

SHORT HIERARCHICALLY HYPERBOLIC GROUPS II: QUOTIENTS AND THE HOPF PROPERTY FOR ARTIN GROUPS

GIORGIO MANGIONI AND ALESSANDRO SISTO

ABSTRACT. We prove that most Artin groups of large and hyperbolic type are Hopfian, meaning that every self-epimorphism is an isomorphism. The class covered by our result is generic, in the sense of Goldsborough-Vaskou. Moreover, assuming the residual finiteness of certain hyperbolic groups with an explicit presentation, we get that all large and hyperbolic type Artin groups are residually finite. We also show that “most” quotients of the five-holed sphere mapping class group are hierarchically hyperbolic, up to taking powers of the normal generators of the kernels.

The main tool we use to prove both results is a Dehn-filling-like procedure for short hierarchically hyperbolic groups (these also include e.g. non-geometric 3-manifolds, and triangle- and square-free RAAGs).

CONTENTS

Introduction	1
1. Combinatorial HHS	6
2. Definition and results on short HHG	9
3. Rotating families and projection graphs, revisited	13
4. A short structure for the quotient	27
5. Hopf property from central quotients	39
6. Hopf property for Artin groups	46
7. Quotients of the five-holed sphere mapping class group	57
References	59

INTRODUCTION

By work of Osin [Osi07], and independently Groves and Manning [GM08], if a group G is hyperbolic relative to a subgroup P , and $N \trianglelefteq P$ is a “sufficiently deep” normal subgroup, then $G/\langle\langle N \rangle\rangle$ is hyperbolic relative to P/N . This quotient procedure is called *Dehn filling*, as it abstracts Thurston’s Dehn filling procedure for cusped hyperbolic 3-manifolds [Thu97]. In this paper, we develop an analogue Dehn filling technique for *short hierarchically hyperbolic groups*, first introduced in [Man24]. This family includes the mapping class group of the five-holed sphere, extra-large Artin groups, fundamental groups of non-geometric 3-manifolds, right-angled Artin groups without triangles and squares in their defining graphs, and many others. Our tools allow us to construct hyperbolic quotients of these groups; in turn this has

several applications, the main one being related to the Hopf property for Artin groups, as we now explain.

Hopf property for Artin groups. The motivation of our Dehn filling machinery is to “lift” properties from the quotients to the groups themselves. In fact, we study when a group G within our class is *Hopfian*, which means that every surjective homomorphism $G \rightarrow G$ is an isomorphism. This property is very useful to produce automorphisms of G , because, while it is usually easy to check that a map is surjective (it is enough that a generating set is contained in the image), proving injectivity can be quite fiddly.

Within our class, we prove the Hopf property for “most” Artin groups of *large and hyperbolic type*, by which we mean that all edge labels in the defining graph are at least 3 and there is no triangle whose labels are all 3. The exact statement is Theorem 6.6, from which we extract some special cases. First, recall that an Artin group is *even* if all edge labels in its defining graph are even.

Theorem A. *Let A_Γ be an even large-type Artin group. Then A_Γ is Hopfian.*

For the second application, we say that an Artin group has a *single odd component* if every two vertices in the defining graph are connected by a combinatorial path with odd labels. This is a natural notion, as two generators of an Artin group are conjugate if and only if the corresponding vertices are connected by an odd path [Par97].

Theorem B. *Let A_Γ be an Artin group of large and hyperbolic type with a single odd component. Then A_Γ is Hopfian.*

As explained in Remark 6.16, the aforementioned class of Artin groups is generic in the sense of [GV25], so we obtain:

Corollary C. *A generic Artin group is Hopfian.*

The Corollary was also very recently obtained in [BMV24], where, by completely different means, the authors identify two other classes which turn out to be generic (see Remark 6.17 for a comparison between the families).

In [Bar24], Barak proved that all HHG satisfying a technical condition are equationally Noetherian, hence Hopfian (see e.g. [GH19, Corollary 3.14 and Theorem D] for the implication). The further requirement however, is not satisfied by our groups, as discussed in Remark 6.20.

Towards residual finiteness. It is a widespread belief that all Artin groups are residually finite [Gol, Problem 8.05]. Residual finiteness, which would imply the Hopf property, has only been established for certain classes of Artin groups; in fact, to the best of our knowledge, it is not known whether generic Artin groups are residually finite (see Remark 6.19 for an overview). In Figure 11 we even exhibit a four-generated Artin group which is Hopfian by Theorem B but is not known to be residually finite.

However, our Theorem 6.6 is proven by constructing “sufficiently many” hyperbolic quotients, of which an explicit presentation is given in Remark 6.18. All hyperbolic groups are Hopfian by e.g. [WR19, Corollary 6.13], and if our quotients were in fact residually finite, then our Artin groups would be as well. We summarise this in the following:

Theorem D. *Let A_Γ be an Artin group of large and hyperbolic type. Then A_Γ is residually hyperbolic. Furthermore, if all the hyperbolic groups from Remark 6.18 are residually finite, then A_Γ is residually finite.*

In the same Remark 6.18, we relate residual finiteness of A_Γ to residual finiteness of certain *Shephard groups*. This class, named after G. Shephard [She52], generalises both Coxeter and Artin groups, as one is allowed to require that some of the generators in an Artin presentation have finite order.

Remark E. In [Gol24] Goldman proved that any large and hyperbolic type Artin group A_Γ is residually hyperbolic. Our work and Goldman's are independent, use different tools, and were carried out simultaneously. However, we both obtain the following intermediate result: for a suitable choice of a positive integer N , the Shephard group obtained by killing the N -th power of all standard generators of A_Γ is hyperbolic relative to the images of the dihedral subgroups. As a consequence, this group admits hyperbolic “Dehn filling” quotients, either by our Theorem H below or by the “standard” relatively hyperbolic Dehn filling Theorem. Furthermore, under the further assumption that Γ is triangle-free, Goldman was able to obtain residual finiteness [Gol24, Corollary G].

Quotients of $\mathcal{MCG}^\pm(S_{0,5})$. Only a handful of mapping class groups of finite-type surfaces are short HHG, including that of the five-holed sphere, which we focus on here. There are various ways to take quotients of mapping class groups that yield hierarchically hyperbolic groups, including quotients by powers of pseudo-Anosovs [BHS17a] and quotients by powers of Dehn twists [BHMS24]. These quotients have been further studied by exploiting hierarchical hyperbolicity to obtain quasi-isometric and algebraic rigidity results [MS25, Man23]. They have also notably been used to relate famous questions about profinite properties of mapping class groups and profinite rigidity of certain 3-manifolds to residual finiteness of certain hyperbolic groups [BHMS24, WS24].

In this context, we previously asked, roughly, whether given any finite collection of elements of a mapping class group one can quotient by suitable powers and obtain a hierarchically hyperbolic group, see [MS25, Question 3]. We provide an almost complete answer for $\mathcal{MCG}^\pm(S_{0,5})$:

Theorem F (see Theorem 7.1). *Let $S = S_{0,5}$, and let $g_1, \dots, g_l \in \mathcal{MCG}^\pm(S)$. Suppose that every partial pseudo-Anosov g_i has no hidden symmetries. Then there exists $N \in \mathbb{N} - \{0\}$ such that, for all $K_1, \dots, K_l \in \mathbb{Z} - \{0\}$ we have that $\mathcal{MCG}^\pm(S)/\langle\langle\{g_i^{K_i N}\}\rangle\rangle$ is hierarchically hyperbolic.*

Here, not having hidden symmetries is a technical condition which only partial pseudo-Anosovs can satisfy, see Definition 3.44. Unfortunately we needed this additional requirement due to fine algebraic reasons, but we do not think that it is necessary, and in fact we believe that there is now enough evidence to upgrade our question to a conjecture:

Conjecture G. *Let S be any surface of finite type, and let $g_1, \dots, g_l \in \mathcal{MCG}^\pm(S)$. There exists $N \in \mathbb{N} - \{0\}$ such that, for all $K_1, \dots, K_l \in \mathbb{Z} - \{0\}$, the group $\mathcal{MCG}^\pm(S)/\langle\langle\{g_i^{K_i N}\}\rangle\rangle$ is hierarchically hyperbolic.*

To the best of our knowledge, the conjecture is open even for quotients by suitable powers of non-separating Dehn Twists. What makes the conjecture interesting is

that tackling it should lead to a more complete theory of Dehn fillings for hierarchically hyperbolic groups, and in turn this should have many applications beyond those presented in this paper and [BHMS24, WS24].

Dehn fillings of short HHG.

Main result. Roughly, a short HHG contains specified subgroups which are \mathbb{Z} -central extensions of hyperbolic groups. We call *cyclic directions* the kernels of these extensions (see Subsection 2.1 for the full definition of a short HHG). For the five-holed sphere mapping class group, the specified extensions are curve stabilisers, and cyclic directions are generated by Dehn twists.

Short HHG are examples of *hierarchically hyperbolic groups*, first introduced in [BHS17b]. The only consequence of hierarchical hyperbolicity that the reader should bear in mind for this Introduction is that a short HHG is hyperbolic if the list of specified subgroups is empty. Hence, in order to make a short HHG “more hyperbolic”, the idea is to take a generator of a cyclic direction and quotient by a power of it, as such an element has “large” centraliser and in a hyperbolic group this is allowed only if the element has finite order. With this in mind, we now present the main technical result of this paper, which should be thought of as an analogue of the relatively hyperbolic Dehn filling theorem:

Theorem H (see Theorem 4.1). *Let G be a short HHG, and let g_1, \dots, g_n be generators of some of its cyclic directions. Then there exists $M \in \mathbb{N} - \{0\}$ such that, for all choices $k_i \in \mathbb{Z} - \{0\}$, the quotient $\overline{G} = G / \langle\langle g_i^{k_i M} \rangle\rangle$ is a short HHG.*

In fact, the Theorem also gives a natural short HHG structure on \overline{G} , where the cyclic directions are images of the cyclic directions of G that do not contain conjugates of the g_i . Thus, a particularly interesting case is where we take quotients by powers of generators of all cyclic directions, hence obtaining a hyperbolic group. A refinement of the above yields the following, which readily implies Theorem D. Recall that a group G is *fully residually P* for some property P if, for every finite subset $F \subset G$, there exists a quotient $G \rightarrow \overline{G}$ where F injects, and such that \overline{G} enjoys P .

Corollary I (see Corollary 4.23). *Short HHG are fully residually hyperbolic.*

Tools and techniques. The proof of Theorem H combines two approaches. Firstly, as explored in [Man24], the hierarchical structure of a short HHG can be modified by constructing suitable quasimorphisms on the specified \mathbb{Z} -central extensions (see Subsection 2.3 for further details). In Subsection 4.1 we take advantage of this flexibility to make the structure of a short HHG “as compatible as possible” with the quotient projection. Secondly, we adapt the machinery of *rotating families*, first introduced in [Dah18], to short HHG. Mimicking arguments from [DHS21] and [BHMS24], these tools allow one to lift certain combinatorial configurations from the quotient \overline{G} to the original group G . Each HHG axiom for \overline{G} follows from the corresponding statement for G , which is already hierarchically hyperbolic. The most novel and difficult part of this construction, compared with [BHMS24] and other papers, is that the HHS structure of the quotients involves quasilines coming from the aforementioned quasimorphisms; therefore, new ideas are required to construct suitable retractions onto those, as detailed in Subsection 4.3.4.

Comparison with known Dehn Filling results. In [BHMS24], the authors study quotients of mapping class groups of arbitrary finite-type surfaces by suitably large powers of *all* Dehn twists, which as mentioned above should be thought of as the equivalent of what we call central directions. On the one hand, our Theorem H applies to quotients by suitably large powers of *any collection* of central directions. On the other, our techniques fail if we consider the quotient of a non-short mapping class group by large enough powers of some Dehn twists, such as only those around non-separating curves. The problem is roughly as follows: in any reasonable candidate HHS structure for the quotient, the image of a curve stabiliser will be a product region, and therefore hierarchically hyperbolic itself. Now, if a curve does not lie in the chosen collection, then the quotient image of its stabiliser will be a \mathbb{Z} -central extension, and by [HRSS25, Corollary 4.3] a necessary condition for it to be hierarchically hyperbolic is that its Euler class is *bounded*. In the case of the five-holed sphere, the quotient extension has hyperbolic base, and so is bounded by [NR97]; however this is not necessarily true for surfaces of higher complexity. Motivated by this setup, in future work we will explore under which conditions a quotient of a \mathbb{Z} -central extension of a group G remains bounded [FFMS25].

A criterion for hopficity. The fundamental tool in our study of the Hopf property for short HHG is the following criterion (stated here in slightly simplified form), whose proof is straightforward but for which we could not find a suitable reference:

Proposition J (see Lemma 5.3). *Suppose that G has enough Hopfian quotients, meaning that, for every surjective homomorphism $\phi : G \rightarrow G$ and $g_0 \in G - \{1\}$, there exists a quotient H of G , say with quotient map q such that:*

- $q(g_0) \neq 1$,
- H is Hopfian,
- ϕ induces a homomorphism $\psi : H \rightarrow H$.

Then G is Hopfian.

Now let G be a short HHG, and let ϕ, g_0 be as above. If ϕ maps central directions to central directions, then it induces a map of some hyperbolic (hence Hopfian [WR19, Corollary 6.13]) Dehn filling quotient H , obtained by annihilating suitable powers of all central directions, and by Corollary I we can also assume that g_0 survives in H . This is not always the case; however, ϕ often preserves “enough” central directions, and the Dehn filling quotient by those directions will still be “hyperbolic enough” to be Hopfian. More precisely, we shall make use of hopficity of certain relatively hyperbolic groups [GH19] and the following fact, which we highlight as it is of independent interest.

Theorem K (see Theorem 5.7). *Let G be hyperbolic relative to \mathbb{Z} -central extensions of hyperbolic groups (including the case that G itself is such an extension). Then G is Hopfian.*

Future directions. It is natural to ask whether, using a version of our argument or otherwise, one can in fact show that all short HHG are Hopfian:

Question L. Are all short HHG Hopfian? More ambitiously, are all HHG Hopfian?

But in fact, this is already a major question in the context of Artin groups:

Question M. Are all large hyperbolic type Artin groups Hopfian?

Going one step further, given the connection to residual finiteness of certain hyperbolic groups, the following question could have important consequences:

Question N. Are all large hyperbolic type Artin groups residually finite?

As explained in Theorem D, a positive answer to Question N would follow from residual finiteness of an explicit collection of hyperbolic groups, so we ask:

Question O. Are all hyperbolic groups from Remark 6.18 residually finite?

Finally, we believe that there are many more examples of short HHG, and in fact we believe that the following question has a positive answer:

Question P. Is a random quotient of a short HHG, for a suitable notion of randomness, again a short HHG?

Outline. Section 1 contains background on combinatorial hierarchically hyperbolic groups, of which short HHG are an instance. In Section 2, we describe the family of short HHG, and collect results about them from [Man24]. The class of Dehn filling quotients we consider is then made precise in Subsection 2.4.

In Section 3 we adapt the machinery of rotating families from [Dah18, DHS21] to our quotients of interest. We then construct a short HHG structure for the quotients in Section 4: see in particular Theorem 4.1, which is Theorem H from the Introduction. As a by-product, in Subsection 4.5 we prove residual hyperbolicity of short HHG: see Corollary 4.23, which is Corollary I above.

Section 5 develops tools to study self-epimorphisms of short HHG. An example of how they are put into practice is presented in Subsection 5.5, where we prove the Hopf property for certain HNN extensions of the direct product of \mathbb{Z} and a free group: see Example 5.19. The same techniques are then pushed further in Section 6, where we prove that many Artin groups of large and hyperbolic type are Hopfian: see Theorem 6.6 for the full statement, which encompasses both Theorems A and B.

In Section 7 we prove hierarchical hyperbolicity of many quotients of the five-holed sphere mapping class group: see Theorem 7.1, which is Theorem F.

Acknowledgements. We are grateful to Oli Jones for suggesting an enlightening way to shorten some proofs of the Hopf property. We also thank Giovanni Sartori and Nicolas Vaskou for answering many questions about Artin groups. Special thanks go to Martín Blufstein and Alexandre Martin for fruitful discussions. Finally, we thank the referee for useful comments, both on the mathematical content and on the exposition.

1. COMBINATORIAL HHS

In this section we recall the definition of a combinatorial HHS and its hierarchically hyperbolic structure, as first introduced in [BHMS24]. The reader might want to refer to [BHMS24, Section 1], which contains discussion of all the various notions we recall below. Also, the reader might find [HMS23] useful, as in there it is explained how to create a combinatorial HHS structure from an HHS structure (in many cases).

Definition 1.1 (Induced subgraph). Let X be a simplicial graph. 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 1.2 (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 *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 1.3 (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 $X^{+\mathcal{W}}$.

Definition 1.4 (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 non-maximal simplices in X .

Definition 1.5 (Complement, link subgraph). Let \mathcal{W} be an X -graph. For each simplex Δ of X , let $\mathcal{C}(\Delta)$ be the induced subgraph of $X^{+\mathcal{W}}$ spanned by $\text{Lk}(\Delta)^{(0)}$, which we call the *augmented link* of Δ . Also, let Y_Δ be the subgraph of $X^{+\mathcal{W}}$ induced by the set of vertices $X^{(0)} - \text{Sat}(\Delta)$.

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.)

Definition 1.6 (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) There exists $n \in \mathbb{N}$ such 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 ;
- (2) There is a constant δ so that for each non-maximal 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 1.5;
- (3) Whenever Δ and Σ are non-maximal simplices for which there exists a non-maximal 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 non-adjacent 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 $\Delta \star \Sigma$.

Definition 1.7 (Nesting, orthogonality, transversality). Let X be a simplicial graph. Let Δ, Δ' be non-maximal 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 $[\Delta] \nabla [\Delta']$.

Definition 1.8 (Projections). Let (X, \mathcal{W}) be a combinatorial HHS. Fix $[\Delta] \in \mathfrak{S}$ and define a map $\pi_{[\Delta]} : \mathcal{W} \rightarrow 2^{\mathcal{C}(\Delta)}$ as follows. Let

$$p : Y_{\Delta} \rightarrow 2^{\mathcal{C}(\Delta)}$$

be the coarse closest point projection, i.e.

$$p(x) = \{y \in \mathcal{C}(\Delta) : d_{Y_{\Delta}}(x, y) \leq d_{Y_{\Delta}}(x, \mathcal{C}(\Delta)) + 1\}.$$

Suppose that w is a vertex of \mathcal{W} , so w corresponds to a unique simplex Σ_w of X . Now, [BHMS24, Lemma 1.15] states that the intersection $\Sigma_w \cap Y_{\Delta}$ is non-empty and has diameter at most 1. Define

$$\pi_{[\Delta]}(w) = p(\Sigma_w \cap Y_{\Delta}).$$

We have thus defined $\pi_{[\Delta]} : \mathcal{W}^{(0)} \rightarrow 2^{\mathcal{C}(\Delta)}$. If $v, w \in \mathcal{W}$ 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 $[\Delta] \nabla [\Delta']$ or $[\Delta'] \sqsubset [\Delta]$. 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 $[\Delta] \sqsubset [\Delta']$. 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 1.9 (HHS structures for X -graphs). *Let (X, \mathcal{W}) be a combinatorial HHS. The above data defines a hierarchically hyperbolic space (HHS) structure on \mathcal{W} , in the sense of [BHS19].*

Moreover, 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 is a hierarchically hyperbolic group (HHG), in the sense of [BHS19].

Definition 1.10. We will say that a group G satisfying the assumptions of Theorem 1.9 is a *combinatorial HHG*. The hierarchically hyperbolic structure comes with a constant E , called the *HHS constant*, which we can assume to be larger than both n and δ .

2. DEFINITION AND RESULTS ON SHORT HHG

This Section recaps the definition and properties of short HHG, as introduced in [Man24].

Definition 2.1 (Blowup graph). Let \overline{X} be a simplicial graph, whose vertices are labelled by graphs $\{L_v\}_{v \in \overline{X}^{(0)}}$. The *blowup* of \overline{X} , with respect to the collection $\{L_v\}$, is the graph X obtained from \overline{X} by replacing every vertex v with the *cone* $\text{Cone}(v) = v * (L_v)^{(0)}$. Two cones $\text{Cone}(v)$ and $\text{Cone}(w)$ span a join in X if and only if v, w are adjacent in \overline{X} , and are disjoint otherwise. In particular, there is a Lipschitz retraction $p: X \rightarrow \overline{X}$ mapping every $\text{Cone}(v)$ to its tip v .




FIGURE 1. The blowup of two adjacent vertices of \overline{X} .

For every simplex $\Delta \subset X$, let $\overline{\Delta} = p(\Delta)$, which we call the *support* of Δ , and for every $v \in \overline{\Delta}$ let $\Delta_v = \Delta \cap \text{Cone}(v)$. When describing a simplex Δ , we shall put vertices belonging to the same Δ_v in parentheses: for example, if the vertices of Δ are $\{v, w, x\}$, where $v, w \in \overline{X}^{(0)}$ and $x \in (L_v)^{(0)}$, then we denote Δ by $\{(v, x), (w)\}$. Then, by inspection of the definition, one gets the following:

Lemma 2.2. *Suppose that \overline{X} is triangle-free, and that no component of \overline{X} is a single point. Then, given a simplex Δ of X , one of the following holds:*

- (1) $\Delta = \emptyset$, and $\text{Lk}_X(\Delta) = X$;
- (2) (*Edge-type*) $\Delta = \{(v, x)\}$, where $v \in \overline{X}^{(0)}$ and $x \in L_v$, and $\text{Lk}_X(\Delta) = p^{-1} \text{Lk}_{\overline{X}}(v)$;
- (3) (*Triangle-type*) $\Delta = \{(v, x), (w)\}$, where $v, w \in \overline{X}^{(0)}$ are \overline{X} -adjacent and $x \in L_v$, and $\text{Lk}_X(\Delta) = (L_w)^{(0)}$;
- (4) $\Delta = \{(v, x), (w, y)\}$ is a maximal simplex, where $v, w \in \overline{X}^{(0)}$ are adjacent, $x \in L_v$ and $y \in L_w$, and $\text{Lk}_X(\Delta) = \emptyset$.
- (5) $\text{Lk}_X(\Delta)$ is a vertex, or a non-trivial join. In particular $\text{diam}(\text{Lk}_X(\Delta)) \leq 2$.

2.1. Definition. Let G be a combinatorial HHG, whose structure comes from the action on the combinatorial HHS (X, \mathcal{W}) . We say that G is *short* if it satisfies Axioms (A)-(B)-(C) below.

Axiom A (Underlying graph). X is obtained as a blowup of some graph \overline{X} , which is triangle- and square-free and such that no connected component of \overline{X} is a point. Moreover, \overline{X} is a G -invariant subgraph of X .

The above Axiom implies that the G -action on X restricts to a G -action on \overline{X} with finitely many G -orbits of edges.




FIGURE 2. A simplex of edge-type (on the left) and a simplex of triangle-type (on the right). Their links are represented by the dashed areas.

Axiom B (Vertex stabilisers are cyclic-by-hyperbolic). For every $v \in \overline{X}^{(0)}$ there is an extension

$$0 \longrightarrow Z_v \longrightarrow \text{Stab}_G(v) \xrightarrow{p_v} H_v \longrightarrow 0$$

where H_v is a finitely generated hyperbolic group and Z_v is a cyclic, normal subgroup of $\text{Stab}_G(v)$ which acts trivially on $\text{Lk}_{\overline{X}}(v)$. We call Z_v the *cyclic direction* associated to v .

Moreover, one requires that the family of such extensions is equivariant with respect to the G -action by conjugation; in particular, $Z_{gv} = gZ_vg^{-1}$ for every $v \in \overline{X}^{(0)}$ and $g \in G$.

Notation 2.3. For every $v \in \overline{X}^{(0)}$, let ℓ_v be the domain associated to any triangle-type simplex whose link is $(L_v)^{(0)}$, and let \mathcal{U}_v be the domain associated to any edge-type simplex supported on v .

Axiom C ((Co)bounded actions). For every $v \in \overline{X}^{(0)}$, the cyclic direction Z_v acts geometrically on $\mathcal{C}\ell_v$ and with uniformly bounded orbits on $\mathcal{C}\ell_w$ for every $w \in \text{Lk}_{\overline{X}}(v)$. In particular, $\mathcal{C}\ell_v$ is a quasiline if Z_v is infinite cyclic, and uniformly bounded otherwise.

We will denote a short HHG, together with its short structure, by $(G, \overline{X}, \mathcal{W})$. The reader can find examples of short HHG in [Man24, Sections 2.3 and 6], among which figure the mapping class group of a five-punctured sphere, Artin groups of large and hyperbolic type, non-geometric graph manifold groups, and RAAGs on triangle- and square-free graphs.

Definition 2.4. A short HHG $(G, \overline{X}, \mathcal{W})$ is *colourable* if there exists a partition of the vertices of \overline{X} into finitely many colours, such that no two adjacent vertices share the same colour, and the G -action on \overline{X} descends to an action on the set of colours.

It is easily seen that the above property coincides with the more general notion of *colourability* of a HHG, which requires the existence of a finite, G -invariant colouring of the whole domain set (see e.g. [Bar24, Definition 4.15]).

2.2. Properties.

Remark 2.5. As a consequence of Lemma 2.2, if a simplex $\Delta \subseteq X$ is such that the associated coordinate space $\mathcal{C}(\Delta)$ is unbounded, then

$$[\Delta] \in \{S\} \cup \{\ell_v\}_{v \in \overline{X}^{(0)}} \cup \{\mathcal{U}_v\}_{v \in \overline{X}^{(0)}, |\text{Lk}_{\overline{X}}(v)|=\infty}$$

where $S = [\emptyset]$ is the \sqsubseteq -maximal domain, and ℓ_v, \mathcal{U}_v are defined as in Notation 2.3.

Furthermore, by how nesting and orthogonality are defined in a combinatorial HHS (Definition 1.7), we see that:

- $\ell_v \perp \ell_w$ whenever $v \neq w$ are adjacent in \overline{X} , and are transverse otherwise;
- $\ell_v \perp \mathcal{U}_v$;
- $\ell_v \subseteq \mathcal{U}_w$ whenever $v \neq w$ are adjacent in \overline{X} ;
- If v has valence greater than one in \overline{X} , and $d_{\overline{X}}(v, w) \geq 2$, then $\mathcal{U}_v \pitchfork \ell_w$.
- If both v and w have valence greater than one in \overline{X} , and $d_{\overline{X}}(v, w) \geq 2$, then $\mathcal{U}_v \pitchfork \mathcal{U}_w$.

We avoided describing the slightly more complicated relations involving \mathcal{U}_v when v has valence one in \overline{X} , as we shall not need them.

Remark 2.6. The main coordinate space $\mathcal{CS} = X^{+\mathcal{W}}$ G -equivariantly retracts onto the *augmented support graph* $\overline{X}^{+\mathcal{W}}$, obtained from \overline{X} by adding an edge between v and w if they belong to \mathcal{W} -adjacent maximal simplices of X . In other words, \mathcal{CS} is G -equivariantly quasi-isometric to a graph with vertex set $\overline{X}^{(0)}$, which contains \overline{X} as a (non-full) G -invariant subgraph.

Similarly, for every $v \in \overline{X}^{(0)}$, \mathcal{CU}_v is $\text{Stab}_G(v)$ -equivariantly quasi-isometric to the *augmented link* $\text{Lk}_{\overline{X}}(v)^{+\mathcal{W}}$, on which Z_v acts trivially.

By [Man24, Lemmas 2.10 and 2.11], which are special cases of a more general phenomenon for combinatorial HHG, a short HHG satisfies several strengthened versions of the bounded geodesic image axiom, which we recall here.

Notation 2.7. Set $\rho_{w'}^w := \rho_{\ell_{w'}}^{\ell_w}$, which is defined as $\rho_{[\Delta']}^{[\Delta]}$ for any two simplices of triangle-type $\Delta = \{(v, x), (w)\}$ and $\Delta' = \{(v', x'), (w')\}$. This is a subset of diameter bounded by the HHS constant E .

Lemma 2.8 (Strong BGI for projections on quasilines). *Whenever $u, v, w \in \overline{X}^{(0)}$, if both ρ_w^u and ρ_w^v are defined and at least $2E$ -apart in \mathcal{Cl}_w , then every geodesic $[u, v] \subset \overline{X}^{+\mathcal{W}}$ must pass through $\text{Star}_{\overline{X}}(w)$.*

Similarly, whenever $u, v, w \in \text{Lk}_{\overline{X}}(z)$, if both ρ_w^u and ρ_w^v are defined and at least $2E$ -apart in \mathcal{Cl}_w , then every geodesic $[u, v] \subset \text{Lk}_{\overline{X}}(z)^{+\mathcal{W}}$ must pass through w .

Notation 2.9. For every $u, v, w \in \overline{X}^{(0)}$ such that $u \neq w$ and $v \neq w$, if w has valence greater than one in \overline{X} set

$$d_{\text{Lk}_{\overline{X}}(w)^{+\mathcal{W}}}(u, v) = d_{\text{Lk}_{\overline{X}}(w)^{+\mathcal{W}}}\left(p(\rho_{\mathcal{U}_w}^{\ell_u}), p(\rho_{\mathcal{U}_w}^{\ell_v})\right),$$

where $p: X \rightarrow \overline{X}$ is the retraction. If instead w has valence one then $\text{Lk}_{\overline{X}}(w)$ is a point, and we set $d_{\text{Lk}_{\overline{X}}(w)^{+\mathcal{W}}}(u, v) = 0$.

