Electronic Colloquium on Computational Complexity, Report No. 180 (2025)

Low-soundness direct-product testers and PCPs from Kaufman–Oppenheim complexes

Ryan O'Donnell*® Noah G. Singer*®

November 13, 2025

Abstract

We study the Kaufman–Oppenheim coset complexes [KO23], which have an elementary and strongly explicit description. Answering an open question of [KOW25], we show that they support sparse direct-product testers in the low soundness regime. Our proof relies on the HDX characterization of agreement testing by [BM24; DD24a], the recent result of [KOW25], and follows techniques from [BLM24; DDL24]. Ultimately, the task reduces to showing dimension-independent coboundary expansion of certain 2-dimensional subcomplexes of the KO complex; following the "Dehn method" of [KO21], we do this by establishing efficient presentation bounds for certain matrix groups over polynomial rings.

As shown in [BMVY25], a consequence of our direct-product testing result is that the Kaufman–Oppenheim complexes can also be used to obtain PCPs with arbitrarily small constant soundness and quasilinear length. Thus the use of sophisticated number theory and algebraic group-theoretic tools in the construction of these PCPs can be avoided.

1 Introduction

The main result of our work is an elementary, strongly explicit, sparse, two-query direct-product tester, based on the high-dimensional expanders of Kaufman, Oppenheim, and Weinberger [KO23; KOW25]:

Theorem 1.1 (Direct-product testing application theorem). *Given $\delta > 0$, provided $k \geq \exp(\mathrm{poly}(1/\delta))$, $n \geq \exp(\mathrm{poly}(k))$, and $p \geq \mathrm{poly}(n)$ is prime, the [KOW25] complexes $\{\mathfrak{A}_n(\mathbb{F}_Q) : Q = p^s,\ s > 3n\}$ yield a strongly explicit family of k-uniform hypergraphs with $N \approx Q^{(n+1)^2-1}$ vertices and $p^{O(n^2)} \cdot N$ hyperedges, such that performing the standard "V-test" with them gives a δ-sound direct-product tester.*

Previously, Bafna–Lifshitz–Minzer [BLM24] and Dikstein–Dinur–Lubotzky [DDL24] independently showed a similar result for complexes constructed as quotients of affine Bruhat–Tits buildings associated to the symplectic group over the p-adics. These complexes require deep number theory and algebraic group theory to construct and analyze, and at a technical level are not known to be strongly explicit. By contrast, in §2.1.1 we will give a completely concrete description of the direct-product tester described in Theorem 1.1; moreover, its analysis, though lengthy, can arguably be described as elementary.

As shown by Bafna, Minzer, and Vyas [BMVY25], direct-product testers as in [BLM24; DDL24] and Theorem 1.1 can be successfully put to use for their original purpose: highly efficient PCPs. More precisely, [BMVY25] constructs from them an efficient PCP system for NP with the shortest known proof lengths, $n \mapsto n \cdot \mathrm{polylog}\, n$, for any small constant soundness $\delta > 0$. Our Theorem 1.1 provides an alternative ingredient that allows their entire construction to be elementary. It is our hope that the improved simplicity of this construction will help lead to quantitative improvements in the future.[1]

*Computer Science Department, Carnegie Mellon University. {odonnell,ngsinger}@cs.cmu.edu.
® Author order randomized.

[1] For example, we are hopeful that Theorem 1.1's strong explicitness property will lead to improved prover/verifier *time-efficiency* in the PCP, but we leave this for future work.

ISSN 1433-8092

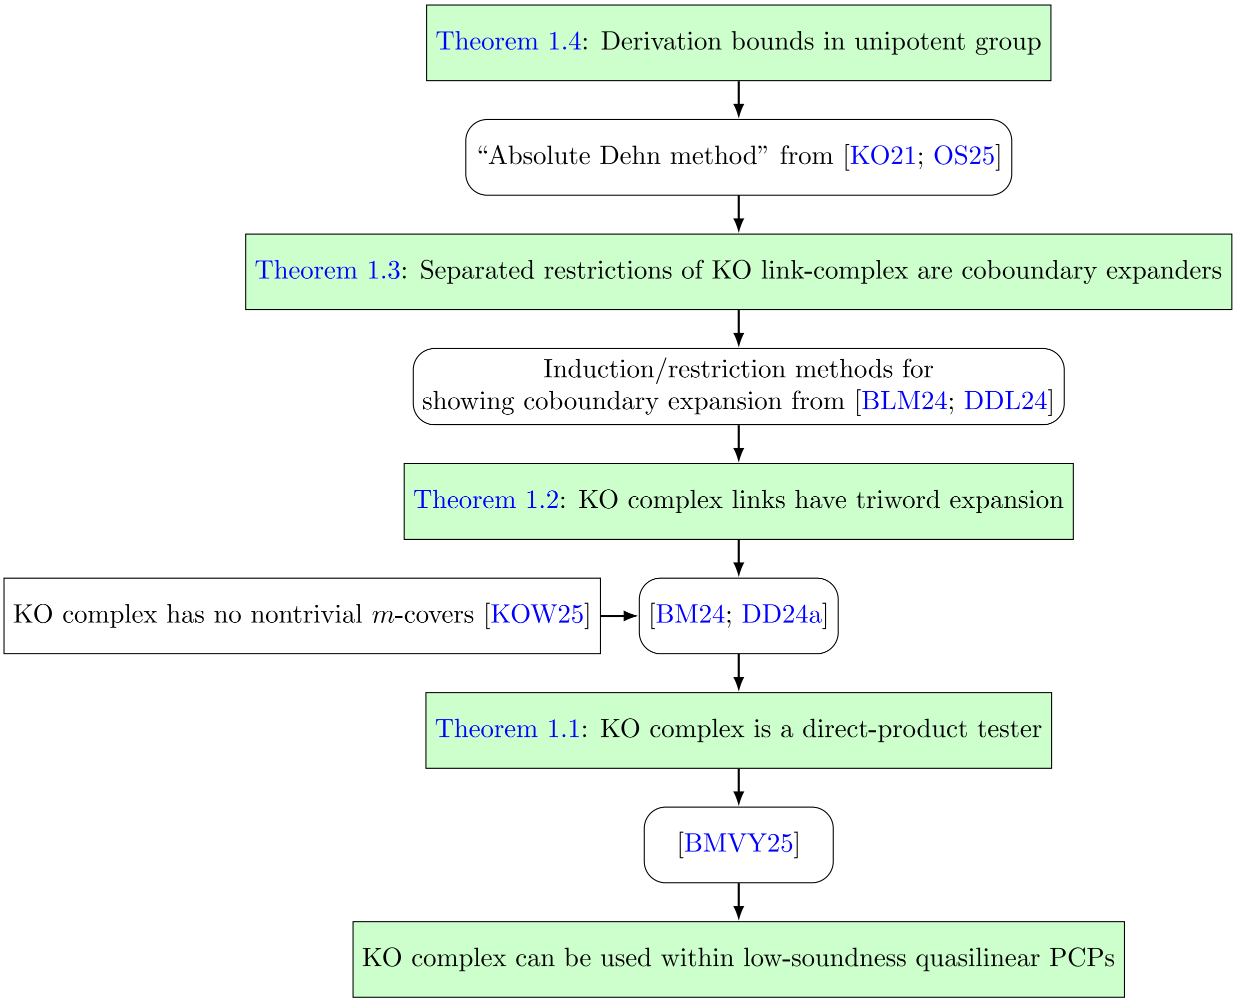

Figure 1: A schematic diagram of our proof; the contributions of our work are shaded. Rounded rectangles are "reduction" theorems.

To establish Theorem 1.1, we use a characterization of the precise high-dimensional expansion (HDX) properties needed to obtain a (two-query, low-soundness) direct-product tester from a simplicial complex. This characterization was determined by Bafna–Minzer [BM24] and Dikstein–Dinur [DD24a] following a sequence of important works including [DM22; GK23; DD24b]. Assuming the k-uniform hypergraph tester is embedded inside an n-dimensional complex $\mathfrak{X}$ (for $n \gg k$) with excellent spectral expansion, roughly speaking the following two properties of $\mathfrak{X}$ are needed (see §2.2 for an explanation of terms):

1. The global hypergraph $\mathfrak{X}$ should have a qualitative connectedness property: no nontrivial m-covers for $m \le \mathrm{poly}(1/\delta)$ (aka "trivial 1-cohomology over the symmetric group $\mathfrak{S}(m)$").
2. The local neighborhood ("link") of each vertex of $\mathfrak{X}$ should have a quantitatively strong form of topological expansion vis-à-vis its size-r sets; in this paper's terminology, it should have "r-triword expansion at least $\exp(-O(r^{.99}))$".[2]

Property (1) was recently shown for the Kaufman–Oppenheim complexes by Kaufman, Oppenheim, and Weinberger [KOW25]. Thus to establish Theorem 1.1, it suffices for us to establish Property (2); i.e., to prove the following:

[2]Two remarks: First, as in Property (1), one only needs this expansion over $\mathfrak{S}(m)$ for $m = \mathrm{poly}(1/\delta)$; however, we get it for all m. Second, "triword expansion" is our own terminology. In [BM24], the condition would be called "$(m, r, \xi, \exp(-O(r^{.99}))\cdot\xi)$-UG-coboundary expansion", and in [DD24a] it would be called "$(\exp(-O(r^{.99})), r)$-swap-coboundary expansion over $\mathfrak{S}(m)$".

Theorem 1.2 (KO link coboundary expansion)**.** *Given $r \in \mathbb{N}^+$, provided $n \gg r^2$ and $p \gg r^2n^2$ is prime, the common vertex-link of the [KOW25] complexes, $\mathfrak{LA}_n(\mathbb{F}_p)$, has r-triword expansion at least* $\exp(-O(r^{.67}))$ *(over every group Γ).*

Before discussing how we prove Theorem 1.2, we recap the prior analogous work. Dikstein and Dinur [DD24c, Theorem 1.4] showed that the links of the "usual" (non-symplectic) [LSV05] building quotient complexes satisfy Property (2) above, but no form of them is known to satisfy Property (1). However Chapman and Lubotzky [CL25] subsequently showed that more exotic, symplectic forms of these HDXs *do* satisfy Property (1) over $\mathfrak{S}(2)$. This paved the way for Bafna–Lifshitz–Minzer [BLM24] and Dikstein–Dinur–Lubotzky [DDL24] to show the same result over $\mathfrak{S}(m)$, thereby establishing the symplectic building quotients satisfy Property (1). Thus to derive low soundness direct-product testers, it remained for them to establish Property (2) for the links of the symplectic building quotients, which [BLM24, Theorem 5.3], [DDL24, Theorem 1.3] did by building on the work of [DD24c]. (We remark that it is the construction and analysis for Property (1) that uses deep number theory and algebraic group theory; [BLM24; DDL24]'s establishment of Property (2), though not easy, uses only elementary tools.)

The KO vertex-link complex $\mathfrak{LA}_n(\mathbb{F}_p)$ is n-partite. Moreover, the parts $P = \{1, \dots, n\}$ have a natural linear ordering; i.e., it is relevant to think of them as lying on an n-path graph and to consider the "distance" $|a-b|$ between parts a and b. Because of this, if w is a set of vertices in $\mathfrak{LA}_n(\mathbb{F}_p)$ from different parts $W \subseteq P$, and we study the "w-pinned subcomplex" (i.e., "w-link") $\mathfrak{LA}_n(\mathbb{F}_p)_w$, the expansion of this subcomplex is affected by the "geometry" of how W sits within the path P. Similarly, if we take, say, 3 parts $a < b < c$ from P and consider the subcomplex $\mathfrak{LA}_n(\mathbb{F}_p)[\{a\}, \{b\}, \{c\}]$ restricted to those parts, then the distances $b-a$ and $c-b$ play a role in the expansion of the subcomplex. (Similar considerations arise in the link-complexes studied in [DD24c; BLM24; DDL24].)

To prove Theorem 1.2, we follow an induction/restriction strategy on parts. Specifically, we follow the strategy from [BLM24] (with some streamlining); [DD24c; DDL24] uses similar-looking ideas, though the technical details seem to diverge. This ultimately lets us reduce Theorem 1.2 to the below theorem about 1-coboundary expansion within pinnings/restrictions of $\mathfrak{LA}_n(\mathbb{F}_p)$, where the three parts to which we restrict are at path-distance $\Omega(n)$.

Theorem 1.3 (Separated restrictions have dimension-independent coboundary expansion)**.** *In $\mathfrak{LA}_n(\mathbb{F})$, let $a < b < c$ be in $[1\,..\,n]$, and assume the "well-separation" condition $b - a, c - b \geq \xi n$. Let $W \subseteq [1\,..\,n] \setminus [a\,..\,c]$, and let w consist of one vertex per part in W. Then the pinned-and-restricted subcomplex $\mathfrak{LA}_n(\mathbb{F}_p)_w[\{a\}, \{b\}, \{c\}]$ has 1-coboundary expansion at least* $\mathrm{poly}(\xi)$ *(over any group Γ).*

We remark that Kaufman and Oppenheim [KO21] (essentially) proved this theorem in the particular case of $n = 3$ (and char $\mathbb{F} \neq 2$). The crucial aspect of our Theorem 1.3 is that $\xi \geq \Omega_{n\to\infty}(1)$ for "most" triples $a < b < c$ in $[1\,..\,n]$, and hence the 1-coboundary expansion is *independent of* n. (Even though n is ultimately a "constant" in our construction, to get Theorem 1.2 we require this quantitative independence from n.)

As in Kaufman and Oppenheim [KO21] (and [DD24c]; see also [Gro10; LMM16]), we establish the coboundary-expansion result Theorem 1.3 using the "cones method". As shown in [KO21], the strong symmetry of the KO complex simplifies what one needs to prove: thinking of $\mathfrak{T} \coloneqq \mathfrak{LA}_n(\mathbb{F}_p)_w[\{a\}, \{b\}, \{c\}]$ as a "triangle-complex" (3-uniform hypergraph), we only need to show:

- its *diameter* is small, $\mathrm{diam}(\mathfrak{T}) \leq C_0 = C_0(\xi)$;
- every loop of length at most $2C_0 + 1$ can be triangulated using at most $C = C(\xi)$ triangles (3-uniform hyperedges).

This is a quantitative form of saying that $\mathfrak{T}$ is "simply connected". It is known [Lan50; Bun52] that for *coset complexes* like $\mathfrak{T}$, there is a *group-theoretic* characterization of simple connectivity, and Kaufman and Oppenheim [KO21] (see also [OS25]) upgraded this to a quantitative form. Thus to show Theorem 1.3, it suffices to show certain quantitative statements about presentations within the group of unipotent[3] $n \times n$ matrices over $\mathbb{F}_p[t]$. To put it very briefly, given a sequence of generators that multiply to $\mathbb{1}$ in the unipotent group, we need to bound how many steps it takes to *prove* this, using a limited set of relations. Such

[3]Upper-triangular, with 1's on the diagonal.

bounds are often expressed using the notion of the "Dehn function", or "area bounds". Ultimately, the main technical theorems we prove in our paper, constituting the bulk of the work, are the following:

Theorem 1.4 (Unipotent group presentation bounds)**.** *For every $\xi > 0$ there exist $C_0, C \in \mathbb{N}$ such that for all sufficiently large n, the following hold: For every field $\mathbb{F}$, parameter $\kappa \in \mathbb{N}^+$ (defining generators), and $\ell_1 < \ell_2 < \ell_3 \in [1\,..\,n]$ with $\ell_2 - \ell_1, \ell_3 - \ell_2 \geq \xi n$,*

$$\operatorname{diam}(\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n); (\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell_i))_{i\in\{1,2\}}) \leq C_0, \tag{1.5}$$

$$\operatorname{area}^{2C_0+1}(\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n); (\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell_i))_{i\in\{1,2,3\}}) \leq C. \tag{1.6}$$

We explain the meaning of this problem in more detail, using the computer science language of *circuits*, in §2.3.

Organization of the rest of the paper

§2 is devoted to further high-level descriptions of the components of our work:

- In §2.1, we review the definition of low-soundness direct-product testers and concretely describe the one we obtain from the Kaufman–Oppenheim complex. We also describe the subcomplexes $\mathfrak{LA}_n(\mathbb{F}_p)$ and $\mathfrak{LA}_n(\mathbb{F}_p)_w[\{a\},\{b\},\{c\}]$ arising in Theorems 1.2 and 1.3.
- In §2.1, we review the definitions of vanishing 1-cohomology and r-triword expansion.
- In §2.3, we give a high-level overview of our group-theoretic proof of Theorem 1.4 and some of the considerations therein.

The remainder of the work is devoted to technical details. In §3, we setup additional necessary definitions regarding simplicial complexes, coset complexes, unipotent groups, and the (local) Kaufman–Oppenheim complex. In §4, we review some useful known facts and prove some (simple) additional ones about unipotent groups and their subgroups. In §5, we prove Theorem 1.2 modulo Theorem 1.4. In §6, we describe a few useful additional technical tools which we will use to prove Theorem 1.4. Theorem 1.4 is then proven in §7 (for Equation (1.5)) and §8 (for Equation (1.6)). Finally, in §9, we collect together all known theorems about the KO complexes, prove our direct-product testing theorem Theorem 1.1, and verify that the KO complexes can be used in the [BMVY25] PCP construction.

Acknowledgments

We would like to thank Mitali Bafna and Dor Minzer for their advice on applying [BLM24] to other complexes. The second author in particular thanks Mitali for discussions about details in [BLM24].

2 High-level description of components

2.1 Direct-product testing

"Direct-product testers" were introduced by Goldreich and Safra [GS00] as a kind of "derandomized parallel repetition" for use in PCPs. At a high level, one imagines that a "prover" claims to have a function $f : \mathcal{U} \to \Sigma$, mapping a large universe $\mathcal{U}$ of size N to some alphabet Σ. The tester wishes to simultaneously get k values $f(u_1), \ldots, f(u_k)$ with one "question" $\{u_1, \ldots, u_k\}$ to the prover — but is concerned that the prover might not honestly give the "direct-product" answer $f(u_1), \ldots, f(u_k)$.

A *two-query agreement test (or "V-test")* [DR06] attempts to solve this by introducing a second prover, who cooperates but cannot communicate with the first prover. The agreement tester is then defined by a k-uniform hypergraph $\mathfrak{X}$ on $\mathcal{U}$.[4] The tester randomly chooses two hyperedges $\boldsymbol{w}, \boldsymbol{w}'$ from $\mathfrak{X}$ conditioned on $|\boldsymbol{w} \cap \boldsymbol{w}'| = \sqrt{k}$, asks the first prover for f's values on $\boldsymbol{w}$, and asks the second prover for f's values on $\boldsymbol{w}'$. Finally, the tester verifies that the provers gave the same answers for the $\sqrt{k}$ elements in common. One says

[4]One can allow for a weighted k-uniform hypergraph, but unlike in [BLM24; DDL24] we will not need to.

the agreement tester defined by $\mathfrak{X}$ and k achieves δ-*soundness* if, whenever the test accepts with probability at least δ, there is an $f : \mathcal{U} \to \Sigma$ that is (nearly) consistent with the provers' answers for at least a $\delta^{O(1)}$ fraction of $\mathfrak{X}$'s hyperedges. Here is a precise definition (which merges the two provers into one function F):

Definition 2.1 (Agreement tester soundness)**.** Let $\mathfrak{X}$ be a k-uniform hypergraph on $\mathcal{U}$. We say the agreement tester based on $\mathfrak{X}$ achieves δ-*soundness* if the following holds: Whenever it accepts $F : \mathfrak{X} \ni w \mapsto \Sigma^w$ with probability least δ, there exists $f : \mathcal{U} \to \Sigma$ such that

$$\Pr_{\boldsymbol{w}\sim\mathfrak{X}}[\mathrm{dist}(F(\boldsymbol{w}), (f(u))_{u\in\boldsymbol{w}}) \leq \exp(-1/\delta^{\Omega(1)})] \geq \delta^{O(1)}, \tag{2.2}$$

where dist denotes fractional Hamming distance.[5] ◊

It is not easy even to show that the *complete* k-uniform hypergraph on $\mathcal{U}$ achieves δ-soundness, but this was done by Dinur and Goldenberg [DG08] for $k = \mathrm{poly}(1/\delta)$ (cf. [IKW12]). However, the fact that the complete hypergraph has size $\Theta(N^k)$ for $|\mathcal{U}| = N$ renders it quantitatively undesirable for applications to PCPs. Consequently, Dinur and Kaufman [DK17] sought *sparse* agreement testers, meaning ones with $O_\delta(N)$ hyperedges. They showed that the high-dimensional expanders (HDXs) of [LSV05] yield sparse agreement testers in the *high soundness* regime when δ is close to 1, and conjectured that HDXs could also be used to get sparse agreement testers in the *low soundness* regime of δ close to 0. As described in §1, this conjecture was ultimately resolved positively by Bafna–Lifshitz–Mizer and Dikstein–Dinur–Lubotzky [BLM24; DDL24], for certain (nontrivial to construct and analyze) quotients of affine Bruhat–Tits buildings associated to the symplectic group over the p-adics.

Subsequently, Bafna, Minzer, and Vyas [BMVY25] used these agreement testers as a key ingredient in their $(n\cdot\mathrm{polylog}\, n)$-length PCPs with arbitrarily small constant soundness $\delta > 0$.[6] Their work also essentially showed that *if* our main Theorem 1.1 is true, then the [KO23; KOW25] complexes can be substituted into their construction in place of the symplectic building quotients. (There are a few minor details to check, which we do in §9.) Thus our work shows that the elementary agreement tester which are described below can be used in their PCP construction.

2.1.1 An elementary, sparse, low-soundness agreement tester

In this section, we give a concrete description of the δ-sound agreement tester in our Theorem 1.1, arising from the [KO23; KOW25] complexes $\mathfrak{A}_n(\mathbb{F}_Q)$. We describe it as a strongly explicit family of sparse, k-uniform hypergraphs ($k = \mathrm{poly}(1/\delta)$) over universes $\mathcal{U}$ of growing size N.

Let δ, k, n, p be as in Theorem 1.1,[7] let $s > 3n$ be a growing parameter, and let $Q = p^s$. The universe $\mathcal{U}$ will be closely related to the set of $(n+1)\times(n+1)$ matrices[8] over the finite field $\mathbb{F}_Q$.; thus very roughly, $N \approx Q^{n^2} = p^{sn^2}$, so $s = \Theta(\log N)$.

Indexing the dimensions of such matrices as $[0\,..\,n]$, we define for $0 \leq i \neq j \leq n$ and $\alpha \in \mathbb{F}_Q$ the *transvection* matrix

$$e_{i,j}(\alpha) = \begin{bmatrix} 1 & & & & \\ & 1 & & \alpha & \\ & & 1 & & \\ & & & \ddots & \\ & & & & 1 \end{bmatrix} \qquad (\alpha \text{ in the } (i,j)\text{-entry}); \tag{2.3}$$

[5] Strictly speaking, it is not sensible to put these $\exp(-1/\delta^{\Omega(1)})$ and $\delta^{O(1)}$ expressions into Equation (2.2): they only make sense in the context of a growing family of direct-product testers, and it's also not particularly natural to insist on certain fixed functions of δ in such a definition. However, we *will* eventually work in the context of a growing family, and the theorems will achieve some unspecified $\exp(-1/\delta^{\Omega(1)})$ and $\delta^{O(1)}$ bounds, so we write the definition this way to minimize the number of parameters that need to be introduced.

[6] This, however, required a nontrivial appendix by Yun to verify that one could allow the "p" parameter to be superconstant; namely $p = \mathrm{polylog}\, N$. Note that no extra effort is required to verify this for the KO complex agreement tester we presently describe.

[7] In this section, we will treat these parameters as "constants", but in actuality the construction remains strongly explicitly even for p as large as $\mathrm{polylog}\, N$, and this is the parameter setting needed for the PCP application. We also mention that p could be a prime power rather than a prime.

[8] The choice of $n+1$ rather than n here is for notational convenience in the main body of the paper.

when multiplying from the right, this is the elementary operation of adding α-times-the-ith-column to the jth column, familiar from Gaussian elimination. It is well known that if one starts with the identity matrix and repeatedly applies transvections, the set of all matrices that can be achieved is precisely $G = \mathrm{SL}_n(\mathbb{F}_Q)$, the matrices in $\mathbb{F}_Q^{n\times n}$ of determinant 1. The size of G is roughly Q^{n^2-1}.[9]

Let us fix a concrete representation of $\mathbb{F}_Q$ as $\mathbb{F}_p[X]/(\mathrm{irred}(X))$, where $\mathrm{irred}(X)$ denotes some irreducible polynomial in $\mathbb{F}_p[X]$ of degree s.[10] We define the following "tiny" set T of p^4 field elements:[11]

$$T \coloneqq \{f(X) \in \mathbb{F}_p[X] : \deg(f) \leq \kappa \coloneqq 3\} \subseteq \mathbb{F}_Q; \tag{2.4}$$

then we define a transvection $e_{i,j}(\alpha)$ to be "tiny" if $\alpha \in T$ and $j = i + 1 \bmod n + 1$.

Next, let us define an equivalence relation $\equiv_0$ on matrices $g \in G$:

$$g \equiv_0 g' \iff g' \text{ can be obtained from } g \text{ by repeatedly applying tiny transvections } e_{i,i+1}(\alpha), \quad i+1 \neq 0 \bmod n+1. \tag{2.5}$$

(In other words, we disallow use of the "wraparound" tiny transvections.) It is not hard to show that the size of every associated equivalence class $[g]_0$ is some $c \coloneqq p^{O(n^2)}$, a (very large) constant. More generally, for $0 \leq a \leq n$, we define the equivalence relation $\equiv_a$ by allowing in Equation (2.5) only the tiny transvections $e_{i,i+1}$ with $i + 1 \neq a$; this yields equivalence classes $[g]_a$ also of size c.

Now for each matrix $g \in G$ we get $n + 1$ equivalence classes $[g]_0, \dots, [g]_n$ of size c. The universe $\mathcal{U}$ for the assignment tester will be $(n + 1)$-partite, with the ath part consisting of all equivalence classes $[g]_a$; thus $N \coloneqq |\mathcal{U}| = (n + 1)|G|/c = \Theta(|G|)$. We can naturally form an $(n + 1)$-uniform hypergraph (i.e., n-dimensional simplicial complex) on $\mathcal{U}$ by taking all $|G|$ hyperedges of the form $\{[g]_a : a \in [0\,..\,n]\}$. For the actual assignment tester, we simply take all the size-k sub-edges $\{[g]_a : a \in A, |A| = k\}$; this yields a sparse agreement tester, because there are $\binom{n}{k}|G| = \Theta(|G|)$ such sets.

Explicitly, the agreement tester for $f : \mathcal{U} \to \Sigma$, when described for two provers, is the following. The tester picks two random size-k subsets $\boldsymbol{A}, \boldsymbol{A}'$ of $\{0, 1, \dots, n\}$ conditioned on $|\boldsymbol{A} \cap \boldsymbol{A}'| = \sqrt{k}$, picks $\boldsymbol{g} \in \mathrm{SL}_{n+1}(\mathbb{F}_Q)$ uniformly at random, and asks the two provers for f's values on $([\boldsymbol{g}]_a)_{a\in\boldsymbol{A}}$ and on $([\boldsymbol{g}]_{a'})_{a'\in\boldsymbol{A}'}$.

The *local KO complex*. Before ending this section, let us also describe the subhypergraphs appearing in Theorems 1.2 and 1.3. We first have what we call the *local KO complex* $\mathfrak{LA}_n(\mathbb{F}_p)$, to which all vertex-neighborhoods (links) are isomorphic. It may be defined by restricting attention to the n-partite hypergraph with just the hyperedges

$$\{[g]_1, \dots, [g]_n\}, \text{ where } g \equiv_0 \mathbb{1}. \tag{2.6}$$

Note that this hypergraph is of "constant" size, only being defined by the c elements of $[\mathbb{1}]_0$. These are all the matrices that can be obtained by multiplying together non-wraparound tiny transvections $e_{i,i+1}(\alpha)$. It is not hard to show (using $s > \kappa n$) that these are all upper-triangular matrices in $\mathbb{F}_Q^{(n+1)\times(n+1)}$ with 1's on the diagonal and with (i, j)-entry β satisfying $\deg(\beta) \leq \kappa(j - i)$. In particular, this set does not depend on s or "Q" per se, which is why we use the notation $\mathfrak{LA}_n(\mathbb{F}_p)$.

As for the further subcomplexes $\mathfrak{LA}_n(\mathbb{F}_p)_w[\{a\}, \{b\}, \{c\}]$ from Theorem 1.3, these are all isomorphic to the tripartite hypergraph with just the hyperedges

$$\{[g]_a, [g]_b, [g]_c\}, \text{ where } g \equiv_i \mathbb{1} \text{ for all } i \in \{0\} \cup W. \tag{2.7}$$

[9]Precisely, its size is $Q^{\binom{n+1}{2}} \cdot Q^n \cdot \sum_{\pi\in\mathfrak{S}(n+1)} Q^{\mathrm{inv}(\pi)}$, where $\mathrm{inv}(\pi)$ denotes the number of inversions in π. This is because there is a 1-1 correspondence between matrices in $\mathrm{SL}_{n+1}(\mathbb{F}_Q)$ and their "Bruhat decompositions" of the form $UDP_\pi L$, where U is upper-triangular with 1's on the diagonal, D is diagonal with entries multiplying to 1, P_π is the permutation matrix for permutation π, and L is lower-triangular with 1's on the diagonal and $L_{i,j} = 0$ whenever $\pi(i) > \pi(j)$. This 1-1 correspondence gives an easy way to draw a matrix from $\mathrm{SL}_{n+1}(\mathbb{F}_Q)$ exactly uniformly at random in $\Theta(\log Q)$ time.

[10][KO23] mod out by X^s, but following [OP22] we find it more natural to mod out by an irreducible; it makes no difference to the result in [KOW25]. Shoup [Sho90] shows how to find such an irreducible in deterministic $\mathrm{poly}(s) = \mathrm{polylog}\, N$ time.

[11]The distinction between this construction from [KOW25] and the original one in [KO23] is that the latter has $\kappa = 1$ in the definition of T. [KOW25] needs $\kappa = 3$ for a minor technical reason.

2.2 Triword expansion

In this section, we define the "topological" HDX properties that go into the Bafna–Minzer/Dikstein–Dinur [BM24; DD24a] characterization of when a hypergraph can serve as a low-soundness agreement tester. (Roughly speaking, Properties (1), (2) from §1.) For the most precise definitions, see §3.

Let $X = (V, E, T)$ be a connected "triangle complex" (pure 2-dimensional simplicial complex); the reader may also think of X as a graph (V, E) formed by the union of a set T of 3-cliques. Let $f : \vec{E} \to \mathfrak{S}(m)$ be a labeling of X's directed edges by permutations of $[m]$ (satisfying $f(u,v) = f(v,u)^{-1}$). The associated *m-lift* X^f of X is the graph on vertex set $V \times [m]$ which has an edge from (u,i) to $(v, f(u,v)(i))$ whenever $(u,v) \in \vec{E}$.

We say the lift is an *m-cover* if every triangle (u,v,w) in T gets lifted to m disjoint triangles in X^f — equivalently, $f(u,v)f(v,w)f(w,v) = \mathbb{1}$. There are certain "trivial" ways to get m-covers; take any vertex-labeling $g : V \to \mathfrak{S}(m)$ and define $f(u,v) = g(u)g(v)^{-1}$. For these "trivial" (sometimes called "balanced") lifts, the lifted graph X^f consists of m disjoint copies of X, and the product of f around *every* cycle is $\mathbb{1}$ (not just around every 3-cycle). When the *only* m-covers of X are these trivial ones, X is said to have *trivial 1-cohomology over* $\mathfrak{S}(m)$. More generally, this can be defined over any group Γ:

Definition 2.8 (Trivial 1-cohomology over Γ)**.** A triangle complex $X = (V, E, T)$ is said to have *trivial 1-cohomology over* Γ if, whenever $f : \vec{E} \to \Gamma$ (with $f(u,v) = f(v,u)^{-1}$) has $f(u,v)f(v,w)f(w,u) = \mathbb{1}$ for all $(u,v,w) \in T$, there exists $g : V \to \Gamma$ such that $f(u,v) = g(u)g(v)^{-1}$ for all $(u,v) \in E$. ♢

The main theorem of Kaufman, Oppenheim, and Weinberger [KOW25] is that the $k = 3$ hypergraph described in the preceding section §2.1.1 indeed has no nontrivial m-covers for $m < p$:

Theorem 2.9. *([KOW25].) Provided Γ has no nontrivial elements of order p, the [KOW25] complexes from Theorem 1.1 have vanishing 1-cohomology over Γ.*

As mentioned in §1, this result verifies that the complexes have Property (1) of the Bafna–Minzer/Dikstein–Dinur HDX characterization of agreement testers.

Let us now define our term "r-triword expansion" from Property (2). First, when the 1-cohomology of X over $\mathfrak{S}(m)$ vanishes, one can ask for a quantitative strengthening:

Definition 2.10 (1-coboundary expansion)**.** X is said to have *1-coboundary expansion at least ϵ (over Γ)* if, whenever $f : \vec{E} \to \Gamma$ (with $f(u,v) = f(v,u)^{-1}$) has $f(u,v)f(v,w)f(w,u) = \mathbb{1}$ for a $1 - \zeta$ fraction of triangles $(u,v,w) \in T$, there exists $g : V \to \Gamma$ such that $f(u,v) = g(u)g(v)^{-1}$ for at least a $1 - \zeta/\epsilon$ fraction of edges $(u,v) \in E$. ♢

At a (very) intuitive level (which is more accurate if X has bounded diameter), to have good 1-coboundary expansion over $\mathfrak{S}(m)$ is to say that whenever X^f is an m-lift in which most triangles get lifted to m disjoint triangles, then most longer loops *also* get lifted to m disjoint copies. One can then intuit that this may be true if loops in X can be triangulated with "few" triangles. This is the essence of the "cones method" described at the end of §1.

This defines 1-coboundary expansion, but to finish explaining terms in Theorem 1.2, we make a generalized notion suitable for n-uniform hypergraphs $\mathfrak{X}$ such as $\mathfrak{LA}_n(\mathbb{F}_p)$. For any such $\mathfrak{X}$ and any parameter $1 \le r \le n/3$, we may turn $\mathfrak{X}$ into a triangle complex $(V_r(\mathfrak{X}), E_r(\mathfrak{X}), T_r(\mathfrak{X}))$ (called the "r-faces complex" in [DD24a]) by specifying the following: the elements of $V_r(\mathfrak{X})$ are all sets of r vertices in $\mathfrak{X}$ that appear together in some hyperedge; and, the elements of $T_r(\mathfrak{X})$ are all triples (S_1, S_2, S_3) from $V_r(\mathfrak{X})$ where $S_1 \cup S_2 \cup S_3$ is a size-$3r$ set of vertices in $\mathfrak{X}$ appearing together in some hyperedge. Finally, we say that $\mathfrak{X}$ has *r-triword expansion at least ϵ over* Γ if $(V_r(\mathfrak{X}), E_r(\mathfrak{X}), T_r(\mathfrak{X}))$ has 1-coboundary expansion at least ϵ over Γ. With this definition in place, we have explained all terms appearing in Theorems 1.2 and 1.3.

2.3 Proofs of identity in the unipotent group

In this subsection, before we finally dive into the main body of the paper, we preview some of the group-theoretic setup and proof strategies needed for our core technical theorem, Theorem 1.4. The discussion is substantially simplified, but we hope it gives a flavor of the considerations which do come up in the actual proof.

2.3.1 Circuits

Our key coboundary expansion result, Theorem 1.3, is concerned with loops and triangles inside the "local KO complex" described at the end of §2.1.1. In turn, this means we need to understand well which matrices over can be obtained from others through multiplication by tiny transvections. For this, we found it extremely helpful to use the computer science language of *circuits*.

Fix any field $\mathbb{F}$. Let $\mathbb{F}^{\leq d}[X]$ denote the set of polynomials in X of degree at most d over $\mathbb{F}$. We will be interested in circuits that model $(n+1)\times(n+1)$ matrices over $\mathbb{F}^{\leq d}[X]$ that are unipotent (upper-triangular with 1's on the diagonal).

Consider circuits with $n+1$ horizontal "wires", numbered 0 through n from top to bottom.[12] We use the notation $[0\,..\,n]=\{0,\dots,n\}$. A circuit will then simply be an ordered sequence of *gates*, each of which represents a(n upper-triangular) transvection. Generalizing the set of "tiny transvections", each gate is specified by a pair $i<j\in[0\,..\,n]$ and a polynomial $f\in\mathbb{F}^{\leq 3(j-i)}[X]$; this gate is written $e_{i,j}(f)$. We think of i as the "head" of the gate, j as the "tail", and $j-i$ as the "height". The head wire is always (strictly) above the tail, and the height is therefore always (strictly) positive. The degree is bounded by the constant 3 times the height. (The particular choice of the constant 3 is not important for the purposes of this subsection.) A circuit C is a sequence of gates, denoted $\langle e_{i_1,j_1}(f_1)\rangle\cdots\langle e_{i_T,j_T}(f_T)\rangle$. Here, we use the notation $\langle\cdot\rangle$ around each individual gate to emphasize that a circuit $\langle e_{i_1,j_1}(f_1)\rangle\cdots\langle e_{i_T,j_T}(f_T)\rangle$ is formally distinct from the (unipotent) matrix product $e_{i_1,j_1}(f_1)\cdots e_{i_T,j_T}(f_T)$. In particular, many distinct circuits may represent the same unipotent matrix.

Although we don't formally *need* to think of these circuits as acting upon anything, it will psychologically helpful to do so. We think of the wires as "registers" that store polynomials $v_0,\dots,v_n\in\mathbb{F}[X]$; collectively, we think of them as storing a row-vector $v=(v_0,\dots,v_n)\in\mathbb{F}[X]^{n+1}$. The action of the gate $e_{i,j}(f)$ is to update the contents of the jth register by setting $v_j\leftarrow v_j+fv_i$. We may also think of a circuit as a "program", with a list of "instructions" (gates), where $e_{i,j}(f)$ is the instruction "$\texttt{add}\ fv_i\ \texttt{to}\ v_j$;".

Linear-algebraically, a circuit C corresponds to an $(n+1)\times(n+1)$ (unipotent) matrix over $\mathbb{F}[X]$, and it acts on a row vector $v\in\mathbb{F}[X]^{n+1}$ to produce $v\cdot C$. We will write $C\equiv C'$ and say that C and C' are *equivalent* if $v\cdot C=v\cdot C'$ for every $v\in\mathbb{F}^{n+1}$; in other words, if C and C' are equal as matrices. We will be particularly interested in the case of $C\equiv\mathbb{1}$, meaning C is equivalent to the "empty" (trivial) circuit with no gates. Note that $C\equiv C'$ iff $C^{-1}C'\equiv\mathbb{1}$, where C^{-1} is the "inverse" circuit which reverses the orders of gates and negates their entries (note that $\langle e_{i,j}(f)\rangle\langle e_{i,j}(-f)\rangle\equiv\mathbb{1}$).

2.3.2 "Relations" among circuits

Consider the following important examples:

Example 2.11 ("Linearity" relation)**.** Let $n=1$ and $C:=\langle e_{0,1}(f)\rangle\langle e_{0,1}(f')\rangle$. If the initial state is (v_0,v_1), then the states after the gates operate are:

$$\begin{pmatrix}v_0\\v_1\end{pmatrix}\overset{e_{0,1}(f)}{\longmapsto}\begin{pmatrix}v_0\\v_1+fv_0\end{pmatrix}\overset{e_{0,1}(f')}{\longmapsto}\begin{pmatrix}v_0\\v_1+(f+f')v_0\end{pmatrix}.$$

Hence, $C\equiv\langle e_{0,1}(f+f')\rangle$. ◊

Example 2.12 ("Nontrivial commutator" relation)**.** Let $n=2$ and $C:=\langle e_{0,1}(f)\rangle\langle e_{1,2}(g)\rangle\langle e_{0,1}(-f)\rangle\langle e_{0,1}(-g)\rangle$. If the initial state is (v_0,v_1,v_2), then the states after the gates operate are:

$$\begin{pmatrix}v_0\\v_1\\v_2\end{pmatrix}\overset{e_{0,1}(f)}{\longmapsto}\begin{pmatrix}v_0\\v_1+fv_0\\v_2\end{pmatrix}\overset{e_{1,2}(g)}{\longmapsto}\begin{pmatrix}v_0\\v_1+fv_0\\v_2+gv_1+fgv_0\end{pmatrix}\overset{e_{0,1}(-f)}{\longmapsto}\begin{pmatrix}v_0\\v_1\\v_2+gv_1+fgv_0\end{pmatrix}\overset{e_{1,2}(-g)}{\longmapsto}\begin{pmatrix}v_0\\v_1\\v_2+fgv_0\end{pmatrix}.$$

Hence, $C\equiv\langle e_{0,2}(fg)\rangle$. ◊

[12] We follow the quantum computing convention of circuits with horizontal wires computing linear transformations; but unlike in that field, matrices will get multiplied in the same left-to-right order as gates, hence we consider action by right-multiplication on row vectors.

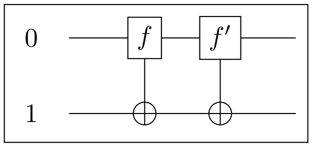

(a) A two-gate circuit ($n = 1$). Written as a program: "`add` fv_0 `to` v_1; `add` $f'v_0$ `to` v_1;".

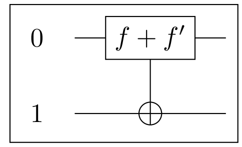

(b) A one-gate circuit ($n = 1$). Written as a program: "`add` $(f + f')v_0$ `to` v_1;".

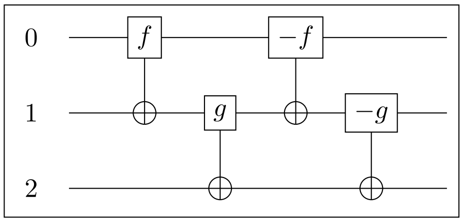

(c) A four-gate circuit ($n = 2$). Written as a program: "`add` fv_0 `to` v_1; `add` gv_1 `to` v_2; `subtract` fv_0 `from` v_1; `subtract` gv_1 `from` v_2;".

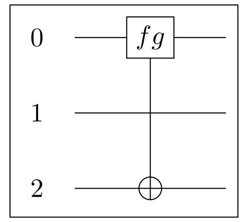

(d) A one-gate circuit ($n = 2$). Written as a program: "`add` fgv_0 `to` v_2;".

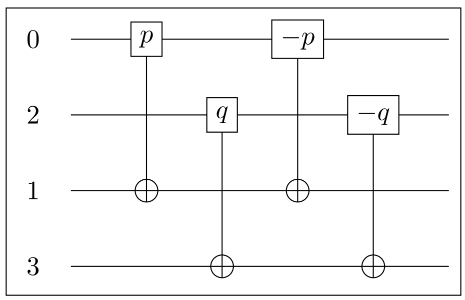

(e) A four-gate circuit ($n = 3$). Written as a program: "`add` pv_0 `to` v_2; `add` qv_1 `to` v_3; `subtract` pv_0 `from` v_2; `subtract` qv_1 `from` v_3;".

Figure 2: Some circuits. The circuits in Figures 2a and 2b are equivalent (the "linearity" relation, Example 2.11). The circuits in Figures 2c and 2d are equivalent (the "nontrivial commutator" relation, Example 2.12). The circuit in Figure 2e is equivalent to the trivial (empty) circuit (the "trivial commutator" relation, Example 2.13).

Example 2.13 ("Trivial commutator" relation)**.** Let $n = 3$ and $C \coloneqq \langle e_{0,2}(p)\rangle\langle e_{1,3}(q)\rangle\langle e_{0,2}(-p)\rangle\langle e_{1,3}(-q)\rangle$. If the initial state is (v_0, v_1, v_2, v_3), then the states after the gates operate are:

$$\begin{pmatrix} v_0 \\ v_1 \\ v_2 \\ v_3 \end{pmatrix} \xmapsto{e_{0,2}(p)} \begin{pmatrix} v_0 \\ v_1 \\ v_2 + pv_0 \\ v_3 \end{pmatrix} \xmapsto{e_{1,3}(q)} \begin{pmatrix} v_0 \\ v_1 \\ v_2 + pv_0 \\ v_3 + qv_1 \end{pmatrix} \xmapsto{e_{0,2}(-p)} \begin{pmatrix} v_0 \\ v_1 \\ v_2 \\ v_3 + qv_1 \end{pmatrix} \xmapsto{e_{1,3}(-q)} \begin{pmatrix} v_0 \\ v_1 \\ v_2 \\ v_3 \end{pmatrix}.$$

Hence, $C \equiv \mathbb{1}$. ◊

(Here, the particular choice of the pairs $(0, 2)$ and $(1, 3)$ was not important. The only important feature is disjointness; $(0, 3)$ and $(1, 2)$, or $(0, 1)$ and $(2, 3)$, would have worked as well.)

These examples generalize easily to the following fact:

Fact 2.14 ("Steinberg relations")**.** *We have:*

1. *For every $i < j \in [0 \,..\, n]$ and $f, f' \in \mathbb{F}^{\leq 3(j-i)}[X]$, $\langle e_{i,j}(f)\rangle\langle e_{i,j}(f')\rangle \equiv \langle e_{i,j}(f + f')\rangle$.*

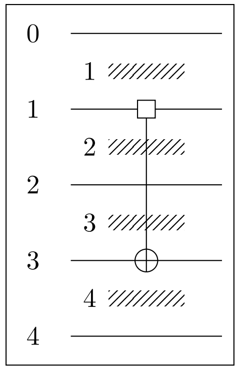

Figure 3: An illustration of the situation when $n = 4$. There are $n + 1 = 5$ wires, labeled $0, \ldots, 4$, and n moats, labeled $1, \ldots, 4$. The "gate" from wire 1 to wire 3 crosses moats 2 and 3 but not moats 1 or 4. This corresponds to the fact that $(1\,..\,3] = \{2, 3\}$.

2. *For every* $i < j < k \in [0\,..\,n]$ *and* $f \in \mathbb{F}^{\leq 3(j-i)}[X], g \in \mathbb{F}^{\leq 3(k-j)}[X]$, $\langle e_{i,j}(f)\rangle\langle e_{j,k}(g)\rangle\langle e_{i,j}(-f)\rangle\langle e_{j,k}(-g)\rangle = \langle e_{i,k}(fg)\rangle$.

3. *For every* $i < j \in [0\,..\,n]$ *and* $k < \ell \in [0\,..\,n]$ *such that* $j \neq k$ *and* $i \neq \ell$, *and* $f \in \mathbb{F}^{\leq 3(j-i)}[X]$ *and* $g \in \mathbb{F}^{\leq 3(k-j)}[X]$, $\langle e_{i,j}(f)\rangle\langle e_{k,\ell}(h)\rangle\langle e_{i,j}(-f)\rangle\langle e_{k,\ell}(-h)\rangle = \mathbb{1}$.

Compare to Definition 3.50 below.

Given a circuit C, it is not hard to check whether $C \equiv \mathbb{1}$; you can simply compute the matrix corresponding to C and check if it is the identity matrix. But we are concerned with the problem of *certifying* that $C \equiv \mathbb{1}$ in particular, limited proof models. We are interested both in the *qualitative* question of "Given C, is it possible to certify that $C \equiv \mathbb{1}$?" and the *quantitative* question of whether this certification can be *efficient*. We first focus on the qualitative question, and then circle back to the quantitative version (which is slightly subtler to define) at the end of this section.

2.3.3 Avoidant transformations

We now turn to defining the particular proof model we will be interested in, which we call the "avoidant transformations" model.

Recall that we have $n+1$ wires labeled by $[0\,..\,n] = \{0, \ldots, n+1\}$. We now also consider the *gaps* between each pair of adjacent wires. We label the gap between wires 0 and 1 by 1, the gap between wires 1 and 2 by 2, and so on; hence there are n gaps, labeled $1, \ldots, n$. Now, let us fix a small set $\mathcal{A} \subseteq [1\,..\,n]$, which we call the *moats*. We say a gate from wire i to wire j *crosses* moat $a \in \mathcal{A}$ if $a \in (i\,..\,j] = \{i + 1, \ldots, j\}$, or equivalently $i + 1 \leq a \leq j$; we say a gate is *a-avoiding* if it does *not* cross a. For $a \in \mathcal{A}$, we say a circuit is *a-avoiding* if every gate in it is a-avoiding, and *avoidant* if every gate in it avoids *some* moat.

Now, we define the concept of an *avoidant transformation* on a circuit C. Suppose C is decomposed into three subcircuits $C = C_1C_2C_3$ (any of which may be empty), and C_2 is a-avoiding for some moat $a \in \mathcal{A}$. Then an *avoidant transformation* replaces C_2 with another a-avoiding circuit C_2' such that $C_2 \equiv C_2'$.

Given this setup, we can now formally define our goal: We want to show that for every set of three moats $\mathcal{A}$ and circuit C such that $C \equiv \mathbb{1}$,[13] C can be transformed to $\mathbb{1}$ using only avoidant transformations. (This is the "qualitative" goal; we address quantitative aspects, which are the main actual contribution of our work, at the end of this section.)

Fix three moats $a < b < c$. We develop some notations for how these moats break up the set of wires $[0\,..\,n]$ into intervals and how these in turn break up the gates into different types. In particular, we use $\mathcal{I}_{\mathrm{A}}$, $\mathcal{I}_{\mathrm{B}}$, $\mathcal{I}_{\mathrm{C}}$, and $\mathcal{I}_{\mathrm{D}}$ to denote, respectively, wires above moat a, wires between moats a and b, wires between moats b and c, and wires below moat d. These four sets of wires, which we call *bands*, partition the full set

[13]Technically, we need to restrict here to only those in which every gate avoids at least one moat, since those gates cannot possibly be changed via avoidant transformations. We elide this distinction in this overview.

$[0\,..\,n]$. We then divide the gates of an arbitrary circuit into ten *types*, corresponding to letters, according to the following table:

Gate type	A	B	C	D	F	G	H	P	Q	Z
Head interval	$\mathcal{I}_\mathrm{A}$	$\mathcal{I}_\mathrm{B}$	$\mathcal{I}_\mathrm{C}$	$\mathcal{I}_\mathrm{D}$	$\mathcal{I}_\mathrm{A}$	$\mathcal{I}_\mathrm{B}$	$\mathcal{I}_\mathrm{C}$	$\mathcal{I}_\mathrm{A}$	$\mathcal{I}_\mathrm{B}$	$\mathcal{I}_\mathrm{A}$
Tail interval	$\mathcal{I}_\mathrm{A}$	$\mathcal{I}_\mathrm{B}$	$\mathcal{I}_\mathrm{C}$	$\mathcal{I}_\mathrm{D}$	$\mathcal{I}_\mathrm{B}$	$\mathcal{I}_\mathrm{C}$	$\mathcal{I}_\mathrm{D}$	$\mathcal{I}_\mathrm{C}$	$\mathcal{I}_\mathrm{D}$	$\mathcal{I}_\mathrm{D}$
Avoids?	a,b,c	a,b,c	a,b,c	a,b,c	b,c	a,c	a,b	c	a	none

Table 1: The ten possible types of gates with respect to the three moats.

See also Figure 7 below for a visual depiction of the situation.

2.3.4 The power of working space

We now consider two concrete examples of circuits which show some important differences under the lens of avoidant transformations.

Example 2.15. Consider the case $n = 3$, $a = 1$, $b = 2$, $c = 3$, and the circuit

$$C \coloneqq \langle e_{0,2}(p)\rangle\langle e_{1,3}(q)\rangle\langle e_{0,2}(-p)\rangle\langle e_{1,3}(-q)\rangle$$

(which is also depicted in Figure 4 below). Here, $p, q \in \mathbb{F}^{\le 6}[X]$ are arbitrary sextic polynomials. Since $0 \in \mathcal{I}_\mathrm{A}$, $1 \in \mathcal{I}_\mathrm{B}$, $2 \in \mathcal{I}_\mathrm{C}$, and $3 \in \mathcal{I}_\mathrm{D}$, the gates $\langle e_{0,2}(\pm p)\rangle$ are "type-P" and the gates $\langle e_{1,3}(\pm q)\rangle$ are "type-Q". By the "trivial commutator" Steinberg relation (i.e., Example 2.13), this circuit $C \equiv \mathbb{1}$. However, it is not completely obvious how to avoidantly transform this circuit into the empty circuit. Doing so turns out to be possible, but it requires a few tricks — firstly, a clever "symmetry-breaking" expansion to solve the problem in the case where p and q are degree-0 polynomials [BD01, §4] — and secondly, a "lifting" argument to generalize to arbitrary $p, q \in \mathbb{F}^{\le 3}[X]$ [KO21, §8]. (Actually, Kaufman and Oppenheim [KO21] only consider the case $p, q \in \mathbb{F}^{\le 1}[X]$, but it is not hard to show that their argument generalizes. See [OS25] for more on lifting arguments.) ◊

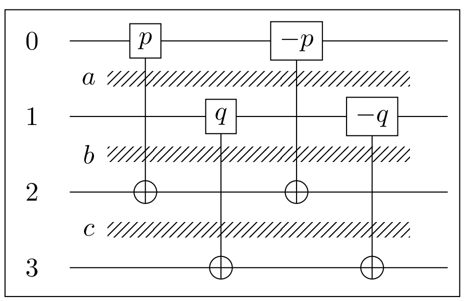

Figure 4: The circuit in Example 2.15 with $n = 3$, $a = 1$, $b = 2$, and $c = 3$. Every gate avoids at least one moat; in particular, the $\pm p$ gates avoid c and the $\pm q$ gates avoid a.

A simple but crucial insight of our group-theoretic work is that spreading out the moats a, b, and c can facilitate the use of avoidant transformations. For instance:

Example 2.16. Consider the case $n = 5$, $a = 1$, $b = 3$, $c = 5$, and the circuit

$$C \coloneqq \langle e_{0,3}(p)\rangle\langle e_{2,5}(q)\rangle\langle e_{0,3}(-p)\rangle\langle e_{2,5}(-q)\rangle$$

(which is also depicted in Figure 5 below). $p, q \in \mathbb{F}^{\le 9}[X]$ are arbitrary nonic polynomials. Since $0 \in \mathcal{I}_\mathrm{A}$, $2 \in \mathcal{I}_\mathrm{B}$, $3 \in \mathcal{I}_\mathrm{C}$, and $5 \in \mathcal{I}_\mathrm{D}$, the gates $e_{0,3}(\pm p)$ are again "type-P" and $e_{2,5}(\pm q)$ "type-Q", and the circuit $C \equiv \mathbb{1}$. This time, we claim it is (relatively) straightforward to certify that $C \equiv \mathbb{1}$, or equivalently that $e_{0,3}(p)e_{2,5}(q) \equiv e_{2,5}(q)e_{0,3}(p)$, avoidantly.

First, we claim that we can reduce to the case where we have factorizations $p = fg^{(1)}$ and $q = g^{(2)}h$, for cubic polynomials $f, h \in \mathbb{F}^{\le 3}[X]$ and sextic polynomials $g^{(1)}, g^{(2)} \in \mathbb{F}[X]$. For instance, write the degree decompositions $p = p_0 + p_1 X + p_2 X^2 + \cdots + p_9 X^9$ and similarly for q. By the linearity relations, to commute $e_{0,3}(p)$ and $e_{2,5}(q)$ using avoidant transformations, it suffices to commute each pair $e_{0,3}(p_i X^i)$ and $e_{2,5}(q_j X^j)$ using avoidant transformations, and for each such pair we have

$$p_i X^i = \underbrace{(p_i X^{\max\{i-3,0\}})}_{=: f_i \in \mathbb{F}^{\le 3}[X]} \cdot \underbrace{(X^{i-\max\{i-3,0\}})}_{=: g_i^{(1)} \in \mathbb{F}^{\le 6}[X]} \text{ and } q_j X^j = \underbrace{(X^{j-\max\{j-3,0\}})}_{=: g_j^{(2)} \in \mathbb{F}^{\le 6}[X]} \cdot \underbrace{(q_j X^{\max\{j-3,0\}})}_{=: h_j \in \mathbb{F}^{\le 3}[X]}.$$

Now, consider the problem of commuting $e_{0,3}(fg^{(1)})$ with $e_{2,5}(g^{(2)}h)$. We have, using (avoidant) nontrivial commutator relations,

$$\langle e_{0,3}(fg^{(1)})\rangle\langle e_{2,5}(g^{(2)}h)\rangle = \langle e_{0,1}(f)\rangle\langle e_{1,2}(g^{(1)})\rangle\langle e_{0,1}(-f)\rangle\langle e_{1,2}(-g^{(1)})\rangle\cdot\langle e_{2,4}(g^{(2)})\rangle\langle e_{4,5}(h)\rangle\langle e_{2,4}(-g^{(2)})\rangle\langle e_{4,5}(-h)\rangle. \tag{2.17}$$

Now, one can check that using (avoidant) trivial commutator relations, each of the first four gates on the right-hand side of Equation (2.17) commutes with each of the last four. See Figure 5 for a visual depiction. ♢

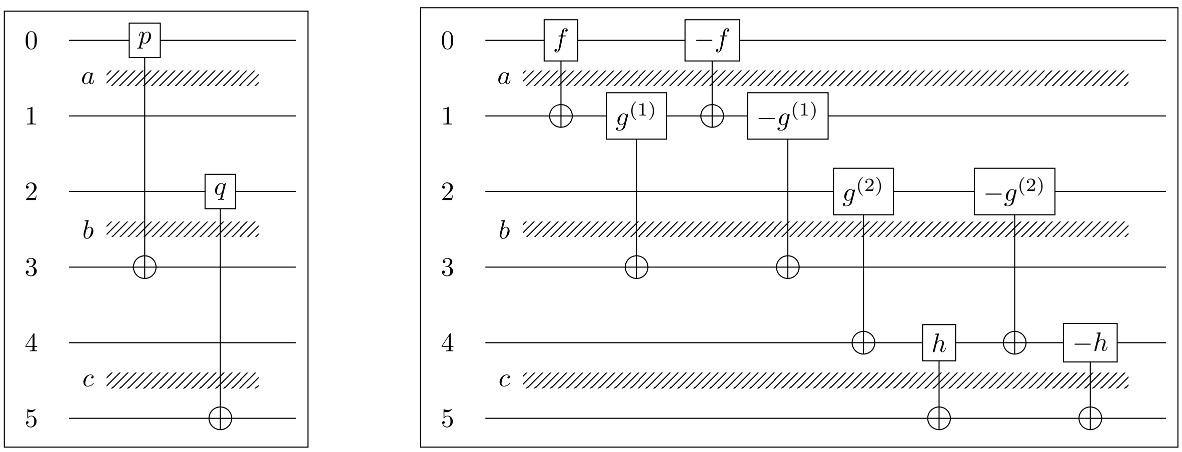

(a) A two-gate circuit C.

(b) Suppose $fg^{(1)} = p$ and $g^{(2)}h = q$; then this is a circuit $C' \equiv C$ with eight gates. C' can be obtained from C avoidantly (by expanding each gate using the Steinberg commutator relation). The $\pm g^{(1)}$ and $\pm g^{(2)}$ gates avoid a and c; $\pm f$ and $\pm g^{(2)}$ avoid c; $\pm g^{(1)}$ and $\pm h$ avoid a; and $\pm f$ and $\pm h$ avoid b. Hence, the first half of the gates can be commuted past the second half using avoidant transformations.

Figure 5: The utility of additional working space.

Intuitively, having two wires in the bands $\mathcal{I}_{\mathrm{B}}$ and $\mathcal{I}_{\mathrm{C}}$ is helpful when certifying this circuit because we have enough "room" for gates to "slide past" one another without interference. We view this extra room as a kind of *working space*, allowing us to, e.g., simultaneously "write" values onto the band $\mathcal{I}_{\mathrm{B}}$ (the output of the $e_{0,1}(\pm f)$ gates) *and* "read" values from the band $\mathcal{I}_{\mathrm{B}}$ (the input of the $e_{2,4}(\pm g^{(2)})$ gates).

2.3.5 The Steinberg relations

In the previous subsection, Example 2.16 showed how having more than one wire in the bands $\mathcal{I}_{\mathrm{B}}$ and $\mathcal{I}_{\mathrm{C}}$ let us easily prove that type-`P` and type-`Q` gates commute using avoidant transformations. This is only the initial step towards our full goal: Certifying triviality of an *arbitrary* circuit (wherein all gates are avoidant).

For this full goal, we need to employ a useful reduction due to [KO21], which we now explain. It is a well-known fact in group theory that the Steinberg relations (Claim 2.14) suffice to "present" the group of unipotent matrices over a finite field [Ste62]. What this means is that — temporarily ignoring the issue of transformations needing to be "avoiding" — given any circuit which is equivalent to the identity, it is possible to transform it into the identity simply by repeated applications of the equations in Claim 2.14 to adjacent pairs of gates. (This can be viewed essentially as an implementation of Gaussian elimination on a unipotent matrix.)

But how can we implement this via avoidant transformations? Indeed, many Steinberg relations will involve type-Z gates (that is, gates $e_{i,j}(z)$ where $i \in \mathcal{I}_\mathrm{A}$ and $j \in \mathcal{I}_\mathrm{D}$); these gates are not avoidant, since they cross all three of the moats a, b, and c. Therefore, these relations cannot be used (or even stated) in the world of avoidant transformations.

The idea of Kaufman and Oppenheim [KO21] (as interpreted also by [OS25]) was to *simulate* these "missing" type-Z gates by an "alias" circuit $\zeta_{i,j}(z)$ to "stand in" for the gate $e_{i,j}(f)$. We should have that $\zeta_{i,j}(f) \equiv \langle e_{i,j}(f)\rangle$ and yet $\zeta_{i,j}(f)$ is composed entirely of avoidant gates. When f admits a nice factorization, the existence of such $\zeta_{i,j}(z)$ is immediate — e.g., set $\zeta_{i,j}(z) := \langle e_{i,k}(f)\rangle\langle e_{i,k}(q)\rangle\langle e_{i,k}(-f)\rangle\langle e_{i,k}(-q)\rangle$ when $z = fq$, $f \in \mathbb{F}^{\leq 3(k-i)}[X]$, and $q \in \mathbb{F}^{\leq 3(j-k)}[X]$, and $k \in \mathcal{I}_\mathrm{B}$ is of our choosing. Then, we again extend to arbitrary f by linearity. It is enough to then implement the remaining Steinberg relations as avoidant transformations while replacing every occurrence of $e_{i,j}(z)$ with $\zeta_{i,j}(z)$. These relations can be partitioned into three classes:

- The "nontrivial commutator relations" producing type-Z elements:

$$\zeta_{i,j}(z) \equiv \langle e_{i,k}(f)\rangle\langle e_{i,k}(-q)\rangle\langle e_{i,k}(-f)\rangle\langle e_{i,k}(-q)\rangle$$

  (for *every* $k \in \mathcal{I}_\mathrm{B}$) and

$$\zeta_{i,j}(z) \equiv \langle e_{i,k}(p)\rangle\langle e_{i,k}(h)\rangle\langle e_{i,k}(-p)\rangle\langle e_{i,k}(-h)\rangle$$

  (for *every* $k \in \mathcal{I}_\mathrm{C}$). We emphasize that only one such relation, for one fixed value of k, can be imposed by fiat in the definition of $\zeta_{i,j}(z)$. (There are also nontrivial commutators with type-A and type-D elements, though we defer these until the end.)

- The "trivial commutator" relations involving type-Z elements: Every gate of type B, C, F, G, F, G, H, P, and Q commutes the alias circuits $\zeta_{i,j}(z)$ for Z.

- The "linearity" relation for type-Z elements: $\zeta_{i,j}(z_1)\zeta_{i,j}(z_2) \equiv \zeta_{i,j}(z_1 + z_2)$.

2.3.6 Quantitative aspects of the proof

So far, we have only discussed the qualitative aspects of whether it is *possible* to certify triviality of a circuit, and not the quantitative question of how *efficiently* this can be done. In particular, we are interested in using a number of operations which is independent of the number of wires n.

An example of quantitative considerations already came up in Example 2.16, when we used the following sort of reduction: For a type-P gate $e_{i,j}(p)$, we have

$$\langle e_{i,j}(p)\rangle = \langle e_{i,j}(p_0X)\rangle\langle e_{i,j}(p_1X)\rangle\langle e_{i,j}(p_2X^2)\rangle \cdots \langle e_{i,j}(p_{3(j-i)}X^{3(j-i)})\rangle.$$

However, $j - i$ could already be as large as n, so we cannot afford to implement this full decomposition. Instead, we will have to break up the terms in p by degree only into $O(1)$ blocks (and similarly for q) and carry the proof through carefully; these considerations create a number of annoying pitfalls later in the proof.

Furthermore, to prove our full theorem, we will actually need to work in a slightly different model. In this model, our circuits are now made up of *supergates* of the form $\langle e_{i_1,j_1}(f_1)\cdots e_{i_T,j_T}(f_T)\rangle$, where $e_{i_1,j_1}(f_1), \ldots, e_{i_T,j_T}(f_T)$ all avoid the same moat.

