



Extra-large type Artin groups are hierarchically hyperbolic

Mark Hagen¹ · Alexandre Martin² · Alessandro Sisto²

Received: 30 September 2021 / Accepted: 14 November 2022 / Published online: 10 December 2022
© The Author(s) 2022

Abstract

We show that Artin groups of extra-large type, and more generally Artin groups of large and hyperbolic type, are hierarchically hyperbolic. This implies in particular that these groups have finite asymptotic dimension and uniform exponential growth. We prove these results by using a combinatorial approach to hierarchical hyperbolicity, via the action of these groups on a new complex that is quasi-isometric both to the coned-off Deligne complex introduced by Martin–Przytycki and to a generalisation due to Morris–Wright of the graph of irreducible parabolic subgroups of finite type introduced by Cumplido–Gebhardt–González–Meneses–Wiest.

Contents

Introduction	868
Hyperbolic features of Artin groups	868
A curve complex for Artin groups	870
Combinatorial approach to HHS, and strategy of proof	871
Outline of the paper	872
1 A combinatorial criterion for hierarchical hyperbolicity	873
2 Background on Artin groups and Deligne complexes	874
2.1 Artin groups	874
2.2 Structure of dihedral Artin groups	875
2.3 The modified Deligne complex	877
2.4 Standard trees and the coned-off Deligne complex	878
2.5 Links of vertices	883
3 The commutation graph of an Artin group	885
3.1 The commutation graph	885

✉ Mark Hagen
markfhagen@posteo.net

Alexandre Martin
alexandre.martin@hw.ac.uk

Alessandro Sisto
a.sisto@hw.ac.uk

¹ School of Mathematics, University of Bristol, Bristol, UK

² Department of Mathematics, Heriot-Watt University, Edinburgh, UK

3.2	The graph of proper irreducible parabolic subgroups of finite type	891
3.3	Adjacency in the commutation graph and intersection of neighbourhoods of cosets	895
4	The blown-up commutation graph	897
4.1	Blow-up data	898
4.2	The blown-up commutation graph	903
4.3	Simplices of the blow-up and their links	904
4.4	From maximal simplices to elements of A_Γ	907
5	The augmented complex	911
5.1	Construction of the augmented complex	911
5.2	Fullness of links	915
5.3	Maps between augmented complexes	916
6	The geometry of the augmented complex	921
6.1	Hyperbolicity of the augmented links	921
6.2	Quasi-isometric embeddings of augmented links	923
6.3	Obtaining hierarchical hyperbolicity	933
	References	937

Introduction

Hyperbolic features of Artin groups

The geometry of Artin groups has seen an explosion of results in the last decade. While Artin groups remain in general much less understood than their Coxeter relatives, a driving theme behind current research has been to show that these groups are as well-behaved as Coxeter groups, and this has indeed been verified for several classes of Artin groups. On the geometric side, a popular theme has been to understand the “hyperbolic features” of groups, and Artin groups are conjectured to display such hyperbolic features. This rather vague notion comes in many flavours, a first one being the notion of *acylindrically hyperbolic* group. Loosely speaking, such groups can be described as having “hyperbolic directions” (see [50] for the precise definition and its many consequences). It is conjectured that for an irreducible Artin group A_Γ , its central quotient $A_\Gamma/Z(A_\Gamma)$ is acylindrically hyperbolic.

This question has been answered positively for most standard classes of Artin groups, such as right-angled Artin groups [19], Artin groups of finite type [17], Artin groups of Euclidean type [16], Artin groups whose underlying presentation graph is not a join [22], and two-dimensional Artin groups [55].

While acylindrical hyperbolicity guarantees the existence of some hyperbolic directions, it does not provide much control on the overall geometry of the group. For instance, the free product $A * B$, where A, B are the worst infinite groups you can think of, is acylindrically hyperbolic.

A notion of non-positive curvature that provides a much stronger control over the coarse geometry of the group is the notion of *hierarchically hyperbolic group* (or HHG). This notion was introduced by Behrstock–Hagen–Sisto [8, 9] and inspired by work of Masur–Minsky [45, 46] as a framework that unifies and generalises the geometry of mapping class groups of hyperbolic surfaces and that of (all known) cocompactly cubulated groups [33]. The idea of hierarchical hyperbolicity is to describe the coarse geometry of a group/space using a “coordinate system” where the coordinates take values in various hyperbolic spaces.

Since this notion gives much better control than acylindrical hyperbolicity, it implies many results expected of non-positively curved groups that are not true for more general notions of non-positive curvature; for example, it gives quadratic isoperimetric inequality [9, 13], solubility of the word and conjugacy problem [9, 31], the Tits alternative [27, 28], finite asymptotic dimension [7], semi-hyperbolicity [29, 31], etc. In particular, since both braid groups and right-angled Artin groups belong to this family [8, 9], the following is a natural question that was for instance raised by Calvez–Wiest [18].

Question *Which Artin groups are hierarchically hyperbolic?*

The above question is also prefigured by a very prescient comment by Kapovich–Schupp [40, p. 155], which predates the notion of hierarchical hyperbolicity: “*It seems plausible that most Artin groups satisfy some sort of ‘nested’ version of the bounded coset penetration property...*”.

Building new examples of hierarchically hyperbolic groups is nontrivial; there are combination theorems [9, 11], and theorems about persistence of hierarchical hyperbolicity under various quotients [6, 7], but the examples coming from “nature” are dominated by compact special groups, mapping class groups, and fundamental groups of certain 3-manifolds. The preceding question has only been previously answered positively for right-angled Artin groups and braid groups.

In this article, we add a large class of Artin groups to the “naturally occurring” hierarchically hyperbolic groups:

Theorem A *Artin groups of large and hyperbolic type are hierarchically hyperbolic.*

Note that Artin groups of large and hyperbolic type contain in particular all Artin groups of extra-large type (see Sect. 2.1 for precise definitions of these classes). We refer to Theorem 6.15 for a description of the HH structure.

The following results are new for Artin groups of large and hyperbolic type, and follow from the cited results about hierarchically hyperbolic groups:

Corollary B *Let A_Γ be an Artin group of large and hyperbolic type, with Γ connected and not a single vertex. Then:*

- (1) A_Γ has finite asymptotic dimension [7, Thm. A],
- (2) A_Γ has uniform exponential growth [3, Cor. 1.3],
- (3) A finitely generated subgroup of A_Γ is stable if and only if the orbit maps to the coned-off Deligne complex are quasi-isometric embeddings [1, Thm. B].

Items 2 and 3 require a short argument to link the cited results to Theorem A; this appears at the end of Sect. 6.3. The coned-off Deligne complex—which is one of our most important tools—was defined and extensively studied in [44]. As a consequence of our construction, this complex is (equivariantly quasi-isometric to) the maximal hyperbolic space in our HHG structure, see Theorem 6.15.

Regarding asymptotic dimension, it follows from Theorem 6.15 and [7, Theorem 5.2, Corollary 3.3] that the asymptotic dimension of the coned-off Deligne complex \hat{D}_Γ is finite (with no explicit bound) and that the asymptotic dimension of A_Γ is at most $asdim(\hat{D}_\Gamma) + 3$. In view of this, it is natural to ask:

Question Given Γ , what is (a good bound on) the asymptotic dimension of the coned-off Deligne complex?

Note that our main Theorem also allows us to recover from a unified perspective several other known results for these groups, such as the above-mentioned Tits alternative (see [44, Corollary B]; this was proved in the extra-large case in [49]), the solubility of the conjugacy problem (also proved in [36, Corollary 1.3] for any two-dimensional Artin group), and control over their quasi-flats (see [35, Theorem 1.1] for the statement about two-dimensional Artin groups, and the closely-related [10, Theorem A] about HHGs). Moreover, we recover the fact that these groups are semi-hyperbolic, as they are coarsely Helly by [31, Thm. A, Cor. F], see also [29]. (The semi-hyperbolicity also follows from the fact that these groups are systolic [37], hence biautomatic [39, 53], hence semi-hyperbolic [4].)

A curve complex for Artin groups

There are earlier results in the direction of hierarchical hyperbolicity for Artin groups, see e.g. [18] for a detailed discussion. The first step in such a result is to introduce an analogue of the curve graph, to play the role of the “maximal” hyperbolic space in the HHS structure, i.e. a space quasi-isometric to the space obtained from the group by coning off the standard product regions. There are various significant results about actions of Artin groups on hyperbolic complexes, notably those of Calvez–Wiest for spherical type Artin groups [17], Calvez in the Euclidean-type case [16], and Martin–Przytycki in the two-dimensional case [44].

A candidate for such a curve complex for Artin groups of finite type has been proposed recently by Cumplido–Gebhardt–Gonzales–Meneses–Wiest [24], and generalised to all Artin groups by Morris–Wright [48] who studied this complex extensively for Artin groups of type FC. This is the notion of *graph of irreducible proper parabolic subgroups of finite type* of an Artin group.

It is conjectured that this complex is an infinite-diameter hyperbolic space (except in degenerate cases) for Artin groups of finite type [24, Conjecture 2.4] and type FC [48, Conjecture 5.6]. We prove the analogous result for Artin groups of large and hyperbolic type as a corollary of Proposition 3.16 and Theorem A.

Corollary C *Let A_Γ an Artin group of large and hyperbolic type on at least three generators. Then A_Γ admits an HHG structure whose maximal hyperbolic space is the graph of irreducible proper parabolic subgroups of finite type.*

Note that the action of A_Γ on this graph is acylindrical. This can be seen by combining the above corollary with [8, Theorem K], or by combining Proposition 3.16 with [44, Theorem A]. This corollary gives hope that this complex not only is hyperbolic for larger classes of Artin groups but also is the maximal hyperbolic space in an HHG structure.

Combinatorial approach to HHS, and strategy of proof

We will not strictly need the full definition of hierarchical hyperbolicity but, roughly, a hierarchically hyperbolic space (HHS) is a space \mathcal{X} that comes with a family of hyperbolic spaces and projections from \mathcal{X} to the various hyperbolic spaces satisfying various conditions [9, Definition 1.1]. Additionally, a group is hierarchically hyperbolic if it acts geometrically on an HHS preserving the HHS structure, in a way expressed most simply in [52, p. 4].

We will not need the full definition because we will verify a more combinatorial criterion from [6] for a group G to be hierarchically hyperbolic. We state this as Theorem 1.4, but for this discussion it suffices to know that criterion involves a hyperbolic simplicial complex X on which G acts, and maximal simplices of X are in bijection with the vertices of a graph quasi-isometric to G . The links of simplices in X give the hyperbolic spaces in the HHS structure, so the fine geometry of X is relevant for this purpose.

We now explain how to come up with a candidate X for a large-type Artin group A_Γ of hyperbolic type. The strategy makes sense more generally and potentially could be applied to other classes of groups as well. First, we need a little more discussion on hierarchical hyperbolicity. An HHS contains a family of so-called *standard product regions*, which are HHS themselves and (coarsely) split as a product of HHS. Moreover, the standard product regions are “arranged hyperbolicly” meaning that coning them off yields a hyperbolic space. (This is the “hierarchical structure” that gives the name.) In the context of the combinatorial criterion, said coned-off space is quasi-isometric to the hyperbolic simplicial complex X .

In view of this, a natural candidate for an action of a given group on a simplicial complex to which Theorem 1.4 can be applied is constructed as follows.

Step 1: Consider the family of subgroups that are maximal virtual products (these should give the standard product regions if the group is hierarchically hyperbolic). For our Artin groups, the tools to identify these subgroups come from the acylindricity of the action of A_Γ on its coned-off Deligne complex and $\text{CAT}(-1)$ geometry [44], see Sect. 2.4.

Step 2: Isolate the subgroups arising as minimal infinite intersections between such virtual product subgroups. Coarse intersections of standard product regions in an HHS are again coarse products of HHS, which are “smaller” than the original product subgroups. This should identify the simplest sub-HHS, if the group is to be hierarchically hyperbolic. For our Artin groups, here we have the cyclic subgroups conjugate to either the subgroup generated by a standard generators or centres of *dihedral parabolic* subgroups. See Sect. 3.1 for more discussion.

Step 3: To construct a hyperbolic space that the group acts on, it is natural to consider the graph encoding the intersections of candidates for the standard product regions. This graph is quasi-isometric to the *commutation graph* of the minimal subgroups above, see Sect. 3.1. This graph has a vertex for each such minimal subgroup and we put an edge when two such subgroups commute. In the case of right-angled Artin groups, we recover the extension graph, which is indeed hyperbolic [41]. In our case, there is a natural map between the commutation graph and the coned-off Deligne

complex. This map is a quasi-isometry with very nice local properties, and this is crucial for our arguments in Sect. 5.

Step 4: The commutation graph as above is expected to have the right coarse geometry, but not the right local geometry, because stabilisers of maximal simplices are usually infinite. We remedy this by a *blow-up* construction, where we replace each vertex of the commutation graph by a quasi-isometric copy of the corresponding subgroup, preserving the A_Γ -action. In our situation, the vertices of the commutation graph need to be blown up to quasi-lines. This creates a somewhat delicate situation where, for a certain cyclic subgroup H , we need to construct an action of $N(H)$ on a quasiline with certain properties. Remarkably, this difficulty is circumvented using quasimorphisms (which are related to actions on quasilines by [2]), in a similar way as in [32].

Once we have our hyperbolic complex X , we can build our quasi-isometry model W of A_Γ . As mentioned above, the vertex set of W is the set of maximal simplices of X . The adjacency relation is defined in such a way as to guarantee that orbit maps $A_\Gamma \rightarrow W$ are quasi-isometries (Sect. 5.1). The remaining work is to verify the technical conditions of Theorem 1.4, and here we once again rely on the CAT(−1) geometry of the coned-off Deligne complex.

In the end, we get an HHS structure on A_Γ where the maximal hyperbolic space—the HHS analogue of the curve graph in the mapping class group setting—is quasi-isometric to X and hence to the commutation graph.

As a final remark, our results do not cover the more general case of 2-dimensional Artin groups of hyperbolic type (that is, the case where edges with label 2 are allowed). In this case, the combinatorics of the commutation graph and associated spaces are more complex, with several statements in Sect. 3 and beyond needing more nuance since more cases have to be considered. However, we still believe that with more sophisticated arguments one can deal with the more general case using the same combinatorial HHS approach.

Outline of the paper

Section 1 contains the definitions and results from [6] that we will need. Section 2 contains background on Artin groups, and in particular the subclass of Artin groups considered in this paper, along with the coned-off Deligne complex. Section 3 is about the commutation graph, and also contains the discussion relating the commutation graph and the graph of irreducible parabolics. Section 4 contains the construction of the simplicial complex X —the blow-up of the commutation graph—along with some purely combinatorial facts about X , and the relationship between maximal simplices in X and coarse points in A_Γ . In Sect. 5, we define the graph of maximal simplices of X , prove that it is quasi-isometric to A_Γ , and study its combinatorial structure. Finally, in Sect. 6, we verify the remaining hypotheses of Theorem 1.4. The final subsection of Sect. 6 assembles the pieces into a proof of Theorem A and Corollary B.

1 A combinatorial criterion for hierarchical hyperbolicity

In this section, we recall the main combinatorial criterion introduced in [6] to show that a group is hierarchically hyperbolic by means of a map to a hyperbolic simplicial complex. We will not require these notions until Sect. 4, but we introduce them now to motivate the constructions in earlier sections. We refer the reader to [6, Section 1.2,1.5] for a more informal discussion of the various hypotheses in the definition of a combinatorial HHS.

Definition 1.1 (*X-graph, augmented complex*) Let X be a flag simplicial complex. An X -**graph** is a graph W whose vertex set is the set of maximal simplices of X . We say that two maximal simplices Δ, Δ' of X are W -**adjacent** if the corresponding vertices of W are adjacent in W .

We denote by X^{+W} the complex obtained from the 1-skeleton of X by adding edges between vertices that belong to W -adjacent maximal simplices. For a subcomplex X_0 of X , we denote by X_0^{+W} the full subcomplex of X^{+W} induced by X_0 . We refer to these various objects X^{+W}, X_0^{+W} , etc. as being **augmented complexes**.

Definition 1.2 (*Link, star, saturation of a simplex*) Let X be a flag simplicial complex, and let Δ be a simplex of X .

The **star** of Δ , denoted $St_X(\Delta)$, is the union of all the simplices of X containing Δ .

The **link** of Δ , denoted $Lk_X(\Delta)$, is the full subcomplex of $St_X(\Delta)$ induced by $St_X(\Delta) - \Delta$.

The **saturation** of Δ is

$$Sat_X(\Delta) = \bigcup_{\{\Sigma:Lk_X(\Delta)=Lk_X(\Sigma)\}} \Sigma^{(0)}.$$

The following is Definition 1.8 in [6].

Definition 1.3 (*Combinatorial HHS*) A *combinatorial HHS* is a pair (X, W) , where X is a simplicial complex and W is an X -graph, such that all of the following hold for some $\delta < \infty, n \in \mathbb{N}$:

- (I) If $\Delta_0, \dots, \Delta_m$ are simplices of X and $Lk_X(\Delta_i) \subsetneq Lk_X(\Delta_{i+1})$ for $0 \leq i \leq m - 1$, then $m \leq n$. This condition will be called *finite complexity*.
- (II) Let Δ be a non-maximal simplex of X . Let $\mathcal{C}(\Delta) = Lk_X(\Delta)^{+W}$ and let $Y_\Delta = (X^{(0)} - Sat(\Delta))^{+W}$. Then $\mathcal{C}(\Delta)$ is δ -hyperbolic and the inclusion $\mathcal{C}(\Delta) \hookrightarrow Y_\Delta$ is a (δ, δ) -quasi-isometric embedding. We will call this condition *hyperbolic links*.
- (III) Let Δ be a non-maximal simplex of X and let $v, w \in Lk_X(\Delta)$ be distinct non-adjacent vertices. Suppose that v, w are contained in W -adjacent maximal simplices of X . Then there exist maximal simplices Σ_v, Σ_w of $Lk_X(\Delta)$, respectively containing v, w , such that $\Sigma_v \star \Delta$ and $\Sigma_w \star \Delta$ are W -adjacent. We will call this condition *fullness of links*.
- (IV) Let Σ, Δ be non-maximal simplices of X such that there exists a non-maximal simplex Γ with $Lk_X(\Gamma) \subseteq Lk_X(\Sigma) \cap Lk_X(\Delta)$ and $diam(\mathcal{C}(\Gamma)) > \delta$. Then there

exists a non-maximal simplex $\Pi \subset \text{Lk}_X(\Sigma)$ such that $\text{Lk}_X(\Sigma \star \Pi) \subseteq \text{Lk}_X(\Delta)$ and all Γ as above satisfy $\text{Lk}_X(\Gamma) \subset \text{Lk}_X(\Sigma \star \Pi)$. We call this condition the *intersection condition*.

The following criterion is immediate from [6, Theorem 1.18, Remark 1.19] and the fact that hierarchical hyperbolicity is a quasi-isometry invariant property [9, Proposition 1.10].

Theorem 1.4 *Let (X, W) be a combinatorial HHS. Then any quasigeodesic space quasi-isometric to W is a hierarchically hyperbolic space. Moreover, suppose that the group G acts by simplicial automorphisms on X , and that the resulting G -action on the set of maximal simplices of X extends to a proper cobounded action of G on W . Suppose moreover that X contains finitely many G -orbits of subcomplexes of the form $\text{Lk}_X(\Delta)$, for Δ a simplex. Then G is a hierarchically hyperbolic group.*

(In the statement, the notion of properness used is sometimes called metric properness, and what we mean is that given a ball in X there are only finitely many elements of G that do not map the ball to a disjoint ball.)

Note that in addition to the properties of combinatorial HHS, Theorem 1.4 requires another condition, namely that there are finitely many orbits of links of simplices. For readers familiar with HHS terminology, we mention that this is so that the action of A_Γ on the index set of the eventual HHG structure is cofinite, as the elements of the index set correspond to the links of the non-maximal simplices.

We will apply the above theorem to an Artin group of large and hyperbolic type A_Γ by explicitly constructing (X, W) with W quasi-isometric to A_Γ , and then verifying each of the properties from Definition 1.3.

2 Background on Artin groups and Deligne complexes

2.1 Artin groups

A **presentation graph** is a finite simplicial graph Γ such that every edge between vertices $a, b \in V(\Gamma)$ is labelled by an integer $m_{ab} \geq 2$. The **Artin group** associated to Γ is the group A_Γ given by the following presentation:

$$A_\Gamma := \langle a \in V(\Gamma) \mid \underbrace{aba \cdots}_{m_{ab}} = \underbrace{bab \cdots}_{m_{ab}} \text{ whenever } a, b \text{ are connected by an edge of } \Gamma \rangle.$$

An Artin group is of **large type** if all coefficients m_{ab} are at least 3, and of **extra-large type** if they all are at least 4. An Artin group on two generators a, b with $m_{ab} < \infty$ is a **dihedral** Artin group.

Given an Artin group A_Γ , the **associated Coxeter group** W_Γ is obtained by further requiring that each generator $a \in V(\Gamma)$ satisfies the relation $a^2 = 1$. An Artin group is said to be of **hyperbolic type** if the associated Coxeter group is hyperbolic, and of **finite type** if the associated Coxeter group is finite.

For a (possibly empty) full subgraph $\Gamma' \subset \Gamma$, the subgroup of A_Γ generated by the vertices of Γ' is called a **standard parabolic subgroup**. Such a subgroup is isomorphic to the Artin group $A_{\Gamma'}$ by a result of Van der Lek [54], and moreover we have $A_{\Gamma_1} \cap A_{\Gamma_2} = A_{\Gamma_1 \cap \Gamma_2}$ for full subgraphs Γ_1, Γ_2 of Γ . Conjugates of standard parabolic subgroups are called **parabolic subgroups**.

2.2 Structure of dihedral Artin groups

Since dihedral Artin groups appear as stabilisers of vertices of dihedral type in the modified Deligne complex and its cone-off, we mention some structural results that will be needed in this article.

A dihedral Artin group on two standard generators a, b will be denoted A_{ab} for simplicity, even though the group depends on the coefficient m_{ab} . Dihedral Artin groups come into two types: If $m_{ab} = 2$, the group is a copy of \mathbb{Z}^2 . The rest of this subsection focuses on the structure of dihedral Artin groups with $m_{ab} \geq 3$. We start by recalling the following definition:

Definition 2.1 For a dihedral Artin group with $m_{ab} \geq 3$, the **Garside element** $\Delta_{ab} \in A_{ab}$ is defined as follows:

$$\Delta_{ab} := \underbrace{aba \cdots}_{m_{ab}} = \underbrace{bab \cdots}_{m_{ab}}.$$

Lemma 2.2 ([15]) *The centre of a dihedral Artin group A_{ab} with $m_{ab} \geq 3$ is infinite cyclic and generated by the element*

$$z_{ab} := \Delta_{ab} \text{ if } m_{ab} \text{ is even, and } z_{ab} := \Delta_{ab}^2 \text{ if } m_{ab} \text{ is odd.}$$

Lemma 2.3 *Let A_{ab} be a dihedral Artin group with $m_{ab} \geq 3$. The central quotient $A_{ab}/\langle z_{ab} \rangle$ is virtually a finitely generated non-trivial free group. In particular, A_{ab} contains a finite-index subgroup that splits as a direct product of the form $\langle z_{ab} \rangle \times K$, where K is a finitely generated free subgroup of A_{ab} .*

This virtual splitting is well-known to experts, see for instance [23]. We give here a geometric proof of this result that uses objects that will be needed in Sect. 5.

Definition 2.4 (*Atoms, left-weighted form*) An **atom** of A_{ab} is a strict subword of $\underbrace{aba \cdots}_{m_{ab}}$ or $\underbrace{bab \cdots}_{m_{ab}}$, that is, an alternating product of a and b with strictly fewer than m_{ab} letters. We denote by M the set of all atoms of A_{ab} . A product of the form $m_1 \cdots m_k$, with each $m_i \in M$, is said to be **left-weighted** if for each i the last letter of m_i coincides with the first letter of m_{i+1} .

It follows from the existence and uniqueness of Garside normal forms (see for instance [42]) that elements of the quotient $A_{ab}/\langle \Delta_{ab} \rangle$ are in bijection with left-weighted elements of the free monoid M^\bullet on M , where $\langle \Delta_{ab} \rangle$ acts on A_{ab} by right multiplication.




Fig. 1 A portion of the quasi-tree \mathcal{T}_{ab} for $m_{ab} = 3$. In red, the full subgraph spanned by the $\langle a \rangle$ -orbit of e . In blue, the full subgraph spanned by the $\langle b \rangle$ -orbit of e

Definition 2.5 We denote by \mathcal{T}_{ab} the full subgraph of the Cayley graph $\text{Cayley}(M^\bullet, M)$ spanned by left-weighted elements. The action of A_{ab} on $A_{ab}/\langle \Delta_{ab} \rangle$ by left multiplication induces an action of A_{ab} on \mathcal{T}_{ab} (seen as an unlabelled graph).

Since M^\bullet is a free monoid, the graph \mathcal{T}_{ab} is a quasi-tree, as already explained in [12]. More precisely, the flag completion of \mathcal{T}_{ab} has a structure of tree of simplices of dimension $m_{ab} - 1$ glued along vertices (see Fig. 1), where the simplices are either in the A_{ab} -orbit of the simplex spanned by

$$e, a, ab, \dots, \underbrace{aba \dots}_{m_{ab}-1},$$

or in the A_{ab} -orbit of the simplex spanned by

$$e, b, ba, \dots, \underbrace{bab \dots}_{m_{ab}-1}.$$

Proof of Lemma 2.3 The group A_{ab} acts by left multiplication on \mathcal{T}_{ab} . Since the element z_{ab} is central and is a power of Δ_{ab} , it follows that $\langle z_{ab} \rangle$ acts trivially on that graph, hence the quotient $A_{ab}/\langle z_{ab} \rangle$ acts by left multiplication on \mathcal{T}_{ab} . The action is cocompact and proper, thus $A_{ab}/\langle z_{ab} \rangle$ is virtually free, and hence A_{ab} is virtually a direct product of the form $\langle z_{ab} \rangle \times K$, where K is a finitely-generated free group. Note

that a and b define elements of infinite order in $A_{ab}/\langle z_{ab} \rangle$ since their orbits in \mathcal{T}_{ab} span embedded lines (see Fig. 1). □

We recall that the **syllabic length** of an element $g \in A_{ab}$ is the smallest non-negative integer n such that g can be written as a product of the form $g = x_1^{k_1} \cdots x_n^{k_n}$ with $k_i \in \mathbb{Z}$ and $x_i \in \{a, b\}$ for all $1 \leq i \leq n$.

Lemma 2.6 ([55, Proposition 4.6]) *Let g be an element of A_{ab} that can be written only with positive letters and that has syllabic length greater than 1. If $m_{ab} \geq 3$, then the syllabic length of g^n goes to infinity as n goes to infinity.*

Corollary 2.7 *In a dihedral Artin group A_{ab} with $m_{ab} \geq 3$, no non-trivial power of z_{ab} is equal to a power of a standard generator.*

Proof By Lemma 2.6, the syllabic length of z_{ab}^n explodes as n grows, while the syllabic length of a^n and b^n is always equal to one. □

Lemma 2.8 *In a dihedral Artin group A_{ab} with $m_{ab} \geq 3$, two distinct conjugates of standard generators never generate a subgroup isomorphic to \mathbb{Z}^2 .*

Proof Let us consider a dihedral Artin group A_{ab} with $m_{ab} \geq 3$. Up to conjugation, we can assume that the element a and a conjugate $x := gcg^{-1}$, with $c \in \{a, b\}$, commute. Since the centraliser of a in A_{ab} is $\langle a, z_{ab} \rangle$ by [23, Lemma 7], it follows that there exist integers ℓ, k such that $x = a^k z_{ab}^\ell$. We claim that necessarily $\ell = 0$. Indeed, if that were not the case, then the syllabic length of the powers of z_{ab}^ℓ would go to infinity by Lemma 2.6, and since a^k and z_{ab}^ℓ commute, so would the syllabic length of the powers of x (since the powers of a all have syllabic length 1). But since x is conjugate to a power of a generator, its powers have a uniformly bounded syllabic length, a contradiction. We thus have $x = a^k$. By using the homomorphism $A_{ab} \rightarrow \mathbb{Z}$ sending both generators to 1, we get that $k = 1$, hence $x = a$.

By taking the contrapositive, distinct conjugates of standard generators of A_{ab} do not commute. □

2.3 The modified Deligne complex

Parabolic subgroups of finite type of an Artin group are used to define a simplicial complex as follows:

Definition 2.9 (*Modified Deligne complex* [21]) The cosets $gA_{\Gamma'}$ of standard parabolic subgroups of finite type of A_Γ form a partially ordered set, for the partial order given by

$$gA_{\Gamma'} < gA_{\Gamma''} \quad \text{if } g \in A_\Gamma \text{ and } \Gamma' \subsetneq \Gamma'' \text{ are full subgraphs of } \Gamma.$$

The **modified Deligne complex** (or Charney–Davis complex) D_Γ of an Artin group A_Γ is the geometric realisation of this poset. That is, vertices of D_Γ correspond to cosets $gA_{\Gamma'}$ of standard parabolic subgroups of finite type, and for every chain of the form

$$gA_{\Gamma_0} < gA_{\Gamma_1} < \cdots < gA_{\Gamma_n},$$

we add an n -simplex spanned by the vertices $gA_{\Gamma_0}, gA_{\Gamma_1}, \dots, gA_{\Gamma_n}$. The group A_Γ acts on its modified Deligne complex by left multiplication on left cosets.

Convention 2.10 From now on, we fix a large-type Artin group A_Γ of hyperbolic type. This implies that its only parabolic subgroups of finite type are its parabolic subgroups on at most two generators, and hence the Deligne complex of A_Γ is a 2-dimensional simplicial complex.

It was shown in [21] that for an Artin group of large and hyperbolic type, there exists an A_Γ -invariant piecewise hyperbolic metric that turns D_Γ into a CAT(-1) space. From now on, we will assume that D_Γ is endowed with such a metric.

Notation 2.11 For simplicity, we will often omit the ‘modified’ from the name and call D_Γ the Deligne complex.

The vertex of D_Γ corresponding to the standard dihedral parabolic subgroup A_{ab} will be denoted v_{ab} . Vertices of D_Γ corresponding to cosets of dihedral parabolic subgroups are said to be of **dihedral type**.

Remark 2.12 We describe the stabilisers of vertices of the Deligne complex. Since a vertex of D_Γ is a left coset of the form $gA_{\Gamma'}$ (with $\Gamma' \subset \Gamma$), its stabiliser is the conjugate $gA_{\Gamma'}g^{-1}$. In particular, we get the following description, for each type of vertices of D_Γ :

- A vertex that corresponds to a left coset of the trivial subgroup has trivial stabiliser.
- A vertex that corresponds to a left coset of the form $g\langle a \rangle$, with $a \in V(\Gamma)$, has a stabiliser that is infinite cyclic.
- A vertex of dihedral type has a stabiliser that is isomorphic to a dihedral Artin group.

