

*Pacific
Journal of
Mathematics*

Volume 341 No. 2

April 2026

PACIFIC JOURNAL OF MATHEMATICS

Founded in 1951 by E. F. Beckenbach (1906–1982) and F. Wolf (1904–1989)

msp.org/pjm

EDITORS

Don Blasius (Managing Editor)
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
blasius@math.ucla.edu

Matthias Aschenbrenner
Fakultät für Mathematik
Universität Wien
Vienna, Austria
matthias.aschenbrenner@univie.ac.at

Vyjayanthi Chari
Department of Mathematics
University of California
Riverside, CA 92521-0135
chari@math.ucr.edu

Atsushi Ichino
Department of Mathematics
Kyoto University
Kyoto 606-8502, Japan
atsushi.ichino@gmail.com

Robert Lipshitz
Department of Mathematics
University of Oregon
Eugene, OR 97403
lipshitz@uoregon.edu

Kefeng Liu
School of Sciences
Chongqing University of Technology
Chongqing 400054, China
liu@math.ucla.edu

Sucharit Sarkar
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
sucharit@math.ucla.edu

Dimitri Shlyakhtenko
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
shlyakht@ipam.ucla.edu

Ruixiang Zhang
Department of Mathematics
University of California
Berkeley, CA 94720-3840
ruixiang@berkeley.edu

PRODUCTION

Silvio Levy, Scientific Editor, production@msp.org


See inside back cover or msp.org/pjm for submission instructions.

The subscription price for 2026 is US \$710/year for the electronic version, and \$965/year for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004, U.S.A. The Pacific Journal of Mathematics is indexed by Mathematical Reviews, Zentralblatt MATH, PASCAL CNRS Index, Referativnyi Zhurnal, Current Mathematical Publications and Web of Knowledge (Science Citation Index).

The Pacific Journal of Mathematics (ISSN 1945-5844 electronic, 0030-8730 printed) at Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004, is published twelve times a year. Periodical rate postage paid at Berkeley, CA 94704, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004.

PJM peer review and production are managed by EditFLOW[®] from Mathematical Sciences Publishers.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<http://msp.org/>

© 2026 Mathematical Sciences Publishers

CHAINMAIL LINKS AND L -SPACES

IAN AGOL

We prove that alternating chainmail links are L -space links. The proof is inspired by corresponding proofs for double branched covers of alternating links. More generally, we show that flat augmented chainmail links are generalized L -space links. Some other properties of these links are also considered.

1. Introduction

In the process of exploring lens space Dehn fillings on knots, Ozsváth and Szabó introduced the notion of L -spaces [OS3, Definition 1.1] and proved that double branched covers of nonsplit alternating links are L -spaces [OS4, Lemma 3.2, Proposition 3.3]. L -spaces are closed rational homology 3-spheres Y whose Heegaard Floer homology is minimal (like a lens space), in the sense that $\text{rank}(\widehat{HF}(Y)) = |H_1(Y; \mathbb{Z})|$. L -spaces are interesting manifolds since they do not admit taut orientable foliations [OS1, Theorem 1.4; KR, Corollary 1.6]; the converse has been conjectured for irreducible closed 3-manifolds by Juhász [Juh, Conjecture 5]. Related concepts have been defined for other versions of Floer homology, for example an *instanton L -space* Y satisfies $\dim I^\#(Y) = |H_1(Y; \mathbb{Z})|$ [BS, Definition 1.13]. It is conjectured that $\dim I^\#(Y) = \dim \widehat{HF}(Y; \mathbb{C})$ [KM, Conjecture 7.24].

Quasi-alternating links [OS4, Definition 3.1] are a class of links whose double branched covers are L -spaces, leading to the question of whether there are L -spaces which are not double branched covers of links? Examples of hyperbolic L -spaces were given which do not admit any symmetries, and hence cannot be double branched covers of links [DHL, Theorem 1.2]. Later asymmetric hyperbolic L -space knots were discovered leading to further examples [BL; ABGK+]. Meanwhile, Boyer, Gordon and Watson conjectured that L -spaces do not have orderable fundamental group [BGW, Conjecture 1]; they proved this for branched double covers of alternating links [BGW, Theorem 4], with a subsequent simpler proof given by Greene [Gre2, Theorem 2.1]. Greene's proof relies on a presentation of the fundamental group associated to a certain Dehn filling on a link described by Ozsváth and Szabó which gives a Kirby diagram for a double branched cover of a link (described in [Gre1, Section 3.1]).

Ian Agol is supported by a Simons Investigator grant #376200.

MSC2020: 57K18.

Keywords: alternating link, Heegaard Floer homology.

Polyak introduced the notion of chainmail graphs and corresponding chainmail links, such that surgery on this class of links gives all closed 3-manifolds [Pol]. He indicated that there exists a finite set of moves on chainmail graphs (analogous to Reidemeister moves) generating the equivalence relation of chainmail graphs giving the same 3-manifold. This is just a restricted class of Kirby diagrams, motivated by an alternative proof of Kirby's theorem given by Matveev and Polyak using a tangle presentation for the mapping class group [MP]. The Ozsváth–Szabó surgery diagrams of for double branched covers of links (mentioned in the last paragraph) are certain surgeries on chainmail links. We learned about these topics from a MathOverflow post by Lucas Culler [hc].

Gorsky and Némethi introduced the terminology of L -space link (links such that sufficiently large positive surgeries are L -spaces) and proved that algebraic links are L -space links [GN2, Theorem 2]. Subsequently Yajing Liu explored the properties of L -space links and introduced the terminology of *generalized L -space links* which also admit infinitely many L -space fillings [Liu, Definition 2.9]. He asked whether every 3-manifold is a surgery on a generalized L -space link [Liu, Question 1.19]? The main result of this paper is that alternating chainmail links are L -space links. Thus each double branched cover associated to an alternating diagram of a link fits into infinite families of L -spaces.

More generally, we prove that fully augmented chainmail links are generalized L -space links, hence answering Liu's question positively by Polyak's observation that every 3-manifold is surgery on a chainmail link (and hence on a fully augmented chainmail link). Since the symmetries of flat augmented links are understood, we may also therefore conclude the existence of asymmetric L -space links with arbitrarily many components.

We show that certain positive surgeries on negative alternating chainmail links do not have orderable fundamental group. The proof follows closely the strategy of Greene, using a related presentation of the fundamental group. Finally, we note that all known examples in the literature of L -space links have complements that fiber over the circle (this is true for L -space knots by [Ni1]). We observe that alternating chainmail links have fibered complements, continuing this trend. In the final section we ask some natural questions stemming from these examples.


Background. Different sections of this paper assume different background knowledge. Because some of the proofs are fairly short modifications of results in the literature, we will use standard terminology as defined in the primary sources for each section to avoid a plethora of definitions which are only used once. Hopefully the reader will be able to black-box as much as possible and refer to the sources for more details when necessary.

2. Chainmail links

2.1. Chainmail graphs. Polyak introduced the terminology of *chainmail link* [Pol] and pointed out that a result of Matveev and Polyak implies that every closed connected orientable 3-manifold is realized by surgery on a chainmail link [MP]. Since Polyak’s results have not been published, we review some of his notation and terminology which, although not strictly necessary, will be a useful way to formalize the proofs in this paper.

Definition 2.1. A *chainmail graph* is a finite planar graph $G \subset S^2$ whose edges and vertices are decorated with weights in \mathbb{Z} . Note that G may have loops and multiedges. Each vertex $v \in V(G)$ is decorated with a weight $v(v) \in \mathbb{Z}$, and each edge $e \in E(G)$ is given a weight $\epsilon(e) \in \mathbb{Z}$, where $V(G)$ and $E(G)$ are the set of vertices and edges of G respectively. Thus a chainmail graph consists of a triple (G, v, ϵ) , but much of the time we will abbreviate this just as G with the weights implicit (see also [AWvW] for this notation).

Given a chainmail graph G , one may define a *chainmail link* L_G in two steps in the following way. For each $v_i \in V(G)$, place a planar unknot L_i . Orient L_i counterclockwise. For each edge e connecting vertices v_i, v_j , introduce a $\epsilon(e)$ -clapped ribbon linking L_i and L_j . Let w_{ij} be the sum of edge weights of edges connecting v_i and v_j , with the convention that if distinct vertices v_i and v_j are not connected by an edge, then define $w_{ij} = 0$.



Now we assign framings to the components L_i with framing $w_{ii} = v(v_i) - \sum_{k \neq i} w_{ik}$. This framing is chosen in such a way that the linking number of the framing slope of L_i with the chainmail link is w_{ii} . In turn this gives a framing matrix $\Lambda(G) = (w_{ij})$. Surgery on L_G with these framings gives a 3-manifold M_G . Moreover, $\det(\Lambda(G)) \neq 0$ if and only if M_G is a rational homology 3-sphere, in which case $|H_1(M_G; \mathbb{Z})| = |\det(\Lambda(G))|$. We may also view $\Lambda(G)$ as the intersection form of $H_2(W_G)$ of the 4-manifold W_G obtained by adding two handles to the 4-ball along the link L_G with given framings.

One special property of chainmail links associated to simplicial graphs is that the crossing number is the minimal possible given the linking numbers between components. Chainmail graphs generalize plumbing diagrams for graph manifolds studied by Scharf [Sch] and Fintushel and Stern [FS], who considered the case of trees with edges of weight ± 1 and allowing rational weights on vertices (note however that the vertex labeling convention here may be different).




Figure 1. Closure of a tangle.

We briefly review some of the notation from [MP]. Consider framed tangles T_{2n} in $\mathbb{R}^2 \times [0, 1]$ with $2n$ framed endpoints on $\mathbb{R}^2 \times \{1\}$ and on $\mathbb{R}^2 \times \{0\}$. We may draw these tangles with blackboard framing, and compose the tangles by stacking.

One may close up a tangle in T_{2n} to get a framed link as in Figure 1 by adding strands on top of $\mathbb{R}^2 \times \{1\}$ and removing the strands meeting the bottom $\mathbb{R}^2 \times \{0\}$, and hence a Kirby diagram for a 3-manifold. Matveev and Polyak prove that one may obtain all genus 2 manifolds by stacking the tangles and their mirror images in Figure 2, then closing. (In the first diagram of figure 2, a positive twist, which increases the framing by 1 in the blackboard convention, is denoted by a little circle to simplify the notation.) In fact they show that these tangles generate the genus two mapping class group when one mods out by Kirby moves, such that the generators in Figure 2 correspond to the Dehn twists about curves in Figure 11. For example, the manifold associated to $\beta_2^{-1} \alpha_1 \alpha_2 \beta_1^{-1} \delta_2 \beta_2^{-1} \alpha_1 \alpha_2 \beta_1^{-1} \beta_2^{-1}$ is shown in Figure 3.

Theorem 2.2 [Pol]. *Any closed oriented 3-manifold $Y \cong M_G$ for some chainmail graph G .*




Figure 2. Generators for the genus 2 mapping class group.




Figure 3. A plat closure associated to the tangle $\beta_2^{-1}\alpha_1\alpha_2\beta_1^{-1}\delta_2\beta_2^{-1}\alpha_1\alpha_2\beta_1^{-1}\beta_2^{-1}$.

Proof. The plat representation of a Kirby diagram for a three-manifold Y given in [MP, Theorem 7.1] is easily seen to be an integral surgery on a chainmail link [Pol] as in Figure 3. □

We remark on a few moves on chainmail graphs that give rise to the same manifold. These may be verified by using flypes of the corresponding diagram keeping track of framings or Kirby calculus. A zero-weighted edge may be erased. Two edges e_1, e_2 connecting the same pair of vertices may be redrawn as one edge of weight $\epsilon(e_1) + \epsilon(e_2)$. A loop may be erased. A vertex v with weight $\nu(v) = 0$ of degree 1 adjacent to the edge e with $\epsilon(e) = \pm 1$ may be removed along with the edge e (see Figure 4; see also [EN, Theorem 18.3]). Even though we may get rid of loops and multiedges, it is convenient to allow them in chainmail graphs since they can get created under edge contraction.

2.2. Alternating chainmail links. When all of the edge weights of a chainmail graph G are of the same sign, then the associated chainmail link is alternating, see Figure 5. If the edge weights are all negative, then we call the associated




Figure 4. Moves on chainmail graphs G preserving the 3-manifold M_G .




Figure 5. From a graph with negative edge weights to an alternating chainmail link.

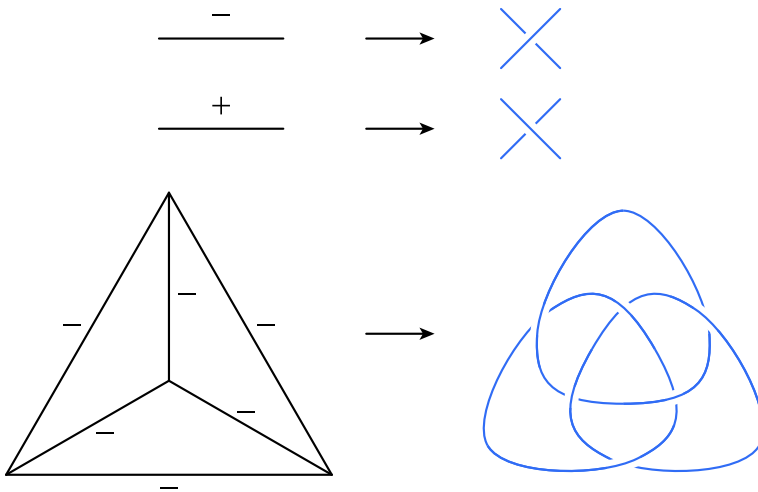


Figure 6. From an edge-weighted graph G to a link K_G .

link a *negative alternating chainmail link*. If the graph is balanced and connected ($v(v) = 0$ for all $v \in V(G)$), then it was shown by Ozsváth and Szabó that M_G is homeomorphic to the connected sum of the 2-fold branched cover of the link K_G and $S^2 \times S^1$, where K_G is obtained from G as in Figure 6 [Gre1, §3.1].

2.3. Augmented chainmail links. Assume the graph G is simplicial (so that there are no loops or edges with the same endpoints). We may insert crossing circles into L_G then twist to get an augmented link (Figure 7), possibly at only some of the edges. If we insert crossing circles associated to every edge, then it is called a *flat fully augmented (chainmail) link*; otherwise it is a partially augmented link. Thus one obtains a flat fully augmented chainmail link for every planar graph (Figure 8). The labels on a chainmail graph then correspond to framing coefficients for the loops corresponding to vertices and the reciprocals edge weights of the flat fully




Figure 7. Augmenting a chainmail link.




Figure 8. Going from a planar graph to a flat augmented link.

augmented chainmail link. Thus every 3-manifold is surgery on a fully augmented (flat) link by Theorem 2.2, since every chainmail link is surgery on such a link.

The following theorem may be deduced from the main result of [MT]:

Theorem 2.3. *A symmetry of a flat fully augmented link complement which sends crossing loops to crossing loops is induced by a symmetry of the link in S^3 .*

This will be used in the proof of Theorem 6.1.

3. Alternating chainmail links are L -space links

In this section, we prove that alternating chainmail links are L -space links. Although Theorem 3.2 is a corollary of Theorem 4.2 by taking the coefficients on each crossing loop to be -1 , we prove it separately as it serves as a base case of the induction for the more general theorem.

Definition 3.1. [GN1, Definition 2.2] An l -component link $L \subset S^3$ is called an (instanton) L -space link if there exist integers p_1, \dots, p_l such that $S^3_{n_1, \dots, n_l}(L)$ is an (instanton) L -space for all n_1, \dots, n_l with $n_i > p_i$ for all $1 \leq i \leq l$.

Given a chainmail graph (G, ν, ϵ) , we may form minors. For an edge $e \in E(G)$ with $\epsilon(e) = -1$, let $G - e$ be the graph obtained by deleting the edge e and keeping the weight functions the same on vertices and edges. If e is not a loop, and e has endpoints v_1, v_2 , let G/e be the graph obtained by quotienting e to a point (which preserves planarity). The weight function on vertices $V(G) - \{v_1, v_2\}$ is the restriction of ν , and the weight function ϵ on edges are the same. There is a new




Figure 9. Minors of G obtained by deleting and contracting the edge e .




Figure 10. Links and surgeries on augmented link corresponding to minors of G .

vertex v' of G/e obtained by identifying v_1 and v_2 . We define $v(v') = v(v_1) + v(v_2)$ (Figure 9).

Let us see how the framed links / 3-manifolds associated to the graphs G , $G - e$, and G/e differ, with $\epsilon(e) = -1$. Consider the augmented link obtained by inserting a crossing loop c about L_1 and L_2 , the components corresponding to v_1 and v_2 , then do a twist to remove the double crossing. Then L_G corresponds to doing -1 surgery on c , L_{G-e} corresponds to ∞ surgery on c (which corresponds to erasing c), and $L_{G/e}$ corresponds to 0 surgery on c (see Figure 10). This last claim follows from a Kirby calculus operation K_3 described in [MP, Figure 3, p. 539].

Thus we see that M_G , M_{G-e} , and $M_{G/e}$ are related by surgeries on the knot c with slopes which have pairwise intersection number 1. Hence we have an exact triangle

$$\widehat{HF}(M_G) \rightarrow \widehat{HF}(M_{G-e}) \rightarrow \widehat{HF}(M_{G/e}) \rightarrow \widehat{HF}(M_G)$$

[OS2, Theorem 9.16].

The following theorem generalizes Example 3.15 of [Liu].

Theorem 3.2. *Negative alternating chainmail links are L -space links. More specifically, for a chainmail graph G with all $v(v) \geq 0$, $v \in V(G)$, and at least one $v(v) > 0$ in each component of G , and $\epsilon(e) < 0$ for all $e \in E(G)$, M_G is an L -space and an instanton L -space.*

Proof. The proof is similar to the proof that double branched covers of (quasi-)alternating links are L -spaces in [OS3, Proposition 3.3]. By replacing edges e

with weight $\epsilon(e) < -1$ with parallel copies, we may assume that $\epsilon(e) = -1$ for every edge $e \in E(G)$, and we may assume that there are no loops by removing them without changing the manifold.

If G is a chainmail graph and $\nu(v) = 0$ for all vertices and $\epsilon(e) = -1$ for all edges, then the linking matrix $\Lambda(G)$ is the matrix of the graph Laplacian of the underlying unweighted graph, hence all eigenvalues are ≥ 0 , and the 0-eigenspace has dimension the number of components of G (essentially the combinatorial Hodge theorem for 0-cocycles). If G has $\nu(v) \geq 0$, and at least one $\nu(v) > 0$ in each component of G and $\epsilon(e) = -1$, then when we add the diagonal matrix with diagonal entries $\nu(v_i)$ to the Laplacian matrix to get $\Lambda(G)$, we get a symmetric positive definite matrix. Thus $\det(\Lambda(G)) > 0$ for such graphs.

The proof is by induction on $|E(G)|$, the number of edges of G . The base case is an edgeless graph G , in which case all vertices have positive weight by hypothesis and M_G is a connected sum of lens spaces.

Now assume that $E(G) = n$, and all such chainmail graphs with fewer edges represent L -spaces. Choose an edge e of G , and delete the edge to get $G - e$ and contract the edge to get G/e , both of which have fewer edges. If there are parallel edges to e (with the same vertices), then G/e may have loops, but we may delete them in order to satisfy the induction hypothesis while still representing the same manifold (see Figure 4).

We may assume that e is not an isthmus, unless G is a tree (equivalently every edge is an isthmus). If G has at least two vertices with $\nu(v) > 0$, then we may choose e so that both components of $G - e$ satisfy the hypothesis that there is at least one vertex v with $\nu(v) > 0$. Otherwise, there is a single vertex v with $\nu(v) > 0$. Choose any leaf vertex v' with $\nu(v') = 0$ (a tree with nonempty edge set must have at least two leaf vertices, so such a vertex exists). Then we may delete v' keeping the same manifold $M_{G-v'}$ by the last move in Figure 4. Thus in the latter case $M_G \cong M_{G-v'}$ will be an L -space by induction, and thus we may assume we are in the former case in which case $G - e$ is a chainmail graph satisfying the hypotheses. We will assume that the endpoints of e are labeled v_1 and v_2 .

Claim. $\det \Lambda(G) = \det \Lambda(G - e) + \det \Lambda(G/e)$.

Proof. Consider $\Lambda(G)$ and $\Lambda(G - e)$ as linear operators on functions on $V(G) = V(G - e)$. We may write the symmetric matrices for these as

$$\Lambda(G) = \begin{pmatrix} w_{11} & w_{12} & x_1^T \\ w_{12} & w_{22} & x_2^T \\ x_1 & x_2 & A \end{pmatrix}, \quad \Lambda(G - e) = \begin{pmatrix} w_{11} - 1 & w_{12} + 1 & x_1^T \\ w_{12} + 1 & w_{22} - 1 & x_2^T \\ x_1 & x_2 & A \end{pmatrix},$$

differing only in the upper left block corresponding to vertices v_1 and v_2 . Since we are only interested in the determinant, we may do elementary row operations

keeping the determinants the same. Add the first row to the second row and first column to the second column (in either order) of $\Lambda(G)$, $\Lambda(G - e)$ to get matrices with the same determinant:

$$\Lambda(G)' = \begin{pmatrix} w_{11} & w_{11}+w_{12} & x_1^T \\ w_{11}+w_{12} & w_{11}+w_{22}+2w_{12} & x_1^T+x_2^T \\ x_1 & x_1+x_2 & A \end{pmatrix},$$

$$\Lambda(G-e)' = \begin{pmatrix} w_{11}-1 & w_{11}+w_{12} & x_1^T \\ w_{11}+w_{12} & w_{11}+w_{22}+2w_{12} & x_1^T+x_2^T \\ x_1 & x_1+x_2 & A \end{pmatrix}.$$

Then we see that

$$\Lambda(G)' - \Lambda(G-e)' = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

By the linearity of determinants (say with respect to the first row), we see that

$$\det(\Lambda(G)) - \det(\Lambda(G-e)) = \begin{vmatrix} 1 & 0 & 0 \\ w_{11}+w_{12} & w_{11}+w_{22}+2w_{12} & x_1^T+x_2^T \\ x_1 & x_1+x_2 & A \end{vmatrix} = \det(\Lambda(G/e)),$$

since the lower block matrix is $\Lambda(G/e)$.

This establishes the claim. One may also prove it using [AWvW, Theorem 4.4]. \square

Hence $|H_1(M_G)| = |H_1(M_{G-e})| + |H_1(M_{G/e})|$. By our induction hypothesis, M_{G-e} and $M_{G/e}$ are L -spaces, and they are related by an exact triangle. Then by [OS3, Proposition 2.1], M_G is also an L -space, where the knot is the crossing loop c . The proof is very similar to the proof of [OS4, Lemma 3.2, Proposition 3.3].

For the framed instanton homology case, the base case is the same (lens spaces are instanton L -spaces). We have $\dim I^\#(Y) \geq |H_1(Y)|$ by [Sca, Corollary 1.4]. There is an exact triangle (with $\lambda = 0$) given in [Sca, Section 7.5]. Applied to the case at hand, we get

$$I^\#(M_G) \rightarrow I^\#(M_{G-e}; \mu) \rightarrow I^\#(M_{G/e}) \rightarrow I^\#(M_G)$$

where μ is a loop representing the core of the Dehn filling of M_{G-e} on the crossing loop c (Figure 10), and similar exact sequences by permuting the role of μ for M_G and $M_{G/e}$. But in $M_G - c$, two out of the three slopes $\{0, -1, \infty\}$ will be nontrivial in $H_1(M_G - c; \mathbb{F}_2)$ and the third will be trivial by half-lives, half-dies. Filling along a homologically nontrivial slope, the homologically trivial slope will be isotopic to the core, and hence $[\mu] = 0 \in H_1(M_{G'}; \mathbb{F}_2)$, where G' is G , $G - e$, or




Figure 11. Generators for the mapping class group of a Heegaard surface Σ for S^3 .

G/e depending on the filling slope. So in this case we get an exact triangle

$$I^\#(M_G) \rightarrow I^\#(M_{G-e}) \rightarrow I^\#(M_{G/e}) \rightarrow I^\#(M_G).$$

Now we apply the same argument as for Heegaard Floer homology to conclude that

$$\begin{aligned} |H_1(M_G)| &\leq \dim I^\#(M_G) \leq \dim I^\#(M_{G-e}) + \dim I^\#(M_{G/e}) \\ &= |H_1(M_{G-e})| + |H_1(M_{G/e})| = |H_1(M_G)|, \end{aligned}$$

and hence M_G is an instanton L -space. □

Remark. Zhenkun Li pointed out that one can take a connected sum of any component of a negative alternating chainmail link with an L -space knot to get another L -space link. The point is that the proof goes through as before, with the vertex label adjusted larger, with the base case of the induction including a split link with one component an L -space knot so that surgery still gives an L -space for large enough coefficients.

One may give a slightly different characterization of some of the L -spaces described in Theorem 3.2 in terms of Heegaard splittings. Consider a set of generators $\alpha_i, \beta_i, \delta_i$ for the mapping class group $\text{Mod}(\Sigma)$ of the Heegaard surface Σ by right-hand Dehn twists about curves a_i, b_i, d_i respectively shown in Figure 11, which represents a standard Heegaard splitting $H_1 \cup_\Sigma H_2$ of S^3 .

Corollary 3.3. *Choose $\sigma \in \{\alpha_i, \beta_i^{-1}, \delta_i\}^*$, the submonoid of products of these elements. Then the resulting manifold $H_1 \cup_\sigma H_2$ is a connected sum of an L -space and $\#^k S^2 \times S^1$ for some k .*

Proof. The Kirby diagram produced from the products of these generators by [MP, Theorem 7.1] is a negative alternating chainmail link corresponding to a chainmail graph with nonnegative vertex weights (see Figure 3 for an example). In each component of the chainmail graph with vertex weights 0, there will be an $S^2 \times S^1$ summand connected sum with a double-branched cover of an alternating link in the corresponding component of the Kirby diagram. Hence the manifold will be of the claimed type in each connect summand, and therefore in total. □

4. Augmented negative alternating chainmail links are generalized L -space links

Next we consider augmented chainmail links.

Definition 4.1. [Liu, Definition 2.9] Let $L = \cup_{i=1}^l L_i \subset S^3$ be a link, and for each i choose a sign $\epsilon_i \in \{\pm\}$. Then L is called a *generalized (instanton) $(\epsilon_1 \cdots \epsilon_l)$ L -space link* if there exist integers p_1, \dots, p_l such that $S_{\epsilon_1 n_1, \dots, \epsilon_l n_l}^3(L)$ is a (instanton) L -space for all n_1, \dots, n_l with $n_i > p_i$ for all $1 \leq i \leq l$. If we do not specify the sign, then L is denoted a generalized L -space link.

Theorem 4.2. *Partially augmented negative alternating chainmail links are generalized L -space links with a $+$ sign associated to each vertex loop and $-$ to each edge (crossing) loop.*

We answer a question of Yajing Liu [Liu, Question 1.9]:

Corollary 4.3. *Any closed orientable connected 3-manifold is obtained by surgery on a generalized L -space link.*

Proof. By Theorem 2.2 any 3-manifold is obtained by Dehn surgery on a chainmail link, which in turn arises from surgery on a flat fully augmented chainmail link. By Theorem 4.2, such a link is a generalized L -space link. \square

Proof of Theorem 4.2. First, we need a lemma that allows us to relate the homologies of manifolds in a Floer exact triangle.

Lemma 4.4. *Consider a partially augmented chainmail link simplicial graph G , with negative weights on the vertices $V_c(G) \subset V(G)$ corresponding to crossing loops and nonnegative weights on the other vertices $V(G) - V_c(G)$. Moreover, the edges between vertices of $V(G) - V_c(G)$ have weight -1 . Then the sign of the determinant of $\Lambda(G)$ is $(-1)^{|V_c(G)|}$ if nonzero.*

Proof. The proof is by induction on the sum of the negative crossing loop weights. If there are no crossing loops, then $\Lambda(G)$ is the linking matrix of a negative alternating chainmail graph, and hence the sign of the determinant is nonnegative by the proof of Theorem 3.2, establishing the base case.

Now consider a single crossing loop of the flat augmented link with weight $-c$, $c > 0$, and its two adjacent loops having framings a' and b' . The matrix $\Lambda(G)$ looks like

$$\begin{pmatrix} -c & 1 & -1 & 0 \\ 1 & a' & 0 & * \\ -1 & 0 & b' & * \\ 0 & * & * & A \end{pmatrix}.$$

When $c = 1$, we do column operations adding the first column to the second and

subtracting the first column from the third to get

$$\begin{vmatrix} -1 & 0 & 0 & 0 \\ 1 & a'+1 & -1 & * \\ -1 & -1 & b'+1 & * \\ 0 & * & * & A \end{vmatrix} = (-1) \begin{vmatrix} a'+1 & -1 & * \\ -1 & b'+1 & * \\ * & * & A \end{vmatrix}.$$

The last determinant is that of a partially augmented chainmail link by doing a -1 twist along the crossing loop (see the left diagrams of Figure 10). So the sign of this determinant is $(-1)^{(|V_c(G)|-1)}$ by induction. Thus it is true by induction for $c = 1$.

When $c > 1$, we have

$$\begin{vmatrix} -c-1 & 1 & -1 & 0 \\ 1 & a' & 0 & * \\ -1 & 0 & b' & * \\ 0 & * & * & A \end{vmatrix} = \begin{vmatrix} -c & 1 & -1 & 0 \\ 1 & a' & 0 & * \\ -1 & 0 & b' & * \\ 0 & * & * & A \end{vmatrix} - \begin{vmatrix} 1 & 0 & 0 & 0 \\ 1 & a' & 0 & * \\ -1 & 0 & b' & * \\ 0 & * & * & A \end{vmatrix}.$$

The matrix on the right has determinant

$$\begin{vmatrix} a' & 0 & * \\ 0 & b' & * \\ * & * & A \end{vmatrix}$$

which is $\det \Lambda(G - v)$. By induction, this matrix has the opposite sign, so we get by induction on c that the sign with weight $-c - 1$ is $(-1)^{|V_c(G)|}$. \square

Returning to the proof of 4.2, let v be a vertex of the weighted graph G corresponding to a partially augmented link. Denote by G the graph with the vertex v with weight $v(v) = -c$ and by G' the graph with the weight $v(v) = -c - 1$. Assume that all vertices of $V(G) - V_c(C)$ have positive weight. The claim is that M_G is an L -space. We prove this by induction on c and on $|V_c(G)|$. If $c = 0$, then this corresponds to the graph $G - v$, and hence by induction on $|V_c(G)|$ this is an L -space. Suppose that M_G is an L -space. There is an exact triangle of Floer homology of $M_G, M_{G'}$ and M_{G-v} . Moreover, we have

$$\det \Lambda(G') = \det \Lambda(G) - \det \Lambda(G - v).$$

By Lemma 4.4, $\text{sgn}(\det \Lambda(G')) = \text{sgn}(\det \Lambda(G)) = -\text{sgn}(\det \Lambda(G - v))$, so $|H_1(M'_G)| = |H_1(M_G)| + |H_1(M_{G-v})|$. It follows by induction that $|H_1(M_{G'})| = \text{rank} \widehat{HF}(M_{G'})$, and we see that $M_{G'}$ is an L -space. Again by induction, the augmented chainmail link associated to G is a generalized L -space link. The framed instanton case is similar to the argument at the end of the proof of Theorem 3.2. \square

5. Nonorderable positive fillings on negative alternating chainmail links

We are not able to show that the L -spaces described in Theorem 4.2 have nonorderable fundamental group, addressing [BGW, Conjecture 1], but we can show nonorderability in special cases.

Theorem 5.1. *For a chainmail graph G with edge weights -1 , and nonnegative vertex weights with at most one vertex weight nonzero, $\pi_1(M_G)$ is nonorderable.*

Proof. The proof is modeled on the proof of Greene that double-branched covers of alternating links have nonorderable fundamental group [Gre2, Theorem 2.1] (proved originally by Boyer, Gordon and Watson [BGW, Theorem 4]).

Each edge of G with a chosen orientation corresponds to an element of the fundamental group of $S^3 - L_G$, as shown in Figure 12, bottom left. If a vertex v_i has weight $\nu(v_i) = 0$, then the relation associated to Dehn filling will be the product of the incoming edges (see Figure 13).




Figure 12. Top: a planar graph. Bottom left: loops associated to oriented edges. Bottom right: acyclic orientation from group ordering.




Figure 13. The zero-framing loop (left) is a product of incoming edge loops.

Suppose that $\pi_1(M_G)$ admits a left ordering. Then we may orient each edge so that the associated element is > 1 (assuming all of these elements are nontrivial first). This gives an acyclic orientation of the graph, seen in Figure 12, bottom right; any oriented loop corresponds to the trivial element but is also a product of elements > 1 , a contradiction. Then there must be a sink vertex and a source vertex. Since only one vertex v has weight $\nu(v) > 0$ by assumption, we may choose a sink or source vertex v' with $\nu(v') = 0$. In this case, the product of the edges oriented into this will be trivial but also > 1 (or < 1), a contradiction.

In the case that some of the edges are trivial, we may still locate a sink or source vertex, allowing that some of the adjacent edges are trivial. If all edges are trivial, then either all vertex weights are zero, and we are in the case considered by Greene [Gre2], or one vertex has nonzero weight. In this case the relator from Dehn filling will give a power of the meridian of that loop is trivial. If the meridian is torsion, then the group cannot be orderable. If the meridian is trivial, the one can see that all meridians are trivial (since the edge loops are products of meridians), and hence the fundamental group is trivial. But by definition the trivial group is not left-orderable. \square

A theorem of Tao Li implies that a 3-manifold Y with Heegaard genus 2 that is an L -space does not have orderable fundamental group [Li, Theorem 1.2]. One may apply this to conclude that the manifolds in Corollary 3.3 whose gluing map lies in the monoid $\{\alpha_1, \alpha_2, \beta_1^{-1}, \beta_2^{-1}, \delta_2\}^*$ do not have orderable fundamental group unless it is a connected sum of copies of $S^1 \times S^2$. Such manifolds are realized by surgery on the framed links made by composing the tangles in Figure 2, such as in Figure 3.

6. Asymmetric L -space links

It was shown that there exist asymmetric L -spaces and asymmetric L -space knots [DHL; BL, Theorem 1.2]. Here we show the existence of asymmetric L -space links.

Theorem 6.1. *There are alternating hyperbolic chainmail links of arbitrarily many components which have trivial symmetry group.*

Proof. Start with a flat augmented chainmail link associated to a graph G such as in Figure 8 with hyperbolic complement (so no parallel crossing loops). Perform $-1/n_i$ surgeries with n_i large on each of the crossing loops c_i to make a negative alternating chainmail link. By choosing the n_i sufficiently large, we may assume that the cores of the Dehn fillings on each of the crossing loops c_i give the $|E(G)|$ shortest loops in the hyperbolic metric on the manifold, and they have distinct lengths [Thu, Theorem 5.8.2]. Thus, any isometry of this link complement must preserve the Margulis tubes about these short loops. Thus it also preserves the




Figure 14. A flat augmented chainlink with two crossing disks must have a separating edge.

complements of these loops which is the initial flat augmented chainmail link complement. By [MT], any homeomorphism of this link complement must be induced by a homeomorphism of the link. Moreover, cusps are taken to cusps, and hence cusps corresponding to crossing loops are fixed.

Choose a simplicial planar graph G which contains no triangles and no vertices of degree 2 and no edge that separates the graph. By [MT, Proposition 2.7], the only thrice punctured spheres in the fully augmented chainmail link complement associated to G are crossing disks. Moreover, each crossing circle bounds a unique twice-punctured disk intersecting two knot loops (one can see that for a flat augmented chainmail link that a second crossing circle would give a separating edge; see Figure 14). We may also assume that these have a unique reflection surface, since our assumptions rule out the examples in that paper with multiple reflection surfaces [MT, Theorem 1.2]. Thus, any isometry of $S^3 - L$ must send the reflection surface to the reflection surface and the crossing disks to crossing disks. Each loop in the reflection surface bounds a disk punctured by crossing loops and intersecting the crossing disks in arcs. One may recover the graph G from this pattern of intersections, as well as its planar embedding from the reflection surface. So the only orientation preserving symmetries of the link complement must preserve the graph fixing each edge and hence are trivial. There is an orientation reversing symmetry by reflecting in the reflection surface, but this will not preserve the slopes of the fillings of the crossing loops, and hence does not extend to a symmetry of the alternating link. Thus we have shown the existence of negative alternating chainmail links which admit no nontrivial symmetry. \square

Remark. Further Dehn fillings on asymmetric negative alternating chain links give rise to more examples of asymmetric L -spaces. Compare to [DHL, Theorem 1.2] and [ABGK+] which require computation to verify the results, and [BL] which uses an intricate construction to give asymmetric L -space knots.

7. Alternating chainmail links are fibered

It was proved that L -space knots are fibered [Ni1]; in fact, all of the examples in the literature of L -space links have fibered complement. Example 3.9 of [Liu] shows that the links do not admit fibered Seifert surfaces; nevertheless the complements fiber.

Theorem 7.1. *Nonsplit negative alternating chainmail links have fibered complement.*

Proof. Take the checkerboard surface for the alternating chainmail link L_G which is a Seifert surface by applying Seifert's algorithm to the orientation of the link in which all of the loops are oriented counterclockwise (Figure 15). The complement of this surface is a sutured handlebody (H, γ) , where γ will be a union of curves parallel to the boundary components of the Seifert surface. Gabai proved that (H, γ) is disc-decomposable [Gab, Theorem 6.1]. Moreover, one sees that the restriction maps $H^1(S^3 - L_G) \rightarrow H^1(H)$ is surjective. Choose a cutset of oriented disks $D \subset H$ giving a taut sutured decomposition, and a cohomology class $\alpha' \in H^1(H)$ Poincaré dual to $[D, \partial D] \subset H_2(H, \partial H)$. Then there exists a cohomology class $\alpha \in H^1(S^3 - L_G)$ so that $\alpha|_H = \alpha'$. Let $\mu \in H^1(S^3 - L_G)$ denote the cohomology class dual to the Seifert surface. Then $k\mu + \alpha$ will lie in a cone over a face of the Thurston norm ball adjacent to that containing μ for $k \in \mathbb{N}$ large enough. By [AZ, Theorem 1.4] the guts $\Gamma(k\mu + \alpha)$ is obtained from $\Gamma(\mu)$ by decomposing along α' represented by D . (See [AZ, Definition 2.17] for the notion of *guts*.) Since (H, γ) is disc-decomposable, this is a product sutured manifold, and hence $k\mu + \alpha$ is a fibered class. \square




Figure 15. A Seifert surface for a chainmail link.

8. Conclusion

We conclude with some questions and challenges stemming from the examples in this paper.

- (1) Are any of the new examples of L -spaces described in this paper strong L -spaces [GL]?
- (2) Are there nonfibered L -space link complements?
- (3) Describe all L -space surgeries on alternating chainmail and augmented chainmail links. Since this includes all 3-manifolds, such a characterization will include a description of all L -spaces and thus is likely to be intricate [GN3].
- (4) Find the foliar surgeries on these links. If the L -space conjecture is true, then this should consist of the complement of the reducible and L -space fillings [Juh, Conjecture 5].
- (5) Prove that the L -space surgeries on alternating and augmented chainmail links described in these theorems have nonorderable fundamental group (Question (1) is relevant if there is a positive answer by [LL]). If the L -space conjecture is true, then one would expect the fundamental groups to be nonorderable.
- (6) Presumably, one could prove Theorem 7.1 by computing the multivariable Alexander polynomial of negative alternating chainmail links, which allows one to detect if the link complement fibers [Ni1, Corollary 1.2; Ni2; OS5, Theorem 1.2]. It might be interesting to find a formula for these polynomials and use it to reprove Theorem 7.1. It would also be interesting to know which augmented alternating links have fibered complement?
- (7) Are there more L -space links that may be found using this approach? Which chainmail links are (generalized) L -space links? Complete this analogy:
 alternating links: alternating chainmail links :: quasi-alternating links: ?

Acknowledgements

We thank Qiuyu Ren, Joshua Greene, Eugene Gorsky, Zhenkun Li and Francesco Lin for helpful comments.

We also thank the referee for many helpful comments. After this paper appeared on the arxiv, Antonio Alfieri contacted us to say that he had discovered some of these results [Alfieri].

References

- [ABGK+] C. Anderson, K. L. Baker, X. Gao, M. Kegel, K. Le, K. Miller, S. Onaran, G. Sangston, S. Tripp, A. Wood, and A. Wright, “ L -space knots with tunnel number > 1 by experiment”, preprint, 2019. arXiv 1909.00790

- [Alfieri] A. Alfieri, “ L -space surgeries on links coming from weighted planar graphs”, unpublished manuscript (8 pp.), 2017.
- [AWvW] F. Aliniaefard, V. Wang, and S. van Willigenburg, “Deletion-contraction for a unified Laplacian and applications”, preprint, 2021. arXiv 2110.13949
- [AZ] I. Agol and Y. Zhang, “Guts in Sutured Decompositions and the Thurston Norm”, preprint, 2022. arXiv 2203.12095
- [BGW] S. Boyer, C. M. Gordon, and L. Watson, “On L -spaces and left-orderable fundamental groups”, *Math. Ann.* **356**:4 (2013), 1213–1245. MR
- [BL] K. L. Baker and J. Luecke, “Asymmetric L -space knots”, *Geom. Topol.* **24**:5 (2020), 2287–2359. MR
- [BS] J. A. Baldwin and S. Sivek, “Instantons and L -space surgeries”, *J. Eur. Math. Soc.* **25**:10 (2023), 4033–4122. MR
- [DHL] N. M. Dunfield, N. R. Hoffman, and J. E. Licata, “Asymmetric hyperbolic L -spaces, Heegaard genus, and Dehn filling”, *Math. Res. Lett.* **22**:6 (2015), 1679–1698. MR
- [EN] D. Eisenbud and W. Neumann, *Three-dimensional link theory and invariants of plane curve singularities*, Annals of Mathematics Studies **110**, Princeton University Press, 1985. MR
- [FS] R. Fintushel and R. J. Stern, “Constructing lens spaces by surgery on knots”, *Math. Z.* **175**:1 (1980), 33–51. MR
- [Gab] D. Gabai, “Foliations and genera of links”, *Topology* **23**:4 (1984), 381–394. MR
- [GL] J. E. Greene and A. S. Levine, “Strong Heegaard diagrams and strong L -spaces”, *Algebr. Geom. Topol.* **16**:6 (2016), 3167–3208. MR
- [GN1] E. Gorsky and A. Némethi, “Lattice and Heegaard Floer homologies of algebraic links”, *Int. Math. Res. Not.* **2015**:23 (2015), 12737–12780. MR
- [GN2] E. Gorsky and A. Némethi, “Links of plane curve singularities are L -space links”, *Algebr. Geom. Topol.* **16**:4 (2016), 1905–1912. MR
- [GN3] E. Gorsky and A. Némethi, “On the set of L -space surgeries for links”, *Advances in Mathematics* **333** (2018), 386–422.
- [Gre1] J. E. Greene, “A spanning tree model for the Heegaard Floer homology of a branched double-cover”, *J. Topol.* **6**:2 (2013), 525–567. MR
- [Gre2] J. E. Greene, “Alternating links and left-orderability”, *Proc. Amer. Math. Soc.* **146**:6 (2018), 2707–2709. MR
- [hc] L. Culler, “Presenting 3-manifolds by planar graphs”, MathOverflow posting, 2018, available at <https://mathoverflow.net/q/315927>.
- [Juh] A. Juhász, “A survey of Heegaard Floer homology”, pp. 237–296 in *New ideas in low dimensional topology*, Ser. Knots Everything **56**, World Scientific, Hackensack, NJ, 2015. MR
- [KM] P. Kronheimer and T. Mrowka, “Knots, sutures, and excision”, *J. Differential Geom.* **84**:2 (2010), 301–364. MR
- [KR] W. H. Kazez and R. Roberts, “ C^0 approximations of foliations”, *Geom. Topol.* **21**:6 (2017), 3601–3657. MR
- [Li] T. Li, “Taut foliations of 3-manifolds with Heegaard genus 2”, *Duke Math. J.* **173**:8 (2024), 1427–1475. MR
- [Liu] Y. Liu, “ L -space surgeries on links”, *Quantum Topol.* **8**:3 (2017), 505–570. MR
- [LL] A. S. Levine and S. Lewallen, “Strong L -spaces and left-orderability”, *Math. Res. Lett.* **19**:6 (2012), 1237–1244. MR

- [MP] S. Matveev and M. Polyak, “A geometrical presentation of the surface mapping class group and surgery”, *Comm. Math. Phys.* **160**:3 (1994), 537–550. MR
- [MT] C. Millichap and R. Trapp, “Flat fully augmented links are determined by their complements”, *Algebr. Geom. Topol.* **25**:8 (2025), 4839–4896. MR
- [Ni1] Y. Ni, “Knot Floer homology detects fibred knots”, *Invent. Math.* **170**:3 (2007), 577–608. MR
- [Ni2] Y. Ni, “Erratum: Knot Floer homology detects fibred knots”, *Invent. Math.* **177**:1 (2009), 235–238. MR
- [OS1] P. Ozsváth and Z. Szabó, “Holomorphic disks and genus bounds”, *Geom. Topol.* **8** (2004), 311–334. MR
- [OS2] P. Ozsváth and Z. Szabó, “Holomorphic disks and three-manifold invariants: properties and applications”, *Ann. of Math. (2)* **159**:3 (2004), 1159–1245. MR
- [OS3] P. Ozsváth and Z. Szabó, “On knot Floer homology and lens space surgeries”, *Topology* **44**:6 (2005), 1281–1300. MR
- [OS4] P. Ozsváth and Z. Szabó, “On the Heegaard Floer homology of branched double-covers”, *Adv. Math.* **194**:1 (2005), 1–33. MR
- [OS5] P. Ozsváth and Z. Szabó, “Link Floer homology and the Thurston norm”, *J. Amer. Math. Soc.* **21**:3 (2008), 671–709. MR
- [Pol] M. Polyak, “From 3-manifolds to planar graphs and cycle-rooted trees”, lecture notes, 2014, available at <https://polyak.net.technion.ac.il/files/2021/05/From-3-manifolds-to-planar-graphs-and-cycle-rooted-trees.pdf>.
- [Sca] C. W. Scaduto, “Instantons and odd Khovanov homology”, *J. Topol.* **8**:3 (2015), 744–810. MR
- [Sch] A. Scharf, “Zur Faserung von Graphenmannigfaltigkeiten”, *Math. Ann.* **215** (1975), 35–45. MR
- [Thu] W. P. Thurston, “The geometry and topology of 3-manifolds”, lecture notes from 1978–80, Princeton University, 1980, available at <https://library.slmath.org/nonmsri/gt3m>.

Received July 20, 2023. Revised June 22, 2025.

IAN AGOL
DEPARTMENT OF MATHEMATICS
UNIVERSITY OF CALIFORNIA, BERKELEY
BERKELEY, CA
UNITED STATES
ianagol@math.berkeley.edu

MEKLER'S CONSTRUCTION AND MURPHY'S LAW FOR 2-NILPOTENT GROUPS

BLAISE BOISSONNEAU, ARIS PAPADOPOULOS AND PIERRE TOUCHARD

Mekler's construction is a powerful technique for building purely algebraic structures from combinatorial ones. Its power lies in the fact that it allows various model-theoretic tameness properties of the combinatorial structure to transfer to the algebraic one. In this paper, we push this ideology much further, describing a broad class of properties that transfer through Mekler's construction. This technique subsumes many well-known results and opens avenues for many more.

As a straightforward application of our methods, we obtain transfer principles for stably embedded pairs of Mekler groups and construct strictly NFOP_k pure groups for all $k \in \mathbb{N}_{>2}$. We also answer a question of Chernikov and Hempel on transfer of burden.

1. Introduction	220
2. Preliminaries	223
2.1. Mekler's construction	223
2.2. Generalised indiscernibles	226
2.3. Model-theoretic dividing lines	228
3. Relative quantifier elimination	230
3.1. First reduction: from G to $\mathcal{F}(G)$	231
3.2. Second reduction: From $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta\}$ to $(\mathcal{V}, +, 0, R)$	234
3.3. Third reduction: From $(\mathcal{V}, +, 0, R)$ to C	237
4. Transfer principles	248
4.1. Model completeness and Stable embeddedness	248
4.2. Characterisation of indiscernibles	258
4.3. Transfers of dividing lines	264
4.4. Resplendence	266
5. Second proof of Theorem A	267
6. Open questions	271
Acknowledgements	272
References	272

Papadopoulos was partially supported through a Leeds Doctoral Scholarship from the University of Leeds. Touchard was partially supported by KU Leuven through IF C16/23/010.

MSC2020: primary 03C10, 03C45, 03C60; secondary 03C50, 20F18.

Keywords: Mekler's construction, nilpotent groups of class 2, model theory, relative quantifier elimination, Shelah's classification.

1. Introduction

Tame classes of groups. Model theorists have a good understanding of the definable structure of modules and, in particular, of abelian groups. It is well known that any module is *stable*, i.e., it does not “code” a linear order. In broader terms, this means that definable sets in pure abelian groups are “combinatorially tame”. It could be reasonable to think that other natural classes of groups are also “tame”. However, as shown by Mekler [Mek81], already the class $\mathbb{G}_{2,p}$ of 2-nilpotent groups of prime exponent p — certainly the easiest class of nonabelian groups that one can think of — can witness all sorts of model-theoretic behaviours; this was extended by many subsequent works [Bau02; CH19; Ahn20], many of which are directly related to Shelah’s classification programme.

Classification theory. An essential aspect of Shelah’s classification theory is, very roughly speaking, based on the idea that the presence/absence of simple combinatorial data (e.g. linear orders, random graphs, and more) can give us a lot of information on the “complexity” of a given theory (e.g. stability, NIP, and more).

The dichotomies given by the presence/absence of combinatorial configurations allow model theorists to divide (hence the name *dividing lines*) the class of all first-order theories into smaller regions, which can then be studied in more detail. Since the various regions of the model-theoretic universe are defined using combinatorial data, it is often the case that building combinatorial examples inhabiting each region is a much easier task than building *purely* algebraic ones.

Mekler’s construction is now a classical technique for building a purely algebraic structure (a 2-nilpotent group) from a purely combinatorial one (a “nice” graph). We will review the basic ideas of Mekler’s construction in Section 2.1, but for now, the reader should keep in mind that this construction is done in a way that has been shown to preserve the presence (and absence) of various kinds of combinatorial data, such as, for instance, λ -stability [Mek81], nonindependence property (of any arity $n \in \mathbb{N}$), the tree property of the second kind [CH19]. See Fact 4.1 for an exhaustive (to our knowledge) list of properties preserved by Mekler’s construction.

Hodges somewhat ironically referred to this fact as *Murphy’s law for 2-nilpotent groups*. Far from being a negative fact, Murphy’s law for 2-nilpotent groups justifies the rather reassuring hope that even the most subtle combinatorial property, if satisfied by an ad-hoc combinatorial example, is probably also satisfied by a pure 2-nilpotent group.

Main results of the paper. Our paper aims to push the ideas discussed in the preceding paragraphs further by developing a method that allows us to prove general transfer principles for groups obtained via Mekler’s construction. In a precise sense, we show that Mekler’s construction preserves the presence/absence of

all interesting combinatorial data characterised through generalised indiscernibles (all these concepts will be explained in detail later in the paper). Let us call a *Mekler group* any group elementarily equivalent to a group obtained by Mekler's construction. The first main result of this paper is the following:

Theorem A. *Let \mathcal{I} be a Ramsey structure, and \mathcal{J} a reduct of \mathcal{I} . Let M be a Mekler group and C its associated nice graph. The following conditions are equivalent:*

- (1) C collapses \mathcal{I} -indiscernibles (resp. to \mathcal{J} -indiscernibles).
- (2) $M(C)$ collapses \mathcal{I} -indiscernibles (resp. to \mathcal{J} -indiscernibles).

Ramsey structures¹ are suitable indexing structures, and in particular, indiscernible sequences indexed by a Ramsey structure satisfy desirable properties similar to (ordered)-indiscernible sequences. A structure \mathcal{M} collapses \mathcal{I} -indiscernibles if all \mathcal{I} -indiscernible sequences in \mathcal{M} are indiscernible in a strictly stronger way (see Definition 2.20). In particular, this theorem shows that the class $\mathbb{G}_{2,p}$ does not lie in the tame side of any dividing lines characterised by the collapsing of generalised indiscernibles.

This result is proved in Section 4.3. As mentioned, Theorem A generalises many previously known results, and we briefly explain this in Section 4.3. In this paper, a novel use of this theorem is the construction of examples of NFOP_k pure groups for all $k \in \mathbb{N}_{>2}$, confirming an expectation of Abd Aldaim, Conant, and Terry [AACT25, Remark 2.13].

We also take the occasion to give a negative answer to a problem of Chernikov and Hempel [CH19, Problem 5.8] in Section 4.4.2, where we observe that if C is any infinite nice graph, then its Mekler group $M(C)$ has arbitrarily large finite burden.

The method we developed also allows us to study pairs of Mekler groups and, in particular, to state a transfer principle for *stably embedded* pairs of models. A stably embedded pair $(\mathcal{M}, \mathcal{M}')$ consists of an elementary extension $\mathcal{M} \preceq \mathcal{M}'$ where all types over \mathcal{M} realised in \mathcal{M}' are definable over \mathcal{M} .

Theorem B. *Let $M \preceq M'$ be an elementary extension of Mekler groups of respective graphs C and C' .*

- (1) M is stably embedded in M' if and only if C is stably embedded in C' .
- (2) M is uniformly stably embedded in M' if and only if C is uniformly stably embedded in C' .

¹For a precise definition, see e.g. [Neš05].

This theorem is proved in Section 4.1. To our knowledge, this is the first transfer principle for Mekler groups for a property that depends on the model itself and not just on its theory. This suggests that other model-theoretic properties of Mekler groups, depending on the models, can be reduced to the study of their graphs.

To prove Theorem A and Theorem B, we give a new *relative quantifier elimination* result from the Mekler group $M(C)$ down to its nice graph C (modulo finite imaginaries), which we believe is of independent interest. This occupies all of Section 3, and is then used in Section 4 to prove the main theorems.

We also give an alternative and faster proof of Theorem A in Section 5, under the additional (but mild) hypothesis of *specific collapsing* (see Definition 2.20). The two proofs are very different in nature, and we consider them both to be of interest. Finally, in Section 6, we conclude with open questions about transfers for Mekler groups, which we think could be answered by an adaptation of our methods.

Related results. The results and methods presented in [dMRS25], which build in part on the work of Baudisch [Bau02], represent a significant breakthrough in the model-theoretic study of nilpotent groups of finite exponent. The authors successfully employ the Lazard correspondence as a model-theoretic bridge between the theory of Lie algebras and the study of nilpotent groups of prime exponent. Among several classification results, they show that $\mathbf{G}_{c,p}$, the Fraïssé limit of the class of Lazard groups of exponent p and nilpotency class c equipped with predicates for the Lazard series, is a strictly NIP_c pure group. Furthermore, they show that $\mathbf{G}_{2,p}$ is NFOP_2 (see [dMRS25, §3.3]).

Our approach in the present paper is rather different, as we restrict to the class of 2-nilpotent groups, and do not fix any particular complete theory of groups. However, the complementary points of view with [dMRS25] may suggest further development: can the various results concerning $\mathbf{G}_{2,p}^L$ be encapsulated in a more general transfer principle? Or, can Mekler’s construction be adapted to groups of higher nilpotency class? We do not address these questions in this paper, leaving them for potential (ambitious) future work.

Notation. Throughout this paper, we assume familiarity with first-order logic and basic model theory (types, saturation, monster models). All of this material can be found in [TZ12, Chapters 1-5]. Our notation is either standard or explained. We denote by \mathcal{L}_{grp} the language of groups with inverse $\{\cdot, {}^{-1}, 1\}$, often identified with $\{+, -, 0\}$, in the context of abelian groups. Not all functions mentioned in the remainder of the paper will be defined everywhere; to remedy this, for notational convenience, languages include by default a constant symbol u , standing for “undetermined”. This will be used (in the obvious way) for functions whose actual domain is smaller than the base set.

2. Preliminaries


To keep this paper self-contained, we will start our preliminaries by giving a brief exposition of Mekler's construction in Section 2.1. This construction originates from [Mek81]. We will then briefly review some basic facts around generalised indiscernibles and Ramsey structures in Section 2.2. Specific examples of generalised indiscernibles (and how collapses thereof characterise model-theoretic dividing lines) will also be discussed in more detail in Section 2.3.

2.1. Mekler's construction. A standard reference for Mekler's construction is [Hod93, Appendix A.3], and we will closely follow the notation and terminology given there. Several refinements (and some corrections) to statements of [Hod93] were given by Chernikov and Hempel in [CH19, Section 2], and we will include some of them here as well. This section is purposefully laconic, and we refer the reader to both [Hod93] and [CH19] for more information and proofs.

Definition 2.1. A graph² $C = (V, E)$ is called *nice* if it satisfies the following properties:

- (1) $|V| \geq 2$.
- (2) For any two distinct vertices $v_1, v_2 \in V$ there is some vertex $v \in V \setminus \{v_1, v_2\}$ such that $\{v_1, v\} \in E$ and $\{v_2, v\} \notin E$.
- (3) There are no triangles or squares in the graph.

Example 2.2. Natural examples that we will consider are infinite, but it can be useful to think of finite ones. The simplest, and only, example with at most 5 vertices is the pentagon, and can be used as a building block to build larger finite examples.



Conditions (1)–(3) above are mild conditions on the graph, in the following sense:

Fact 2.3. *Any structure in a finite relational language is bi-interpretable with a nice graph.*

²A *graph* is a structure in a relational language with a single binary relation E which is assumed to be irreflexive and symmetric.

Many proofs exist of this well-known fact (for example in [Hod93, Theorem 5.5.1, Exercise 5.5.9]), so we have opted not to include our own.

For any graph C and odd prime p , we write $M(C)$ for the 2-nilpotent group of exponent p which is generated freely in the variety of 2-nilpotent groups of exponent p by the vertices of C , with only relations those imposing that two generators commute if and only if they are connected by an edge in C .

Definition 2.4. If C is a nice graph, we call $M(C)$ *the Mekler group of C* . More generally, we say that a group is *a Mekler group* if it is elementarily equivalent to a group $M(C)$, for a nice graph C .

Example 2.5. Consider the pentagon P , and the Mekler group $G := M(P)$. One can show that G is a special p -group, with centre $Z := C_p^5$ and quotient $G/Z := C_p^5$, where C_p is the cyclic group of p elements.

An axiomatisation of the theory of Mekler groups can be found in [Hod93, Appendix A.3]. We now recall some standard terminology.

For the remainder of this section, fix a nice graph C and let $M(C)$ be the 2-nilpotent group of exponent p , constructed as above. We write Z for the centre of $M(C)$.

Definition 2.6 (equivalence relations \sim and \approx). We define the following two equivalence relations on $M(C)$:

- (1) *Centraliser*: $g \sim h$ if and only if $C(g) = C(h)$, where $C(g)$ denotes the centraliser of g .
- (2) *Powers modulo centre*: $g \approx h$ if and only if there is some $c \in Z$ and some $\alpha \in \{1, \dots, p-1\}$ such that $h = g^\alpha c$.

Remark 2.7 [Hod93, Lemma A.3.3]. The equivalence relation \approx refines \sim , that is, for all $g, h \in M(C)$, if $g \approx h$ then $g \sim h$.

Definition 2.8 (types and isolation). Let $g \in M(C) \setminus Z$ and $q \in \mathbb{N}$.

g is said to be of *type q* if $[g]_{\sim}$ splits into exactly q many \approx -classes.

g is called *isolated* if every noncentral element that commutes with g is \approx -equivalent to g .

If g is of type $q \in \mathbb{N}$, it is said to have *type q^l* if it is isolated and *type q^v* otherwise.

By convention, central elements are excluded from this definition.

Remark 2.9. For all $q \in \mathbb{N}$, each of the sets of elements of type q , q^l and q^v in $M(C)$ are \emptyset -definable, since \sim and \approx are \emptyset -definable.

Fact 2.10 [Hod93, Lemmas A.3.2, A.3.6–A.3.10].

- (1) *Every element of $M(C)$ can be written in the form $a_0^{\alpha_0} \cdots a_{m-1}^{\alpha_{m-1}} z$ where $m \in \mathbb{N}$, $a_i \in C$ distinct, $\alpha_i \in \{1, \dots, p-1\}$ and $z \in Z$.*

- (2) *The derived subgroup $[M(C), M(C)]$ is the centre Z .*
- (3) *Elements of the form $a^\alpha z$ where $a \in C$, $\alpha \in \{1, \dots, p-1\}$, $z \in Z$, are of type 1^ν .*
- (4) *Elements of the form $a_0^{\alpha_0} a_1^{\alpha_1} z$ where $\alpha_i \in \{1, \dots, p-1\}$, $a_i \in C$ are distinct and connected, are of type $p-1$.*
- (5) *Elements of the form $g := a_0^{\alpha_0} \cdots a_{m-1}^{\alpha_{m-1}} z$ where $m \geq 2$, $\alpha_i \in \{1, \dots, p-1\}$, and a_i distinct and connected to a **unique** vertex $a \in C$, are of type p . This vertex a is called the **handle** of g .*
- (6) *Elements of the form $a_0^{\alpha_0} \cdots a_{m-1}^{\alpha_{m-1}} z$ with distinct vertices $a_i \in C$ and where no vertex in C is connected to all a_i 's are of type 1^l .*

In particular, noncentral elements of $M(C)$ can only be of type p^ν , $(p-1)^\nu$, 1^ν or 1^l , thus, we will suppress the superscript ν , when discussing elements of type p^ν or type $(p-1)^\nu$. Moreover, elements of type $p-1$ are exactly those elements that can be written as the product of two \sim -inequivalent and connected elements of type 1^ν .

Notation 2.11. We denote by \mathfrak{E}^ν , \mathfrak{E}^p and \mathfrak{E}^l the set of elements of type respectively 1^ν , p and 1^l . Those subsets of $M(C)$ are definable in the language of groups.

Definition 2.12 (graph Γ). Let X be a subset of a group G , such that X is closed under the equivalence relation \sim . We define $\Gamma(X)$ to be the set X/\sim equipped with a graph relation E given by $[a]_{\sim} E [b]_{\sim}$ if and only if a and b commute in G .

By [Hod93, Theorem A.3.10], the graph C is interpreted by $(\Gamma(\mathfrak{E}^\nu), E)$, which is an imaginary sort of M . The main goal of this paper is to describe (generalised) indiscernible sequences in a Mekler group (recall that, in this paper, a Mekler group is a group elementarily equivalent to $M(C)$, for some nice graph C). To this end, we need a general description of the elements of M . Notice that Fact 2.10 fails to give a description of elements in all Mekler groups. Indeed, by compactness, we should expect elements which cannot be written as a finite product of elements of type 1^ν . We will construct *independent* sequences of elements that can be extended to a transversal (definition below). First, we clarify our notion of independence:

Definition 2.13 (independence). Let G be a group of exponent p . Let \mathfrak{E} be a subset of G and \bar{a} a tuple of elements in G . We say that \bar{a} is independent over \mathfrak{E} if for all \mathcal{L}_{grp} -terms $t(\bar{x}, \bar{y})$ and elements $\bar{b} \in \mathfrak{E}$, if $G \models t(\bar{a}, \bar{b}) = 1$, then $G \models (\forall \bar{x}) t(\bar{x}, \bar{b}) = 1$.

When the group G is abelian (of exponent p), this notion of independence coincides with linear independence if we view G as an \mathbb{F}_p -vector space.

Definition 2.14 (transversal). Let $M = (G, \cdot, 1)$ be a Mekler group. A *transversal* of M is a set X which can be written as the union of three disjoint sets X^ν , X^p , and X^l where

- X^ν is a subset of \mathfrak{E}^ν independent over Z and maximal for this property,
- X^p is a subset of \mathfrak{E}^p independent over $\langle Z, \mathfrak{E}^\nu \rangle$ and maximal for this property,
- X^l is a subset of \mathfrak{E}^l independent over $\langle Z, \mathfrak{E}^\nu, \mathfrak{E}^p \rangle$ and maximal for this property.

Every Mekler group M admits a transversal X , and all elements of M can be written as a finite product

$$a_1^{r_1} \cdots a_n^{r_n} g_1^{s_1} \cdots g_m^{s_m} w_1^{t_1} \cdots w_k^{t_k} z,$$

where $a_i \in X^\nu$, $g_i \in X^p$, $w_i \in X^l$, $z \in Z$ pairwise distinct and r_i, s_i, t_i are in $\{1, \dots, p-1\}$. This description is unique once we have fixed an order on the elements of X . We therefore often see X as a tuple rather than a subset.

Fact 2.15 [CH19, Lemma 2.7]. *For every small (possibly infinite) tuple x , there is a partial type $\phi(x)$ such that $M \models \phi(x)$ if and only if x can be extended to a transversal of M .*

This description of elements of M (which is usual in the literature of Mekler groups) can be made more precise: Let X^ζ be a subset of Z independent over $\langle \mathfrak{E}^\nu, \mathfrak{E}^p, \mathfrak{E}^l \rangle$, and maximal with this property. Then an element z of the centre can be written as

$$\prod_{x,y \in X} [x, y]^{n_{x,y}} \prod_{z \in X^\zeta} z^{n_z},$$

where $n_z, n_{x,y} \in \{0, \dots, p-1\}$ are almost all trivial. Again, we have uniqueness if we fix orders on X^ζ and X (and if we consider the lexicographic order on X^2). This follows easily from [Hod93, Axiom 8].

Thus, every element of M can be written in the form

$$\prod_{x \in X} x^{n_x} \prod_{x,y \in X} [x, y]^{n_{x,y}} \prod_{z \in X^\zeta} z^{n_z}.$$

We call $X^\nu X^p X^l X^\zeta$ a *full transversal* of M .

2.2. Generalised indiscernibles. Throughout the paper, if \mathcal{S} is a first-order structure, we let $\text{dom}(\mathcal{S})$ denote the domain of \mathcal{S} . In this section, we fix a language \mathcal{L} and a countable language \mathcal{L}' . Let T be a \mathcal{L} -structure with a monster model \mathbb{M} and let \mathcal{I} be an \mathcal{L}' -structure.

Definition 2.16 (generalised indiscernibles). Given an \mathcal{I} -indexed sequence of tuples $(\bar{a}_i)_{i \in \mathcal{I}}$ from \mathbb{M} , and a small subset A of \mathbb{M} , $(\bar{a}_i)_{i \in \mathcal{I}}$ is called an *\mathcal{I} -indiscernible sequence over A* , if for all positive integers n and all sequences $i_1, \dots, i_n, j_1, \dots, j_n$ from \mathcal{I} we have that if

$$\text{qftp}_{\mathcal{I}}^{\mathcal{L}'}(i_1, \dots, i_n) = \text{qftp}_{\mathcal{I}}^{\mathcal{L}'}(j_1, \dots, j_n)$$

then

$$\text{tp}(\bar{a}_{i_1}, \dots, \bar{a}_{i_n}/A) = \text{tp}(\bar{a}_{j_1}, \dots, \bar{a}_{j_n}/A).$$

If $A = \emptyset$, $(\bar{a}_i)_{i \in \mathcal{I}}$ is called an \mathcal{I} -indiscernible sequence.

We will always assume that the indexing structure \mathcal{I} is a Ramsey structure, i.e., the Fraïssé limit of a *Ramsey class*. The crucial fact about sequences indexed by Ramsey structures is that the conclusion of *the standard lemma* [TZ12, Lemma 5.1.3] holds for them. More precisely, we have the following theorem, originally due to Scow, [Sco15], later generalised in [MP23]:

Theorem 2.17 (generalised standard lemma). *Let \mathcal{L}' be a first-order language and \mathcal{I} an infinite, locally finite \mathcal{L}' -structure. The following conditions are equivalent:*

- (1) $\text{Age}(\mathcal{I})$ is a Ramsey class.³
- (2) \mathcal{I} -indexed indiscernibles have the modelling property.

The precise definitions of a *Ramsey class* and of the *modelling property* will be omitted, but we will recall all the other tools we need in this paper. For a proper introduction to these concepts, we direct the reader to [Bod15; Sco15]. We only recall the notion of sequence locally based on another:

Definition 2.18 (locally based). Let \mathcal{M} be an \mathcal{L} -structure. Given indexing structures \mathcal{I} and \mathbb{I} in the same language, an \mathbb{I} -indexed sequence $(\bar{a}_i : i \in \text{dom}(\mathbb{I}))$ in \mathcal{M} is said (locally) based on $(\bar{b}_i : i \in \text{dom}(\mathcal{I}))$ if for any finite set Δ of \mathcal{L} -formulas and any finite tuple (i_1, \dots, i_n) from \mathbb{I} there is a finite tuple (j_1, \dots, j_n) from \mathcal{I} such that

- $\text{qftp}_{\mathbb{I}}(i_1, \dots, i_n) = \text{qftp}_{\mathcal{I}}(j_1, \dots, j_n)$,
- $\text{tp}^{\Delta}(\bar{a}_{i_1}, \dots, \bar{a}_{i_n}) = \text{tp}^{\Delta}(\bar{b}_{j_1}, \dots, \bar{b}_{j_n})$.

Definition 2.19 (reduct). Let \mathcal{I} and \mathcal{J} be structures on the same domain. Then \mathcal{J} is a (first-order) *reduct of \mathcal{I}* if every relation and every function in \mathcal{J} is definable in \mathcal{I} without parameters. We say that \mathcal{J} is a *strict reduct* if it is a reduct of \mathcal{I} , and \mathcal{J} and \mathcal{I} are not interdefinable.

Definition 2.20. Let \mathcal{I} and \mathcal{J} be two structures and assume that \mathcal{J} is a reduct of \mathcal{I} . A structure \mathcal{M} *collapses indiscernibles from \mathcal{I} to \mathcal{J}* if every \mathcal{I} -indiscernible sequence in \mathcal{M} is a \mathcal{J} -indiscernible sequence. A complete theory T *collapses indiscernibles from \mathcal{I} to \mathcal{J}* if the monster model collapses indiscernibles from \mathcal{I} to \mathcal{J} . We say that T *collapses \mathcal{I} -indiscernibles* if it collapses any \mathcal{I} -indiscernible sequence $(a_i)_{i \in \mathcal{I}}$ to \mathcal{J} -indiscernible sequence, where \mathcal{J} is some strict reduct of \mathcal{I} (depending on the sequence $(a_i)_{i \in \mathcal{I}}$).

³Recall that the *age* of a structure \mathcal{I} is the class of all its finitely generated substructures, denoted here by $\text{Age}(\mathcal{I})$.

2.3. Model-theoretic dividing lines.

2.3.1. \mathcal{K} -configurations. The notion of \mathcal{K} -configurations originated from the work of Guingona and Hill [GH19] and was further developed in [GPS23]. In this section, we will fix a relational first-order language \mathcal{L}' with signature $\text{Sig}(\mathcal{L}')$ and an arbitrary first-order language \mathcal{L} . We abusively use the symbol \mathcal{L} to also denote the set of (parameter-free) \mathcal{L} -formulas. Throughout, we will use \mathcal{K} to denote a class of finite \mathcal{L}' -structures closed under isomorphism.

Definition 2.21 (\mathcal{K} -configuration, [GPS23, Definition 5.1]). An \mathcal{L} -structure \mathcal{M} admits (or codes) a \mathcal{K} -configuration if there are a positive integer n , a function $I : \text{Sig}(\mathcal{L}') \rightarrow \mathcal{L}$ and a sequence of functions $(f_A : A \in \mathcal{K})$ such that

- (1) for all $A \in \mathcal{K}$, $f_A : A \rightarrow \mathcal{M}^n$, and
- (2) for all $R \in \text{Sig}(\mathcal{L}')$, for all $A \in \mathcal{K}$, for all $a \in A^{\text{arity}(R)}$,

$$A \models R(a) \iff \mathcal{M} \models I(R)(f_A(a)).$$

A theory T admits a \mathcal{K} -configuration if there is a model $\mathcal{M} \models T$ which admits a \mathcal{K} -configuration. We denote by $\mathcal{C}_{\mathcal{K}}$ the class of theories that admit a \mathcal{K} -configuration, and by $\text{N}\mathcal{C}_{\mathcal{K}}$ the class of theories that *do not* admit a \mathcal{K} -configuration.

Example 2.22.

- The dense linear order $(\mathbb{Q}, <)$ codes (but does not interpret) the discrete order $(\mathbb{N}, <, P_1)$ where $P_1 := \{(n, n+1)\}$. Indeed consider an increasing sequence $(a_i)_{i \in \mathbb{N}}$. An $\text{Age}(\mathbb{N})$ -configuration is given by:
 - (1) for $J \subseteq \mathbb{N}$, $f_J : n \in J \mapsto (a_n, a_{n+1}) \in \mathbb{Q}^2$,
 - (2) $I(<)(\mathbb{Q}^2) = \{(a, b), (a', b') : a < a'\}$ and $I(P)(\mathbb{Q}^2) = \{(a, b), (a', b') : a = b'\}$.
- A theory T is stable if and only if it does not admit a $\text{Age}(\mathbb{Q}, <)$ -configuration. In other words, if and only if it does not code a linear order.
- A theory T is NIP if and only if it does not admit a \mathcal{K} -configuration, where \mathcal{K} is the class of all finite graphs.

The following theorem originally appeared in [GH19, Theorem 3.14] with the additional assumption that $\text{Age}(\mathcal{I})$ is a Fraïssé class with the *strong amalgamation property*. Later, in [MPT23], it was extended to remove this assumption, and this is the version we present below.

Theorem 2.23 [MPT23, Theorem 3.3]. *Let \mathcal{I} be an \aleph_0 -categorical Fraïssé limit of a Ramsey class. The following conditions are equivalent for a theory T :*

- (1) $T \in \text{N}\mathcal{C}_{\mathcal{I}}$.
- (2) T collapses \mathcal{I} -indiscernibles.

2.3.2. The $NFOP_k$ property. Most of the classical model-theoretic dividing lines (such as stability, NIP and NSOP; see [She90]) are essentially *binary* (in that the local combinatorial configuration whose absence guarantees strong structure results is one given by formulas in two tuples of variables). In recent years, the power of these dividing lines in finite *graph*-combinatorics has become evident, for instance, in the *stable regularity lemma* of Malliaris and Shelah [MS14]. Finding appropriate dividing lines that allow us to prove similar results for *hypergraphs* is one of the most prominent topics in the nexus of model theory and combinatorics.

In recent years, it has become evident that these dividing lines should not be binary. The *k -independence property* is one of the more well-established higher-arity dividing lines, generalising the independence property. In [AACT25], a rather robust higher-arity generalisation of stability, called the *order property*, is developed via the *functional order property*.⁴ As this material is not (yet) standard, we provide a brief summary below and refer the reader to the introduction of [AACT25] for a historical account of the concept.

Definition 2.24. Let T be a complete \mathcal{L} -theory and $k \in \mathbb{N}$. An \mathcal{L} -formula

$$\varphi(x, x_1 \dots, x_k)$$

has the *k -functional order property* (FOP_k) in T if there is a model $\mathcal{M} \models T$ and

- a sequence $(a_f)_{f:\omega^{k-1} \rightarrow \omega}$ of elements in M^x ,
- sequences $(b_i^t)_{i < \omega}$ of elements in M^{x_t} , for $1 \leq t \leq k$,

such that

$$\mathcal{M} \models \varphi(a_f, b_{i_1}^1, \dots, b_{i_k}^k) \iff i_k \leq f(i_1, \dots, i_{k-1}).$$

The theory T is $NFOP_k$ if no $(k+1)$ -partitioned formula has FOP_k in T .

One of the results from [AACT25] that will be important in this paper is a characterisation of FOP_k via collapsing indiscernibles. We will recall this in the sequel, but first, we need to introduce some terminology.

Let k be a positive integer. We denote by \mathcal{L}_k the language

$$\mathcal{L}_k = \{P_1, \dots, P_{k+1}, <, <_k, R\}.$$

We let Q denote the k -ary relation $P_1 \times \dots \times P_k$.

Definition 2.25 [AACT25, Definition 3.12]. Define T_k to be the \mathcal{L}_k -theory consisting of the following axioms:

- (1) P_1, \dots, P_{k+1} is a partition.
- (2) $<$ is a linear order with $P_1 < \dots < P_{k+1}$.

⁴Alternative notions can be found in the literature, such as Takeuchi's NOP_2 dividing line.

- (3) R only holds on $P_1 \times \dots \times P_{k+1}$ (which we also view as $Q \times P_{k+1}$).
- (4) $<_k$ only holds on $Q \times Q$, and is a linear order on Q .
- (5) For any $\bar{x}, \bar{y} \in Q$ and $w, z \in P_{k+1}$, $(\bar{x} \leq_k \bar{y} \wedge R(\bar{y}, w) \wedge w \leq z) \rightarrow R(\bar{x}, z)$.

Definition 2.26. We denote by \mathcal{H}_k the class of finite models of T_k .

Fact 2.27 [AACT25, Corollaries 3.15 and 3.16]. \mathcal{H}_k is a Fraïssé class with the Ramsey property. In particular, the Fraïssé limit \mathcal{H}_k has the modelling property.

Fact 2.28 [AACT25, Theorem 4.15]. Let T be a complete theory with monster model M . The following are equivalent.

- (1) T is NFOP $_k$.
- (2) Every \mathcal{H}_k -indexed indiscernible sequence in M is $(\mathcal{L}_k \setminus \{R\})$ -indiscernible.

3. Relative quantifier elimination

We deal, in this section, with quantifier elimination in Mekler groups. We refer the reader to [Rid17, Appendix A] for basic definitions of relative quantifier elimination and related notions. We recall simply here the key concept of a closed sort:

Definition 3.1 [Rid17, Definition A.7]. Let \mathcal{M} be a multisorted structure. A set of sorts Σ is called *closed* if any predicate involving a sort in Σ and any function with a domain involving a sort in Σ only involves sorts in Σ .

A sort S will be abusively called closed if S^{eq} is closed. A closed sort has good syntactical properties and all results of quantifier elimination presented in this paper will be therefore relative to a closed sort.

Let $C = (V; E)$ be a nice graph. As shown by Mekler, and summarised in the previous section, the graph C is interpretable in the Mekler group $M(C) = (G, \cdot, 1)$. The goal of this section is to provide a detailed description of the “definable structure” of $M(C)$, relative to the “definable structure” of C . We will proceed in a step-by-step fashion, summarised in this reduction diagram:

$$\begin{array}{c}
 M(C) = (G, \cdot, 1) \\
 \downarrow \\
 \mathcal{F}(G) = \{(G/Z, \cdot, 1), (Z, \cdot, 1), [\cdot, \cdot] : (G/Z)^2 \rightarrow Z\} \text{ where } Z := Z(G) \\
 \downarrow \\
 (G/Z, 0, +, R) \text{ where } R := \{(a + Z, b + Z) : [a, b] = 1\} \\
 \downarrow \\
 C = (V, E).
 \end{array}$$

Here a vertical arrow $A \rightarrow B$ should be understood as a relative quantifier elimination

statement of the form

“ A eliminates quantifiers relative to B^{eq} .”

The precise languages in which this happens will be introduced in the following paragraphs. These results will be enough for our practical purposes. As discussed in the introduction, the purpose of each relative quantifier elimination $A \rightarrow B$ is to obtain transfer principles, and notably a characterisation of (generalised) indiscernibles in A relative to (generalised) indiscernibles in B . This will be done in the next section, where we will also point out how these characterisations allow us to establish various transfers of dividing lines.