Two supergates are treated as identical if the underlying products are the same when considered as matrices: $e_{i_1,j_1}(f_1)\cdots e_{i_T,j_T}(f_T) = e_{i'_1,j'_1}(f'_1)\cdots e_{i_{T'},j_{T'}}(f'_{T'})$. Importantly, we will have to deal with the case where T, the number of gates making up a supergate, is allowed to be $\Theta(n^2)$.

In this new model, an avoidant transformation is now allowed to do the following: Suppose we have a decomposition $C = C_1C_2C_3$, where C_1, C_2, C_3 are subcircuits made up of supergates, C_2 contains $O(1)$ supergates and every supergate in C_2 avoids the same moat. Then, we can replace C_2 with an equivalent subcircuit, also of length $O(1)$, avoiding this same moat.

Qualitatively, the old notion of avoidantly certifying the triviality a circuit made up of gates and the new notion of certifying the triviality of a circuit made up of supergates are equivalent; this is because we can transform a supergate $\langle e_{i_1,j_1}(f_1)\cdots e_{i_T,j_T}(f_T)\rangle$ avoidantly into a sequence of gates $\langle e_{i_1,j_1}(f_1)\rangle\cdots\langle e_{i_T,j_T}(f_T)\rangle$ and vice versa. However, quantitatively, this transformation is very expensive: A single supergate can become $\Theta(n^2)$ gates! Hence, we will instead need to operate on the supergates themselves directly.

Our ultimate strategy is very roughly to divide up every supergate into $O(1)$ sub-supergates, based on the wires and monomial degrees which appear. We then systematically re-implement the Steinberg relations in a "block-wise" manner. There are some technical details — the most annoying of which disappear in the case where $c-b$ and $b-c$ have greater common divisor $\Omega(n)$ — but morally, the proof is quite systematic and elementary. We made no (significant) attempt to optimize the final bounds we prove on the cost of avoidant transformations, and view this as an interesting problem for future work.

3 Preliminaries

Notations

For $a \le b \in \mathbb{N}$, we let $[a\,..\,b] \coloneqq \{c \in \mathbb{N} : a \le c \le b\}$, $[a\,..\,b) \coloneqq \{c \in \mathbb{N} : a \le c < b\}$, and $(a\,..\,b) \coloneqq \{c \in \mathbb{N} : a < c < b\}$.

For $r \in \mathbb{N}$, we will write $X \subseteq_r Y$ to denote that $X \subseteq Y$ and $|X| = r$. We will also very frequently need to refer to the situation where we have multiple pairwise disjoint subsets of a given set, so we invent some notation for this. We use $X_1, \dots, X_k \sqsubset Y$ to denote that $X_1, \dots, X_k$ are pairwise disjoint subsets of Y. Further, we use $X_1, \dots, X_k \sqsubset_r Y$ to denote that each $|X_i| = r$, and $X_1, \dots, X_k \sqsubset_{\le r} Y$ to denote that each $|X_i| \le r$. We also use $X \sqcup Y$ to denote the union of two sets X and Y when we want to emphasize that X and Y are disjoint.

3.1 Indexed simplicial complexes

In the body of this paper (i.e., outside the appendices), we only consider the case of n-partite, pure rank-n, weighted simplicial complexes. We term these *indexed simplicial complexes*.[14]

Definition 3.1 (Indexed simplicial complex)**.** Let $\mathcal{I}$ be a set of finite size $n \in \mathbb{N}$. An *$\mathcal{I}$-indexed simplicial complex* $\mathfrak{X}$ is a probability distribution over a product set

$$\Omega^{\mathcal{I}} \coloneqq \prod_{i \in \mathcal{I}} \Omega^i,$$

where each Ω^i is a nonempty finite set that we call a *universe* of $\mathfrak{X}$. We call the elements of Ω^i *symbols*, and the elements of $\Omega^{\mathcal{I}}$ *strings*. $\Diamond$

Definition 3.2 (Substrings)**.** Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex. For $x \in \Omega^{\mathcal{I}}$ and $S \subseteq \mathcal{I}$, we define the *substring* $x_S \coloneqq (x_i)_{i \in S} \in \Omega^S$. We sometimes call x_S an *S-string*. $\Diamond$

Definition 3.3 (Restricted complexes)**.** Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex. For a set $\mathcal{J}$ and a tuple of pairwise disjoint subsets $\mathcal{S} = (S_j \subset \mathcal{I})_{j \in \mathcal{J}}$, we define the *$\mathcal{S}$-restricted* $\mathcal{J}$-indexed simplicial complex $\mathfrak{X}[\mathcal{S}]$. We preliminarily define the universes of this complex to be Ω^{S_j} (for $j \in \mathcal{J}$). Then, the distribution $\mathfrak{X}[\mathcal{S}]$ is given by drawing $\boldsymbol{x} \sim \mathfrak{X}$ and outputting the tuple of substrings $(\boldsymbol{x}_{S_j})_{j \in \mathcal{J}}$. Finally, in case the resulting indexed complex has symbols that never occur in $\operatorname{supp}(\mathfrak{X}[\mathcal{S}])$, we delete these symbols from their universes. $\Diamond$

Remark 3.4. Dikstein and Dinur [DD24c] call the restricted complex $\mathfrak{X}[\mathcal{S}]$ the *colored faces complex*, and the special case of $\mathfrak{X}[S_1, S_2]$ the *colored swap walk*. This walk was also studied in [DD19; GLL22]. $\Diamond$

Definition 3.5 (Words/faces)**.** In the setting of the previous definitions, if $x \in \operatorname{supp}(\mathfrak{X})$, we call x_S a *word*[15] or a *face* of $\mathfrak{X}$. We also call x_S an S-word (to emphasize S) or an $|S|$-word (to emphasize the size of S). Hence, the set of S-faces is precisely $\operatorname{supp}(\mathfrak{X}[S])$. $\Diamond$

Two indexed simplicial complexes are isomorphic if one can be obtained from the other by re-indexing and renaming symbols. Formally:

Definition 3.6 (Isomorphism of indexed simplicial complexes)**.** Let $\mathfrak{X}$ (respectively, $\widetilde{\mathfrak{X}}$) be an $\mathcal{I}$-indexed (respectively, $\widetilde{\mathcal{I}}$-indexed) simplicial complex with universes $(\Omega^i)_{i \in \mathcal{I}}$ (respectively, $(\widetilde{\Omega}^j)_{j \in \widetilde{\mathcal{I}}}$). We say they are

[14] These are also called *numbered* complexes by Lannér [Lan50] and *completely balanced* complexes by Stanley [Sta79].

[15] We apologize for renaming "faces" in this way, particularly since the term could be confused with the group-theoretic notion of "word". Nevertheless, we find it has the right mnemonic value; as in English, every word is a string, but not every string is a word.

isomorphic (notated $\mathfrak{X} \cong \widetilde{\mathfrak{X}}$) if there is a bijection $\iota : \mathcal{I} \to \widetilde{\mathcal{I}}$ and, for each $i \in \mathcal{I}$ a bijection $\phi_i : \Omega^i \to \widetilde{\Omega}^{\iota(i)}$ (from $\{i\}$-words to $\{\iota(i)\}$-words) such that for $\boldsymbol{x} \sim \mathfrak{X}$, the string-valued random variable $(\phi_i(\boldsymbol{x}_i))_{\iota(i)\in\widetilde{\mathcal{I}}}$ is distributed as $\widetilde{\mathfrak{X}}$. $\Diamond$

Definition 3.7 (Clique complex)**.** We say the $\mathcal{I}$-indexed simplicial complex $\mathfrak{X}$ is a *clique complex* provided the following holds for all substrings x_S: if x_E is a word for all $E \subseteq_2 S$, then x_S is a word. $\Diamond$

Remark 3.8. An $\mathcal{I}$-indexed clique complex is equivalent to a (weighted, finite) *incidence geometry*; see, e.g., [BC13]. Almost all indexed complexes we deal with will be clique complexes. $\Diamond$

Definition 3.9 (Links/pinnings)**.** Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex, $S \subseteq \mathcal{I}$, and w an S-word. Write $\overline{S} = \mathcal{I} \setminus S$. We define an associated *pinned* (or *link*) complex $\mathfrak{X}_w$. This is the $\overline{S}$-indexed simplicial complex whose universes are (preliminarily) defined to be $(\Omega^i)_{i\in\overline{S}}$, and whose distribution $\mathfrak{X}_w$ is defined by drawing $\boldsymbol{x} \sim \mathfrak{X}$ conditioned on $\boldsymbol{x}_{\mathcal{I}\setminus S} = w$ and outputting $\boldsymbol{x}_{\overline{S}}$. As in the definition of restricted complexes, in case this indexed complex has any symbols that occur with zero probability, they are deleted from their universe. $\Diamond$

We now effectively give the definitions of "r-faces complex" from [DD24a] (see also [BLM24]):

Definition 3.10 (r-triwords (and biwords))**.** Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex, let $1 \leq r \leq |\mathcal{I}|/3$ and let $S_1, S_2, S_3 \sqsubset_r \mathcal{I}$. We then call any element of $\operatorname{supp}(\mathfrak{X}[S_1, S_2, S_3])$ an *r-triword* of $\mathfrak{X}$. We can alternatively think of an r-triword as an unordered triple of r-words (with disjoint index sets). We write $T_r(\mathfrak{X})$ for the probability distribution on r-triwords obtained by drawing $\boldsymbol{x} \sim \mathfrak{X}$, independently drawing $\boldsymbol{S}_1, \boldsymbol{S}_2, \boldsymbol{S}_3 \sqsubset_r \mathcal{I}$ uniformly at random, and outputting $\{\boldsymbol{x}_{\boldsymbol{S}_1}, \boldsymbol{x}_{\boldsymbol{S}_2}, \boldsymbol{x}_{\boldsymbol{S}_3}\}$. We may analogously define biwords (unordered *pairs* of words) and the biword distribution $E_r(\mathfrak{X})$. And finally, we write $V_r(\mathfrak{X})$ for the analogous distribution on words ("mono-words"). $\Diamond$

Remark 3.11. Identifying $V_r(\mathfrak{X}), E_r(\mathfrak{X}), T_r(\mathfrak{X})$ with their support, we may regard $\{\emptyset, V_r(\mathfrak{X}), E_r(\mathfrak{X}), T_r(\mathfrak{X})\}$ as a *non-indexed* simplicial complex. More precisely, it is a weighted, pure, rank-3 (dimension-2) simplicial complex. We may call it the *r-triword complex* of $\mathfrak{X}$. $\Diamond$

3.2 Spectral expansion

Definition 3.12. Let $\mathfrak{X}$ be an $\{i_1, i_2\}$-indexed edge complex (a.k.a. a joint distribution on $\Omega^{i_1} \times \Omega^{i_2}$). We define the *averaging operator* $A : \mathbb{R}^{\Omega^{i_1}} \to \mathbb{R}^{\Omega^{i_2}}$ via

$$(Af)(v_2) \coloneqq \mathop{\mathbb{E}}_{v_1 \sim \mathfrak{X}_{v_2}[\ell_1]}[f(v_1)].$$

We say $\mathfrak{X}$ is an *ϵ-bipartite (spectral) expander* if the second largest singular value of A is at most ϵ. (The ordering on i_1, i_2 does not matter; the converse averaging operator is the transpose $A^\intercal$, which has the same (nonzero) singular values.) $\Diamond$

Fact 3.13. *In the setup of the preceding definition, $\mathfrak{X}$ is the* uniform *distribution on the product set $\Omega^{i_1} \times \Omega^{i_2}$ if and only if it is a 0-bipartite expander.*

The following notion was studied in [DD19] and formally defined in [GLL22, §3]:

Definition 3.14. Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex and $\epsilon > 0$. $\mathfrak{X}$ is *ϵ-product* if for every $\ell_1 \neq \ell_2 \in \mathcal{I}$, $W \subseteq \mathcal{I} \setminus \{\ell_1, \ell_2\}$, and W-word w, $\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}]$ is an ϵ-bipartite expander. $\Diamond$

Remark 3.15. ϵ-productness, along with cliqueness, are *heritable* properties of simplicial complexes: If a simplicial complex satisfies one of these properties, then so do all its links. $\Diamond$

We also define the following weaker notion:

Definition 3.16. Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex ($|\mathcal{I}| \geq 2$) and $\epsilon > 0$. $\mathfrak{X}$ is an *ϵ-local bipartite expander* if for every $\ell_1 \neq \ell_2 \in \mathcal{I}$, and $(\mathcal{I} \setminus \{\ell_1, \ell_2\})$-word w, $\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}]$ is an ϵ-bipartite expander. $\Diamond$

We then have the following trickle-down type theorem for establishing ϵ-productness:

Theorem 3.17 ([DD19, Corollary 7.6])**.** *Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex, $|\mathcal{I}| = n$. Let $\epsilon < \frac{1}{2}$ and suppose that:*

1. *$\mathfrak{X}$ is an $(\frac{\epsilon}{(n-1)\epsilon+1})$-local bipartite expander.*

2. *For every $W \subseteq_{\leq n-2} \mathcal{I}$ and W-word w, $\mathfrak{X}_w$ is connected.*

Then $\mathfrak{X}$ is ϵ-product.

3.3 Triword expansion

1-cohomology of a triangle complex over a group Γ, and its quantitative form 1-coboundary expansion, are defined in §2.2. In the remainder of the paper, we will only need this definition in the context of the r-triword complex. For clarity, we explicitly give the definition here:

Definition 3.18 (Antisymmetric biword labelings)**.** Let Γ be a group, $\mathfrak{X}$ an $\mathcal{I}$-indexed simplicial complex, and $r \leq \frac{1}{2}|\mathcal{I}|$. A *biword labeling* is a function $f : \vec{E}_r(\mathfrak{X}) \to \Gamma$ satisfying the antisymmetry condition that for every $(s_1, s_2) \in \vec{E}_r(\mathfrak{X})$, $f(s_1, s_2) = f(s_2, s_1)^{-1}$. ♢

Definition 3.19 (Triword expansion)**.** Let Γ be a group, $\mathfrak{X}$ an $\mathcal{I}$-indexed simplicial complex, $r \leq \frac{1}{3}|\mathcal{I}|$, and $\epsilon > 0$. $\mathfrak{X}$ has *r-triword expansion (at least) ϵ over* Γ if for every biword-labeling $f : \vec{E}_r(\mathfrak{X}) \to \Gamma$, there exists a word-labeling $g : V_r(\mathfrak{X}) \to \Gamma$ such that

$$\Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2)\sim\vec{E}_r(\mathfrak{X})}\left[g(\boldsymbol{s}_1)g(\boldsymbol{s}_2)^{-1} \neq f(\boldsymbol{s}_1,\boldsymbol{s}_2)\right] \leq \frac{1}{\epsilon} \cdot \Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2,\boldsymbol{s}_3)\sim\vec{T}_r(\mathfrak{X})}[f(\boldsymbol{s}_1,\boldsymbol{s}_2)f(\boldsymbol{s}_2,\boldsymbol{s}_3)f(\boldsymbol{s}_3,\boldsymbol{s}_1) \neq \mathbb{1}].$$

When r is omitted, we assume it is 1. ♢

Remark 3.20. The property that we call "triword expansion" is called "swap coboundary expansion" in [DD24c; DDL24] and "unique-games coboundary expansion" in [BM24; BLM24]. ♢

3.4 Direct-product testing from triword expansion

Bafna and Minzer [BM24] and Dikstein and Dinur [DD24a] determined "high-dimensional expansion conditions" on $\mathcal{D}$ that sufficed for the associated agreement tester to achieve good soundness. We present the following version of their results, which can be viewed as a version of [DD24a, Theorem 3.1], together with background results; see also [BMVY25, Theorem A.1]. We comment that conditions 4, 5 below correspond to Properties (2), (1) in §1, respectively.

Theorem 3.21 (Sufficient conditions for direct-product testability, combining [DD24a; DD24b])**.** *Given $\delta > 0$, there is $m = \mathrm{poly}(1/\delta)$, such that if $k \geq \exp(\mathrm{poly}(1/\delta))$, $r \geq \mathrm{poly}(k)$, $n \geq \exp(\mathrm{poly}(r))$, and $\gamma < 1/n^2$, then the following holds:*

Suppose $\mathfrak{X}$ is a rank-n complex (not necessarily indexed) satisfying the following:

1. Spectral expansion:
   For any pinning w of size $n-2$, the graph $\mathfrak{X}_w$ has normalized 2nd eigenvalue at most γ.

2. Connectivity:
   For every S-word w with $|S| \leq r$, the graph $(V_r(\mathfrak{X}_w), E_r(\mathfrak{X}_w))$ is connected.

3. Clique complex:
   $\mathfrak{X}$ is a clique complex.

4. Local triword expansion:
   For every 1-word w ("vertex"), $\mathfrak{X}_w$ has r-triword expansion at least $\exp(-r^{.99})$ over $\mathfrak{S}(m_0)$ for all $1 < m_0 \leq m$.

5. Trivial global cohomology:
   $\mathfrak{X}$ has nonzero 1-triword expansion (i.e., trivial 1-cohomology) over $\mathfrak{S}(m_0)$ for all $1 < m_0 \leq m$.

Then the agreement tester with distribution given by $\mathfrak{X}$'s rank-k faces achieves δ-soundness for any alphabet Σ, in the sense of Definition 2.1.

In Appendix B below, we discuss why this theorem follows from the works of [DD24a; DD24b].

3.5 Strongly symmetric complexes

We define a notion of strong symmetry in indexed simplicial complexes, which, in mathematical terms, posits that the automorphism group of a complex acts transitively on words:

Definition 3.22. Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex. We say $\mathfrak{X}$ is *strongly symmetric* if for every pair of words $x = (x_i)_{i\in\mathcal{I}}, x' = (x'_i)_{i\in\mathcal{I}} \in \mathrm{supp}(\mathfrak{X})$, there exist bijections $\phi_i : \Omega^i \to \Omega^i$ such that:

1. $\phi_i(x_i) = x'_i$ for every $i \in \mathcal{I}$.
2. $(\phi_i)_{i\in\mathcal{I}}$ form an isomorphism from $\mathfrak{X}$ to itself (together with the identity map $\iota : \mathcal{I} \to \mathcal{I}$), in the sense of Definition 3.6. That is, for $\boldsymbol{x} = (\boldsymbol{x}_i)_{i\in\mathcal{I}} \sim \mathfrak{X}$, $(\phi_i(\boldsymbol{x}_i))_{i\in\mathcal{I}}$ is also distributed as $\mathfrak{X}$. $\Diamond$

Remark 3.23. This notion of strong symmetry in an indexed simplicial complex is stronger than the standard definition of strong symmetry in simplicial complexes. (The standard definition allows a single bijection $\Phi : V(\mathfrak{X}) \to V(\mathfrak{X})$ that need not preserve the individual parts Ω^i.) We use this condition because coset complexes do satisfy it and it is easier to state. $\Diamond$

Definition 3.24 (Diameter)**.** Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex. For $s, s' \in V(\mathfrak{X})$, the *distance* between s and s' is the minimum $k \in \mathbb{N}$ such that there exist $(s^{(i)} \in V(\mathfrak{X}))_{i\in[0\,..\,k]}$ such that $s^{(0)} = s$, $s^{(k)} = s'$, and for every $i \in [k]$, $(s^{(i-1)}, s^{(i)}) \in \vec{E}(\mathfrak{X})$. The *diameter* $\mathrm{diam}(\mathfrak{X})$ of $\mathfrak{X}$ is the maximum distance between any $s, s' \in V(\mathfrak{X})$. $\Diamond$

3.6 Triword expansion of linearly ordered complexes

We next turn to the question of proving that a simplicial complex is a triword expander. We will need to do this for the local Kaufman–Oppenheim complex in order to fulfill Condition 4 of Theorem 3.21. We present a sufficient condition for triword expansion in complexes with "linearly ordered" geometry in Theorem 3.30 below. This again follows the proof of [BLM24, §3-5] which is specialized to the complex of Chapman and Lubotzky [CL25], but our statement is generic and can be applied to the complex of Kaufman and Oppenheim [KO18] as well.

Definition 3.25 (Productization)**.** Let $\mathfrak{X}$ an $\mathcal{I}$-indexed simplicial complex, and $S_1, S_2 \sqsubset \mathcal{I}$. We say $\mathfrak{X}$ *productizes over* S_1 *and* S_2, denoted $S_1 \perp_{\mathfrak{X}} S_2$, if for $(\boldsymbol{s}_1, \boldsymbol{s}_2) \sim \mathfrak{X}[S_1, S_2]$, $\boldsymbol{s}_1$ and $\boldsymbol{s}_2$ are independent. (Equivalently, the joint distribution factors into a product distribution, $\mathfrak{X}[S_1, S_2] = \mathfrak{X}[S_1] \times \mathfrak{X}[S_2]$.) $\Diamond$

Definition 3.26 (Ordering on sets)**.** Let $n \in \mathbb{N}$. For two sets $S_1, S_2 \subseteq [1\,..\,n]$, we say S_1 *precedes* S_2, denoted $S_1 \prec S_2$, if $\max S_1 < \min S_2$. In particular, $S_1 \prec S_2 \implies S_1 \cap S_2 = \emptyset$ (so that $S_1, S_2 \sqsubset [1\,..\,n]$). We write $S_1 \prec \cdots \prec S_k \sqsubset [1\,..\,n] \setminus W$ if each $S_i \subset [1\,..\,n] \setminus W$ and $S_i \prec S_{i+1}$ for every $1 \le i \le k-1$. $\Diamond$

Definition 3.27 (Separated set)**.** Let $\sigma \in \mathbb{N}$. We say $R \subseteq [1\,..\,n]$ is σ*-separated* if for every $\ell_1 \neq \ell_2 \in R$, we have $|\ell_2 - \ell_1| \ge \sigma$. $\Diamond$

Note that if $S \subseteq R$ and R is σ-separated, then S is also σ-separated.

Definition 3.28 (Productization on the line)**.** Let $\mathfrak{X}$ be an $[1\,..\,n]$-indexed simplicial complex. We say $\mathfrak{X}$ is *productizing on the line (PoL)* if the following is true: For every $W \subseteq [1\,..\,n]$ and $S_1 \prec S_2 \sqsubset [1\,..\,n] \setminus W$, if $W \cap (\max S_1\,..\,\min S_2) \neq \emptyset$, then for every W-word w, $S_1 \perp_{\mathfrak{X}_w} S_2$. $\Diamond$

Remark 3.29. The foregoing three definitions are the main reason why we need to consider simplicial complexes indexed by $[1\,..\,n]$, which has "the geometry of the line". $\Diamond$

For completeness, we prove the following theorem in Appendix A, though it is essentially wholly due to [BLM24]:

Theorem 3.30 (Proving triword expansion for "linearly ordered" complexes, abstracted from [BLM24])**.** *There exist absolute constants $K > 1$ and $\epsilon_0 > 0$ such that for every $r \in \mathbb{N}$, there exist $\xi > 0$ and $n_0 \in \mathbb{N}$ such that for every $n \ge n_0 \in \mathbb{N}$, the following holds. Let $\mathfrak{X}$ be a $[1\,..\,n]$-indexed simplicial complex and Γ an arbitrary group. Suppose that $\mathfrak{X}$ satisfies the following conditions:*

1. Spectral expansion:
   $\mathfrak{X}$ is $\epsilon_0 r^{-2}$-product.

2. Productization:
   $\mathfrak{X}$ productizes on the line, in the sense of Definition 3.28.

3. Symmetry:
   $\mathfrak{X}$ is strongly symmetric, in the sense of Definition 3.22.

4. Coboundary expansion of separated singleton restrictions:
   For every $W \subseteq [1\,..\,n]$ and $\ell_1 < \ell_2 < \ell_3 \in [1\,..\,n]$ such that $W \cap [\ell_1\,..\,\ell_3] = \emptyset$ and $\{\ell_1, \ell_2, \ell_3\}$ is ξn-separated, for every W-word w, the complex $\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}, \{\ell_3\}]$ has 1-triword expansion at least β over Γ.

5. Diameter of separated singleton restrictions:
   For every $W \subseteq [1\,..\,n]$ and $\ell_1 < \ell_2 \in [1\,..\,n]$ such that $W \cap [\ell_1\,..\,\ell_2] = \emptyset$ and $\{\ell_1, \ell_2\}$ is ξn-separated, for every W-word w, the complex $\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}]$ has diameter at most C_0.

Then for every $W \subset_{\leq r} [1\,..\,n]$ and W-word w, $\mathfrak{X}_w$ has r-triword expansion at least $(K\beta/C_0)^{O(r^{0.67})}$ over Γ.

We note that there are a few minor differences in our theorem statement from that of [BLM24]. In particular: we have a quantitatively spectral expansion hypothesis ($O(r^{-2})$ vs. $\exp(-\mathrm{poly}(r))$), have a quantitatively stronger conclusion (exponent of $r^{0.67}$ vs. $r^{0.99}$), and do not need to account for additive error. However, we emphasize that these improvements arise only from some slack in their proof and we do not claim them as contributions.

3.7 Subgroups and cosets

For two groups H and G, $H \leq G$ denotes that H is a subgroup of G. For $g \in G$, $gH \coloneqq \{gh : h \in H\}$ is the *left coset* corresponding to g. $G/H \coloneqq \{gH : g \in G\}$ denotes the set of all left cosets. More generally, for two subgroups $H_1, H_2 \leq G$, we write $H_1H_2 \coloneqq \{h_1h_2 : h_1 \in H_1, h_2 \in H_2\}$.

Definition 3.31. Let G be a finite group and $\mathcal{I}$ a finite set. An *$\mathcal{I}$-indexed subgroup family* for G is a collection $\mathcal{H} = (H_i \leq G)_{i\in\mathcal{I}}$ of subgroups indexed by $\mathcal{I}$. For $S \subseteq \mathcal{I}$, we write $H_S \coloneqq \bigcap_{i\in S} H_i$, with $H_\emptyset \coloneqq G$. An $\mathcal{I}$-indexed subgroup family $\mathcal{H} = (H_i \leq G)_{i\in\mathcal{I}}$ for G and $\mathcal{I}'$-indexed subgroup family $\mathcal{H}' = (H'_i \leq G)_{i\in\mathcal{I}'}$ are *isomorphic* if there exists a re-indexing bijection $\iota : \mathcal{I} \to \mathcal{I}'$ and a group isomorphism $\phi : G \xrightarrow{\sim} G'$ such that $\phi(H_i) = H'_{\iota(i)}$ for every $i \in \mathcal{I}$. ◊

We collect here some trivial but useful facts about cosets. Mutually intersecting cosets of different subgroups are cosets of the intersection subgroup:

Proposition 3.32 ([Rot95, Ex. 2.26])**.** *For a group G, an $\mathcal{I}$-indexed subgroup family $\mathcal{H} = (H_i \leq G)_{i\in\mathcal{I}}$, and a tuple of cosets $(C_i \in G/H_i)_{i\in\mathcal{I}}$, either $\bigcap_i C_i = \emptyset$ or $\bigcap_i C_i$ is a coset of the intersection subgroup $\bigcap_i H_i$.*

Two cosets intersect if their ratio is a product:

Proposition 3.33 ([HS24, Observation 3.2])**.** *For all groups G and subgroups $H_1, H_2 \leq G$ and $g_1, g_2 \in G$, $g_1H_1 = g_2H_2$ iff $g_1^{-1}g_2 \in H_1H_2$.*

Taking $H_1 = H_2$, we get:

Corollary 3.34 (also [Rot95, Thm. 2.8])**.** *For all groups $H \leq G$ and $g_1, g_2 \in G$, $g_1H = g_2H$ iff $g_1^{-1}g_2 \in H$.*

and therefore:

Corollary 3.35. *For all groups $H \leq G$, $g \in G$, and $C \in G/H$, $g \in C$ if and only if $C = gH$.*

Proof. Obviously, if $C = gH$ then $g = g\mathbb{1} \in C$. Conversely, if $C = g'H$ and $g \in C$, there exists $h \in H$ such that $g'h = g$ and therefore $(g')^{-1}g \in H$; then apply the previous proposition. □

Also, the cosets G/H partition G:

Proposition 3.36 ([Rot95, Thm. 2.9])**.** *For all groups $H \leq G$, any two cosets $C, C' \in G/H$ are either identical or disjoint.*

We have the following product formula:

Proposition 3.37 ([Rot95, Thm. 2.20])**.** *For all groups $H_1, H_2 \leq G$, $|H_1H_2| \cdot |H_1 \cap H_2| = |H_1| \cdot |H_2|$.*

G acts on G/H by left multiplication:

Fact 3.38. *For all groups $H \leq G$, cosets $C \in G/H$ of H, and elements $g \in G$, $gC \in G/H$, i.e., gC is again a coset of H. Further, this is a group action, i.e., $\mathbb{1}C = C$ and $g_1(g_2C) = (g_1g_2)C$.*

Proof. We have $g(g'H) = (gg')H$, $\mathbb{1}(g'H) = g'H$, and $(g_1g_2)(g'H) = (g_1g_2g')Hg_1(g_2(g'H))$. ☐

3.8 Coset complexes

We now define a way to construct an indexed simplicial complex from an indexed subgroup family, dating back at least as far as the work of Lannér [Lan50]:

Definition 3.39. Let G be a finite group and $\mathcal{H} = (H_i \leq G)_{i \in \mathcal{I}}$ an $\mathcal{I}$-indexed subgroup family. The *G-coset complex* $\mathfrak{CC}(G; \mathcal{H})$ is the $\mathcal{I}$-indexed simplicial complex with universes $\Omega^i := G/H_i$ and the following distribution $\mathfrak{CC}(G; \mathcal{H})$: Draw $\boldsymbol{g} \sim G$ uniformly and output the string of cosets $(\boldsymbol{g}H_i)_{i \in \mathcal{I}}$. ◊

Remark 3.40. This definition goes under different names in the literature; for instance, if one ignores the probabilistic weighting, it is called a finite *coset incidence system* in [BC13]. The definition in [KO18; KO23] of coset complexes is slightly different than ours (which is consistent with [Lan50; Gar79; Tit86; OP22; HS24; OS25]). The [KO18; KO23] definition takes the result of Definition 3.39, which is not necessarily a clique complex, and then adds extra faces corresponding to all cliques. The definitions coincide when the complex defined in Definition 3.39 is already a clique complex. ◊

We now collect several basic properties of coset complexes which we shall use repeatedly:

Proposition 3.41. *For $S \subseteq \mathcal{I}$, an S-string $w = (C_i \in G/H_i)_{i \in S}$ is an S-word (that is, $w \in \operatorname{supp}((\mathfrak{CC}(G; \mathcal{H}))[S])$) iff the cosets intersect mutually; i.e., $\bigcap_{i \in S} C_i \neq \emptyset$.*

Proof. w is an S-word if and only if there exists $g \in G$ such that $C_i = gH_i$ for every $i \in \mathcal{I}$; then use Corollary 3.35. ☐

Proposition 3.42. $\mathfrak{CC}(G; \mathcal{H})$ *is always strongly symmetric.*

Proof. Let $x = (C_i)_{i \in \mathcal{I}}, x' = (C'_i)_{i \in \mathcal{I}} \in \operatorname{supp}(\mathfrak{CC}(G; \mathcal{H}))$ be words. By Proposition 3.41, $\bigcap_{i \in \mathcal{I}} C_i \neq \emptyset$ and $\bigcap_{i \in \mathcal{I}} C'_i \neq \emptyset$. By Proposition 3.32, there exist $g, g' \in G$ such that $\bigcap_{i \in \mathcal{I}} C_i = gH_{\mathcal{I}}$ and $\bigcap_{i \in \mathcal{I}} C'_i = g'H_{\mathcal{I}}$. (Here, we used the notation $H_{\mathcal{I}} = \bigcap_{i \in \mathcal{I}} H_i$.) Hence $g \in C_i$ and $g' \in C'_i$ for every $i \in \mathcal{I}$. So by Corollary 3.35, $C_i = gH_i$ and $C'_i = g'H_i$ for every $i \in \mathcal{I}$.

We use the bijections $\phi_i : G/H_i \to G/H_i$ defined by $C \mapsto g'g^{-1}C$. (If $C \in G/H_i$, then $g'g^{-1}C \in G/H_i$ by Claim 3.38. Further, ϕ_i is a bijection on G/H_i because its inverse is $C \mapsto g(g')^{-1}C$.) By definition, $\phi_i(C_i) = g'g^{-1}(gH_i) = g'H_i = C'_i$. Finally, these maps preserve the distribution $\mathfrak{CC}(G; \mathcal{H})$ because for $\boldsymbol{g} \sim G$, $\boldsymbol{g}$ and $g'g^{-1}\boldsymbol{g}$ have the same distribution; therefore, $(\boldsymbol{g}H_i)_{i \in \mathcal{I}}$ and $(g'g^{-1}\boldsymbol{g}H_i)_{i \in \mathcal{I}}$ also have the same distribution. ☐

Remark 3.43. A coset clique complex is equivalent to a (finite) *coset geometry*; see, e.g., [BC13]. ◊

A useful property of G-coset complexes is that any (nontrivial) pinning is naturally a G'-coset complex for $G' \leq G$:

Proposition 3.44 (e.g., [HS24, Lemma 3.4])**.** *For any $W \subseteq \mathcal{I}$ and W-word w, we have the isomorphism*

$$(\mathfrak{CC}(G; \mathcal{H}))_w \cong \mathfrak{CC}(H_S; (H_{W \cup \{i\}})_{i \in \mathcal{I} \setminus W}).$$

(The left-hand side is a $\mathcal{I} \setminus W$-indexed simplicial complex and the index mapping $\iota : \mathcal{I} \setminus W \to \mathcal{I} \setminus W$ is the identity function.)

Also, if $N \trianglelefteq G$ is a normal subgroup of G, $\mathfrak{CC}(G; \mathcal{H})$ is isomorphic to a G/N-quotient complex, assuming that $N \leq H_i$ for every $i \in \mathcal{I}$. We present this proposition in the following equivalent form:

Proposition 3.45. *Let G be a group, $L \le G$, $\pi : G \twoheadrightarrow L$ a surjective map onto L, and $\mathcal{H} = (H_i \le G)_{i\in\mathcal{I}}$ an indexed subgroup family satisfying $\ker(\pi) \le H_{\mathcal{I}}$. Then $\mathfrak{CC}(G;\mathcal{H}) \cong \mathfrak{CC}(L;(\pi(H_i))_{i\in\mathcal{I}})$.*

See Appendix D below for a (standard) proof. One useful special case of Proposition 3.45 is when the surjective map is actually an isomorphism:

Corollary 3.46. *Suppose we have isomorphic indexed subgroup families, in the sense of Definition 3.31: G and G' are groups, $\phi : G \xrightarrow{\sim} G'$ is a group isomorphism, $\mathcal{H} = (H_i \le G)_{i\in\mathcal{I}}$ an $\mathcal{I}$-indexed subgroup family, and $\iota : \mathcal{I} \to \mathcal{I}'$ a reindexing function. Then $\mathfrak{CC}(G;\mathcal{H}) \cong \mathfrak{CC}(G';(\phi(H_i))_{i\in\mathcal{I}})$.*

Proof. Use the prior Proposition 3.45 and the fact that $\ker(\phi)$ is trivial. □

This is the only possible application of Proposition 3.45 when the intersection of the subgroups $H_{\mathcal{I}}$ is trivial, which is the typical situation (for instance, this is the case in the local Kaufman–Oppenheim complex). However, cases where the intersection of subgroups is nontrivial do arise. One typical case is when we restrict the subgroup family $(H_i)_{i\in\mathcal{I}}$ to a subfamily $(H_i)_{i\in\mathcal{J}}$ for some $\mathcal{J} \subsetneq \mathcal{I}$. We give a corollary which is useful for this case in Proposition 4.6 below.

Finally, we prove the following simple condition for productization in Appendix D:

Proposition 3.47. *Let G be a group and $\mathcal{H} = (H_i \le G)_{i\in\mathcal{I}}$ an $\mathcal{I}$-indexed subgroup family. Suppose $W, S_1, S_2 \sqsubseteq \mathcal{I}$. Then for every W-word w, $S_1 \perp_{\mathfrak{CC}(G;\mathcal{H})_w} S_2$ if and only if $H_W = H_{W\cup S_1} H_{W\cup S_2}$. In particular, taking $W = \emptyset$, we get that $S_1 \perp_{\mathfrak{CC}(G;\mathcal{H})} S_2$ if and only if $G = H_{S_1} H_{S_2}$.*

All rings are commutative rings with identity. For a field $\mathbb{F}$, $\mathbb{F}[X]$ denotes the ring of polynomials over $\mathbb{F}$ in a single indeterminate X. $\deg(f)$ is the degree of a polynomial $f \in \mathbb{F}[X]$, with the convention $\deg(0) \coloneqq -\infty$.

For a group G and two elements $g, h \in G$, the *commutator* of g and h is

$$[g,h] \coloneqq g^{-1}h^{-1}gh \tag{3.48}$$

so that

$$gh = [g,h]hg. \tag{3.49}$$

(In particular, g and h commute in G iff $[g,h] = 1$.)

3.9 The graded unipotent group and the local Kaufman–Oppenheim complex

For a field $\mathbb{F}$, an indeterminate X, and $d \in \mathbb{N}$, let $\mathbb{F}^{\le n}[X]$ denote the subset of polynomials over X of degree at most d.

Definition 3.50 (Graded unipotent group)**.** Let $n, \kappa \in \mathbb{N}$, and $\mathbb{F}$ be a field. The group $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ is defined by the following presentation:

- The generators are symbols $e_{i,j}(f)$, where $i < j \in [0\,..\,n]$ and $f \in \mathbb{F}^{\le \kappa(j-i)}[X]$.
- The relations are:
  - The *linearity* relations: For every $i < j \in [0\,..\,n]$ and $f, g \in \mathbb{F}^{\le \kappa(j-i)}[X]$,
  $$e_{i,j}(f) \cdot e_{i,j}(g) = e_{i,j}(f+g).$$
  - The *nontrivial commutator* relations: For every $i < j < k \in [0\,..\,n]$ and $f \in \mathbb{F}^{\le \kappa(j-i)}[X], g \in \mathbb{F}^{\le \kappa(k-j)}[X]$,
  $$[e_{i,j}(f), e_{j,k}(g)] = e_{i,k}(fg).$$
  - The *trivial commutator* relations: For every $i < j \in [0\,..\,n]$ and $k < \ell \in [0\,..\,n]$ with $j \ne k$ and $i \ne \ell$ and $f \in \mathbb{F}^{\le \kappa(j-i)}[X], g \in \mathbb{F}^{\le \kappa(\ell-k)}[X]$,
  $$[e_{i,j}(f), e_{k,\ell}(g)] = 1.$$

These relations are called the *Steinberg relations* for $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$. (See also Examples 2.11 to 2.13 and Figure 2.) ◊

Note that the nontrivial commutator relations are well-defined because $\deg(fg) \le \deg(f) + \deg(g) \le \kappa(j-i) + \kappa(k-j) = \kappa(k-i)$.

Remark 3.51. The group $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ is denoted $\mathrm{Unip}_n(\mathbb{F}[t];\kappa)$ in [KOW25, §2.7]. When we omit κ, we assume it is 3; this is the case for which [KOW25] prove trivial cohomology of the global complex, and therefore the only case where we currently get our downstream application of direct-product testing and PCPs. Our results on local expansion work for all values of κ, including the case $\kappa = 1$, which corresponds to the complex originally defined by [KO21; KO23]. We do not currently have direct-product testers in the $\kappa = 1$ case because trivial cohomology of the global complex is not (yet) known. $\Diamond$

We now define a collection of subgroups of the group $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$:

Definition 3.52 (Staircase subgroups). For $n, \kappa \in \mathbb{N}$, $\mathbb{F}$ a field, and $\ell \in [1\,..\,n] = [1\,..\,n]$, we define the "staircase subgroup" $\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell) \le \mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ as the subgroup generated by the elements $e_{i,j}(f)$ ($i < j \in [0\,..\,n]$ and $f \in \mathbb{F}^{\le \kappa(j-i)}[X]$) where $(i\,..\,j] \not\ni \ell$. The staircase subgroups form an $[1\,..\,n]$-indexed subgroup family

$$(\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell) \le \mathrm{Unip}^{\kappa}_{\mathbb{F}}(n))_{\ell \in [1\,..\,n]}.$$

Thus, (as in Definition 3.31,) we can also define the intersection subgroups, for every $S \subseteq [1\,..\,n]$,

$$\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;S) \coloneqq \bigcap_{\ell \in S} \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell).$$

(By a slight abuse of notation, we also refer to all these subgroups as "staircase subgroups".) $\Diamond$

To interpret the condition $(i\,..\,j] \not\ni \ell$, it is helpful to consider the following trivial fact:

$$\forall i < j \in [0\,..\,n], \ell \in [1\,..\,n], \quad (i\,..\,j] \not\ni \ell \iff (j < \ell \vee i \ge \ell). \tag{3.53}$$

If we view the elements of $[0\,..\,n]$ as $n+1$ "wires" and the elements of $[1\,..\,n]$ as n interspersed "moats" (so that moat i is between wires $i-1$ and i), this condition equivalently captures whether a "gate" stretching from wire i to wire j crosses moat ℓ. We call $j - i$ the "height" of the gate (i,j).

See §4 below for intuition for and structure characterizations of the group $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ and its subgroups $\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;a)$.

3.10 The Kaufman–Oppenheim complex

Having the stair allows us to then define the coset complex:

Definition 3.54. For $n, \kappa \in \mathbb{N}$, we define the coset complex $\mathfrak{LA}^{\kappa}_{\mathbb{F}}(n) \coloneqq \mathfrak{CC}(\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n); (\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;a))_{a \in [1\,..\,n]})$. We call this the *local Kaufman–Oppenheim complex* (on $n+1$ "wires"). $\Diamond$

For $n, \kappa, s \in \mathbb{N}$ and finite field $\mathbb{F}$, Kaufman and Oppenheim [KO23] defined a certain coset complex $\mathfrak{A}^{\kappa,s}_n(\mathbb{F})$, which we call a *global Kaufman–Oppenheim complex.*[16] This complex is a coset complex with rank $n+1$ / dimension n. In the interest of notational simplicity, we postpone the explicit definition of the global complex $\mathfrak{A}^{\kappa,s}_n(\mathbb{F})$ to Appendix C; in the body of our paper, we use only the following four statements about this complex, all of which are proven explicitly or implicitly in [KO23]:

Theorem 3.55 (Vertex links of global Kaufman–Oppenheim complex). *For every field $\mathbb{F}$, $n, \kappa \in \mathbb{N}$, $s \ge \kappa n$, 1-word (vertex) $w \in V(\mathfrak{A}^{\kappa,s}_n(\mathbb{F}))$, $v \in \mathrm{supp}([i])$, the link of v in the global Kaufman–Oppenheim complex is isomorphic to the local Kaufman–Oppenheim complex: $\mathfrak{A}^{\kappa,s}_n(\mathbb{F})_v \cong \mathfrak{LA}^{\kappa}_{\mathbb{F}}(n)$.*

Theorem 3.56 (Global Kaufman–Oppenheim complex is clique). *For every field $\mathbb{F}$, $n, \kappa \in \mathbb{N}$, and $s \ge \kappa n$, the global Kaufman–Oppenheim complex $\mathfrak{A}^{\kappa,s}_n(\mathbb{F})$ is a clique complex.*

Theorem 3.57 (Local bipartite expansion of global Kaufman–Oppenheim complex). *For every prime p, $n, \kappa \in \mathbb{N}$, and $s \ge \kappa n$, the global Kaufman–Oppenheim complex $\mathfrak{A}^{\kappa,s}_n(\mathbb{F}_p)$ is a $1/\sqrt{p}$-local bipartite expander.*

Theorem 3.58 ("Strong gallery connectivity" of global Kaufman–Oppenheim complex). *For every prime p, $n, \kappa \in \mathbb{N}$, and $s \ge \kappa n$, and for every $r \le n-1$ and $w \in V_r(\mathfrak{X})$, the link $(\mathfrak{A}^{\kappa,s}_n(\mathbb{F}_p))_w$ is connected.*

[16] Actually, [KO23] only had $\kappa = 1$; the very minor generalization to $\kappa = 3$ is from [KOW25].

We also use the following property, which was proven by Kaufman, Oppenheim, and Weinberger [KOW25]:

Theorem 3.59 (Trivial cohomology of Kaufman–Oppenheim complex, [KOW25])**.** *Let p be prime and Γ a group which has no nontrivial elements of order p. (For instance, this is the case if $|\Gamma| < p$.) Let $\kappa := 3$, $n \in \mathbb{N}$, and $s \geq \kappa n$. Then complex $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ has trivial 1-cohomology over Γ.*

Finally, Theorem 3.17 together with Theorems 3.57 and 3.58 gives:

Corollary 3.60 (Global Kaufman–Oppenheim complex is product)**.** *For every prime p, $n, \kappa \in \mathbb{N}$, and $s \geq \kappa n$, if $\frac{1}{\sqrt{p}-n+1} < \frac{1}{2}$, then $\mathfrak{A}_n^{\kappa,s}(\mathbb{F})$ is $(\frac{1}{\sqrt{p}-n+1})$-product.*

3.11 The absolute Dehn method

Definition 3.61 (Subgroup words)**.** Let G be a group and $\mathcal{H} = (H_i < G)_{i\in\mathcal{I}}$ an $\mathcal{I}$-indexed family of subgroups. A *word over* $\mathcal{H}$ is a string of group elements written $\langle g_1\rangle \cdots \langle g_T\rangle$ such that each $g_t \in \bigcup \mathcal{H}$, i.e., $g_t \in H_{i_t}$ for some $i_t \in \mathcal{I}$. (The notation $\langle\cdot\rangle$ serves to highlight that a subgroup element $g \in \bigcup \mathcal{H}$ is formally different from the corresponding length-1 subgroup word $\langle g\rangle$.) The *length* of a word $w = \langle g_1\rangle \cdots \langle g_T\rangle$ is $|w| := T$. The inverse of a word $w = \langle g_1\rangle \cdots \langle g_T\rangle$ over $\mathcal{H}$ is the word $w^{-1} := \langle g_T^{-1}\rangle \cdots \langle g_1^{-1}\rangle$, also over $\mathcal{H}$. The *concatenation* of two words $w = \langle g_1\rangle \cdots \langle g_T\rangle$ and $w' = \langle g'_1\rangle \cdots \langle g'_{T'}\rangle$ is $ww' := \langle g_1\rangle \cdots \langle g_T\rangle\langle g'_1\rangle \cdots \langle g'_{T'}\rangle$. The empty (length-0) word is denoted e. ♢

Definition 3.62 (Evaluations)**.** Let $w = \langle g_1\rangle \cdots \langle g_T\rangle$ be a word over $\mathcal{H}$. The *evaluation* of w, denoted $\mathrm{eval}(w) \in G$, is the ordered product $g_1 \cdots g_T \in G$. For two words w, w' over $\mathcal{H}$, we write $w \equiv w'$ iff $\mathrm{eval}(w) = \mathrm{eval}(w')$. ♢

Definition 3.63. Let G be a group and $\mathcal{H} = (H_i < G)_{i\in\mathcal{I}}$ an $\mathcal{I}$-indexed subgroup family. Define:

$$\mathrm{diam}(G; \mathcal{H}) := \max_{g\in G} \min\{|w| : w \text{ is a word over } \mathcal{H} \text{ and } \mathrm{eval}(w) = g\}. \qquad ♢$$

Note that if $H_1H_2 = G$ then $\mathrm{diam}(G; (H_1, H_2)) = 2$, while $\mathrm{diam}(G; \mathcal{H}) < \infty$ iff the subgroups $\mathcal{H}$ generate G. The following proposition justifies the use of the notation $\mathrm{diam}(G; \mathcal{H})$:

Proposition 3.64 ([KO21])**.** *Let G be a group and $\mathcal{H} = (H_i < G)_{i\in\mathcal{I}}$ an indexed subgroup family. Then*

$$\mathrm{diam}(\mathfrak{CC}(G; \mathcal{H})) \leq \mathrm{diam}(G; \mathcal{H}).$$

Definition 3.65 (Relators)**.** A word w over $\mathcal{H}$ is a *relator* if its evaluation is $\mathbb{1}$ in G. (Hence, w is a relator iff $w \equiv e$.) ♢

Definition 3.66 (Equations)**.** We define an equivalence relation $\sim$ on words over $\mathcal{H}$ via:

1. $xy \sim yx$,
2. $x \sim x^{-1}$,

If $u \sim v$, then u is a relator iff v is a relator. An *equation* over $\mathcal{H}$ is a string "$y \equiv z$", where y, z are words over $\mathcal{H}$ and yz^{-1} is a relator; we identify the equation with this relator. ♢

Definition 3.67 (Derivations)**.** Let r, x, y be relators over $\bigcup \mathcal{H}$. We say x is *derived from y via r* if (for some words p, q, u, v over $\mathcal{H}$),

$$x \sim puq, \quad y \sim pvq, \quad \text{"}u \equiv v\text{"} \sim r. \qquad ♢$$

Definition 3.68. $\mathcal{R}^{\ell}_{\mathrm{common}}(G; \mathcal{H})$ is the set of relators $\langle g_1\rangle \cdots \langle g_\ell\rangle$ such that there exists a single $i \in \mathcal{I}$ with $g_1, \ldots, g_\ell \in H_i$. ♢

Definition 3.69 (Derivation length)**.** Let w be a relator over $\mathcal{H}$. We write $\mathrm{area}_{\mathcal{H}}(w)$ for the least number of derivation steps, via relators from $\mathcal{R}^3_{\mathrm{common}}(G; \mathcal{H})$, that it takes to reduce w to the word $\mathbb{1}$. (Or, we write $\mathrm{area}_{\mathcal{H}}(w) = \infty$ if this is not possible.) For $\mathcal{R}$ a set of relators over $\mathcal{H}$, we write

$$\mathrm{area}_{\mathcal{H}}(\mathcal{R}) := \max_{w\in\mathcal{R}} \mathrm{area}_{\mathcal{H}}(w)$$

for the maximum area of any relation in $\mathcal{R}$. Finally, for $T \in \mathbb{N}$, we write

$$\mathrm{area}^T(G; \mathcal{H}) := \mathrm{area}_{\mathcal{H}}(\{w : w \text{ a relator and } |w| \leq T\}),$$

for the maximum area of any relator of length T in G. ♢

Remark 3.70. It is not hard to see that for every $\ell \in \mathbb{N}$, $\mathrm{area}_{\mathcal{H}}\big(\mathcal{R}^{\ell}_{\mathrm{common}}(G;\mathcal{H})\big) \le \ell - 3$. $\diamondsuit$

A theorem relating area and coboundary expansion in coset complexes was originally proven in Kaufman and Oppenheim [KO21]; we give here a version with improved parameters due to O'Donnell and Singer [OS25].

Theorem 3.71 ([OS25, §C])**.** *For every group G and subgroup family $\mathcal{H}$, for $R_0 \coloneqq \mathrm{diam}(G;\mathcal{H})$ and $R_1 \coloneqq \mathrm{area}^{2R_0+1}(G;\mathcal{H})$, the coset complex $\mathfrak{CC}(G;\mathcal{H})$ has 1-triword expansion at least $12/R_1$ over every group Γ.*

4 Structure of the local Kaufman–Oppenheim complex

In this section, we prove a number of useful properties of the graded unipotent group, the staircase subgroups, and local Kaufman–Oppenheim complex. These proofs give good intuition for calculating with the Steinberg relations. Many of the properties are well-known (or are obvious when working with matrix realizations of the group) and have appeared (at least for the case $\kappa = 1$) in many previous works.

4.1 Group theory review: Direct products and retractions

Theorem 4.1 (Von Dyck's theorem, [Rot95, p. 346])**.** *Let $G = \langle S \mid R\rangle$ be a presented group (S are the generators and R the relations) and H another group. Suppose $\phi : S \to H$ is a mapping such that the image under ϕ of every relation in R is trivial. Then ϕ extends to a (unique) group homomorphism $\phi : G \to H$.*

Definition 4.2. Let G be a finite group and $\mathcal{H} = (H_i \le G)_{i \in \mathcal{I}}$ an $\mathcal{I}$-indexed subgroup family ($\mathcal{I}$ finite). We say $\mathcal{H}$ *forms an internal direct product* in G if:

1. For every $i \ne i' \in \mathcal{I}$, the subgroups H_i and $H_{i'}$ intersect trivially, i.e., $H_i \cap H_{i'} = \{\mathbb{1}\}$.
2. For every $i \ne i' \in \mathcal{I}$, the subgroups H_i and $H_{i'}$ commute, i.e., $H_i H_{i'} = H_{i'} H_i$.

In this case, we use $\prod \mathcal{H}$ to denote the subgroup of G generated by the subgroups in $\mathcal{H}$. We also say $\prod \mathcal{H}$ *splits as an internal direct product* of the family $\mathcal{H}$. $\diamondsuit$

Fact 4.3. *When G is a finite group, $\mathcal{H}$ an $\mathcal{I}$-indexed subgroup family, and G splits as an internal direct product $G = \prod \mathcal{H}$:*

1. *Every element of G can be written uniquely as a product $\prod_{i \in \mathcal{I}} h_i$ where $h_i \in H_i$ (order does not matter),*
2. *There are homomorphisms $\pi_i : G \twoheadrightarrow H_i$ for every $i \in \mathcal{I}$ satisfying $\pi_i(h) = h$ for every $h \in H_i$, $\pi_i(\pi_i(g)) = \pi_i(g)$ for every $g \in G$, and $H_{i'} \in \ker(\pi_i)$ for every $i \ne i' \in \mathcal{I}$.*

We now define the group-theoretic notion of a "retraction", which generalizes the homomorphisms π_i in the preceding fact:

Definition 4.4 (e.g., [Rot95, Lemma 7.20])**.** Let G be a group and $L \le G$ a subgroup. A *retraction* onto L is a homomorphism $\pi : G \twoheadrightarrow L$ such that for every $\ell \in L$, $\pi(\ell) = \ell$. $\diamondsuit$

Recall that by the first isomorphism theorem, if $\pi : G \twoheadrightarrow L$ is a surjective homomorphism, then $\ker(\pi) \trianglelefteq G$ is a normal subgroup and $L \cong G/\ker(\pi)$ is isomorphic to its quotient group. A retraction onto L is a special kind of a surjective map onto L, wherein this quotient group is itself isomorphic to a subgroup of G.

When a retraction onto L exists, G is called a "semidirect product", written $\ker(\pi) \rtimes L$. The subgroups $\ker(\pi)$ and L need not commute (in which case G is not a direct product of $\ker(\pi)$ and L), but a semidirect product still exhibits product-like properties:

Proposition 4.5. *Let G be a group, $L \le G$ a subgroup, and $\pi : G \twoheadrightarrow L$ a retraction onto L.*

1. *The subgroups $\ker \pi, L \le G$ intersect trivially.*
2. *π is a projection, i.e., for every $g \in G$, $\pi(\pi(g)) = \pi(g)$.*
3. *Every element $g \in G$ can be written uniquely as $g = k\ell$ for $k \in \ker \pi$ and $\ell \in L$ (and as $g = \ell' k'$ for $\ell' \in L$ and $k' \in \ker \pi$).*

4. *Suppose* $\ker(\pi) \le H$. *Then* $\pi(H) = H \cap L$.

5. *Suppose* $M \le \ker(\pi)$. *Then* $\ker(\pi) \cap LM = \ker(\pi) \cap ML = M$.

Proof. We have:

1. Every $g \in \ker \pi$ has $\pi(g) = \mathbb{1}$ and every $g \in L$ has $\pi(g) = g$.

2. Follows from the definition of retraction and that $\pi(g) \in L$.

3. For existence, set $\ell \coloneqq \pi(g)$ and $k \coloneqq g\ell^{-1}$, so that $\pi(k\ell) = \pi(g)(\pi(\ell))^{-1} = \pi(g)(\pi(\pi(g)))^{-1} = \pi(g)(\pi(g))^{-1} = \mathbb{1}$. For uniqueness, note that if $g = k\ell$ then $\pi(g) = \pi(k)\pi(\ell) = \pi(\ell) = \ell$. Hence ℓ is unique, and so so is k. The proof for the other case is symmetric.

4. Let $h \in H$. On one hand, if $h = k\ell$ with $k \in \ker(\pi)$ and $\ell \in L$, then $\pi(h) = \ell$. But $k^{-1}h = \ell \in H$ since $\ker(\pi) \le H$. On the other hand, if $h \in H \cap L$, then by the retraction property, $h = \pi(h) \in \pi(H)$.

5. Obviously, $M \le LM$ and $M \le \ker(\pi)$. Conversely, suppose that $h \in \ker(\pi) \cap LM$. Then on one hand, since $h \in LM$, we have $h = \ell m$ with $\ell \in L$ and $m \in M$, and so $\pi(h) = \pi(\ell) = \ell$ (since π is a retraction onto L and $m \in M \le \ker(\pi)$). At the same time, since $h \in \ker(\pi)$, we have $\pi(h) = \mathbb{1}$. Hence $\ell = 1$ and so $h = m \in M$. The proof for the other case is symmetric. ☐

4.2 "Reducing" a coset complex

The definitions in the preceding section yield the following corollary:

Proposition 4.6. *Let G be a group and $\mathcal{H} = (H_i \le G)_{i \in \mathcal{I}}$ an $\mathcal{I}$-indexed subgroup family. Suppose $W, \mathcal{J} \sqsubseteq \mathcal{I}$. Further, suppose $\pi : G \twoheadrightarrow L$ is a retraction onto L and $H_W = LM$ with $M \le \ker(\pi) \cap H_{\mathcal{J}}$. Then $\mathfrak{CC}(H_W; (H_{W \cup \{a\}})_{a \in \mathcal{J}}) \cong \mathfrak{CC}(L; (L \cap H_a)_{a \in \mathcal{J}})$ (and the index map $\iota : \mathcal{J} \to \mathcal{J}$ is the identity function).*

Proof. We apply Proposition 3.45 with $G' \coloneqq H_W$, $\pi' \coloneqq \pi|_{H_W}$, and $\mathcal{H}' \coloneqq (H_{W \cup \{a\}})_{a \in \mathcal{J}}$. Note that $\widetilde{\pi} : H_W \twoheadrightarrow L$ is still a retraction because $L \le H_W$. It remains to check that $\ker(\pi') \le H_{W \cup \{a\}}$ for every $a \in \mathcal{J}$. We have $\ker(\pi') = \ker(\pi) \cap H_W \le H_W$, so we need only verify that $\ker(\pi) \cap H_W \le H_a$. By the final item of Proposition 4.5, we have $\ker(\pi) \cap H_W = M \le H_{\mathcal{J}} \le H_a$. ☐

4.3 Triangular subgroups of the graded unipotent group

We now define another family of subgroups of $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ which will be of recurring interest to us in the paper. We say a set $\mathcal{I} \subseteq \mathbb{N}$ is an *interval* if $\mathcal{I} = [\min(\mathcal{I}) \,..\, \max(\mathcal{I})]$. Note that in this case, $|\mathcal{I}| = \max(\mathcal{I}) - \min(\mathcal{I}) + 1$.

Definition 4.7 (Triangular subgroups)**.** For $n, \kappa \in \mathbb{N}$, $\mathbb{F}$ a field, and an interval $\mathcal{I} \subseteq [0 \,..\, n]$, we define the "triangular subgroup" $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; \mathcal{I}) \le \mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ as the subgroup generated by the elements $e_{i,j}(f)$, where $i < j \in \mathcal{I}$ and $f \in \mathbb{F}^{\le \kappa(j-i)}[X]$. Note that by definition, $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [0 \,..\, n]) = \mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$, while $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [j-1 \,..\, j]) = \{\mathbb{1}\}$ for every $j \in [1 \,..\, n]$. ♢

We now prove several useful properties of these subgroups:

Proposition 4.8. *Let $n, \kappa \in \mathbb{N}$, $\mathbb{F}$ be a field, and $\mathcal{I}_1 \subseteq \mathcal{I}_2 \subseteq [0 \,..\, n]$ be intervals. We have the containment of subgroups:*

$$\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; \mathcal{I}_1) \le \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; \mathcal{I}_2).$$

Proof. Every generator of the left-hand group is also in the right-hand group. ☐

Proposition 4.9. *Let $n, \kappa \in \mathbb{N}$, $\mathbb{F}$ be a field, and $\mathcal{I} \subseteq [0 \,..\, n]$ be an interval. Consider the groups $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; \mathcal{I})$ and $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(|\mathcal{I}| - 1)$; write the generators of the former as $e_{i,j}(f)$ ($i < j \in \mathcal{I}$, $f \in \mathbb{F}^{\le \kappa(j-i)}[X]$) and of the latter as $\widetilde{e}_{i,j}(f)$ ($i < j \in [0 \,..\, |\mathcal{I}| - 1]$, $f \in \mathbb{F}^{\le \kappa(j-i)}[X]$). Then there is an isomorphism:*

$$\begin{aligned} \Phi^{(n;\mathcal{I})}_{\mathbb{F}} : \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; \mathcal{I}) &\xrightarrow{\sim} \mathrm{Unip}^{\kappa}_{\mathbb{F}}(|\mathcal{I}| - 1), \\ e_{i,j}(f) &\longmapsto \widetilde{e}_{\,i-\min(\mathcal{I}),\, j-\min(\mathcal{I})}(f). \end{aligned}$$

(Note that $\widetilde{e}_{i-\min(\mathcal{I}),j-\min(\mathcal{I})}(f)$ is a valid generator in the codomain $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(|\mathcal{I}|-1)$ because $i \geq \min(\mathcal{I}) \implies i - \min(\mathcal{I}) \geq 0$, $j \leq \max(\mathcal{I}) \implies j - \min(\mathcal{I}) \leq \max(\mathcal{I}) - \min(\mathcal{I})$, and $(j-\min(\mathcal{I})) - (i - \min(\mathcal{I})) = j - i$.) Further, for every interval $\mathcal{I}' \subseteq \mathcal{I}$,

$$\Phi^{(n;\mathcal{I})}_{\mathbb{F}}(\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}')) = \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(|\mathcal{I}|-1;[\min(\mathcal{I}')-\min(\mathcal{I})\,..\,\max(\mathcal{I}')-\min(\mathcal{I})]).$$

Proof. It is easy to check that this map, defined on the generators, is a homomorphism by applying von Dyck's theorem (Theorem 4.1) and the definition of the unipotent group (Definition 3.50). One can similarly verify that there is a homomorphism $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(|\mathcal{I}|-1) \to \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}) : \widetilde{e}_{i,j}(f) \mapsto e_{\min(\mathcal{I})+i,\min(\mathcal{I})+j}(f)$, and finally that this homomorphism is the inverse of $\Phi^{(n;\mathcal{I})}_{\mathbb{F}}$. The "further" statement follows since $\Phi^{(n;\mathcal{I})}_{\mathbb{F}}$, restricted to $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}')$, agrees with $\Phi^{(n;\mathcal{I}')}_{\mathbb{F}}$. ☐

Proposition 4.10. *Let $n,\kappa \in \mathbb{N}$, $\mathbb{F}$ be a field, and intervals $\mathcal{I}_1,\mathcal{I}_2 \subseteq [0\,..\,n]$ satisfy $\mathcal{I}_1 \cap \mathcal{I}_2 = \emptyset$. The subgroups $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1)$ and $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_2)$ of $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ commute.*

Proof. $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1)$ is generated by elements $e_{i,j}(f)$ with $i < j \in \mathcal{I}_1$, and $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_2)$ by elements $e_{i,j}(f)$ with $i < j \in \mathcal{I}_2$. The disjointness assumption and the trivial commutator relations imply that all generators of the former subgroup commute with all generators of the latter subgroup. ☐

In order to prove richer properties of these subgroups, we now define a retraction of the "big triangle" group $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n) = \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;[0\,..\,n])$ onto "small triangle" subgroups $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I})$. Together with the isomorphism of $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I})$ and $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(|\mathcal{I}|-1)$ (Proposition 4.9), the retraction will be useful to us several times in the sequel to analyze the "self-reducibility" of the unipotent group construction.

Proposition 4.11. *Let $n,\kappa \in \mathbb{N}$, $\mathcal{I} \subseteq [0\,..\,n]$ be an interval, and $\mathbb{F}$ be a field. Consider the groups $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n) \geq \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I})$; write the generators of the former as $e_{i,j}(f)$ ($i < j \in [0\,..\,n]$, $f \in \mathbb{F}^{\leq \kappa(j-i)}[X]$) and of the latter as $\widetilde{e}_{i,j}(f)$ ($i < j \in \mathcal{I}$, $f \in \mathbb{F}^{\leq \kappa(j-i)}[X]$). Then there is a retraction:*

$$\pi^{(n;\mathcal{I})}_{\mathbb{F}} : \mathrm{Unip}^{\kappa}_{\mathbb{F}}(n) \twoheadrightarrow \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}),$$
$$e_{i,j}(f) \longmapsto \begin{cases} \widetilde{e}_{i,j}(f) & \text{if } i < j \in \mathcal{I}, \\ \mathbb{1} & \text{otherwise.} \end{cases}$$

(Thus, for every $i < j \in [0\,..\,n]$ with $\{i,j\} \not\subset \mathcal{I}$, $e_{i,j}(f) \in \ker(\pi^{(n;\mathcal{I})}_{\mathbb{F}})$ for every $f \in \mathbb{F}^{\leq \kappa(j-i)}[X]$.)