2.4 Standard trees and the coned-off Deligne complex

The structure of fixed-point sets of parabolic subgroups of A_Γ play a crucial role. We start by a useful result:

Lemma 2.13 *The fixed-point set in D_Γ of a parabolic subgroup on two generators of A_Γ is a single vertex. In particular, such parabolic subgroups are self-normalising.*

Proof If a parabolic subgroup $gA_{ab}g^{-1}$ were to fix two distinct points of D_Γ , then it would fix the unique CAT(0) geodesic of D_Γ between them. Since stabilisers of edges and triangles of D are either trivial or \mathbb{Z} , $gA_{ab}g^{-1}$ would embed in an infinite cyclic group, which is impossible as A_{ab} is either \mathbb{Z}^2 or contains a copy of \mathbb{Z}^2 by Lemma 2.3.

Since an element of the normaliser of $gA_{ab}g^{-1}$ stabilises $\text{Fix}(gA_{ab}g^{-1})$, hence fixes the vertex gv_{ab} , it follows that $gA_{ab}g^{-1}$ is self-normalising. \square

The fixed-point sets of infinite cyclic parabolic subgroup are much more interesting:

Definition 2.14 (*Standard trees* [44, Definition 4.1]) For an element $g \in A_\Gamma$ that is a conjugate of a standard generator, the fixed-point set $\text{Fix}(g)$ is a convex subtree of D_Γ that is contained in the 1-skeleton of D_Γ .

Such subtrees are called **standard trees** of D_Γ .

Note in particular that all edges of a standard tree have the same infinite cyclic stabiliser.

We list here a few immediate results:

Lemma 2.15 *Two edges of D_Γ have stabilisers that either intersect trivially or are equal. Moreover, if two edges of D_Γ have the same non-trivial stabiliser, then they belong to the same standard tree.*

Proof Since a non-trivial element of A_Γ stabilising two points fixes pointwise the unique CAT(0) geodesic between them, this result is a direct consequence of [44, Lemma 4.3]. \square

Corollary 2.16 *For every standard generator a and non-zero integer $k \in \mathbb{Z} - \{0\}$, the trees $\text{Fix}(a)$ and $\text{Fix}(a^k)$ coincide.*

Proof The inclusion $\text{Fix}(a) \subseteq \text{Fix}(a^k)$ is clear, so let us show the other inclusion. Consider any point x in $\text{Fix}(a^k)$, and any edge e of the tree $\text{Fix}(a)$. If x lies in e , we are done, otherwise we can consider a minimal length geodesic from x to e , which is a^k -invariant. Since triangles of D_Γ have trivial stabilisers, this geodesic is contained in the 1-skeleton and it intersects an edge e' in a non-trivial subpath containing x (if x is not a vertex we cannot say that the geodesic contains the edge). We have that e' is also a^k -invariant, and since $k \neq 0$ we have that the stabilisers of e and e' intersect non-trivially, and therefore by Lemma 2.15 e' , whence x , belongs to $\text{Fix}(a)$. \square

Corollary 2.17 *Two points of D_Γ have stabilisers that intersect non-trivially if and only if they are contained in a common standard tree.*

Proof If two points x, y of D_Γ have a nontrivial common stabiliser H , then H fixes pointwise the unique CAT(0) geodesic γ between them. The geodesic γ cannot pass through the interior of a triangle because H is non-trivial, so γ is contained in the 1-skeleton and is contained in a minimal path of edges of the form e_1, \dots, e_n . So H fixes edges e_1, \dots, e_n , and by Lemma 2.15, it follows that e_1, \dots, e_n , hence x and y , are contained in the same standard tree. (If $n = 1$, note that an edge with non-trivial stabiliser belongs to a standard tree.) \square

Corollary 2.18 *Two distinct standard trees intersect in at most one vertex.*

Proof This follows directly from Lemma 2.15 and the convexity of standard trees. \square

Definition 2.19 (*Coned-off Deligne complex* [44, Definition 4.8]) The **coned-off Deligne complex**, denoted \widehat{D}_Γ , is obtained from D_Γ by coning-off each standard tree of D_Γ . That is, for every standard tree T of D_Γ we add a new vertex v_T , which we connect by an edge to every vertex of T . The complex \widehat{D}_Γ is then the flag completion of the resulting complex. The action of A_Γ on D_Γ extends to an action on \widehat{D}_Γ .




Fig. 2 A portion of the coned-off Deligne complex \widehat{D}_Γ for an Artin group on three generators a, b, c . A fundamental domain for the action of A_Γ on D_Γ is represented in grey, and is a subdivided triangle with vertices v_{ab}, v_{bc}, v_{ac} . A portion of the standard tree T_a is represented in green. In the coned-off Deligne complex \widehat{D}_Γ , this tree is the basis of a cone with apex the vertex v_a

Notation 2.20 For a standard generator $a \in \Gamma$, the standard tree $\text{Fix}(a)$ (i.e. the standard tree containing the vertex $\langle a \rangle$ of D_Γ) will be denoted T_a , and the apex of the cone \widehat{T}_a over that tree will be denoted v_a , see Fig. 2. With this notation, the vertex $g v_a$ corresponds to the apex of the cone over the standard tree $\text{Fix}(g a g^{-1})$ containing $g \langle a \rangle$. The apex of a cone over a standard tree will be called a vertex of tree type.

Remark 2.21 By work of Paris [51, Corollary 4.2], two standard generators are conjugated if and only if there is a path in the presentation graph Γ consisting of edges with odd labels connecting the corresponding vertices. Since distinct generators may be conjugated, it may happen that we have an equality of the form $g v_a = h v_b$ for distinct standard generators a, b . The following result shows that this is the only case where such an equality happens.

Lemma 2.22 *Let T be a standard tree of D_Γ , let a, b be two standard generators, and let $g, h \in A_\Gamma$. If the two vertices $g \langle a \rangle$ and $h \langle b \rangle$ of D_Γ are contained in T , then a and b are connected in Γ by a path with odd labels. (In particular, a and b are conjugated.)*

Proof Since T is a connected tree, we consider a geodesic path e_1, \dots, e_k of T from $g \langle a \rangle$ to $h \langle b \rangle$. Each edge e_i joins a coset of a cyclic standard parabolic subgroup and a coset of parabolic subgroup of dihedral type. For each i , let a_i be the unique standard generator such that e_i contains a vertex that is a coset of $\langle a_i \rangle$. It is enough to show that for every $1 \leq i < k$, a_i and a_{i+1} are either equal or connected by an edge of Γ with odd label. Consider the vertex v of D_Γ where e_i and e_{i+1} meet. If v corresponds to a coset of cyclic standard parabolic subgroup, then $a_i = a_{i+1}$. If v is a vertex of dihedral type, then it follows from [44, Lemma 4.3] that either $a_i = a_{i+1}$, or a_i and a_{i+1} are adjacent in Γ (since e_i and e_{i+1} meet along a vertex that is a coset of $A_{a_i, a_{i+1}}$, which must be an Artin group of finite type by construction of D_Γ) and the label of that edge is odd. This concludes the proof. \square

Remark 2.23 It was shown in [44, Proposition 4.8] that there exists an A_Γ -invariant piecewise hyperbolic metric that turns \widehat{D}_Γ into a CAT(-1) space. From now on, we will assume that \widehat{D}_Γ is endowed such a metric from [44].

It should be noted that this metric depends on a constant $\varepsilon > 0$ that can be chosen arbitrary small, see [44, Definition 4.7]. This constant is such that for an edge e of D_Γ contained in a standard tree, the triangle of D_Γ over e has angles at least $\pi/2 - \varepsilon$ at the two vertices of e . In this article, the choice of ε will be mostly irrelevant. We will only need to consider this constant in Lemma 2.25 below, where ε needs to be smaller than a certain constant depending only on the group A_Γ .

We mention here a slight generalisation of the CAT(-1)-ness of \widehat{D}_Γ , which will be used in Sect. 5.

Lemma 2.24 *Let Z be the full subcomplex of \widehat{D}_Γ whose vertex set is obtained from \widehat{D}_Γ by removing some vertices of tree type. Then Z is also CAT(-1) for the induced metric.*

Proof The complex Z can be thought of as being obtained from the simply connected complex D_Γ by coning-off only certain (contractible) standard trees of D_Γ , so Z is also simply connected. Moreover, for a point $x \in Z$, the link $\text{Lk}_Z(x)$ is a subgraph of $\text{Lk}_{\widehat{D}_\Gamma}(x)$, since \widehat{D}_Γ is a two-dimensional complex. Since \widehat{D}_Γ is CAT(-1), links of points are graphs with a systole of at least 2π , so the same is true for the subgraph $\text{Lk}_Z(x)$. Thus, Z is locally CAT(-1). Since Z is simply-connected and locally CAT(-1), it is CAT(-1). □

The standard trees of D_Γ are convex in D_Γ but become bounded in \widehat{D}_Γ . We mention the following intermediate result, which will be used in Sect. 5.1.

Lemma 2.25 *We can choose the constant $\varepsilon > 0$ from Remark 2.23 small enough so that the following holds: Let T be a standard tree of D_Γ , and let Z be the full subcomplex of \widehat{D}_Γ obtained by removing the vertex of tree type associated to T . Then T is convex in Z for the induced metric.*

Proof It follows from Lemma 2.25 that Z is CAT(-1) for the induced metric. Since T is a subtree of the CAT(-1) complex Z , it is enough to show that two edges of T that share a vertex v make an angle of at least π at v . (This follows from the local characterisation of geodesics in a CAT(0) space) This amounts to showing that for distinct vertices w, w' of $T \cap \text{Lk}_Z(v)$, their distance in $\text{Lk}_Z(v)$ is at least π . By construction, the simplicial graph $\text{Lk}_Z(v)$ is obtained from $\text{Lk}_{D_\Gamma}(v)$ by coning-off the sets $T' \cap \text{Lk}_{D_\Gamma}(v)$ for each standard tree T' other than T , i.e. by constructing a simplicial cone over every such sets $T' \cap \text{Lk}_{D_\Gamma}(v)$. Moreover, each of these new edges has length at least $\pi/2 - \varepsilon$ by construction, see Remark 2.23. Since the action of A_Γ on D_Γ is cocompact, there is a uniform lower bound on the length of edges in the link of an arbitrary vertex of D_Γ , and we can choose the constant $\varepsilon < \pi/10$ smaller than half this uniform lower bound, which we now assume. We now claim that the distance in $\text{Lk}_Z(v)$ between w and w' is at least π . Indeed, assume that this were not the case. Since T is a convex subtree of D_Γ by construction, two distinct points w, w' of $\text{Lk}_{D_\Gamma}(v)$ that belong to T are at distance at least π in $\text{Lk}_{D_\Gamma}(v)$.

Therefore, if there was a geodesic of length less than π between w and w' in $\text{Lk}_Z(v)$, it would have to go through at least two of the additional edges of length $\geq \pi/2 - \varepsilon$ coming from the cone-off procedure. Moreover, since w and w' are not in the same standard tree (Corollary 2.18), said geodesic should also contain another edge, which has length at least 2ε . Therefore, the geodesic would have length at least

$$2\varepsilon + 2(\pi/2 - \varepsilon) \geq \pi,$$

a contradiction. Thus, T is convex in Z . □

Convention 2.26 From now on, we will assume that \widehat{D}_Γ is endowed with a $\text{CAT}(-1)$ metric from [44] such that Lemma 2.25 holds.

Moreover, since we are considering a fixed Artin group A_Γ , we will from now on simply denote by D and \widehat{D} the complexes D_Γ and \widehat{D}_Γ .

We now describe the stabiliser of the vertices of \widehat{D} of tree type:

Lemma 2.27 *The stabiliser of the vertex of tree type v_a is exactly the centraliser (and normaliser) of the cyclic subgroup $\langle a \rangle$. Moreover, this centraliser splits as a direct product of the form*

$$\langle a \rangle \times K,$$

where K is a finitely-generated free group. More precisely, the subgroup $\langle a \rangle$ acts trivially on T_a , while K acts cocompactly on it with trivial edge stabilisers.

Proof Let us first note that the normaliser and centraliser of $\langle a \rangle$ coincide. Indeed, if $g \in A_\Gamma$ is such that $g^{-1}ag \in \langle a \rangle$, we apply the homomorphism $A_\Gamma \rightarrow \mathbb{Z}$ sending every generator to 1 and deduce that $g^{-1}ag = a$, hence g centralises $\langle a \rangle$.

The stabiliser of v_a coincides with the global stabiliser of the standard tree T_a . Let us show that an element $g \in A_\Gamma$ stabilises T_a if and only if it normalises $\langle a \rangle$.

If g stabilises T_a , then it sends the vertex $\langle a \rangle \in T_a$ to the vertex $g\langle a \rangle \in T_a$, so in particular we have

$$a \cdot g\langle a \rangle = g\langle a \rangle,$$

or in other words $g^{-1}ag \in \langle a \rangle$, hence g normalises $\langle a \rangle$. Conversely, let us assume that g normalises $\langle a \rangle$, and let x be a point of T_a . Let us show that $gx \in T_a$. We have

$$a \cdot gx = g(g^{-1}ag)x = gx,$$

the last equality following from the fact that x is fixed by $\langle a \rangle$ by definition of T_a . Thus, gx is fixed by a , and it follows that g stabilises T_a .

The decomposition of the stabiliser of T_a as a direct product as in the statement is a consequence of [44, Lemma 4.5]. Observe that $\text{Stab}(v_a)$ acts cocompactly on T_a . Indeed, D contains finitely many A_Γ orbits of edges, and no edge is contained in distinct translates of a standard tree by Corollary 2.18.

The quotient K was defined in the proof of [44, Lemma 4.5] as the fundamental group of a graph of groups over the graph $T_a/\text{Stab}(v_a)$, with trivial edge stabilisers and vertex stabilisers that are trivial or infinite cyclic. Since $T_a/\text{Stab}(v_a)$ is a finite graph by the above, it follows that K is finitely generated. \square

Regarding vertices of dihedral type, we have the following similar result:

Lemma 2.28 *The stabiliser A_{ab} of the vertex of dihedral type v_{ab} is exactly the centraliser of the cyclic subgroup $\langle z_{ab}^p \rangle$ for all $p \neq 0$, which further coincides with the normaliser.*

Proof Since the stabiliser of v_{ab} is equal to A_{ab} by construction, and z_{ab}^p is central in A_{ab} , it is enough to show that an element in the normaliser of z_{ab}^p fixes v_{ab} . Suppose that we have an element h such that $h \in N(z_{ab}^p)$. Then $z_{ab}^q = h z_{ab}^p h^{-1}$ (for some $q \neq 0$) also fixes the vertex $h v_{ab}$, so z_{ab}^q fixes pointwise the unique CAT(0) geodesic of D joining v_{ab} to $h v_{ab}$. If this geodesic is nontrivial, then z_{ab}^q fixes an edge of D containing v_{ab} , since triangles in D have trivial stabilisers. This is impossible, as otherwise z_{ab}^q would be contained in an edge stabiliser, hence would be conjugate in A_{ab} to a power of a or b (since the edge in question contains v_{ab}). Hence z_{ab}^q would be equal to a power of a or b since z_{ab} is central in A_{ab} , contradicting Corollary 2.7. So, $h v_{ab} = v_{ab}$, as required. \square

We will also mention the following lemma about centralisers, which will be used later in this article (see Lemma 4.5):

Lemma 2.29 *Let $c \in A_\Gamma$ and suppose that either c is a standard generator, or $c = z_{ab}$ for some standard generators a and b generating a dihedral Artin group. Then the centraliser $C(c)$ satisfies $C(c) = C(c^p)$ for all $p \in \mathbb{Z} - \{0\}$.*

Proof There are two cases, according to whether c is a standard generator or $c = z_{ab}$ for standard generators a, b .

First consider the case where $c = z_{ab}$, where a, b are standard generators generating a dihedral type Artin subgroup A_{ab} . Recall that $\langle z_{ab} \rangle$ fixes a point in \widehat{D} , namely the vertex v_{ab} . Suppose that, for some $p \in \mathbb{Z} - \{0\}$, we have an element h such that $h \in C(z_{ab}^p)$. By Lemma 2.28, we have that $h v_{ab} = v_{ab}$, whence $h \in A_{ab}$, and this subgroup is equal to $C(z_{ab})$ by Lemma 2.28.

Next consider the case where c is a standard generator $c = a$, and let e be an edge of the corresponding standard tree T_a , which has $\langle a \rangle$ as stabiliser. Suppose that, for some $p \in \mathbb{Z} - \{0\}$, we have an element h such that $h \in C(a^p)$. Then the stabiliser of the edge he is $\text{Stab}(he) = h \langle a \rangle h^{-1} \supset \langle a^p \rangle$. In particular, the stabilisers of e and he intersect non-trivially, and it follows from Lemma 2.15 that they are in the same standard tree T_a . From Corollary 2.18, we get that $h T_a = T_a$, so $h \in C(a)$. \square

2.5 Links of vertices

The local structure of Deligne complexes will play an important role in this article, so we now describe it in further detail.

Lemma 2.30 *Let v_a be a vertex of tree type. The link $\text{Lk}_{\widehat{D}}(v_a)$ is the standard tree T_a , and the action of $\text{Stab}(v_a)$ on it is cocompact.*

Proof The description of the link follows from the construction of the cone-off, and the cocompactness follows from Lemma 2.27. □

Throughout the rest of this subsection, we fix a vertex of D of dihedral type of the form v_{ab} . Before describing the link in the cone-off \widehat{D} , we start by describing the link in the original Deligne complex D . The link $\text{Lk}_D(v_{ab})$ of that vertex has a simple description, which is a direct consequence of the construction of D :

Lemma 2.31 *The link $\text{Lk}_D(v_{ab})$ is A_{ab} -equivariantly isomorphic to the the geometric realisation of the poset of cosets of strict standard parabolic subgroups of A_{ab} . That is, vertices of $\text{Lk}_D(v_{ab})$ correspond to cosets of the form $g\langle a \rangle$, $g\langle b \rangle$, or $g\{1\}$, and for every $g \in A_{ab}$, we add an edge between $g\{1\}$ and $g\langle a \rangle$, as well as an edge between $g\{1\}$ and $g\langle b \rangle$.*

In particular, the action of $\text{Stab}(v_{ab})$ on $\text{Lk}_D(v_{ab})$ is cocompact. □

This link is a particular case of a general construction that we will use again in Sect. 5.

Definition 2.32 (*Graph of orbits*) Let \mathcal{G} be a graph with an A_{ab} -action, such that the subgroups $\langle a \rangle$ and $\langle b \rangle$ act freely on it. We define a new graph encoding the pattern of intersections of orbits of $\langle a \rangle$ and $\langle b \rangle$ as follows: We put a vertex for every $\langle a \rangle$ -orbit and one vertex for every $\langle b \rangle$ -orbit. If two such orbits have a non-empty intersection, we put an edge between them. The **graph of orbits** $\text{Orbit}_{a,b}(\mathcal{G})$ is defined as the first barycentric subdivision of the graph obtained in this way.

Remark 2.33 With that terminology, the link $\text{Lk}_D(v_{ab})$ is A_{ab} -equivariantly isomorphic to the graph of orbits $\text{Orbit}_{a,b}(\text{Cayley}_{a,b}(A_{ab}))$, where $\text{Cayley}_{a,b}(A_{ab})$ denotes the Cayley graph of A_{ab} for the standard generators a, b .

The action of A_{ab} on $\text{Lk}_D(v_{ab})$ has been studied by Vaskou in [55]. In particular, the following result, which is a geometric counterpart of Lemma 2.6 will be useful in Sect. 5.

Lemma 2.34 ([55, Proposition 4.7]) *Let g be an element of A_{ab} that can be written only with positive letters and that has syllabic length greater than 1. If $m_{ab} \geq 3$, then the $\langle g \rangle$ -orbit maps to $\text{Lk}_D(v_{ab})$ are quasi-isometric embeddings.* □

Since the Garside element Δ_{ab} either conjugates the generators a, b (when m_{ab} is odd) or centralises each of them (when m_{ab} is even), the action of $\langle \Delta_{ab} \rangle$ by right multiplication induces an action on the left cosets. Indeed, we have

$$\begin{aligned}
 g\langle a \rangle \Delta_{ab} &= g\Delta_{ab}\langle b \rangle \quad \text{and} \quad g\langle b \rangle \Delta_{ab} = g\Delta_{ab}\langle a \rangle \quad \text{if } m_{ab} \text{ is odd,} \\
 g\langle a \rangle \Delta_{ab} &= g\Delta_{ab}\langle a \rangle \quad \text{and} \quad g\langle b \rangle \Delta_{ab} = g\Delta_{ab}\langle b \rangle \quad \text{if } m_{ab} \text{ is even.}
 \end{aligned}$$

We now describe the link of v_{ab} in the coned-off Deligne complex. There is a simple characterisation of the trace of a standard tree on the link of a vertex of dihedral type. The following is a reformulation of [44, Lemma 4.3].

Lemma 2.35 *Two vertices of the link $\text{Lk}_D(v_{ab})$ correspond to edges in the same standard tree of D if and only if the corresponding cosets are in the same $\langle \Delta_{ab} \rangle$ -orbit (for the multiplication on the right).*

Moreover, two vertices of the link $\text{Lk}_D(v_{ab})$ that are in the same A_{ab} -orbit correspond to edges in the same standard tree of D if and only if the corresponding cosets are in the same $\langle z_{ab} \rangle$ -orbit (for the multiplication on the right). \square

This local characterization of standard trees allows us to simply describe the links of vertices in the coned-off space \widehat{D} :

Corollary 2.36 *The link $\text{Lk}_{\widehat{D}}(v_{ab})$ is obtained from $\text{Lk}_D(v_{ab})$ by coning-off every $\langle \Delta_{ab} \rangle$ -orbit of vertices corresponding to cosets of the form $g\langle a \rangle$ or $g\langle b \rangle$.*

In particular, the action of $\text{Stab}(v_{ab})$ on $\text{Lk}_{\widehat{D}}(v_{ab})$ is cocompact.

Proof The description of the link is a consequence of Lemma 2.35. The cocompactness of the action of $\text{Stab}(v_{ab})$ on $\text{Lk}_D(v_{ab})$ follows from Lemma 2.31. Moreover, since there are finitely many $\text{Stab}(v_{ab})$ -orbits of edges of D containing v_{ab} , there are in particular finitely many orbits of standard trees containing v_{ab} , and hence finitely many orbits of apices of standard trees containing v_{ab} . The cocompactness of the action of $\text{Stab}(v_{ab})$ on $\text{Lk}_{\widehat{D}}(v_{ab})$ now follows. \square

3 The commutation graph of an Artin group

From now on, in line with Convention 2.10, all Artin groups under consideration are assumed to be of large hyperbolic type. In particular, $m_{ab} \geq 3$ for all generators a, b ; we will occasionally emphasise this assumption in the sequel to remind the reader.

3.1 The commutation graph

We now construct a complex to which Theorem 1.4 will be applied, and which turns out to be quasi-isometric to the coned-off Deligne complex \widehat{D} . This complex is obtained via a general construction that encodes the commutation of certain chosen subgroups of a given group that we now describe.

Definition 3.1 (*Commutation graph*) Let G be a group, and let \mathcal{H} be a set of subgroups of G . We define a simplicial graph, called the **commutation graph** of \mathcal{H} and denoted $Y_{\mathcal{H}}$, as follows. The vertex set of $Y_{\mathcal{H}}$ is

$$\bigsqcup_{H \in \mathcal{H}} G/N(H),$$

and we put an edge between $gN(H)$ and $hN(H')$ if gHg^{-1} and $hH'h^{-1}$ commute, that is, every element of one subgroup commutes with every element of the other subgroup. Note that this is independent of the choice of coset representatives g and h . The group G acts on this graph by left multiplication.

Remark 3.2 Commutation graphs have already been considered in the work of Kim–Koberda on right-angled Artin groups under the name of *extension graphs* [41]. Indeed, for a right-angled Artin group G on generators g_1, \dots, g_n , the extension graph they study is exactly the commutation graph for the family $\mathcal{H} = \{\langle g_1 \rangle, \dots, \langle g_n \rangle\}$.

We now discuss the family of subgroups of A_Γ that yield the correct commutation graph for our purposes.

First, the motivation. The acylindricity of the action of A_Γ on its coned-off Deligne complex and the CAT(0) geometry of this complex can be shown to imply that the maximal subgroups of A_Γ that virtually split as products are the dihedral parabolic subgroups (stabilisers of vertices of dihedral type), which are virtual products by Lemma 2.2, and the normalisers of standard generators, which are products by Lemma 2.27. It is an exercise, left to the reader as it is not needed in this article, to identify the minimal infinite subgroups obtained by taking arbitrary intersections of such maximal virtual products. They come in two families:

- The cyclic subgroups generated by a conjugate of a standard generator: These subgroups are obtained by taking the intersection of stabilisers of vertices of dihedral type contained in a common standard tree.
- The centres of dihedral parabolic subgroups with $m_{ab} \geq 3$: These subgroups are obtained by taking the intersection of the stabilisers of two standard trees that share a vertex.

Remark 3.3 We know from the work of Paris [51, Corollary 4.2] that two standard generators are conjugate if and only if there is a path in the presentation graph Γ consisting of edges with odd labels connecting the corresponding vertices. We can thus choose a set of representatives of conjugacy classes of elements of $V(\Gamma)$, which defines a subset $V_{\text{odd}}(\Gamma)$ of $V(\Gamma)$.

This motivates the following definition:

Definition 3.4 We define the following collection of subgroups of the Artin group A_Γ :

$$\begin{aligned} \mathcal{H} := & \{ \langle a \rangle \mid a \text{ a standard generator in } V_{\text{odd}}(\Gamma) \} \\ & \cup \{ \langle z_{ab} \rangle \mid a, b \text{ span an edge of } \Gamma \text{ with } m_{ab} \geq 3 \}. \end{aligned}$$

We will simply denote by $Y = Y_\Gamma = Y_{\mathcal{H}}$ the commutation graph of \mathcal{H} .

Lemma 3.5 *Distinct elements of \mathcal{H} are in different conjugacy classes. Moreover, if two conjugates of elements of \mathcal{H} intersect non-trivially, then they are equal.*

Proof We have to show that for two distinct elements x, y of the form z_{ab} or $a \in V_{\text{odd}}(\Gamma)$, no non-trivial power of x is conjugate to a non-trivial power of y . We treat several different cases.

Let a, b be two standard generators of $V_{\text{odd}}(\Gamma)$. If non-trivial powers a^k and b^ℓ were conjugate, then using the homomorphism $A_\Gamma \rightarrow \mathbb{Z}$ sending every generator to 1 we would obtain that $k = \ell$. By [51, Corollary 5.3] (which says that an element conjugates

a^k to b^k if and only if it conjugates a to b), we get that a and b are conjugate, which is excluded by construction of $V_{\text{odd}}(\Gamma)$.

The centraliser of a standard generator a has the form $K \times \langle a \rangle$ for some free group K by Lemma 2.27, while the centraliser of an element of the form $z_{a'b'}$ is the dihedral Artin group $A_{a'b'}$ by Lemma 2.28. These are not isomorphic for example because the abelianisations can be isomorphic only if K is trivial (consider the cases of $m_{a'b'}$ odd or even), but $A_{a'b'}$ is not isomorphic to \mathbb{Z} by Lemma 2.3.

Finally, suppose that non-trivial powers of z_{ab} and $z_{a'b'}$ are conjugate. Then their centralisers are also conjugate. By Lemma 2.28, these centralisers are A_{ab} and $A_{a'b'}$. In turn, by Lemma 2.13, these have fixed-point sets respectively consisting of v_{ab} and $v_{a'b'}$ only, so that these fixed point sets are not translates of each other. Therefore, the centralisers cannot be conjugate, and the powers of z_{ab} and $z_{a'b'}$ cannot be conjugate. \square

Notation 3.6 For a standard generator $a \in \Gamma$, we denote by u_a the vertex of Y corresponding to $N(\langle a \rangle)$. The A_Γ -translates of such vertices are said to be of **tree type**. Analogously, for standard generators a, b spanning an edge of Γ with $m_{ab} \geq 3$, we denote by u_{ab} the vertex of Y corresponding to $N(\langle z_{ab} \rangle)$. The A_Γ -translates of such vertices are said to be of **dihedral type**.

In the rest of this section, we construct an equivariant quasi-isometry between the commutation graph Y and the coned-off Deligne complex \widehat{D} . We start by defining such a map at the level of vertices:

Lemma 3.7 *We define a map $\iota : Y^{(0)} \rightarrow \widehat{D}^{(0)}$ as follows:*

- For a vertex of Y of dihedral type of the form gu_{ab} , we set $\iota(gu_{ab}) := gv_{ab}$.
- For a vertex of Y of tree type of the form gu_a , we set $\iota(gu_a) := gv_a$.

Then ι is well-defined, is injective, and realises bijections between the vertices of Y of dihedral type and the vertices of \widehat{D} of dihedral type, as well as between the vertices of Y of tree type and the vertices of \widehat{D} of tree type.

Proof We have to show the following properties:

- if $gN(\langle z_{ab} \rangle) = hN(\langle z_{ab} \rangle)$, then $gv_{ab} = hv_{ab}$.
- if $gN(\langle a \rangle) = hN(\langle a \rangle)$, where $a \in V_{\text{odd}}(\Gamma)$, then $gv_a = hv_a$.

If $gN(\langle z_{ab} \rangle) = hN(\langle z_{ab} \rangle)$, then $g^{-1}h \in N(\langle z_{ab} \rangle) = A_{ab}$, the latter equality following from Lemma 2.28, and we know from Lemma 2.28 that $N(\langle z_{ab} \rangle)$ stabilises v_{ab} . It thus follows that $hv_{ab} = g(g^{-1}h)v_{ab} = gv_{ab}$, so ι is well defined on vertices of dihedral type. Moreover, it in fact realises a bijection on vertices of dihedral type as we just saw that in both Y and \widehat{D} these vertices are cosets of the A_{ab} .

If $gN(\langle a \rangle) = hN(\langle a \rangle)$, then $g^{-1}h \in N(\langle a \rangle)$, and we know from Lemma 2.27 that $N(\langle a \rangle)$ stabilises v_a . It thus follows that $hv_a = g(g^{-1}h)v_a = gv_a$, so ι is well defined on vertices of tree type.

Let us show that ι is surjective on vertices of tree type. Let gv_a be a vertex of \widehat{D} of tree type. By construction, there exists a standard generator $b \in V_{\text{odd}}(\Gamma)$ and an

element $x \in A_\Gamma$ such that $xbx^{-1} = a$. It thus follows that $xv_b = v_a$, and in particular $gv_a = gxv_b = \iota(gxu_b)$.

Let us now show that ι is injective on vertices of tree type. Consider $a, b \in V_{\text{odd}}(\Gamma)$ and $g, h \in A_\Gamma$ such that $gv_a = hv_b$. We want to show that $gu_a = hu_b$. By construction of vertices of tree type, this means that the standard tree of D containing the vertex $g\langle a \rangle$ and the standard tree of D containing the vertex $h\langle b \rangle$ coincide. In particular, the vertices $g\langle a \rangle$ and $h\langle b \rangle$ of D are in the same standard tree, and it follows from Lemma 2.22 that a and b are connected by a path of Γ with odd labels. By definition of $V_{\text{odd}}(\Gamma)$, this implies that $a = b$. We thus have $gv_a = hv_a$, and so $g^{-1}h \in \text{Fix}(v_a) = N(a)$, the latter equality following from Lemma 2.27. We thus have $gN(a) = hN(a)$, hence $gu_a = hu_a = hu_b$, which shows injectivity.