Notice that one cannot simply “combine the arrows” of our reduction diagram, as we eliminate quantifiers relatively at the cost of adding (finite) imaginary sorts (which is why an arrow $A \rightarrow B$ indicates that A eliminates quantifiers relative to B^{eq} rather than B). In fact, we will not present a complete quantifier elimination result from the Mekler group $M(C)$ to the nice graph C , as it seems to us that such a result would involve a complicated language and will not reveal much more information than our results. This appears to be due to the fact that our structures do not eliminate *finite imaginaries*,⁵ and we do not attempt to find sorts in which the various structures eliminate (finite) imaginaries.

3.1. First reduction: from G to $\mathcal{F}(G)$. Let G be a 2-nilpotent group of exponent p . For the remainder of this section, we will denote by Z its centre, $Z(G)$. As in [Bau02], $\mathcal{F}(G)$ will denote the structure

$$\mathcal{F}(G) := ((\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta : \mathcal{V} \times \mathcal{V} \rightarrow \mathcal{W}),$$

where

- $(\mathcal{V}, +, 0)$ denotes the \mathbb{F}_p -vector space $(G/Z, \cdot, 1)$ (with additive notation),
- $(\mathcal{W}, +, 0)$ denotes the \mathbb{F}_p -vector space $(Z, \cdot, 1)$ (with additive notation),
- $\beta : \mathcal{V} \times \mathcal{V} \rightarrow \mathcal{W}$ denotes the *commutator map*, that is,

$$\beta : G/Z \times G/Z \rightarrow Z, \quad (a \bmod Z, b \bmod Z) \mapsto [a, b].$$

Notice that these sorts are all interpretable in the language of groups, so we can expand $(G, \cdot, 1)$ to a structure in the (multisorted) language $\mathcal{L}_{G, \mathcal{F}(G)}$ (following the notation above):

$$\{(G, \cdot, {}^{-1}, 1), (\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta : \mathcal{V} \times \mathcal{V} \rightarrow \mathcal{W}, \pi : G \rightarrow \mathcal{V}, \rho : G \rightarrow \mathcal{W}\},$$

where

- $\pi : G \rightarrow \mathcal{V}$ denotes the natural projection map,

⁵See [TZ12, Definition 8.4.9(1)].

- $\rho : G \rightarrow \mathcal{W}$ denotes the “inverse” inclusion map, that is,

$$\rho : G \rightarrow \mathcal{W}, \quad g \mapsto \begin{cases} 0 & \text{if } g \notin Z, \\ g & \text{if } g \in Z. \end{cases}$$

The following fact is a reformulation of [Bau02, Corollary 3.1].

Fact 3.2. *Let G be a 2-nilpotent group of exponent p . Then, $(G, \cdot, 1)$ eliminates quantifiers relative to $\mathcal{F}(G)$, in the language $\mathcal{L}_{G, \mathcal{F}(G)}$.*

Thus any \mathcal{L}_{grp} -formula $\phi(\bar{x})$ is equivalent, modulo $\text{Th}(G)$, to a formula of the form

$$\phi_{\mathcal{F}(G)}(\pi(t(\bar{x})), \rho(t(\bar{x}))),$$

where $\phi_{\mathcal{F}(G)}(\bar{y}, \bar{z})$ is a formula in the language $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta\}$, and t is a tuple of \mathcal{L}_{grp} -terms.

In particular, the formula

$$t(x) = 1$$

for a term in \mathcal{L}_{grp} is equivalent to

$$\pi(t(x)) = 0 \wedge \rho(t(x)) = 0.$$

Sketch of proof of Fact 3.2. We apply the well-known criterion for relative quantifier elimination and argue by back-and-forth and saturation. Let G and G' be 2-nilpotent groups of exponent p , such that G' is $|G|$ -saturated and with isomorphic substructures $\mathcal{H} = (H, H_{\mathcal{F}}) \subseteq (G, \mathcal{F}(G))$ and $\mathcal{H}' = (H', H'_{\mathcal{F}}) \subseteq (G', \mathcal{F}(G'))$. Let $f = (f_G, f_{\mathcal{F}})$ denote the isomorphism between \mathcal{H} and \mathcal{H}' , and assume that $f_{\mathcal{F}} : H_{\mathcal{F}} \rightarrow H'_{\mathcal{F}}$ is elementary. By saturation, $f_{\mathcal{F}}$ can be extended to an embedding $\tilde{f}_{\mathcal{F}}$ of $\mathcal{F}(G)$ into $\mathcal{F}(G')$. Arguing as in [Bau02, Corollaries 3.1 and 3.2], we can extend $f_G \cup \tilde{f}_{\mathcal{F}}$ to an embedding of \tilde{f} of $(G, \mathcal{F}(G))$ into $(G', \mathcal{F}(G'))$. \square

We call the structure $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta : \mathcal{V}^2 \rightarrow \mathcal{W}\}$ an *alternating bilinear system* of \mathbb{F}_p -vector spaces. As shown by Baudisch, in [Bau02, Section 3], \mathcal{F} is a functor from the category $\mathbb{G}_{2,p}$ of 2-nilpotent groups of (finite) exponent p to the category \mathbb{B}_p of alternating bilinear systems of \mathbb{F}_p -vector spaces.

When G is a Mekler group $\mathcal{F}(G)$ enjoys an additional property which will be defined now and will be useful for our analysis.

Definition 3.3. Let $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta : \mathcal{V}^2 \rightarrow \mathcal{W}\}$ be an alternating bilinear system, and let $\mathcal{V}' \subseteq \mathcal{V}$ be a subspace.

- A basis $(v_i)_i$ of \mathcal{V}' is called *separated* if for all sequences $(\alpha_{i,j})_{i < j}$ of scalars with only finitely many $\alpha_{i,j}$ nontrivial, if

$$\sum_{i < j} \alpha_{i,j} \beta(v_i, v_j) = 0,$$

then for all $i < j$, either $\alpha_{i,j} = 0$ or $\beta(v_i, v_j) = 0$.

- $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta : \mathcal{V}^2 \rightarrow \mathcal{W}\}$ is called *separated* if it admits a separated basis.

Example 3.4. Let \mathcal{V} be any vector space.

- (1) The bilinear system $(\mathcal{V}, \mathcal{V} \wedge \mathcal{V}, \wedge)$, where $\mathcal{V} \wedge \mathcal{V}$ is the wedge product, is separated. This follows from the fact that for any basis $(v_i)_i$ of \mathcal{V} , $(v_i \wedge v_j)_{i < j}$ is a basis of $\mathcal{V} \wedge \mathcal{V}$, by the universal property of the wedge product.
- (2) Let v_1, v_2, v_3, v_4 be a basis of \mathcal{V} . Let \mathcal{W} be $\mathcal{V} \wedge \mathcal{V} / \langle v_1 \wedge v_2 - v_3 \wedge v_4 \rangle$ and $\beta : \mathcal{V}^2 \rightarrow \mathcal{W}$ be induced by the wedge product. Then the bilinear system $(\mathcal{V}, \mathcal{W}, \beta)$ is not separated. Indeed, for all linearly independent vectors $v, v' \in \mathcal{V}$, observe that $\beta(v, v') \neq 0$. Therefore, a separated basis of \mathcal{V} would give six vectors linearly independent in \mathcal{W} , but $\dim(\mathcal{W}) = 5$.
- (3) If \mathcal{V} has dimension ≤ 2 , then any bilinear system of the form $(\mathcal{V}, \mathcal{W}, \beta)$ is separated.

If \mathcal{V} is infinite-dimensional, then separatedness of \mathcal{V} is *a priori* not a first-order property. In the case of Mekler groups, we can, in a first-order way, express the stronger property that any finite-dimensional vector subspace is contained in a finite-dimensional separated vector subspace. To this end, we need the following fact, which is a slightly weakened reformulation of [Hod93, Axioms 8 and 9]:

Fact 3.5. *Let k, m, n be integers and let g_0, \dots, g_{n-1} be a sequence of independent elements in a Mekler group M such that*

- g_0, \dots, g_{k-1} are of type 1^v ,
- g_k, \dots, g_{m-1} are of type p , and no product of g_k, \dots, g_{m-1} is a product of at most $k + 1$ elements of type 1^v ,
- g_m, \dots, g_{n-1} are of type 1^l , and no product of g_m, \dots, g_{n-1} is a product of at most $m + 1$ elements of type 1^v or of type p .

Then $g_0 \bmod Z, \dots, g_{n-1} \bmod Z$ is a separated basis of $\langle g_0 \bmod Z, \dots, g_{n-1} \bmod Z \rangle$ in $\mathcal{F}(G)$.

Corollary 3.6. *There is a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that the following holds: For G a Mekler group, let $(\mathcal{V}, \mathcal{W})$ denotes $\mathcal{F}(G)$. Then,*

*every vector subspace of \mathcal{V} of dimension n is included in a separated vector space of dimension at most $f(n)$. (*_f)*

Proof. We show that $f(n) = 2^{2^n}$. Let V be a vector space of dimension n , and let $0 \leq k \leq m \leq n$ be integers and g_0, \dots, g_{n-1} a basis of V such that

- g_0, \dots, g_{k-1} are of type 1^v ,

- g_k, \dots, g_{m-1} are of type p ,
- g_m, \dots, g_{n-1} are of type 1^l .

By induction on $(n - m, m - k, k)$ with the lexicographic order, we show that \mathcal{V} is included in a vector space of dimension at most $(k + 1)2^{(m+1)2^{n-m-k}}$ (which is less than 2^{2^n}). If g_0, \dots, g_{n-1} satisfy the conditions of Fact 3.5, then g_0, \dots, g_{n-1} is separated and there is nothing to show. If not, then at least one of the following holds:

- there is a product of the g_i 's, $m \leq i < n$, which is a product of at most $m + 1$ elements h_0, \dots, h_m of type p or 1^v , or
- there is a product of the g_i 's, $k \leq i < m$, which is a product of at most $k + 1$ elements h_0, \dots, h_k of type 1^v .

In any case, we consider the vector space \mathcal{V}' generated by $g_0, \dots, g_{n-1}, h_0, \dots, h_m$ (resp. $g_0, \dots, g_{n-1}, h_0, \dots, h_k$) and assume that h_0, \dots, h_k are independent. Let g_{i_0}, \dots, g_{i_k} be a proper subset of the g_i 's such that $g_{i_0}, \dots, g_{i_k}, h_0, \dots, h_m$ (resp. $g_{i_0}, \dots, g_{i_k}, h_0, \dots, h_k$) is a basis of \mathcal{V}' . By induction, \mathcal{V} is included in a vector space of dimension at most $(k + 1)2^{(2m+2)2^{n-m-1-k}}$ (resp. $(2k + 2)2^{(m+1)2^{n-m-k-1}}$), which is equal to $(k + 1)2^{(m+1)2^{n-m-k}}$. \square

This bound is probably not optimal. In fact, as far as we know, an optimal bound could be $f(n) = n$, i.e., every vector subspace of \mathcal{V} could be separated.

3.2. Second reduction: From $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta\}$ to $(\mathcal{V}, +, 0, R)$. We now analyse the definable sets in the bilinear system $\mathcal{F}(G)$. We define a binary relation on \mathcal{V} by

$$R := \{(v_0, v_1) \mid \beta(v_0, v_1) = 0\}.$$

Heuristically, by separatedness, one should be able to recover all the first-order expressiveness of $\mathcal{F}(G)$ from the structure $(\mathcal{V}, R, +, 0)$. However, within the structure $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta : \mathcal{V}^2 \rightarrow \mathcal{W}\}$, the sort \mathcal{V} is not closed in the sense of Definition 3.1. This needs to be addressed first. Since vector spaces do not eliminate finite imaginaries, this will come at the cost of introducing sorts B_n , for all positive integers n , which are a certain quotient of $P_{\leq k_n}(\mathcal{V}^n)$, the set of subsets of n -tuples in \mathcal{V} of size less than a integer k_n depending on n .

For $n \in \mathbb{N}$, we denote by $\mathbb{A}_n(\mathbb{F}_p)$ the (finite) set of antisymmetric (i.e., skew-symmetric) $n \times n$ matrices $A = (a_{i,j})$ with coefficients in \mathbb{F}_p , and by $\mathbb{A}(\mathbb{F}_p)$ the union of $\mathbb{A}_n(\mathbb{F}_p)$, for all $n \in \mathbb{N}$.

For $n \in \mathbb{N}$, let \mathcal{W}_n be the set

$$\mathcal{W}_n = \left\{ w \in \mathcal{W} \mid \exists v_0, \dots, v_{n-1} \in \mathcal{V}, \exists (a_{i,j}) \in \mathbb{A}_n(\mathbb{F}_p) \text{ s.t. } w = \sum_{i < j} a_{i,j} \beta(v_i, v_j) \right\},$$

that is, the set of elements of order n . For two positive integers $n > m$, there is a natural inclusion of \mathcal{W}_m in \mathcal{W}_n and for every positive integer n , there is a function $+: (\mathcal{W}_n)^2 \rightarrow \mathcal{W}_{2n}$ which is the trace of the addition in \mathcal{W} .

It follows from separatedness that \mathcal{W}_n , with the structure above, is an imaginary of the structure $(\mathcal{V}, +, 0, R)$, for all $n \in \mathbb{N}$:

Lemma 3.7. *Let $(\mathcal{V}, \mathcal{W}, \beta)$ be an alternating bilinear system satisfying property $(*_f)$ of Corollary 3.6, and denote by \simeq the equivalence relation on $\mathbb{A}_n \times \mathcal{V}^n$ given by*

$$(A, \bar{v}) \simeq (A', \bar{v}') \iff \sum_{i < j} a_{i,j} \beta(v_i, v_j) = \sum_{i < j} a'_{i,j} \beta(v'_i, v'_j).$$

Let $B_n := \mathbb{A}_n \times \mathcal{V}^n / \simeq$. Then:

- (1) *The equivalence relation \simeq is interpretable in $(\mathcal{V}, +, 0, R)$.*
- (2) *The set B_n is an imaginary sort of $(\mathcal{V}, +, 0, R)$ and can be identified with \mathcal{W}_n .*
- (3) *The addition $+: (\mathcal{W}_n)^2 \rightarrow \mathcal{W}_{2n}$ is also interpretable in $(\mathcal{V}, +, 0, R)$.*

Notice that $\bigcup_n \mathcal{W}_n$ is $\langle \beta(\mathcal{V}, \mathcal{V}) \rangle$, the \mathbb{F}_p -vector subspace of \mathcal{W} generated by the image of β , and it is not equal to \mathcal{W} in general.

Notation 3.8. Let $(\mathcal{V}, \mathcal{W}, \beta)$ be a bilinear system satisfying property $(*_f)$, and let \simeq be the equivalence relation defined in the above lemma. For $A \in \mathbb{A}_n(\mathbb{F}_p)$, we write

$$\pi_A : \mathcal{V}^n \rightarrow B_n, \quad \bar{v} \mapsto (A, \bar{v}) / \simeq,$$

for the natural projection, and

$$f_n : \mathcal{W} \rightarrow B_n, \quad w \mapsto \begin{cases} (A, \bar{v}) / \simeq & \text{if } w = \sum_{i < j} a_{i,j} \beta(v_i, v_j), \\ \text{u} & \text{if } w \notin \mathcal{W}_n, \end{cases}$$

for the reverse inclusion. Recall that the constant symbol u is interpreted as “undetermined”.

Proof. (1) By assumption, property $(*_f)$ holds, and the vector space $\mathcal{V} := \langle \bar{v}, \bar{v}' \rangle$ is included in a separated vector space \mathcal{V}' of dimension at most $f(n)$, where n is the dimension of \mathcal{V} . Given a separated basis $\bar{u} = (u_i)$ of \mathcal{V}' , the equality

$$\sum_{i < j} a_{i,j} \beta(v_i, v_j) = \sum_{i < j} a'_{i,j} \beta(v'_i, v'_j)$$

can be rewritten in this basis to give an equation

$$\sum_{i < j} c_{i,j} \beta(u_i, u_j) = 0,$$

for coefficients $c_{i,j} \in \mathbb{F}_p$. Therefore, $\sum_{i < j} a_{i,j} \beta(v_i, v_j) = \sum_{i < j} a'_{i,j} \beta(v'_i, v'_j)$ holds if and only if, in some basis \bar{u} of a vector subspace of dimension at most $f(n)$, for

all $i < j$, $c_{i,j} = 0$ or $(u_i, u_j) \in R$. It follows that the equivalence relation \simeq can be written using only the addition in \mathcal{V} and the predicate R .

(2) Immediate.

(3) By property $(*_f)$, we can work in a finite-dimensional separated vector subspace. Consider two elements w', w'' and an appropriate separated basis $(v_i)_{i < m}$ with $A' = (a'_{i,j})$, $A'' = (a''_{i,j}) \in \mathbb{A}_m(\mathbb{F}_p)$, such that $f_m(w') = \pi_{A'}(v_i)$ and $f_m(w'') = \pi_{A''}(v_i)$.

Suppose that $A = A' + A''$. Then, we simply have $f_m(w' + w'') = \pi_A(v_i)$. Therefore, for $b, b', b'' \in B_n$,

$$\exists w, w' \in \mathcal{W}_n (f_n(w') = b' \wedge f_n(w'') = b'' \wedge f_n(w + w') = b)$$

holds if and only if for some $m \leq 2n$, there are $A', A'' \in \mathbb{A}_m(\mathbb{F}_p)$ and $\bar{v}' \in \mathcal{V}^m$ such that $\pi_{A'}(\bar{v}) = b'$, $\pi_{A''}(\bar{v}) = b''$ and $\pi_{A'+A''}(\bar{v}) = b$. The statement follows. \square

Of course, we can recover the alternating map by considering the function f_A , where $A = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$. Then, for a nontrivial $w \in \mathcal{W}$ we have that $f_2(w) = \pi_A(v_0, v_1)$ if and only if $\beta(v_0, v_1) = w$. In particular:

Corollary 3.9. *The structures*

$$\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta : \mathcal{V}^2 \rightarrow \mathcal{W}\}$$

and

$$\{(\mathcal{V}, +, 0, R)^{\text{eq}}, (\mathcal{W}, +, 0), \pi_A : \mathcal{V}^n \rightarrow B_n, f_n : \mathcal{W} \rightarrow B_n; n \in \mathbb{N}, A \in \mathbb{A}_n(\mathbb{F}_p)\}$$

are bi-interpretable.

Proposition 3.10. *An alternating bilinear system*

$$\{(\mathcal{V}, +, 0, R)^{\text{eq}}, (\mathcal{W}, +, 0), \pi_A, f_n; n \in \mathbb{N}, A \in \mathbb{A}_n\}$$

having property $(*_f)$ of Corollary 3.6 eliminates quantifiers relative to $(\mathcal{V}, +, 0, R)^{\text{eq}}$.

In particular, in such a bilinear system $\{(\mathcal{V}, +, 0, R), (\mathcal{W}, +, 0), \beta\}$, every formula $\phi(w_1, \dots, w_n; v_1, \dots, v_n)$, where $n \in \mathbb{N}$, $v_1, \dots, v_n \in \mathcal{V}$ and $w_1, \dots, w_n \in \mathcal{W}$, is equivalent to a Boolean combination of formulas of the form

- $\sum a_i w_i = 0$ where the a_i 's are in \mathbb{F}_p , and
- $\phi_{\mathcal{V}}(f_n(\sum a_{1,i} w_i), \dots, f_n(\sum a_{k,i} w_i), \bar{v})$ where $\phi_{\mathcal{V}}$ is a formula in the language of $(\mathcal{V}, +, 0, R)^{\text{eq}}$, and the $a_{j,i}$'s are in \mathbb{F}_p .

Thus \mathcal{V} is a stably embedded sort and the induced structure on \mathcal{V} is given by $(\mathcal{V}, +, 0, R)$.

Remark 3.11. Let us briefly argue why the nonelimination of finite imaginaries prevents us from exhibiting a simpler language for quantifier elimination. Consider, for instance, an element of \mathcal{W} of the form

$$\beta(a, b + c) + \beta(d, c) = \beta(a, b) + \beta(a, c) + \beta(d, c) = \beta(a, b) + \beta(d + a, c),$$

where $a, b, c, d \in \mathcal{V}$ form a separated basis. At this level of generality, none of these three forms can be favoured over the others, and thus we do not try to distinguish between all the possibilities.

Remark 3.12. We can deduce the weaker statement that $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta\}$ eliminates \mathcal{W} -quantifiers. For instance, a formula of the form

$$\phi_{\mathcal{V}}(f_n(w), \bar{v}),$$

where $\phi_{\mathcal{V}}$ is a $(\mathcal{V}, +, 0, R)^{\text{eq}}$ -formula, can be rewritten as

$$\exists v'_1, \dots, v'_n \in V^m (\phi'_{\mathcal{V}}(v'_1, \dots, v'_n, \bar{v}) \wedge w = \sum a_{i,j} \beta(v'_i, v'_j))$$

for some appropriate integer m , and an appropriate $\{\mathcal{V}, +, 0, R\}$ -formula $\phi'_{\mathcal{V}}$. The predicate R can then be eliminated using the function β .

Proof of Proposition 3.10. Let

$$\mathcal{M} := \{(\mathcal{V}, +, 0, R)^{\text{eq}}, (\mathcal{W}, +, 0), \beta\} \quad \text{and} \quad \mathcal{N} := \{(\mathcal{V}', +, 0, R)^{\text{eq}}, (\mathcal{W}', +, 0), \beta\}$$

be alternating bilinear systems satisfying property $(*_f)$ of Corollary 3.6. Let $\sigma = (\sigma_{\mathcal{V}}, \sigma_{\mathcal{W}}) : \mathcal{A} \rightarrow \mathcal{B}$ be an isomorphism between two substructures $\mathcal{A} = (A_{\mathcal{V}}, A_{\mathcal{W}})$ and $\mathcal{B} = (B_{\mathcal{V}}, B_{\mathcal{W}})$ of \mathcal{M} and \mathcal{N} , respectively. Suppose furthermore that $\sigma_{\mathcal{V}} : A_{\mathcal{V}} \rightarrow B_{\mathcal{V}}$ is elementary. It suffices to show that we can extend σ to an embedding $\tilde{\sigma}$ of \mathcal{M} into \mathcal{N} . By elementarity, we can extend $\sigma_{\mathcal{V}}$ to $\tilde{\sigma}_{\mathcal{V}} : \mathcal{V}^{\text{eq}} \rightarrow \mathcal{V}'^{\text{eq}}$.

By identifying \mathcal{W}_n with B_n , we may assume that $\mathcal{W}_n \subseteq A_{\mathcal{W}}$ for every n . It remains to show that the map σ can be extended to the elements w with nonfinite order, i.e., such that $f_A(w) = u$ for every n and $A \in \mathbb{A}_n(\mathbb{F}_p)$. Take any vector space \mathcal{W}_{ω} such that $A_{\mathcal{W}} \oplus \mathcal{W}_{\omega} = \mathcal{W}$, and let $(w_i)_{i \in I}$ be a basis of \mathcal{W}_{ω} . By saturation, we can find $|I|$ many vectors $(w'_i)_{i \in I}$ of \mathcal{W}' such that $(w'_i)_{i \in I}$ are linearly independent over $B_{\mathcal{W}} \cup \bigcup_n \mathcal{W}'_n$. Then it follows that mapping $w_i \mapsto w'_i$ gives an extension of the embedding $\sigma : \mathcal{W} \rightarrow \mathcal{W}'$. \square

3.3. Third reduction: From $(\mathcal{V}, +, 0, R)$ to \mathbf{C} . We go back to Mekler groups and multiplicative notation. Let $\mathbf{C} = (V, E)$ be a nice graph, and $G = M(\mathbf{C})$ be its Mekler group. Let Z denote the centre of G . By Corollary 3.6 and Proposition 3.10, we now know the induced structure on G/Z . It is given by

$$(G/Z, \cdot, 1, R),$$

where

$$R := \{(a + Z, b + Z) \mid [a, b] = 1\}.$$

The equivalence relations $a \sim b$ and $a \approx b$ pass to the quotient modulo Z . Abusing notation, we will write

- $(a \bmod Z) \approx (b \bmod Z)$ if $\bigvee_{0 < n < p} b^n \bmod Z = a \bmod Z$,
- $(a \bmod Z) \sim (b \bmod Z)$ if $\forall x \ xRa \leftrightarrow xRb$.

Similarly, abusing terminology, we will also say that an element $a \bmod Z \in G/Z$ is of type 1^ν (resp. 1^l or p) if a is. The subset of elements of type 1^ν is of course still definable and given by the following formula $\varphi(x)$:

$$x \notin Z \wedge \forall y (x \sim y \leftrightarrow x \approx y) \wedge (\exists z [z \notin Z \wedge zRx \wedge \neg(z \approx x)])$$

and we recover the nice graph $C = (V, E)$ the same way by considering the quotient $\phi(G/Z)/\sim$, where the edge relation E is induced by the predicate R on G/Z . We want to analyse the definable structure in $(\mathcal{V}, +, 0, R)$ and reduce it to $C = (V, E)$.

Recall from Fact 2.10(5) that an element a of G/Z can be written $a = a_0 \cdots a_{n'-1}$ where $n' \in \mathbb{N}$ and $a_0, \dots, a_{n'-1}$ are distinct elements of type 1^ν , and this presentation is unique once we fix an order on $\{[a]_\sim : a \in G \text{ of type } 1^\nu\}$. We call the set $\{[a_0]_\sim, \dots, [a_{n'-1}]_\sim\}$ the *support* of a . If C is infinite, the support is not interpretable, in the sense that there is no formula $\psi(x, y)$ that holds precisely when x and y have the same support. However, it becomes definable (with quantifiers) when we bound the size of the support. Therefore, the function $a \mapsto \{[a_0]_\sim, \dots, [a_{n'-1}]_\sim\}$ must be added to the language in order to obtain quantifier elimination. This however is not sufficient. The reason is that, for a general description of elements in an elementary extension, we must use a transcendental basis. We introduce therefore in the next definition a language strong enough to reflect this.

If $g \in G/Z$ is an element of type p , recall that there is a unique class $[a]_\sim$ of an element of type 1^ν , such that gRa . This class $[a]_\sim$ is called *handle* of g , denoted by $h(g)$.

Definition 3.13. We denote by $A_{n,m}$ the set of elements a consisting of a product of at most n elements $a_0, \dots, a_{n'-1}$ of type 1^ν and of at most m elements $g_0, \dots, g_{m'-1}$ of type p such that

- (i) $g_0, \dots, g_{m'-1}$ are not a product of less than $n + m + 1$ elements of type 1^ν ,
- (ii) $[a_0]_\sim, \dots, [a_{n'-1}]_\sim, h(g_0), \dots, h(g_{m'-1})$ are pairwise distinct.
- (iii) $[a_0]_\sim, \dots, [a_{n'-1}]_\sim$ are not connected (according to E) to any of $h(g_0), \dots, h(g_{m'-1})$.

Note that these conditions are first-order.

We define

$$\begin{aligned} S_{n,m} &: A_{n,m} \rightarrow \mathcal{P}_{\leq n}(C), & a &\mapsto \{[a_0]_\sim, \dots, [a_{n'-1}]_\sim\}, \\ S'_{n,m} &: A_{n,m} \rightarrow \mathcal{P}_{\leq m}(C), & a &\mapsto \{h(g_0), \dots, h(g_{m'-1})\}, \end{aligned}$$

where $\mathcal{P}_{\leq n}(C) := \bigcup_{k \leq n} \mathcal{P}_k(C)$ denotes the set of subsets of C with at most n elements.⁶ We set $S_{n,m}(a) = S'_{n,m}(a) = u$ if $a \notin A_{n,m}$.

⁶We will view $\mathcal{P}_{\leq n}(C)$ as an imaginary sort of C , or if C has elimination of finite imaginaries, as a definable subset of C^n .

The sets $A_{n,m}$ don't form a partition: for $(n, m), (k, l) \in \mathbb{N}^2$, $A_{n,m} \cap A_{k,l}$ is not necessarily empty (for example, $A_{n,0} \subseteq A_{n+1,0}$), and in general, $\bigcup_{n,m} A_{n,m} \subsetneq G$. However, for any a there is a pair $(n, m) \in \mathbb{N}^2$ which is minimal for the reverse lexicographic order, such that $A_{n,m}(a)$.

For such a minimal pair (n, m) , (ii) and (iii) will be easy to verify and (i) will hold automatically.

Remark 3.14. Let a, b be elements of a Mekler group M . If (n, m) and (k, l) are minimal such that $A_{n,m}(a)$ and $A_{k,l}(b)$ hold, and $a = a_0 \cdots a_{n-1} \cdot g_0 \cdots g_{m-1}$ is the corresponding decomposition, then

- g_0, \dots, g_{m-1} are not a (finite) product of elements of type 1^ν ,
- $A_{k+n,l+m}(a \cdot b)$ holds.

However, $(k + n, l + m)$ may not be minimal such that $A_{k+n,l+m}(a \cdot b)$ holds. The more complicated cases in the proofs below arise when this is indeed not the case.

Proposition 3.15. *The functions $S_{n,m}$ and $S'_{n,m}$ are well-defined and \emptyset -definable in Mekler groups. We refer to them as the **support functions**.*

Proof. Consider $G := M(C)$ the Mekler group associated with a nice graph C and let $a \in G$. Suppose that we are given two decompositions

$$a = a_0 \cdots a_{n-1} \cdot g_0 \cdots g_{m-1} = a'_0 \cdots a'_{n-1} \cdot g'_0 \cdots g'_{m-1}$$

which satisfy (i), (ii) and (iii) from Definition 3.13. We must show that

$$\begin{aligned} \{[a_0]_\sim, \dots, [a_{n-1}]_\sim\} &= \{[a'_0]_\sim, \dots, [a'_{n-1}]_\sim\}, \\ \{h(g_0), \dots, h(g_{m-1})\} &= \{h(g'_0), \dots, h(g'_{m-1})\}. \end{aligned}$$

We start with the latter.

Suppose that $h(g_0) \notin \{h(g'_0), \dots, h(g'_{m-1})\}$. Since C is triangle-free and square-free, for all i 's, g_0 and g'_i are products of elements of type 1^ν and at most one of these elements is in both products. It follows that g_0 is composed with at most $n + m$ elements of type 1^ν , which contradicts (i). Therefore, we have that $h(g_0) \in \{h(g'_0), \dots, h(g'_{m-1})\}$ and, arguing in the same manner, it follows that $\{h(g_0), \dots, h(g_{m-1})\} = \{h(g'_0), \dots, h(g'_{m-1})\}$.

Now, we have by assumption that $[a_0]_\sim$ is not equal to and does not commute with any element of $\{h(g'_0), \dots, h(g'_{m-1})\}$. Therefore, $[a_0]_\sim$ can't be an element composing any of the g'_i and we must have $[a_0]_\sim = [a'_i]_\sim$ for some $i < n$. More generally, we have $\{[a_0]_\sim, \dots, [a_{n-1}]_\sim\} = \{[a'_0]_\sim, \dots, [a'_{n-1}]_\sim\}$, as wanted. \square

Remark 3.16. Let $a \in G$ satisfy $A_{n,m}$. A decomposition $a = a_0 \cdots a_{n'-1} \cdot g_0 \cdots g_{m'-1}$ such that $S_{n,m}(a) = \{[a_0]_\sim, \dots, [a_{n'-1}]_\sim\}$ and $S'_{n,m}(a) = \{h(g_0), \dots, h(g_{m'-1})\}$ is almost uniquely determined. If $a = a'_0 \cdots a'_{n''-1} \cdot g'_0 \cdots g'_{m''-1}$ is another such

decomposition, then

- $\{a_0, \dots, a_{n'-1}\} = \{a'_0, \dots, a'_{n''-1}\},$
- $\{g_0 \pmod{V_c}, \dots, g_{m'-1} \pmod{V_c}\} = \{g'_0 \pmod{V_c}, \dots, g'_{m''-1} \pmod{V_c}\}$

where $V_c := \langle a_{i,j} : i < j < m \rangle$ is the subgroup generated by the set of connections, i.e., the elements $a_{i,j}$ of type 1^\vee such that $[a_{i,j}]_\sim$ is connected to both $h(g_i)$ and $h(g_j)$ (as the graph C is square free, there are $(p-1)m(m-1)/2$ of them). The connection $a_{i,j}$ can therefore be “included” in the element of type p and handle $h(g_i)$ or in the one of handles $h(g_j)$.

In particular, $n' = n''$ and $m' = m''$, that is, the decompositions have the same length.

The next result lies at the heart of our relative quantifier elimination result. We postpone the proof to the end of the section, in order to first give more context, as well as some essential lemmas.

Proposition 3.17. *The structure*

$$\{(G/Z, \cdot, ^{-1}, (A_{n,m})_{n,m}), C = (V, E)^{\text{eq}}, S_{n,m} : A_{n,m} \rightarrow \mathcal{P}_{\leq n}(C), S'_{n,m} : A_{n,m} \rightarrow \mathcal{P}_{\leq m}(C)\}$$

eliminates quantifiers relative to the graph C .

In particular, every formula $\phi(x)$ in the structure $(G/Z, \cdot, 1, R)$ is equivalent to a Boolean combination of formulas of the form

- $A_{n,m}(t(x))$, where $n, m \in \mathbb{N}$, $t(x)$ is a group term, and
- $\phi_C(S_{n,m}(t(x)), S'_{n,m}(t(x)))$, where $n, m \in \mathbb{N}$, $t(x)$ is a tuple of group terms and $\phi_C(x_C, x'_C)$ is a formula in $C = (V, R)$.

To illustrate this result, we give “main-sorted-quantifier-free” definitions of important sets (proofs are optional and left to the reader).

- $A_{0,0} = \{1\}$, $A_{0,1} \setminus A_{0,0}$ is exactly the set of elements of type p , and $A_{1,0} \setminus A_{0,0}$ the set of elements of type 1^\vee .
- For any element a of type 1^\vee and element g of type p , we have

$$S_{1,0}(a) = \{[a]_\sim\} \quad \text{and} \quad S'_{0,1}(g) = \{h(g)\}.$$

- Elements of type $p-1$ are exactly the elements x with support consisting of two commuting elements:

$$A_{2,0}(x) \wedge \exists \alpha, \beta \in S_{2,0}(x) \alpha E \beta.$$

- Elements of type 1^ι are all the other elements:

$$\neg A_{0,1}(x) \wedge (\neg A_{2,0}(x) \vee (A_{2,0}(x) \wedge \exists \alpha, \beta \in S_{2,0}(x) \neg \alpha E \beta \wedge \alpha \neq \beta)).$$

Recall that $A_{0,0}, A_{1,0} \subseteq A_{2,0}$.

- For $a, b \in G/Z$, the relation aRb induced by commutation on G/Z is given by

$$\begin{aligned} (a \in A_{1,0} \wedge b \in A_{0,1} \wedge S'_{0,1}(b) = S_{1,0}(a)) \\ \vee (a \in A_{0,1} \wedge b \in A_{1,0} \wedge S'_{0,1}(a) = S_{1,0}(b)) \\ \vee (a, b \in A_{2,0} \wedge \forall \alpha \in S_{2,0}(a) \forall \beta \in S_{2,0}(b) \alpha E \beta). \end{aligned}$$

The proof, which we will detail below, uses the criterion of Shoenfield, consisting of extending a partial isomorphism $f : \mathcal{A} \subseteq \mathcal{M} \rightarrow \mathcal{B} \subseteq \mathcal{N}$ between two models with saturation. We will do it in five steps:

Step 0: We enlarge f to all elements of $C(\mathcal{M})^{\text{eq}}$.

Step 1: We enlarge f to elements a of type 1^v occurring in the decomposition of an element in $G_{\mathcal{A}}$.

Step 2: We enlarge f to elements g of type p occurring in the decomposition of an element in $G_{\mathcal{A}}$.

Step 3: We find a transversal X that is “compatible” with \mathcal{A} , and enlarge f to this transversal.

Step 4: We extend f to a full embedding of the graph $\Gamma(X)$ (which contains $C(\mathcal{M})$), then to a full embedding of \mathcal{M} .

We chose the language \mathcal{L} of Proposition 3.17 so that the following improvement of Fact 2.15 holds:

Lemma 3.18. *Given a small set of variables $Y = Y^v \frown Y^p \frown Y^l$, the statement: “The elements of Y^v , Y^p and Y^l are respectively of type 1^v , p and 1^l , and the set Y can be extended to a transversal of G ”*

*is a **quantifier-free** type in the variables Y in the language \mathcal{L} .*

Proof. The quantifier-free type

$$\{\neg A_{n,0}(\prod v_i^{k_i}) \mid (k_1, \dots, k_l) \in \{0, 1, \dots, p-1\} \setminus \{0, \dots, 0\}, n \in \mathbb{N}\}$$

expresses that $v_1, \dots, v_l \in \mathcal{V}$ are independent modulo $\langle \mathfrak{E}^v \rangle$.

Similarly, the quantifier-free type

$$\{\neg A_{n,m}(\prod w_i^{k_i}) \mid (k_1, \dots, k_l) \in \{0, 1, \dots, p-1\} \setminus \{0, \dots, 0\}, n, m \in \mathbb{N}\}$$

expresses that $w_1, \dots, w_l \in V$ are independent modulo $\langle \mathfrak{E}^v, \mathfrak{E}^p \rangle$. The lemma follows immediately. \square

In the next lemma, we characterise finite subsets Δ of elements d which share the same element h_0 of type p . Let $d, d' \in S$ and assume that there are pairs of integers $(k, l), (k', l')$, minimal with the reverse lexicographic order, such that $A_{k,\ell}(d), A_{k',\ell'}(d')$ hold. Write accordingly $d = a_0 \cdots a_k \cdot g_0 \cdots g_l$ and $d' = b_0 \cdots b_{k'} \cdot h_0 \cdots h_{l'}$.

Assume that $g_0 = h_0$. Then, in the product dd'^{-1} , the elements g_0 and h_0 cancel out and

- $A_{k+k', \ell+\ell'-2}(dd'^{-1})$ holds,
- $S_{k+k', \ell+\ell'-2}(dd'^{-1}) \subseteq S_{k, \ell}(d) \cup S_{k', \ell'}(d')$,
- $S'_{k+k', \ell+\ell'-2}(dd'^{-1}) \subseteq S'_{k, \ell}(d) \cup S'_{k', \ell'}(d') \setminus \{[h_0]_{\sim}\}$.

This generalises of course to any finite subset of elements sharing the same element of type p in there decomposition. We need a partial reciprocal of the fact above. It is easier to state if we restrict to the case where only one element of type p “cancels out” in the product dd'^{-1} (i.e., when the inclusions above are strict equalities).

Lemma 3.19. *Let Δ be a finite subset of $\bigcup_{n,m} A_{n,m}$ and $c \in C$. Assume that for all distinct $d, d' \in \Delta$, we have*

- $A_{k+k', \ell+\ell'-2}(dd'^{-1})$ holds,
- $S_{k+k', \ell+\ell'-2}(dd'^{-1}) = S_{k, \ell}(d) \cup S_{k', \ell'}(d')$,
- $S'_{k+k', \ell+\ell'-2}(dd'^{-1}) = S'_{k, \ell}(d) \cup S'_{k', \ell'}(d') \setminus \{c\}$

whenever $(k, \ell), (k', \ell')$ are minimal in the reverse lexicographic order such that $A_{k, \ell}(d), A_{k', \ell'}(d')$ hold.

Then, there is an element h_0 of type p and with handle c such that for all $d \in \Delta$, and for (k, ℓ) minimal such that $A_{k, \ell}(d)$ holds, we have that $A_{k, \ell-1}(dh_0)$ holds, $S'_{k, \ell-1}(dh_0) = S'_k(d)$ and $S'_{t, \ell-1}(dh_0) = S'_{t, \ell-1}(d) \setminus \{c\}$.

The proof follows from the fact that two distinct elements d and d' of such a set Δ must decompose in products of elements of type 1^v and p and with exactly one element (of type p and handle c) appearing in both decomposition. The details of the proof are left to the reader.

Definition 3.20. Let C be an infinite nice graph. A *cover* of C is a graph (Γ, E) containing C as a subgraph such that for every $b \in \Gamma \setminus C$, one of the following two statements holds:

- There is a unique vertex a in Γ connected to b , and moreover a is in C and connected to infinitely many elements in C .
- b is an isolated vertex in Γ .

A cover of C is called *infinite* if moreover

- every vertex a in C with infinitely many neighbours in C has infinitely many neighbours in $\Gamma \setminus C$, and
- there are infinitely many isolated vertices in $\Gamma \setminus C$.

Lemma 3.21. *Let (Γ, E) be an infinite cover of a nice graph C . Then the structure (Γ, C, E) , with a predicate for the subgraph C eliminates quantifiers relative to (C, E) .*

Proof. Let $\mathcal{M} = (\Gamma(\mathcal{M}), C(\mathcal{M}))$ and $\mathcal{N} = (\Gamma(\mathcal{N}), C(\mathcal{N}))$ be two models, and \mathcal{A}, \mathcal{B} be substructures of \mathcal{M} and \mathcal{N} , respectively. Assume \mathcal{N} is $|\mathcal{M}|$ -saturated. Consider a partial isomorphism $f : \mathcal{A} \rightarrow \mathcal{B}$ such that $f|_{C(\mathcal{M})}$ is elementary in C .


Then, $f|_{C(\mathcal{M})}$ can be extended to a full embedding $\tilde{f}_{C(\mathcal{M})} : C(\mathcal{M}) \rightarrow C(\mathcal{N})$ and $\tilde{f}_{C(\mathcal{M})} \cup f$ is a partial isomorphism. We reset the notation by setting $f = \tilde{f}_{C(\mathcal{M})} \cup f$. It remains to extend f to a point b in $\Gamma(\mathcal{M}) \setminus \mathcal{A}$. If b is connected to a unique $a \in C(\mathcal{M})$, then $f(a)$ has infinitely many neighbours. By saturation, there is an element $b' \in \Gamma(\mathcal{N}) \setminus \mathcal{B}$ joined with $f(a)$. If b is isolated, then by saturation we can find an isolated element b' in $\Gamma(\mathcal{M}) \setminus \mathcal{B}$. In any case, we see that $\tilde{f} = f \cup (b, b')$ extends f at b , as wanted. \square

Recall that given a group $(G, \cdot, 1)$ and X a subset of G , we denote by $(\Gamma(X), E)$ or simply $\Gamma(X)$ the graph with vertices the classes of elements in X modulo \sim and edge relation E given by commutation:

$$[a]_{\sim} E [b]_{\sim} \iff ab = ba.$$

Recall also that, in a Mekler group M , if X^v is a choice of representatives of elements of type 1^v modulo \sim , then $\Gamma(X^v) = C$ is the nice graph associated to M . The following lemma is almost immediate:

Lemma 3.22. *Let $M = (G, \cdot, 1)$ be the Mekler group of an infinite nice graph C , and X be a transversal of M . Then $\Gamma(X)$ is a cover of C . Moreover, if M/Z is \aleph_0 -saturated, then $\Gamma(X)$ is an infinite cover of C .*



Proof. That $\Gamma(X)$ is a cover of C is a simple exercise (see [CH19, Fact 2.10]). Assume that M/Z is \aleph_0 -saturated. Let $a \in C$ be a vertex with an infinite set of

neighbours $N_a \subset C$. Then “ x is of type p and has handle a ” is satisfied by products of (at least three) elements of N_a . The type $p(x_i : i \in \mathbb{N})$:

$$\bigcup_{i \in \mathbb{N}} \{x_i \text{ is of type } p, \text{ handle } a \text{ and linearly independent over all elements of type } 1^\nu\}$$

is consistent. As M/Z is \aleph_0 -saturated, M realises this type, and clearly, for any choice of transversal $X = X^\nu \cup X^p \cup X^l$, a is connected to infinitely many elements $[g]_\sim, g \in X^p$. Similarly, if C is infinite, there are infinitely many elements of type 1^l and the type $q(x_i : i \in \mathbb{N})$:

$$\bigcup_{i \in \mathbb{N}} \{x_i \text{ is of type } 1^l, \text{ and linearly independent over all elements of types } 1^\nu \text{ and } p\}$$

is consistent. Again, by \aleph_0 -saturation, there will be infinitely many isolated point in $\Gamma(X) \setminus C$ for any choice of transversal. \square

Proof of Proposition 3.17. We denote, only in this proof, the group sort by G , instead of G/Z , and we use multiplicative notation, despite the group being abelian. In the proof below, we follow the 5-step strategy outlined earlier.

Let \mathcal{M} and \mathcal{N} be Mekler groups of exponent p . Assume that \mathcal{M} is \aleph_0 -saturated and \mathcal{N} is $|\mathcal{M}|^+$ -saturated. Let $\mathcal{A} \subseteq \mathcal{M}, \mathcal{B} \subseteq \mathcal{N}$ and let $f : \mathcal{A} \hookrightarrow \mathcal{B}$ be a partial isomorphism. The embedding f consists of two (compatible) functions $f_G : G_{\mathcal{A}} \rightarrow G_{\mathcal{B}}, f_C : C_{\mathcal{A}} \rightarrow C_{\mathcal{B}}$, for the sorts G and C , respectively. Assume that $f_C : C_{\mathcal{A}} \hookrightarrow C_{\mathcal{B}}$ is elementary. Notice that $G_{\mathcal{A}}$ is already a subgroup since the structure is closed under multiplication and the groups are of exponent p . We will show that f can be extended to a full embedding of \mathcal{M} in \mathcal{N} .

Step 0: By elementarity, we can extend $f_C : C_{\mathcal{A}} \hookrightarrow C_{\mathcal{B}}$ to $C(\mathcal{M})^{\text{eq}} \hookrightarrow C(\mathcal{N})^{\text{eq}}$, and $f \cup f_C$ is still a partial isomorphism, as C is a closed sort, and

$$\bigcup_{n,m} S_{n,m}(G_{\mathcal{A}}) \cup S'_{n,m}(G_{\mathcal{A}}) \subseteq C_{\mathcal{A}}.$$

For simplicity, we update the notation and let f denote the extension.

In the next two steps, we first extend the partial isomorphism f to elements of type 1^ν and p , which appear in the support of an element of $G_{\mathcal{A}}$. Let therefore $g \in (\bigcup_{n,m} A_{n,m}) \cap G_{\mathcal{A}}$. Let (n, m) be minimal for the reverse lexicographic order such that $g \in A_{n,m}$ and write $g = a_0 \cdots a_{n-1} \cdot g_0 \cdots g_{m-1}$ as a product of n elements of type 1^ν and m elements of type p .

Step 1: We reduce to the case where $a_0, \dots, a_{n-1} \in G_{\mathcal{A}}$.

Assume, for example, that $a_0 \notin G_{\mathcal{A}}$. In this case, since f preserves the function $S_{n,m}, f_C([a_0]_\sim) \in A_{n,m}(f_G(g))$, and this means that we can find $b_0 \in G(\mathcal{N})$ of type 1^ν such that

$$[b_0]_\sim = f_C([a_0]_\sim) \quad \text{and} \quad [b_0]_\sim \notin S_{n,m}(f_G(g) \cdot b_0^{-1}).$$

Then we can extend f_G by setting $f_G(a_0) = b_0$. We show that this function preserves quantifier-free formulas. To see this, fix $g' \in G_{\mathcal{A}}$ and $0 < r < p$. We should show that, for every integer s, t ,

- (i) $A_{s,t}(g'a_0^r)$ holds if and only if $A_{s,t}(f(g')b_0^r)$ holds,
- (ii) $f(S_{s,t}(g'a_0^r)) = S_{s,t}(f(g')b_0^r)$ if the first point holds,
- (iii) $f(S'_{s,t}(g'a_0^r)) = S'_{s,t}(f(g')b_0^r)$ if the first point holds.

Indeed, it in particular shows that $f_G \cup \{(a_0, b_0)\}$ preserves equations of the form $g_0 \cdots g_{m-1}a_0^r = 1$ where $r < p$ (since such an equation holds if and only if $A_{0,0}(g_0 \cdots g_{m-1}a_0^r)$ holds). It will follow that for a quantifier-free formula $\phi(x, y, c)$ with parameters c in $C(\mathcal{M})$ and g' a tuple from $G_{\mathcal{A}}$,

$$\mathcal{M} \models \phi(g', a_0, c) \text{ if and only if } \mathcal{N} \models \phi(f_G(g'), f_G(a_0), f_G(c)).$$

We start with a trivial case:

- *Case 0:* $g' \notin \bigcup_{k,l} A_{k,l}$. Then for all integers s, t , $A_{s,t}(g'a_0^r)$ and $A_{s,t}(f(g')b_0^r)$ don't hold and $S_{s,t}(g'a_0^r) = u$, $S_{s,t}(f(g')b_0^r) = u$.

Therefore, we may assume that $g' \in \bigcup_{k,l} A_{k,l}$, and let (k, ℓ) be minimal (in the reverse lexicographic order) such that $g' \in A_{k,\ell}$. We should first find the smallest (k', ℓ') such that $g'a_0^r \in A_{k',\ell'}$ and then express $S_{k',\ell'}(g'a_0^r)$ and $S'_{k',\ell'}(g'a_0^r)$ in terms of g, g' and $[a_0]_{\sim}$. We distinguish cases:

- *Case 1:* $[a_0]_{\sim} \notin S_{k,\ell}(g')$ and $[a_0]_{\sim} \notin S'_{k,\ell}(g')$.
Then $(k', \ell') = (k + 1, \ell)$ and

$$S_{k+1,\ell}(g'a_0^r) = S_{k,\ell}(g') \cup \{[a_0]_{\sim}\}, \quad S'_{k+1,\ell}(g'a_0^r) = S'_{k,\ell}(g').$$

- *Case 2:* $[a_0]_{\sim} \notin S_{k,\ell}(g')$ and $[a_0]_{\sim} E\alpha$ or $[a_0]_{\sim} = \alpha$ for some element $\alpha \in S'_{k,\ell}(g')$.
Then $(k', \ell') = (k, \ell)$ and $S_{k,\ell}(g'a_0^r) = S_{k,\ell}(g')$ and $S'_{k,\ell}(g'a_0^r) = S'_{k,\ell}(g')$.
- *Case 3:* $[a_0]_{\sim} \in S_{k,\ell}(g')$ (and, by minimality of (k, ℓ) , $\neg[a_0]_{\sim} E\alpha$ for all elements $\alpha \in S'_{k,\ell}(g')$).

We now need to consider two subcases:

- 3A: $A_{k+n,\ell+m}(g'g^r)$ and $[a_0]_{\sim} \notin S_{k+n,\ell+m}(g'g^r)$. This is the subcase where the powers of a_0 in g' and in g^r cancel out in the product $g'g^r$.
Then $(k', \ell') = (k - 1, \ell)$ and

$$S_{k-1,\ell}(g'a_0^r) = S_{k,\ell}(g') \setminus \{[a_0]_{\sim}\}, \quad S'_{k-1,\ell}(g'a_0^r) = S'_{k,\ell}(g').$$

- 3B: $A_{k+n,\ell+m}(g'g^r)$ and $[a_0]_{\sim} \in S_{k+n,\ell+m}(g'g^r)$. In this subcase, the powers of a_0 in g' and in g^r don't cancel.

Then we have $A_{k,\ell}(g'a_0^r)$ with $(k', \ell') = (k, \ell)$ minimal, and

$$S_{k,\ell}(g'a_0^r) = S_{k,\ell}(g') \quad \text{and} \quad S'_{k,\ell}(g'a_0^r) = S'_{k,\ell}(g').$$

The same statement holds in \mathcal{N} if we replace a_0 with b_0 , g' with $f_G(g')$ and g with $f_G(g)$. Therefore, conditions (i)–(iii) hold for (k', l') . It will automatically follow for the other integers (s, t) , as (k', l') , $S_{k',l'}$ and $S'_{k',l'}$ determine all the other pairs of integers such that $A_{s,t}(g'a_0^r)$ holds and determine $S_{s,t}(g'a_0^r)$ and $S'_{s,t}(g'a_0^r)$.

This argument shows that $f_G \cup \{(a_0, b_0)\}$ is a partial isomorphism. We can reset the notation by setting $f_G = f_G \cup \{(a_0, b_0)\}$, and therefore assume that $a_0 \in G_{\mathcal{A}}$.

Repeating the argument above, we may assume that for any $g = a_0 \cdots a_{n-1} \cdot g_0 \cdots g_{m-1} \in G_{\mathcal{A}} \cap A_{n,m}$, we have $a_0, \dots, a_{n-1} \in G_{\mathcal{A}}$.

Fix again $g \in (\bigcup_{n,m} A_{n,m}) \cap G_{\mathcal{A}}$ with (n, m) minimal (for the reverse lexicographic order) such that $g = a_0 \cdots a_{n-1} \cdot g_0 \cdots g_{m-1}$ a product of n elements of types 1^v and m elements of type p . By the previous step, we may assume that $n = 0$ and $g = g_0 \cdots g_{m-1}$.

Step 2: We reduce to the case where $g_0, \dots, g_{m-1} \in G_{\mathcal{A}}$.

This is similar to Step 1. Assume, for example, that $g_0 \notin G_{\mathcal{A}}$. By minimality of $(0, m)$, g_0 is not a finite product of elements of type 1^v . Consider the set $V_g^{h(g_0)}$ of elements g' where $h(g_0)$ “vanishes” in the product $g'g$:

$$\{g' \in G \mid \text{for } (k, \ell) \in \mathbb{N}^2 \text{ minimal such that } A_{k,\ell}(g'), \\ A_{k,\ell+m-2}(g'g) \text{ and } h(g_0) \notin S_{k,\ell+m-2}(g'g) = S_{k,\ell}(g') \cup S_{0,m}(g) \setminus \{h(g_0)\}\}.$$

(k, ℓ) depend on g' , but we hide this dependence for the sake of notational simplicity). Then, consider the partial type $p(x)$:

$$\{x \text{ of type } p \text{ and handle } h(g_0) : A_{k,\ell-1}(g'x), S_{k,\ell-1}(g'x) = S_{k,\ell}(g') \\ \text{and } S'_{k,\ell-1}(g'x) = S'_{k,\ell}(g') \setminus \{h(g_0)\}\}.$$

Since f is a partial isomorphism, we can consider the image by f of the type $p(x)$:

$$f(p(x)) := \{x \text{ of type } p \text{ and handle } f(h(g_0)) : A_{k,\ell-1}(f(g')x), \\ S_{k,\ell-1}(f(g')x) = S_{k,\ell}(f(g')) \text{ and } S'_{k,\ell-1}(f(g')x) = S'_{k,\ell}(f(g')) \setminus f(h(g_0))\}.$$

By Lemma 3.19, one can see that this type is consistent. Since \mathcal{N} is saturated, we may find h_0 in $G(\mathcal{N})$ satisfying $f(p(x))$.

Then we extend f_G by setting $f_G(g_0) = h_0$. We show that this function preserves quantifier-free formulas. To show this, we fix $g' \in G_{\mathcal{A}}$ and $0 < r < p$. Let (k, ℓ) be minimal for the reverse lexicographic order such that $g' \in A_{k,\ell}$. To see why can assume that such (k, ℓ) exist, as in Step 1, we have:

- *Case 0:* $g' \notin \bigcup_{k,l} A_{k,l}$. Then for all integers s, t , $A_{s,t}(g'g_0^r)$ and $A_{s,t}(f(g')h_0^r)$ don't hold and $S_{s,t}(g'g_0^r) = u$, $S_{s,t}(f(g')h_0^r) = u$.

Therefore, we may indeed assume that $g' \in \bigcup_{k,l} A_{k,l}$, and we need to find the smallest (k', ℓ') such that $g'g_0^r \in A_{k',\ell'}$. Again, there are several cases to consider:

- *Case 1:* $\neg h(g_0)E\alpha$ for all $\alpha \in S_{k,\ell}(g')$ and $h(g_0) \notin S_{k,\ell}(g')$.
Then $(k', \ell') = (k, \ell + 1)$, $S_{k,\ell+1}(g'g'_0) = S_{k,\ell}(g')$ and $S'_{k,\ell+1}(g'g'_0) = S'_{k,\ell}(g') \cup \{h(g_0)\}$.
- *Case 2:* For some $t > 0$, $h(g_0)$ is related (according to E) with exactly t many elements of $S_{k,\ell}(g')$ (and by minimality of (k, ℓ) , $h(g_0) \notin S'_{k,\ell}(g')$).
Then $(k', \ell') = (k - t, \ell + 1)$, $S_{k-t,\ell+1}(g'g'_0) = S_{k,\ell}(g') \setminus \{\alpha \in S_{k,\ell}(g') \mid \alpha E h(g_0)\}$, and $S'_{k-t,\ell+1}(g'g'_0) = S'_{k,\ell}(g') \cup \{h(g_0)\}$.
- *Case 3:* $h(g_0) \in S'_{k,\ell}(g')$ (then by minimality of (k, ℓ) , $\neg h(g_0)E\alpha$ for all $\alpha \in S_{k,\ell}(g')$).

There are two subcases:

- 3A: There is an integer t such that $A_{t,\ell+m-2}(g'g^r)$ and $h(g_0) \notin S'_{t,\ell+m-2}(g'g^r)$. This means that the handle $h(g_0)$ vanishes in the product g'_0g' . At the cost of multiplying g' with elements of type 1^ν , we may assume that $g' \in V_{g^r}^{h(g_0)}$. We then have $(k', \ell') = (k, \ell - 1)$, $A_{k,\ell-1}(g'g'_0)$ and

$$S'_{k,\ell-1}(g'g'_0) = S'_{k,\ell}(g') \setminus \{h(g_0)\}, \quad S_{k,\ell-1}(g'g'_0) = S_{k,\ell}(g').$$

- 3B: There is no such integer t .

Then we have $(k', \ell') = (k, \ell)$ and

$$S_{k,\ell}(g'g'_0) = S_{k,\ell}(g'), \quad S'_{k,\ell}(g'g'_0) = S'_{k,\ell}(g').$$

By choice of h_0 , we have that same statement holds in \mathcal{N} with h_0 instead of g_0 , $f_G(g')$ instead of g' and $f_G(g)$ instead of g . It follows that $f_G \cup \{(g_0, h_0)\}$ preserves all the predicates $A_{k,l}$ and the functions $S_{k,l}, S'_{k,l}$. This argument shows that $f_G \cup \{(g_0, h_0)\}$ is a partial isomorphism. We can reset the notation by setting $f_G = f_G \cup \{(g_0, h_0)\}$, and therefore assume that $h_0 \in G_{\mathcal{A}}$.

By repeating Step 2, we may assume that for any $g = a_0 \cdots a_{n-1} \cdot g_0 \cdots g_{m-1} \in A_{n,m} \cap G_{\mathcal{A}}$, we have $g_0, \dots, g_{m-1} \in G_{\mathcal{A}}$.

Step 3: We find a transversal X compatible with the substructure \mathcal{A} , i.e., such that $G_{\mathcal{A}} = \langle X_{\mathcal{A}} \rangle$ where $X_{\mathcal{A}} = X \cap G_{\mathcal{A}}$.

We may, indeed, consider first a set $X_{\mathcal{A}}^\nu$ of representatives in \mathcal{A} of the \sim -classes of elements of type 1^ν in \mathcal{A} . Then, $X_{\mathcal{A}}^p$ is the set of representative in \mathcal{A} of \sim -classes of elements of type p independent over $X_{\mathcal{A}}^\nu$. If a product $g_1 \cdots g_m \in \mathcal{A}$ of elements of type p is a finite product $a_1 \cdots a_n \in G$ of elements of type 1^ν , these elements are, by the previous steps, in \mathcal{A} , which is a contradiction. Therefore, they must also be independent over all elements of type 1^ν in G . Finally, $X_{\mathcal{A}}^l$ is the set of representatives of elements in \mathcal{A} of type of 1^l independent over $X_{\mathcal{A}}^\nu, X_{\mathcal{A}}^p$. By the previous steps, and with a similar argument, they are also independent over all elements of type p and 1^ν in G .

Step 4: We extend f to a full embedding of $\Gamma(X)$, then to a full embedding of \mathcal{M} .

By Lemma 3.18, and since f preserves quantifier-free formulas, $f(X_A)$ can also be extended to a transversal Y of B . Denote by $\Gamma(X)$ the graph with vertices $\{[x]_{\sim} : x \in X\}$ and edges $\{([x]_{\sim}, [y]_{\sim}) : xRy \in \mathcal{M}\}$ (this is exactly the graph $\Gamma(\tilde{X})$ where \tilde{X} is a lift of X in the Mekler group).

We therefore have a partial isomorphism f_{Γ} between $\Gamma(X)$ and $\Gamma(Y)$ as graphs. By saturation and Lemma 3.22, these graphs are infinite covers of $C(\mathcal{M})$ and $C(\mathcal{N})$, respectively. By $|\Gamma(X)|$ -saturation of $\Gamma(Y)$ (see [CH19, Fact 2.11(3)]) and by quantifier elimination relative to C (see Lemma 3.21) we can extend it to a full embedding $f_{\Gamma} : \Gamma(X) \hookrightarrow \Gamma(Y)$.

This embedding gives immediately an extension of $f|_{\langle X_A \rangle}$ to

$$f : G(\mathcal{M}) = \langle X \rangle \hookrightarrow G(\mathcal{N}) = \langle Y \rangle$$

as wanted. □

4. Transfer principles

A *transfer principle for Mekler groups* (or, in our context, simply a *transfer principle*) is any statement characterising a model-theoretic property of the Mekler group M at the level of the graph C . Quite a lot of transfer principles for Mekler groups are already known. We recall below the current state of affairs. The reader can refer to the cited sources for relevant definitions and proofs.

Fact 4.1. *A Mekler group M has the property P if and only if its associated graph C has the property P , where P is one of the following properties:*

- λ -stability for every cardinal λ [Mek81; Hod93];
- CM -triviality [Bau02];
- the n -independence property, for every $n \in \mathbb{N}$ [CH19];
- the tree property of the second kind [CH19];
- the first and second strong order properties [Ahn20];
- the antichain tree property [AKL22].

In this section, we prove new transfer principles. Our main tool is the relative quantifier elimination result developed in the previous section, and our method is again a three-step reduction. In the first subsection we discuss pairs of Mekler groups. The second subsection is dedicated to dividing lines and includes some negative results.

4.1. Model completeness and Stable embeddedness. From relative quantifier elimination, one can always deduce a relative model completeness result:

Lemma 4.2. *Assume that a complete theory T eliminates quantifiers relative to a closed sort Σ in a language \mathcal{L} , and let $\mathcal{N}, \mathcal{M} \models T$ be two structures such that*

$$\mathcal{M} \subseteq_{\mathcal{L}} \mathcal{N} \quad \text{and} \quad \Sigma(\mathcal{M}) \preceq \Sigma(\mathcal{N}).$$

Then $\mathcal{M} \preceq \mathcal{N}$.

However, we can often optimise this result by weakening the first condition, considering only a reduct of the language \mathcal{L} and adding a more relevant algebraic requirement. This is what we shall do in the next paragraphs: We will deduce model completeness statements from the three relative quantifier elimination results stated in the previous section (Propositions 4.6, 4.8 and 4.10 below). We will then combine these results in the last part of this subsection to obtain model completeness for Mekler groups relative to their graphs (Proposition 4.14).

We will implicitly use that “elementary extension” is a property of a pair of structures, and does not depend on the language in which the structures are taken:

Fact 4.3. *Let $\mathcal{M} \preceq \mathcal{N}$ be an elementary extension. Then $\mathcal{M}^{\text{eq}} \preceq \mathcal{N}^{\text{eq}}$.*

At the same time, we will look at stably embedded pairs of Mekler groups. Let us now recall the definition:

Definition 4.4 (stable embeddedness). Let T be a complete theory. An elementary extension $\mathcal{M} \preceq \mathcal{N}$ is called *stably embedded* if for every formula $\phi(x, b)$ with $b \in \mathcal{N}$, there is a formula $\psi(x, c)$ with $c \in \mathcal{M}$ such that $\phi(\mathcal{M}, b) = \psi(\mathcal{M}, c)$. In this case, we write $\mathcal{M} \preceq^{\text{st}} \mathcal{N}$.

The pair is called *uniformly stably embedded* if, in addition, the choice of the formula $\psi(x, z)$ depends only on $\phi(x, y)$, and is independent of the parameters b . In this case, we write $\mathcal{M} \preceq^{\text{ust}} \mathcal{N}$.

Stable embeddedness says that every subset of \mathcal{M} which is *externally* definable in \mathcal{N} is *internally* definable. The proof of Theorem B will follow, without much effort, from quantifier elimination, as the characterisation does not require any new assumptions on the pair of Mekler groups. Again, stable embeddedness is a property of a pair of structures and does not depend on the language in which the structures are taken:

Fact 4.5. (See e.g. [Tou23].) *Let $\mathcal{M} \preceq^{\text{st}} \mathcal{N}$ (resp. $\mathcal{M} \preceq^{\text{ust}} \mathcal{N}$) be a (uniformly) stably embedded elementary (resp. **uniformly** stably embedded) pair of structures. Then $\mathcal{M}^{\text{eq}} \preceq^{\text{st}} \mathcal{N}^{\text{eq}}$ (resp. $\mathcal{M}^{\text{eq}} \preceq^{\text{ust}} \mathcal{N}^{\text{eq}}$).*

This fact justifies our stepwise approach and will be used implicitly.

4.1.1. First reduction: From G to $\mathcal{F}(G)$.

Proposition 4.6 (relative model completeness). *Let G and G' be 2-nilpotent groups of exponent p . Assume that G is a subgroup of G' such that $Z(G) \subseteq Z(G')$. Then, $G \preceq G'$ if and only if $\mathcal{F}(G) \preceq \mathcal{F}(G')$, as alternating bilinear systems.*

Proof. Since $Z(G) \subseteq Z(G')$, the following diagram commutes:

$$\begin{array}{ccccc}
 & & \rho_Z & & \\
 & \swarrow & & \searrow & \\
 Z(G) & \longrightarrow & G & \xrightarrow{\pi_{G/Z}} & G/Z(G) \\
 \downarrow & & \downarrow & & \downarrow \\
 Z(G') & \longrightarrow & G' & \longrightarrow & G'/Z(G')
 \end{array}$$

In particular, we have that $G \subseteq_{\mathcal{L}} G'$, where $\mathcal{L} = \mathcal{L}_{G, \mathcal{F}(G)}$ is the language of relative quantifier elimination in Fact 3.2. We may conclude by Lemma 4.2. \square

Proposition 4.7 (relative stable embeddedness). *Let $G \preceq G'$ be 2-nilpotent groups of exponent p . Then*

$$G \preceq^{\text{st}} G' \text{ if and only if } \mathcal{F}(G) \preceq^{\text{st}} \mathcal{F}(G')$$

and

$$G \preceq^{\text{ust}} G' \text{ if and only if } \mathcal{F}(G) \preceq^{\text{ust}} \mathcal{F}(G').$$

Proof. Assume $\mathcal{F}(G) \preceq^{\text{st}} \mathcal{F}(G')$. Let b be a tuple from G' , and $\phi(x, y)$ be an \mathcal{L}_{grp} -formula. We want to define $\phi(G, b)$ using parameters in G . By Fact 3.2, $\phi(x, b)$ is equivalent to a formula of the form

$$\phi_{\mathcal{F}(G)}(\pi(t'(x, b)), \rho(t(x, b)))$$

where $\phi_{\mathcal{F}(G)}(y, z, w)$ is a formula in the language $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta\}$, and t' and t are tuples of \mathcal{L}_{grp} -terms. To simplify the argument, we assume that t consists of only one term.

Since π is a morphism, we have $\pi(t'(x, b)) = t'(\pi(x), \pi(b))$ (where the term t' is understood additively in the right-hand side of the equality).

Therefore, after incorporating t' into $\phi_{\mathcal{F}(G)}$, the formula becomes equivalent to one of the form

$$\phi'_{\mathcal{F}(G)}(\pi(x), \pi(b), \rho(t(x, b))).$$

Notice, however, that ρ is a morphism only when it is restricted to the centre, Z . In particular, if $t(x, b) = z(x, b)$ is a product of commutators, it takes value in Z , and we have

$$\rho(z(x, b)) = z_{\beta}(\pi(x), \pi(b)),$$

where z_{β} is a sum of β -terms (we can separate the variables x from the parameters b using bilinearity of β).

In general, $t(x, b)$ need not be a product of commutators, so we have to reduce to that case. If there is no $a_0 \in G$ such that $t(a_0, b)$ is in Z , then, by definition of ρ , we can replace $\rho(t(x, b))$ by 0 (recall that groups in $\mathcal{L}_{\mathcal{F}(G)}$ are written additively). So, assume that such an element a_0 exists. Since the group G' is 2-nilpotent, we

can rewrite the term $t(x, y)$ as $t_1(x)t_2(y)z(x, y)$ where $t_1(x)$, $t_2(y)$, and $z(x, y)$ are \mathcal{L}_{grp} -terms and $z(x, y)$ is a product of commutators (and therefore takes values in Z). Then, for all $a \in G$, we have

$$\begin{aligned} t(a, b) \in Z &\iff t_1(a)t_2(b)z(a, b) \in Z \\ &\iff t_1(a)t_1(a_0)^{-1}t_1(a_0)t_2(b)z(a_0, b) \in Z \\ &\iff t_1(a)t_1(a_0)^{-1}t(a_0, b) \in Z \\ &\iff t_1(a)t_1(a_0)^{-1} \in Z. \end{aligned}$$

We rewrite $t(x, b)$ as follows:

$$\begin{aligned} t(x, b) &= t_1(x)t_1(a_0)^{-1}t_1(a_0)t_2(b)z(x, b) \\ &= t_1(x)t_1(a_0)^{-1}t(a_0, b)z(a_0, b)^{-1}z(x, b). \end{aligned}$$

Let $\sigma(x, a_0, b)$ be the term

$$\sigma(x, a_0, b) := \rho(t_1(x)t_1(a_0)^{-1}) + \rho(t(a_0, b)) + \rho(z(a_0, b)^{-1}) + \rho(z(x, b)).$$

Then our formula is equivalent to

$$\begin{aligned} (t_1(x)t_1(a_0)^{-1} \in Z \wedge \phi'_{\mathcal{F}(G)}(\pi(x), \pi(b), \sigma(x, a_0, b))) \\ \vee (t_1(x)t_1(a_0)^{-1} \notin Z \wedge \phi'_{\mathcal{F}(G)}(\pi(x), \pi(b), 0)). \end{aligned}$$

As before, we can replace the terms $\rho(z(x, b))$ in $\sigma(x, a_0, b)$ by equivalent terms of the form $z_\beta(\pi(x), \pi(b))$ which are the sum of β -terms. Therefore, we get a formula of the form

$$\phi''_{\mathcal{F}(G)}(\pi(x), \rho(t_1(x)t_1(a_0)^{-1}), \pi(b), \rho(t(a_0, b))).$$

Since $\mathcal{F}(G) \preceq^{\text{st}} \mathcal{F}(G')$, there is a formula $\psi_{\mathcal{F}(G)}(x_{\mathcal{F}}, y_{\mathcal{F}}, a_{\mathcal{F}})$ where $a_{\mathcal{F}} \in \mathcal{F}(G)$ such that

$$\phi''_{\mathcal{F}(G)}(\mathcal{F}(G)^{|x|+1}, \pi(b), \rho(t(a_0, b))) = \psi_{\mathcal{F}(G)}(\mathcal{F}(G)^{|x|+1}, a_{\mathcal{F}}).$$

At the end, we have that $\phi(G, b) = \psi(G, a)$ where $a \in G$ is such that $\pi(a) = a_{\mathcal{F}}$, and where $\psi(x, a)$ is the formula

$$\psi_{\mathcal{F}(G)}(\pi(x), \rho(t_1(x)t_1(a_0)^{-1}), \pi(a)).$$

We have shown that G is stably embedded in G' . The uniform stable embeddedness statement follows the same way once we notice that the formula $\phi''_{\mathcal{F}(G)}$ does not depend on the choice of b —it depends only on whether or not there is $a_0 \in G$ such that $t(a_0, b)$ is in Z , and the case distinction can be handled using parameters in G . \square

4.1.2. Second reduction: From $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta\}$ to $(\mathcal{V}, +, 0, R)$. Recall that we denote by $\bigcup_n \mathcal{W}_n := \langle \beta(\mathcal{V}, \mathcal{V}) \rangle$ the subspace of \mathcal{W} generated by all terms $\beta(v, v')$, where $v, v' \in \mathcal{V}$. We can always find a complement \mathcal{W}_ω such that $\mathcal{W} = \langle \beta(\mathcal{V}, \mathcal{V}) \rangle \oplus \mathcal{W}_\omega$.

We recommend that the reader review Notation 3.8 before delving into the proofs.

Proposition 4.8 (relative model completeness). *Let $(\mathcal{V}, \mathcal{W}, \beta) \subseteq (\mathcal{V}', \mathcal{W}', \beta')$ be two alternating bilinear systems satisfying property $(*_f)$ of Corollary 3.6. Denote by R (resp. R') the relation induced on \mathcal{V} (resp. on \mathcal{V}') by β (resp. β'). Assume that a complement U of \mathcal{W}_ω is disjoint from $\bigcup_n \mathcal{W}'_n = \langle \beta(\mathcal{V}', \mathcal{V}') \rangle$. Then $(\mathcal{V}, \mathcal{W}, \beta) \preceq (\mathcal{V}', \mathcal{W}', \beta')$ if and only if $(\mathcal{V}, +, 0, R) \preceq (\mathcal{V}', +, 0, R')$.*

Proof. Since $\beta'|_{\mathcal{V}^2} = \beta$, the embedding $(\mathcal{V}, \mathcal{W}, \beta) \subseteq (\mathcal{V}', \mathcal{W}', \beta')$ preserves π_A for $A \in \mathbb{A}_n(\mathbb{F}_p)$. It also preserves the functions f_n for $n \in \mathbb{N}$ restricted to $\langle \beta(\mathcal{V}, \mathcal{V}) \rangle$. The condition on the complement U ensures that on $\mathcal{W} \setminus \langle \beta(\mathcal{V}, \mathcal{V}) \rangle$, we have $f_n(w) = f'_n(w) = u$. Therefore, the embedding also preserves the language of relative quantifier elimination stated in Proposition 3.10. Again, we may conclude by Lemma 4.2. \square

Proposition 4.9 (relative stable embeddedness). *Let $(\mathcal{V}, \mathcal{W}, \beta) \preceq (\mathcal{V}', \mathcal{W}', \beta')$ be an elementary extension of bilinear systems satisfying property $(*_f)$ of Corollary 3.6. With the same notation as in the previous proposition, we have*

$$(\mathcal{V}, \mathcal{W}, \beta) \preceq^{\text{st}} (\mathcal{V}', \mathcal{W}', \beta') \text{ if and only if } (\mathcal{V}, +, 0, R) \preceq^{\text{st}} (\mathcal{V}', +, 0, R'),$$

and

$$(\mathcal{V}, \mathcal{W}, \beta) \preceq^{\text{ust}} (\mathcal{V}', \mathcal{W}', \beta') \text{ if and only if } (\mathcal{V}, +, 0, R) \preceq^{\text{ust}} (\mathcal{V}', +, 0, R').$$

Proof. Assume $(\mathcal{V}, +, 0, R) \preceq^{\text{st}} (\mathcal{V}', +, 0, R')$ and let $\phi(x_{\mathcal{V}}, x_{\mathcal{W}}, v', w')$ be a formula in the language of $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta\}$ with parameters $v' \in \mathcal{V}'$ and $w' \in \mathcal{W}'$. We need to find a formula $\psi(x_{\mathcal{V}}, x_{\mathcal{W}}, v, w)$ with parameters v in \mathcal{V} and parameters w in \mathcal{W} such that

$$\phi(\mathcal{V}^{|x_{\mathcal{V}}|}, \mathcal{W}^{|x_{\mathcal{W}}|}, v', w') = \psi(\mathcal{V}^{|x_{\mathcal{V}}|}, \mathcal{W}^{|x_{\mathcal{W}}|}, v, w).$$

By Proposition 3.10, the formula $\phi(x_{\mathcal{V}}, x_{\mathcal{W}}, v', w')$ is equivalent to a Boolean combination of formulas of the following form:

- $a \cdot x_{\mathcal{W}} = w''$ where a is a $x_{\mathcal{W}}$ tuple in \mathbb{F}_p , $w'' \in \mathcal{W}'$, (the operation \cdot denotes here the sum of products of the components a and $x_{\mathcal{W}}$).
- $\phi_{\mathcal{V}}(f_n(a_1 \cdot x_{\mathcal{W}} + w'_1), \dots, f_n(a_k \cdot x_{\mathcal{W}} + w'_k), x_{\mathcal{V}}, v')$ where $\phi_{\mathcal{V}}$ is a formula in the language of $(\mathcal{V}, +, 0, R)^{\text{eq}}$, the a_i 's are $|x_{\mathcal{W}}|$ -tuples in \mathbb{F}_p , and the w'_i 's are parameters in \mathcal{W}' (the operation \cdot is as in the previous point).

The first formula, which is only an equality between terms, is not a problem. If there is no $w \in \mathcal{W}^{|\mathcal{X}\mathcal{W}|}$ satisfying the formula $a \cdot x_{\mathcal{W}} = w''$, we can replace this formula with \perp and we are done. Otherwise, it means that w'' belongs to \mathcal{W} and we don't need to find a new parameter.

We analyse the second formula. To simplify, we assume that there is only one term in f_n , and we assume that the formula $\phi(x_{\mathcal{V}}, x_{\mathcal{W}}, v', \mathcal{W}')$ is given by

$$\phi_{\mathcal{V}}(f_n(a \cdot x_{\mathcal{W}} + w'), x_{\mathcal{V}}, v').$$

The general case follows by repeating the same process for each term.

Let \mathcal{W}'_n be the set of elements of the form $\sum_{i < j < n} \beta(u_i, u'_j)$ with $u_0, \dots, u_{n-1} \in \mathcal{W}'$. If for all $w \in \mathcal{W}^{|\mathcal{X}\mathcal{W}|}$, $a \cdot w + w'$ is not in \mathcal{W}'_n (and thus $f_n(a \cdot w + w') = u$ for all $w \in \mathcal{W}^{|\mathcal{X}\mathcal{W}|}$), then we can replace the term $f_n(a \cdot x_{\mathcal{W}} + w')$ by u .

Assume there is $w_0 \in \mathcal{W}^{|\mathcal{X}\mathcal{W}|}$ such that $a \cdot w_0 + w'$ is in \mathcal{W}'_n , then for all $w \in \mathcal{W}^{|\mathcal{X}\mathcal{W}|}$, we can write

$$a \cdot w + w' = a \cdot (w - w_0) + a \cdot w_0 + w'.$$

It follows that $a \cdot w + w'$ is in \mathcal{W}_n only if $a \cdot (w - w_0)$ is in \mathcal{W}_{2n} . Then, we have that $a \cdot w + w' \in \mathcal{W}_n$ and the value of $f_n(a \cdot w + w')$ can be expressed using $f_{2n}(a \cdot (w - w_0))$ and $f_n(a \cdot w_0 + w')$. Therefore, the formula

$$\phi_{\mathcal{V}}(f_n(a \cdot x_{\mathcal{W}} + w'), x_{\mathcal{V}}, v')$$

is equivalent to

$$\begin{aligned} & (f_{2n}(a \cdot (w - w_0)) \neq u \wedge \phi'_{\mathcal{V}}(f_{2n}(a \cdot (w - w_0)), f_n(a \cdot w_0 + w'), x_{\mathcal{V}}, v')) \\ & \vee (f_{2n}(a \cdot (w - w_0)) = u \wedge \phi_{\mathcal{V}}(u, x_{\mathcal{V}}, v')) \end{aligned}$$

where $\phi'_{\mathcal{V}}$ is some formula in the language $(\mathcal{V}, +, 0, R)^{\text{eq}}$.