Lemma 2.10 (Strong BGI for projections on links). *Let $w \in \overline{X}^{(0)}$. For every $u, v \in \overline{X}^{(0)} - \{w\}$, if $d_{\text{Lk}_{\overline{X}}(w)^{+\mathcal{W}}}(u, v) \geq 2E$, then every geodesic $[u, v] \subset \overline{X}^{+\mathcal{W}}$ must pass through w .*

Finally, recall that, given a hyperbolic group H and an infinite cyclic subgroup $Q \leq H$, there exists a maximal virtually cyclic subgroup $K_H(Q)$ containing Q .

Lemma 2.11 ([Man24, Lemma 2.17]). *Let $(G, \overline{X}, \mathcal{W})$ be a short HHG, $v \in \overline{X}^{(0)}$, and let W be a collection of $\text{Stab}_G(v)$ -orbit representatives of vertices in $\text{Lk}_{\overline{X}}(v)$ with infinite cyclic directions. Then H_v is hyperbolic relative to $\{K_{H_v}(\mathbf{p}_v(Z_w))\}_{w \in W}$.*

2.3. Blowup materials. The main technical result of [Man24] is that short HHG can be equivalently defined as those groups equipped with certain data, called *blowup materials*. We first set up some notation.

Definition 2.12. A finitely generated group G is *weakly hyperbolic* relative to a collection of conjugacy classes of subgroups if, given any finite generating set S for G , and any choice of a subgroup H_i in each conjugacy class, the coned-off Cayley graph $\text{Cay}(G, S \cup_i H_i)$ is hyperbolic.

Notation 2.13 (Product regions). In what follows, G is a group acting on a simplicial graph \overline{X} with finitely many orbits of edges. Fix $V = \{v_1, \dots, v_k\}$ a set of representatives of the G -orbits of vertices in \overline{X} . For every $v_i \in V$ and every $g \in G$, set $P_{gv_i} = g \text{Stab}_G(v_i)$, which we call the *product region* associated to gv_i (with respect to this choice of orbit representatives). We always see P_v as a metric subspace of G , where the latter is equipped with the word metric with respect to any finite generating set.

The following definition is equivalent to [Man24, Definition 3.9], which we rephrase for improved clarity.

Definition 2.14 (Blowup materials). Let G be a finitely generated group, acting on a simplicial graph \overline{X} with finitely many orbits of edges. For every $v \in \overline{X}^{(0)}$ let $Z_v \trianglelefteq \text{Stab}_G(v)$ be a cyclic, normal subgroup; furthermore, whenever Z_v is infinite, let $\phi_v: E_v \rightarrow \mathbb{R}$ be a homogeneous quasimorphism, where $E_v \trianglelefteq \text{Stab}_G(v)$ is some finite-index, normal subgroup centralising $Z_v \cap E_v$. Let V be a set of representatives for the G -orbits of vertices in \overline{X} , and for every $v \in \overline{X}^{(0)}$ let P_v be the associated product region with respect to V , as in Notation 2.13.

The above data defines *blowup materials* on G if the following conditions hold:

- (0) All subgroups and quasimorphisms are equivariant with respect to the G -action by conjugation.
- (1) \overline{X} is triangle- and square-free, and no connected component of \overline{X} is a single point.
- (2) For every $v \in \overline{X}^{(0)}$, the quotient $H_v := \text{Stab}_G(v)/Z_v$ is hyperbolic.
- (3) Whenever $e = \{v, w\}$ is an edge of \overline{X} , Z_v fixes w , and $\text{Stab}_G(v) \cap \text{Stab}_G(w)$ contains $\langle Z_v, Z_w \rangle$ as a subgroup of finite index. Moreover, $\mathfrak{p}_v(Z_w)$ is quasiconvex in H_v .
- (4) G is weakly hyperbolic relative to $\{\text{Stab}_G(v)\}_{v \in \overline{X}}$.
- (5) For every $v \in \overline{X}^{(0)}$ with Z_v infinite, the quasimorphism ϕ_v is unbounded on $Z_v \cap E_v$ and trivial on $Z_w \cap E_v$ for every vertex $w \in \text{Lk}_{\overline{X}}(v)$.
- (6) There exist a constant $B \geq 0$ and, for every $v_i \in V$, a B -coarsely Lipschitz, B -coarse retraction $\mathfrak{g}_{v_i}: G \rightarrow 2^{E_{v_i}}$, with the following properties:
 - (i) If $w \in \text{Lk}_{\overline{X}}(v_i)$ then $\mathfrak{g}_{v_i}(P_w) \subseteq N_B(P_w)$.
 - (ii) For every i there exists a finite collection $\{w_i^1, \dots, w_i^r\} \subseteq \text{Lk}_{\overline{X}}(v_i)$ such that, whenever $d_{\overline{X}}(v_i, u) \geq 2$, there exist $g \in \text{Stab}_G(v_i)$ and $1 \leq j \leq r$ such that $\mathfrak{g}_{v_i}(P_u) \subseteq N_B(gZ_{w_i^j})$.

We call \overline{X} the *support graph*, each Z_v the *cyclic direction* associated to v , and each \mathfrak{g}_{v_i} the *gate* onto E_{v_i} .

Theorem 2.15. ([Man24, Theorem 3.10]) *Let G be a finitely generate group admitting blowup materials, with support graph \overline{X} . Then G admits a short HHG structure, where \overline{X} is the support graph and the cyclic direction associated to each $v \in \overline{X}^{(0)}$ is (a finite-index subgroup of) Z_v .*

Proposition 2.16. ([Man24, Proposition 4.1]) *A short HHG $(G, \overline{X}, \mathcal{W})$ admits blowup materials, whose support graph is \overline{X} and whose extensions $Z_v \rightarrow \text{Stab}_G(v) \rightarrow H_v$ are those from Axiom B.*

In particular, we get:

Corollary 2.17 (Edge groups). *Let $(G, \overline{X}, \mathcal{W})$ be a short HHG. Whenever $\{v, w\}$ is an edge of \overline{X} , the edge group $\text{Stab}_G(v) \cap \text{Stab}_G(w)$ contains $\langle Z_v, Z_w \rangle$ as a subgroup of finite index.*

2.4. Subgroups generated by large cyclic directions.

Notation 2.18 (Kernel of the projection). Let $(G, \overline{X}, \mathcal{W})$ be a colourable short HHG. Let $\mathcal{B} = \{s_1, \dots, s_r\} \subset \overline{X}^{(0)}$ be a (possibly non-maximal) collection of vertices belonging to pairwise different G -orbits. For every $i = 1, \dots, r$ choose a non-zero natural number $M_i \in \mathbb{N} - \{0\}$, and for every $v \in G\{s_i\}$ let $\Gamma_v = M_i Z_v$. From now on, we will focus on all normal subgroups of the form

$$\mathcal{N} := \langle \langle M_1 Z_{s_1}, \dots, M_r Z_{s_r} \rangle \rangle = \langle \Gamma_v \rangle_{v \in G\mathcal{B}}.$$

\mathcal{B} will be called the *base* of \mathcal{N} .

Definition 2.19 (Deep enough quotient). Let $\mathcal{N} = \langle \langle M_1 Z_{s_1}, \dots, M_r Z_{s_r} \rangle \rangle$ be as in Notation 2.18. We will say that a property of \mathcal{N} , or of G/\mathcal{N} , holds *if \mathcal{N} is deep enough* if there exists $D \in \mathbb{N} - \{0\}$ such that the property holds whenever every M_i is a multiple of D .

Remark 2.20. Notice that, if \mathcal{N} is deep enough, one can always assume that:

- Every finite Γ_v is trivial, or in other words \mathcal{B} does not contain vertices with finite cyclic direction;
- Whenever v, w are \overline{X} -adjacent, Z_v commutes with Γ_w (this is because the centraliser of Z_v in $\text{Stab}_G(v)$ has index at most two);
- \mathcal{N} lies in the normal subgroup G_0 of G which acts trivially on the set of colours of \overline{X} .

3. ROTATING FAMILIES AND PROJECTION GRAPHS, REVISITED

In this Section, we prove that the machinery of *composite projection systems*, devised in [Dah18] and then developed further in [DHS21], can be adapted to study quotients of colourable short HHG by powers of cyclic directions. This allows us to describe the quotient graphs \overline{X}/\mathcal{N} and $\overline{X}^{+\mathcal{W}}/\mathcal{N}$, and the vertex stabilisers for the induced G/\mathcal{N} -action on \overline{X}/\mathcal{N} . The reader who is only interested in short HHG is advised to skip to the consequences which are gathered in Subsection 3.3.

3.1. Strong composite projection graphs. We first recall some definitions from [Dah18].

Definition 3.1 ([Dah18, Definition 1.2]). Let \mathbb{Y}_* be the disjoint union of finitely many countable sets $\mathbb{Y}_1, \dots, \mathbb{Y}_m$, called *colours*, and for every $y \in \mathbb{Y}_*$ let $i(y)$ be the index such that $y \in \mathbb{Y}_{i(y)}$. A *composite projection system* on \mathbb{Y}_* consists of

- a constant $\theta \geq 0$;
- a family of subsets $\text{Act}(y) \subset \mathbb{Y}_*$ for $y \in \mathbb{Y}_*$ (the *active set* for y) such that $\mathbb{Y}_{i(y)} \subset \text{Act}(y)$, and such that $x \in \text{Act}(y)$ if and only if $y \in \text{Act}(x)$ (*symmetry in action*),
- and a family of functions $d_y^\pi : (\text{Act}(y) - \{y\})^2 \rightarrow \mathbb{R}_+$, satisfying:
 - **Symmetry:** $d_y^\pi(x, z) = d_y^\pi(z, x)$ for $x, z \in \text{Act}(y) - \{y\}$;
 - **Triangle inequality:** $d_y^\pi(w, x) \leq d_y^\pi(w, z) + d_y^\pi(z, x)$ for all $w, x, z \in \text{Act}(y) - \{y\}$;
 - **Behrstock inequality:** $\min\{d_y^\pi(x, z), d_z^\pi(x, y)\} \leq \theta$ whenever both quantities are defined;
 - **Properness:** $|\{y \in \mathbb{Y}_i, d_y^\pi(x, z) \geq \theta\}| < \infty$ for all $x, z \in \mathbb{Y}_i$;
 - **Separation:** $d_y^\pi(z, z) < \theta$ for $z \in \text{Act}(y) - \{y\}$;
 - **Closeness in inaction:** if $x \notin \text{Act}(z)$ then, for all $y \in \text{Act}(x) \cap \text{Act}(z)$, we have $d_y^\pi(x, z) \leq \theta$;
 - **Finite filling:** for all $\mathcal{Z} \subset \mathbb{Y}_*$, there is a finite collection $x_j \in \mathcal{Z}$ such that $\cup_j \text{Act}(x_j)$ covers $\cup_{x \in \mathcal{Z}} \text{Act}(x)$.

We will also require the following “uniform” version of the properness axiom:

Definition 3.2. A composite projection system on \mathbb{Y}_* is *uniformly proper* if $|\{y \in \mathbb{Y}_*, d_y^\pi(x, z) \geq \theta\}| < \infty$ for all $x, z \in \mathbb{Y}_*$.

Definition 3.3. A group G acts on a composite projection system \mathbb{Y}_* by isomorphisms if:

- G acts on \mathbb{Y}_* , and the action permutes the colours;
- for all $g \in G$ and $y \in \mathbb{Y}_*$, $\text{Act}(gy) = g \text{Act}(y)$;
- for all $g \in G$ and $x, y, z \in \mathbb{Y}_*$ such that $d_y^\pi(x, z)$ is defined, $d_{gy}^\pi(gx, gz) = d_y^\pi(x, z)$.

Remark 3.4. We shall follow the constructions and arguments from [DHS21] as closely as possible, in order to recover their results without altering most of the proofs. To this extent, arguing as in the paragraph below [DHS21, Definition 1.1], we first replace each distance d_y^π by a modified function $d_y^* : (\text{Act}(y) - \{y\})^2 \rightarrow \mathbb{R}_+$, constructed as in [BBF15, Theorem 3.3]. The new functions d_y^* differ from d_y^π by a uniformly bounded amount depending only on θ , and they satisfy almost all properties from Definitions 3.1 and 3.2 with θ replaced by $\Theta = 4\theta + 1$; the only difference is that the triangle inequality is “coarsified”, meaning that there exists a constant $\kappa \geq 0$ such that, for all $w, x, z \in \text{Act}(y) - \{y\}$, $d_y^*(w, x) \leq d_y^*(w, z) + d_y^*(z, x) + \kappa$. Furthermore, if G acts on (\mathbb{Y}_*, d_*^π) by isomorphisms, then the new functions are G -equivariant, meaning that for all $g \in G$ and $x, y, z \in \mathbb{Y}_*$ such that $d_y^*(x, z)$ is defined, $d_{gy}^*(gx, gz) = d_y^*(x, z)$.

Definition 3.5. (Composite rotating family) Consider a composite projection system $(\mathbb{Y}_*, \theta, d_*^\pi)$, endowed with an action of a group G by isomorphisms. Let $\Theta = 4\theta + 1$, and let d_*^* be the modified functions from Remark 3.4.

A *composite rotating family* on (\mathbb{Y}_*, G) , with *rotation control* $\Theta_{rot} > 0$, is a family of subgroups $\{\Gamma_v\}_{v \in \mathbb{Y}_*}$ such that the following holds:

- For all $x \in \mathbb{Y}_*$, $\Gamma_x < \text{Stab}_G(x)$ is an infinite group;
- whenever $y = x$ or $y \notin \text{Act}(x)$, the subgroup Γ_x fixes y and d_y^* ;
- for all $g \in G$, and all $x \in \mathbb{Y}_*$, we have $\Gamma_{gx} = g\Gamma_x g^{-1}$;

- if $x \notin \text{Act}(z)$ then Γ_x and Γ_z commute;
- Γ_x acts with **proper isotropy**, meaning that for all $R > 0$ and $y \in \text{Act}(x)$, the set $\{\gamma \in \Gamma_x, d_x^*(y, \gamma y) < R\}$ is finite;
- for all $i \leq m$ and for all $x, y, z \in \mathbb{Y}_i$, if $d_y^*(x, z) \leq \Theta$, then

$$d_y^*(x, gz) \geq \Theta_{rot}$$

for all $g \in \Gamma_y - \{1\}$.

Set \mathcal{N} as the normal subgroup of G generated by $\bigcup_{v \in \mathbb{Y}_*} \Gamma_v$.

Definition 3.6 (SCPG). Let G be a finitely generated group. A *strong G -composite projection graph* is the data of:

- A hyperbolic graph \mathcal{S} on which G acts by simplicial automorphisms, with finitely many orbits of vertices;
- A structure of a uniformly proper composite projection system on $\mathcal{S}^{(0)}$, with constant θ , on which the induced action of G is by isomorphisms;
- A G -invariant subset $\mathbb{Y}_* \subset \mathcal{S}^{(0)}$, which inherits the structure of a composite projection system;
- A composite rotating family $\{\Gamma_v\}_{v \in \mathbb{Y}_*}$ with rotation control Θ_{rot} .
- *Strong Bounded Geodesic Image*: a constant $C \in \mathbb{R}_+$ so that the following holds. For each $x, y, s \in \mathcal{S}$ so that $d_s^*(x, y)$ is defined and larger than C , every geodesic $[x, y] \subset \mathcal{S}$ contains a vertex w such that $d_{\mathcal{S}}(w, s) \leq C$, and either $w = s$ or $w \notin \text{Act}(s)$. In particular, w is fixed by Γ_s .

We shall denote a SPCG by the tuple $(\mathcal{S}, G, \mathbb{Y}_*, \{\Gamma_v\})$.

Remark 3.7 (Comparison with [DHS21]). The above is a slight variant of the main Definition in [DHS21, Section 2]. There are three notable differences:

- In [DHS21] the authors require that, for every $s \in \mathcal{S}^{(0)}$, Γ_s fixes $\text{Lk}_{\mathcal{S}}(s)$, which is in general not true for short HHG (see Remark 3.29).
- The version of the bounded geodesic image appearing in [DHS21] asks that, whenever x, y project far to s , on any \mathcal{S} -geodesic $[x, y]$ there exists a vertex w which is \mathcal{S} -adjacent to s . However, what is actually needed to run the arguments from [DHS21] is the existence of some $w \in [x, y]$ which is fixed by Γ_s and lies within uniform \mathcal{S} -distance from s . Both these requirements are part of our strong bounded geodesic image assumption.
- Though the definition in [DHS21] involves a sub-projection complex \mathbb{Y}_*^{τ} on which the composite rotating family is defined, all proofs there implicitly assume that \mathbb{Y}_*^{τ} coincides with the whole projection complex. Nonetheless, this does not invalidate the consequences the authors get in [DHS21, Section 5] about mapping class groups of surfaces, since in that setting the composite rotating family is defined on the whole projection complex. In our setting, instead, it is relevant that the composite rotating family \mathbb{Y}_* is *not* defined on the whole $\mathcal{S}^{(0)}$, because we want to be able to quotient by any collection of cyclic directions. We establish a way to pass from the whole projection complex $\mathcal{S}^{(0)}$ to the sub-complex \mathbb{Y}_* in the Transfer-like Lemma 3.9 below, whose main ingredients are uniform properness (Definition 3.2) and the fact that G acts cofinitely on $\mathcal{S}^{(0)}$.

Standing assumption 3.8. For the rest of the Subsection, let \mathcal{S} be a hyperbolic graph, let $(\mathcal{S}^{(0)}, \theta, d_*^{\pi})$ be a uniformly proper composite projection system, and let

$(\mathcal{S}, G, \mathbb{Y}_*, \{\Gamma_v\})$ be a SCPG with rotation control Θ_{rot} . Let \mathcal{N} be the subgroup of G generated by $\{\Gamma_v\}_{v \in \mathbb{Y}_*}$.

Lemma 3.9 (Transfer-like Lemma). *There exists a constant $B \geq 0$, not depending on Θ_{rot} , with the following property. For every $x \in \mathcal{S}^{(0)}$ and every $i = 1, \dots, m$ there exists $t_i(x) \in \mathbb{Y}_i$ such that, for every $y \in \text{Act}(x) \cap \text{Act}(t_i(x))$, we have that $d_y^*(x, t_i(x)) \leq B$.*

Proof. Let $x_1, \dots, x_l \in \mathbb{Y}_*$ be a finite collection of representatives of the G -orbits. For each of these points x_j and for each colour $i = 1, \dots, m$ choose a point $t_i(x_j) \in \mathbb{Y}_i$. By uniform properness (Definition 3.2, together with Remark 3.4), there exists a constant B such that

$$\max_{i,j} \sup_{y \in \text{Act}(x_j) \cap \text{Act}(t_i(x_j))} d_y^*(x_j, t_i(x_j)) \leq B$$

Now, for every $x \in \mathbb{Y}_*$ there exists $j(x) \leq l$ and an element $g \in G$ such that $x = g(x_{j(x)})$. Moreover, since G acts on the set of colours, for every $i = 1, \dots, m$ there exists a unique $i' \in \{1, \dots, m\}$ such that $gt_{i'}(x_{j(x)}) \in \mathbb{Y}_i$, and we can set $t_i(x) = gt_{i'}(x_{j(x)})$. Notice that, since G preserves projection distances,

$$\sup_{y \in \text{Act}(x) \cap \text{Act}(t_i(x))} d_y^*(x, t_i(x)) = \sup_{y \in \text{Act}(x_j) \cap \text{Act}(t_{i'}(x_j))} d_y^*(x_j, t_{i'}(x_j)) \leq B,$$

and this concludes the proof. \square

Corollary 3.10 (cf. [DHS21, Corollary 1.8]). *The following holds if the rotation control Θ_{rot} is large enough. Let $v \in \mathcal{S}^{(0)}$ and $w \in \mathbb{Y}_*$. If v is w -active, then $d_w^*(v, \gamma_w v) > \theta$ for every $\gamma_w \in \Gamma_w - \{0\}$. In particular, $\gamma_w v \neq v$.*

Proof. Let $t = t_{i(w)}(v)$. Since now t and w lie in the same colour, we have that $d_w^*(t, \gamma_w t) > \Theta_{rot}$. Then the coarse triangular inequality yields that $d_w^*(v, \gamma_w v) > \Theta_{rot} - 2B - 2\kappa$, and we can choose the rotation control so that the latter quantity is greater than θ . \square

The following Lemma combines some properties from [DHS21, Section 3], which only use that $\{\Gamma_v\}_{v \in \mathbb{Y}_*}$ is a composite rotating family on a composite projection complex \mathbb{Y}_* .

Lemma 3.11. *There exist a constant $\aleph \in \mathbb{R}$ depending only on \mathbb{Y}_* , a well-order of \mathcal{N} which we call complexity, and an indexing $i: \mathcal{N} \rightarrow \{1, \dots, m\}$ such that the following holds. For every $\gamma \in \mathcal{N} - \{1\}$ and every $t \in \mathbb{Y}_{i(\gamma)}$, there exist $s \in \mathbb{Y}_*$ and $\gamma_s \in \Gamma_s$ such that:*

- $\gamma_s \gamma$ has strictly lower complexity than γ ;
- $d_s^*(t, \gamma t) \geq \Theta_{rot}/2 - \aleph$.

Proof. Set $\aleph = \Theta_0/2 + \kappa$, where κ is the constant from Remark 3.4 and Θ_0 is defined in [DHS21, Remark 1.2 and Standing assumption 1.4] as $\Theta_0 = 8\theta + 2 + 3\kappa$. We stress that both θ and κ only depend on the composite projection system on \mathbb{Y}_* and on the modified distance function from Remark 3.4. The well-order is the lexicographic order $(\alpha(\gamma), n(\gamma))$, where the ordinal α and the integer n are defined in [DHS21, Theorem 3.1 and Definitions 3.2 and 3.3]. The index $i(\gamma) := i(\alpha(\gamma))$ is defined in [DHS21, Theorem 3.1]. The statement about γ and t follows by combining the second and the fourth bullet of [DHS21, Theorem 3.5]. \square

Now set $\Theta_{short} = \Theta_{rot}/2 - \aleph - 2B - 2\kappa$, where B is the constant from the Transfer-like Lemma 3.9.

Lemma 3.12 (Shortening pair). *For all $\gamma \in \mathcal{N} - \{1\}$, and all $x \in \mathcal{S}^{(0)}$, there exist a shortening pair (s, γ_s) (here $s \in \mathbb{Y}_*$ and $\gamma_s \in \Gamma_s$) so that $\gamma_s \gamma$ has strictly lower complexity than γ , and either*

- (1) *at least one of x and γx is s -inactive, or*
- (2) $d_s^*(x, \gamma x) \geq \Theta_{short}$.

Proof. Let $t = t_{i(\gamma)}(x)$, defined as in the Transfer-like Lemma 3.9. By Lemma 3.11, there exist $s \in \mathbb{Y}_*$ and $\gamma_s \in \Gamma_s$ such that $\gamma_s \gamma$ has strictly lower complexity, and $d_s^*(t, \gamma t) \geq \Theta_{rot}/2 - \aleph$. Now, if one between x and γx is s -inactive we are done. Otherwise, the (coarse) triangle inequality yields

$$d_s^*(x, \gamma x) \geq d_s^*(t, \gamma(t)) - 2B - 2\kappa \geq \Theta_{short},$$

where we used that

$$d_s^*(x, t) \leq \sup_{v \in \text{Act}(x) \cap \text{Act}(t)} d_v^*(x, t) \leq B,$$

and therefore also

$$d_s^*(\gamma x, \gamma t) \leq \sup_{v \in \text{Act}(\gamma x) \cap \text{Act}(\gamma t)} d_v^*(\gamma x, \gamma t) = \sup_{v' \in \text{Act}(x) \cap \text{Act}(t)} d_{v'}^*(x, t) \leq B.$$

□

3.1.1. Lifting and projecting. From now on, let \mathcal{S}/\mathcal{N} be the graph whose vertices and edges are \mathcal{N} -orbits of vertices and edges in \mathcal{S} , and let $q: \mathcal{S} \rightarrow \mathcal{S}/\mathcal{N}$ be the quotient projection. Notice that, a priori, \mathcal{S}/\mathcal{N} might not be a simplicial graph. The next lemmas are the analogues of the results in [DHS21, Section 4], whose proofs can be run verbatim in our setting. Indeed, they all rely only on [DHS21, Corollary 3.6] (which is our Lemma 3.12), and the fact that, whenever $x, y \in \mathcal{S}^{(0)}$ have sufficiently large projection on some $s \in \mathcal{S}^{(0)}$, then every geodesic $[x, y]$ contains a point w which is fixed by Γ_s (which for us is a consequence of the strong bounded geodesic image assumption).

Lemma 3.13. *If the rotation control Θ_{rot} is sufficiently large, then \mathcal{S}/\mathcal{N} is a simplicial graph.*

Proof. No double edges. Let $e = \{v, w\}$ be any edge in \mathcal{S} , and suppose there exist $n, n' \in \mathcal{N}$ such that $e' = \{nv, n'w\}$ is also an edge. Up to replacing e' by $(n')^{-1}e'$, assume that $n' = 1$. We now prove that e and e' must lie in the same \mathcal{N} -orbit, and therefore \mathcal{S}/\mathcal{N} has no double edges.

We argue by induction on the complexity of n . If $nv = v$, then the two edges coincide, and we are done. Otherwise, by Lemma 3.12 there exists a shortening pair (s, γ_s) . If v is s -inactive, then we can apply γ_s to both edges, and we get a new concatenation of translates of e and e' with endpoints v and $\gamma_s nv$. Then we can conclude by induction, as $\gamma_s n$ has strictly lower complexity than n . Similarly, if nv is s -inactive, we simply replace n by $\gamma_s n$, and again we get a concatenation of two edges with endpoints v and $\gamma_s nv$. If w is s -inactive, we can apply γ_s to e' ; this yields a new concatenation of translates of e and e' since γ_s fixes w , and again we conclude by induction.

If none of the above happens, we have that $d_s^*(v, nv) \geq \Theta_{short}$, and by the coarse triangle inequality one of $d_s^*(v, w)$ and $d_s^*(w, nv)$ is at least $\Theta_{short}/2 - \kappa$. Notice

that we can always choose Θ_{rot} large enough that the latter quantity is greater than the constant C from the strong BGI assumption. But this is a contradiction, because then one of v, w, nv would have to be fixed by γ_s .

No loops. Towards a contradiction, suppose that there exists an edge e of \mathcal{S} with endpoints v, nv for some $n \in \mathcal{N}$, and choose n of minimal complexity among all such edges. Since $n \neq 1$, there exists a shortening pair (s, γ_s) , as in Lemma 3.12. If v is s -inactive, then we can replace e by $\gamma_s e = \{v, \gamma_s nv\}$, thus violating minimality of n since $\gamma_s n$ has strictly lower complexity. Similarly, if nv is s -inactive, then $e = \{v, \gamma_s nv\}$, again contradicting minimality of n . Otherwise $d_s^*(v, nv) \geq \Theta_{short}$, which we can assume to be bigger than the constant C from the strong BGI assumption. This is a contradiction, since then one of v and nv would have to be s -inactive. \square

For the next lemma, by a *combinatorial path* in a simplicial graph we mean a concatenation of edges.

Lemma 3.14. *For each combinatorial path $\bar{\gamma}$ in \mathcal{S}/\mathcal{N} starting at \bar{x} , and any point x in the preimage of \bar{x} (henceforth: a lift of \bar{x}), there exists a combinatorial path in \mathcal{S} so that $q \circ \gamma = \bar{\gamma}$, which we call a lift of $\bar{\gamma}$. Moreover, if $\bar{\gamma}$ is a geodesic, then so is γ .*

Proof. This follows from the fact that \mathcal{N} acts simplicially on \mathcal{S} , and that q is 1-Lipschitz. \square

Definition 3.15. A k -gon in a simplicial graph Λ is a subgraph of the form

$$Q = \bigcup_{i=0}^{k-1} [x_i, x_{i+1}],$$

where $x_0, \dots, x_k \in \Lambda^{(0)}$, $x_k = x_0$ and each $[x_i, x_{i+1}]$ is a geodesic path.

Lemma 3.16 (cf. [DHS21, Proposition 4.3]). *For every $k \in \mathbb{N}$ there exists a constant $\Theta_{rot}(k)$ such that the following holds if the rotation control Θ_{rot} is larger than $\Theta_{rot}(k)$. For every k -gon $\bar{Q} \subset \mathcal{S}/\mathcal{N}$ there exists a k -gon $Q \subset \mathcal{S}$ such that $q(Q) = \bar{Q}$, which we call a lift of \bar{Q} .*

Proof. Let $\bar{Q} = \bigcup_{i=0}^{k-1} [\bar{x}_i, \bar{x}_{i+1}]$. Lift all segments of \bar{Q} to a (possibly non-closed) chain $\bigcup_{i=0}^{k-1} [x_i, x_{i+1}]$. Let $\gamma \in \mathcal{N}$ be such that $\gamma x_0 = x_k$. We proceed by induction on the complexity of γ . If $\gamma = 1$ then $x_0 = x_k$ and the chain is already a k -gon. Otherwise, by Lemma 3.12 there exists a shortening pair (s, γ_s) . If x_0 is s -inactive, then we can apply γ_s to the whole chain, and we get a new chain with endpoints $x'_0 = \gamma_s x_0 = x_0$ and $x'_k = \gamma_s \gamma x_0$. Then we can conclude by induction, as $\gamma_s \gamma$ has strictly lower complexity than γ . Similarly, if γx_0 is s -inactive, we can apply γ_s^{-1} to the whole chain, and we get a new chain with endpoints $x'_0 = \gamma_s^{-1} x_0$ and $x'_k = \gamma_s \gamma x_0 = \gamma x_0$. Again, the endpoints of the chain are such that $x'_k = \gamma_s \gamma x'_0$, and we conclude by induction.