Proof. This map, if it is a valid homomorphism, is obviously a retraction. To check that is is a homomorphism, we again apply von Dyck's theorem (Theorem 4.1), and verify that the image of every relation defining $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ vanishes. This follows from some simple casework and the fact that commutators with trivial elements are trivial:

- For linearity relations $e_{i,j}(f) \cdot e_{i,j}(g) = e_{i,j}(f+g)$ ($i < j \in [0\,..\,n]$, $f,g \in \mathbb{F}[X]$ with $\deg(f),\deg(g) \leq \kappa(j-i)$), either $i < j \in \mathcal{I}$, in which case the image is the (true) linearity relation $\widetilde{e}_{i,j}(f) \cdot \widetilde{e}_{i,j}(g) = \widetilde{e}_{i,j}(f+g)$, or the image is the trivial relation $\mathbb{1} \cdot \mathbb{1} = \mathbb{1}$.
- For nontrivial commutator relations $[e_{i,j}(f), e_{j,k}(g)] = e_{i,k}(fg)$ ($i < j < k \in [0\,..\,n]$, $f \in \mathbb{F}^{\leq \kappa(j-i)}[X]$, $g \in \mathbb{F}^{\leq \kappa(k-j)}[X]$), either $i,k \in \mathcal{I}$, in which case the image is the (true) nontrivial commutator relation $[\widetilde{e}_{i,j}(f), \widetilde{e}_{j,k}(g)] = \widetilde{e}_{i,k}(fg)$, or (at least) one of $e_{i,j}(f)$ and $e_{j,k}(g)$ has trivial image *and* $e_{i,k}(fg)$ has trivial image.
- Finally, for trivial commutator relations $[e_{i,j}(f), e_{k,\ell}(g)] = \mathbb{1}$ ($i < j \in [0\,..\,n]$, $k < \ell \in [0\,..\,n]$, $j \neq k$, $i \neq \ell$, $f \in \mathbb{F}^{\leq \kappa(j-i)}[X]$, $g \in \mathbb{F}^{\leq \kappa(\ell-k)}[X]$), either $i,j,k,\ell \in \mathcal{I}$, in which case the image is the (true) trivial commutator relation $[\widetilde{e}_{i,j}(f), \widetilde{e}_{k,\ell}(g)] = \mathbb{1}$, or (at least) one of $e_{i,j}(f)$ and $e_{k,\ell}(g)$ has trivial image. ☐

This retraction has the following useful properties:

Proposition 4.12. *Let $n,\kappa \in \mathbb{N}$, $\mathbb{F}$ be a field, and $\mathcal{I}_1,\mathcal{I}_2 \subseteq [0\,..\,n]$ be intervals. Then*

$$\pi^{(n;\mathcal{I}_1)}_{\mathbb{F}}(\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_2)) = \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1 \cap \mathcal{I}_2).$$

Proof. We claim that $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_2)$ is generated by two subgroups, $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1\cap\mathcal{I}_2)$ and another subgroup $L \le \ker(\pi^{(n;\mathcal{I}_1)}_{\mathbb{F}})$. This suffices because $\pi^{(n;\mathcal{I}_1)}_{\mathbb{F}}$ preserves $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1\cap\mathcal{I}_2)$ and kills L.

Indeed, the generators of $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_2)$ are of the form $e_{i,j}(f)$ where $i<j\in\mathcal{I}_2$ and $f\in\mathbb{F}^{\le\kappa(j-i)}[X]$. We set L to be the subgroup generated by the generators $e_{i,j}(f)$ where $\{i,j\}\not\subseteq\mathcal{I}_1$. Such generators are killed by the retraction $\pi^{(n;\mathcal{I}_1)}_{\mathbb{F}}$, and the remaining generators of $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1\cap\mathcal{I}_2)$ are those where $\{i,j\}\subseteq\mathcal{I}_1$ and are therefore contained in $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1\cap\mathcal{I}_2)$. □

Proposition 4.13. *Let $n,\kappa\in\mathbb{N}$, $\mathbb{F}$ be a field, and $\mathcal{I}_1,\mathcal{I}_2\subseteq[0\,..\,n]$ be intervals. Then:*

$$\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1)\cap\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_2)=\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1\cap\mathcal{I}_2).$$

In particular, if $|\mathcal{I}_1\cap\mathcal{I}_2|\le 1$, then $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1)$ and $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_2)$ intersect trivially.

Proof. The $\ge$ direction follows immediately from Proposition 4.8. For the $\le$ direction, consider any element $g\in\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1)\cap\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_2)$. Since $g\in\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1)$, g is fixed by the retraction $\pi^{(n;\mathcal{I}_1)}_{\mathbb{F}}$, while by Proposition 4.12, the same retraction maps g into $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_1\cap\mathcal{I}_2)$. □

Note that Propositions 4.10 and 4.13 together imply that for any family of pairwise disjoint intervals $(\mathcal{I}_t\subseteq[0\,..\,n])_{t\in\mathcal{T}}$, the triangle subgroups $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_t)$ form an internal direct product within $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$.

4.4 Staircase subgroups of the unipotent group

We next turn to describing the staircase subgroups (Definition 3.52) in terms of the triangle subgroups (Definition 4.7).

Proposition 4.14. *In the setup of the preceding definition, $\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell)$ splits as an internal direct product:*

$$\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell)=\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;[0\,..\,\ell))\times\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;[\ell\,..\,n]).$$

Proof. The subgroups $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;[0\,..\,\ell))$ and $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;[\ell\,..\,n])$ commute by Proposition 4.10. Their intersection is trivial by Proposition 4.13. They are contained in and generate $\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell)$ since this group is generated by elements $e_{i,j}(f)$ with either $j<\ell$ or $i\ge\ell$; generators of the former kind are present in and generate the subgroup $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;[0\,..\,\ell))$, while generators of the latter kind are present in and generate the subgroup $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;[\ell\,..\,n])$. □

Now, we prove the following simple fact in Appendix D:

Proposition 4.15. *Let G be a group and $H,K\le G$ subgroups, with $H=H_1\times H_2$ and $K=K_1\times K_2$ splitting as internal direct products. Further, suppose that $K_1\le H_1$ and $H_2\le K_2$. Then $H\cap K$ also splits as an internal direct product:*

$$H\cap K=K_1\times(H_1\cap K_2)\times H_2.$$

Applying this inductively together with Propositions 4.10 and 4.14 gives the following, which we also prove in Appendix D:

Proposition 4.16. *For every $n,\kappa\in\mathbb{N}$ and field $\mathbb{F}$, and $W\subseteq[1\,..\,n]$, the following holds. Suppose $W=\{\ell_1,\dots,\ell_T\}$ with $\ell_1<\cdots<\ell_T$, and define $\ell_0:=0$ and $\ell_{T+1}:=n+1$. Define the family of intervals $(\mathcal{I}_t:=[\ell_{t-1}\,..\,\ell_t)\subseteq[0\,..\,n])_{t\in[1\,..\,T+1]}$. Then $\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;W)$ splits as an internal direct product:*

$$\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;W)=\prod_{t\in[1\,..\,T+1]}\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I}_t).$$

(Recall that $\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n;\mathcal{I})$ is trivial when $|\mathcal{I}|=1$, so these intervals can be ignored.)

This yields several nice corollaries:

Corollary 4.17. *For every $n,\kappa\in\mathbb{N}$, field $\mathbb{F}$, and $W\subseteq[1\,..\,n]$, $\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;W)$ is the subgroup of $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ generated by elements of the form $e_{i,j}(f)$ for $i<j\in[0\,..\,n]$ satisfying $(i\,..\,j]\cap W=\emptyset$ and $f\in\mathbb{F}^{\le\kappa(j-i)}[X]$.*

Proof. Use the prior proposition and the definition of triangle subgroups by generators (Definition 4.7). ☐

Corollary 4.18. *For every $n, \kappa \in \mathbb{N}$, field $\mathbb{F}$, and $W, W_1, W_2 \subseteq [1 \,..\, n]$, the following holds. Suppose that W_1 and W_2 satisfy that $W_1 \supseteq W$, $W_2 \supseteq W$, and for every $i \in W_1 \setminus W$ and $j \in W_2 \setminus W$, $(\min\{i,j\} \,..\, \max\{i,j\}] \cap W \neq \emptyset$. Then:*

$$\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; W) \subseteq \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; W_1) \cdot \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; W_2).$$

Corollary 4.19. *For every $n, \kappa \in \mathbb{N}$, field $\mathbb{F}$, and interval $\mathcal{I} \subseteq [0 \,..\, n]$, we have:*

$$\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; \mathcal{I}) = \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; [1 \,..\, \min(\mathcal{I})] \cup (\max(\mathcal{I}) \,..\, n]).$$

(Alternatively, $[1 \,..\, \min(\mathcal{I})] \cup (\max(\mathcal{I}) \,..\, n] = [1 \,..\, n] \setminus (\mathcal{I} \setminus \{\min(\mathcal{I})\})$.)

Corollary 4.20. *For every $n, \kappa \in \mathbb{N}$ and field $\mathbb{F}$, the staircase subgroup $\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; [1 \,..\, n]) = \bigcap_{\ell=1}^{n} \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; \ell)$ is trivial.*

4.5 Some self-reducibility results

The prior calculations let us easily calculate the retraction of a staircase subgroup:

Proposition 4.21. *For every $n, \kappa \in \mathbb{N}$, field $\mathbb{F}$, interval $\mathcal{I} \subseteq [0 \,..\, n]$, $a \in \mathcal{I}$, we have:*

$$\pi^{(n;\mathcal{I})}_{\mathbb{F}}(\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; a)) = \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; a) \cap \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; \mathcal{I}) = \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [\min(\mathcal{I}) \,..\, a)) \times \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [a \,..\, \max(\mathcal{I})]).$$

Proof. Abbreviate $\pi := \pi^{(n;\mathcal{I})}_{\mathbb{F}}$. Using Propositions 4.12 and 4.14:

$$\begin{aligned} \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [\min(\mathcal{I}) \,..\, a)) \times \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [a \,..\, \max(\mathcal{I})]) &= \pi(\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [0 \,..\, a))) \times \pi(\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [a \,..\, n])) \\ &= \pi(\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [0 \,..\, a)) \times \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [a \,..\, n])) = \pi(\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; a)). \end{aligned}$$

At the same time, using Corollary 4.19 and Proposition 4.16:

$$\begin{aligned} &\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; a) \cap \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; \mathcal{I}) = \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; a) \cap \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; [1 \,..\, \min(\mathcal{I})] \cup (\max(\mathcal{I}) \,..\, n]) \\ &= \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; [1 \,..\, \min(\mathcal{I})] \cup \{a\} \cup (\max(\mathcal{I}) \,..\, n]) = \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [\min(\mathcal{I}) \,..\, a))) \times \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [a \,..\, \max(\mathcal{I})]), \end{aligned}$$

as desired. ☐

Corollary 4.22. *For every $n, \kappa \in \mathbb{N}$, field $\mathbb{F}$, interval $\mathcal{I} \subseteq [0 \,..\, n]$, and $a \in \mathcal{I}$, we have:*

$$\Phi^{(n;\mathcal{I})}_{\mathbb{F}}\left(\pi^{(n;\mathcal{I})}_{\mathbb{F}}(\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; a))\right) = \mathrm{Stair}^{\kappa}_{\mathbb{F}}(|\mathcal{I}| - 1; a - \min(\mathcal{I})).$$

Proof. Abbreviate $\pi := \pi^{(n;\mathcal{I})}_{\mathbb{F}}$ and $\Phi := \Phi^{(n;\mathcal{I})}_{\mathbb{F}}$. We calculate using Propositions 4.9 and 4.14, and the prior proposition:

$$\begin{aligned} \mathrm{Stair}^{\kappa}_{\mathbb{F}}(|\mathcal{I}| - 1; a - \min(\mathcal{I})) &= \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(|\mathcal{I}| - 1; [0 \,..\, a - \min(\mathcal{I}))) \times \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(|\mathcal{I}| - 1; [a - \min(\mathcal{I}) \,..\, |\mathcal{I}| - 1]) \\ &= \Phi(\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [\min(\mathcal{I}) \,..\, a))) \times \Phi(\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [a \,..\, \max(\mathcal{I})])) \\ &= \Phi(\mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [\min(\mathcal{I}) \,..\, a)) \times \mathrm{Triangle}^{\kappa}_{\mathbb{F}}(n; [a \,..\, \max(\mathcal{I})])) = \Phi(\pi(\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; a))), \end{aligned}$$

as desired. ☐

Proposition 4.23. *Let $W \subseteq [1 \,..\, n]$. Let $\ell^{(1)} < \ell^{(2)} \in W \cup \{0, n+1\}$. Let $W_1 := W \cup (\ell \,..\, \ell')$ and $W_2 := W \cup [1 \,..\, \ell) \cup (\ell' \,..\, n]$. Then $\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; W)$ splits as an internal direct product:*

$$\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; W) = \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; W_1) \times \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; W_2).$$

Proof. As in Proposition 4.16, let $W = \{\ell_1, \dots, \ell_T\}$ with $\ell_1 < \cdots < \ell_T$ and $\ell_0 := 0$ and $\ell_{T+1} := n+1$. Suppose $\ell^{(1)} = \ell_{t^{(1)}}$ and $\ell^{(2)} = \ell_{t^{(2)}}$, for some $t^{(1)} < t^{(2)} \in [0\,..\,T+1]$. By Proposition 4.16,

$$\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W) = \prod_{t\in[1\,..\,T+1]} \mathcal{I}_t = \prod_{t\in[t^{(1)}\,..\,t^{(2)})} \mathcal{I}_t \times \prod_{t\in[1\,..\,T+1]\setminus[t^{(1)}\,..\,t^{(2)})} \mathcal{I}_t$$

for the intervals $(\mathcal{I}_t := [\ell_{t-1}\,..\,\ell_t) \subseteq [0\,..\,n])_{t\in[1\,..\,T+1]}$. We claim that the first factor equals $\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W_2)$ and the second $\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W_1)$. Indeed, note that:

$$\begin{aligned} W &= \{\ell_t : t \in [1\,..\,T+1)\}, \\ W_2 &= [1\,..\,\ell_{t^{(1)}} - 1] \cup \{\ell_t : t \in [t^{(1)}\,..\,t^{(2)})\} \cup [\ell_{t^{(2)}}\,..\,n+1), \\ W_1 &= \{\ell_t : t \in [1\,..\,t^{(1)} - 1)\} \cup [\ell_{t^{(1)}}\,..\,\ell_{t^{(2)}}) \cup \{\ell_t : t \in [t^{(2)}\,..\,T+1)\}. \end{aligned}$$

Thus, the nonempty intervals in the decomposition of $\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W_2)$ are $(\mathcal{I}_t)_{t\in[t^{(1)}\,..\,t^{(2)})}$, and for $\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W_1)$ are $(\mathcal{I}_t)_{t\in[1\,..\,T+1]\setminus[t^{(1)}\,..\,t^{(2)})}$. Applying Proposition 4.16 again finishes the proof. □

We also have the following consequence, which we will use to establish productization:

Proposition 4.24. *Let $\kappa, n \in \mathbb{N}$, field $\mathbb{F}$, and $W \subseteq [1\,..\,n]$. If $S_1 \prec S_2 \sqsubset [1\,..\,n] \setminus W$ satisfy that $W \cap (\max S_1\,..\,\min S_2) \neq \emptyset$, then $\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W)$ equals a product of subgroups:*

$$\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W) = \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;S_1 \cup W) \cdot \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;S_2 \cup W).$$

Proof. Pick $\ell \in W \cap (\max S_1\,..\,\min S_2)$. Note that $S_1 \subseteq [1\,..\,\ell)$ and $S_2 \subseteq (\ell\,..\,n]$.

By the prior Proposition 4.23, $\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W)$ splits as an internal direct product $\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W \cup (\ell\,..\,n+1)) \times \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W \cup [1\,..\,\ell) \cup (n+1\,..\,n])$. The first factor is contained in $\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W \cup S_2)$ since $W \cup S_2 \subseteq W \cup (\ell\,..\,n]$ and the second factor is in $\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W \cup S_1)$ since $W \cup S_1 \subseteq W \cup [1\,..\,\ell)$. □

Proposition 4.6 has the following useful corollary which characterizes induced complexes in links of the local Kaufman–Oppenheim complex. Roughly, it says that the induced complex on a set of wires between two moats is itself a (smaller) local Kaufman–Oppenheim complex.

Proposition 4.25 ("Self-reducibility" of the local Kaufman–Oppenheim complex)**.** *For every $n, \kappa \in \mathbb{N}$ and field $\mathbb{F}$, and every $W \subseteq [1\,..\,n]$, the following holds. Suppose $\ell_1 < \ell_2 \in W \cup \{0, n+1\}$, and $W \cap (\ell_1\,..\,\ell_2) = \emptyset$ and w is a W-word. Then the induced complex on $(\ell_1\,..\,\ell_2)$ is isomorphic to a smaller Kaufman–Oppenheim complex:*

$$(\mathfrak{LA}_{\mathbb{F}}^{\kappa}(n))_w[(\ell_1\,..\,\ell_2)] \cong \mathfrak{LA}_{\mathbb{F}}^{\kappa}(\ell_2 - \ell_1 - 1).$$

(The bijection on indices is $\iota : (\ell_1\,..\,\ell_2) \to [1\,..\,\ell_2 - \ell_1 - 1] : \ell \mapsto \ell - \ell_1$.)

Proof. We set $G := \mathrm{Unip}_{\mathbb{F}}^{\kappa}(n)$, $\mathcal{J} := (\ell_1\,..\,\ell_2)$, $H_a := \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;a)$, $L := \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n;[\ell_1\,..\,\ell_2))$, and $M := \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W \cup (\ell_1\,..\,\ell_2))$. By Proposition 4.11, $\pi := \pi_{\mathbb{F}}^{(n;[\ell_1\,..\,\ell_2))}$ is a retraction $G \twoheadrightarrow L$. By Propositions 4.5 and 4.21, $H \cap L_a = \pi(L_a)$.

Now, we apply Proposition 4.6. If we can check that (1) $H_W = LM$ and (2) $M \leq \ker(\pi) \cap H_{\mathcal{J}}$, we then have that $(\mathfrak{LA}_{\mathbb{F}}^{\kappa}(n))_w[(\ell_1\,..\,\ell_2)] \cong \mathfrak{CC}(L; H_a \cap L)$. Then, we can apply the isomorphism $\Phi_{\mathbb{F}}^{(n;[\ell_1\,..\,\ell_2))}$, which carries L to $\mathrm{Unip}_{\mathbb{F}}^{\kappa}(\ell_2 - \ell_1 - 1)$ and, by Corollary 4.22, $L \cap H_a$ to $\mathrm{Unip}_{\mathbb{F}}^{\kappa}(a - \ell_1)$.

First, note that by Corollary 4.19, $L = \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n;[\ell_1\,..\,\ell_2)) = \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;[1\,..\,\ell_1] \cup [\ell_2\,..\,n]) = \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;W \cup [1\,..\,\ell_1) \cup (\ell_2\,..\,n])$ (where the last equality used that $\ell_1, \ell_2 \in W$ and $(\ell_1\,..\,\ell_2) \cap W = \emptyset$). Hence, (1) follows immediately from the prior Proposition 4.23.

For (2), we have $H_{\mathcal{J}} = \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;(\ell_1\,..\,\ell_2))$. Hence $M \leq H_{\mathcal{J}}$ since $(\ell_1\,..\,\ell_2) \subseteq W \cup (\ell_1\,..\,\ell_2)$. On the other hand,

$$\begin{aligned} \pi(M) &\leq \bigcap_{a\in(\ell_1\,..\,\ell_2)} \pi(\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;a)) = \bigcap_{a\in(\ell_1\,..\,\ell_2)} (\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;a) \cap \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n;(\ell_1\,..\,\ell_2))) \\ &= \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;(\ell_1\,..\,\ell_2)) \cap \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n;(\ell_1\,..\,\ell_2)) = \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;(\ell_1\,..\,\ell_2)) \cap L = \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;[1\,..\,n]) = \{\mathbb{1}\}, \end{aligned}$$

and so $M \in \ker(\pi)$, as desired. □

Proposition 4.26 (Productization in the local Kaufman–Oppenheim complex)**.** *For every $n, \kappa \in \mathbb{N}$ and field $\mathbb{F}$, the local Kaufman–Oppenheim complex $\mathfrak{LA}^{\kappa}_{\mathbb{F}}(n)$ productizes on the line (in the sense of Definition 3.28).*

Proof. We want to show that for every $W \subseteq [1\,..\,n]$ and $S_1 \prec S_2 \sqsubset [1\,..\,n] \setminus W$, if $W \cap (\max S_1\,..\,\min S_2) \neq \emptyset$, then for every W-word w, $S_1 \perp_{\mathfrak{LA}^{\kappa}_{\mathbb{F}}(n)_w} S_2$. By definition, $\mathfrak{LA}^{\kappa}_{\mathbb{F}}(n) = \mathfrak{CC}(\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n); (\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell))_{\ell \in [1\,..\,n]})$. By Proposition 3.47, we therefore need to check whether $\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;W) = \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;W \cup S_1) \cdot \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;W \cup S_2)$. But this is precisely Proposition 4.24. □

4.6 Representing $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$

The next important structural result lets us, in a certain sense, write unique expressions for every element in $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$:

Proposition 4.27. *Let $\Phi := \{(i,j) : i < j \in [0\,..\,n]\}$ (so that $|\Phi| = \binom{n+1}{2}$). For every ordering $\pi : [1\,..\,\binom{n+1}{2}] \to \Phi$ and $g \in \mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$, there exist unique polynomials $(f_{i,j} \in \mathbb{F}^{\leq \kappa(j-i)}[X])_{(i,j)\in\Phi}$ such that*

$$g = \prod_{p=1}^{\binom{n+1}{2}} e_{i,j}(f_{i,j}).$$

By the direct product decomposition in Proposition 4.16, we immediately get:

Corollary 4.28 (Essentially [OP22, Lemma 3.14])**.** *Let $W \subseteq [1\,..\,n]$ and $\Phi_W := \{(i,j) : i < j \in [0\,..\,n], W \cap (i\,..\,j] = \emptyset\}$. For every group element $x \in \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;S)$ and bijection $\pi : [|\Phi_W|] \to \Phi_W$, there exist unique polynomials $(f_{i,j})_{(i,j)\in\Phi_W}$ with $\deg(f_{i,j}) \leq \kappa(j-i)$ such that*

$$x = \prod_{p=1}^{|\Phi_W|} e_{\pi(p)}(f_{\pi(p)}).$$

The group $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ has a natural realization as a matrix group. Consider the set of matrices with entries in $\mathbb{F}[X]$:

$$S^{\kappa}_{n}(\mathbb{F}) := \left\{ M \in (\mathbb{F}[X])^{[0\,..\,n]\times[0\,..\,n]} : \begin{array}{ll} \forall j < i \in [0\,..\,n], & M_{i,j} = 0 \\ \forall i \in [0\,..\,n], & M_{i,i} = 1 \\ \forall i < j \in [0\,..\,n], & \deg(M_{i,j}) \leq \kappa(j-i) \end{array} \right\} \tag{4.29}$$

(note that these are $n+1 \times n+1$ matrices).

Proposition 4.30. *Let $n \in \mathbb{N}$ and $\mathbb{F}$ be a field. The set of matrices $S_n(\mathbb{F})$ in Equation (4.29) forms a group under matrix multiplication, and moreover, there is an isomorphism $\phi : \mathrm{Unip}^{\kappa}_{\mathbb{F}}(n) \to S^{\kappa}_{n}(\mathbb{F})$ defined on the generators by*

$$\phi(e_{i,j}(f)) := \mathbb{1} + f \cdot 1_{i,j}.$$

Hence, there is a faithful representation of $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ on the vector space $(\mathbb{F}[X])^{[0\,..\,n]}$, defined on generators as

$$(v_0, \ldots, v_n) \xmapsto{e_{i,j}(f)} (v_0, \ldots, v_{j-1}, v_j + f v_i, v_{j+1}, \ldots, v_n).$$

This gives some useful alternative perspectives on the group $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$:

- Given a "state vector" $v = (v_0, \ldots, v_n)$, $e_{i,j}(f)$ can be viewed as an "instruction" of the form "`add` $f v_i$ `to` v_j;", and a product of generators can be viewed as a "program" consisting of a sequence of instructions.
- Visually, we can view the coordinates $[0\,..\,n]$ as wires storing the current state, the generators $e_{i,j}(f)$ as "gates" adding a higher wire (closer to 0) to a lower wire (closer to 1) with a certain (polynomial) weight.

Each program/circuit *evaluates* to a linear operator in $\mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$ (the evaluation map is many-to-one).

In this view, for $W \subseteq [1\,..\,n]$, elements in the staircase subgroup $\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;W)$ are precisely the evaluations of circuits where, when we place a "moat" between wire $i-1$ and wire i for each $i \in W$, no gate crosses any moat.

4.7 Matrix Steinberg relations

Definition 4.31 ("Triangular" graded matrices)**.** Let $n, \kappa \in \mathbb{N}$, $\mathbb{F}$ be a field, and $\mathcal{I} \subseteq [0\,..\,n]$ be an interval. We define the set of matrices

$$\mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}) := \left\{ A \in \mathbb{F}[X]^{\mathcal{I}\times\mathcal{I}} : \begin{array}{l} \forall i < j \in \mathcal{I},\ \deg(A_{i,j}) \leq \kappa(j-i), \\ \forall i \geq j \in \mathcal{I},\ A_{i,j} = 0. \end{array} \right\}. \qquad \Diamond$$

Note that each matrix $A \in \mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{I})$ satisfies $\deg(A_{i,j}) \leq \kappa(j-i)$ for every $i, j \in \mathcal{I}$ under our standard convention $\deg(0) = -\infty$.

Definition 4.32 ("Rectangular" graded matrices)**.** Let $n, \kappa \in \mathbb{N}$, $\mathbb{F}$ be a field, and $\mathcal{I} \prec \mathcal{J} \subset [0\,..\,n]$ be two intervals. We define the set of matrices

$$\mathrm{RectMat}_{\mathbb{F}}^{\kappa}(\mathcal{I} \times \mathcal{J}) := \{F \in \mathbb{F}[X]^{\mathcal{I}\times\mathcal{J}} : \forall i \in \mathcal{I}, j \in \mathcal{J},\ \deg(F_{i,j}) \leq \kappa(j-i)\}. \qquad \Diamond$$

The following proposition expresses closure conditions for these sets under matrix addition and multiplication:

Fact 4.33. *Let $n, \kappa \in \mathbb{N}$ and $\mathbb{F}$ be a field.*

1. *Let $\mathcal{I} \subseteq [0\,..\,n]$ be an interval. Then for every $A, A' \in \mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{I})$, we have $FF', F+F' \in \mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{I})$.*
2. *Let $\mathcal{I} \prec \mathcal{J} \subset [0\,..\,n]$ be intervals. Then for every $A \in \mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{I})$ and $F \in \mathrm{RectMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}\times\mathcal{J})$, we have $AF \in \mathrm{RectMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}\times\mathcal{J})$. Similarly, for every $B \in \mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{J})$ and $F \in \mathrm{RectMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}\times\mathcal{J})$, we have $FB \in \mathrm{RectMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}\times\mathcal{J})$.*
3. *Let $\mathcal{I} \prec \mathcal{J} \prec \mathcal{K} \subset [0\,..\,n]$. Then for every $F \in \mathrm{RectMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}\times\mathcal{J})$ and $G \in \mathrm{RectMat}_{\mathbb{F}}^{\kappa}(\mathcal{J}\times\mathcal{K})$, we have $FG \in \mathrm{RectMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}\times\mathcal{K})$.*

Proof. We prove each item separately, and only prove one representative inclusion for each item.

1. We expand $(AA')_{i,j} = \sum_{k=\min(\mathcal{I})}^{\max(\mathcal{I})} A_{i,k}A'_{k,j}$. Note that if $i \geq j$, then for every k, either $i \geq k$ (so $A_{i,k} = 0$) or $k \geq i$ (so $A'_{k,j} = 0$). Otherwise, $\deg((AA)_{i,j}) \leq \max_{k=\min(\mathcal{I})}^{\max(\mathcal{I})}(\deg(A_{i,k}) + \deg(A'_{k,j})) \leq \kappa((k-i)+(j-k)) = \kappa(j-i)$, as desired. Similarly, we expand $(A+A')_{i,j} = A_{i,j} + A'_{i,j}$, so if $i \geq j$ all terms vanish and otherwise $\deg((A+A')_{i,j}) \leq \max\{\deg(A_{i,\bar{a}'}), \deg(A'_{i,\bar{a}'})\} \leq \kappa(j-i)$.
2. We prove only the first, as the second is similar. We expand $(AF)_{i,k} = \sum_{k=\min(\mathcal{J})}^{\max(\mathcal{J})} A_{i,j}F_{j,k}$, so $\deg((AF)_{i,k}) \leq \max_{k=\min(\mathcal{J})}^{\max(\mathcal{J})}(\deg(A_{i,j}) + \deg(F_{j,k})) \leq \kappa((j-i)+(k-j)) = \kappa(k-i)$.
3. Similar to the previous item. $\square$

We also have the following fact regarding inverses of strictly upper-triangular matrices with added 1's on the diagonal:

Fact 4.34. *For every $n, \kappa \in \mathbb{N}$, field $\mathbb{F}$, interval $\mathcal{I} \subseteq [0\,..\,n]$, and $A \in \mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{I})$, there exists a (unique) $A^{\dagger} \in \mathrm{Mat}_{\mathtt{A}}$ such that $AA^{\dagger} + A + A^{\dagger} = 0$ (equiv., $(A + \mathbb{1})(A^{\dagger} + \mathbb{1}) = \mathbb{1}$).*

Definition 4.35 (Group elements corresponding to matrices)**.** Let $n, \kappa \in \mathbb{N}$ and $\mathbb{F}$ be a field. For an interval $\mathcal{I} \subseteq [0\,..\,n]$ and $A \in \mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{I})$, we define the group element:

$$\{A\} := \prod_{j=\max(\mathcal{I})}^{\min(\mathcal{I})} \left(\prod_{i=\max(\mathcal{I})}^{j-1} e_{i,j}(A_{i,j}) \right) \in \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; \mathcal{I}).$$

For intervals $\mathcal{I} \prec \mathcal{J} \subset [0\,..\,n]$ and $F \in \mathrm{RectMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}\times\mathcal{J})$, we define the group element:

$$\{F\} := \prod_{i=\min(\mathcal{I})}^{\min(\mathcal{J})} \prod_{j=\min(\mathcal{J})}^{\max(\mathcal{J})} e_{i,j}(F_{i,j}) \in \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [\min(\mathcal{I})\,..\,\max(\mathcal{J})]).$$

We always think of the matrices A and F as "carrying" the appropriate type information (triangle vs. rectangle) and the index sets $\mathcal{I}$ (and $\mathcal{J}$ in the rectangular case), and therefore use the same notation $\{\cdot\}$ to unambiguously denote group elements corresponding to arbitrary matrices. $\Diamond$

Note that by the trivial Steinberg relations, all elements in the definition of $\{F\}$ commute, and therefore their order is not important. (However, the ordering of elements in $\{A\}$ is important.)

We have the following matrix version of the Steinberg relations, applied to the group elements defined in the prior definition:

Proposition 4.36 (Matrix Steinberg relations)**.** *Let $n, \kappa \in \mathbb{N}$; we have the following:*

1. Homogeneous triangle relations: *For every interval $\mathcal{I} \subseteq [0\,..\,n]$,*
$$\forall A, A' \in \mathrm{TriMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}), \quad \{A\} \cdot \{A'\} = \{AA' + A + A'\}.$$
2. Homogeneous rectangle relations: *For all intervals $\mathcal{I} \prec \mathcal{J} \subset [0\,..\,n]$,*
$$\forall F, F' \in \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I} \times \mathcal{J}), \quad \{F\} \cdot \{F'\} = \{F + F'\}.$$
3. Heterogeneous triangle-triangle relations: *For all intervals $\mathcal{I} \prec \mathcal{J} \subset [0\,..\,n]$*
$$\forall A \in \mathrm{TriMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}), B \in \mathrm{TriMat}^{\kappa}_{\mathbb{F}}(\mathcal{J}), \quad [\{A\}, \{B\}] = \mathbb{1}.$$
4. Nontrivial heterogeneous rectangle-rectangle relations: *For all intervals $\mathcal{I} \prec \mathcal{J} \prec \mathcal{K} \subset [0\,..\,n]$,*
$$\forall F \in \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I} \times \mathcal{J}), G \in \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{J} \times \mathcal{K}), \quad [\{F\}, \{G\}] = \{FG\}.$$
5. Trivial heterogeneous rectangle-rectangle relations: *For all intervals $\mathcal{I} \prec \mathcal{J} \subset [0\,..\,n]$ and $\mathcal{K} \prec \mathcal{L} \subset [0\,..\,n]$ such that $\mathcal{I} \cap \mathcal{L} = \mathcal{J} \cap \mathcal{K} = \emptyset$,*
$$\forall F \in \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I} \times \mathcal{J}), H \in \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{K} \times \mathcal{L}), \quad [\{F\}, \{H\}] = \mathbb{1}.$$
6. Nontrivial triangle-rectangle relations: *For all intervals $\mathcal{I} \prec \mathcal{J} \subset [0\,..\,n]$,*
$$\forall A \in \mathrm{TriMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}), F \in \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I} \times \mathcal{J}), \quad [\{A\}, \{F\}] = \{AF\},$$
$$\forall B \in \mathrm{TriMat}^{\kappa}_{\mathbb{F}}(\mathcal{J}), F \in \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I} \times \mathcal{J}), \quad [\{B\}, \{F\}] = \left\{FB^{\dagger}\right\}.$$
7. Trivial triangle-rectangle relations*: For all intervals $\mathcal{I} \prec \mathcal{J} \subset [0\,..\,n]$, $\mathcal{K} \subset [0\,..\,n]$ with $\mathcal{K} \cap (\mathcal{I} \cup \mathcal{J}) = \emptyset$,*
$$\forall C \in \mathrm{TriMat}^{\kappa}_{\mathbb{F}}(\mathcal{K}), F \in \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I} \times \mathcal{J}), \quad [\{C\}, \{F\}] = \mathbb{1}.$$

The proof is just linear algebra and we omit it.

5 Proof of Theorem 1.2

We now turn to establishing Theorem 1.2, which is the main technical result of our paper. We restate it in the following formal form:

Theorem 5.1. *For every $r \in \mathbb{N}$, there exists $n_0 = O(r^2) \in \mathbb{N}$ such that for every $n \geq n_0 \in \mathbb{N}$, there exists $p_0 = O(n^2 r^2) \in \mathbb{N}$ such that for every prime $p \geq p_0 \in \mathbb{N}$ and $\kappa \in \mathbb{N}$, the local Kaufman–Oppenheim complex $\mathfrak{LA}^{\kappa}_{\mathbb{F}_p}(n)$ has r-triword expansion at least $2^{O(-r^{0.99})}$ over every group Γ.*

Our proof combines the induction method of [BLM24] (Theorem 3.30) with the following key lemmas:

Lemma 5.2 (Separated diameter bound)**.** *For every $\xi > 0$, there exist $n_0, C_0 \in \mathbb{N}$ such that the following holds. For every field $\mathbb{F}$, $\kappa \geq 1 \in \mathbb{N}$, $n \geq n_0 \in \mathbb{N}$, $W \subseteq [1\,..\,n]$ and $\ell_1 < \ell_2 \in [1\,..\,n]$ such that $W \cap [\ell_1\,..\,\ell_2] = \emptyset$ and $\ell_2 - \ell_1 \geq \xi n$, for every W-word w, the complex $(\mathfrak{LA}^{\kappa}_{\mathbb{F}}(n))_w[\{\ell_1\}, \{\ell_2\}]$ has diameter at most C_0.*

Lemma 5.3 (Separated coboundary expansion bound)**.** *For every $\xi > 0$, there exist $n_0, C_0 \in \mathbb{N}$ such that the following holds. For every field $\mathbb{F}$, $\kappa \geq 1 \in \mathbb{N}$, $n \geq n_0 \in \mathbb{N}$, $W \subseteq [1\,..\,n]$ and $\ell_1 < \ell_2 < \ell_3 \in [1\,..\,n]$ such that $W \cap [\ell_1\,..\,\ell_3] = \emptyset$ and $\ell_3 - \ell_2, \ell_2 - \ell_1 \geq \xi n$, for every W-word w, the complex $(\mathfrak{LA}^{\kappa}_{\mathbb{F}}(n))_w[\{\ell_1\}, \{\ell_2\}, \{\ell_3\}]$ has diameter at most C_0.*

Lemma 5.3 is the formal version of Theorem 1.3.

These lemmas will both follow from our key group-theoretic theorem:

Theorem 1.4 (Unipotent group presentation bounds)**.** *For every $\xi > 0$ there exist $C_0, C \in \mathbb{N}$ such that for all sufficiently large n, the following hold: For every field $\mathbb{F}$, parameter $\kappa \in \mathbb{N}^+$ (defining generators), and $\ell_1 < \ell_2 < \ell_3 \in [1 \mathinner{..} n]$ with $\ell_2 - \ell_1, \ell_3 - \ell_2 \geq \xi n$,*

$$\operatorname{diam}(\operatorname{Unip}_{\mathbb{F}}^{\kappa}(n); (\operatorname{Stair}_{\mathbb{F}}^{\kappa}(n; \ell_i))_{i \in \{1,2\}}) \leq C_0, \tag{1.4}$$

$$\operatorname{area}^{2C_0+1}(\operatorname{Unip}_{\mathbb{F}}^{\kappa}(n); (\operatorname{Stair}_{\mathbb{F}}^{\kappa}(n; \ell_i))_{i \in \{1,2,3\}}) \leq C. \tag{1.5}$$

The idea behind the proofs of Lemmas 5.2 and 5.3 is is to combine Proposition 3.64 and Theorem 3.71 with Theorem 1.4. However, Theorem 1.4 is only stated for a local Kaufman–Oppenheim complex, while $(\mathfrak{LA}_{\mathbb{F}}^{\kappa}(n))_w$ is not a local Kaufman–Oppenheim complex (merely a link of one). Thus, we first have to apply Proposition 4.25 to pass to a (possibly) smaller local Kaufman–Oppenheim complex.

Proof of Lemma 5.2. Let n_0', C_0 denote the constants guaranteed by Theorem 1.4 for our value of ξ. Set $n_0 \coloneqq n_0'/\xi$. Let $i \coloneqq \max((W \cap [1 \mathinner{..} \ell_1)) \cup \{0\})$ and $j \coloneqq \min((W \cap (\ell_2 \mathinner{..} n]) \cup \{n+1\})$. By construction, $i < j \in W \cup \{0, n+1\}$, $i < \ell_1$, and $j > \ell_2$. Further,

$$W \cap (i \mathinner{..} j) = \emptyset. \tag{5.4}$$

(Indeed, suppose that $a \in W \cap (i \mathinner{..} j)$. Then $a \notin [\ell_1 \mathinner{..} \ell_2]$ by assumption, hence either $a \in (i \mathinner{..} \ell_1) \cup (\ell_2 \mathinner{..} j)$. We would then have $a \in W \cap [1 \mathinner{..} \ell_1)$ (in which case $i \geq a$) or $a \in W \cap (\ell_2 \mathinner{..} n]$ (in which case $j \leq a$).)

By definition of restrictions,

$$(\mathfrak{LA}_{\mathbb{F}}^{\kappa}(n))_w[\{\ell_1\}, \{\ell_2\}] = \big((\mathfrak{LA}_{\mathbb{F}}^{\kappa}(n))_w[(i \mathinner{..} j)]\big)[\{\ell_1\}, \{\ell_2\}]$$

(since $\ell_1, \ell_2 \in (i \mathinner{..} j)$). Hence by Equation (5.4) and Proposition 4.25,

$$(\mathfrak{LA}_{\mathbb{F}}^{\kappa}(n))_w[\{\ell_1\}, \{\ell_2\}] \cong (\mathfrak{LA}_{\mathbb{F}}^{\kappa}(n'))[\{\ell_1'\}, \{\ell_2'\}]$$

where $n' \coloneqq j - i - 1$, $\ell_2' \coloneqq \ell_2 - i$ and $\ell_1' \coloneqq \ell_1 - i$. Now $\ell_2' - \ell_1' = (\ell_2 - i) - (\ell_1 - i) = \ell_2 - \ell_1 \geq \xi n$ and $n' = j - i - 1 \geq \ell_2 - \ell_1 \geq \xi n \geq \xi n_0 \geq n_0'$. So Equation (1.5) gives that

$$\operatorname{diam}((\mathfrak{LA}_{\mathbb{F}}^{\kappa}(n'))_w[\{\ell_1'\}, \{\ell_2'\}]) \leq C_0.$$

Finally, we apply Proposition 3.64. ☐

Similarly:

Proof of Lemma 5.3. Follows from Theorem 3.71 and Eq. (1.6) just as Lemma 5.2 did from Proposition 3.64 and Eq. (1.5). ☐

Given these lemmas, we can now prove:

Proof of Theorem 1.2. We pick the n_0 needed in Theorem 1.4 and ϵ_0 needed in Theorem 3.30. Given r and n, we then set p_0 large enough such that $\frac{1}{\sqrt{p-n+1}} < \epsilon_0 r^{-2}$. Then, we apply Theorem 3.30, and simply need to verify its hypotheses. Condition 1 follows from Corollary 3.60. Condition 2 (productization on the line) follows from Proposition 4.26. Condition 3 (strong symmetry) follows from Proposition 3.42. Condition 5 (diameter bound on separated singletons) is Lemma 5.2 and similarly Condition 4 (coboundary bound on separated singletons) is Lemma 5.3. ☐

6 Miscellaneous technical tools

Before proving Theorem 1.4, we develop a few more useful tools.

6.1 Diameter and area under homomorphisms

We develop some simple machinery to transfer bounds on diameter and area from one indexed subgroup family to another using surjective group homomorphisms.

Proposition 6.1 (Padding homomorphism)**.** *Let $n^{(1)}, n, n^{(2)}, \kappa \in \mathbb{N}$ and $\mathbb{F}$ be any field. Then there exists a surjective homomorphism*

$$\psi : \mathrm{Unip}^{\kappa}_{\mathbb{F}}(n^{(1)} + n + n^{(2)}) \twoheadrightarrow \mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$$

such that for every $\ell \in [1\,..\,n]$, $\psi(\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; n^{(1)} + a)) = \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; a)$.

Proof. We set $\psi \coloneqq \Phi_{\mathbb{F}}^{(n^{(1)}+n+n^{(2)};n^{(1)},n^{(1)}+n)} \circ \pi_{\mathbb{F}}^{(n^{(1)}+n+n^{(2)};n^{(1)},n^{(1)}+n)}$ (recalling Propositions 4.9 and 4.11). By Corollary 4.22, $\psi(\mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; n^{(1)} + a)) = \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n; |\mathcal{I}| - 1)a - \min(\mathcal{I})$ where $\mathcal{I} \coloneqq [n^{(1)}\,..\,n^{(1)} + n]$. Hence, $|\mathcal{I}| = n + 1$ and $\min(\mathcal{I}) = n^{(1)}$. $\square$

Proposition 6.2. *Let G, G' be groups and $\mathcal{H} = (H_i < G)_{i\in\mathcal{I}}, \mathcal{H}' = (H'_i < G')_{i\in\mathcal{I}}$ be $\mathcal{I}$-indexed subgroup families. Let $\psi : G' \to G$ be a group homomorphism such that $\psi(H'_i) \le H_i$ for every $i \in \mathcal{I}$. For every word $w' = \langle g'_1\rangle \cdots \langle g'_T\rangle$ over $\bigcup \mathcal{H}'$, define the word*

$$\Psi(w') \coloneqq \langle\psi(g'_1)\rangle \cdots \langle\psi(g'_T)\rangle$$

over $\mathcal{H}$. Then:

1. *$\Psi(w'_1 w'_2) = \Psi(w'_1)\Psi(w'_2)$ and $\Psi((w')^{-1}) = (\Psi(w'))^{-1}$.*
2. *For every word w' over $\mathcal{H}'$, $|\Psi(w')| = |w'|$.*
3. *For every word w' over $\mathcal{H}'$, $\mathrm{eval}(\Psi(w')) = \psi(\mathrm{eval}(w'))$. Hence, if w' is a relator, then so is $\Psi(w')$.*
4. *For all words x', y' over $\mathcal{H}'$, if $x' \sim y'$, then $\Psi(x') \sim \Psi(y')$.*
5. *For every common-subgroup relator $r' \in \mathcal{R}^{\mathcal{H}'}_{\mathrm{common}}(3; G')$, $\Psi(r')$ is also a common-subgroup relator ($\Psi(r') \in \mathcal{R}^{\mathcal{H}}_{\mathrm{common}}(3; G)$).*
6. *For all words x', y', r' over $\mathcal{H}'$, if y' is derived from x' via r', then $\Psi(y')$ is derived from $\Psi(x')$ via $\Psi(r')$.*

Proof. Items 1, 2 and 4 are immediate by definition. Item 3 is because $\mathrm{eval}(\Psi(w')) = \mathrm{eval}(\langle\psi(g'_1)\rangle \cdots \langle\psi(g'_T)\rangle) = \psi(g'_1)\cdots\psi(g'_T) = \psi(g'_1 \cdots g'_T) = \psi(\mathrm{eval}(w'))$ (where the third equality uses that ψ is a group homomorphism). Item 5 follows immediately from Item 3 and the definition. For Item 6, if y' is derived from x' via r', then there exist words p', q', u', v' over $\mathcal{H}'$ such that $x' \sim p'u'q'$, $y' \sim p'v'q'$, and $r' \sim$ "$u' \equiv v'$". Hence $\Psi(x') \sim \Psi(p')\Psi(u')\Psi(q')$, $\Psi(y') \sim \Psi(p')\Psi(v')\Psi(q')$, and $\Psi(r') \sim$ "$\Psi(u') \equiv \Psi(v')$" by Items 1 and 4. Since $\Psi(r')$ is a common-subgroup relator by Item 5, we conclude that $\Psi(y')$ is derived from $\Psi(x')$ via $\Psi(r')$, as desired. $\square$

Proposition 6.3. *Let G, G' be groups and $\mathcal{H} = (H_i < G)_{i\in\mathcal{I}}, \mathcal{H}' = (H'_i < G')_{i\in\mathcal{I}}$ be $\mathcal{I}$-indexed subgroup families. Let $\psi : G' \twoheadrightarrow G$ be a surjective group homomorphism mapping each H'_i onto H_i, i.e., $\psi(H'_i) \le H_i$. Then*

$$\mathrm{diam}(G; \mathcal{H}) \le \mathrm{diam}(G'; \mathcal{H}').$$

Proof. Let $C'_0 \coloneqq \mathrm{diam}(G'; \mathcal{H}')$. Consider any $x \in G$. By surjectivity of ψ, there exists $x' \in G'$ with $\psi(x') = x$. By assumption, there is some subgroup word w' over $\bigcup \mathcal{H}'$ with $\mathrm{eval}(w') = x'$ and $|w'| \le C'_0$. As in Proposition 6.2, the word $\Psi(w')$ over $\mathcal{H}$ satisfies $\mathrm{eval}(\Psi(w')) = \psi(\mathrm{eval}(w')) = \psi(w') = w$ and $|\Psi(w')| = |w'| \le C'_0$. $\square$

Proposition 6.4. *Let G, G' be groups and $\mathcal{H} = (H_i < G)_{i\in\mathcal{I}}, \mathcal{H}' = (H'_i < G')_{i\in\mathcal{I}}$ be $\mathcal{I}$-indexed subgroup families. Let $\psi : G' \twoheadrightarrow G$ be a surjective group homomorphism mapping each H'_i* surjectively *onto H_i, i.e., $\psi(H'_i) = H_i$. Then for every $C_0 \in \mathbb{N}$,*

$$\mathrm{area}^{C_0}(G; \mathcal{H}) \le \mathrm{area}^{C_0}(G'; \mathcal{H}').$$

Proof. Consider any relator w over $\mathcal{H}$ of length $|w| \le C_0$; we want to show that $\mathrm{area}_{\mathcal{H}}(w) \le \mathrm{area}^{C_0}(G'; \mathcal{H}')$.

Hence, $w = \langle g_1 \rangle \cdots \langle g_T \rangle$, each $g_t \in \bigcup \mathcal{H}$, and $\mathrm{eval}(w) = g_1 \cdots g_T = \mathbb{1} \in G$. By surjectivity, for every $t \in [T]$, there exists $g'_t \in \bigcup \mathcal{H}'$ with $\psi(g'_t) = g_t$. We can therefore define a word $w' := \langle g'_1 \rangle \cdots \langle g'_T \rangle$ over $\mathcal{H}'$. Note that (as in Proposition 6.2), $\Psi(w') = w$ by definition, and w' is a relator with $|w'| = |w| \le C_0$.

We now claim that $\mathrm{area}_{\mathcal{H}}(w) \le R := \mathrm{area}_{\mathcal{H}'}(w')$; this is sufficient because $R \le \mathrm{area}^{C_0}(G'; \mathcal{H}')$ since $|w'| = |w| \le C_0$. Indeed, by definition of area, there are words $(w')^{(0)} = w', \dots, (w')^{(R)} = e$ over $\mathcal{H}'$ and relators $(r')^{(1)}, \dots, (r')^{(R)} \in \mathcal{R}^{\mathcal{H}'}_{\mathrm{common}}(3; G')$ such that $(w')^{(i)}$ is derived from $(w')^{(i-1)}$ via $(r')^{(i)}$ for each $i \in [R]$. By Proposition 6.2, the words $w^{(i)} := \Psi((w')^{(i)})$ and $r^{(i)} := \Psi((r')^{(i)})$ satisfy $w^{(0)} = w$, $w^{(R)} = 0$, $r^{(i)} \in \mathcal{R}^{\mathcal{H}}_{\mathrm{common}}(3; G)$, and each $w^{(i)}$ is derived from $w^{(i-1)}$ via $r^{(i)}$ for each $i \in [R]$. Hence, w can be derived from $\mathcal{R}^{\mathcal{H}}_{\mathrm{common}}(3; G)$ in at most R steps, as desired. ☐

6.2 Basic matrices

We next define a useful special set of matrices satisfying certain degree-bound conditions, which we call *basic matrices*. In the subsequent sections, we will use these basic matrices to decompose operators on large sets of wires into a bounded number of "basic block" operators.

Definition 6.5 (Basic matrix)**.** Let $\mathbb{F}$ be a field, X an indeterminate, $m, s, \kappa \in \mathbb{N}$. We define the set of *$s$-bounded basic matrices*:

$$\mathrm{BasicMat}^{\kappa}(s) := \Big\{ M \in \mathbb{F}[X]^{[0..m) \times [0..m)} : \forall i, j \in [0..m), \quad \deg(M_{i,j}) \le \kappa(s + j - i) \Big\}.$$

(Here, we recall the convention that $\deg(0) = -\infty$. Hence if $s + j - i < 0$, the constraint is that $M_{i,j} = 0$.) In particular, every $M \in \mathrm{BasicMat}^{\kappa}(s)$ satisfies $\deg(M_{0,0}), \deg(M_{m-1,m-1}) \le \kappa s$; $\deg(M_{m-1,0}) \le \kappa(s + m - 1)$, and $\deg(M_{0,m-1}) \le \kappa(s - m + 1)$. ♢

Let $I_m \in \mathbb{F}[X]^{[0..m) \times [0..m)}$ denote the identity matrix.

Proposition 6.6 (Examples of basic matrices)**.** *Let $\kappa, m \in \mathbb{N}$ and $\mathbb{F}$ be any field.*

1. *For every $s \in \mathbb{N}$ and $t \in [0..\kappa s]$, $X^t I_m \in \mathrm{BasicMat}^{\kappa}(s)$. (In particular, $I_m \in \mathrm{BasicMat}^{\kappa}(0)$.)*
2. *For every $s \in \mathbb{N}$, if $M \in \mathbb{F}[X]^{[0..m) \times [0..m)}$ satisfies $\deg_{i,j}(M) \le \kappa(s - (m-1))$ for every $i, j \in [0..m)$, then $M \in \mathrm{BasicMat}^{\kappa}(s)$.*
3. *For every $s \in \mathbb{N}$, if $M, N \in \mathrm{BasicMat}^{\kappa}(s)$, then $M + N \in \mathrm{BasicMat}^{\kappa}(s)$.*
4. *For every $s, s', s \in \mathbb{N}$, if $M \in \mathrm{BasicMat}^{\kappa}(s)$, $N \in \mathrm{BasicMat}^{\kappa}(s')$, then $MN \in \mathrm{BasicMat}^{\kappa}(s + s')$.*

Proof. We calculate:

1. The only nonzero entries of $X^t I_m$ are the diagonal entries. These have degree $\deg(M_{i,i}) = t \le \kappa(s + i - i) = \kappa s$ as long as $t \le \kappa s$.
2. Note that $s - (m-1) \le s + j - i$ since $j - i \ge 0 - (m-1)$.
3. For every $i, j \in [0..m)$, $\deg((M+N)_{i,j}) \le \max\{\deg(M_{i,j}), \deg(N_{i,j})\} \le \kappa(s + j - i)$.
4. For every $i, j \in [0..m)$,

$$\deg((MN)_{i,j}) = \max_{k=0}^{m-1} (\deg(M_{i,k}) + \deg(M_{k,j})) \le \kappa \max_{k=0}^{m-1} (s + (k - i) + s' + (j - k)) = \kappa(s + s' + j - i).$$

☐

We now give a proposition, which we will use several times in the following sections, allowing us to decompose a large-bounded basic matrix into a sum of small-bounded basic matrices.

Proposition 6.7 (Decomposition into basic matrices). *Let $\mathbb{F}$ be a field, X an indeterminate, $m, s, \Delta \in \mathbb{N}$, $s \geq m$. Let $M \in \mathrm{BasicMat}^{\kappa}(\Delta+2s)$ be a $(\Delta+2s)$-bounded basic matrix, and let $M_{i,j} = \sum_{t=0}^{\kappa(\Delta+2s+j-i)} X^t \mu_{i,j}^{(t)}$ be the degree decomposition of $M_{i,j}$. Let $\tau := \lfloor \Delta/s \rfloor$. Define the matrices $M_{\mathrm{body}}^{(k)}$ ($k < \tau$), M_{gap}, and M_{tail} in $\mathbb{F}[X]^{[0\,..\,m)\times[0\,..\,m)}$ via:*

$$(M_{\mathrm{body}}^{(k)})_{i,j} := \sum_{t=0}^{\kappa s-1} X^t \mu_{i,j}^{(t+\kappa\cdot ks)}, (M_{\mathrm{gap}})_{i,j} := \sum_{t=0}^{\kappa(\Delta-\tau s)-1} X^t \mu_{i,j}^{(t+\kappa\cdot\tau s)}, (M_{\mathrm{tail}})_{i,j} := \sum_{t=0}^{\kappa(2s+j-i)} X^t \mu_{i,j}^{(t+\kappa\cdot\Delta)}.$$

Then

$$M = \sum_{k=0}^{\tau-1} X^{\kappa\cdot ks} M_{\mathrm{body}}^{(k)} + X^{\kappa\cdot\tau s} M_{\mathrm{gap}} + X^{\kappa\Delta} M_{\mathrm{tail}}$$

and $M_{\mathrm{body}}^{(k)}, M_{\mathrm{gap}}, M_{\mathrm{tail}} \in \mathrm{BasicMat}^{\kappa}(2s)$.

Proof. For the first equality, note that the decomposition partitions the terms in $M_{i,j}$ by degree: $X^{\kappa\cdot ks}(M_{\mathrm{body}}^{(k)})_{i,j}$ takes the terms with degrees in $[\kappa \cdot ks \,..\, \kappa \cdot (k+1)s - 1)$; $X^{\kappa\cdot\tau s}(M_{\mathrm{gap}})_{i,j}$ takes the terms with degrees in $[\kappa \cdot \tau s \,..\, \kappa\Delta)$; and $X^{\kappa\Delta}(M_{\mathrm{tail}})_{i,j}$ takes the terms with degrees in $[\kappa\Delta \,..\, \kappa(\Delta + 2s + j - i)]$.

For the second claim, note that all entries in $M_{\mathrm{body}}^{(k)}$ and M_{gap} have degree at most $\kappa s - 1 \leq \kappa(s + m - (m-1)) \leq \kappa(2s - (m-1))$ (using the assumption $s \geq m$); then we use Item 2 in Proposition 6.6. For M_{tail}, the (i,j) entry of M_{tail} has degree $\kappa(2s + j - i)$ and so M_{tail} is $2s$-bounded by definition. ☐

Remark 6.8. When we apply this proposition in the sequel, we will always want $s = \Omega_\rho(\Delta)$, where ρ is the "aspect ratio" parameter, so that the decomposition has size $O_\rho(1)$. ◊

6.3 Packing

Let $x, y \in \mathbb{N}$ and $z \in [0\,..\,x+y]$, we define

$$\mathrm{pack}(z; x, y) := \begin{cases} (z, 0) & z \leq x, \\ (x, z - x) & z > x. \end{cases} \tag{6.9}$$

This should be viewed as a packing of z units of an item into two buckets of size x and y, where we pack greedily into the first bucket (of size x) as much as possible.

6.4 Collecting commutators

In this subsection, we write some simple inequalities bounding the area of a "rearrangement" from commutators of products to products of commutators. We give slightly more refined bounds than we will actually need.

We first have the following trivial equivalence: For all words u, v, w over a subgroup family $\mathcal{H}$ of some group G,

$$\mathrm{area}_{\mathcal{H}}([u, v] \equiv w) = \mathrm{area}_{\mathcal{H}}(uv = wvu). \tag{6.10}$$

Fact 6.11. *Let G be a group, $\mathcal{H}$ an indexed subgroup family, and u, $v_1, \ldots, v_\ell$ words over $\mathcal{H}$. Then:*

$$\mathrm{area}_{\mathcal{H}}\left(u\left(\prod_{j=1}^{\ell} v_j\right) \equiv \left(\prod_{j=1}^{\ell} [u, v_j] v_j\right) u\right) \leq \ell|u| + \sum_{j=1}^{\ell} |v_j|.$$

Proof. Induct on ℓ. In the base case $\ell = 0$, $u = u$ is a tautology. For $\ell > 0$, we have $u\left(\prod_{j=1}^{\ell} v_j\right) = uv_1\left(\prod_{j=2}^{\ell} v_j\right)$ and therefore $\mathrm{area}_{\mathcal{H}}\left(u\left(\prod_{j=1}^{\ell} v_j\right) \equiv uv_1u^{-1}v_1^{-1}v_1u\left(\prod_{j=2}^{\ell} v_j\right)\right) \leq \mathrm{area}_{\mathcal{H}}(u^{-1}v_1^{-1}v_1u) \leq |u| + |v_1|$. Finally, $uv_1u^{-1}v_1^{-1}v_1u\left(\prod_{j=2}^{\ell} v_j\right) = [u, v_1]v_1u\left(\prod_{j=2}^{\ell} v_j\right)$ and we continue inductively. ☐

Claim 6.12. *Let G be a group, $\mathcal{H}$ an indexed subgroup family, and u, $v_1, \dots, v_\ell$ words over $\mathcal{H}$. Then:*

$$\mathrm{area}_{\mathcal{H}}\left(\left[u, \prod_{j=1}^{\ell} v_j\right] \equiv \prod_{j=1}^{\ell} [u, v_j]\right) \leq \ell|u| + \sum_{j=1}^{\ell} |v_j| + \sum_{j=1}^{\ell} \sum_{j'=j+1}^{\ell} \mathrm{area}_{\mathcal{H}}([v_j, [u, v_{j'}]]).$$

Note that this lemma essentially presupposes that $[v_{j'}, [u, v_j]] \equiv \mathbb{1}$ for every $j < j' \in [1 \,..\, \ell]$, since otherwise the RHS is ∞.

Proof. We apply Equation (6.10) and again induct on ℓ. In the base case $\ell = 0$, we again get the trivial equality $u = u$. For $\ell > 0$, we have $\prod_{j=1}^{\ell} [u, v_j] v_j = [u, v_1] v_1 \left(\prod_{j=2}^{\ell} [u, v_j] v_j\right)$. We then use:

$$\mathrm{area}_{\mathcal{H}}\left(\prod_{j=1}^{\ell} [u, v_j] v_j \equiv [u, v_1] v_1 \left(\prod_{j=2}^{\ell} [u, v_j]\right)\left(\prod_{j=2}^{\ell} v_j\right)\right) \leq \sum_{j=2}^{\ell} \sum_{j'=j+1}^{\ell} \mathrm{area}_{\mathcal{H}}([v_j, [u, v_{j'}]]),$$

$$\mathrm{area}_{\mathcal{H}}\left([u, v_1] v_1 \left(\prod_{j=2}^{\ell} [u, v_j]\right)\left(\prod_{j=2}^{\ell} v_j\right) \equiv \left(\prod_{j=1}^{\ell} [u, v_j]\right)\left(\prod_{j=1}^{\ell} v_j\right)\right) \leq \sum_{j'=2}^{\ell} \mathrm{area}_{\mathcal{H}}([v_1, [u, v_{j'}]]),$$

where the first inequality is inductive and the second since $\mathrm{area}_{\mathcal{H}}([v_1, [u, v_1]]) = \mathrm{area}_{\mathcal{H}}(v_1 [u, v_j] \equiv [u, v_j] v_1)$. Chaining these together with Claim 6.11 gives the inequality. $\square$

Claim 6.13. *Let G be a group, $\mathcal{H}$ an indexed subgroup family, and $u_1, \dots, u_k$, $v_1, \dots, v_\ell$ words over $\mathcal{H}$. Then:*

$$\begin{aligned}\mathrm{area}_{\mathcal{H}}&\left(\left[\prod_{i=1}^{k} u_i, \prod_{j=1}^{\ell} v_j\right] \equiv \prod_{i=k}^{1}\left(\prod_{j=1}^{\ell} [u_i, v_j]\right)\right) \\ &\leq \ell \sum_{i=1}^{k} |u_i| + k \sum_{j=1}^{\ell} |v_j| + \sum_{i=1}^{k} \sum_{i'=i+1}^{k} \sum_{j=1}^{\ell} \mathrm{area}_{\mathcal{H}}([u_i, [u_{i'}, v_j]]) + \sum_{i=1}^{k} \sum_{j=1}^{\ell} \sum_{j'=j+1}^{\ell} \mathrm{area}_{\mathcal{H}}([v_j, [u_i, v_{j'}]]).\end{aligned}$$

This statement may appear slightly unwieldy, but all we really will need is that the RHS is a "polynomial"; see Corollary 6.14 below.

Proof. We apply Equation (6.10) and again induct on k. The base case $k = 0$ again gives the trivial equality $\prod_{j=1}^{\ell} v_j = \prod_{j=1}^{\ell} v_j$. For $k > 0$, we have $\left(\prod_{i=1}^{k} u_i\right)\left(\prod_{j=1}^{\ell} v_j\right) = u_1 \left(\prod_{i=2}^{k} u_i\right)\left(\prod_{j=1}^{\ell} v_j\right)$. Inductively, we have:

$$\begin{aligned}\mathrm{area}_{\mathcal{H}}&\left(\left(\prod_{i=2}^{k} u_i\right)\left(\prod_{j=1}^{\ell} v_j\right) \equiv \left(\prod_{i=k}^{2}\left(\prod_{j=1}^{\ell} [u_i, v_j]\right)\right)\left(\prod_{j=1}^{\ell} v_j\right)\left(\prod_{i=2}^{k} u_i\right)\right) \\ &\leq \ell \sum_{i=2}^{k} |u_i| + (k-1) \sum_{j=1}^{\ell} |v_j| + \sum_{i=2}^{k} \sum_{i'=i+1}^{k} \sum_{j=1}^{\ell} \mathrm{area}_{\mathcal{H}}([u_i, [u_{i'}, v_j]]) + \sum_{i=2}^{k} \sum_{j=1}^{\ell} \sum_{j'=j+1}^{\ell} \mathrm{area}_{\mathcal{H}}([v_j, [u_i, v_{j'}]]).\end{aligned}$$

At the same time, we have:

$$\mathrm{area}_{\mathcal{H}}\left(u_1 \left(\prod_{i=k}^{2}\left(\prod_{j=1}^{\ell} [u_i, v_j]\right)\right) \equiv \left(\prod_{i=k}^{2}\left(\prod_{j=1}^{\ell} [u_i, v_j]\right)\right) u_1\right) \leq \sum_{i=2}^{k} \sum_{j=1}^{\ell} \mathrm{area}_{\mathcal{H}}([u_1, [u_i, v_j]]),$$

$$\mathrm{area}_{\mathcal{H}}\left(u_1 \left(\prod_{j=1}^{\ell} v_j\right) \equiv \left(\prod_{j=1}^{\ell} [u_1, v_j]\right)\left(\prod_{j=1}^{\ell} v_j\right) u_1\right) \leq \ell|u_1| + \sum_{j=1}^{\ell} |v_j| + \sum_{j=1}^{\ell} \sum_{j'=j+1}^{\ell} \mathrm{area}_{\mathcal{H}}([v_j, [u_1, v_{j'}]]),$$

where the second inequality follows from the previous claim (and the first is trivial). Chaining these together gives the inequality. $\square$

Corollary 6.14. *Let G be a group, $\mathcal{H}$ an indexed subgroup family, and $u_1, \ldots, u_k, v_1, \ldots, v_\ell$ words over $\mathcal{H}$. Then:*

$$\operatorname{area}_{\mathcal{H}}\left(\left[\prod_{i=1}^{k} u_i, \prod_{j=1}^{\ell} v_j\right] \equiv \prod_{i=k}^{1}\left(\prod_{j=1}^{\ell}[u_i, v_j]\right)\right) \leq k\ell(U + V + kA + \ell B),$$

where:

$$U := \max_i u_i, \qquad V := \max_j v_j, \qquad A := \max_{i,i',j} \operatorname{area}_{\mathcal{H}}([u_i, [u_{i'}, v_j]]), \qquad B := \max_{i,j,j'} \operatorname{area}_{\mathcal{H}}([v_j, [u_i, v_{j'}]]).$$

7 Diameter of well-separated restrictions

In this section, we prove Equation (1.5).