Using the fact that $(\mathcal{V}, +, 0, R) \preceq^{\text{st}} (\mathcal{V}', +, 0, R)$, and considering $f_n(a \cdot w_0 + w')$ as parameters in $(\mathcal{V}')^{\text{eq}}$, we can find a formula $\psi_{\mathcal{V}}(x_{B_n}, x_{\mathcal{V}}, v_0)$ with parameters $v_0 \in \mathcal{V}$ such that for all $w \in \mathcal{W}^{|\mathcal{X}\mathcal{W}|}$ and $v \in \mathcal{V}^{|\mathcal{X}\mathcal{V}|}$,

$$(\mathcal{V}, +, 0, R) \models \phi_{\mathcal{V}}(f_n(a \cdot w + w'), v, v') \leftrightarrow \psi_{\mathcal{V}}(f_{2n}(a \cdot (w - w_0)), v, v_0).$$

In other words, if we set $\psi(x_{\mathcal{V}}, x_{\mathcal{W}}, v_0, w_0) = \psi_{\mathcal{V}}(f_{2n}(a \cdot (x_{\mathcal{W}} - w_0)), x_{\mathcal{V}}, v_0)$, we have

$$\phi(\mathcal{V}^{|\mathcal{X}\mathcal{V}|}, \mathcal{W}^{|\mathcal{X}\mathcal{W}|}, v', w') = \psi(\mathcal{V}^{|\mathcal{X}\mathcal{V}|}, \mathcal{W}^{|\mathcal{X}\mathcal{W}|}, v_0, w_0),$$

as wanted.

We have shown that $(\mathcal{V}, \mathcal{W}, \beta) \preceq^{\text{st}} (\mathcal{V}', \mathcal{W}', \beta')$. The uniform stable embeddedness statement follows the same way, as discussed in the proof of Proposition 4.6—the case distinctions can again be encoded with parameters in \mathcal{V} . \square

4.1.3. Third reduction: From $(\mathcal{V}, +, 0, R)$ to (C, R) . Let $M = (G, \cdot)$, $M' = (G', \cdot)$ be Mekler groups of respective graphs (C, E) and (C', E') , denote additively $(\mathcal{V}, +, 0, R)$ (resp. $(\mathcal{V}', +, 0, R')$) the abelian quotient G/Z (resp. G'/Z'), and where R (resp. R') denotes the relation induced by the commutativity relation on G (resp. G') on \mathcal{V} (resp. \mathcal{V}') as well as on C (resp. C').

Proposition 4.10 (relative model completeness). *In the notation above, assume that $(\mathcal{V}, +, 0, R) \subseteq (\mathcal{V}', +, 0, R')$ and that a transversal $X = X^v X^p X^t$ of \mathcal{V} that can be extended to a transversal $X' = X^{v'} X^{p'} X^{t'}$ of \mathcal{V}' (with $X^v \subseteq X^{v'}$, $X^p \subseteq X^{p'}$ and $X^t \subseteq X^{t'}$). Then $(\mathcal{V}, +, 0, R) \preceq (\mathcal{V}', +, 0, R')$ if and only if $(C, E) \preceq (C', E')$ as graphs.*

Proof. It is easy to see that the conditions for being a transversal imply that for all $n, m \in \mathbb{N}^2$, we have

$$S_{n,m}^{\mathcal{V}'}|_{\mathcal{V}} = S_{n,m}^{\mathcal{V}} \quad \text{and} \quad S_{n,m}^{(\mathcal{V}')'}|_{\mathcal{V}} = S_{n,m}^{\mathcal{V}}.$$

We can conclude using Lemma 4.2. □

Proposition 4.11 (relative stable embeddedness). *In the notation above, assume that $G \preceq G'$ is an elementary extension of Mekler groups. Then*

$$\begin{aligned} (\mathcal{V}, +, 0, R) &\preceq^{\text{st}} (\mathcal{V}', +, 0, R') \text{ if and only if } (C, E) \preceq^{\text{st}} (C', E'), \\ (\mathcal{V}, +, 0, R) &\preceq^{\text{ust}} (\mathcal{V}', +, 0, R') \text{ if and only if } (C, E) \preceq^{\text{ust}} (C', E'). \end{aligned}$$

A proof similar to that of Proposition 4.9 will not work, and we have to be more careful, since we lose some information with the support function $S_{n,m}$. We are going to use the predicates $A_{n,m}$ as a pre-distance, and approximate elements from \mathcal{V}' by elements from \mathcal{V} :

Definition 4.12. A best approximation in \mathcal{V} of a singleton $b \in \mathcal{V}'$ is an element $a_0 \in \mathcal{V}$ such that $f_{n,m}(b - a_0)$ holds and (n, m) is minimal for the reverse lexicographic order on \mathbb{N}^2 . If no such element exists, we say that b has no approximation in \mathcal{V} .

Lemma 4.13. *In the notation of Proposition 4.11, let a_0 be a best approximation of $b \in \mathcal{V}'$, with (k, l) minimal such that $A_{k,l}(b - a_0)$ holds. Then:*

- The element $b - a_0$ is the sum of k elements of types 1^v in $\mathcal{V}' \setminus \mathcal{V}$ and l elements of type p of \mathcal{V}' which are independent over $\langle \mathfrak{E}^v(\mathcal{V}') \rangle$.
- For all $a \in \mathcal{V}$, for all $n, m \in \mathbb{N}$, we have

$$\begin{aligned} A_{n+k,m+l}(b - a) &\iff A_{n,m}(a - a_0), \\ S_{n+k,m+l}(b - a) &= S_{k,l}(b - a_0) \sqcup S_{n,m}(a - a_0), \\ S'_{n+k,m+l}(b - a) &= S'_{k,l}(b - a_0) \sqcup S'_{n,m}(a - a_0). \end{aligned}$$

Proof. The first point easily follows from the minimality of (n, m) .

The second follows from the fact that, since the support of $b - a_0$ is in $\mathcal{V}' \setminus \mathcal{V}$, the support of $a - a_0 \in \mathcal{V}$ and $b - a_0$ must be disjoint. \square

Proof of Proposition 4.11. Assume $(C, E) \preceq^{\text{st}} (C', E')$ and let $\phi(x, b)$ be a group formula with parameters $b \in \mathcal{V}'$. By Proposition 3.17, the formula $\phi(x, b)$ in the structure $(G/Z, \cdot, 1, R)$ is equivalent to a Boolean combination of formulas of the following form:

- $t(x, b) = 1$, where $t(x, b)$ is a group term,
- $A_{n,m}(t(x, b))$, where $n, m \in \mathbb{N}$, $t(x)$ is a group term,
- $\phi_C(S_{n,m}(t(x, b)), S'_{n,m}(t(x, b)))$, where $n, m \in \mathbb{N}$, $t(x, y)$ is a tuple of \mathcal{L}_{grp} -terms and $\phi_C(x_C, x'_C)$ is a formula in the language (C, E) .

We handle the first form as before, so let us assume that $\phi(x, b)$ is a conjunction of formulas of the second and of the third form. Since the group is abelian, we can separate the variables x from the parameters b and assume that the terms are of the form $t(x) + b_1$ where $b_1 \in b$.

So, let us assume that our formula is given by

$$\phi(x, b) = \phi_C(S_{n,m}(t(x) + b), S'_{n,m}(t(x) + b)) \wedge A_{n,m}(t(x) + b),$$

and that $t(x)$ is a single term and b is a singleton— of course, the general case with many terms follows the same way.

If b does not have a best approximation in \mathcal{V} , then $A_{n,m}(t(a) + b)$ never holds, and we can replace the formula by \perp .

Assume that there exists a best approximation a_0 of b in \mathcal{V} , and let (k, l) be minimal (for the reverse lexicographic order of \mathbb{N}^2) such that $A_{k,l}(b - a_0)$ holds. Then, by Lemma 4.13, for all $a \in \mathcal{V}$ and $k, l \in \mathbb{N}$, we have

$$A_{n,m}(t(a) + b) \iff A_{n-k,m-l}(t(a) + a_0)$$

and

$$S_{n,m}(t(a) + b) = S_{k,l}(b - a_0) \sqcup S_{n-k,m-l}(t(a) + a_0),$$

and similarly for $S'_{n,m}(t(a) + b)$.

Observe that if $k > n$ or $l > m$, then neither $S_{n,m}(t(a) + b)$ nor $S'_{n,m}(t(a) + b)$ can hold. This means that $\phi(x, b)$ can be rewritten as a formula of the form

$$\phi(x, b) = \phi_C(S_{n-k,m-l}(t(a) + a_0), S'_{k,l}(t(a) + a_0), S_{k,l}(b - a_0), S'_{k,l}(b - a_0)),$$

and the terms $S_{k,l}(b - a_0), S'_{k,l}(b - a_0)$ can be seen as parameters in $(C', R)^{\text{eq}}$.

Using our assumption that $(C, E) \preceq^{\text{st}} (C', E')$, there is a formula $\psi_C(x, y, \alpha)$ in $(C', E')^{\text{eq}}$ with parameters $\alpha \in C$ such that

$$\phi_C(\mathcal{P}_{n-k}(C), \mathcal{P}_{m-l}(C), S_{k,l}(b - a_0), S'_{k,l}(b - a_0)) = \psi_C(\mathcal{P}_{n-k}(C), \mathcal{P}_{m-l}(C), \alpha),$$

where $\mathcal{P}_{n-k}(\mathbb{C})$ denotes the set of all subsets of size $\leq n - k$ of \mathbb{C} .

Let $a \in \mathbb{C}^{|\alpha|}$ be such that $S_{0,1}(a) = \{\alpha\}$, and by $\psi(x, a)$ the formula

$$\psi_{\mathbb{C}}(S_{n-k,m-l}(t(x) + a_0), \{S'_{n-k,m-l}(t(x) + a_0) : k \leq n, l \leq m\}, S_{1,0}(a))$$

(where we identify α with $\{\alpha\} = S_{1,0}(a)$).

We have

$$\phi(G(\mathbb{C}), b) = \psi(G(\mathbb{C}), a),$$

as desired and thus $(\mathcal{V}, +, 0, R) \preceq^{\text{st}} (\mathcal{V}', +, 0, R')$.

Notice that the formula ψ depends on the existence of a minimal (k, l) such that $A_{k,l}(b - a_0)$ holds, but if N is the largest integer that occurs in the formula, one need only consider the case for $k < N$ and $l < N$. Thus, the uniform statement follows in a similar way. \square

4.1.4. Reduction from $(G, \cdot, 1)$ to (\mathbb{C}, R) . Now we combine all the previous results.

Proposition 4.14 (relative model completeness). *Let \mathbb{C} and \mathbb{C}' be nice graphs and let $\mathbb{M} = (G, \cdot, 1)$ and $\mathbb{M}' = (G', \cdot, 1)$ be their respective Mekler groups. Assume that*

- $\mathbb{M} \subseteq \mathbb{M}'$ and $\mathbb{Z} \subseteq \mathbb{Z}'$;
- there is a transversal $X = X^v X^p X^l$ of G which can be extended to a transversal $X' = X^{v'} X^{p'} X^{l'}$ of G' (with $X^v \subseteq X^{v'}$, $X^p \subseteq X^{p'}$ and $X^l \subseteq X^{l'}$);
- there is a subspace H of \mathbb{Z} such that $G = \langle X \rangle \times H$ and such that $\langle X' \rangle \cap H = \emptyset$, where X and X' are transversal as above.

Then $\mathbb{M} \preceq \mathbb{M}'$ if and only if $\mathbb{C} \preceq \mathbb{C}'$ as graphs.

Some similarities may be observed with statements already known, such as [CH19, Lemma 2.14], which we will recall in Section 5. The proposition was perhaps already known, but we take the occasion to give a proof that makes use of our stepwise reductions.


Proof. The embedding of G in G' preserves the centre $\mathbb{Z} \subseteq \mathbb{Z}'$ by definition. The induced embedding $\mathcal{F}(G) \subseteq \mathcal{F}(G')$ preserves β automatically (since it comes from an embedding of groups). The subspace H of \mathbb{Z} does not intersect $\langle \beta(V, V') \rangle$. The set $X \bmod \mathbb{Z}$ is a transversal of G/\mathbb{Z} that extends to the transversal $X' \bmod \mathbb{Z}$ in G'/\mathbb{Z}' . Therefore, this proposition follows immediately from the previous ones, as we have

$$\begin{aligned} \mathbb{M} \preceq \mathbb{M}' &\iff \mathcal{F}(G) \preceq \mathcal{F}(G') && \text{(Proposition 4.6)} \\ &\iff (G/\mathbb{Z}, \cdot, R) \preceq (G'/\mathbb{Z}', \cdot, R') && \text{(Proposition 4.8)} \\ &\iff \mathbb{C} \preceq \mathbb{C}' && \text{(Proposition 4.10),} \end{aligned}$$


and the result follows. \square

The following example shows that it is not enough to assume that G is a subgroup of G' , and that the requirement of preserving a transversal is necessary.

Example 4.15. Consider the nice graphs C given by



and C' given by



Denote by M and M' their Mekler groups. The elementary embedding $C \rightarrow C'$ can easily be extended to an embedding of $M \hookrightarrow M'$ that is elementary by Proposition 4.14.

Now consider an elementary extension $N = \langle G, g \rangle$ of M , generated by M and a new element g of type 1^l and independent over M . Then there is a group embedding f of N in M' such that $f(Z) \subset Z'$ and such that $f(g)$ is of type 1^v and $[f(g)]_{\sim} = b_0$. In particular, for this embedding, we have $N \not\leq M'$, even if the embedding satisfies the first item of Proposition 4.14.

Theorem B (relative stable embeddedness). *Let $M \leq M'$ be an elementary extension of Mekler groups of nice graphs C and C' , respectively. Then*

$$M \leq^{st} M' \text{ if and only if } C \leq^{st} C',$$

$$M \leq^{ust} M' \text{ if and only if } C \leq^{ust} C'.$$

Proof. We simply combine the previous statements. Assume $G \leq G'$. Then, with the notation previously established, we have

$$M \leq^{(u)st} M' \iff \mathcal{F}(G) \leq^{(u)st} \mathcal{F}(G') \quad (\text{Proposition 4.7})$$


$$\iff (G/Z, \cdot, R) \leq^{(u)st} (G'/Z', \cdot, R') \quad (\text{Proposition 4.9})$$

$$\iff C \leq^{(u)st} C' \quad (\text{Proposition 4.11}),$$

and the result follows. □

Examples. Let C be any of the following graphs:

- $C = (V, E)$ with $V = \mathbb{Z}$ infinite and $E = \{(n, n + 1) \mid n \in \mathbb{Z}\}$:



- a regular tree of any degree,
- any nice graph bi-interpretable with $(\mathbb{Z}, <, +, 0)$, $(\mathbb{R}, <)$, $(\mathbb{R}, <, +, 0)$, $(\mathbb{R}, +, \cdot, 0, 1)$.

Then C , and therefore $M(C)$, is stably embedded in every elementary extension.

4.2. Characterisation of indiscernibles. We will characterise indiscernible sequences in a Mekler group M in terms of indiscernible sequences in the graph C . To simplify the presentation, we introduce some new terminology. We fix an indexing structure \mathcal{I} in a language \mathcal{L}' . By an “indiscernible sequence” in this subsection we always mean an \mathcal{I} -indiscernible sequence.

Definition 4.16. Let \mathcal{M} be a structure and $(\bar{a}_i)_{i \in \mathcal{I}}$ and $(\bar{b}_i)_{i \in \mathcal{I}}$ be sequences in \mathcal{M} . We say that the indiscernibility of $(\bar{a}_i)_{i \in \mathcal{I}}$ over $B \subseteq \mathcal{M}$ is *witnessed* by the sequence $(\bar{b}_i)_{i \in \mathcal{I}}$ if there exist parameters \bar{b} and a definable tuple of functions \bar{f} such that

- (1) $(\bar{b}_i)_{i \in \mathcal{I}}$ is indiscernible over $B \cup \bar{b}$,
- (2) $\bar{a}_i = \bar{f}(\bar{b}_i, \bar{b})$ for every $i \in \mathcal{I}$.

(As usual, when $B = \emptyset$, we do not mention it.)

Note that if $(\bar{b}_i)_{i \in \mathcal{I}}$ and $(\bar{a}_i)_{i \in \mathcal{I}}$ are as above, $(\bar{a}_i)_{i \in \mathcal{I}}$ is indiscernible over $B \cup \bar{b}$.

We work in M^{eq} . We shall show that, if \mathcal{I} is Ramsey, then the \mathcal{I} -indiscernibility of a sequence $(\bar{a}_i)_{i \in \mathcal{I}}$ in M^{eq} must be witnessed by an \mathcal{I} -indiscernible sequence $(\bar{b}_i)_{i \in \mathcal{I}}$ in C . In general, the length of the witnessing tuples \bar{b}_i will be (much) greater than that of \bar{a}_i . In any case, we will have $|\bar{b}_i| \leq N|\bar{a}_i|$ for some N depending on $\text{qftp}_{\mathcal{I}}(i)$ and of course on the sequence $(a_i)_i$.

As for relative quantifier elimination, we proceed in three steps. Before we get into that, we give the following useful lemma:

Lemma 4.17. *Let \mathcal{I} be a Ramsey indexing structure. Let D be a \emptyset -definable subset (in the pure language of groups) of a group G of exponent p . The indiscernibility of any sequence $(\bar{g}_i)_{i \in \mathcal{I}}$ in G is witnessed by an indiscernible sequence $(\bar{a}_i \bar{b}_i)_{i \in \mathcal{I}}$ where $\bigcup_{i \in \mathcal{I}} \bar{b}_i \subseteq D$ and $\bigcup_{i \in \mathcal{I}} \bar{a}_i$ is independent (in the sense of Definition 2.13) over $\langle D \rangle$.*

Proof. Since \mathcal{I} is Ramsey, by the generalised standard lemma (Theorem 2.17) it suffices to prove the proposition when \mathcal{I} is the *Fraïssé limit*⁷ of its *age*, that is, when $\mathcal{I} = \text{Flim}(\text{Age}(\mathcal{I}))$. To see this, suppose that the result holds for all $\text{Flim}(\text{Age}(\mathcal{I}))$ -indexed sequences. Then, given an \mathcal{I} -indexed sequence $(\bar{a}_i)_{i \in \mathcal{I}}$, by the generalised standard lemma, we can find a $\text{Flim}(\text{Age}(\mathcal{I}))$ -indexed indiscernible $(\bar{a}'_i)_{i \in \text{Flim}(\text{Age}(\mathcal{I}))}$ which is locally based on it and a $\text{Flim}(\text{Age}(\mathcal{I}))$ -indiscernible of the desired form $(\bar{b}'_i)_{i \in \text{Flim}(\text{Age}(\mathcal{I}))}$ witnessing the indiscernibility of $(\bar{a}'_i)_{i \in \text{Flim}(\text{Age}(\mathcal{I}))}$. Applying the generalised standard lemma again, this time to the sequence $(\bar{a}'_i \bar{b}'_i)_{i \in \text{Flim}(\text{Age}(\mathcal{I}))}$ we can find an \mathcal{I} -indiscernible $(\bar{a}''_i \bar{b}''_i)_{i \in \mathcal{I}}$ locally based on $(\bar{a}'_i \bar{b}'_i)_{i \in \text{Flim}(\text{Age}(\mathcal{I}))}$. It is clear that $(\bar{a}''_i)_{i \in \mathcal{I}}$ is locally based on $(\bar{a}_i)_{i \in \mathcal{I}}$ and that they are both \mathcal{I} -indiscernible, so we can find an automorphism σ sending \bar{a}''_i to \bar{a}_i . Then $(\sigma(\bar{b}''_i))_{i \in \mathcal{I}}$ remains of the desired form and witnesses the indiscernibility of $(\bar{a}_i)_{i \in \mathcal{I}}$. Thus, in

⁷See [Hod93, Theorem 7.1.2] for a statement of Fraïssé’s Theorem. We will denote the *Fraïssé limit* of an amalgamation class \mathcal{C} by $\text{Flim}(\mathcal{C})$.

the remainder of the proof, assume $\mathcal{I} = \text{Flim}(\text{Age}(\mathcal{I}))$. Observe also that since \mathcal{I} is Ramsey, by [MP23, Theorem B], it expands a linear order, which we will assume is already in \mathcal{L}' .

Let us first write $(\bar{g}_i)_{i \in \mathcal{I}}$ as a sequence $(\bar{a}_i \bar{b}_i)_{i \in \mathcal{I}}$ such that $\bar{a}_i \notin \langle D \rangle$ and $\bar{b}_i \in \langle D \rangle$, for all $i \in \mathcal{I}$.

Consider a group term $t(\bar{x}_0, \dots, \bar{x}_{k-1}, \bar{x})$ where one of the variables, say $x^0 \in \bar{x}$ really occurs.⁸ Assume that $t(\bar{a}_{i_0}, \dots, \bar{a}_{i_{n-1}}, \bar{a}_i) \in \langle D \rangle$ for some $i_0 < \dots < i_{n-1} < i$. For every j with $\text{qftp}(j) = \text{qftp}(i)$, we want to replace a_j^0 with a term in $\langle D \rangle$ such that the new sequence witnesses the indiscernibility of the original one.

For every such j , there are j_0, \dots, j_{n-1} such that the term $b_j'' = t(\bar{a}_{j_0}, \dots, \bar{a}_{j_{n-1}}, \bar{a}_j)$ is in $\langle D \rangle$. For j such that $\text{qftp}(j) \neq \text{qftp}(i)$, we set $b_j'' = \emptyset$, and we can consider the sequence $(\bar{a}_i \bar{b}_i b_i'')_{i \in \mathcal{I}}$. At this point, we need the following claim:

Claim 1. *There is an elementary extension $\mathcal{I}' \succ \mathcal{I}$ and elements $j_1 < \dots < j_n \in \mathcal{I}'$ such that for all $k \in \mathcal{I}$ for which there exist $k_1 < \dots < k_n < k$ with*

$$\text{qftp}^{\mathcal{I}}(i_1, \dots, i_n, i) = \text{qftp}^{\mathcal{I}}(k_1, \dots, k_n, k),$$

we have

$$\text{qftp}^{\mathcal{I}'}(i_1, \dots, i_n, i) = \text{qftp}^{\mathcal{I}'}(j_1, \dots, j_n, k).$$

Proof of claim. We use the fact that $\text{Age}(\mathcal{I})$ has the amalgamation property and compactness. Formally, fix an enumeration $(c_i : i \in \mathbb{N})$ of \mathcal{I} , and for every finite $\Delta(y_1, \dots, y_n, y) \subseteq_{\text{fin}} \text{qftp}(i_1, \dots, i_n, i)$, let $\phi_m^\Delta(x_1, \dots, x_n, c_1, \dots, c_m)$ be the formula

$$\bigwedge_{i \leq m} ((\exists z_1, \dots, z_n \Delta(z_1, \dots, z_n, c_i)) \rightarrow \Delta(x_1, \dots, x_n, c_i)).$$

Intuitively, ϕ_m^Δ says that x_1, \dots, x_n is a *common Δ -witness*, that is, it verifies the claim, for formulas in Δ and for the first m elements in the enumeration of \mathcal{I} , so it suffices to show that

$$\Sigma(x_1, \dots, x_n) := \left\{ \phi_m^\Delta(x_1, \dots, x_n, c_1, \dots, c_m) : m \in \mathbb{N}, \Delta \subseteq_{\text{fin}} \text{qftp}(i_1, \dots, i_n, i) \right\}$$

is satisfiable. Clearly, it is enough to show that a single ϕ_m^Δ is satisfiable, and this easily follows from the amalgamation and homogeneity of \mathcal{I} . Indeed, we can build a witness for ϕ_m^Δ inductively: once the witness for the first $l < m$ elements has been constructed, we use amalgamation to find a structure where the $(l + 1)$ -st element has a common witness with the first l elements (if one exists for this element). Then, since \mathcal{I} is homogeneous, we can embed this structure appropriately in \mathcal{I} . \square

⁸We mean that the term is not equal to a term $t'(\bar{x}_0, \dots, \bar{x}_{n-1}, \bar{x}')$ where \bar{x}' is the tuple \bar{x} from which we removed x^0 .

By Claim 1, we find an elementary extension $\mathcal{I}' \succ \mathcal{I}$ and $k_0 < \dots < k_{n-1} \in \mathcal{I}'$ such that for all $j \in \mathcal{I}$ with $\text{qftp}(j) = \text{qftp}(i)$, we have

$$\text{qftp}^{\mathcal{I}'}(k_0, \dots, k_{n-1}, j) = \text{qftp}^{\mathcal{I}}(i_0, \dots, i_{n-1}, i).$$

By the generalised standard lemma, there is an indiscernible sequence $(\tilde{a}_i \tilde{b}_i \tilde{b}_i'')_{i \in \mathcal{I}'}$ indexed by \mathcal{I}' that is locally based on $(\bar{a}_i \bar{b}_i b_i'')_{i \in \mathcal{I}}$. Notice that the sequence $(\tilde{a}_i \tilde{b}_i \tilde{b}_i'')_{i \in \mathcal{I}'}$ restricted to $\mathcal{I} \subseteq \mathcal{I}'$ is indiscernible over $\{\tilde{a}_{k_0}, \dots, \tilde{a}_{k_{n-1}}\}$.

We find an automorphism σ sending $\bar{a}_i \bar{b}_i$ to $\tilde{a}_i \tilde{b}_i$, for all $i \in \mathcal{I}$. Denote by $\bar{b} := \bar{b}_0, \dots, \bar{b}_{n-1}$ the tuple $\sigma^{-1}(\tilde{a}_{k_0}, \dots, \tilde{a}_{k_{n-1}})$.

For $j \in \mathcal{I}$ with $\text{qftp}(i) = \text{qftp}(j)$, denote $c_j = t(\bar{b}_0, \dots, \bar{b}_{n-1}, \bar{a}_j) \in D$ and by \bar{a}'_j the tuple \bar{a}_j from which we removed a_j^0 . For $j \in \mathcal{I}$ with $\text{qftp}(i) \neq \text{qftp}(j)$, we set $\bar{a}'_j = \bar{a}_j$ and $c_j = \emptyset$.

Since for all $i \in \mathcal{I}$, we have that $a_i^0 \in \text{dcl}(c_i, \bar{a}'_i, \bar{b})$, and it follows that the indiscernibility of $(\bar{a}_i \bar{b}_i)_{i \in \mathcal{I}}$ over \bar{b} is witnessed by that of $(\bar{a}'_i \bar{b}_i \frown c_i)_{i > k}$ over \bar{b} . We may repeat the process until, for all $\text{qftp}(j) = \text{qftp}(i)$, \bar{a}_j is independent over $\langle D \rangle$, and we reset the notation. Of course, this process must essentially be carried out for all quantifier-free types in \mathcal{I} , but the transformations for different quantifier-free types do not interact with each other, and hence we can perform them all simultaneously.

It remains to find \bar{b}'_i in D such that $(\bar{a}_i, \bar{b}'_i)_{i \in \mathcal{I}}$ is indiscernible and witnesses the indiscernibility of $(\bar{a}_i, \bar{b}_i)_{i \in \mathcal{I}}$. To simplify the notation, assume $\bar{b}_i = b_i$ is a singleton and $b_i = b_{0,i} \dots b_{n-1,i}$ where $b_{k,i} \in D$.

By the generalised standard lemma, we can now find an indiscernible sequence $(\tilde{a}_i, \tilde{b}_{0,i}, \dots, \tilde{b}_{n-1,i})_{i \in \mathcal{I}}$ that is locally based on $(\bar{a}_i, b_{0,i}, \dots, b_{n-1,i})_{i \in \mathcal{I}}$. For every $i \in \mathcal{I}$, set $\tilde{b}_i = \tilde{b}_{0,i} \dots \tilde{b}_{n-1,i}$.

Finally, let σ be an automorphism which for every i sends $(\tilde{a}_i, \tilde{b}_i)$ to (\bar{a}_i, b_i) . Then $\sigma(\tilde{b}'_{0,i}), \dots, \sigma(\tilde{b}'_{n-1,i})$ are as desired. \square

4.2.1. Reduction from G to $\mathcal{F}(G)$. Let G be a 2-nilpotent group of prime exponent p . We characterise the indiscernibility of sequences in G in terms of indiscernibility in $\mathcal{F}(G)$.

Proposition 4.18.

- (*\mathcal{I} Ramsey*) The indiscernibility of a sequence $(g_i)_{i \in \mathcal{I}}$ from G is witnessed by an indiscernible sequence $(\bar{a}_i \bar{b}_i)_{i \in \mathcal{I}}$ such that $\bar{b}_i \in Z$ and \bar{a}_i are independent over Z .
- (*Arbitrary \mathcal{I}*) Let $(\bar{a}_i \bar{b}_i)_{i \in \mathcal{I}}$ be a sequence in G such that $\bar{b}_i \in Z$ and \bar{a}_i are independent over Z . Then $(\bar{a}_i \bar{b}_i)_{i \in \mathcal{I}}$ is indiscernible if and only if $(\pi(\bar{a}_i), \rho(\bar{b}_i))_{i \in \mathcal{I}}$ is indiscernible in $\mathcal{F}(G)$.

Proof. The first point follows from Lemma 4.17, by setting $D = Z$, the centre of G .

It remains to show the second point. Notice that, since \bar{b}_i are central elements, for every \mathcal{L}_{grp} -term $t(\bar{x}, \bar{y})$, there is an \mathcal{L}_{grp} -term $t''(\bar{x})$ such that

$$\pi(t(\bar{a}_i, b_i)) = \pi(t''(\bar{a}_i)) = t''(\pi(\bar{a}_i)).^9$$

Also, since \bar{a}_i are independent over Z , for every \mathcal{L}_{grp} -term $t(\bar{x}, \bar{y})$ there is an \mathcal{L}_{grp} -term $t'(\bar{y})$ such that

$$\rho(t(\bar{a}_i, \bar{b}_i)) = \rho(t'(\bar{b}_i)) = t'(\rho(\bar{b}_i)).$$

In fact, by the property of independence, we can take $t'(\bar{b}_i) = t(\bar{1}, \bar{b}_i)$.

Then, by quantifier elimination, we know that any formula $\phi(\bar{x}, \bar{y})$ is equivalent to a formula of the form

$$\phi_{\mathcal{F}(G)}(\pi(t(\bar{x}, \bar{y})), \rho(t(\bar{x}, \bar{y}))),$$

where $\phi_{\mathcal{F}(G)}(\bar{y}, \bar{z})$ is a formula in the language $\{(V, +, 0), (W, +, 0), \beta\}$.

By the above, we find group terms $t'(\bar{y})$ and $t''(\bar{x})$ such that $\phi(\bar{a}_i, \bar{b}_i)$ is equivalent to

$$\phi_{\mathcal{F}(G)}(t'(\pi(\bar{a}_i)), t''(\rho(\bar{b}_i))).$$

The statement follows easily. \square

4.2.2. Reduction from $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta\}$ to $(\mathcal{V}, +, 0, R)$. Consider an alternating bilinear system $\{(\mathcal{V}, +, 0), (\mathcal{W}, +, 0), \beta\}$ satisfying property $(*_f)$ of Corollary 3.6. Let $(\bar{v}_i \bar{u}_i \bar{w}_i)_{i \in \mathcal{I}}$ be a sequence with $\bar{v}_i \in \mathcal{V}$, $\bar{u}_i \in \bigcup_n \mathcal{W}_n$, and $\bar{w}_i \notin \bigcup_n \mathcal{W}_n$.

Lemma 4.19. *Let \mathcal{I} be a Ramsey indexing structure. Let $(\bar{v}_i \bar{u}_i \bar{w}_i)_{i \in \mathcal{I}}$ be an indiscernible sequence with $\bar{v}_i \in \mathcal{V}$, $\bar{u}_i \in \bigcup_n \mathcal{W}_n$, and $\bar{w}_i \notin \bigcup_n \mathcal{W}_n$. Then, for every realised quantifier-free type q in \mathcal{I} there exist a positive integer n_q , coefficients $\bar{d}_q^{k,l}$, for $k < l < n_q$ and for every $i \in \mathcal{I}$ with $\text{qftp}(i) = q$, a sequence of tuples $(\bar{v}_q^{0,i}, \dots, \bar{v}_q^{n_q-1,i})$ such that*

- $(\bar{v}_i, \bar{v}_{\text{qftp}(i)}^{0,i}, \dots, \bar{v}_{\text{qftp}(i)}^{n_q-1,i}, \bar{w}_i)_{i \in \mathcal{I}}$ is indiscernible; and
- for all $i \in \mathcal{I}$ we have

$$\bar{u}_i = \sum_{k < l} \bar{d}_q^{k,l} \beta(\bar{v}_{\text{qftp}(i)}^{k,i}, \bar{v}_{\text{qftp}(i)}^{l,i}),^{10}$$

where $q = \text{qftp}(i)$.

Proof. To simplify the notation, assume that $\bar{u}_i = u_i$ is a singleton. For every $i \in \mathcal{I}$, we can rewrite u_i as a sum $\beta(v', v'')$ for v', v'' certain tuples $v_{0,i}, \dots, v_{n-1,i}$. By

⁹The term t'' should be understood multiplicatively in the group G and additively in the sort \mathcal{W} of $\mathcal{F}(G)$.

¹⁰The sum and β are componentwise.

indiscernibility, we can choose n and the coefficients depending only on $\text{qftp}(i)$, so, for all $i \in \mathcal{I}$ with $\text{qftp}(i) = q$ we have

$$u_i = \sum_{k < l} \bar{d}_q^{k,l} \beta(\bar{v}^{k,i}, \bar{v}^{l,i}),$$

for some $\bar{d}_q^{k,l} \in \mathbb{F}_p$. For every realised quantifier-free type q in \mathcal{I} and all $i \in \mathcal{I}$ realising q , we substitute u_i by the tuples $v^{0,1}, \dots, v^{n_q-1,i}$ we found above. The new sequence need not be indiscernible, but by the generalised standard lemma, we can now find an indiscernible sequence

$$(\tilde{v}_i, \tilde{v}^{0,i}, \dots, \tilde{v}^{n_{\text{qftp}(i)}-1,i}, \tilde{w}_i)_{i \in \mathcal{I}}$$

locally based on our new sequence. Set

$$\tilde{u}_i = \sum_{k < l} \bar{d}_q^{k,l} \beta(\tilde{v}^{k,i}, \tilde{v}^{l,i}).$$

Let σ be an automorphism which, for every i , sends $(\tilde{v}_i, \tilde{u}_i, \tilde{w}_i)$ to $(\bar{v}_i, \bar{u}_i, \bar{w}_i)$. Then, for each $i \in \mathcal{I}$, the tuple $\sigma(\tilde{v}^{0,i}), \dots, \sigma(\tilde{v}^{n_{\text{qftp}(i)}-1,i})$ is as desired. \square

Proposition 4.20.

- (\mathcal{I} Ramsey) Any indiscernible sequence $(\bar{u}_i)_{i \in \mathcal{I}}$ is witnessed by a sequence $(\bar{v}_i \bar{w}_i)_{i \in \mathcal{I}}$ such that $\bar{v}_i \in V$ and $\cup \bar{w}_i$ are in \mathcal{W} and linearly independent over $\bigcup_n \mathcal{W}_n$.
- (Arbitrary \mathcal{I}) Let $(\bar{v}_i \bar{w}_i)_{i \in \mathcal{I}}$ be a sequence in $(\mathcal{V}, \mathcal{W}, \beta)$ with $\bar{v}_i \in \mathcal{V}$ and $\bar{w}_i \in \mathcal{W}$. Assume that $\bigcup_i \bar{w}_i$ is a linearly independent subset of \mathcal{W} over $\bigcup_n \mathcal{W}_n$. Then $(\bar{v}_i \bar{w}_i)_{i \in \mathcal{I}}$ is indiscernible if and only if $(\bar{v}_i)_{i \in \mathcal{I}}$ is indiscernible in $(\mathcal{V}, +, 0, R)$.

Proof. By applying Lemma 4.17 with $D = \mathcal{W}_n$ for every n , we may assume that $\bigcup \bar{w}_i \setminus \bigcup_n \mathcal{W}_n$ are linearly independent over $\bigcup_n \mathcal{W}_n$. If $w \in \cup \bar{w}_i$ is in \mathcal{W}_n , we use Lemma 4.19 (for every quantifier-free type) to replace w with n (potentially) new elements $v \in \mathcal{V}$. This gives us the first point.

For the second, let $\phi(v, w)$ be a formula. By relative quantifier elimination, the formula is equivalent to a Boolean combination of formulas for the form

- $\sum a_i w_i = 0$, where a_i 's are in \mathbb{F}_p , and
- $\phi_{\mathcal{V}}(f_n(\sum a_{0,i} w_i), \dots, f_n(\sum a_{k-1,i} w_i), \bar{v})$, where $\phi_{\mathcal{V}}$ is a formula in the language of $(\mathcal{V}, +, 0, R)^{\text{eq}}$, and $a_{j,i}$'s are in \mathbb{F}_p .

By definition of the linear independence of \bar{w}_i 's, a formula of the form $\sum a_i w_i = 0$ never holds unless all coefficients a_i are 0. Also, since nontrivial sums $\sum a_i w_i$ are never in \mathcal{W}_n , we must have $f_n(\sum a_i w_i) = u$ (that is, the function f_n is not defined

on $\sum a_i w_i$). It follows that $\phi((\bar{v}_i)_i, (w_i)_i)$ is equivalent to a formula of the form $\phi'_{\mathcal{V}}((\bar{v}_i)_i)$, and the statement follows. \square

4.2.3. Reduction from $(\mathcal{V}, +, 0, R)$ to (C, E) . Let $M = (G, \cdot, 1)$ be a Mekler group, denote additively $(\mathcal{V}, +, 0)$ the abelian quotient G/Z , and denote by R (resp. E) the relation induced on \mathcal{V} (resp. on C) by the commutativity relation on M .

Proposition 4.21.

- (\mathcal{I} is Ramsey) Any indiscernible sequence $(\bar{u}_i)_{i \in \mathcal{I}}$ in \mathcal{V} is witnessed by a sequence $(\bar{v}_i \bar{v}'_i \bar{v}''_i)_{i \in \mathcal{I}}$ which is part of a transversal, i.e., such that
 - $\bigcup_i \bar{v}_i$ is composed of elements of type 1^v ,
 - $\bigcup_i \bar{v}'_i$ is composed of elements of type p and linearly independent over elements of type 1^v ,
 - $\bigcup_i \bar{v}''_i$ is composed of elements of type 1^t and linearly independent over elements of type 1^v and p .
- (Arbitrary \mathcal{I}) Consider a sequence $(\bar{v}_i, \bar{v}'_i, \bar{v}''_i)_{i \in \mathcal{I}}$ such that
 - $\bigcup_i \bar{v}_i$ is composed of elements of type 1^v ,
 - $\bigcup_i \bar{v}'_i$ is composed of elements of type p and linearly independent over elements of type 1^v ,
 - $\bigcup_i \bar{v}''_i$ is composed of elements of type 1^t and linearly independent over elements of type 1^v and p .

Then $(\bar{v}_i, \bar{v}'_i, \bar{v}''_i)_{i \in \mathcal{I}}$ is indiscernible if and only if, the sequence

$$([\bar{v}_i]_{\sim}, h(\bar{v}'_i))_i,^{11}$$

is indiscernible in (C, E) .

Proof. We deduce the first point directly from Lemma 4.17, and the second follows immediately from relative quantifier elimination. Notice that we don't need to consider terms of the form $S_{n,m}(v_{i_1} \cdots v_{i_n} v'_{j_1} \cdots v'_{j_m}) = \{[v_{i_1}]_{\sim}, \dots, [v_{i_n}]_{\sim}\}$ for $n + m > 1$ since they can be recovered from $S_{1,0}(v_{i_j})$ for $1 \leq j \leq n$. \square

4.2.4. Reduction from $(G, \cdot, 1)$ to (C, E) . Let $\mathcal{M} = (G, \cdot, 1)$ be a monster model of a Mekler group, and let $C = (V, E)$ be the corresponding nice graph.

Lemma 4.22. Let $(\alpha_i \beta_i)_{i \in \mathcal{I}}$ be an indiscernible sequence where for all i 's, $\beta_i \in Z$ is a finite product of commutator. Then, for every realised quantifier-free type q in \mathcal{I} there exist a positive integer n_q , coefficients $\bar{d}_q^{k,l}$, for $k < l < n_q$ and for every $i \in \mathcal{I}$ with $\text{qftp}(i) = q$, a sequence of tuples $(\beta_q^{0,i}, \dots, \beta_q^{n-1,i})$ such that

- the sequence $(\alpha_i, \beta_{\text{qftp}(i)}^{0,i}, \dots, \beta_{\text{qftp}(i)}^{n-1,i})_{i \in \mathcal{I}}$ is indiscernible; and

¹¹This is the sequence of tuples consisting of the class modulo \sim of each component of \bar{v}_i and of the handles of each component of \bar{v}'_i .

- for all $i \in \mathcal{I}$, we have

$$\beta_i = \prod_{k < l} [\beta_{\text{qftp}(i)}^{k,i}, \beta_{\text{qftp}(i)}^{l,i}]^{d_q^{k,l}},^{12}$$

where $q = \text{qftp}(i)$.

Proof. The proof is similar to that of Lemma 4.19. □

Proposition 4.23.

- (\mathcal{I} Ramsey) Any indiscernible sequence $(\bar{g}_i)_{i \in \mathcal{I}}$ in G is witnessed by an indiscernible sequence $(\alpha_i, \beta_i, \gamma_i, \delta_i)_{i \in \mathcal{I}}$ in G which is part of a full transversal, i.e., such that
 - $\bigcup_i (\alpha_i \cup \beta_i \cup \gamma_i)$ is part of a transversal,
 - $\bigcup_i \alpha_i$ is a subset of \mathfrak{E}^v (and is independent over Z),
 - $\bigcup_i \beta_i$ is a subset of \mathfrak{E}^p (and is independent over $\langle Z, \mathfrak{E}^v \rangle$),
 - $\bigcup_i \gamma_i$ is a subset of \mathfrak{E}^t (and is independent over $\langle Z, \mathfrak{E}^v, \mathfrak{E}^p \rangle$),
 - $\bigcup_i \delta_i$ is a subset of Z and is independent over $\langle \mathfrak{E}^v, \mathfrak{E}^p, \mathfrak{E}^t \rangle$.
- (Arbitrary \mathcal{I}) Such a sequence $(\alpha_i, \beta_i, \gamma_i, \delta_i)_{i \in \mathcal{I}}$ which is part of a full transversal — with the same notation as above — is indiscernible if and only if the sequence

$$([\alpha_i]_{\sim}, h(\beta_i))_{i \in \mathcal{I}}$$

is indiscernible in C .

Proof. One can easily see why a sequence of indiscernibles must be witnessed by an indiscernible sequence which is part of a transversal using Lemmas 4.17 and 4.22. The second point follows immediately from Propositions 4.18, 4.20 and 4.21. □

4.3. Transfers of dividing lines.

4.3.1. Collapsing and NC_K . Having established our characterisation of indiscernibles in Mekler groups, we can now deduce (almost immediately) our main theorem:

Theorem A (collapsing transfer). *Let \mathcal{I} be Ramsey and \mathcal{J} a (not necessarily Ramsey) reduct of \mathcal{I} . Let M be a Mekler group of the nice graph C . Then M collapses \mathcal{I} -indiscernibles (resp. collapses \mathcal{I} -indiscernibles to \mathcal{J} -indiscernibles) if and only if C collapses \mathcal{I} -indiscernibles (resp. collapses \mathcal{I} -indiscernibles to \mathcal{J} -indiscernibles).*

Proof. Assume that C collapses \mathcal{I} to \mathcal{J} indiscernibles. Consider an \mathcal{I} -indiscernible sequence $(g_i)_{i \in \mathcal{I}}$ in M . By Proposition 4.23, the \mathcal{I} -indiscernibility of $(g_i)_{i \in \mathcal{I}}$ is witnessed by an indiscernible sequence $(\alpha_i, \beta_i, \gamma_i, \delta_i)_{i \in \mathcal{I}}$, with α_i of type 1^v and β of type p and such that $\bigcup_i (\alpha_i \cup \beta_i \cup \gamma_i)$ is part of a full transversal. In particular,

¹²If β_i is a tuple, all operations are done componentwise.

$([\alpha_i]_{\sim}, h(\beta_i))_{i \in \mathcal{I}}$ is \mathcal{I} -indiscernible, and it is therefore \mathcal{J} -indiscernible, since \mathbf{C} collapses \mathcal{I} -indiscernibles to \mathcal{J} -indiscernibles.

By the equivalence, $(\alpha_i, \beta_i, \gamma_i)_{i \in \text{dom}(\mathcal{I})}$ is \mathcal{J} -indiscernible. As it witnesses the indiscernibility of $(g_i)_{i \in \mathcal{I}}$, the latter is also \mathcal{J} -indiscernible. Notice that no assumption has been made on \mathcal{J} ; therefore, the proof can be adapted to (nonspecific) collapsing. \square

Corollary 4.24 (NC_K transfer). *Let \mathcal{K} be a Ramsey class in a countable language with an \aleph_0 -categorical Fraïssé limit. A Mekler group \mathbf{M} is NC_K if and only if its associated graph \mathbf{C} is NC_K .*

Proof. By Theorem 2.23, we have that for any complete theory T , T is NC_K if and only if T has no uncollapsed $\text{Flim}(\mathcal{K})$ -indiscernible. Therefore, this is only a reformulation of the previous theorem. \square

We can now, for instance, recover the key Theorem 4.7 of [CH19]. For that, we recall a characterisation of NIP_k , for an integer $k > 1$. Let \mathcal{C}_{H_k} be the class of all finite ordered k -partite hypergraphs. This is, of course, a Ramsey class by a classical result of Nešetřil and Rödl [NR77].

Fact 4.25 [CPT19, Proposition 5.4], [GH19, Theorem 3.14]. *The following are equivalent for a first-order theory T and $k \in \mathbb{N}$:*

- (1) T is NIP_k .
- (2) T is $\text{NC}_{\mathcal{C}_{H_k}}$.

Corollary 4.26 [CH19, Corollary 4.8]. *For every $k \geq 1$, a Mekler group \mathbf{M} is NIP_k if and only if its associated graph \mathbf{C} is NIP_k .*

We can also immediately deduce the following, from Fact 2.28:

Corollary 4.27. *For every $k \geq 1$ a Mekler group \mathbf{M} is NFOP_k if and only if its associated graph \mathbf{C} is NFOP_k .*

In particular, for all $k \geq 2$ there is a strictly NFOP_k pure group.

Proof. The first part of the corollary is immediate from Fact 2.28 and Theorem A.

By Fact 2.3, any structure in a finite relational language is bi-interpretable with a nice graph. From [AACT25, Proposition 3.25] we know that any k -ary structure (i.e., any structure which admits quantifier elimination in a relational language where all symbols have arity at most k) is NFOP_k , so in particular, the random k -hypergraph, $k \geq 2$, is NFOP_k . We know from [AACT25, Proposition 2.8] that

$$\text{NFOP}_k \implies \text{NIP}_k,$$

and since the random k -hypergraph has IP_{k-1} it has FOP_{k-1} , is thus strictly NFOP_k . Let \mathbf{C} be any nice graph bi-interpretable with the random k -hypergraph. Then $\mathbf{M}(\mathbf{C})$ is a strictly NFOP_k pure group, by the first part of the corollary. \square

4.4. Resplendence. The quantifier elimination results and the transfer principles obtained from them are all resplendent, in the sense of [Rid17, Appendix A]. This means that Theorems A and B and Proposition 4.23 and 4.14 remain valid if we replace, *mutatis mutandis*, the Mekler groups M with a C -enrichment

$$M^* = \{M = (G, \cdot, 1), C = (V, R, \dots), \pi : E^v \rightarrow C\},$$

where \dots denotes additional structure on C .

To finish off this section, we give some negative results. These are derived from basic observations and do not require the relative quantifier elimination.

4.4.1. Distality. In this paper, a precise definition of distality is not necessary. The reader can find one in [Sim15, Chapter 9]. For our purposes, it suffices to note that a distal theory does not admit any nonconstant totally indiscernible sequences.

Proposition 4.28. *Let M be a Mekler group. Then $\text{Th}(M)$ is nondistal.*

Proof. We show that M admits a nonconstant totally indiscernible sequence. Consider an infinite sequence $(\gamma_i)_{i \in \mathbb{N}}$ of elements of type 1^t which are independent over $\langle Z, \mathcal{E}^v, \mathcal{E}^p \rangle$. By Proposition 4.23, this sequence is totally indiscernible. It only remains to show that such a sequence exists, which we briefly justify for the sake of completeness. Since C is an infinite graph, it has, by Ramsey, an infinite clique (i.e., a subset where every two elements are connected one to the other), or an infinite anticlique (i.e., an infinite subset where no two elements are connected one to the other). Since C is nice, it cannot have an infinite clique and, therefore, it must have an infinite anticlique A . By compactness and considering larger and larger products of elements in A , we can find elements of type 1^t which are independent over $\langle Z, \mathcal{E}^v, \mathcal{E}^p \rangle$. \square

A more relevant question is therefore the following:

Question 4.29. Assume that a nice graph C admits a distal expansion. Does $M(C)$ admit a distal expansion?

It seems that if C admits a linear order $<$ such that the structure $(C, <, E)$ is distal, then one can define a natural valuation on G/Z and Z , and this additional structure would eliminate all traces of stability that Mekler's construction can bring "on top" of the structure C .

4.4.2. Burden. As we have seen, Mekler's construction preserves many dividing lines. However, it seems that it does not (always) preserve notions of dimension. We will treat here the *burden*, i.e., the notion of dimension attached to NTP_2 theories (which coincide with the dp-rank if the theory is NIP).

Proposition 4.30. *If C is an infinite nice graph, $M(C)$ has burden at least \aleph_{0-} .*

We refer the reader to [Adl] for the precise definitions needed for the remainder of this section.

This proposition provides a negative answer to [CH19, Problem 5.8]. One should expect similar results for any other reasonable notion of dimension. The proof below describes how one can obtain an inp-pattern of arbitrary finite size.

Proof. Consider the pattern consisting of an array $(a_{i,j})_{i,j \in \omega}$ of distinct elements of C and the formula $\phi(x, y_C)$:

$$A_{m,0}(\pi(x)) \wedge y_\gamma \in S_{m,0}(x)$$

where x is a group sorted variable, y_γ is a C -sorted variable, π is the projection modulo Z , $A_{m,0}$ is the predicate for elements x which are sum of m elements of type 1^ν , and $S_{m,0}(x)$ is the support of such element x .

Then notice that $\{\phi(x, a_{i,j}) \mid i \in \omega\}$ is $(m + 1)$ -inconsistent (but m -consistent!) and $\{\phi(x, a_{i,f(i)}) \mid i \in \omega\}$ is consistent for all $f : m \rightarrow \omega$.

We thus have an inp-pattern of depth m for every integer m , therefore $M(C)$ has burden at least \aleph_{0-} . □

In [CH19, Remark 5.7], Chernikov and Hempel argue that a Mekler group $M(C)$ is *strong* if and only if C is strong. More generally, one can certainly ask what is the exact burden of $M(C)$. We conjecture the following formula:

Conjecture 4.31. We have $\text{bdn}(M(C)) = \min\{\aleph_{0-}, \text{bdn}(C)\}$.

5. Second proof of Theorem A

We now present a second proof of Theorem A. We follow a proof strategy similar to that of Hempel and Chernikov in [CH19]. It avoids the formal description of definable sets from Section 3, at the cost of introducing a (reasonable) additional assumption on the collapsing property of the nice graph. Let us recall the following key lemma from [CH19]:

Lemma 5.1 [CH19, Lemma 2.14]. *Let C be a nice graph and \mathcal{M} a saturated model of $\text{Th}(M(C))$. Let X be a transversal and $K_X \leq Z(\mathcal{M})$ be such that $\mathcal{M} = \langle X \rangle \times \langle K_X \rangle$. Let W and Y be two small subsets of X and \bar{h}_1, \bar{h}_2 two tuples in K_X . Suppose that the following conditions hold:*

- *There is a bijection $f : W \rightarrow Y$ which respects the 1^ν -, p -, 1^l -parts, the handles, and $\text{tp}_\Gamma(W^\nu) = \text{tp}_\Gamma(f(Y)^\nu)$.*
- $\text{tp}_{K_X}(\bar{h}_1) = \text{tp}_{K_X}(\bar{h}_2)$.

Then, there is an automorphism $\sigma \in \text{Aut}(\mathcal{M})$ extending f such that $\sigma(\bar{h}_1) = \bar{h}_2$.

In [CH19], amongst other results, the authors essentially show the following [CH19, Theorem 4.7]: If, for some nice graph C , every model of $\text{Th}(C)$ collapses

H_k -indiscernibles (where H_k is the random k -partite hypergraph) to P_k -indiscernibles (where P_k is the random k -partite set), then so must every model of $\text{Th}(\mathcal{M}(\mathcal{C}))$. Below, we adapt their argument for the general case of collapsing \mathcal{I} -indiscernibles (in place of H_k -indiscernibles) to \mathcal{J} -indiscernibles (in place of P_k -indiscernibles).

Definition 5.2 (specific collapsing). Let \mathcal{I} be an \aleph_0 -categorical Fraïssé limit of a Ramsey class and \mathcal{J} a proper reduct of \mathcal{I} . We say that \mathcal{I} has *specific collapsing to \mathcal{J}* if for all theories T , and all $\mathcal{M} \models T$, every collapsing \mathcal{I} -indiscernible sequence in \mathcal{M} , is (a possibly collapsing) \mathcal{J} -indiscernible sequence.

Many dividing lines are characterised by specific collapsing to a reduct. Classical examples include stability (where we have specific collapsing of order-indiscernibles to indiscernible sets), NIP_k (here we have specific collapsing of random ordered k -hypergraph indiscernibles to linear orders), and recently NFOP_k (this is more complicated to describe but follows from results in [AACT25], in particular Lemma 4.14 therein) for all $k \in \mathbb{N}$.

We start with an easy lemma:

Lemma 5.3. *Let T be a stable theory, \mathcal{I} an \aleph_0 -categorical Fraïssé limit of a Ramsey class in a finite relational language, and \mathcal{J} a proper reduct of \mathcal{I} . If \mathcal{I} has specific collapsing to \mathcal{J} then every model of T collapses \mathcal{I} -indiscernibles to \mathcal{J} -indiscernibles.*

Proof. Since \mathcal{I} is Ramsey, it expands a linear order (see [Bod15, Corollary 2.26]), and since T is stable, we must have that \mathcal{I} -indiscernibles collapse to the reduct of \mathcal{I} where we forget the order in models of T . By specific collapsing, we then have that \mathcal{I} -indiscernibles collapse to \mathcal{J} -indiscernibles in models of T . \square

Theorem A' (second statement). *Let \mathcal{I} be an \aleph_0 -categorical Fraïssé limit of a Ramsey class in a finite relational language, and \mathcal{J} a proper reduct of \mathcal{I} , also in a finite relational language. For every nice graph \mathcal{C} , if*

- \mathcal{I} has specific collapsing to \mathcal{J} , and
- $\text{Th}(\mathcal{C})$ collapses \mathcal{I} -indiscernibles to \mathcal{J} -indiscernibles,

then

- $\text{Th}(\mathcal{M}(\mathcal{C}))$ collapses \mathcal{I} -indiscernibles to \mathcal{J} -indiscernibles.

Proof. Assume toward a contradiction that $\text{Th}(\mathcal{C})$ collapses \mathcal{I} -indiscernibles to \mathcal{J} -indiscernibles but $\text{Th}(\mathcal{M}(\mathcal{C}))$ does not. Since both \mathcal{I} and \mathcal{J} are finitely homogeneous, they are k -ary, for some $k \in \mathbb{N}$. By appropriately padding the symbols in $\mathcal{L}_{\mathcal{I}}$ and $\mathcal{L}_{\mathcal{J}}$, the respective languages of \mathcal{I} and \mathcal{J} , we may assume that they all have arity exactly k . Moreover, since \mathcal{I} has quantifier elimination we can also assume that $\mathcal{L}_{\mathcal{J}} \subseteq \mathcal{L}_{\mathcal{I}}$.

By assumption, we can find a model $\mathcal{M} \models \text{Th}(\mathcal{M}(\mathcal{C}))$ and an \mathcal{I} -indiscernible sequence $B = (b_i : i \in \text{dom}(\mathcal{I}))$ in \mathcal{M} which is not \mathcal{J} -indiscernible. Since \mathcal{I} has specific collapsing to \mathcal{J} , this must mean that B is noncollapsing. In particular, for all $(i_1, \dots, i_k), (j_1, \dots, j_k) \in \text{dom}(\mathcal{I})^k$ we have that

$$\text{qftp}_{\mathcal{I}}(i_1, \dots, i_k) = \text{qftp}_{\mathcal{I}}(j_1, \dots, j_k) \iff \text{tp}(b_{i_1}, \dots, b_{i_k}) = \text{tp}(b_{j_1}, \dots, b_{j_k}).$$

Claim 1. *There is some $R \in \mathcal{L}_{\mathcal{I}} \setminus \mathcal{L}_{\mathcal{J}}$ for which we can find an \mathcal{L}_{grp} -formula ϕ_R such that for and all $i_1, \dots, i_k \in \mathcal{I}$ we have*

$$\mathcal{M} \models \phi_R(b_{i_1}, \dots, b_{i_k}) \text{ if and only if } \mathcal{I} \models R(i_1, \dots, i_k).$$

Proof of claim. First, since \mathcal{I} is \aleph_0 -categorical there are finitely many quantifier-free k -types in $\text{Th}(\mathcal{I})$, say p_1, \dots, p_n . Every relation symbol $R \in \mathcal{L}_{\mathcal{I}}$ belongs to a finite Boolean combination of these types, and we claim that for some $R \in \mathcal{L}_{\mathcal{I}} \setminus \mathcal{L}_{\mathcal{J}}$ there is a Boolean combination of (isolating formulas of) quantifier-free types ψ such that

$$\mathcal{I} \models \psi(i_1, \dots, i_k) \text{ if and only if } \mathcal{I} \models R(i_1, \dots, i_k).$$

If not, then we would be able to find a reduct \mathcal{J}' of \mathcal{I} which would make B a \mathcal{J}' -indiscernible sequence, contradicting our initial assumption.

Now, since B is an uncollapsed \mathcal{I} -indiscernible, for each quantifier-free k -type p_i of \mathcal{I} there is a unique k -type q_i of \mathcal{M} in B , corresponding to the k -tuples indexed by p_i . For each such type we can find an \mathcal{L}_{grp} -formula ϕ_i separating it from the rest.

Recall that since \mathcal{I} is \aleph_0 -categorical for each quantifier-free n -type $p \in S_n^{\mathcal{I}}(\emptyset)$, there is a formula $\text{iso}(p) \in p$ isolating p . Moreover, for all $n \in \mathbb{N}$, by Ryll-Nardzewski, $|S_n^{\mathcal{I}}(\emptyset)| < \aleph_0$.

Given the above, for a relation symbol $R \in \mathcal{L}_{\mathcal{I}}$, there is a unique Boolean combination of isolating formulas

$$\psi := \bigwedge \bigvee \text{iso}(p_i)^{\epsilon_i},$$

equivalent to R .

For ψ as above, take ϕ_R to be the formula $\bigwedge \bigvee \phi_i^{\epsilon_i}$, i.e., the corresponding Boolean combination of the formulas separating the types q_i in \mathcal{M} . \square

Fix κ to be \aleph_0^+ and $\mathbb{I} \succ \mathcal{I}$ be an elementary extension of \mathcal{I} with $|\mathbb{I}| = \kappa$. Without loss of generality, we may assume that \mathcal{M} is saturated. Since \mathcal{I} is Ramsey, \mathcal{I} -indiscernibles have the modelling property, so by compactness/saturation we can actually find, in \mathcal{M} , an \mathbb{I} -indexed indiscernible sequence $A = (a_i : i \in \text{dom}(\mathbb{I}))$ which is based on B . Let $\mathbb{J} \succ \mathcal{J}$ be the elementary extension of \mathcal{J} given by taking the $\mathcal{L}_{\mathcal{J}}$ -reduct of \mathbb{I} .

Claim 2. *The sequence $(a_i : i \in \text{dom}(\mathbb{I}))$ is not \mathbb{J} -indiscernible.*

Proof of claim. Since $(a_i : i \in \text{dom}(\mathcal{I}))$ is not \mathcal{J} -indiscernible, we can find $i_1, \dots, i_k, j_1, \dots, j_k \in \text{dom}(\mathcal{I})$ such that $\text{qftp}_{\mathcal{J}}(i_1, \dots, i_k) = \text{qftp}_{\mathcal{J}}(j_1, \dots, j_k)$ and $\text{tp}(b_{i_1}, \dots, b_{i_k}) \neq \text{tp}(b_{j_1}, \dots, b_{j_k})$. Let $\Delta_1 \subseteq \text{tp}(b_{i_1}, \dots, b_{i_k})$, and $\Delta_2 \subseteq \text{tp}(b_{j_1}, \dots, b_{j_k})$ be two finite sets of \mathcal{L}_{grp} -formulas such that $\Delta_1 \cup \Delta_2$ is inconsistent.

Since $(a_i : i \in \text{dom}(\mathbb{I}))$ is based on $(b_i : i \in \text{dom}(\mathcal{I}))$ and $\mathbb{I} \succ \mathcal{I}$, given $l_1, \dots, l_k \in \mathbb{I}$ such that $\text{qftp}_{\mathbb{I}}(l_1, \dots, l_k) = \text{qftp}_{\mathcal{I}}(i_1, \dots, i_k)$, we can find a tuple $i'_1, \dots, i'_k \in \mathcal{I}$ such that $\text{qftp}_{\mathbb{I}}(l_1, \dots, l_k) = \text{qftp}_{\mathcal{I}}(i'_1, \dots, i'_k)$ and thus

$$\text{tp}^{\Delta_1}(a_{l_1}, \dots, a_{l_k}) = \text{tp}^{\Delta_1}(b_{i'_1}, \dots, b_{i'_k}) = \text{tp}^{\Delta_1}(b_{i_1}, \dots, b_{i_k}).$$

Similarly, we find m_1, \dots, m_k such that $\text{qftp}_{\mathbb{I}}(m_1, \dots, m_k) = \text{qftp}_{\mathcal{I}}(j_1, \dots, j_k)$ and $\text{tp}^{\Delta_2}(a_{m_1}, \dots, a_{m_k}) = \text{tp}^{\Delta_2}(b_{j_1}, \dots, b_{j_k})$. But, by construction,

$$\text{qftp}_{\mathbb{I}}(l_1, \dots, l_k) = \text{qftp}_{\mathbb{I}}(m_1, \dots, m_k),$$

and thus $(a_i : i \in \text{dom}(\mathbb{I}))$ is not \mathbb{J} -indiscernible. \square

Let us write \mathcal{M} in the form $\langle X \rangle \times \langle H \rangle$, where X is a transversal and $H \subseteq \mathcal{M}$ a set that is linearly independent over $[\mathcal{M}, \mathcal{M}]$.

After fixing an enumeration of $\text{dom}(\mathbb{I})$, and rearranging A to be in the form $(a_i : i < \kappa)$, we express the elements in A as \mathcal{L}_{grp} -terms built up from elements of X and H , as follows: For each $\lambda < \kappa$ let t_λ be an \mathcal{L}_{grp} -term, and $\bar{x}_\lambda, \bar{h}_\lambda$ be finite tuples from X and H , respectively, such that $a_\lambda = t_\lambda(\bar{x}_\lambda, \bar{h}_\lambda)$.

By passing to a cofinal subsequence of A of cardinality κ (in the fixed enumeration), we can find a term $t \in \mathcal{L}_{\text{grps}}$ such that for all $\lambda < \kappa$ we have $t_\lambda = t$. Since each \bar{x}_λ is a tuple from X , we can assume that it is of the form $\bar{x}_\lambda^v \frown \bar{x}_\lambda^p \frown \bar{x}_\lambda^t$, where we simply list all elements of \bar{x}_λ of the corresponding types. We may also append the handles of the elements in the tuple \bar{x}_λ^p to the beginning of \bar{x}_λ^v (so that the handle of the j -th element of \bar{x}_λ^p is the j -th element of \bar{x}_λ^v , and we allow repetition of elements).

Thus, at this point, after rearranging, we have an \mathbb{I} -indiscernible sequence $(t(\bar{x}_i, \bar{h}_i) : i \in \text{dom}(\mathbb{I}))$. Of course, since $\mathcal{I} \preceq \mathbb{I}$, we may actually work with the subsequence $(t(\bar{x}_i, \bar{h}_i) : i \in \text{dom}(\mathcal{I}))$. By construction, this sequence is \mathcal{I} -indiscernible, and arguing as in Claim 2, $(\bar{x}_i \frown \bar{h}_i : i \in \text{dom}(\mathcal{I}))$ is not \mathcal{J} -indiscernible.

By Claim 1 and the fact that A is based on B , we can find some $R \in \Sigma = \mathcal{L}_{\mathcal{I}} \setminus \mathcal{L}_{\mathcal{J}}$ and some \mathcal{L}_{grp} -formula ϕ_R such that

$$\mathcal{M} \models \phi_R(a_{i_1}, \dots, a_{i_k}) \text{ if and only if } \mathcal{I} \models R(i_1, \dots, i_k).$$

Now, let $\Gamma(\mathcal{M})$ be a saturated model of $\text{Th}(\mathcal{C})$, containing all the elements $(\bar{x}_i^v : i \in \text{dom}(\mathcal{I}))$. Since $\Gamma(\mathcal{M})$ is interpretable in \mathcal{M} , this sequence remains \mathcal{I} -indiscernible, and since $\text{Th}(\mathcal{C})$ collapses \mathcal{I} -indiscernibles to \mathcal{J} -indiscernibles this is actually a \mathcal{J} -indiscernible in Γ . Similarly, let \mathcal{H} be a saturated model of $\text{Th}(\langle H \rangle)$, it is easy

to see that $(\bar{h}_i : i \in \text{dom}(\mathcal{I}))$ is \mathcal{I} -indiscernible in \mathcal{H} , and since $\text{Th}(\langle H \rangle)$ is stable, from Lemma 5.3, we have that $(\bar{h}_i : i \in \text{dom}(\mathcal{I}))$ is \mathcal{J} -indiscernible in \mathcal{H} .

Since the sequence $(\bar{x}_i \frown \bar{h}_i : i \in \text{dom}(\mathcal{I}))$ is not \mathcal{J} -indiscernible, there are $i_1, \dots, i_k, j_1, \dots, j_k \in \text{dom}(J)$ such that $\text{qftp}_{\mathcal{J}}(i_1, \dots, i_k) = \text{qftp}_{\mathcal{J}}(j_1, \dots, j_k)$ and $\text{tp}((\bar{x}_{i_1} \frown \bar{h}_{i_1}), \dots, (\bar{x}_{i_k} \frown \bar{h}_{i_k})) \neq \text{tp}((\bar{x}_{j_1} \frown \bar{h}_{j_1}), \dots, (\bar{x}_{j_k} \frown \bar{h}_{j_k}))$. Of course, by \mathcal{I} -indiscernibility of the sequence we must have that $\text{qftp}_{\mathcal{I}}(i_1, \dots, i_k) \neq \text{qftp}_{\mathcal{I}}(j_1, \dots, j_k)$, so there exists some relation symbol $R \in \mathcal{L}_{\mathcal{I}}$ such that

$$\mathcal{I} \models R(i_1, \dots, i_k) \wedge \neg R(j_1, \dots, j_k).$$

In particular, we can find an \mathcal{L}_{grp} -formula ϕ_R such that

$$\mathcal{M} \models \phi_R((\bar{x}_{i_1} \frown \bar{h}_{i_1}), \dots, (\bar{x}_{i_k} \frown \bar{h}_{i_k})) \wedge \neg \phi_R((\bar{x}_{j_1} \frown \bar{h}_{j_1}), \dots, (\bar{x}_{j_k} \frown \bar{h}_{j_k}))$$

At this point we are in the following situation:

- Since $\text{qftp}_{\mathcal{J}}(i_1, \dots, i_k) = \text{qftp}_{\mathcal{J}}(j_1, \dots, j_k)$ and the sequences $(\bar{x}_i^v : i \in \text{dom}(\mathcal{I}))$ and $(\bar{h}_i : i \in \text{dom}(\mathcal{I}))$ are \mathcal{J} -indiscernible we have

$$\text{tp}_{\Gamma}(\bar{x}_{i_1}^v, \dots, \bar{x}_{i_k}^v) = \text{tp}_{\Gamma}(\bar{x}_{j_1}^v, \dots, \bar{x}_{j_k}^v) \quad \text{and} \quad \text{tp}_{\langle H \rangle}(\bar{h}_{i_1}, \dots, \bar{h}_{i_k}) = \text{tp}_{\langle H \rangle}(\bar{h}_{j_1}, \dots, \bar{h}_{j_k}).$$

- The map sending $\bar{x}_{i_l} \frown \bar{h}_{i_l} \mapsto \bar{x}_{j_l} \frown \bar{h}_{j_l}$ respects the 1^v -, p -, 1^l -parts, and the handles.

Thus, by applying Lemma 5.1, we can extend this map to an automorphism $\sigma \in \text{Aut}(\mathcal{M})$ sending $(\bar{x}_{i_l} \frown \bar{h}_{i_l})$ to $(\bar{x}_{j_l} \frown \bar{h}_{j_l})$ for all $l \in \{1, \dots, k\}$, which is a contradiction. \square

6. Open questions

We suggest some possible transfers for Mekler groups that we try to briefly motivate.

Beautiful pairs. Extending Theorem B, one could try to characterise the existence of beautiful pairs of Mekler groups, in the sense of [CKHY23]. In this work, Cubides Kovacsics, Hils and Ye develop a general theory of beautiful pairs, extending the ideas of Poizat. They show, in particular, that the existence of a λ -saturated λ -beautiful pair for $\lambda > |T|^+$ implies strict pro-definability of the space of definable types. Hence saturated beautiful pairs of Mekler groups would give examples of 2-nilpotent groups with a strict pro-definable space of definable types.

Imaginariness. To complement our result of relative quantifier elimination, one could try to characterise imaginary elements of Mekler groups (up to finite imaginaries). Transfer principles for imaginaries are known in the context of valued fields (e.g. [RKV23]). A similar result for Mekler groups could certainly be interesting; we would then have a complete understanding of the first-order structure of Mekler's construction (definable sets, and quotients by definable equivalence relations).

Domination. In a given structure, if the tensor product respects the domination relation between invariant types, the set of invariant types modulo the domination equivalence is a monoid (the *domination monoid*), when equipped with the multiplicative law induced by the tensor product. In [HM24], Mennuni and Hils prove various transfer principles $A \rightarrow B$ for domination: they compute the domination monoid of A in terms of the domination monoid of B . A similar computation $M \rightarrow C$ for Mekler groups would give many examples of 2-nilpotent groups with a well-defined domination monoid.

Acknowledgements

We thank Paolo Marimon for reminding us of the results of Abd-Aldaim, Conant, and Terry. We also thank Raf Cluckers and Dugald Macpherson for their constant and precious support. Finally, we would like to thank the two referees for their very detailed and insightful comments, which greatly increased the quality of this paper.

References

- [AACT25] A. Abd Aldaim, G. Conant, and C. Terry, “Higher arity stability and the functional order property”, *Selecta Math. (N.S.)* **31**:3 (2025), art. id. 59, 79 pp. MR
- [Adl] H. Adler, “Strong theories, burden, and weight”, preprint, 2007, available at <https://tinyurl.com/Adler-StrongTheories-2007>.
- [Ahn20] J. Ahn, “Mekler’s construction and tree properties”, preprint, 2020. arXiv 1903.07087
- [AKL22] J. Ahn, J. Kim, and J. Lee, “On the antichain tree property”, *J. Math. Logic* **23**:2 (2022), art. id. 2250021.
- [Bau02] A. Baudisch, “Mekler’s construction preserves CM-triviality”, *Ann. Pure Appl. Logic* **115**:1-3 (2002), 115–173. MR
- [Bod15] M. Bodirsky, “Ramsey classes: examples and constructions”, pp. 1–48 in *Surveys in Combinatorics 2015*, edited by A. Czumaj et al., London Math. Soc. Lecture Note Series **424**, Cambridge University Press, 2015.
- [CH19] A. Chernikov and N. Hempel, “Mekler’s construction and generalized stability”, *Israel J. Math.* **230**:2 (2019), 745–769.
- [CKHY23] P. Cubides-Kovacsics, M. Hils, and J. Ye, “Beautiful pairs”, preprint, 2023. arXiv 2112.00651
- [CPT19] A. Chernikov, D. Palacin, and K. Takeuchi, “On n -dependence”, *Notre Dame J. Form. Log.* **60**:2 (2019), 195–214.
- [dMRS25] C. d’Elbée, I. Müller, N. Ramsey, and D. Siniora, “Model-theoretic properties of nilpotent groups and Lie algebras”, *J. Algebra* **662** (2025), 640–701.
- [GH19] V. Guingona and C. D. Hill, “On positive local combinatorial dividing-lines in model theory”, *Arch. Math. Logic* **58**:3-4 (2019), 289–323. MR
- [GPS23] V. Guingona, M. Parnes, and L. Scow, “Products of classes of finite structures”, *Notre Dame J. Form. Log.* **64**:4 (2023), 441–469. MR

- [HM24] M. Hils and R. Mennuni, “The domination monoid in henselian valued fields”, *Pacific J. Math.* **328**:2 (2024), 287–323.
- [Hod93] W. Hodges, *Model theory*, Cambridge University Press, 1993.
- [Mek81] A. H. Mekler, “Stability of nilpotent groups of class 2 and prime exponent”, *J. Symb. Logic* **46**:4 (1981), 781–788.
- [MP23] N. Meir and A. Papadopoulos, “Practical and structural infinitary expansions”, preprint, 2023. arXiv 2212.08027
- [MPT23] N. Meir, A. Papadopoulos, and P. Touchard, “Generalised indiscernibles, dividing lines, and products of structures”, preprint, 2023. arXiv 2311.05996
- [MS14] M. Malliaris and S. Shelah, “Regularity lemmas for stable graphs”, *Trans. Am. Math. Soc.* **366**:3 (2014), 1551–1585.
- [Neš05] J. Nešetřil, “Ramsey classes and homogeneous structures”, *Combinatorics, Probability and Computing* **14**:1-2 (2005), 171–189.
- [NR77] J. Nešetřil and V. Rödl, “Partitions of finite relational and set systems”, *J. Combin. Theory A* **22**:3 (1977), 289–312.
- [Rid17] S. Rideau, “Some properties of analytic difference valued fields”, *J. Inst. Math. Jussieu* **16**:3 (2017), 447–499. MR
- [RKV23] S. Rideau-Kikuchi and M. Vicařía, “Imaginaries in equicharacteristic zero henselian fields”, preprint, 2023. arXiv 2311.00657
- [Sco15] L. Scow, “Indiscernibles, EM-types, and Ramsey classes of trees”, *Notre Dame J. Form. Log.* **56**:3 (2015), 429–447.
- [She90] S. Shelah, *Classification theory and the number of nonisomorphic models*, 2nd ed., Studies in Logic and the Foundations of Mathematics **92**, North-Holland, Amsterdam, 1990. MR
- [Sim15] P. Simon, *A guide to NIP theories*, Cambridge University Press, 2015.
- [Tou23] P. Touchard, “Stably embedded submodels of Henselian valued fields”, *Arch. Math. Logic* **63** (2024), 279–315.
- [TZ12] K. Tent and M. Ziegler, *A course in model theory*, Cambridge University Press, 2012.

Received January 20, 2025. Revised November 20, 2025.

BLAISE BOISSONNEAU
 FACULTY OF MATHEMATICS AND NATURAL SCIENCES
 HEINRICH HEINE UNIVERSITY DÜSSELDORF
 DÜSSELDORF
 GERMANY
 blaise.boissonneau@hhu.de

ARIS PAPADOPOULOS
 DEPARTMENT OF MATHEMATICS
 UNIVERSITY OF MARYLAND
 COLLEGE PARK, MD
 UNITED STATES
 aris@umd.edu

PIERRE TOUCHARD
INSTITÜT FÜR ALGEBRA
TECHNISCHE UNIVERSITÄT DRESDEN
DRESDEN
GERMANY
pierre.touchard@tu-dresden.de

AXIAL SYMMETRY IN CONVEX BODIES

RITESH GOENKA, KENNETH MOORE,
WEN RUI SUN AND ETHAN PATRICK WHITE

For a two-dimensional convex body, the Kovner–Besicovitch measure of symmetry is defined as the volume ratio of the largest centrally symmetric body contained inside the body to the original body. A classical result states that the Kovner–Besicovitch measure is at least $\frac{2}{3}$ for every convex body and equals $\frac{2}{3}$ for triangles. Lassak showed that an alternative measure of symmetry, i.e., symmetry about a line (axiality) has a value of at least $\frac{2}{3}$ for every convex body. However, the smallest known value of the axiality of a convex body is around 0.81584, achieved by a convex quadrilateral. We show that every plane convex body has axiality at least $\frac{2}{41}(10 + 3\sqrt{2}) \approx 0.69476$, thereby establishing a separation with the central symmetry measure. Moreover, we find a family of convex quadrilaterals with axiality approaching $\frac{1}{3}(\sqrt{2} + 1) \approx 0.80474$. We also establish improved bounds for a “folding” measure of axial symmetry for plane convex bodies, and for a generalization of axiality to high-dimensional convex bodies.

1. Introduction

Symmetry is central to the study of mathematics. Common examples include the invariance of lattices and tilings under translation, rotation, and reflection. In the context of geometric objects, there are several different notions of symmetry. For example, a three-dimensional body is said to exhibit mirror symmetry if there is a plane, reflecting the body through which yields itself. Symmetry measures are often used to quantify symmetry in convex bodies (compact convex sets). We refer the reader to [11] and [22] (see also [23]) for detailed accounts of various quantitative symmetry measures arising in convex geometry.

For $n \in \mathbb{N}$ and $0 \leq k < n$ integer, Chakerian and Stein [5] define the k -symmetry of a convex body $K \subset \mathbb{R}^n$ to be

$$\text{Sym}_k(K) = \frac{1}{\text{Vol}_n(K)} \left[\max_{k\text{-dim. flats } \mathcal{L}} \text{Vol}_n(K \cap \text{refl}_{\mathcal{L}}(K)) \right],$$

White is supported in part by an NSERC PDF. Moore is supported by ERC Advanced Grants “GeoScape”, no. 882971 and “ERMiD”, no. 101054936.

MSC2020: primary 52A10, 52A38; secondary 52A20, 52A41.

