_\ell}$. Thus, replacing h by their product requires at most five relations from R_I^ℓ.

Applying this to every factor, we obtain a product of at most 6^{m+1} elements, each supported on block size at most $\max\{(b-a+2)(9/10)^{m+1}, 5\}$. The total number of relations used is at most

$$(6^m - 1) + 5 \cdot 6^m = 6^{m+1} - 1.$$

This completes the induction.

Finally, take $m = \lceil \log_{10/9} C \rceil + 1$. The resulting factors have block size at most $\max\{\frac{9}{10}(b-a+2)/C, 5\} \leqslant^{b-a \geqslant 100} \frac{b-a}{C}$, and both their number and the number of relations used are bounded by

$$6^m \leqslant 6^2 C^{18}.$$

Thus we can take $p(C) = 6^2 C^{18}$. □

Proof of Lemma 6.1. Let $w = w_1 w_2 \dots w_s$ be a word over S such that $s \leqslant t$ and $w = 1$ in $U_{[d]\setminus J}$. For every $1 \leqslant j \leqslant s$, choose $I_j \in \mathcal{I} \setminus \{J\}$ such that

$$w_j \in H_{I_j}^{[d]\setminus J}.$$

By the direct product decomposition of $H_{I_j}^{[d]\setminus J}$, we can write

$$w_j = w_j^0 w_j^1 \dots w_j^r$$

with

$$w_j^\ell \in H_{I_j}^{B_\ell}$$

for every $0 \leqslant \ell \leqslant r$.

Every partial product of the factors $w_j^0, w_j^1, \dots, w_j^r$ belongs to $H_{I_j}^{[d]\setminus J}$. Therefore, using the multiplication-table relations of $H_{I_j}^{[d]\setminus J}$, we can replace w_j by $w_j^0 w_j^1 \dots w_j^r$ in $O(r)$ relation steps. Applying this to every letter transforms w, in $O(sr)$ steps, into

$$w_1^0 w_1^1 \dots w_1^r w_2^0 w_2^1 \dots w_2^r \dots w_s^0 w_s^1 \dots w_s^r.$$

We next sort the letters according to their superscripts, obtaining

$$(w_1^0 w_2^0 \dots w_s^0)(w_1^1 w_2^1 \dots w_s^1) \dots (w_1^r w_2^r \dots w_s^r).$$

We use a bubble-sort procedure, repeatedly commuting adjacent letters whose superscripts are in the wrong order. The following claim shows that each such commutation costs only a constant number of relation steps (its proof is in the end of this section).

Claim 6.3. Let $x = w_i^j$ and $y = w_{i'}^{j'}$ for $j \neq j'$. Then one can transform xy into yx using $O(1)$ relations from R.

There are $s(r+1)$ letters, and hence at most $O((s(r+1))^2)$ adjacent commutations. By Claim 6.3, the sorting therefore uses $O((s(r+1))^2)$ relation steps.

For every $0 \leqslant \ell \leqslant r$, denote

$$w^{(\ell)} \coloneqq w_1^\ell w_2^\ell \dots w_s^\ell.$$

Since

$$U_{[d]\setminus J} = \prod_{\ell=0}^{r} U_{B_\ell}$$

and $w = 1$ in $U_{[d]\setminus J}$, each $w^{(\ell)}$ is equal to the identity in U_{B_ℓ}. It is a word of length at most $s \leqslant t$ over S_ℓ, and therefore

$$\mathrm{Area}_{R_\ell}(w^{(\ell)}) \leqslant \mathrm{Dehn}_{Y_\ell}(t) \leqslant At.$$

Every relation in R_ℓ is a multiplication-table relation inside some

$$H_I^{B_\ell} \leqslant H_I^{[d]\setminus J},$$

and hence is also a relation in R. Thus, the same reduction contracts $w^{(\ell)}$ inside the presentation defining Dehn_Y. Contracting all $r+1$ component words costs at most $(r+1)At$ additional relation steps. Altogether,

$$\mathrm{Area}_R(w) \leqslant O((t(r+1))^2) + (r+1)At.$$

Taking the maximum over all trivial words of length at most t proves the lemma.

It remains to prove Claim 6.3.

Proof of Claim 6.3. Write

$$B_j = [a, b], \qquad B_{j'} = [a', b'], \qquad c = b - a, \qquad c' = b' - a'.$$

By Claim 6.2, we can decompose

$$x = x_1 x_2 \dots x_{C_0} \qquad \text{and} \qquad y = y_1 y_2 \dots y_{C_1},$$

where $C_0, C_1 = O(1)$, every x_m is supported on a block of size at most $\frac{c}{20}$, and every $y_{m'}$ is supported on a block of size at most $\frac{c'}{20}$. These decompositions can be performed using $O(1)$ multiplication-table relations in $H_{I_i}^{[d]\setminus J}$ and $H_{I_{i'}}^{[d]\setminus J}$, respectively.

For every x_m, there is some $q \in \{0, 1, 2, 3\}$ such that the indexes of its support are contained in

$$[a + q\lfloor \tfrac{c}{5} \rfloor, \min\{a + (q+2)\lceil \tfrac{c}{5} \rceil + 4, b\}].$$

Similarly, for every $y_{m'}$, there is some $q' \in \{0, 1, 2, 3\}$ such that the indexes of its support are contained in the analogous interval inside $B_{j'}$.

By niceness of $\mathcal{I}$, for every pair of these intervals there exists $K \in \mathcal{I} \setminus \{J\}$ that avoids both intervals. Consequently, the corresponding pair $x_m, y_{m'}$ lies in the same subgroup

$$H_K^{[d]\setminus J} = U_{[d]\setminus(J\cup K)}.$$

Since x_m and $y_{m'}$ belong to distinct direct factors, they commute in $U_{[d]\setminus J}$. We can therefore transform $x_m y_{m'}$ into $y_{m'} x_m$ using two multiplication-table relations of $H_K^{[d]\setminus J}$: first replace $x_m y_{m'}$ by their product, and then expand this product as $y_{m'} x_m$.

Commuting all pairs $x_m, y_{m'}$, and then recombining the factors, transforms xy into yx using $O(C_0 C_1) = O(1)$ relations from R. □

□

7 Bounding the diameter and the Dehn function

The aim of this section is to prove Lemma 3.21 and Lemma 3.26 and thus conculding the proof of Theorem 3.20.

Let $\mathcal{I}$ is a nice set of indices, let $J \in \mathcal{I}$ and let

$$Y \coloneqq \mathcal{CC}\left(U_{[d]\setminus J}; \left(H_I^{[d]\setminus J}\right)_{I\in\mathcal{I}\setminus\{J\}}\right)$$

and for $0 \leqslant \ell \leqslant r$

$$Y_\ell \coloneqq \mathcal{CC}\left(U_{B_\ell}; \left(H_I^{B_\ell}\right)_{I\in\mathcal{I}\setminus\{J\}}\right).$$

In light of Theorem 3.19, in order to bound the cone radii of Y, it is enough to bound the diameter of Y and the Dehn function of $U_{[d]\setminus J}$.

We will start by bounding the diameter of Y (proving Lemma 3.21):

Proposition 7.1. *It holds that* $\operatorname{diam}(Y) \leqslant 6(r+1)+1$.

Proof. By [KO21, Proposition 6.3], it is enough to show that every $g \in U_{[d]\setminus J}$ can be written as a product of at most $6(r+1)$ elements in $\bigcup_{I\in\mathcal{I}\setminus\{J\}} H_I^{[d]\setminus J}$. Since $U_{[d]\setminus J}$ is a direct product of the groups $U_{B_0}, ..., U_{B_r}$ it follows that there are $g_\ell \in U_{B_\ell}, 0 \leqslant \ell \leqslant r$ such that $g = g_0 ... g_r$. Thus it is enough to show that for every $0 \leqslant \ell \leqslant r$ and every $g_\ell \in U_{B_\ell}$, g_ℓ can be written as at most 6 elements of $\bigcup_{I\in\mathcal{I}\setminus\{J\}} H_I^{B_\ell}$.