7.1 Setup

We actually prove the following lemma with additional technical conditions:

Lemma 7.1. *For every field $\mathbb{F}$, $\kappa \geq 1 \in \mathbb{N}$, and $n_{\mathrm{A}}, n_{\mathrm{B}}, n_{\mathrm{C}}, m \in \mathbb{N}$, with $m \geq 1$, $n_{\mathrm{B}} \geq 4m$ and $n_{\mathrm{A}} \equiv n_{\mathrm{C}} \equiv 0 \pmod{m}$, we have*

$$\operatorname{diam}(\mathrm{Unip}_{\mathbb{F}}^{\kappa}(n); (\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n; \ell_i))_{i \in \{1,2\}}) \leq O(\rho^3),$$

where $n := n_{\mathrm{A}} + n_{\mathrm{B}} + n_{\mathrm{C}} - 1$, $\ell_1 := n_{\mathrm{A}}$, $\ell_2 := n_{\mathrm{A}} + n_{\mathrm{B}}$, and $\rho := n/m$.

It is not difficult to see how Equation (1.5) follows from Lemma 7.1:

Proof of Equation (1.5) modulo Lemma 7.1. We set $n_0 := \lceil 4\xi^{-1} \rceil$ and $C_0 := O((4\xi^{-1} + 2)^3)$ (with the same constants as in Lemma 7.1). Now given an input, set $m := \lfloor \frac{1}{4}(\ell_2 - \ell_1) \rfloor$. We therefore have $m \geq \frac{1}{4}(\ell_2 - \ell_1) \geq \frac{1}{4}\xi n \geq \frac{1}{4}\xi n_0 \geq 1$ and $\ell_2 - \ell_1 \geq 4m$.

Next, let $n_{\mathrm{A}} := m\lceil \ell_1/m \rceil$, $n_{\mathrm{B}} := \ell_2 - \ell_1$, and $n_{\mathrm{C}} := m\lceil (n - \ell_2 + 1)/m \rceil$; these fulfill the hypotheses of Lemma 7.1 and therefore

$$\operatorname{diam}(\mathrm{Unip}_{\mathbb{F}}^{\kappa}(n'); (\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n'; \ell_i'))_{i \in \{1,2\}}) \leq O(\rho^3),$$

where $n' := n_{\mathrm{A}} + n_{\mathrm{B}} + n_{\mathrm{C}} - 1$, $\ell_1' := n_{\mathrm{A}}$, $\ell_2' := n_{\mathrm{A}} + n_{\mathrm{B}}$, and $\rho := n'/m$. Letting $n^{(1)} := m\lceil \ell_1/m \rceil - \ell_1$ and $n^{(2)} := (m\lceil (n - \ell_2 + 1)/m \rceil - (n - \ell_2 + 1)$, we have $n' - n = n^{(1)} + n^{(2)} \leq 2m$ and therefore $\rho \leq \frac{n}{m} + 2 \leq 4\xi^{-1} + 2$. Finally, we apply Propositions 6.1 and 6.3 to give the desired bound. $\square$

In the remainder of this section, we fix $\mathbb{F}, \kappa, n, m, n_{\mathrm{A}}, n_{\mathrm{B}}, n_{\mathrm{C}}, \ell_1, \ell_2, \rho$ as in Lemma 7.1, and prove Lemma 7.1. We define

$$\begin{aligned} \mathcal{I}_{\mathrm{A}} &:= [0\,..\,n_{\mathrm{A}}), & & |\mathcal{I}_{\mathrm{A}}| = n_{\mathrm{A}}, \\ \mathcal{I}_{\mathrm{B}} &:= [n_{\mathrm{A}}\,..\,n_{\mathrm{A}} + n_{\mathrm{B}}), & \text{so that} \quad & |\mathcal{I}_{\mathrm{B}}| = n_{\mathrm{B}}, \\ \mathcal{I}_{\mathrm{C}} &:= [n_{\mathrm{A}} + n_{\mathrm{B}}\,..\,n_{\mathrm{A}} + n_{\mathrm{B}} + n_{\mathrm{C}}), & & |\mathcal{I}_{\mathrm{C}}| = n_{\mathrm{C}}, \\ \mathcal{I}_n &:= [0\,..\,n] & & |\mathcal{I}_n| = n + 1 \end{aligned}$$

and $\mathcal{I}_{\mathrm{A}} \sqcup \mathcal{I}_{\mathrm{B}} \sqcup \mathcal{I}_{\mathrm{C}} = \mathcal{I}_n$. We use the convention that indices in $\mathcal{I}_{\mathrm{A}}$, $\mathcal{I}_{\mathrm{B}}$, and $\mathcal{I}_{\mathrm{C}}$ are written as $\bar{a}$, $\bar{b}$, and $\bar{c}$, respectively. We abbreviate $\mathsf{GU} := \mathrm{Unip}_{\mathbb{F}}^{\kappa}(n)$, $\mathsf{GU}_{\mathrm{A}} := \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n; \ell_1) < \mathsf{GU}$, $\mathsf{GU}_{\mathrm{B}} := \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n; \ell_2) < \mathsf{GU}$, and $\mathcal{H} := \{\mathsf{GU}_{\mathrm{A}}, \mathsf{GU}_{\mathrm{B}}\}$.

7.2 Matrices and group elements

We now define the sets of matrices:

$$\begin{aligned} \mathrm{Mat}_{\mathtt{A}} &:= \mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}_{\mathrm{A}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{A}} \times \mathcal{I}_{\mathrm{A}}}, \\ \mathrm{Mat}_{\mathtt{B}} &:= \mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}_{\mathrm{B}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{B}} \times \mathcal{I}_{\mathrm{B}}}, \\ \mathrm{Mat}_{\mathtt{C}} &:= \mathrm{TriMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}_{\mathrm{C}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{C}} \times \mathcal{I}_{\mathrm{C}}}, \\ \mathrm{Mat}_{\mathtt{F}} &:= \mathrm{RectMat}_{\mathbb{F}}^{\kappa}(\mathcal{I}_{\mathrm{A}} \times \mathcal{I}_{\mathrm{A}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{A}} \times \mathcal{I}_{\mathrm{B}}}, \end{aligned}$$

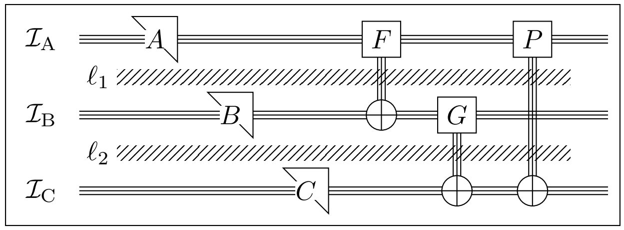

Figure 6: The circuit view of the general form of an operator in GU vis-à-vis the subgroups GU_{A} and GU_{B} (cf. Proposition 7.3): The wires $\mathcal{I}_n$ are bundled into sets $\mathcal{I}_{\mathrm{A}}$, $\mathcal{I}_{\mathrm{B}}$, and $\mathcal{I}_{\mathrm{C}}$ by the "moats" ℓ_1 and ℓ_2. We then divide up the operator's action into "blocks" which perform weighted additions within $\mathcal{I}_{\mathrm{A}}$ ($\{A\}$), within $\mathcal{I}_{\mathrm{B}}$ ($\{B\}$), within $\mathcal{I}_{\mathrm{C}}$ ($\{C\}$), from $\mathcal{I}_{\mathrm{A}}$ to $\mathcal{I}_{\mathrm{B}}$ ($\{F\}$), from $\mathcal{I}_{\mathrm{B}}$ to $\mathcal{I}_{\mathrm{C}}$ ($\{G\}$), or from $\mathcal{I}_{\mathrm{A}}$ to $\mathcal{I}_{\mathrm{C}}$ ($\{P\}$).

$$\begin{aligned}
\mathrm{Mat}_{\mathtt{G}} &:= \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{B}} \times \mathcal{I}_{\mathrm{C}}) \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{B}} \times \mathcal{I}_{\mathrm{C}}}, \\
\mathrm{Mat}_{\mathtt{P}} &:= \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{A}} \times \mathcal{I}_{\mathrm{C}}) \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{A}} \times \mathcal{I}_{\mathrm{C}}}.
\end{aligned}$$

Correspondingly, Definition 4.35 gives us group elements corresponding to such matrices, satisfying:

$$\begin{aligned}
A &\in \mathrm{Mat}_{\mathtt{A}}, & \{A\} &\in \mathsf{GU}_{\mathrm{A}} \cap \mathsf{GU}_{\mathrm{B}}, \\
B &\in \mathrm{Mat}_{\mathtt{B}}, & \{B\} &\in \mathsf{GU}_{\mathrm{A}} \cap \mathsf{GU}_{\mathrm{B}}, \\
C &\in \mathrm{Mat}_{\mathtt{C}}, & \{C\} &\in \mathsf{GU}_{\mathrm{A}} \cap \mathsf{GU}_{\mathrm{B}}, \\
F &\in \mathrm{Mat}_{\mathtt{F}}, & \{F\} &\in \mathsf{GU}_{\mathrm{B}}, \\
G &\in \mathrm{Mat}_{\mathtt{G}}, & \{G\} &\in \mathsf{GU}_{\mathrm{A}}, \\
P &\in \mathrm{Mat}_{\mathtt{P}}, & \{P\} &\in \mathsf{GU}.
\end{aligned}$$

Next, Proposition 4.36 gives:

Proposition 7.2 ((A subset of the) Steinberg relations)**.** *For every $F, F' \in \mathrm{Mat}_{\mathtt{F}}$, $G, G' \in \mathrm{Mat}_{\mathtt{G}}$, and $P, P' \in \mathrm{Mat}_{\mathtt{P}}$, we have the following identities inside GU:*

- *Homogeneous relations:*
  - *Rectangle:*

$$\{F\}\{F'\} = \{F + F'\}, \qquad \{G\}\{G'\} = \{G + G'\}, \qquad \{P\}\{P'\} = \{P + P'\}.$$

- *Heterogeneous relations:*
  - *Rectangle-rectangle:*

$$[\{F\}, \{G\}] = \{FG\}, \qquad [\{F\}, \{P\}] = \mathbb{1}, \qquad [\{G\}, \{P\}] = \mathbb{1}.$$

Proposition 7.3. *The following hold for elements of GU and its subgroups:*

- *Every element of $\mathrm{Unip}^{\kappa}_{F}(n)$ can be written uniquely as $\{A\}\{B\}\{C\}\{F\}\{G\}\{P\}$ for some $A \in \mathrm{Mat}_{\mathtt{A}}, B \in \mathrm{Mat}_{\mathtt{B}}, C \in \mathrm{Mat}_{\mathtt{C}}, F \in \mathrm{Mat}_{\mathtt{F}}, G \in \mathrm{Mat}_{\mathtt{G}}, P \in \mathrm{Mat}_{\mathtt{P}}$.*
- *Every element of GU_{A} can be written uniquely as $\{A\}\{B\}\{C\}\{G\}$ for some $A \in \mathrm{Mat}_{\mathtt{A}}, B \in \mathrm{Mat}_{\mathtt{B}}, C \in \mathrm{Mat}_{\mathtt{C}}, G \in \mathrm{Mat}_{\mathtt{G}}$.*
- *Every element of GU_{B} can be written uniquely as $\{A\}\{B\}\{C\}\{F\}$ for some $A \in \mathrm{Mat}_{\mathtt{A}}, B \in \mathrm{Mat}_{\mathtt{B}}, C \in \mathrm{Mat}_{\mathtt{C}}, F \in \mathrm{Mat}_{\mathtt{F}}$.*

Proof. Follows from Corollary 4.28. ∎

See Figure 6 for a graphical depiction of Proposition 7.3.

In light of Proposition 7.3, Lemma 7.1 follows from the following simpler lemma:

Lemma 7.4. *For every $P \in \mathrm{Mat}_{\mathsf{P}}$, there exists a word w over $\mathcal{H}$ of length $O(\rho^3)$ evaluating to $\{P\}$.*

Proof of Lemma 7.1 modulo Lemma 7.4. We want to show that every element of GU can be expressed a word over $\mathcal{H}$ of length at most $O(\rho^3)$. By Proposition 7.3, every element of GU can be expressed as

$$\underbrace{\{A\}\{B\}\{C\}\{F\}}_{\in \mathsf{GU}_{\mathrm{B}}}\underbrace{\{G\}}_{\in \mathsf{GU}_{\mathrm{A}}}\{P\},$$

and then we apply Lemma 7.4 to $\{P\}$. □

7.3 Decomposing wires into blocks

We next turn to proving Lemma 7.4. We essentially need to divide up the wire-bands $\mathcal{I}_{\mathrm{A}}$ and $\mathcal{I}_{\mathrm{C}}$ into *blocks* of size m. Towards this, we define:

$$\begin{aligned}
\mathcal{S}_{\mathrm{A}} &\coloneqq [0\,..\,n_{\mathrm{A}} - m] && \subset \mathcal{I}_{\mathrm{A}},\\
\mathcal{S}_{\mathrm{B}} &\coloneqq [n_{\mathrm{A}}\,..\,n_{\mathrm{A}} + n_{\mathrm{B}} - m] && \subset \mathcal{I}_{\mathrm{B}},\\
\mathcal{S}_{\mathrm{C}} &\coloneqq [n_{\mathrm{A}} + n_{\mathrm{B}}\,..\,n_{\mathrm{A}} + n_{\mathrm{B}} + n_{\mathrm{C}} - m] && \subset \mathcal{I}_{\mathrm{C}}.
\end{aligned}$$

We typically use the variables a, b, and c to refer to elements of these subsets, respectively.

For $a \in \mathcal{S}_{\mathrm{A}}$, we define $\mathcal{I}_{\mathrm{A}}^a \coloneqq [a\,..\,a+m) \subset \mathcal{I}_{\mathrm{A}}$. We think of this as a contiguous block of m wires starting at wire a. Similarly, we define $\mathcal{I}_b^m \subset \mathcal{I}_{\mathrm{B}}$ for every $b \in \mathcal{S}_{\mathrm{B}}$, and $\mathcal{I}_{\mathrm{C}}^c \subset \mathcal{I}_{\mathrm{C}}$ for every $b \in \mathcal{S}_{\mathrm{C}}$.

Next, we divide up $\mathcal{I}_{\mathrm{A}}$ and $\mathcal{I}_{\mathrm{C}}$ in a canonical way. Recalling that n_{A} and n_{C} are multiples of m, we define:

$$\left.\begin{aligned}
a(0) &\coloneqq 0,\\
a(1) &\coloneqq m,\\
&\dots,\\
a(n_{\mathrm{A}}/m - 1) &\coloneqq (n_{\mathrm{A}}/m - 1)m = n_{\mathrm{A}} - m
\end{aligned}\right\} \in \mathcal{S}_{\mathrm{A}} \quad\text{and}\quad \left.\begin{aligned}
c(0) &\coloneqq n_{\mathrm{A}} + n_{\mathrm{B}},\\
c(1) &\coloneqq n_{\mathrm{A}} + n_{\mathrm{B}} + m,\\
&\dots,\\
c(n_{\mathrm{C}}/m - 1) &\coloneqq n_{\mathrm{A}} + n_{\mathrm{B}} + (n_{\mathrm{C}}/m - 1)m,
\end{aligned}\right\} \in \mathcal{S}_{\mathrm{C}}. \tag{7.5}$$

We view the former as inducing a partition $\mathcal{I}_{\mathrm{A}}^{a(0)}, \mathcal{I}_{\mathrm{A}}^{a(1)}, \dots, \mathcal{I}_{\mathrm{A}}^{a(n_{\mathrm{A}}/m-1)}$ of $\mathcal{I}_{\mathrm{A}}$, and the latter as inducing a partition $\mathcal{I}_{\mathrm{C}}^{c(0)}, \mathcal{I}_{\mathrm{C}}^{c(1)}, \dots, \mathcal{I}_{\mathrm{C}}^{c(n_{\mathrm{C}}/m-1)}$ of $\mathcal{I}_{\mathrm{C}}$. We will only use this decomposition briefly in the proof of Claim 7.9, but it will serve as a useful point of comparison for decompositions we perform in the following section.

7.4 Type-P basic block matrices and decomposing type-P matrices

Definition 7.6 (Type-P basic block matrix)**.** For every $s \in \mathbb{N}$, $a \in \mathcal{S}_{\mathrm{A}}$, $c \in \mathcal{S}_{\mathrm{C}}$, $t_{\mathrm{P}} \in [0\,..\,c - a - s]$, and s-bounded basic matrix $M \in \mathrm{BasicMat}^{\kappa}(s)$, we define the *type-P s-bounded basic block matrix* $P_{a;c}^{t_{\mathrm{P}}}(M) \in \mathbb{F}[X]^{\mathcal{I}_{\mathrm{A}} \times \mathcal{I}_{\mathrm{C}}}$ via

$$(P_{a;c}^{t_{\mathrm{P}}}(M))_{\bar{a},\bar{c}} \coloneqq \begin{cases} X^{\kappa t_{\mathrm{P}}} M_{\bar{a}-a,\bar{c}-c} & (\bar{a},\bar{c}) \in \mathcal{I}_{\mathrm{A}}^a \times \mathcal{I}_{\mathrm{C}}^c,\\ 0 & \text{otherwise.}\end{cases}$$

Further, $\deg((P_{a;c}^{t_{\mathrm{P}}}(M))_{\bar{a},\bar{c}}) \le \kappa t_{\mathrm{P}} + \deg(M_{\bar{a}-a,\bar{c}-c}) \le \kappa(t_{\mathrm{P}} + s + (\bar{c} - c) - (\bar{a} - a)) \le \kappa(\bar{c} - \bar{a})$ and so $P_{a;c}^{t_{\mathrm{P}}}(M) \in \mathrm{Mat}_{\mathsf{P}}$, i.e., $P_{a;c}^{t_{\mathrm{P}}}(M)$ is a type-P matrix. We call the corresponding group element $\left\{P_{a;c}^{t_{\mathrm{P}}}(M)\right\}$ a *type-P s-bounded basic block element.* ◊

Remark 7.7. The most important part of this definition is that t_{P} is assumed to be in the interval $[0\,..\,c-a-s]$. (This assumption is necessary in order to guarantee that $P_{a;c}^{t_{\mathrm{P}}}(M) \in \mathcal{P}$.) In particular, we will use that at $s = 2m$, this interval is always nonempty because $c \ge n_{\mathrm{A}} + n_{\mathrm{B}} \ge n_{\mathrm{A}} + m \ge a + 2m$. ◊

Claim 7.8 (Decomposing type-P block matrices)**.** *Let $a \in \mathcal{S}_{\mathrm{A}}$ and $c \in \mathcal{S}_{\mathrm{C}}$. Suppose $P \in \mathrm{Mat}_{\mathsf{P}}$ is supported only on the $\mathcal{I}_{\mathrm{A}}^a \times \mathcal{I}_{\mathrm{C}}^c$ entries. Then P is a sum of $\lfloor (c-a)/m \rfloor = O(\rho)$ $2m$-bounded basic block matrices (i.e., matrices of the form $P_{a;c}^{t_{\mathrm{P}}}(M)$ for some $t_{\mathrm{P}} \in [0\,..\,c - a - 2m]$, $M \in \mathrm{BasicMat}^{\kappa}(2m)$).*

Proof. Let $M_{i,j} \coloneqq P_{a+i,c+j}$. Hence $\deg(M_{i,j}) = \deg(P_{a+i,c+j}) \le \kappa((c+j)-(a+i))$. Hence M is $(c-a)$-bounded and $P = P^0_{a;c}(M)$.

Now, letting $\Delta \coloneqq c-a-2m$, $s \coloneqq m$, $\tau \coloneqq \lfloor \Delta/m \rfloor$, M is $(\Delta+2s)$-bounded and Proposition 6.7 lets us decompose M as

$$M = \sum_{k=0}^{\tau-1} X^{\kappa\cdot km} M^{(k)}_{\text{body}} + X^{\kappa\cdot\tau m} M_{\text{gap}} + X^{\kappa\Delta} M_{\text{tail}}$$

for $M^{(k)}_{\text{body}}, M_{\text{gap}}, M_{\text{tail}} \in \text{BasicMat}^\kappa(2m)$. Correspondingly,

$$P = P^0_{a;c}(M) = \sum_{k=0}^{\tau-1} P^0_{a;c}\Big(X^{\kappa\cdot km} M^{(k)}_{\text{body}}\Big) + P^0_{a;c}(X^{\kappa\cdot\tau m} M_{\text{gap}}) + P^0_{a;c}\big(X^{\kappa\Delta} M_{\text{tail}}\big)$$
$$= \sum_{k=0}^{\tau-1} P^{km}_{a;c}\Big(M^{(k)}_{\text{body}}\Big) + P^{\tau m}_{a;c}(M_{\text{gap}}) + P^{\Delta}_{a;c}(M_{\text{tail}}).$$

This decomposition uses $\tau + 2 = O(\rho)$ matrices, as desired. ☐

We combine this with the following:

Claim 7.9 (Decomposing type-$\mathtt{P}$ matrices)**.** *Every $P \in \text{Mat}_{\mathtt{P}}$ is the sum of $O(\rho^2)$ type-$\mathtt{P}$ $2m$-bounded basic block matrices.*

Proof. By partitioning P into blocks according to the partition in Equation (7.5), we can write $P = \sum_{\alpha=0}^{n_{\text{A}}/m-1} \sum_{\gamma=0}^{n_{\text{C}}/m-1} P^{(\alpha,\gamma)}$ where $P^{(\alpha,\gamma)}$ is supported only on the $\mathcal{I}^{a(\alpha)}_{\text{A}} \times \mathcal{I}^{c(\gamma)}_{\text{C}}$ entries. ☐

By the commutator relation for type-$\mathtt{F}$ and -$\mathtt{G}$ elements (Proposition 7.2), we deduce:

Corollary 7.10. *For every $P \in \mathcal{P}$, $\{P\}$ can be expressed as a product of at most $O(\rho^3)$ type-$\mathtt{P}$ $2m$-bounded basic block elements (i.e., elements of the form $\big\{P^{t_{\text{P}}}_{a;c}(M)\big\}$ for some $a \in \mathcal{S}_{\text{A}}$, $c \in \mathcal{S}_{\text{C}}$, $t_{\text{P}} \in [0\,..\,c-a-2m]$, $M \in \text{BasicMat}^\kappa(2m)$).*

7.5 Type-$\mathtt{F}$ (and -$\mathtt{G}$) basic block matrices

Definition 7.11 (Type-$\mathtt{F}$ basic block matrix)**.** For every $s \in \mathbb{N}$, $a \in \mathcal{S}_{\text{A}}$, $b \in \mathcal{S}_{\text{B}}$, $t_{\text{F}} \in [0\,..\,b-a-s]$, and $M \in \text{BasicMat}^\kappa(s)$, we define $F^{t_{\text{F}}}_{a;b}(M)$ via

$$(F^{t_{\text{F}}}_{a;b}(M))_{\bar{a},\bar{b}} \coloneqq \begin{cases} X^{\kappa f} M_{\bar{a}-a,\bar{b}-b} & (\bar{a},\bar{b}) \in \mathcal{I}^a_{\text{A}} \times \mathcal{I}^m_b, \\ 0 & \text{otherwise.} \end{cases}$$

Further, $\deg((F^{t_{\text{F}}}_{a;b}(M))_{\bar{a},\bar{b}}) \le \kappa t_{\text{F}} + \deg(M_{\bar{a}-a,\bar{b}-b}) \le \kappa(t_{\text{F}} + s + (\bar{b}-\bar{a}) - (b-a)) \le \kappa(\bar{b}-\bar{a})$ and hence $F^{t_{\text{F}}}_{a;b}(M) \in \text{Mat}_{\mathtt{F}}$. We call the corresponding group element $\Big\{F^{t_{\text{F}}}_{a;b}(M)\Big\}$ a *type-$\mathtt{F}$ s-bounded basic block element.* For $b \in \mathcal{S}_{\text{B}}$, $c \in \mathcal{S}_{\text{C}}$, $t_{\text{G}} \in [0\,..\,c-b-s]$, and $M \in \text{BasicMat}^\kappa(s)$, we similarly define $G^{t_{\text{G}}}_{b;c}(M)$ and $\Big\{G^{t_{\text{G}}}_{b;c}(M)\Big\}$. ◊

Remark 7.12. There is an important difference between this definition and the definition of type-$\mathtt{P}$ bounded basic block elements: E.g. for type-$\mathtt{F}$ elements, at $s = 2m$, the interval $[0\,..\,b-a-s]$ might be empty, since we can have $b = n_{\text{A}}$ and $a = n_{\text{A}} - m$. ◊

Fact 7.13. *For every $s, s' \in \mathbb{N}$, $a \in \mathcal{S}_{\text{A}}$, $b \in \mathcal{S}_{\text{B}}$, $c \in \mathcal{S}_{\text{C}}$, $t_{\text{F}} \in [0\,..\,b-a-s]$, $t_{\text{G}} \in [0\,..\,c-b-s']$, $M \in \text{BasicMat}^\kappa(s)$, and $N \in \text{BasicMat}^\kappa(s')$, we have: $F^{t_{\text{F}}}_{a;b}(M) \cdot G^{t_{\text{G}}}_{b;c}(n) = P^{t_{\text{F}}+t_{\text{G}}}_{a;c}(MN)$.*

Proof. Direct calculation of a product of block matrices. ☐

Lemma 7.14. *For every $a \in \mathcal{S}_{\text{A}}$ and $c \in \mathcal{S}_{\text{C}}$, $t_{\text{P}} \in [0\,..\,c-a-2m]$, and $M \in \text{BasicMat}^\kappa(2m)$, there exist $F \in \text{Mat}_{\mathtt{F}}$ and $G \in \text{Mat}_{\mathtt{G}}$ such that $[\{F\},\{G\}] = \big\{P^{t_{\text{P}}}_{a;c}(M)\big\}$.*

Proof. Fix arbitrary $b' \in [n_{\mathrm{A}} + m \,..\, n_{\mathrm{A}} + n_{\mathrm{B}} - m]$ with $\mathcal{I}_b^m \cap \mathcal{I}_{b'}^m = \emptyset$. (This is possible since $n_{\mathrm{B}} \geq 4m$. For instance, case on whether $b \leq n_{\mathrm{A}} + 2m$. If so, then $b = n_{\mathrm{A}} + 3m$ works, and if not, then $b' = n_{\mathrm{A}} + 2m$ works.) Also, fix arbitrary $t_{\mathrm{F}}, t_{\mathrm{G}} \in \mathbb{N}$ such that

$$\begin{aligned} &t_{\mathrm{F}} \in [0 \,..\, b' - a - 2m], \\ &t_{\mathrm{G}} \in [0 \,..\, c - b'], \\ &t_{\mathrm{P}} = t_{\mathrm{F}} + t_{\mathrm{G}}, \end{aligned}$$

which is possible since

$$\begin{aligned} (b' - a - 2m) + (c - b') = c - a - 2m &\geq t_{\mathrm{P}}, \\ b' - a - 2m \geq (n_{\mathrm{A}} + m) - (n_{\mathrm{A}} - m) - 2m &= 0, \\ c - b' \geq (n_{\mathrm{A}} + n_{\mathrm{B}} + m) - (n_{\mathrm{A}} + n_{\mathrm{B}} - m) = 2m &\geq 0. \end{aligned}$$

For instance, one can take $(t_{\mathrm{F}}, t_{\mathrm{G}}) = \mathrm{pack}(t_{\mathrm{P}}; b' - a - 2m, c - b')$.

Now, define

$$F \coloneqq F_{a;b'}^{t_{\mathrm{F}}}(M) \text{ and } G \coloneqq G_{b';c}^{t_{\mathrm{G}}}(I_m)$$

so that

$$FG = P_{a;c}^{t_{\mathrm{F}} + t_{\mathrm{G}}}(M) = P_{a;c}^{t_{\mathrm{P}}}(M),$$

as desired. □

Finally, we can give:

Proof of Lemma 7.4. Let $P \in \mathrm{Mat}_{\mathtt{P}}$. Let $R = O(\rho^3)$ denote the size of the product decomposition of $\{P\}$ into type-$\mathtt{P}$ $2m$-bounded basic block elements in Corollary 7.10:

$$\{P\} = \prod_{i=1}^{R} \left\{ P_{a^{(i)};c^{(i)}}^{t_{\mathrm{P}}^{(i)}} \left(M^{(i)} \right) \right\},$$

where $a^{(i)} \in \mathcal{S}_{\mathrm{A}}$, $c^{(i)} \in \mathcal{S}_{\mathrm{C}}$, $t_{\mathrm{P}}^{(i)} \in [0 \,..\, c^{(i)} - a^{(i)} - 2m]$, and $M^{(i)} \in \mathrm{BasicMat}^{\kappa}(2m)$. By Lemma 7.14, for each $i \in [R]$, there exists $F^{(i)} \in \mathrm{Mat}_{\mathtt{F}}, G^{(i)} \in \mathrm{Mat}_{\mathtt{G}}$ such that $\left\{ P_{a^{(i)};c^{(i)}}^{t_{\mathrm{P}}^{(i)}} \left(M^{(i)} \right) \right\} = \left[\left\{ F^{(i)} \right\}, \left\{ G^{(i)} \right\} \right]$. Hence

$$\{P\} = \prod_{i=1}^{R} \left[\underbrace{\left\{ F^{(i)} \right\}}_{\ni \mathsf{GU}_{\mathrm{B}}}, \underbrace{\left\{ G^{(i)} \right\}}_{\ni \mathsf{GU}_{\mathrm{A}}} \right],$$

which is a product of length $4R$. □

8 Area of well-separated restrictions

In this section, we prove Equation (1.6). The derivation follows the template, outlined in §2.3, of implementing the Steinberg relations on "blocks".

8.1 Setup

Again, we prove the following "padded" form of Equation (1.6):

Lemma 8.1. *For every field $\mathbb{F}$, $\kappa \geq 1 \in \mathbb{N}$, and $n_{\mathrm{A}}, n_{\mathrm{B}}, n_{\mathrm{C}}, n_{\mathrm{D}}, m, C_0 \in \mathbb{N}$, with $m \geq 1$, $n_{\mathrm{B}}, n_{\mathrm{C}} \geq 4m$, and $n_{\mathrm{A}} \equiv n_{\mathrm{D}} \equiv 0 \pmod m$,*

$$\mathrm{area}^{C_0}(\mathrm{Unip}_{\mathbb{F}}^{\kappa}(n); (\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n; \ell_i))_{i \in \{1,2,3\}}) \leq O(\rho^{17})$$

where $n \coloneqq n_{\mathrm{A}} + n_{\mathrm{B}} + n_{\mathrm{C}} + n_{\mathrm{D}} - 1$, $\ell_1 \coloneqq n_{\mathrm{A}}$, $\ell_2 \coloneqq n_{\mathrm{A}} + n_{\mathrm{B}}$, $\ell_3 \coloneqq n_{\mathrm{A}} + n_{\mathrm{B}} + n_{\mathrm{C}}$, and $\rho \coloneqq n/m$.

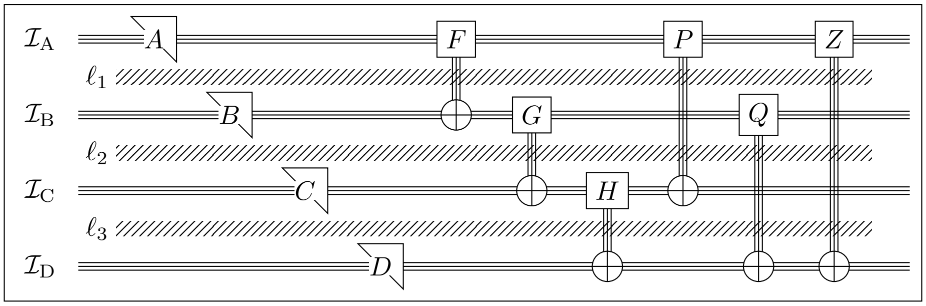

Figure 7: The circuit view of the general form of an operator in GU vis-à-vis the subgroups GU_{A}, GU_{B}, and GU_{C} (cf. Proposition 8.2): The wires $\mathcal{I}_n$ are bundled into sets $\mathcal{I}_{\mathrm{A}}$, $\mathcal{I}_{\mathrm{B}}$, $\mathcal{I}_{\mathrm{C}}$, and $\mathcal{I}_{\mathrm{D}}$ by the "moats" ℓ_1, ℓ_2, and ℓ_3. We then divide up the operator's action into "blocks" which perform weighted additions within $\mathcal{I}_{\mathrm{A}}$ ($\{A\}$), within $\mathcal{I}_{\mathrm{B}}$ ($\{B\}$), within $\mathcal{I}_{\mathrm{C}}$ ($\{C\}$), within $\mathcal{I}_{\mathrm{D}}$ ($\{D\}$), from $\mathcal{I}_{\mathrm{A}}$ to $\mathcal{I}_{\mathrm{B}}$ ($\{F\}$), from $\mathcal{I}_{\mathrm{B}}$ to $\mathcal{I}_{\mathrm{C}}$ ($\{G\}$), from $\mathcal{I}_{\mathrm{C}}$ to $\mathcal{I}_{\mathrm{D}}$ ($\{H\}$), from $\mathcal{I}_{\mathrm{A}}$ to $\mathcal{I}_{\mathrm{C}}$ ($\{P\}$), from $\mathcal{I}_{\mathrm{B}}$ to $\mathcal{I}_{\mathrm{D}}$ ($\{Q\}$), and from $\mathcal{I}_{\mathrm{A}}$ to $\mathcal{I}_{\mathrm{D}}$ ($\{Z\}$).

Proof of Equation (1.6) modulo Lemma 8.1. Same as the proof of Equation (1.5) modulo Lemma 7.1, now using Proposition 6.4. □

In the remainder of this section, we fix the parameters $\mathbb{F}, \kappa, n_{\mathrm{A}}, n_{\mathrm{B}}, n_{\mathrm{C}}, n_{\mathrm{D}}, m, C_0, n, \ell_1, \ell_2, \ell_3$ and prove Lemma 8.1.

We define

$$
\begin{aligned}
\mathcal{I}_{\mathrm{A}} &\coloneqq [0\,..\,n_{\mathrm{A}}), && & |\mathcal{I}_{\mathrm{A}}| &= n_{\mathrm{A}},\\
\mathcal{I}_{\mathrm{B}} &\coloneqq [n_{\mathrm{A}}\,..\,n_{\mathrm{A}}+n_{\mathrm{B}}), && & |\mathcal{I}_{\mathrm{B}}| &= n_{\mathrm{B}},\\
\mathcal{I}_{\mathrm{C}} &\coloneqq [n_{\mathrm{A}}+n_{\mathrm{B}}\,..\,n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}), && \text{so that} & |\mathcal{I}_{\mathrm{C}}| &= n_{\mathrm{C}},\\
\mathcal{I}_{\mathrm{D}} &\coloneqq [n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}\,..\,n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}+n_{\mathrm{D}}), && & |\mathcal{I}_{\mathrm{D}}| &= n_{\mathrm{D}},\\
\mathcal{I}_{n} &\coloneqq [0\,..\,n] && & |\mathcal{I}_{n}| &= n+1,
\end{aligned}
$$

and $\mathcal{I}_{\mathrm{A}} \sqcup \mathcal{I}_{\mathrm{B}} \sqcup \mathcal{I}_{\mathrm{C}} \sqcup \mathcal{I}_{\mathrm{D}} = \mathcal{I}_n$.

We abbreviate $\mathsf{GU} \coloneqq \mathrm{Unip}^{\kappa}_{\mathbb{F}}(n)$, $\mathsf{GU}_{\mathrm{A}} \coloneqq \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell_1) < \mathsf{GU}$, $\mathsf{GU}_{\mathrm{B}} \coloneqq \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell_2) < \mathsf{GU}$, $\mathsf{GU}_{\mathrm{C}} \coloneqq \mathrm{Stair}^{\kappa}_{\mathbb{F}}(n;\ell_3) < \mathsf{GU}$, and $\mathcal{H} \coloneqq \{\mathsf{GU}_{\mathrm{A}}, \mathsf{GU}_{\mathrm{B}}, \mathsf{GU}_{\mathrm{C}}\}$.

8.2 Matrices and group elements

We again define sets of matrices:

$$
\begin{aligned}
\mathrm{Mat}_{\mathtt{A}} &\coloneqq \mathrm{TriMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{A}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{A}}\times\mathcal{I}_{\mathrm{A}}},\\
\mathrm{Mat}_{\mathtt{B}} &\coloneqq \mathrm{TriMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{B}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{B}}\times\mathcal{I}_{\mathrm{B}}},\\
\mathrm{Mat}_{\mathtt{C}} &\coloneqq \mathrm{TriMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{C}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{C}}\times\mathcal{I}_{\mathrm{C}}},\\
\mathrm{Mat}_{\mathtt{D}} &\coloneqq \mathrm{TriMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{D}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{D}}\times\mathcal{I}_{\mathrm{D}}},\\
\mathrm{Mat}_{\mathtt{F}} &\coloneqq \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{A}}\times\mathcal{I}_{\mathrm{A}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{A}}\times\mathcal{I}_{\mathrm{B}}},\\
\mathrm{Mat}_{\mathtt{G}} &\coloneqq \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{B}}\times\mathcal{I}_{\mathrm{C}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{B}}\times\mathcal{I}_{\mathrm{C}}},\\
\mathrm{Mat}_{\mathtt{H}} &\coloneqq \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{C}}\times\mathcal{I}_{\mathrm{D}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{C}}\times\mathcal{I}_{\mathrm{D}}},\\
\mathrm{Mat}_{\mathtt{P}} &\coloneqq \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{A}}\times\mathcal{I}_{\mathrm{C}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{A}}\times\mathcal{I}_{\mathrm{C}}},\\
\mathrm{Mat}_{\mathtt{Q}} &\coloneqq \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{B}}\times\mathcal{I}_{\mathrm{D}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{B}}\times\mathcal{I}_{\mathrm{D}}},\\
\mathrm{Mat}_{\mathtt{Z}} &\coloneqq \mathrm{RectMat}^{\kappa}_{\mathbb{F}}(\mathcal{I}_{\mathrm{A}}\times\mathcal{I}_{\mathrm{D}}) && \subseteq \mathbb{F}[X]^{\mathcal{I}_{\mathrm{A}}\times\mathcal{I}_{\mathrm{D}}}.
\end{aligned}
$$

(See also Table 1.)

and we consider the corresponding group elements:

$$A \in \mathrm{Mat}_{\mathtt{A}}, \quad \{A\} \in \mathsf{GU}_{\mathrm{A}} \cap \mathsf{GU}_{\mathrm{B}} \cap \mathsf{GU}_{\mathrm{C}},$$

$$
\begin{aligned}
B &\in \mathrm{Mat}_{\mathtt{B}}, & \{B\} &\in \mathsf{GU}_{\mathrm{A}} \cap \mathsf{GU}_{\mathrm{B}} \cap \mathsf{GU}_{\mathrm{C}}, \\
C &\in \mathrm{Mat}_{\mathtt{C}}, & \{C\} &\in \mathsf{GU}_{\mathrm{A}} \cap \mathsf{GU}_{\mathrm{B}} \cap \mathsf{GU}_{\mathrm{C}}, \\
D &\in \mathrm{Mat}_{\mathtt{D}}, & \{D\} &\in \mathsf{GU}_{\mathrm{A}} \cap \mathsf{GU}_{\mathrm{B}} \cap \mathsf{GU}_{\mathrm{C}}, \\
F &\in \mathrm{Mat}_{\mathtt{F}}, & \{F\} &\in \mathsf{GU}_{\mathrm{B}} \cap \mathsf{GU}_{\mathrm{C}}, \\
G &\in \mathrm{Mat}_{\mathtt{G}}, & \{G\} &\in \mathsf{GU}_{\mathrm{A}} \cap \mathsf{GU}_{\mathrm{C}}, \\
H &\in \mathrm{Mat}_{\mathtt{H}}, & \{H\} &\in \mathsf{GU}_{\mathrm{A}} \cap \mathsf{GU}_{\mathrm{B}}, \\
P &\in \mathrm{Mat}_{\mathtt{P}}, & \{P\} &\in \mathsf{GU}_{\mathrm{C}}, \\
Q &\in \mathrm{Mat}_{\mathtt{Q}}, & \{Q\} &\in \mathsf{GU}_{\mathrm{A}}, \\
Z &\in \mathrm{Mat}_{\mathtt{Z}}, & \{Z\} &\in \mathsf{GU}.
\end{aligned}
$$

We have the following analogue of Proposition 7.3:

Proposition 8.2. *The following hold for elements of* GU *and its subgroups:*

- *Every element of* GU *can be written uniquely as* $\{A\}\{B\}\{C\}\{D\}\{F\}\{G\}\{H\}\{P\}\{Q\}\{Z\}$ *for some* $A \in \mathrm{Mat}_{\mathtt{A}}, B \in \mathrm{Mat}_{\mathtt{B}}, C \in \mathrm{Mat}_{\mathtt{C}}, D \in \mathrm{Mat}_{\mathtt{D}}, F \in \mathrm{Mat}_{\mathtt{F}}, G \in \mathrm{Mat}_{\mathtt{G}}, H \in \mathrm{Mat}_{\mathtt{H}}, P \in \mathrm{Mat}_{\mathtt{P}}, Q \in \mathrm{Mat}_{\mathtt{Q}}, Z \in \mathrm{Mat}_{\mathtt{Z}}$.
- *Every element of* GU_{A} *can be written uniquely as* $\{A\}\{B\}\{C\}\{D\}\{G\}\{H\}\{Q\}$ *for some* $A \in \mathrm{Mat}_{\mathtt{A}}, B \in \mathrm{Mat}_{\mathtt{B}}, C \in \mathrm{Mat}_{\mathtt{C}}, D \in \mathrm{Mat}_{\mathtt{D}}, G \in \mathrm{Mat}_{\mathtt{G}}, H \in \mathrm{Mat}_{\mathtt{H}}, Q \in \mathrm{Mat}_{\mathtt{Q}}$*;*
- *Every element of* GU_{B} *can be written uniquely as* $\{A\}\{B\}\{C\}\{D\}\{F\}\{H\}$ *for some* $A \in \mathrm{Mat}_{\mathtt{A}}, B \in \mathrm{Mat}_{\mathtt{B}}, C \in \mathrm{Mat}_{\mathtt{C}}, D \in \mathrm{Mat}_{\mathtt{D}}, F \in \mathrm{Mat}_{\mathtt{F}}, H \in \mathrm{Mat}_{\mathtt{H}}$.
- *Every element of* GU_{C} *can be written uniquely as* $\{A\}\{B\}\{C\}\{D\}\{F\}\{G\}\{P\}$ *for some* $A \in \mathrm{Mat}_{\mathtt{A}}, B \in \mathrm{Mat}_{\mathtt{B}}, C \in \mathrm{Mat}_{\mathtt{C}}, D \in \mathrm{Mat}_{\mathtt{D}}, F \in \mathrm{Mat}_{\mathtt{F}}, G \in \mathrm{Mat}_{\mathtt{G}}, P \in \mathrm{Mat}_{\mathtt{P}}$.

See Figure 7 for a graphical depiction of Proposition 8.2. We also have the following analogue of Proposition 7.2:

Proposition 8.3 (Steinberg relations)**.** *For every* $A, A' \in \mathrm{Mat}_{\mathtt{A}}$*,* $B, B' \in \mathrm{Mat}_{\mathtt{B}}$*,* $C, C' \in \mathrm{Mat}_{\mathtt{C}}$*,* $D, D' \in \mathrm{Mat}_{\mathtt{D}}$*,* $F, F' \in \mathrm{Mat}_{\mathtt{F}}$*,* $G, G' \in \mathrm{Mat}_{\mathtt{G}}$*,* $H, H' \in \mathrm{Mat}_{\mathtt{H}}$*,* $P, P' \in \mathrm{Mat}_{\mathtt{P}}$*,* $Q, Q' \in \mathrm{Mat}_{\mathtt{Q}}$*, and* $Z, Z' \in \mathrm{Mat}_{\mathtt{Z}}$*, we have:*

- *Homogeneous relations:*
  - *Rectangle:*

$$
\begin{aligned}
\{F\}\{F'\} &= \{F + F'\}, & \{G\}\{G'\} &= \{G + G'\}, & \{H\}\{H'\} &= \{H + H'\} \\
\{P\}\{P'\} &= \{P + P'\}, & \{Q\}\{Q'\} &= \{Q + Q'\}, & \{Z\}\{Z'\} &= \{Z + Z'\}.
\end{aligned}
$$

  - *Triangle:*

$$
\begin{aligned}
\{A\}\{A'\} &= \{AA' + A + A'\}, & \{B\}\{B'\} &= \{BB' + B + B'\} \\
\{C\}\{C'\} &= \{CC' + C + C'\}, & \{D\}\{D'\} &= \{DD' + D + D'\}.
\end{aligned}
$$

- *Heterogeneous relations:*
  - *Rectangle-rectangle:*

$$
\begin{aligned}
[\{F\}, \{G\}] &= \{FG\}, & [\{F\}, \{P\}] &= \mathbb{1}, & [\{G\}, \{P\}] &= \mathbb{1}, \\
[\{G\}, \{H\}] &= \{GH\}, & [\{G\}, \{Q\}] &= \mathbb{1}, & [\{H\}, \{Q\}] &= \mathbb{1}, \\
[\{F\}, \{H\}] &= \mathbb{1}, & & & & \\
[\{P\}, \{Q\}] &= \mathbb{1}, & [\{F\}, \{Q\}] &= \{FQ\}, & [\{P\}, \{H\}] &= \{PH\}, \\
[\{F\}, \{Z\}] &= \mathbb{1}, & [\{G\}, \{Z\}] &= \mathbb{1}, & [\{H\}, \{Z\}] &= \mathbb{1}, \\
& & [\{P\}, \{Z\}] &= \mathbb{1}, & [\{Q\}, \{Z\}] &= \mathbb{1}.
\end{aligned}
$$

– *Triangle-triangle:*

$$\begin{aligned}
&[\{A\},\{B\}] = \mathbb{1}, && [\{B\},\{C\}] = \mathbb{1}, && [\{C\},\{D\}] = \mathbb{1},\\
&[\{A\},\{C\}] = \mathbb{1}, && [\{B\},\{D\}] = \mathbb{1}, && [\{A\},\{D\}] = \mathbb{1}.
\end{aligned}$$

– *Rectangle-triangle:*

$$\begin{aligned}
&[\{A\},\{F\}] = \{AF\}, && [\{B\},\{F\}] = \left\{FB^\dagger\right\}, && [\{C\},\{F\}] = \mathbb{1}, && [\{D\},\{F\}] = \mathbb{1},\\
&[\{B\},\{G\}] = \{BG\}, && [\{C\},\{G\}] = \left\{GC^\dagger\right\}, && [\{A\},\{G\}] = \mathbb{1}, && [\{D\},\{G\}] = \mathbb{1},\\
&[\{C\},\{H\}] = \{CH\}, && [\{D\},\{H\}] = \left\{HD^\dagger\right\}, && [\{A\},\{H\}] = \mathbb{1}, && [\{B\},\{H\}] = \mathbb{1},\\
&[\{A\},\{P\}] = \{AP\}, && [\{C\},\{P\}] = \left\{PC^\dagger\right\}, && [\{B\},\{P\}] = \mathbb{1}, && [\{D\},\{P\}] = \mathbb{1},\\
&[\{B\},\{Q\}] = \{BQ\}, && [\{D\},\{Q\}] = \left\{QD^\dagger\right\}, && [\{A\},\{Q\}] = \mathbb{1}, && [\{C\},\{Q\}] = \mathbb{1},\\
&[\{A\},\{Z\}] = \{AZ\}, && [\{D\},\{Z\}] = \left\{ZD^\dagger\right\}, && [\{B\},\{Z\}] = \mathbb{1}, && [\{C\},\{Z\}] = \mathbb{1}.
\end{aligned}$$

8.3 Reducing to Steinberg relations

Let $\mathbb{L} \coloneqq \{\mathtt{A},\mathtt{B},\mathtt{C},\mathtt{D},\mathtt{F},\mathtt{G},\mathtt{H},\mathtt{P},\mathtt{Q}\}$. (Note $\mathtt{Z}$ is *not* included.) For each $\lambda \in \mathbb{L}$ and $L \in \mathrm{Mat}_\lambda$, we use $\langle L\rangle$ to denote the subgroup element $\{L\}$ regarded as a length-1 word over $\mathcal{H}$. (The advantage of this notation is that it lets us clearly distinguish between multiplication of words over $\mathcal{H}$, i.e., concatenation, and multiplication of elements within GU.)

In this section, we prove the following reduction lemma, which can be thought of as a "higher-order" analogue of the "bubble sort" argument of [KO21, Lemma 7.3].

Lemma 8.4. *Let $R \in \mathbb{N}$. Suppose that for every $Z \in \mathrm{Mat}_\mathtt{Z}$, there exists a subgroup word* $\mathrm{alias}(Z)$ *of length at most R with* $\mathrm{eval}(\mathrm{alias}(Z)) = \{Z\}$ *such that the following holds. Let $\mathcal{R}_{\mathrm{Steinberg}}$ denote the set of subgroup relators over $\mathcal{H}$ formed by taking all equations in Proposition 8.3 and replacing:*

- *Every equality in GU (the $=$ sign) with equality of evaluations of words (the $\equiv$ sign),*
- *Every instance of $\{L\}$ with $\langle L\rangle$ for $L \in \mathrm{Mat}_\lambda$, $\lambda \in \mathbb{L}$,*
- *Every instance of $\{Z\}$ with* $\mathrm{alias}(Z)$ *for $Z \in \mathrm{Mat}_\mathtt{Z}$, and*
- *Every instance of $\mathbb{1}$ with e.*

If $\mathrm{area}_\mathcal{H}(\mathcal{R}_{\mathrm{Steinberg}}) \le R$*, then*

$$\mathrm{area}^{C_0}(G;\mathcal{H}) \le O(C_0 \cdot R),$$

where $O(\cdot)$ means absolute universal constants.

We call $\mathrm{alias}(Z)$ the *alias word* for the group element $\{Z\}$; when we are performing derivations via words over $\mathcal{H}$, we use it to "stand in" for $\{Z\}$ since $\{Z\}$ itself is not in any of the subgroups in $\mathcal{H}$.

For $S \subseteq \mathbb{L}$, we say a word over $\mathcal{H}$ is an S*-word* if every element in it is of the form $\langle L\rangle$ where $L \in \mathrm{Mat}_\lambda$ for some $\lambda \in S$. In this section $O(\cdot)$ denotes some absolute linear function. Our first lemma lets us convert arbitrary words over $\mathcal{H}$ to $\mathbb{L}$-words:

Claim 8.5. *For every $g \in \mathsf{GU}_\mathrm{A} \cup \mathsf{GU}_\mathrm{B} \cup \mathsf{GU}_\mathrm{C}$, there exists an $\mathbb{L}$-word w such that* $\mathrm{area}_\mathcal{H}(\langle g\rangle \equiv w) \le 1$ *and $|w| \le 7$.*

Note that the finiteness of the area of $\langle g\rangle \equiv w$ in particular implies that $g = \mathrm{eval}(w)$.

Proof. Suppose $g \in \mathsf{GU}_\mathrm{A}$. According to Proposition 8.2, we have $g = \{A\}\{B\}\{C\}\{D\}\{G\}\{H\}\{Q\}$ for some $A \in \mathrm{Mat}_\mathtt{A}, B \in \mathrm{Mat}_\mathtt{B}, C \in \mathrm{Mat}_\mathtt{C}, D \in \mathrm{Mat}_\mathtt{D}, G \in \mathrm{Mat}_\mathtt{G}, H \in \mathrm{Mat}_\mathtt{H}, Q \in \mathrm{Mat}_\mathtt{Q}$. Hence $\langle g\rangle \equiv \langle A\rangle\langle B\rangle\langle C\rangle\langle D\rangle\langle G\rangle\langle H\rangle\langle Q\rangle$. The proofs for the other cases are similar. ☐

Claim 8.6. *For every $\lambda \in \mathbb{L}$,* $\mathrm{area}_\mathcal{H}(\langle 0_\lambda\rangle) \le O(R)$ *(where $0_\lambda \in \mathrm{Mat}_\lambda$ is the all-zeroes type-λ matrix) and further,* $\mathrm{area}_\mathcal{H}(\mathrm{alias}(0_\mathtt{Z})) \le O(R)$.

Proof. For the first claim, we use that $0_\lambda + 0_\lambda = 0$ and therefore by the definition of $\mathcal{R}_{\text{Steinberg}}$,

$$\operatorname{area}_{\mathcal{H}}(\langle 0_\lambda\rangle \cdot \langle 0_\lambda\rangle \equiv \langle 0_\lambda\rangle) \leq R.$$

We similarly have

$$\operatorname{area}_{\mathcal{H}}(\operatorname{alias}(0_{\mathtt{Z}}) \cdot \operatorname{alias}(0_{\mathtt{Z}}) \equiv \operatorname{alias}(0_{\mathtt{Z}})) \leq R.$$

We can cancel quantities from both sides to achieve the desired forms in $O(R)$ steps (since $|\operatorname{alias}(0_{\mathtt{Z}})| \leq R$). $\square$

Next, we define a total ordering $\prec$ on $\mathbb{L}$ as the alphabetical ordering $\mathtt{A} \prec \mathtt{B} \prec \mathtt{C} \prec \mathtt{D} \prec \mathtt{F} \prec \cdots$. For $\lambda \in \mathbb{L}$, we define $\operatorname{succ}_\lambda := \{\lambda' \in \mathbb{L} : \lambda \prec \lambda'\}$ as the set of letters which are lexicographically larger than λ. Note that a $\operatorname{succ}_\lambda$ word does *not* contain any type-λ elements.

Our first claim lets us move a type-λ element from right to left across a $\operatorname{succ}_\lambda$-word w, at the cost of doubling the length of w and possibly creating a type-$\mathtt{Z}$ alias:

Claim 8.7. *Let $\lambda \in \mathbb{L}$ be any letter. For every matrix $L \in \operatorname{Mat}_\lambda$ and "prefix" $\operatorname{Succ}_\lambda$-word w_{pre}, there exists a matrix $L' \in \operatorname{Mat}_\lambda$, a $\operatorname{Succ}_\lambda$-word w'_{pre}, and a matrix $Z' \in \operatorname{Mat}_{\mathtt{Z}}$ such that*

$$\operatorname{area}_{\mathcal{H}}\big(w_{\text{pre}} \cdot \langle L\rangle \equiv \langle L'\rangle \cdot w'_{\text{pre}} \cdot \operatorname{alias}(Z')\big) \leq O(|w_{\text{pre}}|R),$$

and $|w'_{\text{pre}}| \leq 2|w_{\text{pre}}|$.

Proof. We induct on $|w_{\text{pre}}|$. In the base case, $|w_{\text{pre}}| = 0$; we set $L' := L$, $w'_{\text{pre}} := e$, and $Z := 0$, and we have $\operatorname{area}_{\mathcal{H}}(\operatorname{alias}(0) \equiv e) \leq O(R)$ by Claim 8.6.

Otherwise, peel one element off w_{pre}: Decompose $w_{\text{pre}} = y_{\text{pre}} \cdot \langle L_{\text{pre}}\rangle$, where $L_{\text{pre}} \in \operatorname{Mat}_{\lambda_{\text{pre}}}$, $\lambda_{\text{pre}} \in \operatorname{Succ}_\lambda$. (Thus, $|y_{\text{pre}}| = |w_{\text{pre}}| - 1$.) Hence $w_{\text{pre}} \cdot \langle L_{\text{pre}}\rangle = y_{\text{pre}} \cdot \langle L_{\text{pre}}\rangle \cdot \langle L\rangle$. Now, by definition of $\mathcal{R}_{\text{Steinberg}}$ and inspection of Proposition 8.3, there exists a $\operatorname{Succ}_\lambda$-word x which "commutes" $\langle L\rangle$ and $\langle L_{\text{pre}}\rangle$, i.e., such that

$$\operatorname{area}_{\mathcal{H}}(\langle L_{\text{pre}}\rangle \cdot \langle L\rangle \equiv \langle L\rangle \cdot \langle L_{\text{pre}}\rangle \cdot x) \leq R, \tag{8.8}$$

and further, x has one of the following three forms: (i) empty, (ii) $\langle L_{\text{comm}}\rangle$ for some $L_{\text{comm}} \in \operatorname{Mat}_{\lambda_{\text{comm}}}$, $\lambda_{\text{comm}} \in \operatorname{Succ}_\lambda$, or (iii) $\operatorname{alias}(Z)$ for some $Z \in \operatorname{Mat}_{\mathtt{Z}}$. (E.g., case (iii) occurs iff $\{\lambda, \lambda_{\text{pre}}\} \in \{\{\mathtt{F}, \mathtt{Q}\}, \{\mathtt{P}, \mathtt{H}\}\}$.)

Next, apply the inductive hypothesis to $y_{\text{pre}} \cdot \langle L\rangle$. Thus, there is some $L' \in \operatorname{Mat}_\lambda$, $\operatorname{Succ}_\lambda$-word y'_{pre}, and $Z \in \operatorname{Mat}_{\mathtt{Z}}$ such that

$$\operatorname{area}_{\mathcal{H}}\big(y_{\text{pre}} \cdot \langle L\rangle \equiv \langle L'\rangle \cdot y'_{\text{pre}} \cdot \operatorname{alias}(Z)\big) \leq O(|y_{\text{pre}}|R) \tag{8.9}$$

and $|y'_{\text{pre}}| \leq 2|y_{\text{pre}}|$. Consequently, there is a $\operatorname{Succ}_\lambda$-word x' and $Z' \in \operatorname{Mat}_{\mathtt{Z}}$ such that

$$\operatorname{area}_{\mathcal{H}}(x \cdot \operatorname{alias}(Z) \equiv x' \cdot \operatorname{alias}(Z')) \leq R. \tag{8.10}$$

(In cases (i) and (ii), $x' = x$ and $Z' = Z$. In case (iii), x' is empty and $Z' = Z + Z$, and we again use the definition of $\mathcal{R}_{\text{Steinberg}}$ and inspection of Proposition 8.3.) We now have:

$$\begin{aligned}
w \cdot \langle L\rangle &= y_{\text{pre}} \cdot \langle L_{\text{pre}}\rangle \cdot \langle L\rangle \\
&\equiv y_{\text{pre}} \cdot \langle L\rangle \cdot \langle L_{\text{pre}}\rangle \cdot x && \text{(Equation (8.8))} \\
&\equiv \langle L'\rangle \cdot y'_{\text{pre}} \cdot \operatorname{alias}(Z) \cdot \langle L_{\text{pre}}\rangle \cdot x && \text{(Equation (8.9))} \\
&\equiv \langle L'\rangle \cdot y'_{\text{pre}} \cdot \langle L_{\text{pre}}\rangle \cdot x \cdot \operatorname{alias}(Z) && \text{(comm. of type-}\mathtt{Z}\text{ elts. in }\mathcal{R}_{\text{Steinberg}}) \\
&\equiv \langle L'\rangle \cdot \underbrace{y'_{\text{pre}} \cdot \langle L_{\text{pre}}\rangle \cdot x'}_{=: w'_{\text{pre}}} \cdot \operatorname{alias}(Z'), && \text{(Equation (8.10))}
\end{aligned}$$

Further, $|y'_{\text{pre}}| \leq 2|y_{\text{pre}}|$ by inductive hypothesis, so $|w'_{\text{pre}}| \leq |y'_{\text{pre}}| + 2 \leq 2|y_{\text{pre}}| + 2 = 2|w_{\text{pre}}|$. Then, the four steps have costs R, $O((|w_{\text{pre}}| - 1)R)$, $2R$, and R, respectively, for a total cost of $O(|w_{\text{pre}}|R)$. $\square$

We next show how to move a type-λ element from right to left over a $(\operatorname{Succ}_\lambda \cup \{\lambda\})$-word:

Claim 8.11. *Let $\lambda \in \mathbb{L}$ be any letter. For every matrix $L \in \mathrm{Mat}_\lambda$ and "prefix" $(\mathrm{Succ}_\lambda \cup \{\lambda\})$-word w_{pre}, there exists a matrix $L' \in \mathrm{Mat}_\lambda$, a Succ_λ-word w'_{pre}, and a matrix $Z \in \mathrm{Mat}_{\mathtt{Z}}$ such that*

$$\mathrm{area}_{\mathcal{H}}\big(w_{\mathrm{pre}} \cdot \langle L\rangle \equiv \langle L'\rangle \cdot w'_{\mathrm{pre}} \cdot \mathrm{alias}(Z)\big) \le O(|w_{\mathrm{pre}}|R)$$

and $|w'_{\mathrm{pre}}| \le 2|w_{\mathrm{pre}}|$.

(The difference with the preceding proposition is that we do *not* require that w does not contain any elements of type λ.)

Proof. Let $|w_{\mathrm{pre}}|_\lambda$ denote the number of elements of the form $\langle L'\rangle$ ($L' \in \mathrm{Mat}_\lambda$) in w_{pre}. We induct on $|w_{\mathrm{pre}}|_\lambda$. If $|w_{\mathrm{pre}}|_\lambda = 0$, i.e., w_{pre} is already a Succ_λ-word, then we just apply the previous Claim 8.7.

Otherwise, we have $w_{\mathrm{pre}} = x_{\mathrm{pre}} \cdot \langle L_{\mathrm{pre}}\rangle \cdot y_{\mathrm{pre}}$ where x_{pre} is a $(\mathrm{Succ}_\lambda \cup \{\lambda\})$-word, y_{pre} is a Succ_λ-word (i.e., it does not contains elements of type λ), and $|x_{\mathrm{pre}}|_\lambda = |w_{\mathrm{pre}}|_\lambda - 1$. Hence:

$$\begin{aligned}
w \cdot \langle L\rangle &= x_{\mathrm{pre}} \cdot \langle L_{\mathrm{pre}}\rangle \cdot y_{\mathrm{pre}} \cdot \langle L\rangle \\
&\equiv x_{\mathrm{pre}} \cdot \langle L_{\mathrm{pre}}\rangle \cdot \langle L\rangle \cdot y'_{\mathrm{pre}} \cdot \mathrm{alias}(Z_{\mathrm{suff}}) && (\text{Claim 8.7 on } y_{\mathrm{pre}} \cdot \langle L\rangle)\\
&\equiv x_{\mathrm{pre}} \cdot \langle L_{\mathrm{pre}} + L\rangle \cdot y'_{\mathrm{pre}} \cdot \mathrm{alias}(Z_{\mathrm{suff}}) && (\text{lin. of type-}\lambda\text{ elts. in } \mathcal{R}_{\mathrm{Steinberg}})\\
&\equiv \langle L'\rangle \cdot x'_{\mathrm{pre}} \cdot \mathrm{alias}(Z_{\mathrm{pre}}) \cdot y'_{\mathrm{pre}} \cdot \mathrm{alias}(Z_{\mathrm{suff}}) && (\text{ind. hyp. on } x_{\mathrm{pre}} \cdot \langle L_{\mathrm{pre}} + L\rangle)\\
&\equiv \langle L'\rangle \cdot x'_{\mathrm{pre}} \cdot y'_{\mathrm{pre}} \cdot \mathrm{alias}(Z_{\mathrm{pre}}) \cdot \mathrm{alias}(Z_{\mathrm{suff}}) && (\text{comm. of type-}\mathtt{Z}\text{ elts. in } \mathcal{R}_{\mathrm{Steinberg}})\\
&\equiv \langle L'\rangle \cdot \underbrace{x'_{\mathrm{pre}} \cdot y'_{\mathrm{pre}}}_{=:w'_{\mathrm{pre}}} \cdot \mathrm{alias}\Bigg(\underbrace{Z_{\mathrm{pre}} + Z_{\mathrm{suff}}}_{=:Z'}\Bigg). && (\text{lin. of type-}\mathtt{Z}\text{ elts. in } \mathcal{R}_{\mathrm{Steinberg}})
\end{aligned}$$

We have $|y'_{\mathrm{pre}}| \le 2|y_{\mathrm{pre}}|$ and $|x'_{\mathrm{pre}}| \le 2|x_{\mathrm{pre}}|$. The costs of the steps are $O(|y_{\mathrm{pre}}|R)$, R, $O(|x_{\mathrm{pre}}|R)$, $|y'_{\mathrm{pre}}|R \le 2|y_{\mathrm{pre}}|R$, and R, respectively, for a total cost of $O((|x_{\mathrm{pre}}| + |y_{\mathrm{pre}}|)R) \le O(|w_{\mathrm{pre}}|R)$, as desired. □

Claim 8.12. *For every $\lambda \in \mathbb{L}$ and $(\mathrm{Succ}_\lambda \cup \{\lambda\})$-word w, there exists $L \in \mathrm{Mat}_\lambda$, a Succ_λ-word y, and $Z \in \mathrm{Mat}_{\mathtt{Z}}$ such that*

$$\mathrm{area}_{\mathcal{H}}(w \equiv \langle L\rangle \cdot y \cdot \mathrm{alias}(Z)) \le O(|w|R)$$

and $|y| \le 2|w|$.