Keywords: convex body, axial symmetry, axiality, folding symmetry.

where $\text{Vol}_n(\cdot)$ is the n -dimensional Lebesgue measure and $\text{refl}_{\mathcal{L}}(K)$ is the convex body obtained by reflecting K through \mathcal{L} ,

$$\text{refl}_{\mathcal{L}}(K) := \{u + 2(v - u) : u \in K, v = \text{projection of } u \text{ onto } \mathcal{L}\}.$$

It is easy to see that $\text{Sym}_k(K)$ represents the largest k -symmetric body inside of the given set K , scaled by its volume. Let

$$\sigma(n, k) = \inf_{K \subset \mathbb{R}^n \text{ convex body}} \text{Sym}_k(K)$$

be the size of the largest k -symmetric body inside of a convex body in \mathbb{R}^n , minimized over all unit volume convex bodies. Chakerian and Stein [5, Theorem 3] also prove the general lower bound

$$(1) \quad \sigma(n, k) \geq \frac{\max\{k!, (n - k)!\}}{2^{n-k}n!}.$$

The case $k = 0$ corresponds to reflection through points. Thus $\sigma(n, 0)$ is a measure of the lowest amount of central symmetry that an n -dimensional convex body has. The bound $\sigma(n, 0) \geq 2^{-n}$ given by (1) was proved previously by Stein [21]. Improved bounds have been established in low dimensions, including $\sigma(2, 0) = \frac{2}{3}$ [2; 7; 16] and $\sigma(3, 0) \geq \frac{2}{9}$ [3]. The former of these is known as the Kovner–Besicovitch theorem and is notable in that the bound is realized by triangles. Fáry and Rédei [8] computed the central symmetry of the regular n -simplex Δ^n , thereby establishing the upper bound

$$(2) \quad \sigma(n, 0) \leq \text{Sym}_0(\Delta^n) = \frac{1}{(n + 1)^n} \sum_{i=0}^{\lfloor n/2 \rfloor} (-1)^i \binom{n+1}{i} (n + 1 - 2i)^n.$$

Similarly, the case $k = n - 1$ corresponds to reflection through hyperplanes. Thus $\sigma(n, n - 1)$ is a measure of the lowest amount of hyperplane mirror symmetry that an n -dimensional convex body has. A related measure of asymmetry called *chirality coefficient* has been introduced in attempts to quantify chirality in chemistry [9]. There are two differences: (1) it is defined as 1 minus the symmetry measure since it is a measure of asymmetry, and (2) it also allows for rotation and translation of the reflected body to maximize the overlap volume. Gilat and Gordon [10] showed that in two dimensions, the addition of rotation and translation does not improve overlap. The same is true in three dimensions if the body of maximum overlap is unique. Gilat and Gordon also obtained upper bounds on the chirality coefficient using Ball’s volume ratio inequality [1] for John ellipsoids inscribed in convex bodies. In particular, the volume ratio inequality yields

$$\sigma(n, k) \geq \frac{\pi^{n/2}n!}{n^{n/2}(n + 1)^{(n+1)/2}\Gamma(\frac{n}{2} + 1)},$$

since ellipsoids are k -symmetric for all $0 \leq k < n$. The above estimate is weaker than (1) for all sufficiently large values of n .

For plane convex bodies, the measure $\text{Sym}_1(\cdot)$ is often called axiality. Krakowski [13] established a lower bound of $\frac{5}{8}$ on axiality, and Lassak [15] improved it to $\frac{2}{3}$. A better lower bound of $2(\sqrt{2} - 1)$ was proved by Nohl [18] for the axiality of centrally symmetric bodies. Buda and Mislow [4] showed that the same lower bound also holds for triangles. Both these results are tight in the sense that there is a centrally symmetric parallelogram and a sequence of triangles with axiality equal to/approaching $2(\sqrt{2} - 1)$. This stood as the best upper bound on $\sigma(2, 1)$ until Choi improved the upper bound to approximately 0.81584 and conjectured that the true value of $\sigma(2, 1)$ is around 0.81 [6].

Building on the ideas of Lassak [15], we obtain improved lower and upper bounds on $\sigma(2, 1)$.

Theorem 1.1. *Any planar convex body has axiality at least $\frac{2}{41}(10 + 3\sqrt{2})$. There is a sequence of quadrilaterals for which the axiality values approach $\frac{1}{3}(\sqrt{2} + 1)$. Therefore, $0.69476 < \frac{2}{41}(10 + 3\sqrt{2}) \leq \sigma(2, 1) \leq \frac{1}{3}(\sqrt{2} + 1) < 0.80474$.*

Theorem 1.1 shows that there is a strict separation between $\sigma(2, 1)$ and $\sigma(2, 0)$. In particular, $\sigma(2, 1) > \sigma(2, 0)$. Setting $k = n - 1$ in (1) yields the lower bound

$$(3) \quad \sigma(n, n - 1) \geq \frac{1}{2n},$$

which together with (2) implies that $\sigma(n, n - 1) > \sigma(n, 0)$ for all $n \geq 11$. It remains to determine whether a similar separation also holds for intermediate dimensions $3 \leq n \leq 10$.

In high dimensions, no nontrivial upper bounds have been found for $\sigma(n, n - 1)$. We prove the following general theorem, which implies that the upper bound for $\sigma(2, 1)$ carries over to $\sigma(n, n - 1)$ with higher values of n .

Theorem 1.2. *Let $n \geq 2$ and $0 \leq k < n$ be integers. We have*

$$\max\{(2 - 2^{-1/n})^{-n}, \sigma(n, k)\} \geq \sigma(n + 1, k + 1).$$

In particular, $\sigma(n, n - 1) \leq \sigma(2, 1) \leq \frac{1}{3}(\sqrt{2} + 1)$.

The best known lower bound on $\sigma(n, n - 1)$ is given by (3), which decays to zero as $n \rightarrow \infty$. An interesting problem is to determine if $\sigma(n, n - 1)$ remains bounded away from zero as $n \rightarrow \infty$.

Lassak [15] defined a *folding measure* of hyperplane mirror symmetry $\text{Sym}_{\text{fold}}(K)$ as twice the volume ratio of the largest portion of K that can be folded into K by reflection about a hyperplane to the original body K . Let

$$\varphi(n) = \inf_{K \subset \mathbb{R}^n \text{ convex body}} \text{Sym}_{\text{fold}}(K)$$

be the minimum symmetry among all n -dimensional unit volume convex bodies. Lassak [15, Theorem 3] proved that $\varphi(2) \geq \frac{1}{4}$ but did not provide any upper bound. An upper bound of $\frac{1}{2}$ was shown in [19], but as we will discuss later, this result is incorrect. We improve the lower bound and repair the upper bound on $\varphi(2)$. And we also provide an improved lower bound for centrally symmetric bodies.

Theorem 1.3. *Any planar convex body has folding symmetry at least $\frac{3}{8}$. There is a sequence of parallelograms for which the folding symmetry values approach $1/\phi$, where ϕ is the golden ratio. Therefore, $\frac{3}{8} \leq \varphi(2) \leq 1/\phi < 0.61804$. Furthermore, any centrally symmetric planar convex body has folding symmetry at least $\frac{4}{9}$.*

The proofs of the lower and upper bounds in Theorem 1.1 are discussed in Section 2. The bound in high dimensions is proved in Section 4. And the bounds for folding symmetry are proved in Section 3.

2. Axiality in the plane

We prove lower and upper bounds on the lowest value of axiality for plane convex bodies.

2.1. Lower bound. We say a nondegenerate hexagon $ABCDEF$ is *axially regular* if (i) the line CF is an axis of symmetry for the hexagon, (ii) segments \overline{AB} , \overline{CF} , and \overline{DE} are parallel, and (iii) $2|\overline{AB}| = |\overline{CF}| = 2|\overline{DE}|$. Lassak shows that an axially regular hexagon can be inscribed in every plane convex body and that such hexagons have at least $\frac{2}{3}$ the area of the convex body [15, §3]. By gluing triangles to axially regular hexagons, and considering triangles that well-approximate a convex body, we achieve an improved lower bound on planar axial symmetry.

Our method is to write down a convex program (see page 280) that captures our new ideas. Let $K \subset \mathbb{R}^2$ be a convex body. The objective of the optimization will be to minimize the variable λ , representing the area ratio of an axially symmetric convex body contained in K . By the aforementioned result of Lassak, we can find an axially regular hexagon inscribed in K . Let the vertices of the hexagon be A, B, C, D, E, F as shown in Figure 1. We normalize so that the inscribed hexagon has unit area. Extend the edges of the hexagon to form a six-pointed star, with vertices labeled A', B', C', D', E', F' in counterclockwise order as shown. Without loss of generality, assume $|\overline{AA'}| \geq |\overline{BC'}|$. The main variables of the program will be the area of K contained in the 6 points of the star. Let a, b, c, d, e, f be these areas in counterclockwise order, beginning with a as the area of K contained in triangle AFA' . Let $t = a + b + c + d + e + f$ be the total area of K outside of the hexagon $ABCDEF$.

Let the supporting line of K at B intersect $\overline{AB'}$ at A'' and $\overline{CC'}$ at C'' . The sum of the areas $|\Delta AA''B|$ and $|\Delta BCC''|$ is at most $\frac{1}{6}$ since triangles ABB' and BCC'




Figure 1. An axially regular hexagon inscribed in the convex body K .

have areas at most $\frac{1}{6}$. This shows $b + c \leq \frac{1}{6}$. Similarly, we have $d + e \leq \frac{1}{6}$; this gives constraints (4) and (5). The areas of $\triangle AFA'$ and $\triangle BCC'$ sum to $\frac{1}{3}$. We conclude

$$a + b \leq |\triangle AFA'| \leq \frac{1}{3} - |\triangle BCC'| \leq \frac{1}{3} - d.$$

This shows $a + b + d \leq \frac{1}{3}$, and by a similar argument $c + e + f \leq \frac{1}{3}$, verifying constraints (6) and (7).

Now we add isosceles triangles to the hexagon while maintaining the axial symmetry about the line CF . Let the perpendicular bisector of \overline{AB} meet the boundary of K at M and M' as shown. By considering the supporting line of K at M , we see that the area of triangle ABM is at least $b/2$. Similarly, $\triangle DEM'$ has area at least $e/2$. The pair of isosceles triangles with bases \overline{AB} and \overline{DE} , and height the minimum of the heights of triangles ABM and DEM' , are contained in K and symmetric about the line CF . It follows that K contains an axially symmetric convex body of area at least $1 + \min\{b, e\}$. The idea of adding isosceles triangles can also be applied to the pairs of triangles AFA', EFF' and BCC', CDD' . In summary, we can find an axially symmetric convex body of area at least

$$1 + \min\{a, f\} + \min\{b, e\} + \min\{c, d\},$$

inside of K . This proves constraint (8).

If K occupies a large portion of $\triangle A'C'E'$ or $\triangle B'D'F'$, then K has almost as much axial symmetry as a triangle. We prove a short lemma to make this precise.

Lemma 2.1. *Let $K \subseteq L \subset \mathbb{R}^2$ be convex bodies. For any line ℓ , we have*

$$|K \cap \text{refl}_\ell(K)| \geq |L \cap \text{refl}_\ell(L)| - 2|L \setminus K|.$$

Proof. Let $J = L \setminus K$. For notational ease, let J^* , K^* , and L^* denote $\text{refl}_\ell(J)$, $\text{refl}_\ell(K)$, and $\text{refl}_\ell(L)$. We have

$$L \cap L^* = (K \cup J) \cap (K^* \cup J^*) = (K \cap K^*) \cup (K \cap J^*) \cup (J \cap (K^* \cup J^*)).$$

Since $|K \cap J^*| \leq |J|$ and $|J \cap (K^* \cup J^*)| \leq |J|$, we have the desired result. \square

For any triangle $T \subset \mathbb{R}^2$, Buda and Mislow [4] show that $\text{Sym}_1(T) \geq 2(\sqrt{2} - 1)$. The area of triangle $A'C'E'$ is $\frac{3}{2}$; it follows that there is a line ℓ that satisfies $|\Delta A'C'E' \cap \text{refl}_\ell(\Delta A'C'E')| \geq 3(\sqrt{2} - 1)$. The area inside $\Delta A'C'E'$ but outside of K is at most $\frac{1}{2} - a - c - e$. By Lemma 2.1, this implies there is an axially symmetric convex subset of K inside of triangle $A'C'E'$ of size at least

$$3(\sqrt{2} - 1) - 2(\frac{1}{2} - a - c - e) = 3\sqrt{2} - 4 + 2(a + c + e).$$

A similar argument applies to triangle $B'D'F'$, verifying constraints (9) and (10).

AXIAL SYMMETRY PROGRAM

VARIABLES: $\lambda, a, b, c, d, e, f, t$

MINIMIZE: λ

SUBJECT TO: $a, b, c, d, e, f \geq 0,$

$$t = a + b + c + d + e + f,$$

(4) $b + c \leq \frac{1}{6},$

(5) $d + e \leq \frac{1}{6},$

(6) $a + b + d \leq \frac{1}{3},$

(7) $c + e + f \leq \frac{1}{3},$

(8) $(1 + t)\lambda \geq 1 + \min\{a, f\} + \min\{b, e\} + \min\{c, d\},$

(9) $(1 + t)\lambda \geq 3\sqrt{2} - 4 + 2(a + c + e),$

(10) $(1 + t)\lambda \geq 3\sqrt{2} - 4 + 2(b + d + f).$

Proposition 2.2. $\sigma(2, 1) \geq \frac{2}{41}(10 + 3\sqrt{2}).$

Proof. We solve the axial symmetry program by treating it as a linear program parametrized by t , and determining the optimal objective value for the dual linear program in terms of t . We prove the correctness of our solution in Appendix A. \square

2.2. Upper bound. In [4], Choi improved the upper bound on $\sigma(2, 1)$ to approximately 0.81584 using a computer assisted proof, providing the first example with axiality lower than $2\sqrt{2} - 2$. He also conjectured that the true value of $\sigma(2, 1)$ is around 0.81.

We implemented a basic algorithm to iteratively search for low symmetry shapes, using simulated annealing. The program is available on GitHub [17], and can be




Figure 2. A family of quadrilaterals parametrized by ε , which yields a new upper bound for $\sigma(2, 1)$.

used to compute the axiality of polygons, including Choi’s example, and to run our search for lower symmetry ones. With this search we were able to find polygons with axiality < 0.806 , one of which is preloaded as one launches the program.

We initially intended to present a computational proof of a lower symmetry body in a similar fashion to Choi. However, we later discovered a family of polygons which has lower axiality than anything we obtained experimentally. The family becomes flat as it approaches the minimum axiality, which means the program would never obtain this number, unless there is a different better construction.

Proposition 2.3. $\sigma(2, 1) \leq \frac{1}{3}(\sqrt{2} + 1) \approx 0.8047378541$.

Proof. We consider the quadrilateral

$$\left\{ (0, 0), (1, 0), \left(1, \frac{\sqrt{2}}{1+\sqrt{2}}\varepsilon\right), \left(\frac{1}{\sqrt{2}}, \varepsilon\right) \right\}$$

(see Figure 2) and show that by sending $\varepsilon \rightarrow 0$, its axiality approaches the desired number.

To do this, one must compute the area of overlap in every combination of angle and translation of the reflecting line. The computation is cumbersome, but routine. We have placed our analysis in Appendix B. In short, one finds that reflections in the angle bisector of $\angle CAB$ and the vertical line $x = \frac{2}{3}$ both either attain or approach this number, and every other reflection line yields smaller overlap. \square

Propositions 2.3 and 2.2 together yield Theorem 1.1.

3. Folding symmetry in the plane

Lassak defined the folding symmetry of a convex body in [15]. It is an alternative measure of hyperplane symmetry. Define the *folding symmetry* of K by

$$\text{Sym}_{\text{fold}}(K) = \frac{2}{\text{Vol}_n(K)} \left[\max_{\text{halfspaces } \mathcal{H}} \{ \text{Vol}_n(K \cap \mathcal{H}) : \text{refl}_{\partial\mathcal{H}}(K \cap \mathcal{H}) \subset K \} \right].$$

We are interested in the minimal symmetry possible among all bodies, so define

$$\varphi(n) = \inf_{K \subset \mathbb{R}^n \text{ convex body}} \text{Sym}_{\text{fold}}(K).$$

The range of possibilities for $\varphi(n)$ is $[0, 1]$. Considering the case of plane convex bodies, Lassak made an initial improvement on this, showing that $\varphi(2) \in [\frac{1}{4}, 1]$.

3.1. Lower bound. We first improve the lower bound by setting up another optimization problem (see page 284). We begin as before with a convex body K and an inscribed axially regular convex hexagon, normalized to have area 1. Reuse the same labeling of the hexagon and six-pointed star as in Figure 1. It will be convenient to discuss our constraints when the hexagon is embedded in a coordinate plane. Let the vertices $A, B, D,$ and E of the hexagon in Figure 1 have coordinates $(\frac{w}{2}, h), (-\frac{w}{2}, h), (-\frac{w}{2}, -h),$ and $(\frac{w}{2}, -h),$ respectively. Suppose that the midpoint of \overline{CF} is $(uw, 0),$ for some $0 \leq u \leq \frac{1}{2}.$ This determines the coordinates of all other vertices in Figure 1; we display the resulting coordinates in Figure 3. Since the hexagon has area 1, we have $hw = \frac{1}{3}.$ Let a, b, c, d, e, f, t have the same meaning as in Section 2.1. The solution to the axial symmetry program proved in Appendix A gives constraints (13). We denote by $\lambda,$ the halved folding symmetry of $K;$ it is the objective function of our program.

Suppose the rightmost and leftmost part of K have x -coordinates αw and $\beta w,$ respectively; this gives (11). We will construct folds along vertical lines that capture as much area as possible from the right and left side of $K.$ Let P and




Figure 3. Inscribed axially regular hexagon with coordinates for the folding program.

Q be the topmost and bottommost points of K ; suppose their coordinates are $(m_1w, (\phi_B + 1)h)$ and $(m_2w, -(\phi_E + 1)h)$, respectively.

Lemma 3.1. *The area in K to the right of $x = v_1w$ can be folded into the area on the left, where*

$$v_1 = \max \left\{ \frac{2m_1 + 1}{4}, \frac{2m_2 + 1}{4}, \frac{2\alpha - 1}{4} \right\}.$$

Proof. The average of αw and $-w/2$ is $w(2\alpha - 1)/4$. Hence the area in K between $y = -h$ and $y = h$ belonging to the fold will be folded inside of K because $v_1 \geq (2\alpha - 1)/4$. The boundary of K between B and A is decreasing from $x = m_1w$ to $x = w/2$. Hence a fold along $x = (m_1w + w/2)/2 = w(2m_1 + 1)/4$ would fold area inside $K \cap \triangle ABB'$ into area inside $K \cap \triangle ABB'$. A symmetric situation occurs in $\triangle DEE'$. \square

Larger α and smaller β imply better positive lower bounds on $a + f$ and $c + d$; we calculate these next. Without loss of generality, suppose the line $x = \alpha w$ meets the boundary of K inside $\triangle AFA'$. Area a is made smallest when the boundary of K meets line $x = \alpha w$ on segment $\overline{FA'}$, the coordinates of this point are $(\alpha w, \frac{2\alpha - (2u - 3)h}{2u + 1})$. It follows that

$$a \geq \frac{1}{2}(2u + 1)hw - \frac{1}{2}hw(2u + 1) \left(\frac{4u - 2\alpha + 3}{2u + 1} \right) = \frac{\alpha - u - 1}{3}.$$

Similarly, if we suppose that $x = \beta w$ meets the boundary of K inside $\triangle BCC'$, then we obtain $c \geq (-\beta + u - 1)/3$. This proves constraints (14).

It will be helpful to assume $v_1, v_2 \in [-\frac{1}{2}, \frac{1}{2}]$. Since $\beta \geq u - 1 \geq -1$, we immediately obtain $(2\beta + 1)/4 \geq -\frac{1}{4}$ and so $v_2 \in [-\frac{1}{2}, \frac{1}{2}]$. On the other hand, suppose $\alpha > \frac{3}{2}$. Then the area inside K to the right of the folding line $x = (\frac{1}{2}\alpha - \frac{1}{4})w$ is at least

$$hw(u + 1 - \frac{1}{2}) - 2hw(\frac{1}{2}\alpha - \frac{1}{4} - \frac{1}{2}) + \frac{1}{3}(\alpha - u - 1) = \frac{1}{3}.$$

From the axial symmetry program, or constraint (13), we know $|K| = 1 + t < 1.44$. Therefore we obtain the bound $\lambda > 0.23$ when $\alpha > \frac{3}{2}$. It follows that we may assume constraints (12).

Next we calculate the area inside K to the right of the folding line $x = v_1w$; it is the sum of the areas of a rectangle, the triangle AEF , and the part of K inside triangles AFA' and EFF' . The area is at least

$$2h(\frac{1}{2}w - v_1w) + (wu + \frac{1}{2}w)h + a + f = \frac{1}{2} - \frac{2}{3}v_1 + \frac{1}{3}u + a + f.$$

A similar calculation gives a lower bound on the area to the left of $x = v_2w$. This gives constraints (23) and (24).

Note that b is less than equal to area of the trapezium formed by the line parallel to AB passing through P , and the lines AB , AB' , and BB' ; this gives the first part of constraints (15). The second part follows similarly. For constraints (16), note that b and e are greater than equal to the areas of $\triangle APB$ and $\triangle DQE$, respectively. Constraints (18) follow from the fact that P lies below the lines AB' and BB' . Constraints (19) follow similarly.

We assume without loss of generality that the supporting line of K at F makes an acute or right angle from the line DE in the counterclockwise direction. Let

$$k_1 = \sup\{k \in [0, 1] : (w(u + 1), kh) \in \text{int}(K)\},$$

where we define the supremum of an empty set to be 0. Let X be the point on ∂K with coordinates $(w(u + 1), k_1h)$.

Suppose that the tangent to ∂K at C also makes an acute or right angle from the line DE in the anticlockwise direction. Then we may define a point Y with coordinates $(w(u - 1), -k_2h)$ in a manner similar to X . Note that a and d are greater than or equal to the areas of $\triangle AXF$ and $\triangle CYD$, respectively; this gives constraints (17). Constraint (20) arises from the fact that the absolute slope of AX is larger than that of AP . Constraint (21) follows similarly.

Let $y_1 = \max\{k_1, \frac{1}{2}\}$. We claim that it is possible to fold the convex body from the top along the line $y = y_1h$. Let the line $y = y_1h$ intersect ∂K in points M and N on the left and right, respectively. Let M' and N' be the projections of M and N onto the line CF , respectively. The conditions $y_1 \geq k_1$ and $y_1 \geq \frac{1}{2}$ together ensure that the region $MBAN$ (where MB and AN are arcs along ∂K) folds from the top into the rectangle $MNN'M'$ and the region BPA (with arcs BP and PA along ∂K) folds into the rectangle $ABDE$. A similar argument shows that one can fold along the line $y = y_2h$ from the bottom.

Next, the area of the fold obtained along the line $y = y_1h$ is at least the sum of the areas of the region BPA (with BP and PA arcs along ∂K) and the trapezium cut off from the hexagon $ABCDEF$ by the folding line $y = y_1h$. This gives constraint (25). Constraint (26) follows similarly.

FOLDING SYMMETRY PROGRAM

VARIABLES: $\lambda, a, b, c, d, e, f, t, u, m_1, m_2, v_1, v_2, \alpha, \beta, \phi_B, \phi_E, k_1, k_2, y_1, y_2$

MINIMIZE: λ

SUBJECT TO: $a, b, c, d, e, f, u \in [0, \frac{1}{2}]$,

$$(11) \quad \alpha \in [u + 1, 2u + \frac{3}{2}], \quad \beta \in [2u - \frac{3}{2}, u - 1],$$

$$(12) \quad v_1, v_2, m_1, m_2 \in [-\frac{1}{2}, \frac{1}{2}],$$

$$\phi_B, \phi_E, k_1, k_2 \in [0, 1], \quad y_1, y_2 \in [\frac{1}{2}, 1],$$

$$\begin{aligned}
 (13) \quad & t = a + b + c + d + e + f, \quad t \leq \frac{1}{4}(6 - 3\sqrt{2}), \\
 (14) \quad & a + f \geq \frac{1}{3}(\alpha - u - 1), \quad c + d \geq \frac{1}{3}(-\beta + u - 1), \\
 (15) \quad & b \leq \frac{1}{6}(1 - (1 - \phi_B)^2), \quad e \leq \frac{1}{6}(1 - (1 - \phi_E)^2), \\
 (16) \quad & b \geq \frac{1}{6}\phi_B, \quad e \geq \frac{1}{6}\phi_E, \\
 (17) \quad & a \geq \frac{1}{12}k_1(2u + 1), \quad d \geq \frac{1}{12}k_2(1 - 2u), \\
 (18) \quad & (1 - 2m_1) \geq \phi_B(1 + 2u), \quad (2m_1 + 1) \geq \phi_B(1 - 2u), \\
 (19) \quad & (1 - 2m_2) \geq \phi_E(1 + 2u), \quad (2m_2 + 1) \geq \phi_E(1 - 2u), \\
 (20) \quad & (1 - 2m_1)(1 - k_1) \geq \phi_B(2u + 1), \\
 (21) \quad & (1 + 2m_2)(1 - k_2) \geq \phi_E(1 - 2u), \\
 & v_1 = \max\left\{\frac{1}{4}(2m_1 + 1), \frac{1}{4}(2m_2 + 1), \frac{1}{4}(2\alpha - 1)\right\}, \\
 & v_2 = \min\left\{\frac{1}{4}(2m_1 - 1), \frac{1}{4}(2m_2 - 1), \frac{1}{4}(2\beta + 1)\right\}, \\
 (22) \quad & y_1 = \max\left\{k_1, \frac{1}{2}\right\}, \quad y_2 = \max\left\{k_2, \frac{1}{2}\right\}, \\
 (23) \quad & (1 + t)\lambda \geq \frac{1}{2} - \frac{2}{3}v_1 + \frac{1}{3}u + a + f, \\
 (24) \quad & (1 + t)\lambda \geq \frac{1}{2} + \frac{2}{3}v_2 - \frac{1}{3}u + c + d, \\
 (25) \quad & (1 + t)\lambda \geq b + \frac{1}{6}(1 - y_1)(3 - y_1), \\
 (26) \quad & (1 + t)\lambda \geq e + \frac{1}{6}(1 - y_2)(3 - y_2).
 \end{aligned}$$

Finally, we consider the case when the tangent to ∂K at C makes an obtuse angle from the line DE in the anticlockwise direction. In this case, the point Y lies on the arc BC with coordinates $(w(u - 1), k_2h)$. By similar arguments as the previous case, we get the following constraints in lieu of (17), (21), and (22), respectively:

$$\begin{aligned}
 & a \geq \frac{1}{12}k_1(2u + 1), \quad c \geq \frac{1}{12}k_2(1 - 2u), \\
 & (1 + 2m_1)(1 - k_2) \geq \phi_B(1 - 2u), \\
 & y_1 = \max\{k_1, k_2, \frac{1}{2}\}, \quad y_2 = \frac{1}{2}.
 \end{aligned}$$


Proposition 3.2. $\varphi(2) \geq \frac{3}{8}$.

Proof. We used Gurobi [12] to solve the folding symmetry program above. Although nonconvex, the program can be written so that it consists only of linear and quadratic constraints, which makes it susceptible to solution using Gurobi’s mixed-integer quadratically constrained program (MIQCP) solver. Per the previous discussion, we have two different cases based on the slope of the tangent to ∂K at C . In both cases, Gurobi’s MIQCP solver is able to establish a lower bound of 0.18803 on λ .

To establish these lower bounds, Gurobi first translates the program into bilinear form and then uses a spatial branch and bound algorithm to find feasible dual solutions with high objective values. In our computations, we set the constraint violation tolerance to 10^{-9} . Since the combined total number of variables and constraints in our program is bounded above by 100, the resulting error in the lower bound is a few orders of magnitude smaller than 10^{-3} ; hence the desired result. \square

3.2. Upper bound. Lassak did not provide an upper bound for this problem, but it was claimed in [19] that a family of parallelograms can be constructed to show $\varphi(2) \leq \frac{1}{2}$. Unfortunately, this paper contains a critical error, which we will explain in detail during the proof of the following theorem.

Proposition 3.3. *There is a sequence of parallelograms for which the folding symmetry approaches $1/\phi$, where ϕ is the golden ratio. Therefore, $\varphi(2) \leq 1/\phi \approx 0.61803$. For parallelograms, this value is tight, i.e., $\text{Sym}_{\text{fold}}(P) > 1/\phi$ for any parallelogram P .*



Proof. We consider the possible folds in a parallelogram as labeled in the figure just above. For any parallelogram, orienting it such that the longest side is AB will ensure that $d_2 > 0$. Now, there are three possible candidates for the best fold — the folding line ℓ can cut through

1. \overline{AB} and \overline{BD} or 2. \overline{AB} and \overline{CD} or 3. \overline{BD} and \overline{CD} .

Notably, the second case is not a possible fold unless $d_2 \geq d_1$, and this was the location of the error in [19]. The second case is simply omitted, and it is a very large fold in their final construction which is a family of parallelograms that becomes nearly rectangular. The other two cases are computed in their paper.

Case 1. We consider a line ℓ that cuts \overline{AB} at a distance x from B . If the reflection B' of B in ℓ does not lie on \overline{CD} (see Figure 4), the size of the fold can be increased by translating ℓ upwards until it does. It is shown in [19] that twice the area of the




Figure 4. The fold in case 1 (the blue triangular region).

fold in this case is

$$2 \cdot (\text{area of fold}) = \frac{hx^2}{x - d_1 + \sqrt{x^2 - h^2}}.$$

We have $d_1 < x \leq 1$; otherwise ℓ would pass above A or D . Now a quick optimization reveals that the size of the fold is maximized at $x = 1$. Thus the ratio of twice the area of the maximal fold in this case with the area of the parallelogram is

$$\frac{1}{1 - d_1 + \sqrt{1 - h^2}}.$$

Case 2. There is no freedom in choosing the angle of ℓ , so the best fold is along the perpendicular bisector of \overline{CD} (or \overline{AB}). This fold yields an area ratio of

$$\frac{1}{h} \left[\frac{h}{2} (2d_2 - 2d_1) + \frac{h}{2} (2d_1) \right] = 1 - d_1.$$

Case 3. In this case, twice the area of the maximal fold is found to be

$$|\overline{BD}|^2 \sin(\angle BDC) = |\overline{BD}|^2 \frac{h}{\sqrt{d_1^2 + h^2}} = h\sqrt{d_1^2 + h^2}.$$

So in general, for any parallelogram P , we have

$$\text{Sym}_{\text{fold}}(P) = \max \left\{ \frac{1}{1 - d_1 + \sqrt{1 - h^2}}, \sqrt{d_1^2 + h^2}, 1 - d_1 \right\}.$$

For a fixed d_1 , taking $h = 0$ minimizes each term, so to find the lowest possible symmetry we minimize $\max\{1/(2 - d_1), d_1, 1 - d_1\}$. This occurs at $d_1 = 2 - \phi$. A family of parallelograms with $d_1 = 2 - \phi$ and h descending to 0 therefore approaches a symmetry of $1/\phi$, and no parallelogram attains this value or anything lower. \square

3.3. Lower bound in the centrally symmetric case. We had another idea to obtain larger folds, which works for convex sets that are centrally symmetric. We will point out later how this idea might be easily generalizable to all convex bodies. Define $\varphi_{cs}(n)$ to be the minimum folding symmetry, taken only over the centrally symmetric bodies.

Proposition 3.4. $\varphi_{cs}(2) \geq \frac{4}{9}.$

Our proof idea is to consider appropriate folds after obtaining a large inscribed rectangle in the convex body. Radziszewski [20] showed that any convex body has an inscribed rectangle of area at least $\frac{1}{2}$, but as we shall later see, it would be more helpful if we could find smaller inscribed rectangles. In particular, we need the following lemma. If a parallelogram P is inscribed in a convex body K , then $\text{int}(K) \setminus P$ consists of 4 connected components whose closures we call *caps*.

Lemma 3.5. *Any centrally symmetric convex body K of area 1 has an inscribed rectangle of area r for any $r \leq \frac{1}{2}$.*

Proof. First, suppose K is strictly convex. Translate so that the origin is the centroid of K . For a given point $p \in \partial K$, we create an inscribed parallelogram having area r with one vertex at p as follows. Begin with the unique supporting line ℓ at p , and rotate it clockwise, causing it to form a chord in K . Let θ be the angle of rotation that causes ℓ to pass through the origin. For any angle between 0 and θ , there is a unique parallel second chord of equal length, and the four end points of these chords form a parallelogram.

The area of the parallelogram formed as a function of the rotation over $[0, \theta]$ is continuous, and we claim it has a maximum value of at least $\frac{1}{2}$. To see this, label the cap areas counterclockwise a, b, c , and d , where a is the cap bordered by ℓ . Initially, $a + c = 0$ and $b + d = 1$, and this is reversed once we rotate ℓ to angle θ . Thus, we can rotate until $a + c = \frac{1}{4}$, and let r be the area of the parallelogram in this case. Through an area preserving affine transformation, we can map the parallelogram to a square with center at the origin. A computation reveals that the amount of area in the other two caps is at most $\frac{1}{4}$, hence the parallelogram has area at least $\frac{1}{2}$. Said computation is almost identical to the one that proves Lemma 3.6 below, so we omit the details. We remark that one could also use the feature of central symmetry to just apply Lemma 3.6 towards an easy contradiction, but central symmetry is not necessary for the claim to hold.

This means we can rotate until we first reach a parallelogram of area r , which we denote by P_p . This process applied to every point on the boundary creates a continuous family of parallelograms. Now, define the functions

$$\begin{aligned} \tau : \partial K &\rightarrow \partial K, & \tau(p) &= \text{the next vertex after } p \text{ on } P_p \text{ clockwise along } \partial K, \\ \rho : \partial K &\rightarrow \mathbb{R}, & \rho(p) &= \text{the distance of } p \text{ from the origin, and} \\ \kappa : \partial K &\rightarrow \mathbb{R}, & \kappa(p) &= \rho(\tau(p)) - \rho(p). \end{aligned}$$

All these functions are continuous. The function κ effectively measures the “skew” of P_p , namely, $\kappa(p) = 0$ if and only if P_p is a rectangle. Therefore, we assume for contradiction that $\kappa(p)$ is never 0. The function κ is continuous, so without loss of generality we can assume it is always positive.

Since ∂K is compact, κ attains some minimum value $\delta > 0$ and the function ρ is uniformly continuous. Let p_0 be any point on ∂K , and let $p_n = \tau(p_{n-1})$ for $n = 1, 2, \dots$. This creates a countable sequence of points on ∂K , so we can find two elements arbitrarily close together. In particular, there are distinct elements p_i and p_j such that $|\rho(p_i) - \rho(p_j)| \leq \delta/2$. This is a contradiction, since $|\rho(p_i) - \rho(p_j)| \geq |i - j|\delta > \delta/2$.

Extending this to bodies that are not strictly convex is a standard limiting argument. For convenience, we replicate the argument found in Section 5 of [14]. We can approximate any convex body with a sequence of strictly convex sets $K_i \subseteq K$ converging pointwise to K . In each of these, we can inscribe a rectangle R_i of area r from our previous work. By compactness, we can make four selections of finer subsequences to make each of the four vertices converge to a point of ∂K . \square

Lemma 3.5 is the only gap in proving Proposition 3.4 in complete generality. We conjecture that the lemma should be true for all convex bodies, but our method of proof does not seem to generalize.

Now, if we inscribe a rectangle of area $\frac{1}{3}$ for example, we can immediately find a cap of size at least $\frac{1}{6}$, and this is enough to obtain a fold of size $\frac{1}{6}$. However, using the following lemma which holds for general convex sets, we can improve this further.

Lemma 3.6. *Let K be a convex body of area 1, with an inscribed rectangle R of area $r \leq \frac{1}{2}$. There is a cap of area at least $\frac{1}{4}(\sqrt{1 - 2r} + 1 - r)$.*

Proof. There are two opposite caps with areas a and c whose summed area is at least $\frac{1}{2}(1 - r)$. Suppose $a \geq c$, and denote the areas of the smaller pair of caps by b and d . To start, we instead scale the picture so that the area of the rectangle is 4, and translate so that the centroid of the rectangle is the origin. We will modify K , and attempt to put as much area into the other two caps as possible while maintaining the area of $a + c$ and keeping R fixed.

Let the vertices of the rectangle be $A, B, C,$ and D be as shown in Figure 5, where the side cap cut off by AB has area a . First, observe that every cap should be made triangular. Indeed, if the smaller caps are not triangular, there is additional area that may freely be added. If the area a cap is not triangular, we can “squish” the cap, by turning supporting lines at A, B inward onto the cap until their crossing, along with A and B , forms a triangle with area a , and replace the cap with this triangle. This creates additional room for the small caps to occupy.




Figure 5. The caps of a convex body with an inscribed rectangle.




Figure 6. The triangle used to compute the largest cap area.

Next, let P and Q be the extra vertices of the large triangular caps with areas a and c , respectively. These two points now control the areas of all the caps, in the sense that the lines PA , PB , QC , and QD border all of them. Now, by translating Q up and down, we keep a and c constant while changing the areas of the small caps b and d . A routine computation of the areas as a function of Q_y shows that $b+d$ is maximized when $Q_y = P_y$. At this stage, we can translate P and Q together, which does not change $a+c$ and $b+d$. Therefore, to minimize the largest cap, we will set $P_y = Q_y = 0$ and $P_x = -Q_x$.

Finally, we have reached a centrally symmetric rhombus whose largest cap is simple to compute. For the total area, we quarter the diagram, resulting in a triangle with an inscribed rectangle of area 1, as shown in Figure 6.

The ratio between the areas of the rectangle and the triangle is $r' \leq r$, and we compute the area of the cap-portion $\triangle BP'P$ of the triangle to be

$$\frac{a}{2} = \frac{1}{2r'}(\sqrt{1-2r'} + 1 - r') \geq \frac{1}{2r}(\sqrt{1-2r} + 1 - r).$$

Returning to the original scale yields the desired number. □

Proof of Proposition 3.4. Let K be a centrally symmetric plane convex body with unit area. Find an inscribed rectangle R in K with area $\frac{4}{9}$. Let the vertices of the rectangle be A, B, C , and D as shown in Figure 5. Let a, b, c , and d denote the areas of the caps with bases \overline{AB} , \overline{BC} , \overline{CD} , and \overline{DA} , respectively. Without loss of generality, we may assume that a is the largest of these areas. By Lemma 3.6,

$$a \geq \frac{1}{4}(\sqrt{1-2(\frac{4}{9})} + 1 - \frac{4}{9}) = \frac{2}{9}.$$

Let S be the cap with base \overline{AB} and P be a point in S that is farthest from AB . Let Q be the foot of the perpendicular from P to AB . If $|\overline{PQ}| \leq |\overline{BC}|$, then the cap S can be folded along the line AB to stay fully inside K , in which case

$$\text{Sym}_{\text{fold}}(K) \geq 2 \times \text{area}(S) = 2a \geq \frac{4}{9}.$$

Thus, we may henceforth assume $|\overline{PQ}|$ exceeds $|\overline{BC}|$. Let $\alpha > 0$ be such that $|\overline{PQ}| = (1+2\alpha)|\overline{BC}|$. Let ℓ be the line cutting through S and parallel to AB at a perpendicular distance $(1+\alpha)|\overline{BC}|$ from P . Let S' be the smaller cap obtained by the intersection of S with the halfspace defined by ℓ containing P . It is evident that

S' can be folded along ℓ to stay fully inside K . Let X and Y be the intersection points of ℓ with the lines BC and DA , respectively. Then we have

$$\begin{aligned} \text{area}(S') &\geq \text{area of triangle } APB - \text{area of rectangle } ABXY \\ &= \frac{1}{2} \cdot |\overline{AB}| \cdot (1 + 2\alpha)|\overline{BC}| - |\overline{AB}| \cdot \alpha|\overline{BC}| \\ &= \frac{1}{2} |\overline{AB}| |\overline{BC}| = \frac{2}{9}. \end{aligned}$$

This implies that $\text{Sym}_{\text{fold}}(K) \geq \frac{4}{9}$. □

We remark that $\frac{4}{9}$ is the optimal choice for the area of the inscribed rectangle in the above proof. This can be verified by taking an inscribed rectangle of area r and finding the r for which the resulting expression for the lower bound on $\text{Sym}_{\text{fold}}(K)$ is maximized.

Proposition 3.4 together with Propositions 3.2 and 3.3 yields Theorem 1.3.

4. Bounds in higher dimensions

In dimensions 3 and above, we know of no existing upper bounds on $\sigma(n, k)$ for $k > 0$ in the literature. In this section we obtain a bound on $\sigma(n + 1, k + 1)$ in terms of $\sigma(n, k)$, which in particular yields the best known upper bounds on $\sigma(n, n - 1)$ for all $n \geq 3$. We need the following lemma, which helps bound the symmetry of bodies that are very near other bodies in volume.

Lemma 4.1. *Let $K \subset \mathbb{R}^n$ be a convex body, and let $k \in \{0, \dots, n - 1\}$. Define two other convex bodies, $K_1 \subseteq K \subseteq K_2$, such that $\text{Vol}_n(K_1) \geq (1 - \varepsilon) \text{Vol}_n(K_2)$ for a positive constant $\varepsilon < 1$, and $\text{Sym}_k(K_1) = \text{Sym}_k(K_2)$. Then*

$$|\text{Sym}_k(K) - \text{Sym}_k(K_2)| \leq \varepsilon.$$

Proof. Let \mathcal{L} be the maximal k -flat so that $\text{Vol}_n(K \cap \text{refl}_{\mathcal{L}}(K)) = \text{Vol}_n(K)\text{Sym}_k(K)$. $K \cap \text{refl}_{\mathcal{L}}(K)$ is contained in $K_2 \cap \text{refl}_{\mathcal{L}}(K_2)$, so

$$\text{Sym}_k(K_2) \geq \frac{1}{\text{Vol}_n(K_2)} \text{Vol}_n(K)\text{Sym}_k(K) \geq (1 - \varepsilon)\text{Sym}_k(K).$$

Next, let \mathcal{L}' be the k -flat such that $\text{Vol}_n(K_1 \cap \text{refl}_{\mathcal{L}'}(K_1)) = \text{Vol}_n(K_1)\text{Sym}_k(K_1)$. Now $K_1 \cap \text{refl}_{\mathcal{L}'}(K_1)$ is contained in $K \cap \text{refl}_{\mathcal{L}'}(K)$, so

$$\text{Sym}_k(K) \geq \frac{1}{\text{Vol}_n(K)} \text{Vol}_n(K_1)\text{Sym}_k(K_1) \geq (1 - \varepsilon)\text{Sym}_k(K_1) = (1 - \varepsilon)\text{Sym}_k(K_2),$$

and the claim follows. □

In the following, we use $\text{Circ}(K)$ to denote the diameter of the smallest closed sphere fully containing a bounded set K .




Figure 7. The pyramid construction, K_{n+1} .

Proof of Theorem 1.2. Let $\varepsilon > 0$, then take a body $K_n \in \mathbb{R}^n$ with k -symmetry $\leq \sigma(n, k) + \varepsilon$ and rescale so that $\text{Vol}_n(K_n) = n + 1$. Also, translate so the centroid of K_n is the origin. Create a pyramid K_{n+1} in \mathbb{R}^{n+1} by placing a point p orthogonally above the centroid of K_n at a suitable distance D which will be specified later; but for now, assume it is larger than $2 \cdot (n + 1) \cdot \text{Circ}(K_n)$. Denote by ℓ_D the line joining p to the centroid of K_n as shown in Figure 7. The volume of this new shape is $\text{Vol}_{n+1}(K_{n+1}) = \frac{D}{n+1} \text{Vol}_n(K_n) = D$.

We now define θ_D to be the small angle in the right triangle with perpendicular side lengths D and $2(n + 1) \text{Circ}(K_n)$. Given a $(k + 1)$ -dimensional subspace $\mathcal{L} \subseteq \mathbb{R}^{n+1}$, there are three cases for the angle between \mathcal{L} and ℓ_D :

1. \mathcal{L} is within θ_D of ℓ_D .
2. \mathcal{L} is within θ_D of the orthogonal complement of ℓ_D .
3. Neither of the above hold.

Case 1. We will attain a bound on the symmetry realized by such a subspace slightly above $\text{Sym}_{n-1}(K_n)$. Fix a hyperplane \mathcal{L} sufficiently close in angle to ℓ_D . Then the projection ℓ of ℓ_D onto this subspace is at an angle of $\leq \theta_D$ with ℓ_D . Parametrize hyperplanes H_t orthogonal to ℓ , where H_0 passes through p , and as t increases, we move in the direction towards the base of the pyramid. We will let $\tau = \sup_{t>0}(H_t \cap K_{n+1} \neq \emptyset)$. Finally, we treat $K_t := H_t \cap K_{n+1}$ as a subset of \mathbb{R}^n with the origin at the intersection of ℓ_D with K_t and let $f(t) = \text{Vol}_n(K_t) \text{Sym}_{n-1}(K_t)$. We now estimate that

$$(27) \quad (K_{n+1} \cap \text{refl}_{\mathcal{L}}(K_{n+1})) \leq \int_0^\tau f(t) dt,$$

as we suppose that our chosen subspace happened to reflect every slice optimally. Now, increase D until θ_D becomes small enough so there are parameters $t_1 \leq t \leq t_2$ such that $(t_1/D) \cdot K_n \subseteq K_t \subseteq (t_2/D) \cdot K_n$ with t_1 sufficiently close to t_2 to apply Lemma 4.1 using ε . In particular, one can compute that suitable slices parallel to the base can be chosen at a distance at most $D \sec(\theta_D) - D + 2 \tan(\theta_D)$ apart along ℓ_D , and by plugging θ_D into this expression we see it goes to zero as D grows. As a by-product, we also have that $\text{Vol}_n(K_t) \leq \text{Vol}_n(\frac{t}{D} K_n)(1 + \varepsilon)$. Recall that the volume of $\frac{t}{D} K_n$ is $(\frac{t}{D})^n \text{Vol}_n(K_n)$. We show this configuration of slices in Figure 8.




Figure 8. The slice K_t , between two scaled copies of the base of the pyramid.

Now, it is always the case that $\tau \geq D$, but by choosing D large enough, we can make $\tau^{n+1} \leq D^{n+1}(1 + \varepsilon)$. Using this fact with equation (27), we have

$$\begin{aligned} & \frac{1}{\text{Vol}_{n+1}(K_{n+1})} \text{Vol}_{n+1}(K_{n+1} \cap \text{refl}_{\mathcal{L}}(K_{n+1})) \\ & \leq \frac{1}{\text{Vol}_{n+1}(K_{n+1})} \int_0^\tau \left(\text{Vol}_n\left(\frac{t}{D}K_n\right)(1 + \varepsilon) \right) ((\text{Sym}_{n-1}(K_n) + \varepsilon) dt \\ & = (1 + \varepsilon) \frac{\tau^{n+1}}{D^n(n+1)} \frac{\text{Vol}_n(K_n)(\text{Sym}_{n-1}(K_n) + \varepsilon)}{\text{Vol}_{n+1}(K_{n+1})} \\ & \leq (1 + \varepsilon) \frac{D^{n+1}(1 + \varepsilon)}{D^{n+1}} (\text{Sym}_{n-1}(K_n) + \varepsilon) \\ & = (1 + \varepsilon)^2 (\text{Sym}_{n-1}(K_n) + \varepsilon). \end{aligned}$$

Case 2. Take a hyperplane H that contains \mathcal{L} , and whose normal vector makes an angle less than θ_D with ℓ_D . Let τ now be the distance from q to the base of the pyramid. Let H_t be the hyperplanes parallel to H , parametrized so that H_0 passes through p , and $H_{D-\tau}$ passes through q , and let $K_t = H_t \cap K_{n+1}$. A reflection of any slice K_t in \mathcal{L} can be decomposed into a reflection in H and a reflection within H_t .

Taking D to be larger if necessary, we make $\text{Vol}_n(K_t) \leq \text{Vol}_n\left(\frac{t}{D}K_n\right)(1 + \varepsilon)$. Supposing that the reflections of each slice intersect perfectly, we have the upper bound

$$\begin{aligned} \frac{\text{Vol}_{n+1}(K_{n+1} \cap \text{refl}_{\mathcal{L}}(K_{n+1}))}{\text{Vol}_{n+1}(K_{n+1})} & \leq \frac{2}{D} \int_{D-2\tau}^{D-\tau} \text{Vol}_n\left(\frac{t}{D}K_n\right)(1 + \varepsilon) dt \\ & = \frac{2}{D} \int_{D-2\tau}^{D-\tau} \frac{t^n}{D^n}(n + 1)(1 + \varepsilon) dt \\ & = \frac{2(1 + \varepsilon)}{D^{n+1}} [(D - \tau)^{n+1} - (D - 2\tau)^{n+1}] =: *, \end{aligned}$$

which attains its max at

$$\tau = D \left(\frac{1}{2 - 2^{1/n+2}} + \frac{1}{2} \right).$$

Plugging this in and simplifying, we get

$$* \leq 2(1 + \varepsilon) \left[\left(\frac{2^{1/n}}{2^{1/n+1} - 1} \right)^{n+1} - \left(\frac{1}{2^{1/n+1} - 1} \right)^{n+1} \right] = (1 + \varepsilon) \left(\frac{1}{2 - 2^{-1/n}} \right)^n,$$

which creates the second part of the bound in the theorem.

Case 3. Here we can overestimate how much overlap there is by considering a cylinder of diameter $\text{Circ}(K_n)$ and height D . Because of the skew in angle, the reflection of the cylinder can overlap with the original cylinder in only at most $\text{Circ}(K_n) \cot(\theta_D) = \frac{D}{2(n+1)}$ of its height. Thus, the maximal volume of overlap is $\leq \text{Vol}_n(K_n) \cdot \frac{D}{2(n+1)} = \frac{D}{2}$, i.e., at most $\frac{1}{2}$ the area of the pyramid. This bound is weaker than the one in Case 2 for all n .

Altogether, in all cases the most symmetry a subspace can yield is

$$\max \left\{ \frac{1 + \varepsilon}{(2 - 2^{-1/n})^n}, (1 + \varepsilon)^2(\sigma(n, n - 1) + 2\varepsilon) \right\}.$$

Since ε was arbitrary, we are done. □

Appendix A. Analysis of the program for Proposition 2.2

We will consider four cases according to inequalities between a, f and b, e . By symmetry, we may assume $c \leq d$ without loss of generality.

Case 1: $a \leq f, b \geq e$. Sum 3 times constraint (8) and $\frac{3}{2}$ times (10) and obtain

$$(28) \quad \frac{9}{2}(1 + t)\lambda \geq 3 + \frac{3}{2}(3\sqrt{2} - 4) + 3t.$$

This inequality becomes weaker on λ as t grows. By constraints (6) and (7), we see that $t \leq \frac{2}{3}$. Substituting $t = \frac{2}{3}$ into (28) gives the bound $\lambda > 0.715$.

Case 2: $a \geq f, b \geq e$. By constraint (8),

$$\lambda \geq \frac{1 + f + e + c}{1 + t} \geq \frac{1}{1 + a + b + d}.$$

By constraint (6), the above is at least $\frac{3}{4}$.

Case 3: $a \geq f, b \leq e$. The bound in Theorem 1.1 is tight in this case and the next one. Due to the inequalities between a, f and b, e , we only need to consider the following subprogram of the full axial symmetry program.

VARIABLES: $\lambda, a, b, c, d, e, f, t$

MINIMIZE: λ

SUBJECT TO: $a, b, c, d, e, f \geq 0,$

$$a + b + c + d + e + f \geq t,$$

$$d + e \leq \frac{1}{6},$$

$$a + b + d \leq \frac{1}{3},$$

$$(1 + t)\lambda \geq 1 + f + b + c,$$

$$(1 + t)\lambda \geq 3\sqrt{2} - 4 + 2(a + c + e).$$

Our strategy will be to treat t as a constant, so that the above is a linear program. We will write down the corresponding dual program, and give a lower bound on the maximum of the dual for all values of t . The dual program is

VARIABLES: y_1, y_2, y_3, y_4, y_5

MAXIMIZE: $ty_1 - y_2/6 - y_3/3 + y_4 + (3\sqrt{2} - 4)y_5$

SUBJECT TO: $y_1, y_2, y_3, y_4, y_5 \geq 0,$

$$(1 + t)(y_4 + y_5) \leq 1,$$

$$y_1 - y_3 - 2y_5 \leq 0,$$

$$y_1 - y_3 - y_4 - y_5 \leq 0,$$

$$y_1 - y_4 - 2y_5 \leq 0,$$

$$y_1 - y_2 - y_3 \leq 0,$$

$$y_1 - y_2 - 2y_5 \leq 0,$$

$$y_1 - y_4 \leq 0.$$

The dual variables correspond to the equations of the primal in the order written above. We observe two feasible solutions to the dual, the vectors of y_i variables being respectively

$$\frac{1}{1+t}(0, 0, 0, 1, 0) \quad \text{and} \quad \frac{1}{1+t}\left(\frac{4}{5}, \frac{2}{5}, \frac{2}{5}, \frac{4}{5}, \frac{1}{5}\right).$$

This shows that the objective of the modified primal program above satisfies

$$\lambda \geq \max \left\{ \frac{1}{1+t}, \frac{1}{5(1+t)}(4t + 3\sqrt{2} - 1) \right\}.$$

The minimum value is achieved at $t = \frac{6-3\sqrt{2}}{4}$, giving the bound $\lambda \geq \frac{4}{10-3\sqrt{2}}$.

Case 4: $a \leq f, b \leq e$. Similar to the previous case, we only need to consider the following subprogram of the full axial symmetry program.

$$\begin{aligned}
&\text{VARIABLES: } \lambda, a, b, c, d, e, f, t \\
&\text{MINIMIZE: } \lambda \\
&\text{SUBJECT TO: } a, b, c, d, e, f \geq 0, \\
&\quad a + b + c + d + e + f \geq t, \\
&\quad d + e \leq \frac{1}{6}, \\
&\quad c + e + f \leq \frac{1}{3}, \\
&\quad (1 + t)\lambda \geq 1 + a + b + c, \\
&\quad (1 + t)\lambda \geq 3\sqrt{2} - 4 + 2(b + d + f).
\end{aligned}$$

By swapping variables a with f , b with c , and d with e , we obtain exactly the primal subprogram written down in Case 3.

This concludes the proof of Proposition 2.2.

Appendix B. A detailed proof of Proposition 2.3


In this appendix, we prove $\sigma(2, 1) \leq \frac{1}{3}(\sqrt{2} + 1)$ by analyzing the axiality of the ε -parametrized quadrilateral shown in Figure 2. By sending $\varepsilon \rightarrow 0$, the axiality approaches the desired number.

Let α be the argument of a direction vector normal to the reflection line, which we denote ℓ . The vertices of the reflected body are called A' , B' , C' , and D' . To summarize each case in the proof, we break up the choices for the angle and translate of ℓ into the following:

- (1) **Small-angle case:** $0 \leq \alpha \leq \frac{1}{2} \arctan \frac{\sqrt{2}\varepsilon}{1+\sqrt{2}}$. The translate is recorded using the x -intercept of ℓ , and can take values such that
 - (a) ℓ intersects \overline{AC} , or
 - (b) ℓ intersects \overline{CD} .
- (2) **Middle-angle case:** $\frac{\pi}{2} \leq \alpha \leq \frac{1}{2}(\pi + \arctan(\sqrt{2}\varepsilon))$. The translate is recorded using the y -intercept of ℓ , and may take values that are
 - (a) positive and large enough so that C' is above AB ,
 - (b) negative,
 - (c) positive such that ℓ intersects \overline{CD} and C' is below AB , or
 - (d) positive such that ℓ intersects \overline{BD} and C' is below AB .
- (3) $-\frac{\pi}{4} \leq \alpha < \frac{\pi}{4}$, excluding the interval of the small-angle case. All translates are easily ruled out.

- (4) $\frac{1}{2}(\pi + \arctan(\sqrt{2}\varepsilon)) \leq \alpha < \frac{3\pi}{4}$. We check translates such that
 - (a) C' lies above AB , or
 - (b) C' lies below AB .
- (5) $\frac{\pi}{4} \leq \alpha < \frac{\pi}{2}$. All translates are easily ruled out.

B.1. Cases (1) and (3). Reflection lines passing through \overline{CD} are quite poor, so we can rule out case (1.b). Indeed, as ε becomes small, such a line yields symmetry at most $\frac{1}{2}(3 - \sqrt{2}) \approx 0.79$, since it loses the entire triangle cut off by the line $D'B'$. Further, an angle for which B' is above AB , or D' is below AC , is sharply worse than otherwise. B' lying below AB implies $\alpha \geq 0$, and D' lying above AC implies $\alpha \leq \frac{1}{2} \arctan \frac{\sqrt{2}\varepsilon}{1+\sqrt{2}}$, which is because the largest possible such angle comes from the perpendicular bisector of A and D . We have therefore ruled out case (3), and it only remains to check in this section the case (1.a), which has this configuration:



With these considerations, let R and S be the intersections of $\overline{D'B'}$ with \overline{AB} and \overline{AC} , respectively, and P and Q the intersections of ℓ with \overline{AB} and \overline{AC} , respectively. Let t be the distance from A to P .

We compute the locations of the four points,

$$\begin{aligned}
 P &= t(1, 0), \\
 R &= \left(t - \frac{1-t}{\cos 2\alpha}\right)(1, 0), \\
 S &= \frac{\cot 2\alpha}{\sqrt{2}\varepsilon + \cot 2\alpha} \left(t - \frac{1-t}{\cos 2\alpha}\right)(1, \sqrt{2}\varepsilon), \\
 Q &= \frac{t \cot \alpha}{\sqrt{2}\varepsilon + \cot \alpha}(1, \sqrt{2}\varepsilon),
 \end{aligned}$$

and then use the shoelace formula to compute the area:

$$\begin{aligned}
 2 \cdot \text{Area} &= Q_x S_y - Q_y S_x + S_x R_y - S_y R_x + R_x P_y - R_y P_x + P_x Q_y - P_y Q_x \\
 &= -S_y R_x + P_x Q_y \\
 &= \varepsilon\sqrt{2} \left[t^2 \frac{\cot \alpha}{\sqrt{2}\varepsilon + \cot \alpha} - \left(t - \frac{1-t}{\cos 2\alpha}\right)^2 \frac{\cot 2\alpha}{\sqrt{2}\varepsilon + \cot 2\alpha} \right].
 \end{aligned}$$

Expanding this and writing it as a quadratic in t gives

$$\frac{2 \cdot \text{Area}}{\sqrt{2}\varepsilon} = t^2 \left(\frac{\cot \alpha}{\cot \alpha + \sqrt{2}\varepsilon} - \frac{\cot 2\alpha + 2 \csc 2\alpha + \csc 2\alpha \sec 2\alpha}{\cot 2\alpha + \sqrt{2}\varepsilon} \right) + t \cdot \frac{2 \csc 2\alpha + 2 \csc 2\alpha \sec 2\alpha}{\cot 2\alpha + \sqrt{2}\varepsilon} - \frac{\csc 2\alpha \sec 2\alpha}{\cot 2\alpha + \sqrt{2}\varepsilon}.$$

We can find the maximum of a quadratic easily, and it simplifies significantly. The maximum area, altogether, is

$$\frac{\varepsilon\sqrt{2}}{\sqrt{2}\varepsilon \sin \alpha + 2 \cos \alpha + 1}.$$


So we now have the maximal ratio with the original body:

$$\frac{1}{2 - \sqrt{2}} \cdot \frac{\sqrt{2}}{\sqrt{2}\varepsilon \sin \alpha + 2 \cos \alpha + 1} = \frac{1 + \sqrt{2}}{\sqrt{2}\varepsilon \sin \alpha + 2 \cos \alpha + 1}.$$

We need to find for what interval containing α is this ratio at most $\frac{1}{3}(1 + \sqrt{2})$. This is equivalent to finding where $\sqrt{2}\varepsilon \sin \alpha + 2 \cos \alpha \geq 2$, which is the case when $0 \leq \alpha < 2 \arctan \frac{\varepsilon}{\sqrt{2}}$. This in particular is valid for $\alpha \leq \frac{1}{2} \arctan \frac{\sqrt{2}\varepsilon}{1 + \sqrt{2}}$, so case (1.a) is complete. As an aside, one can easily check that for $\alpha = 0$, the optimal translate of the line occurs when $t = \frac{2}{3}$ and attains the value $\frac{1}{3}(1 + \sqrt{2})$ precisely.

B.2. Cases (2), (4), and (5). Here, the translate t is the distance of ℓ along the y -axis. For convenience, we let $\beta = \alpha - \frac{\pi}{2}$. Denote the internal angle bisector of AB and AC by ℓ_0 , and the angle it makes with AB by $\beta_0 = \frac{1}{2} \arctan(\varepsilon\sqrt{2})$. First, suppose $\beta > \beta_0$, so that we are in case (4). Let S and S_0 be the intersections of ℓ and ℓ_0 with \overline{CD} respectively (it will become clear later that ℓ must intersect \overline{CD} for any good translate, so S is well defined), and let C' and C'_0 be the reflections of C in ℓ and ℓ_0 respectively. We denote the quadrilateral (A, C'_0, S_0, C) by \mathcal{A}_0 , which comes from reflection in ℓ_0 .

We first discuss case (4.a), which are reflections having this configuration:



If C' is above C'_0S_0 then ℓ is trivially worse than ℓ_0 , since the area of overlap obtained by reflection in ℓ is strictly contained in \mathcal{A}_0 . If C' is below C'_0S_0 but still above AB , we have two triangles $\triangle C'_0C'M$ and $\triangle SS_0M$, where $M = C'S \cap C'_0S_0$. Observe that $\triangle C'_0C'M$ represents an area gained over \mathcal{A}_0 , and $\triangle SS_0M$ an area lost. We show $|\triangle C'_0C'M| < |\triangle SS_0M|$, by showing that M approaches C'_0 as ε goes to 0.


The lines $C'S$ and C'_0S_0 are given by the equations

$$y = \frac{2\varepsilon}{1-2\varepsilon^2}(\sqrt{2}x - \sqrt{1+2\varepsilon^2}) \quad \text{and} \quad y = \frac{\sqrt{2}\varepsilon \cot 2\beta + 1}{\cot 2\beta - \sqrt{2}\varepsilon}x - \frac{(2\varepsilon-t) \csc 2\beta}{\cot 2\beta - \sqrt{2}\varepsilon} + t,$$

respectively. The worst possible translate (giving the largest $\Delta C'_0C'M$ and smallest ΔSS_0M) is when $\overline{C'}$ lies on \overline{AB} . For fixed β , the appropriate translation required for C' to lie on the \overline{AB} is given by $t = \varepsilon(1 + \frac{1}{2} \sec^2 \beta) - \frac{1}{\sqrt{2}} \tan \beta$. Substituting this value of t into the equations defining C'_0S_0 and $C'S$, then solving simultaneously for M_x , we get


$$M_x = \frac{(\cot 2\beta - \sqrt{2}\varepsilon)(2\varepsilon\sqrt{1+2\varepsilon^2} + (\varepsilon(1 + \frac{1}{2} \sec^2 \beta) - \frac{1}{\sqrt{2}} \tan \beta)(1 - 2\varepsilon^2)) - (1 - 2\varepsilon^2)(\frac{1}{\sqrt{2}} \tan \beta - \frac{1}{2}\varepsilon \sec^2 \beta)}{(1 + 2\varepsilon^2)(\sqrt{2}\varepsilon \cot 2\beta - 1) \sin 2\beta}.$$

As ε goes to 0, one can compute that M_x approaches $\frac{1}{\sqrt{2}}$. Next we address case (4.b), which has this configuration:



We can see that $\overline{PR_0}$ is longer than $\overline{R_0S_0}$, otherwise ΔAPC is easily shown to be too large an area excluded from the overlap. As we translate downwards in this case, we lose a slice of width roughly $|\overline{PR}| \geq |\overline{PR_0}|$, and gain a slice of size roughly $|\overline{RS}| \leq |\overline{R_0S_0}|$. Therefore, one can use an integral argument to see that translation in this direction is a loss, and this case is strictly worse than case (4.a). This concludes all of case (4).

We move on now to case (2). The hardest out of all the cases to analyze is (2.c), which looks like this:



We rule out case (2.a), where C' is below AB , since translating ℓ downwards strictly increases the area. We observe further that by the same argument that disqualifies case (4.b) above, the translate must be positive, so A' sits above A , thereby settling case (2.b). We now make the assumption that \overline{CD} and $\overline{C'D'}$ intersect

at the point we will call S , so we are in case (2.c). Observe that S is still to the right of S_0 , otherwise the overlap region is entirely contained \mathcal{A}_0 . We compute the locations of the four vertices in the figure, which are

$$\begin{aligned} P &= \frac{t}{\sqrt{2\varepsilon} - \tan \beta} (1, \sqrt{2\varepsilon}), \\ Q &= \left(P_x \frac{2\varepsilon^2 + 2\sqrt{2\varepsilon} \cot 2\beta - 1}{(\sqrt{2\varepsilon} \cot 2\beta - 1)}, 0 \right), \\ R &= \left(\frac{(2\varepsilon - t) \sec 2\beta + \sqrt{2\varepsilon} t \tan 2\beta - t}{\sqrt{2\varepsilon} + \tan 2\beta}, 0 \right), \\ S &= \left(\frac{2\varepsilon - t}{\tan \beta + \sqrt{2\varepsilon}}, 2\varepsilon - \sqrt{2\varepsilon} \frac{2\varepsilon - t}{\tan \beta + \sqrt{2\varepsilon}} \right). \end{aligned}$$

This time, there are fewer cancellations in the shoelace formula. We have

$$\begin{aligned} 2 \cdot \text{Area} &= P_x Q_y + Q_x R_y + R_x S_y + S_x P_y - Q_x P_y - R_x Q_y - S_x R_y - P_x S_y \\ &= P_x \cdot 0 + Q_x \cdot 0 + R_x \cdot (-\sqrt{2\varepsilon} S_x + 2\varepsilon) + S_x \cdot \sqrt{2\varepsilon} P_x \\ &\quad - Q_x \cdot \sqrt{2\varepsilon} P_x - R_x \cdot 0 - S_x \cdot 0 - P_x \cdot (-\sqrt{2\varepsilon} S_x + 2\varepsilon) \\ &= \sqrt{2\varepsilon} [\sqrt{2} P_x (\sqrt{2} S_x - 1) + R_x (\sqrt{2} - S_x) - P_x Q_x]. \end{aligned}$$

Plugging those quantities in results in a monstrous expression, which can be expanded and written as a quadratic $at^2 + bt + c$ in the same manner as in case (2.c). For simplicity, let $k = \sqrt{2\varepsilon}$. We computed the coefficients, which are

$$\begin{aligned} a &= (\tan \beta + k)(\tan 2\beta + k)(2k + (k^2 - 1) \tan 2\beta) \\ &\quad + 2(\tan \beta - k)(\tan 2\beta - k)(\tan 2\beta + k) \\ &\quad + (\tan \beta - k)^2 (\tan 2\beta - k)(k \tan 2\beta - \sec 2\beta - 1), \\ b &= \sqrt{2}(\tan \beta - k)^2 (\tan 2\beta - k) \\ &\quad \times (\tan 2\beta + k + (\tan \beta (k \tan 2\beta - \sec 2\beta - 1) + k \sec 2\beta)), \\ c &= 2k(\tan \beta - k)^2 (\tan 2\beta - k) \tan \beta \sec 2\beta. \end{aligned}$$

Maximizing over t gives the value

$$(29) \quad m(\beta) = 2(1 + \sqrt{2})k \frac{\tan \beta g(\beta) - k \sec 2\beta (\tan \beta - k)^2 (\tan 2\beta - k)}{\cos 2\beta (\tan \beta + k)(\tan 2\beta + k)g(\beta)},$$

where

$$\begin{aligned} g(\beta) &= (\tan \beta + k)(\tan 2\beta + k)(2k + (2k - 1) \tan 2\beta) \\ &\quad + 2(\tan \beta - k)(\tan 2\beta - k)(\tan 2\beta + k) \\ &\quad + (\tan \beta - k)^2 (\tan 2\beta - k)(k \tan 2\beta - \sec 2\beta - 1). \end{aligned}$$

This can be moderately simplified to

$$m(\beta) = -4(1 + \sqrt{2})k \frac{(k + 2) \cos \beta + (k - 2) \cos 3\beta - 2 \sin \beta}{h(\beta)},$$

where

$$h(\beta) = -(14k^2 + 4k) \cos \beta - 8k^2 \cos 3\beta + (4k - 2k^2) \cos 5\beta + (2k - 4) \sin \beta + (k^3 - 2k^2 + 9k - 2) \sin 3\beta + (k^3 - 2k^2 - k + 2) \sin 5\beta.$$

When k is small, the important pieces become the constant terms, which only occur on the bottom:


$$-4 \sin \beta - 2 \sin 3\beta + 2 \sin 5\beta.$$

However, this collection of terms vanishes when β is small. In particular, let $\beta \leq \arctan k$ (this value is chosen to avoid singularities which can be found using the first expression for m , equation (29); otherwise any constant times $\arctan k$ would have worked as well). Thus we are interested only in the k terms, which are

$$\frac{1}{4(1 + \sqrt{2})} m(\beta) \approx \frac{-2 \cos \beta + 2 \cos 3\beta + 2 \sin \beta}{-4 \cos \beta + 4 \cos 5\beta + 2 \sin \beta + 9 \sin 3\beta - \sin 5\beta}.$$

Finally, near 0, the cosine terms all cancel, and we are left with the contribution of the sine terms which is $\frac{2}{2+9\cdot 3-5} = \frac{1}{12}$, so that $m(\beta) \rightarrow \frac{1}{3}(1 + \sqrt{2})$.

Case (2.d) is when D' is reflected below D so that \overline{CD} and $\overline{C'D'}$ do not intersect:



This is in fact expected to be the optimal translate for some angles between 0 and $\arctan \frac{h\sqrt{2}}{1+\sqrt{2}}$; however, we can bound it using the previous computation. Our computation in case (2.c) still finds the area of the quadrilateral $PQRS_0$ shown in the figure just above. Since this contains the polygon created from (2.d), we can upper bound it using this area. The ratio from this area approached the correct value for all angles in case (2), so we are done.

Case (5) is easily ruled out as well, noting that the shape is an isosceles triangle with a triangle cut off by the line BD . The reflection lines in this case are repetitions of the ones in case (2) and (4), only now, the larger end of the shape is reflected further apart.

References

- [1] K. Ball, “Volume ratios and a reverse isoperimetric inequality”, *J. London Math. Soc.* (2) **44**:2 (1991), 351–359. MR
- [2] A. S. Besicovitch, “Measure of asymmetry of convex curves”, *J. London Math. Soc.* **23** (1948), 237–240. MR
- [3] A. Bielecki and K. Radziszewski, “Sur les parallélépipèdes inscrits dans les corps convexes”, *Ann. Univ. Mariae Curie-Skłodowska Sect. A* **8** (1954), 97–100. MR
- [4] A. B. Buda and K. Mislow, “On a measure of axiality for triangular domains”, *Elem. Math.* **46**:3 (1991), 65–73. MR
- [5] G. D. Chakerian and S. K. Stein, “On measures of symmetry of convex bodies”, *Canadian J. Math.* **17** (1965), 497–504. MR
- [6] C.-Y. Choi, *Finding the largest inscribed axially symmetric polygon for a convex polygon*, Master’s thesis, EECS Dept., Korea Advanced Institute of Science and Technology, 2006.
- [7] I. Fáry, “Sur la densité des réseaux de domaines convexes”, *Bull. Soc. Math. France* **78** (950), 152–161. MR
- [8] I. Fáry and L. Rédei, “Der zentralsymmetrische Kern und die zentralsymmetrische Hülle von konvexen Körpern”, *Math. Ann.* **122** (1950), 205–220. MR
- [9] G. Gilat, “On the quantification of asymmetry in nature”, *Found. Phys. Lett.* **3** (1990), 189–196.
- [10] G. Gilat and Y. Gordon, “Geometric properties of chiral bodies”, *J. Math. Chem.* **16**:1-2 (1994), 37–48. MR
- [11] B. Grünbaum, “Measures of symmetry for convex sets”, pp. 233–270 in *Convexity*, Proc. Sympos. Pure Math. **7**, Amer. Math. Soc., 1963. MR
- [12] “Gurobi Optimizer reference manual”, online resource, 2023, at <https://docs.gurobi.com/projects/optimizer/en/current/index.html>.
- [13] F. Krakowski, “Bemerkung zu einer Arbeit von W. Nohl”, *Elemente der Mathematik* **18** (1963), 60–61.
- [14] M. Lassak, “Approximation of convex bodies by rectangles”, *Geom. Dedicata* **47**:1 (1993), 111–117. MR
- [15] M. Lassak, “Approximation of convex bodies by axially symmetric bodies”, *Proc. Amer. Math. Soc.* **130**:10 (2002), 3075–3084. MR
- [16] M. A. Lavrentiev and L. A. Lyusternik, *Elements of the calculus of variations*, vol. 1, part 2, Moscow, 1935.
- [17] K. Moore, “Convex symmetry”, online resource, 2023, at <https://github.com/Kenneth-Moore/Axiality>.
- [18] W. Nohl, “Die innere axiale Symmetrie zentrischer Eibereiche der euklidischen Ebene”, *Elemente der Mathematik* **17** (1962), 59–63.
- [19] M. Nowicka, “On the measure of axial symmetry with respect to folding for parallelograms”, *Beitr. Algebra Geom.* **53**:1 (2012), 97–103. MR
- [20] K. Radziszewski, “Sur un problème extrémal relatif aux figures inscrites et circonscrites aux figures convexes”, *Ann. Univ. Mariae Curie-Skłodowska Sect. A* **6** (1952), 5–18. MR
- [21] S. Stein, “The symmetry function in a convex body”, *Pacific J. Math.* **6** (1956), 145–148. MR
- [22] G. Toth, *Measures of symmetry for convex sets and stability*, Springer, 2015. MR
- [23] B. A. de Valcourt, “Measures of axial symmetry for ovals”, *Israel J. Math.* **4** (1966), 65–82. MR

Received May 14, 2024. Revised September 11, 2025.

RITESH GOENKA
MATHEMATICAL INSTITUTE
UNIVERSITY OF OXFORD
OXFORD
UNITED KINGDOM
ritesh.goenka@maths.ox.ac.uk

KENNETH MOORE
HUN-REN ALFRÉD RÉNYI INSTITUTE OF MATHEMATICS
BUDAPEST
HUNGARY
kjmoore@renyi.hu

WEN RUI SUN
DEPARTMENT OF MATHEMATICAL AND STATISTICAL SCIENCES
UNIVERSITY OF ALBERTA
EDMONTON, AB
CANADA
wrsun@ualberta.ca

ETHAN PATRICK WHITE
DEPARTMENT OF MATHEMATICS
UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN
URBANA, IL
UNITED STATES
epw@illinois.edu

A COMBINATORIAL STRUCTURE FOR MANY HIERARCHICALLY HYPERBOLIC SPACES

MARK HAGEN, GIORGIO MANGIONI AND ALESSANDRO SISTO

This has nothing to do with links. – A. S.

The combinatorial hierarchical hyperbolicity criterion is a very useful way of constructing new hierarchically hyperbolic spaces (HHSs). We show that, conversely, HHSs satisfying natural assumptions (satisfied, for example, by mapping class groups) admit a combinatorial HHS structure. This can be useful in constructions of new HHSs, and also our construction clarifies how to apply the combinatorial HHS criterion to suspected examples. We also uncover connections between HHS notions and lattice theory notions.

Introduction	305
1. Background on hierarchical hyperbolicity	310
2. Combinatorial HHSs	315
3. Combinatorial hyperbolicity from hierarchical hyperbolicity	318
4. Construction of the combinatorial HHS	323
5. Proof of the main theorem	327
6. Adding a group action	344
7. Some other hypotheses	348
8. Near equivalence of HHS and combinatorial HHS	354
9. Mapping class groups are combinatorial HHS	357
10. Why orthogonals for nonsplit domains?	363
Acknowledgements	375
References	375

Introduction

Showing that a given space or group is hierarchically hyperbolic yields a lot of information about it ([BHS21; HHP23; ANS+24; HHL23; DMS23] is a partial list), but it is quite challenging to check the definition directly. To remedy this, a criterion was devised in [BHMS24] to show that a space is hierarchically hyperbolic, roughly consisting of checking that a certain simplicial complex has links which are

MSC2020: 20F65.

Keywords: hierarchical hyperbolicity, combinatorial HHS.

hyperbolic, along with some combinatorial conditions. Such complexes, introduced formally in Section 2, are called combinatorial HHSs. The criterion has proven useful to show that various spaces and groups are HHS, see [BHMS24; HMS22; HRSS25; Rus25; DDLS24]. It is natural to wonder whether there is a converse to this, namely whether every HHS admits a combinatorial HHS structure. We show that is true under mild assumptions on the HHS, see Theorem 1 below.

Besides the intrinsic interest of a converse statement, our theorem has proven useful in constructing new examples of HHSs, by introducing a combinatorial structure which is easier to manipulate. For instance, the second author used the flexibility of combinatorial structures coming from the main theorem to construct uncountably many coarse median structures on certain HHGs including the mapping class group of the 5-holed sphere [Man24]. The same flexibility has in turn been used in [MS26] to construct Dehn-filling-like quotients of certain HHGs, leading to a proof of the Hopf property for generic Artin groups. Additionally, our construction provides a blueprint for identifying candidate combinatorial HHS structures, and for instance this is one of the ingredients in [Rag25].

We defer the precise statements of the additional conditions to Section 3; the reader can find a list of the conditions and where to find them below, together with a short discussion. For now, we just mention that the properties hold in many natural examples, including mapping class groups, see Section 9, and we now state (a stripped down version of) our main theorem:

Theorem 1 (see Theorem 3.16). *Let $(\mathcal{Z}, \mathfrak{S})$ be a hierarchically hyperbolic space with weak wedges, clean containers, the orthogonals for nonsplit domains property, and dense product regions. Then there exists a combinatorial HHS (X, \mathcal{W}) such that \mathcal{Z} is quasi-isometric to \mathcal{W} .*

Theorem 1 only contains the nonequivariant part of Theorem 3.16, but our constructions are equivariant in a suitable sense and compatible with the notion of hierarchically hyperbolic group, rather than just space, as would be needed for the application to quotients mentioned above; see Theorem 6.6 for the exact statement.

Our construction also clarifies how to construct, starting with a space or group that one suspects to be hierarchically hyperbolic, a candidate combinatorial HHS structure for it, which can then be used to show that the given space or group is indeed HHS. We explain this in Section 4. We do not know whether the additional conditions we have to impose are necessary, but we provide an example where our construction fails to yield a combinatorial HHS structure in the absence of the additional conditions, in Section 10. In fact, we uncovered intriguing connections with lattice theory that arise from these considerations, see in particular Remark 10.21. Roughly, modifying an HHS structure to ensure the additional condition reduces to a problem in lattice theory, see Question 10.22. To highlight

the connections with lattice theory, we note that the notion of an ortholattice is very closely related to those of wedges and clean containers (which have appeared in the HHS literature, for example in [BR20; ABD21; CRHK24; Rus22; AB23; Hag23]); see Definition 7.10.

To conclude this subsection, we suggest a possible application of the aforementioned construction: showing that mapping class groups of finite type nonorientable surfaces are hierarchically hyperbolic.

A true converse. A second aim of this paper is to clarify how various conditions on an HHS structure relate to each other and to properties of combinatorial HHS structure. As this is fairly technical we do not make precise statements in the introduction, but we refer the reader to Section 7 and in particular Lemma 7.4; we believe this can also be useful for applications. An especially striking output of this study is the following “true converse” theorem, which provides an actual equivalence between combinatorial HHS structures and HHS structures, each satisfying natural conditions (see list below for where to find each condition).