Otherwise, we have that $d_s^*(x_0, \gamma x_0) \geq \Theta_{short}$. If at least one x_i is s -inactive, we can apply γ_s to the chain “after” x_i (meaning, to all geodesics $[x_j, x_{j+1}]$ where $i \leq j$), and again conclude by induction. Otherwise, by the coarse triangle inequality we get that there exists some $i \in \{0, \dots, k-1\}$ such that $d_s^*(x_i, \gamma x_{i+1}) \geq \Theta_{short}/k - k\kappa$. Notice that we can always choose Θ_{rot} large enough that the latter quantity is

greater than the constant C from the bounded geodesic image assumption. Therefore, there exists $w \in [x_i, x_{i+1}]$ such that either $w = s$ or w is s -inactive. In both cases γ_s fixes w , and therefore one can change the lift of the chain “after” w (meaning, one can apply γ_s to $[w, x_{i+1}]$ and to all geodesics $[x_j, x_{j+1}]$ where $i + 1 \leq j$). Again, the conclusion then follows by induction. \square

With the same techniques, one can prove the following two results:

Lemma 3.17 (cf. [DHS21, Proposition 4.3], “moreover” part). *Let \bar{Q} be a geodesic quadrangle in \mathcal{S}/\mathcal{N} . If the geodesics $[\bar{v}_1, \bar{w}_1]$, $[\bar{v}_2, \bar{w}_2]$ of Q have lifts $[v_i, w_i]$ so that $d_s^*(v_i, w_i) \leq \Theta_{\text{rot}}/10$ whenever the quantity is defined, then there exists a lift Q of \bar{Q} such that the lift $[v'_i, w'_i]$ of $[\bar{v}_i, \bar{w}_i]$ contained in Q is an \mathcal{N} -translate of $[v_i, w_i]$.*

Lemma 3.18 (cf. [DHS21, Lemma 4.4]). *Suppose that $x, y \in \mathcal{S}^{(0)}$ have the property that $d_s^*(x, y) \leq \Theta_{\text{rot}}/10$ whenever the quantity is defined. Then $q|_{[x, y]}$ is isometric, for any geodesic $[x, y] \subset \mathcal{S}$.*

We can finally prove the following:

Theorem 3.19 (cf. [DHS21, Theorem 2.1]). *Let $(\mathcal{S}, G, \mathbb{Y}_*, \{\Gamma_v\})$ be a SCPG, and let \mathcal{N} be the subgroup of G generated by $\bigcup_{v \in \mathbb{Y}_*} \Gamma_v$. If the rotation control Θ_{rot} is sufficiently large, then:*

- (1) \mathcal{S}/\mathcal{N} is δ -hyperbolic, where δ is any hyperbolicity constant for \mathcal{S} ;
- (2) If the action of G on \mathcal{S} admits a loxodromic element (resp. loxodromic WPD), then so does the action of G/\mathcal{N} on \mathcal{S}/\mathcal{N} . More precisely, for every loxodromic (resp. loxodromic WPD) element $g \in G$ there exists a constant $\Theta_{\text{loxo}}(g)$ such that, if $\Theta_{\text{rot}} > \Theta_{\text{loxo}}(g)$, then $\bar{g} \in G/\mathcal{N}$ is loxodromic (resp. loxodromic WPD).
- (3) If the action of G on \mathcal{S} is non-elementary, then so is the action of G/\mathcal{N} on \mathcal{S}/\mathcal{N} .

Proof. Item (1) follows from the fact that geodesic triangles in \mathcal{S}/\mathcal{N} lift to geodesic triangles in \mathcal{S} , and that the projection $\mathcal{S} \rightarrow \mathcal{S}/\mathcal{N}$ is 1-Lipschitz. Next, fix a basepoint $x_0 \in \mathcal{S}$, and let $g \in G$ act loxodromically on \mathcal{S} .

Claim 3.20. *There exists $\Theta_{\text{loxo}}(g) > 0$ such that $d_s^*(x_0, g^k x_0) \leq \Theta_{\text{loxo}}(g)/10$ for every $k \in \mathbb{Z}$ and every s for which the projection is defined.*

If the Claim is true, then Lemma 3.18 implies that the projection map is an isometry on $\{g^k x_0\}_{k \in \mathbb{N}}$ whenever $\Theta_{\text{rot}} > \Theta_{\text{loxo}}(g)$, and in particular \bar{g} is loxodromic. If moreover g is WPD, then \bar{g} is WPD as well by [DHS21, Proposition 4.6], whose proof runs verbatim once one replaces [DHS21, Proposition 4.3] by Lemma 3.17.

Proof of Claim 3.20. Let $\mu = \mu(g, \delta)$ be such that each geodesic $[x_0, g^k x_0]$ lies at Hausdorff distance at most μ from the quasigeodesic sequence $x_0, gx_0, \dots, g^k x_0$. Fix $s \in \mathcal{S}^{(0)}$. By strong BGI, either $d_s^*(x_0, g^k x_0) \leq C$, or there exists some vertex on $[x_0, g^k x_0]$ within distance C from s . Hence there exist integers p, q with $1 \leq p \leq q \leq k$ and $|p - q|$ bounded by some constant depending only on g , so that s is not C -close to any vertex of $[x_0, g^i x_0]$ or $[g^i x_0, g^k x_0]$ unless $p \leq i \leq q$. Thus, by strong BGI and the coarse triangle inequality, we get

$$d_s^*(x_0, g^k x_0) \leq d_s^*(g^p x_0, g^q x_0) + 2C + 2\kappa = d_{g^{-p}s}^*(x_0, g^{q-p} x_0) + 2C + 2\kappa.$$

Now notice that the quantity $\sup_{s \in \mathbb{Y}_*} d_{g^{-p}s}^*(x_0, g^{q-p}x_0) = \sup_{s \in \mathbb{Y}_*} d_s^*(x_0, g^{q-p}x_0)$ is bounded above by uniform properness. Moreover, the difference $q - p$ can assume finitely many values, so we see that $d_s^*(x_0, g^k x_0)$ is uniformly bounded. \blacksquare

We finally turn to Item (3). Let g, h be independent loxodromic elements, and suppose that $\Theta_{rot} \geq \max\{\Theta_{loxo}(g), \Theta_{loxo}(h)\}$, so that both \bar{g} and \bar{h} are loxodromic in the quotient.

Claim 3.21. *There exists $\Theta_{ind}(g, h) > 0$ such that $d_s^*(g^k x_0, h^l x_0) \leq \Theta_{ind}(g, h)/10$ for every $k, l \in \mathbb{Z}$ and every s for which the projection is defined.*

Proof. Since g, h are independent, there exists μ such that, for all integers k, l , any geodesic $[g^k x_0, h^l x_0]$ lies at Hausdorff distance at most μ from the sequence $\{g^k x_0, g^{k-1} x_0, \dots, x_0\} \cup \{x_0, h x_0, \dots, h^l x_0\}$. We can now argue as in Claim 3.20, replacing $\{x_0, g x_0, \dots, g^k x_0\}$ by $\{g^k x_0, g^{k-1} x_0, \dots, x_0\} \cup \{x_0, h x_0, \dots, h^l x_0\}$. \blacksquare

By combining Claim 3.21 and Lemma 3.18 we get that the projection map is an isometry on $\langle g \rangle x_0 \cup \langle h \rangle x_0$, and therefore \bar{g} and \bar{h} are again independent. \square

3.1.2. Other consequences. We gather here a collection of facts which are not used in the above proof, but are relevant for our arguments on short HHG. Firstly, the following theorem gives a presentation of \mathcal{N} under the (fairly restrictive) assumption that any two vertices are active. This is stated in the language of the more general [Dah18, Theorem 2.2], though it already follows from [DGO17, Theorem 5.3.(a)].

Theorem 3.22. *Let $\{\Gamma_v\}_{v \in \mathbb{Y}_*}$ be a composite rotating family on (\mathbb{Y}_*, G) . Suppose that $\text{Act}(z) = \mathbb{Y}_*$ for every $z \in \mathbb{Y}_*$. If the rotation control Θ_{rot} is sufficiently large, there exists a (possibly infinite) subset $\mathcal{J} \subset \mathbb{Y}_*$ such that*

$$\mathcal{N} \cong \bigast_{v \in \mathcal{J}} \Gamma_v.$$

Next, we describe the elements of the kernel \mathcal{N} fixing a given vertex;

Lemma 3.23 (cf. [DHS21, Proposition 4.8]). *If the rotation control Θ_{rot} is sufficiently large, then for any vertex $v \in \mathcal{S}$ we have*

$$\text{Stab}_G(v) \cap \mathcal{N} = \langle \Gamma_w \cap \text{Stab}_G(v) \rangle_{w \in \mathbb{Y}_*} = \langle \Gamma_w \rangle_{w \in \mathbb{Y}_* - (\text{Act}(v) - v)}.$$

Proof. One can follow the proof of [DHS21, Proposition 4.8], which only uses the existence of shortening pairs (which is Lemma 3.12) and [DHS21, Corollary 1.8] (which is our Corollary 3.10). \square

We conclude with two lemmas allowing us to inject certain finite configurations in the quotient. First, a consequence of Lemma 3.18:

Corollary 3.24. *For any two distinct vertices $v, w \in \mathcal{S}^{(0)}$ there exists a constant $\Theta_{rot}(v, w)$ such that, if the rotation control Θ_{rot} is greater than $\Theta_{rot}(v, w)$, then $\bar{v} \neq \bar{w}$. As a consequence, for every finite set $\mathcal{W} \subseteq \mathcal{S}^{(0)}$, there exists a constant $\Theta_{rot}(\mathcal{W})$ such that, if the rotation control Θ_{rot} is greater than $\Theta_{rot}(\mathcal{W})$, then \mathcal{W} injects inside $\mathcal{S}^{(0)}/\mathcal{N}$.*

In the same spirit, we show that we can inject finite subsets of G in the quotient group:

Lemma 3.25. *For every $g \in G - \{1\}$ there exists a constant $\Theta_{\text{rot}}(g)$ such that, if the rotation control Θ_{rot} is greater than $\Theta_{\text{rot}}(g)$, then $g \notin \mathcal{N}$.*

As a consequence, for every finite subset $F \subset G$ there exists a constant $\Theta_{\text{rot}}(F)$ such that, if $\Theta_{\text{rot}} \geq \Theta_{\text{rot}}(F)$, then F injects in G/\mathcal{N} .

Proof. Take an element $y_i \in \mathbb{Y}_i$ for every colour $i = 1, \dots, m$. By uniform properness, there exists a constant M such that

$$\max_{i=1, \dots, m} \sup_{s \in \text{Act}(y_i) \cap \text{Act}(gy_i)} d_s^*(y_i, gy_i) \leq M.$$

Choose $\Theta_{\text{rot}}(g)$ such that $\Theta_{\text{rot}}(g)/2 - \aleph > M$, where \aleph is the constant from Lemma 3.11 which only depends on \mathbb{Y}_* . Now let \mathcal{N} have rotation control greater than $\Theta_{\text{rot}}(g)$. If by contradiction $g \in \mathcal{N}$, then we can define its index $i(g) \in \{1, \dots, m\}$. By Lemma 3.11 we can then find $s \in \mathbb{Y}_*$ such that

$$M < \Theta_{\text{rot}}/2 - \aleph \leq d_s^*(y_{i(g)}, gy_{i(g)}) \leq M,$$

and this yields a contradiction. \square

3.2. From a short HHG to a SCPG. In this Subsection we show that short HHG fit the framework of strong composite projection graphs. We break the argument into a series of Propositions.

Let $(G, \overline{X}, \mathcal{W})$ be a colourable short HHG, with HHS constant E . For every $v \in \overline{X}^{(0)}$ set $\text{Act}(v) = \overline{X}^{(0)} - \text{Lk}_{\overline{X}}(v)$, which contains all vertices with the same colour as v since no two \overline{X} -adjacent vertices have the same colour. Furthermore, set $d_v^\pi(x, y) = \text{diam}_{\ell_v}(\rho_v^x \cup \rho_v^y)$ whenever the quantity is defined.

Proposition 3.26. *There exists $\theta \geq 0$ such that $(\overline{X}^{(0)}, \theta, d_*^\pi)$ is a uniformly proper composite projection system.*

Proof. The proof is similar to that of [BHMS24, Lemma 8.21]. Symmetry and the triangle inequality are clear, since d_v^π is defined as a diameter of a union. The Behrstock inequality is implied by the existence of a HHG structure, as it follows by combining the partial realisation axiom [BHS19, Definition 1.1.(8)] with the first consistency inequality [BHS19, Definition 1.1.(4)]. Uniform properness follows from the Distance Formula [BHS19, Theorem 4.5], again together with [BHS19, Definition 1.1.(8)]. Separation corresponds to the fact that, for every two non-adjacent $v, w \in \overline{X}^{(0)}$, $\text{diam} \rho_v^w \leq E$.

Symmetry in action is also clear, since $x \in \text{Act}(y)$ if and only if $x \notin \text{Lk}_{\overline{X}}(y)$. Moving to closeness in inaction, if $x, y, z \in \overline{X}^{(0)}$ are such that x and y are \overline{X} -adjacent while z is disjoint from both, then $\ell_x \perp \ell_y$ and both are transverse to ℓ_z . Then $d_z^\pi(x, y)$ is uniformly bounded by [DHS17, Lemma 1.5].

We finally prove finite filling. Let $\mathcal{Z} \subseteq \overline{X}^{(0)}$. By definition, $y \in \bigcup_{x \in \mathcal{Z}} \text{Act}(x)$ if and only if $y \notin \bigcap_{x \in \mathcal{Z}} \text{Lk}_{\overline{X}}(x)$, so it is enough to show that there does not exist an infinite sequence $\{x_i\}_{i \in \mathbb{N}}$ in $\overline{X}^{(0)}$ such that

$$\bigcap_{i \leq k} \text{Lk}_{\overline{X}}(x_i) \subsetneq \bigcap_{i \leq k+1} \text{Lk}_{\overline{X}}(x_i)$$

for all k . This is because, if x_1 and x_2 are \overline{X} -adjacent, then $\text{Lk}_{\overline{X}}(x_1) \cap \text{Lk}_{\overline{X}}(x_2) = \emptyset$, since \overline{X} is triangle-free; if instead x_1 and x_2 are not \overline{X} -adjacent, then $\text{Lk}_{\overline{X}}(x_1) \cap \text{Lk}_{\overline{X}}(x_2)$ consists of at most one vertex, since \overline{X} is square-free. \square

Now, G acts by isomorphisms on this composite projection system, because it acts by isomorphisms on \overline{X} , the action permutes the colours, and projections in the HHG structure are G -equivariant.

As in Notation 2.18, let $\mathcal{B} = \{s_1, \dots, s_r\}$ be a base, such that every Z_{v_i} is infinite. Choose non-trivial integers M_1, \dots, M_r , and set $\Gamma_{g v_i} = g(M_i Z_{v_i})g^{-1}$ for all $g \in G$ and $i = 1, \dots, r$. Set $\mathbb{Y}_* = G\mathcal{B}$, with the induced set of colours.

Proposition 3.27. *For every $\Theta_{rot} > 0$ there exists $M \in \mathbb{N} - \{0\}$ such that, if every M_i is a non-trivial multiple of M , then $\{\Gamma_v\}_{v \in \mathbb{Y}_*}$ is a composite rotating family on (\mathbb{Y}_*, G) with rotation control Θ_{rot} .*

With the notation from Definition 2.19, the above Proposition states that every rotation control can be achieved if \mathcal{N} is deep enough.

Proof. By construction, every Γ_v is infinite, and it acts trivially on the star of v in \overline{X} by Axiom B. The collection $\{\Gamma_v\}$ is equivariant by construction, and if M is even then Γ_x commutes with Γ_y whenever $x \in \text{Lk}_{\overline{X}}(y)$. Proper isotropy, as well as the last axiom, both follow from the fact that Z_v acts metrically properly and coboundedly on the quasiline $\mathcal{C}\ell_v$. \square

Finally, let $\mathcal{S} = \overline{X}^{+\mathcal{W}}$ be the augmented support graph, which is hyperbolic as it is G -equivariantly quasi-isometric to the main coordinate space \mathcal{CS} (see Remark 2.6). Notice that, by Lemma 2.8, every $\overline{X}^{+\mathcal{W}}$ -geodesic with large projection to some ℓ_v contains a vertex in the \overline{X} -star of v ; hence the strong Bounded Geodesic Image property from Definition 3.6 holds in this setting. Summing up, we proved the following:

Proposition 3.28. *For every $\Theta_{rot} > 0$ there exists $M \in \mathbb{N}$ such that, if every M_i is a non-trivial multiple of M , the collection $(\overline{X}^{+\mathcal{W}}, G, G\mathcal{B}, \{\Gamma_v\}_{v \in G\mathcal{B}})$ is a strong composite projection graph with rotation control Θ_{rot} .*

Remark 3.29. Notice that we could not have used the original definition from [DHS21]. Indeed, in our setting the hyperbolic graph is $\overline{X}^{+\mathcal{W}}$ (not \overline{X}); furthermore, $\text{Lk}_{\overline{X}^{+\mathcal{W}}}(s)$ contains $\text{Lk}_{\overline{X}}(s)$ but they might not coincide, and this means that there might be vertices of $\text{Lk}_{\overline{X}^{+\mathcal{W}}}(s)$ which are *not* fixed by Γ_s .

Remark 3.30 (SCPG for augmented links). For every $v \in \overline{X}^{(0)}$ set $\mathbb{Y}_*^v = G\mathcal{B} \cap \text{Lk}_{\overline{X}}(v)$. Arguing as above, one can check that the collection

$$\left(\text{Lk}_{\overline{X}}(v)^{+\mathcal{W}}, \text{Stab}_G(v), \mathbb{Y}_*^v, \{\Gamma_w\}_{w \in \mathbb{Y}_*^v} \right),$$

defines a strong composite projection graph, whose rotation control can be made arbitrarily large (for all vertices at once) by choosing \mathcal{N} is deep enough. We stress that, since any two vertices $w, w' \in \mathbb{Y}_*^v$ are always disjoint in $\text{Lk}_{\overline{X}}(v)$, one has that $\text{Act}(w) = \mathbb{Y}_*^v$ for every $w \in \mathbb{Y}_*^v$. This is relevant as then, by Lemma 3.22, the subgroup $\langle \Gamma_w \rangle_{w \in \mathbb{Y}_*^v}$, which we shall later denote by \mathcal{N}_v , is a free product of (some of) the Γ_w .

3.3. Summary of the consequences for short HHG. We gather here all implications of the above discussion to our setting. Recall that we have a colourable short HHG $(G, \overline{X}, \mathcal{W})$, with HHS constant E , and we are considering the quotient by a subgroup \mathcal{N} , as in Notation 2.18.

3.3.1. The quotient support graph. Let (Λ, N) be either $(\overline{X}^{+\mathcal{W}}, \mathcal{N})$ or $(\text{Lk}_{\overline{X}}(v)^{+\mathcal{W}}, \mathcal{N}_v)$ for some $v \in \overline{X}^{(0)}$. Define Λ/N as the graph whose vertices and edges are N -orbits of vertices and edges of Λ , and let $q: \Lambda \rightarrow \Lambda/N$ be the quotient projection. We shall denote the N -orbit of a vertex $v \in \Lambda$ by $[v]$.

Corollary 3.31 (of Lemma 3.13). *Λ/N is simplicial.*

Moreover, a plethora of subgraphs of Λ/N lift isometrically to Λ . For example:

Corollary 3.32 (of Lemma 3.14). *For each combinatorial path $\overline{\gamma}$ in Λ/N starting at $[v]$, and any point v in the preimage of $[v]$ (henceforth: a lift of $[v]$), there exists a combinatorial path in Λ so that $q \circ \gamma = \overline{\gamma}$, which we call a lift of $\overline{\gamma}$. Moreover, if $\overline{\gamma}$ is a geodesic, then so is γ .*

Furthermore, we can also lift *geodesic k -gons*:

Definition 3.33. A k -gon in a metric graph Λ is a closed combinatorial path Q made of k geodesic segments. In other words,

$$Q = \bigcup_{i=0}^{k-1} [x_i, x_{i+1}],$$

where $x_0, \dots, x_k \in \Lambda^{(0)}$, $x_k = x_0$ and each $[x_i, x_{i+1}]$ is a geodesic path.

Corollary 3.34 (of Lemma 3.16). *For every $k \in \mathbb{N}$ the following holds if N is deep enough. For every k -gon \overline{Q} inside Λ/N there exists a k -gon Q inside Λ such that $q(Q) = \overline{Q}$. We say that Q is a lift of \overline{Q} .*

Now we focus our attention on $\Lambda = \overline{X}^{+\mathcal{W}}$. Recall that \overline{X} is a (non-full) G -invariant subgraph of $\overline{X}^{+\mathcal{W}}$, as pointed out in Remark 2.6, so the quotient projection restricts to a map $\overline{X} \rightarrow \overline{X}/\mathcal{N}$, which we shall still call q . As \overline{X}/\mathcal{N} is then a G/\mathcal{N} -invariant subgraph of $\overline{X}^{+\mathcal{W}}/\mathcal{N}$, it is itself simplicial; furthermore, Corollary 3.32 implies that, for every $k \in \mathbb{N}$, we can find \mathcal{N} deep enough that every closed combinatorial path $\overline{\gamma} \subset \overline{X}/\mathcal{N}$ of length at most k has a lift *inside* \overline{X} . Then we summarise the properties of \overline{X}/\mathcal{N} below:

Lemma 3.35. *If \mathcal{N} is deep enough, then:*

- \overline{X}/\mathcal{N} is a triangle- and square-free simplicial graph, and none of its connected components is a point.
- For every \mathcal{N} -orbit $[v]$ of a vertex $v \in \overline{X}^{(0)}$, $\text{Lk}_{\overline{X}/\mathcal{N}}([v]) = q(\text{Lk}_{\overline{X}}(v))$.

Proof. We already noticed that \overline{X}/\mathcal{N} is simplicial. Furthermore, any non-degenerate triangle or square in \overline{X}/\mathcal{N} would lift to \overline{X} if \mathcal{N} is deep enough (notice that the lift would again be non-degenerate, as its edges would have different quotient projections), and every vertex of \overline{X}/\mathcal{N} belongs to at least one edge, because this is true in \overline{X} . Moving to the second bullet, Corollary 3.32 implies that every edge $\{[v], [w]\}$ lifts to an edge $\{v, w\}$. \square

3.3.2. Large rotations stabilising a vertex. Given $v \in \overline{X}^{(0)}$, we have the following description of the elements of \mathcal{N} which fix v :

Corollary 3.36 (of Lemma 3.23). *If \mathcal{N} is deep enough, then for every $v \in \overline{X}^{(0)}$ we have that $\mathcal{N} \cap \text{Stab}_G(v) = \langle \Gamma_w \rangle_{w \in \text{Star}_{\overline{X}}(v) \cap G\mathcal{B}}$.*

As Γ_v acts trivially on $\text{Lk}_{\overline{X}}(v)$, we get that

$$\text{Lk}_{\overline{X}}(v) / \langle \Gamma_w \rangle_{w \in \text{Star}_{\overline{X}}(v) \cap G\mathcal{B}} \text{Lk}_{\overline{X}}(v) / \mathcal{N}_v,$$

where

$$\mathcal{N}_v := \langle \Gamma_w \rangle_{w \in \text{Lk}_{\overline{X}}(v) \cap G\mathcal{B}}.$$

Moreover, no two elements of $\text{Lk}_{\overline{X}}(v)$ are adjacent, as \overline{X} is triangle-free, so we get:

Corollary 3.37 (of Lemma 3.22). *If \mathcal{N} is deep enough, then for any $v \in \overline{X}^{(0)}$ there exists $\mathcal{J} \subseteq \text{Lk}_{\overline{X}}(v)^{(0)}$ such that \mathcal{N}_v has a free presentation*

$$\mathcal{N}_v = \bigast_{w \in \mathcal{J}} \Gamma_w.$$

3.3.3. Uniform hyperbolicity of the quotients. By choosing \mathcal{N} to be deep enough, we can ensure that the quotient of the main coordinate space of (X, \mathcal{W}) remains hyperbolic, with the same hyperbolicity constant (which we bound by the HHS constant E):

Corollary 3.38 (of Theorem 3.19, global version). *If \mathcal{N} is deep enough, then $\overline{X}^{+\mathcal{W}}/\mathcal{N}$ is E -hyperbolic. Furthermore, G/\mathcal{N} acts acylindrically on $\overline{X}^{+\mathcal{W}}/\mathcal{N}$, and if G acts non-elementarily on $\overline{X}^{+\mathcal{W}}$ then so does G/\mathcal{N} on $\overline{X}^{+\mathcal{W}}/\mathcal{N}$.*

By Remark 3.30, we can also establish a “local version” of the above result:

Corollary 3.39 (of Theorem 3.19, local version). *If \mathcal{N} is deep enough, then for every $v \in \overline{X}^{(0)}$ we have that $\text{Lk}_{\overline{X}}(v)^{+\mathcal{W}}/\mathcal{N}_v$ is E -hyperbolic.*

3.3.4. Preserving finite data in the quotient.

Corollary 3.40 (of Lemmas 3.24 and 3.25). *Let $\{w_1, \dots, w_k\} \subset \overline{X}^{(0)}$ be a finite collection of vertices, and let $F \subset G - \{1\}$ be a finite subset. If \mathcal{N} is deep enough then*

- $\{w_1, \dots, w_k\}$ injects in the quotient $\overline{X}^{(0)}/\mathcal{N}$;
- $F \cap \mathcal{N} = \emptyset$, so that F injects in the quotient G/\mathcal{N} .

Combining the Corollary with the description of \mathcal{N}_v , we get the following characterisation of which cyclic directions survive in the quotient:

Lemma 3.41. *The following holds whenever \mathcal{N} is deep enough. For any $v \in \overline{X}^{(0)} - G\mathcal{B}$, we have that $\mathcal{N} \cap Z_v = \{0\}$.*

Proof. First, let $F = Z_{x_1} \cup \dots \cup Z_{x_l}$, where x_1, \dots, x_l are representatives of the G -orbits of vertices with finite cyclic directions. By Corollary 3.40, we can choose \mathcal{N} deep enough so that it does not contain any conjugates of elements of F , so that \mathcal{N} intersects trivially every finite cyclic direction.

Thus, it remains to show that $\mathcal{N} \cap Z_v = \{0\}$ whenever $v \notin G\mathcal{B}$ and Z_v is infinite. Notice first that $\mathcal{N} \cap Z_v \leq \mathcal{N} \cap \text{Stab}_G(v) = \mathcal{N}_v$, as v does not belong to $G\mathcal{B}$. Since we assumed \mathcal{N} to be deep enough to satisfy Remark 2.20, we have that Z_v commutes with Γ_w for every $w \in \text{Lk}_{\overline{X}}(v)$, so $\mathcal{N} \cap Z_v$ must lie in the centre of \mathcal{N}_v . Then Corollary 3.37 tells us that \mathcal{N}_v is a free product of (some) Γ_w s, and in particular it has non-trivial centre if and only if $\mathcal{N}_v = \Gamma_w$ for some $w \in \text{Lk}_{\overline{X}}(v)$. In this case, it suffices to notice that $Z_v \cap \Gamma_w \leq Z_v \cap Z_w$ must be trivial, since Z_v acts geometrically on the quasiline $\mathcal{C}\ell_v$ while Z_w acts on it with uniformly bounded orbits. \square

3.3.5. *The quotient extensions.* Recall that, for every $v \in \overline{X}^{(0)}$, we defined $[v] \in \overline{X}^{(0)}/\mathcal{N}$ as its \mathcal{N} -orbit. Let $Z_{[v]} = Z_v/(Z_v \cap \mathcal{N})$ and $H_{[v]} := H_v/\mathfrak{p}_v(\mathcal{N}_v)$.

Lemma 3.42. *For every $v \in \overline{X}^{(0)}$ there is a commutative diagram of group extensions, where the vertical arrows are the restrictions of the quotient projection $G \rightarrow G/\mathcal{N}$:*

$$\begin{array}{ccccccccc} 0 & \longrightarrow & Z_v & \longrightarrow & \text{Stab}_G(v) & \xrightarrow{\mathfrak{p}_v} & H_v & \longrightarrow & 1 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & Z_{[v]} & \longrightarrow & \text{Stab}_G([v]) & \xrightarrow{\mathfrak{p}_{[v]}} & H_{[v]} & \longrightarrow & 0. \end{array}$$

Consequently, $Z_{[v]}$ is a cyclic, normal subgroup of $\text{Stab}_G([v])$ acting trivially on $\text{Lk}_{\overline{X}/\mathcal{N}}([v])$, and the collection of quotient extensions is equivariant with respect to the G/\mathcal{N} -action by conjugation.