Fix $0 \leqslant \ell \leqslant r$. By the assumption that $\mathcal{I}$ is nice, it holds that there are $I_1, I_2 \in \mathcal{I}$ such that $I_1 \cap B_\ell \subseteq [a, a + \lfloor \frac{b-a}{5} \rfloor]$, $I_2 \cap B_\ell \subseteq [b - \lfloor \frac{b-a}{5} \rfloor, b]$. Thus by Proposition 5.14, each element in U_{B_ℓ} can be written as a product of at most 6 elements in $H_{I_1}^{B_\ell} \cup H_{I_2}^{B_\ell}$ as needed. ◻

In light of Lemma 6.1, in order to prove a bound on the Dehn function of $U_{[d]\setminus J}$, it is enough to prove a bound on the Dehn function of $\operatorname{Dehn}_{Y_\ell}$ for every $0 \leqslant \ell \leqslant r$, where we recall that R_I^ℓ is the multiplication-table relations of $H_I^{B_\ell}$ and

$$\operatorname{Dehn}_{Y_\ell} = \operatorname{Dehn}_{S_\ell, R_\ell}, \qquad S_\ell = \bigcup_{I\in\mathcal{I}\setminus\{J\}} H_I^{B_\ell}, \qquad R_\ell = \bigcup_{I\in\mathcal{I}\setminus\{J\}} R_I^\ell.$$

Thus, we can restate and prove Lemma 3.26:

Lemma 7.2. *There is a constant A such that for every $0 \leqslant \ell \leqslant r$ and every $s \in \mathbb{N}$, $\operatorname{Dehn}_{S_\ell, R_\ell}(s) \leqslant As$.*

Proof. Fix $0 \leqslant \ell \leqslant r$ and denote $B_\ell = [a, b], n = b - a$. In the notation of Section 5, U_{B_ℓ} is isomorphic to $\operatorname{Unip}_n(\mathbb{F}_q)$.

By the assumption that $\mathcal{I}$ is nice, it holds that there are $I_1, I_2, I_3 \in \mathcal{I}$ such that $I_1 \cap B_\ell \subseteq [a, a + \lfloor \frac{b-a}{5} \rfloor]$, $I_2 \cap B_\ell \subseteq [b - \lfloor \frac{b-a}{5} \rfloor, b]$ and $I_3 \cap B_\ell \subseteq [a + \lceil \frac{b-a}{5} \rceil, a + 3\lfloor \frac{b-a}{5} \rfloor]$.

Thus, in the notion of Section 5, $S_{\mathrm{Elm}} \subseteq S_\ell$ and every relation in R_{Elm} can be derived from $R_{I_1}^\ell \cup R_{I_2}^\ell \cup R_{I_3}^\ell \subseteq R_\ell$ by applying at most 4 relations in $R_{I_1}^\ell \cup R_{I_2}^\ell \cup R_{I_3}^\ell \subseteq R_\ell$.

By Theorem 5.5, there is a universal constant C' such that $\operatorname{Dehn}_{S_{\mathrm{Elm}}, R_{\mathrm{Elm}}}(m) \leqslant C'm$.

Thus it is enough to show that there are universal constants C_0, C_1 such that given a word $g \in S_\ell$, g can be transformed into $w_1 ... w_m$ where $m \leqslant C_0, w_1, ..., w_m \in S_{\mathrm{Elm}}$ and this requires C_1 relations in R_ℓ. Indeed, assume we showed this for every $g \in S_\ell$. Then given a trivial word $g_1 ... g_s$ in the alphabet S_ℓ, it can be transformed into a word $w_1 ... w_t$ in the alphabet S_{Elm} with $t \leqslant C_0 s$ using $C_1 s$ relations in R_ℓ. Then using Theorem 5.5, $w_1 ... w_t$ can be reduced to the trivial word using at most $C'm$ relations in R_{Elm} and thus by at most $4C'm$ relations in R_ℓ. It follows that

$$\operatorname{Dehn}_{S_\ell, R_\ell}(s) \leqslant (4C_0C' + C_1)s,$$

as needed.

Let $g \in S_\ell$. By Claim 6.2, there are constants C_0', C_1' such that g can be transformed into $w_1' ... w_{m'}'$ where $m' \leqslant C_0'$, and w_i' are of block size $\leqslant \frac{n}{5}$. Thus applying at most 2 relations from R_ℓ, each w_i' can be transformed to a product of at most 3 elements in S_{Elm}. It follows that we can take $C_0 = 3C_0'$ and $C_1 = C_1' + 2C_0'$ (all are universal constants). ◻

With the above Lemma, we can finally restate and prove Theorem 3.20:

Theorem 7.3. *With respect to the coset-complex presentation of Y,*

$$\mathrm{Dehn}_Y(t) \leqslant O((t(r+1))^2)$$

for every $t \in \mathbb{N}$.

Proof. Combine Lemma 6.1 and Lemma 7.2. □

8 AI Methodology

All ideas and proofs are independent work of the authors. We used Open AI's Astra 6 to simplify arguments and optimize constants in the proof of Lemma 3.7 as our original proof had a bound of $r \geqslant 10^{2700}$. We also used various other LLMs to proof read, edit text, create diagrams and search the literature for references. The authors verified the correctness and originality of all content, including references.

References

[Tit86] J. Tits. "Ensembles ordonnés, immeubles et sommes amalgamées." In: *Bulletin de la Société Mathématique de Belgique. Série A* 38 (1986), pp. 367–387. ISSN: 0037-9476.

[AH93] H. Abels and S. Holz. "Higher generation by subgroups." In: *Journal of Algebra* 160.2 (1993), pp. 310–341. ISSN: 0021-8693,1090-266X. DOI: 10.1006/jabr.1993.1190. URL: https://doi.org/10.1006/jabr.1993.1190.

[Abr96] Peter Abramenko. *Twin buildings and applications to S-arithmetic groups*. Vol. 1641. Lecture Notes in Mathematics. Springer-Verlag, Berlin, 1996, pp. x+123. ISBN: 3-540-61973-9. DOI: 10.1007/BFb0094079. URL: https://doi.org/10.1007/BFb0094079.

[RS96] Ronitt Rubinfeld and Madhu Sudan. "Robust Characterizations of Polynomials with Applications to Program Testing." In: *SIAM Journal on Computing* 25.2 (1996), pp. 252–271. DOI: 10.1137/S0097539793255151.

[RS97] Ran Raz and Shmuel Safra. "A Sub-constant Error-probability Low-degree Test, and a Sub-constant Error-probability PCP Characterization of NP." In: *Proceedings of the Twenty-ninth Annual ACM Symposium on Theory of Computing*. STOC '97. El Paso, Texas, USA: ACM, 1997, pp. 475–484. ISBN: 0-89791-888-6. DOI: 10.1145/258533.258641. URL: http://doi.acm.org/10.1145/258533.258641.

[BD01] Daniel K. Biss and Samit Dasgupta. "A presentation for the unipotent group over rings with identity." In: *J. Algebra* 237.2 (2001), pp. 691–707. ISSN: 0021-8693,1090-266X. DOI: 10.1006/jabr.2000.8604. URL: https://doi-org.bengurionu.idm.oclc.org/10.1006/jabr.2000.8604.

[KMS02] Marcos Kiwi, Frédéric Magniez, and Miklos Santha. "Exact and Approximate Testing/Correcting of Algebraic Functions: A Survey." In: *Theoretical Aspects of Computer Science: Advanced Lectures*. Springer, 2002, pp. 30–83. DOI: 10.1007/3-540-45878-6_2.

[AS03] Sanjeev Arora and Madhu Sudan. "Improved Low-Degree Testing and its Applications." In: *Combinatorica* 23.3 (2003), pp. 365–426. DOI: 10.1007/s00493-003-0025-0.

[LM06] Nathan Linial* and Roy Meshulam*. "Homological connectivity of random 2-complexes." In: *Combinatorica* 26.4 (2006), pp. 475–487.

[DM07] A. Devillers and B. Mühlherr. *On the simple connectedness of certain subsets of buildings*. 2007. DOI: 10.1515/FORUM.2007.037. URL: https://doi.org/10.1515/FORUM.2007.037.

[MR08] Dana Moshkovitz and Ran Raz. "Sub-Constant Error Low Degree Test of Almost-Linear Size." In: *SIAM Journal on Computing* 38.1 (2008), pp. 140–180. DOI: 10.1137/060656838.