Proof. If w is already a Succ_λ-word (i.e., it contains no element of the form $\langle L\rangle$ for $L \in \mathrm{Mat}_\lambda$), we set $L \coloneqq 0$, $y \coloneqq w$, and $Z \coloneqq 0$. By Claim 8.6, we already have $\mathrm{area}_{\mathcal{H}}(w \equiv \langle 0\rangle \cdot y \cdot \mathrm{alias}(0)) \le O(R)$.

Otherwise, $w = w_{\mathrm{pre}} \cdot \langle L\rangle \cdot w_{\mathrm{suff}}$ for some $(\mathrm{Succ}_\lambda \cup \{\lambda\})$-word w_{pre} and Succ_λ-word w_{suff}. We apply the prior claim (Claim 8.11) to $w_{\mathrm{pre}} \cdot \langle L\rangle$, so that there exists $L' \in \mathrm{Mat}_\lambda$, y_{pre} a Succ_λ-word, and $Z \in \mathrm{Mat}_{\mathtt{Z}}$ such that

$$\mathrm{area}_{\mathcal{H}}(w_{\mathrm{pre}} \cdot \langle L\rangle \equiv \langle L'\rangle \cdot y_{\mathrm{pre}} \cdot \mathrm{alias}(Z)) \le O(|w_{\mathrm{pre}}|R). \tag{8.13}$$

Now we have the G-equalities

$$\begin{aligned}
w_{\mathrm{pre}} \cdot \langle L\rangle \cdot w_{\mathrm{suff}} &\equiv \langle L'\rangle \cdot y_{\mathrm{pre}} \cdot \mathrm{alias}(Z) \cdot w_{\mathrm{suff}} && (\text{Equation (8.13)})\\
&\equiv \langle L'\rangle \cdot \underbrace{y_{\mathrm{pre}} \cdot w_{\mathrm{suff}}}_{=:y} \cdot \mathrm{alias}(Z). && (\text{comm. of type-}\mathtt{Z}\text{ elts. in } \mathcal{R}_{\mathrm{Steinberg}})
\end{aligned}$$

The cost of the former is $O(|w_{\mathrm{pre}}|R)$ and the latter is $|w_{\mathrm{suff}}|R$. □

With this setup, we can now prove Lemma 8.4:

Proof of Lemma 8.4. We are given a relator w over $\mathcal{H}$ and want to show that $\mathrm{area}_{\mathcal{H}}(w) \le O(|w|R)$.

Our first step is to convert w into an $\mathbb{L}$-word. In particular, for $w = \langle g_1\rangle \cdots \langle g_T\rangle$ with each $g_t \in \bigcup \mathcal{H}$, we apply Claim 8.5 to each element g_t to get $\mathbb{L}$-words $w_1, \ldots, w_T$ such that for every $t \in [T]$, $\mathrm{area}_{\mathcal{H}}(\langle g_t\rangle \equiv w_t) \le 1$ and $|w_t| \le 7$. Chaining these equalities together in any arbitrary order gives that, for the $\mathbb{L}$-relator $w^{\mathtt{A}} \coloneqq w_1 \cdots w_T$, we have $\mathrm{area}_{\mathcal{H}}(w^{\mathtt{A}}) \le |w|$ and $|w^{\mathtt{A}}| \le O(|w|)$.

Next, we progressively eliminate letters from $w^{\mathtt{A}}$ using Claim 8.12. For each letter λ in $\mathbb{L}$, we let $\mathrm{next}(\lambda)$ denote the succeeding letter to λ in the natural ordering on $\mathbb{L} \cup \{\mathtt{Z}\}$. (If $\lambda \prec \mathtt{Q}$, then $\mathrm{next}(\lambda) \in \mathbb{L}$, whereas $\mathrm{next}(\mathtt{Q}) = \mathtt{Z} \notin \mathbb{L}$.) For every $\lambda \in \mathbb{L}$, note that $\mathrm{Succ}_{\mathrm{next}(\lambda)} \cup \{\lambda\} = \mathrm{Succ}_\lambda$ (with the convention that $\mathrm{Succ}_{\mathtt{Z}} := \emptyset$). Consider the following iterative process. For λ initialized to $\mathtt{A}$, w^λ is a $(\mathrm{Succ}_\lambda \cup \{\lambda\})$-word. Hence there exists $L \in \mathrm{Mat}_\lambda$, a Succ_λ-word $w^{\mathrm{next}(\lambda)}$, and $Z^\lambda \in \mathrm{Mat}_{\mathtt{Z}}$ such that $w^{\mathrm{next}(\lambda)} \le 2|w^\lambda|$ and

$$\mathrm{area}_{\mathcal{H}}\Big(w^\lambda \equiv \langle L\rangle \cdot w^{\mathrm{next}(\lambda)} \cdot \mathrm{alias}\big(Z^\lambda\big)\Big) \le O(|w^\lambda|R).$$

We now replace λ with $\mathrm{next}(\lambda)$ and iterate (until $\lambda = \mathtt{Z}$, when we halt). In the end, $w^{\mathtt{Z}}$ is an $\emptyset$-word and therefore must be the empty word e.

Finally, we define

$$Z := Z^{\mathtt{A}} + Z^{\mathtt{B}} + Z^{\mathtt{C}} + Z^{\mathtt{D}} + Z^{\mathtt{F}} + Z^{\mathtt{G}} + Z^{\mathtt{H}} + Z^{\mathtt{P}} + Z^{\mathtt{Q}},$$

so that letting

$$\widetilde{Z} := \mathrm{alias}\big(Z^{\mathtt{A}}\big) \cdot \mathrm{alias}\big(Z^{\mathtt{B}}\big) \cdot \mathrm{alias}\big(Z^{\mathtt{C}}\big) \cdot \mathrm{alias}\big(Z^{\mathtt{D}}\big) \cdot \mathrm{alias}\big(Z^{\mathtt{F}}\big) \cdot \mathrm{alias}\big(Z^{\mathtt{G}}\big) \cdot \mathrm{alias}\big(Z^{\mathtt{H}}\big) \cdot \mathrm{alias}\big(Z^{\mathtt{P}}\big) \cdot \mathrm{alias}\big(Z^{\mathtt{Q}}\big),$$

we have $\mathrm{area}_{\mathcal{H}}\Big(\mathrm{alias}(Z) \equiv \widetilde{Z}\Big) \le 8R$.

Substituting backward, we get that

$$\mathrm{area}_{\mathcal{H}}(w \equiv \langle A\rangle\langle B\rangle\langle C\rangle\langle D\rangle\langle F\rangle\langle G\rangle\langle H\rangle\langle P\rangle\langle Q\rangle \cdot \mathrm{alias}(Z)) \le O(|w|R).$$

Finally, we claim that we must have $A = B = C = D = F = G = H = P = Q = Z = 0$. Indeed, we assumed that $\mathrm{eval}(w) \equiv \mathbb{1}$, and so

$$\{A\}\{B\}\{C\}\{D\}\{F\}\{G\}\{H\}\{P\}\{Q\}\{Z\} = \mathbb{1}$$

(using that e.g. $\mathrm{eval}(\langle A\rangle) = \{A\}$ and $\mathrm{eval}(\mathrm{alias}(Z)) = \{Z\}$). At the same time, by Claim 8.6,

$$\{0_{\mathtt{A}}\}\{0_{\mathtt{B}}\}\{0_{\mathtt{C}}\}\{0_{\mathtt{D}}\}\{0_{\mathtt{F}}\}\{0_{\mathtt{G}}\}\{0_{\mathtt{H}}\}\{0_{\mathtt{P}}\}\{0_{\mathtt{Q}}\}\{0_{\mathtt{Z}}\} = \mathbb{1}.$$

Hence the equalities follow from the uniqueness statement in Proposition 8.2. Hence by Claim 8.6, $\mathrm{area}_{\mathcal{H}}(w) \le O(|w|R)$, as desired. □

8.4 The "missing" Steinberg relations

We now introduce a final notation: For an equation "$x \equiv y$" over $\mathcal{H}$ (or the corresponding relator xy^{-1}), and $i \in \mathbb{N}$, "$\vdash_i x \equiv y$" denotes that $\mathrm{area}_{\mathcal{H}}(x \equiv y) = O(\rho^i)$. We now list the relations $w \in \mathcal{R}_{\mathrm{Steinberg}}$ that already have the common subgroup property, and therefore have $\mathrm{area}_{\mathcal{H}}(w) = 0$:

Lemma 8.14 ("Present" Steinberg relations)**.** *For every $A, A' \in \mathrm{Mat}_{\mathtt{A}}$, $B, B' \in \mathrm{Mat}_{\mathtt{B}}$, $C, C' \in \mathrm{Mat}_{\mathtt{C}}$, $D, D' \in \mathrm{Mat}_{\mathtt{D}}$, $F, F' \in \mathrm{Mat}_{\mathtt{F}}$, $G, G' \in \mathrm{Mat}_{\mathtt{G}}$, $H, H' \in \mathrm{Mat}_{\mathtt{H}}$, $P, P' \in \mathrm{Mat}_{\mathtt{P}}$, $Q, Q' \in \mathrm{Mat}_{\mathtt{Q}}$, we have:*

- *Homogeneous relations:*
  - *Rectangle:*

$$\vdash_0 \begin{cases} \langle F\rangle\langle F'\rangle \equiv \langle F + F'\rangle, & \langle G\rangle\langle G'\rangle \equiv \langle G + G'\rangle, \quad \langle H\rangle\langle H'\rangle \equiv \langle H + H'\rangle \\ \langle P\rangle\langle P'\rangle \equiv \langle P + P'\rangle, & \langle Q\rangle\langle Q'\rangle \equiv \langle Q + Q'\rangle. \end{cases}$$

  - *Triangle:*

$$\vdash_0 \begin{cases} \langle A\rangle\langle A'\rangle \equiv \langle AA' + A + A'\rangle, & \langle B\rangle\langle B'\rangle \equiv \langle BB' + B + B'\rangle \\ \langle C\rangle\langle C'\rangle \equiv \langle CC' + C + C'\rangle, & \langle D\rangle\langle D'\rangle \equiv \langle DD' + D + D'\rangle. \end{cases}$$

- *Heterogeneous relations:*

– *Rectangle-rectangle:*

$$\vdash_0 \begin{cases} [\langle F\rangle,\langle G\rangle] \equiv \langle FG\rangle \text{ and } [\langle F\rangle,\langle P\rangle] \equiv [\langle G\rangle,\langle P\rangle] \equiv e, \\ [\langle G\rangle,\langle H\rangle] \equiv \langle GH\rangle \text{ and } [\langle G\rangle,\langle Q\rangle] \equiv [\langle H\rangle,\langle Q\rangle] \equiv e, \\ [\langle F\rangle,\langle H\rangle] \equiv e. \end{cases}$$

– *Triangle-triangle:*

$$\vdash_0 \begin{cases} [\langle A\rangle,\langle B\rangle] \equiv e, \quad [\langle B\rangle,\langle C\rangle] \equiv e, \quad [\langle C\rangle,\langle D\rangle] \equiv e, \\ [\langle A\rangle,\langle C\rangle] \equiv e, \quad [\langle B\rangle,\langle D\rangle] \equiv e, \quad [\langle A\rangle,\langle D\rangle] \equiv e. \end{cases}$$

– *Rectangle-triangle:*

$$\vdash_0 \begin{cases} [\langle A\rangle,\langle F\rangle] \equiv \langle AF\rangle, & [\langle B\rangle,\langle F\rangle] \equiv \langle FB^\dagger\rangle, & [\langle C\rangle,\langle F\rangle] \equiv e, & [\langle D\rangle,\langle F\rangle] \equiv e, \\ [\langle B\rangle,\langle G\rangle] \equiv \langle BG\rangle, & [\langle C\rangle,\langle G\rangle] \equiv \langle GC^\dagger\rangle, & [\langle A\rangle,\langle G\rangle] \equiv e, & [\langle D\rangle,\langle G\rangle] \equiv e, \\ [\langle C\rangle,\langle H\rangle] \equiv \langle CH\rangle, & [\langle D\rangle,\langle H\rangle] \equiv \langle HD^\dagger\rangle, & [\langle A\rangle,\langle H\rangle] \equiv e, & [\langle B\rangle,\langle H\rangle] \equiv e, \\ [\langle A\rangle,\langle P\rangle] \equiv \langle AP\rangle, & [\langle C\rangle,\langle P\rangle] \equiv \langle PC^\dagger\rangle, & [\langle B\rangle,\langle P\rangle] \equiv e, & [\langle D\rangle,\langle P\rangle] \equiv e, \\ [\langle B\rangle,\langle Q\rangle] \equiv \langle BQ\rangle, & [\langle D\rangle,\langle Q\rangle] \equiv \langle QD^\dagger\rangle, & [\langle A\rangle,\langle Q\rangle] \equiv e, & [\langle C\rangle,\langle Q\rangle] \equiv e. \end{cases}$$

Proof. Follows by inspection of the relators in Proposition 8.3 and the subgroup memberships of the group elements. □

Because of Lemmas 8.4 and 8.14, in order to prove Lemma 8.1, we will just need to prove the following lemma:

Lemma 8.15 ("Missing" Steinberg relations)**.** *For every $P \in \mathrm{Mat}_{\mathtt{P}}, Q \in \mathrm{Mat}_{\mathtt{Q}}$, $\vdash_6 [\langle P\rangle,\langle Q\rangle] \equiv e$. Further, for every $Z \in \mathrm{Mat}_{\mathtt{Z}}$, there exists a subgroup word* $\mathrm{alias}(Z)$ *of length $O(\rho^3)$ with* $\mathrm{eval}(\mathrm{alias}(Z)) = \{Z\}$ *such that the following holds: For every $A \in \mathrm{Mat}_{\mathtt{A}}$, $B \in \mathrm{Mat}_{\mathtt{B}}$, $C \in \mathrm{Mat}_{\mathtt{C}}$, $D \in \mathrm{Mat}_{\mathtt{D}}$, $F \in \mathrm{Mat}_{\mathtt{F}}$, $G \in \mathrm{Mat}_{\mathtt{G}}$, $H \in \mathrm{Mat}_{\mathtt{H}}$, $P \in \mathrm{Mat}_{\mathtt{P}}$, $Q \in \mathrm{Mat}_{\mathtt{Q}}$, and $Z, Z' \in \mathrm{Mat}_{\mathtt{Z}}$,*

$$\vdash_{17} \begin{cases} \mathrm{alias}(Z)\cdot\mathrm{alias}(Z') \equiv \mathrm{alias}(Z+Z'), & & \\ [\langle F\rangle,\langle Q\rangle] \equiv \mathrm{alias}(FQ), & [\langle P\rangle,\langle H\rangle] \equiv \mathrm{alias}(PH), & \\ [\langle F\rangle,\mathrm{alias}(Z)] \equiv e, & [\langle G\rangle,\mathrm{alias}(Z)] \equiv e, & [\langle H\rangle,\mathrm{alias}(Z)] \equiv e, \\ [\langle P\rangle,\mathrm{alias}(Z)] \equiv e, & [\langle Q\rangle,\mathrm{alias}(Z)] \equiv e, & \\ [\langle A\rangle,\mathrm{alias}(Z)] \equiv \mathrm{alias}(AZ), & [\langle D\rangle,\mathrm{alias}(Z)] \equiv \mathrm{alias}(ZD^\dagger), & \\ [\langle B\rangle,\mathrm{alias}(Z)] \equiv e, & [\langle C\rangle,\mathrm{alias}(Z)] \equiv e, & \end{cases}$$

Proof of Lemma 8.1 modulo Lemma 8.15. Follows immediately since Lemmas 8.14 and 8.15 cover the conditions in Lemma 8.1. □

8.5 Lifted proofs

Next, we give two statements which follow from what is essentially a "lifting" argument in [OS25].

Lemma 8.16 (Lifted type-`P` and -`Q` elements commute)**.** *For every $F \in \mathrm{Mat}_{\mathtt{F}}$, $G \in \mathrm{Mat}_{\mathtt{G}}$, and $H \in \mathrm{Mat}_{\mathtt{H}}$,*

$$\vdash_0 [\langle FG\rangle,\langle GH\rangle] \equiv e.$$

Proof. We begin with eight useful equations which are all constant-length in-subgroup relations (Lemma 8.14). The following four express type-`P` and -`Q` elements as commutators:

$$\langle FG\rangle \equiv \langle -G\rangle\langle F\rangle\langle G\rangle\langle -F\rangle, \tag{8.17}$$

$$\langle GH\rangle \equiv \langle -2H\rangle\langle \tfrac{1}{2}G\rangle\langle 2H\rangle\langle -\tfrac{1}{2}G\rangle, \tag{8.18}$$

$$\langle -FG\rangle \equiv \left\langle \tfrac{1}{2}G\right\rangle\langle 2F\rangle\left\langle -\tfrac{1}{2}G\right\rangle\langle -2F\rangle, \tag{8.19}$$
$$\langle -GH\rangle \equiv \langle -H\rangle\langle -G\rangle\langle H\rangle\langle G\rangle. \tag{8.20}$$

and the following four are commutator expressions:

$$\langle -F\rangle\left\langle \tfrac{1}{2}G\right\rangle\langle 2F\rangle \equiv \langle G\rangle\langle F\rangle\left\langle -\tfrac{1}{2}G\right\rangle, \tag{8.21}$$
$$\langle 2H\rangle\left\langle -\tfrac{1}{2}G\right\rangle\langle -H\rangle \equiv \left\langle \tfrac{1}{2}G\right\rangle\langle H\rangle\langle -G\rangle, \tag{8.22}$$
$$\langle G\rangle\langle -2H\rangle\langle G\rangle \equiv \langle -H\rangle\langle 2G\rangle\langle -H\rangle, \tag{8.23}$$
$$\langle -G\rangle\langle -2F\rangle\langle -G\rangle \equiv \langle -F\rangle\langle -2G\rangle\langle -F\rangle. \tag{8.24}$$

The basic plan is now to write down $[\langle FG\rangle, \langle GH\rangle]$ and expand using Equations (8.17) to (8.20). We proceed as follows using in-subgroup relations (Lemma 8.14) and these equations:

$$\begin{aligned}
&[\langle FG\rangle, \langle GH\rangle] \\
&\equiv \langle FG\rangle\langle GH\rangle\langle -FG\rangle\langle -GH\rangle \\
&\equiv \langle FG\rangle\langle -2H\rangle\left\langle \tfrac{1}{2}G\right\rangle\langle 2H\rangle\left\langle -\tfrac{1}{2}G\right\rangle\left\langle \tfrac{1}{2}G\right\rangle\langle 2F\rangle\left\langle -\tfrac{1}{2}G\right\rangle\langle -2F\rangle\langle -GH\rangle && \text{(Equations (8.18) and (8.19))} \\
&\equiv \langle FG\rangle\langle -2H\rangle\left\langle \tfrac{1}{2}G\right\rangle\langle 2H\rangle\langle 2F\rangle\left\langle -\tfrac{1}{2}G\right\rangle\langle -2F\rangle\langle -GH\rangle && \text{(canceling } \mathtt{G}\text{'s)} \\
&\equiv \langle -G\rangle\langle F\rangle\langle G\rangle\langle -F\rangle\langle -2H\rangle\left\langle \tfrac{1}{2}G\right\rangle\langle 2H\rangle\langle 2F\rangle\left\langle -\tfrac{1}{2}G\right\rangle\langle -2F\rangle\langle -H\rangle\langle -G\rangle\langle H\rangle\langle G\rangle && \text{(Equations (8.17) and (8.20))} \\
&\equiv \langle -G\rangle\langle F\rangle\langle G\rangle\langle -2H\rangle\langle -F\rangle\left\langle \tfrac{1}{2}G\right\rangle\langle 2F\rangle\langle 2H\rangle\left\langle -\tfrac{1}{2}G\right\rangle\langle -H\rangle\langle -2F\rangle\langle -G\rangle\langle H\rangle\langle G\rangle && \text{(commuting } \mathtt{F}\text{'s and } \mathtt{H}\text{'s)} \\
&\equiv \langle -G\rangle\langle F\rangle\langle G\rangle\langle -2H\rangle\langle G\rangle\langle F\rangle\left\langle -\tfrac{1}{2}G\right\rangle\left\langle \tfrac{1}{2}G\right\rangle\langle H\rangle\langle -G\rangle\langle -2F\rangle\langle -G\rangle\langle H\rangle\langle G\rangle && \text{(Equations (8.21) and (8.22))} \\
&\equiv \langle -G\rangle\langle F\rangle\langle G\rangle\langle -2H\rangle\langle G\rangle\langle F\rangle\langle H\rangle\langle -G\rangle\langle -2F\rangle\langle -G\rangle\langle H\rangle\langle G\rangle && \text{(canceling } \mathtt{G}\text{'s)} \\
&\equiv \langle -G\rangle\langle F\rangle\langle -H\rangle\langle 2G\rangle\langle -H\rangle\langle F\rangle\langle H\rangle\langle -F\rangle\langle -2G\rangle\langle -F\rangle\langle H\rangle\langle G\rangle && \text{(Equations (8.23) and (8.24))} \\
&\equiv \langle -G\rangle\langle F\rangle\langle -H\rangle\langle 2G\rangle\langle -2G\rangle\langle -F\rangle\langle H\rangle\langle G\rangle && \text{(commuting and cancelling } \mathtt{F}\text{'s and } \mathtt{H}\text{'s)} \\
&\equiv \langle -G\rangle\langle F\rangle\langle -H\rangle\langle -F\rangle\langle H\rangle\langle G\rangle && \text{(cancelling } \mathtt{G}\text{'s)} \\
&\equiv \langle -G\rangle\langle G\rangle && \text{(commuting and cancelling } \mathtt{F}\text{'s and } \mathtt{H}\text{'s)} \\
&\equiv e, && \text{(cancelling } \mathtt{G}\text{'s)}
\end{aligned}$$

as desired. ☐

Lemma 8.25 (Interchange)**.** *For every $F \in \mathrm{Mat}_{\mathtt{F}}$, $G \in \mathrm{Mat}_{\mathtt{G}}$, and $H \in \mathrm{Mat}_{\mathtt{H}}$,*

$$\vdash_0 [\langle FG\rangle, \langle H\rangle] \equiv [\langle F\rangle, \langle GH\rangle].$$

Proof. We calculate using in-subgroup relations (Lemma 8.14) and Lemma 8.16:

$$\begin{aligned}
&\langle F\rangle\langle GH\rangle \\
&\equiv \langle F\rangle\langle G\rangle\langle H\rangle\langle -G\rangle\langle -H\rangle && \text{(commutator of } \mathtt{G} \text{ and } \mathtt{H}) \\
&\equiv \langle FG\rangle\langle G\rangle\langle F\rangle\langle H\rangle\langle -G\rangle\langle -H\rangle && \text{(commutator of } \mathtt{F} \text{ and } \mathtt{G}) \\
&\equiv \langle FG\rangle\langle G\rangle\langle H\rangle\langle F\rangle\langle -G\rangle\langle -H\rangle && \text{(commuting } \mathtt{F} \text{ and } \mathtt{H}) \\
&\equiv \langle FG\rangle\langle G\rangle\langle H\rangle\langle -FG\rangle\langle -G\rangle\langle F\rangle\langle -H\rangle && \text{(commutator of } \mathtt{F} \text{ and } \mathtt{G}) \\
&\equiv \langle FG\rangle\langle G\rangle\langle H\rangle\langle -FG\rangle\langle -G\rangle\langle -H\rangle\langle F\rangle && \text{(commuting } \mathtt{F} \text{ and } \mathtt{H}) \\
&\equiv \langle FG\rangle\langle GH\rangle\langle H\rangle\langle G\rangle\langle -FG\rangle\langle -G\rangle\langle -H\rangle\langle F\rangle && \text{(commutator of } \mathtt{G} \text{ and } \mathtt{H}) \\
&\equiv \langle FG\rangle\langle GH\rangle\langle H\rangle\langle -FG\rangle\langle -H\rangle\langle F\rangle && \text{(commuting } \mathtt{P} \text{ and } \mathtt{G} \text{ and cancelling } \mathtt{G}\text{'s)} \\
&\equiv \langle FG\rangle\langle H\rangle\langle -FG\rangle\langle -H\rangle\langle GH\rangle\langle F\rangle && \text{(commuting } \mathtt{Q} \text{ with } \mathtt{H}\text{'s and Lemma 8.16)} \\
&\equiv [\langle FG\rangle, \langle H\rangle]\langle GH\rangle\langle F\rangle
\end{aligned}$$

as desired. ☐

8.6 Decomposing wires into blocks

Define

$$\begin{aligned}
\mathcal{S}_{\mathrm{A}} &\coloneqq [0\,..\,n_{\mathrm{A}}-m] && \subset \mathcal{I}_{\mathrm{A}},\\
\mathcal{S}_{\mathrm{B}} &\coloneqq [n_{\mathrm{A}}\,..\,n_{\mathrm{A}}+n_{\mathrm{B}}-m] && \subset \mathcal{I}_{\mathrm{B}},\\
\mathcal{S}_{\mathrm{C}} &\coloneqq [n_{\mathrm{A}}+n_{\mathrm{B}}\,..\,n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}-m] && \subset \mathcal{I}_{\mathrm{C}},\\
\mathcal{S}_{\mathrm{D}} &\coloneqq [n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}\,..\,n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}+n_{\mathrm{D}}-m] && \subset \mathcal{I}_{\mathrm{D}}.
\end{aligned}$$

We typically use the variables a, b, c, and d to refer to elements of these subsets, respectively.

Recall a and d are multiples of m. Define

$$\left.\begin{aligned}
a(0) &\coloneqq 0,\\
a(1) &\coloneqq m,\\
&\dots,\\
a(n_{\mathrm{A}}/m-1) &\coloneqq (n_{\mathrm{A}}/m-1)m
\end{aligned}\right\} \in \mathcal{S}_{\mathrm{A}} \quad \text{and} \quad \left.\begin{aligned}
d(0) &\coloneqq n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}},\\
d(1) &\coloneqq n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}+m,\\
&\dots,\\
d(n_{\mathrm{D}}/m-1) &\coloneqq n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}+(n_{\mathrm{D}}/m-1)m.
\end{aligned}\right\} \in \mathcal{S}_{\mathrm{D}}. \tag{8.26}$$

Similarly, define

$$\left.\begin{aligned}
c(0) &\coloneqq n_{\mathrm{A}}+n_{\mathrm{B}},\\
c(1) &\coloneqq n_{\mathrm{A}}+n_{\mathrm{B}}+(n_{\mathrm{C}}-\lfloor n_{\mathrm{C}}/m\rfloor m),\\
c(2) &\coloneqq n_{\mathrm{A}}+n_{\mathrm{B}}+(n_{\mathrm{C}}-\lfloor n_{\mathrm{C}}/m\rfloor m)+m,\\
&\dots,\\
c(\lfloor n_{\mathrm{C}}/m\rfloor) &\coloneqq n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}-m.
\end{aligned}\right\} \in \mathcal{S}_{\mathrm{C}} \quad \text{and} \quad \left.\begin{aligned}
b(0) &\coloneqq n_{\mathrm{A}},\\
b(1) &\coloneqq n_{\mathrm{A}}+m,\\
&\dots,\\
b(\lfloor n_{\mathrm{B}}/m\rfloor-1) &\coloneqq n_{\mathrm{A}}+(\lfloor n_{\mathrm{B}}/m\rfloor-1)m,\\
b(\lfloor n_{\mathrm{B}}/m\rfloor) &\coloneqq n_{\mathrm{A}}+n_{\mathrm{B}}-m.
\end{aligned}\right\} \in \mathcal{S}_{\mathrm{B}}. \tag{8.27}$$

One unfortunate issue is that n_{B} and n_{C} need not be multiples of m, and therefore it may be impossible to partition $\mathcal{I}_{\mathrm{B}}$ and $\mathcal{I}_{\mathrm{C}}$ into blocks of size m. Therefore, we divide them into $\lfloor n_{\mathrm{C}}/m\rfloor+1$ and $\lfloor n_{\mathrm{B}}/m\rfloor+1$ blocks of size m.

8.7 Types-H and -Q basic block matrices

Recall the definition of a basic matrix (Definition 6.5) and of s-bounded basic block matrices $F^{t_{\mathrm{F}}}_{a;b}(M)$ and $G^{t_{\mathrm{G}}}_{b;c}(M)$ (Definition 7.11) and $P^{t_{\mathrm{P}}}_{a;c}(M)$ (Definition 7.6) which are members of $\mathrm{Mat}_{\mathtt{F}}$, $\mathrm{Mat}_{\mathtt{G}}$, and $\mathrm{Mat}_{\mathtt{P}}$, respectively. We also define the s-bounded basic block matrices:

Definition 8.28. We define:

- $H^{t_{\mathrm{H}}}_{c;d}(M) \in \mathrm{Mat}_{\mathtt{H}}$ (for $c \in \mathcal{S}_{\mathrm{C}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{H}} \in [0\,..\,d-c-s]$, and $M \in \mathrm{BasicMat}^{\kappa}(s)$),
- $Q^{t_{\mathrm{Q}}}_{b;d}(M) \in \mathrm{Mat}_{\mathtt{Q}}$ (for $b \in \mathcal{S}_{\mathrm{B}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Q}} \in [0\,..\,d-b-s]$, and $M \in \mathrm{BasicMat}^{\kappa}(s)$),
- $Z^{t_{\mathrm{Z}}}_{a;d}(M) \in \mathrm{Mat}_{\mathtt{Z}}$ (for $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0\,..\,d-a-s]$, and $M \in \mathrm{BasicMat}^{\kappa}(s)$)

analogously to Definitions 7.6 and 7.11. $\Diamond$

The following useful corollary of Lemma 8.25 allows us to pass between type-P/H and type-F/Q basic block element commutators.

Lemma 8.29 (Block interchange)**.** *Let $s, s', s'' \in \mathbb{N}$, $a \in \mathcal{S}_{\mathrm{A}}$, $b \in \mathcal{S}_{\mathrm{B}}$, $c \in \mathcal{S}_{\mathrm{C}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{F}} \in [0\,..\,b-a-s]$, $t_{\mathrm{G}} \in [0\,..\,c-b-s']$, $t_{\mathrm{H}} \in [0\,..\,d-c-s'']$, $M \in \mathrm{BasicMat}^{\kappa}(s)$, $M' \in \mathrm{BasicMat}^{\kappa}(s')$, $M'' \in \mathrm{BasicMat}^{\kappa}(s'')$. Then:*

$$\vdash_0 \left[\left\langle P^{t_{\mathrm{F}}+t_{\mathrm{G}}}_{a;c}(MM')\right\rangle, \left\langle H^{t_{\mathrm{H}}}_{c;d}(M'')\right\rangle\right] \equiv \left[\left\langle F^{t_{\mathrm{F}}}_{a;b}(M)\right\rangle, \left\langle Q^{t_{\mathrm{G}}+t_{\mathrm{H}}}_{b;d}(M'M'')\right\rangle\right].$$

The notation in the lemma is justified as $MM' \in \mathrm{BasicMat}^{\kappa}(s+s')$, $t_{\mathrm{F}}+t_{\mathrm{H}} \le c-a-(s+s')$, $M'M'' \in \mathrm{BasicMat}^{\kappa}(s'+s'')$, and $t_{\mathrm{G}}+t_{\mathrm{H}} \le d-b-(s'+s'')$.

Proof. Follows from the matrix calculations

$$P^{t_F+t_G}_{a;c}(MM') = F^{t_F}_{a;b}(M) \cdot G^{t_G}_{b;c}(M'),$$
$$Q^{t_G+t_H}_{b;d}(M'M'') = G^{t_G}_{b;c}(M') \cdot H^{t_H}_{c;d}(M'')$$

and Lemma 8.25. □

8.8 Commuting type-P and -Q matrices

We next show that type-P and -Q elements commute, beginning with basic block elements and then using decompositions to lift to all elements.

Lemma 8.30 (Type-P and -Q basic block elements commute)**.** *Let $a \in \mathcal{S}_A$, $b \in \mathcal{S}_B$, $c \in \mathcal{S}_C$, $d \in \mathcal{S}_D$, $t_P \in [0\,..\,c-a-2m]$, $t_Q \in [0\,..\,d-b-2m]$, and $M, N \in \mathrm{BasicMat}^\kappa(2m)$. Then*

$$\vdash_0 \left[\left\langle P^{t_P}_{a;c}(M)\right\rangle, \left\langle Q^{t_Q}_{b;d}(N)\right\rangle\right] \equiv e.$$

Proof. As in the proof of Lemma 7.14, we can fix arbitrary $b' \in [n_A + m\,..\,n_A + n_B - m]$ with $\mathcal{I}^{b'}_B \cap \mathcal{I}^{b'}_B = \emptyset$ and $c' \in [n_A + n_B + m\,..\,n_A + n_B + n_C - m]$ with $\mathcal{I}^{c'}_C \cap \mathcal{I}^{c}_C = \emptyset$, and fix arbitrary $t_F, t_G \in \mathbb{N}$ such that

$$t_F \in [0\,..\,b' - a - 2m],$$
$$t_G \in [0\,..\,c - b'],$$
$$t_P = t_F + t_G,$$

and $t_H, t'_G \in \mathbb{N}$ such that

$$t_H \in [0\,..\,d - c' - 2m],$$
$$t'_G \in [0\,..\,c' - b],$$
$$t_Q = t'_G + t_H.$$

Now, define

$$F \coloneqq F^{t_F}_{a;b'}(M), G \coloneqq G^{t_G}_{b';c}(I_m) + G^{t'_G}_{b;c'}(I_m), \text{ and } H \coloneqq H^{t_H}_{c';d}(N),$$

so that

$$FG = P^{t_F+t_G}_{a;c}(M) = P^{t_P}_{a;c}(M) \text{ and } GH = Q^{t'_G+t_H}_{b;d}(N) = Q^{t_Q}_{b;d}(N).$$

and then we can apply Lemma 8.16. □

We have the following analogue of Claim 7.8:

Claim 8.31 (Decomposing type-P and -Q block matrices)**.** *Let $a \in \mathcal{S}_A$ and $c \in \mathcal{S}_C$. Suppose $P \in \mathrm{Mat}_{\mathtt{P}}$ is supported only on the $\mathcal{I}^a_A \times \mathcal{I}^c_C$ entries. Then P is a sum of $\lfloor (a-c)/m \rfloor = O(\rho)$ $2m$-bounded basic block matrices (i.e., matrices of the form $P^{t_P}_{a;c}(M)$ for some $t_P \in [0\,..\,c-a-2m]$, $M \in \mathrm{BasicMat}^\kappa(2m)$). Similarly, for $b \in \mathcal{S}_B$ and $d \in \mathcal{S}_D$ and $Q \in \mathrm{Mat}_{\mathtt{Q}}$ supported only on the $\mathcal{I}^b_B \times \mathcal{I}^d_D$ entries, Q is a sum of $\lfloor (d-b)/m \rfloor = O(\rho)$ $2m$-bounded basic block matrices.*

In turn, we get:

Claim 8.32 (Decomposing type-P and -Q matrices)**.** *Every $P \in \mathrm{Mat}_{\mathtt{P}}$ can be decomposed as*

$$P = \sum_{\alpha=0}^{n_A/m-1} \sum_{\gamma=0}^{\lfloor n_C/m \rfloor} P^{(\alpha,\gamma)}$$

where $P^{(\alpha,\gamma)} \in \mathrm{Mat}_{\mathtt{P}}$ is supported on $\mathcal{I}^{a(\alpha)}_A \times \mathcal{I}^{c(\gamma)}_C$. Similarly, every $Q \in \mathrm{Mat}_{\mathtt{Q}}$ can be decomposed as $Q = \sum_{\beta=0}^{\lfloor n_B/m \rfloor - 1} \sum_{\delta=0}^{n_D/m-1} Q^{(\beta,\delta)}$ where $Q^{(\beta,\delta)} \in \mathrm{Mat}_{\mathtt{Q}}$ is supported on $\mathcal{I}^{b(\beta)}_B \times \mathcal{I}^{d(\delta)}_D$.

This is identical the same as Corollary 7.10 except we have to account for the fact that the size of the "inner" wireset (n_{B} for type-Q, and n_{C} for type-P) may not being a multiple of m.

Lemma 8.33 (Type-P and -Q elements commute)**.** *For every $P \in \mathrm{Mat}_{\mathtt{P}}$ and $Q \in \mathrm{Mat}_{\mathtt{Q}}$,*

$$\vdash_6 [\langle P\rangle, \langle Q\rangle] \equiv e.$$

Proof. By Claims 8.31 and 8.32, P is a sum of at $O(\rho^3)$ $2m$-bounded basic block matrices and Q is a sum of at most $O(\rho^3)$ $2m$-bounded basic block matrices. Hence by in-subgroup relations, $\langle P\rangle$ is a product of $O(\rho^3)$ $2m$-bounded basic block elements and $\langle Q\rangle$ is a product of at most $O(\rho^3)$ $2m$-bounded basic block elements. Finally, by Lemma 8.30, each pair of $2m$-bounded basic block elements commutes with cost $O(1)$. □

8.9 Defining type-Z basic block aliases

Next, we define alias elements corresponding to type-Z basic blocks and prove that these elements can be constructed as commutators of appropriate pairs of type-P/H and type-F/Q basic blocks.

Recall the definition of packing (Equation (6.9)).

Definition 8.34 (Type-Z basic block aliases)**.** Fix $c^* := n_{\mathrm{A}} + n_{\mathrm{B}} + n_{\mathrm{C}} + m \in \mathcal{S}_{\mathrm{C}}$. For $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, and $t_{\mathrm{Z}} \in [0\,..\,d-a-4m]$, let $(t_{\mathrm{P}}, t_{\mathrm{H}}) := \mathrm{pack}(t_{\mathrm{Z}}; c^* - a - 4m, d - c^*)$. Then for $M \in \mathrm{BasicMat}^{\kappa}(4m)$, define

$$\zeta^{t_{\mathrm{Z}}}_{a;d}(M) := \left[\left\langle P^{t_{\mathrm{P}}}_{a;c^*}(M)\right\rangle, \left\langle H^{t_{\mathrm{H}}}_{c^*;d}(I_m)\right\rangle\right]. \qquad \Diamond$$

We have the following normalization lemma for packing:

Proposition 8.35. *For every $t_{\mathrm{P}}, t_{\mathrm{H}}, m_{\mathrm{F}}, m_{\mathrm{G}}, m_{\mathrm{H}} \in \mathbb{N}$, $t_{\mathrm{P}} \le m_{\mathrm{F}} + m_{\mathrm{G}}$, and $t_{\mathrm{H}} \le m_{\mathrm{H}}$, the following holds. Let $t^{(0)}_{\mathrm{P}} := t_{\mathrm{P}}$, $t^{(0)}_{\mathrm{H}} := t_{\mathrm{H}}$, $t_{\mathrm{Z}} := t_{\mathrm{P}} + t_{\mathrm{H}}$, $(t^*_{\mathrm{P}}, t^*_{\mathrm{H}}) := \mathrm{pack}(t_{\mathrm{Z}}; m_{\mathrm{F}} + m_{\mathrm{G}}, m_{\mathrm{H}})$, and $K := \lceil m_{\mathrm{G}}/m_{\mathrm{F}}\rceil$. For $k = 1, \ldots, K$, iterate the following:*

$$(t^{(k)}_{\mathrm{F}}, t^{(k)}_{\mathrm{G}}) := \mathrm{pack}\Big(t^{(k-1)}_{\mathrm{P}}; m_{\mathrm{F}}, m_{\mathrm{G}}\Big), \tag{8.36}$$

$$((t'_{\mathrm{G}})^{(k)}, t^{(k)}_{\mathrm{H}}) := \mathrm{pack}\Big(t^{(k)}_{\mathrm{G}} + t^{(k-1)}_{\mathrm{H}}; m_{\mathrm{G}}, m_{\mathrm{H}}\Big), \tag{8.37}$$

$$t^{(k)}_{\mathrm{P}} := t^{(k)}_{\mathrm{F}} + (t'_{\mathrm{G}})^{(k)}. \tag{8.38}$$

Then $(t^{(K)}_{\mathrm{P}}, t^{(K)}_{\mathrm{H}}) = (t^*_{\mathrm{P}}, t^*_{\mathrm{H}})$.

Proof. We have the invariants $t^{(k)}_{\mathrm{P}} + t^{(k)}_{\mathrm{H}} = t_{\mathrm{Z}}$ (because $t^{(k-1)}_{\mathrm{P}} + t^{(k-1)}_{\mathrm{H}} = t^{(k)}_{\mathrm{F}} + t^{(k)}_{\mathrm{G}} + t^{(k-1)}_{\mathrm{H}} = t^{(k)}_{\mathrm{F}} + (t'_{\mathrm{G}})^{(k)} + t^{(k)}_{\mathrm{H}} = t^{(k)}_{\mathrm{P}} + t^{(k)}_{\mathrm{H}}$ by Equations (8.36) to (8.38), respectively), $t^{(k)}_{\mathrm{P}} \le m_{\mathrm{F}} + m_{\mathrm{G}}$ (since $t^{(k)}_{\mathrm{F}} \le m_{\mathrm{F}}$ and $(t'_{\mathrm{G}})^{(k)} \le m_{\mathrm{G}}$ by Equations (8.36) and (8.37), respectively), and $t^{(k)}_{\mathrm{H}} \le m_{\mathrm{H}}$. By definition of $\mathrm{pack}(t_{\mathrm{Z}}; m_{\mathrm{F}} + m_{\mathrm{G}}, m_{\mathrm{H}})$, these inequalities guarantee that $t^{(k)}_{\mathrm{P}} \le t^*_{\mathrm{P}}$ and $t^{(k)}_{\mathrm{H}} \ge t^*_{\mathrm{H}}$.

Now, we claim that if $t^{(k-1)}_{\mathrm{P}} \ge m_{\mathrm{F}}$, then $t^{(k)}_{\mathrm{P}} = t^*_{\mathrm{P}}$. This is because $t^{(k)}_{\mathrm{F}} = m_{\mathrm{F}}$, $t^{(k)}_{\mathrm{G}} = t^{(k-1)}_{\mathrm{P}} - m_{\mathrm{F}}$, $t^{(k)}_{\mathrm{G}} + t^{(k-1)}_{\mathrm{H}} = t^{(k-1)}_{\mathrm{P}} - m_{\mathrm{F}} + t^{(k-1)}_{\mathrm{H}} = t_{\mathrm{Z}} - m_{\mathrm{F}}$, and so $(t'_{\mathrm{G}})^{(k)} = \min\{t_{\mathrm{Z}} - m_{\mathrm{F}}, m_{\mathrm{G}}\}$. In the former case, $t^{(k)}_{\mathrm{P}} = t^{(k)}_{\mathrm{F}} + (t'_{\mathrm{G}})^{(k)} = t_{\mathrm{Z}}$, while in the latter case, $t^{(k)}_{\mathrm{P}} = t^{(k)}_{\mathrm{F}} + (t'_{\mathrm{G}})^{(k)} = m_{\mathrm{F}} + m_{\mathrm{G}}$. In either case, since $t^{(k)}_{\mathrm{P}} \le t^*_{\mathrm{P}}$, we conclude $t^{(k)}_{\mathrm{P}} = t^*_{\mathrm{P}}$ by definition of t^*_{P}.

So, it suffices to check that at every step k, either $t^{(k)}_{\mathrm{P}} = t^*_{\mathrm{P}}$, $t^{(k)}_{\mathrm{P}} \ge m_{\mathrm{F}}$, or $t^{(k)}_{\mathrm{P}} - t^{(k-1)}_{\mathrm{P}} = m_{\mathrm{G}}$. Indeed, observe that

$$t^{(k)}_{\mathrm{P}} - t^{(k-1)}_{\mathrm{P}} = t^{(k-1)}_{\mathrm{H}} - t^{(k)}_{\mathrm{H}} = t^{(k-1)}_{\mathrm{H}} - ((t^{(k)}_{\mathrm{G}} + t^{(k-1)}_{\mathrm{H}}) - (t'_{\mathrm{G}})^{(k)}) = (t'_{\mathrm{G}})^{(k)} - t^{(k)}_{\mathrm{G}}.$$

Since $t^{(k-1)}_{\mathrm{P}} < m_{\mathrm{F}}$ by assumption, then $t^{(k)}_{\mathrm{F}} = t^{(k-1)}_{\mathrm{P}}$ and $t^{(k)}_{\mathrm{G}} = 0$, and so $(t'_{\mathrm{G}})^{(k)} = \min\{t^{(k-1)}_{\mathrm{H}}, m_{\mathrm{G}}\}$. Hence either $(t'_{\mathrm{G}})^{(k)} = t^{(k-1)}_{\mathrm{H}}$, in which case $t^{(k)}_{\mathrm{P}} = t^{(k)}_{\mathrm{F}} + (t'_{\mathrm{G}})^{(k)} = t^{(k-1)}_{\mathrm{P}} + t^{(k-1)}_{\mathrm{H}} = t_{\mathrm{Z}} = t^*_{\mathrm{P}}$, or $(t'_{\mathrm{G}})^{(k)} = m_{\mathrm{G}}$, in which case $t^{(k)}_{\mathrm{P}} - t^{(k-1)}_{\mathrm{P}} = m_{\mathrm{G}}$. □

Lemma 8.39 (Normalizing the type-`Z` exponent)**.** *Let c^* be as in Definition 8.34. For every $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{P}} \in [0\,..\,c^* - a - 4m]$, $t_{\mathrm{H}} \in [0\,..\,d - c^*]$, and $M \in \mathrm{BasicMat}^{\kappa}(4m)$,*

$$\vdash_1 \zeta_{a;d}^{t_{\mathrm{P}}+t_{\mathrm{H}}}(M) \equiv \Big[\big\langle P_{a;c^*}^{t_{\mathrm{P}}}(M)\big\rangle, \Big\langle H_{c^*;d}^{t_{\mathrm{H}}}(I_m)\Big\rangle\Big].$$

Proof. Pick $b \coloneqq n_{\mathrm{A}} + 3m$ and run the process described in Proposition 8.35 with $m_{\mathrm{F}} = b - a - 3m$, $m_{\mathrm{G}} = c^* - b - m$, and $m_{\mathrm{H}} = d - c^*$. We have iteratively:

$$\begin{aligned}
\Bigg[\Bigg\langle P_{a;c^*}^{t_{\mathrm{P}}^{(k-1)}}(M)\Bigg\rangle, \Bigg\langle H_{c^*;d}^{t_{\mathrm{H}}^{(k-1)}}(I)\Bigg\rangle\Bigg] &= \Bigg[\Bigg\langle P_{a;c^*}^{t_{\mathrm{F}}^{(k)}+t_{\mathrm{G}}^{(k)}}(M\cdot I_m)\Bigg\rangle, \Bigg\langle H_{c^*;d}^{t_{\mathrm{H}}^{(k-1)}}(I_m)\Bigg\rangle\Bigg] && \text{(Equation (8.36))}\\
&\equiv \Bigg[\Bigg\langle F_{a;b}^{t_{\mathrm{F}}^{(k)}}(M)\Bigg\rangle, \Bigg\langle Q_{b;d}^{t_{\mathrm{G}}^{(k)}+t_{\mathrm{H}}^{(k-1)}}(I_m\cdot I_m)\Bigg\rangle\Bigg] && \text{(Lemma 8.29)}\\
&= \Bigg[\Bigg\langle F_{a;b}^{t_{\mathrm{F}}^{(k)}}(M)\Bigg\rangle, \Bigg\langle Q_{b;d}^{(t'_{\mathrm{G}})^{(k)}+t_{\mathrm{H}}^{(k)}}(I_m\cdot I_m)\Bigg\rangle\Bigg] && \text{(Equation (8.37))}\\
&\equiv \Bigg[\Bigg\langle P_{a;c^*}^{t_{\mathrm{F}}^{(k)}+(t'_{\mathrm{G}})^{(k)}}(M\cdot I_m)\Bigg\rangle, \Bigg\langle H_{c^*;d}^{t_{\mathrm{H}}^{(k)}}(M)\Bigg\rangle\Bigg] && \text{(Lemma 8.29)}\\
&= \Bigg[\Bigg\langle P_{a;c^*}^{t_{\mathrm{P}}^{(k)}}(M)\Bigg\rangle, \Bigg\langle H_{c^*;d}^{t_{\mathrm{H}}^{(k)}}(I_m)\Bigg\rangle\Bigg], && \text{(Equation (8.38))}
\end{aligned}$$

and by the conclusion of Proposition 8.35,

$$\Bigg[\Bigg\langle P_{a;c^*}^{t_{\mathrm{P}}^{(K)}}(M)\Bigg\rangle, \Bigg\langle H_{c^*;d}^{t_{\mathrm{H}}^{(K)}}(I_m)\Bigg\rangle\Bigg] = \Big[\Big\langle P_{a;c^*}^{t_{\mathrm{P}}^*}(M)\Big\rangle, \Big\langle H_{c^*;d}^{t_{\mathrm{H}}^*}(I_m)\Big\rangle\Big] = \zeta_{a;c}^{t_{\mathrm{P}}^*+t_{\mathrm{H}}^*}(M),$$

as desired. Also, $m_{\mathrm{G}} \le c^* \le n$ while $m_{\mathrm{F}} \ge (n_{\mathrm{A}} + 3m) - (n_{\mathrm{A}} - m) - 3m = m$ and so $K \le \lceil n/m \rceil = \rho$. ☐

Lemma 8.40 (Commutators of compatible `P-H` and `F-Q` pairs are type-`Z` basic block aliases)**.** *For every $s, s' \in \mathbb{N}$ with $s + s' = 4m$, $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $M \in \mathrm{BasicMat}^{\kappa}(s)$ and $N \in \mathrm{BasicMat}^{\kappa}(s')$,*

- *for every $c \in \mathcal{S}_{\mathrm{C}}$, $t_{\mathrm{P}} \in [0\,..\,c - a - s]$, and $t_{\mathrm{H}} \in [0\,..\,d - c - s']$,*

$$\vdash_1 \Big[\big\langle P_{a;c}^{t_{\mathrm{P}}}(M)\big\rangle, \Big\langle H_{c;d}^{t_{\mathrm{H}}}(N)\Big\rangle\Big] \equiv \zeta_{a;d}^{t_{\mathrm{P}}+t_{\mathrm{H}}}(MN).$$

- *for every $b \in \mathcal{S}_{\mathrm{B}}$, $t_{\mathrm{F}} \in [0\,..\,b - a - s]$, and $t_{\mathrm{Q}} \in [0\,..\,d - b - s']$,*

$$\vdash_1 \Big[\Big\langle F_{a;b}^{t_{\mathrm{F}}}(M)\Big\rangle, \Big\langle Q_{b;d}^{t_{\mathrm{Q}}}(N)\Big\rangle\Big] \equiv \zeta_{a;d}^{t_{\mathrm{F}}+t_{\mathrm{Q}}}(MN).$$

Note that e.g. in the first bullet, if $c = n_{\mathrm{A}} + n_{\mathrm{B}} + n_{\mathrm{C}} - m$, $d = n_{\mathrm{A}} + n_{\mathrm{B}} + n_{\mathrm{C}}$, and $s' > m$, the interval $[0\,..\,d - c - s']$ that contains t_{H} is empty. See the next statement for a nonvacuous statement which we need in this "corner case".

Proof. Let c^* be as in Definition 8.34. We prove the items in sequence.

- Fix $b \coloneqq n_{\mathrm{A}}+n_{\mathrm{B}}+m$. Let $(t_{\mathrm{F}}, t_{\mathrm{G}}) \coloneqq \mathrm{pack}(t_{\mathrm{P}}; b - a - s, c - b)$, and let $(t'_{\mathrm{G}}, t'_{\mathrm{H}}) \coloneqq \mathrm{pack}(t_{\mathrm{G}} + t_{\mathrm{H}}; c^* - b - s', d - c^*)$, $t'_{\mathrm{P}} \coloneqq t_{\mathrm{F}} + t'_{\mathrm{G}}$. Applying Lemma 8.29 twice:

$$\begin{aligned}
&\Big[\big\langle P_{a;c}^{t_{\mathrm{P}}}(M)\big\rangle, \Big\langle H_{c;d}^{t_{\mathrm{H}}}(N)\Big\rangle\Big] = \Big[\big\langle P_{a;c}^{t_{\mathrm{F}}+t_{\mathrm{G}}}(M\cdot I_m)\big\rangle, \Big\langle H_{c;d}^{t_{\mathrm{H}}}(N)\Big\rangle\Big] \equiv \Big[\Big\langle F_{a;b}^{t_{\mathrm{F}}}(M)\Big\rangle, \Big\langle Q_{b;d}^{t_{\mathrm{G}}+t_{\mathrm{H}}}(I_m\cdot N)\Big\rangle\Big]\\
&= \Big[\Big\langle F_{a;b}^{t_{\mathrm{F}}}(M)\Big\rangle, \Big\langle Q_{b;d}^{t'_{\mathrm{G}}+t'_{\mathrm{H}}}(N\cdot I_m)\Big\rangle\Big] \equiv \Big[\Big\langle P_{a;c^*}^{t_{\mathrm{F}}+t'_{\mathrm{G}}}(MN)\Big\rangle, \Big\langle H_{c^*;d}^{t'_{\mathrm{H}}}(I_m)\Big\rangle\Big] = \Big[\Big\langle P_{a;c^*}^{t'_{\mathrm{P}}}(MN)\Big\rangle, \Big\langle H_{c^*;d}^{t'_{\mathrm{H}}}(I_m)\Big\rangle\Big].
\end{aligned}$$

  Finally, we apply the previous lemma (Lemma 8.39).

- Let $(g,h) \coloneqq \mathrm{pack}(t_{\mathrm{Q}}; c^* - b - 2m, d - c^*)$. Hence by Lemma 8.29 again:

$$\left[\left\langle F^{t_{\mathrm{F}}}_{a;b}(M)\right\rangle, \left\langle Q^{t_{\mathrm{Q}}}_{b;d}(N)\right\rangle\right] = \left[\left\langle F^{t_{\mathrm{F}}}_{a;b}(M)\right\rangle, \left\langle Q^{t_{\mathrm{G}}+t_{\mathrm{H}}}_{b;d}(N)\right\rangle\right]$$
$$\equiv \left[\left\langle P^{t_{\mathrm{F}}+t_{\mathrm{G}}}_{a;c^*}(MN)\right\rangle, \left\langle H^{t_{\mathrm{H}}}_{c^*;d}(I_m)\right\rangle\right] = \zeta^{t_{\mathrm{F}}+t_{\mathrm{Q}}}_{a;d}(MN)$$

  by the first item. □

Corollary 8.41 (Type-`Z` basic block aliases as commutators of type-`F` and -`Q` elements)**.** *For every $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0\,..\,d-a-4m]$, and $M \in \mathrm{BasicMat}^{\kappa}(4m)$, there exist $F \in \mathrm{Mat}_{\mathtt{F}}$ and $Q \in \mathrm{Mat}_{\mathtt{Q}}$ such that:*

$$\vdash_1 \zeta^{t_{\mathrm{Z}}}_{a;d}(M) \equiv [\langle F\rangle, \langle Q\rangle].$$

Proof. Fix $b \coloneqq n_{\mathrm{A}} + m$ and $(t_{\mathrm{F}}, t_{\mathrm{Q}}) \coloneqq \mathrm{pack}(t_{\mathrm{Z}}; b-a, d-b-4m)$ and use the previous lemma (Lemma 8.40) to get

$$\left[\left\langle F^{t_{\mathrm{F}}}_{a;b}(I_m)\right\rangle, \left\langle Q^{t_{\mathrm{Q}}}_{b;d}(M)\right\rangle\right] \equiv \zeta^{t_{\mathrm{Z}}}_{a;d}(M).$$

(Note that this is well-defined because $d - b - 4m \geq (n_{\mathrm{A}} + n_{\mathrm{B}} + n_{\mathrm{C}}) - (n_{\mathrm{A}} + m) - 4m = n_{\mathrm{B}} + n_{\mathrm{C}} - 5m \geq 3m \geq 0$.) □

Lemma 8.42 (Corner case)**.** *We have:*

- *For every $a \in \mathcal{S}_{\mathrm{A}}$, $M \in \mathrm{BasicMat}^{\kappa}(2m)$, $N \in \mathrm{BasicMat}^{\kappa}(m)$, $c \coloneqq n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}-m$, $d \coloneqq n_{\mathrm{A}}+n_{\mathrm{B}}+n_{\mathrm{C}}$, $t_{\mathrm{P}} \in [0\,..\,c-a-2m]$,*

$$\vdash_0 \left[\left\langle P^{t_{\mathrm{P}}}_{a;c}(M)\right\rangle, \left\langle H^{0}_{c;d}(N)\right\rangle\right] \equiv \zeta^{t_{\mathrm{P}}}_{a;d}(MN).$$

- *For every $d \in \mathcal{S}_{\mathrm{D}}$, $M \in \mathrm{BasicMat}^{\kappa}(m)$, $N \in \mathrm{BasicMat}^{\kappa}(2m)$, $a \coloneqq n_{\mathrm{A}} - m$, $b \coloneqq n_{\mathrm{A}}$, $t_{\mathrm{Q}} \in [0\,..\,d-b-2m]$,*

$$\vdash_0 \left[\left\langle F^{0}_{a;b}(M)\right\rangle, \left\langle Q^{t_{\mathrm{Q}}}_{b;d}(N)\right\rangle\right] \equiv \zeta^{t_{\mathrm{Q}}}_{a;d}(MN).$$

Proof. We prove the first equality as the second is dual. Let $b \coloneqq n_{\mathrm{A}}+m$. Let $(t_{\mathrm{F}}, t_{\mathrm{G}}) \coloneqq \mathrm{pack}(t_{\mathrm{P}}; b-a-2m, c-b)$. By Lemmas 8.29 and 8.40:

$$\left[\left\langle P^{t_{\mathrm{P}}}_{a;c}(M)\right\rangle, \left\langle H^{0}_{c;d}(N)\right\rangle\right] = \left[\left\langle P^{t_{\mathrm{F}}+t_{\mathrm{G}}}_{a;c}(M\cdot I_m)\right\rangle, \left\langle H^{0}_{c;d}(N)\right\rangle\right]$$
$$\equiv \left[\left\langle F^{t_{\mathrm{F}}}_{a;b}(M)\right\rangle, \left\langle Q^{t_{\mathrm{G}}}_{b;d}(I_m\cdot N)\right\rangle\right] = \zeta^{t_{\mathrm{F}}+t_{\mathrm{G}}}_{a;d}(MN). \quad \square$$

Lemma 8.43 (Incompatible `P`-`H` and `F`-`Q` pairs commute)**.** *For every $a \in \mathcal{S}_{\mathrm{A}}$ and $d \in \mathcal{S}_{\mathrm{D}}$, $M \in \mathrm{BasicMat}^{\kappa}(s)$ and $N \in \mathrm{BasicMat}^{\kappa}(s')$:*

- *for every $c, c' \in \mathcal{S}_{\mathrm{C}}$ such that $\mathcal{I}^{c}_{\mathrm{C}} \cap \mathcal{I}^{c'}_{\mathrm{C}} = \emptyset$, $t_{\mathrm{P}} \in [0\,..\,c-a-s]$, and $t_{\mathrm{H}} \in [0\,..\,d-c'-s']$,*

$$\vdash_0 \left[\left\langle P^{t_{\mathrm{P}}}_{a;c}(M)\right\rangle, \left\langle H^{t_{\mathrm{H}}}_{c';d}(N)\right\rangle\right] \equiv e.$$

- *for every $b, b' \in \mathcal{S}_{\mathrm{B}}$ such that $\mathcal{I}^{b'}_{\mathrm{B}} \cap \mathcal{I}^{b'}_{\mathrm{B}} = \emptyset$, $t_{\mathrm{F}} \in [0\,..\,b-a-s]$, and $t_{\mathrm{Q}} \in [0\,..\,d-b'-s']$,*

$$\vdash_0 \left[\left\langle F^{t_{\mathrm{F}}}_{a;b}(M)\right\rangle, \left\langle Q^{t_{\mathrm{Q}}}_{b';d}(N)\right\rangle\right] \equiv e.$$

Proof. For the first item, pick $b \coloneqq a+m$ and let $(t_{\mathrm{F}}, t_{\mathrm{G}}) \coloneqq \mathrm{pack}(t_{\mathrm{P}}; b-a-s, c-b)$. Then $P^{t_{\mathrm{P}}}_{a;c}(M) = F^{t_{\mathrm{F}}}_{a;b}(M) \cdot G^{t_{\mathrm{G}}}_{b;c}(I_m)$ and so

$$\left\langle P^{t_{\mathrm{P}}}_{a;c}(M)\right\rangle \equiv \left[\left\langle F^{t_{\mathrm{F}}}_{a;b}(M)\right\rangle, \left\langle G^{t_{\mathrm{G}}}_{b;c}(I_m)\right\rangle\right].$$

But $\left\langle H^{t_{\mathrm{H}}}_{c';d}(N)\right\rangle$ commutes with both these elements by constant-length in-subgroup relations (Lemma 8.14): The first since all type-`F` and -`H` elements commute, and the second because

$$\left[\left\langle G^{t_{\mathrm{G}}}_{b;c}(I_m)\right\rangle, \left\langle H^{t_{\mathrm{H}}}_{c';d}(N)\right\rangle\right] \equiv \left\langle G^{t_{\mathrm{G}}}_{b;c}(I_m)\cdot H^{t_{\mathrm{H}}}_{c';d}(N)\right\rangle \equiv \langle 0\rangle \equiv e$$

by the disjointness assumption. The second item is similar. □

8.10 Remaining relations for type-Z basic block aliases

We now prove all remaining relations which show that type-Z basic block aliases commute with other types of elements. These proofs essentially follow from [OS25] but we provide them for completeness.