Theorem 2. *Let $(\mathcal{Z}, \mathfrak{S})$ be a hierarchically hyperbolic space. Then \mathcal{Z} has wedges, clean containers and the strong orthogonal property if and only if there exists a CHHS (X, \mathcal{W}) with simplicial wedges and simplicial containers such that \mathcal{W} is quasi-isometric to \mathcal{Z} .*

We note that the standard HHS structure on mapping class groups does not satisfy the stronger conditions above, but we construct a different structure which does in Section 9. We now state this result, together with the aforementioned result on the standard HHS structure, which in fact confirms the speculations from [BHMS24, Section 1.6].

The following combines Theorems 9.8 and 9.9 (which in fact give combinatorial HHG structures):

Theorem 3. *Let S be obtained from a closed connected oriented surface of finite genus by deleting finitely many points and open discs.*

There exists a combinatorial HHS structure (X, \mathcal{W}) for $\mathcal{MCG}(S)$, where X is the blow-up of the curve graph of S , obtained by replacing every curve with the cone over its annular curve graph.

Moreover, there exists a (different) combinatorial HHS structure for $\mathcal{MCG}(S)$ with simplicial wedges and simplicial containers.

In the new HHS structures, the elements of the index set are still isotopy classes of certain essential subsurfaces, as in [BHS19], and in particular the unbounded hyperbolic spaces involved in the HHS structure all still come from curve graphs of connected subsurfaces. The difference between the HHS structures in Section 9, which are used to construct combinatorial HHS structures, and the “standard” one

is in which subsurfaces with bounded curve graph are present; care in this choice is what allows us to arrange extra combinatorial properties of the index set needed to make a CHHS structure.

As a final note on mapping class groups, we summarise in the following remark the connection between our construction and clean markings:

Remark 1 (relation with clean markings). The reader familiar with Masur and Minsky’s graph of complete clean markings from [MM00] will notice that the combinatorial HHS structure for $\mathcal{MCG}(S)$ that our main theorem provides has the same flavour of the marking graph. Indeed, a maximal simplex of the graph X from Definition 4.2 will correspond to a choice of a maximal collection of disjoint annuli A_1, \dots, A_k (that is, a pants decomposition), plus a choice of a point x_i inside the annular curve graph associated to A_i for every $i = 1, \dots, k$ (that is, a transversal for every curve in the pants decomposition). Hence, a maximal simplex corresponds to a complete marking. Moreover, some of the \mathcal{W} -edges we define in Definition 4.10 correspond to elementary moves. Indeed, let $\Sigma, \Delta \subset X$ be two maximal simplices, and suppose that their supports differ by a single curve (say, the support of Σ is $\alpha \cup P$ and the support of Δ is $\beta \cup P$, for some almost-maximal collection of pairwise disjoint curves P). Then these simplices are \mathcal{W} -adjacent if and only if:

- α and β are close in the curve graph of the subsurface of complexity 1 cut out by P ;
- α projects close to the coordinate prescribed by Δ in the annular curve graph of β ;
- the same holds with α and β swapped.

In other words, our \mathcal{W} -edges detect when one replaces a curve with one of its transversals, and this corresponds to one of the elementary moves from [MM00].

Additional conditions. The additional conditions on an HHS structure in Theorems 1 and 2 are stated precisely in the following places:

- Wedges are defined in Property 3.1.
- Clean containers are defined in Property 3.4.
- Dense product regions are defined in Property 3.10.
- Orthogonals for nonsplit domains are defined in Property 3.9.
- The strong orthogonal property is defined in Property 7.1 (see also Remark 10.21).

Wedges and clean containers are standard assumptions on an HHS structure introduced in [BR20] and [ABD21] respectively — wedges make the nesting poset into a lattice, while clean containers make it into a complemented poset. Dense product regions is a coboundedness assumption automatically satisfied by HHGs.

Strong orthogonality corresponds to *orthomodularity* of the nesting lattice, and is the source of the lattice-theoretic question mentioned earlier. It is one way of verifying orthogonals for nonsplit domains, which is really the main enabling assumption in Theorem 3.16 and is modelled on the role of boundary annuli in the HHS structure on mapping class groups.

The additional conditions on a combinatorial HHS are:

- Simplicial wedges are defined in Definition 8.2.
- Simplicial containers are defined in Definition 8.1.

These conditions are very natural properties one might hope for from a simplicial complex, along the lines that containment of links of simplices corresponds to reverse containment of the simplices.

Outline of the paper. Sections 1 and 2 contain all relevant definitions and facts about (combinatorial) HHS. Section 3 gathers the hypotheses of the main result of this paper, which is Theorem 3.16 and states that a HHS $(\mathcal{Z}, \mathfrak{S})$ with some additional assumption is quasi-isometric to a combinatorial HHS. The actual construction of the candidate combinatorial HHS (X, \mathcal{W}) is in Section 4 (see in particular Definitions 4.2 and 4.10). The quasi-isometry $f : \mathcal{W} \rightarrow \mathcal{Z}$ is constructed in Definition 4.11, and maps any maximal simplex of X to one of its realisation points (in the sense of the partial realisation axiom (8)).

In Section 5 we verify that, under the assumptions from Section 3, the pair (X, \mathcal{W}) is a combinatorial HHS and f is a quasi-isometry (see Assumption 5.1 and Theorem 5.2). In Section 6 we show that our construction is equivariant, meaning that whenever a group G acts on $(\mathcal{Z}, \mathfrak{S})$ by hieromorphisms then it also acts on (X, \mathcal{W}) (see Theorem 6.2). Then we use this fact to prove that, whenever a group acts metrically properly and coboundedly on $(\mathcal{Z}, \mathfrak{S})$ and some other mild assumptions hold, then G has a structure of hierarchically hyperbolic group coming from the action on a combinatorial HHS (see Theorem 6.6).

In Section 7 we present some more “natural” hypotheses that one could require on $(\mathcal{Z}, \mathfrak{S})$, and we show how they relate to each other and to the ones from Section 3. In Section 8 we establish an equivalence between strong orthogonality properties on the HHS structure of $(\mathcal{Z}, \mathfrak{S})$ and some strong intersection properties on the links of the associated combinatorial HHS (X, \mathcal{W}) (see Theorem 8.3).

In Section 9 we apply our results to the mapping class group of a compact orientable surface, with the usual HHS structure (the one from, e.g., [BHS19, Section 11]), showing that it admits a combinatorial HHS structure whose underlying graph is a certain blow-up of the curve graph. This confirms the speculations from [BHMS24, Section 1.6]. Moreover, we show that Theorem 8.3 applies if one adds to the index set some nonessential subsurfaces, including pairs of pants (see Theorem 9.9).

Finally, in Section 10 we illustrate the necessity of the hypotheses of Theorem 3.16 by providing a counterexample of an unbounded space \mathcal{Z} for which the construction from Section 4 can only yield a bounded CHHS. Remarkably, \mathcal{Z} can be chosen to be a CAT(0) cube complex with a factor system, with the usual HHS structure (the one from [BHS17b, Remark 13.2]). Then we speculate on which conditions on the factor system could allow one to modify the HHS structure in order to satisfy our hypotheses.

1. Background on hierarchical hyperbolicity

1.1. Axioms. We recall from [BHS19] the definition of a hierarchically hyperbolic space.

Definition 1.1 (HHS). The q -quasigeodesic space $(\mathcal{Z}, d_{\mathcal{Z}})$ is a *hierarchically hyperbolic space* if there exists $E \geq 0$, called the *HHS constant*, an index set \mathfrak{S} , whose elements will be referred to as *domains*, and a set $\{\mathcal{C}U \mid U \in \mathfrak{S}\}$ of E -hyperbolic spaces $(\mathcal{C}U, d_U)$, called *coordinate spaces*, such that the following conditions are satisfied:

- (1) (**projections**) There is a set $\{\pi_U : \mathcal{Z} \rightarrow 2^{\mathcal{C}U} \mid U \in \mathfrak{S}\}$ of *projections* sending points in \mathcal{Z} to sets of diameter bounded by E in the various $\mathcal{C}U \in \mathfrak{S}$. Moreover, for all $U \in \mathfrak{S}$, the coarse map π_U is (E, E) -coarsely Lipschitz and $\pi_U(\mathcal{Z})$ is E -quasiconvex in $\mathcal{C}U$.
- (2) (**nesting**) \mathfrak{S} is equipped with a partial order \sqsubseteq , and either $\mathfrak{S} = \emptyset$ or \mathfrak{S} contains a unique \sqsubseteq -maximal element, denoted by S . When $V \sqsubseteq U$, we say V is *nested* in U . For each $U \in \mathfrak{S}$, we denote by \mathfrak{S}_U the set of $V \in \mathfrak{S}$ such that $V \sqsubseteq U$. Moreover, for all $U, V \in \mathfrak{S}$ with $V \not\sqsubseteq U$ there is a specified subset $\rho_U^V \subset \mathcal{C}U$ with $\text{diam}_{\mathcal{C}U}(\rho_U^V) \leq E$. There is also a *projection* $\rho_U^V : \mathcal{C}U \rightarrow 2^{\mathcal{C}V}$. (The similarity in notation is justified by viewing ρ_U^V as a coarsely constant map $\mathcal{C}U \rightarrow 2^{\mathcal{C}V}$.)
- (3) (**orthogonality**) \mathfrak{S} has a symmetric and anti-reflexive relation called *orthogonality*: we write $U \perp V$ when U, V are orthogonal. Also, whenever $V \sqsubseteq U$ and $U \perp W$, we require that $V \perp W$. We require that for each $T \in \mathfrak{S}$ and each $U \in \mathfrak{S}_T$ such that $\{V \in \mathfrak{S}_T \mid V \perp U\} \neq \emptyset$, there exists $W \in \mathfrak{S}_T - \{T\}$, which we call a *container* for U inside T , so that whenever $V \perp U$ and $V \sqsubseteq T$, we have $V \sqsubseteq W$. Finally, if $U \perp V$, then U, V are not \sqsubseteq -comparable.
- (4) (**transversality and consistency**) If $U, V \in \mathfrak{S}$ are not orthogonal and neither is nested in the other, then we say U, V are *transverse*, denoted $U \pitchfork V$. In this case there are sets $\rho_U^V \subseteq \mathcal{C}U$ and $\rho_V^U \subseteq \mathcal{C}V$, each of diameter at most E and satisfying:

$$\min \{d_U(\pi_U(z), \rho_U^V), d_V(\pi_V(z), \rho_V^U)\} \leq E$$

for all $z \in \mathcal{Z}$. Furthermore, for $U, V \in \mathfrak{S}$ satisfying $V \sqsubseteq U$ and for all $z \in \mathcal{Z}$, we have:

$$\min \{d_U(\pi_U(z), \rho_U^V), \text{diam}_{\mathcal{C}V}(\pi_V(z) \cup \rho_V^U(\pi_U(z)))\} \leq E.$$

The preceding two inequalities are the *consistency inequalities* for points in \mathcal{Z} .

Finally, if $U \sqsubseteq V$, then $d_W(\rho_W^U, \rho_W^V) \leq E$ whenever $W \in \mathfrak{S}$ satisfies either $V \not\sqsubseteq W$ or $V \pitchfork W$ and $W \not\sqsubseteq U$.

- (5) (**finite complexity**) There exists $n \geq 0$, the *complexity* of \mathcal{Z} (with respect to \mathfrak{S}), so that any set of pairwise- \sqsubseteq -comparable elements has cardinality at most n .
- (6) (**large links**) Let $U \in \mathfrak{S}$, let $z, z' \in \mathcal{Z}$ and let $N = d_U(\pi_U(z), \pi_U(z'))$. Then there exists $\{T_i\}_{i=1, \dots, \lfloor N \rfloor} \subseteq \mathfrak{S}_U - \{U\}$ such that, for any domain $T \in \mathfrak{S}_U - \{U\}$, either $T \in \mathfrak{S}_{T_i}$ for some i , or $d_T(\pi_T(z), \pi_T(z')) < E$. Also, $d_U(\pi_U(z), \rho_U^{T_i}) \leq N$ for each i .
- (7) (**bounded geodesic image**) For all $U \in \mathfrak{S}$, all $V \in \mathfrak{S}_U - \{U\}$, and all geodesics γ of $\mathcal{C}U$, either $\text{diam}_{\mathcal{C}V}(\rho_V^U(\gamma)) \leq E$ or $\gamma \cap \mathcal{N}_E(\rho_V^U) \neq \emptyset$.
- (8) (**partial realisation**) Let $\{V_j\}$ be a family of pairwise orthogonal elements of \mathfrak{S} , and let $p_j \in \pi_{V_j}(\mathcal{Z}) \subseteq \mathcal{C}V_j$. Then there exists $z \in \mathcal{Z}$, which we call a *partial realisation point* for the family, so that:
 - $d_{V_j}(z, p_j) \leq E$ for all j ,
 - for each j and each $V \in \mathfrak{S}$ with $V_j \sqsubseteq V$, we have $d_V(z, \rho_V^{V_j}) \leq E$, and
 - for each j and each $V \in \mathfrak{S}$ with $V_j \pitchfork V$, we have $d_V(z, \rho_V^{V_j}) \leq E$.
- (9) (**uniqueness**) For each $\kappa \geq 0$, there exists $\theta_u = \theta_u(\kappa)$ such that if $x, y \in \mathcal{Z}$ and $d_{\mathcal{Z}}(x, y) \geq \theta_u$, then there exists $V \in \mathfrak{S}$ such that $d_V(x, y) \geq \kappa$.

We often refer to \mathfrak{S} , together with the nesting and orthogonality relations, and the projections as a *hierarchically hyperbolic structure* for the space \mathcal{Z} . Observe that \mathcal{Z} is hierarchically hyperbolic with respect to $\mathfrak{S} = \emptyset$, i.e., hierarchically hyperbolic of complexity 0, if and only if \mathcal{Z} is bounded. Similarly, \mathcal{Z} is hierarchically hyperbolic of complexity 1 with respect to the index set $\mathfrak{S} = \{\mathcal{Z}\}$, if and only if \mathcal{Z} is hyperbolic.

Notation 1.2. Where it will not cause confusion, given $U \in \mathfrak{S}$, we will often suppress the projection map π_U when writing distances in $\mathcal{C}U$, i.e., given $x, y \in \mathcal{Z}$ and $p \in \mathcal{C}U$ we write $d_U(x, y)$ for $d_U(\pi_U(x), \pi_U(y))$ and $d_U(x, p)$ for $d_U(\pi_U(x), p)$. Note that when we measure distance between a pair of sets (typically both of bounded diameter) we are taking the minimum distance between the two sets. Given $A \subset \mathcal{Z}$ and $U \in \mathfrak{S}$ we set

$$\pi_U(A) = \bigcup_{a \in A} \pi_U(a).$$

1.2. Useful facts about HHS. We now recall results from [BHS19] that will be useful later on.

Lemma 1.3 [DHS17, Lemma 1.5]. *Let $U, V, W \in \mathfrak{S}$ satisfying $U \perp V$, and $U, V \not\sqsubseteq W$, and $W \not\sqsubseteq U, V$. Then $d_V(\rho_W^U, \rho_W^V) \leq 2E$.*

Remark 1.4 (normalisation). As argued in [BHS19, Remark 1.3], it is always possible to assume that the HHS structure is *normalised*, that is, for every $U \in \mathfrak{S}$ the projection $\pi_U : \mathcal{Z} \rightarrow \mathcal{C}U$ is uniformly coarsely surjective. In order to do so, one roughly replaces every $\mathcal{C}U$ with $\pi_U(\mathcal{Z})$, which is itself hyperbolic since it is quasiconvex in $\mathcal{C}U$, and then replaces every projection ρ_U^V with the composition $p_U \circ \rho_U^V$, where $p_U : \mathcal{C}U \rightarrow \pi_U(\mathcal{Z})$ is the coarse closest point retraction. The resulting space, which is again hierarchically hyperbolic, has the same set of domains \mathfrak{S} with the same relations of nesting and orthogonality.

Assumption 1.5. In view of the remark above, we will always assume that the HHS structures we consider are normalised.

Definition 1.6 (consistent tuple). Let $\kappa \geq 1$ and let $(b_U)_{U \in \mathfrak{S}} \in \prod_{U \in \mathfrak{S}} 2^{\mathcal{C}U}$ be a tuple such that for each $U \in \mathfrak{S}$, the U -coordinate b_U has diameter $\leq \kappa$. Then $(b_U)_{U \in \mathfrak{S}}$ is κ -consistent if for all $V, W \in \mathfrak{S}$, we have

$$\min\{d_V(b_V, \rho_V^W), d_W(b_W, \rho_W^V)\} \leq \kappa$$

whenever $V \pitchfork W$ and

$$\min\{d_W(b_W, \rho_W^V), \text{diam}_V(b_V \cup \rho_V^W(b_W))\} \leq \kappa$$

whenever $V \not\sqsubseteq W$.

Theorem 1.7 (realisation [BHS19, Theorem 3.1]). *Let $(\mathcal{Z}, \mathfrak{S})$ be a hierarchically hyperbolic space. Then for each $\kappa \geq 1$, there exists $\theta = \theta(\kappa)$ so that, for any κ -consistent tuple $(b_U)_{U \in \mathfrak{S}}$, there exists $x \in \mathcal{Z}$ such that $d_V(x, b_V) \leq \theta$ for all $V \in \mathfrak{S}$.*

The uniqueness axiom (Definition (9)) implies that the *realisation point* x for $(b_U)_{U \in \mathfrak{S}}$ provided by Theorem 1.7 is coarsely unique.

Definition 1.8 (product regions and factors). Fix a constant $\kappa \geq 1$. For any domain U , let F_U be the set of κ -consistent tuples for U , that is, all tuples $(b_V)_{V \in \mathfrak{S}_U}$ that satisfy the consistency inequalities involving only domains nested in U . Similarly, one can define E_U as the set of κ -consistent tuples of the form $(b_V)_{V \perp U}$.

Now let $P_U = F_U \times E_U$, which we call the *product region* associated to U . By the realisation Theorem 1.7 there is a coarsely well-defined map $\phi : P_U \rightarrow \mathcal{Z}$. If we fix $e \in E_U$, the image of the *factor* $F_U \times \{e\}$, which we will still denote by F_U when the dependence on e is irrelevant, can be endowed with the subspace metric, which makes it a sub-HHS of \mathcal{Z} with domain set $\mathfrak{S}_U = \{V \in \mathfrak{S} \mid V \sqsubseteq U\}$. Two parallel

copies $F_U \times \{e\}$ and $F_U \times \{e'\}$ are quasi-isometric (see e.g. [DHS20, Section 2.2]); thus the metric structure on F_U is well-defined up to quasi-isometry.

A similar argument holds for E_U . For more details on product regions, see [CRHK24, Section 15].

It will often be convenient to think of F_U as an abstract space, instead of as a subspace of \mathcal{Z} . This way, whenever $V \sqsubseteq U$, we have a (nonunique) embedding $F_V \rightarrow F_U$, given as follows. Let E_V^U be the set of κ -consistent tuples of the form

$$(b_W)_{W \not\sqsubseteq U, W \perp V_i \forall i=1, \dots, k},$$

and choose $e \in E_V^U$. Then define a map $F_V \rightarrow F_U$ by sending a tuple $(y_W)_{W \sqsubseteq V}$ to the tuple $(x_W)_{W \sqsubseteq U}$, defined as follows:

$$x_W = \begin{cases} y_W & \text{if } W \sqsubseteq V; \\ \rho_W^V & \text{if } V \not\sqsubseteq W \text{ or } V \pitchfork W; \\ e_W & \text{if } W \perp V. \end{cases}$$

In other words, we extend the tuple y “naturally” whenever we have a well-defined projection from V to W , and then we choose consistent coordinates whenever $W \perp V$. This kind of argument will recur throughout the paper.

Definition 1.9 (relative product regions). Fix a constant $\kappa \geq 0$. Let $U, V_1, \dots, V_k \in \mathfrak{S}$ be such that $V_i \sqsubseteq U$ and $V_i \perp V_j$ for every $i, j \leq k$. The *relative product region* associated to V_1, \dots, V_k inside U is defined as

$$P_{\{V_i\}}^U = F_{V_1} \times \dots \times F_{V_k} \times E_{\{V_i\}}^U \subset F_U,$$

where $E_{\{V_i\}}^U$ is the set of κ -consistent tuples of the form

$$(b_W)_{W \not\sqsubseteq U, W \perp V_i \forall i=1, \dots, k}.$$

With a slight abuse of notation, whenever the ambient domain U is clear we will drop the superscript and refer to the relative product region simply as $P_{\{V_i\}}$.

Remark 1.10. For the rest of the paper, unless otherwise stated, all factors and relative product regions are with respect to $\kappa = 20E$, where E is the HHS constant. This specific choice of κ will be motivated in Assumption 5.1.

Theorem 1.11 (distance formula [BHS19, Theorem 4.5]). *Let $(\mathcal{Z}, \mathfrak{S})$ be a hierarchically hyperbolic space. Then there exists s_0 such that for all $s \geq s_0$, there exist C, K so that for all $x, y \in \mathcal{Z}$,*

$$d(x, y) \asymp_{K,C} \sum_{U \in \mathfrak{S}} \{d_U(x, y)\}_s.$$

(The notation $\{A\}_B$ denotes the quantity which is A if $A \geq B$ and 0 otherwise. The notation $A \asymp_{K,C} B$ means $A \leq KB + C$ and $B \leq KA + C$.)

1.3. Groups acting on HHS. First, we need to discuss which group actions we allow on a hierarchically hyperbolic space. The following are some definitions from [BHS19] and [BHMS24]:

Definition 1.12 (automorphism). Let $(\mathcal{Z}, \mathfrak{S})$ be a HHS. An *automorphism* consists of a map $g : \mathcal{Z} \rightarrow \mathcal{Z}$, a bijection $g^\sharp : \mathfrak{S} \rightarrow \mathfrak{S}$ preserving nesting and orthogonality, and, for each $U \in \mathfrak{S}$, an isometry $g^\diamond(U) : \mathcal{C}U \rightarrow \mathcal{C}(g^\sharp(U))$ for which the following two diagrams commute for all $U, V \in \mathfrak{S}$ such that $U \not\sqsubset V$ or $U \pitchfork V$:

$$\begin{array}{ccc}
 \mathcal{Z} & \xrightarrow{g} & \mathcal{Z} \\
 \downarrow \pi_U & & \downarrow \pi_{g^\sharp(U)} \\
 \mathcal{C}U & \xrightarrow{g^\diamond(U)} & \mathcal{C}(g^\sharp(U))
 \end{array}
 \quad \text{and} \quad
 \begin{array}{ccc}
 \mathcal{C}U & \xrightarrow{g^\diamond(U)} & \mathcal{C}(g^\sharp(U)) \\
 \downarrow \rho_V^U & & \downarrow \rho_{g^\sharp(U)}^{g^\sharp(U)} \\
 \mathcal{C}V & \xrightarrow{g^\diamond(V)} & \mathcal{C}(g^\sharp(V))
 \end{array}$$

Notice that g must be a uniform quasi-isometry by the distance formula, Theorem 1.11. Whenever it will not cause ambiguity, we will abuse notation by dropping the superscripts and just calling all maps g .

We say that two automorphisms g, g' are *equivalent*, and we write $g \sim g'$, if $g^\sharp = (g')^\sharp$ and $g^\diamond(U) = (g')^\diamond(U)$ for each $U \in \mathfrak{S}$. Given an automorphism g , a quasi-inverse \bar{g} for g is an automorphism with $\bar{g}^\sharp = (g^\sharp)^{-1}$ and such that, for every $U \in \mathfrak{S}$, $\bar{g}^\diamond(U) = g^\diamond(U)^{-1}$. Since the composition of two automorphisms is an automorphism, the set of equivalence classes of automorphisms forms a group, denoted $\text{Aut}(\mathfrak{S})$.

Definition 1.13. A finitely generated group G *acts* on a HHS $(\mathcal{Z}, \mathfrak{S})$ by automorphisms if there is a homomorphism $G \rightarrow \text{Aut}(\mathfrak{S})$.

Remark 1.14. The original definition of an automorphism, which is [BHS19, Definition 1.20], only requires the diagrams from Definition 1.12 to coarsely commute (with uniform constants). However, as shown in [DHS20, Section 2.1], if G acts on $(\mathcal{Z}, \mathfrak{S})$ in the sense of [BHS19] then one can ensure that the diagrams genuinely commute by perturbing every $\pi_U : \mathcal{Z} \rightarrow \mathcal{C}U$ and every ρ_V^U , whenever the quantity is defined, by a uniformly bounded amount. This way, up to a single initial change in the constant E , the HHS structure is unaffected, meaning that the new structure has the same domain set \mathfrak{S} with the same relations and the same coordinate spaces.

Definition 1.15 (HHG). A finitely generated group G is *hierarchically hyperbolic* if there exists a hierarchically hyperbolic space $(\mathcal{Z}, \mathfrak{S})$ and an action $G \rightarrow \text{Aut}(\mathfrak{S})$ so that the uniform quasi-action of G on \mathcal{Z} is metrically proper and cobounded and \mathfrak{S} contains finitely many G -orbits. Then we can equip G with a HHS structure, whose domains and coordinate spaces are the same as the ones for \mathcal{Z} and whose projections are obtained by precomposing the projections for $(\mathcal{Z}, \mathfrak{S})$ with the G -equivariant quasi-isometry $G \rightarrow \mathcal{Z}$ given by the Milnor–Schwarz lemma.

2. Combinatorial HHS

In this section we recall the definition of a combinatorial HHS and its hierarchically hyperbolic structure, as first introduced in [BHMS24].

2.1. Basic definitions. Let X be a simplicial graph.

Definition 2.1 (induced subgraph). Given a subset $S \subseteq X^{(0)}$ of the set of vertices of X , the subgraph *spanned* by S is the complete subgraph of X with vertex set S .

Definition 2.2 (join, link, star). Given disjoint simplices Δ, Δ' of X , we let $\Delta \star \Delta'$ denote the simplex spanned by $\Delta^{(0)} \cup \Delta'^{(0)}$, if it exists.

For each simplex Δ , the *link* $\text{Lk}(\Delta)$ is the union of all simplices Σ of X such that $\Sigma \cap \Delta = \emptyset$ and $\Sigma \star \Delta$ is a simplex of X . Observe that $\text{Lk}(\Delta) = \emptyset$ if and only if Δ is a maximal simplex. The link of a subgraph of X is the intersection of the links of its vertices.

The *star* of Δ is $\text{Star}(\Delta) := \text{Lk}(\Delta) \star \Delta$, i.e., the union of all simplices of X that contain Δ .

Definition 2.3 (X -graph, \mathcal{W} -augmented graph). An X -graph is a graph \mathcal{W} whose vertex set is the set of all maximal simplices of X .

For a simplicial graph X and an X -graph \mathcal{W} , the \mathcal{W} -augmented graph $X^{+\mathcal{W}}$ is the graph defined as follows:

- the 0-skeleton of $X^{+\mathcal{W}}$ is $X^{(0)}$;
- if $v, w \in X^{(0)}$ are adjacent in X , then they are adjacent in $X^{+\mathcal{W}}$;
- if two vertices in \mathcal{W} are adjacent, then we consider σ, ρ , the associated maximal simplices of X , and in $X^{+\mathcal{W}}$ we connect each vertex of σ to each vertex of ρ .

We equip \mathcal{W} with the usual path-metric, in which each edge has unit length, and do the same for $X^{+\mathcal{W}}$.

2.2. HHS structure. [BHMS24, Theorem 1.18] states that, under some assumptions on the pair (X, \mathcal{W}) , \mathcal{W} has the hierarchically hyperbolic structure described below. First, we define what will be the index set.

Definition 2.4 (equivalence between simplices, saturation). For Δ, Δ' simplices of X , we write $\Delta \sim \Delta'$ to mean $\text{Lk}(\Delta) = \text{Lk}(\Delta')$. We denote the \sim -equivalence class of Δ by $[\Delta]$. Let $\text{Sat}(\Delta)$ denote the set of vertices $v \in X$ for which there exists a simplex Δ' of X such that $v \in \Delta'$ and $\Delta' \sim \Delta$, i.e.,

$$\text{Sat}(\Delta) = \left(\bigcup_{\Delta' \in [\Delta]} \Delta' \right)^{(0)}.$$

We denote by \mathfrak{S} the set of \sim -classes of *nonmaximal* simplices in X .

Next we introduce the candidate coordinate spaces:

Definition 2.5 (complement, link subgraph). Let \mathcal{W} be an X -graph. For each simplex Δ of X , we define Y_Δ to be the subgraph of $X^{+\mathcal{W}}$ induced by the set $(X^{+\mathcal{W}})^{(0)} - \text{Sat}(\Delta)$ of vertices.

Let $\mathcal{C}(\Delta)$ be the induced subgraph of Y_Δ spanned by $\text{Lk}(\Delta)^{(0)}$. Note that $\mathcal{C}(\Delta) = \mathcal{C}(\Delta')$ whenever $\Delta \sim \Delta'$. (We emphasise that we are taking links in X , not in $X^{+\mathcal{W}}$, and then considering the subgraphs of Y_Δ induced by those links.)

The following is the equivalent of the finite complexity Axiom (5) in the combinatorial framework:

Definition 2.6 (finite complexity). The simplicial complex X has *finite complexity* if there exists $n \in \mathbb{N}$ so that any chain $\text{Lk}(\Delta_1) \subsetneq \cdots \subsetneq \text{Lk}(\Delta_i)$, where each Δ_j is a simplex of X , has length at most n ; the minimal such n is the *complexity* of X .

The following is the main definition from [BHMS24]:

Definition 2.7 (combinatorial HHS). A *combinatorial HHS* (X, \mathcal{W}) consists of a simplicial graph X and an X -graph \mathcal{W} satisfying the following conditions:

- (1) X has complexity $n < +\infty$, as in Definition 2.6;
- (2) There is a constant δ so that for each nonmaximal simplex Δ , the subgraph $\mathcal{C}(\Delta)$ is δ -hyperbolic and (δ, δ) -quasi-isometrically embedded in Y_Δ , where Y_Δ is as in Definition 2.5;
- (3) Whenever Δ and Σ are nonmaximal simplices for which there exists a nonmaximal simplex Γ such that $\text{Lk}(\Gamma) \subseteq \text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$, and $\text{diam}(\mathcal{C}(\Gamma)) \geq \delta$, then there exists a simplex Π which extends Σ such that $\text{Lk}(\Pi) \subseteq \text{Lk}(\Delta)$, and all Γ as above satisfy $\text{Lk}(\Gamma) \subseteq \text{Lk}(\Pi)$;
- (4) If v, w are distinct nonadjacent vertices of $\text{Lk}(\Delta)$, for some simplex Δ of X , contained in \mathcal{W} -adjacent maximal simplices, then they are contained in \mathcal{W} -adjacent simplices of the form $\Delta \star \Sigma$.

In order to complete the HHS structure on \mathcal{W} we are left to define nesting and orthogonality relations on \mathfrak{S} , and projections between coordinate spaces.

Definition 2.8 (nesting, orthogonality, transversality, complexity). Let X be a simplicial graph. Let Δ, Δ' be nonmaximal simplices of X . Then:

- $[\Delta] \sqsubseteq [\Delta']$ if $\text{Lk}(\Delta) \subseteq \text{Lk}(\Delta')$;
- $[\Delta] \perp [\Delta']$ if $\text{Lk}(\Delta') \subseteq \text{Lk}(\text{Lk}(\Delta))$.

If $[\Delta]$ and $[\Delta']$ are neither \perp -related nor \sqsubseteq -related, we write $[\Delta] \pitchfork [\Delta']$.

Note that $[\emptyset]$ is the unique \sqsubseteq -maximal \sim -class of simplices in X and that \sqsubseteq is a partial ordering on the set of \sim -classes of simplices in X . Notice that the simplicial

graph X has finite complexity, in the sense of Definition 2.6, if there exists $n \in \mathbb{N}$ so that any \sqsubseteq -chain has length at most n ; the minimal such n is the complexity of X .

Remark 2.9. The definition of \perp says that any vertex in the link of Δ' is connected to any vertex in the link of Δ .

One might be tempted to think of nesting as being equivalent to inclusion of simplices, but this only works in one direction, namely:

Remark 2.10. Let Δ, Δ' be simplices of X . If $\Delta \subseteq \Delta'$, then $[\Delta'] \sqsubseteq [\Delta]$.

Definition 2.7(3) can be rephrased as follows:

- (4) Whenever Δ and Σ are nonmaximal simplices for which there exists a non-maximal simplex Γ such that $[\Gamma] \sqsubseteq [\Delta]$, $[\Gamma] \sqsubseteq [\Sigma]$, and $\text{diam}(\mathcal{C}(\Gamma)) \geq \delta$, then there exists a simplex Π such that:
- Π extends Σ (in particular $[\Pi] \sqsubseteq [\Sigma]$);
 - $[\Pi] \sqsubseteq [\Delta]$;
 - all $[\Gamma]$ as above satisfy $[\Gamma] \sqsubseteq [\Pi]$.

Our next goal is to define projections from \mathcal{W} to $\mathcal{C}([\Delta])$ for $[\Delta] \in \mathfrak{S}$.

Definition 2.11 (projections). Let (X, W, δ, n) be a combinatorial HHS.

Fix $[\Delta] \in \mathfrak{S}$ and define a map $\pi_{[\Delta]} : W \rightarrow 2^{\mathcal{C}([\Delta])}$ as follows. Let

$$p : Y_\Delta \rightarrow 2^{\mathcal{C}([\Delta])}$$

be the coarse closest point projection, i.e.

$$p(x) = \{y \in \mathcal{C}([\Delta]) : d_{Y_\Delta}(x, y) \leq d_{Y_\Delta}(x, \mathcal{C}([\Delta])) + 1\}.$$

As explained in [BHMS24, Remark 1.17], p is a coarsely Lipschitz, coarse map. Roughly, this is because either $\text{diam}(\mathcal{C}([\Delta])) \leq \delta$, and the conclusion is immediate, or Y_Δ is uniformly hyperbolic, and therefore p is the coarse retraction onto the quasiconvex subset $\mathcal{C}([\Delta])$.

Now suppose that w is a vertex of \mathcal{W} , so w corresponds to a unique simplex Σ_w of X . By [BHMS24, Lemma 1.15], the intersection $\Sigma_w \cap Y_\Delta$ is nonempty and has diameter at most 1. Define

$$\pi_{[\Delta]}(w) = p(\Sigma_w \cap Y_\Delta).$$

We have thus defined $\pi_{[\Delta]} : W^{(0)} \rightarrow 2^{\mathcal{C}([\Delta])}$. If $v, w \in W$ are joined by an edge e of \mathcal{W} , then Σ_v, Σ_w are joined by edges in $X^{+\mathcal{W}}$, and we let

$$\pi_{[\Delta]}(e) = \pi_{[\Delta]}(v) \cup \pi_{[\Delta]}(w).$$

Now let $[\Delta], [\Delta'] \in \mathfrak{S}$ satisfy $[\Delta] \pitchfork [\Delta']$ or $[\Delta'] \not\sqsubseteq [\Delta]$. Let

$$\rho_{[\Delta]}^{[\Delta']} = p(\text{Sat}(\Delta') \cap Y_\Delta),$$

where $p : Y_\Delta \rightarrow \mathcal{C}([\Delta])$ is coarse closest-point projection.

Let $[\Delta] \not\sqsubseteq [\Delta']$. Let $\rho_{[\Delta]}^{[\Delta']} : \mathcal{C}([\Delta']) \rightarrow \mathcal{C}([\Delta])$ be defined as follows. On $\mathcal{C}([\Delta']) \cap Y_\Delta$, it is the restriction of p to $\mathcal{C}([\Delta']) \cap Y_\Delta$. Otherwise, it takes the value \emptyset .

We are finally ready to state the main theorem of [BHMS24]:

Theorem 2.12 (HHS structures for X -graphs). *Let (X, W) be a combinatorial HHS. Let \mathfrak{S} be as in Definition 2.4, define nesting and orthogonality relations on \mathfrak{S} as in Definition 2.8, let the associated hyperbolic spaces be as in Definition 2.7, and define projections as in Definition 2.11.*

Then (W, \mathfrak{S}) is a hierarchically hyperbolic space, and the HHS constants only depend on δ, n as in Definition 2.7.

The aim of the present paper is, morally, to establish a “converse” of the previous result, by showing that any HHS satisfying reasonable hypotheses has a hierarchically hyperbolic structure that comes from a combinatorial HHS.

3. Combinatorial hyperbolicity from hierarchical hyperbolicity

Fix a hierarchically hyperbolic space $(\mathcal{Z}, \mathfrak{S})$. The goal of this section is to construct a combinatorial HHS structure (X, \mathcal{Z}) for the space \mathcal{Z} . The exact statement is Theorem 3.16, which will require the additional mild assumptions on $(\mathcal{Z}, \mathfrak{S})$ that we now present.

3.1. (Weak) wedges. The following property was first articulated in [BR20]. It is a fairly natural requirement, satisfied by all reasonable naturally occurring examples.

Property 3.1 (wedges). The HHS $(\mathcal{Z}, \mathfrak{S})$ has *wedges* if for all $U, V \in \mathfrak{S}$, one of the following holds:

- there exists a unique \sqsubseteq -maximal $T \in \mathfrak{S}$ such that $T \sqsubseteq U$ and $T \sqsubseteq V$, and we write $T = U \wedge V$;
- there does not exist $T \in \mathfrak{S}$ with $T \sqsubseteq U$ and $T \sqsubseteq V$, and we formally write $U \wedge V = \emptyset$.

What we will actually need is the following weak version of the wedge property:

Property 3.2 (weak wedges). The HHS $(\mathcal{Z}, \mathfrak{S})$ has *weak wedges* if for all $U, V \in \mathfrak{S}$, one of the following holds:

- (1) there exists a $T \in \mathfrak{S}$ such that $T \sqsubseteq U$, $T \sqsubseteq V$ and whenever $W \in \mathfrak{S}$ is a \sqsubseteq -minimal domain that is nested in both U and V then $W \sqsubseteq T$;
- (2) there does not exist $T \in \mathfrak{S}$ with $T \sqsubseteq U$ and $T \sqsubseteq V$.

Remark 3.3. If $(\mathcal{Z}, \mathfrak{S})$ has weak wedges and $U, V \in \mathfrak{S}$ share a common nested domain then we can find a unique \bar{T} satisfying the properties of item (1) which is

\sqsubseteq -minimal among all domains with the same properties. To do so, let $\mathfrak{T} = \{T_i\}_{i \in I}$ be the family of domains which are nested in both U and V and that contain every \sqsubseteq -minimal domain W which is nested in both U and V . For any two $T_i, T_j \in \mathfrak{T}$ there exists some $T \sqsubseteq T_i, T \sqsubseteq T_j$ which satisfies the properties in Property 3.2 for T_i and T_j . Hence T is again an element of \mathfrak{T} , since it must contain all \sqsubseteq -minimal domains which are nested in both T_i and T_j (and therefore in both U and V). Now, there must be an element $\bar{T} \in \mathfrak{T}$ which is nested in all elements of the family, because otherwise, by the previous observation, we could find an infinite chain $T_1 \not\sqsupseteq T_2 \not\sqsupseteq \dots$, which would contradict the finite complexity of the HHS. Hence, we say that \bar{T} is *the* weak wedge of U and V , and denote it by $U \wedge_{\min} V$.

3.2. Clean containers. The second main property was first articulated in [ABD21], and is still very natural.

Property 3.4 (clean containers). The HHS $(\mathcal{Z}, \mathfrak{S})$ has *clean containers* if the following holds. Let $T \in \mathfrak{S}$. Suppose that $U \not\sqsupseteq T$ and

$$A = \{V \in \mathfrak{S} : V \not\sqsupseteq T, V \perp U\} \neq \emptyset.$$

Then there exists $U_T^\perp \not\sqsupseteq T$, which we call the *orthogonal complement* of U inside T , such that $U_T^\perp \perp U$ and $V \sqsubseteq U_T^\perp$ for each $V \in A$.

In other words, the clean container property states that there exists a unique container for U inside T , as in Definition (3), and it is actually orthogonal to U . When the ambient domain T coincides with S we simply write $U^\perp := U_S^\perp$.

3.3. Orthogonals for nonsplit domains. For the next property, which is the first real requirement on $(\mathcal{Z}, \mathfrak{S})$, we first recall a definition from [BHS17a]:

Definition 3.5 (friendly). Let $V, W \in \mathfrak{S}$. Then W is *friendly* to V if $W \sqsubseteq V$ or $W \perp V$.

Notice that, as often happens in life, friendship is not always a symmetric relation.

Definition 3.6 (split). A domain $U \in \mathfrak{S}$ is *split* if there exists a \sqsubseteq -minimal domain $W \sqsubseteq U$ such that, for every $V \sqsubseteq U$, we have $W \sqsubseteq V$ or $W \perp V$. We say that W is a *Samaritan* for U , since it is friendly to every other $V \sqsubseteq U$.

An example of a split domain in the usual HHS structure of the mapping class group is as follows. Let U be a subsurface given by the disconnected union of an annulus W and another subsurface. Then W is a Samaritan for U , since any subsurface $V \sqsubseteq U$ which does not contain W must be nested in $U - W$. We postpone the details to Section 9.

Remark 3.7. Notice that, if W is a Samaritan for U and $W \sqsubseteq U' \sqsubseteq U$, then by definition U' is also split with Samaritan W .

Remark 3.8. If U is \sqsubseteq -minimal then it is trivially split, since it coincides with its unique Samaritan.

Property 3.9. A hierarchically hyperbolic space has *orthogonals for nonsplit domains* if for every two domains $U \not\sqsubseteq V$ either U is split or there exists $W \not\sqsubseteq V$ such that $W \perp U$.

This property feels like it should be related to the unbounded product property from [ABD21], but in practice the connection seems to not be very strong.

3.4. (Everywhere) dense product regions. Finally, for our construction to work, we must require that every coordinate space $\mathcal{C}U$ can be reconstructed from the projections coming from the domains nested inside U .

Property 3.10 (DPR). A hierarchically hyperbolic space $(\mathcal{Z}, \mathfrak{S})$ has *dense product regions* if there exists a constant M_0 such that, whenever $U \in \mathfrak{S}$ is not \sqsubseteq -minimal, for any $p \in \mathcal{C}U$ there exists $V \not\sqsubseteq U$ such that $d_U(p, \rho_U^V) \leq M_0$.

Notice that, up to choosing a bigger constant M_0 , we may always find a domain $V \not\sqsubseteq U$ as in the previous property which is also \sqsubseteq -minimal, since if $V \not\sqsubseteq V' \not\sqsubseteq U$ then $d_U(\rho_U^V, \rho_U^{V'}) \leq 10E$ by the consistency axiom (4).

Remark 3.11. In the study of HHS, it is often the case that one can remove minimal domains with bounded coordinate spaces. This is, however, not possible in our setting, since we might need these domains to witness DPR.

Actually, we will use a seemingly stronger, yet equivalent version of the DPR property, which we now state.

Property 3.12 (EDPR). A hierarchically hyperbolic space $(\mathcal{Z}, \mathfrak{S})$ has *everywhere dense product regions* if there exists a constant C_0 such that the following holds. For every $U \in \mathfrak{S}$ and every $x \in F_U$ there exists a maximal family $V_1, \dots, V_k \not\sqsubseteq U$ of \sqsubseteq -minimal, pairwise orthogonal domains such that x is C_0 -close to the relative product region $P_{\{V_i\}_{i=1,\dots,k}}^U$.

Clearly, Property 3.12 is stronger than Property 3.10. The converse also holds:

Lemma 3.13. *A HHS with the DPR (3.10) also has the EDPR (3.12).*

Proof. Let $(\mathcal{Z}, \mathfrak{S})$ be a HHS with the DPR property, with some constant M_0 , and let E be a HHS constant for $(\mathcal{Z}, \mathfrak{S})$. We shall prove the existence of a constant $C_0 = C_0(M_0, E)$ such that $(\mathcal{Z}, \mathfrak{S})$ also enjoys the EDPR property with respect to C_0 .

We will prove the lemma by induction on the *level* l of U , that is, the maximum k such that there exists a chain $U_0 \not\sqsubseteq \dots \not\sqsubseteq U_k = U$. If $l = 0$ then U is \sqsubseteq -minimal and the EDPR (3.12) clearly holds.

Now suppose the theorem holds for every domain of level strictly less than l , and let U be a domain of level l . Before going on with the proof we recall some definitions from [BHS19].

We will say that the collection \mathfrak{U} of elements of \mathfrak{S}_U is *totally orthogonal* if any pair of distinct elements of \mathfrak{U} are orthogonal. Given a totally orthogonal family \mathfrak{U} , we say that $W \sqsubseteq U$ is \mathfrak{U} -*generic* if there exists $V \in \mathfrak{U}$ so that W is not orthogonal to V .

Now fix $x \in F_U$ that we want to realise with minimal domains. A totally orthogonal collection \mathfrak{S} is C -*good* if any E -partial realisation point y for \mathfrak{U} , as defined in the partial realisation axiom (8), has the property that for each $W \sqsubseteq U$ we have $d_W(x_w, y_w) \leq C$. Notice that our goal is to find a maximal family \mathfrak{U} which is made of minimal supports and C -good for some uniform constant C . Notice moreover that, if \mathfrak{U} is C -good but not maximal, then we can add \sqsubseteq -minimal domains to \mathfrak{U} and complete it to a maximal totally orthogonal family. The latter will again be C -good, because a partial realisation point for the larger family is also a partial realisation point for \mathfrak{U} .

A totally orthogonal collection \mathfrak{U} is C -*generically good* if any E -partial realisation point y for \mathfrak{U} has the property that for each \mathfrak{U} -generic W we have $d_W(x_w, y_w) \leq E$.

We allow that \mathfrak{U} can be empty. In this case, we say that a C -partial realisation point for \emptyset is simply a point y such that $d_U(x_U, y_U) \leq C$. Notice that no W is \emptyset -generic.

Lemma 3.14. *For every $C \geq 100E^3$ the following holds. Let \mathfrak{U} be totally orthogonal and C -generically good but not C -good. Then there exists a totally orthogonal, $10C$ -generically good collection \mathfrak{U}' with $\mathfrak{U} \subsetneq \mathfrak{U}'$, obtained by adding \sqsubseteq -minimal domains.*

This fact is proven exactly as [BHS19, Lemma 3.3], whose proof runs verbatim in our case. The only difference is that our inductive hypothesis, which replaces that of [BHS19, Theorem 3.1], allows us to assume that the additional domains are all \sqsubseteq -minimal.

Now we can prove Lemma 3.13. Recall that we want to realise a point $x \in F_U$. If $\mathfrak{U} = \emptyset$ is already M_0 -good we can choose any maximal family of pairwise orthogonal, minimal domains $V_1, \dots, V_n \sqsubseteq U$ such that $d_U(x_U, \rho_U^{V_1}) \leq M_0$, whose existence is granted by the DPR (3.10). Then any realisation point y for $\{(V_i, x_{V_i})\}$ is also a realisation point for $\mathfrak{U} = \emptyset$, since $d_U(x_U, y_U) \leq M_0$ by construction, and therefore x and y are M_0 -close in every coordinate space.

Otherwise, since no W is \emptyset -generic, we can apply Lemma 3.14 and find a larger \mathfrak{U}_1 which is $10M_0$ -generically good. If \mathfrak{U}_1 is $10M_0$ -good we can complete it to a maximal family of pairwise orthogonal, \sqsubseteq -minimal domains which is again $10M_0$ -good, and we are done. Otherwise, we can repeat the process with \mathfrak{U}_1 . Since

there is a bound on the cardinality of totally orthogonal sets, in finitely many steps we necessarily get a good totally orthogonal set made of minimal supports, and this concludes the proof. \square

3.5. *The main theorem.*

Remark 3.15 (normalisation preserves our hypotheses). Before stating the main theorem we notice that, if $(\mathcal{Z}, \mathfrak{S})$ has one of the properties defined in the previous section, then so does the normalised structure, as defined in Remark 1.4. This is because the new structure has the same set of domains \mathfrak{S} with the same relations of nesting and orthogonality, thus all combinatorial assumptions on the domain set (wedges, clean containers and Property 3.9) are preserved under the normalisation procedure. The DPR (3.10) still holds as well, since the coarse closest point projection is coarsely Lipschitz. Hence, our Assumption 1.5 that the HHS structure is normalised does not lose any generality.

Theorem 3.16. *Let $(\mathcal{Z}, \mathfrak{S})$ be a normalised hierarchically hyperbolic space with weak wedges, clean containers, the orthogonals for nonsplit domains property (3.9) and the DPR (3.10). There exists a combinatorial HHS (X, \mathcal{W}) such that \mathcal{Z} is quasi-isometric to \mathcal{W} .*

Moreover, let G be a finitely generated group which acts on $(\mathcal{Z}, \mathfrak{S})$ by automorphisms. Then G acts on (X, \mathcal{W}) , and the quasi-isometry $f : \mathcal{W} \rightarrow \mathcal{Z}$ is coarsely G -equivariant.

Outline of the proof of Theorem 3.16. The graphs X and \mathcal{W} are constructed in Sections 4.2 and 4.3, respectively. The four conditions of Definition 2.7 are verified in Section 5, and more precisely:

- Condition 1 is Corollary 5.10.
- The two parts of Condition 2 are proved in Sections 5.4 and 5.5.
- Condition 3 is Lemma 5.7, which is implied by Lemma 5.8.
- Condition 4 is Lemma 5.11.

In Definition 4.11 we define a map $f : \mathcal{W} \rightarrow \mathcal{Z}$, which we prove to be a quasi-isometry in Lemma 5.19. Finally, the “moreover” part of the statement is proved in Theorem 6.2. \square

Remark 3.17. Towards relative HHS versions of our results we remark that the proof of Theorem 3.16 features \sqsubseteq -minimal domains only in Lemma 5.14 and Claim 5.17. The proof of the former shows that certain augmented links are quasi-isometric to the coordinate spaces of \sqsubseteq -minimal domains, with hyperbolicity of these only used to then conclude that those augmented links are hyperbolic. The proof of the latter does not use hyperbolicity of coordinate spaces of \sqsubseteq -minimal domains.

4. Construction of the combinatorial HHS

In this section we construct the pair (X, \mathcal{W}) and the map $f : \mathcal{W} \rightarrow \mathcal{Z}$.

Remark 4.1. The construction of (X, \mathcal{W}) and f will only require $(\mathcal{Z}, \mathfrak{S})$ to have clean containers. The other hypotheses of Theorem 3.16 will be needed later, to ensure that (X, \mathcal{W}) is actually a combinatorial HHS and that f is a quasi-isometry.

4.1. Moral compass. Before going into the actual details, we explain the idea of the construction, and why it should work (at least morally).

First, we consider the graph \bar{X} whose vertices are all \sqsubseteq -minimal domains of \mathfrak{S} (see Definition 4.2). Now, whenever $U \in \mathfrak{S}$ is not \sqsubseteq -minimal, its coordinate space $\mathcal{C}U$ can be reconstructed by just looking at the projections ρ_U^V coming from the \sqsubseteq -minimal domains, by the dense product regions property (3.10), and such projections are close whenever the \sqsubseteq -minimal domains are orthogonal, by Lemma 1.3. Hence, in a sense, the graph \bar{X} will contain all information about the HHS structure coming from the non- \sqsubseteq -minimal domains.

However, \bar{X} does not see the coordinate spaces of \sqsubseteq -minimal domains. Therefore, for every vertex V of \bar{X} we consider its coordinate space $\mathcal{C}V$, which we may assume to be a simplicial graph up to quasi-isometry (see for example [CdH16, Lemma 3.B.6]), and we replace V with the cone over the 0-skeleton $\mathcal{C}V^{(0)}$. This way, the apex of the cone, call it v_V , will have $\mathcal{C}V^{(0)}$ inside its link, and after adding the right \mathcal{W} -edges we will be able to see $\mathcal{C}V$ as the augmented link of some simplex. Call X the graph obtained after this “blow-up” procedure (again, see Definition 4.2).

Hence, a vertex of \mathcal{W} , which corresponds to a maximal simplex of X , is the data of a collection V_1, \dots, V_k of pairwise orthogonal and \sqsubseteq -minimal domains, and a point $x_i \in \mathcal{C}V_i$ for all $i = 1, \dots, k$. Such a collection $\{(V_i, x_i)\}$ admits a unique realisation point, in the sense of Theorem 1.7, thus we can define a map f between vertices of \mathcal{W} and points in \mathcal{Z} (see Definition 4.11).

As to the edges of \mathcal{W} , morally we would like to say that, if $\Sigma = \{(V_i, x_i)\}_{i=1, \dots, k}$ and $\Delta = \{(W_j, x_j)\}_{j=1, \dots, l}$ are two maximal simplices, then they are \mathcal{W} -adjacent if and only if their realisation points are close, so that f is a quasi-isometry almost by definition. In turn, such realisation points are close if and only if their coordinates are close in every coordinate space. For some technical reasons (mainly appearing in the proof of Lemma 5.11), the exact definition of the edges of \mathcal{W} will also take into account the supports $\{V_i\}_{i=1, \dots, k}$ and $\{W_j\}_{j=1, \dots, l}$ of the two simplices: the more these supports intersect, the further we allow the coordinates of the realisation points to be (see Definition 4.10).

4.2. The minimal orthogonality graph and its blow-up. The first step in the proof of Theorem 3.16 is to construct the simplicial complex X , which will heavily depend




Figure 1. U and V are orthogonal, \sqsubseteq -minimal domains. Therefore, the cone over $(CU)^{(0)}$ (here, in red) and the cone over $(CV)^{(0)}$ (in blue) form a join. Notice that any two points in $(CU)^{(0)}$ are *not* adjacent.

on both \mathcal{Z} and the actual HHS structure. For the purpose of this subsection we do not need to assume any property on $(\mathcal{Z}, \mathfrak{S})$.

Let \mathfrak{S}_{\min} be the set of \sqsubseteq -minimal elements of \mathfrak{S} . Let $\bar{X}^{(1)}$ be the graph with vertex-set \mathfrak{S}_{\min} , with $U, V \in \mathfrak{S}_{\min}$ joined by an edge when $U \perp V$. Let \bar{X} be the flag complex on $\bar{X}^{(1)}$.

For each $U \in \mathfrak{S}_{\min}$ we can assume, up to quasi-isometry, that $\mathcal{C}U$ is a graph. Thus let $L(U)$ be the cone on $(\mathcal{C}U)^{(0)}$, and denote by v_U the cone-vertex. Let $X^{(1)}$ be the graph formed from $\bigsqcup_{U \in \mathfrak{S}_{\min}} L(U)$ by joining each vertex of $L(U)$ to each vertex of $L(V)$ whenever $U \perp V$ (i.e., whenever U, V are adjacent in $\bar{X}^{(1)}$). Let X be the flag complex on $X^{(1)}$.

Definition 4.2 (blow-up). Define the retraction $p : X \rightarrow \bar{X}$ by collapsing each subcomplex $L(U)$ to the vertex U . We will refer to \bar{X} as the *minimal orthogonality graph* of the structure $(\mathcal{Z}, \mathfrak{S})$, and to X as the *blow-up* of \bar{X} .

For each simplex Δ of X , let $\bar{\Delta} = p(\Delta)$ be the image simplex in \bar{X} . We will say that Δ is *supported* in $\bar{\Delta}$.

Given a simplex Δ of X and a vertex $U \in \bar{\Delta}$, let $\Delta_U = \Delta \cap p^{-1}(U)$. Note that Δ_U is either a vertex of Δ or an edge of Δ . Moreover, we have

$$\Delta = \star_{U \in \bar{\Delta}^{(0)}} \Delta_U.$$

A careful inspection of the construction yields the following (compare with [HMS22, Lemma 4.12]):

Lemma 4.3 (decomposition of links). *Let Δ be a simplex of X . Then*

$$\text{Lk}(\Delta) = p^{-1}(\text{Lk}_{\bar{X}}(\bar{\Delta})) \star (\star_{U \in \bar{\Delta}^{(0)}} \text{Lk}_{p^{-1}(U)}(\Delta_U)).$$

Corollary 4.4. *Let Δ be a simplex of X . Then one of the following holds:*

- (1) $\text{Lk}(\Delta)$ is either a single vertex or a nontrivial join.
- (2) For each $U \in \bar{\Delta}^{(0)}$, we have that Δ_U is an edge.




Figure 2. All cases in which $\text{Lk}(\Delta)$ (represented by the dashed ellipses) is not a nontrivial join.

- (3) *The simplex $\bar{\Delta}$ is maximal inside \bar{X} and the following holds. There exists $U \in \bar{\Delta}^{(0)}$ such that $\Delta_U = v_U$, and for every $V \in \bar{\Delta}^{(0)} - \{U\}$ we have that Δ_V is an edge.*

Proof. If at least two of the terms of the join from Lemma 4.3 are nonempty then $\text{Lk}(\Delta)$ is a nontrivial join, and item (1) holds. Thus suppose that exactly one of the terms of the join is nonempty, so that $\text{Lk}(\Delta)$ coincides with that term.

If $p^{-1}(\text{Lk}_{\bar{X}}(\bar{\Delta})) \neq \emptyset$ then for every $U \in \bar{\Delta}^{(0)}$ we have $\text{Lk}_{p^{-1}(U)}(\Delta_U) = \emptyset$. Hence Δ_U is an edge, and item (2) is satisfied (see the picture on the left in Figure 2).

Otherwise, suppose that $p^{-1}(\text{Lk}_{\bar{X}}(\bar{\Delta})) = \emptyset$ (that is, that $\bar{\Delta}$ is a maximal simplex) and that $\text{Lk}_{p^{-1}(U)}(\Delta_U) = \emptyset$ for all domains except one, call it U_0 . In particular, we have

$$\text{Lk}(\Delta) = \text{Lk}_{p^{-1}(U_0)}(\Delta_{U_0}).$$

If $\Delta_{U_0} = v_{U_0}$ then we are in the case of item (3) (see the central picture in Figure 2). Otherwise Δ_{U_0} is a point in $(CU)^{(0)}$, and therefore $\text{Lk}_{p^{-1}(U_0)}(\Delta_{U_0}) = v_{U_0}$ is a single point, thus again giving item (1) (see the picture on the right in Figure 2). \square

Remark 4.5. By definition, the maximum cardinality of a simplex in X is twice the maximum cardinality of a family of \sqsubseteq -minimal, pairwise orthogonal domains, which in turn is bounded above by the *complexity* of the HHS structure by [BHS19, Lemma 2.1]. Thus X has finite dimension.

4.3. Edges between maximal simplices. Now we define the graph \mathcal{W} for the combinatorial HHS (X, \mathcal{W}) . The construction of this subsection will only require $(\mathcal{Z}, \mathfrak{S})$ to have clean containers.

Let $\mathfrak{M}(X)$ be the set of maximal simplices of X . The vertex set of the graph \mathcal{W} is $\mathfrak{M}(X)$. Now, given $\sigma \in \mathfrak{M}(X)$, let $U_1, \dots, U_n \in \mathfrak{S}_{\min}$ be the vertices of $p(\sigma)$ (a maximal collection of pairwise orthogonal elements), so that $\sigma = \star_{i=1}^n \{v_{U_i}, x_i\}$, where $x_i \in CU_i$ and $\{v_{U_i}, x_i\}$ is the edge of X joining v_{U_i} to x_i . We will refer to the domains $U_1, \dots, U_n \in \mathfrak{S}_{\min}$ as the *support* of σ , and to the points $x_i \in CU_i$ as the *coordinates* prescribed by σ .

Definition 4.6. Let $\sigma = \star_{i=1}^n \{v_{U_i}, x_i\}$ be a maximal simplex of X . Define the tuple $b(\sigma) = (b_V)_{V \in \mathfrak{G}} \in \prod_{V \in \mathfrak{G}} 2^{\mathcal{C}V}$ as follows. Let $V \in \mathfrak{G}$. Note that V is not properly nested in any U_i , since each $U_i \in \mathfrak{G}_{\min}$. Moreover, V cannot be orthogonal to all of the U_i ; otherwise, we could choose $V' \sqsubseteq V$ with $V' \in \mathfrak{G}_{\min}$ and observe that $V' \perp U_i$ for all i , contradicting maximality of σ . So, for some $i \leq n$, we have either $U_i \pitchfork V$ or $U_i \not\sqsubseteq V$ or $U_i = V$. In particular, if $V \notin \{U_1, \dots, U_n\}$, then for some i we have a set $\rho_V^{U_i} \subset \mathcal{C}V$ of diameter at most E . Moreover, by Lemma 1.3 and the fact that $U_i \perp U_j$ for all $i \neq j$, we have $\text{diam}(\bigcup_{i \leq n, U_i \not\sqsubseteq V} \rho_V^{U_i}) \leq 10E$.

Then:

- If $V = U_i$ for some $i \leq n$ set $b_V = x_i$. This is why we called x_i the coordinate prescribed by σ .
- Otherwise, let $b_V \in \mathcal{C}V$ be $b_V = \bigcup_{i \leq n, U_i \not\sqsubseteq V} \rho_V^{U_i}$, which is a nonempty set of diameter at most $10E$.

Definition 4.7 (colevel). Let $U \in \mathfrak{G}$. We define the *colevel* of U , denoted by $\text{co-lv}(U)$, as the maximum k such that there exists a chain of the form $U = U_0 \not\sqsubseteq \dots \not\sqsubseteq U_k = S$.

Notice that if $U \not\sqsubseteq V$ then $\text{co-lv}(U) \geq \text{co-lv}(V)$. Moreover, the maximum colevel is $n - 1$, where n is the complexity of the HHS structure.

Definition 4.8 (orthogonal complement). Let $\bar{\Delta} = \{U_1, \dots, U_k\}$ be a nonempty simplex inside \bar{X} and let V be a domain such that $U_i \sqsubseteq V$ for all i . If there exists $T \not\sqsubseteq V$ that is orthogonal to all U_i then the *orthogonal complement* of $\bar{\Delta}$ inside V , which we will denote as $\bar{\Delta}_V^\perp$, is constructed inductively as follows:

- set $\{U_1\}_V^\perp$ as the orthogonal complement of U_1 inside V , as in Property 3.4;
- If $\{U_1, \dots, U_i\}_V^\perp$ has already been defined, then $\{U_1, \dots, U_{i+1}\}_V^\perp$ is the orthogonal complement of U_{i+1} inside $\{U_1, \dots, U_i\}_V^\perp$.

If $V = S$ we denote the orthogonal complement of $\bar{\Delta}$ in the maximal domain simply as $\bar{\Delta}^\perp$.

Notice that the definition is independent of the order of the vertices, because by construction $\bar{\Delta}_V^\perp$ is also the unique \sqsubseteq -maximal element T which is nested inside V and is orthogonal to all vertices of $\bar{\Delta}$.

Notation 4.9. With an innocent abuse of notation, we could say that:

- the orthogonal complement of the empty simplex is the \sqsubseteq -maximal element S , and $\text{co-lv}(S) = 0$;
- the orthogonal complement of a maximal simplex is empty, and $\text{co-lv}(\emptyset) = n$.

Now we can finally define the edges of \mathcal{W} .

Definition 4.10 (\mathcal{W} -edges). Let Σ, Δ be two maximal simplices of X and let $\bar{\Sigma}, \bar{\Delta}$ be their supports. Let $b(\Sigma) = (b_U)_{U \in \mathfrak{G}}$ and $b(\Delta) = (c_U)_{U \in \mathfrak{G}}$. Let n be the complexity of the HHS structure. Let $\lambda \geq 0$ be some constant.

Let $W = (\bar{\Sigma} \cap \bar{\Delta})^\perp$ be the orthogonal complement of the intersection, and let k be the clevel of W (with the Notation 4.9 for the exceptional cases). Then Σ and Δ are \mathcal{W} -adjacent if and only if, for every $U \in \mathfrak{G}$,

$$d_{c_U}(b_U, c_U) \leq (k + 1)\lambda.$$

In other words, the more two adjacent simplices share their supports, the further away we let them be in the fewer and fewer domains where b_U and c_U may actually differ. Notice that the definition depends on the constant λ , which we will later choose to be large enough.

4.4. The realisation map. Finally, we define the map $f : \mathcal{W} \rightarrow \mathcal{Z}$ that will be the required quasi-isometry.

Definition 4.11 (realisation map). For every simplex $\sigma = \star_{i=1}^n \{v_{U_i}, x_i\}$, the partial realisation axiom (8) provides the existence of a *realisation point* z such that:

- $d_{U_i}(x_i, \pi_{U_i}(z)) \leq E$;
- for every $V \in \mathfrak{G}$ such that $U_i \not\sqsubset V$ or $U_i \pitchfork V$ for some $i \leq n$ we have $d_V(\rho_V^{U_i}, \pi_V(z)) \leq E$.

In other words, the coordinates of z are E -close to the tuple $b(\sigma)$ from Definition 4.6, and we say that z *realises* $b(\sigma)$. By the uniqueness axiom (9), z is uniformly coarsely unique, and the bound only depends on E . Hence, setting $f(\sigma) = z$ gives a well-defined coarse map $f : \mathcal{W} \rightarrow \mathcal{Z}$.

Remark 4.12 (consistency of $b(\sigma)$). The existence of a z that realises $b(\sigma)$ also shows that the latter is a $20E$ -consistent tuple. Indeed, every coordinate of $b(\sigma)$ has diameter at most $10E$ and is E -close to the corresponding coordinates of z , which satisfies the consistency axiom (4).

5. Proof of the main theorem

Unless otherwise stated, in this section we will work under the following assumption:

Assumption 5.1. $(\mathcal{Z}, \mathfrak{G})$ is a normalised hierarchically hyperbolic space with weak wedges (3.2), clean containers (3.4), orthogonals for nonsplit domains (3.9) and the DPR property (we will work with its strong form, which is Property 3.12). Let E be the HHS constant for $(\mathcal{Z}, \mathfrak{G})$, and for every $U \in \mathfrak{G}$ let F_U be the space of $20E$ -consistent partial tuples. Let X be the graph from Section 4.2, and let \mathcal{W} be the graph from Section 4.3, whose edges depend on the constant λ from Definition 4.10,

that we will later choose (see Lemma 5.11, Claim 5.15 and Lemma 5.19). Finally, let $f : \mathcal{W} \rightarrow \mathcal{Z}$ be the realisation map from Definition 4.11.

We choose F_U to be the space of $20E$ -consistent tuples, in order to include the tuples of the form $b(\sigma)$ from Definition 4.6.

Our first goal is to prove the following, which readily implies the first half of Theorem 3.16:

Theorem 5.2. *Under Assumption 5.1 there exists $\tilde{\lambda} \geq 0$ such that, whenever $\lambda \geq \tilde{\lambda}$, the pair (X, \mathcal{W}) is a combinatorial HHS, and f is a quasi-isometry.*

5.1. Weak orthogonal complements. Before getting into the proof, we develop some more technical notation and lemmas regarding the interaction between weak wedges and clean containers.

Definition 5.3 (weak orthogonal complement). Let $\bar{\Delta} = \{U_1, \dots, U_k\}$ be a simplex inside \bar{X} . Then its *weak orthogonal complement* is

$$\bar{\Delta}_{\min}^{\perp} := \bar{\Delta}^{\perp} \wedge_{\min} \bar{\Delta}^{\perp}.$$

In other words, $\bar{\Delta}_{\min}^{\perp}$ is the unique domain $T \in \mathfrak{S}$ such that:

- $T \perp U$ for every $U \in \bar{\Delta}$;
- T contains every $V \in \mathfrak{S}_{\min}$ which is orthogonal to $\bar{\Delta}$ (that is, every $V \in \text{Lk}_{\bar{X}}(\bar{\Delta})$);
- T is \sqsubseteq -minimal among domains with the previous properties.

The weak orthogonal complement is uniquely determined by the link of $\bar{\Delta}$, in the following sense:

Lemma 5.4. *Let $\bar{\Delta}, \bar{\Sigma}$ be two simplices inside \bar{X} . Then the following are equivalent:*

- (1) $\text{Lk}(\bar{\Delta}) \subseteq \text{Lk}(\bar{\Sigma})$;
- (2) $\bar{\Delta}_{\min}^{\perp} \sqsubseteq \bar{\Sigma}_{\min}^{\perp}$.

Proof. (1 \Rightarrow 2) Let $T := \bar{\Delta}_{\min}^{\perp} \wedge_{\min} \bar{\Sigma}_{\min}^{\perp} \sqsubseteq \bar{\Sigma}_{\min}^{\perp}$. Since $T \sqsubseteq \bar{\Delta}_{\min}^{\perp}$, we have that $T \perp U$ for every $U \in \bar{\Delta}$. Moreover, by definition of weak wedge, T contains every $V \in \mathfrak{S}_{\min}$ which is nested inside both $\bar{\Delta}_{\min}^{\perp}$ and $\bar{\Sigma}_{\min}^{\perp}$, that is, every $V \in \text{Lk}_{\bar{X}}(\bar{\Delta}) \cap \text{Lk}_{\bar{X}}(\bar{\Sigma}) = \text{Lk}_{\bar{X}}(\bar{\Delta})$. Therefore, by minimality of $\bar{\Delta}_{\min}^{\perp}$, we must have $T = \bar{\Delta}_{\min}^{\perp}$, hence $\bar{\Delta}_{\min}^{\perp} \sqsubseteq \bar{\Sigma}_{\min}^{\perp}$.

(2 \Rightarrow 1) If V is \sqsubseteq -minimal, then $V \in \text{Lk}(\bar{\Delta})$ if and only if $V \sqsubseteq \bar{\Delta}_{\min}^{\perp}$, by definition of $\bar{\Delta}_{\min}^{\perp}$, and the same holds for $\bar{\Sigma}$. \square

Corollary 5.5. *Let $\bar{\Delta}, \bar{\Sigma}$ be two simplices inside \bar{X} . Then $\text{Lk}(\bar{\Delta}) = \text{Lk}(\bar{\Sigma})$ if and only if $\bar{\Delta}_{\min}^{\perp} = \bar{\Sigma}_{\min}^{\perp}$.*

Lemma 5.6. *Let $\bar{\Delta}$ be a simplex of \bar{X} . Then exactly one of the following holds:*

- (1) $\bar{\Delta}_{\min}^\perp = \bar{\Delta}^\perp$;
- (2) $\bar{\Delta}_{\min}^\perp$ is split.

Moreover, in the second case $\text{Lk}(\bar{\Delta})$ is the cone with cone point $V \in \text{Lk}(\bar{\Delta})$, where V is any Samaritan for $\bar{\Delta}_{\min}^\perp$.

Proof. Clearly $\bar{\Delta}_{\min}^\perp \sqsubseteq \bar{\Delta}^\perp$, and if they do not coincide then the orthogonals for nonsplit domains property (3.9) states that either $\bar{\Delta}_{\min}^\perp$ is split or there exists $V \sqsubseteq \bar{\Delta}^\perp$ such that $V \perp \bar{\Delta}_{\min}^\perp$. But then there exists some \sqsubseteq -minimal domain $V' \sqsubseteq V$ which lies in $\bar{\Delta}^\perp$ but not in $\bar{\Delta}_{\min}^\perp$, contradicting the definition of the latter.

For the “moreover” part, just notice that, if V is a Samaritan for $\bar{\Delta}_{\min}^\perp$, then V is orthogonal to any other vertex of $\text{Lk}(\bar{\Delta})$. □

5.2. Intersection of links and finite complexity. Now we turn to the proof of Theorem 5.2. First, we check the parts of Definition 2.7 that depend on X only.

Lemma 5.7 (verification of Definition 2.7(3)). *Let Σ, Δ be nonmaximal simplices of X and suppose that there exists a nonmaximal simplex Γ such that $[\Gamma] \sqsubseteq [\Sigma]$, $[\Gamma] \sqsubseteq [\Delta]$ and $\text{diam}(\mathcal{C}([\Gamma])) \geq 3$. Then there exists a nonmaximal simplex Π which extends Σ such that $[\Pi] \sqsubseteq [\Delta]$ and all Γ as above satisfy $[\Gamma] \sqsubseteq [\Pi]$.*

Arguing as in the proof of [BHMS24, Theorem 6.4] (at the beginning of the paragraph titled “ (X, W) is a combinatorial HHS”), one sees that Lemma 5.7 is implied by the following, which is [BHMS24, Condition 6.4.B] there:

Lemma 5.8. *Under Assumption 5.1, let Σ, Δ be nonmaximal simplices of X . Then there exist two (possibly empty or maximal) simplices $\Pi, \Psi \subset X$ such that $\Sigma \subset \Pi$ and*

$$\text{Lk}(\Sigma) \cap \text{Lk}(\Delta) = \text{Lk}(\Pi) \star \Psi.$$

Proof of Lemma 5.8. We subdivide the proof into two major steps.

Finding the support of the extended simplex: Let $\bar{\Sigma}$ and $\bar{\Delta}$ be the supports of Σ, Δ , respectively, and let $\bar{\Sigma}^\perp, \bar{\Delta}^\perp \in \mathfrak{G}$ be their orthogonal complements. Let $\bar{\Phi} = \bar{\Delta} \cap \text{Lk}(\bar{\Sigma})$, and let Y_0 be the orthogonal complement of $\bar{\Phi}$ inside $\bar{\Sigma}^\perp$, that is, $Y_0 = (\bar{\Sigma} \star \bar{\Phi})^\perp$. Finally, set $\bar{\Theta} = \bar{\Psi} = \emptyset$. We will progressively add vertices to these simplices, which will form the supports of the simplices Π and Ψ we are looking for.

If $\bar{\Sigma}^\perp$ and $\bar{\Delta}^\perp$ have no common nested domain we formally set $W_0 = \emptyset$. Otherwise, let $W_0 = \bar{\Sigma}^\perp \wedge_{\min} \bar{\Delta}^\perp$ be the weak wedge of the orthogonal complements. Since by construction W_0 is nested in $\bar{\Delta}^\perp$, every vertex of $\bar{\Phi}$ is orthogonal to W_0 . Thus $W_0 \sqsubseteq Y_0$, since W_0 is also nested in $\bar{\Sigma}^\perp$ and Y_0 is a clean container.

Now we do the following procedure, which is divided into three parts.

Part 1: If $W_0 = \emptyset$ or W_0 is nonsplit then we can set $W' = W_0$ and $Y'_0 = Y_0$ and skip to Part 2. Otherwise, there exists a Samaritan $U_1 \sqsubseteq W_0$ such that $U_1 \perp V$ for every other \sqsubseteq -minimal domain $V \not\sqsubseteq W_0$, and we add U_1 to $\bar{\Psi}$. Then let $W_1 = \{U_1\}_{W_0}^\perp$ (which might be empty), and similarly let $Y_1 = \{U_1\}_{Y_0}^\perp$. Clearly $W_1 \sqsubseteq Y_1$, thus if W_1 is again split we can repeat this argument with Y_1 and W_1 . This procedure, which adds one vertex at a time to $\bar{\Psi}$, must end after at most n steps by the finite complexity axiom (5), since every new W_i is properly nested into W_{i-1} for all i . Moreover, this procedure stops when $W' = \bar{\Psi}_{W_0}^\perp$ is either empty or nonsplit. Set $Y'_0 = \bar{\Psi}_{Y_0}^\perp$ (which again might be empty).

Part 2: Now, if $W' = \emptyset$ we choose a simplex $\bar{\Theta} = \{V_1, \dots, V_k\}$ of pairwise orthogonal, \sqsubseteq -minimal domains inside Y'_0 , and we skip to Part 3. Otherwise $W' \sqsubseteq W_0$ is a nonsplit domain which is nested inside $Y'_0 \sqsubseteq Y_0$. If $W' = Y'_0$ we set $\bar{\Theta} = \emptyset$. Otherwise, by the orthogonals for nonsplit domains property (3.9) there exists a \sqsubseteq -minimal domain $V_1 \sqsubseteq Y'_0$ such that $V_1 \perp W'$. Then let $Y'_1 = \{V_1\}_{Y'_0}^\perp$, which contains W' and is properly nested into Y'_0 . Now we can iterate this construction with W' and Y'_1 , and the procedure has to stop after at most n steps since Y'_i is properly nested inside Y'_{i-1} for all i . Thus, in the end we find a simplex $\bar{\Theta} = \{V_1, \dots, V_k\}$ of \sqsubseteq -minimal and pairwise orthogonal domains, which are nested in Y'_0 and whose orthogonal complement inside Y'_0 is W' .

Part 3: Summing up, we have defined some (possibly formally empty) domains and two simplices $\bar{\Psi}, \bar{\Theta} \subset \bar{X}$ such that

$$\begin{cases} Y_0 = (\bar{\Sigma} \star \bar{\Phi})^\perp; \\ Y'_0 = \bar{\Psi}_{Y_0}^\perp; \\ W_0 = \bar{\Sigma}^\perp \wedge_{\min} \bar{\Delta}^\perp; \\ W' = \bar{\Psi}_{W_0}^\perp = \bar{\Theta}_{Y'_0}^\perp. \end{cases}$$

Hence

$$(\bar{\Sigma} \star \bar{\Phi} \star \bar{\Psi} \star \bar{\Theta})^\perp = W' \sqsubseteq W_0 = \bar{\Sigma}^\perp \wedge_{\min} \bar{\Delta}^\perp.$$

Since $\text{Lk}(\bar{\Sigma}) \cap \text{Lk}(\bar{\Delta})$ is the subgraph of X spanned by all \sqsubseteq -minimal domain that are nested inside W_0 , the previous nesting can be restated as

$$\text{Lk}(\bar{\Sigma}) \cap \text{Lk}(\bar{\Delta}) \supseteq \text{Lk}(\bar{\Sigma} \star \bar{\Phi} \star \bar{\Psi} \star \bar{\Theta}),$$

and since the domains of $\bar{\Psi}$ lie in W_0 by the construction from Part 1 we also have

$$\text{Lk}(\bar{\Sigma}) \cap \text{Lk}(\bar{\Delta}) \supseteq \text{Lk}(\bar{\Sigma} \star \bar{\Phi} \star \bar{\Psi} \star \bar{\Theta}) \star \bar{\Psi}.$$

The converse inclusion is also true, since, by the construction of Part 1, if a minimal domain V is nested in W_0 , then either V is one of the vertices of $\bar{\Psi}$ or

Π_U	$\bar{\Sigma}$	$\text{Lk}(\bar{\Sigma})$
$\bar{\Delta}$	extend Σ_U if needed	Δ_U for every $U \in \bar{\Phi}$
$\text{Lk}(\bar{\Delta})$	Σ_U	v_U for every $U \in \bar{\Psi}$
$\bar{X} - \text{Star}(\bar{\Delta})$	complete Σ_U to an edge	choose an edge for every $U \in \bar{\Theta}$

Table 1. Schematic representation of the simplex Π . Each cell describes how Π_U is defined whenever the domain U belongs to the area given by the intersection between the row label and the column label (for example, if $U \in \bar{\Sigma} \cap \text{Lk}(\bar{\Delta})$ we have $\Pi_U = \Sigma_U$).