[Gro10] M. Gromov. "Singularities, Expanders and Topology of Maps. Part 2: from Combinatorics to Topology Via Algebraic Isoperimetry." In: *Geom. Funct. Anal.* 20 (2010), pp. 416–526. DOI: 10.1007/s00039-010-0073-8.

[EK16] Shai Evra and Tali Kaufman. "Bounded degree cosystolic expanders of every dimension." In: *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 2016, pp. 36–48.

[KKL16] Tali Kaufman, David Kazhdan, and Alexander Lubotzky. "Isoperimetric inequalities for Ramanujan complexes and topological expanders." In: *Geometric and Functional Analysis* 26.1 (2016), pp. 250–287.

[LMM16] Alexander Lubotzky, Roy Meshulam, and Shahar Mozes. "Expansion of building-like complexes." In: *Groups, Geometry, and Dynamics* 10.1 (2016), pp. 155–175.

[CM17] Matt Clay and Dan Margalit, eds. *Office Hours with a Geometric Group Theorist*. Princeton, NJ: Princeton University Press, 2017. ISBN: 978-0-69115-866-2.

[KM18] Tali Kaufman and David Mass. "Good Distance Lattices from High Dimensional Expanders." (manuscript). 2018. eprint: 1803.02849.

[KO18] Tali Kaufman and Izhar Oppenheim. "Construction of new local spectral high dimensional expanders." In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 2018, pp. 773–786.

[HS19] Prahladh Harsha and Ramprasad Saptharishi. "A note on the elementary construction of high-dimensional expanders of Kaufman and Oppenheim." In: *arXiv preprint arXiv:1912.11225* (2019).

[KM19] Dmitry N Kozlov and Roy Meshulam. "Quantitative aspects of acyclicity." In: *Research in the Mathematical Sciences* 6.4 (2019), pp. 1–32.

[KM21] Tali Kaufman and David Mass. "Unique-Neighbor-Like Expansion and Group-Independent Cosystolic Expansion." In: *32nd International Symposium on Algorithms and Computation (ISAAC)*. 2021, 56:1–56:17. DOI: 10.4230/LIPIcs.ISAAC.2021.56.

[KO21] Tali Kaufman and Izhar Oppenheim. "Coboundary and Cosystolic Expansion from Strong Symmetry." In: *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*. Vol. 198. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 84:1–84:16. DOI: 10.4230/LIPIcs.ICALP.2021.84. URL: https://doi.org/10.4230/LIPIcs.ICALP.2021.84.

[DM22] Irit Dinur and Roy Meshulam. "Near coverings and cosystolic expansion." In: *Archiv der Mathematik* 118.5 (May 2022), pp. 549–561. ISSN: 1420-8938. DOI: 10.1007/s00013-022-01720-6. URL: https://doi.org/10.1007/s00013-022-01720-6.

[GLL22] Tom Gur, Noam Lifshitz, and Siqi Liu. "Hypercontractivity on high dimensional expanders." In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 2022, pp. 176–184.

[KM22] Tali Kaufman and David Mass. "Double Balanced Sets in High Dimensional Expanders." In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*. 2022, 3:1–3:17. DOI: 10.4230/LIPIcs.APPROX/RANDOM.2022.3.

[OP22] Ryan O'Donnell and Kevin Pratt. "High-Dimensional Expanders from Chevalley Groups." In: *37th Computational Complexity Conference, CCC 2022, July 20-23, 2022, Philadelphia, PA, USA*. Ed. by Shachar Lovett. Vol. 234. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 18:1–18:26. DOI: 10.4230/LIPIcs.CCC.2022.18. URL: https://doi.org/10.4230/LIPIcs.CCC.2022.18.

[BM23] Mitali Bafna and Dor Minzer. "Characterizing Direct Product Testing via Coboundary Expansion." 2023.

[CL23] Michael Chapman and Alexander Lubotzky. *Stability of Homomorphisms, Coverings and Cocycles I: Equivalence*. 2023. eprint: arXiv:2310.17474.

[DD23a] Yotam Dikstein and Irit Dinur. "Agreement theorems for high dimensional expanders in the small soundness regime: the role of covers." In: *CoRR* abs/2308.09582 (2023). DOI: 10.48550/ARXIV.2308.09582. arXiv: 2308.09582. URL: https://doi.org/10.48550/arXiv.2308.09582.

[DD23b] Yotam Dikstein and Irit Dinur. "Coboundary and cosystolic expansion without dependence on dimension or degree." In: *CoRR* abs/2304.01608 (2023). DOI: 10.48550/arXiv.2304.01608. arXiv: 2304.01608. URL: https://doi.org/10.48550/arXiv.2304.01608.

[DD23c] Yotam Dikstein and Irit Dinur. "Swap Cocycle Expansion." In: (2023). eprint: arXiv:2312.15325.

[KO23] Tali Kaufman and Izhar Oppenheim. "High dimensional expanders and coset geometries." In: *European Journal of Combinatorics* 111 (2023), Paper No. 103696, 31. ISSN: 0195-6698,1095-9971. DOI: 10.1016/j.ejc.2023.103696. URL: https://doi.org/10.1016/j.ejc.2023.103696.

[BLM24] Mitali Bafna, Noam Lifshitz, and Dor Minzer. *Constant Degree Direct Product Testers with Small Soundness*. 2024. eprint: arXiv:2402.00850.

[DDL24] Yotam Dikstein, Irit Dinur, and Alexander Lubotzky. *Low Acceptance Agreement Tests via Bounded-Degree Symplectic HDXs*. 2024. URL: https://eccc.weizmann.ac.il/report/2024/019/.

[Har+24] Prahladh Harsha, Mrinal Kumar, Ramprasad Saptharishi, and Madhu Sudan. "An Improved Line-Point Low-Degree Test." In: *Proceedings of the 65th IEEE Symposium on Foundations of Computer Science (FOCS)*. 2024, pp. 1883–1892. DOI: 10.1109/FOCS61266.2024.00113.

[Baf+25] Mitali Bafna, Dor Minzer, Nikhil Vyas, and Zhiwei Yun. "Quasi-Linear Size PCPs with Small Soundness from HDX." In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC)*. 2025, pp. 45–53. DOI: 10.1145/3717823.3718197.

[DLZ25] Irit Dinur, Siqi Liu, and Rachel Yun Zhang. "New Codes on High Dimensional Expanders." In: *40th Computational Complexity Conference (CCC)*. 2025, 27:1–27:42. DOI: 10.4230/LIPIcs.CCC.2025.27.

[GV25] Laura Grave de Peralta and Inga Valentiner-Branth. "High-dimensional expanders from Kac–Moody–Steinberg groups." In: *European Journal of Combinatorics* 126 (2025), p. 104131. DOI: 10.1016/j.ejc.2025.104131.

[KOW25] Tali Kaufman, Izhar Oppenheim, and Shmuel Weinberger. "Coboundary Expansion of Coset Complexes." In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC)*. 2025, pp. 1722–1731. DOI: 10.1145/3717823.3718136.

[LO25] Jonathan Leake and Shayan Oveis Gharan. *Trickle-down Theorems via $\mathcal{C}$-Lorentzian Polynomials II: Pairwise Spectral Influence and Improved Dobrushin's Condition*. Version 3, April 2026. 2025. arXiv: 2510.06549 [math.CO].

[OS25] Ryan O'Donnell and Noah G. Singer. *Low-soundness direct-product testers and PCPs from Kaufman–Oppenheim complexes*. 2025. arXiv: 2511.10514.

[OV25] Izhar Oppenheim and Inga Valentiner-Branth. *Non-Abelian expansion of congruence KMS complexes*. 2025. arXiv: 2510.18351.

[BDL26] Inbar Ben Yaacov, Yotam Dikstein, and Siqi Liu. "Sparsified low degree tests in the low acceptance regime." Unpublished manuscript; working title. 2026.

[Dik+26a] Yotam Dikstein, Irit Dinur, Max Hopkins, and Noah Singer. "Sounder, shorter, and simpler PCPs from HDX neighborhoods." Unpublished, submitted to STOC 2027. 2026.

[Dik+26b] Yotam Dikstein, Max Hopkins, Russell Impagliazzo, and Toniann Pitassi. "High Rate Efficient Local List Decoding from HDX." In: *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC)*. 2026, pp. 1789–1799. DOI: 10.1145/3798129.3800886.