Lemma 8.44 (Type-B elements commute with type-Z basic block aliases)**.** *For every $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0\,..\,d-a-4m]$, $M \in \mathrm{BasicMat}^{\kappa}(4m)$, and $B \in \mathrm{Mat}_{\mathtt{B}}$,*

$$\vdash_0 \left[\langle B\rangle, \zeta^{t_{\mathrm{Z}}}_{a;d}(M)\right] \equiv e.$$

Proof. By definition (Definition 8.34), $\zeta^{t_{\mathrm{Z}}}_{a;d}(M) = [\langle P\rangle, \langle H\rangle]$ for some $P \in \mathrm{Mat}_{\mathtt{P}}$ and $H \in \mathrm{Mat}_{\mathtt{H}}$. Both parts commute with $\langle B\rangle$ by constant-length in-subgroup relations (Lemma 8.14). □

Lemma 8.45 (Type-C elements commute with type-Z basic block aliases)**.** *For every $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0\,..\,d-a-4m]$, $M \in \mathrm{BasicMat}^{\kappa}(4m)$, and $C \in \mathrm{Mat}_{\mathtt{C}}$,*

$$\vdash_1 \left[\langle C\rangle, \zeta^{t_{\mathrm{Z}}}_{a;d}(M)\right] \equiv e.$$

Proof. By Corollary 8.41, $\vdash_1 \zeta^{t_{\mathrm{Z}}}_{a;d}(M) \equiv [\langle F\rangle, \langle Q\rangle]$ for some $F \in \mathrm{Mat}_{\mathtt{F}}$ and $Q \in \mathrm{Mat}_{\mathtt{Q}}$. Both parts commute with $\langle C\rangle$ by constant-length in-subgroup relations (Lemma 8.14). □

Lemma 8.46 (Type-F elements commute with type-Z basic block aliases)**.** *For every $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0\,..\,d-a-4m]$, $M \in \mathrm{BasicMat}^{\kappa}(4m)$, and $F \in \mathrm{Mat}_{\mathtt{F}}$,*

$$\vdash_0 \left[\langle F\rangle, \zeta^{t_{\mathrm{Z}}}_{a;d}(M)\right] \equiv e.$$

Proof. By definition (Definition 8.34), $\zeta^{t_{\mathrm{Z}}}_{a;d}(M) = [\langle P\rangle, \langle H\rangle]$ for some $P \in \mathrm{Mat}_{\mathtt{P}}$ and $H \in \mathrm{Mat}_{\mathtt{H}}$. Both parts commute with $\langle F\rangle$ by constant-length in-subgroup relations (Lemma 8.14). □

Lemma 8.47 (Type-H elements commute with type-Z basic block aliases)**.** *For every $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0\,..\,d-a-4m]$, $M \in \mathrm{BasicMat}^{\kappa}(4m)$, and $H \in \mathrm{Mat}_{\mathtt{H}}$,*

$$\vdash_1 \left[\langle H\rangle, \zeta^{t_{\mathrm{Z}}}_{a;d}(M)\right] \equiv e.$$

Proof. By Corollary 8.41, $\vdash_1 \zeta^{t_{\mathrm{Z}}}_{a;d}(M) \equiv [\langle F\rangle, \langle Q\rangle]$ for some $F \in \mathrm{Mat}_{\mathtt{F}}$ and $Q \in \mathrm{Mat}_{\mathtt{Q}}$. Both parts commute with $\langle H\rangle$ by constant-length in-subgroup relations (Lemma 8.14). □

Lemma 8.48 (Type-G elements commute with type-Z basic block aliases)**.** *For every $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0\,..\,d-a-4m]$, $M \in \mathrm{BasicMat}^{\kappa}(4m)$, and $G \in \mathrm{Mat}_{\mathtt{G}}$,*

$$\vdash_6 \left[\langle G\rangle, \zeta^{t_{\mathrm{Z}}}_{a;d}(M)\right] \equiv e.$$

Proof. By definition (Definition 8.34), $\zeta^{t_{\mathrm{Z}}}_{a;d}(M) = [\langle P\rangle, \langle H\rangle]$ for some $P \in \mathrm{Mat}_{\mathtt{P}}$ and $H \in \mathrm{Mat}_{\mathtt{H}}$. Hence using constant-length in-subgroup relations (Lemma 8.14):

$$\begin{aligned}
&\langle G\rangle\zeta^{t_{\mathrm{Z}}}_{a;d}(M) \\
&= \langle G\rangle\langle P\rangle\langle H\rangle\langle -P\rangle\langle -H\rangle && \text{(def. of } \mathtt{Z} \text{ basic block alias, Definition 8.34)} \\
&\equiv \langle P\rangle\langle G\rangle\langle H\rangle\langle -P\rangle\langle -H\rangle && (\mathtt{G} \text{ and } \mathtt{P} \text{ commute}) \\
&\equiv \langle P\rangle\langle H\rangle\langle GH\rangle\langle G\rangle\langle -P\rangle\langle -H\rangle && \text{(commutator of } \mathtt{G} \text{ and } \mathtt{H}) \\
&\equiv \langle P\rangle\langle H\rangle\langle GH\rangle\langle -P\rangle\langle G\rangle\langle -H\rangle && (\mathtt{G} \text{ and } \mathtt{P} \text{ commute}) \\
&\equiv \langle P\rangle\langle H\rangle\langle GH\rangle\langle -P\rangle\langle -GH\rangle\langle -H\rangle\langle G\rangle && \text{(commutator of } \mathtt{G} \text{ and } \mathtt{H}) \\
&\equiv \langle P\rangle\langle H\rangle\langle -P\rangle\langle -H\rangle\langle G\rangle && (\mathtt{P} \text{ and } \mathtt{Q} \text{ commute (Lemma 8.33) and linearity of } \mathtt{Q}) \\
&= \zeta^{t_{\mathrm{Z}}}_{a;d}(M)\langle G\rangle, && \text{(def. of } \mathtt{Z} \text{ basic block alias, Definition 8.34)}
\end{aligned}$$

as desired. □

Lemma 8.49 (Type-$\mathtt{P}$ elements commute with type-$\mathtt{Z}$ basic block aliases)**.** *For every $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0 \,..\, d-a-4m]$, $M \in \mathrm{BasicMat}^{\kappa}(4m)$, and $P \in \mathrm{Mat}_{\mathtt{P}}$,*

$$\vdash_1 \left[\langle P\rangle, \zeta^{t_{\mathrm{Z}}}_{a;d}(M)\right] \equiv e.$$

Proof. By Corollary 8.41, $\vdash_1 \zeta^{t_{\mathrm{Z}}}_{a;d}(M) \equiv [\langle F\rangle, \langle Q\rangle]$ for some $F \in \mathrm{Mat}_{\mathtt{F}}$ and $Q \in \mathrm{Mat}_{\mathtt{Q}}$. Both parts commute with $\langle P\rangle$, the first by in-subgroup relations (Lemma 8.14) and the second by Lemma 8.30. □

Lemma 8.50 (Type-$\mathtt{Q}$ elements commute with type-$\mathtt{Z}$ basic block aliases)**.** *For every $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0 \,..\, d-a-4m]$, $M \in \mathrm{BasicMat}^{\kappa}(4m)$, and $Q \in \mathrm{Mat}_{\mathtt{Q}}$,*

$$\vdash_0 \left[\langle Q\rangle, \zeta^{t_{\mathrm{Z}}}_{a;d}(M)\right] \equiv e.$$

Proof. By definition (Definition 8.34), $\zeta^{t_{\mathrm{Z}}}_{a;d}(M) = [\langle P\rangle, \langle H\rangle]$ for some $P \in \mathrm{Mat}_{\mathtt{P}}$ and $H \in \mathrm{Mat}_{\mathtt{H}}$. Both parts commute with $\langle Q\rangle$, the first by Lemma 8.30 and the second by in-subgroup relations (Lemma 8.14). □

Lemma 8.51 (Type-$\mathtt{Z}$ basic block aliases commute)**.** *For every $a, a' \in \mathcal{S}_{\mathrm{A}}$, $d, d' \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0 \,..\, d-a-4m]$, $t'_{\mathrm{Z}} \in [0 \,..\, d'-a'-2m]$, and $M, M' \in \mathrm{BasicMat}^{\kappa}(4m)$,*

$$\vdash_1 \left[\zeta^{t_{\mathrm{Z}}}_{a;d}(M), \zeta^{t'_{\mathrm{Z}}}_{a';d'}(M')\right] \equiv e.$$

Proof. By definition (Definition 8.34), $\zeta^{t_{\mathrm{Z}}}_{a;d}(M) \equiv [\langle P\rangle, \langle H\rangle]$ for some $P \in \mathrm{Mat}_{\mathtt{P}}$ and $H \in \mathrm{Mat}_{\mathtt{H}}$. Both parts commute with $\zeta^{t'_{\mathrm{Z}}}_{a';d'}(M')$ by Lemmas 8.47 and 8.49. □

Lemma 8.52 (Linearity for type-$\mathtt{Z}$ basic block aliases)**.** *For every $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0 \,..\, d-a-4m]$, and $M, M' \in \mathrm{BasicMat}^{\kappa}(4m)$:*

$$\vdash_1 \zeta^{t_{\mathrm{Z}}}_{a;d}(M) \cdot \zeta^{t_{\mathrm{Z}}}_{a;d}(M') \equiv \zeta^{t_{\mathrm{Z}}}_{a;d}(M+M').$$

Proof. Using (Definition 8.34), we have:

$$\begin{aligned}
&\zeta^{t_{\mathrm{Z}}}_{a;d}(M) \cdot \zeta^{t_{\mathrm{Z}}}_{a;d}(M') \\
&= \left\langle P^{t_{\mathrm{P}}}_{a;c}(M)\right\rangle\left\langle H^{t_{\mathrm{H}}}_{c;d}(I_m)\right\rangle\left\langle P^{t_{\mathrm{P}}}_{a;c}(-M)\right\rangle\left\langle H^{t_{\mathrm{H}}}_{c;d}(-I_m)\right\rangle\zeta^{t_{\mathrm{Z}}}_{a;d}(M') \\
&\hspace{15em}\text{(def. of type-}\mathtt{Z}\text{ basic block elt., Definition 8.34)} \\
&\equiv \left\langle P^{t_{\mathrm{P}}}_{a;c}(M)\right\rangle \cdot \zeta^{t_{\mathrm{Z}}}_{a;d}(M') \cdot \left\langle H^{t_{\mathrm{H}}}_{c;d}(I_m)\right\rangle\left\langle P^{t_{\mathrm{P}}}_{a;c}(-M)\right\rangle\left\langle H^{t_{\mathrm{H}}}_{c;d}(-I_m)\right\rangle \\
&\hspace{6em}\text{(type-}\mathtt{Z}\text{ basic block elts. commute with type-}\mathtt{P}\text{ and -}\mathtt{H}\text{ elts., Lemmas 8.47 and 8.49)} \\
&\equiv \left\langle P^{t_{\mathrm{P}}}_{a;c}(M)\right\rangle\left\langle P^{t_{\mathrm{P}}}_{a;c}(M')\right\rangle\left\langle H^{t_{\mathrm{H}}}_{c;d}(I_m)\right\rangle\left\langle P^{t_{\mathrm{P}}}_{a;c}(-M')\right\rangle\left\langle H^{t_{\mathrm{H}}}_{c;d}(-I_m)\right\rangle\left\langle H^{t_{\mathrm{H}}}_{c;d}(I_m)\right\rangle\left\langle P^{t_{\mathrm{P}}}_{a;c}(-M)\right\rangle\left\langle H^{t_{\mathrm{H}}}_{c;d}(-I_m)\right\rangle \\
&\hspace{15em}\text{(def. of type-}\mathtt{Z}\text{ basic block elt., Definition 8.34)} \\
&\equiv \left\langle P^{t_{\mathrm{P}}}_{a;c}(M+M')\right\rangle\left\langle H^{t_{\mathrm{H}}}_{c;d}(I_m)\right\rangle\left\langle P^{t_{\mathrm{P}}}_{a;c}(-(M+M'))\right\rangle\left\langle H^{t_{\mathrm{H}}}_{c;d}(-I_m)\right\rangle \\
&\hspace{15em}\text{(lin. of type-}\mathtt{P}\text{ basic block elts., Proposition 8.3)} \\
&= \zeta^{t_{\mathrm{Z}}}_{a;d}(M+M') \hspace{10em}\text{(def. of type-}\mathtt{Z}\text{ basic block elt., Definition 8.34)}
\end{aligned}$$

as desired. □

8.11 Defining type-Z aliases

Having defined type-$\mathtt{Z}$ basic block aliases and shown they commute with other relevant elements, we next turn to defining aliases corresponding to general type-$\mathtt{Z}$ matrices. We emphasize that starting from this subsection, we will begin considering words of non-constant (i.e., ρ-dependent) length.

Definition 8.53 (Decomposition of type-Z matrices)**.** For $Z \in \mathrm{Mat}_{\mathtt{Z}}$, for each $\alpha \in [0\,..\,n_{\mathrm{A}}/m)$, $\delta \in [0\,..\,n_{\mathrm{D}}/m)$, let $Z^{(\alpha,\delta)}$ denote Z restricted to the entries $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$, so that $Z = \sum_{\alpha=0}^{n_{\mathrm{A}}/m-1} \sum_{\delta=0}^{n_{\mathrm{D}}/m-1} Z^{(\alpha,\delta)}$.

For each α, δ, let $M_{i,j}^{(\alpha,\delta)} = Z_{a(\alpha)+i,d(\delta)+j}$ so that $Z^{(\alpha,\delta)} = Z^{0}_{a(\alpha);d(\delta)}(M^{(\alpha,\delta)})$. Letting $\Delta^{(\alpha,\delta)} \coloneqq d(\delta) - a(\alpha) - 4m$, $s = 2m$, and $\tau^{(\alpha,\delta)} \coloneqq \lfloor \Delta^{(\alpha,\delta)}/(2m) \rfloor$, Proposition 6.7 lets us decompose $M^{(\alpha,\delta)}$ as

$$M^{(\alpha,\delta)} = \sum_{k=0}^{\tau^{(\alpha,\delta)}-1} X^{\kappa\cdot 2km} M_{\mathrm{body}}^{(\alpha,\delta,k)} + X^{\kappa\cdot 2\tau^{(\alpha,\delta)}m} M_{\mathrm{gap}}^{(\alpha,\delta)} + X^{\kappa\Delta} M_{\mathrm{tail}}^{(\alpha,\delta)}$$

for $M_{\mathrm{body}}^{(\alpha,\delta,k)}, M_{\mathrm{gap}}^{(\alpha,\delta)}, M_{\mathrm{tail}}^{(\alpha,\delta)} \in \mathrm{BasicMat}^{\kappa}(4m)$. Correspondingly,

$$Z^{(\alpha,\delta)} = Z^{0}_{a(a(\alpha));d(d(\delta))}\left(M^{(\alpha,\delta)}\right) = \sum_{k=0}^{\tau^{(\alpha,\delta)}-1} Z^{k\cdot 2m}_{a(\alpha);d(\delta)}\left(M_{\mathrm{body}}^{(\alpha,\delta,k)}\right) + Z^{2\tau^{(\alpha,\delta)}m}_{a(\alpha);d(\delta)}\left(M_{\mathrm{gap}}^{(\alpha,\delta)}\right) + Z^{\Delta^{(\alpha,\delta)}}_{\alpha;\delta}\left(M_{\mathrm{tail}}^{(\alpha,\delta)}\right). \ \Diamond$$

Definition 8.54 (“Alias” word corresponding to type-Z matrix)**.** For every $Z \in \mathrm{Mat}_{\mathtt{Z}}$, we define:

$$\mathrm{alias}(Z) \coloneqq \prod_{\alpha=0}^{n_{\mathrm{A}}/m-1} \prod_{\delta=0}^{n_{\mathrm{D}}/m-1} \left(\left(\prod_{k=0}^{\tau^{(\alpha,\delta)}-1} \zeta^{km}_{a;d}\left(M_{\mathrm{body}}^{(\alpha,\delta,k)}\right) \right) \cdot \zeta^{\tau^{(\alpha,\delta)}-1}_{a;d}\left(M_{\mathrm{gap}}^{(\alpha,\delta)}\right) \cdot \zeta^{\Delta^{(\alpha,\delta)}}_{a;d}\left(M_{\mathrm{tail}}^{(\alpha,\delta)}\right) \right)$$

where $\Delta^{(\alpha,\delta)}, \tau^{(\alpha,\delta)}, M_{\mathrm{body}}^{(\alpha,\delta,k)}, M_{\mathrm{gap}}^{(\alpha,\delta)}, M_{\mathrm{tail}}^{(\alpha,\delta)}$ are as in Definition 8.53. This is a subgroup word over $\mathcal{H}$ of length $(n_{\mathrm{A}}/m - 1)(n_{\mathrm{D}}/m - 1)(\tau^{(\alpha,\delta)} - 1 + 2) \le O(\rho^3)$, since $a, d, \Delta^{(\alpha,\delta)} \le n$. $\Diamond$

Lemma 8.55 (Linearity for type-Z aliases)**.** *For every $Z, Z' \in \mathrm{Mat}_{\mathtt{Z}}$,*

$$\vdash_7 \mathrm{alias}(Z) \cdot \mathrm{alias}(Z') \equiv \mathrm{alias}(Z + Z').$$

Proof. Abbreviate $\Delta \coloneqq \Delta^{(\alpha,\delta)}$ and $\tau \coloneqq \tau^{(\alpha,\delta)}$. By construction, the decomposition in Definition 8.53 is linear, in the sense that:

$$\begin{aligned} M_{\mathrm{body}}^{(\alpha,\delta,k)} + N_{\mathrm{body}}^{(\alpha,\delta,k)} &= (M+N)_{\mathrm{body}}^{(\alpha,\delta,k)}, \\ M_{\mathrm{gap}}^{(\alpha,\delta)} + N_{\mathrm{gap}}^{(\alpha,\delta)} &= (M+N)_{\mathrm{gap}}^{(\alpha,\delta)}, \\ M_{\mathrm{tail}}^{(\alpha,\delta)} + N_{\mathrm{tail}}^{(\alpha,\delta)} &= (M+N)_{\mathrm{tail}}^{(\alpha,\delta)} \end{aligned}$$

(the first for every $k \in [0\,..\,\tau)$).

By the definition of type-Z aliases (Definition 8.54), type-Z basic block alias linearity (Lemma 8.52) and autocommutation (Lemma 8.51), we have

$$\begin{aligned} &\mathrm{alias}(Z) \cdot \mathrm{alias}(Z') \\ &= \prod_{\alpha=0}^{n_{\mathrm{A}}/m-1} \prod_{\delta=0}^{n_{\mathrm{D}}/m-1} \left(\left(\prod_{k=0}^{\tau-1} \zeta^{km}_{a;d}\left(M_{\mathrm{body}}^{(\alpha,\delta,k)}\right) \right) \cdot \zeta^{\tau}_{a;d}\left(M_{\mathrm{gap}}^{(\alpha,\delta)}\right) \cdot \zeta^{\Delta}_{a;d}\left(M_{\mathrm{tail}}^{(\alpha,\delta)}\right) \right) \\ &\qquad \cdot \prod_{\alpha=0}^{n_{\mathrm{A}}/m-1} \prod_{\delta=0}^{n_{\mathrm{D}}/m-1} \left(\left(\prod_{k=0}^{\tau-1} \zeta^{km}_{a;d}\left(N_{\mathrm{body}}^{(\alpha,\delta,k)}\right) \right) \cdot \zeta^{\tau}_{a;d}\left(N_{\mathrm{gap}}^{(\alpha,\delta)}\right) \cdot \zeta^{\Delta}_{a;d}\left(N_{\mathrm{tail}}^{(\alpha,\delta)}\right) \right) \\ &\equiv \prod_{\alpha=0}^{n_{\mathrm{A}}/m-1} \prod_{\delta=0}^{n_{\mathrm{D}}/m-1} \left(\left(\prod_{k=0}^{\tau-1} \zeta^{km}_{a;d}\left(M_{\mathrm{body}}^{(\alpha,\delta,k)} + N_{\mathrm{body}}^{(\alpha,\delta,k)}\right) \right) \cdot \zeta^{\tau}_{a;d}\left(M_{\mathrm{gap}}^{(\alpha,\delta)} + N_{\mathrm{gap}}^{(\alpha,\delta)}\right) \cdot \zeta^{\Delta}_{a;d}\left(M_{\mathrm{tail}}^{(\alpha,\delta)} + N_{\mathrm{tail}}^{(\alpha,\delta)}\right) \right) \\ &\equiv \left(\left(\prod_{k=0}^{\tau-1} \zeta^{km}_{a;d}\left((M+N)_{\mathrm{body}}^{(\alpha,\delta,k)}\right) \right) \cdot \zeta^{\tau}_{a;d}\left((M+N)_{\mathrm{gap}}^{(\alpha,\delta)}\right) \cdot \zeta^{\Delta}_{a;d}\left((M+N)_{\mathrm{tail}}^{(\alpha,\delta)}\right) \right) \\ &= \mathrm{alias}(Z + Z'). \end{aligned}$$

The number of steps is $O(\rho^7)$ since the subgroup words $\mathrm{alias}(Z)$ and $\mathrm{alias}(Z')$ each have length $O(\rho^3)$ and each application of Lemmas 8.51 and 8.52 uses $O(\rho)$ derivation steps. $\square$

Lemma 8.56. *For every $Z \in \mathrm{Mat}_{\mathtt{Z}}$, $\lambda \in \{\mathtt{B}, \mathtt{C}, \mathtt{F}, \mathtt{G}, \mathtt{H}, \mathtt{P}, \mathtt{Q}\}$, and $L \in \mathrm{Mat}_\lambda$,*

$$\vdash_9 [\langle L\rangle, \mathrm{alias}(Z)] \equiv e.$$

Proof. By Definition 8.54, alias(Z) is a product of $O(\rho^3)$ type-Z basic block aliases. Each basic block alias commutes with type-B, -C, -F, -G, -H, -P, and -Q elements by Lemmas 8.44 to 8.50, respectively. alias(Z) has length $O(\rho^3)$ and each application of these lemmas requires $O(\rho^6)$ derivation steps. ☐

8.12 Normalizing type-Z basic block aliases

In the previous section, we defined words corresponding to general type-Z matrices. We now show how to relate the type-Z basic block aliases we defined in Definition 8.54 to this new definition.

Lemma 8.57 (Shift)**.** *For every $s \in [0\,.\,.\,4m]$, $a \in \mathcal{S}_{\mathrm{A}}$, $d \in \mathcal{S}_{\mathrm{D}}$, $t_{\mathrm{Z}} \in [0\,.\,.\,d-a-4m-s]$, if $M \in \mathrm{BasicMat}^\kappa(4m-s)$ (equiv., $X^{\kappa s}M \in \mathrm{BasicMat}^\kappa(4m)$), then*

$$\vdash_1 \zeta_{a;d}^{t_{\mathrm{Z}}+s}(M) \equiv \zeta_{a;d}^{t_{\mathrm{Z}}}(X^{\kappa s}M).$$

Proof. By definition, for some fixed c^* and $(t_{\mathrm{P}}, t_{\mathrm{H}}) \coloneqq \mathrm{pack}(t_{\mathrm{Z}}; c^*-a-4m, d-c^*)$,

$$\zeta_{a;d}^{t_{\mathrm{Z}}}(X^{\kappa s}M) = \left[\left\langle P_{a;c^*}^{t_{\mathrm{P}}}(X^{\kappa s}M)\right\rangle, \left\langle H_{c^*;d}^{t_{\mathrm{H}}}(I_m)\right\rangle\right].$$

Note that we actually have $t_{\mathrm{H}} \in [0\,.\,.\,d-c^*-s]$ by assumption on t_{Z}. Now note that $\left\langle P_{a;c^*}^{t_{\mathrm{P}}}(X^{\kappa s}M)\right\rangle = \left\langle P_{a;c^*}^{t_{\mathrm{P}}+s}(M)\right\rangle$ and $M \in \mathrm{BasicMat}^\kappa(4m-s)$ while $I_m \in \mathrm{BasicMat}^\kappa(s)$. Hence by Lemma 8.40, we get:

$$\left[\left\langle P_{a;c^*}^{t_{\mathrm{P}}+s}(M)\right\rangle, \left\langle H_{c^*;d}^{t_{\mathrm{H}}}(I_m)\right\rangle\right] \equiv \zeta_{a;d}^{t_{\mathrm{Z}}+s}(M),$$

as desired. ☐

Lemma 8.58 (Normalization)**.** *Let $\alpha \in [0\,.\,.\,n_{\mathrm{A}}/m)$, $\delta \in [0\,.\,.\,n_{\mathrm{D}}/m)$, $t_{\mathrm{Z}} \in [0\,.\,.\,d(\delta)-a(\alpha)-4m]$, and $M \in \mathrm{BasicMat}^\kappa(4m)$. Then*

$$\vdash_2 \zeta_{a(\alpha);d(\delta)}^{t_{\mathrm{Z}}}(M) \equiv \mathrm{alias}\Big(Z_{a(\alpha);d(\delta)}^{t_{\mathrm{Z}}}(M)\Big).$$

Proof. We abbreviate $\Delta \coloneqq \Delta^{(\alpha,\delta)} = d(\delta)-a(\alpha)-4m$ and $\tau \coloneqq \tau^{(\alpha,\delta)} = \lfloor \Delta^{(\alpha,\delta)}/(2m)\rfloor$ as in Definition 8.53.

For $(i,j) \in [0\,.\,.\,m) \times [0\,.\,.\,m)$, let $M_{i,j} = \sum_{t=0}^{\kappa(4m+j-i)} X^t \mu_{i,j}^{(t)}$ be the degree decomposition of the (i,j) entry of M, and $\mu_{i,j}^{(t)} = 0$ for every $t < 0$ or $t > \kappa(4m+j-i)$. Correspondingly, define $N_{\mathrm{body}}^{(k)}, N_{\mathrm{gap}}, N_{\mathrm{tail}} \in \mathbb{F}[X]^{[0\,.\,.\,m)\times[0\,.\,.\,m)}$ via:

$$\begin{aligned}
(N_{\mathrm{body}}^{(k)})_{i,j} &\coloneqq \sum_{t=0}^{\kappa\cdot 2m-1} X^{t+\kappa(k\cdot 2m-t_{\mathrm{Z}})}\mu_{i,j}^{(t+\kappa(k\cdot 2m-t_{\mathrm{Z}}))},\\
(N_{\mathrm{gap}})_{i,j} &\coloneqq \sum_{t=0}^{\kappa(\Delta-\tau\cdot 2m)-1} X^{t+\kappa(\tau\cdot 2m-t_{\mathrm{Z}})}\mu_{i,j}^{(t+\kappa(\tau\cdot 2m-t_{\mathrm{Z}}))},\\
(N_{\mathrm{tail}})_{i,j} &\coloneqq \sum_{t=0}^{\kappa(4m+j-i)} X^{t+\kappa(\Delta-t_{\mathrm{Z}})}\mu_{i,j}^{(t+\kappa(\Delta-t_{\mathrm{Z}}))}.
\end{aligned}$$

Note that there may be negative powers of X here, but by definition of μ, only terms with degrees in $[0\,.\,.\,4m+j-i]$ have nonzero μ values. Indeed, this means that $N_{\mathrm{body}}^{(k)}, N_{\mathrm{gap}}, N_{\mathrm{tail}}$ are all in $\mathrm{BasicMat}^\kappa(4m)$.

Since $M = \sum_{k=0}^{\tau-1} N_{\mathrm{body}}^{(k)} + N_{\mathrm{gap}} + N_{\mathrm{tail}}$ (because the terms in the expansion exactly partition the terms in $M_{i,j}$), and since $\Delta \le n$, by block linearity (Lemma 8.52),

$$\vdash_2 \zeta_{a;d}^{t_{\mathrm{Z}}}(M) \equiv \prod_{k=0}^{\tau-1} \zeta_{a;d}^{t_{\mathrm{Z}}}\Big(N_{\mathrm{body}}^{(k)}\Big) \cdot \zeta_{a;d}^{t_{\mathrm{Z}}}(N_{\mathrm{gap}}) \cdot \zeta_{a;d}^{t_{\mathrm{Z}}}(N_{\mathrm{tail}}).$$

Recalling the definition of type-Z alias (Definition 8.54), it suffices to verify that

$$\vdash_1 \begin{cases} \zeta^{t_Z}_{a;d}\left(N^{(k)}_{\text{body}}\right) \equiv \zeta^{k\cdot 2m}_{a;d}\left(M^{(k)}_{\text{body}}\right), \\ \zeta^{t_Z}_{a;d}(N_{\text{gap}}) \equiv \zeta^{\tau\cdot 2m}_{a;d}(M_{\text{gap}}), \\ \zeta^{t_Z}_{a;d}(N_{\text{tail}}) \equiv \zeta^{\Delta}_{a;d}(M_{\text{tail}}). \end{cases}$$

Each of these follows by Lemma 8.57. For instance, for the first equality, we observe that if $k \cdot 2m - t_Z \geq 0$, then $N^{(k)}_{\text{body}} = X^{\kappa(k\cdot 2m - t_Z)} M^{(k)}_{\text{body}}$ and otherwise $M^{(k)}_{\text{body}} = X^{\kappa(t_Z - k\cdot 2m)} N^{(k)}_{\text{body}}$. In the former case, we use $s = k \cdot 2m - t_Z$ and have $\zeta^{t_Z+s}_{a;d}\left(M^{(k)}_{\text{body}}\right) \equiv \zeta^{t_Z}_{a;d}\left(N^{(k)}_{\text{body}}\right)$ (using that $t_Z \leq \Delta - s$ because $t_Z + s = k \cdot 2m \leq \Delta$). In the latter case, we use $s = t_Z - k\cdot 2m$ and have $\zeta^{k\cdot 2m+s}_{a;d}\left(N^{(k)}_{\text{body}}\right) \equiv \zeta^{k\cdot 2m}_{a;d}\left(M^{(k)}_{\text{body}}\right)$ (using that $k\cdot 2m \leq \Delta - s$ because $k \cdot 2m + s = t_Z \leq \Delta$). The proofs for the other two equalities are the same, with either $\tau \cdot 2m$ or Δ itself playing the role of $k \cdot 2m$ here. $\square$

8.13 Establishing type-Z elements

We now turn to proving the following lemma:

Lemma 8.59 (F-Q and P-H commutators)**.** *For every $P \in \mathrm{Mat}_{\mathtt{P}}, H \in \mathrm{Mat}_{\mathtt{H}}$, $\vdash_{15} [\langle P\rangle, \langle H\rangle] \equiv \mathrm{alias}(PH)$. Similarly, for every $F \in \mathrm{Mat}_{\mathtt{F}}, Q \in \mathrm{Mat}_{\mathtt{Q}}$, $\vdash_{15} [\langle F\rangle, \langle Q\rangle] \equiv \mathrm{alias}(FQ)$.*

We focus on proving the first equality, since the latter is dual. Specifically, all proofs in this subsection will work if we replace P with Q, F with H, γ with $\lfloor n_B/m \rfloor - \beta$, α with δ, and δ with $n_A/m - \alpha - 1$.

Claim 8.60 (Decomposing type-H block matrices)**.** *For $\gamma \in [0\,..\,\lfloor n_C/m \rfloor]$ and $\delta \in [0\,..\,n_D/m)$, and $H \in \mathrm{Mat}_{\mathtt{H}}$ supported only on the $\mathcal{I}^{c(\gamma)}_{C} \times \mathcal{I}^{d(\delta)}_{D}$ entries:*

- *If $\gamma = \lfloor n_C/m \rfloor$ and $\delta = 0$, then H is an m-bounded basic block matrix, i.e., a matrix of the form $H^{t_H}_{c(\gamma);d(\delta)}(M)$ for some $t_H \in [0\,..\,d(\delta) - c(\gamma) - m]$, $M \in \mathrm{BasicMat}^{\kappa}(m)$.*
- *Otherwise, H is a sum of $\lfloor (d(\delta) - c(\gamma))/m \rfloor = O(\rho)$ $2m$-bounded basic block matrices (i.e., matrices of the form $H^{t_H}_{c(\gamma);d(\delta)}(M)$ for some $t_H \in [0\,..\,d(\delta) - c(\gamma) - 2m]$, $M \in \mathrm{BasicMat}^{\kappa}(2m)$).*

Proof. Define $M \in \mathbb{F}[X]^{[0\,..\,m)\times[0\,..\,m)}$ via $M_{i,j} \coloneqq H_{c(\gamma)+i,d(\delta)+j}$. Then $\deg(M_{i,j}) = \deg(H_{i,j}) \leq \kappa((d(\delta) + j) - (c(\gamma) + i))$, and so M is $(d(\delta) - c(\gamma))$-bounded. Therefore $H = H^{0}_{c(\gamma);d(\delta)}(M)$.

Now we consider two cases. Firstly, if $d(\delta) - c(\gamma) \leq 2m$, which (by Equations (8.26) and (8.27)) happens iff $\gamma = \lfloor n_C/m \rfloor$ and $\delta = 0$, i.e., $d(\delta) - c(\gamma) = m$, then H is already an m-bounded basic block matrix. Otherwise, let $\Delta \coloneqq d(\delta) - c(\gamma) - 2m \geq 0$, $s \coloneqq m$, $\tau \coloneqq \lfloor \Delta/m \rfloor$, so that M is $(\Delta + 2s)$-bounded and by Proposition 6.7,

$$M = \sum_{k=0}^{\tau-1} X^{\kappa\cdot km} M^{(k)}_{\text{body}} + X^{\kappa\cdot\tau m} M_{\text{gap}} + X^{\kappa\Delta} M_{\text{tail}}$$

for $M^{(k)}_{\text{body}}, M_{\text{gap}}, M_{\text{tail}} \in \mathrm{BasicMat}^{\kappa}(2m)$. Correspondingly,

$$\begin{aligned} H = H^{0}_{c(\gamma);d(\delta)}(M) &= \sum_{k=0}^{\tau-1} H^{0}_{c(\gamma);d(\delta)}\left(X^{\kappa\cdot km} M^{(k)}_{\text{body}}\right) + H^{0}_{c(\gamma);d(\delta)}(X^{\kappa\cdot\tau m} M_{\text{gap}}) + H^{0}_{c(\gamma);d(\delta)}\left(X^{\kappa\Delta} M_{\text{tail}}\right) \\ &= \sum_{k=0}^{\tau-1} H^{km}_{c(\gamma);d(\delta)}\left(M^{(k)}_{\text{body}}\right) + H^{\tau m}_{c(\gamma);d(\delta)}(M_{\text{gap}}) + H^{\Delta}_{c(\gamma);d(\delta)}(M_{\text{tail}}), \end{aligned}$$

as desired. $\square$

Claim 8.61. *Let $\alpha \in [0\,..\,n_{\mathrm{A}}/m)$, $\gamma, \gamma' \in [0\,..\,\lfloor n_{\mathrm{C}}/m \rfloor]$, and $\delta \in [0\,..\,n_{\mathrm{D}}/m)$. Suppose $P \in \mathrm{Mat}_{\mathtt{P}}$ is supported only on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times \mathcal{I}_{\mathrm{C}}^{c(\gamma)}$ and $H \in \mathrm{Mat}_{\mathtt{H}}$ is supported only on $\mathcal{I}_{\mathrm{C}}^{c(\gamma')} \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$.*

Let $\rho_{\mathrm{P}} := \lfloor (c(\gamma) - a(\alpha))/m \rfloor + 2$ and $\rho_{\mathrm{H}} := \lfloor (d(\delta) - c(\gamma'))/m \rfloor + 2 \le O(\rho)$. Then P is a sum of ρ_{P} $2m$-bounded basic block matrices, $P = \sum_{k=1}^{\rho_{\mathrm{P}}} P_{a(\alpha);c(\gamma)}^{t_{\mathrm{P}}^{(k)}}\big(M^{(k)}\big)$ where $M^{(k)} \in \mathrm{BasicMat}^{\kappa}(2m)$ and $t_{\mathrm{P}}^{(k)} \in [0\,..\,c(\gamma) - a(\alpha) - 2m]$. Correspondingly

$$\vdash_1 \langle P \rangle = \left\langle \sum_{k=1}^{\rho_{\mathrm{P}}} P_{a(\alpha);c(\gamma)}^{t_{\mathrm{P}}^{(k)}}\Big(M^{(k)}\Big) \right\rangle \equiv \prod_{k=1}^{\rho_{\mathrm{P}}} \left\langle P_{a(\alpha);c(\gamma)}^{t_{\mathrm{P}}^{(k)}}\Big(M^{(k)}\Big) \right\rangle. \tag{8.62}$$

H is either an m-bounded basic block matrix $H^0_{c(\lfloor n_{\mathrm{C}}/m \rfloor);d(0)}(N)$ (for some $N \in \mathrm{BasicMat}^{\kappa}(m)$) if $\gamma' = \lfloor n_{\mathrm{C}}/m \rfloor$ and $\delta = 0$, and otherwise a sum of ρ_{H} $2m$-bounded basic block matrices, $H = \sum_{\ell=1}^{\rho_{\mathrm{H}}} P_{c(\gamma');d(\delta)}^{t_{\mathrm{H}}^{(\ell)}}\big(N^{(\ell)}\big)$ where $N^{(\ell)} \in \mathrm{BasicMat}^{\kappa}(2m)$ and $t_{\mathrm{H}}^{(\ell)} \in [0\,..\,d(\delta) - c(\gamma') - 2m]$. Correspondingly, either

$$\langle H \rangle = \left\langle H^0_{c(\lfloor n_{\mathrm{C}}/m \rfloor);d(0)}(N) \right\rangle \qquad \text{if } \gamma' = \lfloor n_{\mathrm{C}}/m \rfloor \wedge \delta = 0, \tag{8.63a}$$

$$\vdash_1 \langle H \rangle = \left\langle \sum_{\ell=1}^{\rho_{\mathrm{H}}} H_{c(\gamma');d(\delta)}^{t_{\mathrm{H}}^{(\ell)}}\Big(N^{(\ell)}\Big) \right\rangle \equiv \prod_{\ell=1}^{\rho_{\mathrm{H}}} \left\langle H_{c(\gamma');d(\delta)}^{t_{\mathrm{H}}^{(\ell)}}\Big(N^{(\ell)}\Big) \right\rangle \qquad \text{otherwise.} \tag{8.63b}$$

Proof. Follows immediately from Claims 8.31 and 8.60 and Lemma 8.14. ☐

Claim 8.64. *Let $\alpha \in [0\,..\,n_{\mathrm{A}}/m)$, $\gamma, \gamma' \in [0\,..\,\lfloor n_{\mathrm{C}}/m \rfloor]$, and $\delta \in [0\,..\,n_{\mathrm{D}}/m)$. Suppose $P \in \mathrm{Mat}_{\mathtt{P}}$ is supported only on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times \mathcal{I}_{\mathrm{C}}^{c(\gamma)}$ and $H \in \mathrm{Mat}_{\mathtt{H}}$ is supported only on $\mathcal{I}_{\mathrm{C}}^{c(\gamma')} \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$. Further, suppose that $\mathcal{I}_{\mathrm{C}}^{c(\gamma)} \cap \mathcal{I}_{\mathrm{C}}^{c(\gamma')} = \emptyset$. Then*

$$\vdash_2 [\langle P \rangle, \langle H \rangle] \equiv e.$$

Proof. We apply Claim 8.61. By Lemma 8.43, each element in the RHS of Equation (8.62) commutes with each element the RHS of Equation (8.63a) or Equation (8.63b). Each invocation of Lemma 8.43 requires $O(1)$ derivation steps and we do so $\rho_{\mathrm{P}} \cdot \rho_{\mathrm{H}} \le O(\rho^2)$ times. ☐

Claim 8.65. *Let $\alpha \in [0\,..\,n_{\mathrm{A}}/m)$, $\gamma := \lfloor n_{\mathrm{C}}/m \rfloor \in [0\,..\,\lfloor n_{\mathrm{C}}/m \rfloor]$, and $\delta \in [0\,..\,n_{\mathrm{D}}/m)$. Suppose $P \in \mathrm{Mat}_{\mathtt{P}}$ is supported only on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times \mathcal{I}_{\mathrm{C}}^{c(\gamma)}$ and $H \in \mathrm{Mat}_{\mathtt{H}}$ is supported only on $\mathcal{I}_{\mathrm{C}}^{c(\gamma)} \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$. Then*

$$\vdash_8 [\langle P \rangle, \langle H \rangle] \equiv \mathrm{alias}(PH).$$

Proof. We have, by Equations (8.62) and (8.63a):

$$[\langle P \rangle, \langle H \rangle] = \left[\prod_{k=1}^{\rho_{\mathrm{P}}} \underbrace{\left\langle P_{a(\alpha);c(\lfloor n_{\mathrm{C}}/m \rfloor)}^{t_{\mathrm{P}}^{(k)}}\Big(M^{(k)}\Big) \right\rangle}_{=:u_k}, \underbrace{\left\langle H^{t_{\mathrm{H}}}_{c(\lfloor n_{\mathrm{C}}/m \rfloor);d(0)}(N) \right\rangle}_{=:v} \right].$$

We apply Corollary 6.14. The pairwise commutators are

$$[u_k, v] = \left[\left\langle P_{a(\alpha);c(\lfloor n_{\mathrm{C}}/m \rfloor)}^{t_{\mathrm{P}}^{(k)}}\Big(M^{(k)}\Big) \right\rangle, \left\langle H^{t_{\mathrm{H}}}_{c(\lfloor n_{\mathrm{C}}/m \rfloor);d(0)}(N) \right\rangle \right] \equiv \zeta_{a(\alpha);d(0)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}}\Big(M^{(k)}N\Big)$$

via Lemma 8.42 (each application costs $O(1)$). The triple commutators are

$$\begin{aligned}[u_{k'}, [u_k, v]] &= \left[\left\langle P_{a(\alpha);c(\lfloor n_{\mathrm{C}}/m \rfloor)}^{t_{\mathrm{P}}^{(k)}}\Big(M^{(k)}\Big) \right\rangle, \left[\left\langle P_{a(\alpha);c(\lfloor n_{\mathrm{C}}/m \rfloor)}^{t_{\mathrm{P}}^{(k)}}\Big(M^{(k)}\Big) \right\rangle, \left\langle H^{t_{\mathrm{H}}}_{c(\lfloor n_{\mathrm{C}}/m \rfloor);d(0)}(N) \right\rangle \right] \right] \\ &\equiv \left[\left\langle P_{a(\alpha);c(\lfloor n_{\mathrm{C}}/m \rfloor)}^{t_{\mathrm{P}}^{(k)}}\Big(M^{(k)}\Big) \right\rangle, \zeta_{a(\alpha);d(0)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}}\Big(M^{(k)}N\Big) \right] \equiv e\end{aligned}$$

by Lemma 8.49 (each application costs $O(1)$). All factors have size $O(1)$. Thus,

$$[\langle P\rangle,\langle H\rangle] \equiv \prod_{k=\rho_{\mathrm{P}}}^{1} \zeta_{a(\alpha);d(0)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}}\left(M^{(k)}N\right) \equiv \prod_{k=\rho_{\mathrm{P}}}^{1} \operatorname{alias}\left(Z_{a(\alpha);d(0)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}}\left(M^{(k)}N\right)\right)$$

$$\equiv \operatorname{alias}\left(\sum_{k=\rho_{\mathrm{P}}}^{1} Z_{a(\alpha);d(0)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}}\left(M^{(k)}N\right)\right) = \operatorname{alias}(Z),$$

where the equivalences use Corollary 6.14 and Lemmas 8.55 and 8.58. The last equation follows because

$$Z_{a(\alpha);d(0)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}}\left(M^{(k)}N\right) = P_{a(\alpha);c(\lfloor n_{\mathrm{C}}/m\rfloor)}^{t_{\mathrm{P}}^{(k)}}\left(M^{(k)}\right)\cdot H_{c(\lfloor n_{\mathrm{C}}/m\rfloor);d(0)}^{t_{\mathrm{H}}}(N)$$

and therefore

$$\sum_{k=1}^{\rho_{\mathrm{P}}} Z_{a(\alpha);d(0)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}}\left(M^{(k)}N\right) = \left(\sum_{k=1}^{\rho_{\mathrm{P}}} P_{a(\alpha);c(\lfloor n_{\mathrm{C}}/m\rfloor)}^{t_{\mathrm{P}}^{(k)}}\left(M^{(k)}\right)\right)\cdot H_{c(\lfloor n_{\mathrm{C}}/m\rfloor);d(0)}^{t_{\mathrm{H}}}(N) = PH.$$

Each invocation of Lemmas 8.42, 8.47 and 8.49 takes $O(1)$ derivation steps. Lemma 8.58 takes $O(\rho^2)$, and Lemma 8.55 takes $O(\rho^7)$. The invocation of Lemma 8.55 $O(\rho)$ times dominates. $\square$

Claim 8.66. *Let $\alpha \in [0\,..\,n_{\mathrm{A}}/m)$, $\gamma \in [0\,..\,\lfloor n_{\mathrm{C}}/m\rfloor)$, and $\delta \in [0\,..\,n_{\mathrm{D}}/m)$. Suppose $P \in \mathrm{Mat}_{\mathtt{P}}$ is supported only on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times \mathcal{I}_{\mathrm{C}}^{c(\gamma)}$ and $H \in \mathrm{Mat}_{\mathtt{H}}$ is supported only on $\mathcal{I}_{\mathrm{C}}^{c(\gamma)} \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$. Then*

$$\vdash_9 [\langle P\rangle,\langle H\rangle] \equiv \operatorname{alias}(PH).$$

Proof. By Equations (8.62) and (8.63b):

$$[\langle P\rangle,\langle H\rangle] = \left[\prod_{k=1}^{\rho_{\mathrm{P}}} \underbrace{\left\langle P_{a(\alpha);c(\gamma)}^{t_{\mathrm{P}}^{(k)}}\left(M^{(k)}\right)\right\rangle}_{=:u_k}, \prod_{\ell=1}^{\rho_{\mathrm{H}}} \underbrace{\left\langle H_{c(\gamma);d(\delta)}^{t_{\mathrm{H}}^{(\ell)}}\left(N^{(\ell)}\right)\right\rangle}_{=:v_\ell}\right].$$

We again apply Corollary 6.14. The pairwise commutators are

$$[u_k, v_\ell] = \left[\left\langle P_{a(\alpha);c(\gamma)}^{t_{\mathrm{P}}^{(k)}}\left(M^{(k)}\right)\right\rangle, \left\langle H_{c(\gamma);d(\delta)}^{t_{\mathrm{H}}^{(\ell)}}\left(N^{(\ell)}\right)\right\rangle\right] \equiv \zeta_{a(\alpha);d(\delta)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}^{(\ell)}}\left(M^{(k)}N^{(\ell)}\right)$$

via Lemma 8.42 (cost is $O(1)$). The triple commutators are

$$[u_{k'},[u_k,v_\ell]] = \left[\left\langle P_{a(\alpha);c(\gamma)}^{t_{\mathrm{P}}^{(k')}}\left(M^{(k')}\right)\right\rangle, \left[\left\langle P_{a(\alpha);c(\gamma)}^{t_{\mathrm{P}}^{(k)}}\left(M^{(k)}\right)\right\rangle, \left\langle H_{c(\gamma);d(\delta)}^{t_{\mathrm{H}}^{(\ell)}}\left(N^{(\ell)}\right)\right\rangle\right]\right]$$

$$\equiv \left[\left\langle P_{a(\alpha);c(\gamma)}^{t_{\mathrm{P}}^{(k')}}\left(M^{(k')}\right)\right\rangle, \zeta_{a(\alpha);d(\delta)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}^{(\ell)}}\left(M^{(k)}N^{(\ell)}\right)\right] \equiv e$$

by Lemma 8.49, and similarly for $[v_{\ell'},[u_k,v_\ell]]$ by Lemma 8.47 (cost is $O(1)$). Thus,

$$[\langle P\rangle,\langle H\rangle] \equiv \prod_{k=\rho_{\mathrm{P}}}^{1}\prod_{\ell=1}^{\rho_{\mathrm{H}}} \zeta_{a(\alpha);d(\delta)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}^{(\ell)}}\left(M^{(k)}N^{(\ell)}\right) \equiv \prod_{k=\rho_{\mathrm{P}}}^{1}\prod_{\ell=1}^{\rho_{\mathrm{H}}} \operatorname{alias}\left(Z_{a(\alpha);d(\delta)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}^{(\ell)}}\left(M^{(k)}N^{(\ell)}\right)\right)$$

$$\equiv \operatorname{alias}\left(\sum_{k=1}^{\rho_{\mathrm{P}}}\sum_{\ell=1}^{\rho_{\mathrm{H}}} Z_{a(\alpha);d(\delta)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}^{(\ell)}}\left(M^{(k)}N^{(\ell)}\right)\right) = \operatorname{alias}(Z),$$

where the equivalences use Corollary 6.14 and Lemmas 8.55 and 8.58. The last equation follows because

$$Z_{a(\alpha);d(\delta)}^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}^{(\ell)}}\left(M^{(k)}N^{(\ell)}\right) = P_{a(\alpha);c(\gamma)}^{t_{\mathrm{P}}^{(k)}}\left(M^{(k)}\right)\cdot H_{c(\gamma);d(\delta)}^{t_{\mathrm{H}}^{(\ell)}}\left(N^{(\ell)}\right)$$

and therefore

$$\sum_{k=1}^{\rho_{\mathrm{P}}}\sum_{\ell=1}^{\rho_{\mathrm{H}}} Z^{t_{\mathrm{P}}^{(k)}+t_{\mathrm{H}}^{(\ell)}}_{a(\alpha);d(\delta)}\left(M^{(k)}N^{(\ell)}\right) = \left(\sum_{k=1}^{\rho_{\mathrm{P}}} P^{t_{\mathrm{P}}^{(k)}}_{a(\alpha);c(\gamma)}\left(M^{(k)}\right)\right)\cdot\left(\sum_{\ell=1}^{\rho_{\mathrm{H}}} H^{t_{\mathrm{H}}^{(\ell)}}_{c(\gamma);d(\delta)}\left(N^{(\ell)}\right)\right) = PH.$$

Each invocation of Lemmas 8.42, 8.47 and 8.49 takes $O(1)$ derivation steps, Lemma 8.58 takes $O(\rho^2)$, and Lemma 8.55 takes $O(\rho^7)$. The invocation of Lemma 8.55 $O(\rho^2)$ times dominates. ☐

Claim 8.67. *Let $\alpha \in [0\,..\,n_{\mathrm{A}}/m)$, $\gamma, \gamma'\lfloor n_{\mathrm{C}}/m\rfloor \in [0\,..\,\lfloor n_{\mathrm{C}}/m\rfloor]$, and $\delta \in [0\,..\,n_{\mathrm{D}}/m)$. Suppose $P \in \mathrm{Mat}_{\mathtt{P}}$ is supported only on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times \mathcal{I}_{\mathrm{C}}^{c(\gamma)}$ and $H \in \mathrm{Mat}_{\mathtt{H}}$ is supported only on $\mathcal{I}_{\mathrm{C}}^{c(\gamma')} \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$. Then*

$$\vdash_9 [\langle P\rangle, \langle H\rangle] \equiv \mathrm{alias}(PH).$$

Proof. We consider cases based on $\alpha, \gamma, \gamma', \delta$.

Case 1: $\mathcal{I}_{\mathrm{C}}^{c(\gamma)} \cap \mathcal{I}_{\mathrm{C}}^{c(\gamma')} = \emptyset$. We immediately apply Claim 8.64.

Case 2: $\gamma = \gamma'$. We split into two further subcases.

Case 2a: $\gamma = \lfloor n_{\mathrm{C}}/m\rfloor$ and $\delta = 0$. We apply Claim 8.65.

Case 2b: $\gamma < \lfloor n_{\mathrm{C}}/m\rfloor$ or $\delta > 0$. We apply Claim 8.66.

Case 3 (leftover). Since we fell through the first two cases, $\mathcal{I}_{\mathrm{C}}^{c(\gamma)}$ and $\mathcal{I}_{\mathrm{C}}^{c(\gamma')}$ intersect but $\gamma \neq \gamma'$. The only way that this can happen, by construction of the intervals $\mathcal{I}_{\mathrm{C}}^{c(\cdot)}$ (Equation (8.27)), is that $\gamma = 0$ and $\gamma' = 1$ or $\gamma = 1$ and $\gamma' = 0$. Hence, P is supported on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times (\mathcal{I}_{\mathrm{C}}^{c(0)} \cup \mathcal{I}_{\mathrm{C}}^{c(1)})$ and H on $(\mathcal{I}_{\mathrm{C}}^{c(0)} \cup \mathcal{I}_{\mathrm{C}}^{c(1)}) \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$. Unfortunately, the intervals $\mathcal{I}_{\mathrm{C}}^{c(0)}$ and $\mathcal{I}_{\mathrm{C}}^{c(1)}$ overlap: Since $c(0) = n_{\mathrm{A}} + n_{\mathrm{B}}$ and $c(1) = n_{\mathrm{A}} + n_{\mathrm{B}} + (n_{\mathrm{C}} - (\lfloor n_{\mathrm{C}}/m\rfloor)m)$,

$$\begin{aligned}
\mathcal{I}_{\mathrm{C}}^{c(0)} &= [n_{\mathrm{A}} + n_{\mathrm{B}}\,..\,n_{\mathrm{A}} + n_{\mathrm{B}} + m),\\
\mathcal{I}_{\mathrm{C}}^{c(1)} &= [n_{\mathrm{A}} + n_{\mathrm{B}} + (n_{\mathrm{C}} - \lfloor n_{\mathrm{C}}/m\rfloor m)\,..\,n_{\mathrm{A}} + n_{\mathrm{B}} + (n_{\mathrm{C}} - \lfloor n_{\mathrm{C}}/m\rfloor m) + m),\\
\mathcal{I}_{\mathrm{C}}^{c(0)} \cap \mathcal{I}_{\mathrm{C}}^{c(1)} &= [n_{\mathrm{A}} + n_{\mathrm{B}} + (n_{\mathrm{C}} - \lfloor n_{\mathrm{C}}/m\rfloor m)\,..\,n_{\mathrm{A}} + n_{\mathrm{B}} + m).
\end{aligned}$$

(Note that $0 \le n_{\mathrm{C}} - \lfloor n_{\mathrm{C}}/m\rfloor m < m$.) Define $c' \coloneqq n_{\mathrm{A}} + n_{\mathrm{B}} + m$, so that

$$\mathcal{I}_{\mathrm{C}}^{c'} = [n_{\mathrm{A}} + n_{\mathrm{B}} + m\,..\,n_{\mathrm{A}} + n_{\mathrm{B}} + 2m), \quad \mathcal{I}_{\mathrm{C}}^{c(0)} \cap \mathcal{I}_{\mathrm{C}}^{c'} = \emptyset, \quad \text{and} \quad \mathcal{I}_{\mathrm{C}}^{c(1)} \subset \mathcal{I}_{\mathrm{C}}^{c(0)} \cup \mathcal{I}_{\mathrm{C}}^{c'}.$$

In this case, decompose P as a sum $P_0 + P'$, where P_0 is supported on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times \mathcal{I}_{\mathrm{C}}^{c(0)}$ and P' is supported on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times (\mathcal{I}_{\mathrm{C}}^{c'} \cap \mathcal{I}_{\mathrm{C}}^{c(1)})$. Similarly, write H as a sum $H_0 + H'$, where H_0 is supported on $\mathcal{I}_{\mathrm{C}}^{c(0)} \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$ and H' is supported on $(\mathcal{I}_{\mathrm{C}}^{c'} \cap \mathcal{I}_{\mathrm{C}}^{c(1)}) \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$. By disjointness, we have the matrix calculation $PH = P_0H_0 + P'H'$.

With an eye towards applying Corollary 6.14, Claim 8.64 gives, using that H' is supported on $\mathcal{I}_{\mathrm{C}}^{c(1)} \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$, and P' on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times \mathcal{I}_{\mathrm{C}}^{c(1)}$:

$$[\langle P_0\rangle, \langle H'\rangle] \equiv [\langle P'\rangle, \langle H_0\rangle] \equiv \mathbb{1}.$$

Similarly, Claim 8.66 gives, using that H' is supported on $\mathcal{I}_{\mathrm{C}}^{c'} \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$, and P' on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times \mathcal{I}_{\mathrm{C}}^{c'}$, we calculate the double commutators:

$$[\langle P_0\rangle, \langle H_0\rangle] \equiv \mathrm{alias}(P_0H_0) \quad \text{and} \quad [\langle P'\rangle, \langle H'\rangle] \equiv \mathrm{alias}(P'H'). \tag{8.68}$$

All triple commutators vanish, since the double commutators are either $\mathbb{1}$ or type-$\mathtt{Z}$ elements, and we can use Lemma 8.56. We get:

$$[\langle P\rangle, \langle H\rangle] \equiv [\langle P_0\rangle\langle P'\rangle, \langle H_0\rangle\langle H'\rangle] \qquad \text{(Lemma 8.14)}$$

$$\equiv \text{alias}(P_0H_0)\cdot \text{alias}(P'H') \qquad \text{(Corollary 6.14)}$$
$$\equiv \text{alias}(P_0H_0 + P'H') = \text{alias}(PH), \qquad \text{(Lemma 8.55)}$$

as desired. □

Finally, we prove:

Proof of Lemma 8.59. We prove the first equality, as the second is dual. Claim 8.32 gives us a decomposition $P = \sum_{\alpha=0}^{n_{\mathrm{A}}/m-1}\sum_{\gamma=0}^{\lfloor n_{\mathrm{C}}/m\rfloor} P^{(\alpha,\gamma)}$ where $P^{(\alpha,\gamma)}$ is supported only on $\mathcal{I}_{\mathrm{A}}^{a(\alpha)} \times \mathcal{I}_{\mathrm{C}}^{c(\gamma)}$. The same argument gives an analogous decomposition $H = \sum_{\gamma'=0}^{\lfloor n_{\mathrm{C}}/m\rfloor}\sum_{\delta=0}^{n_{\mathrm{D}}/m-1} H^{(\gamma',\delta)}$, where $H^{(\gamma',\delta)}$ is supported only on $\mathcal{I}_{\mathrm{C}}^{c(\gamma')} \times \mathcal{I}_{\mathrm{D}}^{d(\delta)}$. Hence

$$[\langle P\rangle,\langle H\rangle] = \left[\left\langle \sum_{\alpha=0}^{n_{\mathrm{A}}/m-1}\sum_{\gamma=0}^{\lfloor n_{\mathrm{C}}/m\rfloor} P^{(\alpha,\gamma)}\right\rangle, \left\langle \sum_{\gamma'=0}^{\lfloor n_{\mathrm{C}}/m\rfloor}\sum_{\delta=0}^{n_{\mathrm{D}}/m-1} H^{(\gamma',\delta)}\right\rangle\right]$$
$$\equiv \left[\prod_{\alpha=0}^{n_{\mathrm{A}}/m-1}\prod_{\gamma=0}^{\lfloor n_{\mathrm{C}}/m\rfloor} \underbrace{\left\langle P^{(\alpha,\gamma)}\right\rangle}_{=:u_{\alpha,\gamma}}, \prod_{\gamma'=0}^{\lfloor n_{\mathrm{C}}/m\rfloor}\prod_{\delta=0}^{n_{\mathrm{D}}/m-1} \underbrace{\left\langle H^{(\gamma',\delta)}\right\rangle}_{=:v_{\gamma',\delta}}\right].$$

We again apply Corollary 6.14. By Claim 8.67, the pairwise commutators are

$$[u_{\alpha,\gamma}, v_{\gamma',\delta}] = \left[\left\langle P^{(\alpha,\gamma)}\right\rangle, \left\langle H^{(\gamma',\delta)}\right\rangle\right] \equiv \text{alias}\left(P^{(\alpha,\gamma)}H^{(\gamma',\delta)}\right)$$

and by Claim 8.67 and Lemma 8.56, the triple commutators are

$$[u_{\alpha',\gamma''},[u_{\alpha,\gamma}, v_{\gamma',\delta}]] = \left[\left\langle P^{(\alpha',\gamma'')}\right\rangle, \left[\left\langle P^{(\alpha,\gamma)}\right\rangle, \left\langle H^{(\gamma',\delta)}\right\rangle\right]\right] \equiv \left[\left\langle P^{(\alpha',\gamma'')}\right\rangle, \text{alias}\left(P^{(\alpha,\gamma)}H^{(\gamma',\delta)}\right)\right] \equiv e$$

and similarly for $[v_{\gamma'',\delta'},[u_{\alpha,\gamma}, v_{\gamma',\delta}]]$. Hence by Corollary 6.14,

$$[\langle P\rangle,\langle H\rangle] \equiv \prod_{\alpha=n_{\mathrm{A}}/m-1}^{0}\prod_{\gamma=\lfloor n_{\mathrm{C}}/m\rfloor}^{0}\prod_{\gamma'=0}^{\lfloor n_{\mathrm{C}}/m\rfloor}\prod_{\delta=0}^{n_{\mathrm{D}}/m-1} \text{alias}\left(P^{(\alpha,\gamma)}\cdot H^{(\gamma',\delta)}\right)$$
$$\equiv \text{alias}\left(\sum_{\alpha=0}^{n_{\mathrm{A}}/m-1}\sum_{\gamma=0}^{\lfloor n_{\mathrm{C}}/m\rfloor}\sum_{\gamma'=0}^{\lfloor n_{\mathrm{C}}/m\rfloor}\sum_{\delta=0}^{n_{\mathrm{D}}/m-1} P^{(\alpha,\gamma)}\cdot H^{(\gamma',\delta)}\right)$$
$$= \text{alias}\left(\left(\sum_{\alpha=0}^{n_{\mathrm{A}}/m-1}\sum_{\gamma=0}^{\lfloor n_{\mathrm{C}}/m\rfloor} P^{(\alpha,\gamma)}\right)\cdot\left(\sum_{\gamma'=0}^{\lfloor n_{\mathrm{C}}/m\rfloor}\sum_{\delta=0}^{n_{\mathrm{D}}/m-1} H^{(\gamma',\delta)}\right)\right) = \text{alias}(PH),$$

as desired. Each invocation of Claim 8.67 takes $O(\rho^9)$ derivation steps, Lemma 8.56 $O(\rho^9)$ steps, and Lemma 8.55 $O(\rho^7)$ steps. The dominant cost is invoking Claim 8.67 in the triple commutator calculation (which happens $O(\rho^6)$ times). □

8.14 Remaining relations

Finally, we verify the remaining Steinberg relations.

Claim 8.69. *For every $A \in \mathrm{Mat}_{\mathtt{A}}$, $P \in \mathrm{Mat}_{\mathtt{P}}$, and $H \in \mathrm{Mat}_{\mathtt{H}}$,*

$$\vdash_{15} [\langle A\rangle,[\langle P\rangle,\langle H\rangle]] \equiv \text{alias}(APH).$$

Dually, for every $D \in \mathrm{Mat}_{\mathtt{D}}$, $F \in \mathrm{Mat}_{\mathtt{F}}$, and $Q \in \mathrm{Mat}_{\mathtt{Q}}$,

$$\vdash_{15} [\langle D\rangle,[\langle F\rangle,\langle Q\rangle]] \equiv \text{alias}(D^\dagger FQ).$$

Proof. We prove the first equality, as the second is dual. We have:

$$\begin{aligned}
\langle A\rangle[\langle P\rangle,\langle H\rangle] &\equiv \langle A\rangle\langle P\rangle\langle H\rangle\langle -P\rangle\langle -H\rangle && \text{(inv. of } \mathtt{P}\text{'s and } \mathtt{H}\text{'s, Lemma 8.14)}\\
&\equiv \langle AP\rangle\langle P\rangle\langle A\rangle\langle H\rangle\langle -P\rangle\langle -H\rangle && \text{(Equation (3.49) and comm. of } \mathtt{A}\text{'s and } \mathtt{P}\text{'s, Lemma 8.14)}\\
&\equiv \langle AP\rangle\langle P\rangle\langle H\rangle\langle A\rangle\langle -P\rangle\langle -H\rangle && \text{(comm. of } \mathtt{A}\text{'s and } \mathtt{H}\text{'s, Lemma 8.14)}\\
&\equiv \langle AP\rangle\langle P\rangle\langle H\rangle\langle -AP\rangle\langle -P\rangle\langle A\rangle\langle -H\rangle \\
& && \text{(Equation (3.49) and comm. of } \mathtt{A}\text{'s and } \mathtt{P}\text{'s, Lemma 8.14)}\\
&\equiv \langle AP\rangle\langle P\rangle\langle H\rangle\langle -AP\rangle\langle -P\rangle\langle -H\rangle\langle A\rangle && \text{(comm. of } \mathtt{A}\text{'s and } \mathtt{H}\text{'s, Lemma 8.14)}\\
&\equiv \langle P\rangle[\langle AP\rangle,\langle H\rangle]\langle H\rangle\langle AP\rangle\langle -AP\rangle\langle -P\rangle\langle -H\rangle\langle A\rangle && \text{(Equation (3.49))}\\
&\equiv \langle P\rangle\cdot \mathsf{alias}(APH)\cdot\langle H\rangle\langle AP\rangle\langle -AP\rangle\langle -P\rangle\langle -H\rangle\langle A\rangle && \text{(Lemma 8.59)}\\
&\equiv \mathsf{alias}(APH)\cdot[\langle P\rangle,\langle H\rangle]\langle A\rangle. && \text{(Lemma 8.49)}
\end{aligned}$$

The dominant cost is applying Lemma 8.59. ☐

Lemma 8.70 (Remaining Steinberg relations)**.** *For every* $Z \in \mathrm{Mat}_{\mathtt{Z}}$*,* $A \in \mathrm{Mat}_{\mathtt{A}}$*, and* $D \in \mathrm{Mat}_{\mathtt{D}}$*:*

$$\vdash_{17} \begin{cases} [\langle A\rangle, \mathsf{alias}(Z)] \equiv \mathsf{alias}(AZ), \\ [\langle D\rangle, \mathsf{alias}(Z)] \equiv \mathsf{alias}\big(ZD^{\dagger}\big). \end{cases}$$

Proof. We prove the first equation, as the second is dual. By Definition 8.54, Z is a sum of $\rho_{\mathtt{Z}} = O(\rho^3)$ type-$\mathtt{Z}$ basic block matrices:

$$Z = \sum_{i=1}^{\rho_{\mathtt{Z}}} Z^{t_{\mathtt{Z}}^{(i)}}_{a(\alpha^{(i)});d(\delta^{(i)})}\Big(M^{(i)}\Big) = \sum_{i=1}^{\rho_{\mathtt{Z}}} P^{(i)}\cdot H^{(i)}.$$

Therefore by Lemma 8.55,

$$\mathsf{alias}(Z) = \prod_{i=1}^{\rho_{\mathtt{Z}}}\Big[P^{(i)}, H^{(i)}\Big].$$

We now apply Corollary 6.14 a final time to:

$$[\langle A\rangle, \mathsf{alias}(Z)] = \left[\underbrace{\langle A\rangle}_{=:g}, \prod_{i=1}^{\rho_{\mathtt{Z}}} \underbrace{\Big[\Big\langle P^{(i)}\Big\rangle, \Big\langle H^{(i)}\Big\rangle\Big]}_{=:v_i}\right].$$

By Claim 8.69, the pairwise commutators are:

$$[g, v_i] = \Big[\langle A\rangle, \Big[\Big\langle P^{(i)}\Big\rangle, \Big\langle H^{(i)}\Big\rangle\Big]\Big] \equiv \mathsf{alias}\Big(AP^{(i)}H^{(i)}\Big).$$

and by Claim 8.69 and Lemma 8.59, the triple commutators are:

$$[v_{i'}, [g, v_i]] = \Big[\Big[\Big\langle P^{(i')}\Big\rangle, \Big\langle H^{(i')}\Big\rangle\Big], \Big[\langle A\rangle, \Big[\Big\langle P^{(i)}\Big\rangle, \Big\langle H^{(i)}\Big\rangle\Big]\Big]\Big] \equiv \Big[\mathsf{alias}\Big(P^{(i')}H^{(i')}\Big), \mathsf{alias}\Big(AP^{(i)}H^{(i)}\Big)\Big] \equiv e.$$

Hence

$$[\langle A\rangle, \mathsf{alias}(Z)] = \left[g, \prod_{i=1}^{\rho_{\mathtt{Z}}} v_i\right] \equiv \prod_{i=1}^{\rho_{\mathtt{Z}}} \mathsf{alias}\Big(AP^{(i)}H^{(i)}\Big) \equiv \mathsf{alias}\left(A\sum_{i=1}^{\rho_{\mathtt{Z}}} P^{(i)}H^{(i)}\right) \equiv \mathsf{alias}(AZ).$$

Each invocation of Lemma 8.59 takes $O(\rho^{15})$ derivation steps, Claim 8.69 $O(\rho^{15})$, and Lemma 8.55 $O(\rho^9)$. The dominant cost is invoking Claim 8.69, which happens $O(\rho^2)$ times. ☐

8.15 Proof of Lemma 8.15

Given all this setup, we can now prove Lemma 8.15, thereby establishing Lemma 8.1.