Proof. It is easy to see that $\text{Stab}_G([v])$ is the quotient projection of $\text{Stab}_G(v)$. Furthermore, define a map $\mathfrak{p}_{[v]}: \text{Stab}_G([v]) \rightarrow H_{[v]}$ by sending the coset $g(\mathcal{N} \cap \text{Stab}_G(v))$ to the coset $\mathfrak{p}_v(g)\mathfrak{p}_v(\mathcal{N}_v)$, for every $g \in \text{Stab}_G(v)$. This map is well-defined, as

$$\mathfrak{p}_v(\mathcal{N} \cap \text{Stab}_G(v)) = \mathfrak{p}_v(\langle \Gamma_w \rangle_{w \in \{\text{Star}_{\overline{X}}(v)\} \cap G\mathcal{B}}) = \mathfrak{p}_v(\mathcal{N}_v),$$

and it is a group homomorphism with kernel $Z_{[v]}$. Then one can easily see that the above diagram commutes, using that both $\mathcal{N} \cap \text{Stab}_G(v)$ and Z_v are normal in $\text{Stab}_G(v)$. Furthermore, by Lemma 3.35, $\text{Lk}_{\overline{X}/\mathcal{N}}([v]) = q(\text{Lk}_{\overline{X}}(v))$, so $Z_{[v]}$ acts trivially on $\text{Lk}_{\overline{X}/\mathcal{N}}([v])$. The remaining properties of the quotient extension follow from the corresponding features of the top row of the diagram. \square

Next, we show that the quotient $H_{[v]}$ is hyperbolic:

Lemma 3.43. *Let $v \in \overline{X}^{(0)}$, and fix a collection $[W]$ of $\text{Stab}_G([v])$ -orbit representatives of vertices in $\text{Lk}_{\overline{X}/\mathcal{N}}([v])$ with infinite cyclic direction.*

If \mathcal{N} is deep enough, then:

- (1) $H_{[v]}$ is hyperbolic;
- (2) The collection $\{\mathfrak{p}_{[v]}(Z_{[w]})\}_{[w] \in [W]}$ is independent in $H_{[v]}$, that is, for every distinct $[w], [w'] \in [W]$ and every $\bar{h} \in H_{[v]}$, $\mathfrak{p}_{[v]}(Z_{[w]})^{\bar{h}} \cap \mathfrak{p}_{[v]}(Z_{[w']}) = \{1\}$;
- (3) $\mathfrak{p}_{[v]}(Z_{[w']})$ is quasiconvex in $H_{[v]}$ for every $[w'] \in \text{Lk}_{\overline{X}/\mathcal{N}}([v])$.

Proof. Recall from Lemma 2.11 that H_v is hyperbolic relative to the collection $\{K_{H_v}(\mathfrak{p}_v(Z_w))\}_{w \in W}$, for any collection W of $\text{Stab}_G(v)$ -orbit representatives of vertices in $\text{Lk}_{\overline{X}}(v)$ with infinite cyclic direction. We can choose W such that its quotient projection contains $[W]$. Now, since $\mathfrak{p}_v(\mathcal{N}_v) = \langle \mathfrak{p}_v(\Gamma_w) \rangle_{w \in \text{Lk}_{\overline{X}}(v) \cap G\mathcal{B}}$, by the relative Dehn Filling Theorem [Osi07, Theorem 1.1] there exists a finite set $F_v \subset H_v - \{1\}$ such that, if $\mathfrak{p}_v(\Gamma_w) \cap F_v = \emptyset$ for every $w \in W$, then $H_{[v]}$ is hyperbolic relative to

$$\{Q_{[w]} := K_{H_v}(\mathfrak{p}_v(Z_w))/\mathfrak{p}_v(Z_w \cap \mathcal{N})\}_{w \in W}.$$

By choosing \mathcal{N} deep enough, we can ensure that $\mathfrak{p}_v(\Gamma_w)$ avoids F_v for every $w \in W$. Indeed, $\mathfrak{p}_v(\Gamma_w) = M_i \mathfrak{p}_v(Z_w)$, where M_i is one of the integers from Notation 2.18. Since both F_v and W are finite, we can find $M \in \mathbb{N} - \{0\}$ such that $M \mathfrak{p}_v(Z_w) \cap F_v =$

\emptyset for every $w \in W$, and then choose each M_i to be a multiple of M . In turn, since there are finitely many G -orbits of vertices in \overline{X} , if \mathcal{N} is deep enough then $H_{[v]}$ admits the above relative hyperbolic structure for every $v \in \overline{X}^{(0)}$.

Now, since peripherals are quotients of virtually cyclic subgroups, $H_{[v]}$ is hyperbolic. Moreover, for every $w \in W$ such that $[w] \in [W]$, we have that $\mathfrak{p}_{[v]}(Z_{[w]}) \leq Q_{[w]}$, so Item (2) in the statement follows from the fact that conjugates of different peripherals have finite intersection in a relatively hyperbolic group. Finally, every infinite $\mathfrak{p}_{[v]}(Z_{[w']})$ is conjugate to $\mathfrak{p}_{[v]}(Z_{[w]})$ for some $[w] \in [W]$, and the latter is a finite-index subgroup of $Q_{[w]}$. Item (3) then comes from the fact that conjugates of peripherals in a relatively hyperbolic group are quasiconvex. \square

3.3.6. Absence of hidden symmetries. For later use, we also point out that cyclic directions in the kernel have no *hidden symmetries*, in the following sense.

Definition 3.44 (No hidden symmetries). Let $0 \rightarrow \mathbb{Z} \rightarrow G \xrightarrow{\pi} H \rightarrow 1$ be a group extension, with H hyperbolic, and let $C \leq G$ be a cyclic subgroup. We say that C has *no hidden symmetries* if $\pi(C)$ is infinite and C contains a finite-index subgroup which is normal in $\pi^{-1}(K_H(\pi(C)))$. We say that an element $g \in G$ has no hidden symmetries if $\langle g \rangle$ has no hidden symmetries.

Lemma 3.45. *The following holds if \mathcal{N} is deep enough. For any two adjacent vertices $v, w \in \overline{X}^{(0)} - G\mathcal{B}$ with infinite cyclic directions, $Z_{[w]}$ has no hidden symmetries in $\text{Stab}_{G/\mathcal{N}}([v])$, in the sense of Definition 3.44.*

Proof. Let D_0 be such that $\mathcal{N} = \langle \langle M_x Z_x \rangle \rangle_{x \in \mathcal{B}}$ satisfies all properties from Subsection 3.1 whenever each M_x is a non-trivial multiple of D_0 . Furthermore, by Corollary 3.37 applied to the composite rotating family on the whole $\overline{X}^{(0)}$, there exists D_1 such that, for every $v \in \overline{X}^{(0)}$, $R_v := \langle M_y Z_y \rangle_{y \in \text{Lk}_{\overline{X}}(v)}$ is a (possibly infinitely generated) free group whenever each M_y is a non-trivial multiple of D_1 (notice that R_v does not necessarily coincide with \mathcal{N}_v , as \mathcal{B} might not contain a representative for every G -orbit). Then let $\mathcal{N} = \langle \langle M_x Z_x \rangle \rangle_{x \in \mathcal{B}}$, where each M_x is a multiple of $D_0 D_1$.

Now let $v, w \in \overline{X}^{(0)}$ be adjacent vertices, and let $\bar{g} \in \mathfrak{p}_{[v]}^{-1}(K_{H_{[v]}}(\mathfrak{p}_{[v]}(Z_{[w]})))$. We want to show that \bar{g} normalises a finite-index subgroup of $Z_{[w]}$. Since \bar{g} projects to $K_{H_{[v]}}(\mathfrak{p}_{[v]}(Z_{[w]}))$ in $H_{[v]}$, it must normalise a finite-index subgroup of $\mathfrak{p}_{[v]}(Z_{[w]})$, so there exists $m \in \mathbb{N} - \{0\}$, $\varepsilon \in \{\pm 1\}$, and $k \in \mathbb{Z}$ such that

$$z_{\bar{g}[w]}^{mD_1} = \bar{g} z_{[w]}^{mD_1} \bar{g}^{-1} = z_{[w]}^{\varepsilon m D_1} z_{[v]}^k.$$

If $k = 0$ we are done, so suppose otherwise. Let $g \in \text{Stab}_G(v)$ project to \bar{g} . By the above equality, there exists $n \in \mathcal{N}$ such that

$$z_{gw}^{mD_1} = z_w^{\varepsilon m D_1} z_v^k n.$$

In particular $n \in \mathcal{N} \cap \text{Stab}_G(v) = \mathcal{N}_v$. We can rewrite the above equality as

$$z_v^k = z_w^{-\varepsilon 2mD_1} z_{gw}^{2mD_1} n^{-1}.$$

By our choice of how deep \mathcal{N} is, the right hand side belongs to R_v , which is a free group. Now, z_v^2 commutes with z_x for every $x \in \text{Lk}_{\overline{X}}(v)$, so $\langle z_v^{2k} \rangle$, which is an infinite cyclic group because $k \neq 0$, must belong to the centre of R_v . The latter

is trivial unless $R_v \cong \mathbb{Z}$, which happens exactly when $\text{Lk}_{\overline{X}}(v)$ consists of a single vertex. In the latter case we would have that $w = gw$, and therefore

$$z_{[w]}^{mD_1} = z_{\overline{g}[w]}^{mD_1} = \overline{g} z_{[w]}^{mD_1} \overline{g}^{-1},$$

proving that \overline{g} normalises $\langle z_{[w]}^{mD_1} \rangle$. \square

4. A SHORT STRUCTURE FOR THE QUOTIENT

The main aim of this paper is to prove the following theorem, which roughly states that the class of short HHG is stable under taking quotients by large enough cyclic directions. Recall Notation 2.18, where we defined normal subgroups of the form $\mathcal{N} = \langle M_i Z_{s_i} \rangle$, that is, generated by cyclic subgroups of the cyclic directions of a short HHG.

Theorem 4.1. *Let $(G, \overline{X}, \mathcal{W})$ be a short HHG, as in Definition 2.1, and let \mathcal{N} be the normal subgroup as in Notation 2.18. If \mathcal{N} is deep enough, then $(G/\mathcal{N}, \overline{X}/\mathcal{N})$ is a short HHG, where the cyclic direction associated to each $[v] \in (\overline{X}/\mathcal{N})^{(0)}$ is (a finite-index subgroup of) $Z_{[v]} := Z_v/(Z_v \cap \mathcal{N})$. Furthermore, if $v \notin G\mathcal{B}$ then $Z_{[v]} \cong Z_v$.*

Outline of the proof. In Subsection 4.1 we modify the short HHG structure for G , in such a way that the kernel \mathcal{N} acts with uniformly bounded orbits on the quasilines. We then set $\hat{X} = X/\mathcal{N}$, which we prove to be the blowup of \overline{X}/\mathcal{N} in Lemma 4.3, and we consider the \hat{X} -graph $\widehat{\mathcal{W}}$, obtained from the (possibly non simplicial) graph \mathcal{W}/\mathcal{N} by removing loop edges and double edges. Then we prove that $(\hat{X}, \widehat{\mathcal{W}})$ is a combinatorial HHS. Here is where we verify each axiom from Definition 1.6:

- Axioms (1) and (3) both follow from the fact that \hat{X} is a blowup.
- Axiom (2) is split between Subsections 4.3.3 and 4.3.4.
- Axiom (4) is Lemma 4.11.

In Subsection 4.3.5 we notice that the G -action on X induces a G/\mathcal{N} -action on \hat{X} , with finitely many orbits of links, and a geometric action on $\widehat{\mathcal{W}}$. Therefore $(\hat{X}, \widehat{\mathcal{W}})$ is a combinatorial HHG structure for G/\mathcal{N} . Furthermore, in Lemma 4.20 we check that it is actually a short HHG structure, with the required cyclic directions. The “furthermore” part of the statement is simply Lemma 3.41. \square

4.1. Preparing the structure above. We first need to tweak the short HHG structure for G , to make it as “compatible” with \mathcal{N} as possible. This will later allow us to define a combinatorial HHG structure for G/\mathcal{N} by taking the quotient by \mathcal{N} of the refined structure.

First, fix a short HHG structure $(G, \overline{X}, \mathcal{W}_0)$, coming from the action on the combinatorial HHS (X_0, \mathcal{W}_0) . One can use the G -action on $\overline{X}^{+\mathcal{W}_0}$ to build a strong composite projection graph (SCPG), and let \mathcal{N}_0 be deep enough to satisfy all properties from Subsection 3.3 *with respect to this SCPG*. Later we shall choose a deeper subgroup $\mathcal{N} \leq \mathcal{N}_0$, so to avoid confusion we shall denote by $[v]_0 \in (\overline{X}/\mathcal{N}_0)^{(0)}$ the \mathcal{N}_0 -orbit of a vertex $v \in \overline{X}^{(0)}$.

Now let $v \in \overline{X}^{(0)} - G\mathcal{B}$ with infinite cyclic direction, so that $Z_{[v]_0} \cong Z_v \cong \mathbb{Z}$ by Lemma 3.41. By Lemma 3.43, the quotient $H_{[v]_0} = \text{Stab}_{G/\mathcal{N}_0}([v]_0)/Z_{[v]_0}$ is hyperbolic. Furthermore, if we fix a collection $[W]_0$ of $\text{Stab}_{G/\mathcal{N}_0}([v]_0)$ -orbits of vertices in $\text{Lk}_{\overline{X}/\mathcal{N}_0}([v]_0)$ with infinite cyclic directions, then the subgroups $\{Z_{[w]_0}\}_{[w]_0 \in [W]_0}$

have no hidden symmetries by Lemma 3.45, and they project to an independent collection in $H_{[w]_0}$ by Lemma 3.43. Hence, if we denote the centraliser of $Z_{[v]_0}$ in $\text{Stab}_{G/\mathcal{N}_0}([v]_0)$ by $E_{[v]_0}$, Corollary A.4 from the Appendix of [Man24] produces a homogeneous quasimorphism $\psi_{[v]_0} : E_{[v]_0} \rightarrow \mathbb{R}$ which is unbounded on $Z_{[v]_0}$ and is trivial on $Z_{[u]_0}$ for every $u \in \text{Lk}_{\overline{X}}(v)$.

Next, applying Proposition 2.16 to $(G, \overline{X}, \mathcal{W}_0)$ produces blowup materials, and in particular, for every $v \in \overline{X}^{(0)}$, a subgroup $E_v \leq \text{Stab}_G(v)$ centralising Z_v . Whenever $v \in \overline{X}^{(0)} - G\mathcal{B}$ has infinite cyclic direction, we can precompose the quasimorphism $\psi_{[v]_0}$ constructed above with the quotient projection $E_v \rightarrow E_{[v]_0}$. This gives a new homogeneous quasimorphism $\psi_v : E_v \rightarrow \mathbb{R}$, which is unbounded on $Z_v \cap E_v$ and trivial on $Z_w \cap E_v$ whenever $w \in \text{Lk}_{\overline{X}}(v)$. Most importantly, by construction ψ_v also vanishes on the intersection $\mathcal{N}_0 \cap E_v$. If we replace the quasimorphism ϕ_v associated to v by ψ_v , we get new blowup materials for G , and Theorem 2.15 then yields *another* short HHG structure (X, \mathcal{W}) on G , with the same support graph \overline{X} . Using the G -action on $\overline{X}^{+\mathcal{W}}$, one can build *another* strong composite projection graph, and let $\mathcal{N} \leq \mathcal{N}_0$ be deep enough to satisfy all properties of Section 2.4 *with respect to the new SCPG*.

Remark 4.2 (Motivation of the above construction). Earlier we mentioned that we wanted the short HHG structure to be “compatible” with the quotient. The improvement when passing from (X_0, \mathcal{W}_0) to (X, \mathcal{W}) is that, for every $v \in \overline{X}^{(0)}$ with infinite cyclic direction, $\mathcal{N} \cap E_v$ now acts with uniformly bounded orbits on the quasiline L_v associated to v , as its construction involved collapsing the “coarse level sets” of the quasimorphism ψ_v (see [Man24, Definition 3.12] for further details). This will be one of the key ingredients of the proof, in particular when we verify that augmented links in the quotient are hyperbolic (see Lemma 4.14 below).

4.2. The candidate combinatorial structure. From now on, let \mathcal{N} be deep enough to satisfy all properties of Section 2.4. We now move to the description of the combinatorial HHG structure for G/\mathcal{N} . We first recall from [Man24] that the underlying graph X of the short HHG structure for G , granted by Theorem 2.15, is a blowup of \overline{X} , where, given any $v_i \in V$ and any $g \in G$, the vertex $v = gv_i$ is blown up to the cone over $(L_v)^{(0)} = g \text{Stab}_G(v_i)$.

Now let X/\mathcal{N} be the graph whose vertices and edges are \mathcal{N} -orbits of vertices and edges in X . Notice that \overline{X}/\mathcal{N} is a full, simplicial subgraph of X/\mathcal{N} .

Lemma 4.3. *The graph X/\mathcal{N} is isomorphic to the blowup \widehat{X} of \overline{X}/\mathcal{N} with respect to the family $\{L_{[v]} := L_v/(\mathcal{N} \cap \text{Stab}_G(v))\}_{[v] \in \overline{X}/\mathcal{N}^{(0)}}$. In particular, \widehat{X} is simplicial.*

Proof. It suffices to notice that, for every $v \in \overline{X}^{(0)}$, $p \in (L_v)^{(0)}$, and $n \in \mathcal{N}$, np is only adjacent to nv . In other words, each \mathcal{N} -translate of p belongs to a single edge, and all these edges are in the same \mathcal{N} -orbit. Then by definition X/\mathcal{N} is the desired blowup. \square

Let $\widehat{p} : \widehat{X} \rightarrow \overline{X}/\mathcal{N}$ be the retraction mapping every cone to its apex. Given a simplex $\widehat{\Delta}$ of \widehat{X} , we shall call $\widehat{p}(\widehat{\Delta})$ its support. We also denote a maximal simplex of \widehat{X} by $\Delta([x], [y])$, where $[x] \in (L_{[w]})$, $[y] \in (L_{[w]})$, and $[v], [w]$ are \overline{X}/\mathcal{N} -adjacent.

Given a simplex Δ of X , we will denote its projection to \widehat{X} as $\widehat{\Delta}$, and we will say that Δ is a *lift* of $\widehat{\Delta}$. In particular, a lift of a maximal simplex $\Delta([x], [y])$ is of the form $\Delta(x, y)$, where $x \in [x]$ and $y \in [y]$. Conversely:

Lemma 4.4. *Let $\Delta(x, y)$ be a maximal simplex of X . If \mathcal{N} is deep enough, then $\widehat{\Delta} = \Delta([x], [y])$ is a maximal simplex of \widehat{X} .*

Proof. Let $v = p(x)$ and $w = p(y)$. As \overline{X} is a G -invariant subgraph of X , v and x cannot be in the same \mathcal{N} -orbit, and similarly for all other pairs of vertices of Δ where one is in \overline{X} and the other is not. Moreover, if x and y are in the same \mathcal{N} -orbit, then so are v and w . Thus, it suffices to exclude that v and w are in the same \mathcal{N} -orbit, which is true as they must have different colours, and we assumed \mathcal{N} to be deep enough to preserve each colour. \square

As every simplex of X can be completed to a maximal simplex, we get that:

Corollary 4.5. *Every simplex Δ of X injects inside \widehat{X} . In particular, if $a, b \in X^{(0)}$ are X -adjacent vertices, then their projections $[a], [b]$ are distinct.*

Definition 4.6 ($\widehat{\mathcal{W}}$ -edges). Let $\widehat{\mathcal{W}}$ be the \widehat{X} -graph where two maximal simplices $\Delta([x], [y]), \Delta([x'], [y'])$ are adjacent if and only if they admit lifts $\Delta(x, y), \Delta(x', y')$ which are adjacent in \mathcal{W} .

In other words, $\widehat{\mathcal{W}}$ is obtained from \mathcal{W}/\mathcal{N} after collapsing double edges and loops, in order to get a simplicial graph.

4.3. Checking the combinatorial HHG axioms. We now check that the pair $(\widehat{X}, \widehat{\mathcal{W}})$ is a combinatorial HHG structure for G . The recurring theme will be that one is often allowed to lift combinatorial configurations from $\widehat{X}^{+\widehat{\mathcal{W}}}$ to $X^{+\mathcal{W}}$. This way, all properties of $(\widehat{X}, \widehat{\mathcal{W}})$ can be deduced from the corresponding statements about (X, \mathcal{W}) , which we already know to be a combinatorial HHG structure for G .

For convenience, we recall the notion of a shortening pair, which will be a key ingredient for our lifting procedures:

Corollary 4.7 (of Lemma 3.12). *The following holds if \mathcal{N} is deep enough. Let (Λ, N) be either $(\overline{X}^{+\mathcal{W}}, \mathcal{N})$ or $(\text{Lk}_{\overline{X}}(v)^{+\mathcal{W}}, \mathcal{N}_v)$ for some $v \in \overline{X}^{(0)}$. There exists a well-order on N , called *complexity*, such that the minimum element is the identity 1. Moreover, for all $\gamma \in N - \{1\}$ and all $x \in \Lambda^{(0)}$, there exist a shortening pair (s, γ_s) (here $s \in G\mathcal{B} \cap \Lambda^{(0)}$ and $\gamma_s \in \Gamma_s$) so that $\gamma_s \gamma$ has strictly lower complexity than γ , and either*

- (1) *one between x and γx is fixed by Γ_s , or*
- (2) *$d_s(x, \gamma x) \geq 100E$.*

Remark 4.8 (Dependence on \mathcal{N}). From now on, we shall say that a quantity is *depth-resistant* if it does not depend on the choice of powers $\{M_1, \dots, M_k\}$ used to define \mathcal{N} , as in Notation 2.18, but only on the fact that each M_i is a multiple of a large enough integer (that is, the quantity is the same for every deep enough \mathcal{N}). If we could prove that all constants in the proofs below were depth-resistant, then the combinatorial structure $(\widehat{X}, \widehat{\mathcal{W}})$ for G/\mathcal{N} would be *uniformly* hierarchically hyperbolic, i.e. the HHS constant and the uniqueness function would be depth-resistant. This is not the case, but the only exception is that, whenever $w \in G\mathcal{B}$, the diameter of $L_{[w]}$ depends on the index of Γ_w inside Z_w (see Lemma 4.14).

4.3.1. *Finite chains, and intersection of links.* As \hat{X} is a blowup of a triangle- and square-free graph and none of its connected components is a single point, one can argue exactly as in [Man24, Subsection 3.3.2] to get the first and third requirements of Definition 1.6, with depth-resistant constants:

Corollary 4.9 (Verification of Definition 1.6.(1)). *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 25.*

Corollary 4.10 (Verification of Definition 1.6.(3)). *Let Σ, Δ be non-maximal simplices of X , and suppose that there exists a non-maximal simplex Γ such that $[\Gamma] \subseteq [\Sigma]$, $[\Gamma] \subseteq [\Delta]$ and $\text{diam}(\mathcal{C}(\Gamma)) \geq 3$. Then there exists a non-maximal simplex Π which extends Σ such that $[\Pi] \subseteq [\Delta]$ and all Γ as above satisfy $[\Gamma] \subseteq [\Pi]$.*

4.3.2. *Fullness of links.*

Lemma 4.11 (Verification of Definition 1.6.(4)). *Let $\hat{\Delta}$ be a non-maximal simplex of \hat{X} . Suppose that $[a], [b] \in \text{Lk}(\hat{\Delta})$ are distinct, non-adjacent vertices which are contained in $\widehat{\mathcal{W}}$ -adjacent maximal simplices $\hat{\Sigma}_a, \hat{\Sigma}_b$. Then there exist $\widehat{\mathcal{W}}$ -adjacent maximal simplices $\hat{\Pi}_a, \hat{\Pi}_b$ of \hat{X} such that $\hat{\Delta} \star [a] \subseteq \hat{\Pi}_a$ and $\hat{\Delta} \star [b] \subseteq \hat{\Pi}_b$.*

The following proof is prototypical of how to use Corollary 4.7, together with the strong bounded geodesic image Lemma 2.8, in order to lift combinatorial configurations from $\hat{X}^{+\mathcal{W}}$ to $X^{+\mathcal{W}}$.

Proof. Suppose first that $\hat{p}([a]) = \hat{p}([b]) = [v]$. Then $[a], [b] \in (L_{[v]})^{(0)}$, as they are not \hat{X} -adjacent. Let $a, b \in X$ be \mathcal{W} -adjacent lifts of $[a]$ and $[b]$; moreover, let $v = p(a)$ and $v' = p(b)$, which are \mathcal{W} -adjacent as well and in the same \mathcal{N} -orbit. By Corollary 3.31, $\overline{X}^{+\mathcal{W}}/\mathcal{N}$ is simplicial, so v must be equal to v' or $\overline{X}^{+\mathcal{W}}/\mathcal{N}$ would have an edge with the same endpoints. This means that a and b belong to the same $(L_v)^{(0)}$. Let Δ be a lift of $\hat{\Delta}$ inside $\text{Lk}_X(a) = \text{Lk}_X(b)$. Then, as (X, \mathcal{W}) is a combinatorial HHS, there exist \mathcal{W} -adjacent maximal simplices Π_a, Π_b such that $\Delta \star a \subseteq \Pi_a$ and $\Delta \star b \subseteq \Pi_b$. Thus, the required simplices $\hat{\Pi}_a$ and $\hat{\Pi}_b$ are the quotient projections of Π_a and Π_b .

Thus suppose that $\hat{p}([a]) = [w]$ and $\hat{p}([b]) = [w']$ are different. In particular $[w]$ and $[w']$ are not \overline{X}/\mathcal{N} -adjacent, or $[a]$ and $[b]$ would be joined by an edge of \hat{X} . This forces the support of $\hat{\Delta}$ to be a single vertex $[v]$, which is \overline{X}/\mathcal{N} -adjacent to both $[w]$ and $[w']$. Let $[y] = \Sigma_a \cap (L_{[w]})^{(0)}$, so that $[a]$ is either $[y]$ or $[w]$, and $[y'] = \Sigma_b \cap (L_{[w']})^{(0)}$. Now take lifts y of $[y]$ and y' of $[y']$ which are \mathcal{W} -adjacent. Let $w = p(y)$, $w' = p(y')$, and let $v \in \text{Lk}_{\overline{X}}(w')$ be a lift of $[v]$. There exists $n \in \mathcal{N}$ such that $nw \in \text{Lk}_{\overline{X}}(v)$. Hence the situation in $X^{+\mathcal{W}}$ is as in Figure (3) below.

Our goal is to show that there exist \mathcal{W} -adjacent lifts of $[y]$ and $[y']$ which are also X -adjacent to some lift of $[v]$. Then we will lift $\hat{\Delta}$ to some Δ supported on v , and we will conclude as above that $[y]$ and $[y']$ belong to $\widehat{\mathcal{W}}$ -adjacent maximal simplices containing $\hat{\Delta}$.

By Corollary 4.7, \mathcal{N} is equipped with a well-order, called complexity, whose minimum element is the identity, so we proceed by induction on the complexity of n . If $n = 1$ then $w = nw$, and both y and y' are already X -adjacent to v . Otherwise, let (s, γ_s) be a shortening pair, as in Corollary 4.7. Using γ_s , we want to replace some lifts, without breaking the configuration from Figure (3), in such a way that the two new lifts of $[w]$ are w and $\gamma_s nw$. Then we shall conclude by induction, as




FIGURE 3. The full lines represent X -edges, while the dashed lines represent \mathcal{W} -edges.

$\gamma_s n$ has strictly lower complexity than n by the defining properties of a shortening pair.

There are some cases to consider, depending on how γ_s acts on our configuration.

- If γ_s fixes w , then we apply γ_s to all lifts. As γ_s acts by isometries on $X^{+\mathcal{W}}$, the configuration from Figure (3) is preserved; moreover, every lift is mapped to a lift of the same point, as $\gamma_s \in \mathcal{N}$. Now $\gamma_s nw$ differs from $\gamma_s w = w$ by $\gamma_s n$.
- If γ_s fixes nw then $nw = \gamma_s nw$ already differs from w by $\gamma_s n$. Then we leave the configuration untouched.
- Otherwise, Corollary 4.7 implies that $d_s(w, nw) \geq 100E$. If neither v nor w' belonged to $\text{Star}_{\bar{X}}(s)$, then $d_s(w, w')$, $d_s(w', v)$, and $d_s(v, nw)$ would all be well-defined, and by triangle inequality at least one of them would be greater than $33E$. Without loss of generality, say $d_s(w, w') \geq 33E$. However, this would contradict the strong bounded geodesic image Lemma 2.8, because the edge $\{w, w'\}$ of $\bar{X}^{+\mathcal{W}}$ would be a geodesic disjoint from $\text{Star}_{\bar{X}}(s)$.
Thus, suppose first that γ_s fixes v . If we apply γ_s to nw then $v = \gamma_s v$ is again \bar{X} -adjacent to $\gamma_s nw$, so we can replace the lifts without breaking the configuration.
If instead γ_s fixes w' , we apply γ_s to both v and nw . This way $\gamma_s v$ is still X -adjacent to w' , and therefore to y' .

The proof of Lemma 4.11 is now complete. \square

4.3.3. Hyperbolicity of augmented links. Let $\hat{\Delta}$ be a simplex of \hat{X} . We want to show that $\text{Lk}(\hat{\Delta})^{+\mathcal{W}}$ is uniformly hyperbolic, in order to verify the first half of Definition 1.6.(2). As \hat{X} is a blowup graph, we only need to focus on the cases when $\text{Lk}(\hat{\Delta})$ is unbounded, listed in Lemma 2.2.

