Quasi-isometric classification of right-angled Artin groups I: The finite out case

JINGYIN HUANG

Let G and G' be two right-angled Artin groups. We show they are quasi-isometric if and only if they are isomorphic, under the assumption that the outer automorphism groups $\mathrm{Out}(G)$ and $\mathrm{Out}(G')$ are finite. If we only assume $\mathrm{Out}(G)$ is finite, then G' is quasi-isometric to G if and only if G' is isomorphic to a subgroup of finite index in G. In this case, we give an algorithm to determine whether G and G' are quasi-isometric by looking at their defining graphs.

1 Introduction

1.1 Backgrounds and summary of results

Given a finite simplicial graph Γ with vertex set $\{v_i\}_{i \in I}$, the right-angled Artin group (RAAG) with defining graph Γ, denoted by $G(\Gamma)$, is given by the following presentation:

$$\{v_i \text{ for } i \in I \mid [v_i, v_j] = 1 \text{ if } v_i \text{ and } v_j \text{ are joined by an edge}\}.$$

We call $\{v_i\}_{i \in I}$ a *standard generating set* for $G(\Gamma)$; see Section 2.4.

The class of RAAGs enjoys a balance between simplicity and complexity. On one hand, RAAGs have many nice geometric, combinatorial and group theoretic properties (see Charney [16] for a summary); on the other hand, this class inherits the full complexity of the collection of finite simplicial graphs, and even a single RAAG could have very complicated subgroups (see, for example, Bestvina and Brady [8]).

In recent years, RAAGs have become important models to understand other unknown groups, either by (virtually) embedding the unknown groups into some RAAGs (such a program is outlined in Wise [61, Section 6]; see also Agol [2], Hagen and Wise [31; 32], Haglund and Wise [36], Ollivier and Wise [52], Przytycki and Wise [54; 55] and Wise [60; 62]), or by finding embedded copies of RAAGs in the unknown groups (see Baik, Kim and Koberda [3], Clay, Leininger and Mangahas [19], Kim and Koberda [44], Koberda [48] and Taylor [59]).

In this paper, we study the asymptotic geometry of RAAGs and classify a particular class of RAAGs by their quasi-isometric types. Previously, the quasi-isometric classification of RAAGs has been done for the following two classes:

(1) **Tree groups** It is shown by Behrstock and Neumann [7] that for any two trees Γ_1 and Γ_2 with diameter ≥ 3, the RAAGs $G(\Gamma_1)$ and $G(\Gamma_2)$ are quasi-isometric. Higher-dimensional analogs of tree groups are studied by Behrstock, Januszkiewicz and Neumann [5].

(2) **Atomic groups** A RAAG is atomic if its defining graph Γ is connected and does not contain valence-one vertices, cycles of length < 5 and separating closed stars. It is shown by Bestvina, Kleiner and Sageev [9] that two atomic RAAGs are quasi-isometric if and only if they are isomorphic.

Note that atomic groups are much more "rigid" than tree groups. We define the *dimension* of $G(\Gamma)$ to be the maximal n such that $G(\Gamma)$ contains a $\mathbb{Z}^n$ subgroup, and it coincides with the cohomological dimension of $G(\Gamma)$. All atomic groups are 2–dimensional; hence it is natural to ask what higher-dimensional RAAGs satisfy similar rigidity properties as atomic RAAGs. This is the starting point of the current paper.

Since we are looking for RAAGs which are rigid, those with small quasi-isometry groups would be reasonable candidates. However, even in the atomic case, the quasi-isometry group $\mathrm{QI}(G(\Gamma))$ is huge; see the discussion of quasi-isometry flexibility in [9, Section 11]. Then we turn to the outer automorphism group $\mathrm{Out}(G(\Gamma))$ for guidance.

Now we ask whether those RAAGs with "small" outer automorphism groups are also geometrically rigid in an appropriate sense. Actually, "small" outer automorphism groups and (quasi-isometric or commensurability) rigidity results come together in several other cases: for example, higher rank lattices (Eskin [25], Eskin and Farb [26], Kleiner and Leeb [45] and Mostow [51]), mapping class groups (Behrstock, Kleiner, Minsky and Mosher [6] and Hamenstädt [37]), $\mathrm{Out}(F_n)$ (Farb and Handel [27]), etc. Our first result is about the quasi-isometric classification for RAAGs with finite outer automorphism group.

Theorem 1.1 *Pick $G(\Gamma_1)$ and $G(\Gamma_2)$ such that $\mathrm{Out}(G(\Gamma_i))$ is finite for $i = 1, 2$. Then they are quasi-isomeric if and only if they are isomorphic.*

This theorem is proved in Section 4. See Theorem 4.13 for a more detailed version.

The collection of RAAGs with finite outer automorphism group is a reasonably large class. Recall that there is a one-to-one correspondence between finite simplicial graphs and RAAGs (see Droms [23]); thus it makes sense to talk about a random RAAG by

considering the Erdös–Rényi model for random graphs. If the parameters of the model are in the right range, then almost all RAAGs have finite outer automorphism group; see Charney and Farber [18] and Day [22].

The class of 2–dimensional RAAGs with finite outer automorphism group is strictly larger than the class of atomic RAAGs; moreover, there are plenty of higher-dimensional RAAGs with finite outer automorphism group.

Whether $\mathrm{Out}(G(\Gamma))$ is finite or not can be easily read from Γ. We defined the *closed star* of a vertex v in Γ, denoted by $\mathrm{St}(v)$, to be the full subgraph (see Section 2.1) spanned by v and vertices adjacent to v. Similarly, $\mathrm{lk}(v)$ is defined to be the full subgraph spanned by vertices adjacent to v. Note that this definition is slightly different from the usual one.

By results from Laurence [49] and Servatius [57], $\mathrm{Out}(G(\Gamma))$ is generated by the following four types of elements (we identify the vertex set of Γ with a standard generating set of $G(\Gamma)$):

(1) Given a vertex $v \in \Gamma$, the group automorphism defined by sending $v \to v^{-1}$ and fixing all other generators.

(2) Graph automorphisms of Γ.

(3) If $\mathrm{lk}(w) \subset \mathrm{St}(v)$ for vertices $w, v \in \Gamma$, sending $w \to wv$ and fixing all other generators induces a group automorphism, called a *transvection*. It is an *adjacent transvection* when $d(v, w) = 1$, and a *nonadjacent transvection* otherwise.

(4) Suppose $\Gamma \setminus \mathrm{St}(v)$ is disconnected. Then one obtains a group automorphism by picking a connected component C and sending $w \to vwv^{-1}$ for each vertex $w \in C$ (all other generators are fixed). It is called a *partial conjugation*.

Elements of type (3) or (4) have infinite order in $\mathrm{Out}(G(\Gamma))$ while elements of type (1) or (2) are of finite order. $\mathrm{Out}(G(\Gamma))$ is finite if and only if Γ does not contain any separating closed star and there do not exist distinct vertices $v, w \in \Gamma$ such that $\mathrm{lk}(w) \subset \mathrm{St}(v)$.

Theorem 1.2 *Suppose* $\mathrm{Out}(G(\Gamma_1))$ *is finite. Then the following are equivalent:*

(1) $G(\Gamma_2)$ *is quasi-isometric to* $G(\Gamma_1)$.

(2) $G(\Gamma_2)$ *is isomorphic to a subgroup of finite index in* $G(\Gamma_1)$.

(3) Γ_2^e *is isomorphic to* Γ_1^e.

Here Γ^e denotes the *extension graph* introduced by Kim and Koberda in [42]; see Definition 2.11. Extension graphs can be viewed as "curve graphs" for RAAGs; see Kim and Koberda [43]. This analog carries on to the aspect of quasi-isometric rigidity. Namely, if G is a mapping class group and $q: G' \to G$ is a quasi-isometry, then it is shown in [6] that G' naturally acts on the curve graph associated with G. This is still true if G is a RAAG with some restriction on its outer automorphism group, for example, if $\mathrm{Out}(G)$ is finite.

However, in general, there exists a pair of commensurable RAAGs with different extension graphs; see Example 3.22. There also exists a pair of RAAGs, not quasi-isometric, with isomorphic extension graphs; see Huang [38, Section 5.3].

Motivated by Theorem 1.2(2), we now look at finite-index *RAAG subgroups* (ie subgroups which are also RAAGs) of $G(\Gamma_1)$.

Given a RAAG $G(\Gamma)$ (not necessarily having a finite outer automorphism group) with a standard generating set S, let d_S be the word metric on $G(\Gamma)$ with respect to S. A subset $K \subset G(\Gamma)$ is S–*convex* if for any three points $x, y \in K$ and $z \in G(\Gamma)$ such that $d_S(x, y) = d_S(x, z) + d_S(z, y)$, we must have $z \in K$. Every finite S–convex subset K naturally gives rise to a finite-index RAAG subgroup $G \leq G(\Gamma)$ such that K is the fundamental domain of the left action $G \curvearrowright G(\Gamma)$. For example, if $G(\Gamma) = \mathbb{Z} \oplus \mathbb{Z}$ and K is a rectangle of size n by m, then the corresponding subgroup is of the form $n\mathbb{Z} \oplus m\mathbb{Z}$. The detailed construction in the more general case is given in Section 6.1. G is called an S–*special* subgroup of $G(\Gamma)$. A subgroup of $G(\Gamma)$ is *special* if it is S–special for some standard generating set S. A similar construction in the case of right-angled Coxeter groups is in Haglund [34].

Here is an alternative description in terms of the canonical completion introduced by Haglund and Wise [35]. Let $S(\Gamma)$ be the Salvetti complex of $G(\Gamma)$ (see Section 2.4) and let $X(\Gamma)$ be the universal cover. We pick an identification between $G(\Gamma)$ and the 0–skeleton of $X(\Gamma)$. The above subset K gives rise to a convex subcomplex $\bar{K} \subset X(\Gamma)$. Then the corresponding special subgroup is the fundamental group of the canonical completion with respect the local isometry $\bar{K} \to S(\Gamma)$.

Our next result says if $\mathrm{Out}(G(\Gamma))$ is finite, then this is the only way to obtain finite-index RAAG subgroups of $G(\Gamma)$.

Theorem 1.3 *Suppose* $\mathrm{Out}(G(\Gamma))$ *is finite, and let* S *be a standard generating set for* $G(\Gamma)$. *Then all finite-index RAAG subgroups are* S–*special. Moreover, there is a one-to-one correspondence between nonnegative finite* S–*convex subsets of* $G(\Gamma)$ *based at the identity and finite-index RAAG subgroups of* $G(\Gamma)$.

See Theorem 6.13 for a slight reformulation of Theorem 1.3.

We need to explain two terms: nonnegative and based at the identity. For example, take $G = n\mathbb{Z} \oplus m\mathbb{Z}$ inside $\mathbb{Z} \oplus \mathbb{Z}$; then any n by m rectangle could be the fundamental domain for the action of G. We naturally require the rectangle to be in the first quadrant and contain the identity, which would give us a unique choice. Similar things can be done in all RAAGs, and these two terms will be defined precisely in Section 6.

The most simple example is when $G(\Gamma) = \mathbb{Z}$; we have a one-to-one correspondence between finite-index subgroups of the form $n\mathbb{Z}$ and intervals of the form $[0, n-1]$.

Corollary 1.4 *If* $\mathrm{Out}(G(\Gamma_1))$ *is finite, then* $G(\Gamma_2)$ *is quasi-isometric to* $G(\Gamma_1)$ *if and only if* $G(\Gamma_2)$ *is isomorphic to a special subgroup of* $G(\Gamma_1)$.

It turns out that there is an algorithm to enumerate the defining graphs of all special subgroups of a RAAG:

Theorem 1.5 *If* $\mathrm{Out}(G(\Gamma))$ *is finite, then* $G(\Gamma')$ *is quasi-isometric to* $G(\Gamma)$ *if and only if* Γ' *can be obtained from* Γ *by finitely many GSEs. In particular, there is an algorithm to determine whether* $G(\Gamma')$ *and* $G(\Gamma)$ *are quasi-isometric by looking at the graphs* Γ *and* Γ'.

A GSE is a generalized version of a star extension in [9, Example 1.4]; see also [42, Lemma 50]. It will be defined in Section 6.

A question motivated by Theorem 1.2 is the following:

Question 1.6 *Let* $G(\Gamma)$ *be a RAAG such that* $\mathrm{Out}(G(\Gamma))$ *is finite, and let* H *be a finitely generated group quasi-isometric to* $G(\Gamma)$. *What can we say about* H *?*

As a partial answer to this question, we have the following:

Theorem 1.7 (Huang and Kleiner [40]) *Let* $G(\Gamma)$ *and* H *be as in Question 1.6. Then the induced quasi-action* $H \curvearrowright X(\Gamma)$ *is quasi-isometrically conjugate to a geometric action* $H \curvearrowright X'$. *Here* X' *is a* CAT(0) *cube complex which is closely related to* $X(\Gamma)$.