[HR26] Max Hopkins and Arka Ray. "A Simple Sub-Polynomial Degree Coboundary Expander." In: *41st Computational Complexity Conference (CCC 2026)*. Ed. by Dana Moshkovitz. Vol. 383. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2026, 26:1–26:41. DOI: 10.4230/LIPIcs.CCC.2026.26.

[OV26] Izhar Oppenheim and Inga Valentiner-Branth. "New cosystolic high-dimensional expanders from KMS groups." In: *Mathematische Annalen* 396.3 (Oct. 2026), p. 52. ISSN: 1432-1807. DOI: 10.1007/s00208-026-03570-x. URL: https://doi.org/10.1007/s00208-026-03570-x.

[Val26] Inga Valentiner-Branth. *High-dimensional expanders from Kac–Moody–Steinberg groups*. PhD thesis. 2026.

A PCPs with improved proof-length dependence

We show that substituting Theorem 1.2 into the construction of Bafna, Minzer, Vyas and Yun improves its dependence on the soundness parameter. We follow their construction and only record the parameter changes.

A.1 Adapting [Baf+25, Theorem 1.2]

Following [Baf+25] we prove the following theorem.

Theorem (Restatement of Theorem 1.3)**.** *There exists an absolute constant $a > 0$ such that, for every fixed $\varepsilon \in (0, 1/2)$ and all sufficiently large n, the PCP construction of [Baf+25], instantiated using Theorem 1.2, gives a polynomial-time reduction from a 3-SAT instance φ of size n to a two-query PCP with perfect completeness, soundness at most ε, and proof length at most*

$$n(\log n)^{\exp(O(\varepsilon^{-a}))}.$$

The proof alphabet has size at most $\exp\left((1/\varepsilon)^{O(\log(1/\varepsilon))}\right)$.

We first establish the agreement theorem and the quantitative complex construction, and then apply the embedding and composition arguments of [Baf+25].

A.1.1 Complexes for the PCP construction

In this subsection we give a quantitative variant of [Baf+25, Theorem 2.13]. The improvement comes from substituting Theorem 1.2 into the agreement-testing argument of [DD23a]. Throughout this section, all polynomial dependencies have absolute constant exponents, independent of the soundness parameter.

Agreement from swap coboundary expansion. The V-test on an ensemble $F = \{F_s : s \to \Sigma\}_{s \in X(k)}$ samples two k-faces s, s' whose intersection has size $\lfloor \sqrt{k+1} \rfloor$, according to the natural distribution induced by X, and accepts if $F_s|_{s \cap s'} = F_{s'}|_{s \cap s'}$. Equivalently, we first sample their union according to the natural face distribution and then sample the pair uniformly subject to the prescribed intersection size.

We use the following straightforward consequence of the proof of [DD23a, Theorem 3.1] (a similar theorem with slightly different guarantees was proven in [BM23]). Here and below, taking a parameter to be a sufficiently large polynomial means that the polynomial is fixed independently of the complex and the alphabet.

Theorem A.1 (Agreement from swap coboundary expansion)**.** *Let $\varepsilon \in (0, 1/2)$, let $k \geqslant \exp(\mathrm{poly}(1/\varepsilon))$, and let d be a sufficiently large polynomial in k. Let X be a d-dimensional clique complex satisfying the following:*

1. *X is a $1/d^2$-one-sided local spectral expander.*
2. *For a large intermediate dimension $r = \mathrm{poly}(k)$, with d polynomially in r, the faces complex $\mathcal{F}^r X$ is a $(\mathrm{poly}(1/r), m)$-coboundary expander for $m = \mathrm{poly}(1/\varepsilon)$.*

Let Σ be a finite alphabet and let $F = \{F_s : s \to \Sigma\}_{s \in X(k)}$. If the V-test accepts F with probability at least ε, then there exists $f : X(0) \to \Sigma$ such that

$$\mathop{\mathbb{P}}_{s \in X(k)}\left[\mathrm{dist}(F_s, f|_s) \leqslant \gamma\right] \geqslant \varepsilon^{O(1)},$$

where $\gamma = \exp(\mathrm{poly}(1/\varepsilon))\,\mathrm{poly}(1/k)$.

Remark A.2. Note that [DD23a] additionally requires a property called well connectedness, which is used to descend a cover of the faces complex to a cover of X. This requirement can be removed under our stronger assumption of coboundary expansion: the corrected cochain is a coboundary, so we can synchronize the lists directly without constructing or descending a cover. See also [BM23].

Proof sketch. We explain the quantitative improvement over the argument of [DD23a, Section 3]. In that proof, a cochain representing the agreement problem is constructed on a sufficiently large faces complex of X. Coboundary expansion is then used to correct this cochain to a nearby coboundary, from which a global assignment is constructed. We refer the reader to [DD23a, Section 3] for the details of this argument.

The fraction of triangles violated by the constructed cochain is at most $\mu = \exp(-\operatorname{poly}(\varepsilon) r/k) + O(r/d)$. In [DD23a], the available coboundary expansion was $\exp(-O(\sqrt{r}))$, and therefore the argument required $d \gg r\exp(O(\sqrt{r}))$ to ensure that μ is sufficiently small for the correction to be close to the original cochain. Since $r > \operatorname{poly}(\varepsilon^{-1}, k)$, this resulted in a superpolynomial dependence of d on k. Using our improved coboundary expansion of $1/\operatorname{poly}(r)$, it is enough to take $d = \operatorname{poly}(r) = \operatorname{poly}(k)$, while the rest of the argument remains unchanged.

The bound $d = \operatorname{poly}(k)$ follows by tracking the dimension requirements throughout the argument of [DD23a, Section 3]; we omit the full parameter bookkeeping for brevity. □

The quantitative complex construction. We now give the variant of [Baf+25, Theorem 2.13] used for the PCP construction. For a target size n, we construct a complex with between n and $nq^{O(d^2)}$ vertices, where q is the field size. Since q is polylogarithmic in n, this additional factor is absorbed into the claimed proof-length bound.

We will use the following fact about the links of KMS complexes. We sketch its proof after the proof of Theorem A.4.

Fact A.3. *There is an absolute constant $C_0 > 0$ such that the following holds. Let $X = X(A_d, \mathbb{F}_q, f)$, where f is irreducible of degree at least two and $q \geqslant C_0(d+1)$.*

1. *For every face link L, including X itself, the bipartite graph between any two distinct colors is uniformly weighted over its edges and has second singular value at most $4/\sqrt{q}$.*
2. *For every vertex link L, order its colors as $1, \dots, d$ according to the subspace dimensions in its opposition-complex description. For $1 \leqslant a < b \leqslant d$, the bipartite graph between colors a, b has diameter $O(d/(b-a))$, with an absolute implied constant. In particular, this diameter is constant when $b - a = \Omega(d)$.*

Theorem A.4. *For every $\varepsilon \in (0, 1/2)$, there are parameters*

$$k = \exp(\operatorname{poly}(1/\varepsilon)), \qquad d = \operatorname{poly}(k), \qquad q_0 = \exp(\operatorname{poly}(1/\varepsilon))$$

with the following property. For every fixed $C > 0$, all sufficiently large n, and every prime $q = \Theta((\log n)^C)$ with $q \geqslant q_0$, one can construct in polynomial time a d-dimensional KMS complex X such that

$$n \leqslant |X(0)| \leqslant nq^{O(d^2)}$$

and:

1. *X is a $(d+1)$-partite clique complex.*
2. *Every vertex belongs to at most $q^{O(d^2)}$ top-dimensional faces.*
3. *For every vertex link L, there is a group of automorphisms acting transitively on $L(d-1)$.*
4. *For every vertex link L, order its colors as $1, \dots, d$ according to the subspace dimensions in its opposition-complex description. For distinct colors i, j in this ordering, the bipartite graph between $L[i]$ and $L[j]$ is uniformly weighted over its edges and has diameter $O(d/|i-j|)$.*
5. *For every face link L, including X itself, the bipartite graph between any two distinct colors has second singular value at most $4/\sqrt{q}$.*

6. *The normalized adjacency operator of the graph $(X(0), X(1))$ has second singular value at most*

$$\frac{1}{d} + \frac{4}{\sqrt{q}} \leqslant \frac{2}{d}.$$

7. *Let Σ be a finite alphabet and let $F = \{F_s : s \to \Sigma\}_{s \in X(k)}$. Consider the following test:*

   (a) *Sample $t \in X(d)$.*

   (b) *Sample a uniform subset $b \subseteq t$ of size $\lfloor \sqrt{k+1} \rfloor$.*

   (c) *Independently sample two uniform k-faces $s, s' \subseteq t$ containing b.*

   (d) *Accept if $F_s|_b = F_{s'}|_b$.*

   If this test accepts with probability at least ε, then there exists $f : X(0) \to \Sigma$ such that

$$\mathop{\mathbb{P}}_{s \in X(k)} [\mathrm{dist}(F_s, f|_s) \leqslant \varepsilon] \geqslant \varepsilon^{O(1)}.$$

All constants in the displayed parameter bounds are independent of ε.