Firstly, $\text{Lk}(\emptyset)^{+\hat{\mathcal{W}}}$ retracts onto $(\bar{X}/\mathcal{N})^{+\hat{\mathcal{W}}}$, which coincides with $\bar{X}^{+\mathcal{W}}/\mathcal{N}$ as, by construction, two vertices $[v], [w] \in (\bar{X}/\mathcal{N})^{(0)}$ are $\hat{\mathcal{W}}$ -adjacent if and only if they have \mathcal{W} -adjacent lifts $v, w \in \bar{X}^{(0)}$. Then Corollary 3.38 tells us that $\bar{X}^{+\mathcal{W}}/\mathcal{N}$ is E -hyperbolic. Thus we get:

Lemma 4.12. $\text{Lk}(\emptyset)^{+\hat{\mathcal{W}}}$ is hyperbolic, and the hyperbolicity constant is depth-resistant.

We can argue similarly if $\hat{\Delta} = \{([v], [x])\}$ is of edge-type. Indeed, $\text{Lk}(\hat{\Delta})^{+\hat{\mathcal{W}}}$ retracts onto $(\text{Lk}_{\bar{X}/\mathcal{N}}([v]))^{+\hat{\mathcal{W}}} = (\text{Lk}_{\bar{X}}(v)^{+\mathcal{W}})/\mathcal{N}_v$, and the latter is E -hyperbolic by Corollary 3.39. Hence:

Lemma 4.13 (Edge-type). *If $\hat{\Delta}$ is of edge-type then $\text{Lk}(\hat{\Delta})^{+\widehat{\mathcal{W}}}$ is hyperbolic, and the hyperbolicity constant is depth-resistant.*

We are now left with the triangle-type case, which we split into Lemmas 4.14 and 4.15. For every $v \in \overline{X}^{(0)}$ set $L_{[v]} := L_v/(\mathcal{N} \cap \text{Stab}_G(v))$. The definition does not depend on the choice of $v \in [v]$, as $nL_v = L_{nv}$ for every $n \in \mathcal{N}$.

Lemma 4.14. *The following holds if \mathcal{N} is deep enough.*

- *If $Z_{[v]}$ is finite then $L_{[v]}$ is uniformly bounded, and the bound depends on \mathcal{N} .*
- *If instead $Z_{[v]}$ is infinite, then the quotient map $L_v \rightarrow L_{[v]}$ is a $\text{Stab}_G(v)$ -equivariant quasi-isometry, whose constants are depth-resistant. As a consequence, $L_{[v]}$ is a quasiline on which $Z_{[v]}$ acts geometrically, while $Z_{[w]}$ acts with uniformly bounded orbits whenever $[w] \in \text{Lk}_{\overline{X}/\mathcal{N}}([v])$.*

Proof. If Z_v is finite then L_v was already uniformly bounded. Thus assume that Z_v is infinite, so that L_v is a quasiline on which Z_v acts geometrically. If $Z_v \cap \mathcal{N} \neq \{0\}$, then $L_{[v]}$ is bounded, and the bound is uniform as there are finitely many G -orbits of vertices in \overline{X} . Thus suppose instead that $Z_v \cap \mathcal{N} = \{0\}$. Recall that, in Subsection 4.1, we constructed a quasimorphism $\psi_v: E_v \rightarrow \mathbb{R}$ which is trivial on $\mathcal{N} \cap E_v$. As each E_v has finite index in $\text{Stab}_G(v)$, we can assume that \mathcal{N} is deep enough that every Γ_w is contained in E_v whenever $w \in \text{Star}_{\overline{X}}(v) \cap G\mathcal{B}$, so that $\mathcal{N} \cap \text{Stab}_G(v) = \mathcal{N} \cap E_v$. Now, by Remark 4.2, every subgroup of E_v on which ψ_v vanishes (such as $\mathcal{N} \cap \text{Stab}_G(v)$) acts with uniformly bounded orbits on L_v . Thus, the quotient map $L_v \rightarrow L_{[v]}$ is a $\text{Stab}_G(v)$ -equivariant quasi-isometry, whose constants are depth-resistant. \square

Lemma 4.15 (Triangle-type). *The following holds if \mathcal{N} is deep enough. Let $\hat{\Delta} = \{([v], [x]), ([w])\}$ be of triangle-type. Then there is a $\text{Stab}_G([w])$ -equivariant quasi-isometry $\text{Lk}(\hat{\Delta})^{+\widehat{\mathcal{W}}} \rightarrow L_{[w]}$, whose constants are depth-resistant.*

Proof. Let $\Delta = \{(v, x), (w)\}$ be any lift of $\hat{\Delta}$. By [Man24, Lemma 3.34], for every $w \in \overline{X}^{(0)}$ there is a $\text{Stab}_G(w)$ -equivariant (K, K) -quasi-isometry $L_w \rightarrow \text{Lk}(\Delta)^{+\mathcal{W}}$, for some uniform constant $K \geq 0$. Moreover, $L_{[w]} = L_w/(\mathcal{N} \cap \text{Stab}_G(w))$, so it suffices to show that $\text{Lk}(\hat{\Delta})^{+\widehat{\mathcal{W}}} = (\text{Lk}(\Delta)^{+\mathcal{W}})/(\mathcal{N} \cap \text{Stab}_G(w))$, as this shall imply the existence of a $\text{Stab}_G([w])$ -equivariant (K, K) -quasi-isometry $\text{Lk}(\hat{\Delta})^{+\widehat{\mathcal{W}}} \rightarrow L_{[w]}$. It is clear that the quotient projection of $\text{Lk}(\Delta)$ is contained in $\text{Lk}(\hat{\Delta})$, and that if $y, y' \in \text{Lk}(\Delta)$ are \mathcal{W} -adjacent then their projections $[y], [y']$ are $\widehat{\mathcal{W}}$ -adjacent by construction. Conversely, let $[y], [y'] \in \text{Lk}(\hat{\Delta})$ be $\widehat{\mathcal{W}}$ -adjacent. Lift $[y]$ to $y \in \text{Lk}(\Delta)$, and lift $[y']$ to some y' which is \mathcal{W} -adjacent to y . Let $n \in \mathcal{N}$ be such that $y' \in nw$. Then we must have that $nw = w$, or w and nw would be \mathcal{W} -adjacent and there would be an edge in the simplicial graph \overline{X}/\mathcal{N} connecting $[w]$ to itself. Thus $y' \in \text{Lk}(\Delta)$ as well, and we are done. \square

4.3.4. *Quasi-isometric embeddings.* We move on to show that the augmented link of a simplex $\hat{\Delta}$ of \hat{X} is quasi-isometrically embedded in $Y_{\hat{\Delta}}$, thus proving the second part of Axiom (2). Again, we look at all possible shapes of $\text{Lk}(\hat{\Delta})$, according to Lemma 2.2. If $\text{Lk}(\hat{\Delta})$ has diameter at most 2, or if $\hat{\Delta} = \emptyset$, then clearly $\mathcal{C}(\hat{\Delta})$ is

quasi-isometrically embedded in $Y_{\hat{\Delta}}$. Then we only need to deal with the following cases:

- $\hat{\Delta} = \{([v], [x])\}$ of edge-type, where $[v]$ has valence greater than one in \overline{X}/\mathcal{N} ;
- $\hat{\Delta} = \{([v], [x]), ([w])\}$ of triangle-type.

Lemma 4.16 (Edge-type). *The following holds if \mathcal{N} is deep enough. Let $\hat{\Delta} = \{([v], [x])\}$ be a simplex of edge-type, where $[v]$ has valence greater than one in \overline{X}/\mathcal{N} . Then there exists a coarsely Lipschitz retraction from $Y_{\hat{\Delta}} = \hat{p}^{-1}(\overline{X}/\mathcal{N} - \{[v]\})^{+\widehat{\mathcal{W}}}$ to $\text{Lk}(\hat{\Delta})^{+\widehat{\mathcal{W}}}$, whose constants are depth-resistant.*

Proof. The retraction \hat{p} maps $Y_{\hat{\Delta}}$ onto $(\overline{X}/\mathcal{N} - \{[v]\})^{+\widehat{\mathcal{W}}}$ and $\text{Lk}(\hat{\Delta})^{+\widehat{\mathcal{W}}}$ onto $\text{Lk}_{\overline{X}/\mathcal{N}}([v])^{+\widehat{\mathcal{W}}}$, so it is enough to build a retraction

$$\rho: (\overline{X}/\mathcal{N} - \{[v]\})^{+\widehat{\mathcal{W}}} \rightarrow \text{Lk}_{\overline{X}/\mathcal{N}}([v])^{+\widehat{\mathcal{W}}}.$$

For every $[u] \in (\overline{X}/\mathcal{N} - \{[v]\})^{(0)}$, pick any geodesic $\bar{\gamma}$ in $(\overline{X}/\mathcal{N})^{+\widehat{\mathcal{W}}}$ from $[u]$ to a closest point inside $\text{Lk}_{\overline{X}/\mathcal{N}}([v])$, and let $\varrho([u])$ be the endpoint of this geodesic. Notice that $[v]$ does not belong to $\bar{\gamma}$. Indeed, suppose that this is not the case, and let $[t]$ be the vertex of $\bar{\gamma}$ which comes right before $[v]$. If $[t]$ and $[v]$ are $(\overline{X}/\mathcal{N})$ -adjacent, then we would contradict the fact that γ connects $[u]$ to a closest point in $\text{Lk}_{\overline{X}/\mathcal{N}}([v])$. If instead $[t]$ and $[v]$ are $\widehat{\mathcal{W}}$ -adjacent, then $[t]$ is also $\widehat{\mathcal{W}}$ -adjacent to some $[w] \in \text{Lk}_{\overline{X}/\mathcal{N}}([v])$, and we could find a path from $[u]$ to $[w]$ which is shorter than $\bar{\gamma}$, again finding a contradiction.

Now we want to show that ρ is both coarsely well-defined and coarsely Lipschitz, for some depth-resistant constants. Let $[u], [u'] \in (\overline{X}/\mathcal{N} - \{[v]\})^{(0)}$ be such that $d_{(\overline{X}/\mathcal{N})^{+\widehat{\mathcal{W}}}}([u], [u']) \leq 1$, and let $\bar{\gamma}$ (resp. $\bar{\gamma}'$) be a geodesic from $[u]$ (resp. $[u']$) to $\text{Lk}_{\overline{X}/\mathcal{N}}([v])$. The configuration in Figure (4) is therefore a geodesic pentagon inside $(\overline{X}/\mathcal{N})^{+\widehat{\mathcal{W}}}$, which by Corollary 3.34 we can lift to $\overline{X}^{+\mathcal{W}}$ if \mathcal{N} is deep enough.




FIGURE 4. The pentagon inside $(\overline{X}/\mathcal{N})^{+\widehat{\mathcal{W}}}$, where the full lines represent \overline{X}/\mathcal{N} -edges, the dashed lines are geodesics of $(\overline{X}/\mathcal{N})^{+\widehat{\mathcal{W}}}$, and the dash-and-dot line means that $d_{(\overline{X}/\mathcal{N})^{+\widehat{\mathcal{W}}}}([u], [u']) \leq 1$. The configuration lifts to a pentagon inside $\overline{X}^{+\mathcal{W}}$, with vertices v, w, u, u', w' .

Let γ (resp. γ') be the lift of $\bar{\gamma}$ (resp. $\bar{\gamma}'$), with endpoints u and $w \in \varrho([u])$ (resp. u' and $w' \in \varrho([u'])$). Notice that neither γ nor γ' contain any lift of $[v]$, as pointed out above. Let v be the lift of $[v]$ which is adjacent to both w and w' .

Now we claim that $d_{\text{Lk}_{\bar{X}}(v)+\mathcal{W}}(w, w') \leq 6E$, which then implies that $\varrho([u])$ and $\varrho([u'])$ are $6E$ -close in $\text{Lk}_{\bar{X}/\mathcal{N}}([v])^{+\widehat{\mathcal{W}}}$. Indeed, if $d_{\text{Lk}_{\bar{X}}(v)+\mathcal{W}}(w, w') > 6E$, then by triangle inequality one between $d_{\text{Lk}_{\bar{X}}(v)+\mathcal{W}}(w, u)$, $d_{\text{Lk}_{\bar{X}}(v)+\mathcal{W}}(u, u')$, and $d_{\text{Lk}_{\bar{X}}(v)+\mathcal{W}}(u', w')$ is at least $2E$. But this contradicts the strong bounded geodesic image, Lemma 2.10, as neither γ nor γ' can pass through v . \square

The proof in the triangle-type case is similar, but to build the retraction we need something more sophisticated than a geodesic, which we call an *approach path*. Our approach paths are not the same as their homonyms from [BHMS24, Definition 8.36], although they play a similar role.

Lemma 4.17 (Triangle-type). *The following holds if \mathcal{N} is deep enough. Let $\hat{\Delta} = \{([v], [x]), ([w])\}$ be a simplex of triangle-type. There exists a coarsely Lipschitz retraction*

$$\varrho_{[w]}: Y_{\hat{\Delta}} \rightarrow \text{Lk}(\hat{\Delta})^{+\widehat{\mathcal{W}}},$$

whose constants are depth-resistant.

Proof. First, notice that $Y_{\hat{\Delta}} = \left((L_{[w]})^{(0)} \cup \hat{\mathbf{p}}^{-1}(\bar{X}/\mathcal{N} - \text{Star}_{\bar{X}/\mathcal{N}}([w])) \right)^{+\widehat{\mathcal{W}}}$. Now, for every $[y] \in (L_{[w]})^{(0)}$ set $\varrho_{[w]}([y]) = [y]$. For every $[u] \in (\bar{X}/\mathcal{N} - \text{Star}_{\bar{X}/\mathcal{N}}([w]))^{(0)}$ we define the value of the retraction on $\text{Cone}([u])$ as follows. Pick a geodesic $\bar{\gamma}$ in $\bar{X}^{+\mathcal{W}}/\mathcal{N}$ from $[u]$ to $[w]$. Let $[r] \in \bar{\gamma}$ be the last point before $[w]$.

- (1) Suppose first that $[r] \notin \text{Lk}_{\bar{X}/\mathcal{N}}([w])$, i.e. the last edge of $\bar{\gamma}$ is a $\widehat{\mathcal{W}}$ -edge. Then choose any $[y] \in (L_{[w]})^{(0)}$ which is $\widehat{\mathcal{W}}$ -adjacent to $[r]$, and set $\varrho_{[w]}(\text{Cone}([u])) = [y]$. Moreover, let $\bar{\lambda}$ be the subpath of $\bar{\gamma}$ between $[u]$ and $[r]$. What we get is the configuration in Figure (5), which we call an *approach path of type \mathcal{W}* .
- (2) Suppose instead that $[r] \in \text{Lk}_{\bar{X}/\mathcal{N}}([w])$, and let $[t]$ be the last point of $\bar{\gamma}$ before $[r]$. There exists $[a] \in \text{Lk}_{\bar{X}/\mathcal{N}}([r])$ which is within distance 1 from $[t]$ inside $\bar{X}^{+\mathcal{W}}/\mathcal{N}$ ($[a]$ might be $[t]$ itself, if the edge of $\bar{\gamma}$ between $[t]$ and $[r]$ comes from \bar{X}/\mathcal{N}). Notice that $[a] \neq [w]$, because otherwise $[t]$ would be $\bar{X}^{+\mathcal{W}}/\mathcal{N}$ -adjacent to $[w]$ and this would contradict the fact that $\bar{\gamma}$ is a geodesic. Now pick a geodesic from $[a]$ to $[w]$ inside $\text{Lk}_{\bar{X}/\mathcal{N}}([r])^{+\mathcal{W}}$, and let $[b]$ be the second-to-last point of such geodesic. Then $[b]$ is $\widehat{\mathcal{W}}$ -adjacent to $[w]$, and therefore also to some $[y] \in (L_{[w]})^{(0)}$, so we set $\varrho_{[w]}(\text{Cone}([u])) = [y]$. For further reference, let $\bar{\eta}_1$ be the subpath of $\bar{\gamma}$ from $[u]$ to $[t]$, and let $\bar{\eta}_2$ be the subgeodesic from $[a]$ to $[b]$ inside $\text{Lk}_{\bar{X}/\mathcal{N}}([r])^{+\mathcal{W}}$. We get the configuration in Figure (6), which we call an *approach path of type \bar{X}* .

First, we prove that both types of approach paths lift to $X^{+\mathcal{W}}$, meaning that all vertices and geodesics involved in the definition admit lifts which are arranged in the same configuration.

Claim 4.18. *An approach path of type \mathcal{W} lifts to $X^{+\mathcal{W}}$.*




FIGURE 5. An approach path of type \mathcal{W} . Here the full arc is an edge of \widehat{X} , the dashed lines are $\widehat{\mathcal{W}}$ -edges, and the dash-and-dot line is a geodesic inside $\overline{X}^{+\mathcal{W}}/\mathcal{N}$, which does not intersect $\text{Star}_{\overline{X}/\mathcal{N}}([w])$.




FIGURE 6. An approach path of type \overline{X} . Here the full arcs represent edges of \widehat{X} ; the dashed lines are $\widehat{\mathcal{W}}$ -edges; the black dash-and-dot lines are geodesics inside $\overline{X}^{+\mathcal{W}}/\mathcal{N}$; and the red dash-and-dot line is a geodesic inside $\text{Lk}_{\overline{X}/\mathcal{N}}([r])^{+\mathcal{W}}$. It will be relevant that $\bar{\eta}_1$ does not intersect $\text{Star}_{\overline{X}/\mathcal{N}}([w])$.


Proof of Claim 4.18. Lift $\bar{\gamma}$ to a geodesic $\gamma \in \overline{X}^{+\mathcal{W}}$, with endpoints $u \in [u]$ and $r \in [r]$. Then there exists $y \in [y]$ which is \mathcal{W} -adjacent to r , by how $\widehat{\mathcal{W}}$ -edges are defined, and set $w = p(y)$. Then the configuration is as in Figure (5) (notice that w and r are not \overline{X} -adjacent, or their projections $[w]$ and $[r]$ would be adjacent as well). \blacksquare

Claim 4.19. *An approach path of type \overline{X} lifts to $X^{+\mathcal{W}}$.*

Proof of Claim 4.19. Let $\eta_1 \subset \overline{X}^{+\mathcal{W}}$ be a lift of $\bar{\eta}_1$, with endpoints $u \in [u]$ and $t \in [t]$. By Corollary 3.34, we can lift the triangle with vertices $[r], [a], [t]$ to a triangle in $\overline{X}^{+\mathcal{W}}$ with vertices $r \in [r], a \in [a], t' \in [t]$, and up to the action of \mathcal{N} we can assume that $t' = t$. Now, by Lemma 3.32 we can lift $\bar{\eta}_2$ to a geodesic inside $\text{Lk}_{\overline{X}}(r)$, with endpoints $a \in [a]$ and $b \in [b]$. Let $y \in [y]$ be \mathcal{W} -adjacent to b , let $w = p(y)$, and let $r' = nr \in \text{Lk}_{\overline{X}}(w)$ be a lift of $[r]$, for some $n \in \mathcal{N}$. The whole configuration is as in Figure (7).

If $n = 1$, then $r = r'$, and the configuration we get is a lift of the subgraph in Figure (6). Otherwise, we proceed by induction on the complexity of n . Let (s, γ_s) be a shortening pair, as in Corollary 4.7. We want to replace some lifts from Figure (6), without breaking the configuration, in such a way that the new lifts of $[r]$ will differ by $\gamma_s n$, which has strictly less complexity than n .


- If γ_s fixes r then we apply γ_s to all lifts from Figure (7).
- If γ_s fixes r' then $r' = \gamma_s nr$ already differs from r by $\gamma_s n$, so we do nothing.
- Otherwise, $d_s(r, r') \geq 100E$. Arguing as in Lemma 4.11, the strong bounded geodesic image 2.8 tells us that γ_s must fix either b or w . If $\gamma_s w = w$, then

FIGURE 7. The (unclosed) lift of an approach path of type \overline{X} .

we just replace r' by $\gamma_s r'$, which is still X -adjacent to w and therefore to y . If instead $\gamma_s b = b$, then we apply γ_s to r' , w , and y . Notice that $\gamma_s y$ is still \mathcal{W} -adjacent to $b = \gamma_s b$. In both cases, after the replacement r and $\gamma_s r'$ differ by $\gamma_s n$.

Proceeding by induction, we can find lifts such that $r = r'$, as required. \blacksquare

Finally, we shall prove with a single argument that the map ϱ is both coarsely well-defined and coarsely Lipschitz with depth-resistant constants. Let $[u], [u'] \in (\overline{X}/\mathcal{N} - \text{Star}_{\overline{X}/\mathcal{N}}([w]))^{(0)}$ be such that $d_{(\overline{X}/\mathcal{N})+\overline{\mathcal{W}}}([u], [u']) \leq 1$, and consider two approach paths, one from $[u]$ to $[w]$ and one from $[u']$ to $[w]$. Let $u \in [u]$ and $u' \in [u']$ be such that $d_{\overline{X}+\mathcal{W}}(u, u') \leq 1$. Now lift both approach paths, starting from u and u' , respectively, to get the configuration from Figure (8), where both w and $w' = nw$ belong to $[w]$ and $n \in \mathcal{N}$. To illustrate the process, we assume that the path starting at u is of type \overline{X} , while the path starting at u' is of type \mathcal{W} (the two other cases are dealt with analogously).

FIGURE 8. Two lifts of approach paths starting at $\overline{X}^{+\mathcal{W}}$ -adjacent vertices. The blue path is a concatenation of three geodesics of $\overline{X}^{+\mathcal{W}}$.

We split the argument into two steps.

Step 1: Gluing w to w' . We first prove that we can change the lifts, without breaking the configuration from Figure (8), until $w = w'$. This will again be a combination of Corollary 4.7 and the strong bounded geodesic image Lemma 2.8.

We proceed by induction on the complexity of n . If $n = 1$ then we have nothing to prove; otherwise, let (s, γ_s) be a shortening pair, as in Corollary 4.7, so that $\gamma_s n$

has strictly less complexity than n . We want to replace the lifts in such a way that the two new lifts of $[w]$ differ by $\gamma_s n$, in order to conclude by induction.

If γ_s fixes w , we apply γ_s to the whole diagram. If γ_s fixes w' then we do nothing, as $w' = \gamma_s w'$ already differs from w by $\gamma_s n$.

Otherwise, we have that $d_s(w, w') \geq 100E$. Now look at the blue path from Figure (8), which is a concatenation of three geodesics of $\bar{X}^{+\mathcal{W}}$. If there is a point z on the blue path such that $d_{\bar{X}}(z, s) \leq 1$, then $\gamma_s z = z$ and we can apply γ_s to w' , y' , and the subpath of the blue path between z and r' . The new blue path is again a concatenation of three geodesics, since we did not change its projection to \bar{X}/\mathcal{N} and therefore is still a lift of three geodesics. Then we can conclude by induction.

We are left with the case when $d_s(w, w') \geq 100E$, but no point on the blue path belongs to $\text{Star}_{\bar{X}}(s)$. In particular, by the strong bounded geodesic image Lemma 2.8, the distances $d_s(w', r')$, $d_s(r', u')$, $d_s(u', u)$, $d_s(u, t)$ are all well-defined and bounded above by $2E$. Moreover, the triangle inequality yields

$$d_s(w, t) \geq d_s(w, w') - d_s(w', r') - d_s(r', u') - d_s(u', u) - d_s(u, t) \geq 92E.$$

If $r \notin \text{Star}_{\bar{X}}(s)$, then by the triangle inequality, at least one of $d_s(w, r)$ and $d_s(r, t)$ would be at least $46E > 2E$, again contradicting the strong bounded geodesic image Lemma 2.8. Hence γ_s must fix r . Furthermore, if no point on the red path η_2 belongs to $\text{Star}_{\bar{X}}(s)$, then by triangle inequality one between $d_s(w, a)$, $d_s(a, b)$, and $d_s(b, t)$ would be greater than $30E > 2E$. This would again contradict the strong bounded geodesic image Lemma 2.8, either inside $\text{Lk}_{\bar{X}}(r)^{+\mathcal{W}}$ (in the first two cases) or inside $\bar{X}^{+\mathcal{W}}$ (in the last case). Then let $k \in \eta_2$ be fixed by γ_s . If we apply γ_s to everything beyond r and k (meaning, to w' , y' , the blue path, and the subpath of η_2 between k and b), then we do not break the configuration, and we can conclude by induction.

Step 2: Bounding $d_{L_w}(y, y')$. After the previous step, our configuration looks as in Figure (9):




FIGURE 9. Now the lifts of both approach paths terminate at w .

Our final goal is to show that $d_{L_w}(y, y')$ is bounded in terms of E . This will then imply that $d_{L_{[w]}}(y, y')$ as well is bounded in terms of the depth-resistant constant E , concluding the proof.

Firstly, we argue that $y' \subseteq \rho_w^{r'}$. Indeed, with our Notation 2.7, $\rho_w^{r'}$ was defined as $\rho_{[\Delta]}^{[\Delta']}$, where $\Delta' = \{(s', x'), (r')\}$ is any simplex of triangle-type containing r' but no point in $(L_{r'})^{(0)}$. Moreover, by Definition 1.8, $\rho_{[\Delta]}^{[\Delta']} := p(\text{Sat}(\Delta') \cap Y_\Delta) \supseteq p(r')$ is obtained by applying the coarse closest point projection $p: Y_\Delta \rightarrow \mathcal{C}(\Delta)$ to r' , which

is \mathcal{W} -adjacent (that is, adjacent in Y_Δ) to y' . Similarly, $y \subseteq \rho_w^b$, so it suffices to bound the distance $d_w(b, r')$.

Now, notice that no point on the red path η_2 is \overline{X} -adjacent to w , as $\eta_2 \subset \text{Lk}_{\overline{X}}(r)$ and \overline{X} is triangle-free. Moreover, recall that by construction no point on the blue path belongs to $\text{Star}_{\overline{X}}(w)$. Thus both the red and the blue path have well-defined projections on L_w . Now, by the triangle inequality

$$d_w(b, r') \leq d_w(b, a) + d_w(a, t) + d_w(t, u) + d_w(u, u') + d_w(u, r').$$

The first term is at most $2E$, by the strong bounded geodesic image Lemma 2.8, applied inside $\text{Lk}_{\overline{X}}(r)^{+\mathcal{W}}$. All other terms are at most $2E$ each, again by strong BGI applied inside $\overline{X}^{+\mathcal{W}}$. Thus $d_w(b, r') \leq 10E$, and this concludes the proof of Lemma 4.17. \square

4.3.5. G/\mathcal{N} -action. The G -action on X induces a G/\mathcal{N} action on \widehat{X} , which has finitely many G/\mathcal{N} -orbits of links of simplices, and this action extends to \widehat{W} , as each edge of \widehat{W} lifts to some edge of \mathcal{W} .

Moreover, if one fixes a generating set S for G , there is a G -equivariant (K, K) -quasi-isometry $f := \text{Cay}(G, S) \rightarrow \mathcal{W}$, for some K depending on S . By taking the quotient by \mathcal{N} , we get a G/\mathcal{N} -equivariant map

$$\widehat{f} := \text{Cay}(G/\mathcal{N}, S\mathcal{N}) \rightarrow \widehat{\mathcal{W}},$$

which is again a (K, K) -quasi-isometry (notice that K is depth-resistant). Then $(\widehat{X}, \widehat{W})$ is a combinatorial HHG structure for G/\mathcal{N} , in view of the “moreover” part of Theorem 1.9.

4.4. The quotient is short(er). To conclude the proof of Theorem 4.1, we finally check that the combinatorial structure for G/\mathcal{N} , coming from the action on $(\widehat{X}, \widehat{W})$, is short:

Lemma 4.20. *G/\mathcal{N} admits a short HHG structure $(G/\mathcal{N}, \overline{X}/\mathcal{N})$, whose extensions are defined as in Lemma 3.42.*

Proof. Axiom (A) is clear, as \widehat{X} is a blowup of \overline{X}/\mathcal{N} by Lemma 4.3, and by Lemma 3.35 the latter is triangle- and square-free with no connected components which are points. Moreover, Axiom (B) is a combination of Lemmas 3.42 and 3.43. Regarding Axiom (C), the properties of the action on $\mathcal{C}\ell_{[v]} \cong L_{[v]}$ were proved in Lemmas 4.14 and 4.15. \square

4.5. Residual hyperbolicity.

Definition 4.21. We say that $\mathcal{N} = \langle\langle \Gamma_1, \dots, \Gamma_k \rangle\rangle$ is a *full kernel* if, for every $v \in \overline{X}^{(0)}$, there exists i such that Γ_i is conjugated into Z_v .

Corollary 4.22. *If \mathcal{N} is a full, deep-enough kernel, then G/\mathcal{N} is hyperbolic. Furthermore, if G is not virtually cyclic and the main coordinate space CS is unbounded, then G/\mathcal{N} is also non-elementary hyperbolic.*

Proof. Since \mathcal{N} is full, every cyclic direction $Z_{[v]}$ for the quotient is bounded, and therefore so is every $\mathcal{C}\ell_{[v]}$. Then Remark 2.5 shows that no two orthogonal domains in the structure have unbounded coordinate spaces, so G/\mathcal{N} is hyperbolic by [BHS21, Corollary 2.16].

In the “furthermore” setting, [BHS17b, Corollary 14.4] implies that G acts non-elementarily on \mathcal{CS} , and therefore on $\overline{X}^{+\mathcal{W}}$. Then Lemma 3.38 tells us that G/\mathcal{N} still acts non-elementarily on $\overline{X}^{+\mathcal{W}}/\mathcal{N}$, whenever \mathcal{N} is deep enough. \square

Recall that a group G is *fully residually* P for some property P if, for every finite subset $F \subset G$, there exists a quotient $G \rightarrow \overline{G}$ where F injects, and such that \overline{G} enjoys P .