1.2 Comments on the proof

1.2.1 Proof of Theorem 1.1 We start by introducing notation. The Salvetti complex of $G(\Gamma)$ is denoted by $S(\Gamma)$, the universal cover of $S(\Gamma)$ is denoted by $X(\Gamma)$, and flats in $X(\Gamma)$ that cover standard tori in $S(\Gamma)$ are called *standard flats*. See Section 2.4 for precise definition of these terms.

Let $q\colon X(\Gamma) \to X(\Gamma')$ be a quasi-isometry. The proof of Theorem 1.1 follows the scheme of the proof of the main theorem in [9]. Similar schemes can also be found

in [6; 45]. There are three steps in [9]. First they show that q maps top-dimensional flats to top-dimensional flats up to finite Hausdorff distance. However, the collection of all top-dimensional flats is too large to be linked directly to the combinatorics of RAAGs, so the second step is to show that quasi-isometries preserve standard flats up to finite Hausdorff distance. The third step is to straighten the quasi-isometry such that it actually maps standard flats to standard flats exactly, not just up to finite Hausdorff distance, and the conclusion follows automatically.

In our case, the first step has been done in Huang [39], where we show q still preserves top-dimensional flats up to finite Hausdorff distance in the higher-dimensional case. No assumption on the outer automorphism group is needed for this step.

The second step consists of two parts. First we show q preserves certain top-dimensional maximal products up to finite Hausdorff distance. Then one wishes to pass to standard flats by intersecting these top-dimensional objects. However, in the higher-dimensional case, a lower-dimensional standard flat may not be the intersection of top-dimensional objects, and even when it is an intersection, one may not be able to read this information directly from the defining graph Γ. This is quite different from the 2–dimensional situation in [9] and relies on several new ingredients.

A necessary condition for q to preserve the standard flats is that every element in $\mathrm{Out}(G(\Gamma))$ does so, which implies there could not be any transvections in $\mathrm{Out}(G(\Gamma))$. This condition is also sufficient.

Theorem 1.8 *Suppose* $\mathrm{Out}(G(\Gamma))$ *is transvection-free. Then there exists a positive constant* $D = D(L, A, \Gamma)$ *such that for any standard flat* $F \subset X(\Gamma)$*, there exists a standard flat* $F' \subset X(\Gamma')$ *such that* $d_H(q(F), F') < D$*.*

Here d_H denotes the Hausdorff distance.

In Step 3, we introduce an auxiliary simplicial complex $\mathcal{P}(\Gamma)$, which serves as a link between the asymptotic geometry of $X(\Gamma)$ and the combinatorial structure of $X(\Gamma)$. More precisely, on one hand, $\mathcal{P}(\Gamma)$ can be viewed as a simplified Tits boundary for $X(\Gamma)$; on the other hand, one can read certain information about the wall space structure of $X(\Gamma)$ from $\mathcal{P}(\Gamma)$. This complex turns out to coincide with the extension graph introduced in [42], where it was motivated from the viewpoint of the mapping class group.

Denote the Tits boundary of $X(\Gamma)$ by $\partial_T(X(\Gamma))$, and let $T(\Gamma) \subset \partial_T(X(\Gamma))$ be the union of Tits boundaries of standard flats in $X(\Gamma)$. Then $T(\Gamma)$ has a natural simplicial structure. However, $T(\Gamma)$ contains redundant information; this can be seen in the similar situation where the link of the base point of $S(\Gamma)$ looks more complicated than Γ, but they essentially contain the same information.

This redundancy can be resolved by replacing the spheres in $T(X)$ that arise from standard flats by simplexes of the same dimension. This gives rise to a well-defined simplicial complex $\mathcal{P}(\Gamma)$ since for any standard flats F_1 and F_2 with $\partial_T F_1 \cap \partial_T F_2 \neq \varnothing$, there exists a standard flat F such that $\partial_T F = \partial_T F_1 \cap \partial_T F_2$. See Section 4.1 for more properties of $\mathcal{P}(\Gamma)$.

By Theorem 1.8, if both $\mathrm{Out}(G(\Gamma))$ and $\mathrm{Out}(G(\Gamma'))$ are transvection-free, then q induces a boundary map $\partial q\colon \mathcal{P}(\Gamma) \to \mathcal{P}(\Gamma')$, which is a simplicial isomorphism. Next we want to consider the converse and reconstruct a map $X(\Gamma) \to X(\Gamma')$ from the boundary map ∂q in the following sense. Pick a vertex $p \in X(\Gamma)$, and let $\{F_i\}_{i=1}^n$ be the collection of maximal standard flats containing p. By Theorem 1.8, for each i, there exists a unique maximal standard flat $F_i' \subset X(\Gamma')$ such that $d_H(q(F_i), F_i') < \infty$. One may wish to map p, which turns out to be the intersection of the F_i, to the intersection of all the F_i'. However, in general, $\bigcap_{i=1}^n F_i'$ may be empty, or it may contain more than one point; hence our map may not be well defined.

It turns out that if we also rule out partial conjugations in $\mathrm{Out}(G(\Gamma))$, then $\bigcap_{i=1}^n F_i'$ is exactly a point. This give rises to a well-defined map $\overline{q}\colon X(\Gamma)^{(0)} \to X(\Gamma')^{(0)}$ which maps vertices in a standard flat to vertices in a standard flat. If $\mathrm{Out}(G(\Gamma'))$ is also finite, then we can define an inverse map of $\overline{q}$, and this is enough to deduce Theorem 1.1.

1.2.2 Proof of Theorem 1.2 If only $\mathrm{Out}(G(\Gamma))$ is assumed to be finite, we can still recover the fact that ∂q is a simplicial isomorphism (this is nontrivial, since Theorem 1.8 does not say that for any standard flat $F' \subset X(\Gamma')$, we can find a standard flat $F \subset X(\Gamma)$ such that $d_H(q(F), F') < \infty$). Hence we can define $\overline{q}$ as before. However, the inverse of $\overline{q}$ does not exist in general.

The next step is to trying to extend $\overline{q}$ to a cubical map (Definition 2.1) from $X(\Gamma)$ to $X(\Gamma')$. There are obvious obstructions: though $\overline{q}$ maps vertices in a standard geodesic to vertices in a standard geodesic, $\overline{q}$ may not preserve the order of these vertices. A typical example is given in Figure 1, where one can permute the green level and the red level in a tree; then the order of vertices in the black line is not preserved.

A remedy is to "flip backwards". Namely we will precompose $\overline{q}$ with a sequence of permutations of "levels" such that the resulting map restricted to each standard geodesic respects the order. Then we can extend $\overline{q}$ to a cubical map. This argument relies on the understanding of quasi-isometric flexibility, namely how much room we have to perform these flips. One formulation of this aspect is the following.

Theorem 1.9 *If $\mathrm{Out}(G(\Gamma))$ is finite, then $\mathrm{Aut}(\mathcal{P}(\Gamma)) \cong \mathrm{Isom}(G(\Gamma), d_r)$.*

Here d_r denotes the syllable metric, defined in Section 4.3; see also [43, Section 5.2].

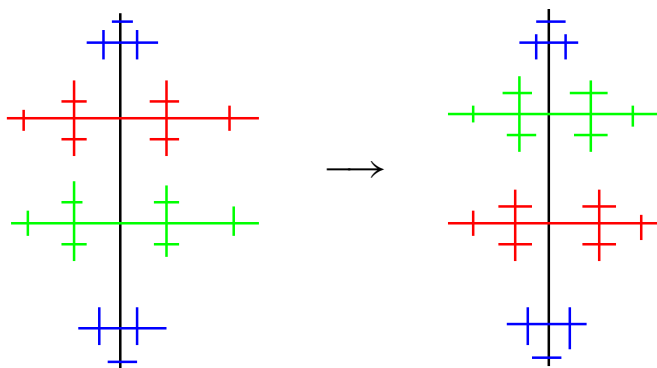Figure 1: Flipping of levels

Theorem 1.2, Theorem 1.3, Corollary 1.4 and Theorem 1.5 rely on the cubical map $\overline{q}$. In particular, $\overline{q}^{-1}(x)$ ($x \in X(\Gamma')$ is a vertex) is a compact convex subcomplex, and this is how we obtain the S–convex subset in Theorem 1.3.

1.3 Organization of the paper

Section 2 contains basic notation used in this paper and some background material about CAT(0) cube complexes and RAAGs. In particular, Section 2.3 collects several technical lemmas about CAT(0) cube complexes. One can skip Section 2.3 on first reading and come back when needed.

In Section 3, we prove Theorem 1.8. Section 3.1 is about the stability of top-dimensional maximal product subcomplexes under quasi-isometries, and Section 3.2 deals with lower-dimensional standard flats. In Section 4, we prove Theorem 1.1. We will construct the extension complex from our viewpoint in Section 4.1 and explain how is this object is related to Tits boundary, flat space and contact graph. In Section 4.2, we describe how to reconstruct the quasi-isometry.

Sections 5.1 and 5.2 are devoted to proving Theorem 1.2. We prove Theorem 1.3, Corollary 1.4 and Theorem 1.5 in Section 6.

Acknowledgements I would like to thank Bruce Kleiner and Robert Young for reading parts of the paper and giving valuable comments. I thank Jason Behrstock, Ruth Charney and Walter Neumann for suggesting interesting questions, which led to Section 6.3 of this paper. I thank Pallavi Dani, Mark Hagen, Thomas Koberda, Do-gyeom Kim and Sang-hyun Kim for related discussions. I also thank the referee for carefully reading the paper and providing helpful comments.

2 Preliminaries

2.1 Notation and conventions

All graphs in this paper are simple.

The flag complex of a graph Γ is denoted by $F(\Gamma)$; ie $F(\Gamma)$ is a flag complex whose 1–skeleton is Γ.

A subcomplex K' in a combinatorial polyhedral complex K is *full* if K' contains all the subcomplexes of K which have the same vertex set as K'. If K is 1–dimensional, then we also call K' a *full subgraph*.

We use "$*$" to denote the join of two simplicial complexes and "$\circ$" to denote the join of two graphs. Let K be a simplicial complex or a graph. By viewing the 1–skeleton as a metric graph with edge length $=1$, we obtain a metric defined on the 0–skeleton, which we denote by d. Let $N \subset K$ be a subcomplex. We define the *orthogonal complement* of N, denoted by $N^\perp$, to be the set $\{w \in K^{(0)} \mid d(w, v) = 1$ for any vertex $v \in N\}$; define the *link* of N, denoted by $\mathrm{lk}(N)$, to be the full subcomplex spanned by $N^\perp$; and define the *closed star* of N, denoted by $\mathrm{St}(N)$, to be the full subcomplex spanned by $N \cup \mathrm{lk}(N)$. Suppose L is a subcomplex such that $N \subset L \subset K$. We denote the closed star of N in L by $\mathrm{St}(N, L)$. If L is a full subcomplex, then $\mathrm{St}(N, L) = \mathrm{St}(N) \cap L$. We can define $\mathrm{lk}(N, L)$ in a similar way. Let $M \subset K$ be an arbitrary subset. We denote the collection of vertices inside M by $v(M)$.

We use id to denote the identity element of a group, and we use Id to denote the identity map from a space to itself.

Let (X, d) be a metric space. The open ball of radius r centered at p in X will be denoted by $B(p, r)$. Given subsets $A, B \subset X$, the open r–neighborhood of a subset A is denoted by $N_r(A)$. The diameter of A is denoted by $\mathrm{diam}(A)$. The Hausdorff distance between A and B is denoted by $d_H(A, B)$. We will also use the adapted notation of coarse set theory introduced in [50], displayed in the following table:

symbol	meaning
$A \subset_r B$	$A \subset N_r(B)$
$A \subset_\infty B$	$\exists r > 0$ such that $A \subset N_r(B)$
$A \overset{r}{=} B$	$d_H(A, B) \leq r$
$A \overset{\infty}{=} B$	$d_H(A, B) < \infty$
$A \cap_r B$	$N_r(A) \cap N_r(B)$

2.2 CAT(0) space and CAT(0) cube complex

The standard reference for CAT(0) spaces is [13].

Let (X, d) be a CAT(0) space. Pick $x, y \in X$, we denote the unique geodesic segment joining x and y by $\overline{xy}$. For $y, z \in X \setminus \{x\}$, denote the comparison angle between $\overline{xy}$ and $\overline{xz}$ at x by $\overline{\angle}_x(y, z)$ and the Alexandrov angle by $\angle_x(y, z)$.