$V \sqsubseteq \bar{\Psi}_{W_0}^\perp = W'$. Then we have proved that

$$(1) \quad \text{Lk}(\bar{\Sigma}) \cap \text{Lk}(\bar{\Delta}) = \text{Lk}(\bar{\Sigma} \star \bar{\Phi} \star \bar{\Psi} \star \bar{\Theta}) \star \bar{\Psi},$$

where

$$\begin{cases} \bar{\Phi} = \text{Lk}(\bar{\Sigma}) \cap \bar{\Delta}; \\ \bar{\Psi} \subseteq \text{Lk}(\bar{\Delta}) \cap \text{Lk}(\bar{\Sigma}); \\ \bar{\Theta} \subseteq \text{Lk}(\bar{\Sigma}) - \text{Star}(\bar{\Delta}). \end{cases}$$

Finding the extension of Σ : Let $\bar{\Lambda} = \bar{\Phi} \star \bar{\Psi} \star \bar{\Theta}$, and let Π be the simplex defined as follows:

- $p(\Pi) = \bar{\Sigma} \star \bar{\Lambda}$.
- If $U \in \bar{\Sigma}$ does not belong to $\text{Star}(\bar{\Delta})$ then Π_U is an edge containing Σ_U , so that $\text{Lk}_{p^{-1}(U)}(\Pi_U) = \emptyset$.
- If $U \in \bar{\Sigma} \cap \text{Lk}(\bar{\Delta})$ then $\Pi_U = \Sigma_U$.
- If $U \in \bar{\Sigma} \cap \bar{\Delta}$ then $\Pi_U = \Sigma_U$ whenever Σ_U and Δ_U are single vertices “of the same kind” (that is, either they are both the vertex of the cone or they are both points in the base); otherwise Π_U is an edge containing Σ_U . In other words, we choose Π_U so that

$$\text{Lk}_{p^{-1}(U)}(\Pi_U) = \text{Lk}_{p^{-1}(U)}(\Sigma_U) \cap \text{Lk}_{p^{-1}(U)}(\Delta_U).$$

- If $U \in \bar{\Phi}$ then $\Pi_U = \Delta_U$.
- If $U \in \bar{\Psi}$ then Π_U is the cone point v_U .
- If $U \in \bar{\Theta}$ then Π_U is an edge, so that $\text{Lk}_{p^{-1}(U)}(\Pi_U) = \emptyset$.

Moreover, we define the simplex Ψ such that $p(\Psi) = \bar{\Psi}$ and that, for every $U \in \bar{\Psi}$, $\Psi_U = v_U$ (this is exactly how we defined Π_U for $U \in \bar{\Psi}$).

We are finally ready to prove that $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) = \text{Lk}(\Pi) \star \Psi$. First, we notice that $\Psi \subseteq \text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$ since its support is $\bar{\Psi}$, whose vertices lie in W_0 . Next, we argue that $\text{Lk}(\Pi) \subseteq \text{Lk}(\Delta)$. Let $u \in \text{Lk}(\Pi)$ and let $U = p(u)$. If $U \in \bar{\Sigma} \star \bar{\Lambda}$ then a careful inspection of how we defined Π shows that $u \in \text{Lk}(\Delta)$.

Otherwise $U \in \text{Lk}(\bar{\Sigma} \star \bar{\Lambda}) \subset \text{Lk}(\bar{\Delta})$, and therefore $u \in \text{Lk}(\Delta)$. Thus we showed that $\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) \supseteq \text{Lk}(\Pi) \star \Psi$.

For the converse inclusion, let $u \in \text{Lk}(\Delta) \cap \text{Lk}(\Sigma)$, so that $U = p(u)$ belongs to $\text{Star}(\bar{\Sigma}) \cap \text{Star}(\bar{\Delta})$ and $u \in (\text{Lk}_{p^{-1}(U)} \Sigma_U) \cap (\text{Lk}_{p^{-1}(U)} \Delta_U)$. There are four possible cases:

- If $U \in \bar{\Sigma} \cap \bar{\Delta}$ then $(\text{Lk}_{p^{-1}(U)} \Sigma_U) \cap (\text{Lk}_{p^{-1}(U)} \Delta_U) = \text{Lk}_{p^{-1}(U)}(\Pi_U)$, as we already noticed.
- If $U \in \bar{\Sigma} \cap \text{Lk}(\bar{\Delta})$ then $\Delta_U = \emptyset$, and again by construction we have

$$(\text{Lk}_{p^{-1}(U)} \Sigma_U) \cap (\text{Lk}_{p^{-1}(U)} \Delta_U) = \text{Lk}_{p^{-1}(U)} \Sigma_U = \text{Lk}_{p^{-1}(U)}(\Pi_U).$$

- Symmetrically, if $U \in \text{Lk}(\bar{\Sigma}) \cap \bar{\Delta} = \bar{\Phi}$ then $\Sigma_U = \emptyset$, and we have

$$(\text{Lk}_{p^{-1}(U)} \Sigma_U) \cap (\text{Lk}_{p^{-1}(U)} \Delta_U) = \text{Lk}_{p^{-1}(U)} \Delta_U = \text{Lk}_{p^{-1}(U)}(\Pi_U).$$

- Finally, suppose $U \in \text{Lk}(\bar{\Sigma}) \cap \text{Lk}(\bar{\Delta})$, that is, U is nested in W_0 . Then by construction either U is nested in W' or $U \in \bar{\Psi}$. In the former case $U \in \text{Lk}(\bar{\Sigma} \star \bar{\Lambda})$, hence $u \in \text{Lk}(\Pi)$. In the latter case, either $u = v_U$ is the cone point, which belongs to Ψ , or $u \in \text{Lk}_{p^{-1}(U)}(v_U)$, and since $\Pi_U = v_U$ we have $u \in \text{Lk}(\Pi)$.

This concludes the proof. \square

We point out the following byproduct of the proof:

Corollary 5.9. *Under Assumption 5.1, let $\bar{\Sigma}, \bar{\Delta}$ be two simplices of \bar{X} . Then there exist two simplices $\bar{\Lambda}, \bar{\Psi} \subset \text{Lk}(\bar{\Sigma})$ such that*

$$\text{Lk}(\bar{\Sigma}) \cap \text{Lk}(\bar{\Delta}) = \text{Lk}(\bar{\Sigma} \star \bar{\Lambda}) \star \bar{\Psi}.$$

Furthermore, we can assume that $\text{Lk}(\bar{\Sigma} \star \bar{\Lambda})$ is not the cone with cone point $V \in \text{Lk}(\bar{\Sigma} \star \bar{\Lambda})$.

Proof. The first part of the statement is just equation (1). For the “furthermore” part, if $\text{Lk}(\bar{\Sigma} \star \bar{\Lambda})$ is the cone with cone point $V \in \text{Lk}(\bar{\Sigma} \star \bar{\Lambda})$, then

$$\text{Lk}(\bar{\Sigma} \star \bar{\Lambda}) \star \bar{\Psi} = \text{Lk}(\bar{\Sigma} \star \bar{\Lambda} \star V) \star \bar{\Psi} \star V.$$

In this case we can set $\bar{\Lambda}' = \bar{\Lambda} \star V$ and $\bar{\Psi}' = \bar{\Psi} \star V$, and then we check again if $\text{Lk}(\bar{\Sigma} \star \bar{\Lambda}')$ is a cone with cone point V' . This process must end after finitely many steps, since \bar{X} has finite dimension by Remark 4.5. \square

As a consequence of the previous lemma we can also verify another axiom:

Corollary 5.10 (verification of Definition 2.7(1)). *Under Assumption 5.1, X has finite complexity in the sense of Definition 2.6.*

Proof. One can argue exactly as in the proof of [BHMS24, Claim 6.9], which only uses that X has finite dimension, as pointed out in Remark 4.5, and [BHMS24, Condition 6.4.B], which is our Lemma 5.8. \square

5.3. Fullness of links.

Lemma 5.11 (verification of Definition 2.7(4)). *Under Assumption 5.1, there exists a constant $\lambda_0 = \lambda_0(E)$ such that \mathcal{W} has the following property whenever $\lambda \geq \lambda_0$. Let Δ be a nonmaximal simplex of X . Suppose that $v, w \in \text{Lk}(\Delta)$ are distinct, nonadjacent vertices contained in \mathcal{W} -adjacent maximal simplices σ_v, σ_w . Then there exist maximal simplices Π_v, Π_w of X such that $\Delta \star v \subseteq \Pi_v$ and $\Delta \star w \subseteq \Pi_w$.*

Proof. Recall that $p : X \rightarrow \bar{X}$ is the retraction from Definition 4.2 that maps every vertex of the blow-up to its support. Moreover, for every maximal simplex $\sigma = \star_{i=1}^k \{v_{U_i}, x_i\}$ of X , where $U_i \in \mathfrak{S}$ and $x_i \in \mathcal{C}U_i$, let $(b(\sigma)_W)_{W \in \mathfrak{S}}$ be the tuple from Definition 4.6, which was defined by

$$b(\sigma)_W = \begin{cases} x_i & \text{if } W = U_i; \\ \bigcup_{U_i \not\perp W} \rho_W^{U_i} & \text{otherwise.} \end{cases}$$

Set $V = p(v)$, $W = p(w)$ and $\bar{\Delta} = p(\Delta)$. Let σ_v, σ_w be the two \mathcal{W} -adjacent simplices containing v and w , respectively, and let $\bar{\Sigma} = \bar{\sigma}_v \cap \bar{\sigma}_w$ (which is possibly empty). Let $k = \text{co-lv}(\bar{\Sigma}_{\min}^\perp)$. Recall that, by Definition 4.10 of the edges of \mathcal{W} , we have, for every $U \in \mathfrak{S}$,

$$d_{\mathcal{C}U}(b(\sigma_v)_U, b(\sigma_w)_U) \leq (k + 1)\lambda.$$

If $V = W$ then we can complete Δ to two simplices Π_v, Π_w with the same support (so that, in the sense of Notation 4.9, the colevel is n) and such that $v \in \Pi_v, w \in \Pi_w$ and these simplices coincide away from V . Since

$$d_V(b(\Pi_v)_V, b(\Pi_w)_V) = d_V(b(\sigma_v)_V, b(\sigma_w)_V) \leq (k + 1)\lambda \leq (n + 1)\lambda,$$

while by construction $d_U(b(\Pi_v)_U, b(\Pi_w)_U) = 0$ whenever $U \neq V$, we are done.

Moreover, notice that $V \not\perp W$, otherwise v and w would be adjacent in X by how we constructed the blow-up graph.

Thus we are left with the case when $V \neq W$ and $V \not\perp W$. In particular, none of them lies inside either $\bar{\Delta}$ or $\bar{\Sigma}$. Therefore $V, W \in \text{Lk}(\bar{\Sigma}) \cap \text{Lk}(\bar{\Delta})$. By Corollary 5.9 we can find two simplices $\bar{\Phi}, \bar{\Phi}' \subset \text{Lk}(\bar{\Sigma})$ such that

$$\text{Lk}(\bar{\Sigma}) \cap \text{Lk}(\bar{\Delta}) = \text{Lk}(\bar{\Delta} \star \bar{\Phi}) \star \bar{\Phi}'.$$

Moreover, we can assume that $\text{Lk}(\bar{\Delta} \star \bar{\Phi})$ is not a cone with any of its vertices as cone point, which implies that $(\bar{\Delta} \star \bar{\Phi})_{\min}^\perp = (\bar{\Delta} \star \bar{\Phi})^\perp$ by Lemma 5.6.

Now, V cannot lie in $\bar{\Phi}'$, since otherwise it would be orthogonal to every other vertex of $\text{Lk}(\bar{\Sigma}) \cap \text{Lk}(\bar{\Delta})$, including W . Hence V must lie in $\text{Lk}(\bar{\Delta} \star \bar{\Phi})$, and the same holds for W . Now set $\bar{\Psi}_v = \bar{\sigma}_v \cap \text{Lk}(\bar{\Delta} \star \bar{\Phi})$, which contains V as we just argued, and similarly $\bar{\Psi}_w = \bar{\sigma}_w \cap \text{Lk}(\bar{\Delta} \star \bar{\Phi})$. The situation in \bar{X} is therefore as in Figure 3.




Figure 3. The simplices involved in the construction of Lemma 5.11, where edges denote joins in \bar{X} . Actually, $\bar{\Sigma}$ and $\bar{\Delta} \star \bar{\Phi}$ need not be disjoint, but none of them can contain V or W .

Now, $\text{Lk}(\bar{\Delta} \star \bar{\Phi}) \subseteq \text{Lk}(\bar{\Sigma})$ by construction, or equivalently $(\bar{\Delta} \star \bar{\Phi})_{\min}^\perp \subseteq \bar{\Sigma}_{\min}^\perp$ by Lemma 5.4. Hence

$$(2) \quad (\bar{\Delta} \star \bar{\Phi})^\perp = (\bar{\Delta} \star \bar{\Phi})_{\min}^\perp \subseteq \bar{\Sigma}_{\min}^\perp \subseteq \bar{\Sigma}^\perp.$$

There are two possible cases, depending on whether the two orthogonal complements coincide or not.

Case 1: Suppose first that $(\bar{\Delta} \star \bar{\Phi})^\perp = \bar{\Sigma}^\perp$. This means that $\bar{\sigma}_v = \bar{\Sigma} \star \bar{\Psi}_v$, because a support of $\bar{\sigma}_v$ which is orthogonal to $\bar{\Sigma}$ must also be orthogonal to $\bar{\Delta} \star \bar{\Phi}$. Therefore $\bar{\Delta} \star \bar{\Phi} \star \bar{\Psi}_v$ is already a maximal simplex, and the same is true for $\bar{\Delta} \star \bar{\Phi} \star \bar{\Psi}_w$. Now complete Δ to maximal simplices Π_v, Π_w as follows:

- Π_v is supported on $\bar{\Delta} \star \bar{\Phi} \star \bar{\Psi}_v$, and similarly Π_w is supported on $\bar{\Delta} \star \bar{\Phi} \star \bar{\Psi}_w$;
- if $U \in \bar{\Delta}$ and $\Delta_U \cap (\mathcal{C}U)^{(0)} = \emptyset$ (that is, if Δ does not prescribe the coordinate for U) then choose the same coordinate both for Π_v and Π_w ;
- if $U \in \bar{\Delta}$ and $\Delta_U \cap (\mathcal{C}U)^{(0)} = \{p_U\}$ (that is, if Δ already prescribes the coordinate for U) then set $\Pi_v = \Pi_w = \{v_U, p_U\}$;
- if $U \in \bar{\Phi}$, choose the same coordinate both for Π_v and Π_w ;
- if $U \in \bar{\Psi}_v$ choose for Π_v the coordinate prescribed by σ_v , and similarly if $U \in \bar{\Psi}_w$ choose for Π_w the coordinate prescribed by σ_w .

Now, since $\bar{\Sigma}^\perp = (\bar{\Delta} \star \bar{\Phi})^\perp$ they have the same colevel k . Thus, in order to show that Π_v and Π_w are \mathcal{W} -adjacent, it is enough to prove that, for any $U \in \mathfrak{S}$, we have $d_{\mathcal{C}U}(b(\Pi_v)_U, b(\Pi_w)_U) \leq (k + 1)\lambda$, because $p(\Pi_v)$ and $p(\Pi_w)$ coincide at least on $\bar{\Delta} \star \bar{\Phi}$.

If $U \notin (\bar{\Delta} \star \bar{\Phi})$ then clearly $d_{\mathcal{C}U}(b(\Pi_v)_U, b(\Pi_w)_U) = 0$. Otherwise U is also orthogonal to $\bar{\Sigma}$, and by maximality of σ_v it cannot be orthogonal to every vertex of $\bar{\sigma}_v - \bar{\Sigma} = \bar{\Psi}_v$. This means that $b(\Pi_v)_U = b(\sigma_v)_U$, since they both depend only on the coordinates over $\bar{\Psi}_v$, which are the same for both Π_v and σ_v by construction.

For the same reason $b(\Pi_w)_U = b(\sigma_w)_U$. This in turn means that

$$d_{CU}(b(\Pi_v)_U, b(\Pi_w)_U) = d_{CU}(b(\sigma_v)_U, b(\sigma_w)_U) \leq (k + 1)\lambda.$$

Case 2: Now we are left to deal with the case when $(\bar{\Delta} \star \bar{\Phi})^\perp \not\sqsubseteq \bar{\Sigma}^\perp$. We will find two maximal simplices Π_v, Π_w whose supports extend $\bar{\Delta} \star \bar{\Phi} \star \bar{\Psi}_v$ and $\bar{\Delta} \star \bar{\Phi} \star \bar{\Psi}_w$, respectively, and then it will suffice to prove that, for every domain $U \in \mathfrak{G}$, we have

$$d_{CU}(b(\Pi_v)_U, b(\Pi_w)_U) \leq (k + 2)\lambda.$$

In other words, it will be enough to loosen the threshold just by adding a single λ . Since $\text{co-lv}((\bar{\Delta} \star \bar{\Phi} \star \bar{\Psi}_v)^\perp) \not\sqsubseteq \text{co-lv}(\bar{\Sigma}^\perp) = k$ we will then have Π_v and Π_w are \mathcal{W} -adjacent.

Let $\bar{\Theta}_v = \bar{\sigma}_v - (\bar{\Sigma} \star \bar{\Psi}_v)$ be the simplex spanned by all remaining vertices of $\bar{\sigma}_v$. Moreover, let $R_v = (\bar{\Delta} \star \bar{\Phi} \star \bar{\Psi}_v)^\perp$, if it exists (if not, the following construction is unnecessary because our simplex is already maximal). Notice that R_v is also orthogonal to $\bar{\Sigma}$, hence every domain $U \sqsubseteq R_v$ is orthogonal to every domain in $\bar{\Sigma} \star \bar{\Psi}_v$, and therefore it cannot be orthogonal to every vertex of $\bar{\Theta}_v$ by maximality of $\bar{\sigma}_v = \bar{\Sigma} \star \bar{\Psi}_v \star \bar{\Theta}_v$.

Then let $r^v = (r^v_U)_{U \sqsubseteq R_v} \in F_{R_v}$ be the tuple defined as follows:

- if $U \in \bar{\Theta}_v$ then r^v_U is the coordinate prescribed by σ_v ;
- otherwise $r^v_U = \bigcup \rho_U^{U'}$ where the union varies among all $U' \in \bar{\Theta}_v$ whose projection to U is defined.

By the previous argument, r^v_U is well-defined for any $U \sqsubseteq R_v$. Moreover, arguing exactly as in Remark 4.12, one sees that r^v is indeed a $20E$ -consistent tuple, and therefore an element of F_{R_v} .

Now, if R_v is \sqsubseteq -minimal then r^v is just a point in $\mathcal{C}R_v$, and we set $\bar{\Omega}_v = R_v$ and Ω_v as the edge $\{v_{R_v}, r^v\}$. Otherwise, by the EDPR (3.12) there exist a maximal family of pairwise orthogonal, \sqsubseteq -minimal domains $\bar{\Omega}_v = \{O_1, \dots, O_l\}$ whose realisation point is C_0 -close to the realisation of r in F_{R_v} . This means that, if we define the simplex Ω_v , supported in $\bar{\Omega}_v$, by choosing for every $I \in \bar{\Omega}_v$ the coordinate r^v_I , then for every $U \sqsubseteq R_v$ and every maximal simplex Ω'_v containing Ω_v , the U -coordinate of the realisation tuple $b(\Omega'_v)$ is M -close to r^v_U , where $M = M(C_0, E)$ is a constant coming from the distance formula, Theorem 1.11.

Define $\bar{\Theta}_w, R_w, \bar{\Omega}_w$ analogously, so that the situation looks like in Figure 4.

Then complete Δ to maximal simplices Π_v, Π_w as follows:

- Π_v is supported on $\bar{\Delta} \star \bar{\Phi} \star \bar{\Psi}_v \star \bar{\Omega}_v$, and similarly Π_w is supported on $\bar{\Delta} \star \bar{\Phi} \star \bar{\Psi}_w \star \bar{\Omega}_w$;
- if $U \in \bar{\Delta} \star \bar{\Phi}$, choose the same coordinates both for Π_v and Π_w ;




Figure 4. The simplices of \bar{X} involved in the second case of the proof of Lemma 5.11, where edges represent joins. By construction $\bar{\sigma}_v = \bar{\Sigma} \star \bar{\Psi}_v \star \bar{\Theta}_v$, and similarly for $\bar{\sigma}_w = \bar{\Sigma} \star \bar{\Psi}_w \star \bar{\Theta}_w$.

- if $U \in \bar{\Psi}_v$ choose for Π_v the coordinate coming from σ_v , and similarly if $U \in \bar{\Psi}_w$ choose for Π_w the coordinate coming from σ_w
- if $U \in \bar{\Omega}_v$ choose the coordinate coming from Ω_v , and similarly if $U \in \bar{\Omega}_w$ choose the coordinate coming from Ω_w .

Now we show that for every $U \in \mathfrak{S}$ we have $d_{CU}(b(\Pi_v)_U, b(\Pi_w)_U) \leq (k + 2)\lambda$.

If $U \not\perp (\bar{\Delta} \star \bar{\Phi})$ then clearly $d_{CU}(b(\Pi_v)_U, b(\Pi_w)_U) = 0$. Otherwise U is also orthogonal to $\bar{\Sigma}$, since $(\bar{\Delta} \star \bar{\Phi})^\perp \subseteq \bar{\Sigma}^\perp$.

Claim 5.12. *If $U \perp (\bar{\Delta} \star \bar{\Phi})$ then $d_{CU}(b(\Pi_v)_U, b(\sigma_v)_U) \leq M$.*

Proof of Claim 5.12. First suppose that $U \not\perp \bar{\Psi}_v$. Then $b(\Pi_v)_U$ and $b(\sigma_v)_U$ both contain the set

$$\alpha_U^v = \begin{cases} (\Psi_v)_U & \text{if } U \in \bar{\Psi}_v; \\ \bigcup_{T \in \bar{\Psi}_v, T \perp U} \rho_U^T & \text{otherwise.} \end{cases}$$

We are left with the case when $U \perp \bar{\Delta} \star \bar{\Phi} \star \bar{\Phi}_v$. In this case $b(\Pi_v)_U$ contains the set

$$\beta_U^v = \begin{cases} (\Psi_v)_U & \text{if } U \in \bar{\Omega}_v; \\ \bigcup_{T \in \bar{\Omega}_v, T \perp U} \rho_U^T & \text{otherwise.} \end{cases}$$

But by our definition of Ω_v , β_v is M -close to the coordinate

$$r_U^v = \begin{cases} (\Psi_v)_U & \text{if } U \in \bar{\Theta}_v; \\ \bigcup_{T \in \bar{\Theta}_v, T \perp U} \rho_U^T & \text{otherwise,} \end{cases}$$

which is a subset of $b(\sigma_v)_U$. □

The proof of the claim also applies to Π_w and σ_w . Hence

$$d_{CU}(b(\Pi_v)_U, b(\Pi_w)_U) \leq d_{CU}(b(\sigma_v)_U, b(\sigma_w)_U) + 2M \leq (k + 1)\lambda + 2M.$$

Therefore it is enough to choose $\lambda_0 = 2M$, so that $(k + 1)\lambda + 2M \leq (k + 2)\lambda$ whenever $\lambda \geq \lambda_0$. \square

5.4. Hyperbolic links. Recall that, for any domain U , we defined F_U as the space of $20E$ -consistent tuples for U , which is a sub-HHS of $(\mathcal{Z}, \mathfrak{S})$ with maximal domain U . For convenience, up to quasi-isometry we can assume that F_U is a graph (again, by [CdH16, Lemma 3.B.6]). Moreover, for every $V \not\sqsubseteq U$ one has the relative product region $P_V^U \subset F_U$, which we will often refer to as P_V when the ambient domain will be clear.

Definition 5.13. The *factored space* \hat{F}_U associated to U is the graph obtained from F_U by coning off the product region P_V for every $V \not\sqsubseteq U$.

We will denote by H_V the vertex of the cone over P_V . By [BHS17a, Corollary 2.9 and Remark 2.10], if \mathcal{Z} is a normalised HHS then \hat{F}_U is uniformly quasi-isometric to $\mathcal{C}U$, and therefore uniformly hyperbolic. More precisely, the quasi-isometry is induced by the projection $\pi_U : F_U \rightarrow \mathcal{C}U$. Then the strategy to prove that links of simplices inside X are hyperbolic will be to show that each of them is either bounded, quasi-isometric to a coordinate space or to a factored space.

Lemma 5.14. *Under Assumption 5.1, there exists $\lambda_1 \geq \lambda_0$ such that the following holds whenever $\lambda \geq \lambda_1$. There exists δ , depending both on λ and on the HHS structure, such that, for every nonmaximal simplex $\Delta \subset X$, the associated coordinate space $\mathcal{C}([\Delta])$ is δ -hyperbolic.*

Proof. We consider all possible shapes of $\text{Lk}(\Delta)$, according to Corollary 4.4.

If $\text{Lk}(\Delta)$ is a point or a nontrivial join: In this case $\mathcal{C}([\Delta])$, which is obtained by adding edges to $\text{Lk}(\Delta)$, has diameter at most 2, and therefore it is 2-hyperbolic.

If Δ is almost-maximal and $\Delta_U = v_U$ for some $U \in \bar{\Delta}$: In this case $\text{Lk}(\Delta)$ is the base of the cone over $\mathcal{C}U$. Now, by construction two points $p, q \in \text{Lk}(\Delta)$ belong to \mathcal{W} -adjacent maximal simplices if and only if

$$d_{\mathcal{C}U}(p, q) \leq (n + 1)\lambda.$$

This shows that $\mathcal{C}([\Delta])$ and $\mathcal{C}U$ are quasi-isometric with uniform constants, and since $\mathcal{C}U$ is E -hyperbolic then $\mathcal{C}([\Delta])$ is δ -hyperbolic for some constant $\delta = \delta(\lambda, E, n)$.

If Δ_W is an edge for every $W \in \bar{\Delta}$: In this case $\mathcal{C}([\Delta])$ is $(2, 2)$ -quasi-isometric to the subgraph of $\bar{X}^{+\mathcal{W}}$ spanned by $\text{Lk}(\bar{\Delta})$, via the retraction $p : X \rightarrow \bar{X}$. Let $\text{Lk}(\bar{\Delta})^{+\mathcal{W}}$ be this subgraph, and let $U = \bar{\Delta}^\perp$.

Now, if U is \sqsubseteq -minimal then $\text{Lk}(\bar{\Delta})^{+\mathcal{W}}$ consists only of U , hence is uniformly bounded. Then suppose that U is non- \sqsubseteq -minimal. At the level of vertices we can define a map $\psi : \text{Lk}(\bar{\Delta})^{+\mathcal{W}} \rightarrow \hat{F}_U$ by sending every \sqsubseteq -minimal domain $V \not\sqsubseteq U$ to the cone point H_V over the corresponding product region P_V . Our goal is to

prove that ψ is a quasi-isometry with uniform constants, and therefore $\text{Lk}(\bar{\Delta})^{+\mathcal{W}}$ is uniformly hyperbolic.

First we show that ψ is coarsely surjective. For every $x \in F_U$, by the EDPR (3.12) we can find a maximal family $V_1, \dots, V_k \not\sqsubseteq U$ of \sqsubseteq -minimal domains whose product region is C_0 -close to x . Then in particular the product region P_{V_1} is C_0 -close to x in F_U , which means that the corresponding cone point $H_{V_1} = \psi(V_1)$ is $(C_0 + 1)$ -close to x in \hat{F}_U .

Next, we prove that the map ψ is uniformly Lipschitz, by showing that adjacent vertices in $\text{Lk}(\bar{\Delta})^{+\mathcal{W}}$ map to uniformly close points inside \hat{F}_U . If $V, V' \not\sqsubseteq U$ are \sqsubseteq -minimal domains which are joined by an edge of \mathcal{W} then one of the following must hold:

- $V \perp V'$: in this case, the product regions P_V and $P_{V'}$ are already uniformly close inside F_U . Indeed, if one chooses two coordinates $p \in CV$ and $p' \in CV'$ then, by the partial realisation axiom (8) for the sub-HHS F_U , one can find an element $x \in F_U$ whose coordinates must satisfy the following properties:

$$\begin{cases} d_V(x_V, p) \leq E; \\ d_{V'}(x_{V'}, p') \leq E; \\ d_U(x_U, \rho_U^V) \leq E \text{ whenever } U \triangleleft V; \\ d_U(x_U, \rho_U^{V'}) \leq E \text{ whenever } U \triangleleft V'. \end{cases}$$

Therefore, by definition, x belongs to both P_V and $P_{V'}$. This means that the cone points H_V and $H_{V'}$ are at distance at most 2.

- There exist two maximal simplices Π, Π' which extend Δ and such that $V \in \bar{\Pi}$ and $V' \in \bar{\Pi}'$, and for every $W \in \mathfrak{S}$ we have

$$d_{CW}(b(\Pi)_W, b(\Pi')_W) \leq (1 + \text{co-lv}(U))\lambda \leq (1 + n)\lambda.$$

In this case let $x = (b(\Pi)_W)_{W \sqsubseteq U}$ and $x' = (b(\Pi')_W)_{W \sqsubseteq U}$ be the corresponding tuples inside F_U . By the distance formula (Theorem 1.11) for the sub-HHS F_U , the distance of these points is bounded by some constant D depending only on λ, n, E . Moreover $x \in P_V$ and $x' \in P_{V'}$ by construction, and therefore x is adjacent to H_V inside \hat{F}_U and similarly for x' and $H_{V'}$. Therefore $d_{\hat{F}_U}(H_V, H_{V'}) \leq 2 + D$.

In order to complete the proof that ψ is a quasi-isometry we need the following, which is the only spot where we have to choose λ_1 carefully:

Claim 5.15. *Under Assumption 5.1, and with the notation of Lemma 5.14, there exists $\lambda_1 \geq \lambda_0$ such that the following holds whenever $\lambda \geq \lambda_1$. Let $V, W \not\sqsubseteq U$ be two \sqsubseteq -minimal domains. If the product regions P_V and P_W lie within distance at most $2C_0 + 1$ in F_U , where C_0 is the constant from the EDPR (3.12), then V and W are \mathcal{W} -adjacent in $\text{Lk}(\bar{\Delta})^{+\mathcal{W}}$.*

Proof of Claim 5.15. Let $y \in P_V, z \in P_W$ be two points that are $(2C_0 + 1)$ -close in F_U . Let $\bar{\Theta} = \{T_1, \dots, T_l\}$ be a family of pairwise orthogonal, \sqsubseteq -minimal domains such that $T_i \perp U$ for all $i = 1, \dots, l$. Our goal is to complete $\bar{\Theta}$ to a maximal family $V = V_0, V_1, \dots, V_k$ of pairwise orthogonal, \sqsubseteq -minimal elements whose product region is uniformly close to y , and similarly for W and z .

If $V = V_0$ has no orthogonal inside U then we have nothing to do. Otherwise, define $a = (y_I)_{I \in V^\perp} \in F_{V_U^\perp}$ as the coordinates of y in the domains of the orthogonal complement. If V_U^\perp is not itself minimal, by the EDPR there exist a maximal family $V_1, \dots, V_k \not\sqsubseteq V_U^\perp$ of \sqsubseteq -minimal and pairwise orthogonal domains whose product region inside $F_{V_U^\perp}$ is C_0 -close to a . Either way, there exist $p_i \in CV_i$, for $i = 0, \dots, k$, such that the realisation point $b(\Sigma)$ of the simplex $\Sigma = \star_{i=0}^k \{v_{V_i}, p_i\} \subset X$ is C_0 -close to y . Arguing the same way for W we can find a simplex Σ' , whose support contains W and whose realisation point is C_0 -close to w . Therefore the realisation points of these two simplices are $(4C_0 + 1)$ -close, and since projections to coordinate spaces are uniformly coarsely Lipschitz there exists $M = M(C_0)$ such that

$$d_{CI}(b(\Sigma), b(\Sigma')) \leq M \quad \text{for every domain } I \in \mathfrak{S}.$$

Therefore, if we set $\lambda_1 = M$, then whenever $\lambda \geq \lambda_1$, for every $I \in \mathfrak{S}$ we have

$$d_{CI}(b(\Sigma), b(\Sigma')) \leq \lambda.$$

Hence Σ and Σ' are \mathcal{W} -adjacent by definition, and hence V and W are \mathcal{W} -adjacent in $\text{Lk}(\bar{\Delta})^{+\mathcal{W}}$. □

Now, we claim that $d_{\text{Lk}(\bar{\Delta})^{+\mathcal{W}}}(V, V') \leq 4d_{\hat{F}_U}(H_V, H_{V'}) + 2$ for any two \sqsubseteq -minimal domains $V, V' \not\sqsubseteq U$, and this will complete the proof that ψ is a quasi-isometric embedding. Consider a geodesic $\gamma \subset \hat{F}_U$ from H_V to $H_{V'}$. The vertices of γ can either be tuples of F_U or cone points associated to product regions. Let x_1, \dots, x_k be the vertices of γ belonging to F_U , and for every x_i consider a \sqsubseteq -minimal domain I_i whose product region is C_0 -close to x_i in F_U , which exists by the EDPR. Then the situation in \hat{F}_U is as shown in Figure 5.




Figure 5. The continuous segments represent the path $\gamma \subset \hat{F}_U$. For each vertex x_i of γ which lies inside F_U we choose a \sqsubseteq -minimal domain $I_i \in \mathfrak{S}$ whose product region is C_0 -close to x_i in F_U .

Now, Claim 5.15 implies that V is \mathcal{W} -adjacent to I_1 (since $x_1 \in P_V$ by construction), and similarly that V' is \mathcal{W} -adjacent to I_k . Then the proof is complete if we show that, for every $i = 1, \dots, k - 1$, $d_{\text{Lk}(\bar{\Delta})+\mathcal{W}}(I_i, I_{i+1}) \leq 4$, because then $d_{\text{Lk}(\bar{\Delta})+\mathcal{W}}(V, V') \leq 4L(\gamma) + 2$ where $L(\gamma)$ is the length of γ .

Thus let x_i, x_{i+1} be two consecutive vertices of $\gamma \cap F_U$. If they are joined by an edge of F_U (in Figure 5 this happens to x_1 and x_2), then P_{I_i} and $P_{I_{i+1}}$ are at most $(2C_0 + 1)$ -close; thus I_i and I_{i+1} are \mathcal{W} -adjacent by Claim 5.15.

Otherwise γ might contain a segment of the form $\{x_i, H_W, x_{i+1}\}$, where $W \not\sqsubseteq U$ is a domain such that $x_i, x_{i+1} \in P_W$ (in Figure 5 this happens to x_2 and x_3).

If W is \sqsubseteq -minimal then by Claim 5.15 we see that I_i and W are \mathcal{W} -adjacent, and similarly for I_{i+1} and W . Therefore $d_{\text{Lk}(\bar{\Delta})+\mathcal{W}}(I_i, I_{i+1}) \leq 2$.

Otherwise suppose that W is not \sqsubseteq -minimal, and let $y_i = ((x_i)_T)_{T \sqsubseteq W} \in F_W$, which is again $20E$ -consistent since it is a sub-tuple of a consistent tuple. By the EDPR there exists a minimal domain $R_i \not\sqsubseteq W$ such that y_i is C_0 -close to the relative product region $P_{R_i}^W$ inside F_W . This also means that the whole tuple x_i is C_0 -close to the relative product region $P_{R_i} = P_{R_i}^U$ inside F_U , since any tuple in $P_{R_i}^W$ can be completed to a tuple in $P_{R_i}^U$ by choosing some coordinates for the domains $T \sqsubseteq W_U^\perp$. Therefore I_i and R_i are \mathcal{W} -adjacent, by Claim 5.15, since their product regions in F_U are $(2C_0 + 1)$ -close.

Now, if W is split then pick a Samaritan $Q \sqsubseteq W$. Hence either $R_i = Q$, or $R_i \perp Q$ and therefore they are \mathcal{W} -adjacent. Then $d_{\text{Lk}(\bar{\Delta})+\mathcal{W}}(I_i, Q) \leq 2$, and by repeating this procedure with I_{i+1} and the same Samaritan Q we also get $d_{\text{Lk}(\bar{\Delta})+\mathcal{W}}(I_{i+1}, Q) \leq 2$. Therefore $d_{\text{Lk}(\bar{\Delta})+\mathcal{W}}(I_i, I_{i+1}) \leq 4$.

If otherwise W is nonsplit, by Property 3.9 there exists a minimal domain $Q' \not\sqsubseteq U$ which is orthogonal to W . Then Q' and R_i are \mathcal{W} -adjacent since they are orthogonal, and arguing as before we get

$$d_{\text{Lk}(\bar{\Delta})+\mathcal{W}}(I_i, I_{i+1}) \leq d_{\text{Lk}(\bar{\Delta})+\mathcal{W}}(I_i, Q') + d_{\text{Lk}(\bar{\Delta})+\mathcal{W}}(Q', I_{i+1}) \leq 4.$$

This concludes the proof. □

5.5. *QI-embedding of coordinate spaces.*

Lemma 5.16. *Under Assumption 5.1, whenever $\lambda \geq \lambda_1$ there exists a constant δ' , depending both on λ and on the HHS structure, such that, for every nonmaximal simplex $\Delta \subset X$, the associated coordinate space $\mathcal{C}([\Delta])$ is (δ', δ') -quasi-isometrically embedded inside Y_Δ .*

Proof. Again, we look at all possible shapes of $\text{Lk}(\Delta)$. If it is a single point or a nontrivial join then of course it is $(2, 2)$ -quasi-isometrically embedded inside Y_Δ . In all other cases, since $\mathcal{C}([\Delta])$ is a subgraph of Y_Δ , it will be enough to construct a coarsely Lipschitz retraction from Y_Δ to $\text{Lk}(\Delta)$.

If Δ is almost maximal: Let U be the \sqsubseteq -minimal domain such that $\Delta_U = v_U$. Then $\text{Lk}(\Delta)^{(0)}$ is the copy of $(\mathcal{C}U)^{(0)}$ inside X , and since in Lemma 5.14 we showed that $\mathcal{C}([\Delta])$ is uniformly quasi-isometric to $\mathcal{C}(U)$, we will often replace distances in $\mathcal{C}([\Delta])$ with distances in $\mathcal{C}(U)$ without explicit mention.

Claim 5.17. Y_Δ is the subgraph spanned by $\text{Lk}(\Delta)^{(0)} = (\mathcal{C}U)^{(0)}$ and the cones over $\mathcal{C}V$ for all \sqsubseteq -minimal domains $V \pitchfork U$. In other words,

$$Y_\Delta = \text{span}_{X+\mathcal{W}} \left\{ \mathcal{C}U^{(0)} \cup \bigcup_{p(v)\pitchfork U} \{v\} \right\}.$$

Proof of Claim 5.17. Since, by definition $Y_\Delta = \text{span}_{X+\mathcal{W}} \{X - \text{Sat}(\Delta)\}^{(0)}$, we will equivalently prove that

$$\text{Sat}(\Delta)^{(0)} = \{v_U\} \cup \bigcup_{p(v)\perp U} \{v\}.$$

Indeed, clearly $v_U \in \text{Sat}(\Delta)^{(0)}$, while $\mathcal{C}U^{(0)} \cap \text{Sat}(\Delta) = \emptyset$. Moreover, whenever $V = p(v)$ is orthogonal to U , it is possible to complete $\{U, V\}$ to a maximal family of pairwise orthogonal domains $\bar{\Sigma} = \{U, V, V_1, \dots, V_k\}$, and we can find an almost-maximal simplex Σ supported in $\bar{\Sigma}$ which contains v and such that $\text{Lk}(\Sigma) = \text{Lk}(\Delta) = \mathcal{C}U^{(0)}$. Hence $v \in \text{Sat}(\Delta)$. Conversely, if $v \in \text{Sat}(\Delta)$ then $\text{Lk}(v) \supseteq \text{Lk}(\Delta)$ by definition, therefore $V = p(v)$ either coincides with, or is orthogonal to U . \square

Now, at the level of vertices, we can define a retraction $r : Y_\Delta \rightarrow \mathcal{C}([\Delta])$ by

$$r(v) = \begin{cases} v & \text{if } v \in (\mathcal{C}([\Delta]))^{(0)}; \\ \rho_U^{p(v)} & \text{otherwise.} \end{cases}$$

This is well-defined, as a consequence of Claim 5.17. We are left to prove that this retraction is coarsely Lipschitz. Let v, v' be two adjacent vertices in Y_Δ , and we will show that $d_{\mathcal{C}(\Delta)}(r(v), r(v'))$ is uniformly bounded from above.

- If both v, v' belong to $(\mathcal{C}U)^{(0)}$ then they are adjacent in $\mathcal{C}([\Delta])$, by how the latter is defined.
- If $v \in (\mathcal{C}U)^{(0)}$ but $v' \notin (\mathcal{C}U)^{(0)}$ then, setting $V' = p(v')$ we must have $U \pitchfork V'$. Moreover, by definition of \mathcal{W} -edges, there must be two simplices σ, σ' such that $v \in \sigma, v' \in \sigma'$ and the corresponding realisation tuples $b(\sigma), b(\sigma')$ are at least $(n + 1)\lambda$ -close in every coordinate space. Now $b(\sigma)_U = v$ by construction, while $b(\sigma')_U$ is a set of diameter $10E$ which contains that $\rho_U^{V'}$. Hence

$$d_U(r(v), r(v')) = d_U(v, \rho_U^{V'}) = d_U(b(\sigma)_U, b(\sigma')_U) + 10E \leq (n + 1)\lambda + 10E.$$

- We are left with the case where both v, v' do not belong to $(\mathcal{C}U)^{(0)}$, and we want to find an upper bound for $d_U(r(v), r(v')) = d_U(\rho_U^V, \rho_U^{V'})$, where $V = p(v)$ and

$V' = p(v')$. If $V = V'$ then we have nothing to prove. Otherwise V and V' may be adjacent in $\bar{X}^{+\mathcal{W}}$ for two different reasons. If $V \perp V'$, then by Lemma 1.3 we see that $d_U(\rho_U^V, \rho_U^{V'}) \leq 2E$. Otherwise there are two simplices σ, σ' such that $v \in \sigma, v' \in \sigma'$ and the corresponding realisation tuples $b(\sigma), b(\sigma')$ are at least $(n + 1)\lambda$ -close in every coordinate space. Then, similarly to the previous case, we have

$$d_U(\rho_U^V, \rho_U^{V'}) \leq d_U(b(\sigma), b(\sigma')) + 20E \leq (n + 1)\lambda + 20E.$$

If Δ_V is an edge for every $V \in \bar{\Delta}$: Let $\bar{\Delta} = \{V_1, \dots, V_l\}$ and let U be the orthogonal complement of $\bar{\Delta}$. If U is split then $\text{Lk}(\Delta)$ is a join, and therefore it is quasi-isometrically embedded into Y_Δ since it is uniformly bounded.

Thus we can assume that U is nonsplit. The next step is the following:

Claim 5.18. *If a vertex $v \in X$ belongs to Y_Δ then $V = p(v)$ is not orthogonal to U , and therefore $\rho_U^{p(v)}$ is well-defined.*

Proof of Claim 5.18. We prove the contrapositive of the statement, that is, we show that if $V = p(v)$ is orthogonal to U then $v \in \text{Sat}(\Delta)$. We have $U \sqsubseteq V^\perp$, and either they coincide or there exists a \sqsubseteq -minimal domain V_1 inside V^\perp such that $V_1 \perp U$, by Property 3.9 (which applies since we are in the case when U is nonsplit). Then after finitely many steps we can find a simplex $\bar{\Sigma} = \{V = V_0, \dots, V_k\}$ containing V and whose orthogonal complement is U . Then $U = \bar{\Sigma}^\perp = \bar{\Delta}^\perp$, and therefore $\text{Lk}_{X+\mathcal{W}}(\bar{\Sigma}) = \text{Lk}_{X+\mathcal{W}}(\bar{\Delta})$ because they are both spanned by the \sqsubseteq -minimal domains inside U . Hence, since $p(v) = V \in \bar{\Sigma}$ we have $v \in \text{Sat}(\Delta)$. \square

Now, the proof of Lemma 5.14 gives a uniform quasi-isometry $\mathcal{C}([\Delta]) \mapsto \hat{F}_U$, mapping the cone over a \sqsubseteq -minimal domain $V \not\sqsubseteq U$ to the cone point H_V . In turn, F_U is uniformly quasi-isometric to $\mathcal{C}U$ via the projection map. Hence the composition of these maps is a quasi-isometry $\psi : \mathcal{C}([\Delta]) \rightarrow \mathcal{C}U$, which maps the cone over a \sqsubseteq -minimal domain $V \not\sqsubseteq U$ to ρ_U^V .

Now define $r : Y_\Delta \rightarrow \mathcal{C}([\Delta])$ by mapping every vertex $v \in Y_\Delta$ to $\rho_U^{p(v)} \in \mathcal{C}U$ (which is well-defined by Claim 5.18) and then applying the inverse quasi-isometry $\psi^{-1} : \mathcal{C}U \rightarrow \mathcal{C}([\Delta])$. Notice that if $p(v) \sqsubseteq U$ then $r(v) = p(v)$, and therefore r is a coarse retraction onto $\mathcal{C}([\Delta])$.

We are left to prove that, whenever $v, v' \in Y_\Delta$ are adjacent vertices, then $d_U(\rho_U^{p(v)}, \rho_U^{p(v')})$ is uniformly bounded from above, and therefore r is coarsely Lipschitz as it is the composition of a Lipschitz map and the uniform quasi-isometry ψ^{-1} . There are three possibilities:

- If $p(v) = p(v')$ then we have nothing to prove.
- If $p(v) \perp p(v')$ then by Lemma 1.3 we have $d_U(\rho_U^{p(v)}, \rho_U^{p(v')}) \leq 2E$.

- If v, v' lie in \mathcal{W} -adjacent simplices σ, σ' , respectively, then

$$d_U(\rho_U^{p(v)}, \rho_U^{p(v')}) \leq d_U(b(\sigma), b(\sigma')) + 20E \leq (n + 1)T + 20E$$

This concludes the proof. □

5.6. The realisation map is a quasi-isometry. We are left to prove the following lemma, which concludes the proof of Theorem 5.2:

Lemma 5.19. *Under Assumption 5.1 there exists $\lambda_2 \geq \lambda_1$ such that, whenever $\lambda \geq \lambda_2$, the map $f : \mathcal{W} \rightarrow \mathcal{Z}$ from Definition 4.11 is a quasi-isometry.*

Proof. First we show that f is Lipschitz. Given two \mathcal{W} -adjacent simplices Σ, Δ , we have $d_{CU}(b(\Sigma)_U, b(\Delta)_U) \leq (n + 1)\lambda$ for every $U \in \mathfrak{S}$, because n is the maximum colevel, so gives the highest threshold in the Definition 4.10 of \mathcal{W} -edges. Moreover, since $f(\Sigma)$ realises $b(\Sigma)$ and $f(\Delta)$ realises $b(\Delta)$, for every $U \in \mathfrak{S}$ we have

$$d_{CU}(f(\Sigma), f(\Delta)) \leq 2E + d_{CU}(b(\Sigma)_U, b(\Delta)_U) \leq 2E + (n + 1)\lambda.$$

Thus, by the distance formula (Theorem 1.11), there exists $M = M(E, \lambda, n)$ such that $d_{\mathcal{Z}}(f(\Sigma), f(\Delta)) \leq M$. This proves that f is M -Lipschitz.

Furthermore, f is coarsely surjective. Indeed, the whole HHS \mathcal{Z} coincides with F_S , where $S \in \mathfrak{S}$ is the \sqsubseteq -maximal element. Hence, by the EDPR there exists a constant C_0 such that every $z \in \mathcal{Z}$ is C_0 -close to the product region $P_{\{V_i\}_{i=1,\dots,k}}$ associated to a maximal family V_1, \dots, V_k of \sqsubseteq -minimal, pairwise orthogonal domains. In particular, z is C_0 -close to some realisation point z' for some simplex Δ , and such a point uniformly coarsely coincides with $f(\Delta)$ by the uniqueness axiom (9).

We are left to prove that, for every two maximal simplices Σ, Δ of X , their distance in \mathcal{W} is bounded above in terms of $d_{\mathcal{Z}}(f(\Sigma), f(\Delta))$. Now, \mathcal{Z} is a K -quasigeodesic metric space for some $K \geq 0$, therefore it is possible to find a (K, K) -quasigeodesic path

$$\gamma = \{x_0 = f(\Sigma), x_1, \dots, x_{l-1}, x_l = f(\Delta)\}$$

from $f(\Sigma)$ to $f(\Delta)$. Thus the number l of vertices of this path is bounded above by $Kd_{\mathcal{Z}}(f(\Sigma), f(\Delta)) + K$, and for every $i = 0, \dots, l - 1$ we have $d_{\mathcal{Z}}(x_i, x_{i+1}) \leq 2K$.

Moreover, by coarse surjectivity of f , for every $i = 1, \dots, l - 1$ we can find a simplex Σ_i such that $f(\Sigma_i)$ is C_0 -close to x_i . Hence, setting $\Sigma_0 = \Sigma$ and $\Sigma_l = \Delta$, for every $i = 0, \dots, l - 1$ we have $d_{\mathcal{Z}}(f(\Sigma_i), f(\Sigma_{i+1})) \leq D$, where $D = 2C_0 + 2K$.

Claim 5.20. *If λ_2 is large enough, each two consecutive simplices Σ_i and Σ_{i+1} must be \mathcal{W} -adjacent.*

If this is true we are done, since then

$$d_{\mathcal{W}}(\Sigma, \Delta) \leq \sum_{i=0}^{l-1} d_{\mathcal{W}}(\Sigma_i, \Sigma_{i+1}) = l \leq K d_{\mathcal{Z}}(f(\Sigma), f(\Delta)) + K.$$

We prove the claim. By the distance formula Theorem 1.11, there exists a constant $M_0 = M_0(D)$ such that, for every $U \in \mathfrak{S}$, we have $d_{CU}(f(\Sigma_i), f(\Sigma_{i+1})) \leq M_0$. Then in turn $d_{CU}(b(\Sigma_i), b(\Sigma_{i+1})) \leq M_0 + 2E$, and if we choose $\lambda_2 \geq M_0 + 2E$ then Σ_i and Σ_{i+1} are \mathcal{W} -adjacent whenever $\lambda \geq \lambda_2$ (regardless of how their supports intersect, because λ is the tightest threshold for the definition of a \mathcal{W} -edge). \square

6. Adding a group action

Here we make some remarks on the construction of the combinatorial HHS for the case when \mathcal{Z} is acted on by some finitely generated group G , in the sense of Section 1.3. First, we want to show that the construction from Section 4 is G -equivariant, that is, the action of G on \mathcal{Z} induces a “compatible” action on (X, \mathcal{W}) . This will prove the “Moreover” statement of Theorem 3.16. Then, in Theorem 6.6 we will prove that, if G has a HHG structure coming from the action on \mathcal{Z} then it will also have a HHG structure coming from the action on (X, \mathcal{W}) .

The following is the combinatorial counterpart of Definition 1.13:

Definition 6.1. We say that a group G acts on the pair (X, \mathcal{W}) , where \mathcal{W} is an X -graph, if G acts on X by simplicial automorphisms, and the G -action on the set of maximal simplices of X extends to an action on \mathcal{W} .

Here we show that if we start with a G -action on $(\mathcal{Z}, \mathfrak{S})$ then the pair (X, \mathcal{W}) inherits a G -action:

Theorem 6.2. Let $(\mathcal{Z}, \mathfrak{S})$ be a HHS with clean containers. Let G be a finitely generated group acting on an HHS $(\mathcal{Z}, \mathfrak{S})$ by automorphisms. Let (X, \mathcal{W}) be the graphs constructed in Sections 4.2 and 4.3. Then G acts on (X, \mathcal{W}) . Moreover, the realisation map $f : \mathcal{W} \rightarrow \mathcal{Z}$ from Definition 4.11 is coarsely G -equivariant, meaning that for every $g \in G$ the following diagram coarsely commutes, with constants independent of g :

$$\begin{array}{ccc} \mathcal{W} & \xrightarrow{f} & \mathcal{Z} \\ \downarrow g & & \downarrow g \\ \mathcal{W} & \xrightarrow{f} & \mathcal{Z} \end{array}$$

Remark 6.3. The space $(\mathcal{Z}, \mathfrak{S})$ is only required to have clean containers since this is the only requirement to build the two graphs (X, \mathcal{W}) , as pointed out in Remark 4.1.

Proof of Theorem 6.2. The group G acts on the domain set \mathfrak{S} , preserving nesting and orthogonality, therefore there is an induced action on the minimal orthogonality graph \bar{X} , mapping every vertex U to $g^\sharp(U)$. This action extends to the blow-up graph X if, for every \sqsubseteq -minimal domain U , every point $p \in \mathcal{C}U$ and every $g \in G$ we set

$$g(p) := g^\diamond(U)(p) \in \mathcal{C}(g^\sharp(U)).$$

With a slight abuse of notation, we will denote the image of a vertex $v \in X$ under the action of an element $g \in G$ simply by gv . Similarly, we will drop the superscript for the action on \bar{X} and set $gU := g^\sharp(U)$.

Now we show that the induced action on $\mathfrak{M}(X)$ preserves \mathcal{W} -edges. Let Σ, Δ be two maximal simplices and let $\bar{\Sigma}, \bar{\Delta}$ be their supports. Let $W = (\bar{\Sigma} \cap \bar{\Delta})^\perp$, in the sense of Notation 4.9 for the exceptional cases, and for every $g \in G$ let $W' = (g\bar{\Sigma} \cap g\bar{\Delta})^\perp$. Clearly $gW = W'$ since the action preserves orthogonality, and the colevel $k = \text{co-lv}(W)$ coincides with the colevel of gW since the action preserves nesting.

Now, if Σ and Δ are \mathcal{W} -adjacent then the realisation tuples $b(\Sigma), b(\Delta)$ are $(k + 1)\lambda$ -close in every coordinate space. Thus, in order to prove that $g(\Sigma)$ and $g(\Delta)$ are again \mathcal{W} -adjacent, one needs to show that also $b(g\Sigma), b(g\Delta)$ are $(k + 1)\lambda$ -close in every coordinate space. Notice that we need the same threshold on distances, since the colevels coincide.

Recall that, if $\bar{\Delta} = \{U_1, \dots, U_k\}$ then the realisation tuple of Δ is

$$b(\Delta)_V = \begin{cases} x_i & \text{if } V = U_i \in \bar{\Delta}; \\ \bigcup_{V \not\subseteq U_i} \rho_V^{U_i} & \text{otherwise.} \end{cases}$$

If we apply g to all coordinates, we get

$$g(b(\Delta)_V) = \begin{cases} g(x_i) & \text{if } gV = gU_i \in \bar{\Delta}; \\ \bigcup_{gV \not\subseteq gU_i} g\rho_V^{U_i} = \bigcup_{gV \not\subseteq gU_i} \rho_{gV}^{gU_i} & \text{otherwise,} \end{cases}$$

where we used that, as discussed in Remark 1.14, we can assume that $\rho_{gV}^{gU} = g\rho_V^U$ for every g, U, V . On the other hand the latter expression is the realisation tuple of $g\Delta$ by definition, thus we just showed that for every $V \in \mathfrak{S}$ we have $b(g\Delta)_{gV} = g(b(\Delta)_V)$. Therefore

$$\begin{aligned} d_{gV}(b(g\Delta)_{gV}, b(g\Sigma)_{gV}) &= d_{gV}(g(b(\Delta)_V), g(b(\Sigma)_V)) \\ &= d_{CV}(b(\Delta)_V, b(\Sigma)_V), \end{aligned}$$

where we used that the map $\mathcal{C}V \rightarrow \mathcal{C}(gV)$ induced by g is an isometry. This means that if $b(\Sigma)$ and $b(\Delta)$ are $(k + 1)\lambda$ -close in every coordinate space then so are $b(g\Sigma)$ and $b(g\Delta)$, and therefore g preserves \mathcal{W} -adjacency.

Finally, in order to prove that the realisation map $f : \mathcal{W} \rightarrow \mathcal{Z}$ is coarsely G -equivariant we just note that, as proved above, for every $g \in G$ and every maximal simplex $\Delta \in \mathcal{W}^{(0)}$, the tuple $b(g\Delta)$, which coarsely coincides with the coordinates of $f(g\Delta)$, is equal to $gb(\Delta)$. Moreover, as discussed in Remark 1.14, we can assume that for every $g \in G$, $V \in \mathfrak{S}$ and $p \in \mathcal{Z}$ we have $\pi_{gV}(gp) = g\pi_V(p)$. This, applied to $p = f(\Delta)$, tells us that

$$\pi_{gV}(gf(\Delta)) = g\pi_V(f(\Delta)) \sim g(b(\Delta)_V) = b(g\Delta)_{gV} \sim \pi_{gV}(f(g\Delta)),$$

where \sim denotes equality up to a bounded error. Thus the coordinates of $gf(\Delta)$ and $f(g\Delta)$ coarsely coincide in every coordinate space, and by the uniqueness axiom we have $d_{\mathcal{Z}}(gf(\Delta), f(g\Delta)) \leq E$. \square

Next we turn our attention to actions with more structure:

Theorem 6.4 [BHMS24]. *Let (X, \mathcal{W}) be a combinatorial HHS, and let G be a group acting on X with finitely many orbits of subcomplexes of the form $\text{Lk}(\Delta)$, where Δ is a simplex of X . Suppose moreover that the action on maximal simplices of X extends to an action on \mathcal{W} , which is metrically proper and cobounded. Then G acts metrically properly and coboundedly on \mathcal{W} and with finitely many G -orbits of domains, and therefore it is a HHG.*

Proof. This is the “moreover” part of [BHMS24, Theorem 1.18]. As stated there, the theorem requires the G -action on X to be cocompact, but as discussed in [BHMS24, Remark 1.19] the proof only uses that there are finitely many G -orbits of links of simplices. \square

Definition 6.5. We will say that a group G satisfying the hypotheses of Theorem 6.4 is a *combinatorial hierarchically hyperbolic group*, meaning that the HHS structure from Definition 1.15 is inherited from the action on a combinatorial HHS.

Here we show that, under a mild assumption on the action of G on \mathfrak{S} , every HHG whose underlying HHS satisfies the hypotheses of Theorem 5.2 is a combinatorial HHG.

Theorem 6.6. *Let G be a hierarchically hyperbolic group, and let $(\mathcal{Z}, \mathfrak{S})$ be a hierarchically hyperbolic space on which G acts metrically properly and coboundedly. Suppose that $(\mathcal{Z}, \mathfrak{S})$ has weak wedges, clean containers, the orthogonals for nonsplit domains Property 3.9, and the DPR. Moreover, suppose that the action $G \curvearrowright \mathfrak{S}$ has finitely many orbits of unordered tuples $\{V_1, \dots, V_k\}$ of pairwise orthogonal elements, for every $k \leq n$. Then G acts on the pair (X, \mathcal{W}) constructed in subsections 4.2 and 4.3, and the action satisfies the hypotheses of Theorem 6.4. Hence G is a combinatorial HHG.*

Proof. Combining Theorems 6.2 and 5.2 we get that (X, \mathcal{W}) is a combinatorial HHS, that G acts on (X, \mathcal{W}) , and that the realisation map $f : \mathcal{W} \rightarrow \mathcal{Z}$ is a coarsely

G -equivariant quasi-isometry. The latter fact already implies that G acts metrically properly and coboundedly on \mathcal{W} , because the same properties hold for the G -action on \mathcal{Z} .

Then we are left to prove that the G -action on X has a finite number of orbits of the form $\text{Lk}(\Delta)$, where Δ is a simplex of X . Recall that, by Lemma 4.3, the link of Δ is given by

$$\text{Lk}(\Delta) = p^{-1}(\text{Lk}_{\bar{X}}(\bar{\Delta})) \star (\star_{U \in \bar{\Delta}^{(0)}} \text{Lk}_{p^{-1}(U)}(\Delta_U)).$$

Let $\bar{\Delta} = \{U_1, \dots, U_k\}$. Then $\text{Lk}(\Delta)$ is uniquely determined by the tuple $\{U_1, \dots, U_k\}$ and by the choice of $\text{Lk}_{p^{-1}(U_i)}(\Delta_{U_i})$ for every i . The assumption on the action of G on \mathfrak{S} tells us that there is a finite number of G -orbits of tuples of the form $\{U_1, \dots, U_k\}$, because these elements are pairwise orthogonal by definition. Moreover, given such a tuple, for every i we have $\text{Lk}_{p^{-1}(U_i)}(\Delta_{U_i})$ can only be one of the following:

- If $\Delta_{U_i} = v_{U_i}$ is the tip of the cone then $\text{Lk}_{p^{-1}(U_i)}(\Delta_{U_i}) = (CU_i)^{(0)}$ is the base of the cone.
- If Δ_{U_i} is a point in the base then $\text{Lk}_{p^{-1}(U_i)}(\Delta_{U_i}) = v_{U_i}$.
- If Δ_{U_i} is an edge then $\text{Lk}_{p^{-1}(U_i)}(\Delta_{U_i}) = \emptyset$.

Therefore there are three possible choices for every U_i , and this concludes the proof that $G \circlearrowleft X$ has finitely many orbits of the form $\text{Lk}(\Delta)$. □

Remark 6.7. The requirement that there are finitely many G -orbit of pairwise orthogonal domains (which we just refer to as *the requirement* in this Remark) is necessary for our combinatorial HHS structure to be a hierarchically hyperbolic group structure. To see this, let U_1, \dots, U_n be a collection of pairwise orthogonal, \sqsubseteq -minimal domains. Then, by how the combinatorial HHS (X, \mathcal{W}) was defined in Section 4, we see that the simplex $\Sigma = \{v_{U_1}, \dots, v_{U_n}\}$ is the link of any simplex Δ such that $\Sigma \star \Delta$ is maximal, and in particular Σ is a domain in the combinatorial HHS structure. Moreover, by definition a hierarchically hyperbolic group structure for G has finitely many G -orbits of domains, so there must be finitely many orbits of simplices corresponding to \sqsubseteq -minimal, pairwise orthogonal domains.

The requirement is easily satisfied, in the sense that, by [BGGH25, Lemma 6.6], it holds when (G, \mathfrak{S}) has the natural property that $\text{Stab}_G(V)$ acts on P_V coboundedly for all $V \in \mathfrak{S}$. On the other hand, the latter condition is stronger and is not a consequence of the definition of an HHG since, for example, one can put exotic HHG structures on a free group where this fails.

If we remove the requirement, the same proof endows a HHG (G, \mathfrak{S}) with a combinatorial G -HHS structure, which is defined as a combinatorial HHG structure except that we only require finitely many orbits of *unbounded* domains. Indeed, by

inspection of Corollary 4.4, we see that unbounded augmented links of simplices of X correspond to either \sqsubseteq -minimal domains in \mathfrak{S} , or to the orthogonal complements of certain collections of domains in \mathfrak{S} . In particular, unbounded domains in (X, \mathcal{W}) correspond to certain domains in \mathfrak{S} , of which there are finitely many G -orbits since (G, \mathfrak{S}) is a HHG.

7. Some other hypotheses

The hypotheses of Theorem 3.16 are rather technical, so in this section we present some more “natural” ones, and we describe how they relate to each other and to the original ones. We also show that, in certain cases, one can match the requirements of Theorem 3.16 by adding (quite a lot of) bounded coordinate spaces.

We start with some possible requirements on the domain set.

Property 7.1. A hierarchically hyperbolic space has the *strong orthogonal property* if for every two domains $U \not\sqsubseteq V$ there exists $W \not\sqsubseteq V$ such that $U \perp W$.

Property 7.2. A hierarchically hyperbolic space has the *weak orthogonal property* if for every two domains $U \not\sqsubseteq V$, if U is non- \sqsubseteq -minimal there exists $W \not\sqsubseteq V$ such that $U \perp W$.

Both these properties imply orthogonals for nonsplit domains (Property 3.9), since every \sqsubseteq -minimal domain is trivially split (it coincides with its unique Samaritan).

Property 7.3. A hierarchically hyperbolic space $(\mathcal{Z}, \mathfrak{S})$ has *bounded split coordinate spaces* if there exists $c \geq 0$ such that, for every $U \in \mathfrak{S} - S$ which is split and non- \sqsubseteq -minimal, the corresponding coordinate space $\mathcal{C}U$ has diameter at most c .

Lemma 7.4 (comparison between properties). *Let $(\mathcal{Z}, \mathfrak{S})$ be a normalised hierarchically hyperbolic space with wedges and clean containers. Then each of the following properties implies the lower ones, meaning that if $(\mathcal{Z}, \mathfrak{S})$ has some property (i) from the list, and $j > i$, then there exists a normalised HHS structure $(\mathcal{Z}, \mathfrak{S}')$, with $\mathfrak{S} \subset \mathfrak{S}'$ and satisfying wedges, clean containers and property (j):*

1. $(\mathcal{Z}, \mathfrak{S})$ has the *strong orthogonal property* (7.1).
2. $(\mathcal{Z}, \mathfrak{S})$ has the *weak orthogonal property* (7.2).
3. $(\mathcal{Z}, \mathfrak{S})$ has the *weak orthogonal property* (7.2) and *dense product regions* (3.10).
4. $(\mathcal{Z}, \mathfrak{S})$ has *orthogonals for nonsplit domains* (3.9) and *bounded split coordinate spaces* (7.3).
5. $(\mathcal{Z}, \mathfrak{S})$ has *orthogonals for nonsplit domains* (3.9) and *dense product regions* (3.10).

Proof of Lemma 7.4. The implication $1 \Rightarrow 2$ is trivial. $3 \Rightarrow 4$ follows from Lemma 7.5 and the fact that \sqsubseteq -minimal domains are split. The implication $2 \Rightarrow 3$ is Lemma 7.7, while Lemma 7.6 is $4 \Rightarrow 5$. \square

Lemma 7.5. *If a normalised HHS $(\mathcal{Z}, \mathfrak{S})$ has the DPR then it has the bounded split coordinate spaces property (7.3).*

We should think of the bounded split coordinate space property as the fact that product regions are dense in split domains. This is, in essence, why Lemma 7.5 holds.

Proof. Let U be a nonminimal split domain and let W be one of its Samaritans. By the DPR property, for every $q \in \mathcal{C}U$ there exists a minimal domain $V \not\sqsubseteq U$ such that $d_{\mathcal{C}U}(q, \rho_U^V) \leq M_0$. Now it suffices to notice that, since W is a Samaritan, we must have either $V = W$ or $V \perp W$, and by Lemma 1.3 we have $\text{diam}_{\mathcal{C}U}(\rho_U^W, \rho_U^V) \leq 10E$. Therefore q is $(M_0 + 10E)$ -close to ρ_U^W , and therefore $\mathcal{C}U$ has diameter at most $2(M_0 + 10E)$. \square

The following lemmas show that, if one starts with more general hypotheses, there is often a way to modify the HHS structure to ensure the DPR property.

Lemma 7.6. *Let $(\mathcal{Z}, \mathfrak{S})$ be a normalised HHS with wedges, clean containers, the bounded split coordinate spaces property (7.3) and the orthogonals for nonsplit domains property (3.9). Then there exists a normalised HHS structure $(\mathcal{Z}, \mathfrak{S}')$ such that $\mathfrak{S} \subseteq \mathfrak{S}'$ and $(\mathcal{Z}, \mathfrak{S}')$ has wedges, clean containers, the DPR and the orthogonals for nonsplit domains property.*

Lemma 7.7. *Let $(\mathcal{Z}, \mathfrak{S})$ be a normalised HHS with wedges, clean containers and the weak orthogonal property (7.2). Then there exists a normalised HHS structure $(\mathcal{Z}, \mathfrak{S}')$ such that $\mathfrak{S} \subseteq \mathfrak{S}'$ and $(\mathcal{Z}, \mathfrak{S}')$ has wedges, clean containers, the weak orthogonal property and the DPR.*

The strategy for proving these two lemmas is the same, so we present an extensive proof only of Lemma 7.6, which is more complicated. Here is a list of the changes needed to prove Lemma 7.7:

- A domain T_x^U must be added for every nonminimal U (thus not only if U is nonsplit).
- The argument below to show the DPR property for wide domains will apply to all non- \sqsubseteq -minimal domains.
- Since \mathfrak{S} and \mathfrak{S}' will have the same non- \sqsubseteq -minimal domains, the weak orthogonal property will be preserved.

Proof of Lemma 7.6. What we will actually prove is that, if a normalised HHS $(\mathcal{Z}, \mathfrak{S})$ has the bounded split coordinate space property (7.3), then we can find a structure $(\mathcal{Z}, \mathfrak{S}')$ with the DPR property. Moreover, if $(\mathcal{Z}, \mathfrak{S})$ has wedges, clean containers or the orthogonals for nonsplit domains property then the procedure will preserve these properties.

We will say that a domain $U \in \mathfrak{S}$ is *wide* if U is nonsplit (and in particular non- \sqsubseteq -minimal) or $U = S$. For every wide domain U we do the following. Recall that we defined F_U as the space of all $20E$ -consistent tuples. For every $x = (x_V)_{V \sqsubseteq U} \in F_U$ we define a domain T_x^U whose coordinate space $\mathcal{C}T_x^U$ is a point, and we let \mathfrak{S}' be the union of \mathfrak{S} and these new domains. Now we show that $(\mathcal{Z}, \mathfrak{S}')$ is an HHS, defining new projections, relative projections and relations when needed.

Projections: The projection $\pi_{T_x^U} : \mathcal{Z} \rightarrow \mathcal{C}T_x^U$ is just the constant map, while all other projections are inherited from the original structure.

Nesting: The domains T_x^U are \sqsubseteq -minimal, and $T_x^U \sqsubseteq V$ if and only if $U \sqsubseteq V$. If $U \not\sqsubseteq V$ then we define the projection $\rho_V^{T_x^U} = \rho_V^U$; moreover we set $\rho_U^{T_x^U} = x_U$. The projections in the opposite direction (namely, $\rho_{T_x^U}^V$ whenever $U \sqsubseteq V$) are just the constant maps.

Finite complexity: Since we just added \sqsubseteq -minimal domains inside non- \sqsubseteq -minimal ones, the complexity of the HHS structure remains the same.

Orthogonality and clean containers: If $U, V \in \mathfrak{S}$, U is wide and $x \in F_U$, we say that $V \perp T_x^U$ if and only if $U \perp V$. If V is also wide and $y \in F_V$ we say that $T_x^U \perp T_y^V$ if and only if $U \perp V$. (See Figure 6.)

Notice that containers already exist for any situation involving elements of $\mathfrak{S}' - \mathfrak{S}$. Indeed, suppose $T_x^U \sqsubseteq W$ for some $W \in \mathfrak{S}$, which implies that either $U = W$ or $U \not\sqsubseteq W$. In the first case, no container is needed, since T_x^U is transverse to every $V \not\sqsubseteq U$ (and therefore also to every T_y^V if V is wide). In the second case, every $V \in \mathfrak{S}$ which is properly nested inside W and orthogonal to T_x^U is also orthogonal to U , and therefore already nested inside the container for U inside W . If $T_y^V \not\sqsubseteq W$ and $T_x^U \perp T_y^V$ then $V \not\sqsubseteq W$ and $U \perp W$, which means that T_y^V is already nested inside the container for U inside W .

Conversely, if $V \in \mathfrak{S}$ and T_x^U are orthogonal and properly nested in W then $U \not\sqsubseteq W$ and $U \perp V$, thus T_x^U is already nested inside the container for V inside W .

Now, if $(\mathcal{Z}, \mathfrak{S})$ has clean containers then so does $(\mathcal{Z}, \mathfrak{S}')$. This is because, as




Figure 6. Illustrating the only way two newly added domains T_x^U and T_y^V can be orthogonal inside some $W \in \mathfrak{S}$. (Arrows denote nesting.)

argued above, the container for T_x^U inside some W is the container for U inside W , and if this container is orthogonal to U then it is also orthogonal to T_x^U .

Transversality: If $U, V \in \mathfrak{S}$, U is wide and $x \in F_U$, then by construction $T_x^U \pitchfork V$ if and only if one of the following holds:

- $U \pitchfork V$: in this case we define $\rho_{V^x}^{T_x^U} = \rho_V^U$.
- $V \not\pitchfork U$: in this case we set $\rho_{V^x}^{T_x^U} = x_V$.

Moreover, whenever $Y \in \mathfrak{S}'$ is transverse to T_x^U we set $\rho_{T_x^U}^Y = \mathcal{C}T_x^U$.

Consistency: Since the only elements of \mathfrak{S}' whose coordinate spaces are not points are in \mathfrak{S} , the first two consistency inequalities are trivial. The final clause of the consistency axiom holds, since by definition $\rho_{V^x}^{T_x^U} = \rho_V^U$ whenever $U \not\pitchfork V$ or $U \pitchfork V$.

Uniqueness, BGI, large links: Since the only elements of \mathfrak{S}' whose coordinate spaces are not points are in \mathfrak{S} , these axioms for $(\mathcal{Z}, \mathfrak{S}')$ follow from the corresponding ones for $(\mathcal{Z}, \mathfrak{S})$.

Partial realisation: Let $V_1, \dots, V_k \in \mathfrak{S}'$ be pairwise orthogonal elements, and let $p_i \in \mathcal{C}V_i$. We show that we can find a partial realisation point for $\{(V_i, p_i)\}$. Up to permutation we can assume that $V_1, \dots, V_l \in \mathfrak{S}$, for some $l \leq k$, and $V_i = T_{x_i}^{U_i}$ for every $l < i \leq k$. Moreover, we can assume that the family V_1, \dots, V_k is maximal, up to adding domains belonging to \mathfrak{S} , since a realisation point for a bigger family is also a realisation point for the original one.

Now, since \mathcal{Z} is normalised, for every $i \leq l$ we can find $y_i \in F_{V_i}$ such that $d_{V_i}(\pi_{V_i}(y_i), p_i)$ is uniformly bounded. Then, for any $U \in \mathfrak{S}'$ set

$$(3) \quad q_U = \begin{cases} (y_i)_U & \text{if } U \sqsubseteq V_i, i \leq l; \\ (x_i)_U & \text{if } U \sqsubseteq U_i, l < i \leq k; \\ \bigcup_{V_i \pitchfork U \text{ or } V_i \not\pitchfork U} \rho_U^{V_i} & \text{otherwise.} \end{cases}$$

Since V_1, \dots, V_k is a maximal family, the coordinate q_U is always well-defined. Moreover, it is easy to see that $(q_U)_{U \in \mathfrak{S}'}$ is $20E$ -consistent, thus by the realisation Theorem 1.7 there exists $z \in \mathcal{Z}$ whose coordinates are uniformly close to $(q_U)_{U \in \mathfrak{S}'}$. It is also easy to verify that z is a partial realisation point for $\{(V_i, p_i)\}$, since by equation (3) it has the right coordinates whenever $U = V_i$, $V_i \pitchfork U$ or $V_i \not\pitchfork U$ for some $i \leq k$.

DPR: Let $U \in \mathfrak{S}$ be a nonminimal domain and let $p \in \mathcal{C}U$. If U is nonsplit then there is $z \in \mathcal{Z}$ such that $\pi_U(z)$ is uniformly close to p , by our normalisation assumption. Thus, if we set $x = (\pi_V(z))_{V \sqsubseteq U}$, then T_x^U projects uniformly close to p in $\mathcal{C}U$, and the DPR property holds with the same constant as the one coming from the normalisation assumption. On the other hand, if U is split then $\mathcal{C}U$ is uniformly bounded by the bounded split coordinate space Property 7.3. Hence $p \in \mathcal{C}U$ is uniformly close to ρ_U^V for any $V \not\pitchfork U$.

Additional properties: We now check that if $(\mathcal{Z}, \mathfrak{S})$ has one of the properties below then so does $(\mathcal{Z}, \mathfrak{S}')$.

- Wedges: The new domains are all \sqsubseteq -minimal; therefore for every $T_x^U, V' \in \mathfrak{S}'$ we can set

$$T_x^U \wedge V' = \begin{cases} T_x^U & \text{if } T_x^U \sqsubseteq V'; \\ \emptyset & \text{otherwise.} \end{cases}$$

Then we just need to verify that, if $V, W \in \mathfrak{S}$ and there exists $R \in \mathfrak{S}'$ which is nested inside both, then the wedge of V and W exists in \mathfrak{S}' . What we will actually show is that V and W already have a well-defined wedge inside \mathfrak{S} , call this wedge Q , and that whenever $R \in \mathfrak{S}'$ is nested inside both V and W then $R \sqsubseteq Q$. Therefore the wedge in \mathfrak{S}' will coincide with the wedge in \mathfrak{S} .

Now let $R \in \mathfrak{S}'$ as above. If $R \in \mathfrak{S}$ then Q exists and $R \sqsubseteq Q$, by the wedge property for \mathfrak{S} . Otherwise $R = T_x^U \not\sqsubseteq U$ for some U and some $x \in F_U$, and by definition of the new nesting relations U must be already nested in both V and W . Hence again Q exists and $R \sqsubseteq U \sqsubseteq Q$.

- Orthogonals for nonsplit domains property (3.9): Let $U \not\sqsubseteq V$, with $V \in \mathfrak{S}$ since the new domains are all minimal. If U is one of the new domains then U is \sqsubseteq -minimal, and therefore split. Hence, suppose that $U \in \mathfrak{S}$. If U is nonsplit in \mathfrak{S} then Property 3.9, which holds for \mathfrak{S} , ensures the existence of some $W \perp U$ inside V . Otherwise U is split in \mathfrak{S} with some Samaritan W , and we claim that U is again split in \mathfrak{S}' with the same Samaritan. Let $Q \in \mathfrak{S}'$ be such that $Q \sqsubseteq U$. If $Q \in \mathfrak{S}$ then $W \sqsubseteq Q$ or $W \perp Q$, by definition of Samaritan, and we have nothing to prove. Otherwise we have $Q = T_x^R \sqsubseteq R \sqsubseteq U$ for some nonsplit domain $R \in \mathfrak{S}$ and some $x \in F_R$. We cannot have $W \sqsubseteq R$, since then R would be split in \mathfrak{S} by Remark 3.7; thus R (and therefore Q) is orthogonal to W .

This concludes the proof. □

Remark 7.8. The weak wedge property is not preserved by the procedure of Lemma 7.6 to ensure DPR, since we are adding many “loose” \sqsubseteq -minimal domains that will not be nested in the original weak wedges. This is why in this section we are assuming the “strong” wedge property.

Remark 7.9 (DPR for HHG with cobounded actions). If G is a HHG, but the underlying space $(\mathcal{Z}, \mathfrak{S})$ does not have the DPR property already, we cannot argue as in Lemma 7.6 to enforce it. This is because, if we add a domain of the form T_x^U whenever $x \in F_U$ and we define the G -action on \mathfrak{S}' in the obvious way (that is, by setting $gT_x^U = T_{gx}^{gU}$), then this action cannot have finitely many G -orbits of domains, since G is countable while F_U might be uncountable.

Nonetheless, the DPR property is fairly common for HHG. For instance, it is implied by the fact that, for every $U \in \mathfrak{S}$, the stabiliser $\text{Stab}_G(U)$ acts coboundedly

on \mathcal{CU} . Indeed, if V is \sqsubseteq -minimal and $V \not\sqsubseteq U$, then the collection $\{\rho_U^{gV}\}_{g \in \text{Stab}_G(U)}$ coincides with the $\text{Stab}_G(U)$ -orbit of ρ_U^V , by definition of a HHG, and the latter is coarsely dense in \mathcal{CU} by coboundedness.

In turn, if the HHS structure $(\mathcal{Z}, \mathfrak{S})$ is normalised, we can further reduce to checking the weaker requirement that, for every $U \in \mathfrak{S}$, the stabiliser $\text{Stab}_G(U)$ acts coboundedly on the product region P_U . Indeed, the projection $\pi_U : P_U \rightarrow \mathcal{CU}$ is coarsely surjective (by normalisation), coarsely Lipschitz (by Definition 1.1 of a HHS) and $\text{Stab}_G(U)$ -equivariant (by Definition 1.13 of the G -action), therefore if $\text{Stab}_G(U)$ acts coboundedly on P_U then it also acts coboundedly on \mathcal{CU} .

In practice, the latter requirement is not very restrictive, since all “reasonable” HHGs have cobounded actions on their product regions, and all known methods of producing new HHGs tend to preserve this property (unless one comes up with some very artificial structures).