Proof sketch. Choose $k = \exp(\mathrm{poly}(1/\varepsilon))$, $d = \mathrm{poly}(k)$, and $q_0 = \exp(\mathrm{poly}(1/\varepsilon))$ sufficiently large that Theorem A.1 applies to the family of KMS complexes for every prime $q \geqslant q_0$, with acceptance threshold $\varepsilon/2$ and error $\gamma \leqslant \varepsilon$. We also take q_0 large enough to satisfy the hypotheses of Fact A.3 and the numerical bound in the sixth item.

The first three items and the sixth follow from basic properties of the KMS complexes and their links [GV25; Dik+26b]. For the clique-complex property, we take $\deg(f) \geqslant 4$ and use the opposition-complex description in [GV25, Section 3.2].[4] The fourth and fifth items follow from Fact A.3, which we prove below. We omit the other standard verifications.

For size selection, vary $s = \deg(f) \geqslant 4$. Each vertex stabilizer has order $q^{\binom{d+1}{2}}$, so the number of vertices is

$$N_s = (d+1)\frac{|\mathrm{SL}_{d+1}(\mathbb{F}_{q^s})|}{q^{\binom{d+1}{2}}}.$$

The standard order formula for SL_{d+1} gives $N_{s+1}/N_s \leqslant 2q^{(d+1)^2-1} = q^{O(d^2)}$. Choose the least $s \geqslant 4$ for which $N_s \geqslant n$. For sufficiently large n, we have $N_4 < n$, since d is fixed and q is polylogarithmic in n. Minimality then gives $n \leqslant N_s \leqslant nq^{O(d^2)}$. The resulting complex can be constructed in polynomial time by [Dik+26b].

Note that the test in the last item differs slightly from the V-test, since the sampled faces may intersect outside b. However, this occurs with probability at most $O(k^2/d)$. Taking d sufficiently large ensures that acceptance probability ε in our test implies acceptance probability at least $\varepsilon/2$ in the V-test, so the agreement theorem applies. □

Proof sketch of Fact A.3. By [GV25; Dik+26b], the two-color graphs in vertex links are symmetric induced subgraphs of the corresponding subspace-incidence graphs, retaining a $1 - O(1/q)$ fraction of the vertices and neighbors. Symmetry gives uniform edge weights. The full incidence graphs have second singular value at most $1/\sqrt{q}$ and diameter $O(d/(b-a))$. For sufficiently large q relative to d, the standard comparisons for these dense symmetric restrictions increase both bounds by at most a factor of two.

The spectral bound also holds in every proper link, since these are joins of smaller opposition complexes. Finally, partite trickle-down [Baf+25, Lemma 2.10] gives the bound $4/\sqrt{q}$ for X itself. □

[4] Before taking the congruence quotient, the complex is a full opposition subcomplex of a building and hence a clique complex. Products of three local subgroup matrices have entries of degree at most three, so reduction modulo f creates no new triangles when $\deg(f) \geqslant 4$. Every triangle in the quotient therefore spans a face. Since its vertex links are clique complexes, the same holds for every larger clique.

For these parameters, the number of top-dimensional faces is at most

$$|X(d)| \leqslant |X(0)| q^{O(d^2)} \leqslant n q^{O(d^2)}.$$

In particular, when $q = \Theta((\log n)^C)$ for fixed C, this is at most $n(\log n)^{\exp(\mathrm{poly}(1/\varepsilon))}$, for sufficiently large n. This gives the quantitative complex input for the [Baf+25] construction; the remaining parameter calculation follows their embedding, soundness amplification and alphabet reduction.

A.2 The PCP construction and its parameters

Proof of Theorem 1.3. We follow the construction of [Baf+25, Section 8], replacing their Theorem 2.13 by Theorem A.4. We explain the resulting parameter improvement and the minor adjustment required in the embedding step.

As in [Baf+25, Lemma 8.8 and Theorem 8.9], it suffices to start with an agreement test of soundness δ^2, where $\delta = \varepsilon^{O(1)}$. Applying Theorem A.4, we obtain a family of complexes with parameters $k = \exp(\mathrm{poly}(1/\varepsilon))$, $d = \mathrm{poly}(k)$, and $q = \Theta((\log n)^C)$, for a sufficiently large absolute constant $C > 0$. In particular, we may assume that q is larger than the threshold required by Theorem A.4 by taking n sufficiently large.

We apply the embedding construction of [Baf+25, Lemma 8.6], starting with a constant-gap CSP. To accommodate the size of the chosen complex, we take sufficiently many disjoint copies of the initial CSP so that its number of vertices matches that of the auxiliary routing graph up to fewer than one copy, as in [Baf+25, Lemma 5.2]. This preserves a constant soundness gap. In each vertex link, we use the color ordering from Fact A.3 when applying the local routing argument; the remaining robust routing argument is unchanged. We then apply [Baf+25, Lemma 6.2] to obtain a generalized label-cover instance with perfect completeness and soundness at most δ.

The quantitative difference is in the size of this instance. By Theorem A.4, the chosen complex has at most $n(\log n)^{\exp(\mathrm{poly}(1/\varepsilon))}$ vertices, and every vertex belongs to at most $q^{O(d^2)}$ top-dimensional faces. Consequently, the number of top-dimensional faces, and hence the size of the resulting generalized label-cover instance, is at most $n(\log n)^E$, where $E = \exp(O(\varepsilon^{-a}))$ for a sufficiently large absolute constant $a > 0$. Its logarithmic alphabet sizes and decision complexity are also bounded by $(\log n)^E$. Passing to a robust PCP using [Baf+25, Lemma 7.3], as in [Baf+25, Lemma 8.8], gives proof length at most $n(\log n)^E$, randomness complexity at most $\log_2 n + E \log_2 \log n$, and decision complexity at most $(\log n)^E$, after increasing the absolute constants in E.

The remaining construction is exactly that of [Baf+25, Lemma 8.8 and Theorem 8.9]: two compositions with the Reed–Muller based decodable PCP, followed by the Hadamard based composition and the conversion back to label cover. Their parameter analysis shows that these operations preserve the proof-length bound $n(\log n)^{O(E)}$, while achieving perfect completeness and soundness at most ε. The final robust PCP has alphabet size $(1/\varepsilon)^{O(1)}$ and query complexity $(1/\varepsilon)^{O(\log(1/\varepsilon))}$. Thus the conversion to a two-query PCP gives a proof alphabet of size at most $\exp((1/\varepsilon)^{O(\log(1/\varepsilon))})$, as required. Increasing the absolute constant in the definition of E completes the proof. □

B Equivalence between cones and bounding the Dehn function

In this subsection we establish a quantitative connection between the geometry of a coset complex and the presentation of its underlying group. More precisely, we show that the diameter of a coset complex is comparable to the word diameter of the group, and that filling loops in the complex is equivalent, up to a linear distortion, to reducing words using the multiplication tables of its subgroups. In particular, in the end of this section we prove Theorem 3.19.

The implication from the Dehn function to fillings of the coset complex was established in a different quantitative form in [KO21, Theorem 6.6]. We give a proof of both directions. Our proof is based on a general comparison between a simplicial complex and the nerve of a cover by subcomplexes.

For a loop P in a simplicial complex Z, let $\mathrm{Area}_Z(P)$ be the minimum number of $\overset{(TR)}{\sim}$ moves needed to contract P, where $\overset{(BT)}{\sim}$ moves are free, as in Subsection 2.6. We set $\mathrm{Area}_Z(P) = \infty$ if P is not contractible,

and write

$$\mathrm{Area}_Z(t) = \sup\{\mathrm{Area}_Z(P) : P \text{ is a loop of length at most } t\}.$$

In other words, we count the triangles used in a filling, with backtracking canceled freely.