The boundary of X, denoted by ∂X, is the collection of asymptotic classes of geodesic rays. The boundary ∂X has an angular metric, which is defined by

$$\angle(\xi_1, \xi_2) = \lim_{t, t' \to \infty} \overline{\angle}_p(l_1(t), l_2(t')),$$

where l_1 and l_2 are unit speed geodesic rays emanating from a base point p such that $l_i(\infty) = \xi_i$ for $i = 1, 2$. This metric does not depend on the choice of p, and the length metric associated to the angular metric, denoted by d_T, is called the Tits metric. The Tits boundary, denoted by $\partial_T X$, is the CAT(1) space $(\partial X, d_T)$; see [13, Chapters II.8 and II.9].

Given two metric spaces (X_1, d_1) and (X_2, d_2), denote the *Cartesian product* of X_1 and X_2 by $X_1 \times X_2$, ie $d = \sqrt{d_1^2 + d_2^2}$ on $X_1 \times X_2$. If X_1 and X_2 are CAT(0), then so is $X_1 \times X_2$.

An n–*flat* in a CAT(0) space X is the image of an isometric embedding $\mathbb{E}^n \to X$. Note that any flat is convex in X.

Pick a convex subset $C \subset X$; then C is also CAT(0). We use π_C to denote the nearest point projection from X to C; it is well defined and 1–Lipschitz. Moreover, pick $x \in X \setminus C$; then $\angle_{\pi_C(x)}(x, y) \geq \frac{\pi}{2}$ for any $y \in C$ such that $y \neq \pi_C(x)$; see [13, Proposition II.2.4].

If $C' \subset X$ is another convex set, then C' is *parallel* to C if $d(\cdot, C)|_{C'}$ and $d(\cdot, C')|_C$ are constant functions. There is a natural isomorphism between $C \times [0, d(C, C')]$ and the convex hull of C and C' in this case. We define the *parallel set* of C, denoted by P_C, to be the union of all convex subsets of X parallel to C. If C has the geodesic extension property, or more generally, C is boundary-minimal (see [14, Section 3.C]), then P_C is a convex subset in X. Moreover, P_C admits a canonical splitting $P_C = C \times C^\perp$, where $C^\perp$ is also a CAT(0) space.

Now we turn to CAT(0) cube complexes. All cube complexes in this paper are assumed to be finite dimensional.

A cube complex X is obtained by gluing a collection of unit Euclidean cubes isometrically along their faces, see [13, Definition II.7.32] for a precise definition. Then the

cube complex has a natural piecewise Euclidean metric. This metric is complete and geodesic if X is finite dimensional [13, I.7.19] and is nonpositively curved if the link of each vertex is a flag complex [29]. If in addition X is simply connected, then this metric is CAT(0) and X is said to be a CAT(0) *cube complex*. We can put a different metric on the 1–skeleton $X^{(1)}$ by considering it as a metric graph with all edge lengths 1. This is called the ℓ^1 *metric*. We use d for the CAT(0) metric on X and d_{ℓ^1} for the ℓ^1 metric on $X^{(1)}$. The natural injection $(X^{(1)}, d_{\ell^1}) \hookrightarrow (X, d)$ is a quasi-isometry; see [13, I.7.31] or [15, Lemma 2.2]. In this paper, we will mainly use the CAT(0) metric unless otherwise specified. Also any notions which depend on the metric, like geodesic, convex subset, convex hull etc, will be understood automatically with respect to the CAT(0) metric unless otherwise specified.

Definition 2.1 [15, Section 2.1] A cellular map between CAT(0) cube complexes is *cubical* if its restriction $\sigma \to \tau$ between cubes factors as $\sigma \to \eta \to \tau$, where the first map $\sigma \to \eta$ is a natural projection onto a face of σ and the second map $\eta \to \tau$ is an isometry.

A *geodesic segment, geodesic ray* or *geodesic line* in X is an isometric embedding of $[a, b]$, $[0, \infty)$ or $\mathbb{R}$ into X with respect to the CAT(0) metric. A *combinatorial geodesic segment, combinatorial geodesic ray* or *combinatorial geodesic* is an ℓ^1–isometric embedding of $[a, b]$, $[0, \infty)$ or $\mathbb{R}$ into $X^{(1)}$ such that its image is a subcomplex.

Let X be a CAT(0) cube complex and let $Y \subset X$ be a subcomplex. Then the following are equivalent (see [34]):

(1) Y is convex with respect to the CAT(0) metric.

(2) Y is a full subcomplex and $Y^{(1)} \subset X^{(1)}$ is convex with respect to the ℓ^1 metric.

(3) $\mathrm{Lk}(p, Y)$ (the link of p in Y) is a full subcomplex of $\mathrm{Lk}(P, X)$ for every vertex $p \in Y$.

The collection of convex subcomplexes in a CAT(0) cube complex enjoys the following version of Helly's property [28]:

Lemma 2.2 Let X be as above, and let $\{C_i\}_{i=1}^k$ be a collection of convex subcomplexes. If $C_i \cap C_j \neq \varnothing$ for any $1 \leq i \neq j \leq k$, then $\bigcap_{i=1}^k C_i \neq \varnothing$.

Lemma 2.3 Let X_1 and X_2 be two CAT(0) cube complexes, and let $K \subset X_1 \times X_2$ be a convex subcomplex. Then K admits a splitting $K = K_1 \times K_2$, where K_i is a convex subcomplex of X_i for $i = 1, 2$.

The lemma is clear when $X_1 \cong [0, 1]$, and the general case follows from this special case.

Now we come to the notion of hyperplane, which is the cubical analog of "track" introduced in [24]. A *hyperplane* h in a cube complex X is a subset such that:

(1) h is connected.

(2) For each cube $C \subset X$, either $h \cap C$ is empty or it is a union of mid-cubes of C.

(3) h is minimal; ie if there exists $h' \subset h$ satisfying (1) and (2), then $h = h'$.

Recall that a *mid-cube* of $C = [0, 1]^n$ is a subset of the form $f_i^{-1}\left(\frac{1}{2}\right)$, where f_i is one of the coordinate functions.

If X is a CAT(0) cube complex, then the following are true (see [56]):

(1) Each hyperplane is embedded; ie $h \cap C$ is either empty or a mid-cube of C (in more general cube complexes, it is possible that $h \cap C$ contains two or more mid-cubes of C).

(2) h is a convex subset in X, and h with the induced cell structure from X is also a CAT(0) cube complex.

(3) $X \setminus h$ has exactly two connected components; they are called *halfspaces*. The closure of a halfspace is called *closed halfspace*, which is also convex in X with respect to the CAT(0) metric.

(4) Let N_h be the smallest subcomplex of X that contains h. Then N_h is a convex subcomplex of X, and there is a natural isometry $i\colon N_h \to h \times [0, 1]$ such that $i(h) = h \times \{\frac{1}{2}\}$. N_h is called the *carrier* of h.

(5) For every edge $e \subset X$, there exists a unique hyperplane h_e which intersects e in its midpoint. In this case, we say h_e is the hyperplane dual to e and e is an edge dual to the hyperplane h_e.

(6) Lemma 2.2 is also true for a collection of hyperplanes.

Now it is easy to see an edge path $\omega \subset X$ is a combinatorial geodesic segment if and only if there do not exist two different edges of ω such that they are dual to the same hyperplane. Moreover, for two vertices $v, w \in X$, their ℓ^1 distance is exactly the number of hyperplanes that separate v from w.

Pick an edge $e \subset X$, and let $\pi_e \colon X \to e \cong [0, 1]$ be the CAT(0) projection. Then:

(1) The hyperplane dual to e is exactly $\pi_e^{-1}\left(\frac{1}{2}\right)$.

(2) $\pi_e^{-1}(t)$ is convex in X for any $0 \le t \le 1$; moreover, if $0 < t < t' < 1$, then $\pi_e^{-1}(t)$ and $\pi_e^{-1}(t')$ are parallel.

(3) Let N_{h_e} be the carrier of the hyperplane dual to e. Then N_{h_e} is the closure of $\pi_e^{-1}(0, 1)$. Alternatively, we can describe N_{h_e} as the parallel set of e.

2.3 Coarse intersections of convex subcomplexes

Lemma 2.4 [39, Lemma 2.10] *Let X be a CAT(0) cube complex of dimension n, and let C_1, C_2 be convex subcomplexes. Put $\Delta = d(C_1, C_2)$. Let $Y_1 = \{y \in C_1 \mid d(y, C_2) = \Delta\}$ and $Y_2 = \{y \in C_2 \mid d(y, C_1) = \Delta\}$. Then:*

(1) *Y_1 and Y_2 are not empty.*

(2) *Y_1 and Y_2 are convex; π_{C_1} maps Y_2 isometrically onto Y_1 and π_{C_2} maps Y_1 isometrically onto Y_2; the CAT(0) convex hull of $Y_1 \cup Y_2$ is isometric to $Y_1 \times [0, \Delta]$ (since we are taking the CAT(0) convex hull, it does not has to be a subcomplex).*

(3) *Y_1 and Y_2 are subcomplexes, and $\pi_{C_2}|_{Y_1}$ is a cubical isomorphism from Y_1 to Y_2 with its inverse given by $\pi_{C_1}|_{Y_2}$.*

(4) *For any $\epsilon > 0$, there exists $A = A(\Delta, n, \epsilon)$ such that if $d(p_1, Y_1) \geq \epsilon > 0$ and $d(p_2, Y_2) \geq \epsilon > 0$ for $p_1 \in C_1$ and $p_2 \in C_2$, then*

$$(2\text{-}1) \qquad d(p_1, C_2) \geq \Delta + A d(p_1, Y_1) \quad \text{and} \quad d(p_2, C_1) \geq \Delta + A d(p_2, Y_2).$$

Remark 2.5 Equation (2-1) implies for any $r > 0$, we have $(C_1 \cap_r C_2) \subset_{r'} Y_i$ $(i = 1, 2)$, where $r' = \min(1, (2r - \Delta)/A) + r$ and $A = A(\Delta, n, 1)$. Moreover, $\partial_T C_1 \cap \partial_T C_2 = \partial_T Y_1 = \partial_T Y_2$.

The remark implies $Y_1 \overset{\infty}{=} Y_2 \overset{\infty}{=} (C_1 \cap_r C_2)$ for r large enough. We use $\mathcal{I}(C_1, C_2) = (Y_1, Y_2)$ to describe this situation, where $\mathcal{I}$ stands for the word "intersect". The next lemma gives a combinatorial description of Y_1 and Y_2.

Lemma 2.6 *Let X, C_1, C_2, Y_1 and Y_2 be as above. Pick an edge e in Y_1 (or Y_2), and let h be the hyperplane dual to e. Then $h \cap C_i \neq \varnothing$ for $i = 1, 2$. Conversely, if a hyperplane h' satisfies $h' \cap C_i \neq \varnothing$ for $i = 1, 2$, then h' is the dual hyperplane of some edge e' in Y_1 (or Y_2). Moreover, $\mathcal{I}(h' \cap C_1, h' \cap C_2) = (h' \cap Y_1, h' \cap Y_2)$.*

Proof The first part follows from Lemma 2.4. Let $\mathcal{I}(h' \cap C_1, h' \cap C_2) = (Y_1', Y_2')$. Pick $x \in Y_1'$ and let $x' = \pi_{h' \cap C_2}(x) \in Y_2'$. Then $\pi_{h' \cap C_1}(x') = x$. Let $N_{h'} = h' \times [0, 1]$ be the carrier of h'. Then $(h' \cap C_i) \times (\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon) = C_i \cap (h' \times (\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon))$ for $i = 1, 2$ and $\epsilon < \frac{1}{2}$. Thus for any $y \in C_2$, we have $\angle_{x'}(x, y) \geq \frac{\pi}{2}$, which implies $x' = \pi_{C_2}(x)$. Similarly, $x = \pi_{C_1}(x') = \pi_{C_1} \circ \pi_{C_2}(x)$; hence $x \in Y_1$ and $Y_1' \subset Y_1$. By the same argument, $Y_2' \subset Y_2$; thus $Y_i' = Y_i \cap h'$ for $i = 1, 2$, and the lemma follows. $\square$

Lemma 2.4, Remark 2.5 and Lemma 2.6 can also be applied to CAT(0) rectangle complexes of finite type, whose cells are of the form $\prod_{i=1}^n [0, a_i]$. "Finite type" means there are only finitely many isometry types of rectangle cells in the rectangle complex.

Lemma 2.7   Let X, C_1, C_2, Y_1 and Y_2 be as above. If h is a hyperplane separating C_1 from C_2, then there exists a convex set $Y \subset h$ such that Y is parallel to Y_1 (or Y_2).