7.1. Orthogonal sets. Another property that one could require on the index set is that orthogonal complementation is an involution. Such a property is satisfied by CAT(0) cube complexes with (weak) factor systems (see Section 10.1) and is implied by the strong orthogonality property (7.1) (Lemma 7.12). However, it goes in a somewhat different direction than the orthogonals for nonsplit domains property (3.9), in the sense that is not enough to prove Theorem 3.16. Indeed, in Section 10.2 we will provide an example of an HHS \mathcal{Z} with wedges, clean containers, dense product regions and where orthogonal complementation is an involution, but such that the graph X from Definition 4.2 cannot be the support of a CHHS structure for \mathcal{Z} .

Definition 7.10. A partially ordered set $(\mathfrak{F}, \sqsubseteq)$ is called *orthogonal* if there exists a symmetric relation \perp on \mathfrak{F} such that the following hold for all $U, V, W \in \mathfrak{F}$:

- $U \not\perp U$.
- \mathfrak{F} has a unique \sqsubseteq -maximal element S .
- If $U \sqsubseteq V$ and $V \perp W$ then $U \perp W$.
- (**wedges**) If $W \sqsubseteq U, V$, then there exists $U \wedge V \in \mathfrak{F}$ such that $U \wedge V \sqsubseteq U, V$, and for all $W \sqsubseteq U, V$ we have that $W \sqsubseteq U \wedge V$.
- (**clean containers**) For all U such that there exists $V \perp U$, there exists $U^\perp \in \mathfrak{F}$ such that, for all $V \perp U$ we have $V \sqsubseteq U^\perp$, and $W \perp U^\perp$ if and only if $W \sqsubseteq U$.
- (**orthogonality determines nesting**) $U \sqsubseteq V$ (resp. $U \not\sqsubseteq V$) if and only if the set of W for which $V \perp W$ is contained (resp. properly contained) in the set of W' for which $U \perp W'$. In particular, if nothing is \perp -related to V then V is the unique \sqsubseteq -maximal element, while if there exists $W \perp V$ then $V^\perp \sqsubseteq U^\perp$ (resp. $V^\perp \not\sqsubseteq U^\perp$) if and only if $U \sqsubseteq V$ (resp. $U \not\sqsubseteq V$).

Lemma 7.11 (orthogonality determines nesting for HHS). *Let $(\mathcal{Z}, \mathfrak{S})$ be an HHS with wedges and clean containers. Then the following are equivalent:*

- **(complementation is an involution)** *For all $U \in \mathfrak{S} - \{S\}$, U^\perp is defined and $U^{\perp\perp} = U$.*
- **(orthogonality determines nesting)** *For all $U, V \in \mathfrak{S} - \{S\}$, we have $U \not\sqsubseteq V$ if and only if $V^\perp \not\sqsubseteq U^\perp$.*

If one (hence both) of the previous holds, then \mathfrak{S} is an orthogonal set.

Proof. The first part of the lemma is proven exactly as [CRHK24, Proposition 6.1] (the statement there is for real cubings, but as pointed out in [CRHK24, Remark 6.2] the same argument works for HHSs). The second part follows from the properties of the domain set \mathfrak{S} of a HHS (see Definition 1.1). □

Lemma 7.12. *Let $(\mathcal{Z}, \mathfrak{S})$ be an HHS with wedges, clean containers and the strong orthogonal Property 7.1. Then \mathfrak{S} is an orthogonal set.*

Proof. For every $U \in \mathfrak{S} - \{S\}$, the strong orthogonal property grants the existence of U^\perp ; moreover $U = U^{\perp\perp}$, because otherwise we could find a $V \sqsubseteq U^{\perp\perp}$ which is orthogonal to U , and this would contradict the definition of U^\perp . Now the conclusion follows from Lemma 7.11. □

8. Near equivalence of HHS and combinatorial HHS

In this section we show that, if the hypotheses on $(\mathcal{Z}, \mathfrak{S})$ are the strongest possible, then the combinatorial HHS (X, \mathcal{W}) arising from the construction has the following two nice properties:

Definition 8.1. A combinatorial HHS (X, \mathcal{W}) has *simplicial containers* if for any simplex $\Delta \subset X$ there exists a simplex $\Phi \subset X$ such that

$$\text{Lk}(\text{Lk}(\Delta)) = \text{Lk}(\Phi).$$

Definition 8.2. A combinatorial HHS (X, \mathcal{W}) has *simplicial wedges* if for any two simplices $\Delta, \Sigma \subset X$ there exists a simplex Π which extend Σ such that

$$\text{Lk}(\Delta) \cap \text{Lk}(\Sigma) = \text{Lk}(\Pi).$$

Theorem 8.3. *Let $(\mathcal{Z}, \mathfrak{S})$ be a normalised hierarchically hyperbolic space. Then \mathcal{Z} has wedges, clean containers and the strong orthogonal property (7.1) if and only if there exists a CHHS (X, \mathcal{W}) with simplicial wedges and simplicial containers such that \mathcal{W} is quasi-isometric to \mathcal{Z} .*

Proof. First we show that, if (X, \mathcal{W}) has simplicial wedges and simplicial containers then \mathcal{W} , with the HHS structure described in Section 2.2, has the following properties:

- (**wedges**) Given two nonmaximal simplices Σ, Δ , if there exists a simplex Γ such that $[\Gamma] \sqsubseteq [\Sigma]$ and $[\Gamma] \sqsubseteq [\Delta]$, then

$$\text{Lk}(\Gamma) \sqsubseteq \text{Lk}(\Sigma) \cap \text{Lk}(\Delta) = \text{Lk}(\Pi)$$

for some Π depending only on Σ, Δ . Therefore $[\Sigma] \wedge [\Delta] = [\Pi]$.

- (**strong orthogonal property**) Let $[\Delta] \not\sqsubseteq [\Delta']$. Now

$$\text{Lk}(\text{Lk}(\Delta)) \cap \text{Lk}(\Delta') = \text{Lk}(\Phi) \cap \text{Lk}(\Delta') = \text{Lk}(\Pi)$$

for some simplices Φ, Π whose existence is granted by simplicial containers and wedges, respectively. Thus $[\Pi] \sqsubseteq [\Delta']$, and since $\text{Lk}(\Pi) \subseteq \text{Lk}(\text{Lk}(\Delta))$ we have $[\Pi] \perp [\Delta]$.

- (**clean containers**) Let $[\Delta] \not\sqsubseteq [\Delta']$ and suppose that there exists $[\Sigma] \not\sqsubseteq [\Delta']$ which is orthogonal to $[\Delta]$. By definition

$$\text{Lk}(\Sigma) \subseteq \text{Lk}(\text{Lk}(\Delta)) \cap \text{Lk}(\Delta') = \text{Lk}(\Pi),$$

where Π is the simplex defined above. Then $[\Sigma] \not\sqsubseteq [\Pi]$, which means that $[\Pi]$ is the container for $[\Delta]$ inside $[\Delta']$. Since $[\Pi] \perp [\Delta]$, this container is also clean.

Now we turn our attention to the converse statement. Let $(\mathcal{Z}, \mathfrak{S})$ be a HHS with wedges, clean containers and the strong orthogonal property. By Lemma 7.7 we can find a structure $(\mathcal{Z}, \mathfrak{S}')$ with wedges, clean, containers, the weak orthogonal property and the DPR property. Then Theorem 3.16 applies to $(\mathcal{Z}, \mathfrak{S}')$ and outputs a combinatorial HHS (X, \mathcal{W}) where \mathcal{W} is quasi-isometric to \mathcal{Z} .

- (**(X, \mathcal{W}) has simplicial wedges**) Let $\bar{\Sigma}$ and $\bar{\Delta}$ be the supports of Σ and Δ , respectively, and let $\bar{\Sigma}^\perp, \bar{\Delta}^\perp \in \mathfrak{S}$ be their orthogonal complements. Let $\bar{\Phi} = \bar{\Delta} \cap \text{Lk}(\bar{\Sigma})$, and let Y_0 be the orthogonal complement of $\bar{\Phi}$ inside $\bar{\Sigma}^\perp$, that is, $Y_0 = (\bar{\Sigma} \star \bar{\Phi})^\perp$. Notice that Y_0 cannot be one of the minimal domains T_x^U that were added in Lemma 7.7 to ensure the DPR property. This is because, if $Y_0 = T_x^U$, then T_x^U is orthogonal to $\bar{\Sigma} \star \bar{\Phi}$, and by construction U is orthogonal to the same simplex. But by the definition of Y_0 we must have $U \sqsubseteq Y_0 = T_x^U \not\sqsubseteq U$, which is a contradiction.

Part 1: If $(\bar{\Sigma})^\perp$ and $(\bar{\Delta})^\perp$ don't have any nested domain in common, we formally set $W_0 = \emptyset$ and we skip to Part 2. Otherwise we can consider the wedge $W_0 = (\bar{\Sigma})^\perp \wedge (\bar{\Delta})^\perp$. Notice that W_0 cannot be one of the minimal domains from Lemma 7.7. This is because, if $W_0 = T_x^U$, then T_x^U is orthogonal to both $\bar{\Sigma}$ and $\bar{\Delta}$, and by construction U is also orthogonal to the two simplices. But by definition of W_0 as a wedge we must have $U \sqsubseteq W_0 = T_x^U \not\sqsubseteq U$, which is a contradiction. This means that $W_0 \in \mathfrak{S}$ as well.

Part 2: Now suppose that W_0 has been defined as in Part 1. If $W_0 = Y_0$ then we set $\bar{\Theta} = \emptyset$. Otherwise, by the strong orthogonal property, which holds for every two elements of the original domain set \mathfrak{S} , we can find a \sqsubseteq -minimal domain $V_0 \sqsubseteq Y_0$ such that $V_0 \perp W_0$ (if $W_0 \neq \emptyset$). Now let $Y_1 = \{V_0\}_{Y_0}^\perp$, which is again in \mathfrak{S} and contains W_0 , and we can argue as above.

In both cases we can find a (possibly empty) simplex $\bar{\Theta} = \{V_0, \dots, V_k\}$ of \bar{X} such that $W_0 = (\bar{\Sigma} \star \bar{\Phi} \star \bar{\Theta})^\perp$, and this readily implies that

$$(4) \quad \text{Lk}(\bar{\Sigma}) \cap \text{Lk}(\bar{\Delta}) = \text{Lk}(\bar{\Sigma} \star \bar{\Phi} \star \bar{\Theta}).$$

From now on we can argue exactly as in Lemma 5.7 (in the paragraph “Finding the extension of Σ ”), and find a simplex supported in $\bar{\Sigma} \star \bar{\Phi} \star \bar{\Theta}$ which extends Σ and whose link is $\text{Lk}(\Sigma) \cap \text{Lk}(\Delta)$.

• ***((x, W) has simplicial containers)*** Let $\Delta \subseteq X$ be a simplex and let $\bar{\Delta}$ be its support. Let $\bar{\Delta}_1 \subset \bar{\Delta}$ be the domains U such that Δ_U is a single point, and let $\bar{\Delta}_2 \subset \bar{\Delta}$ be the domains such that Δ_U is an edge. Taking the link of the expression in Lemma 4.3, which described the shape of the link of Δ inside X , we have

$$\text{Lk}(\text{Lk}(\Delta)) = p^{-1}(\text{Lk}(\text{Lk}(\bar{\Delta}))) \cap \bigcap_{U \in \bar{\Delta}_1} \text{Lk}_X(\text{Lk}_{p^{-1}(U)}(\Delta_U)).$$

Thus $v \in X$ belongs to $\text{Lk}(\text{Lk}(\Delta))$ if its support V lies in $(\text{Lk}(\text{Lk}(\bar{\Delta})))$, and either $V \perp \bar{\Delta}_1$, or $V \in \bar{\Delta}_1$ and $v \in \text{Lk}_{p^{-1}(V)}(\text{Lk}_{p^{-1}(V)}(\Delta_V))$. Therefore, let $W = (\bar{\Delta}^\perp)^\perp$. Again, W does not coincide with any T_x^U , because if $T_x^U = W$ then $T_x^U \perp \bar{\Delta}^\perp$, and therefore also $U \perp \bar{\Delta}^\perp$. Then again $W \in \mathfrak{S}$, and the strong orthogonal property implies that there exists a simplex $\bar{\Theta}$ inside \bar{X} such that $W = \bar{\Theta}^\perp$. Notice that every $U \in \bar{\Delta}$ is nested in W by construction. Now define a simplex $\bar{\Phi}$ with support $\bar{\Theta} \star \bar{\Delta}_1$ by choosing an edge for every domain U which is a vertex of $\bar{\Theta}$, and a point $q_U \in \text{Lk}_{p^{-1}(U)}(\Delta_U)$ for every $U \in \bar{\Delta}_1$. Thus by construction

$$\begin{aligned} \text{Lk}(\bar{\Phi}) &= p^{-1}(\text{Lk}(\bar{\Theta} \star \bar{\Delta}_1)) \star \bigcup_{U \in \bar{\Delta}_1} \text{Lk}_{p^{-1}(U)}(q_U) \\ &= p^{-1}(\text{Lk}(\text{Lk}(\bar{\Delta}))) \cap \bigcap_{U \in \bar{\Delta}_1} \text{Lk}_X(\text{Lk}_{p^{-1}(U)}(\Delta_U)) = \text{Lk}(\text{Lk}(\Delta)), \end{aligned}$$

and we are done. □

Remark 8.4. If in the previous proof the original structure $(\mathcal{Z}, \mathfrak{S})$ already has the DPR then there is no need to invoke Lemma 7.7, and the whole proof of Theorem 8.3 works with \mathfrak{S} instead of \mathfrak{S}' . In other words, if $(\mathcal{Z}, \mathfrak{S})$ already has the DPR property then the combinatorial HHS (X, \mathcal{W}) is exactly the one constructed in Section 4.

9. Mapping class groups are combinatorial HHS

Throughout this section, let S be a surface obtained from a closed, connected, oriented surface after removing a finite number of points and open disks; we call such an S a surface of finite-type with boundary.

It was proven in e.g. [BHS19, Theorem 11.1] that, if S has no boundary, then it admits a HHG structure. In this section we first extend this result to surfaces of finite-type with boundary (see Remark 9.3); then we apply our main theorems to produce two combinatorial HHG structures, one whose underlying graph is a blow-up of the curve graph (Theorem 9.8), and one with combinatorial wedges and combinatorial containers (Theorem 9.9).

9.1. On the meaning of subsurface. The “usual” HHG structure for a mapping class group involves *open* subsurfaces, but it will be convenient to consider a more general type of subsurfaces; here we present the two notions and compare them.

Definition 9.1 [BKMM12, Section 2.1.3]. A subsurface $Y \subset S$ is *essential* if it is the disjoint union of some components of the complement of a collection of disjoint simple closed curves, so that no component is a pair of pants and no two annular components are isotopic.

Recall that nesting of subsurfaces is defined as follows: U is nested in V if U is contained in V (up to isotopy) and no isotopy class representative of U is disjoint from an isotopy class representative of V . (This last clause is only relevant for annuli and unions of annuli; an annulus might be isotopic to a nonessential annulus of another subsurface.)

Theorem 9.2 [BHS19, Theorem 11.1]. *Let S be a surface of finite-type with boundary. Its mapping class group $\mathcal{MCG}(S)$ is an HHG with the following structure:*

- \mathfrak{S} is the collection of isotopy classes of essential subsurfaces.
- For each $U \in \mathfrak{S}$ the space \mathcal{CU} is its curve graph.
- The relation \sqsubseteq is nesting, \perp is disjointness and \pitchfork is overlapping.
- For each $U \in \mathfrak{S}$, the projection $\pi_U : \mathcal{MCG}(S) \rightarrow \mathcal{CU}$ is constructed using the subsurface projection.
- For $U, V \in \mathfrak{S}$ satisfying either $U \not\sqsubseteq V$ or $U \pitchfork V$, the projection is $\rho_V^U = \pi_V(\partial U) \subset \mathcal{CV}$, while for $V \sqsubseteq U$ the map $\rho_V^U : \mathcal{CU} \rightarrow 2^{\mathcal{CV}}$ is the subsurface projection.

Remark 9.3. In previous literature, the HHG structure above is only considered for surfaces without boundary, but everything goes through as above for surfaces with boundary (including braid groups). There are a few ways to see this, besides inspecting [BHS19, Section 11]. One is to regard $\mathcal{MCG}(S)$ as above as a subgroup of the mapping class group of the double of S along all boundary components,

by extending mapping classes to be the identity on the complement of S . In this setting, $\mathcal{MCG}(S)$ acts properly and coboundedly on F_S (with finitely many orbits of subsurfaces nested into S), giving the required structure.

We now describe a different HHS structure, whose index set is made of subsurfaces which might include some of their boundary components, and then discuss how it relates to the one from Theorem 9.2.

The new index set.

Definition 9.4 (block). A *block* is (the isotopy class of) a subsurface of one of the following types:

- (a) a closed annulus which does not bound a disk or a single puncture;
- (b) a connected, nonannular subsurface of complexity at least 1, with some (possibly none) of its boundary components included.

The *included boundary* of a block U of type (b) is the set of curves in its topological boundary (relative to S) which belong to U . With a slight abuse of notation, we say that the included boundary of an annulus is its core curve.

Given two blocks U and V , we say $U \sqsubseteq V$ if U can be isotoped inside V , and every included boundary component of U which is homotopic to a component of the topological boundary of V is also contained in the included boundary of V . Let \perp be disjointness of blocks (up to isotopy), but with these conventions:

- if γ is a curve in the topological boundary of U , then $\gamma \perp U$ if and only if γ does not belong to the included boundary of U ;
- if U and V are blocks of type (b), and they share a component of the included boundary, then they are not orthogonal.

This way, nesting and orthogonality are mutually exclusive, meaning that if two blocks are orthogonal then they are not \sqsubseteq -related.

Definition 9.5 (admissible collection). A collection of pairwise orthogonal blocks U_1, \dots, U_k is *admissible* if, whenever two blocks U_i and U_j (which might coincide) have two topological boundary components which are isotopic (and are distinct if $U_i = U_j$), then none of these components belongs to the included boundary of the respective block. See Figure 7 to understand the forbidden cases.

Let \mathfrak{S}' be the set of all admissible collections of blocks, which we see as subsurfaces of S (up to isotopy). Extend nesting and orthogonality to \mathfrak{S}' , with the same conventions about boundary curves. Note that \sqsubseteq -minimal elements of \mathfrak{S}' are exactly annuli of type (a).

Define the *interior* of an admissible collection as the union of the interiors of its blocks of type (b). Moreover, define the *included boundary* of an admissible




Figure 7. Forbidden and allowed collections of blocks. Dashed lines represent open boundaries (those which do not contain the boundary curve), while full lines represent included boundaries. The forbidden cases are as follows: no two blocks (which might coincide) have isotopic boundary components, unless both of them are not included (top row); no annulus is isotopic to a boundary component of another block (bottom row).

collection as the union of the included boundaries of its blocks. Notice that if $U, V \in \mathfrak{S}'$ then $U \sqsubseteq V$ if and only if the interior of U is nested in the interior of V , and the included boundary of U is nested in the included boundary of V .

Coordinate spaces. Let V be a block. Define $\mathcal{C}V$ as follows:

- If V is a closed, essential annulus, let $\mathcal{C}V$ be its annular curve graph.
- If V is a connected, open subsurface of complexity at least 1, let $\mathcal{C}V$ be its curve graph.
- If V is of type (b) and its included boundary is nonempty, let $\mathcal{C}V$ be the join of the curve graph of its interior and the included boundary.

The coordinate space of an admissible collection is, by definition, the join of the coordinate spaces of its blocks.

Remark 9.6. If $U \in \mathfrak{S}'$, then either U is an annulus or a connected, open subsurface of complexity at least 1, or $\mathcal{C}U$ is uniformly bounded since it is a join.

Projections. For every $U \in \mathfrak{S}'$, let $\pi_U : \mathcal{MCG}(S) \rightarrow \mathcal{C}U$ be the subsurface projection. Analogously, for every $U, V \in \mathfrak{S}'$ such that $U \not\sqsubseteq V$, let $\rho_U^V : \mathcal{C}V \rightarrow \mathcal{C}U$ be the subsurface projection. Moreover, if $\mathcal{C}V$ is bounded then define ρ_V^U by choosing any point in $\mathcal{C}V$; otherwise V must be a connected, essential subsurface, and therefore at least one of the topological boundary curves of U is nested in V , so one can

set $\rho_V^U = \partial U \cap V$. Finally, for every $U, V \in \mathfrak{S}'$ such that $U \pitchfork V$ let ρ_V^U be the subsurface projection of the topological boundary of U inside \mathcal{CV} .

Remark 9.7 (comparison with the index set from [BKMM12]). To pass from our new index set \mathfrak{S}' to \mathfrak{S} , one must exchange every subsurface U with included boundary for the disjoint union of the interior of U and open annuli corresponding to the included boundary curves. This procedure preserves nesting, orthogonality, curve graphs, and projections. Therefore, $(\mathcal{MCG}(S), \mathfrak{S})$ being a HHG implies that $(\mathcal{MCG}(S), \mathfrak{S}')$ is a HHG. The main point here is that the two structures have the same unbounded coordinate spaces, namely the curve graphs of connected open subsurfaces that are not pairs of pants, which takes care of most of the axioms (e.g. uniqueness, which constitutes the majority of the work in [BHS19, Section 11]).

9.2. First CHHS structure. Our next goal is to present the first CHHS structures for $\mathcal{MCG}(S)$, whose underlying graph X is a blow-up of the coordinate space \mathcal{CS} . This will answer a question from [BHMS24, Section 1.6]. As mentioned in Remark 1, the graph \mathcal{W} of this structure will be very similar to the graph of complete clean markings from [MM00].

Theorem 9.8. *Let S be a surface of finite-type with boundary. There exists a combinatorial HHG structure (X, \mathcal{W}) for $\mathcal{MCG}(S)$, where X is the blow-up of the curve graph of S , obtained by replacing every curve with the cone over its annular curve graph.*

Proof. It is enough to show that $(\mathcal{MCG}(S), \mathfrak{S}')$, with the coordinate spaces and the projections defined as above, satisfies the hypotheses of Theorem 6.6.

- (**wedges**) Let $U, V \in \mathfrak{S}'$, and let W be the intersection of their interiors, which is a disjoint union of open pairs of pants and open subsurfaces of complexity at least 1. First, for every curve γ in the topological boundary of W which is nested in both U and V , glue to W the closed annulus with core γ , so that one (resp. both) of the boundary curves of the annulus is identified with one (resp. two) curves in the topological boundary of W . Then replace every component of W which is a pair of pants with its included boundary (in particular, one has to remove open pairs of pants). Let W' the surface obtained after this procedure, and let W'' be the disjoint union of W' and every annulus which is a block of both U and V . Now, by construction W'' is in \mathfrak{S}' (since the gluing procedure prevents the forbidden cases from Figure 7 from appearing), and it is nested in both U and V . Moreover W'' is the wedge of U and V , since $T \in \mathfrak{S}'$ is nested in both U and V if and only if its interior is nested in the interior of W (which is the intersection of the interiors, without the pants components), and its included boundary is nested in the intersection of the included boundaries.

• **(clean containers)** Whenever $U \not\sqsubset V \in \mathfrak{S}'$ and there exists $W \in \mathfrak{S}'$ which is nested in V and orthogonal to U , consider the subsurface obtained from $V - U$ after replacing every pair of pants with its included boundary (in particular, one has to remove open pairs of pants). Let Y be the resulting subsurface, which is nonempty since it contains W . First we prove that the connected components of Y are blocks. Indeed, the connected surfaces which are not blocks are annuli with at most one boundary included, or pairs of pants with some boundary components included. Now, one of the connected components of Y is an open annulus if and only if two blocks of U share a common curve in their included boundary, and this would mean that these blocks are not orthogonal by our convention. Moreover, one of the connected components of Y is an annulus with exactly one boundary curve included if and only if two blocks of U fall in one of the forbidden cases from Figure 7. Finally, we manually replaced every pair of pants with its included boundary. Hence Y is a disjoint union of blocks. Similar arguments show that the blocks of Y cannot share curves in their included boundaries (otherwise U would contain an open annulus), and cannot fall in one of the forbidden cases from Figure 7 (otherwise U would contain an annulus with exactly one boundary component included). This shows that Y is an admissible collection, i.e., an element of \mathfrak{S}' , and by construction it is orthogonal to U and nested in V . Being the maximal subsurface with these properties, Y is also the clean container for U inside V .

• **(orthogonals for nonsplit domains property (3.9))** Let $U \not\sqsubset V$ be two subsurfaces. If U has a connected component which is an annulus then U is split, and such annulus is one of its Samaritans. Otherwise U is the disjoint union of finitely many subsurfaces of complexity at least 1, possibly with boundary. If there exists a connected component of V whose intersection with U is trivial, then such component is orthogonal to U . Otherwise there must be a curve γ in the boundary of U relative to V , which must therefore be essential in V . Then the associated annular domain A_γ is nested in V ; moreover, A_γ is either nested in U or disjoint from U , depending on whether the boundary curve is included in U or not. In the former case, U is a split domain, and A_γ is one of its Samaritans; in the latter, A_γ is orthogonal to U .

• **(dense product regions (3.10))** Let $U \in \mathfrak{S}'$ be a nonminimal domain. Then either U is a connected, open subsurface of complexity at least 1, and its curve graph is covered by the projections of the annuli it contains, or $\mathcal{C}U$ is uniformly bounded, as pointed out in Remark 9.6.

• **(cofinite action)** The action of the mapping class group on \mathfrak{S} has finitely many orbits of pairwise disjoint subsurfaces. This is a consequence of the change of coordinate principle (see e.g. [FM12, Section 1.3.3]).

This proves the theorem. □

9.3. Second CHHS structure, with better properties. Notice that the HHS structure from Theorem 9.8 does not satisfy the weak orthogonal Property 7.2. To see this, let V be an open, connected subsurface of complexity 2 and let U be the union of two disjoint, essential annuli inside V . Then U is nonminimal, but $V - U$ is a disjoint union of open pairs of pants and thus cannot contain any element of \mathfrak{S}' .

However, we can find a larger index set to ensure even the strong orthogonal Property 7.1:

Theorem 9.9. *Let S be a surface of finite-type with boundary. There exists a combinatorial HHG structure for $\mathcal{MCG}(S)$ with simplicial wedges (Definition 8.2) and simplicial containers (Definition 8.1).*

Proof. The proof is very similar to that of Theorem 9.8. First, we weaken Definition 9.4, by allowing a block to be also a pair of pants, with some (possibly none) of its boundary components included. Define nesting and orthogonality between blocks as before, with the same conventions about the included boundary components. This allows one to define the collection \mathfrak{S}'' of admissible blocks, as in Definition 9.5. Notice that now the \sqsubseteq -minimal elements of \mathfrak{S}'' are all closed annuli of type (a), and all open pants.

Define the interior and the included boundary of an admissible collection as before, but now the interior also includes the interior of the pants components. Again, notice that if $U, V \in \mathfrak{S}'$ then $U \sqsubseteq V$ if and only if the interior of U is nested in the interior of V , and the included boundary of U is nested in the included boundary of V .

Define the coordinate spaces as before, and set the coordinate space of an open pair of pants to be a point. Again, the only elements of \mathfrak{S}' with unbounded coordinate spaces are annuli and connected, open subsurfaces of complexity at least 1, because all other coordinate spaces are either points or joins.

Finally, define the projections as above, using subsurface projections. With the same techniques of [BHS19, Theorem 11.1], one can then show that $(\mathcal{MCG}(S), \mathfrak{S}'')$ is a HHG.

Next, we observe that $(\mathcal{MCG}(S), \mathfrak{S}'')$ has the following properties. The proofs are very similar to those which appear in Theorem 9.8 (and even easier since we do not have to remove pairs of pants), but we put them here for clarity:

- (**wedges**) Let $U, V \in \mathfrak{S}''$, and let W be the intersection of their interiors, which is a disjoint union of open pants and open subsurfaces of complexity at least 1. Then, for every curve γ in the topological boundary of W which is nested in both U and V , glue to W the closed annulus with core γ , so that one (resp. both) of the boundary curves of the annulus is identified with one (resp. two) curves in the topological boundary of W . Let W' the surface obtained after the gluing, and let W'' be the disjoint union of W' and every annulus which is a block of both U

and V . Now, by construction W'' is in \mathfrak{S}'' (since the gluing procedure prevents the forbidden cases from Figure 7 from appearing), and it is nested in both U and V . Moreover, the interior of W'' is the intersection of the interiors, and the included boundary of W'' is the intersection of the included boundaries. This shows that W'' is the wedge of U and V , since $T \in \mathfrak{S}'$ is nested in both U and V if and only if its interior is nested in the intersection of the interiors, and its included boundary is nested in the intersection of the included boundaries.

- **(clean containers and the strong orthogonal property (7.1))** For $U \not\sqsupseteq V \in \mathfrak{S}''$, consider the subsurface $V - U$. First notice that the connected components of $V - U$ are blocks. Indeed, the only connected subsurfaces which are not blocks are annuli with at most one of the two boundary curves included. Now, one of the connected components of $V - U$ is an open annulus if and only if two blocks of U share a common curve in their included boundary, and this would mean that these blocks are not orthogonal by our convention. Moreover, one of the connected components of $V - U$ is an annulus with exactly one boundary curve included if and only if two blocks of U fall in one of the forbidden cases from Figure 7. Hence $V - U$ is a disjoint union of blocks. The same argument with U and $V - U$ swapped shows that the blocks of $V - U$ cannot share curves in their included boundaries, and cannot fall in one of the forbidden cases from Figure 7. This shows that $V - U$ is an admissible collection, i.e., an element of \mathfrak{S}' , and by construction it is orthogonal to U and nested in V . Being the maximal subsurface with these properties, $V - U$ is also the clean container for U inside V .

- **(dense product regions)** Let $U \in \mathfrak{S}''$ be a nonminimal domain. Then either U is an essential open subsurface, and its curve graph is covered by the projections of the annuli it contains, or $\mathcal{C}U$ is uniformly bounded, as pointed out above.

- **(cofinite action)** The $\mathcal{MCG}(S)$ -action on \mathfrak{S}'' has finitely many orbits of tuples of pairwise orthogonal domains, again thanks to the change of coordinates principle.

Now the hypotheses of Theorem 8.3 are satisfied, and we can find a combinatorial HHS (X, \mathcal{W}) with simplicial wedges and simplicial containers. Since $(\mathcal{MCG}(S), \mathfrak{S}')$ already has dense product region, the pair (X, \mathcal{W}) is exactly the one constructed in Section 4 from $(\mathcal{MCG}(S), \mathfrak{S}')$, as pointed out in Remark 8.4. Hence, by Theorem 6.6, (X, \mathcal{W}) inherits an action of $\mathcal{MCG}(S)$ that makes the latter into a combinatorial HHG, and we are done. □

10. Why orthogonals for nonsplit domains?

The main goal in this section is to use a simple example of an HHS — a CAT(0) cube complex with a factor system — to illustrate the necessity of the orthogonals for nonsplit domains hypothesis 3.9.

Factor systems yield examples of HHS structures where the index set is an *orthogonal set* in the sense of Definition 7.10. It would be illuminating to find conditions on an orthogonal set allowing one to modify the HHS/HHG structure so that some version of the orthogonality properties (3.9, 7.2 or 7.1) hold. We speculate on this below, and in particular on intriguing relations with problems in lattice theory, namely embedding a complete ortho-lattice inside an orthomodular one, see Remark 10.21.

In this section, we use notation from [CRHK24]; see also [HS20] and [BHS17b, Section 8].

10.1. Background on factor systems for CAT(0) cube complexes. For the rest of the section, let \mathcal{Z} be a CAT(0)-cube complex.

Definition 10.1 (hyperplane, carrier, combinatorial hyperplane). A *midcube* in the unit cube $c = [-\frac{1}{2}, \frac{1}{2}]^n$ is a subspace obtained by restricting exactly one coordinate to 0. A *hyperplane* in \mathcal{Z} is a connected subspace H with the property that, for all cubes c of \mathcal{Z} , either $H \cap c = \emptyset$ or $H \cap c$ consists of a single midcube of c . The *carrier* $\mathcal{N}(H)$ of the hyperplane H is the union of all closed cubes c of \mathcal{Z} with $H \cap c \neq \emptyset$. The inclusion $H \rightarrow \mathcal{Z}$ extends to a combinatorial embedding $H \times [-\frac{1}{2}, \frac{1}{2}] \xrightarrow{\cong} \mathcal{N}(H) \hookrightarrow \mathcal{X}$ identifying $H \times \{0\}$ with H . Now, H is isomorphic to a CAT(0) cube complex whose cubes are the midcubes of the cubes in $\mathcal{N}(H)$. The subcomplexes H^\pm of $\mathcal{N}(H)$ which are the images of $H \times \{\pm\frac{1}{2}\}$ under the above map are isomorphic as cube complexes to H , and are *combinatorial hyperplanes* in \mathcal{Z} . Thus each hyperplane of \mathcal{Z} is associated to two combinatorial hyperplanes in $\mathcal{N}(H)$.

Definition 10.2 (gate maps). For any convex subcomplex $\mathcal{Y} \subseteq \mathcal{Z}$ there is a *gate map* $g_{\mathcal{Y}} : \mathcal{Z} \rightarrow \mathcal{Y}$ such that, for any other convex subcomplex $\mathcal{Y}' \subseteq \mathcal{Z}$, the hyperplanes crossing $g_{\mathcal{Y}}(\mathcal{Y}')$ are precisely the hyperplanes which cross both \mathcal{Y} and \mathcal{Y}' .

Gate maps are fundamental in the study of cube complexes and median spaces; see, for instance, [BHS17b, Section 2] for additional background.

Definition 10.3 (parallelism). The convex subcomplexes F and F' are *parallel*, written $F \parallel F'$, if for each hyperplane H of \mathcal{Z} , we have $H \cap F \neq \emptyset$ if and only if $H \cap F' \neq \emptyset$.

Definition 10.4 (orthogonality, orthogonal complement). The convex subcomplexes F and F' are *orthogonal*, written $F \perp F'$, if the inclusions $F \rightarrow \mathcal{Z}$ and $F' \rightarrow \mathcal{Z}$ extend to a convex embedding $F \times F' \rightarrow \mathcal{Z}$.

Given a convex subcomplex F , let P_F be the smallest subcomplex containing the union of all subcomplexes in the parallelism class of F . By [HS20, Lemma 1.7], for example, there is a cubical isomorphism $P_F \rightarrow F \times F^\perp$, where F^\perp is a CAT(0) cube complex which we call the *abstract orthogonal complement of F* . For any

$f \in F^{(0)}$, the inclusion $P_F \rightarrow \mathcal{Z}$ induces an isometric embedding $\{f\} \times F^\perp \rightarrow \mathcal{Z}$ whose image is a convex subcomplex that we call the *orthogonal complement of F at f* and denote $\{f\} \times F^\perp$. Observe that $\{f\} \times F^\perp$ and $\{f'\} \times F^\perp$ are parallel for all $f, f' \in F^{(0)}$ (see [HS20, Lemma 1.11]). When the base point is unimportant, we sometimes abuse notation and write F^\perp to refer to one of these parallel copies.

Lemma 10.5. *Let $F \subset \mathcal{Z}$ be a convex subcomplex, let $f \in F^{(0)}$, and write $F^\perp = \{f\} \times F^\perp$. Suppose that F' is a convex subcomplex such that $F \perp F'$. Then F' is parallel to a subcomplex of F^\perp . Conversely, if F' is a convex subcomplex of F^\perp , then there exists F'' parallel to F' with $F'' \perp F$.*

Proof. This follows easily from [HS20, Lemma 1.11]; see for instance the proof of [HS20, Theorem C]. □

Definition 10.6 (candidate factor system). A *candidate factor system* \mathfrak{h} is a collection of nontrivial convex subcomplexes of \mathcal{Z} (where “nontrivial” means that we exclude singletons) that satisfy the following properties:

- (1) $\mathcal{Z} \in \mathfrak{h}$, and for all combinatorial hyperplanes H of \mathcal{Z} , we have $H \in \mathfrak{h}$.
- (2) If $F, F' \in \mathfrak{h}$ then $\mathfrak{g}_F(F')$ $\in \mathfrak{h}$.
- (3) If $F \in \mathfrak{h}$ and F' is parallel to F , then $F' \in \mathfrak{h}$.

Definition 10.7 (hyperclosure). The *hyperclosure* $\bar{\mathfrak{h}}$ of \mathfrak{h} is the intersection of all candidate factor systems, and therefore the unique candidate factor system which is minimal by inclusion. In other words,

$$\bar{\mathfrak{h}} = \left(\bigcup_{i=1}^{\infty} \mathfrak{h}_i \right) - \{\text{singletons}\},$$

where \mathfrak{h}_1 is the collection of all subcomplexes parallel to combinatorial hyperplanes, together with the whole space \mathcal{Z} , and where

$$\mathfrak{h}_{i+1} = \{F \mid F \parallel \mathfrak{g}_{F_1}(F_2), F_1, F_2 \in \mathfrak{h}_i\} \quad \text{for all } i \geq 1.$$

Lemma 10.8 (characterisation of $\bar{\mathfrak{h}}$ [HS20, Theorem 3.3]). *Let \mathcal{Z} be a locally finite CAT(0) cube complex, and let $\bar{\mathfrak{h}}$ be its hyperclosure. Then a convex subcomplex F belongs to $\bar{\mathfrak{h}}$ if and only if there exists a compact, convex subcomplex C such that $F = C^\perp$.*

Definition 10.9 (weak factor system). Let \mathfrak{h} be a candidate factor system, and let $\mathfrak{h}_{/\sim}$ the set of parallelism classes of subcomplexes in \mathfrak{h} . If there exists $N \in \mathbb{N}$ such that N bounds the length of chains in the partial order of \mathfrak{h} given by inclusion, then $\mathfrak{h}_{/\sim}$ is a *weak factor system*.

We denote the parallelism class of the subcomplex $F \in \mathfrak{h}$ by $[F] \in \mathfrak{h}_{/\sim}$. Two elements $[F], [F'] \in \mathfrak{h}_{/\sim}$ are *nested* (resp. *orthogonal*), and we write $[F] \sqsubseteq [F']$

(resp. $[F] \perp [F']$) if there exists two representatives F, F' such that $F \subseteq F'$ (resp. $F \perp F'$).

The class of CAT(0) cube complexes which admit weak factor systems is quite large. For example, virtually special groups, in the sense of Haglund and Wise [HW08], act geometrically on CAT(0) cube complexes with weak factor systems by [BHS17b, Proposition B]. The more general class of CAT(0) cube complexes with geometric group actions and weak factor systems is characterised in [HS20], and includes some notable nonspecial examples, like irreducible lattices in products of trees, and certain amalgams of these [Hag23]. There are other amalgams of such lattices that provide the first examples of proper cocompact CAT(0) cube complexes *not* admitting any weak factor system [She25].

Following [BHS17b], one can endow a CAT(0) cube complex with an HHS structure when the hyperclosure gives a weak factor system; this is analysed in more detail in [CRHK24, Section 20]. We summarise this here in order to connect the existing results more explicitly to the hypotheses of Theorem 3.16.

The first lemma is needed to verify the clean containers property. It is proved in [HS20, Proposition 5.1] as part of a more complicated statement whose other parts rely on the presence of a cocompact group action; see [CRHK24, Remark 20.7]. So, for the avoidance of doubt, we extract the exact statement (with the same proof as in [HS20]) here:

Lemma 10.10. *Let \mathcal{Z} be a CAT(0) cube complex and let $\bar{\mathfrak{h}}$ be its hyperclosure. If $\bar{\mathfrak{h}}_{/\sim}$ is a weak factor system then, for every $F \in \bar{\mathfrak{h}} - \{\mathcal{Z}\}$ and every $x \in F^{(0)}$, the subcomplex $F^\perp := \{x\} \times F^\perp$ belongs to $\bar{\mathfrak{h}}$.*

Proof. We will use [HS20, Lemma 5.2] after some preliminary setup. Let $\{H_i\}_{i \in I}$ be the set of hyperplanes that are dual to edges of F incident to x . For $i \in I$, let H_i^+ be the combinatorial hyperplane in which the carrier of $N(H_i)$ intersects the H_i -halfspace of \mathcal{Z} containing x ; in particular, x lies in H_i^+ . Let $Y = \bigcap_{i \in I} H_i^+$. Observe that for any finite $I' \subset I$, we have $Y(I') := \bigcap_{i \in I'} H_i^+ \in \bar{\mathfrak{h}}$, unless $Y(I')$ is a single point (recall that we do not allow single points in the hyperclosure). Assume the former.

By the assumption that the set of parallelism classes represented in the hyperclosure is a weak factor system, together with the observation that $[Y(I')] \sqsubseteq [Y(I'')]$ when $I'' \subset I'$, we see that there exists a finite subset $I_0 \subset I$ such that $Y = Y(I_0)$. Hence $Y \in \bar{\mathfrak{h}}$, or Y consists of a single vertex.

Now let \mathcal{S} be the set of all combinatorial hyperplanes H^\pm such that the associated hyperplane H (i.e., the hyperplane H such that the usual identification of $N(H)$ with $H \times [-\frac{1}{2}, \frac{1}{2}]$ identifies $H \times \{\pm \frac{1}{2}\}$ with H^\pm) crosses F . By [HS20, Lemma 5.2], $F^\perp = \bigcap_{H^\pm \in \mathcal{S}} \mathfrak{g}_Y(H^\pm)$. We now argue as before. First, note that if $S' \subset \mathcal{S}$ is finite, then $A(S') = \bigcap_{H^\pm \in S'} \mathfrak{g}_Y(H^\pm)$ belongs to $\bar{\mathfrak{h}}$, or it is a single point.

As before, if $S' \subset S''$ are finite subsets of S , we have $A(S'') \subset A(S')$, so our assumption that \bar{h} gives a weak factor system again implies that there is a finite S' such that $F^\perp = A(S')$. Hence, either $F^\perp \in \bar{h}$, or F^\perp is a single point.

To complete the proof, we rule out the latter possibility as follows. By hypothesis, $F \neq \mathcal{Z}$. By Lemma 10.8, there exists a (compact) convex subcomplex $C \subset \mathcal{Z}$ with $C^\perp = F$, since $F \in \bar{h}$. This means that $F \perp C$, so Lemma 10.5 implies that, up to parallelism, $C \subset F^\perp$, so it remains to show that C is nontrivial. But if C is trivial, then $C \perp \mathcal{Z}$, so by Lemma 10.5, $\mathcal{Z} \sqsubseteq C^\perp = F$, contradicting that F is a proper subcomplex. \square

Remark 10.11 (Why are singletons excluded?) Let $\bar{h}_{j\sim}$ denote the set of parallelism classes in the hyperclosure \bar{h} . In general, one gets from a weak factor system to an HHS structure using Theorem 10.12 below. The additional property that the index set is an orthogonal set can be arranged by using the weak factor system provided by the hyperclosure, when it exists — see Theorem 10.13.

Recall that we have defined \bar{h} so as to exclude subcomplexes consisting of a single vertex. Note that any two such subcomplexes $\{x\}, \{y\}$ are parallel. Moreover, $\{\{x\}\} \sqsubseteq [F]$, for any $F \in \bar{h}$. On the other hand, we also have $\{\{x\}\} \perp [F]$. This would not be allowed in an HHS structure, so we avoid the issue by excluding $\{x\}$ from the hyperclosure. For nontrivial subcomplexes, this problem does not occur, by, for instance, [CS11, Proposition 2.5] and the fact that hyperplanes do not cross themselves.

In [BHS17b], something more radical is allowed: one can exclude all subcomplexes below some fixed diameter, and still get an HHS structure. But this may cease to satisfy the conclusion of Lemma 10.10. This is not a problem from the point of view of clean containers, but it can break “orthogonality determines nesting” by creating non- \sqsubseteq -maximal $[F]$ that are not orthogonal to anything in the index set.

Theorem 10.12 (see e.g. [CRHK24, Proposition 20.4]). *Let \mathcal{Z} be a CAT(0) cube complex with a weak factor system $h_{j\sim}$. Then $(\mathcal{Z}, h_{j\sim})$ is a hierarchically hyperbolic space with wedges, where the coordinate spaces $C[F], [F] \in h_{j\sim}$ and the projections $\pi_{[F]} : \mathcal{Z} \rightarrow C[F]$ are as in [BHS17b, Remark 13.2].*

The next theorem refines the previous one by strengthening the conclusions about the index set. It should be compared to the very similar [CRHK24, Proposition 20.6], where, however, it is not explicit that the HHS structure comes from the hyperclosure. Note, also, that the clean containers property of the HHS structure from Theorem 10.12 is observed (in the presence of a group action) in [HS20].

Theorem 10.13. *Let \mathcal{Z} be a CAT(0) cube complex which admits a weak factor system, and let \bar{h} be the hyperclosure of \mathcal{Z} . Then $\bar{h}_{j\sim}$ is a weak factor system and an orthogonal set. In particular, $(\mathcal{Z}, \bar{h}_{j\sim})$ is a HHS with wedges, clean containers, and where orthogonality implies nesting.*

Proof. In the proof of [CRHK24, Proposition 20.6] it is shown that, if \mathcal{Z} admits a weak factor system, then there exists a weak factor system $\mathfrak{h}'_{/\sim}$ which consists of all equivalence classes of subcomplexes of the form

$$\mathfrak{g}_{H_1}(\dots(\mathfrak{g}_{H_{n-1}}(H_n))\dots),$$

where H_1, \dots, H_n are combinatorial hyperplanes and $n \geq 0$. By Definition 10.7 of the hyperclosure $\bar{\mathfrak{h}}$, we see that $\mathfrak{h}' \subseteq \bar{\mathfrak{h}}$, and they must coincide since the hyperclosure is the minimal candidate factor system. Hence, the quotient of the hyperclosure by parallelism is a weak factor system, and therefore $(\mathcal{Z}, \bar{\mathfrak{h}}_{/\sim})$ is a HHS with wedges by Theorem 10.12.

Now, by Lemma 10.10, if $\bar{\mathfrak{h}}_{/\sim}$ is a weak factor system then, for every $F \in \bar{\mathfrak{h}}$, we have $F^\perp \in \bar{\mathfrak{h}}$. In particular, this shows that $\bar{\mathfrak{h}}_{/\sim}$ has clean containers, since if $[F], [F'], [C] \in \bar{\mathfrak{h}}_{/\sim}$ are such that $[F], [F'] \sqsubseteq [C]$ and $[F] \perp [F']$, then the clean container for $[F]$ inside $[C]$ is $[C] \wedge [F]^\perp$, where $[F]^\perp = [F^\perp]$, as provided by Lemma 10.5.

By [HS20, Corollary 3.4] we have $F^{\perp\perp} = F$ whenever $F \in \bar{\mathfrak{h}}$, and by Lemma 7.11 this is equivalent to the fact that orthogonality determines nesting in $\bar{\mathfrak{h}}_{/\sim}$. \square

Remark 10.14. The HHS structure $(\mathcal{Z}, \bar{\mathfrak{h}}_{/\sim})$ need not have the DPR property. When it does, this can be verified as follows. For each $[F] \in \bar{\mathfrak{h}}_{/\sim}$, the HHS structure from [BHS17b] has as the coordinate space $\mathcal{C}[F]$ the *factored contact graph*, which can be viewed as a copy of F with additional edges added to cone off subcomplexes of the form $\mathfrak{g}_F(F')$, where $\mathfrak{g}_F : \mathcal{Z} \rightarrow F$ is the gate map and $F' \in \bar{\mathfrak{h}}$. Now, the hyperplanes of F have the form $H \cap F$, where H is a hyperplane of \mathcal{Z} crossing F . This shows that subcomplexes of F of the form $\mathfrak{g}_F(F')$ cover F . This implies DPR provided sufficiently many of those subcomplexes to coarsely cover F are actually in $\bar{\mathfrak{h}}$, i.e., they are not singletons. So, for example, $(\mathcal{Z}, \bar{\mathfrak{h}}_{/\sim})$ has the DPR property provided there exists a constant K such that for all non \sqsubseteq -minimal $F \in \bar{\mathfrak{h}}$ and all $x \in F$, there exists $F' \in \bar{\mathfrak{h}}$ such that $d_{\mathcal{Z}}(x, \mathfrak{g}_F(F')) \leq K$ and $|\mathfrak{g}_F(F')| > 1$.

10.2. The counterexample. In this subsection we present a HHS $(\mathcal{Z}, \mathfrak{S})$ that satisfies all hypotheses of the main Theorem 3.16 except the orthogonals for nonsplit domains Property 3.9, and we prove that the graph X from Definition 4.2, constructed using the coordinate spaces in the HHS structure, cannot be the underlying graph of any combinatorial HHS structure for \mathcal{Z} .

Remarkably, \mathcal{Z} is a CAT(0) cube complex admitting a weak factor system, and $\mathfrak{S} = \mathfrak{h}_{/\sim}$ is the quotient of its hyperclosure by parallelism, which is also an orthogonal set by Theorem 10.13. This shows that the orthogonals for nonsplit domains property, which is the most obscure among the properties of our main theorem, is essential for the construction from Section 4 to work, even if we assume that orthogonality determines nesting.




Figure 8. The space \mathcal{Z} is obtained by gluing the three pieces along the red and blue arrows. The vertical edges are labelled by numbers, while the horizontal edges are labelled by Greek capital letters. Notice that Γ_1 is orthogonal to all edges with positive odd labels, which form an infinite ray, and similarly Γ_2 is orthogonal to all edges with positive even labels. The yellow edges are representatives for the parallelism classes of edges which are in the hyperclosure.

Consider the two-dimensional cube complex \mathcal{Z} obtained by gluing the three infinite complexes in Figure 8 along the red and blue arrows, respecting the numerical labels.

Lemma 10.15. *The complex \mathcal{Z} is CAT(0).*

Proof. This follows from various versions of the principle that gluing simply connected, non-positively-curved spaces along convex subspaces using isometries yields a simply connected non-positively-curved result. More precisely:

Combinatorial version: A cube complex is CAT(0) if and only if its 1-skeleton is a median graph [Che00, Theorem 6.1]. Products of median graphs are median, and a graph decomposing as the union of two median graphs intersecting along a subgraph that is convex in each piece is median (see e.g. [Isb80; Che00]).

CAT(0) version: A complete geodesic metric space is CAT(0) if it is the union of two CAT(0) spaces whose intersection is convex in each piece, and products of CAT(0) spaces are CAT(0) [BH99, Theorem 11.1, Exercise 1.16.(2)].

The part of \mathcal{Z} at left is obtained from two copies of $[0, 2] \times [0, \infty)$, glued along a point, by gluing squares along a collection of edges. A similar observation applies to the two pieces at right. Apply the gluing principle once more, using that the red/blue line is convex in each piece where it appears. \square

Next, let us check quickly that \mathcal{Z} admits a factor system:

Lemma 10.16. *Let \bar{h} be the hyperclosure in \mathcal{Z} . Then $\bar{h}_{/\sim}$ is a weak factor system. More strongly, \bar{h} is a factor system in the sense of [BHS17b].*

Finally, the HHS structure $(\mathcal{Z}, \bar{h}_{/\sim})$ from Theorem 10.13 has the DPR property.

Proof. Suppose that $[F_1], \dots, [F_n] \in \bar{\mathfrak{h}}_{/\sim}$ satisfy $[F_i] \not\sqsubseteq [F_{i+1}]$ for all i , with each $F_i \subsetneq \mathcal{Z}$. For each i consider the parallelism class $[F_i^\perp]$ of its abstract orthogonal, which, by Lemma 10.5, coincides with the parallelism class $[F_i]^\perp$ of the maximal subcomplex orthogonal to F_i . Since $[F_i] \not\sqsubseteq [F_{i+1}]$ we have $[F_i]^\perp \sqsubseteq [F_{i-1}]^\perp$; moreover the nesting is proper, since by [HS20, Corollary 3.4] we have $[F_i]^{\perp\perp} = [F_i]$.

Now, if $\bar{\mathfrak{h}}_{/\sim}$ had arbitrarily large \sqsubseteq -chains, then for any $R \geq 0$, we could choose n as above so that for some $m \leq n$, the subcomplexes F_m and F_m^\perp both have at least R vertices (take n much larger than R , and consider $m = \lfloor n/2 \rfloor$).

Next observe that there is a uniform bound on the degrees of vertices in \mathcal{Z} . Hence, for any $R_1 \geq 0$, we can choose R and thus the F_i so that F_m and F_m^\perp , chosen as above for the given R , have diameter more than R_1 .

Thus \mathcal{Z} contains a convex subcomplex isometric to $F_m \times F_m^\perp$, with each factor having diameter at least R_1 . Now observe that the inclusion into \mathcal{Z} of the union of the red and blue rays is a quasi-isometry $\mathbb{R} \rightarrow \mathcal{Z}$, so \mathcal{Z} is hyperbolic. By taking R_1 sufficiently large in terms of the hyperbolicity constant, we contradict, say, [Hag14, Theorem 7.6] or [CDE+08].

Now, \mathcal{Z} is uniformly locally finite (i.e., the number of 0-cubes in a ball is bounded in terms of the radius of the ball only), and $\bar{\mathfrak{h}}$ is closed under taking intersections (since the projection of A to B is $A \cap B$ when A, B are convex subcomplexes with $A \cap B \neq \emptyset$), the bound on the length of \sqsubseteq -chains implies a bound on the number of elements of the hyperclosure that can contain a given 0-cube, as required by the definition of a factor system in [BHS17b, Section 8].

Finally, the DPR follows from Remark 10.14 in this example, since each element of the hyperclosure is uniformly coarsely covered by edges coloured yellow in Figure 8. □

Crucially, the orthogonal for nonsplit property does not hold in $\bar{\mathfrak{h}}_{/\sim}$. We give here a concrete motivation, which we shall then revisit under a more conceptual light in Remark 10.20.

Lemma 10.17. *$\bar{\mathfrak{h}}_{/\sim}$ does not have the orthogonal for nonsplit domains property.*

Proof. Let F be the hyperplane dual to the edge Δ , and let F' be the subcomplex obtained by projecting to F the hyperplane dual to Γ_1 . Notice that $[n] \sqsubseteq [F]$ for all $n \in \mathbb{N}_{>0}$ and $[F'] \sqsubseteq [F]$. However F is 1-dimensional, and as a consequence the orthogonal complements in F of these subcomplex are points and therefore do not belong to $\bar{\mathfrak{h}}_{/\sim}$. Hence $(\mathcal{Z}, \bar{\mathfrak{h}}_{/\sim})$ fails to have orthogonals for nonsplit domains. □

10.2.1. Failure of the blow-up construction. We now prove that the combinatorial HHS (X, \mathcal{W}) associated to $(\mathcal{Z}, \bar{\mathfrak{h}}_{/\sim})$ cannot be quasi-isometric to \mathcal{Z} . We start by describing the minimal orthogonality graph:




Figure 9. The minimal orthogonality graph \bar{X} .

Lemma 10.18. *The minimal orthogonality graph of $\bar{h}_{/\sim}$ is the graph shown in Figure 9.*

Proof. Each edge labelled by an integer $n \geq -1$ in Figure 8 is parallel to a (yellow) combinatorial hyperplane, and hence its parallelism class $[h]$ belongs to $\bar{h}_{/\sim}$. The same is true of the edges Γ_1, Γ_2 .

Notice moreover that the edge Δ is (parallel to) the gate from the hyperplane dual to -1 to the hyperplane dual to 1 . Hence $[\Delta] \in \bar{h}_{/\sim}$, and the same argument with 0 replacing 1 shows that $[\Sigma] \in \bar{h}_{/\sim}$ as well.

Now recall that \bar{h} excludes singletons, so any element of $\bar{h}_{/\sim}$ represented by a subcomplex consisting of a single edge is \sqsubseteq -minimal; this holds in particular for $[\Gamma_1], [\Gamma_2], [\Sigma], [\Delta]$ and $[n], n \geq -1$. The other parallelism classes of edges do not belong to \bar{h} . Moreover, any element of \bar{h} contains a parallel copy of one of the above-named edges.

Therefore, the minimal orthogonality graph \bar{X} has vertex set

$$\{[\Sigma], [\Delta], [\Gamma_1], [\Gamma_2]\} \cup \{[n]\}_{n \geq -1}.$$

From the definition of orthogonality, vertices of \bar{X} are adjacent if and only if the corresponding edges of \mathcal{Z} have parallel copies spanning a square, and thus \bar{X} is as in Figure 9. □

Now recall how to construct the simplicial complex X from \bar{X} : for each \sqsubseteq -minimal $[F] \in \bar{h}_{/\sim}$, the corresponding vertex of \bar{X} is blown up to a cone over the vertex set of the coordinate space $\mathcal{C}[F]$ from the HHS structure of Theorem 10.12. From [BHS17b, Remark 13.2], $\mathcal{C}[F]$ is the *factored contact graph* of the CAT(0) cube complex F . In the present example, each such F is an edge (one of $\Delta, \Sigma, \Gamma_i, i \in \{1, 2\}$, or $n, n \geq -1$), so $\mathcal{C}[F]$ is a single vertex. Hence X is obtained from \bar{X} by blowing up each vertex to an edge (the cone over a vertex); each edge of \bar{X} therefore blows up to a 3-simplex.

We now argue that this \bar{X} cannot support a combinatorial HHS structure for \mathcal{Z} .

Proposition 10.19. *Suppose that \mathcal{W} is any X -graph such that (X, \mathcal{W}) is a combinatorial HHS. Then \mathcal{W} is bounded, and so cannot be quasi-isometric to \mathcal{Z} .*

Proof. For each vertex v of \bar{X} , let \hat{v} be the 1-simplex supported on v , and let v, Cv be the two 0-simplices supported on v . We use the same notation for a parallelism class of subcomplexes in \bar{h} as for the corresponding vertex of \bar{X} . So, for instance, in X , $[3]$ and $C[3]$ are 0-simplices, and $[\widehat{3}]$ and $\{[3], C[\Gamma_1]\}$ are 1-simplices.

If (X, \mathcal{W}) is a CHHS with \mathcal{W} unbounded, then by Theorem 2.12, the augmented links of simplices in X must have arbitrarily large diameter. However, their diameters are uniformly bounded, in terms of the constant δ from Definition 2.7, which exists by hypothesis. For example:

- $\text{Lk}_X(\emptyset)^{+\mathcal{W}} = X^{+\mathcal{W}}$ has diameter bounded independently of δ , by inspection.
- Let $[\widehat{n}]$ be the edge of X projecting to the vertex $[n]$ of \bar{X} . For $n \geq 1$, the saturation of $[\widehat{n}]$ consists of those $[\widehat{m}]$ with $m = n \bmod 2$, and the link consists of $\widehat{\Sigma} \cup \widehat{\Delta} \cup \Gamma_i$ for one of the values of i . For $n \leq 0$, the link is $\widehat{\Sigma}$ or $\widehat{\Delta}$. In any case, $\text{Lk}_X([\widehat{n}])^{+\mathcal{W}}$ is connected, by Definition 2.7(2), and there are only finitely many such graphs, each of which has finitely many vertices, so these links are uniformly bounded.
- $\text{Lk}_X(\widehat{\Gamma}_1)$ is the union of the edges $[\widehat{2k+1}]$, $k \geq 0$, and $\text{Sat}(\widehat{\Gamma}_1) = \widehat{\Gamma}_1^{(0)} = \{[\Gamma_1], C[\Gamma_1]\}$. Moreover $Y_{\widehat{\Gamma}_1} = (X - \text{Sat}(\widehat{\Gamma}_1))^{+\mathcal{W}}$ is uniformly bounded, and so must be $\text{Lk}_X(\widehat{\Gamma}_1)^{+\mathcal{W}}$ which is quasi-isometrically embedded in $Y_{\widehat{\Gamma}_1}$ by Definition 2.7(2). The same holds for $\widehat{\Gamma}_2$.

Similar arguments give boundedness of the remaining links (some are bounded by construction, recall Corollary 4.4, while the links of almost maximal simplices are points). Hence, \mathcal{W} is bounded since it is a HHS whose coordinate spaces are all bounded (e.g. by the distance formula 1.11). □

Remark 10.20 (What’s wrong with this example?). The HHS structure on \mathcal{Z} has wedges, clean containers and the dense product region property by Theorem 10.13 and Lemma 10.16. Moreover, orthogonality determines nesting, i.e., $[F] \sqsubset [F']$ if and only if $[F']^\perp \not\sqsubset [F]^\perp$.

This illustrates that an HHS in which orthogonality determines nesting need not have the orthogonals for nonsplit domains property. This is a typical phenomenon in CAT(0) cube complexes. Indeed, if $[F] \in \bar{h}_{j\sim}$ (in an arbitrary CAT(0) cube complex), then as long as $[F]$ is not the unique \sqsubseteq -maximal element, $[F]^\perp$ is defined and belongs to $\bar{h}_{j\sim}$. However, the existence of orthogonals is not inherited by the sub-HHS structure on F , in general. More precisely, we could consider the hyperclosure of the CAT(0) cube complex F , called \bar{h}^F . Since hyperplanes of F are of the form $F \cap H$, where H is a hyperplane of \mathcal{Z} intersecting F , by e.g. [CS11, Lemma 3.1], and $F \cap H = \mathfrak{g}_F(H)$ by e.g. [BHS17b, Lemma 2.6], we see that \bar{h}^F naturally embeds in \bar{h} (preserving parallelism), yielding a set of domains in $\bar{h}_{j\sim}$ nested in $[F]$. However, the set of all $[F'] \in \bar{h}_{j\sim}$ which are nested inside $[F]$ is in general larger than this, since it contains subcomplexes of the form $\mathfrak{g}_F(F'')$,

where F'' is an element of the hyperclosure not parallel to any subcomplex in the “intrinsic” hyperclosure of F . Hence, given an element $[F'] \sqsubseteq [F]$, its orthogonal inside $[F]$ might not exist.

Concretely, in our counterexample, we can translate Lemma 10.17 in the above language as follows. Let F be the hyperplane dual to the edge Δ . Since F is 1-dimensional, its hyperplanes are points and so its intrinsic hyperclosure is empty. But in $\bar{h}_{/\sim}$, we have, say, $[n] \sqsubseteq [F]$ and $[F'] \sqsubseteq [F]$, where $[F']$ is the subcomplex obtained by projecting to F the hyperplane dual to Γ_1 . Again since F is 1-dimensional, these subcomplexes do not have nontrivial orthogonal complements that belong to the hyperclosure but are nested in F . Hence $(\mathcal{Z}, \bar{h}_{/\sim})$ fails to have orthogonals for nonsplit domains.

Absent orthogonals and boundedness of \mathcal{W} : Let us see how the failure of orthogonals for nonsplit domains caused problems in the example. Recall that the simplex $\widehat{\Gamma}_1$ had bounded link *because* its saturation failed to contain, say, $\widehat{\Sigma}$, which then acted as a cone-point in $X - \text{Sat}(\widehat{\Gamma}_1)$ over $\text{Lk}_X(\widehat{\Gamma}_1)$. Consider the hyperplanes H_{Γ_1} and H_{Σ} in \mathcal{Z} dual to the edges Γ_1 and Σ , and let F be the projection of the former onto the latter, so $F \in \bar{h}$ and F is the ray consisting of the red edges in Figure 8. Then F is nonsplit, since any two \sqsubseteq -minimal domains it contains (i.e., every two red edges) are transverse. Moreover $F \sqsubseteq H_{\Sigma}$, but $F^{\perp} = \Gamma_1 \cup \Sigma \cup \Delta$, so F is not orthogonal to anything properly nested in H_{Σ} . Back in X , the role of $\mathcal{C}[F]$ (which is a ray) should be played by the link of $\widehat{\Gamma}_1$, which is contained in the link of $\widehat{\Sigma}$. So it seems reasonable that by adding the “missing” orthogonal domain, one could add a vertex w to X in such a way that $w \star \Sigma$ is defined and has the same link as $\widehat{\Gamma}_1$. This way, removing $\text{Sat}(\widehat{\Gamma}_1)$ would now remove Σ , which as we mentioned is an obstruction to having an unbounded augmented link. Of course, one would also need to deal with Δ similarly. But then, if $\widehat{\Gamma}_1$ and $\widehat{\Gamma}_2$ were made to have unbounded links, a new difficulty would arise, because Definition 2.7(3) would then demand that $\widehat{\Sigma}$ and $\widehat{\Delta}$ have a common nested simplex whose link contains those of $\widehat{\Gamma}_1$ and $\widehat{\Gamma}_2$. This would presumably be addressed by adding a single w as above, joined to both Σ and Δ and Γ_1 , and corresponding to an orthogonal complement of F inside the wedge of H_{Σ} and H_{Δ} .

Remark 10.21 (connection to lattice theory). Let $(\mathfrak{S}, \sqsubseteq, \perp)$ be an orthogonal set and let \emptyset be a symbol distinct from all elements of \mathfrak{S} . Then the operation \wedge which, by Definition 7.10, was partially defined on \mathfrak{S} , extends to a binary operation on $\mathfrak{S} \sqcup \{\emptyset\}$ if one sets $U \wedge \emptyset = \emptyset \wedge U = \emptyset$ for $U \in \mathfrak{S} \sqcup \{\emptyset\}$, and $U \wedge V = \emptyset$ if $U, V \in \mathfrak{S}$ do not have any common nested elements of \mathfrak{S} . We extend \sqsubseteq so that \emptyset is the unique \sqsubseteq -minimal element. Assuming that the complexity is finite, we also have a join operation: $U \vee V$ is the unique \sqsubseteq -minimal W such that $U, V \sqsubseteq W$; that this is well-defined is an easy exercise using wedges and finite complexity. In

fact, the poset $(\mathfrak{S} \sqcup \{\emptyset\}, \sqsubseteq)$, equipped with the operations \wedge and \vee , is a *complete lattice*, as in e.g. [BS81, Definition 4.1], where completeness also follows from finite complexity. Moreover, the clean container assumption, together with Lemma 7.11, gives an involution $^\perp : \mathfrak{S} - \{S\} \rightarrow \mathfrak{S} - \{S\}$, where S is the unique \sqsubseteq -maximal element, and we extend this to $\mathfrak{S} \cup \{\emptyset\}$ by declaring $S^\perp = \emptyset$ and $\emptyset^\perp = S$.¹ This makes $\mathfrak{S} \sqcup \{\emptyset\}$, with the lattice relations \wedge, \vee and the orthogonal complementation operator $^\perp$, an *ortholattice*, defined in e.g. [Ste99, Section 1.2].

Strong orthogonality (Property 7.1) then becomes: for all $U, V \in \mathfrak{S}$ such that $U \not\sqsubseteq V$, we have that $U^\perp \wedge V \neq \emptyset$. Another formulation is: for any $V \in \mathfrak{S}$, the order ideal $\{U \in \mathfrak{S} : U \sqsubseteq V\}$ is again an orthogonal set, with the involution $U_V^\perp = U^\perp \wedge V$.

Now, if $U \not\sqsubseteq V$, then $(U^\perp \wedge V) \vee U \sqsubseteq V$. If for some $W \not\sqsubseteq V$ we have $U \sqsubseteq W$ and $U^\perp \wedge V \sqsubseteq W$, then strong orthogonality provides $A := W^\perp \wedge V \neq \emptyset$. But then $(U^\perp \wedge V) \perp A$, since $A \perp W$, while on the other hand $W^\perp \wedge V = A \sqsubseteq U^\perp \wedge V$ since $U \sqsubseteq W$. This is a contradiction. We have showed that strong orthogonality implies the identity $(U^\perp \wedge V) \vee U = V$ whenever $U \sqsubseteq V$. Ortholattices satisfying this identity have a name: they are *orthomodular* [Ste99].

From the point of view of HHS structures, factor systems in CAT(0) cube complexes provide the main motivating examples of orthogonal sets, and we saw earlier that orthomodularity fails in general; in fact, one can already see this failure in HHS structures on, say, right-angled Artin groups (see [BHS17b, Section 8]). This raises the following:

Question 10.22. Let $(\mathfrak{L}, \wedge, \vee, ^\perp, \emptyset, S)$ be an ortholattice such that \sqsubseteq -chains in \mathfrak{L} have length at most $N < \infty$, where $U \sqsubseteq V$ means $U \wedge V = U$. Write $U \perp V$ to mean $U \sqsubseteq V^\perp$.

- Is there an *orthomodular* ortholattice \mathfrak{L}_1 and an injective map $\mathfrak{L} \rightarrow \mathfrak{L}_1$ that preserves the relations \sqsubseteq and \perp , as well as the negations of those relations? Under what conditions can \mathfrak{L}_1 be chosen so that chains in \mathfrak{L}_1 also have length at most N ?
- If a group G acts on \mathfrak{L} cofinitely, preserving the relations $\sqsubseteq, \perp, \not\sqsubseteq, \not\perp$, when can \mathfrak{L}_1 be chosen as above so that the G -action extends to \mathfrak{L}_1 and $|G \backslash \mathfrak{L}_1| < \infty$?

The goal would be to begin with an HHS/ G whose index set is an orthogonal set (e.g., a compact special group) and produce a new HHS/ G structure to which Theorem 8.3 applies. Answers to the above questions are not quite sufficient but appear necessary, and also of independent interest. We suspect that a sufficient condition for constructing \mathfrak{L}_1 will involve the existence of an order-preserving, \perp -preserving map from \mathfrak{L} to a finite boolean lattice. It is also possible that this sort of construction is known to lattice theorists, in which case we would be grateful for a reference.

¹On \mathfrak{S} , the orthogonality relation $U \perp V$ is still equivalent to $U \sqsubseteq V^\perp$, i.e., $U \wedge V^\perp = U$. However, \sqsubseteq and \perp are not mutually exclusive on $\mathfrak{S} \cup \{\emptyset\}$, but the only failure is $\emptyset \sqsubseteq S$ and $\emptyset \perp S$.

Acknowledgements

We thank Carolyn Abbott and Alexandre Martin for helpful discussions. Hagen thanks Montserrat Casals-Ruiz for being a strong proponent of the “orthogonal set” viewpoint during work on [CRHK24], which influenced the ideas in Section 10 here. We thank Jason Behrstock for some useful comments on an earlier version, and the referee for numerous very helpful comments.

References

- [AB23] C. Abbott and J. Behrstock, “Conjugator lengths in hierarchically hyperbolic groups”, *Groups Geom. Dyn.* **17**:3 (2023), 805–838. MR
- [ABD21] C. Abbott, J. Behrstock, and M. G. Durham, “Largest acylindrical actions and stability in hierarchically hyperbolic groups”, *Trans. Amer. Math. Soc. Ser. B* **8** (2021), 66–104. MR
- [ANS+24] C. R. Abbott, T. Ng, D. Spriano, R. Gupta, and H. Petyt, “Hierarchically hyperbolic groups and uniform exponential growth”, *Math. Z.* **306**:1 (2024), art.id. 18, 33 pp. MR
- [BGGH25] E. Bongiovanni, P. Ghosh, F. Gültepe, and M. Hagen, “Characterizing hierarchically hyperbolic free by cyclic groups”, preprint, 2025. arXiv 2508.15738
- [BH99] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, Grundle. Math. Wissen. **319**, Springer, 1999. MR
- [BHMS24] J. Behrstock, M. Hagen, A. Martin, and A. Sisto, “A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups”, *J. Topol.* **17**:3 (2024), art.id. e12351, 94 pp. MR
- [BHS17a] J. Behrstock, M. F. Hagen, and A. Sisto, “Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups”, *Proc. Lond. Math. Soc.* (3) **114**:5 (2017), 890–926. MR
- [BHS17b] J. Behrstock, M. F. Hagen, and A. Sisto, “Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups”, *Geom. Topol.* **21**:3 (2017), 1731–1804. MR
- [BHS19] J. Behrstock, M. Hagen, and A. Sisto, “Hierarchically hyperbolic spaces II: Combination theorems and the distance formula”, *Pacific J. Math.* **299**:2 (2019), 257–338. MR
- [BHS21] J. Behrstock, M. F. Hagen, and A. Sisto, “Quasiflats in hierarchically hyperbolic spaces”, *Duke Math. J.* **170**:5 (2021), 909–996. MR
- [BKMM12] J. Behrstock, B. Kleiner, Y. Minsky, and L. Mosher, “Geometry and rigidity of mapping class groups”, *Geom. Topol.* **16**:2 (2012), 781–888. MR
- [BR20] F. Berlai and B. Robbio, “A refined combination theorem for hierarchically hyperbolic groups”, *Groups Geom. Dyn.* **14**:4 (2020), 1127–1203. MR
- [BS81] S. Burris and H. P. Sankappanavar, *A course in universal algebra*, Graduate Texts in Mathematics **78**, Springer, 1981. MR
- [CDE+08] V. Chepoi, F. F. Dragan, B. Estellon, M. Habib, and Y. Vaxès, “Diameters, centers, and approximating trees of δ -hyperbolic geodesic spaces and graphs”, pp. 59–68 in *Computational geometry* (SCG’08), ACM, New York, 2008. MR
- [CdIH16] Y. Cornuier and P. de la Harpe, *Metric geometry of locally compact groups*, EMS Tracts in Mathematics **25**, European Mathematical Society, Zürich, 2016. MR
- [Che00] V. Chepoi, “Graphs of some CAT(0) complexes”, *Adv. in Appl. Math.* **24**:2 (2000), 125–179. MR

- [CRHK24] M. Casals-Ruiz, M. Hagen, and I. Kazachkov, “Real cubings and asymptotic cones of hierarchically hyperbolic groups”, working draft, 2024, available at https://www.wescac.net/cones_july_2024-public.pdf.
- [CS11] P.-E. Caprace and M. Sageev, “Rank rigidity for CAT(0) cube complexes”, *Geom. Funct. Anal.* **21**:4 (2011), 851–891. MR
- [DDLS24] S. Dowdall, M. G. Durham, C. J. Leininger, and A. Sisto, “Extensions of Veech groups, II: Hierarchical hyperbolicity and quasi-isometric rigidity”, *Comment. Math. Helv.* **99**:1 (2024), 149–228. MR
- [DHS17] M. G. Durham, M. F. Hagen, and A. Sisto, “Boundaries and automorphisms of hierarchically hyperbolic spaces”, *Geom. Topol.* **21**:6 (2017), 3659–3758. MR
- [DHS20] M. G. Durham, M. F. Hagen, and A. Sisto, “Correction to the article Boundaries and automorphisms of hierarchically hyperbolic spaces”, *Geom. Topol.* **24**:2 (2020), 1051–1073. MR
- [DMS23] M. G. Durham, Y. N. Minsky, and A. Sisto, “Stable cubulations, bicombings, and barycenters”, *Geom. Topol.* **27**:6 (2023), 2383–2478. MR
- [FM12] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton Mathematical Series **49**, Princeton University Press, 2012. MR
- [Hag14] M. F. Hagen, “Weak hyperbolicity of cube complexes and quasi-arboreal groups”, *J. Topol.* **7**:2 (2014), 385–418. MR
- [Hag23] M. Hagen, “Non-colorable hierarchically hyperbolic groups”, *Internat. J. Algebra Comput.* **33**:2 (2023), 337–350. MR
- [HHL23] C. Horbez, J. Huang, and J. Lécureux, “Proper proximality in non-positive curvature”, *Amer. J. Math.* **145**:5 (2023), 1327–1364. MR
- [HHP23] T. Haettel, N. Hoda, and H. Petyt, “Coarse injectivity, hierarchical hyperbolicity and semihyperbolicity”, *Geom. Topol.* **27**:4 (2023), 1587–1633. MR
- [HMS22] M. Hagen, A. Martin, and A. Sisto, “Extra-large type Artin groups are hierarchically hyperbolic”, *Math. Ann.* **388**:1 (2024), 867–938. MR
- [HRSS25] M. Hagen, J. Russell, A. Sisto, and D. Spriano, “Equivariant hierarchically hyperbolic structures for 3-manifold groups via quasimorphisms”, *Ann. Inst. Fourier (Grenoble)* **75**:2 (2025), 769–828. MR
- [HS20] M. F. Hagen and T. Susse, “On hierarchical hyperbolicity of cubical groups”, *Israel J. Math.* **236**:1 (2020), 45–89. MR
- [HW08] F. Haglund and D. T. Wise, “Special cube complexes”, *Geom. Funct. Anal.* **17**:5 (2008), 1551–1620. MR
- [Isb80] J. R. Isbell, “Median algebra”, *Trans. Amer. Math. Soc.* **260**:2 (1980), 319–362. MR
- [Man24] G. Mangioni, “Short hierarchically hyperbolic groups, I: uncountably many coarse median structures”, preprint, 2024. arXiv 2410.09232
- [MM00] H. A. Masur and Y. N. Minsky, “Geometry of the complex of curves, II: Hierarchical structure”, *Geom. Funct. Anal.* **10**:4 (2000), 902–974. MR
- [MS26] G. Mangioni and A. Sisto, “Short hierarchically hyperbolic groups II: Quotients and the Hopf property for Artin groups”, *Adv. Math.* **486** (2026), art. id. 110736, 71 pp. MR
- [Rag25] K. Ragosta, “A marking graph for finite-type Artin groups”, preprint, 2025. arXiv 2508.10394
- [Rus22] J. Russell, “From hierarchical to relative hyperbolicity”, *Int. Math. Res. Not.* **2022**:1 (2022), 575–624. MR
- [Rus25] J. Russell, “Extensions of multicurve stabilizers are hierarchically hyperbolic”, *Geom. Topol.* **29**:6 (2025), 3187–3240. MR

[She25] S. Shepherd, “A cubulation with no factor system”, *Algebr. Geom. Topol.* **25**:1 (2025), 255–266. MR

[Ste99] M. Stern, *Semimodular lattices: theory and applications*, Encyclopedia of Mathematics and its Applications **73**, Cambridge University Press, 1999. MR

Received September 13, 2023. Revised January 15, 2026.

MARK HAGEN
SCHOOL OF MATHEMATICS
UNIVERSITY OF BRISTOL
BRISTOL
UNITED KINGDOM
markfhagen@posteo.net

GIORGIO MANGIONI
MAXWELL INSTITUTE AND DEPARTMENT OF MATHEMATICS
HERIOT-WATT UNIVERSITY
EDINBURGH
UNITED KINGDOM
gm2070@hw.ac.uk

ALESSANDRO SISTO
MAXWELL INSTITUTE AND DEPARTMENT OF MATHEMATICS
HERIOT-WATT UNIVERSITY
EDINBURGH
UNITED KINGDOM
a.sisto@hw.ac.uk

THE BÉNARD–CONWAY INVARIANT OF TWO-COMPONENT LINKS

ZEDAN LIU AND NIKOLAI SAVELIEV

The Bénard–Conway invariant of links in the 3-sphere is a Casson–Lin type invariant defined by counting irreducible $SU(2)$ -representations of the link group with fixed meridional traces. For two-component links with linking number one, the invariant has been shown to equal a symmetrized multivariable link signature. We extend this result to all two-component links with nonzero linking number. A key ingredient in the proof is an explicit calculation of the Bénard–Conway invariant for $(2, 2\ell)$ -torus links with the help of Chebyshev polynomials.

1. Introduction

The practice of defining invariants of links in 3-manifolds using unitary representations of the link group has a long history. Xiao-Song Lin [13] defined an invariant of knots in S^3 by counting irreducible $SU(2)$ -representations of the knot group sending the meridians to zero-trace matrices. Herald [10] and Heusener and Kroll [11] extended this construction by allowing matrices of a fixed trace which is not necessarily zero. The construction was further extended to links of more than one component by Harper and Saveliev [9] and Boden and Harper [2] by counting *projective* unitary representations. These invariants are closely related to gauge theory: for example, Floer homology theories of Daemi–Scaduto [7] and Kronheimer–Mrowka [12] can be viewed as categorifying the invariants of Lin and Harper–Saveliev, respectively.