Corollary 4.23. *A short HHG G is fully residually hyperbolic. If moreover G is not virtually cyclic and the main coordinate space \mathcal{CS} is unbounded, then G is fully residually non-elementary hyperbolic.*

Proof. Fix a finite set $F \subset G - \{1\}$, and let \mathcal{N} be a full kernel. By Corollary 3.40, we can choose \mathcal{N} to be deep enough that F injects in G/\mathcal{N} . Furthermore, since \mathcal{N} is full, by Corollary 4.22 we have that G/\mathcal{N} is hyperbolic (and non-elementary in the “moreover” setting). \square

Corollary 4.24. *If all hyperbolic groups are residually finite then all short HHG are residually finite.*

5. HOPF PROPERTY FROM CENTRAL QUOTIENTS

Recall that a group G is *Hopfian*, or has the *Hopf property*, if every surjective homomorphism $\phi: G \rightarrow G$ is an isomorphism. In this Section we develop some tools to study self-epimorphisms of short HHG, with the aim of then proving the Hopf property for most large hyperbolic type Artin groups. We expect that one could treat other short HHG similarly, and indeed in Subsection 5.5 we prove the Hopf property of certain HNN extensions of free groups.

5.1. A criterion. Firstly, we state a simple criterion for a group to be Hopfian.

Definition 5.1. We say that a group G has *enough Hopfian quotients* if for every surjective homomorphism $\phi: G \rightarrow G$ and non-trivial $g_0 \in G$ there exists a quotient H of G , say with quotient map q , and $n \geq 1$ such that:

- $q(g_0) \neq 1$,
- H is Hopfian,
- ϕ^n induces a homomorphism $\psi: H \rightarrow H$, which is necessarily surjective.

Remark 5.2. Note that the third bullet holds if and only if $\phi^n(\ker(q)) \leq \ker(q)$.

Lemma 5.3. *If G has enough Hopfian quotients then it is Hopfian.*

Proof. Given a surjective homomorphism ϕ and $g_0 \neq 1$ we have to argue that $\phi(g_0)$ is non-trivial. In the setting of Definition 5.1, this will follow if we show that $\phi^n(g_0)$ is non-trivial. But we have $\psi(q(g_0)) \neq 1$, and therefore $\phi^n(g_0) \neq 1$, as required. \square

Remark 5.4. We stress that, despite the name “enough Hopfian quotients”, the property from Definition 5.1 is stronger than being residually Hopfian, and indeed there are famous examples of residually finite, non-finitely-generated groups which do not satisfy the conclusion of Lemma 5.3. The key point of Definition 5.1 is the third bullet, which requires the quotient H to be “almost characteristic” in some sense.

5.2. Preliminary lemmas on central extensions.

Lemma 5.5. *Let G be a short HHG, and let $H \leq G$ be a subgroup isomorphic to a \mathbb{Z} -central extension $1 \rightarrow Z \rightarrow H \rightarrow K \rightarrow 1$.*

- (1) *If K is infinite then H is virtually contained in $\text{Stab}_G(v)$ for some $v \in \overline{X}^{(0)}$.*
- (2) *If moreover K is not virtually cyclic then Z is virtually contained in Z_v .*

Proof. We first notice that, if K is infinite, then the centraliser of any element of H is not virtually cyclic, and therefore, H cannot contain any element acting loxodromically on the top-level coordinate space for G (since this action is acylindrical, by [BHS17b, Corollary 14.4]).

Now, by the Omnibus subgroup theorem [DHS17, Theorem 9.20], together with [DHS17, Theorem 7.1], there exists $h \in H$ and a finite collection \mathcal{V} of pairwise orthogonal, unbounded domains, such that H permutes \mathcal{V} , and a power of h acts loxodromically on $\mathcal{C}V$ for every $V \in \mathcal{V}$. By the above discussion, \mathcal{V} does not coincide with the \sqsubseteq -maximal element. Hence, by our description of unbounded domains in a short HHG (Remark 2.5), we see that an index-two subgroup of H must fix a domain $V \in \mathcal{V}$ of the form ℓ_v or \mathcal{U}_v . In the first case H virtually fixes v , so (1) follows. In the second scenario, since $\mathcal{C}\mathcal{U}_v$ is unbounded, v is the only vertex $w \in \overline{X}^{(0)}$ such that $\ell_w \perp \mathcal{U}_v$, and again we get (1).

Towards proving (2), let H_0 be a finite-index subgroup of H contained in $\text{Stab}_G(v)$. The group $H_1 = H_0/(H_0 \cap Z_v)$ embeds in a hyperbolic group, and either Z is virtually contained in Z_v , or H_1 has infinite centre. The latter can only happen if H_1 is virtually cyclic, but then K would also be virtually cyclic, which is not possible under our assumption. Therefore Z is virtually contained in Z_v , as required. \square

We will also need the following support lemma. Recall that a group extension $1 \rightarrow Z \rightarrow H \rightarrow H/Z \rightarrow 1$ is *virtually trivial* if there exists a finite-index subgroup $H' \leq H$ and a group retraction $H' \rightarrow H' \cap Z$, which we call a *virtual retraction* from H to Z .

Lemma 5.6. *Let $1 \rightarrow Z \rightarrow H \rightarrow H/Z \rightarrow 1$ be a non-virtually-trivial extension, and let $\phi : H \rightarrow K$ be a surjective homomorphism whose kernel intersects Z trivially. Then the extension $1 \rightarrow \phi(Z) \rightarrow K \rightarrow K/\phi(Z) \rightarrow 1$ is non-virtually-trivial.*

Proof. Towards a contradiction, assume that the latter central extension is virtually trivial, so that there exist a finite-index subgroup $K' \leq K$ and a group retraction $\rho : K' \rightarrow K' \cap \phi(Z)$. Let $\psi : \phi(Z) \rightarrow Z$ be an inverse of ϕ . Then the composition $\psi \circ \rho \circ \phi$, defined on $\phi^{-1}(K') \leq H$, is a virtual retraction from H to Z , against the hypothesis that the first extension is non-virtually-trivial. \square

5.3. Certain relatively hyperbolic groups are Hopfian.

Theorem 5.7. *Let G be hyperbolic relative to \mathbb{Z} -central extensions of hyperbolic groups (including the case that G itself is such an extension). Then G is Hopfian.*

Proof. First, we note that if a group H is hyperbolic relative to subgroups which are virtually a direct product of \mathbb{Z} and a hyperbolic group (for short, virtual products), then it is Hopfian. Indeed, the peripheral subgroups are equationally Noetherian by [WR19, Corollary 6.13] for hyperbolic groups and [Val21, Theorem E] plus [BMR97, Theorem 1] for finite extensions of direct products of hyperbolic groups, so that H is Hopfian by [GH19, Corollary 3.14 and Theorem D].

We now proceed by induction on the number k of peripheral subgroups which are not virtual products (for short, *twisted*), the case $k = 0$ being what we discussed above.

Suppose that the statement holds when there are at most k twisted peripheral subgroups, and consider G having $k + 1$ twisted peripheral subgroups. Fix a self-epimorphism ϕ of G and $g_0 \neq 1$. We will use the criterion provided by Lemma 5.3, constructing a quotient H which will be hyperbolic relative to \mathbb{Z} -central extensions of hyperbolic groups with k twisted peripheral subgroups.

We first claim that, up to passing to a power of ϕ , there is some twisted peripheral P , with cyclic direction generated by z_P , and some positive integer N_P such that $\phi(z_P^{N_P})$ is conjugated into $\langle z_P^{N_P} \rangle$. Indeed, let P be any twisted peripheral. If $\phi^n(z_P)$ is a torsion element for some $n \in \mathbb{N}$ then we can take N_P to be its order; thus we can assume that, for any twisted peripheral P and any $n \in \mathbb{N}$, we have that $\phi^n(z_P)$ has infinite order. In this case, $\phi(P)$ is a \mathbb{Z} -central extension, and we claim that it cannot be an extension of a virtually cyclic group. Indeed, any such extension is virtually trivial, so this would contradict Lemma 5.6. Since G is a short HHG by [Man24, Proposition 6.3], we are now in a position to apply Lemma 5.5 to $\phi(P)$, and conclude that it is virtually contained in a conjugate P' of some peripheral subgroup, which must be twisted itself (again as a consequence of Lemma 5.6). Moreover, $P' \cap \phi(z_P)P'\phi(z_P)^{-1}$ contains the infinite subgroup $\langle \phi(z_P) \rangle \cap P'$, therefore $\phi(z_P) \in P'$ or we would contradict almost malnormality of peripheral subgroups. Notice also that the centraliser of $\phi(z_P)$ in P' , which contains $\phi(P) \cap P'$, cannot be virtually Abelian; hence $\phi(z_P)$ must be contained in the centre of P' .

Considering the directed graph with vertices the twisted peripherals and a directed edge from P to P' with $\phi(P)$ virtually contained in a conjugate of P' , we see that, up to passing to an iterated of ϕ (which is allowed by Lemma 5.3) there exists a twisted peripheral P such that $\phi(P)$ is virtually contained in a conjugate of P . Moreover, $\phi(z_P)$ is virtually contained in the relevant conjugate of the centre of P .

Therefore, by the relatively hyperbolic Dehn filling theorem, there exists a non-trivial multiple $N = N(g_0)$ of N_P such that:

- the group $H = G/\langle\langle z_P^N \rangle\rangle$ is hyperbolic relative to
 - (1) virtual products and k twisted peripheral subgroups (coming from peripherals of G), and
 - (2) the hyperbolic group $P/\langle\langle z_P^N \rangle\rangle$.
- the image of g_0 in H is non-trivial.

Note that we can drop the subgroup from Item 2 from the list of peripherals, so that H is hyperbolic relative to virtual products and at most k twisted peripheral subgroups, and it is therefore Hopfian by induction. Furthermore, the fact that $\phi(z_P^{N_P})$ is conjugate into $\langle z_P^{N_P} \rangle$ ensures that ϕ induces a homomorphism of H , so that we checked all conditions from Lemma 5.3, and the proof is complete. \square

5.4. The product region graph.

Definition 5.8. Let G be a short HHG with support graph \overline{X} . Let \overline{X}' be the full subgraph of \overline{X} spanned by all vertices v with Z_v infinite. The *product region graph* of G , denoted by $\mathcal{PR}(G)$, is the simplicial graph whose vertex set is $(\overline{X}')^{(0)}/G$, and

where two vertices are adjacent if and only if they admit adjacent representatives in \overline{X}' .

Remark 5.9. The product region graph has the following interpretation. The vertices are conjugacy classes of vertex stabilisers (which are HHS product regions), and two vertices are adjacent if there exist conjugacy representatives that intersect along an edge group.

Definition 5.10. A short HHG $(G, \overline{X}, \mathcal{W})$ has *central cyclic directions* if, for every $v \in \overline{X}^{(0)}$, either Z_v is finite or it lies in the centre of $\text{Stab}_G(v)$ whenever $v \in \overline{X}^{(0)}$. In other words, whenever Z_v is infinite, $\text{Stab}_G(v)$ is a \mathbb{Z} -central extension of a hyperbolic group.

Lemma 5.11. *If a short HHG $(G, \overline{X}, \mathcal{W})$ has discrete product region graph then it is hyperbolic relative to $\{\text{Stab}_G(v)\}_{v \in V}$, where V is a collection of representatives for some G -orbits of vertices with infinite cyclic directions. In particular, if G furthermore has central cyclic directions then it is Hopfian by Theorem 5.7.*

Proof. The product region graph of (G, \overline{X}) being discrete is equivalent to no two vertices v of the support graph \overline{X} with infinite Z_v being connected to each other. Let (X, \mathcal{W}) be a combinatorial HHG structure for G , where X is a blowup of \overline{X} . With the aim of using [Rus22, Corollary 5.2], we now modify the HHS structure, by removing various bounded domains. Namely, we only keep the following:

- The maximal domain S ;
- For every $v \in \overline{X}^{(0)}$ of valence greater than one, the domain $\mathcal{U}_v = \text{Lk}(\{(v, x)\})$, where x is any point in $(L_v)^{(0)}$;
- For every $v \in \overline{X}^{(0)}$ for which L_v is infinite, the domain ℓ_v ;
- For every $v \in \overline{X}^{(0)}$ for which L_v is infinite, the domain $I_v := [\Sigma]$ corresponding to the link of the simplex $\Sigma = \{(v)\}$.

Let $\mathfrak{S}_{\text{keep}} \subset \mathfrak{S}$ be the G -invariant subset containing the above domains. Notice that I_v contains both ℓ_v and \mathcal{U}_v , and has no orthogonal domain in $\mathfrak{S}_{\text{keep}}$.

Now, in view of Remark 2.5, every $U \in \mathfrak{S} - \mathfrak{S}_{\text{keep}}$ has bounded coordinate space. Hence, by inspection of the definition of a HHS (see e.g. [BHS19, Definition 1.1]), removing these domains can only affect the existence of containers and the validity of the large link axiom, so we must check that both still hold in $\mathfrak{S}_{\text{keep}}$.

Containers: By inspection of Remark 2.5, combined with the fact that the product region graph is discrete, the only pairs of orthogonal domains in $\mathfrak{S}_{\text{keep}}$ are of the form \mathcal{U}_v and ℓ_v , for $v \in \overline{X}^{(0)}$ of valence greater than one and with infinite cyclic direction. Thus, whenever U and V are both nested in some $T \in \mathfrak{S}_{\text{keep}}$, the container for U inside T is V , and vice versa.

Large links: Let $U \in \mathfrak{S}_{\text{keep}}$ which is not \sqsubseteq -minimal, and let $z, z' \in G$. We want to prove that, if one sets $N = 2\text{Ed}_U(z, z') + 2E$, there exist $\{T_1, \dots, T_{[N]}\} \subseteq \mathfrak{S}_{\text{keep}}$ properly nested in U and such that, whenever $V \in \mathfrak{S}_{\text{keep}}$ is properly nested in U and $d_V(z, z') > E$, then $V \sqsubseteq T_i$ for some i .

We first notice that, if U contains only finitely many domains of $\mathfrak{S}_{\text{keep}}$ with unbounded coordinate spaces, then the large link axiom holds trivially (possibly after enlarging the HHS constant E). In particular, this happens if $U = I_v$, as it only contains ℓ_v and possibly \mathcal{U}_v .

Moreover, suppose that U only contained \sqsubseteq -minimal domains, already in \mathfrak{S} (this is the case if $U = \mathcal{U}_v$). Then the large link axiom for \mathfrak{S} produces a collection $\{T_1, \dots, T_{[N]}\} \subseteq \mathfrak{S}$, and one can simply intersect such collection with \mathfrak{S}_{keep} to get the required property.

The only case which is not covered by the above is when $U = S$ is the maximal domain. Let $\mathcal{T} = \{T_1, \dots, T_{[N]}\} \subseteq \mathfrak{S}$ the collection granted by the large link axiom for \mathfrak{S} . If $\mathcal{T} \subseteq \mathfrak{S}_{keep}$ we have nothing to prove; otherwise let $T \in \mathcal{T} - \mathfrak{S}_{keep}$. If no $V \in \mathfrak{S}_{keep}$ is properly nested in T we can simply remove T from the collection; otherwise we need to replace T with some finite collection inside \mathfrak{S}_{keep} . Suppose that $T = [\Delta]$ for some simplex $\Delta \subseteq X$. There are several cases to consider, according to the shape of Δ for which $[\Delta] \notin \mathfrak{S}_{keep}$. In what follows, let $\{v, w\}$ be an edge of \overline{X} containing $\overline{\Delta}$, and let $x \in (L_v)^{(0)}$ and $y \in (L_w)^{(0)}$.

- Suppose that $\Delta = \{(v)\}$, where v has finite cyclic direction. If there exists $V \in \mathfrak{S}_{keep}$ which is nested in T , then $V \sqsubseteq \mathcal{U}_v$. If $\mathcal{U}_v \in \mathfrak{S}_{keep}$ then we replace T by \mathcal{U}_v ; otherwise $\text{Lk}_{\overline{X}}(v) = \{w\}$, so that V can only be ℓ_w , and we replace T by ℓ_w .
- Suppose that $\Delta = \{(x)\}$. Again, if there exists $V \in \mathfrak{S}_{keep}$ which is nested in T , then $V \sqsubseteq \mathcal{U}_v$, and we can argue as above.
- Suppose that $\Delta = \{(v, x)\}$, so that $T = \mathcal{U}_v$. As $T \notin \mathfrak{S}_{keep}$, we must have that $\text{Lk}_{\overline{X}}(v) = \{w\}$, so the only $V \in \mathfrak{S}_{keep}$ which is nested in T is $V = \ell_w$. Thus we replace T by ℓ_w .
- Suppose that $\Delta = \{(v, w)\}$. The only unbounded domains which are nested in T can be ℓ_v and ℓ_w , so we can replace T by $\{\ell_v, \ell_w\} \cap \mathfrak{S}_{keep}$.
- Suppose that $\Delta = \{(v, y)\}$. The only unbounded domain which is nested in T can be ℓ_v , so we can replace T by ℓ_v .
- Finally, if $\Delta = \{(x, y)\}$ then no $V \in \mathfrak{S}$ is nested in T , so we can simply remove the latter.

This concludes the verification of the large link axiom.

It is now readily seen that the HHG structure (G, \mathfrak{S}_{keep}) has isolated orthogonality in the sense of [Rus22], specifically isolated by the set of domains $\{I_v\}$ as above. Then [Rus22, Theorem 4.3] states that G , equipped with a fixed word metric, is hyperbolic relative to the product regions $\{P_{I_v}\}_{v \in \overline{X}^{(0)}}$. Furthermore, since G permutes the product regions, [Dru09, Proposition 5.1] yields that G is hyperbolic relative to a collection \mathcal{Q} of subgroups, where for every $Q \in \mathcal{Q}$ there exists a unique P_{I_v} within finite Hausdorff distance. Notice that P_{I_v} is exactly P_{ℓ_v} , so it lies within finite Hausdorff distance from a coset of a vertex stabiliser; hence there are $x \in G$ and a vertex $w \in \overline{X}^{(0)}$ with infinite cyclic direction such that the Hausdorff distance between Q and $x \text{Stab}_G(w)$ is finite. In turn, $d_{Haus}(x \text{Stab}_G(w), \text{Stab}_G(xw)) \leq |x|$, where $|x|$ is the norm of x in the fixed word metric; thus let $w' = xw$, and let $R = d_{Haus}(\text{Stab}_G(w'), Q)$.

We now claim that $\text{Stab}_G(w') = Q$, thus completing the proof. To see this, let $K = Q \cap \text{Stab}_G(w')$, and notice that, by [HW09, Lemma 4.5], there exists a constant $R' > 0$ such that

$$Q \subseteq Q \cap N_R(\text{Stab}_G(w')) \subseteq N_{R'}(Q \cap \text{Stab}_G(w')) = N_{R'}(K)$$

so K has finite index in Q . The same argument shows that K has finite index in $\text{Stab}_G(w')$.

Now, given $g \in \text{Stab}_G(w')$, notice that $K \cap gKg^{-1}$ has finite index in K , as it is the intersection of two finite-index subgroups of $\text{Stab}_G(w')$. Hence $K \cap gKg^{-1}$ has finite index in Q as well, since K has finite index in Q . Thus $Q \cap gQg^{-1}$, which contains $K \cap gKg^{-1}$, has finite index in Q , and since Q is infinite and almost malnormal we must have that $g \in Q$. As g was any element of $\text{Stab}_G(w')$, this proves that $\text{Stab}_G(w') \leq Q$.

For the reverse inclusion let $g \in Q$. The same argument shows that $\text{Stab}_G(w') \cap g\text{Stab}_G(w')g^{-1} = \text{Stab}_G(w') \cap \text{Stab}_G(gw')$ has finite-index in $\text{Stab}_G(w')$. In particular some non-trivial subgroup of $Z_{gw'}$ fixes w' , which means that $d_{\overline{X}}(w', gw') \leq 1$. But since G has discrete product region graph and $Z_{gw'}$ is infinite, gw' must actually coincide with w' , i.e., $g \in \text{Stab}_G(w')$, as required. \square

Definition 5.12. A short HHG $(G, \overline{X}, \mathcal{W})$ has *clean intersections* if, for every \overline{X} -adjacent vertices v, w , the edge group $\text{Stab}_G(v) \cap \text{Stab}_G(w)$ coincides with $\langle Z_v, Z_w \rangle$.

Definition 5.13. We say that a short HHG has *stable product regions* if the following strengthening of Lemma 5.5 holds for G . Let $H \leq G$ be a subgroup isomorphic to a \mathbb{Z} -central extension $1 \rightarrow Z \rightarrow H \rightarrow K \rightarrow 1$, and suppose that H is virtually contained in some $\text{Stab}_G(v)$.

- (1) If K is infinite then $\text{Stab}_G(v)$ actually contains H . If, in addition, K is not virtually cyclic then Z is contained in Z_v .
- (2) There exists I , depending on G only, such that if K is finite then H has an index- $\leq I$ subgroup contained in $\text{Stab}_G(v)$.

Definition 5.14. Let $(G, \overline{X}, \mathcal{W})$ be a short HHG, let $P = \text{Stab}_G(v)$ for some $v \in \overline{X}^{(0)}$, and let $\phi: G \rightarrow G$ be a homomorphism. P is *always restrained* with respect to ϕ if, for every $k \in \mathbb{N}$, either:

- $\phi^k(P)$ is a \mathbb{Z} -central extension of a non-elementary hyperbolic group, or
- $\phi^k(P)$ is virtually \mathbb{Z}^2 , and virtually contained in an edge group.

Lemma 5.15. Let $(G, \overline{X}, \mathcal{W})$ be a colourable short HHG with stable product regions, central cyclic directions, and clean intersections. Let $g_0 \in G - \{1\}$, let $\phi: G \rightarrow G$ be a homomorphism, and let $P_i = \text{Stab}_G(v_i)$ be always restrained vertex stabilisers with respect to ϕ , for $i = 1, \dots, r$. Then there exists a kernel \mathcal{N} , as in Notation 2.18, such that:

- $\phi^M(\mathcal{N}) \leq \mathcal{N}$ for some $M \in \mathbb{N}_{>0}$;
- $Z_{v_i} \cap \mathcal{N} \neq \{1\}$ for every $i = 1, \dots, r$;
- $g_0 \notin \mathcal{N}$;
- G/\mathcal{N} is a colourable short HHG.

Remark 5.16. Notice that, in the above Lemma, the product region graph of G/\mathcal{N} injects in the graph obtained from $\mathcal{PR}(G)$ after removing the open stars of the vertices corresponding to v_1, \dots, v_r .

Now, suppose that G has central cyclic directions. In view of the above discussion, if removing all always restrained stabilisers makes $\mathcal{PR}(G)$ discrete, then G/\mathcal{N} is Hopfian by Lemma 5.11. In other words, the quotient G/\mathcal{N} satisfies all requirements of our criterion, Lemma 5.3.

Proof of Lemma 5.15. We proceed by induction on r , the base case $r = 0$ being trivial. For the inductive step, let \mathcal{N}' be a kernel satisfying the statement, for the collection $\{P_1, \dots, P_{r-1}\}$. Up to replacing ϕ by a power, we can assume that

$\phi(\mathcal{N}') \leq \mathcal{N}'$. Set $P = P_r$, and choose a generator z of Z_r . As P is always restrained, we are in one of the three situations below.

(1) Suppose first that $\phi^k(z^n) \in \mathcal{N}'$ for some $k, n \in \mathbb{N}$ and $n \neq 0$. Then set $\mathcal{N} = \langle\langle \mathcal{N}', z^n \rangle\rangle$, which is preserved by ϕ^k . Working in the short HHG G/\mathcal{N}' , we see that, up to replacing n by a non-trivial multiple, we can assume that $g_0 \notin \mathcal{N}$, and that G/\mathcal{N} is again a colourable short HHG.

(2) Suppose now that, for every $k \in \mathbb{N}$, $\phi^k(P)$ is a \mathbb{Z} -central extension of a hyperbolic group. By stability of product regions, this means that $\phi^k(P)$ is conjugated into some vertex stabiliser Q_k ; moreover $\phi(Q_k)$ is again a \mathbb{Z} -central extension of a non-elementary hyperbolic group, as it contains $\phi^{k+1}(P)$, and the stability assumption implies that the ϕ -image of the centre of Q_k is conjugated into the centre of Q_{k+1} . Now, there are finitely many cyclic directions up to conjugation, so we can find $n \in \mathbb{N}_{>0}$ and a cyclic direction Z' such that $\phi^n(\langle z \rangle) \leq Z'$, and $\phi^n(Z')$ is conjugated inside Z' . Notice that neither $\langle z \rangle$ nor Z' intersect \mathcal{N}' trivially, as this case was covered by point (1). Then set $\mathcal{N} = \langle\langle \mathcal{N}', z^t, tZ' \rangle\rangle$, which is preserved by ϕ^n for any choice of $t \in \mathbb{N}_{>0}$. Again, one can choose t in such a way that $g_0 \notin \mathcal{N}$, and that G/\mathcal{N} is a colourable short HHG.

(3) Finally, suppose that there exists $k_0 \in \mathbb{N}_{>0}$ such that, for every $k \geq k_0$, $\phi^k(P)$ is virtually \mathbb{Z}^2 and is virtually contained in some edge group E_k . Without loss of generality, we can replace ϕ by ϕ^{k_0} and assume that $k_0 = 1$. By definition, each E_k is the intersection of two vertex stabilisers, so the stability of product regions implies that $\phi^k(P)$ is actually a finite-index subgroup of E_k .

Notice that central cyclic directions and clean intersection imply that each E_k is isomorphic to \mathbb{Z}^2 , so that $\phi(E_k)$ is also Abelian. Moreover, $\phi(E_k)$ is a finite-index overgroup of $\phi^{k+1}(P)$, which is virtually \mathbb{Z}^2 and is contained inside E_{k+1} ; hence, by stability of product regions applied to the central extension $\phi(E_k)$, we must have that $\phi(E_k) \leq E_{k+1}$.

As in point (2), the existence of finitely many edge groups allows one to find $n \in \mathbb{N}_{>0}$ and an edge group $E' = \text{Stab}_G(v) \cap \text{Stab}_G(w)$ such that both $\phi^n(P)$ and $\phi^n(E')$ are conjugated inside E' . As G has clean intersections, $E' = \langle Z_v, Z_w \rangle$, so for every $t \in \mathbb{N}_{>0}$ the subgroup $tE' = \langle tZ_v, tZ_w \rangle$ is preserved by ϕ^n (up to conjugation). Then set $\mathcal{N} = \langle\langle \mathcal{N}, z^t, tE' \rangle\rangle$, which is preserved by ϕ^n (notice that $\phi^n(z^t)$ is conjugated inside tE'). If, say, Z_v already intersected \mathcal{N}' , we choose t in such a way that $tZ_v \leq \mathcal{N}'$, and similarly for Z_w . Then again a suitable choice of t grants the required properties of the quotient. \square

5.5. Hopf property for admissible HNN extensions. Before focusing on Artin groups, we provide an easy example of how one can establish the Hopf property for certain short HHG, which will serve as a blueprint for many arguments in the next Section. The additional hypotheses we will assume on the short HHG rule out the difficulties that appear for Artin groups. We start with a general Lemma.

Lemma 5.17. *Let G be a finitely generated group, and let $\phi: G \rightarrow G$ be a surjective homomorphism. Then ϕ induces an automorphism of the abelianisation G^{ab} of G .*

Proof. The commutator subgroup, which is the kernel of the abelianisation, is preserved by any endomorphism, so ϕ descends to a self-epimorphism of G^{ab} . Furthermore, the latter is a finitely generated Abelian group, so it is Hopfian (as it is residually finite, for instance). \square

Proposition 5.18. *Let $(G, \overline{X}, \mathcal{W})$ be a short HHG with central cyclic directions. Suppose that:*

- *All vertex stabilisers are conjugate;*
- *The image of a vertex stabiliser in G^{ab} has torsion-free rank at least 3;*
- *The image of a cyclic direction in G^{ab} has infinite order.*

Then G is Hopfian.

Proof. Let $\phi: G \rightarrow G$ be a surjective homomorphism, and let $g_0 \in G - \{1\}$. In order to apply our criterion, Lemma 5.3, we must produce a Hopfian quotient H of G , such that the image of g_0 is non-trivial, and that some iterate of ϕ induces a self-epimorphism of H .

Let $P = \text{Stab}_G(v)$ for some $v \in \overline{X}^{(0)}$, and let $\langle z \rangle$ be its cyclic direction. As z has infinite order in G^{ab} , $\langle \phi(z) \rangle$ is infinite cyclic; furthermore, since the image of P in G^{ab} has torsion-free rank at least three, the same must be true for $\phi(P)$, which means that the latter must be a \mathbb{Z} -central extension of a non-elementary hyperbolic group. In turn, Lemma 5.5 implies that $\langle \phi(z) \rangle$ is virtually contained in some cyclic direction, which is conjugate to $\langle z \rangle$ by assumption. Hence, up to post-composing ϕ by an inner automorphism, we can assume the existence of some $n \in \mathbb{N}_{>0}$ and $m \in \mathbb{Z} - \{0\}$ such that $\phi(z^n) = z^m$.