Definition B.1 (Nerve)**.** Let $\mathcal{Y}$ be an indexed family of nonempty subcomplexes of Z that covers every vertex and every edge. Its *nerve* $\mathcal{A}$ has one vertex for each member of $\mathcal{Y}$, and $\{Y_0, \dots, Y_j\}$ is a face if these subcomplexes have a common vertex. We retain the indices: two members of the family may be equal as subcomplexes and still give distinct vertices of $\mathcal{A}$.

The application we have in mind is a coset complex $Y = \mathcal{CC}(G; (H_i)_{i\in I})$. Define an auxiliary simplicial complex Z whose vertices are the elements of G, and whose nonempty faces are the finite subsets contained in a single coset gH_i. Its 1-skeleton is the Cayley graph of G with generating set $\bigcup_{i\in I} H_i \setminus \{1\}$. Note that this is not necessarily a clique complex.

The simplices on the cosets gH_i, indexed by both the coset and its color, cover Z, and their nerve is Y. Conversely, cover Y by the chambers

$$Y_g = \{gH_i : i \in I\}, \qquad g \in G,$$

viewed as simplices and indexed by g. Their nerve is Z: the chambers $Y_{g_0}, \dots, Y_{g_j}$ have a common vertex if and only if $g_0, \dots, g_j$ belong to a common coset of one of the H_i. In both covers, every member and every nonempty intersection is a simplex.

We first establish two general lemmas about nerves, and then apply them to these covers.

Lemma B.2. *Let Z be a connected simplicial complex, and let $\mathcal{Y}$ be an indexed cover of its vertices and edges by nonempty connected subcomplexes of diameter at most D_1, where $D_1 \geqslant 1$. If $\mathcal{A}$ is the nerve, then $\mathcal{A}$ is connected and*

$$\mathrm{diam}(\mathcal{A}) \leqslant \mathrm{diam}(Z) + 1, \qquad \mathrm{diam}(Z) \leqslant D_1(\mathrm{diam}(\mathcal{A}) + 1).$$

Proof. Let $Y, Y' \in \mathcal{Y}$, choose $u \in Y$ and $v \in Y'$, and take a shortest path from u to v in Z. Choose a member of $\mathcal{Y}$ containing each edge of the path. Consecutive members intersect, the first intersects Y, and the last intersects Y'. Thus Y, Y' are joined in $\mathcal{A}$ by a path of length at most $\mathrm{dist}_Z(u, v) + 1$. If the chosen path has length zero, Y and Y' already intersect. This proves connectivity and the first bound.

Conversely, let $(Y_0; \dots; Y_m)$ be a path in $\mathcal{A}$, and choose any vertices $u \in Y_0$ and $v \in Y_m$. Choose one vertex in every consecutive intersection $Y_{j-1} \cap Y_j$. Connecting these vertices, together with u and v, by paths inside the respective Y_j gives a path in Z of length at most $(m+1)D_1$. Taking $m \leqslant \mathrm{diam}(\mathcal{A})$ proves the second bound. □

Lemma B.3. *Let Z be connected, and let $\mathcal{Y}$ be an indexed cover of its vertices and edges by nonempty connected subcomplexes. Suppose that every $Y \in \mathcal{Y}$ has diameter at most D_1, and that every nonempty intersection $Y \cap Y'$ is connected and has diameter at most D_2, where $D_1, D_2 \geqslant 1$. Suppose further that every loop of length at most $2D_1 + 2D_2 + 1$ in any $Y \in \mathcal{Y}$ can be contracted inside Y using at most R triangle moves. If $\mathcal{A}$ is the nerve, then for every $t \geqslant 1$,*

$$\mathrm{Area}_Z(t) \leqslant 3R\big(\mathrm{Area}_{\mathcal{A}}(t) + t\big).$$

Proof. For each $Y \in \mathcal{Y}$, choose a vertex v_Y and, for every $u \in Y$, a path $P_{Y,u}$ from v_Y to u inside Y of length at most D_1. For every edge $\{Y, Y'\}$ of $\mathcal{A}$, choose $u_{Y,Y'} = u_{Y',Y} \in Y \cap Y'$. We lift the oriented edge (Y, Y') to the path

$$L_{Y,Y'} = P_{Y,u_{Y,Y'}} \circ P^{-1}_{Y',u_{Y,Y'}}$$

from v_Y to $v_{Y'}$. In particular, $L_{Y',Y} = L^{-1}_{Y,Y'}$. For a path $P' = (Y_0; Y_1; \dots; Y_m)$ in $\mathcal{A}$, its *canonical lift* is the concatenation

$$L(P') = L_{Y_0,Y_1} \circ \cdots \circ L_{Y_{m-1},Y_m}.$$

When P' is a loop, so is $L(P')$.

We will prove that every loop P of length n in Z can first be transformed, using at most $3Rn$ triangle moves, into a canonical lift of a loop P' in $\mathcal{A}$ of length at most n, up to conjugation by a path. We will then

show that each elementary move in a contraction of P' can be performed on its canonical lift using at most $3R$ triangle moves. Importantly, we only count triangle moves: the lengths of intermediate paths need not remain bounded by n.

Step 1: Canonicalizing loops. Let $P = (u_0; u_1; \dots; u_n = u_0)$ be a loop in Z. For each edge (u_i, u_{i+1}), choose $Y_i \in \mathcal{Y}$ containing it, with indices interpreted cyclically. Since $u_i \in Y_{i-1} \cap Y_i$, the cyclic sequence $(Y_0; Y_1; \dots; Y_{n-1}; Y_0)$ is a closed walk in $\mathcal{A}$, possibly with consecutive repetitions. Removing these repetitions gives a loop P' of length at most n.

We first replace each edge (u_i, u_{i+1}) of P by

$$P^{-1}_{Y_i,u_i} \circ P_{Y_i,u_{i+1}}.$$

The original edge and its replacement form a loop inside Y_i of length at most $2D_1 + 1$. By assumption, this replacement can be performed using at most R triangle moves, so the total cost is at most Rn. For an illustration of this part see Figure 3.

After these replacements, the path travels through the base vertices of the sets Y_i. More precisely, define

$$T_i = P_{Y_{i-1},u_i} \circ P^{-1}_{Y_i,u_i},$$

which is a path from $v_{Y_{i-1}}$ to v_{Y_i}. Up to backtracking, the resulting loop is

$$C \circ T_1 \circ T_2 \circ \cdots \circ T_{n-1} \circ T_0 \circ C^{-1}, \qquad C = P^{-1}_{Y_0,u_0}.$$

Thus, it remains to replace each transition T_i by the canonical path L_{Y_{i-1},Y_i}.

If $Y_{i-1} = Y_i$, then T_i cancels by backtracking. Otherwise, let $w = u_{Y_{i-1},Y_i}$. Since $u_i, w \in Y_{i-1} \cap Y_i$, we may choose a path Q from u_i to w inside this intersection, of length at most D_2. Inserting the backtracking $Q \circ Q^{-1}$, we write

$$T_i = P_{Y_{i-1},u_i} \circ Q \circ Q^{-1} \circ P^{-1}_{Y_i,u_i}.$$

We now perform two contractions. First, inside Y_{i-1}, replace $P_{Y_{i-1},u_i} \circ Q$ by $P_{Y_{i-1},w}$. Second, inside Y_i, replace $Q^{-1} \circ P^{-1}_{Y_i,u_i}$ by $P^{-1}_{Y_i,w}$. Each replacement is justified by a loop of length at most $2D_1 + D_2$ inside the corresponding covering set, and therefore costs at most R triangle moves. The resulting path is precisely

$$P_{Y_{i-1},w} \circ P^{-1}_{Y_i,w} = L_{Y_{i-1},Y_i}.$$

There are at most n transitions, so this second part costs at most $2Rn$ triangle moves. We have therefore transformed P into

$$C \circ L(P') \circ C^{-1}$$

using at most $3Rn$ triangle moves. Conjugation by C will not affect the number of triangle moves needed for the remaining contraction. For an illustration of this part see Figure 4.