Proof of Lemma 8.15. Lemma 8.33 gives that $\langle P\rangle$ and $\langle Q\rangle$ commute. Type-Z aliases (i.e., alias(Z)) is defined in Definition 8.54. Linearity for type-Z aliases is Lemma 8.55. The commutation relation of type-Z aliases with all λ elements for $\lambda \in \{\mathtt{B}, \mathtt{C}, \mathtt{F}, \mathtt{G}, \mathtt{H}, \mathtt{P}, \mathtt{Q}\}$ is in Lemma 8.56. The commutators of type-Z aliases with type-A and -D elements are given in Lemma 8.70. ☐

9 The Kaufman–Oppenheim complex: all the theorems

In the following theorem we collect many properties of the KO complexes, including that they yield sparse low-soundness agreement testers and can be used in the [BMVY25] PCP system. Many of the properties were previously known or essentially known; it is the last ones, Items 12 to 16, that are new to this paper.

Theorem 9.1. *For sufficiently large n and $p \geq \mathrm{poly}(n)$ prime, the [KO23; KOW25] complexes $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ ($\kappa = 3$) for $s > 3n$ satisfy the following properties:*

1. *$\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ is an $(n+1)$-partite, pure rank-$(n+1)$ clique coset complex.*
2. *All links of $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ are connected.*
3. *Defining $Q = p^s$, the complex $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ has N vertices, for*
$$N = \frac{n+1}{p^{2n(n+1)}} \cdot |\mathrm{SL}_{n+1}(\mathbb{F}_Q)|, \quad |\mathrm{SL}_{n+1}(\mathbb{F}_Q)| = \tfrac{1}{Q-1}\prod_{i=0}^{n}(Q^{n+1} - Q^i). \tag{9.2}$$
4. *For a fixed n, and given targets N_0 and $p_0 \leq \mathrm{polylog}\, N_0$, one can in deterministic $\mathrm{polylog}\, N_0$ time find values $p_0 \leq p \leq O(p_0)$ and s and hence N such that $N_0 \leq N \leq N_0^{1+O(1/p_0^{.4})}$, the latter bound being $O(N_0)$ if $p_0 \geq \log^{2.5} N$.*
5. *For fixed n, and assuming $p \leq \mathrm{polylog}\, N$, the complexes $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ are strongly explicit, in the sense that the following things can be done in $\mathrm{polylog}\, N$ time: Computing N, picking a uniformly random vertex, computing the parts and neighborhoods of a vertex.*
6. *$\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ is regular: every vertex participates in $\frac{p^{2n(n+1)}}{n+1}$ maximal cliques.*
7. *Every vertex-link is isomorphic to $\mathfrak{L}\mathfrak{A}_{\mathbb{F}_p}^{\kappa}(n)$, an n-partite, pure rank-n clique coset complex.*
8. *The automorphism group of $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ acts transitively on the maximal cliques, and the same is true of $\mathfrak{L}\mathfrak{A}_{\mathbb{F}_p}^{\kappa}(n)$.*
9. *Every link $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)_w$ with $|w| \leq n-1$ has normalized 2nd eigenvalue at most $\frac{1}{\sqrt{p}-n}$.*
10. *Every pinned bipartite restriction $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)_w[\{i\},\{j\}]$ has normalized 2nd eigenvalue at most $\frac{2}{\sqrt{p}}$.*
11. *The graph underlying $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ has all nontrivial normalized eigenvalues at most $\frac{2}{d}$ in absolute value.*
12. *Within $\mathfrak{L}\mathfrak{A}_{\mathbb{F}_p}^{\kappa}(n)$, the bipartite graph between parts a and b has diameter $\mathrm{poly}(1/\xi)$ provided $|b-a| \geq \xi n$.*
13. *Within any pinning (link) $\mathfrak{L}\mathfrak{A}_{\mathbb{F}_p}^{\kappa}(n)_w$, provided $a < b < c$ are inside a contiguous sequence of unpinned parts, the triangle complex on parts a, b, c has 1-coboundary expansion at least $\mathrm{poly}(\xi)$ over any group, provided $b - a, c - b \geq \xi n$.*
14. *Provided $r \ll \sqrt{n}$, $\mathfrak{L}\mathfrak{A}_{\mathbb{F}_p}^{\kappa}(n)$ has r-triword expansion at least $\exp(-O(r^{.67}))$ over any group (i.e., its r-faces complex has 1-coboundary expansion at least $\exp(-O(r^{.67}))$).*

15. *For $\delta > 0$, provided $k \geq \exp(\mathrm{poly}(1/\delta))$ and $n \geq \exp(\mathrm{poly}(k))$, the V-test performed with the size-k sub-edges of $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ is a δ-sound agreement tester in the sense of Definition 2.1.*

16. *$\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ may be used for the direct-product testing component of the Bafna–Minzer–Vyas PCP system (specifically, in [BMVY25, Lemma 5.3]). This achieves their arbitrary constant $\delta > 0$ soundness with $N \mapsto N \operatorname{polylog} N$ proof length using only elementary ingredients.*

Proof. Item 1 is mostly by definition; the fact that $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ is a clique complex is Theorem 3.56. Item 2 is Theorem 3.58. Item 3 arises from the definition of $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}_p)$ and is discussed in §2.1.1. Item 4 takes a little work. Treating n as constant, note that $N = p^{\Theta(s)}$. Thus given p_0 and N_0, the desired s will satisfy $s = \Theta(\frac{\log N_0}{\log p_0})$. Adjusting s by ± 1 changes N by a factor of $p^{\Theta(1)}$; so assuming $p_0 \leq \operatorname{polylog} N$, an algorithm can (in $\operatorname{polylog} N$ time) determine an s that gets N into the range $[\frac{N_0}{\operatorname{polylog} N_0} \,..\, N_0]$. Next, the algorithm can consider adjusting the prime p_0 upward. It is known [BHP01] that for any prime p, there is another that is at most $p + O(p^{.6})$. Thus (in $\operatorname{polylog} N$ time) an algorithm can increase p by a factor of $1 + O(1/p^{.4}) = \exp(O(1/p^{.4}))$, which increases N by a factor of $\exp(O(s/p^{.4})) \leq N^{O(1/p^{.4})}$. At most $O(p_0^{.4})$ such repeated adjustments get N and p into the claimed ranges.

The preceding analysis covers the "computing N" portion of Item 5, and the fact that computing the incidence structure is also computable in $\operatorname{polylog} N$ time is clear from the description in §2.1.1. (The equivalence classes are of size $\mathrm{poly}(p)$ and can be maintained explicitly; the computations over $\mathbb{F}_Q$ can also be done in $\operatorname{polylog} N$ time.) The regularity and degree facts of Item 6 are in §2.1.1. Item 7 is Theorem 3.55, and Item 8 is Proposition 3.42.

Moving on to eigenvalue bounds, Item 9 is from [KOW25, Theorem 2.8]. Item 10 is Corollary 3.60. Item 11 follows from the one-sided spectral bound in Item 9, together with $(n+1)$-partiteness, trickling down [Opp18], and $p \gg n^2$; see the discussion of trickling down in Appendix B.

We come to the new results in our work. Item 12 is from Lemma 5.2, Item 13 is Theorem 1.3, and Item 14 is Theorem 1.2. With these and preceding results in hand, we can now finally prove that the KO complex is a δ-sound agreement tester (i.e., establish Theorem 1.1 and Item 15) by appealing to Bafna–Minzer/Dikstein–Dinur's Theorem 3.21; the only piece missing for that theorem is its trivial global cohomology condition, which is provided by Kaufman–Oppenheim–Weinberger's Theorem 2.9. At last, to verify the PCP result Item 16, we observe [BMVY25, Lemma 5.3] requires a complex that satisfies the long list of properties in [BMVY25, Theorem 2.13]; this theorem shows that all those properties are satisfied except one. The lacuna is that [BMVY25, Theorem 2.13] formally requires the diameter bound from Item 12 to be $O(1/\xi)$, whereas we have only $\mathrm{poly}(1/\xi)$. However it is not hard to check that having this slightly worse bound makes no material difference in the one place it is used, [BMVY25, Section 4.2]. □

References

[BC13] Francis Buekenhout and Arjeh M. Cohen. *Diagram Geometry: Related to Classical Groups and Buildings*. Vol. 57. Ergebnisse Der Mathematik Und Ihrer Grenzgebiete. 3. Folge / A Series of Modern Surveys in Mathematics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013. DOI: 10.1007/978-3-642-34453-4.

[BD01] Daniel K. Biss and Samit Dasgupta. "A Presentation for the Unipotent Group over Rings with Identity". In: *Journal of Algebra* 237.2 (Mar. 2001), pp. 691–707. DOI: 10.1006/jabr.2000.8604.

[BHP01] R. C. Baker, G. Harman, and J. Pintz. "The Difference Between Consecutive Primes, II". In: *Proceedings of the London Mathematical Society* 83.3 (Nov. 2001), pp. 532–562. DOI: 10.1112/plms/83.3.532.

[BLM24] Mitali Bafna, Noam Lifshitz, and Dor Minzer. "Constant Degree Direct Product Testers with Small Soundness". In: *65th Annual Symposium on Foundations of Computer Science*. FOCS 2024 (Chicago, IL, USA, Oct. 27–30, 2024). IEEE Computer Society, July 17, 2024, pp. 862–869. DOI: 10.1109/FOCS61266.2024.00059.

[BM24] Mitali Bafna and Dor Minzer. "Characterizing Direct Product Testing via Coboundary Expansion". In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. STOC 2024 (Vancouver, BC, Canada, June 24–28, 2024). Association for Computing Machinery, Feb. 1, 2024, pp. 1978–1989. DOI: 10.1145/3618260.3649714.

[BMVY25] 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 (Prague, Czechia, June 23–27, 2025). New York, NY, USA: Association for Computing Machinery, June 15, 2025, pp. 45–53. DOI: 10.1145/3717823.3718197.

[Bun52] Svend Bundgaard. "On a Kind of Homotopy in Regular Numbered Complexes". In: *Meddelanden från Lunds Universitets Matematiska Seminarium* Tome Suppl. (1952), pp. 35–46.

[Chv79] Václav Chvátal. "The Tail of the Hypergeometric Distribution". In: *Discrete Mathematics* 25.3 (1979), pp. 285–287. DOI: 10.1016/0012-365X(79)90084-0.

[CL25] Michael Chapman and Alexander Lubotzky. "Stability of Homomorphisms, Coverings and Cocycles II: Examples, Applications and Open Problems". In: *Advances in Mathematics* 463 (Mar. 2025), p. 110117. DOI: 10.1016/j.aim.2025.110117.

[DD19] Yotam Dikstein and Irit Dinur. "Agreement Testing Theorems on Layered Set Systems". In: *2019 IEEE 60th Annual Symposium on Foundations of Computer Science*. FOCS 2019 (Baltimore, MD, USA, Nov. 9–12, 2019). IEEE Computer Society, Nov. 2019, pp. 1495–1524. DOI: 10.1109/FOCS.2019.00088.

[DD24a] Yotam Dikstein and Irit Dinur. "Agreement Theorems for High Dimensional Expanders in the Low Acceptance Regime: The Role of Covers". In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. STOC 2024. Association for Computing Machinery, June 10, 2024, pp. 1967–1977. DOI: 10.1145/3618260.3649685.

[DD24b] Yotam Dikstein and Irit Dinur. "Coboundary and Cosystolic Expansion Without Dependence on Dimension or Degree". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. RANDOM 2024 (London, UK, Aug. 28–30, 2024). Ed. by Amit Kumar and Noga Ron-Zewi. Vol. 317. LIPIcs. 2024, 62:1–62:24. DOI: 10.4230/LIPICS.APPROX/RANDOM.2024.62.

[DD24c] Yotam Dikstein and Irit Dinur. "Swap Cosystolic Expansion". In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. STOC 2024 (Vancouver, BC, Canada, June 24–28, 2024). Association for Computing Machinery, June 10, 2024, pp. 1956–1966. DOI: 10.1145/3618260.3649780.

[DDL24] Yotam Dikstein, Irit Dinur, and Alexander Lubotzky. "Low Acceptance Agreement Tests via Bounded-Degree Symplectic HDXs". In: *65th Annual Symposium on Foundations of Computer Science*. FOCS 2024 (Chicago, IL, USA, Oct. 27–30, 2024). IEEE Computer Society, Oct. 27, 2024, pp. 826–861. DOI: 10.1109/FOCS61266.2024.00058.

[DG08] Irit Dinur and Elazar Goldenberg. "Locally Testing Direct Product in the Low Error Range". In: *49th Annual IEEE Symposium on Foundations of Computer Science*. FOCS 2008. IEEE Computer Society, Oct. 2008, pp. 613–622. DOI: 10.1109/FOCS.2008.26.

[DK17] Irit Dinur and Tali Kaufman. "High Dimensional Expanders Imply Agreement Expanders". In: *IEEE 58th Annual Symposium on Foundations of Computer Science*. FOCS 2017. Oct. 2017, pp. 974–985. DOI: 10.1109/FOCS.2017.94.

[DM22] Irit Dinur and Roy Meshulam. "Near Coverings and Cosystolic Expansion". In: *Archiv der Mathematik* 118.5 (May 2022), pp. 549–561. DOI: 10.1007/s00013-022-01720-6.

[DR06] Irit Dinur and Omer Reingold. "Assignment Testers: Towards a Combinatorial Proof of the PCP Theorem". In: *SIAM Journal on Computing* 36.4 (Jan. 2006). Conference version in FOCS 2004, pp. 975–1024. DOI: 10.1137/S0097539705446962.

[EK23] Shai Evra and Tali Kaufman. "Bounded Degree Cosystolic Expanders of Every Dimension". In: *Journal of the American Mathematical Society* 37.1 (July 21, 2023). Conference version in STOC 2016, pp. 39–68. DOI: 10.1090/jams/1019.

[Gar79] Peter Freedman Garst. "Cohen-Macaulay Complexes and Group Actions". PhD thesis. University of Wisconsin-Madison, May 17, 1979. 144 pp.

[GK23] Roy Gotlib and Tali Kaufman. "List Agreement Expansion from Coboundary Expansion". In: *14th Innovations in Theoretical Computer Science Conference*. ITCS 2023 (Cambridge, MA, USA, Jan. 10–13, 2023). Ed. by Yael Tauman Kalai. Vol. 251. LIPIcs. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 61:1–61:23. DOI: 10.4230/LIPICS.ITCS.2023.61.

[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*. STOC 2022 (Rome, Italy, June 20–24, 2022). Association for Computing Machinery, June 9, 2022, pp. 176–184. DOI: 10.1145/3519935.3520004.

[Gro10] Mikhail Gromov. “Singularities, Expanders and Topology of Maps. Part 2: From Combinatorics to Topology Via Algebraic Isoperimetry”. In: *Geometric and Functional Analysis* 20.2 (Aug. 2010), pp. 416–526. DOI: 10.1007/s00039-010-0073-8.

[GS00] Oded Goldreich and Shmuel Safra. “A Combinatorial Consistency Lemma with Application to Proving the PCP Theorem”. In: *SIAM Journal on Computing* 29.4 (Jan. 2000). Conference version in RANDOM 1997, pp. 1132–1154. DOI: 10.1137/S0097539797315744.

[HO89] Alexander J. Hahn and O. Timothy O’Meara. *The Classical Groups and K-Theory*. Vol. 291. Grundlehren Der Mathematischen Wissenschaften. Springer Berlin Heidelberg, 1989. DOI: 10.1007/978-3-662-13152-7.

[HS24] Prahladh Harsha and Ramprasad Saptharishi. “On the Elementary Construction of High-Dimensional Expanders by Kaufman and Oppenheim”. In: *Theory of Computing* 20.1 (2024), pp. 1–22. DOI: 10.4086/toc.2024.v020a005.

[IKW12] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. “New Direct-Product Testers and 2-Query PCPs”. In: *SIAM Journal on Computing* 41.6 (Jan. 2012), pp. 1722–1768. ISSN: 0097-5397. DOI: 10.1137/09077299X.

[KKL16] Tali Kaufman, David Kazhdan, and Alexander Lubotzky. “Isoperimetric Inequalities for Ramanujan Complexes and Topological Expanders”. In: *Geometric and Functional Analysis* 26.1 (Feb. 1, 2016). Conference version in FOCS 2014, pp. 250–287. DOI: 10.1007/s00039-016-0362-y.

[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*. STOC 2018 (Los Angeles, CA, USA, June 25–29, 2018). Association for Computing Machinery, June 20, 2018, pp. 773–786. DOI: 10.1145/3188745.3188782.

[KO21] Tali Kaufman and Izhar Oppenheim. “Coboundary and Cosystolic Expansion from Strong Symmetry”. In: *48th International Colloquium on Automata, Languages, and Programming*. ICALP 2021 (Glasgow, Scotland / virtual, July 12–16, 2021). Ed. by Nikhil Bansal, Emanuela Merelli, and James Worrell. Vol. 198. LIPIcs. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021, 84:1–84:16. DOI: 10.4230/LIPICS.ICALP.2021.84.

[KO23] Tali Kaufman and Izhar Oppenheim. “High Dimensional Expanders and Coset Geometries”. In: *European Journal of Combinatorics* 111 (June 2023), p. 103696. ISSN: 01956698. DOI: 10.1016/j.ejc.2023.103696.

[KOW25] Tali Kaufman, Izhar Oppenheim, and Shmuel Weinberger. “Coboundary Expansion of Coset Complexes”. In: *57th Annual ACM Symposium on Theory of Computing*. STOC 2025 (Prague, Czechia, June 23–27, 2025). Association for Computing Machinery, June 15, 2025, pp. 1722–1731. DOI: 10.1145/3717823.3718136.

[Lan50] Folke Lannér. “On Complexes with Transitive Groups of Automorphisms”. In: *Meddelanden från Lunds Universitets Matematiska Seminarium* 11 (1950).

[LMM16] Alexander Lubotzky, Roy Meshulam, and Shahar Mozes. “Expansion of Building-like Complexes”. In: *Groups, Geometry, and Dynamics* 10.1 (Feb. 9, 2016), pp. 155–175. DOI: 10.4171/ggd/346.

[LSV05] Alexander Lubotzky, Beth Samuels, and Uzi Vishne. “Explicit Constructions of Ramanujan Complexes of Type $\widetilde{a}_d$”. In: *European Journal of Combinatorics*. Combinatorics and Representation Theory 26.6 (Aug. 1, 2005), pp. 965–993. DOI: 10.1016/j.ejc.2004.06.007.

[OP22] Ryan O’Donnell and Kevin Pratt. “High-Dimensional Expanders from Chevalley Groups”. In: *LIPIcs, Volume 234, CCC 2022*. 37th Computational Complexity Conference (Philadelphia, PA, USA, July 20–23, 2022). Ed. by Shachar Lovett. Vol. 234. LIPIcs. CCC 2022: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022, 18:1–18:26. DOI: 10.4230/LIPICS.CCC.2022.18.

[Opp18] Izhar Oppenheim. “Local Spectral Expansion Approach to High Dimensional Expanders Part I: Descent of Spectral Gaps”. In: *Discrete & Computational Geometry* 59.2 (Mar. 2018), pp. 293–330. DOI: 10.1007/s00454-017-9948-x.

[OS25] Ryan O’Donnell and Noah Singer. “The Relative Dehn Method for Coboundary Expansion”. 2025. URL: https://noahsinger.org/static/pdfs/relative%20dehn.pdf.

[Rot95] Joseph J. Rotman. *An Introduction to the Theory of Groups*. 4th ed. Vol. 148. Graduate Texts in Mathematics. New York, NY: Springer New York, 1995. DOI: 10.1007/978-1-4612-4176-8.

[Sho90] Victor Shoup. "New Algorithms for Finding Irreducible Polynomials over Finite Fields". In: *Mathematics of Computation* 54.189 (1990). Conference version in FOCS 1988, pp. 435–447. DOI: 10.1090/S0025-5718-1990-0993933-0.

[Sta79] Richard P. Stanley. "Balanced Cohen-Macaulay Complexes". In: *Transactions of the American Mathematical Society* 249.1 (1979), pp. 139–157. DOI: 10.1090/S0002-9947-1979-0526314-6.

[Ste62] Robert Steinberg. "Générateurs, Relations et Revêtements de Groupes Algébriques". In: *Centre Belge Rech. math., Colloq. Théorie des Groupes Algébriques* (1962), pp. 113–127.

[Tit86] Jacques Tits. *Buildings of Spherical Type and Finite BN-Pairs*. Lecture Notes in Mathematics Ser v.386. Berlin, Heidelberg: Springer Berlin / Heidelberg, 1986. DOI: 10.1007/978-3-540-38349-9.

A Coboundary expansion of "linearly ordered" complexes

In this appendix, we prove Theorem 3.30, which we restate here for convenience:

Theorem 3.30 (Proving triword expansion for "linearly ordered" complexes, abstracted from [BLM24])**.** *There exist absolute constants $K > 1$ and $\epsilon_0 > 0$ such that for every $r \in \mathbb{N}$, there exist $\xi > 0$ and $n_0 \in \mathbb{N}$ such that for every $n \geq n_0 \in \mathbb{N}$, the following holds. Let $\mathfrak{X}$ be a $[1\,..\,n]$-indexed simplicial complex and Γ an arbitrary group. Suppose that $\mathfrak{X}$ satisfies the following conditions:*

1. Spectral expansion:
   $\mathfrak{X}$ is $\epsilon_0 r^{-2}$-product.

2. Productization:
   $\mathfrak{X}$ productizes on the line, in the sense of Definition 3.28.

3. Symmetry:
   $\mathfrak{X}$ is strongly symmetric, in the sense of Definition 3.22.

4. Coboundary expansion of separated singleton restrictions:
   For every $W \subseteq [1\,..\,n]$ and $\ell_1 < \ell_2 < \ell_3 \in [1\,..\,n]$ such that $W \cap [\ell_1\,..\,\ell_3] = \emptyset$ and $\{\ell_1, \ell_2, \ell_3\}$ is ξn-separated, for every W-word w, the complex $\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}, \{\ell_3\}]$ has 1-triword expansion at least β over Γ.

5. Diameter of separated singleton restrictions:
   For every $W \subseteq [1\,..\,n]$ and $\ell_1 < \ell_2 \in [1\,..\,n]$ such that $W \cap [\ell_1\,..\,\ell_2] = \emptyset$ and $\{\ell_1, \ell_2\}$ is ξn-separated, for every W-word w, the complex $\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}]$ has diameter at most C_0.

Then for every $W \subset_{\leq r} [1\,..\,n]$ and W-word w, $\mathfrak{X}_w$ has r-triword expansion at least $(K\beta/C_0)^{O(r^{0.67})}$ over Γ.

We emphasize that the proof of Theorem 3.30 is extracted from the analysis of the triword expansion of the [CL25] complex in the work of Bafna, Lifshitz, and Minzer [BLM24]. Our contribution in this appendix is to restate their method in a generic way which can be applied to other complexes (and may be useful to future work) and achieves a slightly better final expansion bound (though this does not matter for applications). (We do also make a slight technical adjustment to their lemma, since we prove triword expansion without additive error, though this is also essentially already done in the work of [DD24c].)

A.1 Coboundary expansion from pinned triword expansion

We now state a lemma due to Bafna, Lifshitz, and Minzer [BLM24] which can be used to prove triword expansion of an induced triangle complex via triword expansion of "smaller" complexes. The lemma is not stated explicitly in the work of [BLM24], but it is present as equation (8) within the proof of Lemma 3.6 in that paper.

The general setup is as follows: Suppose we have $\mathfrak{X}$ an $\mathcal{I}$-indexed simplicial complex, subsets $R_1, R_2, R_3 \sqsubset_{\leq r} \mathcal{I}$, and subsets $S_1, S_2, S_3 \sqsubset R_1 \cup R_2 \cup R_3$. We relate the triword expansion of the "global" triangle complex $\mathfrak{T} := \mathfrak{X}[R_1, R_2, R_3]$ to the triword expansion of the "central" triangle complex $\mathfrak{T}^* := \mathfrak{X}[S_1, S_2, S_3]$ as well as the "peripheral" complexes $\mathfrak{T}_{s_j} := \mathfrak{X}_{s_j}[R_1 \setminus S_j, R_2 \setminus S_j, R_3 \setminus S_j]$, where $j \in \{1, 2, 3\}$ and $s_j \in \operatorname{supp} \mathfrak{X}[S_j]$ is a

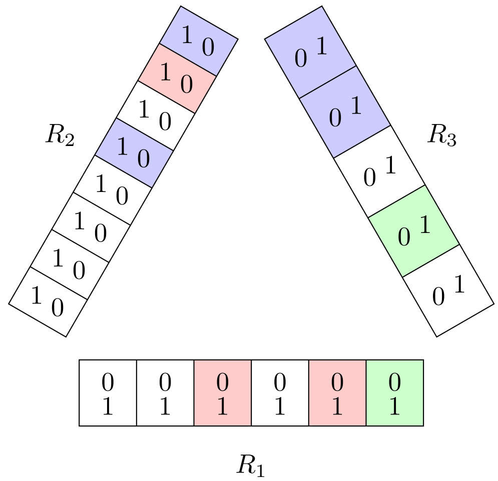

Figure 8: The situation posited in Lemma A.1, in the special case where all the underlying vertex-sets $\mathfrak{X}[i]$ are $\{0,1\}$. That is, let $\mathfrak{X}$ be a distribution on $\{0,1\}^n$. Let $R_1, R_2, R_3 \sqsubset \mathcal{I}$ and $S_1, S_2, S_3 \sqsubset R_1 \cup R_2 \cup R_3$. The overall triangle complex $\mathfrak{T}$ is the distribution on "triwords" $\mathfrak{T} = \mathfrak{X}[R_1, R_2, R_3]$, i.e., on triples in $\{0,1\}^{R_1} \times \{0,1\}^{R_2} \times \{0,1\}^{R_3}$. The sets S_1, S_2, S_3 are depicted in red, green, and blue, respectively. The central complex $\mathfrak{T}^* = \mathfrak{X}[S_1, S_2, S_3]$ is just a distribution on triples in $\{0,1\}^{S_1} \times \{0,1\}^{S_2} \times \{0,1\}^{S_3}$. Each peripheral complex $\mathfrak{T}_{s_j} = \mathfrak{X}_{s_j}[R_1 \setminus S_j, R_2 \setminus S_j, R_3 \setminus S_j]$ is a distribution on $\{0,1\}^{R_1 \setminus S_j} \times \{0,1\}^{R_2 \setminus S_j} \times \{0,1\}^{R_3 \setminus S_j}$ resulting from conditioning on the S_j-coordinates equaling s_j.

valid S_j-word. The peripheral complex $\mathfrak{T}_{s_j}$ can be viewed as an "induced subcomplex" of $\mathfrak{T}$ only on words consistent with s_j on S_j. In particular, every word/biword/triword in the peripheral complex $\mathfrak{T}_{s_j}$ can also be viewed as a word/biword/triword in $\mathfrak{T}$. ($\mathfrak{T}^*$ can't be viewed this way.) See Figure 8 below for a depiction of this situation. Note that, although S_1, S_2, and S_3 need to be pairwise disjoint, they need not interact nicely with R_1, R_2, and R_3's; each R_i could intersect each S_j. They also need not have the same size.

Lemma A.1 (from [BLM24, §3])**.** *There exist absolute constants $K_{\text{pin}} > 1, \epsilon_{\text{pin}} > 0$ such that the following holds. Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex, Γ be any coefficient group, $r \in \mathbb{N}$, $R_1, R_2, R_3 \sqsubset_{\leq r} \mathcal{I}$, and $S_1, S_2, S_3 \sqsubset R_1 \cup R_2 \cup R_3$. Let $0 < \beta, \beta^* < 1$, and suppose that $\mathfrak{X}$ satisfies the following conditions:*

1. Spectral expansion:
   $\mathfrak{X}$ *is* $\epsilon_{\text{pin}} r^{-2}$*-product.*

2. "Central" triword expansion:
   $\mathfrak{T}^* \coloneqq \mathfrak{X}[S_1, S_2, S_3]$ *has* 1*-triword expansion at least* β^* *over* Γ*.*

3. "Peripheral" triword expansion:
   For each $j \in \{1,2,3\}$ *and* $s_j \in \operatorname{supp}(\mathfrak{X}[S_j])$, $\mathfrak{T}_{s_j} \coloneqq \mathfrak{X}_{s_j}[R_1 \setminus S_j, R_2 \setminus S_j, R_3 \setminus S_j]$ *has* 1*-triword expansion at least* β *over* Γ*.*[17]

Then $\mathfrak{T} \coloneqq \mathfrak{X}[R_1, R_2, R_3]$ *has* 1*-triword expansion at least* $K_{\text{pin}} \beta^* \beta$ *over* Γ*.*

The constant K_{pin} is reasonable (at least 0.01).

Remark A.2. The proof of this lemma due to Bafna, Lifshitz, and Minzer [BLM24] also allows additive loss in the triword expansion, but we will not need this. $\Diamond$

A.2 More conditions for triword expansion of induced complexes

Next, we state another lemma due to Bafna, Lifshitz, and Minzer [BLM24] which repeatedly applies their pinning lemma (Lemma A.1 in this paper) to prove coboundary expansion of induced complexes when the

[17] The subscript-j is included to emphasize which coordinates are pinned.

corresponding restrictions satisfy certain "spreadness" conditions. In particular, we make the following definition, which is a generic condition of [BLM24, Item 3 in Assumption 1]. (Similar notions of spreadness are considered in [DD24c, Def. 8.3] and [DDL24, Def. 6.12].)

Definition A.3 (Pseudorandom spreadness)**.** Let $r \le n \in \mathbb{N}$, $R \subset_r [1\,..\,n]$, and $T, \eta \in \mathbb{N}$. R is (T,η)-*pseudorandomly spread* if the following holds: Define $\mathcal{I}_t := ((t-1)\frac{n}{T}\,..\,t\frac{n}{T}]$ for $t \in [1\,..\,T]$, so that each $\mathcal{I}_t$ has size $\frac{n}{T}$ and $\mathcal{I}_1, \dots, \mathcal{I}_T$ partition $[1\,..\,n]$. Then for every $t \in [1\,..\,T]$,

$$\frac{r}{T} - \eta \le |R_i \cap \mathcal{I}_t| \le \frac{r}{T} + \eta. \qquad \Diamond$$

Note that each R_i has size r, so by linearity of expectation, $\mathbb{E}_{\boldsymbol{t}\sim\mathbf{Unif}([1..T])}[|R_i \cap \mathcal{I}_{\boldsymbol{t}}|] = r \cdot \frac{1}{T} = \frac{r}{T}$.

Lemma A.4 ("Lopsided induction", [BLM24, Lemma 3.5])**.** *Let $K_{\text{pin}} > 1$, $\epsilon_{\text{pin}} > 0$ be the absolute constants from Lemma A.1. Let Γ be any group. Let $r \le n \in \mathbb{N}$, let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex. Let $\beta^* \le 1$ and suppose that $\mathfrak{X}$ satisfies the following conditions:*

1. Spectral expansion:
   $\mathfrak{X}$ is an $\epsilon_{\text{pin}} r^{-2}$-local spectral expander.

2. Singleton coboundary expansion:
   For every distinct $i_1, i_2, i_3 \in \mathcal{I}$, the triangle complex $\mathfrak{X}[\{i_1\},\{i_2\},\{i_3\}]$ has 1-triword expansion at least β^* over Γ.

Then for every $k_1, k_2, k_3 \in \mathbb{N}$ with $k_1 + k_2 + k_3 \le r$, $R_1, R_2, R_3 \sqsubset \mathcal{I}$ with $|R_1| = k_1$, $|R_2| = k_2$, $|R_3| = k_3$, the triangle complex $\mathfrak{T} := \mathfrak{X}[R_1, R_2, R_3]$ is has 1-triword expansion at least $(K_{\text{pin}}\beta^*)^{k_1+k_2+k_3}$ over Γ.

The following lemma reduces proving triword expansion of a pseudorandomly spread restriction to triword expansion of ordered restrictions:

Lemma A.5 ("Non-lopsided induction", [BLM24, Lemma 3.6])**.** *Let $K_{\text{pin}} > 1$, $\epsilon_{\text{pin}} > 0$ be the absolute constants from Lemma A.1. Let $\mathfrak{X}$ be an $[1\,..\,n]$-indexed simplicial complex and Γ be a coefficient group. Let $r \le n \in \mathbb{N}$, $R_1, R_2, R_3 \sqsubset_r [1\,..\,n]$, and $k, T, \eta \in \mathbb{N}$. Let $\beta^* < 1$. Suppose that $\mathfrak{X}$ satisfies the following conditions:*

1. Spectral expansion:
   $\mathfrak{X}$ is $\epsilon_0 r^{-2}$-product.

2. Coboundary expansion of ordered restrictions:
   For every $W \subset_{\le r-3k} R_1 \cup R_2 \cup R_3$, restriction $w \in \operatorname{supp}(\mathfrak{X}[W])$, $S_1 \prec S_2 \prec S_3 \sqsubset_k R_1 \cup R_2 \cup R_3 \setminus W$, the complex $\mathfrak{X}_w[S_1, S_2, S_3]$ has 1-triword expansion at least β^* over Γ.

3. Pseudorandom spreadness:
   For each $i \in \{1,2,3\}$, R_i is (T,η)-pseudorandomly spread.

Then the triangle complex $\mathfrak{X}[R_1, R_2, R_3]$ has 1-triword expansion at least β' over Γ, where

$$\beta' := (K_{\text{pin}}\beta^*)^{3(T(k+2\eta)+2(r/T))+r/k}.$$

Remark A.6. The proof of this lemma in [BLM24] sets $k := r^{0.01}$, $T := r/k^{10} = r^{0.9}$, $\eta := k^6 = r^{0.06}$, and $\tau := k^9 = r^{0.09}$, giving an exponent $O(r^{0.9}(r^{0.09} + r^{0.06}) + r^{0.1} + r^{0.99}) = O(r^{0.99})$. Instead, we will set $k := r^{\frac{1}{3}+\delta}$, $T := r^{\frac{1}{3}+\delta}$, and $\eta := r^{\frac{1}{3}+\delta}$. This gives an exponent of $O(r^{\frac{2}{3}+\delta} + r^{\frac{2}{3}}) = O(r^{\frac{2}{3}+\delta})$. We therefore give a reproof of the lemma so that we can account for the exponent with generic parameters, which we can then optimize (though it does not matter for the final testing result). We stress that our proof exactly follows theirs. $\Diamond$

Proof. Set $\tau := k + \eta$, so that $\tau - \eta = k$. By assumption, for every $i \in \{1,2,3\}, t \in [1\,..\,T]$, we have $|R_i \cap \mathcal{I}_t| = \frac{r}{T} + c_{i,t}$ for some $|c_{i,t}| \le \eta$, and for each $i \in \{1,2,3\}$, $\sum_{t=1}^{T} c_{i,t} = 0$. The argument is recursive.

The recursion. Every node of the recursion tree is of the form $N = (R'_1, R'_2, R'_3; W)$ maintaining the following invariants:

1. $W \subseteq [1\,..\,n]$ and for every $i \in \{1,2,3\}$, $R'_i \subseteq R_i \setminus W$. (Hence, R'_1, R'_2, and R'_3 are pairwise disjoint.)
2. $|R'_1| = |R'_2| = |R'_3| = r - kD$, where D is the depth of the node N.
3. For every $t \in [1\,..\,T]$, there exists $0 \le r_t \le \frac{r}{T}$ such that for every $i \in \{1,2,3\}$, $|R'_i \cap \mathcal{I}_t| = r_t + c_{i,t}$.
4. N is an *internal node* if and only if there exist $t_1 < t_2 < t_3 \in [1\,..\,T]$ with $r_{t_j} \ge \tau$ for each $j \in \{1,2,3\}$.

The root of the tree is $N^0 \coloneqq (R_1, R_2, R_3; \emptyset)$. Given an internal node $N = (R'_1, R'_2, R'_3; W)$, we recurse to three nodes N_j, one for each $j \in \{1,2,3\}$. Fix $j \in \{1,2,3\}$. Since R'_1, R'_2, R'_3 are pairwise disjoint, and $|R'_i \cap \mathcal{I}_{t_j}| = r_{t_j} + c_{i,t_j} \ge \tau - \eta = k$ for each $i \in \{1,2,3\}$, we can pick $S_j \subseteq (R'_1 \cup R'_2 \cup R'_3) \cap \mathcal{I}_{t_j}$ such $|R'_i \cap S_j| = k$ for each $i \in \{1,2,3\}$, and then $N_j \coloneqq (R'_1 \setminus S_j, R'_2 \setminus S_j, R'_3 \setminus S_j; W \cup S_j)$.

In any leaf node $N = (R_1, R_2, R_3; W)$, each $r_t < \tau$ for all but (at most) two intervals $t \in [1\,..\,T]$, and therefore for every $i \in \{1,2,3\}$,

$$|R'_i| = \sum_{t=1}^{T} |R'_i \cap \mathcal{I}_t| = \sum_{t=1}^{T} (r_t + c_{i,t}) \le T\eta + \sum_{t=1}^{T} r_t \le (T-2)\tau + 2\Big(\frac{r}{T}\Big) = T(\tau+\eta) + 2\Big(\frac{r}{T} - \tau\Big) =: \rho. \quad \text{(A.7)}$$

Inductive argument. We claim that inductively, every node $N = (R'_1, R'_2, R'_3; W)$ of height h satisfies the following:

> For every $w \in \operatorname{supp}(\mathfrak{X}[W])$, the complex $\mathfrak{X}_w[R'_1, R'_2, R'_3]$ has 1-triword expansion at least $(K_{\text{pin}}\beta^*)^{3\rho+h}$ over Γ.

Noting that the root height is at most r/k (since at a leaf, the sets must be nonempty, they start at size r, and decrease by k at each iteration), this gives the desired bound.

Consider a node $N = (R'_1, R'_2, R'_3; W)$. If N is a leaf node, then we apply the lopsided induction lemma (Lemma A.4). Indeed, for every $i_1 \in R'_1, i_2 \in R'_2, i_3 \in R'_3$, letting i'_1, i'_2, i'_3 be the sorted order of i_1, i_2, i_3 (so that $\{i_1, i_2, i_3\} = \{i'_1, i'_2, i'_3\}$ and $i'_1 < i'_2 < i'_3$), and every $w \in \operatorname{supp}(\mathfrak{X}[W])$, $\{i'_1\} \prec \{i'_2\} \prec \{i'_3\}$, so the assumption gives that $\mathfrak{X}_w[\{i_1\},\{i_2\},\{i_3\}] = \mathfrak{X}_w[\{i'_1\},\{i'_2\},\{i'_3\}]$ has 1-triword expansion at least β^*. Hence by Lemma A.4, for every $w \in \operatorname{supp}(\mathfrak{X}_w)$, $\mathfrak{X}_w[R'_1, R'_2, R'_3]$ has 1-triword expansion at least $(K_{\text{pin}}\beta^*)^{3\rho}$ (using that $|R'_1| + |R'_2| + |R'_3| \le 3\rho$ by Equation (A.7)).

Otherwise, N is an internal node, say of height h. Its three children N_1, N_2, N_3 are all nodes of height (at most) $h-1$. Recall that each $N_j = (R'_1 \setminus S_j, R'_2 \setminus S_j, R'_3 \setminus S_j; W \cup S_j)$. By the inductive hypothesis, for each $j \in \{1,2,3\}$, $(w, s_j) \in \operatorname{supp}(\mathfrak{X}[W, S_j])$, the peripheral complex

$$\mathfrak{T}_{s_j} \coloneqq \mathfrak{X}_{w,s_j}[R'_1 \setminus S_j, R'_2 \setminus S_j, R'_3 \setminus S_j]$$

has 1-triword expansion at least $(K_{\text{pin}}\beta^*)^{3\rho+(h-1)}$. Further, $S_1 \prec S_2 \prec S_3$ by construction (since each $S_j \subseteq \mathcal{I}_{t_j}$); hence, by Condition 2, the central complex $\mathfrak{X}_w[S_1, S_2, S_3]$ has 1-triword expansion at least β^*. Together, Lemma A.1 therefore gives that $\mathfrak{X}_w[R'_1, R'_2, R'_3]$ has 1-triword expansion at least β coboundary expander for

$$\beta \coloneqq K_{\text{pin}}\beta^* \cdot (K_{\text{pin}}\beta^*)^{3\rho+(h-1)} = (K_{\text{pin}}\beta^*)^{3\rho+h},$$

as desired. $\square$

A.3 Diameter bound from pinned diameter bound

We now prove the following analogue of Lemma A.1 for the case of "rank-two" induced complexes (i.e., bipartite graphs). This lemma seems to be used implicitly in the proof of [BLM24, Lemma 4.19]; we make it explicit here.

Lemma A.8. *There exist absolute constants $K > 1, \epsilon_0 > 0$ such that the following holds. Let $r \in \mathbb{N}$ and $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex. Let $R_1, R_2 \sqsubset \mathcal{I}$ and $S_1, S_2 \sqsubset R_1 \cup R_2$. Suppose that $\mathfrak{X}$ satisfies the following conditions:*

1. "Central" diameter bound:
$\mathfrak{A}^* \coloneqq \mathfrak{X}[S_1, S_2]$ *has* $\operatorname{diam}(\mathfrak{A}^*) \le C_0^*$.

2. "Peripheral" diameter bound:
For each $j \in [2]$ *and* $s_j \in \operatorname{supp}(\mathfrak{X}[S_j])$, $\mathfrak{A}_{s_j} \coloneqq \mathfrak{X}_{s_j}[R_1 \setminus S_j, R_2 \setminus S_j]$ *has* $\operatorname{diam}(\mathfrak{A}_{s_j}) \le C_0$.

Then $\mathfrak{A} \coloneqq \mathfrak{X}[R_1, R_2]$ *has* $\operatorname{diam}(\mathfrak{A}) \le 6C_0^* C_0$.

Proof. Fix any $r_1 \in \operatorname{supp}(\mathfrak{X}[R_1])$ and $r_2 \in \operatorname{supp}(\mathfrak{X}[R_2])$; we show that their distance in $\mathfrak{A}$ is at most $3C_0^* C_0$. This gives the desired bound by the triangle inequality.

Pick an arbitrary S_1-word $s_1 \in \operatorname{supp}(\mathfrak{X}_{r_1}[S_1])$ and S_2-word $s_2 \in \operatorname{supp}(\mathfrak{X}_{r_2}[S_2])$. Since s_1 is an S_1-word and s_2 is an S_2-word, there exists a path between s_1 and s_2 them of alternating S_2- and S_1-words of length $k \le C_0^*$. That is, there exists, for each $i \in [k]$, an S_1-word $s_1^{(i)} \in \operatorname{supp}(\mathfrak{X}[S_1])$ and S_2-word $s_2^{(i)} \in \operatorname{supp}(\mathfrak{X}[S_2])$, such that: $s_1^{(1)} = s_1$, $s_2^{(k)} = s_2$, $(s_1^{(i)}, s_2^{(i)}) \in \operatorname{supp}(\mathfrak{X}[S_1, S_2])$ is an (S_1, S_2)-biword (for each $i \in [k]$), and $(s_1^{(i+1)}, s_2^{(i)}) \in \operatorname{supp}(\mathfrak{X}[S_1, S_2])$ is an (S_1, S_2)-biword (for each $i \in [k-1]$).

Next, we "lift" this path of S_1- and S_2-words to a path of R_1- and R_2-words in $\mathfrak{A}$. For each $i \in [k]$, pick an (R_1, R_2)-biword $(u_1^{(i)}, u_2^{(i)}) \in \operatorname{supp}(\mathfrak{X}_{s_1^{(i)}, s_2^{(i)}}[R_1, R_2])$ consistent with $s_1^{(i)}$ and $s_2^{(i)}$. Similarly, for $i \in [k-1]$, pick an (R_1, R_2)-biword $(v_1^{(i+1)}, v_2^{(i)}) \in \operatorname{supp}(\mathfrak{X}_{s_1^{(i+1)}, s_2^{(i)}}[R_1, R_2])$ consistent with $s_1^{(i+1)}$ and $s_2^{(i)}$. We also define $v_1^{(1)} \coloneqq r_1$ and $v_2^{(k)} \coloneqq r_2$.

Finally, we claim that: For every $j \in [2]$ and $i \in [k]$, there is a path of length C_0 in $\mathfrak{A}$ between $u_j^{(i)}$ and $v_j^{(i)}$. Note that these are both R_j-words consistent with the S_j-word $s_j^{(i)}$. (I.e., both are in $\operatorname{supp}(\mathfrak{X}_{s_j^{(i)}}[R_j])$.) Incorporating the (R_1, R_2)-biwords $(u_1^{(i)}, u_2^{(i)})$ and $(v_1^{(i+1)}, v_2^{(i)})$ gives an overall path

$$r_1 = v_1^{(1)} \rightsquigarrow u_1^{(1)} \to u_2^{(1)} \rightsquigarrow v_2^{(1)} \to v_1^{(2)} \rightsquigarrow u_1^{(2)} \to v_1^{(2)} \to \cdots \to u_2^{(k)} \rightsquigarrow v_2^{(k)} = r_2.$$

The total length of this path is

$$C_0(2k+1) + (2k-1) \le C_0(2C_0^* + 1) + (2C_0^* - 1) \le 3C_0^* C_0.$$

Now, we prove the claim for the R_j-words $u_j^{(i)}$ and $v_j^{(i)}$. Recalling that $u_j^{(i)}, v_j^{(i)} \in \operatorname{supp}(\mathfrak{X}_{s_j^{(i)}}[R_j])$, we recall that there is a path of length at most C_0 between the $(R_j \setminus S_j)$-words $u_j^{(i)}|_{R_j \setminus S_j}$ and $v_j^{(i)}|_{S_j \setminus S_j}$ in the link $\mathfrak{X}_{s_1}[R_1 \setminus S_1, R_2 \setminus S_2]$. This path lifts to a path of length C_0 between $u_j^{(i)}$ and $v_j^{(i)}$ by appending s_j to every word. $\square$

A.4 Coboundary expansion of ordered restrictions

The following lemma due to Bafna, Lifshitz, and Minzer [BLM24] allows us to prove triword expansion in symmetric complexes which productize:

Lemma A.9 ([BLM24, Lem. 4.18])**.** *There exists an absolute constant* $K > 0$ *such that the following holds. Let* $n \in \mathbb{N}$ *and* $\mathfrak{X}$ *an* $\mathcal{I}$*-indexed simplicial complex. If* $\mathfrak{X}$ *is a strongly symmetric clique complex, then for every* $S_1, S_2, S_3 \sqsubset \mathcal{I}$ *such that* $S_1 \perp_{\mathfrak{X}} S_j$ *for both* $j \in \{2,3\}$, $\mathfrak{X}[S_1, S_2, S_3]$ *has* 1*-triword expansion at least* $K/\operatorname{diam}(\mathfrak{X}[S_2, S_3])$ *over every coefficient group* Γ.

(This lemma is very similar to [DDL24, Claim 2.11]. The hypothesis that $\mathfrak{X}$ is a clique complex appears to be necessary, but is omitted from both [BLM24, Lem. 4.18] and [DDL24, Claim 2.11].)

Proposition A.10. *Let* $\mathfrak{X}$ *be an* $\mathcal{I}$*-indexed simplicial complex. If* $S_1, S_2 \sqsubset \mathcal{I}$ *such that* $S_1 \perp_{\mathfrak{X}} S_2$, *then* $\operatorname{diam}(\mathfrak{X}[S_1, S_2]) \le 2$.

Proof. We claim that for every S_1-word $s_1 \in \operatorname{supp}(\mathfrak{X}[S_1])$ and S_2-word $s_2 \in \operatorname{supp}(\mathfrak{X}[S_2])$, (s_1, s_2) is a valid (S_1, S_2)-biword, i.e., $(s_1, s_2) \in \operatorname{supp}(\mathfrak{X}[S_1, S_2])$. Thus, there is a path of length 1 between s_1 and s_2, and therefore a path of length 2 between every pair of words in $\mathfrak{X}[S_1, S_2]$. The claim follows immediately from productization: We have

$$\Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2)\sim\mathfrak{X}[S_1,S_2]}[(\boldsymbol{s}_1, \boldsymbol{s}_2) = (s_1, s_2)] = \Pr_{\boldsymbol{s}_1\sim\mathfrak{X}[S_1]}[\boldsymbol{s}_1 = s_1] \cdot \Pr_{\boldsymbol{s}_2\sim\mathfrak{X}[S_2]}[\boldsymbol{s}_2 = s_2] > 0 \cdot 0 = 0,$$

where the inequality uses that both s_1 and s_2 are valid words. □

We also have the following trivial upper bound on the diameter in productizing complexes:

Proposition A.11 (Diameter bound for separated-and-ordered restrictions)**.** *There exists a universal constant $K > 0$ such that the following holds. Let $n \in \mathbb{N}$ and $\mathfrak{X}$ be an $[1\,..\,n]$-indexed simplicial complex. Let σ be a parameter. Suppose that $\mathfrak{X}$ satisfies the following conditions:*

1. Productization:
   $\mathfrak{X}$ is productizing on the line (in the sense of Definition 3.28).
2. Diameter bound for separated singleton restrictions:
   For every $W \subset [1\,..\,n]$ and $\ell_1 < \ell_2 \in [1\,..\,n]$ such that $\{\ell_1, \ell_2\}$ is σ-separated and $W \cap [\ell_1\,..\,\ell_2] = \emptyset$, for every $w \in \mathrm{supp}(\mathfrak{X}[W])$, $\mathrm{diam}(\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}]) \le C_0$.

Then for every $W \subset [1\,..\,n]$, $S_1 \prec S_2 \sqsubset [1\,..\,n] \setminus W$ such that $S_1 \cup S_2$ is σ-separated, and $w \in \mathrm{supp}(\mathfrak{X}[W])$, $\mathrm{diam}(\mathfrak{X}_w[S_1, S_2]) \le KC_0$.

Using these, we can prove the following:

Proof. Let $\ell_1 := \max S_1$ and $\ell_2 := \min S_2$. We apply Lemma A.8 with $R_1 := S_1$, $R_2 := S_2$, $S_1 := \{\ell_1\}$, and $S_2 := \{\ell_2\}$. Hence, we need to bound the "central" diameter $\mathrm{diam}(\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}])$ and the "peripheral" diameters $\mathrm{diam}(\mathfrak{X}_{w,s_1}[S_1 \setminus \{\ell_1\}, S_2])$ and $\mathrm{diam}(\mathfrak{X}_{w,s_2}[S_1, S_2 \setminus \{\ell_2\}])$ (where $s_1 \in \mathfrak{X}_w[\ell_1]$ and $s_2 \in \mathfrak{X}_w[\ell_2]$). The former is bounded by Condition 2 and the latter by Proposition A.10 and the PoL assumption. □

Given this, we now state a lemma which is extracted from the proofs of the "extended base cases" in [BLM24, §4.3].

Lemma A.12 (Coboundary expansion for separated-and-ordered restrictions, due to [BLM24])**.** *There exist universal constants $\epsilon_0, K > 0$ such that the following holds. Let $r \le n \in \mathbb{N}$. Let $\mathfrak{X}$ be an $[1\,..\,n]$-indexed simplicial complex and Γ a coefficient group. Let σ be a parameter. Let $\beta^* < 1$ and $C_0 \in \mathbb{N}$. Suppose that $\mathfrak{X}$ satisfies the following conditions:*

1. Spectral expansion:
   $\mathfrak{X}$ is $\epsilon_0 r^{-2}$-product.
2. Strong symmetry:
   $\mathfrak{X}$ is strongly symmetric.
3. Productization:
   $\mathfrak{X}$ is productizing on the line (in the sense of Definition 3.28).
4. Coboundary expansion of separated singleton restrictions:
   For every $W \subset [1\,..\,n]$ and $\ell_1 < \ell_2 < \ell_3 \in [1\,..\,n]$ such that $\{\ell_1, \ell_2, \ell_3\}$ is σ-separated and $W \cap [\ell_1\,..\,\ell_3] = \emptyset$, for every $w \in \mathrm{supp}(\mathfrak{X}[W])$, the complex $\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}, \{\ell_3\}]$ has 1-triword expansion at least β^* over Γ.
5. Diameter bound for separated singleton restrictions:
   For every $W \subset [1\,..\,n]$ and $\ell_1 < \ell_2 \in [1\,..\,n]$ such that $\{\ell_1, \ell_2\}$ is σ-separated and $W \cap [\ell_1\,..\,\ell_2] = \emptyset$, for every $w \in \mathrm{supp}(\mathfrak{X}[W])$, $\mathrm{diam}(\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}]) \le C_0$.

Then for every $W \subset [1\,..\,n]$, $S_1 \prec S_2 \prec S_3 \sqsubset_r [1\,..\,n] \setminus W$ such that $S_1 \cup S_2 \cup S_3$ is σ-separated, and $w \in \mathrm{supp}(\mathfrak{X}[W])$, the complex $\mathfrak{X}_w[S_1, S_2, S_3]$ has 1-triword expansion at least $K\beta^*/C_0$ over Γ.

Proof. We consider cases based on $r = |S_1| = |S_2| = |S_3|$. Fix W, w, S_1, S_2, and S_3; we bound the triword expansion of $\mathfrak{T} := \mathfrak{X}_w[S_1, S_2, S_3]$ over a fixed Γ.

Case 1: $r = 1$**.** In this case, $S_1 = \{\ell_1\}$, $S_2 = \{\ell_2\}$, and $S_3 = \{\ell_3\}$, with $\ell_1 < \ell_2 < \ell_3$, and $\{\ell_1, \ell_2, \ell_3\}$ is σ-separated by construction.

Subcase 1a: $W \cap [\ell_1 \mathinner{..} \ell_3] = \emptyset$**.** In this subcase, we directly use the assumed Condition 4 on $\mathfrak{T}$, which has 1-triword expansion at least β^* by assumption.

Subcase 1b: $W \cap [\ell_1 \mathinner{..} \ell_3] \neq \emptyset$**.** Let $\ell \in W \cap [\ell_1 \mathinner{..} \ell_3]$. WLOG, $\ell < \ell_2$. By the PoL assumption (Condition 3), $\{\ell_1\} \perp_{\mathfrak{X}_w} \{\ell_j\}$ for both $j \in \{2,3\}$. So by Lemma A.9, $\mathfrak{T}$ has 1-triword expansion at least $K \cdot \operatorname{diam}(\mathfrak{X}[\{\ell_2\}, \{\ell_3\}])$. By Condition 5, and since $\{\ell_2\} \prec \{\ell_3\}$ and $\{\ell_2, \ell_3\}$ is σ-separated, $\operatorname{diam}(\mathfrak{X}[\{\ell_2\}, \{\ell_3\}]) \leq C_0$. Hence $\mathfrak{T}$ has 1-triword expansion at least K/C_0.

Case 2: $r > 2$**.** Let $\ell_1 \coloneqq \max S_1, \ell_2 \coloneqq \min S_2, \ell_3 \coloneqq \min S_3$. We apply Lemma A.1 to $\mathfrak{T}$ with $R_1 = S_1, R_2 = S_2, R_3 = S_3, S_1 = \{\ell_1\}, S_2 = \{\ell_2\}, S_3 = \{\ell_3\}$. By the first case ($r = 1$), the central triangle complex $\mathfrak{T}^* \coloneqq \mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}, \{\ell_3\}]$ has 1-triword expansion at least β^*. On the other hand, consider the peripheral complexes

$$\begin{aligned}
\mathfrak{T}_{s_1} &\coloneqq \mathfrak{X}_{w,s_1}[S_1 \setminus \{\ell_1\}, S_2, S_3], \\
\mathfrak{T}_{s_2} &\coloneqq \mathfrak{X}_{w,s_2}[S_1, S_2 \setminus \{\ell_2\}, S_3], \\
\mathfrak{T}_{s_3} &\coloneqq \mathfrak{X}_{w,s_3}[S_1, S_2, S_3 \setminus \{\ell_3\}]
\end{aligned}$$

(for $s_j \in \operatorname{supp}(\mathfrak{X}_w[\{\ell_j\}])$). We have by the PoL assumption (Condition 3):

- For every $j \in \{2,3\}$, $S_1 \setminus \{\ell_1\} \prec S_j$. Further, since $\min S_1 \setminus \{\ell_1\} = \min S_1$, we have $\ell_1 \in [\min S_1 \setminus \{\ell_1\} \mathinner{..} \max S_2] \subseteq [\min S_1 \setminus \{\ell_1\} \mathinner{..} \max S_3]$. Hence for every $s_1 \in \operatorname{supp}(\mathfrak{X}_w[S_1])$ and $j \in \{2,3\}$, $S_1 \setminus \{\ell_1\} \perp_{\mathfrak{T}_{s_1}} S_j$.
- Dually, for every $j \in \{1,2\}$, $S_j \prec S_3 \setminus \{\ell_3\}$. Further, since $\max S_3 \setminus \{\ell_3\} = \max S_3$, we have $\ell_3 \in [\min S_2 \mathinner{..} \max S_3 \setminus \{\ell_3\}] \subseteq [\min S_1 \mathinner{..} \max S_3 \setminus \{\ell_3\}]$. Hence for every $s_3 \in \operatorname{supp}(\mathfrak{X}_w[S_3])$ and $j \in \{1,2\}$, $S_j \perp_{\mathfrak{T}_{s_3}} S_3 \setminus \{\ell_3\}$.
- $S_1 \prec S_2 \setminus \{\ell_2\} \prec S_3$. Further, since $\max S_2 \setminus \{\ell_2\} = \max S_2$, we have $\ell_2 \in [\min S_1 \mathinner{..} \max S_2 \setminus \{\ell_2\}] \subseteq [\min S_1 \mathinner{..} \max S_3]$. Hence for every $s_2 \in \operatorname{supp}(\mathfrak{X}_w[S_2])$, $S_1 \perp_{\mathfrak{T}_{s_2}} S_3$ and $S_1 \perp_{\mathfrak{T}_{s_2}} S_2 \setminus \{\ell_2\}$.

Note that the S_2 case had to be handled slightly differently because of the asymmetry in the choice of ℓ_2.[18] In all cases, the relevant restrictions are ordered and separated, so we can apply Proposition A.11 (along with our Conditions 3 and 5) and deduce that all the peripheral complexes have 1-triword expansion at least K/C_0. $\square$

A.5 Coboundary expansion of random restrictions

Lemma A.13. *Let $r \leq n \in \mathbb{N}$, $T, \eta \geq 1 \in \mathbb{N}$, $W \subset_{\leq r} [1 \mathinner{..} n]$. Suppose that $r^2/\eta \leq \frac{1}{4}n$ and $Tr \leq \frac{1}{2}n$. Sample $\boldsymbol{R} \subset_r [1 \mathinner{..} n] \setminus W$ uniformly. Then:*

$$\Pr[\boldsymbol{R} \textit{ is } (T,\eta)\textit{-pseudorandomly spread}] \geq 1 - 2T \exp\left(\frac{-\eta^2 T}{24r}\right).$$

Proof. Generate $\boldsymbol{R}$ via the following process: Sample $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_r$ uniformly from $[1 \mathinner{..} n] \setminus W$ without replacement and set $\boldsymbol{R} \coloneqq \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_r\}$. For $\boldsymbol{r}_t \coloneqq |\boldsymbol{R} \cap \mathcal{I}_t|$, we therefore have $\boldsymbol{r}_t = \sum_{i=1}^r 1_{\boldsymbol{x}_i \in \mathcal{I}_t}$. Marginally,

$$p \coloneqq \Pr[\boldsymbol{x}_i \in \mathcal{I}_t] = \frac{|\mathcal{I}_t \setminus W|}{n - |W|}.$$

[18] When $t \geq 3$, this issue can be resolved more symmetrically: We pick $t \in S_2 \setminus \{\min S_2, \max S_2\}$. Then $\ell_2 \in [\min S_1 \mathinner{..} \max S_2 \setminus \{\ell_2\}] \cap [\min S_2 \setminus \{\ell_2\} \mathinner{..} \max S_3]$, so for every $j \in \{1,3\}$ and $s_2 \in \operatorname{supp}(\mathfrak{X}_w[S_2])$, $S_j \perp_{\mathfrak{T}_{s_2}} S_2$.

Using the trivial bounds $\frac{n}{T} - |W| = |\mathcal{I}_t| - |W| \le |\mathcal{I}_t \setminus W| \le |\mathcal{I}_t| = \frac{n}{T}$, we get:

$$\frac{1}{T} - \frac{2r}{n} \le \frac{\frac{n}{T} - r}{n - r} \le \frac{\frac{n}{T} - |W|}{n - |W|} \le p \le \frac{\frac{n}{T}}{n - r} \le \frac{1}{T} + \frac{2r}{n}$$

and consequently $|p - \frac{1}{T}| \le \frac{2r}{n}$ and so $|rp - \frac{r}{T}| \le \frac{2r^2}{n} \le \frac{1}{2}\eta$. We also use $p \le \frac{1}{T} + \frac{2r}{n} \le \frac{1}{T} + \frac{1}{T} = \frac{2}{T}$.

The Chernoff bound for sampling without replacement [Chv79] gives

$$\Pr[|\boldsymbol{r}_t - rp| \ge \tfrac{1}{2}\eta] \le 2\exp\left(\frac{-(\frac{1}{2}\eta)^2}{3rp}\right) = 2\exp\left(\frac{-(\frac{1}{2}\eta)^2}{3rp}\right) \le 2\exp\left(\frac{-\eta^2 T}{24r}\right).$$

Finally, take a union bound over all T blocks. $\square$

Lemma A.14. *Let $r \le \frac{1}{2}n \in \mathbb{N}$ and $W \subset_{\le r} [1\,..\,n]$, Sample $\boldsymbol{R} \subset_r [1\,..\,n] \setminus W$ uniformly. For $\sigma \in \mathbb{N}$, we have:*

$$\Pr[\boldsymbol{R} \text{ is } \sigma\text{-separated}] \ge 1 - \sigma \cdot \frac{2r^2}{n}.$$

Proof. We first claim the following:

$$\Pr_{\boldsymbol{x}_1, \boldsymbol{x}_2 \sim \mathbf{Unif}([1\,..\,n] \setminus W)}[|\boldsymbol{x}_1 - \boldsymbol{x}_2| \le \sigma \mid \boldsymbol{x}_1 \ne \boldsymbol{x}_2] \le \frac{2\sigma}{\frac{1}{2}n - 1}. \tag{A.15}$$

This implies the lemma by a union bound if we generate $\boldsymbol{R}$ by sampling $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_r$ uniformly from $[1\,..\,n]$ without replacement.

It remains to prove Equation (A.15). Let $m := |[1\,..\,n] \setminus W| = n - |W| \ge n - r \ge \frac{1}{2}n$. First, we argue that the probability is maximized when $W = (n - m\,..\,n]$; that is,

$$\Pr_{\boldsymbol{x}_1, \boldsymbol{x}_2 \sim \mathbf{Unif}([1\,..\,n] \setminus W)}[|\boldsymbol{x}_1 - \boldsymbol{x}_2| \le \sigma \mid \boldsymbol{x}_1 \ne \boldsymbol{x}_2] \le \Pr_{\boldsymbol{i}, \boldsymbol{j} \sim \mathbf{Unif}([m])}[|\boldsymbol{i} - \boldsymbol{j}| \le \sigma \mid \boldsymbol{i} \ne \boldsymbol{j}].$$

Indeed, enumerate the elements of $[1\,..\,n] \setminus W$ as $\ell_1 < \cdots < \ell_m$. For $i, j \in [m]$, we have $|\ell_j - \ell_i| \ge |j - i|$ (since if e.g. $j \ge i$ then $\ell_j > \cdots > \ell_i$). This implies this inequality by coupling, since the distribution of $(\ell_{\boldsymbol{i}}, \ell_{\boldsymbol{j}})$ for $\boldsymbol{i}, \boldsymbol{j} \sim \mathbf{Unif}([m]) \mid \boldsymbol{i} \ne \boldsymbol{j}$ is the same as the distribution of $(\boldsymbol{x}_1, \boldsymbol{x}_2)$ for $\boldsymbol{x}_1, \boldsymbol{x}_2 \sim \mathbf{Unif}([m]) \mid \boldsymbol{x}_1 \ne \boldsymbol{x}_2$.