Note that the injectivity of $\iota : Y^{(0)} \rightarrow \widehat{D}^{(0)}$ is now straightforward, since the sets of vertices of tree type and dihedral type of \widehat{D} are disjoint. □

Lemma 3.8 *The map ι is well-defined, is injective, and realises bijections between the vertices of Y of dihedral type and the vertices of \widehat{D} of dihedral type, as well as between the vertices of Y of tree type and the vertices of \widehat{D} of tree type.*

Moreover, two vertices v, v' of Y are adjacent if and only if the following occurs: One of them (say v) is of dihedral type, the other (say v') is of tree type, and $\iota(v)$ is contained in the standard tree having $\iota(v')$ as apex.

In particular, Y is a bipartite graph with respect to the type of vertices.

Proof The first statement is exactly Lemma 3.7. Let us now characterise the edges of Y .

Consider two vertices $v = gN(H), v' = hN(H')$ of Y that are connected by an edge. The subgroups $gHg^{-1}, hH'h^{-1}$ are infinite cyclic, and we denote by $z_v, z_{v'}$ the associated generators. The elements $z_v, z_{v'}$ commute by definition of Y , and they generate a \mathbb{Z}^2 subgroup of A_Γ by Lemma 3.5. Since the action of A_Γ on \widehat{D} is acylindrical and \widehat{D} is CAT(0), this subgroup must fix a point. In particular, the fixed-point sets in \widehat{D} of z_v and $z_{v'}$ have a non-trivial intersection. We will need the following standard result from group actions on trees, whose proof we omit, to show that certain configurations are impossible:

Claim: Let G be a group acting on a simplicial tree T by isometries, let w, w' be two distinct vertices of T , and let g, g' be two elements of G such that for all non-zero $k \in \mathbb{Z}$, we have $\text{Fix}_T(g^k) = \{w\}$ and $\text{Fix}_T((g')^k) = \{w'\}$. Then g, g' generate a non-abelian free subgroup.

We now consider several cases, depending on the type (dihedral or tree) of the corresponding vertices.

Case 1: Suppose by contradiction that v and v' are adjacent vertices of Y of dihedral type. Recall that for an element of the form z_{ab} , we have $\text{Fix}_D(z_{ab}) = \{v_{ab}\}$ by Lemma 2.28. Since $\text{Fix}_{\widehat{D}}(z_v) \cap \text{Fix}_{\widehat{D}}(z_{v'})$ is non-empty, these fixed-point sets intersect along at least (at least) one vertex w , which must be of tree type since $\text{Fix}_D(z_v) \cap \text{Fix}_D(z_{v'}) = \emptyset$. This vertex correspond to a standard tree T stabilised by $\langle z_v, z_{v'} \rangle$. Moreover, T contains both vertices $\iota(v)$ and $\iota(v')$: Indeed, z_v stabilises the unique CAT(0) geodesic between $\iota(v)$ and T , and since z_v is not conjugate to a power of a standard generator by Lemma 3.5, z_v cannot stabilise an edge of D , and it follows that this geodesic is reduced to a point (and similarly for $z_{v'}$). Thus, $\langle z_v, z_{v'} \rangle$ acts on

T , and since no non-trivial power of z_v or $z_{v'}$ is conjugate to a standard generator by Lemma 3.5, it follows that for all non-zero $k \in \mathbb{Z}$, we have $\text{Fix}_T(z_v^k) = \{\iota(v)\}$ and $\text{Fix}_T((z_{v'})^k) = \{\iota(v')\}$. It now follows from the Claim that z_v and $z_{v'}$ generate a non-abelian free subgroup, a contradiction.

Case 2: Let us now assume that v, v' are two adjacent vertices of Y of tree type. By the above argument, z_v and $z_{v'}$ fix a vertex w of \widehat{D} .

If w is a vertex of D , then it belongs to $\text{Fix}_D(z_v) \cap \text{Fix}_D(z_{v'}) = T_{\iota(v)} \cap T_{\iota(v')}$. Since distinct standard trees of D meet in at most one vertex by Corollary 2.18, we can assume that w is that common vertex. Up to conjugation, we can thus assume that z_v and $z_{v'}$ are two distinct commuting conjugates of standard generators of a dihedral Artin group A_{ab} , which contradicts Lemma 2.8.

Case 3: Finally, we assume that v is of dihedral type and v' of tree type. Since $\text{Fix}_{\widehat{D}}(z_v) \cap \text{Fix}_{\widehat{D}}(z_{v'})$ is non-empty, let w be a vertex of \widehat{D} in that intersection. If w belongs to D , then in particular w belongs to $\text{Fix}_D(z_v) \cap \text{Fix}_D(z_{v'}) = \{\iota(v)\} \cap T_{\iota(v')}$, so $\iota(v)$ is contained in $\text{Fix}_D(z_{v'}) = T_{\iota(v')}$, as required.

Let us show that this is the only possibility. By contradiction, if $\text{Fix}_{\widehat{D}}(z_v) \cap \text{Fix}_{\widehat{D}}(z_{v'})$ does not contain any vertex of D , then these fixed-point sets intersect along (at least) one vertex, which must be of tree type since $\text{Fix}_D(z_v) \cap \text{Fix}_D(z_{v'}) = \emptyset$. This vertex corresponds to a standard tree T . Moreover, the same reasoning as in the previous cases shows that T contains $\iota(v)$ and intersects $T_{\iota(v')}$. We thus have that $\langle z_v, z_{v'} \rangle$ induces an action on T , z_v fixes the vertex $w := \iota(v)$, $z_{v'}$ fixes the vertex $w := T_{\iota(v')} \cap T$, w, w' are distinct since $\text{Fix}_D(z_v) \cap \text{Fix}_D(z_{v'})$ is empty by assumption. Moreover, for every non-zero integer k , we have $\text{Fix}_T(z_v^k) = \{w\}$ since no non-trivial power of an element of the form z_{ab} is conjugate to a standard generator by Lemma 3.5, and for every non-zero integer k , we have that $\text{Fix}_T(z_{v'}^k) = \{w'\}$ by Lemma 2.15 since T is a standard tree distinct from $T_{\iota(v')}$. It now follows from the Claim that z_v and $z_{v'}$ generate a non-abelian free subgroup, a contradiction. \square

The previous lemma is useful in understand the structure of the graph Y . In particular, we have the following results:

Lemma 3.9 *The graph Y does not contain any triangle or square.*

Proof Since Y is bipartite with respect to the type of vertices by Lemma 3.8, it does not contain triangles.

Let us now show by contradiction that Y does not contain squares. Since Y is bipartite with respect to the type of vertices by Lemma 3.8, such a square would contain exactly two opposite vertices of dihedral type and two opposite vertices of tree type.

In D_Γ , this would correspond to two distinct standard trees intersecting in at least two different vertices, which contradicts Corollary 2.18. \square

Lemma 3.10 *If Γ is connected, then the commutation graph Y is connected.*

Proof Let us first show that any two vertices of Y of tree type are connected by a path of Y . Since A_Γ is generated by its standard generators, it is enough to show that every pair of vertices of tree type of the form u_a and u_b are connected by a path of Y . Let

a, b be two vertices of Γ . Since Γ is connected, let $c_1 = a, \dots, c_n = b$ be a sequence of vertices of Γ such that c_i and c_{i+1} are adjacent for every i . Then the sequence of vertices

$$u_{c_1}, u_{c_1, c_2}, u_{c_2}, \dots, u_{c_i}, u_{c_i, c_{i+1}}, u_{c_{i+1}}, \dots, u_{c_n}$$

defines a combinatorial path of Y between u_a and u_b .

Moreover, since every vertex of dihedral type of Y is connected by an edge to some vertex of Y of tree type, it follows that Y is connected. □

Lemma 3.11 *The action of A_Γ on Y is cocompact.*

Proof There are finitely many orbits of vertices of Y by construction, since the family \mathcal{H} is finite. It follows from Lemma 3.8 and the characterisation of the edges of Y that the set of edges of Y equivariantly embeds into the set of edges of \widehat{D} . Since the action of A_Γ on \widehat{D} is cocompact, it follows that there are only finitely many A_Γ -orbits of edges of Y , hence the result. □

Lemma 3.8 allows us to extend the map ι to a map from Y to \widehat{D} :

Definition 3.12 We extend the map $\iota : Y^{(0)} \rightarrow \widehat{D}^{(0)}$ into a simplicial and A_Γ -equivariant map $\iota : Y \rightarrow \widehat{D}$ as follows: An edge between a vertex u of dihedral type and a vertex u' of tree type is sent to the corresponding edge of \widehat{D} between $\iota(u)$ and $\iota(u')$.

Lemma 3.13 *Suppose that the graph Γ is connected. Then the map $\iota : Y \rightarrow \widehat{D}$ is an A_Γ -equivariant quasi-isometry. More precisely, ι embeds Y as a coarsely dense subgraph of \widehat{D} .*

Proof We construct a quasi-inverse $\bar{\iota} : \widehat{D} \rightarrow Y$ as follows. First notice that since \widehat{D} has finitely many isometry types of simplices, it is enough to define $\bar{\iota}$ at the level of vertices, where the distance between vertices is defined as the length of a minimal path in the 1-skeleton.

By Lemma 3.8, for a vertex v of \widehat{D} that is either of dihedral or tree type, we define $\bar{\iota}(v)$ to be the unique vertex $u \in Y$ such that $\iota(u) = v$. Let us define $\bar{\iota}$ on the remaining vertices. A vertex v of \widehat{D} corresponding to the coset of a cyclic group generated by a standard generator belongs to a unique standard tree T , and we set $\bar{\iota}(v) = \bar{\iota}(v')$, where v' is the vertex of \widehat{D} that is the apex corresponding to T . For a vertex v corresponding to a coset of the trivial subgroup, we pick a vertex v' of dihedral type of D adjacent to it, and we set $\bar{\iota}(v) = \bar{\iota}(v')$.

Lemma 3.8 implies that $\bar{\iota} \circ \iota$ is the identity on the vertices of Y , and that $\iota \circ \bar{\iota}$ is the identity on the vertices of \widehat{D} that are either of dihedral type or of tree type. For a vertex v of \widehat{D} corresponding either to the coset of a trivial, cyclic, or \mathbb{Z}^2 -subgroup, the construction implies that $\iota \circ \bar{\iota}(v)$ and v are at distance 1 in the 1-skeleton of D , which concludes the proof. □

Note that if Γ is disconnected and can be written as the disjoint union of two full subgraphs Γ_1, Γ_2 , then A_Γ splits as the free product $A_{\Gamma_1} * A_{\Gamma_2}$. In particular, since

a free product of hierarchically hyperbolic groups is itself hierarchically hyperbolic (see [9, Corollary 8.24] or [9, Theorem 9.1]), it is enough to consider the case of a connected graph Γ . This motivates the following convention:

Convention 3.14 In the rest of this article (except the proof of Theorem A in Sect. 6.3), we will assume that the underlying presentation graph Γ is connected and not a single vertex.

3.2 The graph of proper irreducible parabolic subgroups of finite type

As an aside, we highlight the connection between the commutation graph Y and the graph of proper irreducible parabolic subgroups of finite type introduced by Morris-Wright [48], and generalizing a construction in the spherical type of Cumplido et al. [24]. This graph was proposed as an analogue of the curve graph for all Artin groups. We show that this indeed the case for Artin groups of large and hyperbolic type. We emphasise that this subsection is not needed in the rest of the paper and can be omitted in a first reading.

Definition 3.15 ([48]) The **graph of irreducible proper parabolic subgroups of finite type** P of A_Γ is the simplicial graph defined as follows. Vertices correspond to the proper parabolic subgroups of finite type that are *irreducible* (that is, they do not decompose as a direct product of proper standard parabolic subgroups). Two vertices H, H' are connected by an edge when either:

- There is a strict inclusion $H \subsetneq H'$, or
- We have $H \cap H' = \{1\}$ and H, H' commute.

The Artin group A_Γ acts on P by conjugation.

While this definition makes sense for all Artin groups, we recall that we are only dealing in this article (and in particular in this section) with the case of Artin groups A_Γ that are large-type and of hyperbolic type.

Proposition 3.16 *Assume that Γ is a connected graph not reduced to a single edge. Then the graph of irreducible parabolic subgroups of finite type of A_Γ is equivariantly isomorphic to the commutation graph.*

Under such conditions on Γ , the action of A_Γ on the coned-off Deligne complex is acylindrical and universal by [44, Theorem A]. The following is thus a direct consequence of Lemma 3.13.

Corollary 3.17 *Assume that Γ is a connected graph not reduced to a single edge. Then the graph of irreducible proper parabolic subgroups of finite type of A_Γ is hyperbolic of infinite diameter, and the action of A_Γ on it is acylindrical and universal. \square*

In order to prove Proposition 3.16, we will need the following result:

Lemma 3.18 *Assume that Γ is a connected graph not reduced to a single edge. Then distinct elements of \mathcal{H} have normalisers that are in different conjugacy classes.*

Proof The normaliser of a standard generator is of the form $\mathbb{Z} \times F_k$ for F_k a finitely-generated free group by Lemma 2.27, while the normaliser of an element of the form z_{ab} is equal to the dihedral Artin group A_{ab} by Lemma 2.28. We show that these groups are non-isomorphic. This is clear if $k = 1$ as a dihedral Artin group with $m_{ab} \geq 3$ is not isomorphic to \mathbb{Z}^2 since it virtually contains a non-abelian free group by Lemma 2.3. For $k \geq 2$, this follows for instance from the description of their abelianisations: The abelianisation of $\mathbb{Z} \times F_k$ is a free abelian group of rank $k + 1 \geq 3$, while the abelianisation of A_{ab} is generated by two elements since A_{ab} is. Thus, the normaliser of a standard generator and the normaliser of an element of the form z_{ab} are not conjugated.

Let $z_{ab}, z_{a'b'}$ be two elements associated to two distinct edges of Γ , and let $g \in A_\Gamma$. By Lemma 2.28, the normalisers of z_{ab} and $gz_{a'b'}g^{-1}$ are the dihedral Artin groups A_{ab} and $gA_{a'b'}g^{-1}$ and their fixed-point sets in D are v_{ab} and $gv_{a'b'}$ respectively. As these points are in different A_Γ -orbits, it follows that $N(z_{ab})$ and $N(z_{a'b'})$ are not conjugate.

Finally, let a, b two distinct standard generators in $V_{\text{odd}}(\Gamma)$, and suppose by contradiction that $N(a) = gN(b)g^{-1}$ for some $g \in A_\Gamma$. By Lemma 2.27, the normalisers a and b have the form $\langle a \rangle \times F_k$ and $\langle b \rangle \times F_k$ respectively (note that the rank of the free group factors need to coincide since the subgroups are conjugated). If $k \geq 2$, then since $N(a)$ and $N(b)$ are conjugated, so are their centres $\langle a \rangle$ and $\langle b \rangle$. Using the homomorphism $A_\Gamma \rightarrow \mathbb{Z}$ sending each standard generator to 1, it follows that a and b are conjugated. By construction of $V_{\text{odd}}(\Gamma)$, it follows that $a = b$, a contradiction.

Assume now that $k \leq 1$. We thus have two standard generators a, b such that $N(a), N(b)$ are isomorphic to \mathbb{Z} or \mathbb{Z}^2 . We will use the following general claim:

Claim: Let x be a standard generator. If the normaliser $N(x)$ is isomorphic to \mathbb{Z} , then x corresponds to an isolated vertex of Γ . If the normaliser $N(x)$ is isomorphic to \mathbb{Z}^2 , then x corresponds to a leaf of Γ .

Let us prove the Claim. By Lemma 2.27, the normaliser $N(x)$ is of the form $\langle x \rangle \times F$ for some finitely-generated group F . An explicit basis \mathcal{B} of F was given in [44, Remark 4.6]. In particular, the following holds:

- (i) If x is an isolated vertex of Γ , then F is trivial.
- (ii) The basis \mathcal{B} contains the element z_{xy} for every vertex y of Γ adjacent to x .
- (iii) Suppose that y is adjacent to x and the edge of Γ between them has an odd label. If w is a vertex of Γ adjacent to y and distinct from x , then \mathcal{B} contains a suitable conjugate of z_{yw} . We note that such an element cannot be equal to an element described in (ii) by Lemma 3.5.

If $N(x)$ is isomorphic to \mathbb{Z} , then F is trivial, and by item (ii) in the above description of a basis of F , x must be an isolated vertex. If $N(x)$ is isomorphic to \mathbb{Z}^2 , then F is isomorphic to \mathbb{Z} , and by item (ii) in the above can have at most one neighbour in Γ . It cannot be an isolated vertex for otherwise F would be trivial by item (i) above. Thus x is a leaf of Γ . This proves the Claim.

Since $N(a), N(b)$ are isomorphic to \mathbb{Z} or \mathbb{Z}^2 , it follows from the above Claim that a and b are leaves of Γ (since Γ is connected by assumption). Moreover, since $N(a), N(b) \cong \mathbb{Z}^2$ and Γ is not reduced to a single edge, item (iii) above implies that the corresponding edges of Γ containing a and b have even labels.

By Lemma 2.27, $N(a) = gN(b)g^{-1}$ stabilises the standard tree containing the vertex $\langle a \rangle$ and the standard tree containing $g\langle b \rangle$. Because a, b are leaves of Γ with corresponding edges having even labels, the corresponding standard trees T_a and gT_b have a very simple structure, which is a direct consequence of [44, Lemma 4.3] (see also [44, Example 4.7] for a concrete example): Let a' be the unique neighbour of a in Γ , and let e be the edge of D between the parabolic subgroups $\langle a \rangle$ and $A_{aa'}$. Then T_a is the union of the $\langle z_{aa'} \rangle$ -translates of e (such a description was for instance given in [44, Example 4.7]). In particular, T_a (and similarly for gT_b) is isomorphic to a cone over a countably infinite discrete set.

Let us show that the standard trees T_a and gT_b intersect. If these two convex trees were disjoint, then $N(a) = gN(b)g^{-1}$ would also stabilise the unique $\text{CAT}(-1)$ geodesic of D between them, and so both normalisers would be contained in an edge-stabiliser or triangle-stabiliser. This is impossible as such stabilisers do not contain \mathbb{Z}^2 , as they are infinite cyclic or trivial respectively.

Thus these two standard trees meet. The vertices in T_a that are not the central vertex of dihedral type correspond to cosets of $\langle a \rangle$ (and we have an analogous result for b). Since $\langle a \rangle \neq \langle b \rangle$ by assumption, the trees T_a and gT_b necessarily meet at the central vertex of dihedral type. Let a' (respectively b') be the unique neighbour of a (respectively b) in Γ . The unique vertex of T_a of dihedral type corresponds to a coset of $A_{aa'}$. Similarly, the unique vertex of gT_b of dihedral type corresponds to a coset of $A_{bb'}$. As these vertices agree, it follows that $\{a, a'\} = \{b, b'\}$, and since $a \neq b$ this implies a and b are adjacent in Γ . Since a and b are leaves of Γ by the above, it follows that Γ is a single edge, contradicting our assumption on Γ . □

Lemma 3.19 *Assume that Γ is a connected graph not reduced to a single edge. Consider the map $\varphi : Y^{(0)} \rightarrow P^{(0)}$ defined as follows:*

- For a vertex gA_{ab} of Y of dihedral type, we set $\varphi(gA_{ab}) = gA_{ab}g^{-1}$.
- For a vertex $gN(a)$ of Y of tree type, we set $\varphi(gN(a)) = g\langle a \rangle g^{-1}$.

Then φ is well-defined and realises an A_Γ -equivariant bijection between the vertex sets of Y and P .

Proof The map φ is well-defined on vertices of dihedral type, but we need to check that the definition makes sense for vertices of tree type. It is enough to show that if a, b are standard generators in $V_{\text{odd}}(\Gamma)$ with $N(a) = N(b)$, then $a = b$. This is a direct consequence of Lemma 3.18.

Let us show that φ is surjective. It is clear from the construction that every conjugate of the form $gA_{ab}g^{-1} = \varphi(gA_{ab})$ is in the image of φ . Let a be a standard generator of A_Γ . By construction of $V_{\text{odd}}(\Gamma)$, there exists a standard generator $b \in V_{\text{odd}}(\Gamma)$ and an element $h \in A_\Gamma$ such that $a = hbh^{-1}$. Thus, any conjugate of the form $g\langle a \rangle g^{-1}$ can be written as $g\langle a \rangle g^{-1} = (gh)\langle b \rangle (gh)^{-1} = \varphi(ghN(b))$. Thus, φ is surjective.

Let us show that φ is injective. For a given standard generator $a \in V(\Gamma)$, the map φ induces a bijection between the conjugates of $\langle a \rangle$ and the cosets of $N(a)$. Similarly, since dihedral parabolic subgroups of A_Γ are self-normalizing by Lemma 2.13, we have that for a given edge a, b of Γ with $m_{ab} \geq 3$, the map φ induces a bijection between the conjugates of A_{ab} and the cosets of A_{ab} . To show that φ is injective, it remains to show that for H, H' different elements of \mathcal{H} , we cannot have an equality of

the form $\varphi(gN(H)) = \varphi(g'N(H'))$ for $g, g' \in A_\Gamma$. Such an equality would amount to an equality between parabolic subgroups. Note that a conjugate of the form $g\langle a \rangle g^{-1}$, which is isomorphic to \mathbb{Z} , cannot be isomorphic to a conjugate of a dihedral parabolic subgroup by Lemma 2.3. There are thus two cases to consider:

If H, H' correspond to cyclic subgroups of the form $\langle a \rangle, \langle a' \rangle$ respectively, then the equality $\varphi(gN(H)) = \varphi(g'N(H'))$ yields the equality $g\langle a \rangle g^{-1} = g'\langle a' \rangle (g')^{-1}$. Since $a \neq a'$, this equality is impossible by Lemma 3.5.

If H, H' correspond to cyclic subgroups of the form $\langle z_{ab} \rangle, \langle z_{a'b'} \rangle$ respectively, then the equality $\varphi(gN(H)) = \varphi(g'N(H'))$ and the fact that $A_{ab} = N\langle z_{ab} \rangle$ by Lemma 2.28 yield the equality $gN\langle z_{ab} \rangle g^{-1} = g'N\langle z_{a'b'} \rangle (g')^{-1}$. Since $H \neq H'$, this equality is impossible by Lemma 3.18. \square

Proof of Proposition 3.16 We know from Lemma 3.19 that the map φ realises an A_Γ -equivariant bijection between the vertex sets of Y and P . Let us show that this bijection extends to an equivariant isomorphism between P and Y .

First, it follows from Lemma 3.8 that two vertices that are adjacent in Y are also adjacent in P , and we now show the converse.

Let Q and Q' be two proper irreducible parabolic subgroups of finite type, i.e. parabolic subgroups on one or two generators with $m_{ab} \geq 3$, and assume that Q and Q' are connected in P . There are two cases to consider, depending on the two types of edges of P :

Let us first assume that Q, Q' intersect trivially and commute. Since A_Γ is of cohomological dimension 2 by [21] and dihedral parabolic subgroups contain a copy of \mathbb{Z}^2 , this can only happen if both Q and Q' are infinite cyclic, in which case there is a Y -edge between them by construction of Y .

Suppose now that $Q \subsetneq Q'$. First notice that Q and Q' cannot be both dihedral parabolic subgroups by Lemma 2.13. We can thus assume that Q is infinite cyclic. Let e be an edge in the standard tree T associated to Q , and let σ' be a maximal simplex of $\text{Fix}(Q')$ (i.e. σ' is a vertex or an edge depending on whether Q' is of dihedral type or infinite cyclic respectively). Since Q fixes both v and σ' by assumption, Lemma 2.17 implies that v and σ' are contained in a common standard tree, namely the standard tree T .

Let us show by contradiction that Q' cannot be infinite cyclic. Let T' be the standard tree corresponding to Q' . Then since T and T' both contain e , it follows from Corollary 2.18 that $T = T'$. Since Q, Q' is the pointwise stabiliser of T, T' respectively by construction, we get $Q = Q'$, contradicting the assumption that $Q \subsetneq Q'$.

Thus, Q' is a parabolic subgroup of dihedral type. We thus have that the vertex of D corresponding to Q' is contained in the standard tree associated to Q . By Lemma 3.8, this implies that Q and Q' are connected by an edge of Y . \square

Remark 3.20 We note that the alternative description of P given by Proposition 3.16 provides an alternative characterisation of the edges of P : Two proper irreducible parabolic subgroups of finite type are joined by an edge of P if and only if their centres commute. This generalises results known for Artin groups of finite type [24, Theorem 2.2] and of FC type [48, Lemma 4.2].

3.3 Adjacency in the commutation graph and intersection of neighbourhoods of cosets

In this section, we study the interactions between cosets associated to adjacent vertices of the commutation graph Y . Namely, we show that two such adjacent cosets have neighbourhoods that intersect along a quasi-flat (see Lemma 3.23). We start with the following result:

Lemma 3.21 *Let $H_1, H_2 \in \mathcal{H}$ and suppose that $gN(H_1), hN(H_2)$ are adjacent vertices of Y for some $g, h \in A_\Gamma$. Then we have*

$$gN(H_1)g^{-1} \cap hN(H_2)h^{-1} = \langle gH_1g^{-1}, hH_2h^{-1} \rangle.$$

Moreover, $\langle gH_1g^{-1}, hH_2h^{-1} \rangle$ is naturally isomorphic to $gH_1g^{-1} \times hH_2h^{-1} \cong \mathbb{Z}^2$.

Proof By Lemma 3.8, we can assume that $gN(H_1)$ is mapped under ι to the apex of a standard tree T_1 and $hN(H_2)$ is mapped under ι to a vertex of dihedral type contained in T_1 . Notice that, up to a conjugation, it is enough to prove the desired equality when h is trivial, so we will assume that this is the case in the rest of the proof. The subgroup $H_2 \in \mathcal{H}$ is of the form $\langle z_{ab} \rangle$ for some standard generators a, b , and so the vertex v corresponds to the trivial coset A_{ab} by construction of ι . Since the standard tree T_1 contains the vertex A_{ab} , it also contains a vertex of the form $x\langle a \rangle$ or $x\langle b \rangle$ for some $x \in A_{ab}$. Without loss of generality, we can assume that T_1 contains a vertex of the form $x\langle a \rangle$ for some chosen $x \in A_{ab}$.

We have $hN(H_2)h^{-1} = N(z_{ab}) = A_{ab}$ by Lemma 2.28. Moreover, since two edges of T_1 have the same stabiliser by construction, we have $gH_1g^{-1} = x\langle a \rangle x^{-1}$, hence $gN(H_1)g^{-1} = xN(a)x^{-1} = xC(a)x^{-1}$, the latter equality following from Lemma 2.27. Thus, the intersection $gN(H_1)g^{-1} \cap hN(H_2)h^{-1}$ can be identified with the centraliser of the element xax^{-1} in A_{ab} . By [23, Lemma 7.(ii)], this is equal to $\langle xax^{-1}, z_{ab} \rangle = \langle gH_1g^{-1}, hH_2h^{-1} \rangle$.

For the “moreover” part, note that the intersection is naturally isomorphic to the product by definition of edges of Y and Lemma 3.5. □

Definition 3.22 Let $B(e, r)$ denote the closed ball of radius r in the standard Cayley graph of A_Γ .

We fix a constant $r \geq 0$ so that whenever a, b are conjugate standard generators there exists $x \in B(e, r)$ so that $xax^{-1} = b$. The r -neighbourhood $N(H)^{+r}$ of $N(H)$ is defined as

$$N(H)^{+r} := \bigcup_{c \in B(e, r)} N(H)c.$$

Lemma 3.23 *There exists a constant $\ell \geq 0$ such that the following holds:*

Let $H_1, H_2 \in \mathcal{H}$ and suppose that $gN(H_1), hN(H_2)$ are adjacent vertices of Y for some $g, h \in A_\Gamma$.

Then $gN(H_1)^{+r} \cap hN(H_2)^{+r}$ is nonempty. Also, the subgroup $gN(H_1)g^{-1} \cap hN(H_2)h^{-1}$ (which is isomorphic to \mathbb{Z}^2 by Lemma 3.21) acts coboundedly on the

intersection $gN(H_1)^{+r} \cap hN(H_2)^{+r}$, and the quotient of this action has diameter at most ℓ .

Remark 3.24 We advise the reader to focus on the case of *even* Artin groups (i.e. m_{ab} is even for all standard generators) in a first reading. In that particular case, we have $V_{\text{odd}}(\Gamma) = V(\Gamma)$ and we can take $r = \ell = 0$ in the above statement. In particular, we are simply looking at intersections of cosets.

Let us explain briefly the need to consider neighbourhoods of cosets in the general case. If x and y are two standard generators connected by an edge of Γ with an odd label, then we have $\Delta_{xy}x\Delta_{xy}^{-1} = y$, where Δ_{xy} is the Garside element of A_{xy} . In particular, we have $g\Delta_{xy}N(x)\Delta_{xy}^{-1} = gN(y)$, and so the cosets $g\Delta_{xy}N(x)$ and $gN(y)$ are at bounded Hausdorff distance of one another (they are in a sense “parallel”). Let us assume that we are looking at vertices $gN(H)$ and $hN(H')$ where H, H' correspond to standard generators a, b . As we saw in the proof of Lemma 3.21, the intersection $gN(H)g^{-1} \cap hN(H')h^{-1}$ stabilises a vertex of dihedral type. However, while H, H' correspond to standard generators a, b in $V_{\text{odd}}(\Gamma)$, it may happen that this vertex of dihedral type correspond to two different standard generators a', b' that are conjugated to a, b respectively. Thus, while it may not be clear how to study the intersection $gN(a) \cap hN(b)$, it is much simpler to study the corresponding intersection of parallel copies corresponding to cosets of $N(a')$ and $N(b')$. This is why we do not simply consider the cosets $gN(a), hN(b)$ but consider suitable neighbourhoods that contain those parallel copies corresponding to cosets of $N(a')$ and $N(b')$.

Proof of Lemma 3.23 Let $H_1, H_2 \in \mathcal{H}$ and let $gN(H_1), hN(H_2)$ be cosets corresponding to adjacent vertices of Y .

By Lemma 3.8, we can assume that H_1 correspond to a standard generator a_1 , H_2 corresponds to an element of the form $z_{b_1b_2}$, and the standard tree T_1 with apex $\iota(gN(H_1))$ contains the vertex of dihedral type $\iota(hN(H_2))$. This vertex corresponds to a coset of the form $hA_{b_1b_2}$. Since the tree T_1 contains the vertex $hA_{b_1b_2}$, it also contains a vertex corresponding to a coset of the form $hx\langle b_1 \rangle$ or $hx\langle b_2 \rangle$ for some $x \in A_{b_1b_2}$, so without loss of generality we will assume that it contains a vertex corresponding to a coset of the form $hx\langle b_1 \rangle$. Since $N(H_2) = A_{b_1b_2}$ by Lemma 2.28, we have that $hxN(H_2) = hN(H_2)$, so without loss of generality we will assume, up to replacing the element h by hx , that T_1 contains the vertex $h\langle b_1 \rangle$. Moreover, since T_1 contains vertices corresponding to cosets of $\langle a_1 \rangle$ and $\langle b_1 \rangle$, it follows from Lemma 2.22 that a_1 and b_1 are conjugated: By construction of r , we choose an element $x_1 \in B(e, r)$ so that $x_1a_1x_1^{-1} = b_1$.