Proof   Let $\Delta = d(C_1, C_2)$, and let $M = Y_1 \times [0, \Delta]$ be the convex hull of Y_1 and Y_2. We want to prove $M \cap h = Y_1 \times \{t\} \subset Y_1 \times [0, \Delta]$ for some $t \in [0, \Delta]$. It suffices to show for any edge $e \subset Y_1$, we have $(e \times [0, \Delta]) \cap h = e \times \{t\}$ for some t.

Pick a point $x \in e$, and consider the point $\{x\} \times \{t\}$ in $M = Y_1 \times [0, \Delta]$. Since $e \times \{t\}$ and e are parallel, $e \times \{t\}$ sits inside a cube and is parallel to an edge of this cube. Thus either $e \times \{t\} \subset h$ or $e \times \{t\}$ is parallel to some edge dual to h. But the second case implies that h is dual to e and $h \cap Y_1 \neq \varnothing$, which is impossible, so $e \times \{t\} \subset (e \times [0, \Delta]) \cap h$. Now we are done since $(\{x\} \times [0, \Delta]) \cap h$ is exactly one point for each $x \in e$. $\square$

2.4   Right-angled Artin groups

Pick a finite simplicial graph Γ. Let $G(\Gamma)$ be a RAAG. A generating set $S \subset G(\Gamma)$ is called a *standard generating set* if all relators in the associated group presentation are commutators. Each standard generating set S determines a graph Γ_S whose vertices are elements in S, and two vertices are adjacent if the corresponding group elements commute. It follows from [23] that the isomorphism type of Γ_S does not depend on the choice of the standard generating set S; in particular, Γ_S is isomorphic to Γ.

Let S be a standard generating set for $G(\Gamma)$. We label the vertices of Γ by elements in S. The RAAG $G(\Gamma)$ has a nice Eilenberg–Mac Lane space $S(\Gamma)$, called the Salvetti complex; see [17; 16]. This is a nonpositively curved cube complex. The 2–skeleton of $S(\Gamma)$ is the usual presentation complex of $G(\Gamma)$. If the presentation complex contains a copy of 2–skeleton of a 3–torus, then we attach a 3–cell to obtain a 3–torus. We can build $S(\Gamma)$ inductively in this manner, and this process will stop after finitely many steps. The closure of each k–cell in $S(\Gamma)$ is a k–torus. A torus of this kind is called a *standard torus*. There is a one-to-one correspondence between the k–cells (or standard tori of dimension k) in $S(\Gamma)$ and k–cliques (complete subgraphs of k vertices) in Γ; thus $\dim(S(\Gamma)) = \dim(F(\Gamma)) + 1$. We define the *dimension* of $G(\Gamma)$ to be the dimension of $S(\Gamma)$.

Denote the universal cover of $S(\Gamma)$ by $X(\Gamma)$, which is a CAT(0) cube complex. Our previous labeling of vertices of Γ induces a labeling of the standard circles of $S(\Gamma)$, which lifts to a labeling of edges of $X(\Gamma)$. We choose an orientation for each standard circle of $S(\Gamma)$, and this gives us a directed labeling of the edges in $X(\Gamma)$. If we pick a base point $v \in X(\Gamma)$ (v is a vertex), then there is a one-to-one correspondence between words in $G(\Gamma)$ and edge paths in $X(\Gamma)$ which start at v.

Each full subgraph $\Gamma' \subset \Gamma$ gives rise to a subgroup $G(\Gamma') \hookrightarrow G(\Gamma)$. A subgroup of this kind is called a S–*standard subgroup* and a left coset of an S–standard subgroup is called an S–*standard coset* (we will omit S when the generating set is clear). There is also an embedding $S(\Gamma') \hookrightarrow S(\Gamma)$ which is locally isometric. Let $p\colon X(\Gamma) \to S(\Gamma)$ be the covering map. Then each connected component of $p^{-1}(S(\Gamma'))$ is a convex subcomplex isometric to $X(\Gamma')$. We will call these components *standard subcomplexes with defining graph* Γ'. A *standard k–flat* is a standard complex which covers a standard k–torus in $S(\Gamma)$. When $k = 1$, we also call it a *standard geodesic*.

We pick an identification of the Cayley graph of $G(\Gamma)$ with the 1–skeleton of $X(\Gamma)$; hence $G(\Gamma)$ is identified with the vertices of $X(\Gamma)$. Let $v \in X(\Gamma)$ be the base vertex which corresponds to the identity in the Cayley graph of $G(\Gamma)$. Then for any $h \in G(\Gamma)$, the convex hull of $\{hgv\}_{g \in G(\Gamma')}$ is a standard subcomplex associated with Γ'. Thus there is a one-to-one correspondence between standard subcomplexes with defining graph Γ' in $X(\Gamma)$ and left cosets of $G(\Gamma')$ in $G(\Gamma)$.

Note that for every edge $e \in X(\Gamma)$, there is a vertex in Γ which shares the same label as e, and we denote this vertex by V_e. If $K \subset X(\Gamma)$ is a subcomplex, we define V_K to be $\{V_e \mid e$ is an edge in $K\}$ and Γ_K to be the full subgraph spanned by V_K. This subgraph is called the *support* of K. In particular, if K is a standard subcomplex, then the *defining graph* of K is Γ_K.

Every finite simplicial graph Γ admits a canonical join decomposition

$$\Gamma = \Gamma_1 \circ \Gamma_2 \circ \cdots \circ \Gamma_k,$$

where Γ_1 is the maximal clique join factor and Γ_i does not allow any nontrivial join decomposition and is not a point, for $2 \le i \le k$. The graph Γ is *irreducible* if this join decomposition is trivial. This decomposition induces a product decomposition $X(\Gamma) = \mathbb{E}^n \times \prod_{i=2}^{k} X(\Gamma_i)$, which is called the *De Rahm decomposition* of $X(\Gamma)$. This is consistent with the canonical product decomposition of CAT(0) cube complex discussed in [15, Section 2.5].

We turn to the asymptotic geometry of RAAGs. A right-angled Artin group $G(\Gamma)$ is one-ended if and only if Γ is connected. Moreover, the n–connectivity at infinity of $G(\Gamma)$ can be read off from Γ; see [11]. In order to classify all RAAGs up to quasi-isometry, it suffices to consider those one-ended RAAGs. This follows from the main results in [53]. Moreover, we deduce the following lemma from [53, Lemma 3.2].

Lemma 2.8 *If* $q\colon X(\Gamma) \to X(\Gamma')$ *is an* (L, A)–*quasi-isometry, then there exists* $D = D(L, A) > 0$ *such that for any connected component* $\Gamma_1 \subset \Gamma$ *where* Γ_1 *is not a point and any standard subcomplex* $K_1 \subset X(\Gamma)$ *with defining graph* Γ_1, *there is a unique connected component* $\Gamma_1' \subset \Gamma'$ *and a unique standard subcomplex* $K_1' \subset X(\Gamma')$ *with defining graph* Γ_1' *such that* $d_H(q(K_1), K_1') < D$.

It is shown in [4] and [1] that $G(\Gamma)$ has linear divergence if and only if Γ is either a join or a point, which implies Γ being a join is a quasi-isometric invariant. Moreover, their results together with Theorem B of [41] implies that the De Rahm decomposition is stable under quasi-isometry:

Theorem 2.9 *Given $X = X(\Gamma)$ and $X' = X(\Gamma')$, let $X = \mathbb{R}^n \times \prod_{i=1}^{k} X(\Gamma_i)$ and $X' = \mathbb{R}^{n'} \times \prod_{j=1}^{k'} X(\Gamma'_j)$ be the corresponding De Rahm decompositions. If $\phi \colon X \to X'$ is an (L, A)–quasi-isometry, then $n = n'$, $k = k'$ and there exist constants $L' = L'(L, A)$, $A' = A'(L, A)$ and $D = D(L, A)$ such that after reindexing the factors in X', we have (L', A')–quasi-isometry $\phi_i \colon X(\Gamma_i) \to X(\Gamma'_i)$ with $d(p' \circ \phi, \prod_{i=1}^{k} \phi_i \circ p) < D$, where $p \colon X \to \prod_{i=1}^{k} X(\Gamma_i)$ and $p' \colon X' \to \prod_{i=1}^{k} X(\Gamma'_i)$ are the projections.*

Thus in order to study the quasi-isometric classification of RAAGs, it suffices to study those RAAGs which are one-ended and irreducible, but this will rely on finer quasi-isometric invariants of RAAGs.

Recall that in the case of Gromov hyperbolic spaces, quasi-isometries map geodesics to geodesics up to finite Hausdorff distance, hence induce a well-defined boundary map. The analog of this fact for 2–dimensional RAAGs has been established in [10], ie quasi-isometries map 2–flats to 2–flats up to finite Hausdorff distance. The following is a higher-dimensional generalization of [10, Theorem 3.10].

Theorem 2.10 [39, Theorem 5.20] *If $\phi \colon X(\Gamma_1) \to X(\Gamma_2)$ is an (L, A)–quasi-isometry, then $\dim(X(\Gamma_1)) = \dim(X(\Gamma_2))$, and there is a constant $D = D(L, A)$ such that for any top-dimensional flat $F_1 \subset X(\Gamma_1)$, there is a unique flat $F_2 \subset X(\Gamma_2)$ with $d_H(\phi(F_1), F_2) < D$.*

For each right-angled Artin group $G(\Gamma)$, there is a simplicial graph Γ^e, called the *extension graph*, which is introduced in [42]. Extension graphs can be viewed as "curve graphs" for RAAGs [43].

Definition 2.11 [42, Definition 1] *The vertex set of Γ^e consists of words in $G(\Gamma)$ that are conjugate to elements in S (recall S is a standard generating set for $G(\Gamma)$), and two vertices are adjacent in Γ^e if and only if the corresponding words commute in $G(\Gamma)$.*

The flag complex of the extension graph is called the *extension complex*.

Since the curve graph captures the combinatorial pattern of how Dehn twist flats intersect in a mapping class group, it plays an important role in the quasi-isometric rigidity of a mapping class group [37; 6]. Similarly, we will see in Section 4 that the extension graph captures the combinatorial pattern of the coarse intersection of a certain collection of flats in a RAAG, and it is a quasi-isometric invariant for certain classes of RAAGs.

3 Stable subgraph

We now study the behavior of certain standard subcomplexes under quasi-isometries.

3.1 Coarse intersection of standard subcomplexes and flats

Lemma 3.1 *Let Γ be a finite simplicial graph and let K_1, K_2 be two standard subcomplexes of $X(\Gamma)$. If $(Y_1, Y_2) = \mathcal{I}(K_1, K_2)$, then Y_1 and Y_2 are also standard subcomplexes.*

Proof The lemma is clear if $K_1 \cap K_2 \neq \varnothing$. Now we assume $d(K_1, K_2) = c > 0$. Pick a vertex $v_1 \in K_1$. By Lemma 2.4, there exists a vertex $v_2 \in K_2$ such that $d(v_1, v_2) = c$. Let $l \colon [0, c] \to X(\Gamma)$ be the unit speed geodesic from v_1 to v_2. We can find a sequence of cubes $\{B_i\}_{i=1}^N$ and $0 = t_0 < t_1 < \cdots < t_{N-1} < t_N = c$ such that each B_i contains $\{l(t) \mid t_{i-1} < t < t_i\}$ as interior points.

Let $V_l = \bigcup_{i=1}^N V_{B_i}$ (recall that V_{B_i} is the collection of the labels of edges in B_i; see Section 2.4) and let $V_i = V_{K_i}$ for $i = 1, 2$. Put $V' = V_1 \cap V_2 \cap V_l^\perp$ (recall that $V_l^\perp$ denotes the orthogonal complement of V_l; see Section 2.1) and let Γ' be the full subgraph spanned by V'. Let Y_1' be the standard subcomplex that has defining graph Γ' and contains v_1 (if V' is empty, then $Y_1' = v_1$). We claim $Y_1 = Y_1'$.

Pick an edge $e \subset K_1$ such that $v_1 \in e$ and $V_e \in V'$. Let h be the hyperplane dual to e and $N_h \cong h \times [0, 1]$ be the carrier of h. Since $d(V_e, w) = 1$ for any $w \in V_l$, we can assume $l \subset h \times \{1\} \subset N_h$. By our definition of V', there is an edge $e' \in K_2$ with $v_2 \in e'$ and h dual to e'; thus e and e' cobound an isometrically embedded flat rectangle (one side of the rectangle is l), which implies $e \subset Y_1$. Let l' be the side of the rectangle opposite to l. We can define $V_{l'}$ similarly as we define V_l; then $V_{l'} = V_l$. Now let ω be any edge path starting at v_1 such that $V_{e'} \in V'$ for any edge $e' \subset \omega$. Then it follows from the above argument and induction on the combinatorial length of ω that $\omega \subset Y_1$, thus $Y_1' \subset Y_1$.