We now claim that n divides m . Let $\overline{\phi}: G^{ab} \rightarrow G^{ab}$ be the induced map, and let \overline{z} be the image of z in the abelianisation. Since $\overline{\phi}$ maps the torsion subgroup to itself, and the latter does not contain \overline{z} , up to taking a further quotient we can assume that $G^{ab} \cong \mathbb{Z}^r$ for some $r \in \mathbb{N}_{>0}$. Choose a base e_1, \dots, e_r of \mathbb{Z}^r such that $\overline{z} = ke_1$ for some $k \in \mathbb{N}_{>0}$. Then $mke_1 = m\overline{z} = \overline{\phi}(n\overline{z}) \leq nk\mathbb{Z}^r$, hence n divides m .

The above discussion implies that $\phi(z^n) \leq \langle z^n \rangle$, so ϕ induces an automorphism of $H := G / \langle\langle z^{nK} \rangle\rangle$ for every $K \in \mathbb{N}$. Since there is a unique conjugacy class of cyclic directions, Corollary 4.23 grants the existence of some K such that H is hyperbolic (hence Hopfian), and g_0 survives in H , as required. \square

Just to give a concrete example for the Proposition:

Example 5.19. Let $G_0 = \mathbb{Z} \times F_4$, where the centre is generated by $z \in \mathbb{Z}$ and the free factor has a basis e_1, \dots, e_4 . Let $H_1 = \langle z, e_1 \rangle$ and $H_2 = \langle z, e_2 \rangle$. Consider the HNN extension $G: G_0 *_\psi$, where $\psi: H_1 \rightarrow H_2$ maps z to e_2 and e_1 to z . The HNN extension is an admissible graph of groups, in the sense of [CK02, Definition 3.1], and therefore it admits a short HHG structure with support graph the Bass-Serre tree, as argued in [Man24, Subsection 2.3.3]. In particular, all vertex stabilisers are conjugates of G_0 . Furthermore, G^{ab} is the free Abelian group generated by the images of z, e_3, e_4, t ; in particular, the image of G_0 in the abelianisation is isomorphic to \mathbb{Z}^3 , and the image of z has infinite order. Hence G satisfies the requirements of Proposition 5.18, and is therefore Hopfian.

6. HOPF PROPERTY FOR ARTIN GROUPS

The goal of this section is to prove Theorem 6.6 about the Hopf property for Artin groups. We will use our "Dehn filling" quotients to do so. We start with some basic definitions.

Definition 6.1. Let Γ be a simplicial graph, with edge set E , and let $m: E \rightarrow \mathbb{N}_{\geq 2}$ be a labelling of the edges of Γ with positive integers greater than or equal to 2.

Recall that the *Artin group* A_Γ is the group with the following presentation:

$$A_\Gamma = \langle \Gamma^{(0)} \mid \text{prod}(a, b, m_{ab}) = \text{prod}(b, a, m_{ab}) \forall \{a, b\} \in E \rangle,$$

where $m_{ab} := m(\{a, b\})$ and $\text{prod}(u, v, n)$ denotes the prefix of length n of the infinite alternating word $uvuvuv \dots$.

An Artin group A_Γ is of *large type* if all edge labels are at least 3, and it is of *hyperbolic type* if, for every triangle with vertices a, b, c inside Γ , the sum of the inverses of the edge labels is strictly less than one:

$$\frac{1}{m_{ab}} + \frac{1}{m_{bc}} + \frac{1}{m_{ac}} < 1.$$

Definition 6.2 (Odd components). Given a labelled graph Γ , we say that two vertices a, b are in the same *odd component* if there exists a combinatorial path between them, all whose edges have odd labels (i.e. if a and b are in the same connected component after we remove all even edges). The *odd component graph*, denoted Γ_{OC} , is the simplicial graph whose vertices are odd components, and where two odd components C, C' are adjacent if there exist vertices $a \in C, a' \in C'$ which are joined by an even edge.

Remark 6.3. A result of Paris [Par97, Corollary 4.2] states that two standard generators are conjugate if and only if they lie in the same odd component, and this is why odd components will be relevant.

Remark 6.4 (Short HHG structure). In [HMS22], the authors produce a combinatorial HHG structure (X, \mathcal{W}) for A_Γ which, as noticed in [Man24], is a short HHG structure. Here we point out some of its properties.

- Fix a representative vertex for every odd component, and let V be the union of such vertices. Let \mathcal{H} be the collection of all cyclic subgroups generated by either a standard generator $s \in V$, or by the centre z_{ab} of a standard dihedral subgroup $A_{ab} := \langle a, b \rangle$, for every two Γ -adjacent a, b . For every $H \in \mathcal{H}$, let $N(H)$ be its normaliser. Then X is a blowup of the *commutation graph* Y , whose vertices are the cosets of the $N(H)$, and two cosets $gN(H)$ and $hN(H')$ are adjacent if and only if gHg^{-1} commutes with $hH'h^{-1}$. By [HMS22, Lemma 3.10], Y is connected if the defining graph Γ is.
- By construction, Y is bipartite, as two different conjugates of the standard generators never commute, nor do two different conjugates of centres of dihedral subgroups. This gives an A_Γ -invariant colouring of Y .
- In [HMS22, Lemmas 2.27 and 2.28], the authors describe $N(H)$ as follows: if $H = \langle a \rangle$ then $N(H) = C(a)$ is the centraliser of a , which is the direct product of $\langle a \rangle$ and a finitely generated free group; if $H = \langle z_{ab} \rangle$ then $N(H) = A_{ab}$ is the corresponding dihedral subgroup, which is a central extension with kernel $\langle z_{ab} \rangle$ and quotient a free product of two cyclic groups.

For the next definition, recall that a *leaf* of a simplicial graph is an edge with an endpoint of valence one, which we call the *tip* of the leaf. A leaf in Γ is said to be even or odd according to its edge label.

Definition 6.5 (Hanging component). A *hanging component* is an odd component C which is a leaf of Γ_{OC} . A hanging component C is *broad* if $|C| > 1$. A hanging component is a *needle* if $C = \{v\}$ is a single vertex, and v is a (necessarily even) leaf of Γ .

We devote the rest of the Section to the proof of the following:

Theorem 6.6. *Let A_Γ be an Artin group of large and hyperbolic type, such that every hanging component is either broad or a needle. Then every surjective homomorphism $\phi: A_\Gamma \rightarrow A_\Gamma$ is an isomorphism.*

Outline of the proof. Firstly, it is enough to consider the case where Γ is connected, since a free product of (finitely many, finitely generated) Hopfian groups is Hopfian by [DN70, Theorem 1.1]. Thus we are in the setting of Subsection 6, and A_Γ is a short HHG. We can also assume that $|\Gamma| \geq 3$, as \mathbb{Z} and dihedral Artin groups are known to be residually finite and therefore Hopfian. The proof is then split between Propositions 6.11 to 6.15, depending on the number of odd components of Γ . \square




FIGURE 10. There are three hanging components in this graph: $\{d, e\}$ is broad, $\{f\}$ is a needle, and $\{g\}$ is what we forbid in Theorem 6.6.

6.1. Pruning leaves. Recall that, if $p \in \Gamma^{(0)}$ is the tip of an even leaf, then its centraliser is the \mathbb{Z}^2 subgroup generated by p and the centre of the dihedral A_{pq} corresponding to the leaf (see e.g. [CMV23, Corollary 34]). In particular $C(p) \leq A_{pq}$, so the product region associated to p is somewhat redundant. This is made clearer in the next Lemma.

Lemma 6.7. *Let A_Γ be an Artin group of large hyperbolic type. Suppose that Γ is connected and has at least three vertices. There exists a short HHG structure (A_Γ, \overline{X}) , where \overline{X} is the full subgraph of the commutation graph whose vertices are cosets of normalisers of*

- *centres of standard dihedral parabolics, or*
- *cyclic subgroups generated by standard generators which are not the tips of even leaves.*

Proof. Let Y be the commutation graph, let p be the tip of an even leaf $\{p, q\}$ of Γ . For every $g \in A_\Gamma$, the coset $gN(\langle p \rangle)$ is only adjacent to gA_{pq} in Y , and is therefore a vertex of valence one of the commutation graph. Now let \overline{X} be the full, A_Γ -invariant subgraph of Y defined above, which is still triangle- and square-free, as so is Y , and none of its connected components is a point.

Now, by Proposition 2.16, A_Γ admits blowup materials with support graph Y , which we can restrict to \overline{X} by forgetting the data associated to the cosets $gN(\langle p \rangle)$. It is easily seen that the restriction gives blowup materials for A_Γ , as all the requirements of Definition 2.14 are already satisfied in the bigger graph Y . The only non-trivial observation is that point (4) still holds. Indeed, A_Γ is weakly hyperbolic relative to the collection

$$\{A_{ab}\}_{\{a,b\} \in \Gamma^{(1)}} \cup \{N(c)\}_{c \in \Gamma^{(0)}}.$$

However $N(\langle p \rangle)$ is contained inside A_{pq} , so A_Γ is also weakly hyperbolic relative to

$$\{A_{ab}\}_{\{a,b\} \in \Gamma^{(1)}} \cup \{N(c)\}_{c \in \Gamma^{(0)}, |\text{Lk}_\Gamma(c)| > 1}.$$

Then Theorem 2.15 yields the required short HHG structure (A_Γ, \overline{X}) . \square

6.2. Some properties to check. We now argue that the short structure (A_Γ, \overline{X}) defined above fits the framework of Subsection 5.4. For the rest of the Section, by *vertex stabiliser* we will always mean the stabiliser of a vertex of \overline{X} , with respect to the action of A_Γ . Firstly, by inspection of vertex stabilisers, we see that cyclic directions are central, Definition 5.10. Next, an easy observation, which we prove for completeness:

Lemma 6.8. (A_Γ, \overline{X}) has clean intersections, in the sense of Definition 5.12.

Proof. Let $\{v, w\}$ be an edge of \overline{X} . Up to the action of the group, we can assume that $v = A_{ab}$ is a standard dihedral and $w = C(a)$ is the centraliser of a . To prove that $C(a) \cap A_{ab} = \langle a, z_{ab} \rangle$, it is enough to notice that $(C(a) \cap A_{ab})/\langle a \rangle$ must centralise the non-trivial projection of z_{ab} to the free group $C(a)/\langle a \rangle$. \square

Finally, we move to stability of product regions:

Lemma 6.9. (A_Γ, \overline{X}) has stable product regions, in the sense of Definition 5.13.

Proof. Let $H \leq A_\Gamma$ be a subgroup isomorphic to a \mathbb{Z} -central extension of the form $1 \rightarrow Z \rightarrow H \rightarrow K \rightarrow 1$, and suppose that H is virtually contained in $\text{Stab}_{A_\Gamma}(v)$ for some $v \in \overline{X}^{(0)}$.

(1) First, we assume that K is infinite, and we want to show that $\text{Stab}_{A_\Gamma}(v)$ actually contains H . Indeed, up to conjugation, $\text{Stab}_{A_\Gamma}(v)$ is either a standard dihedral A_{ab} or the centraliser $C(a)$ of a standard generator a . In the first case, H is contained in $\text{Stab}_{A_\Gamma}(v)$ because parabolics are root-closed by [CMV23, Theorem D]. In the second case, any element h of H is contained in some subgroup H' of H isomorphic to \mathbb{Z}^2 , since A_Γ is torsion-free (parabolics being root-closed implies this), and H' has a finite-index subgroup H'_0 contained in $C(a)$. We have that H'_0 needs to contain a non-trivial power a^k of a , for otherwise it would embed in the free group $C(a)/\langle a \rangle$. Since H' is Abelian, it is contained in the centraliser $C(a^k)$. By [Par97, Corollary 5.3], this coincides with $C(a)$, so that H' , whence h , is contained in $C(a)$.

Now assume in addition that K is not virtually cyclic, so that Lemma 5.5 tells us that Z is virtually contained in Z_v , and we claim that $Z \leq Z_v$. If $\text{Stab}_{A_\Gamma}(v) = C(a)$ then $H_1 := H/(\langle a \rangle \cap H)$ embeds in a free group and is not virtually cyclic (as otherwise H would be virtually \mathbb{Z}^2). This means that H_1 must have trivial centre, that is, Z must be contained in $\langle a \rangle$. If instead $\text{Stab}_{A_\Gamma}(v) = A_{ab}$, then [MV24, Remark 3.6.(2)] tells us that a proper root of z_{ab} has cyclic centraliser, and then again we must have that $Z \leq \langle z_{ab} \rangle$.

(2) Finally, suppose that K is finite, and we claim that there exist $I \in \mathbb{N}_{>0}$, only depending on Γ , such that H has an index- $\leq I$ subgroup contained in $\text{Stab}_{A_\Gamma}(v)$. As the ambient group A_Γ is torsion-free, H must be infinite cyclic, say generated by g (see e.g. [Mac96, Lemma 3.2]). If $\text{Stab}_{A_\Gamma}(v)$ is conjugate to a dihedral parabolic, then $g \in \text{Stab}_{A_\Gamma}(v)$, because parabolics are root-closed; hence we are left to consider the case where $\text{Stab}_{A_\Gamma}(v) = C(a)$ is the centraliser of some standard generator a . Let $n \in \mathbb{N}_{>0}$ be such that g^n is contained in $C(a)$, which in turn means that $a \in C(g^n)$. Notice that, if $C(g^n)$ coincides with $C(g)$, then $g \in C(a)$ and we are done; so suppose that this is not the case. Centralisers of elements of large-type Artin groups are analysed in detail in [MV24, Section 3], see in particular [MV24, Remark 3.6]. An inspection of all the various possibilities reveals that, if $C(g^n)$ strictly contains $C(g)$, then g lies in a conjugate of a dihedral subgroup A_{bc} , and g^n belongs to the centre of such conjugate. For simplicity, we can assume that g is contained in A_{bc} , as opposed to a conjugate. We now argue that in this case $C(g^n) = A_{bc}$ coincides with $C(g^I)$ for some uniform I , which suffices for our purposes. We have that g maps to a torsion element of $A_{bc}/Z(A_{bc})$, which has bounded torsion as it is a free product of cyclic groups. Therefore, a uniform power of g maps to the trivial element of $A_{bc}/Z(A_{bc})$, that is, said uniform power is contained in the centre of A_{bc} and its centraliser is A_{bc} , as required. \square

6.3. Proof of Theorem 6.6. We finally move to the core of the argument, which we split into three subcases, depending on whether Γ has one, two, or at least three odd components. We start with an observation.

Remark 6.10 (Abelianisation of an Artin group). Let A_Γ be an Artin group. The abelianisation A_Γ^{ab} of A_Γ is the free Abelian group with one generator for every odd component, and the abelianisation map sends each standard generator to its component. In particular, both standard generators and centres of dihedrals have non-trivial image in the abelianisation, so Lemma 5.17 implies that, for every epimorphism $\phi: A_\Gamma \rightarrow A_\Gamma$, their ϕ -images must have infinite order. Now, let P be a vertex stabiliser, and we look at its image inside A_Γ^{ab} . If P is conjugate to a standard dihedral subgroup A_{bc} , then its image has rank 1 if b and c are in the same odd component, and 2 otherwise. If instead P is conjugate to the centraliser $C(a)$ of some standard generator, then the rank of $C(a)$ in the abelianisation is

- 2 if a belongs to a hanging component;
- at least 3 otherwise.

This is because $C(a)$ contains some conjugate of z_{bc} for every dihedral A_{bc} where b is in the same odd component as a and c is not.

6.3.1. One odd component.

Proposition 6.11. *Let A_Γ be a large Artin group of hyperbolic type. Assume further that Γ is a connected graph on at least three vertices, and has a single odd component. Then A_Γ is Hopfian.*

Proof. Let ϕ be an epimorphism, and let $g_0 \in \ker \phi - \{1\}$. Our goal is to produce a Hopfian quotient $A_\gamma \rightarrow G$ where the image of g_0 is non-trivial, and such that ϕ induces a map on G . Then we will conclude by Lemma 5.3. The crucial feature of this case is that all standard generators are conjugate. This means that, given any $a \in \Gamma^{(0)}$, removing the class of a makes the product region graph $\mathcal{PR}(A_\gamma)$ discrete.

Since $\phi^n(a)$ has infinite order for every $n \in \mathbb{N}$, $\phi^n(C(a))$ is a \mathbb{Z} -central extension, say with base B_n . There are three possibilities, **A**, **B**, and **C** below, depending on the isomorphism type of B_n . The second scenario is furthermore split into two possibilities, **B1** and **B2**.

A. Suppose first that B_n is always non-elementary. In this case $C(a)$ is always restrained with respect to ϕ , and by Remark 5.16 we get a Hopfian quotient G satisfying the requirements.

B1. Suppose now that B_n is definitely virtually cyclic. By Lemma 6.9, $\phi^n(C(a))$ must be contained in some $P = \text{Stab}_{A_\Gamma}(v)$. If $\phi^n(C(a))$ is always virtually contained in an edge group then $C(a)$ is always restrained, and we conclude as above. Otherwise, up to replacing ϕ by an iterated, assume that $\phi(C(a))$ is virtually \mathbb{Z}^2 but is not contained in an edge group, so there exists a unique P containing $\phi(C(a))$.

Suppose first that P is conjugate to $C(a)$. Up to composition with an inner automorphism, we can actually assume that $H := \phi(C(a)) \subseteq C(a)$. We have $\phi(H) \leq \phi(C(a)) \leq H$. For $\psi = \phi|_H: H \rightarrow H$ and H_0 a finite-index subgroup of H isomorphic to \mathbb{Z}^2 let $H_1 = \bigcap_{i \geq 0} \psi^{-i}(H_0)$. It is readily checked that H_1 is ψ -invariant, whence ϕ -invariant. We claim that H_1 is a finite-index subgroup of H_0 , which in turn implies that H_1 is isomorphic to \mathbb{Z}^2 . This is because the index of each $\psi^{-i}(H_0)$ in H is at most the index of H_0 in H , and since there are only finitely many subgroups of H of index bounded by a given constant, the intersection defining H_1 is actually equal to a finite intersection of finite-index subgroups. Furthermore, H_1 needs to contain a non-trivial power a^k (for otherwise it would embed in the hyperbolic group $C(a)/\langle a \rangle$); pick $k > 0$ minimal. We can find g such that $\{g, a^k\}$ is a basis of $H_1 \cong \mathbb{Z}^2$. Now, every virtually cyclic subgroup of the free group $C(a)/\langle a \rangle$ is cyclic, so g has no hidden symmetries in $C(a)$ by [Man24, Remark 5.4]. Then by [Man24, Proposition 5.2], there exist $p \in \mathbb{N}_{>0}$ and a short HHG structure in which $g' := g^p$ spans a cyclic direction.

We have that g^{kp} lies in H_1 , and together with a^{k^2p} it generates kpH_1 . Since $\phi(H_1) \leq H_1$, we have that $\phi(\langle\langle g^{Mkp}, a^{Mk^2p} \rangle\rangle) \leq \langle\langle g^{Mkp}, a^{Mk^2p} \rangle\rangle$ for any integer M . By Theorem 4.1, for a suitable M the quotient $G := A_\Gamma / \langle\langle g^{Mkp}, a^{Mk^2p} \rangle\rangle$ is a colourable short HHG, where the image of g_0 is non-trivial. Moreover, as $\mathcal{PR}(G)$ is obtained by removing the class of a from $\mathcal{PR}(A_\Gamma)$, G is Hopfian by Lemma 5.11, so G satisfies the requirements.

B2. In the same setting as above, suppose now that $\phi(C(a))$ is conjugated into a (unique) dihedral parabolic P . Up to conjugation, we can assume that $P = A_{bc}$ for some Γ -adjacent generators b, c . We now argue that $\phi(A_\Gamma)$ is contained in A_{bc} , thus contradicting the surjectivity of ϕ . Say that $C(a)$ is the stabiliser of the vertex $v \in \bar{X}$. For any vertex w of \bar{X} adjacent to v we have that $P_w = \text{Stab}_{A_\Gamma}(w)$ is a conjugate of a dihedral group, and the centraliser of a generator z_w of its centre. Since z_w has infinite-order image in the abelianisation, it maps under ϕ to a non-trivial element of A_{bc} .

Notice that $\phi(z_w)$ is not conjugate to any power of either b or c . If this was not true, say without loss of generality that $\phi(z_w) = b^k$ for some non-trivial $k \in \mathbb{Z}$. Then $\phi(a)$ would lie in the centraliser of b^k inside A_{bc} , which is the edge group $\langle b, z_{bc} \rangle$. If $\phi(a^r) \in \langle b \rangle$ for some $r \geq 0$, then the whole $\phi(C(a))$ would lie in $\langle b, z_{bc} \rangle$, as every element of $\phi(C(a))$ must centralise $\phi(a)$. If instead $\phi(a) \notin \langle b \rangle$, then

$\langle \phi(a), \phi(z_w) \rangle \cong \mathbb{Z}^2$ would coarsely coincide with $\phi(C(a))$. In both cases, one would contradict the fact that $\phi(C(a))$ is not virtually contained in an edge group.

As $\phi(z_w)$ is not conjugate to any power of either b or c , by [MV24, Remark 3.6]) its centraliser is entirely contained in A_{bc} . Hence $\phi(P_w)$, which centralises $\phi(z_w)$, is contained in A_{bc} . Now, given a vertex v' of \overline{X} adjacent to w , $P_{v'} = \text{Stab}_{A_\Gamma}(v')$ contains z_w , so $\phi(P_{v'})$ contains an element in A_{bc} which is not conjugate to any power of either b or c . Moreover, $P_{v'}$ is a conjugate of $C(a)$, and as such its image must be contained in a conjugate of A_{bc} . By [CMV23, Theorem A], the intersection of A_{bc} and one of its conjugates is either trivial, the whole A_{bc} , or is conjugate to either $\langle b \rangle$ or $\langle c \rangle$; hence we must have that in fact $\phi(P_{v'})$ is contained in A_{bc} . We can then proceed inductively on the distance in \overline{X} from a , and as the support graph \overline{X} is connected (see Subsection 6) we eventually get that ϕ maps every vertex stabiliser Q into A_{bc} , as required.

C. Suppose finally that B_n is definitely finite. Up to replacing ϕ by a power, we can assume that $\phi(C(a))$ is virtually infinite cyclic. We first claim that, for every two Γ -adjacent vertices c and d , $\phi(a)$ and $\phi(z_{cd})$ have a non-trivial common power. Indeed, first consider $b \in \text{Lk}_\gamma(a)$. The fact that $\phi(C(a))$ is virtually \mathbb{Z} ensures that $\phi(a)$ and $\phi(z_{ab})$ have a non-trivial common power. In turn, as $\phi(C(b))$ is conjugate to $\phi(C(a))$ and contains $\phi(z_{ab})$, the same must hold for $\phi(b)$ and $\phi(z_{ab})$. Iterating this procedure, we eventually get that, for each z_{cd} , there exist N_{cd}, M_{cd} such that $\phi(z_{cd}^{N_{cd}}) = \phi(a^{M_{cd}})$, that is, $\phi(z_{cd}^{N_{cd}} a^{-M_{cd}}) = 1$. We can in fact take multiples to ensure that all M_{cd} coincide, say $M_{cd} = M$. For $\mathcal{N}_K = \langle\langle z_{cd}^{KN_{cd}} a^{-KM} \rangle\rangle$ we have $\phi(\mathcal{N}_K) \leq \mathcal{N}_K$, so ϕ induces a homomorphism of $A(K) = A_\Gamma / \mathcal{N}_K$. We claim that we can choose K so that $A(K)$ is a \mathbb{Z} -central extension of a hyperbolic group (hence Hopfian by Theorem 5.7) and the image of g_0 in it is non-trivial. If this is true then $G = A(K)$ satisfies all requirements.

In order to do so, we consider the auxiliary group $A'(K) = A_\Gamma / \langle\langle \mathcal{N}_K, a^{KM} \rangle\rangle$. Since $\langle\langle \mathcal{N}_K, a^{KM} \rangle\rangle = \langle\langle z_{cd}^{KN_{cd}}, a^{KM} \rangle\rangle$, by Theorem 4.1, for suitable values of K , $A'(K)$ is a HHG; furthermore, from the description of its structure, it is clear that $A'(K)$ has bounded orthogonality, and is therefore a hyperbolic group by e.g. [BHS21, Corollary 2.14]. By Corollary 3.40 we can further arrange that the image of g_0 is non-trivial in $A'(K)$, and therefore also in its extension $A(K)$. We are left to prove that the natural projection $A(K) \rightarrow A'(K)$ is a \mathbb{Z} -central extension. By construction, the kernel is normally generated by the image of a^{KM} in $A(K)$, so in turn it suffices to prove that said image commutes with a generating set of $A(K)$. The reason for this is that $A(K)$ is obtained from A_Γ by imposing the relations $z_{cd}^{KN_{cd}} = a^{KM}$. As z_{cd} commutes with c , the relations make a^{KM} commute with all the standard generators, as required. \square

6.3.2. At least three odd components. We consider now the general case, postponing the study of Artin groups with two odd components as it is more involved and reuses some techniques from this paragraph.

Proposition 6.12. *Let A_Γ be a large Artin group of hyperbolic type, where Γ is a connected graph on at least three vertices. Suppose that Γ has at least three odd components, and every hanging component is either broad or a needle. Then A_Γ is Hopfian.*

Proof. As usual, given an epimorphism ϕ , we want to find a collection of always restrained vertex stabilisers whose removal makes $\mathcal{PG}(A_\Gamma, \overline{X})$ discrete, and then

conclude by Lemma 5.15. Let C_1, \dots, C_k be the hanging components of Γ , and let Γ_{core} be the subgraph of Γ spanned by all other odd components.

In view of Remark 6.10, if P is a vertex stabiliser for the action on \overline{X} , then its image in the abelianisation of A_Γ has rank at least 3 if and only if P is conjugate to $C(a)$ for some $a \in \Gamma_{core}$. For any such P , $\phi(P)$ must be a \mathbb{Z} -central extension of a non-elementary hyperbolic group, as its projection to A_Γ^{ab} must have rank at least 3; moreover $\phi(P) \leq Q$ for some stabiliser Q , and the projection of Q to A_Γ^{ab} must have rank at least 3 as well. This, together with stability of product regions, implies that, for every $a \in \Gamma_{core}^{(0)}$, $\phi(a)$ is conjugated into $\langle b \rangle$ for some $b \in \Gamma_{core}^{(0)}$. Notice that the above argument also tells us that $C(a)$ is always restrained for every $a \in \Gamma_{core}^{(0)}$.

Now consider the retraction

$$r := A_\Gamma \rightarrow A_{C_1} * \dots * A_{C_k},$$

defined by mapping every generator in Γ_{core} to the identity and all other generators to themselves. This map is well-defined, because any edge connecting Γ_{core} to any hanging component is even; furthermore, ϕ induces a self-epimorphism $\overline{\phi}$ of the quotient, because $\ker(r)$ is normally generated by $\Gamma_{core}^{(0)}$. Notice that, by Proposition 6.11 plus the aforementioned [DN70, Theorem 1.1], the quotient is Hopfian, so $\overline{\phi}$ is an isomorphism. Now, r is injective on every dihedral A_{bc} where b and c belong to the same hanging component, thus $\phi(A_{bc})$ must be a \mathbb{Z} -central extension of a non-elementary hyperbolic group because its r -projection $r(\phi(A_{bc})) = \overline{\phi}(r(A_{bc}))$ is again isomorphic to A_{bc} . As this argument works for every iterate of ϕ , we get that every dihedral in a hanging component is always restrained.

At this point, removing centralisers of core vertices and dihedrals contained in hanging components is still not enough to make the product region graph discrete, so we need to find more always restrained stabilisers. By inspection of $\mathcal{PR}(A_\Gamma, \overline{X})$, it is enough to prove the following:

Claim 6.13. *Let $a, b, c \in \Gamma^{(0)}$ be such that $a \in \Gamma_{core}^{(0)}$, b, c are adjacent vertices in the same hanging component, and a is adjacent to b . Then A_{ab} is always restrained.*

Proof of Claim 6.13. By contradiction, up to passing to an iterated of ϕ , assume that $\phi(A_{ab})$ is virtually \mathbb{Z}^2 and is not contained in an edge group (notice that $\phi(A_{ab})$ cannot be virtually cyclic, as the image of A_{ab} in the abelianisation is isomorphic to \mathbb{Z}^2). By stability of product regions, there exists some vertex stabiliser P , say with centre generated by z , such that $\phi(A_{ab}) \leq P$. Since we already know that $\phi(a)$ is contained in the centre of some vertex stabiliser, we must have that $\phi(a) \in \langle z \rangle$. Moreover, $\phi(b)$ cannot belong to any edge group $E \leq P$, as otherwise $\phi(A_{ab}) = \langle \phi(a), \phi(b) \rangle \leq E$ as well. Thus $\phi(z_{cb})$, which commutes with $\phi(b)$ and is contained in a centre, must belong to $\langle z \rangle$. This contradicts the fact that a and z_{bc} are non-commensurable in the abelianisation, so their ϕ -images cannot lie in the same cyclic subgroup. ■

The proof of Proposition 6.12 is now done. □

6.3.3. Two odd components. We finally move to the case where Γ has two odd components, which are therefore both hanging components. According to their shapes, we split the Proposition into two sub-lemmas.

Proposition 6.14. *Let A_Γ be a large Artin group of hyperbolic type, where Γ is a connected graph on at least three vertices. Suppose that Γ has two odd components, one of which is a needle. Then A_Γ is Hopfian.*

Proof. Pick a self-epimorphism ϕ , and let $g_0 \in \ker \phi - \{1\}$. Again, the goal is to find a quotient $A_\gamma \rightarrow G$ satisfying the requirements of Lemma 5.3. Let C and C' be the two odd components, and assume without loss of generality that $|C| > 1$. As C' is a needle, the only vertex of C' , call it b , is adjacent to a unique vertex $a \in C$. Now, if $C(a)$ is always restrained, then Lemma 5.15 produces a Hopfian quotient with the required properties (notice that, as b is a leaf, its centraliser is not a vertex of \overline{X} , so removing $C(a)$ makes $\mathcal{PR}(A_\Gamma, \overline{X})$ discrete).