By construction of x_1 , we have $hx_1v_{a_1} = hv_{b_1}$. Thus, the apex of the standard tree T_1 is the coset $gN(a_1) = hx_1N(a_1)$, so we can assume without loss of generality that $g = hx_1$. We thus have $gN(H_1) = hx_1N(a_1)$ and $hN(H_2) = hA_{b_1b_2}$. Since $x_1 \in B(e, r)$, the intersection $gN(H_1)^{+r} \cap hN(H_2)^{+r} = hx_1N(a_1)^{+r} \cap hN(z_{b_1b_2})^{+r}$ contains the intersection

$$hx_1N(a_1)x_1^{-1} \cap hN(z_{b_1b_2}) = hN(b_1) \cap hN(z_{b_1b_2}) = h\langle b_1, z_{b_1b_2} \rangle,$$

the latter equality following from Lemma 3.21.

Since $B(e, r)$ contains only finitely many elements of A_Γ and \mathcal{H} is finite, it follows from [34, Proposition 9.4] that there exists a constant $\ell > 0$ (that does not depend on the choice of standard generators a_1, b_1, b_2 and element x_1) such that

$$\begin{aligned} x_1 N(a_1)^{+r} \cap N(z_{b_1 b_2})^{+r} &\subset (N(x_1 a_1 x^{-1}) \cap N(z_{b_1 b_2}))^{+\ell} \\ &= (N(b_1) \cap N(z_{b_1 b_2}))^{+\ell} = \langle b_1, z_{b_1 b_2} \rangle^{+\ell}, \end{aligned}$$

and in particular

$$gN(H_1)^{+r} \cap hN(H_2)^{+r} = hx_1 N(a_1)^{+r} \cap hN(z_{b_1 b_2})^{+r} \subset h\langle b_1, z_{b_1 b_2} \rangle^{+\ell}.$$

Thus, the intersection $gN(H_1)^{+r} \cap hN(H_2)^{+r}$ is at Hausdorff distance at most ℓ from $h\langle b_1, z_{b_1 b_2} \rangle$.

The conjugate $h\langle b_1, z_{b_1 b_2} \rangle h^{-1}$ is equal to $\langle ga_1 g^{-1}, hz_{b_1 b_2} h^{-1} \rangle = gN(H_1)g^{-1} \cap hN(H_2)h^{-1}$ by Lemma 3.21. Since $k\langle b_1, z_{b_1 b_2} \rangle k^{-1}$ acts coboundedly on $k\langle b_1, z_{b_1 b_2} \rangle$ and $gN(H_1)^{+r} \cap hN(H_2)^{+r}$ is at Hausdorff distance at most ℓ from $k\langle b_1, z_{b_1 b_2} \rangle$, it follows that $gN(H_1)g^{-1} \cap hN(H_2)h^{-1}$ acts coboundedly on $gN(H_1)^{+r} \cap hN(H_2)^{+r}$ and the quotient space has diameter at most ℓ . □

Definition 3.25 Let $gN(H)$ be a vertex of the commutation graph Y . For each Y -adjacent coset $hN(H')$, the $hH'h^{-1}$ -orbit of a point of $gN(H)^{+r} \cap hN(H')^{+r}$ is called an **incident direction** at $gN(H)$.

Fixing $gN(H)$ and letting $hN(H')$ vary over the vertices of Y that are adjacent to $gN(H)$, we denote by $\mathcal{I}(gN(H))$ the set of all such incident directions.

We advise the reader to look at Fig. 3 for an illustration of this notion. For later use, we observe:

Lemma 3.26 *For each vertex $gN(H)$ of the commutation graph Y , there are finitely many $gN(H)g^{-1}$ -orbits of incident directions at $gN(H)$.*

Proof By Lemma 2.30 and Corollary 2.36, there are finitely many $gN(H)g^{-1}$ -orbits of vertices of Y adjacent to the vertex $gN(H)$. Thus, it is enough to show that for a fixed vertex $hN(H')$ adjacent to $gN(H)$, there are only finitely many $gN(H)g^{-1}$ -orbits of incident directions corresponding to $hH'h^{-1}$ -orbits. This is indeed the case, since $gN(H)g^{-1}$ contains $gN(H)g^{-1} \cap hN(H')h^{-1}$, which acts coboundedly, hence cofinitely (since the Cayley graph of A_Γ is locally finite) on the intersection $gN(H)^{+r} \cap hN(H')^{+r}$ by Lemma 3.23. □

For later purposes we also note the following consequence:

Corollary 3.27 *There are finitely many orbits of edges in Y .*

4 The blown-up commutation graph

We now introduce a variant of the commutation graph of a group. This is this complex X to which we apply Theorem 1.4 for Artin groups of large and hyperbolic type.




Fig. 3 A simplified picture illustrating two adjacent edges of Y , with above each vertex a depiction of the corresponding cosets $N(H)$, $N(H')$, $N(H'')$. Each such coset is quasi-isometric to a product of a tree with a line, and the directions H , H' , H'' are represented vertically in each coset. Two cosets $N(H)$ and $N(H')$ that are adjacent in Y have uniform neighbourhoods that meet along a subset with a cobounded action of the \mathbb{Z}^2 -subgroup $H \times H'$ (shaded regions in the picture). The H' - and H'' -orbits in the corresponding shaded regions have a certain “slope” when seen inside $N(H)$. The action of $N(H)$ on the quasiline $\Lambda_{N(H)}$ from the blow-up data is chosen so that H acts with unbounded orbits on it, but H' and H'' act with bounded orbits on $\Lambda_{N(H)}$

Recall that we have already introduced the standing assumption that our Artin groups are of large hyperbolic type.


Moreover, as per Convention 3.14, we are considering an Artin group (of large and hyperbolic type) A_Γ with Γ connected, and we let \mathcal{H} be as in Definition 3.4. For each $H \in \mathcal{H}$, let $N(H)$ denote its normaliser. Recall that $\bigsqcup_{H \in \mathcal{H}} A_\Gamma / N(H)$ is the vertex set of the commutation graph Y . Our goal is to produce X by modifying Y . Roughly speaking, each vertex of Y will be “blown up” to a cone on a certain discrete space quasi-isometric to \mathbb{Z} , and each edge of Y will be blown up to the simplicial join of the blow-ups of its vertices. Figure 4 and the discussion following Definition 4.7 may help the reader visualise the construction.

In what follows, we equip A_Γ with the word metric d_{A_Γ} coming from the standard generating set given by the vertices of Γ . When we refer to distances in subspaces of A_Γ (e.g. cosets of subgroups), we are using the subspace metric inherited from (A_Γ, d_{A_Γ}) .

4.1 Blow-up data

The action of A_Γ on the commutation graph Y hides too much information to use it as the underlying simplicial complex in a combinatorial HHS structure. For example, subgroups of the form $N(H)$, $H \in \mathcal{H}$, fix vertices in the commutation graph, so by

Fig. 4 A portion of the blow-up X over an edge of Y . The fibres X_v and $X_{v'}$ span a simplicial join in X . The figure illustrates how an edge of X_v and an edge of $X_{v'}$ span a tetrahedron of X



Lemma 3.21, edges have \mathbb{Z}^2 stabilisers. By Lemma 3.9, edges are maximal simplices. Thus any Y -graph (see Definition 1.1) has \mathbb{Z}^2 vertex stabilisers, and so it is not a quasi-isometry model for A_Γ . We remedy this by replacing the commutation graph by a complex X in which the vertices of the commutation graph have been “blown up”.

We now describe the geometric data necessary to perform such a blow-up construction and show that it exists in the case of Artin groups of large and hyperbolic type.

Definition 4.1 (*Blow-up data*) Let $B_0, L_0 \in \mathbb{N}$ be fixed constants, and let r be as in Definition 3.22. Then *blow-up data* for A_Γ and Y consists of the following:

- **(Quasilines associated to cosets.)** For each coset $gN(H)$, $H \in \mathcal{H}$, $g \in A_\Gamma$, let $\Lambda_{gN(H)}$ be a discrete metric space that is (B_0, B_0) -quasi-isometric to \mathbb{Z} . Moreover, $gN(H)g^{-1}$ acts by isometries on $\Lambda_{gN(H)}$ with at most B_0 orbits of points, fixing the Gromov boundary of $\Lambda_{gN(H)}$ pointwise.
- **(Projections to quasilines.)** For each coset $gN(H)$, $H \in \mathcal{H}$, $g \in A_\Gamma$, let

$$\phi_{gN(H)} : gN(H)^{+r} \rightarrow \Lambda_{gN(H)}$$

be a $gN(H)g^{-1}$ -equivariant, (L_0, L_0) -coarsely Lipschitz map.

Blow-up data must satisfy the following additional conditions:

- (A) **(Equivariance.)** For each $g \in A_\Gamma$ and each coset $hN(H)$, there is an isometry $g : \Lambda_{hN(H)} \rightarrow \Lambda_{ghN(H)}$ such that $\phi_{ghN(H)}(gx) = g \cdot \phi_{hN(H)}(x)$ for all $x \in \Lambda_{hN(H)}$.
- (B) **(Unbounded orbits.)** For each coset $gN(H)$, the action of gHg^{-1} on $\Lambda_{gN(H)}$ has unbounded orbits.
- (C) **(Bounded incident directions.)** For each incident direction $hH'h^{-1} \cdot x \in \mathcal{I}(gN(H))$, we have $\text{diam}(\phi_{gN(H)}(hH'h^{-1} \cdot x)) \leq B_0$. (Incident directions are defined in Definition 3.25.)

Some intuition for the definition of blow-up data is provided by Fig. 3.

The reader may want to read the statement of Lemma 4.19 already, in order to understand the importance of the previous conditions, especially condition (C). While the equivariance condition (A) will be necessary for our constructions, it is slightly redundant in the previous definition, as we can start from a “non-equivariant” blow-up data and extend it equivariantly:

Lemma 4.2 *Suppose that for every $H \in \mathcal{H}$, we have:*

- *a discrete metric space $\Lambda_{N(H)}$ that is (B_0, B_0) -quasi-isometric to \mathbb{Z} , with an action of $N(H)$ on it by isometries with at most B_0 orbits of points and which fixes the Gromov boundary pointwise.*
- *a $N(H)$ -equivariant, (L_0, L_0) -coarsely Lipschitz map*

$$\phi_{N(H)} : N(H)^{+r} \rightarrow \Lambda_{N(H)}$$

that satisfy the following additional conditions:

- (B') *For each $H \in \mathcal{H}$, the action of H on $\Lambda_{N(H)}$ has unbounded orbits.*
- (C') *For each coset $kH'k^{-1} \cdot x \in \mathcal{I}(N(H))$, we have $\text{diam}(\phi_{N(H)}(kH'k^{-1} \cdot x)) \leq B_0$.*

Then it is possible to extend equivariantly this data into a blow-up data for Y .

Proof Fix a favourite representative g of the coset $gN(H)$. Let $\Lambda_{gN(H)}$ be a copy of $\Lambda_{N(H)}$. We define $\phi_{gN(H)}(gx) = \phi(x)$ for $x \in N(H)^{+r}$, and the action of $gN(H)g^{-1}$ on $\Lambda_{gN(H)}$ is defined as $gng^{-1} \cdot \lambda = n\lambda$ for $n \in N(H)$ and $\lambda \in \Lambda_{gN(H)}$. (Note that the coset representative g is fixed throughout.)

We have that $\Lambda_{gN(H)}$ is (B_0, B_0) -quasi-isometric to \mathbb{Z} , $gN(H)g^{-1}$ acts by isometries, it has at most B_0 orbits, and the orbits of gHg^{-1} are unbounded. The fact that $\phi_{gN(H)}$ is coarsely Lipschitz is also straightforward.

Let us check equivariance (which is not immediate as we chose coset representatives to define the ϕ maps). Let $g \in A_\Gamma$, and consider a coset $hN(H)$. We can assume that h is the favourite representative of its coset, and let ghn be the favourite representative of $ghN(H)$, where $n \in N(H)$. The isometry $g : \Lambda_{hN(H)} \rightarrow \Lambda_{ghN(H)}$ is defined as $g(x) = n^{-1}x$ (so that we are using the action of $N(H)$ on $\Lambda_{N(H)}$). For $x \in \Lambda_{hN(H)}$ we have

$$\begin{aligned} \phi_{ghN(H)}(gx) &= \phi_{ghN(H)}(ghn n^{-1}h^{-1}x) \\ &= \phi_{N(H)}(n^{-1}h^{-1}x) \\ &= n^{-1}\phi_{N(H)}(h^{-1}x) \\ &= n^{-1}\phi_{hN(H)}(hh^{-1}x) \\ &= n^{-1}\phi_{hN(H)}(x), \end{aligned}$$

as required.

To show C, we note that by equivariance, for $x \in gN(H)^{+r}$ and $hN(H')$ adjacent to $gN(H)$ in Y , we have $\phi_{gN(H)}(hH'h^{-1} \cdot x) = g\phi_{N(H)}(g^{-1}(hH'h^{-1} \cdot x))$.

Note that $g^{-1}(hH'h^{-1} \cdot x) = g^{-1}hH'h^{-1}g \cdot (g^{-1}x) \in N(H)^{+r}$ and $g^{-1}hH'h^{-1}g$ commutes with H , so that $N(H)$ is adjacent to $g^{-1}hN(H')$. In particular $g^{-1}hH'h^{-1} \cdot x$ is in $\mathcal{I}(N(H))$, so that $\phi_{N(H)}(g^{-1}(hH'h^{-1} \cdot x))$ has diameter at most B_0 . Since g is an isometry, the same holds for $\phi_{gN(H)}(hH'h^{-1} \cdot x)$, as required. \square

Existence of blow-up data. We now verify that blow-up data exists for A_Γ and Y . The rest of this subsection is independent of the rest of the article, as the arguments in the next sections rely on the existence of blow-up data, but not on the specific choice of blow-up data. The reader may thus want to skip the rest of this subsection in a first reading.

As we will see, the key idea is to transform quasimorphisms into actions on quasilines. We use a lemma from [2] to do so (the idea behind it can be traced further back, see e.g. [43, Proposition 4.4]). We need the following two general lemmas, which also underpin the strategy followed in the forthcoming paper [32].

Lemma 4.3 (Initial quasimorphism) *Let $1 \rightarrow \mathbb{Z} \xrightarrow{\iota} G \xrightarrow{\pi} F \rightarrow 1$ be a central extension corresponding to a bounded class in $H^2(F, \mathbb{Z})$. Then there exists a quasimorphism $\phi : G \rightarrow \mathbb{Z}$ which is unbounded on $\iota(\mathbb{Z})$.*

Proof Since the central extension corresponds to a bounded class, there is a set-theoretic section $s : F \rightarrow G$ of π such that the element $s(f_1)s(f_2)s(f_1 f_2)^{-1}$ takes finitely many values as f_1, f_2 vary in F . Since s is a section of π , we have $s(f_1)s(f_2)s(f_1 f_2)^{-1} \in \ker(\pi) = \iota(\mathbb{Z})$ for all $f_1, f_2 \in F$, so for each $f_1, f_2 \in F$ we can choose an integer $c(f_1, f_2)$ such that $\iota(c(f_1, f_2)) = s(f_1)s(f_2)s(f_1 f_2)^{-1}$. Note that $|c(f_1, f_2)|$ is bounded independently of f_1, f_2 , because $s(f_1)s(f_2)s(f_1 f_2)^{-1}$ takes only finitely many possible values.

We now define ϕ . Let $x \in G$. Then there exist unique $f_x \in F, t_x \in \mathbb{Z}$ such that $x = s(f_x)\iota(t_x)$. We set $\phi(x) = t_x$. We first verify that ϕ is a quasimorphism. Indeed, let $x, y \in G$. Then, since $\iota(\mathbb{Z})$ is central, we have

$$xy = s(f_x)\iota(t_x)s(f_y)\iota(t_y) = s(f_x)s(f_y)\iota(t_x + t_y) = s(f_x t_y)\iota(c(f_x, f_y) + t_x + t_y).$$

So

$$\phi(xy) = t_x + t_y + c(f_x, f_y) = \phi(x) + \phi(y) + c(f_x, f_y),$$

whence ϕ is a quasimorphism because of the uniform bound on $|c(f_x, f_y)|$.

Finally, observe that for all $t \in \mathbb{Z}$, we have $\phi(\iota(t)) = t$, proving the last part of the lemma. □

Lemma 4.4 (Action on a quasiline) *Let $1 \rightarrow \mathbb{Z} \xrightarrow{\iota} G \xrightarrow{\pi} F \rightarrow 1$ be a central extension with F a hyperbolic group. Let $\{C_\alpha\}_{\alpha \in I}$ be a finite collection of infinite cyclic subgroups of G such that $\{\pi(C_\alpha)\}_{\alpha \in I}$ is a malnormal collection of infinite cyclic subgroups of F . Then G has an (infinite) generating set S such that:*

- $\text{Cay}(G, S)$ is quasi-isometric to \mathbb{Z} ;
- $\iota(\mathbb{Z})$ is unbounded in $\text{Cay}(G, S)$;
- there exists $D \geq 0$ so that for each $x, g \in G$ and each α the diameter of $gC_\alpha g^{-1}x$ is bounded by D .

Proof By Lemma 4.15 in [2], it suffices to produce a quasimorphism $\psi : G \rightarrow \mathbb{R}$ such that ψ is unbounded on $\iota(\mathbb{Z})$, and each $\psi(gC_\alpha g^{-1}x)$ is contained in an interval of uniformly bounded size.

For each $\alpha \in I$, let a_α generate the infinite cyclic group $\pi(C_\alpha)$. So, $\{ \langle a_\alpha \rangle \}_{\alpha \in I}$ is a malnormal collection in F , and each a_α has infinite order.

We now construct a collection of quasimorphisms; the quasimorphism ψ will be a linear combination of these.

First, let $\phi : G \rightarrow \mathbb{Z}$ be the quasimorphism provided by Lemma 4.3, which applies since F is hyperbolic and therefore every cohomology class is bounded [47, Theorem 15]. Let ψ_0 be the homogenisation of ϕ .

Second, for each $\alpha \in I$, we can choose a homogeneous quasimorphism $\psi'_\alpha : F \rightarrow \mathbb{R}$ such that $\psi'_\alpha(a_\alpha) = 1$ and $\psi'_\alpha(a_{\alpha'}) = 0$ for $\alpha' \in I - \{\alpha\}$. This is possible because of malnormality of the collection $\{ \langle a_\alpha \rangle \}_{\alpha \in I}$ of infinite cyclic subgroups (using either the construction in [30] or combining [38, Theorem 4.2] with [25, Corollary 6.6, Theorem 6.8]). So, $\psi_\alpha = \psi'_\alpha \circ \pi : G \rightarrow \mathbb{R}$ is a quasimorphism for each α .

Now let

$$\psi = \psi_0 - \sum_{\alpha \in I} \psi_0(a_\alpha) \cdot \psi_\alpha.$$

Then we claim that ψ is a quasimorphism on G with the required properties. Indeed, $\psi|_{\iota(\mathbb{Z})} = \psi_0|_{\iota(\mathbb{Z})}$, so ψ is unbounded on $\iota(\mathbb{Z})$. Next, note that $\psi(C_\alpha) = \{0\}$, so that for any $c \in C_\alpha$ and $g, x \in G$ we have that $|\psi(gcg^{-1}x) - \psi(x)|$ is bounded by 3 times the defect of ψ , showing that $\psi(gC_\alpha g^{-1}x)$ is contained in an interval of uniformly bounded size. □

We now use the two preceding lemmas to produce blow-up data.

Lemma 4.5 (Existence of blow-up data) *There exists blow-up data for A_Γ and Y .*

Proof Fix $H \in \mathcal{H}$. We have a central extension $H \xrightarrow{\iota} N(H) \xrightarrow{\pi} F$, where $H \cong \mathbb{Z}$ and F is virtually free (and in particular, hyperbolic). This follows from Definition 3.4 and either Lemma 2.27 or Lemma 2.3, according to whether H is generated by a standard generator or by an element of the form z_{ab} .

Choose one element of each $N(H)$ -orbit in $\mathcal{I}(N(H))$; there are finitely many by Lemma 3.26. Each of these is an orbit $g_i H_i g_i^{-1} \cdot x_i$ in $N(H)^{+r}$.

Each $g_i H_i g_i^{-1}$ is infinite cyclic. Moreover, $g_i H_i g_i^{-1} \cap H = \{1\}$ by Lemma 3.21, so $\pi(g_i H_i g_i^{-1})$ is again infinite cyclic.

Observe that $\{\pi(g_i H_i g_i^{-1}) : \alpha \in I\}$ is a malnormal collection of subgroups of F . Indeed, suppose not, so that for distinct i, j , the following holds. Choose elements c, c_i, c_j (either standard generators or elements of the form z_{ab}) generating H, H_i, H_j respectively. Then, for some $k \in N(H)$, there exist positive integers p, q and an integer r such that $kc_j^p k^{-1} = c_i^q c^r$. Now, $c_i^q c^r$ centralises c_i , whence $kc_j^p k^{-1}$ centralises c_i . So, c_i centralises $kc_j k^{-1}$, by Lemma 2.29. Hence $c_i, kc_j k^{-1}$, and c pairwise commute, contradicting Lemma 3.9.

So, Lemma 4.4 yields a generating set S of $N(H)$. Recall that $N(H)^{+r}$ is a finite disjoint union of right cosets $N(H)c_i$ of $N(H)$. We can then consider the graph with vertex set $N(H)^{+r}$ with vertices g, h connected by an edge if $g^{-1}h$ or $h^{-1}g$ belong to $S \cup \{c_i\}$. By [20, Theorem 5.1] (or a direct argument) the vertex set $\Lambda_{N(H)}$ endowed

with the induced metric is $N(H)$ -equivariantly quasi-isometric to the quasiline $\text{Cay}(G, S)$.

We can define $\phi_{N(H)} : N(H) \rightarrow \Lambda_{N(H)}$ to be the identity map, and we observe that $\Lambda_{N(H)}$ and $\phi_{N(H)}$ satisfies the conditions in Lemma 4.2.

Since there are finitely many subgroups $H \in \mathcal{H}$, we can choose constants B_0, L_0 such that each $\Lambda_{N(H)}$ is (B_0, B_0) -quasi-isometric to \mathbb{Z} , the maps $\Phi_{N(H)}$ are (L_0, L_0) -coarsely lipschitz (where $N(H)$ carries the word metric from the fixed finite generating set of A_Γ), and each element of $\mathcal{I}(N(H))$ maps to a set in $\Lambda_{N(H)}$ of diameter at most B_0 , for all H . By construction, there are boundedly many $N(H)$ -orbits in $\Lambda_{N(H)}$.

We thus have that $\{(\Lambda_{N(H)}, \phi_{N(H)})\}$ satisfies conditions (B') and (C') from Lemma 4.2. By Lemma 4.2, it possible to extend the construction equivariantly to obtain condition (A). □

4.2 The blown-up commutation graph

We now define the simplicial complex X . Fix a blow-up data as in Definition 4.1.

Definition 4.6 (*Blown-up commutation graph*) We define a simplicial graph, called the **blown-up commutation graph**, as follows:

- The graph has a vertex for each element of

$$\bigsqcup_{H \in \mathcal{H}} G/N(H),$$

and for each coset $gN(H)$ in the above set, we add $\Lambda_{gN(H)}$ to the vertex set.

- For every $g \in G$ and $H \in \mathcal{H}$, we put an edge between $gN(H)$ and each vertex of $\Lambda_{gN(H)}$. Moreover, whenever $g, h \in G$ and $H, H' \in \mathcal{H}$ are such that gHg^{-1} and $hH'h^{-1}$ are distinct and commute (i.e. when $gN(H)$ and $hN(H')$ are Y -adjacent), we join every vertex in $\{gN(H)\} \cup \Lambda_{gN(H)}$ to every vertex of $\{hN(H')\} \cup \Lambda_{hN(H')}$.

We denote by X the flag completion of this graph.

Observe that A_Γ acts on X by simplicial automorphisms, because of Definition 4.1.(A).

Definition 4.7 (*Projection, fibre, roots and leaves*) There is a natural A_Γ -equivariant simplicial projection $p : X \rightarrow Y$ obtained as follows. At the level of vertices, for $H \in \mathcal{H}$ and $g \in G$, we set

$$p(gN(H)) = gN(H), \quad p(\Lambda_{gN(H)}) = gN(H).$$

Note that two adjacent vertices of X are either sent to the same vertex or to two adjacent vertices of Y . Thus, this definition extends to a map from X to Y . For a simplex Δ of X , we denote by $\overline{\Delta}$ the image simplex $p(\Delta) \subset Y$.

Note that for each vertex $v := gN(H) \in Y$, the preimage $X_v := p^{-1}(v)$, called the **fibre of v** , is a tree that is the simplicial cone over $\Lambda_{gN(H)}$:

$$X_v = \{gN(H)\} * \Lambda_{gN(H)},$$

where “ $*$ ” denotes the simplicial join between two sub-complexes. The **root** of X_v will be the vertex $gN(H)$, while the other vertices of X_v will be called **leaves**.

Thus, X can be thought of as being obtained from Y by blowing-up each vertex $gN(H)$ of Y into a cone over a quasiline, in such a way that fibres X_v and $X_{v'}$ over adjacent vertices v, v' of Y span a simplicial join $X_v * X_{v'}$ in X .

Definition 4.8 We define the map $q : X \rightarrow \widehat{D}$ as $q := \iota \circ p$, where $\iota : Y \rightarrow \widehat{D}$ is the map to the coned-off Deligne complex defined in Lemma 3.7 and Definition 3.12.

Lemma 4.9 *The map $q : X \rightarrow \widehat{D}$ is an A_Γ -equivariant simplicial quasi-isometry.*

Proof The map q is simplicial, being a composition of simplicial maps, and is equivariant by construction. We have that p is a surjective simplicial whose fibres have uniformly bounded diameter (the bound being 2), so that p is a quasi-isometry. Also, ι is a quasi-isometry by Lemma 3.13, and hence q is also a quasi-isometry. \square

4.3 Simplices of the blow-up and their links

In this subsection, we describe the simplices of X and their links, and prove that X satisfies several of the properties from Definition 1.3 and conditions from Theorem 1.4.

Definition 4.10 (*Join decomposition of simplices*) A simplex Δ of X naturally decomposes as a join as follows. For each vertex $v \in \overline{\Delta}$, we define the **fibre** of Δ at v by $\Delta_v := \Delta \cap X_v \subset \Delta$, which is either a single vertex or a single edge of Δ . Then Δ is naturally the join of all its fibres:

$$\Delta = *_{v \in V(\overline{\Delta})} \Delta_v.$$

Lemma 4.11 *Let Δ be a simplex of X . Then its projection $\overline{\Delta}$ is either a vertex or an edge. In particular, X is a 3-dimensional simplicial complex.* \square

Proof Note that the projection $\overline{\Delta}$ is a simplex of Y . It thus follows from Lemma 3.9 that $\overline{\Delta}$ is either a vertex or an edge. Since the simplices of X_v are either vertices or edges, the result follows. \square

Lemma 4.12 *Let Δ be a simplex of X . The link $Lk_X(\Delta)$ decomposes as the following simplicial join:*

$$Lk_X(\Delta) = p^{-1}(Lk_Y(\overline{\Delta})) * (*_{v \in V(\overline{\Delta})} Lk_{X_v}(\Delta_v)).$$

Proof A vertex in $Lk_X(\Delta)$ is adjacent to every vertex of Δ by definition. Since edges of X either project to vertices or edges of Y , we have that $p(Lk_X(\Delta)) \subset St_Y(\overline{\Delta})$. For two adjacent vertices v, v' of Y , every vertex of X_v is adjacent to every vertex of $X_{v'}$. It follows that for $v \in St_Y(\overline{\Delta})$, a vertex $w \in X_v$ is in $Lk_X(\Delta)$ if and only if it is in $Lk_{X_v}(\Delta \cap X_v)$. The result follows. \square

Theorem 1.4 requires that the action of A_Γ on X has finitely many orbits of links, which we now verify.

Lemma 4.13 (Finite index set) *The action of A_Γ on X has finitely many orbits of subcomplexes of the form $\text{Lk}_X(\Delta)$, where Δ is a simplex of X .*

In the proof of the lemma we will give a classification of links which will also be useful later on.

Proof Let Δ be a non-maximal simplex of X and let $\bar{\Delta}$ be the simplex of Y to which Δ projects. We divide into cases according to $\bar{\Delta}$ and the fibers in Δ over the vertices of $\bar{\Delta}$.

- (1) $\bar{\Delta}$ is a vertex $gN(H)$, and Δ is the single vertex $gN(H)$. There are finitely many A_Γ orbits of such Δ , since \mathcal{H} is finite. Hence there are finitely many orbits of links of such Δ .
- (2) $\bar{\Delta}$ is a vertex $gN(H)$, and Δ is a vertex $\lambda \in \Lambda_{gN(H)}$. There are boundedly many $gN(H)g^{-1}$ orbits of λ for the given $gN(H)$, by Definition 4.1, so there are finitely many orbits of such Δ .
- (3) $\bar{\Delta}$ is a vertex $gN(H)$, and Δ is spanned by $\{gN(H), \lambda\}$ for some $\lambda \in \Lambda_{gN(H)}$. Since $gN(H)g^{-1}$ acts on $\Lambda_{gN(H)}$ with boundedly many orbits, and there are finitely many orbits of vertices $gN(H)$, there are finitely many orbits of such simplices Δ .
- (4) $\bar{\Delta}$ is an edge joining $gN(H)$ to $hN(H')$, and Δ is the 1-simplex spanned by $\{gN(H), hN(H')\}$. By Corollary 3.27, there are finitely many orbits of such simplices.
- (5) $\bar{\Delta}$ is an edge joining $gN(H)$ to $hN(H')$, and Δ joins $gN(H)$ to some $\mu \in \Lambda_{hN(H')}$. Observe that all simplices that join $gN(H)$ to a point in $\Lambda_{hN(H')}$ have the same link, so there are at most as many A_Γ orbits of links of such Δ as there are orbits of edges in Y , i.e. finitely many by Corollary 3.27.
- (6) $\bar{\Delta}$ is an edge joining $gN(H)$ to $hN(H')$, and Δ joins $\lambda \in \Lambda_{gN(H)}$ to $\mu \in \Lambda_{hN(H')}$. Then $\text{Lk}_X(\Delta)$ is the edge joining $gN(H)$, $hN(H')$; there are finitely many orbits of these.
- (7) $\bar{\Delta}$ is an edge joining $gN(H)$, $hN(H')$ and Δ is spanned by $\{gN(H), \lambda, hN(H')\}$, where $\lambda \in \Lambda_{gN(H)}$. Then $\text{Lk}_X(\Delta) = \Lambda_{hN(H')}$, and there are finitely many orbits of such subcomplexes since there are finitely many orbits of vertices in Y .
- (8) $\bar{\Delta}$ is an edge joining $gN(H)$, $hN(H')$, and Δ is spanned by $\{gN(H), \lambda, \mu\}$, where $\lambda \in \Lambda_{gN(H)}$ and $\mu \in \Lambda_{hN(H')}$. Then $\text{Lk}_X(\Delta) = hN(H')$, and there are finitely many orbits of such links because Y has finitely many orbits of vertices.
- (9) $\Delta = \emptyset$, in which case $\text{Lk}_X(\Delta) = X$, and there is a single orbit.