For the other direction, since Y_1 is a convex subcomplex by Lemma 2.4, it suffices to prove every vertex of Y_1 belongs to Y_1'. By the induction argument as above, we only need to show that, for an edge e_1 with $v_1 \in e_1$, if $e_1 \subset Y_1$, then $e_1 \subset Y_1'$. Lemma 2.4 implies that there exists an edge $e_2 \subset Y_2$ such that e_1 and e_2 cobound an isometrically embedded flat rectangle (one side of the rectangle is l). So l is in the carrier of the hyperplane dual to e_1. It follows that $V_{e_1} \in V'$ and $e_1 \subset Y_1'$. □

Corollary 3.2 *Let K_1, K_2, Y_1 and Y_2 be as above. Let h be a hyperplane separating K_1 and K_2 and let e be an edge dual to h. Then $V_e \in V_{Y_1}^\perp = V_{Y_2}^\perp$. In particular, pick a vertex $v \in \Gamma$. Then $v \in V_{Y_1}$ if and only if*

(1) $v \in V_{K_1} \cap V_{K_2}$, and

(2) for any hyperplane h' separating K_1 from K_2 and any edge e' dual to h', $d(v, V_{e'}) = 1$.

Proof Let l and V_l be as in the proof of Lemma 3.1. Let V'_l be a collection of vertices of Γ such that $v \in V'_l$ if and only if $v = V_{e'}$ for some edge $e' \subset X(\Gamma)$ satisfying (2). It suffices to prove $V'_l = V_l$.

It is clear that $V'_l \subset V_l$ since if a hyperplane h separates K_1 from K_2, then l intersects h transversally at one point. To see $V_l \subset V'_l$, it suffices to show $h \cap K_i = \varnothing$ for $i = 1, 2$, where h is a hyperplane that intersects l transversally. Let $x = l \cap h$. Suppose $h \cap K_1 \neq \varnothing$ and let $x' = \pi_{h \cap K_1}(x)$. Now consider the triangle $\Delta(v_1, x, x')$ (recall that $v_1 = l(0)$). We have $\angle_{v_1}(x, x') \geq \frac{\pi}{2}$ (since $\pi_{K_1}(x) = v_1$), $\angle_{x'}(v_1, x) \geq \frac{\pi}{2}$ (see the proof of Lemma 2.6) and $\angle_x(v_1, x') > 0$, which is a contradiction, so $h \cap K_1 = \varnothing$. Similarly, $h \cap K_2 = \varnothing$. $\square$

Remark 3.3 Recall that a standard coset of $G(\Gamma)$ is a left coset of a standard subgroup of $G(\Gamma)$. Lemma 3.1 implies that for each pair of standard cosets of $G(\Gamma)$, we can associated another standard coset which captures the coarse intersection of the pair. Moreover, we can also define a notion of distance between two standard cosets, which takes values in $G(\Gamma)$.

Recall that Γ_K is the support of K (see Section 2.4), and that $\mathrm{lk}(\Gamma_K)$ is the full subgraph spanned by $V_K^{\perp}$ (see Section 2.1).

Lemma 3.4 *Let $K \subset X(\Gamma)$ be a convex subcomplex and let $\Gamma' = \mathrm{lk}(\Gamma_K)$. Then the parallel set P_K of K is a convex subcomplex and canonically splits as $K \times X(\Gamma')$.*

Note that we do not require K to satisfy the geodesic extension property.

Proof Pick a vertex $v \in K$. Let $\Gamma'' = \Gamma_K$ and let P_1 be the unique standard subcomplex that passes through v and has defining graph $\Gamma' \circ \Gamma''$ (recall that $\circ$ denotes the graph join). Then $K \subset P_1$. Let P' be the natural copy of $K \times X(\Gamma')$ inside P_1. It is clear that $P' \subset P_K$.

Let K' be a convex subset parallel to K, and let $\phi \colon K \to K'$ be the isometry induced by CAT(0) projection onto K'. Pick a vertex $v \in K$, and let l be the geodesic segment connecting v and $\phi(v)$. We define V_l as in the proof of Lemma 3.1 (note that $\phi(v)$ is not necessarily a vertex). Let e be any edge such that $v \in e \subset K$. Then there is a flat rectangle with e, $\phi(e)$ and l as its three sides. Thus l is contained in the carrier of the hyperplane dual to e, and $V_l \subset V_e^{\perp}$. Note that if l' is the side opposite to l, then $V_{l'} = V_l$. For any given edge $e' \subset K$, we can find an edge path $\omega \subset K$ such that e is

the first and e' is the last edge in ω. By induction on the combinatorial length of w and the same argument as above, we can show $V_l \subset V_{e'}^{\perp}$, thus $V_l \subset V_K^{\perp}$ and $K' \subset P'$. It follows that $P_K \subset P'$, so $P_K = P'$. $\qquad\qquad\square$

Remark 3.5 The following is a generalization of the above lemma for general CAT(0) cube complexes. Let X be a CAT(0) cube complex. A convex set $K \subset X$ is *regular* if for any $x \in K$, the space of direction $\Sigma_x K$ of x in K [13, Chapter II.3] satisfies:

(1) $\Sigma_x K$ is a subcomplex of $\Sigma_x X$ with respect to the canonical all-right spherical complex structure on $\Sigma_x X$.

(2) There exists $r > 0$ such that $B(x, r) \cap K$ is isometric to the r–ball centered at the cone point in the Euclidean cone over $\Sigma_x K$.

If $K \subset X$ is a regular convex subset, then P_K is convex and admits a splitting $P_K \cong K \times N$, where N has an induced cubical structure from X (N is CAT(0)).

Lemma 3.6 Let $q \colon X(\Gamma_1) \to X(\Gamma_2)$ be an (L, A)–quasi-isometry and let $F \subset X(\Gamma_1)$ be a subcomplex isometric to $\mathbb{E}^k$. Suppose $n = \dim(X(\Gamma_1)) = \dim(X(\Gamma_2))$. If there exist $R_1, R_2 > 0$ and top-dimensional flats F_1 and F_2 such that

$$F \overset{R_2}{=} F_1 \cap_{R_1} F_2 \quad \text{and} \quad F \overset{\infty}{=} F_1 \cap_R F_2$$

for any $R \geq R_1$, then:

(1) There exist a constant $D = D(L, A, R_1, R_2, n)$ and a subcomplex $F' \subset X(\Gamma_2)$ isometric to $\mathbb{E}^k$ such that $q(F) \overset{D}{=} F'$.

(2) There exists a constant $D' = D'(L, A)$ such that $q(P_F) \overset{D'}{=} P_{F'}$.

Proof By Theorem 2.10, there exist top-dimensional flats $F_1', F_2' \subset X(\Gamma_2)$ such that $q(F_i) \overset{D_1}{=} F_i'$ for $D_1 = D_1(L, A)$ and $i = 1, 2$. Thus there exist $R' = R'(L, A, R_1, R_2)$ and $R_3 = R_3(L, A, R_1, R_2) > R_1$ such that $q(F_1 \cap_{R_1} F_2) \subset F_1' \cap_{R'} F_2' \subset q(F_1 \cap_{R_3} F_2)$; this and Remark 2.5 imply

(3-1) $$q(F_1 \cap_{R_1} F_2) \overset{D_2}{=} F_1' \cap_{R'} F_2'$$

for $D_2 = D_2(n, d(F_1, F_2)) = D_2(L, A, R_1, R_2, n)$.

Let $(Y_1, Y_2) = \mathcal{I}(F_1', F_2')$. Then there exists $D_3 = D_3(L, A, R_1, R_2, n)$ such that

(3-2) $$Y_1 \overset{D_3}{=} F_1' \cap_{R'} F_2'.$$

From (3-1) and (3-2), we have

(3-3) $$q(F) \overset{D_4}{=} Y_1$$

for $D_4 = D_4(L, A, R_1, R_2, n)$. By Lemma 2.4, Y_1 is a convex subcomplex of F_1'. This together with (3-3) implies $Y_1 = F' \times \prod_{i=1}^{k'} I_i$, where F' is isometric to $\mathbb{E}^k$ and $\{I_i\}_{i=1}^{k'}$ are finite intervals. Moreover, by (3-3), $\mathrm{diam}\left(\prod_{i=1}^{k'} I_i\right)$ must be bounded in terms of D_4, L and A; thus (1) follows.

Let $\{F_\lambda\}_{\lambda \in \Lambda}$ be the collection of top-dimensional flats in $X(\Gamma_1)$ which are contained in the parallel set P_F of F. Lemma 3.4 implies

$$(3\text{-}4) \qquad d_H\left(\bigcup_{\lambda \in \Lambda} F_\lambda, P_F\right) \leq 1.$$

For $\lambda \in \Lambda$, there exists $R_\lambda > 0$ such that $F \subset_{R_\lambda} F_\lambda$. Let F_λ' be the top-dimensional flat in $X(\Gamma_2)$ such that $q(F_\lambda) \overset{D_1}{=} F_\lambda'$. Then by (1), there exists $R_\lambda' > 0$ such that $F' \subset_{R_\lambda'} (F_\lambda')$. This and Lemma 2.4 imply $F_\lambda' \subset P_{F'}$ for any $\lambda \in \Lambda$. Thus by (3-4), there exists $D' = D'(L, A)$ such that $q(P_F) \subset_{D'} P_{F'}$. And (2) follows by running the same argument for the quasi-isometry inverse of q. □

A *tree product* is a convex subcomplex $K \subset X(\Gamma)$ such that K splits as a product of trees, ie there exists a cubical isomorphism $K \cong \prod_{i=1}^n T_i$ where the T_i are trees. A *standard tree product* is a tree product which is also a standard subcomplex.

One can check that K is a standard tree product if and only if the defining graph Γ_K of K has a join decomposition $\Gamma_K = \Gamma_1 \circ \Gamma_2 \circ \cdots \circ \Gamma_n$, where each Γ_i is discrete. Thus one can choose the above T_i to be standard subcomplexes of K. Note that every standard flat is a standard tree product, and every subcomplex isometric to $\mathbb{E}^k$ is a tree product.

Lemma 3.7 *Suppose* $\dim(X(\Gamma)) = n$. *Let* $q\colon X(\Gamma) \to X(\Gamma')$ *be a quasi-isometry. Let* $K = \prod_{i=1}^n T_i$ *be a top-dimensional tree product with its tree factors. Then there exists a standard tree product* K' *in* $X(\Gamma')$ *such that* $q(K) \subset_\infty K'$.

The proof essentially follows [9, Theorem 4.2].

Proof For $1 \leq i \leq n$, let $V_i = V_{T_i} \in \Gamma$ be the collection of labels of edges in T_i. The case where all the V_i are consist of one point follows from Theorem 2.10. If each V_i contains at least two points, then by Lemma 3.6, for any geodesic $l \subset T_i$, there exists a subcomplex $l' \subset X(\Gamma')$ isometric to $\mathbb{R}$ such that $q(l) \overset{\infty}{=} l'$. Since l' is unique up to parallelism, the collection of labels of edges in l' does not depend on the choice of l' and will be denoted by $V_{q(l)}$. For $1 \leq i \leq n$, define $V_i' = \bigcup_{l \subset T_i} V_{q(l)}$ where l varies among all geodesics in T_i.

We claim $V_i' \subset (V_j')^{\perp}$ for $i \neq j$. To see this, pick geodesic $l_i \in T_i$ and let $F = \prod_{i=1}^{n} l_i$. Then there exist top-dimensional flat F' and geodesic lines $\{l_i'\}_{i=1}^{n}$ (each l_i' is a subcomplex) in $X(\Gamma')$ such that $q(F) \overset{\infty}{\cong} F'$ and $q(l_i) \overset{\infty}{\cong} l_i'$. Since $l_i' \subset_{\infty} F'$, by Lemma 2.4, we can assume l_i' is a subcomplex of F'. Pick $i \neq j$. Since l_i and l_j are orthogonal, they have infinite Hausdorff distance. Thus l_i' and l_j' have infinite Hausdorff distance. By our assumption, l_i' and l_j' are isometric to $\mathbb{E}^1$, and they are convex subcomplexes of $F' \cong \mathbb{E}^n$. Thus either l_i' and l_j' are parallel, or they are orthogonal. The former is impossible since l_i' and l_j' have infinite Hausdorff distance. Thus $\{l_i'\}_{i=1}^{n}$ is a mutually orthogonal collection.