Finally, we show that

$$\Pr_{\boldsymbol{i}, \boldsymbol{j} \sim \mathbf{Unif}([m])}[|\boldsymbol{i} - \boldsymbol{j}| \le \sigma \mid \boldsymbol{i} \ne \boldsymbol{j}] \le \frac{2\sigma}{m - 1}.$$

Indeed, for fixed $i \in [m]$, we have $\Pr_{\boldsymbol{j} \sim \mathbf{Unif}([m] \setminus \{i\})}[|i - \boldsymbol{j}| \le \sigma] \le \frac{2\sigma}{m-1}$. (The factor of 2 is necessary when $\sigma \le i \le m - \sigma$.) $\square$

This lets us give:

Lemma A.16 (Coboundary expansion of typical restrictions)**.** *There exist absolute constants $K > 1$ and $\epsilon_0 > 0$ such that the following holds. Let $r \le n \in \mathbb{N}$, $\mathfrak{X}$ be a $[1\,..\,n]$-indexed simplicial complex and Γ a coefficient group. Let $\sigma, \eta, k, T \in \mathbb{N}$ be parameters. Suppose that $\mathfrak{X}$ satisfies the following conditions:*

1. Spectral expansion:
   $\mathfrak{X}$ is $\epsilon_0 r^{-2}$-product.

2. Productization:
   $\mathfrak{X}$ productizes on the line, in the sense of Definition 3.28.

3. Symmetry:
   $\mathfrak{X}$ is strongly symmetric, in the sense of Definition 3.22.

4. Coboundary expansion of separated singleton restrictions:
   For every $W \subset [1\,..\,n]$ and $\ell_1 < \ell_2 < \ell_3 \in [1\,..\,n]$ such that $W \cap [\ell_1\,..\,\ell_3] = \emptyset$ and $\{\ell_1, \ell_2, \ell_3\}$ is σ-separated, for every $w \in \operatorname{supp}(\mathfrak{X}[W])$, the complex $\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}, \{\ell_3\}]$ has 1-triword expansion at least β^* over Γ.

5. Diameter of separated singleton restrictions:
For every $W \subset [1\,..\,n]$ and $\ell_1 < \ell_2 \in [1\,..\,n]$ such that $W \cap [\ell_1\,..\,\ell_2] = \emptyset$ and $\{\ell_1, \ell_2\}$ is σ-separated, for every $w \in \mathrm{supp}(\mathfrak{X}[W])$, the complex $\mathfrak{X}_w[\{\ell_1\}, \{\ell_2\}]$ has diameter at most C_0.

Then for every $W \subset_{\leq r} [1\,..\,n]$ and $w \in \mathrm{supp}(\mathfrak{X}[W])$,

$$\Pr_{\boldsymbol{R}_1, \boldsymbol{R}_2, \boldsymbol{R}_3 \sqsubset_r [1\,..\,n] \setminus W}[\mathfrak{X}_w[\boldsymbol{R}_1, \boldsymbol{R}_2, \boldsymbol{R}_3] \text{ has 1-triword expansion at least } K\beta' \text{ over } \Gamma]$$
$$\geq 1 - 3\left(2T \exp\left(\frac{-\eta^2 T}{24r}\right) + \sigma \cdot \frac{2r^2}{n}\right),$$

where $\beta' \coloneqq (K\beta^*/C_0)^{3(T(k+2\eta)+2(r/T))+r/k}$.

Proof. Let $\mathcal{S} \coloneqq \{(R_1, R_2, R_3) : R_1, R_2, R_3 \sqsubset_r [1\,..\,n] \setminus W\}$. We claim that for fixed $(R_1, R_2, R_3) \in \mathcal{S}$ such that (i) R_i is (T, η)-pseudorandomly spread for each $i \in \{1, 2, 3\}$ and (ii) $R_1 \cup R_2 \cup R_3$ is ξn-separated, $\mathfrak{X}[R_1, R_2, R_3]$ is a triword expander with the desired parameters. This implies the lemma when we sample $(\boldsymbol{R}_1, \boldsymbol{R}_2, \boldsymbol{R}_3) \sim \mathbf{Unif}(\mathcal{S})$, since marginally, each $\boldsymbol{R}_i$ is a uniformly random r-set in $[1\,..\,n] \setminus W$; we can apply Lemmas A.13 and A.14 and take a union bound.

To prove the claim, we apply Lemma A.5; we need to check its three conditions. Spectral expansion holds by our assumption, and R_1, R_2, R_3 are pseudorandomly spread by construction. Thus, it suffices to therefore check Condition 2 of Lemma A.5 (“ordered triword expansion”): That is, for every $W \subseteq R_1 \cup R_2 \cup R_3$, $w \in \mathrm{supp}(\mathfrak{X}[W])$, and $S_1 \prec S_2 \prec S_3 \sqsubset_k R_1 \cup R_2 \cup R_3 \setminus W$, $\mathfrak{X}_w[S_1, S_2, S_3]$ is a triword expander. For this, we apply Lemma A.12; we have all the necessary conditions, and get triword expansion at least $K\beta^*/C_0$. ☐

A.6 r-triword expansion from typical triword expansion

Now, we state and prove the following theorem, which is based on techniques from [DD24c, §8]:

Theorem A.17. *There exists absolute constants $K, \epsilon_0 > 0$ such that the following holds. Let $r, n \in \mathbb{N}$ with $r < \frac{1}{10}\sqrt{n} \in \mathbb{N}$. Let $\mathcal{I}$ be a set of size n, $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex which is $\epsilon_0 r^{-2}$-product, and $p, \beta < 1$. Let Γ be any coefficient group and suppose that*

$$\Pr_{\boldsymbol{R}_1, \boldsymbol{R}_2, \boldsymbol{R}_3 \sqsubset_r \mathcal{I}}[\mathfrak{X}[\boldsymbol{R}_1, \boldsymbol{R}_2, \boldsymbol{R}_3] \text{ has 1-triword expansion at least } \beta \text{ over } \Gamma] \geq p.$$

Then $\mathfrak{X}$ has r-triword expansion at least $K\beta p^2$ over Γ.

In order to prove this, we use the following lemma:

Lemma A.18. *There exists an absolute constant $\epsilon_0 > 0$ such that the following holds. Let $r, \ell, n \in \mathbb{N}$ with $\ell \geq 3r$ and $n \geq (\ell + 2)k$. Let $\mathcal{J}$ be a set of size n, let $\mathfrak{Y}$ be a $\mathcal{J}$-indexed simplicial complex which is $\epsilon_0 r^{-2}$-product. Let $\boldsymbol{\mathcal{R}}$ be a set of ℓ uniformly random pairwise disjoint r-sets in $\mathcal{J}$. Let Γ any coefficient group and suppose that:*

$$\Pr_{\boldsymbol{\mathcal{R}}}[\mathfrak{Y}[\boldsymbol{\mathcal{R}}] \text{ has 1-triword expansion at least } \beta \text{ over } \Gamma] \geq p.$$

Then $\mathfrak{Y}$ has 1-triword expansion $\Omega(\beta p)$ over Γ.

This lemma is a variant of lemmas in [DD24b, §4] and [BLM24, §5]. As we have not been able to locate a self-contained proof, we give a proof in Appendix E below.

The $\ell = 1, r = 3$ case of this lemma is the following, which can also be viewed as a direct consequence of [DD24b, Lemma 4.1]:

Corollary A.19. *There exists an absolute constant $\epsilon_0 > 0$ such that the following holds. Let $\mathfrak{Y}$ be a $\mathcal{J}$-indexed simplicial complex which is ϵ_0-product. Let $\boldsymbol{j}_1, \boldsymbol{j}_2, \boldsymbol{j}_3$ be three uniformly random distinct indices in $\mathcal{J}$. Let Γ be any coefficient group and suppose that:*

$$\Pr_{\boldsymbol{j}_1, \boldsymbol{j}_2, \boldsymbol{j}_3}[\mathfrak{Y}[\boldsymbol{j}_1, \boldsymbol{j}_2, \boldsymbol{j}_3] \text{ has 1-triword expansion at least } \beta \text{ over } \Gamma] \geq p.$$

Then $\mathfrak{Y}$ has 1-triword expansion $\Omega(\beta p)$ over Γ.

We also use:

Proposition A.20 ([GLL22, Lemma 3.5])**.** *Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex which is ϵ-product. Then for every $R_1, R_2 \sqsubset \mathcal{I}$, the edge complex $\mathfrak{X}[R_1, R_2]$ is an $\epsilon|R_1||R_2|$-bipartite expander.*

Corollary A.21. *Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex which is ϵ-product. Let $\mathcal{R}$ be a set of pairwise disjoint k-sets in $\mathcal{I}$. Then the $\mathcal{R}$-indexed simplicial complex $\mathfrak{X}[\mathcal{R}]$ is ϵr^2-product.*

Proof. We need to show that for every $R_1 \neq R_2 \in \mathcal{R}$, $W \subseteq \mathcal{R} \setminus \{R_1, R_2\}$, and W-word w, $(\mathfrak{X}_w[\mathcal{R}])[R_1, R_2]$ is an ϵr^2-bipartite expander. But this graph is identical to the graph $\mathfrak{X}_w[R_1, R_2]$ (now viewing w as a word on $\bigcup \mathcal{R}$). Since $\mathfrak{X}_w$ is ϵ-product (by heritability on links), we can therefore apply Proposition A.20. □

Claim A.22. *Suppose $\mathfrak{X}$ is an $\mathcal{I}$-indexed simplicial complex satisfying the hypotheses of Theorem A.17. Let $\boldsymbol{\mathcal{R}}$ be a set of $3r$ random pairwise disjoint r-sets in $\mathcal{I}$. Then*

$$\Pr_{\boldsymbol{\mathcal{R}}}[\mathfrak{X}[\boldsymbol{\mathcal{R}}] \text{ has 1-triword expansion at least } \tfrac{1}{2}\beta p] \geq \tfrac{1}{2}p.$$

Proof. First, fix $\mathcal{R}$, a collection of $3r$ disjoint r-sets, and define

$$\gamma(\mathcal{R}) \coloneqq \Pr_{\boldsymbol{R}_1, \boldsymbol{R}_2, \boldsymbol{R}_3 \text{ distinct} \in \mathcal{R}}[\mathfrak{X}[\boldsymbol{R}_1, \boldsymbol{R}_2, \boldsymbol{R}_3] \text{ has 1-triword expansion at least } \beta \text{ over } \Gamma].$$

$\gamma(\mathcal{R})$ is a deterministic function of $\mathcal{R}$ and satisfies $0 \leq \gamma(\mathcal{R}) \leq 1$.

By Corollary A.21, $\mathfrak{X}[\mathcal{R}]$ is ϵ_0-product. We can therefore apply Corollary A.19 (with $\mathfrak{Y} \coloneqq \mathfrak{X}$, $r \coloneqq 1$, $\ell \coloneqq 3$) to deduce that $\mathfrak{X}[\mathcal{R}]$ has 1-triword expansion at least $\beta\gamma(\mathcal{R})$ over Γ.

Next, for random $\boldsymbol{\mathcal{R}}$, $\mathbb{E}_{\boldsymbol{\mathcal{R}}}\, \gamma(\boldsymbol{\mathcal{R}}) \geq p$ by assumption. (This is because the marginal distribution of three random distinct r-sets inside of $\boldsymbol{\mathcal{R}}$ is the same as the marginal distribution of three random pairwise disjoint r-sets.) Hence by Markov's inequality,

$$\Pr_{\boldsymbol{\mathcal{R}}}[\gamma(\boldsymbol{\mathcal{R}}) \leq \tfrac{1}{2}p] = \Pr_{\boldsymbol{\mathcal{R}}}[1 - \gamma(\boldsymbol{\mathcal{R}}) \geq 1 - \tfrac{1}{2}p] \leq \frac{1-p}{1 - \frac{1}{2}p} \leq 1 - \tfrac{1}{2}p$$

(just using that $p < 2$), and therefore $\Pr_{\boldsymbol{\mathcal{R}}}[\gamma(\boldsymbol{\mathcal{R}}) \geq \frac{1}{2}p] \geq \frac{1}{2}p$. □

Proof of Theorem A.17. We use Lemma A.18 again and set $\ell \coloneqq 3r$. Applying Claim A.22, the probability in the hypothesis of Lemma A.18 is at least $\frac{1}{4}p$ (using $\beta' \coloneqq \frac{1}{2}\beta p$). This gives a final coboundary expansion bound of $\frac{1}{8}\beta p^2$, as desired. □

Theorem 3.30 now follows immediately from Theorem A.17, Lemma A.16, and Remark A.6 (setting $\xi \coloneqq 100/r^2$).

B Sufficient conditions for direct-product testing

In this section, we discuss the proof of Theorem 3.21, restated here:

Theorem 3.21 (Sufficient conditions for direct-product testability, combining [DD24a; DD24b])**.** *Given $\delta > 0$, there is $m = \mathrm{poly}(1/\delta)$, such that if $k \geq \exp(\mathrm{poly}(1/\delta))$, $r \geq \mathrm{poly}(k)$, $n \geq \exp(\mathrm{poly}(r))$, and $\gamma < 1/n^2$, then the following holds:*

Suppose $\mathfrak{X}$ is a rank-n complex (not necessarily indexed) satisfying the following:

1. Spectral expansion:
   For any pinning w of size $n-2$, the graph $\mathfrak{X}_w$ has normalized 2nd eigenvalue at most γ.
2. Connectivity:
   For every S-word w with $|S| \leq r$, the graph $(V_r(\mathfrak{X}_w), E_r(\mathfrak{X}_w))$ is connected.
3. Clique complex:
   $\mathfrak{X}$ is a clique complex.

4. Local triword expansion:
For every 1*-word* w *("vertex"),* $\mathfrak{X}_w$ *has* r*-triword expansion at least* $\exp(-r^{.99})$ *over* $\mathfrak{S}(m_0)$ *for all* $1 < m_0 \le m$*.*

5. Trivial global cohomology:
$\mathfrak{X}$ *has nonzero* 1*-triword expansion (i.e., trivial* 1*-cohomology) over* $\mathfrak{S}(m_0)$ *for all* $1 < m_0 \le m$*.*

Then the agreement tester with distribution given by $\mathfrak{X}$*'s rank-*k *faces achieves* δ*-soundness for any alphabet* Σ*, in the sense of Definition 2.1.*

In particular, Theorem 3.21 follows from [DD24a, Theorem 3.1] (henceforth "3.1"), for the following reasons:

- We are specializing 3.1 to the "V-test" (this is the "$\mathcal{D}$" of 3.1; it is acceptable per [DD24a, Example 2.11]).
- Given $\mathfrak{X}$ as in our Theorem 3.21, we do not apply 3.1 to it directly. Rather, we apply 3.1 to the d-skeleton $\mathfrak{X}'$ of $\mathfrak{X}$, for a certain $d \ge \exp(\mathrm{poly}(r))$ (and then we require $n \ge 2d^2$, say). This small trick is for achieving two-sided spectral expansion; see below. Note that $\mathfrak{X}'$ is a rank-d clique complex, and it has Conditions 2, 4 and 5 if (and only if) $\mathfrak{X}$ has (since these properties only depend on the $4r$-skeleton).
- We should verify that Conditions 1, 2 and 4 imply that $\mathfrak{X}'$ is "(d, k, α)-suitable" in the sense of 3.1:
  - First, recalling that 3.1's parameter "α" is $\mathrm{poly}(1/\delta)$, by requiring $r \ge \mathrm{poly}(k)$ and then $d \ge \exp(\mathrm{poly}(r))$ large enough, we can let the "suitable" d_1 parameter equal r, while also arranging for the expression $\alpha\frac{d_1}{k}$ to be at least $r^{.991}$.
  - Next we will verify that $\mathfrak{X}'$ is a $\frac{1}{d^2}$-two-sided local spectral expander. This uses the two-sided "trickling down" results of [Opp18]. By [DDL24, Corollary 2.4] (or see, e.g., [HS24, §A] for a fuller treatment, where one uses the trivial bound $\eta = -1$ to control negative eigenvalues), we get that $\mathfrak{X}'$ is a $\max\{\frac{\gamma}{1-(n-1)\gamma}, \frac{1}{n-d+1}\}$-two-sided local expander. Using $\gamma < 1/n^2$ and recalling $n \ge 2d^2$, the expansion parameter is at most $\frac{1}{d^2}$ as needed.
  - $\mathfrak{X}'$ is required to be "r-well-connected", which is precisely the conjunction of Condition 2 and "simple-connectedness" of the r-triword complex $\mathfrak{X}'_w$ for every vertex w, which we discuss next.
  - This simple-connectedness is equivalent to *nonzero* r-triword expansion of $\mathfrak{X}'_w$ over every group, which is not quite implied by Condition 4.[19] However, in the one place where simple-connectedness is actually used (end of first paragraph of proof of "Claim 5.1" in [DD24a]) it is used to show that an ℓ-cover of $\mathfrak{X}'_w$, with $\ell \le \mathrm{poly}(1/\delta)$ must in fact be ℓ disjoint copies of $\mathfrak{X}'_w$. But Condition 4, in its weaker nonzero expansion form, is precisely equivalent to saying that the only m-covers of $\mathfrak{X}'_w$ are trivial ones, given by m disjoint copies of $\mathfrak{X}'_w$. (This will be discussed again in the final bullet-point.) Thus Condition 4 suffices, provided $m \ge \mathrm{poly}(1/\delta)$ is taken larger than ℓ.
  - Finally, we show that Condition 4 implies the last needed condition for "suitability", namely that $\mathfrak{X}'$ has r-triword "cosystolic" expansion $\exp(-r^{.991})$ over $\mathfrak{S}(m_0)$ for all $1 < m_0 \le m$. (In fact, 3.1 formally requires the cosystolic expansion over every group Γ. However, it was written this way just for brevity; when the condition is used in the proof — going from "Lemma 3.6" to "Corollary 3.7" — only $\Gamma = \mathfrak{S}(m)$ is needed. Cf. [BM24, Theorem 4.1].) We apply the local-to-global method of [KKL16; EK23; DD24b], specifically [DD24b, Theorem 1.2] (with its $k = 1$ and $d = 3$): by combining the r-triword ("coboundary") expansion of $\exp(-r^{.99})$ over $\mathfrak{S}(m_0)$ for each $\mathfrak{X}'_w$, with the one-sided spectral expansion of $\mathfrak{X}'$ with parameter $\frac{1}{d^2} \ll \exp(-r^{.99})$, we conclude that $\mathfrak{X}'$ has r-triword cosystolic expansion $\Omega(\exp(-r^{.99})) \ge \exp(-r^{.991})$, as needed.
- The last distinction is that 3.1 only concludes the agreement tester achieves "δ-cover soundness", rather than δ-soundness; i.e., rather than $F(\boldsymbol{S})$ being often close to $f^{\wedge k}(\boldsymbol{S})$ for some $f : \mathcal{U} \to \Sigma$, there is

[19] In fact, for the complex we study in this paper, $\mathfrak{X}'_w$ *will* have nonzero r-triword expansion over every group, rendering this discussion moot.

merely an explanatory f on the vertices of an m-*cover* $\mathfrak{Y}$ of $\mathfrak{X}'$ (See [DD24a, Definition 2.26] for the precise definition.) However, similar to the prior discussion of simple-connectedness, we have that Condition 5 is precisely equivalent to saying that the only m-covers of $\mathfrak{X}'$ are trivial ones (given by m disjoint copies of $\mathfrak{X}'$). In this case, it is easy to see (and the argument written in [DDL24, Part 2 of the proof of Theorem 4.11]) that the explanatory f on $\mathfrak{Y}$ must yield an equally explanatory f on $\mathfrak{X}'$, as desired for δ-soundness.

C The global Kaufman–Oppenheim complex

In this subsection, we describe a general way to construct Kaufman–Oppenheim-type complexes and some useful properties of them. This construction generalizes both the original complex defined in [KO18; KO23], and also its modification in [KOW25, §2.5].

C.1 Defining the group

For a field $\mathbb{F}$, indeterminate X, and $s \in \mathbb{N}$, define the quotient ring $\mathbb{F}^{(s)} := \mathbb{F}[X]/(X^s)$. We use $\deg(f)$ to refer to the degree of the (unique) minimal-degree representative of an element of $\mathbb{F}^{(s)}$.

Remark C.1. For concreteness, here is a direct definition of this ring, without using the quotient ring construction: $\mathbb{F}^{(s)}$ consists of the set of elements $\{f \in \mathbb{F}[X] : \deg(f) < s\}$ equipped with the standard addition operation and the "truncated" multiplication operation. That is, for $f = \sum_{i=0}^{s-1} f_i X^i$ and $g = \sum_{i=0}^{s-1} g_i X^i$, $fg := \sum_{i=0}^{s-1} \sum_{j=0}^{i} f_j g_{i-j} X^i$. Hence, $\deg(fg) = \min\{\deg(f) + \deg(g), s-1\}$. $\Diamond$

Definition C.2 (Special linear group)**.** Let $n, s \in \mathbb{N}$ and $\mathbb{F}$ be a field. $\mathsf{SL}_{n+1}(\mathbb{F}^{(s)})$ is the group of $(n+1) \times (n+1)$ matrices (indexed by $[0\,.\,.\,n]$) with entries in the ring $\mathbb{F}^{(s)}$ and determinant 1. For $i \neq j \in [0\,.\,.\,n]$ and $f \in \mathbb{F}^{(s)}$, we let $e_{i,j}(f) \in \mathsf{SL}_{n+1}(\mathbb{F}^{(s)})$ denote the matrix with 1's on the diagonal and f in the (i,j) entry. These matrices are known variously as *elementary matrices*, *transvections*, or *shear matrices*. $\Diamond$

Remark C.3. One can more generally define $\mathsf{SL}_{n+1}(R)$ when R is any (commutative, unital) ring. For general R, there is a technical distinction between the group $\mathsf{SL}_{n+1}(R)$ and its subgroup $\mathsf{EL}_{n+1}(R)$, which is the subgroup generated by the elementary matrices $e_{i,j}(r)$. Some sources, in particular [KO23], deal with the groups $\mathsf{EL}_{n+1}(R)$ instead of $\mathsf{SL}_{n+1}(R)$. However, $\mathsf{EL}_{n+1}(R)$ and $\mathsf{SL}_{n+1}(R)$ coincide (i.e., the elementary matrices generate $\mathsf{SL}_{n+1}(R)$) whenever $\mathbb{F}$ is a field and $R = \mathbb{F}^{(s)}$ (our case); indeed, they coincide under the (much) more general condition that R is a *semilocal ring* (see e.g. [HO89, Theorem 4.3.9]). Hence, we are justified in using these notations interchangeably in this appendix. $\Diamond$

Remark C.4. This is an instantiation of a more general construction in [KO23, §3.1]; specifically, we apply that construction with the ring $R = \mathbb{F}$, finitely generated algebra $\mathcal{R} = \mathbb{F}[X]$, and generating set $T = \{1, \ldots, X^\kappa\}$. We can therefore invoke some known properties of the constructions in [KO23] in the sequel. In particular, Theorem 3.56 follows immediately from [KO23, Theorem 3.5]. $\Diamond$

We use $\mu_{n+1}(\cdot) : \mathbb{Z} \to [0\,.\,.\,n]$ to denote the function reducing an integer to its residue in $[0\,.\,.\,n]$ modulo $n+1$.

Definition C.5 (Unipotent subgroups of special linear group)**.** For $n, s, \kappa \in \mathbb{N}$ and $\ell \in [0\,.\,.\,n]$, the subgroup $\mathsf{Unip}_{n,\ell}^{\kappa,(s)}(\mathbb{F}) < \mathsf{SL}_n(\mathbb{F}^{(s)})$ is the subgroup generated by elements of the form $e_{j,\mu_{n+1}(j+1)}(f)$ for $j \in [0\,.\,.\,n] \setminus \{\ell\}$ and $f \in \mathbb{F}^{(s)}$ with $\deg(f) \leq \kappa$. $\Diamond$

Remark C.6. Note that if $s \leq \kappa$ then the condition $\deg(f) \leq \kappa$ is superfluous. $\Diamond$

These subgroups form an $[0\,.\,.\,n]$-indexed subgroup family $(\mathsf{Unip}_{n,\ell}^{\kappa,(s)}(\mathbb{F}) < \mathsf{SL}_n(\mathbb{F}^{(s)}))_{\ell \in [0\,.\,.\,n]}$. This family has many nice properties; for instance, it has "dihedral symmetry" because there are automorphisms which cyclically shift and/or reverse the order of the subgroups (see [KO23, §3.2]).

Hence, as in Definition 3.31, for $W \subseteq [0\,.\,.\,n]$ we can consider the intersection subgroup

$$\mathsf{Unip}_{n,W}^{\kappa,(s)}(\mathbb{F}) := \bigcap_{\ell \in W} \mathsf{Unip}_{n,\ell}^{\kappa,(s)}(\mathbb{F}).$$

(Again, $\mathsf{Unip}_{n,\emptyset}^{\kappa,(s)}(\mathbb{F}) = \mathsf{SL}_n(\mathbb{F}^{(s)})$.)

We have the following, which essentially follows from [KOW25, Theorem 2.20]:

Proposition C.7. *Let $n, s, \kappa \in \mathbb{N}$ and $\mathbb{F}$ be a field. Suppose that $s \geq \kappa n$. For every $\ell \in [0..n]$, let $K_\ell := \mathsf{Unip}_{n,\ell}^{\kappa,(s)}(\mathbb{F})$. Then for every $\ell \in [0..n]$, the subgroup family $(K_\ell \cap K_a \leq K_\ell)_{a\in[0..n]\setminus\{\ell\}}$ is isomorphic to the subgroup family $(\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;b) \leq \mathrm{Unip}_{\mathbb{F}}^{\kappa}(n))_{b\in[1..n]}$ defined in Definition 3.50.*

C.2 Defining the complex

Given the definitions of groups in the previous subsection, it is now possible to define the *global* Kaufman–Oppenheim complex:

Definition C.8 (The global Kaufman–Oppenheim complex)**.** Let $n, s, \kappa \in \mathbb{N}$ and $\mathbb{F}$ be a field. The (global) Kaufman–Oppenheim complex, denoted $\mathfrak{A}_n^{\kappa,s}(\mathbb{F})$, is the coset complex defined by the group $\mathsf{SL}_n(\mathbb{F}^{(s)})$ and its subgroup family $(\mathsf{Unip}_{n,\ell}^{\kappa,(s)}(\mathbb{F}))_{\ell\in[0..n]}$. ◊

Then, we can easily prove Theorem 3.55, which states that vertex-links in the global Kaufman–Oppenheim complex are isomorphic to the local Kaufman–Oppenheim complex (as long as $s \geq \kappa n$):

Proof of Theorem 3.55. First, we define shorthands for our groups: Let $G := \mathsf{SL}_n(\mathbb{F}^{(s)})$. For $\ell \in [0..n]$, let $K_\ell := \mathsf{Unip}_{n,\ell}^{\kappa,(s)}(\mathbb{F}) \leq G$. Let $\mathcal{K} := (K_\ell \leq G)_{\ell\in[0..n]}$ denote the corresponding indexed subgroup family of G, and, for each $\ell \in [0..n]$, let $\mathcal{H}^\ell := (K_\ell \cap K_a \leq K_\ell)_{a\in[0..n]\setminus\{\ell\}}$ denote the corresponding indexed subgroup family of K_ℓ.

By definition, the global Kaufman–Oppenheim complex is $\mathfrak{A}_n^{\kappa,s}(\mathbb{F}) = \mathfrak{CC}(G;\mathcal{K})$. By Proposition 3.44, its vertex-links are: $(\mathfrak{CC}(G;\mathcal{K}))_v \cong \mathfrak{CC}(K_\ell;(K_\ell \cap K_a))_{a\in[0..n]\setminus\ell}) = \mathfrak{CC}(K_\ell;\mathcal{H}^\ell)$. By Proposition C.7, $\mathcal{H}^\ell$ is isomorphic to the subgroup family $(\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n;b) \leq \mathrm{Unip}_{\mathbb{F}}^{\kappa}(n))_{b\in[1..n]}$ defined in Definition 3.50. But this is precisely the local Kaufman–Oppenheim complex, defined in Definition 3.54. □

D Additional proofs regarding groups and coset complexes

D.1 Proof of Proposition 3.45

Proposition D.1. *For $H \leq G$ groups and $\boldsymbol{g} \sim \mathbf{Unif}(G)$, the coset-valued random variable $\boldsymbol{g}H$ is uniformly distributed on G/H.*

Proof. We have $\Pr[\boldsymbol{g}H = C] = \frac{1}{|G|}|\{g \in G : gH = C\}|$. Then, $\{g \in G : gH = C\} = C$, and $|C| = |H|$, so that $\Pr[\boldsymbol{g}H = C] = \frac{|H|}{|G|} = \frac{1}{[G:H]} = \frac{1}{|G/H|}$ by Lagrange's theorem. (Here, $[G : H]$ is the *index* of H in G, a.k.a., the number of cosets in G/H.) □

Proof of Proposition 3.45. We define the mapping

$$\begin{aligned}\phi_i : G/H_i &\to L/(\pi(H_i)) \\ gH_i &\mapsto \pi(g)(\pi(H_i)).\end{aligned}$$

This is well-defined because if $gH_i = g'H_i$, then $g^{-1}g' \in H_i$, so $\pi(g)^{-1}\pi(g') = \pi(g^{-1}g') \in \pi(H_i)$, so $\pi(g)(\pi(H_i)) = \pi(g')(\pi(H_i))$.

ϕ_i is obviously surjective: For any $\ell \in L$, by surjectivity, there is some $g \in G$ with $\pi(g) = \ell$; then $\phi(gH_i) = \pi(g)(\pi(H_i)) = \ell(\pi(H_i))$.

Conversely, ϕ_i is also injective: If $\phi_i(gH_i) = \phi_i(g'H_i)$, then $\pi(g)(\pi(H_i)) = \pi(g')(\pi(H_i))$, so $\pi(g^{-1}g') = \pi(g)^{-1}\pi(g') \in \pi(H_i)$, hence there exists $h \in H_i$ with $\pi(g^{-1}g') = \pi(h)$, hence $\pi(g^{-1}g'h^{-1}) = \mathbb{1}$, hence $g^{-1}g'h^{-1} \in \ker(\pi) \leq H_i$, so finally $g^{-1}g' \in H_i$ and so $gH_i = g'H_i$.

Finally, we show that $(\phi_i)_{i\in\mathcal{I}}$ is an isomorphism of coset complexes. The coset complex $\mathfrak{CC}(G;K\mathcal{H})$ samples $\boldsymbol{g}$ and outputs the cosets $(\boldsymbol{g}H_i)_{i\in\mathcal{I}}$. The image under $(\phi_i)_{i\in\mathcal{I}}$ of this distribution samples $\boldsymbol{g} \sim G$ uniformly and outputs $(\pi(\boldsymbol{g})\pi(H_i))_{i\in\mathcal{I}}$. The marginal distribution of $\pi(\boldsymbol{g})$ is uniform[20] on $\pi(G) = L$ and therefore that this is the same as sampling $\boldsymbol{l} \sim L$ and outputting $(\boldsymbol{l}\pi(H_i))_{i\in\mathcal{I}}$, which is the definition of $\mathfrak{CC}(L;(\pi(H_i))_{i\in\mathcal{I}})$. □

[20]This is because the fibers $\pi^{-1}(\ell)$ for $\ell \in L$ are in bijection with the cosets $G/\ker(\pi)$, which partition G.

D.2 Proof of Proposition 3.47

We now prove the following strong form of Proposition 3.47:

Proposition D.2. *Let G be a group and $H_1, H_2 \le G$ two subgroups. The following are equivalent:*

1. *$H_1H_2 = G$, i.e., eveAry $g \in G$ can be written $g = h_1h_2$ for $h_1 \in H_1$, $h_2 \in H_2$.*
2. *Every $C_1 \in G/H_1$ and $C_2 \in G/H_2$ have nonempty intersection.*
3. *For $\boldsymbol{g} \sim G$ uniform, the coset-valued random variables $\boldsymbol{g}H_1$ and $\boldsymbol{g}H_2$ are independent.*
4. *For $\boldsymbol{h}_1 \sim H_1$ and $\boldsymbol{h}_2 \sim H_2$ uniform and independent, $\boldsymbol{h}_1\boldsymbol{h}_2$ is distributed uniformly on G.*

Proof. (Item 1 $\Longrightarrow$ Item 2.) Pick $a \in C_1$ and $b \in H_2$. Then $a^{-1}b \in H_1H_2$ so $a^{-1}b = h_1h_2$ for some $h_1 \in H_1$ and $h_2 \in H_2$, hence $bh_2^{-1} = ah_1 \in aH_1 \cap bH_2 \ne \emptyset$.

(Item 2 $\Longrightarrow$ Item 3.) Note that $gH_i = C_i \iff g \in C_i$. Hence $gH_1 = C_1 \wedge gH_2 = C_2 \iff g \in C_1 \cap C_2$. By Proposition 3.32, any nonempty intersection of two cosets is a coset of $H_1 \cap H_2$, so for every C_1, C_2,

$$\Pr[\boldsymbol{g}H_1 = C_1,\ \boldsymbol{g}H_2 = C_2] = \frac{|H_1 \cap H_2|}{|G|}.$$

Hence the joint distribution of $(\boldsymbol{g}H_1, \boldsymbol{g}H_2)$ is uniform.

(Item 3 $\Longrightarrow$ Item 4.) Independence of $\boldsymbol{g}H_1$ and $\boldsymbol{g}H_2$ gives that for every coset pair $C_1 \in G/H_1,\ C_2 \in G/H_2$,

$$\Pr[\boldsymbol{g}H_1 = C_1,\ \boldsymbol{g}H_2 = C_2] = \Pr[\boldsymbol{g}H_1 = C_1]\Pr[\boldsymbol{g}H_2 = C_2] = \frac{|H_1||H_2|}{|G|^2}.$$

But as in the last implication, this probability (being nonzero) must equal $|H_1 \cap H_2|/|G|$, and so

$$|H_1||H_2| = |G|\,|H_1 \cap H_2|.$$

Hence by Proposition 3.37, $|H_1H_2| = |G|$ and thus $H_1H_2 = G$.

(Item 1 $\Longrightarrow$ Item 4.) For fixed $g \in G$, the pairs $(h_1, h_2) \in H_1 \times H_2$ such that $g = h_1h_2$ correspond bijectively to the set

$$\{h_1 \in H_1 :\ h_1^{-1}g \in H_2\} = H_1 \cap gH_2.$$

This set is nonempty (since $H_1H_2 = G$) and so by Item 2 it is a coset of $H_1 \cap H_2$ and has size $|H_1 \cap H_2|$. Therefore,

$$\Pr[\boldsymbol{h}_1\boldsymbol{h}_2 = g] = \frac{|H_1 \cap H_2|}{|H_1||H_2|} = \frac{1}{|G|},$$

so $\boldsymbol{h}_1\boldsymbol{h}_2 \sim \mathbf{Unif}(G)$.

(Item 4 $\Longrightarrow$ Item 1.) If $\boldsymbol{h}_1\boldsymbol{h}_2$ is uniform on G, then every $g \in G$ occurs with positive probability as a product h_1h_2, so $g \in H_1H_2$. Hence $H_1H_2 = G$. □

D.3 Proof of Proposition 4.15

Proof of Proposition 4.15. First, we claim that all three factors are contained in the intersection $H \cap K$. For the middle factor, we have $H_1 \subseteq H_1H_2 = H$ and $K_2 \subseteq K_1K_2 = K$. For e.g. the left factor, we have $K_1 \le K_1 \cap H_1 \le K \cap H$ and similarly for the right factor.

Next, we verify that the three factors are pairwise commuting and trivially intersecting. Indeed, K_1 and K_2 commute and intersect trivially by assumption, and therefore so do K_1 and $H_1 \cap K_2$. The proof for H_2 and $H_1 \cap K_2$ is similar. Finally, K_1 and H_2 commute and intersect trivially because $K_1 \le H_1$ and K_1 and K_2 commute and intersect trivially. (Alternatively, we could use that $H_2 \le K_2$ and that H_1 and H_2 commute and intersect trivially.)

So, it remains to verify that the three factors together generate $H \cap K$. Let $s \in H \cap K$. Hence $s = h_1h_2 = k_1k_2$ for some $h_i \in H_i, k_i \in K_i$. In particular, we can define $t \coloneqq h_2k_2^{-1} = h_1^{-1}k_1$. We have $t = h_2k_2^{-1} \in H_2K_2 = K_2$ (since $H_2 \le K_2$), while also $t = h_1^{-1}k_1 \in H_1K_1 = H_1$ (since $K_1 \le H_1$). Hence $t \in H_1 \cap K_2$. Finally, we can write $s = h_1h_2 = k_1t^{-1}h_2$, which has the desired form.[21] □

[21] This statement can also been seen via the so-called "Dedekind law" or "modular law", which states that if $A \le B$ then $A(B \cap C) = B \cap AC$ ([Rot95, Exercise 2.49]). In particular, our claim follows from applying this statement twice, using the two assumed containments.

D.4 Proof of Proposition 4.16

Proof of Proposition 4.16. Induct on t. In the base case, $t = 0$, we have that $\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n; \emptyset) = \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [0\,..\,n]) = \mathrm{Unip}_{\mathbb{F}}^{\kappa}(n)$, so the claim holds trivially. Otherwise, let $W' = \{\ell_1, \dots, \ell_{t-1}\}$ so that $W = W' \cup \{\ell\}$. Inductively,

$$\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n; W') = \underbrace{\left(\mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [0\,..\,\ell_1)) \times \left(\prod_{i=1}^{t-2} \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [\ell_i\,..\,\ell_{i+1}))\right)\right)}_{=:K_1} \times \underbrace{\mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [\ell_{t-1}\,..\,n])}_{=:K_2}.$$

while by Proposition 4.14,

$$\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n; \ell_t) = \underbrace{\mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [0\,..\,\ell_t - 1])}_{=:H_1} \times \underbrace{\mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [\ell_t\,..\,n])}_{=:H_2}.$$

Note that by Proposition 4.8, $K_1 \le H_1$ while $H_2 \le K_2$. Also, by Proposition 4.13,

$$H_1 \cap K_2 = \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [\ell_{t-1}\,..\,\ell_t - 1]).$$

Hence by the prior Proposition 4.15,

$$\begin{aligned}
&\mathrm{Stair}_{\mathbb{F}}^{\kappa}(n; W) = \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n; W') \cap \mathrm{Stair}_{\mathbb{F}}^{\kappa}(n; \ell_t) \\
&= \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [0\,..\,\ell_1)) \times \left(\prod_{i=1}^{t-2} \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [\ell_i\,..\,\ell_{i+1}))\right) \times \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [\ell_{t-1}\,..\,\ell_t)) \times \mathrm{Triangle}_{\mathbb{F}}^{\kappa}(n; [\ell_t\,..\,n]).
\end{aligned}$$

This is precisely the desired result. $\square$

E Local-to-global theorem for triword expansion

In this section, we prove Lemma A.18. To simplify notations, we replace $\mathfrak{Y}$ with $\mathfrak{X}$, $\mathcal{J}$ with $\mathcal{I}$, and r with k. We remark that when reading the proof, it may be helpful to first consider the case where $k = 1$. In this case, for $S_1, S_2 \sqsubset_k \mathcal{I}$, $S_1 \cap S_2 = \emptyset \iff S_1 \neq S_2$, and our argument recovers the proof in [DD24b, §4]. For general k, we have to incorporate a few additional inequalities.

E.1 Spectral expansion

The following simple proposition expresses edge expansion of induced complexes assuming productness:

Proposition E.1. *Let $\mathfrak{X}$ be an $\mathcal{I}$-indexed simplicial complex which is ϵk^{-2}-product and suppose that $\mathcal{R}$ is a set of pairwise disjoint k-sets in $\mathcal{I}$ and S is a k-set in $\mathcal{I}$ disjoint from all sets in $\mathcal{R}$. Let $s \in \mathfrak{X}[S]$ be an S-word, Γ be a finite set, and $h : V(\mathfrak{X}_s[\mathcal{R}]) \to \Gamma$ a word-labeling by Γ. Then there exists $x \in \Gamma$ such that*

$$\Pr_{\boldsymbol{R} \in \mathcal{R}, \boldsymbol{r} \sim \mathfrak{X}_s[\boldsymbol{R}]}[h(\boldsymbol{r}) \neq x] \le \frac{1}{1-\epsilon} \Pr_{\boldsymbol{R}_1 \neq \boldsymbol{R}_2 \in \mathcal{R}, (\boldsymbol{r}_1, \boldsymbol{r}_2) \sim \mathfrak{X}_s[\boldsymbol{R}_1, \boldsymbol{R}_2]}[h(\boldsymbol{r}_1) \neq h(\boldsymbol{r}_2)].$$

Proof. Let G denote the 1-skeleton of the induced complex $\mathfrak{X}_s[\mathcal{R}]$. By Corollary A.21, $\mathfrak{X}_s[\mathcal{R}]$ is ϵ-product. Hence the induced subgraph on each pair of parts in G (a.k.a. the bipartite graph $\mathfrak{X}_s[R_1, R_2]$ for $R_1, R_2 \in \mathcal{R}$) is an ϵ-bipartite expander and therefore a $1/(1-\epsilon)$-edge expander. Hence G itself is also an $1/(1-\epsilon)$-edge expander by an averaging argument (and since each pairwise induced subgraph has the same total weight), as desired. $\square$

In the sequel, we set $\epsilon_0 := \frac{1}{2}$ so that for $\epsilon \le \epsilon_0$, $\frac{1}{1-\epsilon} \le 2$.

E.2 Choosing a good center $\mathcal{R}$

We first define a parameter γ, depending only on f and $\mathfrak{X}$, which we call the *triword inconsistency* of f. Consider the following experiment: Sample three uniformly random pairwise disjoint k-sets $\boldsymbol{S}_1, \boldsymbol{S}_2, \boldsymbol{S}_3 \sqsubset_k \mathcal{I}$, sample an $(\boldsymbol{S}_1, \boldsymbol{S}_2, \boldsymbol{S}_3)$-triword $(\boldsymbol{s}_1, \boldsymbol{s}_2, \boldsymbol{s}_3) \sim \mathfrak{X}[\boldsymbol{S}_1, \boldsymbol{S}_2, \boldsymbol{S}_3]$, and define $\gamma := \Pr[f(\boldsymbol{s}_1, \boldsymbol{s}_2)f(\boldsymbol{s}_2, \boldsymbol{s}_3)f(\boldsymbol{s}_3, \boldsymbol{s}_1) \neq \mathbb{1}]$.

Further, in this setup, for any fixed set $\mathcal{R}$ of ℓ distinct k-sets in $\mathcal{I}$, and $j \in \{0,1,2,3\}$, we let $\gamma^{(\mathcal{R},j)} := \Pr[f(\boldsymbol{s}_1, \boldsymbol{s}_2)f(\boldsymbol{s}_2, \boldsymbol{s}_3)f(\boldsymbol{s}_3, \boldsymbol{s}_1) \neq \mathbb{1} \mid \boldsymbol{S}_{1:j} \in \mathcal{R}, \boldsymbol{S}_{j+1:3} \notin \mathcal{R}]$. (Here, e.g. $\boldsymbol{S}_{1:j} \in \mathcal{R}$ denotes that $\boldsymbol{S}_i \in \mathcal{R}$ for every $i \in [1\,..\,j]$.) Note that for any fixed j and uniformly random $\boldsymbol{\mathcal{R}}$, the marginal distribution of $(\boldsymbol{S}_1, \boldsymbol{S}_2, \boldsymbol{S}_3)$ conditioned on $\boldsymbol{S}_{1:j} \in \boldsymbol{\mathcal{R}}, \boldsymbol{S}_{j+1:3} \notin \boldsymbol{\mathcal{R}}$ is uniformly random, and so $\mathbb{E}_{\boldsymbol{\mathcal{R}}}[\gamma^{(\boldsymbol{\mathcal{R}},j)}] = \gamma$.

By the law of total probability, we can therefore fix $\mathcal{R}$ such that $\mathfrak{X}[\mathcal{R}]$ has triword expansion at least β over Γ, and $\max\{\gamma^{(\mathcal{R},1)}, \gamma^{(\mathcal{R},2)}, \gamma^{(\mathcal{R},3)}\} \leq 3\gamma/p$. For the remainder of this subsection, we treat $\mathcal{R}$ as fixed. We write $\gamma^{(j)} := \gamma^{(\mathcal{R},j)}$.

E.3 The construction

Now we define a word-labeling $g : V_k(\mathfrak{X}) \to \Gamma$ in two steps:

1. Consider the restriction $f^* : \vec{E}(\mathfrak{X}[\mathcal{R}]) \to \Gamma$, which labels the (R_1, R_2)-biwords for $R_1 \neq R_2 \in \mathcal{R}$. The triword inconsistency of f^* is $\gamma^{(\mathcal{R},3)} \leq p\gamma$.

   By the assumed triword expansion of $\mathfrak{X}[\mathcal{R}]$, there is word-labeling $g^* : V(\mathfrak{X}[\mathcal{R}]) \to \Gamma$, which labels k-words for $R \in \mathcal{R}$, in the following sense: Suppose that we sample $\boldsymbol{R}_1 \neq \boldsymbol{R}_2 \in \mathcal{R}$ uniformly, sample $(\boldsymbol{r}_1, \boldsymbol{r}_2) \sim \mathfrak{X}[\boldsymbol{R}_1, \boldsymbol{R}_2]$, and set $\eta^* := \Pr[g^*(\boldsymbol{r}_1)g^*(\boldsymbol{r}_2)^{-1} \neq f(\boldsymbol{r}_1, \boldsymbol{r}_2)]$. Then we have $\eta^* \leq \gamma^{(3)}/\beta \leq 3\gamma/(\beta p)$.

2. Now, for every $S \subset_k \mathcal{I}$ with $S \notin \mathcal{R}$, for every S-word $s \in \mathfrak{X}[S]$, we define a word-labeling $h_s : V(\mathfrak{X}_s[\{R \in \mathcal{R} : R \cap S = \emptyset\}]) \to \Gamma$, by setting $h_s(r) := f(s, r) \cdot g^*(r)$.

   Finally, we do the plurality decoding: $g(s) := \arg\max_{\gamma \in \Gamma} \Pr[h_s(r) = \gamma]$ (ties broken arbitrarily). (We also define $g(r) := g^*(r)$ for $r \in \mathfrak{X}[R]$, $R \in \mathcal{R}$).

We now consider the biword inconsistency η of g. Specifically, consider sampling $\boldsymbol{S}_1, \boldsymbol{S}_2 \sqsubset_k \mathcal{I}$ uniformly at random and $(\boldsymbol{s}_1, \boldsymbol{s}_2) \sim \mathfrak{X}[\boldsymbol{S}_1, \boldsymbol{S}_2]$ and set $\eta := \Pr[g(\boldsymbol{s}_1)g(\boldsymbol{s}_2)^{-1} \neq f(\boldsymbol{s}_1, \boldsymbol{s}_2)]$. We want to upper-bound η multiplicatively relative to γ.

To do so, we consider the following quantity $\eta^{(1)}$: Sample $\boldsymbol{R}, \boldsymbol{S} \sqsubset_k \mathcal{I}$ uniformly, sample $(\boldsymbol{r}, \boldsymbol{s}) \sim \mathfrak{X}[\boldsymbol{R}, \boldsymbol{S}]$, and set $\eta^{(1)} := \Pr[g(\boldsymbol{r})g(\boldsymbol{s})^{-1} \neq f(\boldsymbol{r}, \boldsymbol{s}) \mid \boldsymbol{R} \in \mathcal{R}, \boldsymbol{S} \notin \mathcal{R}]$.

Similarly, define $\eta^{(2)}$ via: Sample $\boldsymbol{S}_1, \boldsymbol{S}_2 \sqsubset_k \mathcal{I}$ uniformly conditioned on $\boldsymbol{S}_1, \boldsymbol{S}_2 \notin \mathcal{R}$, sample $(\boldsymbol{s}_1, \boldsymbol{s}_2) \sim \mathfrak{X}[\boldsymbol{S}_1, \boldsymbol{S}_2]$, and then set $\eta^{(2)} := \Pr[g(\boldsymbol{s}_1)g(\boldsymbol{s}_2)^{-1} \neq f(\boldsymbol{s}_1, \boldsymbol{s}_2)]$.

Now we claim that η is a convex combination of η^*, $\eta^{(1)}$, and $\eta^{(2)}$; indeed, we can condition on the overlap size $|\{\boldsymbol{S}_1, \boldsymbol{S}_2\} \cap \mathcal{R}|$, and the probabilities conditioned on overlap sizes 0, 1, and 2 are precisely $\eta^{(2)}$, $\eta^{(1)}$, and η^*, respectively. We have already bounded η^* relative to γ. In the sequel, we bound $\eta^{(1)}$ and $\eta^{(2)}$.

E.4 Qualitative analysis

We first establish some simple implications involving the labelings h_s and g.

Claim E.2. *Let $R_1, R_2, S \sqsubset_k \mathcal{I}$ with $R_1, R_2 \in \mathcal{R}$, $S \notin \mathcal{R}$, and suppose that $(r_1, r_2, s) \in \mathfrak{X}[R_1, R_2, S]$ is a valid (R_1, R_2, S)-triword. Then:*

$$f(r_1, r_2)f(r_2, s)f(s, r_1) = \mathbb{1} \wedge g^*(r_1)g^*(r_2)^{-1} = f(r_1, r_2) \quad \Longrightarrow \quad h_s(r_1) = h_s(r_2).$$

Proof. By definition, $h_s(r_1) = h_s(r_2) \iff f(s, r_1) \cdot g^*(r_1) = f(s, r_2) \cdot g^*(r_2) \iff f(r_2, s) \cdot f(s, r_1) \cdot g^*(r_1) \cdot g^*(r_2)^{-1} = \mathbb{1}$. The two hypotheses finish the proof. $\square$

Claim E.3. *Let $R, S \sqsubset_k \mathcal{I}$ with $R \in \mathcal{R}$ and $S \notin \mathcal{R}$, and suppose that $(r, s) \in \mathfrak{X}[R, S]$ is a valid (R, S)-biword. Then*

$$g(s) = h_s(r) \quad \Longrightarrow \quad g(r)g(s)^{-1} = f(r, s).$$

Proof. By definition, $g(r) = g^*(r)$ and $h_s(r) = f(s,r) \cdot g^*(r)$. Hence by hypothesis, $g(r)g(s)^{-1} = g(r)h_s(r)^{-1} = f(r,s)$, as desired. ☐

Claim E.4. *Suppose $R, S_1, S_2 \sqsubset_k \mathcal{I}$ with $R \in \mathcal{R}$ and $S_1, S_2 \notin \mathcal{R}$, and suppose that $(r, s_1, s_2) \in \mathfrak{X}[R, S_1, S_2]$ is a valid (R, S_1, S_2)-triword. Then:*

$$g(s_1) = h_{s_1}(r) \wedge g(s_2) = h_{s_2}(r) \wedge f(s_1,s_2)f(s_2,r)f(r,s_1) = \mathbb{1} \quad \Longrightarrow \quad g(s_1)g(s_2)^{-1} = f(s_1,s_2).$$

Proof. By the first hypothesis, $g(s_1)g(s_2)^{-1} = h_{s_1}(r)h_{s_2}(r)^{-1}$. By definition, we have $h_{s_1}(r)h_{s_2}(r)^{-1} = f(s_1,r)g^*(r)g^*(r)^{-1}f(r,s_2) = f(s_1,r)f(r,s_2)$. The second hypothesis completes the proof. ☐

E.5 Sampling sets

Now, we turn to proving some simple lemmas about sampling k-sets conditioned on disjointness hypotheses. We stress that this section is unnecessary in the $k = 1$ case since disjointness and distinctness are equivalent.

Given two distributions X and Y on a finite set Ω and $\lambda \geq 1$, we say X and Y are *λ-multiplicatively close* if for every $\omega \in \Omega$, $\lambda^{-1} \Pr_{\boldsymbol{\omega} \sim Y}[\boldsymbol{\omega} = \omega] \leq \Pr_{\boldsymbol{\omega} \sim X}[\boldsymbol{\omega} = \omega] \leq \lambda \Pr_{\boldsymbol{\omega} \sim Y}[\boldsymbol{\omega} = \omega]$. Note that this is a symmetric relationship, and that it implies that for every nonnegative function $f : \Omega \to \mathbb{R}_{\geq 0}$, we have $\lambda^{-1} \mathbb{E}_{\boldsymbol{\omega} \sim Y}[f(\boldsymbol{\omega})] \leq \mathbb{E}_{\boldsymbol{\omega} \sim X}[f(\boldsymbol{\omega})] \leq \lambda \mathbb{E}_{\boldsymbol{\omega} \sim Y}[f(\boldsymbol{\omega})]$.

Claim E.5. *Let $G = (L, R, E)$ be an (unweighted) bipartite graph. Consider the following two distributions over E:*

1. *Sample $(\boldsymbol{u}, \boldsymbol{v}) \in E$ uniformly.*
2. *Sample $\boldsymbol{v} \in L$ uniformly, then sample $\boldsymbol{u} \in L$ uniformly conditioned on $(\boldsymbol{u}, \boldsymbol{v}) \in E$.*

For $v \in R$, let $\deg(v) \coloneqq |\{u \in L : (u,v) \in E\}|$. The two above distributions are $(\max_v \deg(v))/(\min_v \deg(v))$-multiplicatively close.

Proof. The probability of an edge (u,v) in the first distribution is $1/|E|$ and the probability in the second distribution is $1/(|R| \cdot \deg(v))$. Note also that $|E| = \sum_{v \in R} \deg(v)$. Hence the ratio of probabilities is $|L| \cdot \deg(v)/(\sum_{v' \in R} \deg(v'))$, which is e.g. uniformly upper-bounded by $|R| \cdot (\max_v \deg(v))/(\sum_{v \in R} \min_{v'} \deg(v')) = (\max_v \deg(v))/(\min_v \deg(v))$. ☐

Claim E.6. *Consider the following distributions on pairs (R, S) with $R, S \sqsubset_k \mathcal{I}$, $R \in \mathcal{R}$, and $S \notin \mathcal{R}$:*

1. *Sample $\boldsymbol{R}, \boldsymbol{S} \subset_k \mathcal{I}$ uniformly and independently, conditioned on $\boldsymbol{R} \in \mathcal{R}$, $\boldsymbol{S} \notin \mathcal{R}$, and $\boldsymbol{R} \cap \boldsymbol{S} = \emptyset$.*
2. *First, sample $\boldsymbol{S} \subset_k \mathcal{I}$ uniformly conditioned on $\boldsymbol{S} \notin \mathcal{R}$. Then, sample $\boldsymbol{R} \in \mathcal{R}$ uniformly conditioned on $\boldsymbol{R} \cap \boldsymbol{S} = \emptyset$.*

Then these two distributions are $3/2$-multiplicatively close.

Proof. Form the bipartite graph where the left-vertices are $\mathcal{R}$, the right-vertices are $S \subset_k \mathcal{I}$ with $S \notin \mathcal{R}$, and there is an edge between R and S if $R \cap S = \emptyset$. The maximum right-degree is at most ℓ (the total number of left-vertices) and the minimum right-degree is at least $\ell - k$ since any k-set can intersect at most k sets in $\mathcal{R}$ (since the sets in $\mathcal{R}$ are pairwise disjoint). Finally, $\ell/(\ell - k) \leq 3/2$ by the assumption that $\ell \geq 3k$. ☐

Claim E.7. *Consider the following distributions on pairs (R_1, R_2, S) with $R_1, R_2, S \sqsubset_k \mathcal{I}$, $R_1, R_2 \in \mathcal{R}$, $S \notin \mathcal{R}$:*

1. *Sample $\boldsymbol{R}_1, \boldsymbol{R}_2, \boldsymbol{S} \subset_k \mathcal{I}$ uniformly and independently, conditioned on $\boldsymbol{R}_1, \boldsymbol{R}_2 \in \mathcal{R}$, $\boldsymbol{S} \notin \mathcal{R}$, and $\boldsymbol{R}_1, \boldsymbol{R}_2, \boldsymbol{S}$ are pairwise disjoint.*
2. *First, sample $\boldsymbol{S} \subset_k \mathcal{I}$ uniformly conditioned on $\boldsymbol{S} \notin \mathcal{R}$. Then, sample $\boldsymbol{R}_1, \boldsymbol{R}_2 \in \mathcal{R}$ uniformly conditioned on $\boldsymbol{R}_1 \neq \boldsymbol{R}_2$, $\boldsymbol{R}_1 \cap \boldsymbol{S} = \emptyset$, and $\boldsymbol{R}_2 \cap \boldsymbol{S} = \emptyset$.*

Then these two distributions are 3-multiplicatively close.

Proof. Form the bipartite graph where left-vertices are ordered pairs (R_1, R_2) with $R_1 \neq R_2 \in \mathcal{R}$, right-vertices are $S \subset_k \mathcal{I}$ with $S \notin \mathcal{R}$, and there is an edge between (R_1, R_2) and S if $(R_1 \cup R_2) \cap S = \emptyset$. The maximum right-degree is at most $\ell(\ell-1)$ while the minimum right-degree $(\ell-k)(\ell-k-1)$, and the ratio is $(\ell(\ell-1))/((\ell-k)(\ell-k-1)) \leq (3k(3k-1))/(2k(2k-1)) \leq 3$. ☐

Claim E.8. *Consider the following distributions on triples (R, S_1, S_2) with $R, S_1, S_2 \sqsubset_k \mathcal{I}$, $R \in \mathcal{R}$, and $S_1, S_2 \notin \mathcal{R}$:*

1. *Sample $\boldsymbol{R}, \boldsymbol{S}_1, \boldsymbol{S}_2 \sqsubset_k \mathcal{I}$ uniformly conditioned on $\boldsymbol{R} \in \mathcal{R}, \boldsymbol{S}_1, \boldsymbol{S}_2 \notin \mathcal{R}$.*
2. *Sample $\boldsymbol{S}_1, \boldsymbol{S}_2 \sqsubset_k \mathcal{I}$ uniformly conditioned on $\boldsymbol{S}_1, \boldsymbol{S}_2 \notin \mathcal{R}$. Then, sample $\boldsymbol{R} \in \mathcal{R}$ uniformly conditioned on $\boldsymbol{R} \cap (\boldsymbol{S}_1 \cup \boldsymbol{S}_2) = \emptyset$.*

Then these two distributions are 3*-multiplicatively close.*

Proof. Form the bipartite graph where the left-vertices are $\mathcal{R}$, the left-vertices are (S_1, S_2) with $S_1, S_2 \sqsubset_k \mathcal{I}$ and $S_1, S_2 \notin \mathcal{R}$, and there is an edge between R and (S_1, S_2) if $R \cap (S_1 \cup S_2) = \emptyset$. The maximum right-degree is at most ℓ and the minimum right-degree is at least $\ell - 2k$, which yields a ratio of $\ell/(\ell - 2k) \leq 3$. ☐

E.6 Quantitative analysis: $\eta^{(1)}$

To analyze $\eta^{(1)}$, we define a few more quantities. For fixed $S \subset_k \mathcal{I}$ with $S \notin \mathcal{R}$, we consider:

- Sample $\boldsymbol{R} \in \mathcal{R}$ uniformly conditioned on $\boldsymbol{R} \cap S = \emptyset$, sample $(\boldsymbol{r}, \boldsymbol{s}) \sim \mathfrak{X}[\boldsymbol{R}, S]$, and let $\eta_S := \Pr[g(\boldsymbol{r})g(\boldsymbol{s})^{-1} \neq f(\boldsymbol{r}, \boldsymbol{s})]$ and $\psi_S := \Pr[g(\boldsymbol{s}) \neq h_{\boldsymbol{s}}(\boldsymbol{r})]$.
- Sample $\boldsymbol{R}_1 \neq \boldsymbol{R}_2 \in \mathcal{R}$ uniformly conditioned on $(\boldsymbol{R}_1 \cup \boldsymbol{R}_2) \cap S = \emptyset$, sample $(\boldsymbol{r}_1, \boldsymbol{r}_2, \boldsymbol{s}) \sim \mathfrak{X}[\boldsymbol{R}_1, \boldsymbol{R}_2, S]$, and let $\theta_S := \Pr[h_{\boldsymbol{s}}(\boldsymbol{r}_1) \neq h_{\boldsymbol{s}}(\boldsymbol{r}_2)]$ and $\gamma_S := \Pr[f(\boldsymbol{r}_1, \boldsymbol{r}_2)f(\boldsymbol{r}_2, \boldsymbol{s})f(\boldsymbol{s}, \boldsymbol{r}_1) \neq \mathbb{1}]$.
- Sample $\boldsymbol{R}_1 \neq \boldsymbol{R}_2 \in \mathcal{R}$ uniformly conditioned on $(\boldsymbol{R}_1 \cup \boldsymbol{R}_2) \cap S = \emptyset$, sample $(\boldsymbol{r}_1, \boldsymbol{r}_2) \sim \mathfrak{X}[\boldsymbol{R}_1, \boldsymbol{R}_2]$, and let $\eta_S^* := \Pr[g(\boldsymbol{r}_1)g(\boldsymbol{r}_2)^{-1} \neq f(\boldsymbol{r}_1, \boldsymbol{r}_2)]$.

By Claim E.3, $\eta_S \leq \psi_S$. By Proposition E.1, $\psi_S \leq 2\theta_S$. By Claim E.2, $\theta_S \leq \gamma_S + \eta_S^*$. Hence $\eta_S \leq 2(\gamma_S + \eta_S^*)$.

Now, consider sampling $\boldsymbol{S} \subset_k \mathcal{I}$ uniformly conditioned on $\boldsymbol{S} \notin \mathcal{R}$. By Claim E.6, $\eta^{(1)} \leq \frac{3}{2} \mathbb{E}\, \eta_{\boldsymbol{S}}$. By Claim E.7, $\mathbb{E}\, \gamma_{\boldsymbol{S}} \leq 3\gamma^{(2)}$ and $\mathbb{E}\, \eta_{\boldsymbol{S}}^* \leq 3\eta^*$. These combine to give $\mathbb{E}\, \eta_{\boldsymbol{S}} \leq 6(\gamma^{(2)} + \eta^*)$ and therefore $\eta^{(1)} \leq 9(\gamma^{(2)} + \eta^*)$.

E.7 Quantitative analysis: $\eta^{(2)}$

Again, we define some quantities. For fixed $S_1, S_2 \sqsubset_k \mathcal{I}$ with $S_1, S_2 \notin \mathcal{R}$:

- Sample $(\boldsymbol{s}_1, \boldsymbol{s}_2) \sim \mathfrak{X}[S_1, S_2]$, and let $\eta_{S_1,S_2} := \Pr[g(\boldsymbol{s}_1)g(\boldsymbol{s}_2)^{-1} \neq f(\boldsymbol{s}_1, \boldsymbol{s}_2)]$.
- Sample $\boldsymbol{R} \sim \mathcal{R}$ conditioned on $\boldsymbol{R} \cap (S_1 \cup S_2) = \emptyset$, sample $(\boldsymbol{s}_1, \boldsymbol{s}_2, \boldsymbol{r}) \sim \mathfrak{X}[S_1, S_2, \boldsymbol{R}]$, and let $\gamma_{S_1,S_2} := \Pr[f(\boldsymbol{s}_1, \boldsymbol{s}_2)f(\boldsymbol{s}_2, \boldsymbol{r})f(\boldsymbol{r}, \boldsymbol{s}_1) \neq \mathbb{1}]$.

By Claim E.4, $\eta_{S_1,S_2} \leq \eta_{S_1} + \eta_{S_2} + \gamma_{S_1,S_2}$.

Now, consider sampling $\boldsymbol{S}_1, \boldsymbol{S}_2 \sqsubset_k \mathcal{I}$ conditioned on $\boldsymbol{S}_1, \boldsymbol{S}_2 \notin \mathcal{R}$. By definition, $\eta^{(2)} = \mathbb{E}\, \eta_{\boldsymbol{S}_1,\boldsymbol{S}_2}$. Hence $\eta^{(2)} \leq \mathbb{E}\, \eta_{\boldsymbol{S}_1} + \mathbb{E}\, \eta_{\boldsymbol{S}_2} + \mathbb{E}\, \gamma_{\boldsymbol{S}_1,\boldsymbol{S}_2}$. By Claim E.7, $\mathbb{E}\, \gamma_{\boldsymbol{S}_1,\boldsymbol{S}_2} \leq 3\gamma^{(1)}$. Further, noting that the marginal distribution of each $\boldsymbol{S}_i$ is the same as sampling $\boldsymbol{S}_i \subset_k \mathcal{I}$ conditioned on $\boldsymbol{S}_i \notin \mathcal{R}$, we get $\mathbb{E}\, \eta_{\boldsymbol{S}_i} \leq 6(\gamma^{(2)} + \eta^*)$. Altogether, we get $\eta^{(2)} \leq 12(\gamma^{(2)} + \eta^*) + 3\gamma^{(1)}$.

E.8 Finishing the proof

Finally, we give:

Proof of Lemma A.18. By convexity, $\eta \leq \max\{\eta^*, \eta^{(1)}, \eta^{(2)}\}$. We upper-bounded the second term by $9(\gamma^{(2)} + \eta^*)$ and the third term by $12(\gamma^{(2)} + \eta^*) + 3\gamma^{(1)}$. Using that $\eta^* \leq \gamma^{(3)}/\beta$ and $\max\{\gamma^{(1)}, \gamma^{(2)}, \gamma^{(3)}\} \leq 3\gamma/p$, we get a final bound of $9\gamma(5 + 4/\beta)/p \leq 117\gamma/(\beta p)$. This gives a final expansion bound of $\beta p/117$, as desired. ☐

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il