This exhausts all the cases. □

The list in the previous proof allows to prove that the blow-up X satisfies the finite complexity condition from Definition 1.3.

Proposition 4.14 (Finite complexity) *There exists N such that the following holds. Let $\Delta_1, \dots, \Delta_n$ be non-maximal simplices of X such that $\text{Lk}_X(\Delta_i) \subsetneq \text{Lk}_X(\Delta_{i+1})$ for $1 \leq i \leq n - 1$. Then $n \leq N$.*

Proof Suppose that Σ, Δ are non-maximal simplices with $\text{Lk}_X(\Sigma) \subsetneq \text{Lk}_X(\Delta)$. Considering the nine types of simplex explained in the proof of Lemma 4.13, we see that links of the same type cannot be properly contained in each other. Hence the length of a chain as in the statement is at most 9. \square

In the rest of this subsection is to prove the following:

Proposition 4.15 (Lattice-ness of links) *The poset of links of simplices of X , ordered by inclusion, has finite height. Moreover, this poset satisfies the following.*

For all simplices Δ_1, Δ_2 of X such that there exists a simplex Σ of X with unbounded link and with the property that $\text{Lk}_X(\Sigma) \subset \text{Lk}_X(\Delta_1) \cap \text{Lk}_X(\Delta_2)$ the following holds. There exists a simplex Δ of X containing Δ_1 , and such that for every simplex Σ of X with unbounded link and such that $\text{Lk}_X(\Sigma) \subset \text{Lk}_X(\Delta_1) \cap \text{Lk}_X(\Delta_2)$, we have $\text{Lk}_X(\Sigma) \subset \text{Lk}_X(\Delta)$.

Proof The finite height condition follows from Proposition 4.14.

Note that $\text{Lk}(\Sigma)$ is unbounded if and only if Σ is either the empty simplex, an edge contained in a fibre of X , or a triangle of X containing exactly two roots, by the classification given in the proof of Lemma 4.13.

Let $\Delta_1, \Delta_2, \Sigma$ be as in the statement.

The case $\Sigma = \emptyset$ is immediate since in that case $\text{Lk}_X(\Sigma) = \text{Lk}_X(\Delta_1) = \text{Lk}_X(\Delta_2) = X$, and we can take $\Delta = \emptyset$. We thus focus on the other two cases.

Case 1: Suppose that there exists such a simplex Σ such that $\text{Lk}_X(\Sigma) \subset \text{Lk}_X(\Delta_1) \cap \text{Lk}_X(\Delta_2)$, and such that Σ is an edge contained in a fibre of X . Then its projection is a single vertex $v \in Y$, and by Lemma 4.12 we have $\text{Lk}_X(\Sigma) = p^{-1}(\text{Lk}_Y(v))$. Since $\text{Lk}_X(\Sigma)$ is contained in $\text{Lk}_X(\Delta_1) \cap \text{Lk}_X(\Delta_2)$, it follows from Lemma 3.9 that Δ_1, Δ_2 are also contained in X_v . Since Δ_1 is either a vertex or an edge of X_v , we can choose an edge Δ contained in $\text{St}_{X_v}(\Delta_1)$. We thus have $\text{Lk}_X(\Delta) = p^{-1}(\text{Lk}_Y(v))$ by Lemma 4.12, and $\text{Lk}_X(\Delta)$ satisfies the required property. Indeed, any link contained in $\text{Lk}_X(\Delta_1) \cap \text{Lk}_X(\Delta_2)$ and strictly larger than $\text{Lk}_X(\Delta)$ must contain a vertex of X_v , and so this link is a simplicial join, hence bounded.

Case 2: Suppose that for every simplex Σ such that $\text{Lk}_X(\Sigma) \subset \text{Lk}_X(\Delta_1) \cap \text{Lk}_X(\Delta_2)$, Σ is a triangle of X with exactly two roots. Choose such a simplex Σ . Then there exists a vertex $v \in Y$ such that $\text{Lk}_X(\Sigma)$ consists of all the leaves of X_v . Moreover, it follows from Lemma 3.9 that for every other simplex Σ' with unbounded augmented link such that $\text{Lk}_X(\Sigma') \subset \text{Lk}_X(\Delta_1) \cap \text{Lk}_X(\Delta_2)$, we have $\text{Lk}_X(\Sigma') \subset X_v$, and in particular $\text{Lk}_X(\Sigma') = \text{Lk}_X(\Sigma)$. Now since $\text{Lk}_X(\Sigma) \subset \text{Lk}_X(\Delta_1)$, then $\overline{\Delta_1}$ is either a vertex or an edge of $\text{St}_Y(v)$.

If $\overline{\Delta_1} = \{v\}$, then we choose a vertex u adjacent to v , an edge f_u of X_u , and we define Δ as the triangle of X spanned by f_u and the root of X_v . We thus have $\Delta \subset \text{St}_X(\Delta_1)$, and Δ satisfies the required conditions.

If $\overline{\Delta_1}$ is an edge containing v , we denote by u the other vertex of $\overline{\Delta_1}$. We choose an edge f_u of X_u that belongs to $\text{St}_{X_u}((\Delta_1)_u)$, and we define Δ as the triangle of X spanned by f_u and the root of X_v . We thus have $\Delta \subset \text{St}_X(\Delta_1)$, and Δ satisfies the required conditions. \square

4.4 From maximal simplices to elements of A_Γ

In this subsection, we associate to each maximal simplex of X a uniformly bounded subset of A_Γ . We first describe the maximal simplices of X :

Remark 4.16 (Description of maximal simplices) Let Δ be a maximal simplex of X . Then Δ has the following form. First, $\overline{\Delta}$ is an edge of Y joining a vertex $gN(H)$ to a vertex $hN(H')$. Then, there are vertices $\lambda \in \Lambda_{gN(H)}$ and $\mu \in \Lambda_{hN(H')}$ such that Δ is the join of the edges $\{gN(H), \lambda\}$ and $\{hN(H'), \mu\}$.

Definition 4.17 For simplices σ_1, σ_2 of X projecting to vertices of Y and spanning a simplex of X , we denote that simplex by $\Delta(\sigma_1, \sigma_2)$.

Recall that one of our goals is to define an X -graph W quasi-isometric to A_Γ , and whose vertex set is the set of maximal simplices of X , see Definition 1.3. We start by defining the vertices of this graph and we construct a map to A_Γ . The edges of W and the quasi-isometry $W \rightarrow A_\Gamma$ will be constructed in Sect. 5.

Definition 4.18 We denote by $W^{(0)}$ the set of maximal simplices of X .

The blow-up data was constructed to obtain the following crucial lemma:

Lemma 4.19 (Bounded sets in A_Γ) *There exists $B_1 \geq 0$ such that the following holds for all $B \geq B_1$. Let $gN(H), hN(H')$ be Y -adjacent cosets. Let $\lambda \in \Lambda_{gN(H)}, \mu \in \Lambda_{hN(H')}$. Then the subset*

$$\phi_{gN(H)}^{-1}(N_B(\lambda)) \cap \phi_{hN(H')}^{-1}(N_B(\mu))$$

of A_Γ is nonempty and has diameter bounded in terms of B . Moreover, $\phi_{gN(H)}(gN(H)^{+r} \cap hN(H')^{+r})$ is B_1 -dense.

We prove the lemma after the following two auxiliary lemmas (coboundedness in the first lemma will only be needed in the next section).

Lemma 4.20 *Suppose C_0, C_1 are infinite cyclic groups and Λ_0 and Λ_1 are quasilines. Suppose that $C_0 \times C_1$ acts on both Λ_i with the property that C_i acts with unbounded orbits on Λ_i and with bounded orbits on Λ_{i+1} (where the action of C_i is the restriction of the action to the subgroup C_i of $C_0 \times C_1$). Then the diagonal action of $C_0 \times C_1$ on $\Lambda_0 \times \Lambda_1$ is proper and cobounded (where the product is given the ℓ^1 metric).*

Proof Fix $\mu_0 \in \Lambda_0, \mu_1 \in \Lambda_1$. Let $L \geq 0$ be given and suppose that $(c_0, c_1), (d_0, d_1) \in C_0 \times C_1$ have the property that $d_{\Lambda_0 \times \Lambda_1}((c_0, c_1) \cdot (\mu_0, \mu_1), (d_0, d_1) \cdot (\mu_0, \mu_1)) \leq L$, i.e.

$$d_{\Lambda_0}((c_0, c_1) \cdot \mu_0, d_0 \cdot \mu_0) + d_{\Lambda_1}((c_0, c_1) \cdot \mu_1, d_1 \cdot \mu_1) \leq L.$$

Let B be the bound on C_i orbits in Λ_{i+1} . Then $d_{\Lambda_0}((c_0, c_1) \cdot \mu_0, (c_0, 1) \cdot \mu_0) \leq B$ and $d_{\Lambda_1}((c_0, c_1) \cdot \mu_1, (1, c_1) \cdot \mu_1) \leq B$. So

$$d_{\Lambda_0}((c_0, 1) \cdot \mu_0, \mu_0) \leq L + B$$

and

$$d_{\Lambda_0}((1, c_1) \cdot \mu_1, \mu_1) \leq L + B.$$

Since C_i has unbounded orbits on the quasiline Λ_i , there are finitely many such c_i , for $i \in \{0, 1\}$. So (c_0, c_1) is one of finitely many elements of $C_0 \times C_1$, as required.

Regarding coboundedness, the action of the cyclic group C_i on the quasiline Λ_i is, say, B_i -cobounded, which easily implies that the action of $C_0 \times C_1$ is $(B_0 + B_1 + 2B)$ -cobounded. □

Lemma 4.21 *Let X and Y be metric spaces, and let $f : X \rightarrow Y$ be a coarsely Lipschitz map. Suppose that a group G acts metrically properly on Y and coboundedly on X , and that f is G -equivariant. Then f is uniformly metrically proper, that is, for every L there exists D such that the preimage under f of any ball of radius L in Y has diameter at most D .*

Proof Let K be such that f is (K, K) -coarsely Lipschitz. Fix $x \in X$ and $R \geq 0$ such that $G \cdot B_R^X(x) = X$. Let $L \geq 0$. By properness, there exists a finite set $F_L \subset G$ such that $d_Y(hf(x), f(x)) \leq L + 2(KR + K)$ implies $h \in F_L$.

We will show that if two points of X map L -close in Y then they are $D(L)$ -close in X ; this is easily seen to imply the lemma (after increasing D).

Suppose $x_0, x_1 \in X$ satisfy $d_Y(f(x_0), f(x_1)) \leq L$. Choose $g_0, g_1 \in G$ such that $d_X(x_i, g_i x) \leq R$ for $i \in \{0, 1\}$. Then $d_Y(g_0 f(x), g_1 f(x)) \leq L + 2(KR + K)$, so $d_Y(g_1^{-1} g_0 f(x), f(x)) \leq L + 2(KR + K)$. Thus $g_1^{-1} g_0 \in F_L$, so

$$d_X(x_0, x_1) \leq 2R + d_X(g_0 x, g_1 x) \leq 2R + \max_{h \in F_L} d_X(x, hx),$$

as required. □

Proof of Lemma 4.19 Fix Y -adjacent cosets $gN(H), hN(H')$. Recall from Lemma 3.21 that $\langle gHg^{-1}, hH'h^{-1} \rangle$ is naturally isomorphic to $gHg^{-1} \times hH'h^{-1}$. By definition of the blow-up data, we have a $gHg^{-1} \times hH'h^{-1}$ -equivariant map

$$\phi_{gN(H)} \times \phi_{hN(H')} : gN(H)^{+r} \cap hN(H')^{+r} \rightarrow \Lambda_{gN(H)} \times \Lambda_{hN(H')}.$$

Equipping $gN(H)^{+r} \cap hN(H')^{+r}$ with the word metric and $\Lambda_{gN(H)} \times \Lambda_{hN(H)}$ with the ℓ_1 -metric, we see that $\phi_{gN(H)} \times \phi_{hN(H')}$ is $(2L_0, 2L_0)$ -coarsely Lipschitz.

Since gHg^{-1} stabilises $hN(H')^{+r}$ (because it is contained in $hN(H')h^{-1}$ by definition of the edges of Y), we have that it also acts on $gN(H)^{+r} \cap hN(H')^{+r}$. Since gHg^{-1} has unbounded orbits in $\Lambda_{gN(H)}$ by Condition (B), we have that $\phi_{gN(H)}$ restricts to a B_1 -coarsely surjective map on $gN(H)^{+r} \cap hN(H')^{+r}$ for some B_1 (proving the moreover part). By equivariance (Condition (A)) and the fact that there are finitely many orbits of edges of Y (Corollary 3.27), the constant B_1 can be chosen independently of the cosets in question.

Let us prove the “nonempty” part. Fix $\lambda \in \Lambda_{gN(H)}$ and $\mu \in \Lambda_{hN(H')}$. By coarse surjectivity, there exists $x \in gN(H)^{+r} \cap hN(H')^{+r}$ such that $d_{\Lambda_{gN(H)}}(\phi_{gN(H)}(x), \lambda) \leq$

B_1 . In view of Condition (C), the $hH'h^{-1}$ -orbit of $\phi_{gN(H)}(x)$ is contained in the $(B_0 + B_1)$ -neighbourhood of λ . Since the action of $hH'h^{-1}$ on $\Lambda_{hN(H)}$ is cobounded (with uniform constant), up to enlarging B_1 we can find an element in $hH'h^{-1}$ -orbit of x in $gN(H)^{+r} \cap hN(H')^{+r}$ that maps B_1 -close to both λ and μ , as required.

We now make the following claim, which will also be used later.

Claim 4.22 *Fix Y -adjacent cosets $gN(H), hN(H')$. Then the map*

$$\phi_{gN(H)} \times \phi_{hN(H')} : gN(H)^{+r} \cap hN(H')^{+r} \rightarrow \Lambda_{gN(H)} \times \Lambda_{hN(H')}$$

is proper. More precisely, there exists $f_0 : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ independent of $gN(H), hN(H')$ such that, for all $s \in \mathbb{R}_{\geq 0}$, the preimage under $\phi_{gN(H)} \times \phi_{hN(H')}$ of any s -ball in $\Lambda_{gN(H)} \times \Lambda_{hN(H')}$ has diameter at most $f_0(s)$. Moreover, the action of $\langle gHg^{-1}, hH'h^{-1} \rangle$ on $\Lambda_{gN(H)} \times \Lambda_{hN(H')}$ is cobounded.

Proof By Lemma 4.20 and Conditions (B) and (C), the action of $gHg^{-1} \times hH'h^{-1}$ on $\Lambda_{gN(H)} \times \Lambda_{hN(H')}$ is proper and cobounded. Then, by Lemma 4.21 (with $f = \phi_{gN(H)} \times \phi_{hN(H')}$) and Lemma 3.23 we have there exists $f_0 : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ such that, for all $s \in \mathbb{R}_{\geq 0}$, the preimage under $\phi_{gN(H)} \times \phi_{hN(H')}$ of any s -ball in $\Lambda_{gN(H)} \times \Lambda_{hN(H')}$ has diameter at most $f_0(s)$. Finiteness of \mathcal{H} and condition (A) imply that this f_0 can be chosen independently of the cosets $gN(H), hN(H')$. \square

Now, if $x, y \in gN(H)^{+r} \cap hN(H')^{+r}$ map B -close to the image of x under both $\phi_{gN(H)}$ and $\phi_{hN(H')}$, then $d_{A_\Gamma}(x, y) \leq f_0(2B)$, yielding the required bound on the diameter of $\phi_{gN(H)}^{-1}(N_B(\lambda)) \cap \phi_{hN(H')}^{-1}(N_B(\mu))$. \square

We are almost ready to fix constants for the rest of the paper, but we first need lemma about the geometry of quasilines.

Lemma 4.23 *Given a quasiline Λ , there exists a constant M_0 such that for all $M \geq M_0$ the following holds. There exists a partial order $<$ on Λ with the following properties:*

- whenever $d_\Lambda(x, y) \geq M$ then x and y are $<$ -comparable.
- for each z there exist z'_1 and z'_2 such that for all z_i with $d_\Lambda(z'_i, z_i) \leq M/10$ we have $z_1 < z < z_2$ and $d_\Lambda(z, z_i) \in [M, 2M]$.
- if $x < y < z$ and $d_\Lambda(x, y), d_\Lambda(y, z) \geq M$ then any geodesic from x to z passes within $M/10$ of y .

Proof Consider a quasi-isometry $f : \Lambda \rightarrow \mathbb{R}$ and define $x < y$ if $f(x) < f(y)$. The details are left to the reader. \square

Convention 4.24 We fix B_0 as in the definition of blow-up-data, then choose B_1 as in Lemma 4.19, with the following additional properties:

- for each Y -adjacent pair $gN(H), hN(H')$ the action of each gHg^{-1} on $\Lambda_{gN(H)}$ is B_1 -cobounded (see Claim 4.22),
- $\phi_{gN(H)}(gN(H)^{+r} \cap hN(H')^{+r})$ is $B_1/100$ -dense,
- $B_1 \geq 10M_0$, where M_0 satisfies Lemma 4.23 for all quasilines $\Lambda_{gN(H)}$.

When $B = B_1$, we denote by B_2 the bound provided by Lemma 4.19.

Fix $gN(H)$. Recall that any $kH'k^{-1} \cdot x \in \mathcal{I}(gN(H))$ has image in $\Lambda_{gN(H)}$ of diameter bounded by B_0 (and B_0 does not depend on H, H', g, k, x), by Definition 4.1.(C).

Let $\delta \in W^{(0)}$ be a maximal simplex spanned by vertices $gN(H), hN(H')$ and $\lambda \in \Lambda_{gN(H)}, \mu \in \Lambda_{hN(H')}$. Recall from Lemma 4.19 that

$$\phi_{gN(H)}^{-1}(N_{B_1}(\lambda)) \cap \phi_{hN(H')}^{-1}(N_{B_1}(\mu))$$

is a nonempty subset of A_Γ , and the corresponding subset with B_1 replaced by $B_1 + B_0$ has diameter (in the word metric) at most B_2 .

Definition 4.25 Let $\delta \in W^{(0)}$ be a maximal simplex spanned by vertices $gN(H), hN(H')$ and $\lambda \in \Lambda_{gN(H)}, \mu \in hN(H')$. Let

$$\sigma(\lambda) = \phi_{gN(H)}^{-1}(N_{B_1}(\lambda))$$

and

$$w(\delta) = \sigma(\lambda) \cap \sigma(\mu).$$

This defines a coarse map $w : W^{(0)} \rightarrow A_\Gamma$.

Lemma 4.26 (Basic properties of w) *The coarse map w is A_Γ -equivariant and sends each maximal simplex to a nonempty subset of A_Γ of diameter at most B_2 . Moreover, if $\delta \in W^{(0)}$ projects to the edge of Y joining cosets $gN(H), hN(H')$, then $w(\delta) \subset gN(H)^{+r} \cap hN(H')^{+r}$.*

Proof Let δ be a maximal simplex, as above. By construction, $w(\delta)$ is

$$\phi_{gN(H)}^{-1}(N_{B_1}(\lambda)) \cap \phi_{hN(H')}^{-1}(N_{B_1}(\mu)),$$

which is nonempty by our choice of B_1 (which we made in Convention 4.24).

Also, by the choice of B_2 , $w(\delta)$ has diameter at most B_2 .

The “moreover” assertion holds by construction, as does A_Γ -equivariance (see Condition A). □

We also have:

Lemma 4.27 (Surjection $W \rightarrow A_\Gamma$) *Each $x \in A_\Gamma$ satisfies $x \in w(\delta)$ for some $\delta \in W^{(0)}$.*

Proof This follows since w is A_Γ -equivariant and the action of A_Γ on itself has one orbit. □

5 The augmented complex

We are considering an Artin group A_Γ of large and hyperbolic type; as per Convention 3.14, we assume Γ is connected. We work with the commutation graph Y and blow-up X constructed in the preceding sections. In this section we construct a combinatorial HHS with underlying simplicial complex X .

5.1 Construction of the augmented complex

Let $r_1, r_2 \geq 0$ be constants to be determined. The first step is to define a graph W_{r_1, r_2} whose vertex set is the set $W^{(0)}$ of maximal simplices of X ; recall that Remark 4.16 and Definition 4.17 describe maximal simplices. The edges of W_{r_1, r_2} are as follows.

Definition 5.1 (*W_{r_1, r_2} -edges*) Let δ_0, δ_1 be maximal simplices of X , where $\delta_i = \Delta(\alpha_i, \beta_i)$. Here, α_i is an edge joining $g_i N(H_i)$ to $\lambda_i \in \Lambda_{g_i N(H_i)}$ and β_i is an edge joining $h_i N(H'_i)$ to $\mu_i \in \Lambda_{h_i N(H'_i)}$.

- **W -edges of type 1:** Suppose $g_0 N(H_0) = g_1 N(H_1)$ and $h_0 N(H'_0) = h_1 N(H'_1)$, and $\mu_0 = \mu_1$. We declare δ_0 to be W_{r_1, r_2} -adjacent to δ_1 if

$$d_{\Lambda_{g_0 N(H_0)}}(\lambda_0, \lambda_1) \leq r_1.$$

Note that simplices joined by an edge of type 1 differ on a single vertex and project to the same edge of Y .

- **W -edges of type 2:** In this case, $h_0 N(H'_0) = h_1 N(H'_1)$ and $\mu_0 = \mu_1$, while $g_0 N(H_0) \neq g_1 N(H_1)$. (So, $p(\delta_0 \cup \delta_1)$ is the union of two distinct edges of Y that share a vertex.) In this case, we declare δ_0 and δ_1 to be W_{r_1, r_2} -adjacent if

$$d_{A_\Gamma}(\sigma(\lambda_0), \sigma(\lambda_1)) \leq r_2.$$

Note that simplices joined by an edge of type 2 intersect in an edge of the fibre over $h_0 N(H'_0) = h_1 N(H'_1)$.

The action of A_Γ on $W^{(0)}$ extends to an action of A_Γ on W by graph automorphisms.

The (finitely many) conditions needed on r_1, r_2 for our constructions to work appear in the proofs of the various results below.

Morally, the definition of W -edges of type 2 should be the inequality in the following lemma. However, we need the definition exactly as stated to show fullness of links (see Proposition 5.6).

Lemma 5.2 *For every r_2 there exists r_3 with the following property. Fix the notation of the “ W -edges of type 2” part of Definition 5.1. Then*

$$d_{A_\Gamma}(w(\delta_0), w(\delta_1)) \leq r_3.$$

Proof We first make two preliminary claims.

Claim 5.3 *There exists a function s such that the following holds. Given cosets $g_0N(H_0)$, $hN(H)$, and $g_1N(H_1)$ that, as vertices of Y , form a path, and $x_i \in g_iN(H_i)^{+r}$ we have*

$$d_{A_\Gamma}(x_0, g_1N(H_1)^{+r} \cap hN(H)^{+r}) \leq s(d_{A_\Gamma}(x_0, x_1)).$$

Proof Set $d = d_{A_\Gamma}(x_0, x_1)$. Up to multiplying all objects on the left by x_0^{-1} , we can assume that x_0 is the identity. Note that there are $N = N(d) < +\infty$ cosets $gN(H')$ with $H' \in \mathcal{H}$ within distance $d + r$ of x_0 . Therefore, it suffices to consider fixed cosets $g_0N(H_0)$, $g_1N(H_1)$ as in the statement, and prove the claimed inequality for x_i in those particular $g_iN(H_i)^{+r}$. Note that since Y does not contain triangles (Lemma 3.9), for each pair $g_0N(H_0)$, $g_1N(H_1)$ there is at most one $hN(H)$ at distance 1 in Y from both. In particular, there are at most finitely many $hN(H)$ that can occur, so that once again we can fix one. But at this point, the left hand side is a number, so we are done. □

Claim 5.4 *There exists a function t such that the following holds. Suppose that $gN(H')$, $hN(H)$ are adjacent vertices of Y and $\lambda \in \Lambda_{gN(H')}$. Given $x \in \sigma(\lambda)$ there exists $y \in \sigma(\lambda) \cap hN(H)^{+r}$ such that $d_{A_\Gamma}(x, y) \leq t(d_{A_\Gamma}(x, gN(H')^{+r} \cap hN(H)^{+r}))$.*

Proof Similarly to the proof of the previous claim, we can assume that x is the identity and that $gN(H')$, $hN(H)$ are fixed. However, λ is not; there are infinitely many possible ones. To overcome this, we will find two candidates for y , and we will see that at least one of them is in $\sigma(\lambda) \cap hN(H)^{+r}$, which suffices.

Towards this, recall the order \prec from Lemma 4.23. From Convention 4.24 we know that the lemma applies with $M = B_1/10$, and that $\phi_{gN(H')}(gN(H')^{+r} \cap hN(H)^{+r})$ is $B_1/100$ -dense. Hence, we can find $y_1, y_2 \in gN(H')^{+r} \cap hN(H)^{+r}$ such that $z_i = \phi_{gN(H')}(y_i) \in \Lambda_{gN(H')}$ satisfy the following:

- $d_{\Lambda_{gN(H')}}(z, z_i) \in [B_1/10, B_1/5]$ where $z = \phi_{gN(H')}(x)$,
- $z_1 \prec z \prec z_2$.

There are now two cases. If $d_{\Lambda_{gN(H')}}(z, \lambda) \leq B_1/2$, then either z_i lies in $N_{B_1}(\lambda)$, that is, either y_i lies in $\sigma(\lambda)$, and we are done. If not, any two points λ, z, z_1, z_2 are either \prec -comparable by construction, or at distance at least $B_1/10$, whence, again, \prec -comparable. If $\lambda \prec z_1$, then it is readily seen from the last item of Lemma 4.23 that $z_1 \in N_{B_1}(\lambda)$, since z lies in said ball and a geodesic from λ to z passes within distance $B_1/100$ of z_1 , so that

$$d_{\Lambda_{gN(H')}}(\lambda, z_1) \leq d_{\Lambda_{gN(H')}}(\lambda, z) - d_{\Lambda_{gN(H')}}(z, z_1) + 2B_1/100 \leq B_1 - B_1/10 + B_1/50 \leq B_1.$$

Similarly, if $z_2 \prec \lambda$ then $z_2 \in N_{B_1}(\lambda)$. On the other hand, we cannot have $z_1 \prec \lambda \prec z$ or $z \prec \lambda \prec z_2$, for otherwise we would have $d_{\Lambda_{gN(H')}}(z, \lambda) < B_1/2$. We now covered all cases. □

We now argue that for every r_2 there exists $r_{2.5}$ so that if $d_{A_\Gamma}(\sigma(\lambda_0), \sigma(\lambda_1)) \leq r_2$ then

$$d_{A_\Gamma}(\sigma(\lambda_0) \cap h_0N(H'_0)^{+r}, \sigma(\lambda_1) \cap h_0N(H'_0)^{+r}) \leq r_{2.5}$$

(That is, if the σ come close, then they come close in the subspace of A_Γ corresponding to the middle vertex $h_0N(H'_0) = p(\delta_0) \cap p(\delta_1)$.)

Indeed, for $x_i \in \sigma(\lambda_i)$ such that $d_{A_\Gamma}(x_0, x_1) \leq r_2$, we can find y_i as in the second claim above, and in view of both claims we have $d_{A_\Gamma}(y_1, y_2) \leq 2t(s(r_2)) + r_2$.

Pick $p \in \sigma(\lambda_0) \cap h_0N(H'_0)^{+r}$ and $q \in \sigma(\lambda_1) \cap h_0N(H'_0)^{+r}$ with $d_{A_\Gamma}(p, q) \leq r_{2.5}$. By translating the pair p, q by an element $h_0H'_0h^{-1}$, we find points p', q' so that

- $d_{A_\Gamma}(p', q') \leq r_{2.5}$ (since we multiplied by an element of A_Γ ,
- $d_{\Lambda_{g_0N(H_0)}}(\phi_{g_0N(H_0)}(p), \phi_{g_0N(H_0)}(p')) \leq B_0$ and $d_{\Lambda_{g_1N(H_1)}}(\phi_{g_1N(H_1)}(q), \phi_{g_1N(H_1)}(q')) \leq B_0$ (Condition C),
- $d_{\Lambda_{h_0N(H'_0)}}(\mu_0, \phi_{h_0N(H'_0)}(p')) \leq B_1$ (Convention 4.24, coboundedness property of B_1).

Furthermore, by Definition 4.1 (blow-up data), the first and third items yield $d_{\Lambda_{h_0N(H'_0)}}(\mu_0, \phi_{h_0N(H'_0)}(q)) \leq L_0r_{2.5} + L_0 + B_1$.

Pick $p'' \in w(\delta_0)$ and $q'' \in w(\delta_1)$. Then we see that p' and p'' are both points in $g_0N(H_0)^{+r} \cap h_0N(H'_0)^{+r}$ that map within bounded distance of $(\lambda_0, \mu_0) \in \Lambda_{g_0N(H_0)} \times \Lambda_{h_0N(H'_0)}$ under $\phi_{g_0N(H_0)} \times \phi_{h_0N(H'_0)}$. By the properness part of Claim 4.22, there is a bound on $d_{A_\Gamma}(p'', p')$. A similar argument applies to q'' , and therefore we get a bound on $d_{A_\Gamma}(p'', q'')$, as required. □

To use Theorem 1.4, we will need the following:

Proposition 5.5 *There exist $R_1, R_2 \geq 0$ such that the following holds for all $r_1 \geq R_1, r_2 \geq R_2$. Let $W = W_{r_1, r_2}$. Then W is connected and the action of A_Γ on W is proper and cobounded. In particular, any orbit map $A_\Gamma \rightarrow W$ is a quasi-isometry.*

Proof The action is proper. Recall the A_Γ -equivariant coarse map $w : W^{(0)} \rightarrow A_\Gamma$ from Definition 4.25. Using Lemma 4.27, choose $\delta_0 \in W^{(0)}$ such that $1 \in w(\delta_0)$. We show that w is coarsely Lipschitz, for any r_1, r_2 , which proves properness of the action since it provides a bound on $d_{A_\Gamma}(w(\gamma_1\delta_0), w(\gamma_2\delta_0)) \geq d_{A_\Gamma}(\gamma_1, \gamma_2) - 2B_2$ when $d_W(\gamma_1\delta_0, \gamma_2\delta_0)$ is bounded.

Let $r_1, r_2 \geq 0$. Let δ_0, δ_1 be W -adjacent maximal simplices of X , where $\delta_i = \Delta(\alpha_i, \beta_i)$ and α_i, β_i are as in Definition 5.1.

First suppose that δ_0, δ_1 are joined by a W -edge of type 2. Then, by Lemma 5.2, we have $d_{A_\Gamma}(w(\delta_0), w(\delta_1)) \leq r_3$.

Second, suppose that δ_0, δ_1 are joined by a W -edge of type 1. Then, by definition, we have the following: associated to the coset $h_0N(H'_0) = h_1N(H'_1)$, we have $\mu_0 = \mu_1$, and associated to the coset $g_0N(H_0) = g_1N(H_1)$, we have $\lambda_0, \lambda_1 \in \Lambda_{g_nN(H_0)}$ with $d_{\Lambda_{g_0N(H_0)}}(\lambda_0, \lambda_1) \leq r_1$.