Step 2: Contracting canonical lifts. Consider a contraction of P' in $\mathcal{A}$ using a triangle moves, together with free backtracking moves. We show that each of these moves can be performed on the corresponding canonical lifts with at most $3R$ triangle moves in Z. This applies equally to moves that remove a triangle or backtrack and to their inverses.

A backtracking move inserts or removes two consecutive edges (Y, Y') and (Y', Y). Their canonical lifts are $L_{Y,Y'}$ and $L_{Y',Y} = L^{-1}_{Y,Y'}$, so the corresponding modification in Z consists only of backtracking moves and costs nothing.

Now consider a triangle move that replaces the path $(Y_0; Y_1; Y_2)$ by the edge $(Y_0; Y_2)$, where $\{Y_0, Y_1, Y_2\}$ is a triangle of $\mathcal{A}$. We need to replace $L_{Y_0,Y_1} \circ L_{Y_1,Y_2}$ by L_{Y_0,Y_2}, or equivalently contract

$$L_{Y_0,Y_1} \circ L_{Y_1,Y_2} \circ L_{Y_2,Y_0}.$$

For brevity, write $u_{ij} = u_{Y_i,Y_j}$ and $P_{i,u} = P_{Y_i,u}$. Since the three covering sets form a triangle in the nerve, there is a vertex $u \in Y_0 \cap Y_1 \cap Y_2$. For each unordered pair $\{i, j\}$, choose a path Q_{ij} from u_{ij} to u inside $Y_i \cap Y_j$, of length at most D_2.

Expanding the definitions of the canonical lifts, the loop we wish to contract is

$$P_{0,u_{01}} \circ \left(P^{-1}_{1,u_{01}} \circ P_{1,u_{12}}\right) \circ \left(P^{-1}_{2,u_{12}} \circ P_{2,u_{20}}\right) \circ P^{-1}_{0,u_{20}}.$$

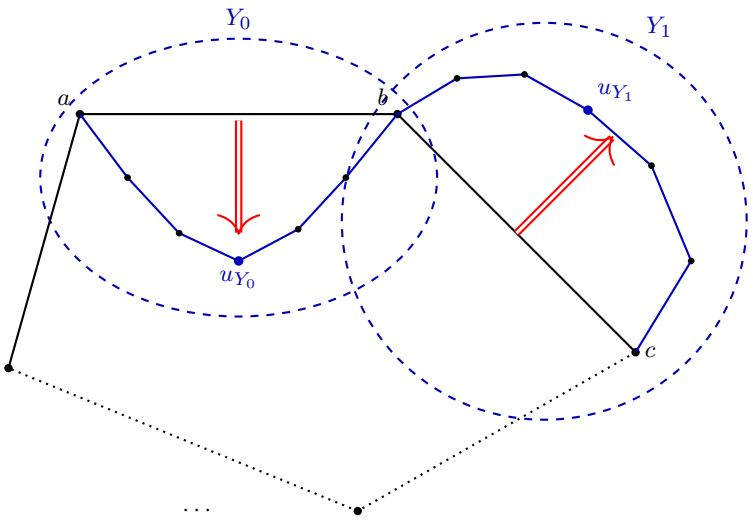


Figure 3: Replacing the edges (a, b) and (b, c) by paths through the chosen base vertices of Y_0 and Y_1, respectively.

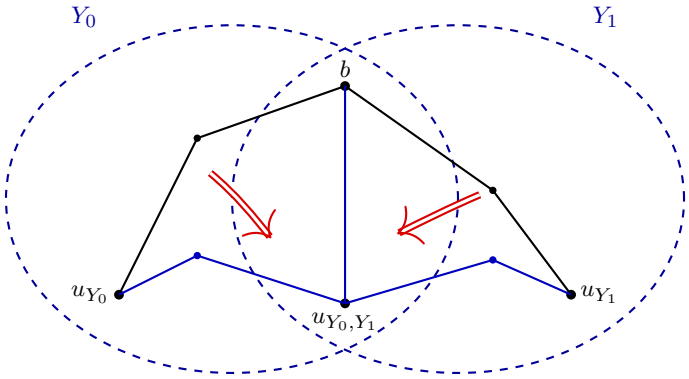


Figure 4: Second part of of the canonicalization. The original transition path from u_{Y_0} to u_{Y_1} passes through b. We replace its left part by a path from u_{Y_0} to b through the chosen intersection vertex u_{Y_0,Y_1} inside Y_0, and similarly replace its right part inside Y_1. The two occurrences of the segment between b and u_{Y_0,Y_1} are then removed by backtracking moves.

First, inside Y_1, replace

$$P_{1,u_{01}}^{-1} \circ P_{1,u_{12}} \quad \text{by} \quad Q_{01} \circ Q_{12}^{-1}.$$

Both paths connect u_{01} to u_{12} inside Y_1, and together form a loop of length at most $2D_1 + 2D_2$. Thus, the replacement costs at most R triangle moves.

Similarly, inside Y_2, replace

$$P_{2,u_{12}}^{-1} \circ P_{2,u_{20}} \quad \text{by} \quad Q_{12} \circ Q_{20}^{-1},$$

at a cost of at most R triangle moves. After these replacements, our original loop has become

$$P_{0,u_{01}} \circ Q_{01} \circ Q_{12}^{-1} \circ Q_{12} \circ Q_{20}^{-1} \circ P_{0,u_{20}}^{-1}.$$

Canceling the backtracking $Q_{12}^{-1} \circ Q_{12}$ leaves

$$P_{0,u_{01}} \circ Q_{01} \circ Q_{20}^{-1} \circ P_{0,u_{20}}^{-1}.$$

This is a loop inside Y_0 of length at most $2D_1 + 2D_2$, so it contracts using at most R additional triangle moves. See Figure 5 for an illustration.

We have contracted the lifted boundary of the triangle using at most $3R$ triangle moves. Reversing these contractions also allows us to perform the reverse triangle move at the same cost.

Consequently, every elementary move in the contraction of P' can be simulated on its canonical lift, and the entire contraction of $L(P')$ requires at most $3Ra$ triangle moves. This bound does not depend on the lengths of the intermediate paths in the contraction.

Combining the two steps, and taking $a \leqslant \mathrm{Area}_{\mathcal{A}}(n)$, gives

$$\mathrm{Area}_Z(P) \leqslant 3Rn + 3R\,\mathrm{Area}_{\mathcal{A}}(n).$$

Taking the supremum over all loops P of length at most t proves

$$\mathrm{Area}_Z(t) \leqslant 3R\big(\mathrm{Area}_{\mathcal{A}}(t) + t\big),$$

as required.

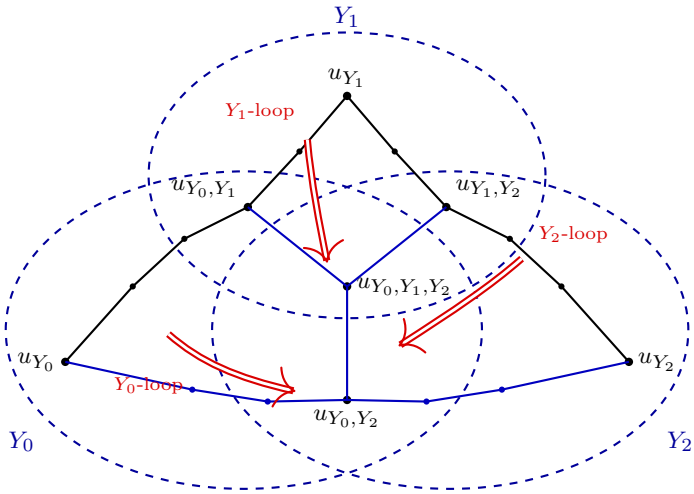


Figure 5: The triangle move on canonical lifts. The original canonical path from u_{Y_0} to u_{Y_2} passes through u_{Y_0,Y_1}, u_{Y_1}, and u_{Y_1,Y_2}. We connect each pairwise intersection point to the triple-intersection point u_{Y_0,Y_1,Y_2}. This yields three loops, contained respectively in Y_1, Y_2, and Y_0, which are contracted separately. The result is the canonical path from u_{Y_0} to u_{Y_2} through u_{Y_0,Y_2}.

□

We apply the two lemmas to the covers of the coset complex Y and the auxiliary complex Z described above.