Let $\Gamma_1' = V_1' \circ V_2' \circ \cdots \circ V_n' \subset \Gamma'$. Then each V_i' has to be a discrete full subgraph by our dimension assumption. Let $\{F_\lambda\}_{\lambda \in \Lambda}$ be the collection of top-dimensional flats in K and let F_λ' be the unique flat such that $q(F_\lambda) \overset{\infty}{\cong} F_\lambda'$. Note that for arbitrary F_{λ_1} and F_{λ_2}, there exists a finite chain in $\{F_\lambda\}_{\lambda \in \Lambda}$ which starts with F_{λ_1} and ends with F_{λ_2} such that the intersection of adjacent elements in the chain contains a top-dimensional orthant. Thus the collection $\{F_\lambda'\}_{\lambda \in \Lambda}$ also has this property. Then $\bigcup_{\lambda \in \Lambda} F_\lambda'$ is contained in a standard subcomplex of $X(\Gamma')$ with defining graph Γ_1'.

It remains to deal with the case where there exist $i \neq j$ such that $|V_i| = 1$ and $|V_j| \geq 2$. We suppose $|V_i| = 1$ for $1 \leq i \leq m$ and $|V_i| \geq 2$ for $i > m$. By applying Lemma 3.6 with $F = \prod_{i=1}^{m} T_i$, we can reduce to a lower-dimensional case, and the lemma follows by induction on dimension. $\square$

Corollary 3.8 *Let $q \colon X(\Gamma) \to X(\Gamma')$ be a quasi-isometry, and let K be a top-dimensional maximal standard tree product; ie K is not properly contained in another tree product. Then there exists a standard tree product $K' \subset X(\Gamma')$ such that $q(K) \overset{\infty}{\cong} K'$.*

3.2 Standard flats in transvection-free RAAGs

Up to now, we have only dealt with top-dimensional standard subcomplexes. The next goal is to study those standard subcomplexes which are not necessarily top dimensional. In particular, we are interested in whether quasi-isometries will preserve standard flats up to finite Hausdorff distance. The answer turns out to be related to the outer automorphism group of $G(\Gamma)$.

One direction is obvious: namely, if every quasi-isometry $q \colon X(\Gamma) \to X(\Gamma')$ maps any standard flat in $X(\Gamma)$ to a standard flat in $X(\Gamma')$ up to finite Hausdorff distance, then $\mathrm{Out}(G(\Gamma))$ must be transvection-free (ie $\mathrm{Out}(G(\Gamma))$ does not contain any transvections). The converse is also true. Now we set up several necessary tools to prove the converse.

In this section, Γ will always be a finite simplicial graph.

Definition 3.9 A subgraph $\Gamma_1 \subset \Gamma$ is *stable* in Γ if the following are true:

(1) Γ_1 is a full subgraph.

(2) Let $K \subset X(\Gamma)$ be a standard subcomplex such that $\Gamma_K = \Gamma_1$, and let Γ' be a finite simplicial graph such that, for some L and A, there is an (L, A)–quasi-isometry $q \colon X(\Gamma) \to X(\Gamma')$. Then there exists $D = D(L, A, \Gamma_1, \Gamma) > 0$ and a standard subcomplex $K' \subset X(\Gamma')$ such that $d_H(q(K), K') < D$.

For simplicity, we will also say the pair (Γ_1, Γ) is *stable* in this case. A standard subcomplex $K \subset X(\Gamma)$ is *stable* if it arises from a stable subgraph of Γ.

We claim the defining graph $\Gamma_{K'}$ of K' is stable in Γ'. To see this, pick any graph Γ'' so that there is an (L, A)–quasi-isometry $q' \colon X(\Gamma') \to X(\Gamma'')$, and pick any standard subcomplex $K_1' \subset X(\Gamma')$ with defining graph $\Gamma_{K'}$. Note that there is an isometry $i \colon X(\Gamma') \to X(\Gamma')$ such that $i(K') = K_1'$. Since the map $q' \circ i \circ q$ is a quasi-isometry from $X(\Gamma)$ to $X(\Gamma'')$, we have that $q' \circ i \circ q(K)$ is Hausdorff close to a standard subcomplex in $X(\Gamma'')$ by the stability of Γ_1; hence the same is true for $q'(K_1')$. It follows from this claim that one can obtain quasi-isometric invariants by identifying certain classes of stable subgraphs.

It is immediate from the definition that for finite simplicial graphs $\Gamma_1 \subset \Gamma_2 \subset \Gamma_3$, if (Γ_1, Γ_2) is stable and (Γ_2, Γ_3) is stable, then (Γ_1, Γ_3) is stable. However, it is not necessarily true that if (Γ_1, Γ_3) and (Γ_2, Γ_3) are stable, then (Γ_1, Γ_2) is stable. In the sequel, we will investigate several other properties of stable subgraphs. The following lemma is an easy consequence of Lemma 3.1 and Remark 2.5:

Lemma 3.10 *Suppose* Γ_1 *and* Γ_2 *are stable in* Γ. *Then* $\Gamma_1 \cap \Gamma_2$ *is also stable in* Γ.

The following result follows from Lemma 2.8.

Lemma 3.11 *If* Γ_1 *is stable in* Γ, *then every connected component of* Γ_1 *that contains more than one point is also stable in* Γ.

Lemma 3.12 *Suppose* Γ_1 *is stable in* Γ. *Let* V *be the vertex set of* Γ_1 *and let* Γ_2 *be the full subgraph spanned by* V *and the orthogonal complement* $V^\perp$. *Then* Γ_2 *is also stable in* Γ.

Proof Let $K_2 \subset X(\Gamma)$ be a standard subcomplex with defining graph $\Gamma_{K_2} = \Gamma_2$, and let $K_1 \subset K_2$ be any standard subcomplex satisfying $\Gamma_{K_1} = \Gamma_1$. Lemma 3.4 implies $K_2 = P_{K_1} = K_1 \times K_1^\perp$. For a vertex $x \in K_1^\perp$, let $K_x = K_1 \times \{x\}$. Let $q \colon X(\Gamma) \to X(\Gamma')$

be an (L, A)–quasi-isometry. Then there exists standard subcomplex K'_x such that $d_H(q(K_x), K'_x) < D = D(L, A, \Gamma_1, \Gamma)$. Thus $K'_x \overset{\infty}{\cong} K'_y$ for vertices $x, y \in K_1^\perp$. It follows from Lemma 3.1 that K'_x and K'_y are parallel. Thus $q(P_{K_1}) \subset_R P_{K'_x}$ for $R = D + L + A$. Moreover, $P_{K'_x}$ is also a standard subcomplex by Lemma 3.4. By considering the quasi-isometry inverse and repeating the previous argument, we know $q(P_{K_1}) \overset{\infty}{\cong} P_{K'_x}$; thus Γ_2 is also stable in Γ. □

Lemma 3.13 *Suppose Γ_1 is stable in Γ. Pick a vertex $v \notin \Gamma_1$. Then the full subgraph spanned by $v^\perp \cap \Gamma_1$ is stable in Γ.*

Proof We use Γ_2 to denote the full subgraph spanned by $v^\perp \cap \Gamma_1$. Let $K_2 \subset X(\Gamma)$ be a standard subcomplex such that $\Gamma_{K_2} = \Gamma_2$, and let $K_1 \subset X(\Gamma)$ be the unique standard subcomplex such that $\Gamma_{K_1} = \Gamma_1$ and $K_2 \subset K_1$. Pick a vertex $x \in K_2$, and let $e \subset X(\Gamma)$ be the edge such that $V_e = v$ and $x \in e$. Suppose $\bar{x}$ is the other end point of e. Let $\bar{K}_i$ be the standard subcomplex that contains $\bar{x}$ and has defining graph Γ_i for $i = 1, 2$. Denote the hyperplane dual to e by h. Since $v \notin \Gamma_1$, we have $h \cap K_1 = \varnothing$ and $h \cap \bar{K}_1 = \varnothing$; thus h separates K_1 and $\bar{K}_1$, and $d(K_1, \bar{K}_1) = 1$. It follows from Corollary 3.2 that $\mathcal{I}(K_1, \bar{K}_1) = (K_2, \bar{K}_2)$; in particular $K_2 \overset{D}{\cong} K_1 \cap_R \bar{K}_1$ for D depending on R and the dimension of $X(\Gamma)$. Now the lemma follows since Γ_1 is stable. □

The next result is a direct consequence of Corollary 3.8.

Lemma 3.14 *If Γ_1 is stable in Γ, then there exists Γ_2 which is stable in Γ_1 such that*

(1) *Γ_2 is a graph join $\bar{\Gamma}_1 \circ \bar{\Gamma}_2 \circ \cdots \circ \bar{\Gamma}_k$, where $\bar{\Gamma}_i$ is discrete for $1 \le i \le k$;*

(2) *$k = \dim(X(\Gamma_1))$.*

Lemma 3.15 *Let Γ be a finite simplicial graph such that there do not exist vertices $v \ne w$ of Γ with $v^\perp \subset \mathrm{St}(w)$. Then every stable subgraph of Γ contains a stable vertex.*

Proof Let Γ_1 be a minimal stable subgraph; ie it does not properly contain any stable subgraph of Γ. It suffices to show Γ_1 is a point. We argue by contradiction and assume Γ_1 contains more than one point.

First we claim Γ_1 cannot be discrete. Suppose the contrary is true. Pick vertices $v, w \in \Gamma_1$ and pick a vertex $u \in v^\perp \setminus \mathrm{St}(w)$. By Lemma 3.13, $u^\perp \cap \Gamma_1$ is also stable. Note that $v \in u^\perp \cap \Gamma_1$ and $w \notin u^\perp \cap \Gamma_1$, which contradicts the minimality of Γ_1.

We claim Γ_1 must be a clique. Since Γ_1 is not discrete, by Lemma 3.14, we can find a stable subgraph

(3-5) $$\Gamma_2 = \bar{\Gamma}_1 \circ \bar{\Gamma}_2 \circ \cdots \circ \bar{\Gamma}_m \subset \Gamma_1,$$

where $\{\overline{\Gamma}_i\}_{i=1}^m$ are discrete full subgraphs and $m \geq 2$. Then $\Gamma_2 = \Gamma_1$. Suppose some $\overline{\Gamma}_i$ contains more than one point, and let Γ_3 be the join of the remaining join factors. Then Theorem 2.9 implies that Γ_3 is stable, contradicting the minimality of Γ_1. Therefore, Γ_1 is a clique.

Pick distinct vertices $v_1, v_2 \in \Gamma_1$. By our assumption, there exists a vertex $w \in v_1^\perp \setminus \mathrm{St}(v_2)$. Since Γ_1 is a clique, $\Gamma_1 \subset \mathrm{St}(v_2)$, so $w \notin \Gamma_1$. Let Γ_4 be the full subgraph spanned by $w^\perp \cap \Gamma_1$. Then Γ_4 is stable by Lemma 3.13. Moreover, $\Gamma_4 \subsetneq \Gamma_1$ (since $v_2 \notin \Gamma_4$), which yields a contradiction.   □

Lemma 3.16 *Let Γ be as in Lemma 3.15 and let Γ_1 be a stable subgraph of Γ. Then for any vertex $w \in \Gamma_1$, there exists a stable vertex $v \in \Gamma_1$ such that $d(v, w) \leq 1$.*

Proof Denote the combinatorial distances in Γ and Γ_1 by d and d_1, respectively. Since Γ_1 is a full subgraph, $d(x, y) = 1$ if and only if $d_1(x, y) = 1$, and $d(x, y) \geq 2$ if and only if $d_1(x, y) \geq 2$, for vertices $x, y \in \Gamma_1$. If w is isolated in Γ_1, then we can use the argument in the second paragraph of the proof of Lemma 3.15 to get rid of all vertices in Γ_1 except w, which implies w is a stable vertex. If w is not isolated, we can assume Γ_1 is connected by Lemma 3.11.

By Lemma 3.15, there exists a stable vertex $u \in \Gamma_1$. If $d_1(u, w) \leq 1$, then we are done. Otherwise, let ω be a geodesic in Γ_1 connecting u and w (note that ω might not be a geodesic in Γ), and let $\{v_i\}_{i=0}^n$ be the consecutive vertices in ω; here $v_0 = w$, $v_n = u$ and $n = d_1(w, u)$.