Consider the map $\phi_{g_0N(H_0)} \times \phi_{h_0N(H'_0)} : g_0N(H_0) \cap h_0N(H'_0) \rightarrow \Lambda_{g_0N(H_0)} \times \Lambda_{h_0N(H'_0)}$, and let $f_0 : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be the (properness) function as in Claim 4.22 (which says that $\phi_{g_0N(H_0)} \times \phi_{h_0N(H'_0)}$ is proper).

Letting $x \in w(\delta_0)$ and $y \in w(\delta_1)$, we have by Lemma 4.26:

$$d_{\Lambda_{g_0N(H_0)}}(\phi_{g_0N(H_0)}(x), \lambda_0) \leq B_1 + B_0$$

and

$$d_{\Lambda_{g_0N(H_0)}}(\phi_{g_0N(H_0)}(y), \lambda_1) \leq B_1 + B_0.$$

In $\Lambda_{h_0N(H'_0)}$, the images of x, y are both $(B_1 + B_0)$ -close to $\mu_0 = \mu_1$.
So,

$$d_{A_\Gamma}(x, y) \leq f_0(4B_1 + 4B_0 + r_1).$$

This completes the proof that w is coarsely lipschitz. As a side remark, the constants depend on A_Γ , the blow-up data, and the choice of r_1, r_2 , but we have not yet needed to impose any restriction on r_1, r_2 .

Connectedness and coboundedness. We will show at the same time that, for sufficiently large r_1, r_2 , W_{r_1, r_2} is connected and the action of A_Γ is cobounded. Note that a set of representatives for the A_Γ -orbits can be taken to be $\{\delta'_i = \Delta(\alpha_i, \beta_i)\}$ where

- the set $\{p(\delta'_i)\}$ of edges of Y is finite (we can arrange this since there are finitely many orbits of edges of Y by Corollary 3.27), and
- there exists a constant C so that if $p(\delta'_i) = p(\delta'_j)$ then the following holds. Suppose that the endpoints of $p(\delta'_i)$ are $gN(H)$ and $hN(H')$ then $d_{\Lambda_{gN(H)}}(\lambda_i, \lambda_j) \leq C$, where $\alpha_k = \{gN(H), \lambda_k\}$, and similarly in $\Lambda_{hN(H')}$.

The second item can be arranged because the action of $\langle gHg^{-1}, hH'h^{-1} \rangle$ on $\Lambda_{gN(H)} \times \Lambda_{hN(H')}$ is cobounded by Claim 4.22 (“moreover” part).

Since Y is connected by Lemma 3.10 (recalling that we are assuming that Γ is connected throughout), to show connectedness and coboundedness it suffices to show that there is a path in W_{r_1, r_2} connecting δ_0 to δ_1 when either

- (1) δ_0, δ_1 are maximal simplices with $p(\delta_0) = p(\delta_1)$, or
- (2) δ_0, δ_1 are maximal simplices such that $p(\delta_0)$ intersects $p(\delta_1)$ at a vertex $gN(H)$,

and that moreover in the first case the length of the path is bounded if we are in the situation of the second bullet point above.

In the first case, it is easily seen that we can in fact connect the maximal simplices only using edges of type 1 (roughly, changing one vertex at a time moving a bounded amount in one of the relevant quasilines). This requires r_1 to be sufficiently large to move in the quasilines.

Since there are only finitely many $gN(H)g^{-1}$ -orbits of edges of Y emanating from $gN(H)$ and $N(H)$ is finitely generated, there is a sequence of edges $p(\delta_0) = e_0, \dots, e_n = p(\delta_n)$ emanating from $gN(H)$ so that the possible pairs (e_i, e_{i+1}) belong to a fixed set of A_Γ -orbits of pairs of edges of Y . From here, we see that we can change the orders of quantifiers, meaning that it suffices to show that given maximal simplices δ_0, δ_1 of X as in case 2, there exist r_1 and r_2 such that δ_0 and δ_1 are W_{r_1, r_2} -adjacent.

To show that this holds, denote $h_i N(H'_i)$ the other endpoint of $p(\delta_i)$, and observe that since the distance between the sets $gN(H)^{+r} \cap h_i N(H'_i)^{+r}$ is at most some R_2 , for any $r_2 \geq R_2$ there is a W -edge of type 2 connecting maximal simplices δ'_0, δ'_1 with $p(\delta_i) = \delta'_i$. We know already that δ_i can be connected to δ'_i , so we are done. \square

To save notation, we often denote W_{r_1, r_2} by W , even though we will not fix r_1, r_2 until later.

5.2 Fullness of links

On our way to proving that (X, W) is a combinatorial HHS, so as to apply Theorem 1.4, we need to verify condition (3) from Definition 1.3, which we do in the following proposition.

Proposition 5.6 (*W-fullness of links*) *Let Δ be a non-maximal simplex of X , and let v_1, v_2 be two distinct non-adjacent vertices of $\text{Lk}_X(\Delta)$ that are contained in W -adjacent maximal simplices. Then v_1, v_2 are contained in W -adjacent maximal simplices of $\text{St}_X(\Delta)$.*

Since W -edges come in two types, we split Proposition 5.6 in two.

Lemma 5.7 (Type 1 fullness) *Let Δ be a non-maximal simplex of X , and let v_1, v_2 be two distinct non-adjacent vertices of $\text{Lk}_X(\Delta)$. Suppose that there exist maximal simplices Δ_1, Δ_2 of X such that $v_1 \in \Delta_1, v_2 \in \Delta_2$, and Δ_1, Δ_2 are connected by a W -edge of type 1. Then v_1, v_2 are contained in W -adjacent maximal simplices of $\text{St}_X(\Delta)$.*

Proof Recall that W -edges of type 1 connect maximal simplices with the same projection in Y . Let $e = \bar{\Delta}_1 = \bar{\Delta}_2$ be the edge of Y to which Δ_1, Δ_2 project, and denote the endpoints of e by u and v . Then $\Delta_i = \Delta(f_u^i, f_v^i)$, where f_u^i is an edge with an endpoint $\lambda_u^1 \in \Lambda_u$, and similarly for v . Also, we have say, $f_u^1 = f_u^2$ and $d_{\Lambda_v}(\lambda_v^1, \lambda_v^2) \leq r_1$.

Since v_1, v_2 are not adjacent in X , we have $v_1 = \lambda_v^1$ and $v_2 = \lambda_v^2$, since those are the only non-common vertices of Δ_1 and Δ_2 . Since $v_1, v_2 \in \text{Lk}_X(\Delta)$, Lemma 4.12 implies that there exist a vertex u' of Y adjacent to v and an edge $f_{u'}$ of $X_{u'}$ such that $f_{u'}$ is contained in $\text{St}_X(\Delta)$. By Definition 5.1, the maximal simplices $\Delta(f_v^1, f_{u'})$ and $\Delta(f_v^2, f_{u'})$ are joined by a type 1 edge in W . These simplices respectively contain v_1, v_2 and lie in the star of Δ , as required. \square

Lemma 5.8 (Type 2 fullness) *Let Δ be a non-maximal simplex of X , and let v_1, v_2 be two distinct non-adjacent vertices of $\text{Lk}_X(\Delta)$. Suppose that there exist maximal simplices Δ_1, Δ_2 of X such that $v_1 \in \Delta_1, v_2 \in \Delta_2$, and Δ_1, Δ_2 are connected by a W -edge of type 2. Then v_1, v_2 are contained in W -adjacent maximal simplices of $\text{St}_X(\Delta)$.*

Proof Since Δ_1 and Δ_2 are connected by a W -edge of type 2, the edges $\bar{\Delta}_1$ and $\bar{\Delta}_2$ share exactly one vertex u of Y . Let us denote by u_1, u_2 the vertices of $\bar{\Delta}_1, \bar{\Delta}_2$, respectively, that differ from u .

By the definition of W -edges of type 2, there is an edge f_u of X_u and edges f_{u_1}, f_{u_2} of X_{u_1}, X_{u_2} such that $\Delta_i = \Delta(f_{u_i}, f_u)$, for $i = 1, 2$. Letting λ_i be the endpoint of f_{u_i} which is not u_i , we also have $d_{A_\Gamma}(\sigma(\lambda_1), \sigma(\lambda_2)) \leq r_2$.

Candidate simplices: Now, since v_1, v_2 are not adjacent in X , we have $v_i \in f_{u_i}$ for $i = 1, 2$. Since Y does not contain triangles, by Lemma 3.9, the vertices u_1, u_2 are at distance 2 in Y . Since $v_1, v_2 \in \text{Lk}_X(\Delta)$, we therefore have by Lemma 4.12 that $\bar{\Delta}$ is a vertex. Since, by Lemma 3.9, Y does not contain squares, it follows that $\bar{\Delta} = u$.

Let f'_u be an edge of X_u containing Δ . Let $\delta_i = \Delta(f'_u, f_{u_i})$ for $i = 1, 2$. So, δ_1, δ_2 are maximal simplices that share the edge f'_u of X_u and project to $\bar{\Delta}_1, \bar{\Delta}_2$ respectively. Note that δ_1 and δ_2 contain f'_u and hence contain Δ . Thus $\delta_1, \delta_2 \subset \text{St}_X(\Delta)$.

W --adjacency: We are only left to observe that there is a W -edge of type 2 connecting δ_1 and δ_2 , since the σ sets associated to δ_i and Δ_i coincide. (That is, roughly, in the definition of edges of type 2 the “middle” edge does not play a role.) \square

Now the proposition is immediate:

Proof of Proposition 5.6 Combine Lemmas 5.7 and 5.8. \square

5.3 Maps between augmented complexes

Recall that the augmented complex X^{+W} was defined in Definition 1.1. We will also use a corresponding construction in Y :

Definition 5.9 We define the **augmented graph** Y^{+W} which is obtained from Y by adding a W -edge between two vertices v, v' of Y whenever their fibres $X_v, X_{v'}$ contain vertices that are connected by a W -edge. In that case, we say that v, v' are **connected by a W -edge**. Similarly, for a full subgraph Y_0 of Y , we denote by Y_0^{+W} the full subgraph of Y^{+W} induced by Y_0 .

Remark 5.10 Note that vertices of Y connected by a W -edge lie at distance at most 2 from each other, by definition of the W -edges.

The simplicial map p mentioned below is defined in Definition 4.7, while ι is defined in Lemma 3.7 and Definition 3.12. Recall from Lemma 3.13 that ι is a simplicial quasi-isometry, and recall also that q is just the composition of p and ι .

Definition 5.11 The simplicial maps

$$p : X \rightarrow Y, \quad q : X \rightarrow \widehat{D}, \quad \iota : Y \rightarrow \widehat{D}$$

extend to maps that we still denote by

$$p : X^{+W} \rightarrow Y^{+W}, \quad q : X^{+W} \rightarrow \widehat{D}, \quad \iota : Y^{+W} \rightarrow \widehat{D}.$$

The extended map $p : X^{+W} \rightarrow Y^{+W}$ is still simplicial, while the other two can be taken to 2-Lipschitz by Remark 5.10

Again by Remark 5.10 the inclusion $Y \hookrightarrow Y^{+W}$ is an A_Γ -equivariant quasi-isometry, so $\iota : Y^{+W} \rightarrow \widehat{D}$ is an A_Γ -equivariant quasi-isometry. Similarly, $p : X^{+W} \rightarrow Y^{+W}$ is an A_Γ -equivariant quasi-isometry. Hence, so is $q : X^{+W} \rightarrow \widehat{D}$.

We will need the following preliminary lemma, in a similar spirit to Proposition 5.5.

Lemma 5.12 *For any sufficiently large r_2 the following holds. Let v be a vertex of Y and let Δ be an edge of X contained in a fibre X_v . Then the augmented links $\text{Lk}_Y(v)^{+W}$ and $\text{Lk}_X(\Delta)^{+W}$ are connected, and p restricts to a quasi-isometry of said augmented links.*

Proof Connectedness: We start by showing that $\text{Lk}_Y(v)^{+W}$ is connected.

Set $v' := \iota(v)$. By Lemmas 2.30 and 2.31, the action of $\text{Stab}(v')$ on $\text{Lk}_{\widehat{D}}(v')$ is cocompact. We choose a finite connected subcomplex K of $\text{Lk}_{\widehat{D}}(v')$ such that its $\text{Stab}(v')$ -orbits cover $\text{Lk}_{\widehat{D}}(v')$. We can further assume that K satisfies the following additional property: Whenever K contains a vertex w that is not of tree or dihedral type, then it also contains every vertex of tree or dihedral type adjacent to it. Indeed, vertices of \widehat{D} that are not of dihedral or tree type are cosets of the form $g\{1\}$ or $g\langle a \rangle$, and they have only finitely many neighbours in D that are vertices of dihedral type by construction of the modified Deligne complex. Moreover, the cone-off procedure adds no additional neighbour for vertices of the form $g\{1\}$ and add only one additional neighbour for vertices of the form $g\langle a \rangle$ (namely, the apex of the standard tree gT_a). Without loss of generality, we can thus assume that K contains all such neighbours.

Let S be the set of vertices of K that are of dihedral or tree type. The set $\iota^{-1}(S)$ forms a finite set of vertices of Y that are adjacent to v by Lemma 3.8. Let us choose an edge e of X_v . For each $w \in \iota^{-1}(S)$, we choose a maximal simplex Δ_w of X above the edge vw of Y and whose fibre over v is the edge e . Since the action of A_Γ on Y is cocompact by Lemma 3.11, we can now require the constant r_2 in the definition of W -edges (see Definition 5.1) to be large enough so that for every distinct w, w' in $\iota^{-1}(S)$, we have that Δ_w and $\Delta_{w'}$ are connected by a W -edge of type 2. In particular, all pairs of vertices of $\iota^{-1}(S)$ are connected by an edge of Y^{+W} .

Let us now show that $\text{Lk}_Y(v)^{+W}$ is connected. Indeed, if x, x' are two vertices of $\text{Lk}_Y(v)^{+W}$, then by definition of the subcomplex K we can find a finite sequence of elements $g_1, \dots, g_n \in A_\Gamma$ such that $\iota(x) \in g_1K, \iota(x') \in g_nK$ and for every $1 \leq i < n, g_iK \cap g_{i+1}K \neq \emptyset$. By the additional property imposed on K , it follows that each such $g_iK \cap g_{i+1}K$ contains a vertex v'_i of \widehat{D} that is either of tree or dihedral type. By Lemma 3.8, we can define $v_i := \iota^{-1}(v'_i)$. We thus have a sequence of vertices $v_0 := x, v_1, \dots, v_{n-1}, v_n := x'$ and by construction, for every $0 \leq i < n$ we have that $g_{i+1}^{-1}v_i$ and $g_{i+1}^{-1}v_{i+1}$ are in $\iota^{-1}(S)$, hence v_i and v_{i+1} are either equal or connected by an edge of Y^{+W} . This show that $\text{Lk}_Y(v)^{+W}$ is connected.

Let us now show that $\text{Lk}_X(\Delta)^{+W}$ is connected. Since Δ is an edge contained in the fibre X_v , it follows from Lemma 4.12 that $\text{Lk}_X(\Delta)$ is the disjoint union of all the fibres X_w for $w \in \text{Lk}_Y(v)$. For every vertex w of $\text{Lk}_Y(v)$, the fibre X_w is connected by construction. Since in addition $\text{Lk}_Y(v)^{+W}$ is connected by the above, it is now straightforward to check that $\text{Lk}_X(\Delta)^{+W}$ is also connected.

Quasi-isometry: By construction of Y^{+W} , for every e edge of $\text{Lk}_Y(v)^{+W}$, there exists an edge e' in $\text{Lk}_X(\Delta)^{+W}$ such that $p(e') = e$. Thus, we can construct a set-theoretic section $\bar{p} : \text{Lk}_Y(v)^{+W} \rightarrow \text{Lk}_X(\Delta)^{+W}$ by choosing for every point $x \in \text{Lk}_Y(v)^{+W}$ a point in $p^{-1}(x)$. Let us show that \bar{p} is a quasi-inverse of p .

We have $p \circ \bar{p} = Id$ by construction. Since all fibres of the form X_w are connected of diameter 2 by construction, it follows that for every vertex $w \in \text{Lk}_Y(v)$ and every point $x \in X_w$, we have $d(x, \bar{p} \circ p(x)) \leq 2$. Moreover, for every other point $x \in \text{Lk}_X(\Delta)^{+W}$, x is thus contained in an edge with endpoints in two distinct fibres X_w and $X_{w'}$, so we have that $\bar{p} \circ p(x)$ is contained in a possibly different edge between X_w and $X_{w'}$. It follows that

$$d(x, \bar{p} \circ p(x)) \leq 2 + \text{diam}(X_w) \leq 4.$$

Moreover, both p and \bar{p} are coarsely Lipschitz. That p is coarsely Lipschitz follows from the fact that it sends an edge of $\text{Lk}_X(\Delta)^{+W}$ to an edge or a vertex of $\text{Lk}_Y(v)^{+W}$. A similar argument as above shows that if x, y are two points belonging to some edge of Y , then $d(\bar{p}(x), \bar{p}(y)) \leq 4$. Thus, \bar{p} is coarsely Lipschitz, and it then follows that p is a quasi-isometry. □

The following lemma is crucial as it allows us to compare links in X (which are the ones we are interested in) to links in \widehat{D} (which are the ones we understand).

Lemma 5.13 *For any sufficiently large r_1, r_2 the following holds for $W = W_{r_1, r_2}$: Let v be a vertex of Y . Then the quasi-isometry*

$$\iota : Y^{+W} \rightarrow \widehat{D}$$

induces a quasi-isometry between $\text{Lk}_Y(v)^{+W}$ and $\text{Lk}_{\widehat{D}}(\iota(v))$.

Moreover, let Δ be an edge of X contained in a fibre X_v , for some vertex v of Y . Then the quasi-isometry

$$q : X^{+W} \rightarrow \widehat{D}$$

restricts to a quasi-isometry $q : (\text{Lk}_X(\Delta)^{+W})^{(0)} \rightarrow \text{Lk}_{\widehat{D}}(\iota(v))$.

Proof Using the actions of A_Γ , we can and will assume that $v = N(H)$ for some $H \in \mathcal{H}$. We first notice that the map ι induces a $\text{Stab}(v)$ -equivariant bijection between the vertices of $\text{Lk}_Y(v)^{+W}$ and the vertices of $\text{Lk}_{\widehat{D}}(\iota(v))$ of tree or dihedral type. Indeed, this follows from Lemma 3.8, the fact that the bijection between tree-type/dihedral-type sets of vertices of Y and \widehat{D} stated in Lemma 3.8 passes to links is a direct consequence of the characterisation of edges of Y therein.

We remark that $\text{Lk}_{\widehat{D}}(\iota(v))$ is connected, since it is either a tree when v is of tree type or we can use Remark 2.33 and Corollary 2.36. We can thus extend the restriction $\iota : \text{Lk}_Y(v) \rightarrow \text{Lk}_{\widehat{D}}(\iota(v))$ into a map $\iota : \text{Lk}_Y(v)^{+W} \rightarrow \text{Lk}_{\widehat{D}}(\iota(v))$ by sending a W -edge between two distinct vertices $w, w' \in \text{Lk}_Y(v)$ to a geodesic path between $\iota(w)$ and $\iota(w')$. Let us show that this extension is coarsely Lipschitz.

Since ι is $\text{Stab}(v)$ -equivariant, it suffices to show that there are finitely many $\text{Stab}(v)$ -orbits of W -edges in $\text{Lk}_Y(v)^{+W}$ (note that a G -equivariant map between graphs $\Gamma_1 \rightarrow \Gamma_2$ is coarsely Lipschitz provided that Γ_2 is connected and there are finitely many G -orbits of edges in Γ_1). Since there are finitely many $\text{Stab}(v)$ -orbits of vertices in $\text{Lk}_Y(v)$ (Lemma 3.26), it suffices to fix a vertex $z = h_1 N(H_1) \in \text{Lk}_Y(v)$ and show that there

are finitely many $(\text{Stab}(v) \cap \text{Stab}(z))$ -orbits of W -edges in $\text{Lk}_Y(v)^{+W}$ containing z . Consider some vertex $z' = h_2N(H_2) \in \text{Lk}_Y(v)$ connected by a W -edge to z , and notice that the W -edge needs to be of type 2 (since W -edges of type 1 connect maximal simplices that project to the same edge in Y). By definition of W -edges, we must have

$$d_{A_\Gamma}(N(H)^{+r} \cap h_1N(H_1)^{+r}, N(H)^{+r} \cap h_2N(H_2)^{+r}) \leq r_2. \quad (*)$$

Since $K = \text{Stab}(v) \cap \text{Stab}(z) = N(H) \cap h_1N(H_1)h_1^{-1}$ acts coboundedly on $N(H)^{+r} \cap h_1N(H_1)^{+r}$ (Lemma 3.23) and \mathcal{H} is finite, there are finitely many K -orbits of cosets $h_2N(H_2)$ satisfying $(*)$, there must be finitely many K -orbits of vertices z' as required.

Let us now construct a coarsely Lipschitz quasi-inverse. By Lemma 2.31 or Lemma 2.30, the action of $\text{Stab}(\iota(v))$ on $\text{Lk}_{\widehat{D}}(\iota(v))$ has finitely many orbits of edges and $\text{Lk}_Y(v)$ is connected by Lemma 5.12, and as above this suffices.

It thus follows that $\iota : \text{Lk}_Y(v)^{+W} \rightarrow \text{Lk}_{\widehat{D}}(\iota(v))$ is a quasi-isometry.

Regarding the moreover clause, in view of Lemma 5.12, $q = \iota \circ p$ restricts to a composition of quasi-isometries of the relevant links. □

In this section we will often make the following abuse of notation:

Convention 5.14 Let A be a subcomplex of a simplicial complex B . Then by $B - A$ we mean the full subcomplex of B with vertex set $B^{(0)} - A^{(0)}$. Note for example that with this abuse we have $(X - \text{Sat}_X(\Delta))^{+W} = (X^{(0)} - \text{Sat}_X(\Delta))^{+W}$.

We also mention the following related result that will be needed in the Sect. 6.2.

Lemma 5.15 *For any sufficiently large r_1, r_2 the following holds. Let Δ be an edge of X contained in a fibre X_v , for some vertex v of Y . If the presentation graph Γ is a single edge, we have $X - \text{Sat}_X(\Delta) = \text{Lk}_X(\Delta)$. Otherwise, $X - \text{Sat}_X(\Delta)$ is the full subcomplex of X spanned by $X - X_v$, and the restriction*

$$q : X - \text{Sat}_X(\Delta) \rightarrow \widehat{D} - \iota(v)$$

is well-defined and extends to a coarsely Lipschitz map

$$\bar{q} : (X - \text{Sat}_X(\Delta))^{+W} \rightarrow \widehat{D} - \iota(v).$$

Proof Case where Γ is a single edge: Let us first consider the case where Γ is a single edge between the two standard generators a, b . Then by construction, Y contains a unique vertex of dihedral type, namely u_{ab} .

Let us first assume that Δ is contained in the fibre of a vertex of tree type. Since Y is bipartite with respect to the vertex type (tree or dihedral) by Lemma 3.8, it follows that for every vertex u of Y of tree type, its link consists of the single vertex u_{ab} . It now follows from Lemma 4.12 that the link of any edge contained in the fibre of a vertex of tree type is equal to the fibre $X_{u_{ab}}$. Moreover, no vertex of $X_{u_{ab}}$ has $X_{u_{ab}}$ in its link, so we get that

$$X - \text{Sat}_X(\Delta) = X_{u_{ab}} = \text{Lk}_X(\Delta).$$

Let us now assume that Δ is contained in the fibre $X_{u_{ab}}$. It follows from Lemma 4.12 that the link of any edge of $X_{u_{ab}}$ is exactly $X - X_{u_{ab}}$. It also follows from the above discussion that no other simplex of X has link $X - X_{u_{ab}}$, hence

$$X - \text{Sat}_X(\Delta) = X - X_{u_{ab}} = \text{Lk}_X(\Delta).$$

General case: Let us now consider the case where Γ is not a single edge. Let us first show that no other vertex of Y has the same link as v . Assume by contradiction that there exists another such vertex v' . Since Y does not contain squares by Lemma 3.9, we have that $\text{Lk}_Y(v) = \text{Lk}_Y(v')$ is a single vertex. Since a vertex of Y of dihedral type gv_{ab} contains at least two vertices in its link, namely gv_a and gv_b by Lemma 3.8, it follows that v and v' are of tree type, say $v = gv_a$ and $v' = hv_b$. This now forces a and b to be leaves of Γ : Indeed, if a had at least two neighbours c, d in Γ , then $\text{Lk}_Y(gv_a)$ would contain the two distinct vertices gv_{ac} and gv_{ad} . Thus, a and b are leaves of Γ . Moreover, since gv_a and hv_b have a common neighbour that is of dihedral type, this implies that this common neighbour is of the form kv_{ab} , and in particular a and b are adjacent in Γ . Since Γ is connected and a and b are two adjacent leaves of Γ , it follows that Γ is a single edge, a contradiction. Thus, no other vertex of Y has the same link as v .

By Lemma 4.12, the link of Δ is the disjoint union of all the fibres of the form X_w for w a vertex of $\text{Lk}_Y(v)$. Moreover, since distinct vertices of Y do not have the same link by the above, a simplex of X has the same link as Δ if and only if it is an edge of X_v . Thus, $X - \text{Sat}_X(\Delta)$ is the full subcomplex of X obtained by removing the fibre X_v .

It follows from the above description of $X - \text{Sat}_X(\Delta)$ that the restriction map

$$q : X - \text{Sat}_X(\Delta) \rightarrow \widehat{D} - \iota(v)$$

is well-defined (recall that q is simplicial on X , so it is enough to check that the restriction is well-defined on the 0-skeleton).

Since $\widehat{D} - \iota(v)$ is connected (since the link of $\iota(v)$ is connected as explained in the proof of Lemma 5.13), we extend this map to a map

$$\bar{q} : (X - \text{Sat}_X(\Delta))^{+W} \rightarrow \widehat{D} - \iota(v)$$

by sending edges to geodesics in $\widehat{D} - \iota(v)$.

To show that this map is coarsely Lipschitz, it is enough to find a uniform constant C such that for two distinct vertices w, w' of $Y - v$ connected by a W -edge, we have $d_{\widehat{D} - \iota(v)}(\iota(w), \iota(w')) \leq C$. Note that since w and w' are distinct, this W -edge must be of type 2, and in particular w and w' are at distance at most 2 in Y . Since Y does not contain triangles and squares by Lemma 3.9, this distance is exactly 2 and we can consider the midpoint w'' .

If $w'' \in Y - v$, then $\iota(w'') \in \widehat{D} - \iota(v)$ since ι is injective by Lemma 3.8, and we thus have $d_{\widehat{D} - \iota(v)}(\iota(w), \iota(w'')) \leq 2$.

Otherwise, $w'' = v$ and so w, w' are adjacent vertices of $\text{Lk}_Y(v)^{+W}$. It follows from Lemma 5.13 that the maps $p : \text{Lk}_X(\Delta)^{+W} \rightarrow \text{Lk}_Y(v)^{+W}$ and $q :$

$\text{Lk}_X(\Delta)^{+W} \rightarrow \text{Lk}_{\widehat{D}}(\iota(v))$ are quasi-isometries. Since w, w' are adjacent vertices of $\text{Lk}_Y(v)^{+W}$, it follows that there exists a constant C (depending only on A_Γ) such that $d_{\text{Lk}_{\widehat{D}}(\iota(v))}(\iota(w), \iota(w')) \leq C$, and hence $d_{\widehat{D}-\iota(v)}(\iota(w), \iota(w')) \leq C$. This shows that the map $\tilde{q} : (X - \text{Sat}_X(\Delta))^{+W} \rightarrow \widehat{D} - \iota(v)$ is coarsely Lipschitz. \square

We will later prove a similar statement for simplices Δ that are triangles of X , see Lemma 6.11.

6 The geometry of the augmented complex

In this section we finish the proof of Theorem A, by verifying that all conditions of Theorem 1.4 apply to the pair (X, W) where X is the blown-up commutation graph (Definition 4.6) and $W = W_{r_1, r_2}$ is described in Sect. 5.1. Up until Sect. 6.3 (the final argument) we work under Convention 3.14, that is, we work with an Artin group of large and hyperbolic type A_Γ with Γ connected and not a single vertex.

6.1 Hyperbolicity of the augmented links

The goal of this section is to prove the following proposition, verifying part of condition (2) from Definition 1.3 for the candidate combinatorial HHS (X, W) :

Proposition 6.1 *For all sufficiently large r_1, r_2 the following holds. The augmented links $\text{Lk}_X(\Delta)^{+W}$ of all non-maximal simplices Δ of X are connected and hyperbolic.*

Before the proof, we assemble the ingredients.

Let Δ be a simplex of X . If $\text{Lk}_X(\Delta)$ is connected and of bounded diameter, so is its augmented link and there is nothing to prove. Thus, we only need focus on unbounded links.

Lemma 6.2 *Let Δ be a non-maximal simplex of X . Then its link (and hence its augmented link) is connected and of bounded diameter, except possibly when Δ is one of the following:*

- the empty simplex,
- an edge of X contained in a fibre,
- a triangle of X containing exactly two roots.

Proof Let Δ be a non-empty simplex of X . Note that if a link decomposes as a non-trivial join, it is automatically connected and bounded. In view of Lemma 4.12, $\text{Lk}_X(\Delta)$ is connected and bounded, by virtue of decomposing non-trivially as a join, except possibly in one of the following two situations.

Case 1: $\text{Lk}_Y(\overline{\Delta}) \neq \emptyset$ and all the links of fibres are empty. In that case, $\overline{\Delta}$ is a single vertex v of Y , and its associated fibre $\Delta_v = \Delta$ is an edge (and this case is listed in the statement).

Case 2: $\text{Lk}_Y(\overline{\Delta}) = \emptyset$ and exactly one of the fibres of Δ has an empty link. In that case, $\overline{\Delta}$ is an edge, and Δ is a triangle (since it is non-maximal). If Δ contains only

one root, then $\text{Lk}_X(\Delta)$ is a single vertex, so it is bounded. (The remaining case of two root is listed in the statement.) \square

We now treat the remaining cases separately.

Lemma 6.3 *The augmented link of the empty simplex of X is equivariantly quasi-isometric to \widehat{D} , and in particular is hyperbolic.*

Proof By definition, the link of the empty simplex of X is X itself, which is quasi-isometric to \widehat{D} by Lemma 4.9. Note that X^{+W} is quasi-isometric to X since W -edges connect vertices at distance at most 2 in X by Definition 5.1. \square

Lemma 6.4 *Suppose $r_1 \geq 1$. Let Δ be a triangle of X containing exactly two roots. Then its augmented link $\text{Lk}_X(\Delta)^{+W}$ is hyperbolic, and in fact quasi-isometric to \mathbb{Z} .*

Proof Let v be the vertex of Y such that Δ_v is a single root. By Lemma 4.12, the link of Δ is the disjoint union of the leaves of X_v . By Proposition 5.6, the augmented link $\text{Lk}_X(\Delta)^{+W}$ is obtained from $\text{Lk}_X(\Delta)$ by adding an edge between all vertices of $\text{Lk}_X(\Delta)$ that are connected by a W -edge coming from maximal simplices of $\text{St}_X(\Delta)$. By construction of W -edges of type 1, this graph is connected and quasi-isometric to \mathbb{Z} , provided $r_1 \geq 1$. Indeed, given a graph and $C \geq 1$, the graph with the same vertex set and edges connecting pairs of vertices within distance C of each other in the original graph is quasi-isometric to the original graph. \square

Lemma 6.5 *Let Δ be an edge of X contained in a fibre X_v for some vertex v of Y . Then its augmented link $\text{Lk}_X(\Delta)^{+W}$ is a connected quasi-tree. Moreover, it is unbounded except when v corresponds to a standard generator of A_Γ that comes from a leaf of Γ contained in an edge of Γ with even label.*

Proof By Lemma 5.13 we have that $\text{Lk}_X(\Delta)^{+W}$ is quasi-isometric to $\text{Lk}_{\widehat{D}}(t(v))$ where $v := \overline{\Delta}$. Therefore, we have to prove that $\text{Lk}_{\widehat{D}}(t(v))$ is a quasi-tree. There are two cases.