The latest in this line of link invariants is the multivariable Casson–Lin type invariant h_L of Bénard and Conway [1]. It is defined for colored links $L \subset S^3$ by counting irreducible $SU(2)$ -representations of the link group sending the meridians to matrices of a fixed trace away from the roots of the multivariable Alexander polynomial. While this invariant is defined for links L of any number of components, it is best studied for two-component links. In particular, for an oriented ordered link $L = L_1 \cup L_2 \subset S^3$ with linking number $\text{lk}(L_1, L_2) = 1$, Bénard and Conway

Both authors were partially supported by the NSF Grant DMS-1952762.

MSC2020: primary 57K10; secondary 57K31.

Keywords: Casson–Lin invariant, character variety, multivariable link signature.

[1, Theorem 1.1] identify h_L with a symmetrized multivariable link signature σ_L of Cimasoni and Florens [5].

The purpose of this paper is to extend the result of B enard and Conway to arbitrary oriented ordered links $L_1 \cup L_2 \subset S^3$ of two components with linking number $\text{lk}(L_1, L_2) \neq 0$. To be precise, we prove the following theorem.

Theorem 1.1. *Let $L = L_1 \cup L_2 \subset S^3$ be a two-component oriented ordered link with $\text{lk}(L_1, L_2) \neq 0$ and, for any $(\alpha_1, \alpha_2) \in (0, \pi)^2$, write $(\omega_1, \omega_2) = (e^{2i\alpha_1}, e^{2i\alpha_2})$. If the multivariable Alexander polynomial of L satisfies $\Delta_L(\omega_1^{\epsilon_1}, \omega_2^{\epsilon_2}) \neq 0$ for all possible $\epsilon_1, \epsilon_2 = \pm 1$ then*

$$(1) \quad h_L(\alpha_1, \alpha_2) = -\frac{1}{2}(\sigma_L(\omega_1, \omega_2) + \sigma_L(\omega_1, \omega_2^{-1})).$$

The proof consists of two parts, just like the proof of [1, Theorem 1.1]. The first part shows that a crossing change within an individual component of L changes both sides of the formula of Theorem 1.1 by the same amount. This fact was only proved in [1, Theorem 1.1] for links with $\text{lk}(L_1, L_2) = 1$, but that proof easily extends to links with $\text{lk}(L_1, L_2) \neq 0$, as we explain in Section 3. After changing enough crossings within individual components of L , we only need to check that (1) holds for just one representative in each link homotopy class of L . According to Milnor [14], link homotopy classes of two component links are completely characterized by $\text{lk}(L_1, L_2)$; therefore, it is sufficient to prove Theorem 1.1 for $(2, 2\ell)$ -torus links L_ℓ with $\text{lk}(L_1, L_2) = \ell \neq 0$.

This second part of the proof occupies Sections 4 and 5 of the paper, where we compute the B enard–Conway invariant h_{L_ℓ} directly from its definition. The invariant h_{L_ℓ} , whose definition we recall in Section 2, is in an intersection number of two oriented curves in a 2-dimensional orbifold, traditionally referred to as a pillowcase. We come up with a parameterization of the pillowcase, in which the intersecting curves are given by explicit equations in terms of the Chebyshev polynomials; see Theorem 5.4. Checking the transversality and computing the intersection signs are then accomplished by a straightforward calculation.

This argument fails for links $L = L_1 \cup L_2$ with $\text{lk}(L_1, L_2) = 0$ because the base case, which is the link L_ℓ with $\ell = 0$, has a vanishing Alexander polynomial and hence its B enard–Conway invariant is not defined.

The next result follows from Theorem 1.1 and the properties of the Cimasoni–Florens signature [5]. It is proved in Section 6 together with Theorem 1.1.

Theorem 1.2. *For any two-component link $L = L_1 \cup L_2$ as in the statement of Theorem 1.1, the invariant $h_L(\alpha_1, \alpha_2)$ is independent of the orientation of the link L . Moreover, $h_L(\pi/2, \pi/2)$ equals minus the Murasugi signature [17] of the link L .*

Finally, we wish to mention how our work is related to that of Daemi and Scaduto [8]. Let L be a link in S^3 of any number of components with nonzero determinant.

Daemi and Scaduto [8] define irreducible instanton homology $I(L)$ as a $\mathbb{Z}/4$ graded abelian group which, in favorable circumstances, is generated at the chain level by the conjugacy classes of irreducible $SU(2)$ -representations of $\pi_1(S^3 - L)$ sending meridians to zero-trace matrices. In general, a perturbation may be necessary to achieve transversality. One expects that the Euler characteristic of $I(L)$ equals, up to an overall constant, the Bénard–Conway invariant $h_L(\pi/2, \dots, \pi/2)$. Daemi and Scaduto [8] further show that the Euler characteristic of $I(L)$ is proportional to the Murasugi signature of the link L , which matches our Theorem 1.2 in the special case of two-component links L with $\alpha_1 = \alpha_2 = \pi/2$.

2. Preliminaries

In this section, we will recall the definition of the Bénard–Conway invariant [1] for oriented links in the 3-sphere.

2.1. Colored links and colored braids. A μ -colored link $L \subset S^3$ is an oriented link whose components are partitioned into sublinks $L = L_1 \cup \dots \cup L_\mu$. Following Murakami [16], we will interpret these as the closures of colored braids.

Recall that the (Artin) braid group B_n on n strands is the finitely presented group with $n-1$ generators $\sigma_1, \sigma_2, \dots, \sigma_{n-1}$ subject to the relations $\sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}$ for each i , and $\sigma_i \sigma_j = \sigma_j \sigma_i$ for $|i - j| > 1$. Geometrically, a generator σ_i can be viewed as the isotopy class of the braid whose $(i + 1)$ -st strand crosses over the i -th strand. The closure $\widehat{\beta}$ of a braid $\beta \in B_n$ is the link obtained from β by connecting the lower endpoints of the braid and its upper endpoints with parallel strands. The link $\widehat{\beta}$ is canonically oriented by choosing the downward orientations on the strands of β .

A μ -colored braid is a braid $\beta \in B_n$ together with an assignment to each of its strands of an integer (called the *color*) in $\{1, 2, \dots, \mu\}$ via a surjective map. A μ -colored braid induces μ -colorings c and c' on the upper and lower endpoints of the braid, which are n -tuples of integers in $\{1, 2, \dots, \mu\}$. For any μ -coloring c , the μ -colored braids with $c' = c$ form a *colored braid group* $B_c \subset B_n$. For example, if $\mu = 1$, then $c = (1, \dots, 1)$ and B_c is simply the braid group B_n , and if $\mu = n$ and $c = (1, 2, \dots, n)$ then B_c is the pure braid group on n strands.

The closure $\widehat{\beta}$ of a μ -colored braid $\beta \in B_c$, obtained from β by connecting the lower endpoints of the braid with the upper endpoints with colored parallel strands, is a μ -colored link. A colored version of Alexander’s theorem states that every μ -colored link is the closure $\widehat{\beta}$ of a μ -colored braid $\beta \in B_c$, and a colored version of Markov’s theorem determines when two μ -colored braids have isotopic closures.

2.2. The Bénard–Conway invariant. Consider the Lie group $SU(2)$ of two-by-two unitary matrices with determinant one. We will be identifying it with the Lie group

Sp(1) of unit quaternions via

$$(2) \quad \begin{pmatrix} a & b \\ -\bar{b} & \bar{a} \end{pmatrix} \mapsto a + b \mathbf{j}$$

and using the language of matrices and unit quaternions interchangeably.

Let F_n be a free group on n generators x_1, \dots, x_n . The group B_n acts naturally on F_n via

$$(3) \quad x_j \sigma_i = \begin{cases} x_i x_{i+1} x_i^{-1} & \text{if } j = i, \\ x_i & \text{if } j = i + 1, \\ x_j & \text{otherwise.} \end{cases}$$

This action induces an action on the representation space

$$R(F_n) = \text{Hom}(F_n, \text{SU}(2)) = \text{SU}(2)^n.$$

Concretely, every braid $\beta \in B_n$ induces a homeomorphism $\beta : \text{SU}(2)^n \rightarrow \text{SU}(2)^n$ by the rule $(X_1, \dots, X_n)\beta = (X_1\beta, \dots, X_n\beta)$. Thus, for example, $(X_1, X_2)\sigma_1 = (X_1 X_2 X_1^{-1}, X_1)$ for the generator $\sigma_1 \in B_2$.

Let now $L = L_1 \cup \dots \cup L_\mu$ be an oriented μ -colored link. Represent it as the closure of a μ -colored braid $\beta \in B_c$ on n strands. Given a μ -tuple $\alpha = (\alpha_1, \dots, \alpha_\mu) \in (0, \pi)^\mu$, consider the representation space

$$R_n^{\alpha,c} = \{(X_1, X_2, \dots, X_n) \in \text{SU}(2)^n \mid \text{tr}(X_i) = 2 \cos(\alpha_{c_i}) \text{ for } i = 1, \dots, n\}.$$

Since the trace is preserved by conjugation, the homeomorphism $\beta : \text{SU}(2)^n \rightarrow \text{SU}(2)^n$ constructed above restricts to a homeomorphism $\beta : R_n^{\alpha,c} \rightarrow R_n^{\alpha,c}$ with the graph

$$\Gamma_\beta^{\alpha,c} = \{(X_1, X_2, \dots, X_n, X_1\beta, X_2\beta, \dots, X_n\beta) \in R_n^{\alpha,c} \times R_n^{\alpha,c}\}.$$

Note that the trivial braid in B_c gives rise to the graph which is just the diagonal

$$\Lambda_n^{\alpha,c} = \{(X_1, X_2, \dots, X_n, X_1, X_2, \dots, X_n) \in R_n^{\alpha,c} \times R_n^{\alpha,c}\}.$$

Since the product $X_1 \cdots X_n$ is preserved by the action of β (see formula (3)), one immediately concludes that both $\Gamma_\beta^{\alpha,c}$ and $\Lambda_n^{\alpha,c}$ are subspaces of the ambient space

$$H_n^{\alpha,c} = \{(X_1, X_2, \dots, X_n, Y_1, Y_2, \dots, Y_n) \in R_n^{\alpha,c} \times R_n^{\alpha,c} \mid X_1 \cdots X_n = Y_1 \cdots Y_n\}.$$

The group $\text{SO}(3)$ acts via conjugation on the spaces $\Gamma_\beta^{\alpha,c}$, $\Lambda_n^{\alpha,c}$ and $H_n^{\alpha,c}$. We restrict this action to irreducible representations, where it is free, and denote the quotient spaces by $\widehat{\Gamma}_\beta^{\alpha,c}$, $\widehat{\Lambda}_n^{\alpha,c}$ and $\widehat{H}_n^{\alpha,c}$. These are smooth open manifolds of dimensions

$$\dim \widehat{\Gamma}_\beta^{\alpha,c} = 2n - 3, \quad \dim \widehat{\Lambda}_n^{\alpha,c} = 2n - 3, \quad \dim \widehat{H}_n^{\alpha,c} = 4n - 6;$$

see [1, Lemma 3.4].

The key observation now is that the points of the intersection $\widehat{\Lambda}_n^{\alpha,c} \cap \widehat{\Gamma}_\beta^{\alpha,c} \subset \widehat{H}_n^{\alpha,c}$ are precisely the conjugacy classes of irreducible $SU(2)$ -representations of the link group

$$(4) \quad \pi_1(S^3 - L) = \langle x_1, \dots, x_n \mid x_i = x_i \beta \text{ for } i = 1, \dots, n \rangle$$

sending the meridians x_i to matrices $X_i \in SU(2)$ of trace $\text{tr}(X_i) = 2 \cos(\alpha_{c_i})$. From this point on, the definition of the Bénard–Conway invariant proceeds by making sense of the intersection number of $\widehat{\Lambda}_n^{\alpha,c}$ and $\widehat{\Gamma}_\beta^{\alpha,c}$ in $\widehat{H}_n^{\alpha,c}$. We will briefly outline the procedure and refer to [1] for detailed proofs.

Let $\Delta_L(t_1, \dots, t_\mu)$ be the multivariable Alexander polynomial of L , and, given a μ -tuple $\alpha = (\alpha_1, \dots, \alpha_\mu) \in (0, \pi)^\mu$, consider the finite set

$$(5) \quad S(\alpha) = \{(e^{\epsilon_1 2i\alpha_1}, \dots, e^{\epsilon_\mu 2i\alpha_\mu}) \mid \epsilon_i = \pm 1 \text{ for } i = 1, \dots, \mu\}.$$

Proposition 2.1 [1, Proposition 3.5]. *If $\Delta_\beta(\omega) \neq 0$ for all μ -tuples $\omega \in S(\alpha)$ then the intersection $\widehat{\Lambda}_n^{\alpha,c} \cap \widehat{\Gamma}_\beta^{\alpha,c} \subset \widehat{H}_n^{\alpha,c}$ is compact.*

Let $\alpha = (\alpha_1, \dots, \alpha_\mu)$ satisfy the condition of Proposition 2.1. Since $\widehat{\Lambda}_n^{\alpha,c} \cap \widehat{\Gamma}_\beta^{\alpha,c} \subset \widehat{H}_n^{\alpha,c}$ is compact, the graph $\widehat{\Gamma}_\beta^{\alpha,c}$ can be perturbed if necessary using a perturbation with compact support to make the intersection $\widehat{\Lambda}_n^{\alpha,c} \cap \widehat{\Gamma}_\beta^{\alpha,c} \subset \widehat{H}_n^{\alpha,c}$ transversal and hence a compact 0-dimensional manifold.

Next, we will orient all the manifolds in question. Denote by $\mathbb{S}(\theta)$ the conjugacy class of matrices in $SU(2)$ with the trace $2 \cos \theta$. Assuming that $0 < \theta < \pi$, the conjugacy classes $\mathbb{S}(\theta)$ are naturally homeomorphic to each other and to the standard 2-sphere S^2 . Choose an (arbitrary) orientation on S^2 . The space $R_n^{\alpha,c}$ is a product of the spheres $\mathbb{S}(\alpha_{c_i})$, $i = 1, \dots, n$, and we will endow it with the product orientation. The spaces $\Lambda_n^{\alpha,c}$ and $\Gamma_\beta^{\alpha,c}$, which are diffeomorphic to $R_n^{\alpha,c}$ via projection onto the first n factors, will be endowed with the induced orientations. To orient $H_n^{\alpha,c}$, consider the map $f_n : R_n^{\alpha,c} \times R_n^{\alpha,c} \rightarrow SU(2)$ given by $f_n(X_1, \dots, X_n, Y_1, \dots, Y_n) = X_1 \cdots X_n (Y_1 \cdots Y_n)^{-1}$. Observe that $H_n^{\alpha,c} = f_n^{-1}(1)$, so that we can pull back the canonical orientation of $SU(2)$ to obtain an orientation on $H_n^{\alpha,c}$. Since the adjoint action of $SO(3)$ on each $\mathbb{S}(\theta)$ is orientation preserving, we can endow $\widehat{\Lambda}_n^{\alpha,c}$, $\widehat{\Gamma}_\beta^{\alpha,c}$, and $\widehat{H}_n^{\alpha,c}$ with the induced quotient orientation.

The intersection number of $\widehat{\Lambda}_n^{\alpha,c}$ and $\widehat{\Gamma}_\beta^{\alpha,c}$ will be denoted by $\langle \widehat{\Lambda}_n^{\alpha,c}, \widehat{\Gamma}_\beta^{\alpha,c} \rangle_{\widehat{H}_n^{\alpha,c}}$. The following result, which is proved in [1, Section 3.3], ensures that it only depends on the isotopy class of the closure of β .

Proposition 2.2. *Under the assumptions of Proposition 2.1, the intersection number $\langle \widehat{\Lambda}_n^{\alpha,c}, \widehat{\Gamma}_\beta^{\alpha,c} \rangle_{\widehat{H}_n^{\alpha,c}}$ is preserved by the Markov moves.*

We will summarize the above construction in the following definition.

Definition 2.3. Let L be a μ -colored link in S^3 and $\alpha \in (0, \pi)^\mu$ a μ -tuple. Let $\beta \in B_c$ be a colored braid of n strands whose closure is L . Suppose that $\Delta_L(\omega_\epsilon) \neq 0$ for all $\omega_\epsilon \in S(\alpha)$. Then the *Bénard–Conway invariant* of L is well-defined by the formula

$$(6) \quad h_L(\alpha) = \langle \widehat{\Lambda}_n^{\alpha,c}, \widehat{\Gamma}_\beta^{\alpha,c} \rangle_{\widehat{H}_n^{\alpha,c}}.$$

3. Inductive step for two-component links

From now on, we will restrict ourselves to 2-colored links L of two components, which are just ordered two-component links. In this section, we will investigate what happens to the two sides of the formula (1) under a crossing change within a single component of L .

Theorem 3.1. *Under the assumptions of Theorem 1.1, the two sides of formula (1) stay well-defined and change by the same amount when a crossing change occurs within one of the components of the link.*

The rest of this section will be devoted to the proof of Theorem 3.1. Given a pair $\alpha = (\alpha_1, \alpha_2) \in (0, \pi)^2$, the set defined in (5) becomes

$$S(\alpha) = \{(e^{\epsilon_1 2i\alpha_1}, e^{\epsilon_2 2i\alpha_2}) \mid \epsilon_1 = \pm 1, \epsilon_2 = \pm 1\}.$$

In addition, define the sets

$$S_j(\alpha) = \{(e^{\epsilon_1 2i\alpha_1}, e^{\epsilon_2 2i\alpha_2}) \in S(\alpha) \mid \epsilon_j = 1\}, \quad j = 1, 2,$$

and denote by $\nabla_L(t_1, t_2)$ the multivariable Conway potential function [4]. Recall that $\nabla_L(t_1, t_2)$ equals $\Delta_L(t_1^2, t_2^2)$ up to multiplication by the powers of $\pm t_1$ and $\pm t_2$.

Lemma 3.2 [1, Proposition 5.10 and Remark 5.11]. *Let $L = L_1 \cup L_2$ be a two-component oriented ordered link and denote by L_+ the link obtained from L by a negative crossing change within a component L_j of the link L . Suppose that $\alpha = (\alpha_1, \alpha_2) \in (0, \pi)^2$ is such that, for all $\omega = (\omega_1, \omega_2) \in S_j(\alpha)$, one has $\omega_1^2 \neq 1$, $\omega_2^2 \neq 1$ and $\omega_1 \omega_2 \neq 1$. If $\nabla_L(\omega^{1/2}) \neq 0$ and $\nabla_{L_+}(\omega^{1/2}) \neq 0$ then*

$$h_{L_+}(\alpha) - h_L(\alpha) = \#\{\omega \in S_j(\alpha) \mid \nabla_{L_+}(\omega^{1/2})\nabla_L(\omega^{1/2}) < 0\}.$$

Lemma 3.3 [1, Lemma 6.2]. *Let $L = L_1 \cup L_2$ be an ordered two-component link with $\text{lk}(L_1, L_2) \neq 0$ and suppose that $\omega \in (S^1 \setminus \{1\})^2$ is not a root of $\Delta_L(t_1, t_2)$. Then the Cimasoni–Florens signature $\sigma_L(\omega)$ is well defined and*

$$\sigma_L(\omega) \equiv 2 + \text{lk}(L_1, L_2) + \text{sign}(\nabla_L(\omega^{1/2})) \pmod{4}.$$

In addition, suppose that L_+ is obtained from L by a negative crossing change within one of the components of L , and that ω is not a root of either $\Delta_{L_+}(t_1, t_2)$ or $\Delta_L(t_1, t_2)$. Then $\sigma_{L_+}(\omega) - \sigma_L(\omega)$ is either 0 or -2 .

Corollary 3.4. *Let L be an ordered two-component link with $\text{lk}(L_1, L_2) \neq 0$ and L_+ a link obtained from L by a negative crossing change within one of its components. Suppose that $\omega \in (S^1 \setminus \{1\})^2$ is such that $\nabla_L(\omega^{1/2}) \neq 0$ and $\nabla_{L_+}(\omega^{1/2}) \neq 0$. Then*

$$\sigma_{L_+}(\omega) - \sigma_L(\omega) = \begin{cases} 0 & \text{if } \nabla_{L_+}(\omega^{1/2})\nabla_L(\omega^{1/2}) > 0, \\ -2 & \text{if } \nabla_{L_+}(\omega^{1/2})\nabla_L(\omega^{1/2}) < 0. \end{cases}$$

We are now ready to sketch a proof of Theorem 3.1, which will be a straightforward modification of the proof of Theorem 6.4 in [1]. That theorem is first proved under the additional assumption that α_1 and α_2 are transcendental. This assumption is then removed by showing that the invariant $h_L(\alpha)$ is locally constant in α . Only the first part of the proof needs to be modified.

Observe first that a crossing change within one component of L does not make the multivariable Alexander polynomial vanish, which ensures that the Bénard–Conway invariant of L_+ is well-defined. This follows from the Torres formula

$$\Delta_{L_1 \cup L_2}(t_1, 1) \doteq \frac{t_1^{\text{lk}(L_1, L_2)} - 1}{t_1 - 1} \Delta_{L_1}(t_1)$$

(see [18]) and our assumption that $\text{lk}(L_1, L_2) \neq 0$. Next, we will show that the two sides of (1) change by the same amount under the crossing change. Assume without loss of generality that L_+ is obtained from $L = L_1 \cup L_2$ by a negative crossing change in L_1 . According to Lemma 3.2, we have

$$h_{L_+}(\alpha) - h_L(\alpha) = \#\{\omega \in S_1(\alpha) \mid \nabla_{L_+}(\omega^{1/2})\nabla_L(\omega^{1/2}) < 0\}.$$

Note that $S_1(\alpha)$ has exactly two elements, (ω_1, ω_2) and $(\omega_1, \omega_2^{-1})$, where $\omega_1 = e^{2i\alpha_1}$ and $\omega_2 = e^{2i\alpha_2}$. It now follows from Corollary 3.4 that

$$\begin{aligned} h_{L_+}(\alpha) - h_L(\alpha) &= -\frac{1}{2}(\sigma_{L_+}(\omega_1, \omega_2) - \sigma_L(\omega_1, \omega_2)) - \frac{1}{2}(\sigma_{L_+}(\omega_1, \omega_2^{-1}) - \sigma_L(\omega_1, \omega_2^{-1})) \\ &= -\frac{1}{2}(\sigma_{L_+}(\omega_1, \omega_2) + \sigma_{L_+}(\omega_1, \omega_2^{-1})) + \frac{1}{2}(\sigma_L(\omega_1, \omega_2) + \sigma_L(\omega_1, \omega_2^{-1})), \end{aligned}$$

which completes the proof.

4. SU(2)-representations of torus links

To complete the proof of Theorem 1.1, it is sufficient to verify the formula (1) for $(2, 2\ell)$ -torus links L_ℓ with $\ell \neq 0$. Our convention here is that $\ell > 0$ gives the right-handed torus link and $\ell < 0$ the left-handed one. The verification will take up the rest of the paper. We begin in this section by describing the irreducible SU(2)-representations of the link group of L_ℓ with fixed meridional traces.

4.1. Geometry of SU(2). We will continue to identify SU(2) matrices with unit quaternions as in (2). Any $q \in \text{SU}(2)$ can then be written in the form $q = \cos \alpha + \sin \alpha Q = e^{\alpha Q}$, where $\alpha \in [0, \pi]$ and Q is a purely imaginary unit quaternion. This expression is unique except when $q = \pm 1$. Using all unit quaternions, $q = \cos \alpha + \sin \alpha Q$ can be conjugated to $\cos \alpha + \sin \alpha \mathbf{i} = e^{\alpha \mathbf{i}}$. Using only unit complex numbers, $q = \cos \alpha + \sin \alpha Q$ can be conjugated to $\cos \alpha + \sin \alpha (\cos \beta \mathbf{i} + \sin \beta \mathbf{j})$ for some $\beta \in [0, \pi]$. Alternatively, it can be expressed as $\cos \gamma (\cos \beta + \sin \beta \mathbf{i}) + \sin \gamma \mathbf{j}$ for $\gamma \in [0, \pi]$ and $\beta \in [0, 2\pi]$ with $\cos \gamma \cos \beta = \cos \alpha$ since the real part is conjugation invariant.

4.2. Counting the representations. Let us first assume that $\ell > 0$ and consider the presentation of the link group given by

$$\pi_1(S^3 \setminus L_\ell) = \langle x_1, x_2 \mid (x_1 x_2)^\ell = (x_2 x_1)^\ell \rangle,$$

where x_1 and x_2 are the meridians of the two components of L_ℓ . For a fixed choice of $(\alpha_1, \alpha_2) \in (0, \pi)^2$, we wish to describe the conjugacy classes of irreducible representations $\rho : \pi_1(S^3 \setminus L_\ell) \rightarrow \text{SU}(2)$ sending the meridians of the components of L_ℓ to unit quaternions with respective real parts $\cos(\alpha_j)$, $j = 1, 2$. Since

$$(x_2 x_1)^\ell = (x_2 x_1)^\ell x_2 x_2^{-1} = x_2 (x_1 x_2)^\ell x_2^{-1},$$

the relation $(x_1 x_2)^\ell = (x_2 x_1)^\ell$ is equivalent to $(x_1 x_2)^\ell$ commuting with x_2 , and by symmetry, also with x_1 . Therefore, $(x_1 x_2)^\ell$ belongs to the center of $\pi_1(S^3 \setminus L_\ell)$. Since ρ is irreducible, we must have $(\rho(x_1)\rho(x_2))^\ell = \pm 1$. We end up looking for noncommuting unit quaternions $\rho(x_1)$ and $\rho(x_2)$ with prescribed real parts such that $\rho(x_1)\rho(x_2)$ is an ℓ -th root of ± 1 different from ± 1 (because otherwise ρ is reducible). Conjugate ρ so that

$$\rho(x_1) = \cos \alpha_1 + \sin \alpha_1 (\cos \varphi \mathbf{i} + \sin \varphi \mathbf{j}) \quad \text{and} \quad \rho(x_2) = \cos \alpha_2 + \sin \alpha_2 \mathbf{i}$$

for some $\varphi \in (0, \pi)$. The condition above then means that the real part of $\rho(x_1)\rho(x_2)$ equals $\cos(\pi m/\ell)$ or, equivalently, that

$$(7) \quad \cos \alpha_1 \cos \alpha_2 - \sin \alpha_1 \sin \alpha_2 \cos \varphi = \cos(\pi m/\ell), \quad m = 1, 2, \dots, \ell - 1.$$

Proposition 4.1. *Let L_ℓ be a $(2, 2\ell)$ -torus link, $\ell \neq 0$. Given $(\alpha_1, \alpha_2) \in (0, \pi)^2$, the number of conjugacy classes of irreducible representations $\pi_1(S^3 \setminus L_\ell) \rightarrow \text{SU}(2)$ sending the meridians to unit quaternions with real parts $\cos \alpha_j$, $j = 1, 2$, equals the number of integers $m \in \{1, \dots, |\ell| - 1\}$ such that*

$$\cos(\pi m/\ell) \in (\cos(\alpha_1 + \alpha_2), \cos(\alpha_1 - \alpha_2)).$$




Figure 1. Counting conjugacy classes, $\ell = 3$.

Proof. Let us assume that $\ell > 0$ since the argument for $\ell < 0$ is similar. Use trigonometric identities to express (7) in terms of $\alpha_1 - \alpha_2$ and $\alpha_1 + \alpha_2$ as

$$\cos(\alpha_1 - \alpha_2) \sin^2(\varphi/2) + \cos(\alpha_1 + \alpha_2) \cos^2(\varphi/2) = \cos(\pi m/\ell),$$

or, equivalently,

$$\cos(\alpha_1 - \alpha_2)(1 - t) + \cos(\alpha_1 + \alpha_2)t = \cos(\pi m/\ell),$$

where $t = \cos^2(\varphi/2) \in [0, 1]$. The latter equation has a unique solution $t \in [0, 1]$, and therefore (7) has a unique solution $\varphi \in (0, \pi)$ if and only if $\cos(\pi m/\ell) \in (\cos(\alpha_1 + \alpha_2), \cos(\alpha_1 - \alpha_2))$. \square

Example 4.2. The cases of $\ell = 1$ and $\ell = -1$ correspond to the two oriented Hopf links. In both cases, the link group has no irreducible $SU(2)$ -representations. The former case served as the base of induction in [1]. The count of conjugacy classes for $\ell = 3$ is depicted in Figure 1. The diamond shapes in the figure have to do with the Alexander polynomial of the link as discussed in detail in the following subsection.

4.3. The Alexander polynomial. The count of irreducible representations in Proposition 4.1 may jump as α_1 and α_2 change. The following proposition shows that these jumps occur at the roots of the multivariable Alexander polynomial, that is, exactly where the Bénard–Conway invariant is not defined.

Proposition 4.3. *Let L_ℓ be a $(2, 2\ell)$ -torus link with $\ell \neq 0$ and $\Delta_{L_\ell}(t_1, t_2)$ its multivariable Alexander polynomial. The number of conjugacy classes of irreducible representations $\pi_1(S^3 \setminus L_\ell) \rightarrow SU(2)$ sending the meridians to unit quaternions with real parts $\cos(\alpha_j)$, $j = 1, 2$, is a locally constant function on the complement of the set given by the equation $\Delta_{L_\ell}(\omega_1^{\pm 1}, \omega_2^{\pm 1}) = 0$, where $(\omega_1, \omega_2) = (e^{2i\alpha_1}, e^{2i\alpha_2})$.*

Proof. The jumps in the count of irreducible representations in Proposition 4.1 occur exactly when $\cos(\alpha_1 \pm \alpha_2) = \cos(\pi m/\ell)$ with $m = 1, \dots, |\ell| - 1$, that is,

when

$$(8) \quad \alpha_1 + \alpha_2 = \frac{\pi m}{|\ell|} \quad \text{or} \quad \alpha_1 + \alpha_2 = \pi + \frac{\pi m}{|\ell|} \quad \text{or} \quad \alpha_1 - \alpha_2 = \pm \frac{\pi m}{|\ell|}.$$

These (α_1, α_2) are precisely the solutions of the equation $\Delta_{L_\ell}(\omega_1^{\pm 1}, \omega_2^{\pm 1}) = 0$, which can be easily verify using the formula

$$\Delta_{L_\ell}(t_1, t_2) \doteq \frac{(t_1 t_2)^{|\ell|} - 1}{t_1 t_2 - 1}$$

for the multivariable Alexander polynomial of L_ℓ due to Milnor [15]. □

Remark 4.4. When $\ell = 0$, the torus link L_ℓ is just the unlink of two components. Its multivariable Alexander polynomial is identically zero hence its B enard–Conway invariant is not defined. This is the reason behind our assumption that $\ell \neq 0$.

5. The pillowcase and intersection theory

In this section, we will compute the B enard–Conway invariants of $(2, 2\ell)$ -torus links as the intersection number of the manifolds $\widehat{\Lambda}_n^{\alpha, c}$ and $\widehat{\Gamma}_\beta^{\alpha, c}$ inside of $\widehat{H}_n^{\alpha, c}$. This will involve counting the representations described in Proposition 4.1 with plus/minus signs, after making sure that the intersection in question is transversal.

5.1. The setup. Let L_ℓ be a $(2, 2\ell)$ -torus link and assume that ℓ is positive. The case of negative ℓ can be treated in a similar manner, and both cases will be discussed in detail when we perform explicit calculations in Section 5.5. Let $\alpha = (\alpha_1, \alpha_2)$ be an arbitrary point in $(0, \pi)^2$ away from the set (8). This guarantees that the B enard–Conway invariant of L_ℓ is well-defined; see Proposition 4.3. We will view the link L_ℓ as a 2-colored link which is the closure of the 2-colored braid $\beta = \sigma_1^{2\ell} \in B_c$ with the 2-coloring $c = (1, 2)$. The B enard–Conway invariant of L_ℓ is then the intersection number of $\widehat{\Lambda}_2^{\alpha, c}$ and $\widehat{\Gamma}_\beta^{\alpha, c}$ inside of $\widehat{H}_2^{\alpha, c}$. Our first goal will be to parameterize these manifolds.

5.2. The pillowcase. Recall that $\widehat{H}_2^{\alpha, c}$ is an open 2-manifold obtained by removing the conjugacy classes of reducible representations from the orbifold $H_2^{\alpha, c} / \text{SO}(3)$, where

$$H_2^{\alpha, c} = \{(X_1, X_2, Y_1, Y_2) \in R_2^{\alpha, c} \times R_2^{\alpha, c} \mid X_1 X_2 = Y_1 Y_2\}.$$

In the special case of $\alpha_1 = \alpha_2 = \pi/2$, the orbifold $H_2^{\alpha, c} / \text{SO}(3)$ is referred to as a pillowcase; see for instance Lin [13]. We will extend this terminology to the general case.

After conjugation, we may assume that $X_2 = e^{\alpha_2 i}$ and $X_1 = e^{\alpha_1 P_1}$, where $P_1 = i e^{-k\varphi}$ with $\varphi \in [0, \pi]$; see Section 4.2. Write $Y_1 = e^{\alpha_1 Q_1}$ and $Y_2 = e^{\alpha_2 Q_2}$, where Q_1 and Q_2 are purely imaginary unit quaternions. For any given Y_1 , the

equation $X_1 X_2 = Y_1 Y_2$ can be uniquely solved for Y_2 if and only if the real part of $Y_1^{-1} X_1 X_2$ matches that of Y_2 , that is,

$$\operatorname{Re}(e^{-\alpha_1} Q_1 e^{\alpha_1} i e^{-k\varphi} e^{\alpha_2} i) = \cos \alpha_2.$$

Write $Q_1 = u \mathbf{i} + v \mathbf{j} + w \mathbf{k}$, where $u^2 + v^2 + w^2 = 1$. The above equation is equivalent to

$$(9) \quad (\sin \alpha_1 \cos \alpha_2 \cos \varphi + \cos \alpha_1 \sin \alpha_2)u + \sin \alpha_1 \cos \alpha_2 \sin \varphi v - \sin \alpha_1 \sin \alpha_2 \sin \varphi w \\ = \sin \alpha_1 \cos \alpha_2 + \cos \alpha_1 \sin \alpha_2 \cos \varphi.$$

We will handle the case of $\varphi \in (0, \pi)$ first. In this case, (9) has the form $\mathbf{n} \cdot (u, v, w) = d$, where $d = \sin \alpha_1 \cos \alpha_2 + \cos \alpha_1 \sin \alpha_2 \cos \varphi$ and

$$\mathbf{n} = \begin{pmatrix} \sin \alpha_1 \cos \alpha_2 \cos \varphi + \cos \alpha_1 \sin \alpha_2 \\ \sin \alpha_1 \cos \alpha_2 \sin \varphi \\ -\sin \alpha_1 \sin \alpha_2 \sin \varphi \end{pmatrix}$$

is a nonzero vector (because its third coordinate is never zero). Therefore, (9) describes a plane in the space \mathbb{R}^3 of purely imaginary quaternions, and our task becomes describing the intersection of this plane with the unit sphere S^2 of purely imaginary unit quaternions given by the equation $u^2 + v^2 + w^2 = 1$. Note that the point on the plane closest to the origin is given by $d\mathbf{n}/|\mathbf{n}|^2$. One can easily see that the distance from this point to the origin equals

$$(10) \quad \frac{d}{|\mathbf{n}|} = \frac{d}{\sqrt{d^2 + \sin^2 \alpha_2 \sin^2 \varphi}}$$

and that this distance is strictly less than one, making the intersection of the plane with the unit sphere S^2 a circle S_φ^1 for any choice of $\varphi \in (0, \pi)$. Therefore, away from the points with $\varphi = 0$ and $\varphi = \pi$, the orbifold $H_2^{\alpha,c}/\text{SO}(3)$ is homeomorphic to a cylinder $S^1 \times (0, \pi)$. It can be parameterized as follows.

The circle S_φ^1 contains the vector $P_1 = i e^{-k\varphi} = \cos \varphi \mathbf{i} + \sin \varphi \mathbf{j}$ corresponding to the solution with (X_1, X_2, X_1, X_2) of (9). Let

$$\mathbf{v}_1 = P_1 - \frac{d\mathbf{n}}{|\mathbf{n}|^2} \quad \text{and} \quad \mathbf{v}_2 = \frac{\mathbf{n}}{|\mathbf{n}|} \times \mathbf{v}_1.$$

Since the vectors \mathbf{v}_1 and \mathbf{v}_2 are orthogonal to each other, we can write

$$(11) \quad Q_1 = \frac{d\mathbf{n}}{|\mathbf{n}|^2} + \cos \theta \mathbf{v}_1 + \sin \theta \mathbf{v}_2$$

for some $\theta \in [0, 2\pi]$; see Figure 2. Therefore, away from the points with $\varphi = 0$ and $\varphi = \pi$, $H_2^{\alpha,c}/\text{SO}(3)$ is a cylinder parameterized by $(\varphi, \theta) \in (0, \pi) \times [0, 2\pi]$ with the intervals $\theta = 0$ and $\theta = 2\pi$ identified.




Figure 2. Vectors v_1, v_2 and points P_1, Q_1 on the circle S^1_φ .

We now turn to the full space $H_2^{\alpha,c} / \text{SO}(3)$. It is a compactification of the cylinder described above by the points with $\varphi = 0$ and $\varphi = \pi$. The type of compactification one obtains depends on α_1 and α_2 :

- (1) The special case of $\alpha_1 = \alpha_2 = \pi/2$ is the original case investigated by Lin [13]. The traceless condition forces the representations with $\varphi = 0$ and $\varphi = \pi$ to be irreducible for all $\theta \in (0, 2\pi)$. The resulting $\hat{H}_2^{\alpha,c}$ is a sphere with four points removed.
- (2) The case of $\alpha_1 = \alpha_2 \neq \pi/2$, in which $\sin(\alpha_1 - \alpha_2) = 0$ but $\sin(\alpha_1 + \alpha_2) \neq 0$, was originally investigated by Heusener and Kroll [11] and used in [1]. In this case, all representations with $\varphi = \pi$ and $\theta \in (0, 2\pi)$ are irreducible and the points (π, θ) and $(\pi, 2\pi - \theta)$ are identified; see Figure 3. Topologically, $\hat{H}_2^{\alpha,c}$ is a 2-sphere with three points removed.
- (3) In the special case of $\alpha_1 \neq \alpha_2$ and $\alpha_1 + \alpha_2 = \pi$ we have $\sin(\alpha_1 + \alpha_2) = 0$ but $\sin(\alpha_1 - \alpha_2) \neq 0$. This time around, the representations with $\varphi = 0$ and $\theta \in (0, 2\pi)$ are irreducible and $\hat{H}_2^{\alpha,c}$ is a mirror image of Figure 3.

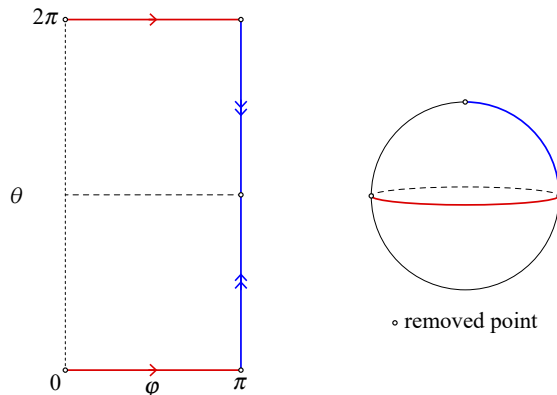


Figure 3. Special case of the parametrization when $\alpha_1 = \alpha_2 \neq \pi/2$.




Figure 4. Generic case with $\sin(\alpha_1 \pm \alpha_2) \neq 0$.

(4) For a generic choice of α_1 and α_2 with $\sin(\alpha_1 \pm \alpha_2) \neq 0$, the above calculation extends to the points with $\varphi = 0$ and $\varphi = \pi$ giving rise in each case to a single reducible representation. Topologically, the cylinder compactifies to a 2-sphere, the reducible locus collapses along one direction, thus $\widehat{H}_2^{\alpha,c}$ is a 2-sphere with two points removed; see Figure 4.

All of the resulting orbifolds $H_2^{\alpha,c} / \text{SO}(3)$ will be referred to as pillowcases. This includes the novel cases (3) and (4), which arise because we allow distinct α_1 and α_2 . The calculation that follows will be the same in all four cases because it happens away from the points with $\varphi = 0$ and $\varphi = \pi$.

5.3. Intersection theory in the pillowcase. The diagonal $\widehat{\Lambda}_2^{\alpha,c}$ is a subspace of $\widehat{H}_2^{\alpha,c}$ given by the equation $X_1 = Y_1$. In our parameterization of the pillowcase, $\widehat{\Lambda}_2^{\alpha,c}$ is then exactly the subspace where $P_1 = Q_1$ or $\theta = 0$. It is shown in red in Figure 4.

Next, we will parametrize $\widehat{\Gamma}_\beta^{\alpha,c}$. Since θ is the angle between the vectors $P_1 - d\mathbf{n}/|\mathbf{n}|^2$ and $Q_1 - d\mathbf{n}/|\mathbf{n}|^2$, we have

$$\cos \theta = \frac{(P_1 - d\mathbf{n}/|\mathbf{n}|^2) \cdot (Q_1 - d\mathbf{n}/|\mathbf{n}|^2)}{1 - d^2/|\mathbf{n}|^2}.$$

Using formula (10), this simplifies to

$$(12) \quad \cos \theta = \frac{(|\mathbf{n}|^2 P_1 - d\mathbf{n}) \cdot Q_1}{\sin^2 \alpha_2 \sin^2 \varphi}.$$

The lemmas that follow simplify this formula further and eventually lead to Theorem 5.4, which identifies the right-hand side of (12) as a Chebyshev polynomial.

Lemma 5.1. *The right-hand side of formula (12) is a polynomial in $\cos \varphi$, which will be denoted by $P(\cos \varphi)$.*

Proof. We will proceed by simplifying (12) while keeping track of its dependence on φ . By a direct calculation,

$$(13) \quad |\mathbf{n}|^2 P_1 - d\mathbf{n} = \begin{pmatrix} \sin^2 \varphi A_1 \\ \sin \varphi A_2 \\ \sin \varphi A_3 \end{pmatrix},$$

where A_1, A_2 and A_3 are real-valued polynomials in $\cos \varphi$ of degrees $\deg A_1 = 1$, $\deg A_2 = 2$ and $\deg A_3 = 1$. To compute Q_1 , recall that Y_1 is the image of X_1 under the action of the braid $\beta = \sigma_1^{2\ell}$, where $(X_1, X_2)\sigma_1 = (X_1 X_2 X_1^{-1}, X_1)$; see Section 2.2. Therefore,

$$(X_1, X_2)\sigma_1^{2\ell} = ((X_1 X_2)^\ell X_1 (X_1 X_2)^{-\ell}, (X_1 X_2)^\ell X_2 (X_1 X_2)^{-\ell})$$

and

$$Q_1 = (X_1 X_2)^\ell P_1 (X_1 X_2)^{-\ell}.$$

Write

$$\begin{aligned} X_1 X_2 &= (\cos \alpha_1 + \sin \alpha_1 \cos \varphi \mathbf{i} + \sin \alpha_1 \sin \varphi \mathbf{j})(\cos \alpha_2 + \sin \alpha_2 \mathbf{i}) \\ &= (\cos \alpha_1 \cos \alpha_2 - \sin \alpha_1 \sin \alpha_2 \cos \varphi) \\ &\quad + (\cos \alpha_1 \sin \alpha_2 + \sin \alpha_1 \cos \alpha_2 \cos \varphi) \mathbf{i} \\ &\quad + \sin \alpha_1 \cos \alpha_2 \sin \varphi \mathbf{j} - \sin \alpha_1 \sin \alpha_2 \sin \varphi \mathbf{k} \\ &= U_1 + \sin \varphi V_1 \mathbf{j}, \end{aligned}$$

where U_1 and V_1 are complex-valued polynomials in $\cos \varphi$ of degrees $\deg U_1 = 1$ and $\deg V_1 = 0$. Using induction on ℓ one can easily see that

$$(X_1 X_2)^\ell = U_\ell + \sin \varphi V_\ell \mathbf{j},$$

where U_ℓ and V_ℓ are complex-valued polynomials in $\cos \varphi$ of degrees $\deg U_\ell = \ell$ and $\deg V_\ell = \ell - 1$. A straightforward calculation then shows that

$$(14) \quad Q_1 = \begin{pmatrix} B_1 \\ \sin \varphi B_2 \\ \sin \varphi B_3 \end{pmatrix},$$

where B_1, B_2 and B_3 are real-valued polynomials in $\cos \varphi$ of degrees $\deg B_1 = 2\ell + 1$, $\deg B_2 = 2\ell$ and $\deg B_3 = 2\ell$. By taking the dot product of (13) and (14) we conclude that the numerator of (12) has the form $\sin^2 \varphi$ times a polynomial in $\cos \varphi$ of degree at most $2\ell + 2$. The factors of $\sin^2 \varphi$ in the numerator and the denominator of (12) cancel, thereby finishing the proof. \square

Lemma 5.2. *The polynomial $P(\cos \varphi)$ has degree 2ℓ and its leading coefficient is $2^{2\ell-1} \sin^{2\ell} \alpha_1 \sin^{2\ell} \alpha_2$.*

Proof. We will follow the proof of the previous lemma but make our calculations more precise. One can check using induction on ℓ that

$$(15) \quad \mathbf{i} \bar{U}_\ell = \cos \varphi V_\ell + \mathbf{i} \cot \alpha_1 V_\ell + \dots,$$

where the dots stand for a polynomial in $\cos \varphi$ of degree at most $\ell - 2$. A tedious but direct calculation of the polynomials in formulas (13) and (14) then yields the formulas

$$\begin{aligned} A_1 &= \sin^2 \alpha_1 \sin^2 \alpha_2 \cos \varphi - \sin \alpha_1 \sin \alpha_2 \cos \alpha_1 \cos \alpha_2, \\ A_2 &= -\sin^2 \alpha_1 \sin^2 \alpha_2 \cos^2 \varphi + \sin \alpha_1 \sin \alpha_2 \cos \alpha_1 \cos \alpha_2 \cos \varphi + \sin^2 \alpha_2, \\ A_3 &= \sin \alpha_1 \cos \alpha_1 \sin^2 \alpha_2 \cos \varphi + \sin^2 \alpha_1 \sin \alpha_2 \cos \alpha_2, \end{aligned}$$

and

$$\begin{aligned} B_1 &= [(|U_\ell|^2 + |V_\ell|^2 \cos^2 \varphi) \cos \varphi - 2 \operatorname{Im}(U_\ell \bar{V}_\ell) \cos^2 \varphi]_{2\ell+1} \\ &\quad - [|V_\ell|^2 \cos \varphi - 2 \operatorname{Im}(U_\ell \bar{V}_\ell)]_{2\ell-1}, \\ B_2 &= [-\operatorname{Re} V_\ell^2 \cos^2 \varphi + 2 \operatorname{Im}(U_\ell V_\ell) \cos \varphi + \operatorname{Re} U_\ell^2]_{2\ell} + [\operatorname{Re} V_\ell^2]_{2\ell-2}, \\ B_3 &= [-\operatorname{Im} V_\ell^2 \cos^2 \varphi - 2 \operatorname{Re}(U_\ell V_\ell) \cos \varphi + \operatorname{Im} U_\ell^2]_{2\ell-1} + [\operatorname{Im} V_\ell^2]_{2\ell-2}. \end{aligned}$$

The brackets in the formulas for B_1 , B_2 and B_3 contain polynomials whose degrees are indicated by the subscripts. Note that the first bracket in the formula for B_3 has degree $2\ell - 1$ rather than 2ℓ , which follows from (15). Another lengthy calculation shows that

$$\begin{aligned} A_1 B_1 + A_2 B_2 + A_3 B_3 &= 2 \sin^2 \alpha_1 \sin^2 \alpha_2 \cos^2 \varphi (\operatorname{Im} U_\ell - \cos \varphi \operatorname{Re} V_\ell)^2 \\ &\quad - 2 \sin \alpha_1 \sin \alpha_2 \cos \alpha_1 \cos \alpha_2 \cos \varphi (\operatorname{Im} U_\ell - \cos \varphi \operatorname{Re} V_\ell)^2 \\ &\quad + 2 \sin^2 \alpha_1 \cos^2 \varphi (\operatorname{Im} V_\ell)^2 + \dots, \end{aligned}$$

where the dots stand for a polynomial in $\cos \varphi$ of degree at most $2\ell - 1$. Using formula (15), this further simplifies to

$$A_1 B_1 + A_2 B_2 + A_3 B_3 = 2 \cos^2 \varphi \sin^2 \alpha_2 (\operatorname{Im} V_\ell)^2 + \dots,$$

which is a polynomial of degree 2ℓ . It follows that $P(\cos \varphi)$ is a polynomial of degree 2ℓ whose leading term matches that of the polynomial $2 \cos^2 \varphi (\operatorname{Im} V_\ell)^2$. Another induction shows that the leading term of $\operatorname{Im} V_\ell$ equals

$$(-1)^\ell 2^{\ell-1} \cos^{\ell-1} \varphi \sin^\ell \alpha_1 \sin^\ell \alpha_2,$$

which completes the proof. □

Lemma 5.3. *The polynomial $P(\cos \varphi)$ evaluates to, respectively, $\cos(2\ell(\alpha_1 + \alpha_2))$ and $\cos(2\ell(\alpha_1 - \alpha_2))$ at $\varphi = 0$ and $\varphi = \pi$.*

Proof. We will use formula (12), which defines $P(\cos \varphi)$ for $\varphi \in (0, \pi)$, to show that $P(\cos \varphi)$ tends to, respectively, $\cos(2\ell(\alpha_1 + \alpha_2))$ and $\cos(2\ell(\alpha_1 - \alpha_2))$ as $\varphi \rightarrow 0$ and $\varphi \rightarrow \pi$. We first calculate the limit as $\varphi \rightarrow 0$. Once we eliminate the factors $\sin^2 \varphi$ as in Lemma 5.1, the calculation reduces to evaluating the expression

$$\frac{A_1 B_1 + A_2 B_2 + A_3 B_3}{\sin^2 \alpha_2}$$

at $\varphi = 0$. Here, we use the notation from the proof of Lemma 5.2. It is easy to see that A_1, A_2, A_3 evaluate to

$$\begin{aligned} A_1 &= -\sin \alpha_1 \sin \alpha_2 \cos(\alpha_1 + \alpha_2), \\ A_2 &= \cos \alpha_1 \sin \alpha_2 \sin(\alpha_1 + \alpha_2), \\ A_3 &= \sin \alpha_1 \sin \alpha_2 \sin(\alpha_1 + \alpha_2). \end{aligned}$$

To evaluate B_1, B_2 and B_3 , we will keep track of $\sin \varphi$ in the formula $Q_1 = (X_1 X_2)^\ell P_1(X_1 X_2)^{-\ell}$ while setting $\cos \varphi$ equal to one. An induction on ℓ can be used to show that

$$(16) \quad (X_1 X_2)^\ell = q^\ell + \sin \varphi \sin \alpha_1 e^{i\alpha_1} \left(q^{-\ell} \frac{1 - q^{2\ell}}{1 - q^2} \right) j + \dots,$$

where $q = e^{i(\alpha_1 + \alpha_2)}$ and the dots stand for higher-degree polynomials in $\sin \varphi$. Using the identity

$$\frac{1 - q^{2\ell}}{1 - q^2} = e^{i(\ell-1)(\alpha_1 + \alpha_2)} \frac{\sin(\ell(\alpha_1 + \alpha_2))}{\sin(\alpha_1 + \alpha_2)},$$

the above formulas can be written in trigonometric form. It now follows from the formula

$$Q_1 = (X_1 X_2)^\ell P_1(X_1 X_2)^{-\ell} = \begin{pmatrix} B_1 \\ \sin \varphi B_2 \\ \sin \varphi B_3 \end{pmatrix}$$

that, when $\varphi = 0$, we have

$$\begin{aligned} B_1 &= 1, \\ B_2 &= \cos(2\ell(\alpha_1 + \alpha_2)) \\ &\quad + 2 \sin \alpha_1 (-\sin \alpha_2 \cos(\ell(\alpha_1 + \alpha_2)) + \cos \alpha_2 \sin(\ell(\alpha_1 + \alpha_2))) \frac{\sin(\ell(\alpha_1 + \alpha_2))}{\sin(\alpha_1 + \alpha_2)}, \\ B_3 &= \sin(2\ell(\alpha_1 + \alpha_2)) \\ &\quad - 2 \sin \alpha_1 (\cos \alpha_2 \cos(\ell(\alpha_1 + \alpha_2)) + \sin \alpha_2 \sin(\ell(\alpha_1 + \alpha_2))) \frac{\sin(\ell(\alpha_1 + \alpha_2))}{\sin(\alpha_1 + \alpha_2)}. \end{aligned}$$

Finally, a tedious but straightforward trigonometric calculation using the formulas

above shows that

$$A_1 B_1 + A_2 B_2 + A_3 B_3 = \sin^2 \alpha_2 \cos(2\ell(\alpha_1 + \alpha_2)),$$

which immediately implies that $P(\cos \varphi)$ limits to $\cos(2\ell(\alpha_1 + \alpha_2))$ as $\varphi \rightarrow 0$. The calculation of the limit as $\varphi \rightarrow \pi$ is similar. \square

Recall that the Chebyshev polynomial of the first kind $T_m(x)$ is the unique polynomial of degree m satisfying $T_m(\cos(\psi)) = \cos(m\psi)$ for $m = 0, 1, 2, \dots$. The following theorem is the main result of this section.

Theorem 5.4. *Let $T_{2\ell}(x)$ be the Chebyshev polynomial of the first kind of degree 2ℓ . Then the formula (12) is equivalent to*

$$(17) \quad \cos \theta = T_{2\ell}(\cos \alpha_1 \cos \alpha_2 - \cos \varphi \sin \alpha_1 \sin \alpha_2).$$

Proof. The solutions of the equation $P(\cos \varphi) = 1$ with $0 < \varphi < \pi$ are precisely the intersection points of $\widehat{\Lambda}_2^{\alpha,c}$ and $\widehat{\Gamma}_\beta^{\alpha,c}$ in the pillowcase $\widehat{H}_2^{\alpha,c}$. According to (4), these correspond to the conjugacy classes of irreducible representations $\rho: \pi_1(S^3 \setminus L_\ell) \rightarrow \text{SU}(2)$, which we described as the values of φ solving (7). The function $P(\cos \varphi)$ achieves its absolute maximum at all such φ ; hence we also know that $P'(\cos \varphi) = 0$.

Let us first assume that the values of (α_1, α_2) are in a sufficiently small neighborhood of $(\pi/2, \pi/2)$ chosen so that (7) has the maximal possible number of solutions, which is $\ell - 1$. Consider the polynomial

$$R(x) = P(x) - T_{2\ell}(\cos \alpha_1 \cos \alpha_2 - x \sin \alpha_1 \sin \alpha_2).$$

It has $\ell - 1$ roots $x = \cos(\pi m/\ell)$, $m = 1, \dots, \ell - 1$, each of multiplicity at least two. In addition, it has roots $x = \pm 1$ by Lemma 5.3. On the other hand, the degree of $R(x)$ is at most 2ℓ by Lemma 5.2, while the leading coefficients of $P(x)$ and $T_{2\ell}(\cos \alpha_1 \cos \alpha_2 - x \sin \alpha_1 \sin \alpha_2)$ match by Lemma 5.2. Therefore, the degree of $R(x)$ is at most $2\ell - 1$ so $R(x)$ must vanish.

To conclude, we observe that both $P(x)$ and $T_{2\ell}(\cos \alpha_1 \cos \alpha_2 - x \sin \alpha_1 \sin \alpha_2)$ are analytic functions in (α_1, α_2) . Since we proved that they equal each other in an open neighborhood of $(\pi/2, \pi/2)$, they must equal each other for all (α_1, α_2) . \square

5.4. Transversality. In general, $\widehat{\Gamma}_\beta^{\alpha,c}$ needs to be perturbed for the intersection number (6) to make sense. However, in the case of $(2, 2\ell)$ -torus links, no perturbation is necessary as the intersection is automatically transversal.

Proposition 5.5. *Let $\beta = (\sigma_1)^{2\ell}$ be the braid whose closure is the $(2, 2\ell)$ -torus link with $\ell \neq 0$. Then the intersection of $\widehat{\Gamma}_\beta^{\alpha,c}$ and $\widehat{\Lambda}_2^{\alpha,c}$ in $\widehat{H}_2^{\alpha,c}$ is transversal, and all the intersections points contribute with the same sign into the algebraic count.*




Figure 5. An example of $\widehat{\Gamma}_\beta^{\alpha,c}$.

Proof. We will assume that $\ell > 0$ since the case of $\ell < 0$ can be handled similarly. The curves $\widehat{\Gamma}_\beta^{\alpha,c}$ and $\widehat{\Lambda}_2^{\alpha,c}$ intersect exactly at the points $(\varphi, 0)$, where φ solves the equation

$$\cos(0) = 1 = T_{2\ell}(\cos \alpha_1 \cos \alpha_2 - \cos \varphi \sin \alpha_1 \sin \alpha_2),$$

which we have shown happens exactly when

$$\cos \alpha_1 \cos \alpha_2 - \cos \varphi \sin \alpha_1 \sin \alpha_2 = \cos \frac{m\pi}{\ell} \quad \text{for } m = 1, 2, \dots, \ell - 1.$$

The curve $\widehat{\Gamma}_\beta^{\alpha,c}$ is smooth near each of the intersection points. It is parameterized by (φ, θ) , where θ is a smooth function of φ found by solving (11). The intersection points give local maxima of the Chebyshev polynomial $T_{2\ell}$ hence the curve $\widehat{\Gamma}_\beta^{\alpha,c}$ must be decreasing near these points as shown in Figure 5. Therefore, all the intersection numbers will be of the same sign once we prove that the intersections are transversal.

To prove transversality, it suffices to show that the derivative of θ with respect to φ is not zero near the intersection points. Differentiating the formula (17) and keeping in mind that $T'_n(x) = nU_{n-1}(x)$, where $U_{n-1}(x)$ is the Chebyshev polynomial of the second kind of degree $n - 1$, we obtain

$$\frac{d\theta}{d\varphi} = -2\ell \sin \alpha_1 \sin \alpha_2 \sin \varphi \frac{U_{2\ell-1}(x)}{\sqrt{1 - T_{2\ell}^2(x)}}.$$

We wish to show that, as $x \rightarrow \cos(\pi m/\ell)$ from the left, the limit of the right-hand side of this equation is nonzero. This is true because of the well known fact that each $\cos(\pi m/\ell)$ is a simple root of $U_{2\ell-1}(x)$ and a double root of $1 - T_{2\ell}^2(x)$. \square

5.5. Fixing the overall sign. An immediate corollary of Propositions 4.1 and 5.5 is that, up to an overall sign, the Bénard–Conway invariant $h_{L_\ell}(\alpha)$ of the $(2, 2\ell)$ -torus link equals the number of integers $m \in \{1, \dots, |\ell| - 1\}$ such that $\cos(\pi m/\ell)$ lies in the interval $(\cos(\alpha_1 + \alpha_2), \cos(\alpha_1 - \alpha_2))$.

To determine the sign, we need to get specific about the orientations of $\widehat{\Gamma}_\beta^{\alpha,c}$, $\widehat{\Lambda}_2^{\alpha,c}$ and $\widehat{H}_2^{\alpha,c}$ and to compute at least one intersection number explicitly. We will only compute the intersection number in the special case of $\alpha = (\alpha_1, \alpha_2) = (\pi/2, \pi/2)$. The calculation for general $\ell > 0$ and α will be similar, and the case of negative ℓ will be addressed later in this section. The following argument is a modification of the argument of Boden and Herald [3].

In the case at hand, $\widehat{\Gamma}_\beta^{\alpha,c}$ intersects $\widehat{\Lambda}_2^{\alpha,c}$ at one point $(\varphi, \theta) = (\pi/2, 0)$, which corresponds to the conjugacy class of $(\mathbf{j}, \mathbf{i}, \mathbf{j}, \mathbf{i})$ in $\widehat{H}_2^{\alpha,c}$. Consider the function $f : R_2^{\alpha,c} \times R_2^{\alpha,c} \rightarrow \text{SU}(2)$ given by $f(X_1, X_2, Y_1, Y_2) = X_1 X_2 Y_2^{-1} Y_1^{-1}$ and the function $g : (0, \pi) \times [0, 2\pi) \rightarrow R_2^{\alpha,c} \times R_2^{\alpha,c}$ given by

$$g(\varphi, \theta) = (\mathbf{i} e^{-k\varphi}, \mathbf{i}, \mathbf{i} e^{-k(\varphi-\theta)}, \mathbf{i} e^{k\theta}).$$

The latter is exactly the quadruple $(e^{\alpha_1 P_1}, e^{\alpha_2 \mathbf{i}}, e^{\alpha_1 Q_1}, e^{\alpha_2 Q_2})$ used to parameterize the pillowcase $\widehat{H}_2^{\alpha,c}$, with $\alpha_1 = \alpha_2 = \pi/2$ substituted in the equation. Notice that $\widehat{H}_2^{\alpha,c}$ is the quotient of $f^{-1}(1)$ by conjugation. We will orient $f^{-1}(1)$ and $\widehat{H}_2^{\alpha,c}$ by applying the base-fiber rule.

First, we will consider the map g . Two tangent vectors that span the tangent space to $\widehat{H}_2^{\alpha,c}$ at $(\mathbf{j}, \mathbf{i}, \mathbf{j}, \mathbf{i})$ are given by

$$u_1 = \left. \frac{\partial g}{\partial \varphi} \right|_{(\varphi, \theta) = (\pi/2, 0)} = (-\mathbf{i}, 0, -\mathbf{i}, 0), \quad u_2 = \left. \frac{\partial g}{\partial \theta} \right|_{(\varphi, \theta) = (\pi/2, 0)} = (0, 0, \mathbf{i}, -\mathbf{j}).$$

The tangent space to the orbit through $(\mathbf{j}, \mathbf{i}, \mathbf{j}, \mathbf{i})$ is spanned by the vectors

$$\begin{aligned} v_1 &= \left. \frac{\partial}{\partial t} e^{it} (\mathbf{j}, \mathbf{i}, \mathbf{j}, \mathbf{i}) e^{-it} \right|_{t=0} = (2\mathbf{k}, 0, 2\mathbf{k}, 0), \\ v_2 &= \left. \frac{\partial}{\partial t} e^{jt} (\mathbf{j}, \mathbf{i}, \mathbf{j}, \mathbf{i}) e^{-jt} \right|_{t=0} = (0, -2\mathbf{k}, 0, -2\mathbf{k}), \\ v_3 &= \left. \frac{\partial}{\partial t} e^{kt} (\mathbf{j}, \mathbf{i}, \mathbf{j}, \mathbf{i}) e^{-kt} \right|_{t=0} = (-2\mathbf{i}, 2\mathbf{j}, -2\mathbf{i}, 2\mathbf{j}). \end{aligned}$$

The vectors $\{u_1, u_2, v_1, v_2, v_3\}$ form a basis of $T_{(\mathbf{j}, \mathbf{i}, \mathbf{j}, \mathbf{i})} f^{-1}(1)$. Complete it to a basis for $T_{(\mathbf{j}, \mathbf{i}, \mathbf{j}, \mathbf{i})}(R_2^{\alpha,c} \times R_2^{\alpha,c})$ using the vectors $\{w_1, w_2, w_3\}$, where we can choose $w_1 = (\mathbf{k}, 0, 0, 0)$, $w_2 = (0, \mathbf{k}, 0, 0)$, and $w_3 = (0, \mathbf{j}, 0, 0)$. Notice that the orientation of the ordered triple $\{df(w_1), df(w_2), df(w_3)\} = \{\mathbf{i}, -\mathbf{j}, -\mathbf{k}\}$ is consistent with that of the standard basis of Lie algebra $\mathfrak{su}(2)$.

The two oriented bases $\{w_1, w_2, w_3, u_1, u_2, v_1, v_2, v_3\}$ and

$$\{(\mathbf{i}, 0, 0, 0), (\mathbf{k}, 0, 0, 0), (0, \mathbf{j}, 0, 0), (0, -\mathbf{k}, 0, 0), (0, 0, \mathbf{i}, 0), (0, 0, \mathbf{k}, 0), (0, 0, 0, \mathbf{j}), (0, 0, 0, -\mathbf{k})\}$$

of the space $T_{(j,i,j,i)}(R_2^{\alpha,c} \times R_2^{\alpha,c})$ are related by the matrix

$$\begin{pmatrix} 0 & 0 & 0 & -1 & 0 & 0 & 0 & -2 \\ 1 & 0 & 0 & 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 2 \\ 0 & -1 & 0 & 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & -1 & 1 & 0 & 0 & -2 \\ 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 \end{pmatrix}$$

of determinant -8 . This implies that the ordered pair $\{u_2, u_1\}$ gives a positively oriented basis in $T_{(j,i,j,i)}\widehat{H}_2^{\alpha,c}$.

Next, consider the one-variable parameterizations $\widehat{\Lambda}_2^{\alpha,c} = \{(ie^{-k\varphi}, i, ie^{-k\varphi}, i)\}$ and $\widehat{\Gamma}_\beta^{\alpha,c} = \{(ie^{-k\varphi}, i, (ie^{-k\varphi}, i))\sigma_1^4\}$, where $(X_1, X_2)\sigma_1 = (X_1X_2X_1^{-1}, X_1)$ is the homeomorphism induced by the braid. Compute the tangent vector

$$\psi_1 = \frac{\partial}{\partial\varphi}(ie^{-k\varphi}, i, ie^{-k\varphi}, i)\Big|_{\varphi=\frac{\pi}{2}} = (-i, 0, -i, 0)$$

to $\widehat{\Lambda}_2^{\alpha,c}$ and the tangent vector

$$\psi_2 = \frac{\partial}{\partial\varphi}(ie^{-k\varphi}, i, \sigma_1^4(ie^{-k\varphi}, i))\Big|_{\varphi=\frac{\pi}{2}} = (-i, 0, -5i, 4j)$$

to $\widehat{\Gamma}_\beta^{\alpha,c}$. The bases $\{\psi_1, \psi_2\}$ and $\{u_2, u_1\}$ in the tangent space $T_{(j,i,j,i)}\widehat{H}_2^{\alpha,c}$ are related by the matrix

$$\begin{pmatrix} 0 & -4 \\ 1 & 1 \end{pmatrix}.$$

Since the determinant of this matrix is positive, we conclude that the intersection number $h_L(\alpha) = \langle \widehat{\Lambda}_2^{\alpha,c}, \widehat{\Gamma}_\beta^{\alpha,c} \rangle_{\widehat{H}_2^{\alpha,c}}$ equals $+1$.

Corollary 5.6. *For any integer $\ell > 0$ and for any choice of (α_1, α_2) away from the set (8), the B enard–Conway invariant $h_{L_\ell}(\alpha)$ of the $(2, 2\ell)$ -torus link is well defined and is equal to the number of integers $m \in \{1, \dots, |\ell| - 1\}$ such that*

$$\cos(\pi m/\ell) \in (\cos(\alpha_1 + \alpha_2), \cos(\alpha_1 - \alpha_2)).$$

For $\ell = 1$, we obviously obtain $h_{L_1}(\alpha) = 0$, as established by B enard and Conway [1]. For the first few $\ell \geq 1$, the invariant $h_{L_\ell}(\alpha)$ equals m in each of the respective diamond-shaped regions R_m depicted in Figure 6. The boundaries of R_m consist of the values of (α_1, α_2) from the set (8), where the multivariable Alexander polynomial of L_ℓ vanishes and hence the B enard–Conway invariant is not defined.




Figure 6. Regions R_m for small ℓ .

If ℓ is negative, we perform a similar calculation for $\ell = -2$ and $\alpha = (\alpha_1, \alpha_2) = (\pi/2, \pi/2)$ to determine the overall sign. The only difference is the parametrization of $\widehat{\Gamma}_\beta^{\alpha,c}$, which is now $\{(ie^{-\kappa\varphi}, i, (ie^{-\kappa\varphi}, i))\sigma_1^{-4}\}$, reflecting the change in braid representative from σ_1^4 to σ_1^{-4} . The new tangent vector ψ_2 is $(-i, 0, 3i, -4j)$ and the resulting change-of-basis matrix has negative determinant.

Corollary 5.7. *For any integer $\ell > 0$ and for any choice of (α_1, α_2) away from the set (8) the invariants $h_{L_\ell}(\alpha)$ and $h_{L_{-\ell}}(\alpha)$ are well defined and related by*

$$h_{L_{-\ell}}(\alpha) = -h_{L_\ell}(\alpha).$$

6. Proof of Theorems 1.1 and 1.2

According to [14], the linking number is a complete invariant of link homotopy for two-component links. Therefore, any two-component oriented link with linking number $\ell \neq 0$ can be obtained from the $(2, 2\ell)$ -torus link via a sequence of crossing changes within individual components. The Cimasoni–Florens signature for $(2, 2\ell)$ -torus links can be easily computed directly from its definition. Assuming that ℓ is positive, it is given by the formulas

$$\sigma_{L_\ell}(\omega_1, \omega_2) = \begin{cases} \ell - 2i - 1 & \text{if } i\pi/\ell < \alpha_1 + \alpha_2 < (i+1)\pi/\ell \text{ and } 0 \leq i \leq \ell - 1, \\ -3\ell + 2i + 1 & \text{if } i\pi/\ell < \alpha_1 + \alpha_2 < (i+1)\pi/\ell \text{ and } \ell \leq i \leq 2\ell - 1, \end{cases}$$

and $\sigma_{L_{-\ell}}(\omega_1, \omega_2) = -\sigma_{L_\ell}(\omega_1, \omega_2)$; see [6, Example 2.5]. From $\sigma_{L_\ell}(\omega_1, \omega_2)$ we obtain $\sigma_{L_\ell}(\omega_1, \omega_2^{-1})$ by a flip across the axis $\alpha_2 = \pi/2$, and the average of the two quantities is then easily calculated. Comparing the answer with the calculation of Section 4 and Corollary 5.6, we conclude that

$$h_{L_\ell}(\alpha) = h_{L_\ell}(\alpha_1, \alpha_2) = -\frac{1}{2}(\sigma_{L_\ell}(\omega_1, \omega_2) + \sigma_{L_\ell}(\omega_1, \omega_2^{-1})).$$

According to Theorem 3.1, the equality (1) remains true with each crossing change. Theorem 1.1 now follows.

In order to prove Theorem 1.2, we will use a formula of Cimasoni and Florens,

$$\sigma_{L_1 \cup L_2}(\omega_1, \omega_2^{-1}) = \sigma_{L_1 \cup -L_2}(\omega_1, \omega_2),$$

where $-L_2$ stands for the component L_2 with reversed orientation; see Proposition 2.8 in [5]. Together with the formula

$$\sigma_{L_1 \cup L_2}(\omega_1^{-1}, \omega_2^{-1}) = \sigma_{L_1 \cup L_2}(\omega_1, \omega_2),$$

which easily follows from the definition of the Cimasoni–Florens signature, this implies that

$$h_L(\alpha_1, \alpha_2) = -\frac{1}{2}(\sigma_{L_1 \cup L_2}(\omega_1, \omega_2) + \sigma_{L_1 \cup -L_2}(\omega_1, \omega_2))$$

is independent of the choice of orientation on the link L . If $\omega_1 = \omega_2 = \omega$, one can use the formula

$$\sigma_{L_1 \cup L_2}(\omega) = \sigma_{L_1 \cup L_2}(\omega, \omega) - \text{lk}(L_1, L_2)$$

from [5, Proposition 2.5], relating the multivariable signature $\sigma_L(\omega, \omega)$ with the Levine–Tristram signature $\sigma_L(\omega)$, to obtain

$$\begin{aligned} h_L(\alpha, \alpha) &= -\frac{1}{2}(\sigma_{L_1 \cup L_2}(\omega, \omega) + \sigma_{L_1 \cup -L_2}(\omega, \omega)) \\ &= -\frac{1}{2}(\sigma_{L_1 \cup L_2}(\omega) + \text{lk}(L_1, L_2) + \sigma_{L_1 \cup -L_2}(\omega) + \text{lk}(L_1, -L_2)) \\ &= -\frac{1}{2}(\sigma_{L_1 \cup L_2}(\omega) + \sigma_{L_1 \cup -L_2}(\omega)). \end{aligned}$$

In the special case of $\omega = -1$, this implies that $h_L(\pi/2, \pi/2)$ equals minus the Murasugi signature of L .

Remark 6.1. There does not appear to be a name in the literature for the quantity $\frac{1}{2}(\sigma_{L_1 \cup L_2}(\omega) + \sigma_{L_1 \cup -L_2}(\omega))$, nor for an analogous quantity for links with more than two components, when $\omega \neq -1$. Perhaps it should be called the equivariant Murasugi signature.

Acknowledgments

We thank Hans Boden, Anthony Conway, Daniel Ruberman, and Chris Scaduto for useful discussions and sharing their expertise.

References

- [1] L. Benard and A. Conway, “A multivariable Casson–Lin type invariant”, *Ann. Inst. Fourier (Grenoble)* **70**:3 (2020), 1029–1084. MR
- [2] H. U. Boden and E. Harper, “The $SU(N)$ Casson–Lin invariants for links”, *Pacific J. Math.* **285**:2 (2016), 257–282. MR
- [3] H. U. Boden and C. M. Herald, “The $SU(2)$ Casson–Lin invariant of the Hopf link”, *Pacific J. Math.* **285**:2 (2016), 283–288. MR
- [4] D. Cimasoni, “A geometric construction of the Conway potential function”, *Comment. Math. Helv.* **79**:1 (2004), 124–146. MR

- [5] D. Cimasoni and V. Florens, “Generalized Seifert surfaces and signatures of colored links”, *Trans. Amer. Math. Soc.* **360**:3 (2008), 1223–1264. MR
- [6] D. Cimasoni, M. Markiewicz, and W. Politarczyk, “Torres-type formulas for link signatures”, preprint, 2023. arXiv 2304.02347
- [7] A. Daemi and C. Scaduto, “Equivariant aspects of singular instanton Floer homology”, *Geom. Topol.* **28**:9 (2024), 4057–4190. MR
- [8] A. Daemi and C. Scaduto, “Unoriented skein exact triangles in equivariant singular instanton Floer theory”, preprint, 2024. arXiv 2409.16390
- [9] E. Harper and N. Saveliev, “A Casson–Lin type invariant for links”, *Pacific J. Math.* **248**:1 (2010), 139–154. MR
- [10] C. M. Herald, “Flat connections, the Alexander invariant, and Casson’s invariant”, *Comm. Anal. Geom.* **5**:1 (1997), 93–120. MR
- [11] M. Heusener and J. Kroll, “Deforming abelian $SU(2)$ -representations of knot groups”, *Comment. Math. Helv.* **73**:3 (1998), 480–498. MR
- [12] P. B. Kronheimer and T. S. Mrowka, “Khovanov homology is an unknot-detector”, *Publ. Math. Inst. Hautes Études Sci.* **113** (2011), 97–208. MR
- [13] X.-S. Lin, “A knot invariant via representation spaces”, *J. Differential Geom.* **35**:2 (1992), 337–357. MR
- [14] J. Milnor, “Link groups”, *Ann. of Math. (2)* **59** (1954), 177–195. MR
- [15] J. Milnor, *Singular points of complex hypersurfaces*, Annals of Mathematics Studies **61**, Princeton University Press, 1968. MR
- [16] J. Murakami, “A state model for the multivariable Alexander polynomial”, *Pacific J. Math.* **157**:1 (1993), 109–135. MR
- [17] K. Murasugi, “On the signature of links”, *Topology* **9** (1970), 283–298. MR
- [18] G. Torres, “On the Alexander polynomial”, *Ann. of Math. (2)* **57** (1953), 57–89. MR

Received December 12, 2024. Revised December 2, 2025.

ZEDAN LIU
DEPARTMENT OF MATHEMATICS
UNIVERSITY OF MIAMI
CORAL GABLES, FL
UNITED STATES
zedan.liu@miami.edu

NIKOLAI SAVELIEV
DEPARTMENT OF MATHEMATICS
UNIVERSITY OF MIAMI
CORAL GABLES, FL
UNITED STATES
saveliev@math.miami.edu

Guidelines for Authors

Authors may submit articles at msp.org/pjm/about/journal/submissions.html and choose an editor at that time. Exceptionally, a paper may be submitted in hard copy to one of the editors; authors should keep a copy.

By submitting a manuscript you assert that it is original and is not under consideration for publication elsewhere. Instructions on manuscript preparation are provided below. For further information, visit the web address above or write to journals+pjm@msp.org or to Pacific Journal of Mathematics, University of California, Los Angeles, CA 90095–1555. Correspondence by email is requested for convenience and speed.

Manuscripts must be in English, French or German. A brief abstract of about 150 words or less in English must be included. The abstract should be self-contained and not make any reference to the bibliography. Also required are keywords and subject classification for the article, and, for each author, postal address, affiliation (if appropriate) and email address if available. A home-page URL is optional.

At submission time a PDF file is required, and it must be derived from a \LaTeX source file in which the command `\documentclass[12pt]{amsart}` is active. No changes to horizontal or vertical margins, or line spacing, should be made. Please include in the preamble the command `\usepackage[pagewise]{lineno}\linenumbers` — this will spare referees tedious effort in line counting.

Carefully preserve all related files, such as \LaTeX sources and individual files for each figure; you will be asked to submit them upon acceptance of the paper.

Bibliographical references should be listed alphabetically at the end of the paper. All references in the bibliography should be cited in the text. Use of \BibTeX is preferred but not required. Any bibliographical citation style may be used but tags will be converted to the house format (see a current issue for examples).

Figures, whether prepared electronically or hand-drawn, must be of publication quality. Figures prepared electronically should be submitted in Encapsulated PostScript (EPS) or in a form that can be converted to EPS, such as GnuPlot, Maple or Mathematica. Many drawing tools such as Adobe Illustrator and Aldus FreeHand can produce EPS output. Figures containing bitmaps should be generated at the highest possible resolution. If there is doubt whether a particular figure is in an acceptable format, the authors should check with production by sending an email to journals+pjm@msp.org.

Each figure should be captioned and numbered, so that it can float. Small figures occupying no more than three lines of vertical space can be kept in the text (“the curve looks like this:”). It is acceptable to submit a manuscript with all figures at the end, if their placement is specified in the text by means of comments such as “Place Figure 1 here”. The same considerations apply to tables, which should be used sparingly.

Forced line breaks or page breaks should not be inserted in the document. There is no point in your trying to optimize line and page breaks in the original manuscript. The manuscript will be reformatted to use the journal’s preferred fonts and layout.

Page proofs will be made available to authors (or to the designated corresponding author) at a website in PDF format. Failure to acknowledge the receipt of proofs or to return corrections within the requested deadline may cause publication to be postponed.

PACIFIC JOURNAL OF MATHEMATICS

Volume 341 No. 2 April 2026

Chainmail links and L -spaces	199
IAN AGOL	
Mekler's construction and Murphy's law for 2-nilpotent groups	219
BLAISE BOISSONNEAU, ARIS PAPADOPOULOS and PIERRE TOUCHARD	
Axial symmetry in convex bodies	275
RITESH GOENKA, KENNETH MOORE, WEN RUI SUN and ETHAN PATRICK WHITE	
A combinatorial structure for many hierarchically hyperbolic spaces	305
MARK HAGEN, GIORGIO MANGIONI and ALESSANDRO SISTO	
The B�enard–Conway invariant of two-component links	379
ZEDAN LIU and NIKOLAI SVELIEV	