Since u is stable, by Lemma 3.12, $\mathrm{St}(u)$ is also stable. Note that $d_1(v_{n-2}, u) = 2$, so $d(v_{n-2}, u) = 2$ and $v_{n-2} \notin \mathrm{St}(u)$. Lemma 3.13 implies $v_{n-2}^\perp \cap \mathrm{St}(u)$ is stable, and by Lemma 3.10, $v_{n-2}^\perp \cap \mathrm{St}(u) \cap \Gamma_1$ is also stable. Note that $v_{n-2}^\perp \cap \mathrm{St}(u) \cap \Gamma_1 \neq \varnothing$ since it contains v_{n-1}. Lemma 3.15 implies there is a stable vertex $u' \in v_{n-2}^\perp \cap \mathrm{St}(u) \cap \Gamma_1$, and it is easy to see $d_1(w, u') = n - 1$. Now the lemma follows by induction.   □

Lemma 3.17 *Let Γ be as in Lemma 3.15. Then every vertex of Γ is stable.*

Proof Let Γ_w be the intersection of all the stable subgraphs that contain w. By Lemma 3.10, Γ_w is the minimal stable subgraph that contains w. It suffices to prove $\Gamma_w = \{w\}$. We argue by contradiction and denote the vertices in $\Gamma_w \setminus \{w\}$ by $\{v_i\}_{i=1}^k$. The minimality of Γ_w implies we cannot use Lemma 3.13 to get rid of some v_i while keeping w; thus $w^\perp \setminus \mathrm{St}(v_i) \subset \{v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_k\}$ for any i. In other words,

$$(3\text{-}6) \qquad w^\perp \subset \mathrm{St}(v_i) \cup \{v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_k\}$$

for $1 \leq i \leq k$. Then there does not exist i such that $\Gamma_w \subset \mathrm{St}(v_i)$: otherwise, we would have $w^\perp \subset \mathrm{St}(v_i)$ by (3-6).

On the other hand, Lemma 3.16 implies there exists a stable vertex $u \in \Gamma_w$ with $d(w, u) = 1$. Then $\mathrm{St}(u)$ is stable (Lemma 3.12) and $\mathrm{St}(u) \cap \Gamma_w$ is stable (Lemma 3.10). Note that $w \in \mathrm{St}(u) \cap \Gamma_w$. By the minimality of Γ_w, we have $\Gamma_w \subset \mathrm{St}(u)$, which yields a contradiction. $\qquad \square$

Lemma 3.18 *Let Γ be a finite simplicial graph, and pick stable subgraphs Γ_1, Γ_2 of Γ. Let $\overline{\Gamma}$ be the full subgraph spanned by V and $V^\perp$, where $V = V_{\Gamma_1}$. If $\Gamma_2 \subset \overline{\Gamma}$, then the full subgraph spanned by $\Gamma_1 \cup \Gamma_2$ is stable in Γ.*

To simplify notation, in the following proof, we will write $q(K) \approx K'$, where q, K and K' are as in Definition 3.9. We will also assume without loss of generality that $q(K) \subset K'$.

Proof Let $q \colon X(\Gamma) \to X(\Gamma')$ be an (L, A)–quasi-isometry. Suppose K_1 and K are standard subcomplexes in $X(\Gamma)$ such that $\Gamma_{K_1} = \Gamma_1$, $\Gamma_K = \overline{\Gamma}$ and $K_1 \subset K$. Put $K' \approx q(K)$, $K'_1 \approx q(K_1)$, $K = K_1 \times K_1^\perp$ and $K' = K'_1 \times K_1'^\perp$. The proof of Lemma 3.12 implies there exist a quasi-isometry $q' \colon K_1^\perp \to K_1'^\perp$ and a constant $D_1 = D_1(L, A, \Gamma_1, \Gamma)$ such that

$$(3\text{-}7) \qquad\qquad d(q' \circ p_2(x), p'_2 \circ q(x)) < D_1$$

for any $x \in K$, where $p_2 \colon K \to K_1^\perp$ and $p'_2 \colon K' \to K_1'^\perp$ are projections.

Let $\Gamma_2 = \Gamma_{21} \circ \Gamma_{22}$, where $\Gamma_{21} = \Gamma_1 \cap \Gamma_2$, and let K_{22}, K_2 be standard subcomplexes such that $\Gamma_{K_{22}} = \Gamma_{22}$, $\Gamma_{K_2} = \Gamma_2$ and $K_{22} \subset K_2 \subset K$. By (3-7), it suffices to prove there exist a standard subcomplex $K'_{22} \subset K'$ and a constant $D = D(L, A, \Gamma_1, \Gamma_2, \Gamma)$ such that $d_H(p'_2 \circ q(K_{22}), K'_{22}) < D$. Let $K'_2 \approx q(K_2)$. Then $K'_2 \subset K'$, and $p'_2(K'_2)$ is a standard subcomplex. By (3-7), $p'_2 \circ q(K_{22}) \overset{\infty}{\cong} p'_2 \circ q(K_2) \overset{\infty}{\cong} p'_2(K'_2)$; thus we can take $K'_{22} = p'_2(K'_2)$. $\qquad \square$

Remark 3.19 In general, the full subgraph spanned by $\Gamma_1 \cup \Gamma_2$ is not necessarily stable even if Γ_1 and Γ_2 are stable; see Remark 3.26.

The next theorem follows from Lemma 3.17 and Lemma 3.18.

Theorem 3.20 *Given a finite simplicial graph Γ, the following are equivalent:*

(1) $\mathrm{Out}(G(\Gamma))$ *is transvection-free.*

(2) *For any (L, A)–quasi-isometry $q \colon X(\Gamma) \to X(\Gamma')$, there exists a positive constant $D = D(L, A, \Gamma)$ such that for any standard flat $F \subset X(\Gamma)$, there exists a standard flat $F' \subset X(\Gamma')$ with $d_H(q(F), F') < D$.*

3.3 Standard flats in general RAAGs

At this point, we have the following natural questions:

(1) In Theorem 3.20, is it true that every standard flat in $X(\Gamma')$ comes from some standard flat in $X(\Gamma)$? A related question could be, is condition (1) in Theorem 3.20 a quasi-isometric invariant for right-angled Artin groups?

(2) What can we say about the stable subgraphs of Γ if we drop condition (1) in Theorem 3.20?

We will first give a negative answer to question (1) in Example 3.22 below. Then we will prove Theorem 3.21, which answers question (2). Section 4 and, in particular, the proof of Theorem 1.1 will not depend on this subsection. However, we will need Theorem 3.21 and Lemma 3.23 for Section 5.

Theorem 3.21 Let Γ be an arbitrary finite simplicial graph. A clique $\Gamma_1 \subset \Gamma$ is stable if and only if there do not exist vertices $w \in \Gamma_1$ and $v \in \Gamma \setminus \Gamma_1$ such that $w^\perp \subset \mathrm{St}(v)$.

In other words, the clique Γ_1 is stable if and only if the corresponding $\mathbb{Z}^n$ subgroup of $G(\Gamma_1)$ is invariant under all transvections.

Example 3.22 Let Γ_1 and Γ_2 be as indicated in Figure 2. It is easy to see $\mathrm{Out}(G(\Gamma_1))$ is transvection-free while $\mathrm{Out}(G(\Gamma_2))$ contains nontrivial transvection (Γ_2 has a dead end at vertex u). We claim $G(\Gamma_1)$ and $G(\Gamma_2)$ are commensurable and, in particular, quasi-isometric.
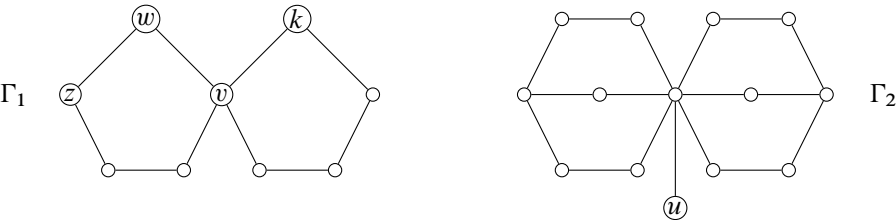

Figure 2: $\mathrm{Out}(G(\Gamma_1))$ is transvection-free while $\mathrm{Out}(G(\Gamma_2))$ contains nontrivial transvection.

Let $\Gamma \subset \Gamma_1$ be the pentagon on the left side and let Y be the Salvetti complex of Γ. Suppose $X_1 = Y \sqcup Y \sqcup (\mathbb{S}^1 \times [0, 1])/\sim$; here the two boundary circles of the annulus are identified with two standard circles which are in different copies of Y. Then $\pi_1(X_1) = G(\Gamma_1)$. Define the homomorphism $h_1 \colon G(\Gamma) \to \mathbb{Z}/2$ by sending w to the nontrivial element in $\mathbb{Z}/2$ and other generators to the identity element. Let Y' be the 2–sheeted cover of Y with respect to $\ker(h_1)$.

Figure 3

Define the homomorphism $h_2\colon G(\Gamma_1) \to \mathbb{Z}/2$ by sending w and k to the nontrivial element in $\mathbb{Z}/2$ and other generators to the identity element. Let X be t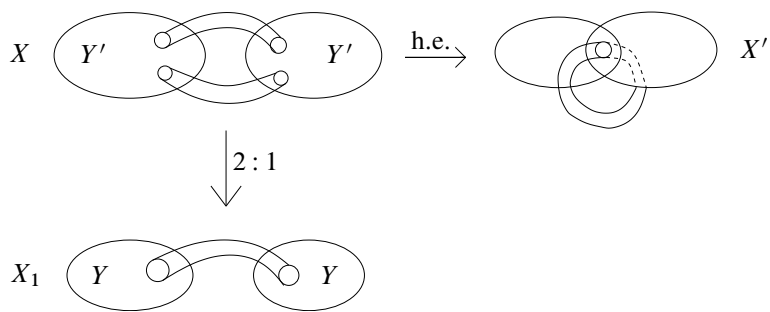he 2–sheeted cover of X_1 with respect to $\ker(h_2)$. Then X is made of two copies of Y' and two annuli with the boundaries of the annuli identified with the v–circles in Y' (each Y' has two v–circles which cover the v–circle in Y); see Figure 3.

The cover X is homotopy equivalent to a Salvetti complex. To see this, let S_w be the circle in Y' which covers the w–circle in Y two times and let $S_z \vee S_v$ be a wedge of the two circles in Y' which covers the wedge of the z–circle and the v–circle in Y. There is a copy of $S_w \times (S_z \vee S_v)$ inside Y'. Let I be a segment in S_w such that its end points are mapped to the base point of Y under the covering map. We collapse $I \times (S_z \vee S_v)$ to $\{pt\} \times (S_z \vee S_v)$ inside each copy of Y' in X, and collapse one of the annuli in X to a circle by killing the interval factor. Denote the resulting space by X'. Then X' is homotopy equivalent to X, and the uncollapsed annulus in X becomes a torus in X'. It is not hard to see X' is a Salvetti complex with defining graph Γ_2.

Any standard geodesic in $X(\Gamma_2)$ which comes from vertex u is not Hausdorff close to a quasi-isometric image of some standard geodesic in $X(\Gamma_1)$, since u is not a stable vertex while every vertex in Γ_1 is stable.

Here is a generalization of the above example. Suppose Γ is a finite simplicial graph with vertices $v_1, v_2 \in \Gamma$ such that $d(v_1, v_2) = 2$ and they are separated by the intersection of links $\mathrm{lk}(v_1) \cap \mathrm{lk}(v_2)$. Define a homomorphism $h\colon G(\Gamma) \to \mathbb{Z}/2$ by sending v_1 and v_2 to the nontrivial element in $\mathbb{Z}/2$ and killing all other generators. Then $\ker(h)$ is also a right-angled Artin group by the same argument as before. To find its defining graph, let $\{C_i\}_{i=1}^{n}$ be the components of $\Gamma \setminus (\mathrm{lk}(v_1) \cap \mathrm{lk}(v_2))$, and suppose $v_1 \in C_1$. Define $\Gamma_1 = C_1 \cup (\mathrm{lk}(v_1) \cap \mathrm{lk}(v_2))$ and $\Gamma_2 = \left(\bigcup_{i=2}^{n} C_i\right) \cup (\mathrm{lk}(v_1) \cap \mathrm{lk}(v_2))$. Then Γ_1 and Γ_2 are full subgraphs of Γ; moreover, $\mathrm{St}(v_i) \in C_i$. For $i = 1, 2$, let Γ_i' be the graph obtained by gluing two copies of Γ_i along $\mathrm{St}(v_i)$, and let Γ_3' be the join of one point and $\mathrm{lk}(v_1) \cap \mathrm{lk}(v_2)$. Then the defining graph of $\ker(h)$ can be obtained by gluing Γ_1', Γ_2' and Γ_3' along $\mathrm{lk}(v_1) \cap \mathrm{lk}(v_2)$.

Note that we are taking advantage of separating closed stars while constructing the counterexample. If separating closed stars are not allowed in Γ, then we have a positive answer to question (1); see Section 5.

In the rest of this subsection, we will prove Theorem 3.21. Γ will be an arbitrary finite simplicial graph in the rest of this subsection. Theorem 3.21 is actually a consequence of the following more general result.