Tree type: Suppose that v is a vertex of Y of tree type. By construction of \widehat{D} (or Lemma 2.30) we have that $\text{Lk}_{\widehat{D}}(t(v))$ is a standard tree of D , and the characterisation of when standard trees are unbounded follows directly from the graph of groups structure given in [44, Remark 4.6]. In particular, it follows from that description that when v corresponds to a standard generator of A_Γ that comes from a leaf of Γ contained in an edge of Γ with even label, then the corresponding standard tree has diameter 2.

Dihedral type: Suppose that v is the vertex v_{ab} of Y of dihedral type for some $a, b \in \Gamma$.

By Corollary 2.36, the action of $\text{Stab}(v) = A_{ab}$ on $\text{Lk}_{\widehat{D}}(t(v))$ is cocompact. Let us describe the maximal stabilisers for this action. Since the action is without inversion, it is enough to describe vertex stabilisers. By Lemma 2.31, the stabiliser of a vertex of $\text{Lk}_D(t(v))$ is either trivial or conjugate to $\langle a \rangle$ or $\langle b \rangle$. Moreover, it follows from Lemma 2.27 that the stabiliser of the vertex of tree type corresponding to the standard tree T_a is equal to $A_{ab} \cap C(a)$, i.e. equal to the centraliser of a in A_{ab} . By [23, Lemma 7], this subgroup is equal to $\langle a, z_{ab} \rangle$. Thus, the stabilisers of vertices of tree type in $\text{Lk}_{\widehat{D}}(t(v))$ are conjugate (again in A_{ab}) to $\langle a, z_{ab} \rangle$ or $\langle b, z_{ab} \rangle$.

It follows from the above that the maximal stabilisers for the cocompact action of A_{ab} on $\text{Lk}_{\widehat{D}}(t(v))$ are the conjugates of $\langle a, z_{ab} \rangle$ and $\langle b, z_{ab} \rangle$. It now follows from [20, Thm 5.1] that $\text{Lk}_{\widehat{D}}(t(v))$ is quasi-isometric to the Cayley graph $\text{Cayley}(A_{ab}; \langle a, z_{ab} \rangle, \langle b, z_{ab} \rangle)$. Note that this Cayley graph is quasi-isometric to $\text{Cayley}(A_{ab}/\langle z_{ab} \rangle; \langle a \rangle, \langle b \rangle)$. Moreover, since A_{ab} contains a finite-index subgroup isomorphic to $\langle z_{ab} \rangle \times F$ with F a finitely-generated free group by Lemma 2.3, the central quotient $A_{ab}/\langle z_{ab} \rangle$ is virtually free. By [5, Theorem 2.19] (which applies since $\langle a \rangle$ and $\langle b \rangle$ are quasiconvex in the Cayley graph of $A_{ab}/\langle z_{ab} \rangle$), the Cayley graph $\text{Cayley}(A_{ab}/\langle z_{ab} \rangle; \langle a \rangle, \langle b \rangle)$ is a quasitree, hence so are $\text{Lk}_{\widehat{D}}(t(v))$ and $\text{Lk}_X(\Delta)^{+W}$.

Unboundedness follows from [50, Lemma 5.12] and results in [25] using the same argument as in [5, Corollary 3.22]. □

Corollary 6.6 *The simplices of X that have an unbounded augmented link are exactly:*

- *the empty simplex of X ,*
- *the triangles of X containing exactly two roots.*
- *the edges contained in a fibre X_v of X , except when v corresponds to a standard generator of A_Γ that comes from a leaf of Γ contained in an edge of Γ with even label.*

All other simplices of X have an augmented link of diameter at most 3. □

6.2 Quasi-isometric embeddings of augmented links

It remains to verify that (X, W) satisfies the quasi-isometric embedding part of condition (2) from Definition 1.3. This is achieved by the following proposition:

Proposition 6.7 *For every non-maximal simplex Δ of X , the augmented link $\text{Lk}_X(\Delta)^{+W}$ is quasi-isometrically embedded in the augmented complex $(X^{(0)} - \text{Sat}_X(\Delta))^{+W}$.*

The proposition is clear when the augmented link is bounded. We treat the remaining cases, listed in Corollary 6.6, separately. Note that there is nothing to do in the case of the empty simplex, as the inclusion $\text{Lk}_X(\emptyset)^{+W} \hookrightarrow (X - \text{Sat}_X(\emptyset))^{+W}$ is the identity map of X^{+W} .

Overview. Let us give an overview of the general strategy. We will use the quasi-isometries between X^{+W} , Y^{+W} and \widehat{D} (see Lemma 3.13 and Definition 5.11) to construct models for $\text{Lk}_X(\Delta)^{+W}$ and $(X - \text{Sat}_X(\Delta))^{+W}$ that are easier to work with. More precisely, we will construct a diagram of the form

$$\begin{array}{ccc}
 \text{Lk}_X(\Delta)^{+W} & \xrightarrow{\quad\quad\quad} & U_\Delta \\
 \downarrow & & \downarrow \\
 (X - \text{Sat}_X(\Delta))^{+W} & \xrightarrow{\quad\quad\quad} & V_\Delta
 \end{array}$$

where

- the complexes U_Δ, V_Δ are built out of subcomplexes of \widehat{D} ,
- the top horizontal map is a quasi-isometry (in the spirit of Lemma 5.13),
- the bottom horizontal map is coarsely lipschitz (in the spirit of Lemma 5.15).
- both horizontal maps coincide with the map q when restricted to vertex sets,
- the diagram commutes up to bounded error. This is due to some bounded choices appearing when constructing the various horizontal arrows (and more precisely extending the map q to W -edges), see the proof of Lemma 5.13, 5.15, and 6.11.

Then, it will be enough to show that $U_\Delta \hookrightarrow V_\Delta$ is a quasi-isometric embedding. This in turn will be done using either the CAT(0) geometry of \widehat{D} to construct a coarsely Lipschitz retraction $V_\Delta \rightarrow U_\Delta$ (in Lemma 6.10), or by using additional properties of dihedral Artin groups when the complexes U_Δ, V_Δ are defined using neighbourhoods of vertices of dihedral type (in Lemmas 6.12 and 6.13).

We start by introducing the main tool for constructing coarsely Lipschitz retractions in \widehat{D} , and more generally in CAT(0) complexes:

Definition 6.8 (*Combinatorial neighbourhood*) Let C be a subcomplex of a complex \mathcal{X} . The **combinatorial neighbourhood** of C , denoted $N(C)$, is the subcomplex consisting of the union of all the simplices of \mathcal{X} intersecting C . The **combinatorial sphere** around C , denoted $\partial N(C)$, is the full subcomplex with vertex set $N(C)^{(0)} - C^{(0)}$. Finally, we let $\mathring{N}(C) = N(C) - \partial N(C)$.

For C a subcomplex of a metric complex \mathcal{X} , note that $\mathcal{X} - \mathring{N}(C)$ is the full subcomplex spanned by $\mathcal{X} - C$, meaning the full subcomplex with vertex set the vertices of \mathcal{X} that are not contained in C .

In the statement below we regard $\mathcal{X} - \mathring{N}(C)$ as endowed with its intrinsic path metric, and similarly for $\partial N(C)$.

Lemma 6.9 *Let \mathcal{X} be a CAT(0) simplicial complex with finitely many isometry types of simplices, and let C be a convex subcomplex of \mathcal{X} . Then there exists a coarsely Lipschitz retraction from $\mathcal{X} - \mathring{N}(C)$ to $\partial N(C)$. In particular, $\partial N(C)$ is quasi-isometrically embedded in $\mathcal{X} - \mathring{N}(C)$.*

Proof We construct a projection from $\mathcal{X} - \mathring{N}(C)$ to $\partial N(C)$ as follows. Since the CAT(0) space \mathcal{X} has finitely many isometry types of simplices, we can choose a number $\varepsilon > 0$ such that the closed metric neighbourhood $\mathcal{N}(C, \varepsilon)$ (for the CAT(0) metric) is contained in $N(C)$. Since C , whence $\mathcal{N}(C, \varepsilon)$, is convex for the CAT(0) metric, the closest-point projection $\pi : \mathcal{X} - \mathring{N}(C) \rightarrow \mathcal{N}(C, \varepsilon)$ is well-defined and 1-Lipschitz. Given a point $x \in \mathcal{X} - \mathring{N}(C)$, we first compute the projection $\pi(x)$, we choose a simplex Δ of $N(C)$ containing $\pi(x)$ (such a simplex cannot be contained in C), and we define $\pi'(x)$ to be any vertex of Δ not in C . This defines the map $\pi' : \mathcal{X} - \mathring{N}(C) \rightarrow \partial N(C)$.

Let $x, y \in \mathcal{X} - \mathring{N}(C)$ be two points at distance at most $\varepsilon/3$ in $\mathcal{X} - \mathring{N}(C)$, and therefore in \mathcal{X} . Their closest-point projections $\pi(x), \pi(y)$ are also at distance at most $\varepsilon/3$. Thus, the CAT(0) geodesic γ between $\pi(x)$ and $\pi(y)$ is of length at most $\varepsilon/3$, and in particular γ does not intersect C since its endpoints are at distance ε from C . Also, γ is contained in the combinatorial neighbourhood $N(C)$ since it is even contained in the convex subspace $\mathcal{N}(C, \varepsilon)$. By [14, Corollary 7.30], there is a uniform

bound k depending only on \mathcal{X} and ε but not on x, y such that γ intersects at most k simplices of $N(C)$ (none of them contained in C). This now implies that there is a path in $\partial N(C)$ connecting $\pi'(x)$ and $\pi'(y)$ consisting of a concatenation of at most k paths each contained in a simplex of $\partial N(C)$ (indeed, note that simplices Δ_1, Δ_2 of $N(C)$ that intersect outside of C also satisfy $\Delta_1 \cap \Delta_2 \cap \partial N(C) \neq \emptyset$). Since \mathcal{X} has finitely many isometry types of simplices, the k paths can be taken to have uniformly bounded length, and this implies that the map $\pi' : \mathcal{X} - \overset{\circ}{N}(C) \rightarrow \partial N(C)$ is coarsely Lipschitz.

Moreover, again since \mathcal{X} has finitely many isometry types of simplices, a similar reasoning implies that there exists a constant k' such that for any vertex v of $\partial N(C)$, the CAT(0) geodesic between v and $\pi(v)$ meets at most k' simplices, which implies that v and $\pi'(v)$ lie within bounded distance in $\mathcal{X} - \overset{\circ}{N}(C)$. Thus, π' is a coarse retraction. □

Lemma 6.10 *Let Δ be an edge of X contained in the fibre of a vertex of Y . Then the augmented link $\text{Lk}_X(\Delta)^{+W}$ is quasi-isometrically embedded in $(X - \text{Sat}_X(\Delta))^{+W}$.*

Proof If Γ is a single edge, then we have $\text{Lk}_X(\Delta) = X - \text{Sat}_X(\Delta)$ by Lemma 5.15, hence $\text{Lk}_X(\Delta)^{+W} = (X - \text{Sat}_X(\Delta))^{+W}$, and the inclusion is a quasi-isometric embedding.

Let us now consider the case where Γ is not reduced to a single edge. The edge Δ is contained in the fibre over a vertex that we call $u \in Y$, and we set $v := \iota(u) \in \widehat{D}$. It follows from Lemma 5.15 that $X - \text{Sat}_X(\Delta)$ is the full subcomplex of X obtained by removing the fibre X_u . Recall also from the same lemma the map $\bar{q} : (X - \text{Sat}(\Delta))^{+W} \rightarrow \widehat{D} - \{v\}$. Also, Lemma 5.13 gives us a quasi-isometry $q' : \text{Lk}_X(\Delta)^{+W} \rightarrow \text{Lk}_{\widehat{D}}(v)$ (note that, in the statement of the lemma, the map q is restricted to the 0-skeleton of $\text{Lk}_X(\Delta)^{+W}$, and q' here is an extension of q across all edges).

Consider now the following diagram:

$$\begin{array}{ccc}
 \text{Lk}_X(\Delta)^{+W} & \xrightarrow{q'} & \partial N(v) \\
 \downarrow & & \downarrow \\
 (X - \text{Sat}_X(\Delta))^{+W} & \xrightarrow{\bar{q}} & \widehat{D} - \{v\}
 \end{array}$$

Note that the diagram commutes at the level of vertices, so that it commutes up to bounded arrow since all maps are coarsely Lipschitz.

By Lemma 6.9, there exists a coarsely Lipschitz retraction $\pi : \widehat{D} - \{v\} \rightarrow \partial N(v)$. Thus, the inclusion $\partial N(v) \hookrightarrow \widehat{D} - \{v\}$ is a quasi-isometric embedding. Since the top and right arrows are quasi-isometric embeddings, so is their composition. Since the left arrow is coarsely Lipschitz, and so is the bottom arrow, the left arrow is a quasi-isometric embedding because the diagram commutes up to bounded error. In other words, $\text{Lk}_X(\Delta)^{+W} \hookrightarrow (X - \text{Sat}_X(\Delta))^{+W}$ is also a quasi-isometric embedding. □

Before considering the final two cases where Δ is a triangle of X , we prove the following result that will allow us to show that the bottom map of our (almost) commutative diagram is coarsely Lipschitz. Recall that we have simplicial maps $p : X \rightarrow Y$, and $\iota : Y \rightarrow \widehat{D}$, and we set $q = \iota \circ p$.

Lemma 6.11 *Let Δ be a triangle of X that contains an edge in the fibre of a vertex of Y and an additional root u . Then we have*

$$X - \text{Sat}_X(\Delta) = p^{-1}(Y - \text{St}_Y(u)) \sqcup \Lambda_u.$$

Moreover, let $Z = Z_u$ be the graph obtained from the disjoint union

$$(\widehat{D} - \text{St}_{\widehat{D}}(\iota(u))) \sqcup \Lambda_u^{+W}$$

as follows: We add an edge between a vertex $\lambda \in \Lambda_u$ and a vertex $v \in \widehat{D} - \text{St}_{\widehat{D}}(\iota(u))$ if there exists a W -edge of X^{+W} between λ and a vertex of X_v .

Then the restriction

$$q : X - \text{Sat}_X(\Delta) \rightarrow Z$$

is well-defined and extends to a coarsely Lipschitz map

$$q' : (X - \text{Sat}_X(\Delta))^{+W} \rightarrow Z.$$

Proof It follows from Lemma 4.12 that $\text{Sat}_X(\Delta)$ consists of all the triangles Δ' of X such that $(\Delta')_u = \Delta_u$ (that is, Δ' is the join of the root u and some edge). The structure of $X - \text{Sat}_X(\Delta)$ follows immediately. Since ι is injective by Lemma 3.8, as well as simplicial, we have $\iota^{-1}(\text{St}_{\widehat{D}}(\iota(u))) \subseteq \text{St}_Y(u)$, and therefore $q^{-1}(\text{St}_{\widehat{D}}(\iota(u))) \subseteq p^{-1}(\text{St}_Y(u))$. Therefore, the restriction of q as in the statement is well-defined.

The required coarsely Lipschitz extension exists provided that any two distinct vertices $v, v' \in X - \text{Sat}_X(\Delta)$ joined by an edge get mapped to vertices of Z that lie within uniformly bounded (in particular, finite) distance.

First notice that, by construction, edges that are not W -edges between two distinct vertices of $X - \text{Sat}_X(\Delta)$ are either mapped to an edge of Z or collapsed to a point. Moreover, if $v \in X - \text{Sat}_X(\Delta) - \Lambda_u$ and a point $\lambda \in \Lambda_u$ are joined by a W -edge, then their images under q are joined by an edge in Z by definition of Z . Thus, it is enough to show that there exists a uniform constant C such that if $v, v' \in (X - \text{Sat}_X(\Delta)) - \Lambda_u$ are distinct vertices joined by a W -edge, then $d_Z(q(v), q(v')) \leq C$.

Let $v, v' \in (X - \text{Sat}_X(\Delta)) - \Lambda_u$ be two distinct vertices joined by a W -edge. First notice that $p(v), p(v')$ belong to $Y - \text{St}_Y(u)$ by the first part of the statement. It is enough to consider the case where $p(v)$ and $p(v')$ are not connected by an edge in Y . Since $p(v)$ and $p(v')$ are connected by a W -edge, they are either equal (for W -edges of type 1) or at distance 2 in Y (for W -edges of type 2). It is enough to consider the second case, so let δ, δ' be maximal simplices of X , containing v, v' respectively, that are connected by a W -edge of type 2.

Let u' be the midpoint between $p(v)$ and $p(v')$. If $u' \in Y - \text{St}_Y(u)$, then $q(v), \iota(u'), q(v')$ is a combinatorial path of length 2 in Z , so $d_Z(q(v), q(v')) \leq 2$.

Otherwise, u' must be a vertex of $\text{Lk}_Y(u)$ (since $u' \neq u$ for otherwise v and v' would lie in $\text{Sat}(\Delta)$). Up to the action of A_Γ , we can assume that u, u' are of the form u_a, u_{ab} (up to permutation) and in particular there are finitely many possible pairs to consider. Hence, it suffices to fix u, u' and produce a constant for those. By construction, we have that $w(\delta), w(\delta')$ are contained in the “thickened” coset $N(z_{u'})^{+r}$. Moreover, since δ and δ' are connected by a W -edge, it follows from Lemma 5.2 that $d_{A_\Gamma}(w(\delta), w(\delta')) \leq r_3$.

To show that $q(v), q(v')$ are at uniformly bounded distance in Z , it will be enough to show that they are at uniformly bounded distance in the full subgraph Z' of Z spanned by $(\text{Lk}_{\widehat{D}}(\iota(u')) - \text{St}_{\widehat{D}}(\iota(u))) \cup \Lambda_u$.

We modify the complex Z' in an equivariant way to obtain a graph on which $N(z_{u'})$ acts naturally. Let Z'' be the full subcomplex of X^{+W} spanned by the quasi-lines Λ_w^{+W} for $w \in \text{Lk}_Y(u')$. By construction of the augmented link $\text{Lk}_Y(u')^{+W}$, the projection $p : X^{+W} \rightarrow Y^{+W}$ induces a surjective simplicial map $p : Z'' \rightarrow \text{Lk}_Y(u')^{+W}$. Moreover, since the link $\text{Lk}_Y(u')^{+W}$ is connected by Lemma 5.12 and the quasi-lines Λ_w^{+W} (that is, the fibers of p over the vertices) are connected by Lemma 6.4, it follows that Z'' is connected. Since $\text{Stab}_Y(u') = N(z_{u'})$ stabilises $\text{Lk}_Y(u')^{+W}$ and preserves the family of quasi-lines Λ_w by construction of the action $A_\Gamma \curvearrowright X^{+W}$, it follows that there is an induced action of $N(z_{u'})$ on the connected graph Z'' .

In particular, for $x \in Z''$, the orbit map

$$\omega_x : N(z_{u'}) \rightarrow Z'', g \mapsto g \cdot x$$

is coarsely Lipschitz (with respect to any given word metric on $N(z_{u'})$), and the Lipschitz constant does not depend on $x \in Z''$ since the action is cobounded, see Lemmas 3.26 and 4.5.

We can extend ω_x to a map $\widehat{\omega}_x : N(z_{u'})^{+2r} \rightarrow Z''$ which is still $N(z_{u'})$ -equivariant and coarsely Lipschitz when $N(z_{u'})^{+2r}$ is endowed with the metric inherited from A_Γ .

Moreover, by construction of Z , the projection $q : X^{+W} \rightarrow \widehat{D}^{+W}$ induces a Lipschitz map $q : Z'' \rightarrow Z'$ obtained by collapsing each quasi-line Λ_w^{+W} other than Λ_u to the corresponding vertex $\iota(w)$. Thus, the composition map $N(z_{u'}) \xrightarrow{\widehat{\omega}_x} Z'' \xrightarrow{q} Z'$ is coarsely Lipschitz, where the coarse Lipschitz constant can be chosen to be uniform.

Let us now bound above the distance $d_{Z'}(q(v), q(v'))$. After perturbing $w(\delta)$ up to distance r , we can assume that it is contained in $yN(z) \cap N(z_{u'})^{+2r}$, where $yN(z) = q(v)$ for some element z belonging to the finite set of elements of the form a or z_{ab} (for some $a, b \in \Gamma$). Similarly, we perturb $w(\delta')$ and define y', z' analogously for δ' . We pick as basepoint for the orbit map the point $x := N(z)$. The coarsely Lipschitz map $q \circ \widehat{\omega}_x$ sends $w(\delta)$ to a set (uniformly bounded by 4.26) containing $yN(z) = q(v)$. Moreover, if we denote by k a uniform upper bound on the distance in $\text{Lk}_Y(u')$ between vertices of the form $N(z_1)$ and $N(z_2)$, for z_1, z_2 of the form a or z_{ab} ($a, b \in \Gamma$), then the map $q \circ \widehat{\omega}_x$ sends $w(\delta')$ to a set (again uniformly bounded by 4.26) containing $y'N(z)$, which is at distance at most k from $y'N(z') = q(v')$. Moreover, since $q \circ \widehat{\omega}_x$ is coarsely Lipschitz and $d_{A_\Gamma}(w(\delta), w(\delta')) \leq r_3 + 4r$ (recall the r -perturbation),

it follows that there exists a constant C' such that $d_Z(q(v), q(v')) \leq C'$, hence $d_Z(q(v), q(v')) \leq C'$. This constant C' depends on k, r, r_3 and the coarse Lipschitz constant for $\hat{\omega}_x$ (in particular, it is independent of v, v', u').

We thus have that $d_Z(q(v), q(v'))$ is uniformly bounded above, and it follows that the map q' is coarsely Lipschitz. □

Lemma 6.12 *Let Δ be a triangle of X that contains two roots, and an edge in the fibre of a vertex of tree type of Y . Then the augmented link $\text{Lk}_X(\Delta)^{+W}$ is quasi-isometrically embedded in the augmented complex $(X - \text{Sat}_X(\Delta))^{+W}$.*

Proof Up to the action of A_Γ , we can assume that there exist adjacent generators $a, b \in \Gamma$ such that $\bar{\Delta}$ is the edge between u_a and u_{ab} . It follows from Lemma 4.12 that $\text{Sat}_X(\Delta)$ consists of all the triangles Δ' of X such that $(\Delta')_{u_{ab}} = \Delta_{u_{ab}}$ (that is, Δ' is a join of u_{ab} and some edge). Moreover, by Proposition 5.6, the augmented link $\text{Lk}_X(\Delta)^{+W}$ consists of the quasi-line $\Lambda_{u_{ab}}$, equipped with its W -edges (this augmented link is still quasi-isometric to \mathbb{Z} by Lemma 6.4).

Let $q' : (X - \text{Sat}_X(\Delta))^{+W} \rightarrow Z$ be the coarsely Lipschitz map constructed in Lemma 6.11.

We have the following diagram that commutes up to bounded error (as it commutes exactly on vertices and all maps are coarsely Lipschitz):

$$\begin{array}{ccc}
 \text{Lk}_X(\Delta)^{+W} & \xrightarrow{=} & \Lambda_{u_{ab}} \\
 \downarrow & & \downarrow \\
 (X - \text{Sat}_X(\Delta))^{+W} & \xrightarrow{q'} & Z
 \end{array}$$

In particular, to show that $\text{Lk}_X(\Delta)^{+W} \hookrightarrow (X - \text{Sat}_X(\Delta))^{+W}$ is a quasi-isometric embedding, it is enough to show that the inclusion $\Lambda_{u_{ab}} \hookrightarrow Z$ is a quasi-isometric embedding. By construction of the W -edges, the inclusion $\Lambda_{u_{ab}} \hookrightarrow Z$ is a quasi-isometric embedding provided the $\langle z_{ab} \rangle$ -orbits are quasi-isometrically embedded in Z .

Let Z' be the full subcomplex of \hat{D} generated by $Z \cup \{v_{ab}\}$. By construction of Z (see Lemma 6.11), the subcomplex Z' is obtained from \hat{D} by removing the apices of all the standard trees containing v_{ab} . Since Z' is obtained from \hat{D} by removing a family of open cones, it is still CAT(0) for the induced metric by Lemma 2.24. In particular, it follows from Lemma 6.9 that there is a coarsely Lipschitz retraction

$$Z = Z' - \{v_{ab}\} \rightarrow \text{Lk}_{Z'}(v_{ab}).$$

Thus, it is enough to show that the $\langle z_{ab} \rangle$ -orbits are quasi-isometrically embedded in $\text{Lk}_{Z'}(v_{ab})$. But by construction of Z' , $\text{Lk}_{Z'}(v_{ab})$ is A_{ab} -equivariantly isomorphic to $\text{Lk}_D(v_{ab})$. Thus, it is enough to show that the $\langle z_{ab} \rangle$ -orbits are quasi-isometrically embedded in $\text{Lk}_D(v_{ab})$, where D is the original Deligne complex. Since A_{ab} is of large type, this now follows from Lemma 2.34 applied to $g = z_{ab}$. □

Lemma 6.13 *Let Δ be a triangle of X that contains two roots, and an edge in the fibre of a dihedral vertex of Y . Then the augmented link $\text{Lk}_X(\Delta)^{+W}$ is quasi-isometrically embedded in the augmented complex $(X - \text{Sat}_X(\Delta))^{+W}$.*

The proof of this lemma is rather long, but follows a similar strategy of reduction. The rest of this section will be devoted to it.

Up to the action of the group, we can assume that the projection $\bar{\Delta}$ is an edge of Y between the vertices u_a and u_{ab} . Since the fibre of Δ_{u_a} is the apex of X_{u_a} , the saturation $\text{Sat}_X(\Delta)$ consists of the union of all the triangles of X that have the same fibre at u_a as Δ . In particular, we have

$$X - \text{Sat}(\Delta) = \Lambda_{u_a} \sqcup q^{-1}(\widehat{D} - \widehat{T}_a),$$

and $(X - \text{Sat}(\Delta))^{+W}$ is obtained from the previous graph by adding all the W -edges whose endpoints are in $X - \text{Sat}(\Delta)$.

Step 1: Let $q' = \pi_1 : (X - \text{Sat}_X(\Delta))^{+W} \rightarrow Z$ be the coarsely Lipschitz map constructed in Lemma 6.11. The the following diagram commutes up to bounded error:

$$\begin{array}{ccc} \text{Lk}_X(\Delta)^{+W} & \xrightarrow{=} & \Lambda_a \\ \downarrow & & \downarrow \\ (X - \text{Sat}_X(\Delta))^{+W} & \xrightarrow{\pi_1} & Z \end{array}$$

In particular, to show that the inclusion

$$\text{Lk}_X(\Delta)^{+W} \hookrightarrow (X - \text{Sat}_X(\Delta))^{+W}$$

is a quasi-isometric embedding, it is enough to show that the inclusion

$$\Lambda_a \hookrightarrow Z$$

is a quasi-isometric embedding.

We reduce the problem further, by using the CAT(0) geometry of \widehat{D} to show that Z coarsely retracts onto a neighbourhood of \widehat{T}_a . Consider the combinatorial neighbourhood $\partial N(\widehat{T}_a) = N(T_a) - v_a$. We define the complex Z_a as the full subcomplex of Z spanned by $\Lambda_a \sqcup \partial N(\widehat{T}_a)$.

Step 2: There exists a coarsely Lipschitz retraction $\pi_2 : Z \rightarrow Z_a$ such that the following diagram commutes:

$$\begin{array}{ccc} \Lambda_a & \xrightarrow{=} & \Lambda_a \\ \downarrow & & \downarrow \\ Z & \xrightarrow{\pi_2} & Z_a \end{array}$$

In particular, to show that the inclusion $\Lambda_a \hookrightarrow Z$ is a quasi-isometric embedding, it is enough to show that the inclusion $\Lambda_a \hookrightarrow Z_a$ is a quasi-isometric embedding.

Proof of Step 2 Since T_a is convex in $\widehat{D} - v_a$ by Lemma 2.25, there exists a coarse retraction of $\widehat{D} - \widehat{T}_a$ on $\partial N(\widehat{T}_a)$ by Lemma 6.9. We define the map $\pi_2 : Z \rightarrow Z_a$ as follows: π_2 is the identity on Z_a and is the coarse retraction $\widehat{D} - \widehat{T}_a \rightarrow \partial N(\widehat{T}_a)$ on $Z - \Lambda_a$. This definition makes sense as $Z_a \cap (Z - \Lambda_a) = \partial N(\widehat{T}_a)$ and π_2 is defined as the identity on $\partial N(\widehat{T}_a)$ in both cases. To conclude that this defines a coarse retraction $Z \rightarrow Z_a$, let us check that an edge between a vertex x of $Z - \Lambda_a$ and a vertex y of Z_a is either contained in $Z - \Lambda_a$ or in Z_a . If $y \notin \Lambda_a$, then x, y belong to $Z - \Lambda_a$, and so does the edge between them. If $y \in \Lambda_a$, then by construction the edge e between x and y comes from a W -edge of type 2 of X^{+W} . But by construction of W -edges of type 2, this forces x to belong to $\partial N(T_a)$, and thus x, y , and e belong to Z_a . \square

We now want to show that the inclusion $\Lambda_a \hookrightarrow Z_a$ is a quasi-isometric embedding. For a vertex v of T_a , we denote by Z_v the subgraph of Z_a spanned by Λ_a and $\partial N(\widehat{T}_a) \cap N(v)$. In particular, Z_a is the union of all the Z_v , where v runs among the vertices of T_a . For a generator $b \in \Gamma$ adjacent to a , we simply denote by Z_{ab} the graph $Z_{v_{ab}}$. To show that the inclusion $\Lambda_a \hookrightarrow Z_a$ is a quasi-isometric embedding, we will first show that each inclusion $\Lambda_a \hookrightarrow Z_v$ is a quasi-isometric embedding (Steps 3 and 4) and then combine these quasi-isometric embeddings to show that $\Lambda_a \hookrightarrow Z_a$ is a quasi-isometric embedding (Step 5).

Before reducing the situation further, we need further information about the graph of orbits introduced in Definition 2.32 and some variants. As the graphs we are looking at have been obtained from $\text{Lk}_D(v_{ab})$ by removing the vertices corresponding to the tree T_a , we introduce a variant of the graph of orbits:

Definition 6.14 Let \mathcal{G} be either $\text{Cayley}_{a,b}(A_{a,b})$ or $\text{Cayley}_{a,b}(A_{a,b}) / \langle \Delta_{ab} \rangle$ (where as usual, $\langle \Delta_{ab} \rangle$ acts by multiplication on the right). We define a new graph, called the **blown-up graph of orbits** and denoted $\text{Orbit}_{a,b}^*(\mathcal{G})$ as follows.