Thus suppose that $C(a)$ is not always restrained. Up to replacing ϕ with a power, we can assume that $\phi(C(a))$ is virtually \mathbb{Z}^2 but not contained in an edge group (notice that $\phi(C(a))$ cannot be virtually cyclic, as the image of $C(a)$ in the abelianisation is isomorphic to \mathbb{Z}^2). By stability of product regions, $\phi(C(a))$ must be contained in a unique vertex stabiliser P . Furthermore, we claim that the whole component C is mapped inside P . Indeed, let $a' \in C$ be connected to a by an odd edge. Then $\phi(A_{aa'}) \leq P$, as it must centralise $\phi(z_{aa'})$ which lies in a non-edge \mathbb{Z}^2 subgroup of P . But then, since a is conjugate to a' by an element of $A_{aa'}$, we get that $\phi(C(a')) \leq P$ as well, and it is again a virtually \mathbb{Z}^2 subgroup not contained in any edge group. As C is an odd component, any two vertices are connected by a path with odd labels, so we get that $\phi(a'') \in P$ for every $a'' \in C$.

Similarly, notice that $\phi(A_{ab}) \leq P$ as well, as it must be contained in the centraliser of $\phi(z_{ab}) \leq \phi(C(a))$. But this violates surjectivity, as then $\phi(A_\Gamma)$ is totally contained in P . \square

Proposition 6.15. *Let A_Γ be a large Artin group of hyperbolic type, where Γ is a connected graph on at least three vertices. Suppose that Γ has two odd components, which are both broad. Then A_Γ is Hopfian.*

Proof. Let $C = \{a_1, \dots, a_k\}$ and $C' = \{b_1, \dots, b_r\}$ be the odd components, let ϕ be an epimorphism, and let $g_0 \in \ker \phi - \{1\}$. Again, if both $C(a_1)$ and $C(b_1)$ are always restrained, then we can invoke Remark 5.16 and conclude. So suppose that $C(a_1)$ is not always restrained, so there exists a vertex stabiliser P , say with centre generated by z , such that $\phi(C(a_1))$ is a virtually \mathbb{Z}^2 subgroup of P not contained in any edge group. Arguing as in Proposition 6.14, one gets that $\phi(C) \leq P$, and $\phi(A_{a_i b_j}) \leq P$ for every i, j such that a_i and b_j are joined by an even edge. We now consider three scenarios, A, B, and C, depending on the shape of $\phi(C(b_1))$.

A. Suppose first that $\phi(C(b_1))$ is a \mathbb{Z} -central extension of a non-elementary hyperbolic group, so b_1 must be sent inside some centre. As standard generators have primitive image in the abelianisation while centres of dihedrals do not, we can assume up to conjugation that $\phi(b_1) = b_1$ or $\phi(b_1) = a_1$. In the former case, ϕ preserves the normal closure of C' , so it induces a self-map $\overline{\phi}$ of the quotient A_C , obtained by retracting onto C . Then one can run the proof of Proposition 6.12 verbatim, to get that A_Γ has “enough” always restrained vertex stabilisers.

In the latter case, $\phi^2(C(b_1)) \leq \phi(C(a_1)) \leq P$. It now matters what P is, taking into account that it cannot be an odd dihedral as $\phi(C(a_1))$ must map to \mathbb{Z}^2 in the abelianisation. If P is a conjugate of $C(b_1)$ then we can argue as above, with ϕ^2 replacing ϕ . If instead P is conjugated to $Q \in \{C(a_1), A_{a_i b_j}\}$ by some $g \in G$, up to

composing ϕ with the conjugation by g^{-1} we can assume that $\phi(Q) \leq Q$. Then we have that:

- $\phi^2(C) \leq \phi(Q) \leq Q$;
- $\phi^2(C') \leq Q$, because $\phi^2(C(b_1)) \leq \phi(C(a_1))$ is contained in a \mathbb{Z}^2 subgroup which is not an edge group.

Hence $\phi^2(A_\Gamma) \leq Q$, violating surjectivity.

B. Suppose now that $\phi(C(b_1))$ is virtually \mathbb{Z}^2 , but not contained in an edge group. Since $\phi(C(b_1))$ contains $\phi(z_{a_1 b_1})$, which belongs to the non-edge \mathbb{Z}^2 subgroup $\phi(C(a_1))$ of P , we must have that $\phi(C(b_1)) \leq P$. Then again, using that $\phi(C(b_1))$ is not contained in an edge group, we get that $\phi(C') \leq P$, which combined with $\phi(C) \leq P$ violates surjectivity.

C. Suppose finally that $\phi(C(b_1))$ is a finite-index subgroup of an edge group. As above, since $\phi(C(b_1))$ contains $\phi(z_{a_1 b_1})$ we must have that $\phi(C(b_1)) \leq P$. We now consider the possible conjugacy types of P .

- Suppose first that P is conjugate to $C(a_1)$, and up to composing ϕ with an inner automorphism we can indeed assume that $P = C(a_1)$. Then $\phi^2(C(b)) \leq \phi(C(a)) \leq P$ is a non-edge, virtually \mathbb{Z}^2 subgroup, and again this implies that $\phi(C') \leq P$, contradicting surjectivity.
- Suppose now that P is conjugate to some dihedral, which must be of the form $A_{a_i b_j}$ (again because any other dihedral has cyclic image in A_Γ^{ab}). Pick $b_2 \in C'$ which is connected to b_1 by an odd edge. Then $\phi(z_{b_1 b_2}) \in P$, and as parabolics are root closed we must have that $\phi(b_1 b_2) \in P$. Hence $\phi(C(b_2))$, which is obtained by conjugating $\phi(C(b_1))$ by a power of $\phi(b_1 b_2)$, is also nested in the subgroup P . Proceeding this way, we eventually get that $\phi(C') \leq P$, and once more ϕ could not be surjective.
- The only case left is when P is conjugate to $C(b_1)$, and again we can indeed assume that $P = C(b_1)$ up to composing ϕ with a conjugation. Say $C(b_1) = \text{Stab}_{A_\Gamma}(v)$ for some $v \in \overline{X}^{(0)}$. As $\phi(C(b_1))$ lies in some edge group, there must be some $w \in \text{Lk}_{\overline{X}}(v)$ such that, if we set $Q = \text{Stab}_{A_\Gamma}(w)$, then $\phi(C(b_1)) \leq C(b_1) \cap Q$. Since \overline{X} is bipartite, Q must be conjugate to a dihedral, so the same trick as above shows that $\phi(C') \leq Q$. If we show that $\phi(Q) \leq C(b_1)$ then $\phi^2(C') \leq C(b_1)$, and $\phi^2(C) \leq \phi(C(b_1)) \leq C(b_1)$. This would then again contradict surjectivity.

To prove that $\phi(Q) \leq C(b_1)$, we first notice that the image of Q in the abelianisation must have rank 2, or it could not contain $\phi(C(b_1))$; hence there exists i, j such that Q is conjugate to $A_{a_i b_j}$. In turn, since $\phi(A_{a_i b_j}) \leq C(b_1)$, there must be some P' , which is a conjugate of $C(b_1)$, such that $\phi(Q) \leq P'$. But $\phi(Q)$ contains $\phi^2(C(b_1))$, which is virtually \mathbb{Z}^2 and lies inside $\phi(C(b_1)) \leq C(b_1)$. As \overline{X} is bipartite, any two different conjugates of $C(b_1)$ intersect along a virtually cyclic subgroup, so we must have that $P' = C(b_1)$, as required.

This concludes the proof of Proposition 6.15, and in turn of Theorem 6.6. \square

6.4. Comments and previous results.

Remark 6.16 (Generic Artin groups are Hopfian). In [GV25], Goldsborough and Vaskou devised a model of random Artin groups, where, given a complete graph on n vertices, each edge label is chosen with uniform probability from the set

$\{\infty, 2, \dots, f(n)\}$, for some non-decreasing divergent function $f: \mathbb{N} \rightarrow \mathbb{N}$. A property of Artin groups is *generic* if there exists a function $f_0: \mathbb{N} \rightarrow \mathbb{N}$ such that, for every choice of function $f \geq f_0$, the property holds with probability approaching 1 as $n \rightarrow +\infty$. In the same paper, the authors prove that the class of extra-large Artin groups is generic. Moreover, let $p(n) = e(n)/f(n)$, where $e(n)$ is the cardinality of odd numbers in the set $\{\infty, 2, \dots, f(n)\}$. For any choice of f , the probability that a random Artin group has a single odd component is the same as the probability that a random (unlabelled) graph on n vertices, where each edge exists with probability $p(n)$, is connected. Such probability is known to approach 1 as $n \rightarrow +\infty$ (see e.g. [ER61]); hence, as our Theorem 6.6 applies to XL Artin groups with a single odd component, we get that a generic Artin group is Hopfian, thus proving Corollary C from the Introduction.

Remark 6.17 (Other generic classes). In [BMV24], Blufstein, Martin, and Vaskou established the Hopf property for large hyperbolic type Artin groups which are either *free-of-infinity* (the defining graph is complete) or *XXXL* (all edge labels are at least 6). Both classes are generic, in the sense of Remark 6.16. Figure 11 provides examples of Artin groups covered by our result, by theirs, and by none of them. We stress that the techniques from [BMV24] are very different from ours, as they involve a full description of all homomorphisms between groups in their families; this also allows them to determine when such groups are co-Hopfian (every *injective* homomorphism is an isomorphism).




FIGURE 11. From left to right, an Artin group which is Hopfian by [BMV24] (it is XXXL), an Artin group which is Hopfian by our Theorem 6.6 (it has a single odd component), and an Artin group which is not covered by either methods. Notice that none of these Artin groups is known to be residually finite (see Remark 6.19 below).

Remark 6.18 (Explicit residual hyperbolicity). As a special case of Corollary 4.24, if all hyperbolic groups are residually finite, then every Artin group A_Γ of large and hyperbolic type is residually finite, hence Hopfian. On the one hand, it is common belief that there exists a non-residually finite hyperbolic group. On the other, the proof of Corollary 4.23 shows that it would suffice that “enough” hyperbolic quotients of A_Γ are residually finite, namely those with the following presentation, for a suitable choice of N :

$$\langle \Gamma^{(0)} \mid \forall c \in \Gamma^{(0)}, \forall \{a, b\} \in E, \text{prod}(a, b, m_{ab}) = \text{prod}(b, a, m_{ab}), c^N = (ab)^{m_{ab}N} = 1 \rangle.$$

An intermediate quotient, falling in the family of *Shephard groups*, is the following:

$$S_\Gamma^N := \langle \Gamma^{(0)} \mid \forall c \in \Gamma^{(0)}, \forall \{a, b\} \in E, \text{prod}(a, b, m_{ab}) = \text{prod}(b, a, m_{ab}), c^N = 1 \rangle.$$

The latter groups were studied in [Gol24], where the author proved that, if Γ is triangle-free and large type, then S_Γ^N is residually finite for all large enough N (see [Gol24, Corollary F]). In turn, this is used in [Gol24, Theorem G] to prove that the corresponding Artin group A_Γ is residually finite (notice that this now follows easily from the fact that every $g \in A_\Gamma$ survives in some S_Γ^N , as a consequence of Corollary 4.23).

Remark 6.19 (Overview on residual finiteness for Artin groups). Few classes of Artin groups are known to be residually finite, among which:

- Artin groups whose defining graph is triangle-free and contains no square whose edge labels are all 2 (this is the full statement of the aforementioned [Gol24, Theorem G]);
- even Artin groups of FC type (including RAAGs, see [BGMPP19]);
- spherical Artin groups (because they are linear, by e.g. [CW02] or [Dig03]);
- certain 2-dimensional Artin groups, including most Artin groups on three generators and even XXXL Artin groups on graphs admitting a “partial orientation” (see [Jan22]);
- “forests” of residually finite parabolic subgroups (see [BGJP18] for details).

Remarkably, none of the above families is generic in the sense of [GV25].

Remark 6.20 (Equational Noetherianity). Barak [Bar24] recently established that, if G is a colourable, *strictly acylindrical* HHG, then G is *equationally Noetherian*, hence Hopfian by e.g. [GH19, Corollary 3.14 and Theorem D]. For our purposes, the only consequence of strict acylindricity to keep in mind is that, for every $U \in \mathfrak{S}$, its stabiliser acts acylindrically on \mathcal{CU} , as a corollary of [BBFS19, Theorem 6.3]; in particular, if \mathcal{CU} is unbounded, then $\text{Stab}_G(U)$ is either virtually cyclic or acylindrically hyperbolic.

In our setting, it is clear that an Artin group A_Γ of large and hyperbolic type is colourable, as it is colourable as a short HHG. However, if Γ has at least one edge, then the short HHG structure is not strictly acylindrical, as the centraliser of a vertex has infinite centre and is therefore not acylindrically hyperbolic.

7. QUOTIENTS OF THE FIVE-HOLED SPHERE MAPPING CLASS GROUP

Let $S = S_{0,5}$ be a five-punctured sphere, and let $\mathcal{MCG}^\pm(S)$ be its extended mapping class group, which is a short HHG as pointed out in [Man24, Subsection 2.3.1]. Our last theorem proves that almost every way of adding finitely many relations to $\mathcal{MCG}^\pm(S)$ results in a hierarchically hyperbolic group, up to stabilising by taking sufficient powers of the relators. This provides an almost complete answer to [MS25, Question 3], in the case of a five-punctured sphere.

Theorem 7.1. *Let $S = S_{0,5}$, and let $g_1, \dots, g_l \in \mathcal{MCG}^\pm(S)$. Suppose the collection $\{g_i\}$ contains no pseudo-Anosovs with hidden symmetries, in the sense of Definition 3.44. Then there exists $N \in \mathbb{N} - \{0\}$ such that, for all $K_1, \dots, K_l \in \mathbb{Z} - \{0\}$, we have that $\mathcal{MCG}^\pm(S)/\langle\langle\{g_i^{K_i N}\}\rangle\rangle$ is hierarchically hyperbolic.*

Proof. Fix a curve γ on S . Given $g_1, \dots, g_l \in \mathcal{MCG}^\pm(S)$, it suffices to prove the statement replacing each g_i with g_i^K for some $K \neq 0$. Therefore, as a consequence of Nielsen-Thurston classification (see [FM12, Corollary 13.3]), up to conjugation we can assume that each g_i is of one of the following types:

- (1) A power of the Dehn twist τ_γ around γ ;

- (2) A power of a multitwist with associated multicurve $\{\gamma, \beta\}$, for a fixed curve β which is disjoint from γ ;
- (3) A partial pseudo-Anosov without hidden symmetries, supported on the unique component Y of $S - \gamma$ which is homeomorphic to $S_{0,4}$;
- (4) A pseudo-Anosov.

In particular, every g_i has infinite order. Up to taking further powers, we can make the following modifications to the collection of elements under consideration:

- Suppose that two elements are commensurable up to conjugation, say for simplicity of notation g_1 and g_2 . Then, up to taking a common power of all the elements in our collection, we can find $g \in \mathcal{MCG}^\pm(S)$ and integers a, b such that $g_1 = g^a$ and g_2 is conjugate to g^b . We can then replace g_1 and g_2 with $g^{\gcd(|a|, |b|)}$, without changing the subgroup normally generated by the collection.
- Similarly, suppose that two elements of type (3) are not commensurable in $\mathcal{MCG}^\pm(S)$ but have commensurable images in $\mathcal{MCG}^\pm(Y)$. Again, denote them by g_1 and g_2 . Up to taking powers of all elements, and replacing g_2 with a conjugate, this means that there exists $g \in \mathcal{MCG}^\pm(S)$ of type (3) such that $g_1 = g^a$ and $g_2 = g^b \tau_\gamma^c$, for some integers a, b, c . The subgroup $\langle g_1, g_2 \rangle$ coincides with $\langle \tau_\gamma^r, g^s \tau_\gamma^t \rangle$ for some integers r, s, t , since any subgroup of $\mathbb{Z}^2 \cong \langle \tau_\gamma, g \rangle$ is of that form. We can then replace g_1, g_2 with τ_γ^r and $g^s \tau_\gamma^t$, and then possibly repeat the procedure in the previous bullet if the new collection contains commensurable elements.

Assume first that there are no elements of type (2). By [Man24, Proposition 5.1], there exists a colourable short HHG structure for $\mathcal{MCG}^\pm(S)$ where every g_i of type (4) generates a cyclic direction, up to taking a suitable power. Furthermore, by [Man24, Proposition 5.2], we can also assume that every element of type (3) has a power that generates a cyclic direction. Then Theorem 4.1 ensures that we can find some integer $N > 0$ such that $\mathcal{MCG}^\pm(S) / \langle\langle \{g_i^{K_i N}\} \rangle\rangle$ is a short HHG, and in particular hierarchically hyperbolic.

Now assume that there are elements of type (2), say $g_1 = \tau_\gamma^{a_1} \tau_\beta^{b_1}, \dots, g_r = \tau_\gamma^{a_r} \tau_\beta^{b_r}$, where all a_i and b_i are non-zero integers. Consider some choice of integers $N \neq 0$ and K_i . If some element, say g_{r+1} , is of the form τ_γ^c , set

$$d = \gcd(a_1 K_1, b_1 K_1, \dots, a_r K_r, b_r K_r, c K_{r+1}),$$

otherwise set

$$d = \gcd(a_1 K_1, b_1 K_1, \dots, a_r K_r, b_r K_r).$$

Since τ_β is conjugate to τ_γ , we have that $\mathcal{N} = \langle\langle g_i^{N K_i} \rangle\rangle \leq \langle\langle \tau_\gamma^{dN}, g_{r+1}^{N K_{r+1}}, \dots, g_l^{N K_l} \rangle\rangle$.

Let G_1 and G_2 be the quotients of $\mathcal{MCG}^\pm(S)$ by the first and second group respectively, so that, similarly to above, G_2 is a short HHG if we choose N suitably. In fact, in the short HHG structure for $\mathcal{MCG}^\pm(S)$ given by [Man24, Propositions 5.1 and 5.2], the cyclic directions are conjugate to powers of $\tau_\gamma, g_{r+1}, \dots, g_l$, so G_2 is hyperbolic because we modded out all cyclic directions. Also, by the inclusion of kernels, we have a surjective homomorphism $\phi: G_1 \rightarrow G_2$, and the kernel is normally generated by the image τ of τ_γ^{dN} in G_1 .

We claim that τ is a central element of G_1 . To see this, using that all Dehn twists are conjugate (in the mapping class groups of $S_{0,4}$, as well as of $S_{0,5}$), we first notice that there exists an integer d' such that \mathcal{N} contains $\tau_\gamma^{dN} \tau_\alpha^{d'}$, for any curve

α disjoint from γ . In terms of G_1 , this implies that τ coincides with the image of $\tau_\alpha^{-d'}$, and in particular τ commutes with the image of the half-twist around α or any curve disjoint from α . Varying α , the images of said half-twists generate G_1 , so that τ commutes with a generating set of G_1 , and is therefore central. Hence G_1 is a cyclic central extension of a hyperbolic group, and therefore it is a HHG by [HRSS25, Corollary 4.3], as required. \square

Example 7.2. We now exhibit a partial pseudo-Anosov without hidden symmetries, in order to show that the hypothesis of Theorem 7.1 is not vacuous. Let γ be a curve on $S = S_{0,5}$, let α, β be disjoint from γ , and let $\tau_\alpha, \tau_\beta, \tau_\gamma$ be the associated Dehn twists. Let Y be the connected component of $S - \gamma$ which is homeomorphic to $S_{0,4}$, let p be the puncture of Y coming from γ , and let q be the puncture which is separated from p by both α and β . By e.g. [FM12, Proposition 3.19], the quotient $\text{Stab}_{\mathcal{MCG}^\pm(S)}(\gamma)/\langle \tau_\gamma \rangle$ is an index two overgroup of $\mathcal{MCG}^\pm(Y, \{p\})$, that is, the subgroup of $\mathcal{MCG}^\pm(Y)$ spanned by all elements that fix the puncture p . Let $i \in \text{Stab}_{\mathcal{MCG}^\pm(S)}(\gamma)$ be an-orientation-preserving mapping class whose image $\bar{i} \in \mathcal{MCG}^\pm(Y, \{p\})$ swaps α and β (for example, i could be a rotation of angle π around the axis passing through q). Let $g = \tau_\gamma \tau_\alpha \tau_\beta^{-1}$, which has a hidden symmetry because $\bar{i} g \bar{i}^{-1} = \bar{\tau}_\beta \bar{\tau}_\alpha^{-1} = \bar{g}^{-1}$ while $i g i^{-1} = \tau_\gamma \tau_\beta \tau_\alpha^{-1} = \tau_\gamma^2 g^{-1}$.

This example should make the terminology "hidden symmetry" clearer, as an axis for \bar{g} is "flipped" by the conjugation by \bar{i} but this symmetry is not witnessed in the extension. Furthermore, the problem persists even if we only consider the orientable mapping class group $\mathcal{MCG}(S)$, as both i and g are orientation-preserving.

REFERENCES

- [Bar24] Ohana Barak. On equational noetherianity of colorable hierarchically hyperbolic groups. *arXiv preprint arXiv:2410.00977*, 2024.
- [BBF15] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publ. Math. Inst. Hautes Études Sci.*, 122:1–64, 2015.
- [BBFS19] Mladen Bestvina, Ken Bromberg, Koji Fujiwara, and Alessandro Sisto. Acylindrical actions on projection complexes. *Enseign. Math.*, 65(1-2):1–32, 2019.
- [BGJP18] Rubén Blasco-García, Arye Juhász, and Luis Paris. Note on the residual finiteness of Artin groups. *J. Group Theory*, 21(3):531–537, 2018.
- [BGMPP19] Rubén Blasco-García, Conchita Martínez-Pérez, and Luis Paris. Poly-freeness of even Artin groups of FC type. *Groups Geom. Dyn.*, 13(1):309–325, 2019.
- [BHMS24] Jason Behrstock, Mark Hagen, Alexandre Martin, and Alessandro Sisto. A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups. *J. Topol.*, 17(3):Paper No. e12351, 94, 2024.
- [BHS17a] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proc. Lond. Math. Soc. (3)*, 114(5):890–926, 2017.
- [BHS17b] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.
- [BHS19] Jason Behrstock, Mark Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces II: Combination theorems and the distance formula. *Pacific J. Math.*, 299(2):257–338, 2019.
- [BHS21] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.*, 170(5):909–996, 2021.
- [BMR97] Gilbert Baumslag, Alexei Myasnikov, and Vitaly Roman’kov. Two theorems about equationally Noetherian groups. *J. Algebra*, 194(2):654–664, 1997.
- [BMV24] Martín Blufstein, Alexandre Martin, and Nicolas Vaskou. Homomorphisms between large-type Artin groups, 2024.

- [CK02] C. B. Croke and B. Kleiner. The geodesic flow of a nonpositively curved graph manifold. *Geom. Funct. Anal.*, 12(3):479–545, 2002.
- [CMV23] María Cumplido, Alexandre Martin, and Nicolas Vaskou. Parabolic subgroups of large-type Artin groups. *Math. Proc. Cambridge Philos. Soc.*, 174(2):393–414, 2023.
- [CW02] Arjeh M. Cohen and David B. Wales. Linearity of Artin groups of finite type. *Israel J. Math.*, 131:101–123, 2002.
- [Dah18] François Dahmani. The normal closure of big Dehn twists and plate spinning with rotating families. *Geom. Topol.*, 22(7):4113–4144, 2018.
- [DGO17] F. Dahmani, V. Guirardel, and D. Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Amer. Math. Soc.*, 245(1156):v+152, 2017.
- [DHS17] Matthew Gentry Durham, Mark F. Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.*, 21(6):3659–3758, 2017.
- [DHS21] François Dahmani, Mark F. Hagen, and Alessandro Sisto. Dehn filling Dehn twists. *Proc. Roy. Soc. Edinburgh Sect. A*, 151(1):28–51, 2021.
- [Dig03] François Digne. On the linearity of Artin braid groups. *J. Algebra*, 268(1):39–57, 2003.
- [DN70] I. M. S. Dey and Hanna Neumann. The Hopf property of free products. *Math. Z.*, 117:325–339, 1970.
- [Dru09] Cornelia Druţu. Relatively hyperbolic groups: geometry and quasi-isometric invariance. *Comment. Math. Helv.*, 84(3):503–546, 2009.
- [ER61] P. Erdős and A. Rényi. On the evolution of random graphs. *Bull. Inst. Internat. Statist.*, 38:343–347, 1961.
- [FFMS25] Francesco Fournier-Facio, Giorgio Mangioni, and Alessandro Sisto. Bounded cohomology, quotient extensions, and hierarchical hyperbolicity. Preprint, arXiv:2505.20462 [math.GR] (2025), 2025.
- [FM12] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.
- [GH19] D. Groves and M. Hull. Homomorphisms to acylindrically hyperbolic groups I: Equationally noetherian groups and families. *Trans. Amer. Math. Soc.*, 372(10):7141–7190, 2019.
- [GM08] Daniel Groves and Jason Fox Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.
- [Gol] Katherine Goldman, editor. *AimPL: Geometry and topology of Artin groups*. Available at <http://aimpl.org/geomartingp>.
- [Gol24] Katherine Goldman. 2-dimensional Shephard groups, 2024.
- [GV25] Antoine Goldsborough and Nicolas Vaskou. Random Artin groups. *Algebr. Geom. Topol.*, 25(3):1523–1544, 2025.
- [HMS22] Mark Hagen, Alexandre Martin, and Alessandro Sisto. Extra-large type Artin groups are hierarchically hyperbolic. *Mathematische Annalen*, pages 1–72, 2022.
- [HMS23] Mark Hagen, Giorgio Mangioni, and Alessandro Sisto. A combinatorial structure for many hierarchically hyperbolic spaces, 2023.
- [HRSS25] Mark Hagen, Jacob Russell, Alessandro Sisto, and Davide Spriano. Equivariant hierarchically hyperbolic structures for 3-manifold groups via quasimorphisms. *Ann. Inst. Fourier (Grenoble)*, 75(2):769–828, 2025.
- [HW09] G. Christopher Hruska and Daniel T. Wise. Packing subgroups in relatively hyperbolic groups. *Geom. Topol.*, 13(4):1945–1988, 2009.
- [Jan22] Kasia Jankiewicz. Residual finiteness of certain 2-dimensional Artin groups. *Adv. Math.*, 405:Paper No. 108487, 37, 2022.
- [Mac96] Dugald Macpherson. Permutation groups whose subgroups have just finitely many orbits. In *Ordered groups and infinite permutation groups*, volume 354 of *Math. Appl.*, pages 221–229. Kluwer Acad. Publ., Dordrecht, 1996.
- [Man23] Giorgio Mangioni. Random quotients of mapping class groups are quasi-isometrically rigid, 2023.
- [Man24] Giorgio Mangioni. Short hierarchically hyperbolic groups I: uncountably many coarse median structures. *arXiv preprint arXiv:2410.09232*, 2024.

- [MS25] Giorgio Mangioni and Alessandro Sisto. Rigidity of mapping class groups mod powers of twists. *Proceedings of the Royal Society of Edinburgh: Section A Mathematics*, pages 1–71, 2025.
- [MV24] Alexandre Martin and Nicolas Vaskou. Characterising large-type Artin groups. *Bull. Lond. Math. Soc.*, 56(11):3346–3357, 2024.
- [NR97] Walter D. Neumann and Lawrence Reeves. Central extensions of word hyperbolic groups. *Ann. of Math. (2)*, 145(1):183–192, 1997.
- [Osi07] Denis V. Osin. Peripheral fillings of relatively hyperbolic groups. *Invent. Math.*, 167(2):295–326, 2007.
- [Par97] Luis Paris. Parabolic subgroups of Artin groups. *J. Algebra*, 196(2):369–399, 1997.
- [Rus22] Jacob Russell. From hierarchical to relative hyperbolicity. *Int. Math. Res. Not. IMRN*, 1:575–624, 2022.
- [She52] G. C. Shephard. Regular complex polytopes. *Proceedings of the London Mathematical Society*, s3-2(1):82–97, 1952.
- [Thu97] William P. Thurston. *Three-dimensional geometry and topology. Vol. 1*, volume 35 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1997. Edited by Silvio Levy.
- [Val21] Motiejus Valiunas. Acylindrical hyperbolicity of groups acting on quasi-median graphs and equations in graph products. *Groups Geom. Dyn.*, 15(1):143–195, 2021.
- [WR19] Richard Weidmann and Cornelius Reinfeldt. Makanin-Razborov diagrams for hyperbolic groups. *Ann. Math. Blaise Pascal*, 26(2):119–208, 2019.
- [WS24] Henry Wilton and Alessandro Sisto. The congruence subgroup property for mapping class groups and the residual finiteness of hyperbolic groups, 2024.

(GIORGIO MANGIONI) MAXWELL INSTITUTE AND DEPARTMENT OF MATHEMATICS, HERIOT-WATT UNIVERSITY, EDINBURGH, UK
Email address: `gm2070@hw.ac.uk`

(ALESSANDRO SISTO) MAXWELL INSTITUTE AND DEPARTMENT OF MATHEMATICS, HERIOT-WATT UNIVERSITY, EDINBURGH, UK
Email address: `a.sisto@hw.ac.uk`