Corollary B.4. *Let G be a finite group with subgroups H_i, $i \in I$, such that $Y = \mathcal{CC}(G; (H_i)_{i\in I})$ is connected. Let L be the maximum, over $g \in G$, of the minimum number of letters from $\bigcup_i H_i$ whose product is g. Then*

$$L - 1 \leqslant \mathrm{diam}(Y) \leqslant L + 1.$$

In particular, $\mathrm{diam}(Y) \leqslant D$ *implies that every $g \in G$ is a product of at most $D+1$ elements from the subgroups.*

Proof. The Cayley graph underlying Z has diameter L. The cover of Z by its coset simplices has nerve Y, and each member has diameter at most 1. Apply Lemma B.2 with $D_1 = 1$. □

We next compare fillings of loops in Z with the Dehn function associated with Y. We use the definition of the Dehn function given in Subsection 2.6, which counts applications of relators in the free group $\mathrm{Free}(S)$, where $S = \bigcup_i H_i \setminus \{1\}$.

There is a minor distinction between words in $\mathrm{Free}(S)$ and paths in Z. For $g \in S$, write $[g]$ for the corresponding generator of $\mathrm{Free}(S)$. The formal inverse $[g]^{-1}$ and the generator $[g^{-1}]$ indexed by the inverse

group element are different symbols in $\mathrm{Free}(S)$. However, both represent the same element of G. In particular, the complete multiplication table of any subgroup containing g includes the relation

$$[g][g^{-1}] = 1.$$

Thus, a geometric backtracking move in Z need not correspond to a free reduction in $\mathrm{Free}(S)$, but can be realized using one multiplication-table relation. We keep track of this distinction in the following lemma.

Lemma B.5. *Let $Y = \mathcal{CC}(G; (H_i)_{i\in I})$ be 1-connected, and let Z be the auxiliary complex defined above. With the complete multiplication-table presentation of G and the Dehn function Dehn_Y defined in the preliminaries, for every $t \geqslant 1$,*

$$\mathrm{Area}_Z(t) \leqslant \mathrm{Dehn}_Y(t) \leqslant 3\,\mathrm{Area}_Z(t) + 2t.$$

Proof. An oriented edge (x, xg) of Z is labeled by an element $g \in S$. Thus a word in the alphabet $S \cup S^{-1}$ whose product is the identity defines a loop in Z, using the corresponding elements of G as edge increments. A free reduction $[g][g]^{-1}$ becomes a backtracking move, while an application of a multiplication-table relation is realized by at most one triangle move in a coset simplex. In particular, a two-letter inverse relation is realized by a backtracking move and requires no triangle moves. Therefore, a reduction of a word gives a contraction of the corresponding loop in Z with no more triangle moves than relation applications. Taking the positive edge-label word of an arbitrary loop in Z proves $\mathrm{Area}_Z(t) \leqslant \mathrm{Dehn}_Y(t)$.

For the other inequality, consider a word $w \in \mathrm{Free}(S)$ of length $n \leqslant t$ representing the identity in G. First, replace each occurrence of a formal inverse $[g]^{-1}$ by the positive generator $[g^{-1}]$. This replacement uses one application of the relation $[g][g^{-1}] = 1$. Thus, after at most n relation applications, we obtain a positive word w' of length n whose letters are the edge labels of a loop P in Z.

Let $a = \mathrm{Area}_Z(P)$. A contraction of P with a triangle moves gives a disk diagram consisting of a triangles, possibly together with backtracking edges. We may remove redundant backtracking branches, so that the remaining diagram consists of the a triangles and the edges needed to connect them to the original boundary. Each triangle contributes three edge occurrences, while the original boundary contributes n. Consequently, the diagram requires at most $(3a + n)/2$ identifications of oppositely oriented edge occurrences.

Each triangle has all its vertices in a common coset xH_i, so its oriented boundary labels form a three-letter multiplication-table relation in H_i. Thus the triangles contribute at most a relation applications. Along an identification of oppositely oriented edges, the positive labels are $[g]$ and $[g^{-1}]$, which can be canceled using the relation $[g][g^{-1}] = 1$. Hence these identifications contribute at most $(3a + n)/2$ additional relation applications.

Including the initial replacement of formal inverses, we obtain

$$\mathrm{Area}_R(w) \leqslant n + a + \frac{3a + n}{2} \leqslant 3a + 2n.$$

Taking the maximum over all words of length at most t proves

$$\mathrm{Dehn}_Y(t) \leqslant 3\,\mathrm{Area}_Z(t) + 2t.$$

□

Corollary B.6. *Let $Y = \mathcal{CC}(G; (H_i)_{i\in I})$ be a 1-connected coset complex with $|I| \geqslant 3$, and let Dehn_Y be as in Definition 2.35. Then, for every $t \geqslant 1$,*

$$\mathrm{Area}_Y(t) \leqslant 15\big(\mathrm{Dehn}_Y(t) + t\big),$$

and

$$\mathrm{Dehn}_Y(t) \leqslant 45\,\mathrm{Area}_Y(t) + 47t.$$

Proof. In both covers of Y and Z, all members and their nonempty intersections are simplices, so we may take $D_1 = D_2 = 1$. Every loop of length at most 5 in a simplex contracts with at most 5 triangle moves. Applying Lemma B.3 to the two covers gives

$$\mathrm{Area}_Y(t) \leqslant 15\big(\mathrm{Area}_Z(t) + t\big), \qquad \mathrm{Area}_Z(t) \leqslant 15\big(\mathrm{Area}_Y(t) + t\big).$$

The first inequality of the corollary follows from $\mathrm{Area}_Z(t) \leqslant \mathrm{Dehn}_Y(t)$. For the second, we have

$$\mathrm{Dehn}_Y(t) \leqslant 3\,\mathrm{Area}_Z(t) + 2t \leqslant 45\,\mathrm{Area}_Y(t) + 47t,$$

as required. □

We conclude by translating the area bounds into bounds on cone radii, and conversely translating cone radii into bounds on the Dehn function. This forward direction of this theorem is the statement in Theorem 3.19.

Theorem B.7. *Let G be a finite group with subgroups H_i, $i \in I$, where $|I| \geqslant 3$, and assume that the coset complex $Y = \mathcal{CC}(G; (H_i)_{i\in I})$ is 1-connected. Write $D = \mathrm{diam}(Y)$. Then Y admits a 1-cone Cone_Y satisfying*

$$\mathrm{Rad}_0(\mathrm{Cone}_Y) \leqslant D, \qquad \mathrm{Rad}_1(\mathrm{Cone}_Y) \leqslant 15\big(\mathrm{Dehn}_Y(2D+1) + 2D + 1\big).$$

Conversely, if Y admits a 1-cone Cone_Y with $\mathrm{Rad}_1(\mathrm{Cone}_Y) \leqslant A$, then for every $t \geqslant 1$,

$$\mathrm{Dehn}_Y(t) \leqslant (45A + 47)t.$$

Proof. Choose an apex v_0 and shortest paths P_u from v_0 to every vertex u, taking P_{v_0} to be the trivial path. Then $\mathrm{Rad}_0(\mathrm{Cone}_Y) \leqslant D$. For each ordered edge (u, v), the loop $P_u \circ (u; v) \circ P_v^{-1}$ has length at most $2D + 1$. By Corollary B.6, we can choose a contraction T_{uv} with at most $15(\mathrm{Dehn}_Y(2D + 1) + 2D + 1)$ triangle moves. These paths and contractions constitute a 1-cone in the sense of Definition 2.27 and prove the first implication.

Conversely, let $P = (u_0; u_1; \ldots; u_m = u_0)$ be any loop of length $m \leqslant t$. Insert the paths P_{u_i} and their inverses between consecutive edges. Using only backtracking moves, we obtain a conjugate of the concatenation of the m loops $P_{u_i} \circ (u_i; u_{i+1}) \circ P_{u_{i+1}}^{-1}$. The cone contracts each of these using at most A triangle moves, so $\mathrm{Area}_Y(P) \leqslant mA$ and hence $\mathrm{Area}_Y(t) \leqslant At$. Applying Corollary B.6 gives

$$\mathrm{Dehn}_Y(t) \leqslant 45\,\mathrm{Area}_Y(t) + 47t \leqslant (45A + 47)t,$$

as required. □