We denote by L_a the simplicial line in \mathcal{G} spanned by the $\langle a \rangle$ -orbit of the identity element. We put a vertex for every $\langle a \rangle$ -orbit of \mathcal{G} except the one corresponding to L_a , one vertex for every $\langle b \rangle$ -orbit in \mathcal{G} , and one vertex for each vertex of L_a . If two such subsets of \mathcal{G} have a non-empty intersection, we put an edge between them. Moreover, for every vertex v of L_a , we also put an edge between v and av .

This construction can be thought of as obtained from the original graph of orbits $\text{Orbit}_{a,b}^*(\mathcal{G})$ by blowing up the vertex corresponding to the coset $\langle a \rangle$, and replacing it with a copy of L_a .

Step 3: There exists a coarsely Lipschitz map

$$\pi_3 : Z_{ab} \rightarrow \text{Orbit}_{a,b}^*(\text{Cayley}_{a,b}(A_{a,b})) / \langle \Delta_{ab} \rangle$$

and a quasi-isometry $\Lambda_a \rightarrow L_a$ such that the following diagram commutes up to bounded error and the inclusion $\Lambda_a \hookrightarrow Z_{ab}$ coarsely Lipschitz:

$$\begin{array}{ccc}
 \Lambda_a & \xrightarrow{\cong} & L_a \\
 \downarrow & & \downarrow \\
 Z_{ab} & \xrightarrow{\pi_3} & \text{Orbit}_{a,b}^*(\text{Cayley}_{a,b}(A_{a,b})) / \langle \Delta_{ab} \rangle
 \end{array}$$

In particular, to show that the inclusion $\Lambda_a \hookrightarrow Z_{ab}$ is a quasi-isometric embedding, it is enough to show that the inclusion

$$L_a \hookrightarrow \text{Orbit}_{a,b}^*(\text{Cayley}_{a,b}(A_{a,b})) / \langle \Delta_{ab} \rangle$$

is a quasi-isometric embedding.

Proof of Step 3 First notice that one can identify the subgraph $Z_{ab} - \Lambda_a$ with the subgraph $\text{Lk}_{\widehat{D}}(v_{ab}) - \text{Lk}_{\widehat{D} \cap T_a}(v_{ab})$.

Now consider the (coarsely Lipschitz) quotient map $Z_{ab} \rightarrow Z_{ab} / \langle \Delta_{ab} \rangle$ for the action of $\langle \Delta_{ab} \rangle$ on the right. The image in that quotient of the apex of a standard tree is a vertex of valence 1. The graph $Z_{ab} / \langle \Delta_{ab} \rangle$ thus retracts via a Lipschitz map onto the subgraph obtained by removing all these valence 1 vertices. Using the identification of $\text{Lk}_D(v_{ab})$ with the graph of orbits $\text{Orbit}_{a,b}(\text{Cayley}(A_{ab}))$, one sees that the graph obtained is the blown-up graph of orbits $\text{Orbit}_{a,b}^*(\text{Cayley}_{a,b}(A_{a,b})) / \langle \Delta_{ab} \rangle$. The composition

$$\pi_3 : Z_{ab} \rightarrow Z_{ab} / \langle \Delta_{ab} \rangle \rightarrow \text{Orbit}_{a,b}^*(\text{Cayley}_{a,b}(A_{a,b})) / \langle \Delta_{ab} \rangle$$

satisfies the required property. Moreover, π_3 is $\langle a \rangle$ -equivariant.

Pick a bounded set B whose $\langle a \rangle$ -translates cover Λ_a containing exactly one element of each orbit. There is a unique $\langle a \rangle$ -equivariant map $\Lambda_a \rightarrow L_a$ that sends B to the identity element. This is a quasi-isometry since $\langle a \rangle$ acts coboundedly on both Λ_a and L_a . We use this as the top arrow in the diagram. Since the other three maps are also $\langle a \rangle$ -equivariant, the diagram commutes up to bounded error.

Note that the map $\Lambda_a \hookrightarrow Z_{ab}$ is simplicial, hence coarsely Lipschitz. □


Step 4: The inclusion

$$L_a \hookrightarrow \text{Orbit}_{a,b}^*(\text{Cayley}_{a,b}(A_{a,b})) / \langle \Delta_{ab} \rangle$$

is a quasi-isometric embedding.

Proof of Step 4 Let us start with a preliminary remark. Starting from the Cayley graph of A_{ab} , there are two operations one can perform: We can mod out by the right-action of $\langle \Delta_{ab} \rangle$, or we can construct the aforementioned blown-up graph of orbits $\text{Orbit}_{a,b}^*(\text{Cayley}_{a,b}(A_{a,b}))$. These two constructions commute. More precisely, there

exists an isomorphism (in dotted lines below) that makes the following diagram commute:



Thus, it is enough to show that the inclusion

$$L_a \hookrightarrow \text{Orbit}_{a,b}^*(\text{Cayley}_{a,b}(A_{a,b})/\langle \Delta_{ab} \rangle)$$

is a quasi-isometric embedding. The graph $\text{Cayley}_{a,b}(A_{a,b})/\langle \Delta_{ab} \rangle$ (for the action of $\langle \Delta_{ab} \rangle$ on the right) is a tree, as it is obtained from the quasi-tree \mathcal{T}_{ab} defined in Sect. 2.2 by removing the edges of \mathcal{T}_{ab} labelled by an atom on more than one generator. Moreover, the various gL_a and gL_b are quasi-lines in the tree $\text{Cayley}_{a,b}(A_{a,b})/\langle \Delta_{ab} \rangle$ that are $g\langle a \rangle g^{-1}$ -invariant and $g\langle b \rangle g^{-1}$ -invariant respectively. Moreover, for any $r \geq 0$ there is a constant $f(r)$ such that the intersection of the r -neighbourhoods of any two such distinct quasi-lines has diameter at most $f(r)$. The fact that $L_a \hookrightarrow \text{Orbit}_{a,b}^*(\text{Cayley}_{a,b}(A_{a,b})/\langle \Delta_{ab} \rangle)$ is a quasi-isometric embedding is now a consequence of Step 3 and the following Claim, applied to the case $Z = \text{Cayley}_{a,b}(A_{a,b})/\langle \Delta_{ab} \rangle$, $L_{i_0} = L_a$, and the L_i are the gL_b and the gL_a other than L_a .

Claim. For all δ, κ and functions $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, there exists $D \geq 0$ such that the following holds. Let Z be a δ -hyperbolic space and let $\{L_i\}$ be a family of (κ, κ) -quasilines in Z such that for all distinct i, j and all r , the intersection of the r -neighbourhoods of L_i and L_j has diameter at most $f(r)$. Then the following holds for all L_{i_0} . Let \widehat{Z} be obtained from Z by coning off each $L_i, i \neq i_0$. Then the image of L_{i_0} under the inclusion $Z \rightarrow \widehat{Z}$ is a (D, D) -quasi-line.

Proof of Claim. The closest-point projection $proj : Z \rightarrow L_{i_0}$ has the property that the image of all $L_i, i \neq i_0$ is uniformly bounded. Therefore, $proj$ gives a coarsely Lipschitz coarse retraction of \widehat{Z} onto the image of L_{i_0} in \widehat{Z} . The existence of a coarse retraction implies that L_{i_0} is quasi-isometrically embedded in \widehat{Z} , as required. \square

This concludes the proof of Step 4. \square

It follows from Steps 3 and 4 that the various inclusions $\Lambda_a \hookrightarrow Z_v$ are quasi-isometric embeddings, for vertices v of T_a . Note that because the action of $\text{Stab}(T_a)$ on T_a is cocompact by Lemma 2.30, we can choose a constant $K \geq 0$ such that all such inclusions $\Lambda_a \hookrightarrow Z_v$ are (K, K) -quasi-isometric embeddings. We now want to combine these various quasi-isometric embeddings into a global quasi-isometric embedding $\Lambda_a \hookrightarrow Z_a$.

Step 5: The inclusion $\Lambda_a \hookrightarrow Z_a$ is a quasi-isometric embedding.

Proof of Step 5 For every distinct vertices v, w of T_a , notice that vertices in $Z_v \cap Z_w$ are at (combinatorial) distance at most 1 from Λ_a . In particular, given a finite combinatorial geodesic γ in Z_a between two points x, x' of Λ_a , there exist points $x_0 := x, x_1, \dots, x_n := x'$ of Λ_a , and distinct points $y_0 := x, y_1, \dots, y_n := x'$ of $\gamma \subset Z_a$ such that the portion of γ between y_i and y_{i+1} is contained in some Z_v , and each y_i is either equal or adjacent to x_i . From Steps 3 and 4 and the definition of the x_i and y_i , it follows that for each i , we get the following inequality between combinatorial length:

$$d_{\Lambda_a}(x_i, x_{i+1}) \leq K d_{Z_v}(y_i, y_{i+1}) + K + 2 \leq (2K + 2) d_{Z_v}(y_i, y_{i+1}),$$

the last inequality following from the fact that $d_{Z_v}(y_i, y_{i+1}) \geq 1$. We thus have

$$d_{\Lambda_a}(x, x') \leq \sum_i d_{\Lambda_a}(x_i, x_{i+1}) \leq (2K + 2) \sum_i d_{Z_v}(y_i, y_{i+1}) = (2K + 2) d_{Z_a}(x, x').$$

This shows that the inclusion $\Lambda_a \hookrightarrow Z_a$ is a quasi-isometric embedding. □

Combining all the previous steps, it follows that the inclusion $\text{Lk}_X(\Delta)^{+W} \hookrightarrow (X - \text{Sat}(\Delta))^{+W}$ is a quasi-isometric embedding, which concludes the proof of Lemma 6.13.

Proof of Proposition 6.7 The preceding sequence of Lemmas 6.10, 6.12, 6.13, together with Corollary 6.6, imply the proposition. □

6.3 Obtaining hierarchical hyperbolicity

We are now ready to prove the main theorem of this article, Theorem A, which says that Artin groups of large and hyperbolic type are hierarchically hyperbolic. First, we just prove that a structure exists, and later on we give a more detailed description for the interested reader.

Proof of Theorem A There are three cases to consider.

Single vertex: If Γ is a single vertex, then A_Γ is isomorphic to \mathbb{Z} , which is an HHG.

Connected case: We consider the case where Γ is connected and contains at least two vertices, so that we are in the case covered by Convention 3.14. Let us first check that (X, W) satisfy the conditions of Definition 1.3: It satisfies (I) by Proposition 4.14, condition (II) by Propositions 6.1 and 6.7, condition (III) by Proposition 5.6, and condition (IV) by Proposition 4.15 and Corollary 6.6. Moreover, the action of A_Γ is proper and cobounded by Proposition 5.5 and there are finitely many orbits of links by Proposition 4.13. It thus follows from Theorem 1.4 that A_Γ is hierarchically hyperbolic.

General case: In general, A_Γ is the free product of Artin groups on connected graphs, which are hierarchically hyperbolic by the above. Taking free products preserves the property of being a hierarchically hyperbolic group, by [9, Corollary 8.24] or [9, Theorem 9.1]. □

Description of the HHG structure We will now describe the HHG structure, and we introduce some notation to this end; we will use standard HHS terminology from [9, Definition 1.1].

Given an HHS (X, \mathfrak{S}) , we can write $\mathfrak{S} = \mathfrak{S}_{fin} \sqcup \mathfrak{S}_\infty$, where \mathfrak{S}_∞ is the set of indices for which the corresponding hyperbolic space is unbounded. We typically have various bounded hyperbolic spaces in an HHS structure which need to be there because of the orthogonality axiom, but do not otherwise play a major role. The interesting part of the structure is \mathfrak{S}_∞ .

Recall from Definition 3.4 that given an Artin group A_Γ we denote

$$\mathcal{H} := \{ \langle a \rangle \mid a \text{ a standard generator of } A_\Gamma \} \cup \{ \langle z_{ab} \rangle \mid a, b \text{ span an edge of } \Gamma \text{ with } m_{ab} \geq 3 \}.$$

We denote two distinct copies of $\bigsqcup_{H \in \mathcal{H}} G/N(H)$ by \mathfrak{S}^{line} and \mathfrak{S}^{tree} . We denote the copy of $gN(H)$ in \mathfrak{S}^{line} by $gN(H)^{line}$, and similarly for \mathfrak{S}^{tree} .

Finally, we denote by \mathfrak{S}^{leaf} the subset of \mathfrak{S}^{tree} consisting of all cosets of the form $gN(\langle a \rangle)^{tree}$, where a is a standard generator corresponding to a valence-1 vertex of Γ whose incident edge has even label.

Theorem 6.15 *Let A_Γ be an Artin group of large and hyperbolic type, with Γ connected and not a single vertex. Then there is an HHG structure \mathfrak{S} on A_Γ with the following properties. We have $\mathfrak{S}_\infty = \{S\} \sqcup \mathfrak{S}^{line} \sqcup (\mathfrak{S}^{tree} - \mathfrak{S}^{leaf})$, and all of the following hold:*

- (1) S is the \sqsubseteq -maximal element of \mathfrak{S} and $\mathcal{C}(S)$ can be taken to be either the coned-off Deligne complex or the complex of irreducible parabolics of finite type.
- (2) For each $gN(H) \in \mathfrak{S}^{line}$, $\mathcal{C}(gN(H))$ is a quasiline on which $gN(H)g^{-1}$ acts, with gHg^{-1} having unbounded orbits.
- (3) For each $gN(H) \in \mathfrak{S}^{tree} - \mathfrak{S}^{leaf}$, the space $\mathcal{C}(gN(H))$ is an unbounded quasitree.
- (4) For each $gN(H)$, we have $gN(H)^{line} \perp gN(H)^{tree}$.
- (5) $gN(H)^{line}$ is nested in $hN(H')^{tree}$ if and only if gHg^{-1} and $hH'h^{-1}$ are distinct and commute.
- (6) $gN(H)^{line}$ is orthogonal to $hN(H')^{line}$ if and only if gHg^{-1} and $hH'h^{-1}$ are distinct and commute.
- (7) The items above completely determine the nesting and orthogonality relation within \mathfrak{S}_∞ .
- (8) No element of \mathfrak{S} is orthogonal to any of its A_Γ -translates.

Proof The HHG structure can be obtained from the description of the HHS structure for a combinatorial HHS [6, Theorem 1.18, Remark 1.19]. The index set of the HHS structure is in bijection with the set of links of simplices, and the hyperbolic spaces in the HHS structure are the augmented links. The identification of the unbounded augmented links is done in Corollary 6.6, with the (equivariant quasi-isometry types of the) augmented links being described in Lemma 6.3 (see also Proposition 3.16), 6.4, and Lemma 6.5.

Orthogonality and nesting of a combinatorial HHS are described in [6, Definition 1.11], and they correspond to links that form joins and links that are nested into each other. The link of the empty simplex will be the \sqsubseteq -maximal element, and a

straightforward case-by-case analysis with the other two types of links yields the relations are in our case (in view of how adjacency in the commutation graph is defined).

Finally, let us prove the last item; for this we argue with the HHS structure of the combinatorial HHS (X, W) that we used throughout the section. Consider by contradiction non-maximal simplices $\Delta, h\Delta$ such that their links form a join. Consider any $x \in \text{Lk}(\Delta)$; in particular, x and hx are adjacent in X . There are two cases. If $x = gN(H)$ for some $H \in \mathcal{H}$ and $g \in A_\Gamma$, then x and hx are also adjacent vertices of the commutation graph Y . However, Y is bipartite (with respect to dihedral and tree type vertices, see Lemma 3.8) and no adjacent vertices are in the same orbit. On the other hand, if $x \in \Lambda_{gN(H)}$, then we can repeat the same argument with $gN(H)$. \square

Remark 6.16 The last item of Theorem 6.15 implies that the HHG structure is A_Γ -colourable in the sense of [29, Definition 2.8].

Remark 6.17 In the dihedral case, it does not suffice to simply observe that A_Γ is virtually free-times- \mathbb{Z} to show that it is a hierarchically hyperbolic group. Indeed, while the property of being a hierarchically hyperbolic space is preserved by quasi-isometries, the extra equivariance properties needed to be a hierarchically hyperbolic group fail badly to be preserved even by commensurability. Indeed, while \mathbb{Z}^2 is a hierarchically hyperbolic group, there are virtually \mathbb{Z}^2 groups, like the $(3, 3, 3)$ triangle group, that are not [52, Corollary 4.5].

We now turn to the consequence of hierarchical hyperbolicity stated in Corollary B.

Proof of Corollary B Assertion 1 follows directly from [7, Thm. A].

All of the assertions follow from Theorem A and general results about hierarchically hyperbolic groups, as mentioned in the statement, except (uniform exponential growth) and (3) (on stable subgroups).

We now handle assertion (2) (uniform exponential growth). In the case where A_Γ is not dihedral, we have that A_Γ is torsion-free by [21] and acylindrically hyperbolic by [44], so by Corollary 1.3 of [3], A_Γ has uniform exponential growth. The remaining case is where A_Γ is dihedral, i.e. Γ is a single edge (and we are assuming that the label is not 2). In this case, Lemma 2.3 implies that A_Γ fits into an exact sequence $1 \rightarrow \mathbb{Z} \rightarrow A_\Gamma \rightarrow F \rightarrow 1$, where F is virtually a finite-rank non-abelian free group. Hence A_Γ has uniform exponential growth by [26, Proposition 2.3]. This proves assertion (2).

Next, we deal with assertion (3) (stable subgroups). Here, we use two additional facts about the hierarchically hyperbolic structure \mathfrak{S} . The first one is that the hyperbolic space associated to the \sqsubseteq -maximal element is A_Γ -equivariantly quasi-isometric to the coned-off Deligne complex \widehat{D} , as stated in Theorem 6.15. Also, we will use the fact that for each non- \sqsubseteq -maximal $V \in \mathfrak{S}$ with unbounded $\mathcal{C}V$ there exists $U \in \mathfrak{S}$ with $U \perp V$ and unbounded $\mathcal{C}U$. This readily follows from Theorem 6.15.

According to [1, Theorem B], there is an HHS structure on A_Γ such that a subgroup $H \leq A_\Gamma$ is stable if and only if H -orbit maps to the \sqsubseteq -maximal hyperbolic space are quasi-isometric embeddings. We observe below (Lemma 6.18) that under our assumptions the \sqsubseteq -maximal hyperbolic space constructed in [1] is A_Γ -equivariantly quasi-isometric to \widehat{D} , which proves the assertion. \square

In [1], Theorem B follows combining [1, Corollary 6.2], which characterises stability for any HHS satisfying an additional condition (having *unbounded products*), and [1, Theorem 3.7], which says that one can change an HHG structure into one satisfying the additional condition.

What we used in the proof of Corollary B is that this change does not affect the equivariant quasi-isometry type of the \sqsubseteq -maximal hyperbolic space, as we now justify.

Lemma 6.18 *Let (G, \mathfrak{S}) be an HHG with the property that for each non- \sqsubseteq -maximal $V \in \mathfrak{S}$ with unbounded $\mathcal{C}V$ there exists $U \in \mathfrak{S}$ with $U \perp V$ and unbounded $\mathcal{C}U$. Then G admits an HHG structure (G, \mathfrak{T}) with unbounded products where the \sqsubseteq -maximal hyperbolic space is G -equivariantly quasi-isometric to that of (G, \mathfrak{S}) .*

Proof We claim that the HHG structure (G, \mathfrak{T}) constructed in the second paragraph of the proof of [1, Theorem 3.7] has \sqsubseteq -maximal hyperbolic space \mathcal{T}_S with the required property; we now recall the construction. First, for a constant M let \mathfrak{S}^M be the set of all $U \in \mathfrak{S}$ for which there exist $U \sqsubseteq V$ and $V \perp W$ with $\text{diam } \mathcal{C}V > M$ and $\text{diam } \mathcal{C}W > M$. We can choose M such that $\text{diam } \mathcal{C}V > M$ is equivalent to $\text{diam } \mathcal{C}V = \infty$ (because the G -action on \mathfrak{S} is cofinite by definition of HHG). Then \mathcal{T}_S is defined as the factored space $\hat{G}_{\mathfrak{S}^M}$. As a set, this is G , however it is endowed with a different metric. Roughly, we start with G and cone-off certain product regions corresponding to the elements of \mathfrak{S}^M , but we will not need the exact definition here.

By [7, Proposition 2.4], $\hat{G}_{\mathfrak{S}^M}$ is an HHS with index set $\mathfrak{S} - \mathfrak{S}^M$, and with the same associated hyperbolic spaces and projections as in (G, \mathfrak{S}) . By the distance formula for HHS [9, Theorem 4.5], for any suitably large constant L there are constants K, C so that for all $x, y \in \hat{G}_{\mathfrak{S}^M}$ we have.

$$d_{\hat{G}_{\mathfrak{S}^M}}(x, y) \approx_{K,C} \sum_{Y \in \mathfrak{S} - \mathfrak{S}^M} [d_{\mathcal{C}Y}(\pi_Y(x), \pi_Y(y))]_L,$$

where $\approx_{K,C}$ denotes quantities that differ by multiplicative constant at most K and additive constant at most C , and $[A]_L$ denotes A if $A \geq L$ and 0 otherwise.

Now it is time to use our hypothesis. In our case \mathfrak{S}^M contains all $U \in \mathfrak{S}$ with $\text{diam } \mathcal{C}U > M$ except the \sqsubseteq -maximal element S . If we take the threshold L to be larger than M , the distance formula states that π_S is a quasi-isometry from $\mathcal{T}_S = \hat{G}_{\mathfrak{S}^M}$ to $\mathcal{C}S$. Since π_S is G -equivariant by definition of HHG, we are done. \square

Acknowledgements We thank Daniel Berlyne for a helpful comment and Jason Behrstock for pointing out the reference [40]. We are grateful for the support of the International Centre for Mathematical Sciences, Edinburgh, which hosted the authors for a week in February 2020 during a Research in Groups where much of the work on this project was done. Hagen was partially supported by EPSRC New Investigator Award EP/R042187/1. Martin was partially supported by the EPSRC New Investigator Award EP/S010963/1.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Abbott, C., Behrstock, J., Durham, M.G.: Largest acylindrical actions and Stability in hierarchically hyperbolic groups. *Trans. Am. Math. Soc. Ser. B* **8**, 66–104 (2021)
2. Abbott, C., Balasubramanya, S.H., Osin, D.: Hyperbolic structures on groups. *Algebr. Geom. Topol.* **19**(4), 1747–1835 (2019)
3. Abbott, C., Ng, T., Spriano, D.: Hierarchically hyperbolic groups and uniform exponential growth. arXiv preprint [arXiv:1909.00439](https://arxiv.org/abs/1909.00439) (2019)
4. Alonso, J.M., Bridson, M.R.: Semihyperbolic groups. *Proc. Lond. Math. Soc.* **3**(1), 56–114 (1995)
5. Balasubramanya, S.: Acylindrical group actions on quasi-trees. *Algebr. Geom. Topol.* **17**(4), 2145–2176 (2017)
6. Behrstock, J., Hagen, M., Martin, A., Sisto, A.: A combinatorial take on hierarchical hyperbolicity and applications to quotients of mapping class groups. arXiv preprint [arXiv:2005.00567](https://arxiv.org/abs/2005.00567) (2020)
7. Behrstock, J., Hagen, M.F., Sisto, A.: Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups. *Proc. Lond. Math. Soc.* (3) **114**(5), 890–926 (2017)
8. Behrstock, J., Hagen, M.F., Sisto, A.: Hierarchically hyperbolic spaces, I: curve complexes for cubical groups. *Geom. Topol.* **21**(3), 1731–1804 (2017)
9. Behrstock, J., Hagen, M., Sisto, A.: Hierarchically hyperbolic spaces II: combination theorems and the distance formula. *Pac. J. Math.* **299**(2), 257–338 (2019)
10. Behrstock, J., Hagen, M.F., Sisto, A.: Quasiflats in hierarchically hyperbolic spaces. *Duke Math. J.* **170**(5), 909–996 (2021)
11. Berlai, F., Robbio, B.: A refined combination theorem for hierarchically hyperbolic groups. *Groups Geom. Dyn.* **14**(4), 1127–1203 (2020)
12. Bestvina, M.: Non-positively curved aspects of Artin groups of finite type. *Geom. Topol.* **3**(1), 269–302 (1999)
13. Bowditch, B.H.: Coarse median spaces and groups. *Pac. J. Math.* **261**(1), 53–93 (2013)
14. Bridson, M.R., Haefliger, A.: *Metric spaces of non-positive curvature*, Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319. Springer, Berlin (1999)
15. Brieskorn, E., Saito, K.: Artin-Gruppen und Coxeter-gruppen. *Invent. Math.* **17**(4), 245–271 (1972)
16. Calvez, M.: Euclidean Artin–Tits groups are acylindrically hyperbolic. arXiv preprint [arXiv:2010.13145](https://arxiv.org/abs/2010.13145) (2020)
17. Calvez, M., Wiest, B.: Acylindrical hyperbolicity and Artin–Tits groups of spherical type. *Geom. Dedicata* **191**(1), 199–215 (2017)
18. Calvez, M., Wiest, B.: Hyperbolic structures for Artin–Tits groups of spherical type. To appear in *Contemp. Math.* arXiv preprint [arXiv:1904.02234](https://arxiv.org/abs/1904.02234) (2019)
19. Caprace, P.-E., Sageev, M.: Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.* **21**(4), 851 (2011)
20. Charney, R., Crisp, J.: Relative hyperbolicity and Artin groups. *Geom. Dedicata* **129**, 1–13 (2007)
21. Charney, R., Davis, M.W.: The $k(\pi, 1)$ -problem for hyperplane complements associated to infinite reflection groups. *J. Am. Math. Soc.* **8**(3), 597–627 (1995)
22. Charney, R., Morris-Wright, R.: Artin groups of infinite type: trivial centers and acylindrical hyperbolicity. *Proc. Am. Math. Soc.* **147**(9), 3675–3689 (2019)
23. Crisp, J.: Automorphisms and abstract commensurators of 2-dimensional Artin groups. *Geom. Topol.* **9**(3), 1381–1441 (2005)
24. Cumplido, M., Gebhardt, V., González-Meneses, J., Wiest, B.: On parabolic subgroups of Artin–Tits groups of spherical type. *Adv. Math.* **352**, 572–610 (2019)
25. Dahmani, F., Guirardel, V., Osin, D.: Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Mem. Am. Math. Soc.* **245**(1156), v+152 (2017)
26. de la Harpe, P.: Uniform growth in groups of exponential growth. In: *Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part II (Haifa, 2000)*, vol. 95, pp. 1–17 (2002)
27. Durham, M., Hagen, M., Sisto, A.: Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.* **21**(6), 3659–3758 (2017)
28. Durham, M.G., Hagen, M.F., Sisto, A.: Correction to the article Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geom. Topol.* **24**(2), 1051–1073 (2020)
29. Durham, M.G., Minsky, Y.N., Sisto, A.: Stable cubulations, bicombings, and barycenters. To appear in *Geom. Topol.* arXiv preprint [arXiv:2009.13647](https://arxiv.org/abs/2009.13647) (2020)

30. Epstein, D.B.A., Fujiwara, K.: The second bounded cohomology of word-hyperbolic groups. *Topology* **36**(6), 1275–1289 (1997)
31. Haettel, T., Hoda, N., Petyt, H.: The coarse Helly property, hierarchical hyperbolicity, and semihyperbolicity. To appear in *Geom. Topol.* arXiv preprint [arXiv:2009.14053](https://arxiv.org/abs/2009.14053) (2020)
32. Hagen, M., Russell, J., Sisto, A., Spriano, D.: Equivariant hierarchically hyperbolic structures for graph manifold groups (**in preparation**) (2021)
33. Hagen, M.F., Susse, T.: On hierarchical hyperbolicity of cubical groups. *Isr. J. Math.* **236**(1), 45–89 (2020)
34. Hruska, G.C.: Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.* **10**(3), 1807–1856 (2010)
35. Huang, J., Osajda, D.: Quasi-euclidean tilings over 2-dimensional Artin groups and their applications. To appear in *Mem. Am. Math. Soc.* arXiv preprint [arXiv:1711.00122](https://arxiv.org/abs/1711.00122) (2017)
36. Huang, J., Osajda, D.: Metric systolicity and two-dimensional Artin groups. *Math. Ann.* **374**(3–4), 1311–1352 (2019)
37. Huang, J., Osajda, D.: Large-type Artin groups are systolic. *Proc. Lond. Math. Soc.* **120**(1), 95–123 (2020)
38. Hull, M., Osin, D.: Induced quasicocycles on groups with hyperbolically embedded subgroups. *Algebr. Geom. Topol.* **13**(5), 2635–2665 (2013)
39. Januszkiewicz, T., Świątkowski, J.: Simplicial nonpositive curvature. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques* **104**(1), 1–85 (2006)
40. Kapovich, I., Schupp, P.: Relative hyperbolicity and Artin groups. *Geom. Dedicata* **107**, 153–167 (2004)
41. Kim, S.-H., Koberda, T.: The geometry of the curve graph of a right-angled Artin group. *Int. J. Algebra Comput.* **24**(2), 121–169 (2014)
42. Mairesse, J., Mathéus, F.: Growth series for Artin groups of dihedral type. *Int. J. Algebra Comput.* **16**(6), 1087–1107 (2006)
43. Manning, J.F.: Geometry of pseudocharacters. *Geom. Topol.* **9**, 1147–1185 (2005)
44. Martin, A., Przytycki, P.: Acylindrical actions for two-dimensional Artin groups of hyperbolic type. *Int. Math. Res. Not.* (**in press**) (2021)
45. Masur, H.A., Minsky, Y.N.: Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.* **138**(1), 103–149 (1999)
46. Masur, H.A., Minsky, Y.N.: Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.* **10**(4), 902–974 (2000)
47. Mineyev, I.: Straightening and bounded cohomology of hyperbolic groups. *Geom. Funct. Anal.* **11**(4), 807–839 (2001)
48. Morris-Wright, R.: Parabolic subgroups in FC-type Artin groups. *J. Pure Appl. Algebra* **225**(1), 106468, 13 (2021)
49. Osajda, D., Przytycki, P.: Tits alternative for groups acting properly on 2-dimensional recurrent complexes (with an appendix written jointly with Jon McCammond). arXiv preprint [arXiv:1904.07796](https://arxiv.org/abs/1904.07796) (2019)
50. Osin, D.: Acylindrically hyperbolic groups. *Trans. Am. Math. Soc.* **368**(2), 851–888 (2016)
51. Paris, L.: Parabolic subgroups of Artin groups. *J. Algebra* **196**(2), 369–399 (1997)
52. Petyt, H., Spriano, D.: Unbounded domains in hierarchically hyperbolic groups. To appear in *Groups Geom. Dyn.* arXiv preprint [arXiv:2007.12535](https://arxiv.org/abs/2007.12535) (2020)
53. Świątkowski, J.: Regular path systems and (bi) automatic groups. *Geom. Dedicata* **118**(1), 23–48 (2006)
54. van der Lek, H.: The homotopy type of complex hyperplane complements (PhD thesis). University of Nijmegen (1983)
55. Vaskou, N.: Acylindrical hyperbolicity for Artin groups of dimension 2. *Geom. Dedicata* **216**, 7 (2022)