Lemma 3.23 *Pick a vertex $w \in \Gamma$, and let Γ_w be the intersection of all stable subgraphs of Γ that contain w. Define $W = \{w' \in \Gamma \mid w^\perp \subset \mathrm{St}(w')\}$. Then Γ_w is the full subgraph spanned by W.*

In other words, $G(\Gamma_w) \leq G(\Gamma)$ is the minimal standard subgroup containing w with the property that $G(\Gamma_w)$ is invariant under any transvection.

Now we show how to deduce Theorem 3.21 from Lemma 3.23

Proof of Theorem 3.21 The "only if" part can be proved by contradiction (choose a transvection which does not preserve the subgroup $G(\Gamma_1)$). For the converse, let $\{v_i\}_{i=1}^n$ be the vertex set of Γ_1, and let Γ_{v_i} be the minimal stable subgraph that contains v_i for $1 \leq i \leq n$. By our assumption and Lemma 3.23, $\Gamma_{v_i} \subset \Gamma_1$. Thus the full subgraph spanned by $\bigcup_{i=1}^n \Gamma_{v_i}$ is stable by Lemma 3.18, which means Γ_1 is stable. $\square$

It remains to prove Lemma 3.23. We first set up two auxiliary lemmas.

Lemma 3.24 *Let $v \in \Gamma$ be a vertex which is not isolated. Then at least one of the following is true:*

(1) *v is contained in a stable discrete subgraph with more than one vertex.*

(2) *v is contained in a stable clique subgraph.*

(3) *There is a stable discrete subgraph with more than one vertex whose vertex set is in $v^\perp$.*

(4) *There is a stable clique subgraph whose vertex set is in $v^\perp$.*

Proof Since v is not isolated, we can assume Γ is connected by Lemma 3.11. By Lemma 3.14, we can find a stable subgraph $\Gamma_1 = \overline{\Gamma}_1 \circ \overline{\Gamma}_2 \circ \cdots \circ \overline{\Gamma}_n$ where $\{\overline{\Gamma}_i\}_{i=1}^n$ are discrete full subgraphs and $n = \dim(X(\Gamma))$. If $v \in \Gamma_1$, then by the third paragraph of the proof of Lemma 3.15, we know either (1), (2) or (4) is true.

Suppose $d(v, \Gamma_1) = 1$. Let Γ_2 be the full subgraph spanned by $v^{\perp} \cap \Gamma_1$. Then Γ_2 is stable by Lemma 3.13. The proof of Lemma 3.15 implies every stable subgraph of Γ contains either a stable discrete subgraph or a stable clique subgraph (this does not depend on the $v^{\perp} \not\subseteq \mathrm{St}(w)$ assumption); thus either (3) or (4) is true.

Suppose $d(v, \Gamma_1) \geq 2$. Pick vertex $u \in \Gamma_1$ such that $d(v, u) = d(v, \Gamma_1) = n$, and let ω be a geodesic connecting v and u. Suppose $\{v_i\}_{i=0}^{n}$ are the consecutive vertices in ω such that $v_0 = v$ and $v_n = u$. Let Γ' be the full subgraph spanned by $v_{n-1}^{\perp} \cap \Gamma$, and let Γ'' be the full subgraph spanned by V and $V^{\perp}$, where $V = V_{\Gamma'}$ (the vertex set of Γ'). Then Γ' is stable by Lemma 3.13, and Γ'' is stable by Lemma 3.12. Note that $d(v, x) \geq n$ for any vertex $x \in V$, so $d(v, y) \geq n - 1$ for any vertex $y \in V^{\perp}$. Thus $d(v, \Gamma'') \geq n - 1$. However, $v_{n-1} \in \Gamma''$. So $d(v, \Gamma'') = n - 1$. Now we can induct on n and reduce to the $d(v, \Gamma_1) = 1$ case. $\qquad\square$

It is interesting to see that if Γ has large diameter, then there are a lot of nontrivial stable subgraphs.

We record the following lemma which is an easy consequence of Theorem 2.9.

Lemma 3.25 *Suppose* $\Gamma = \Gamma_1 \circ \Gamma_2$, *where* Γ_1 *is the maximal clique join factor of* Γ. *If* Γ_2' *is stable in* Γ_2, *then* $\Gamma_1 \circ \Gamma_2'$ *is stable in* Γ.

Now we are ready to prove Lemma 3.23.

Proof of Lemma 3.23 By Lemma 3.10, Γ_w is the minimal stable subgraph that contains w. If there exists vertex $w' \in W$ such that $w' \notin \Gamma_w$, then sending $w \to ww'$ and fixing all other vertices would induce a group automorphism, which gives rise to a quasi-isometry from $X(\Gamma)$ to $X(\Gamma)$. The existence of such a quasi-isometry would contradict the stability of Γ_w; thus $W \subset \Gamma_w$.

Let W' be the vertex set of Γ_w. It remains to prove $W' \subset W$. Suppose $W \subsetneq W'$ and let $u \in W' \setminus W$. Then $\varnothing \neq w^{\perp} \setminus \mathrm{St}(u)$. The minimality of Γ_w implies we cannot use Lemma 3.13 to get rid of u while keeping w; thus $w^{\perp} \setminus \mathrm{St}(u) \subset W' \setminus \{u, w\}$. In summary,

$$(3\text{-}8) \qquad\qquad \varnothing \neq w^{\perp} \setminus \mathrm{St}(u) \subset W' \setminus \{u, w\}.$$

In particular, w is not isolated in Γ_w, and

$$(3\text{-}9) \qquad\qquad \Gamma_w \not\subseteq \mathrm{St}(u).$$

Now we apply Lemma 3.24 to Γ_w and w, and recall that if a subgraph is stable in Γ_w, then it is stable in Γ. If case (1) in Lemma 3.24 is true, then we will get a contradiction since w is not isolated in Γ_w. If case (2) is true, then Γ_w sits inside some clique, which is contradictory to (3-9).
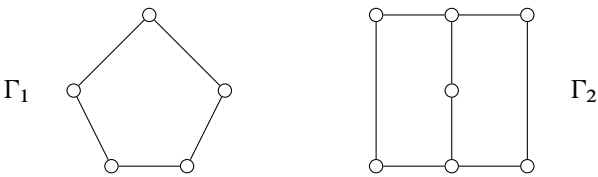
Figure 4: Failure of Theorem 3.21 in the more general case; see Remark 3.26

If case (3) is true, let $\Gamma_1 \subset \Gamma_w$ be the corresponding stable discrete subgraph. Let $V_1 = V_{\Gamma_1}$, and let $V_1' = \{u \in \Gamma_w \mid d(u,v) = 1 \text{ for any } v \in V_1\}$. Suppose Γ_w' is the full subgraph spanned by V_1 and V_1'. Then Γ_w' is stable by Lemma 3.12; hence $\Gamma_w' = \Gamma_w$. Let $\Gamma_w = \overline{\Gamma}_1 \circ \overline{\Gamma}_2 \circ \cdots \circ \overline{\Gamma}_k$ be the join decomposition induced by the De Rham decomposition of $X(\Gamma_w)$. Then $k \geq 2$ and u does not sit inside the clique factor by (3-9).

If there is no clique factor, then each join factor is stable by Theorem 2.9, and w is inside one of the join factors, which contradicts the minimality of Γ_w. If the clique factor exists and w sits inside the clique factor, then by Theorem 2.9, the clique factor is stable, and we have the same contradiction as before. If the clique factor exists and w sits outside the clique factor, this reduces to the next case.

If case (4) is true, let $\Gamma_2 \subset \Gamma_w$ be the corresponding stable clique subgraph. We can also assume without loss of generality that w is not contained in a stable clique. Let $V_2 = V_{\Gamma_2}$ and $V_2' = \{u \in \Gamma_w \mid d(u,v) = 1 \text{ for any } v \in V_2\}$. Suppose Γ_w'' is the full subgraph spanned by V_2 and V_2'. Then $\Gamma_w'' = \Gamma_w$ as before. Let $\Gamma_w = \Gamma_1' \circ \Gamma_2'$ where Γ_1' corresponds to the Euclidean De Rham factor of $X(\Gamma_w)$. Note that Γ_2' is nontrivial, and $w, u \in \Gamma_2'$ as in the discussion of case (3). Equation (3-8) implies that $w^\perp \not\subseteq \mathrm{St}(u)$ is still true if we take the orthogonal complement of w and the closed star of u in Γ_2'; in particular, w is not isolated in Γ_2'. Moreover, $\dim(X(\Gamma_2')) < \dim(X(\Gamma_w)) \leq \dim(X(\Gamma))$.

If $\dim(X(\Gamma)) = 2$, then Γ_2' has to be discrete, which is contradictory to the fact that w is not isolated in Γ_2'. If $\dim(X(\Gamma)) = n > 2$, then by induction, we can assume the lemma is true for all lower-dimensional graphs. Then there exists $\overline{\Gamma}_w$ stable in Γ_2' such that $w \in \overline{\Gamma}_w$ and $u \notin \overline{\Gamma}_w$. By Lemma 3.25, $\overline{\Gamma}_w \circ \Gamma_1'$ is stable in Γ_w, hence in Γ, which contradicts the minimality of Γ_w. $\qquad \square$

Remark 3.26 It is nature to ask whether Theorem 3.21 is still true if we do not require Γ_1 to be a clique. It turns out there are counterexamples. Let Γ_1 and Γ_2 be as indicated in Figure 4. Then $G(\Gamma_1)$ is quasi-isometric to $G(\Gamma_2)$ by the discussion in Section 11 of [9]. Let $q\colon X(\Gamma_2) \to X(\Gamma_1)$ be a quasi-isometry, and let K be a standard subcomplex in $X(\Gamma_2)$ such that its defining graph Γ_K is a pentagon in Γ_2. Suppose $q(K)$ is

Hausdorff close to a standard subcomplex K' in $X(\Gamma)$. Then $\Gamma_{K'}$ must be a connected proper subgraph of Γ_1, hence a tree. But this is impossible by the results in [7].

4 From quasi-isometries to isomorphisms

4.1 The extension complexes

4.1.1 Extension complexes and standard flats Let $q\colon X(\Gamma) \to X(\Gamma')$ be a quasi-isometry. Usually q does not induce a well-defined boundary map; see [20]. However, Theorem 3.20 implies we still have control on a subset of the Tits boundaries when $\mathrm{Out}(G(\Gamma))$ and $\mathrm{Out}(G(\Gamma'))$ are transvection-free. In this subsection, we will reorganize this piece of information in terms of extension complexes.

Recall that we identify the vertex set of Γ with a standard generating set S of $G(\Gamma)$. We also label the standard circles in the Salvetti complex by elements in S. By choosing an orientation for each standard circle, we obtain a directed labeling of edges in $X(\Gamma)$.

Denote the extension complex of Γ by $\mathcal{P}(\Gamma)$. We give an alternative definition of $\mathcal{P}(\Gamma)$ here, which is natural for our purposes. The vertices of $\mathcal{P}(\Gamma)$ are in one-to-one correspondence with the parallel classes of standard geodesics in $X(\Gamma)$ (two standard geodesics are in the same parallel class if they are parallel). Two distinct vertices $v_1, v_2 \in \mathcal{P}(\Gamma)$ are connected by an edge if and only if we can find standard geodesics l_i in the parallel classes associated with v_i ($i = 1, 2$) such that l_1 and l_2 span a standard 2–flat. The next observation follows from Lemmas 3.1 and 2.4:

Observation 4.1 *If $v_1 \neq v_2$, then v_1 and v_2 are joined by an edge if and only if there exist l_i' in the parallel classes associated with v_i ($i = 1, 2$) and $R > 0$ such that $l_1' \subset N_R(P_{l_2'})$.*

We define $\mathcal{P}(\Gamma)$ to be the flag complex of its 1–skeleton.

Lemma 4.2 *$\mathcal{P}(\Gamma)$ is isomorphic to the extension complex of Γ.*

Proof It suffices to show the 1–skeleton of $\mathcal{P}(\Gamma)$ is isomorphic to the extension graph Γ^e. Pick vertex $v \in \mathcal{P}(\Gamma)$, and let l be a standard geodesic in the parallel class associated with v. We identify l with $\mathbb{R}$ in an orientation-preserving way (the orientation in l is induced by the directed labeling). Recall that $G(\Gamma) \curvearrowright X(\Gamma)$ by deck transformations. Let $\alpha_v \in G(\Gamma)$ be the element such that $\alpha_v(l) = l$ and $\alpha_v(x) = x + 1$ for any $x \in l$. It is easy to see α_v is conjugate to an element in S; thus α_v gives rise to a vertex $\alpha_v \in \Gamma^e$ by Definition 2.11. Note that α_v does not depend the choice of l in the parallel class, so we have a well-defined map from the vertex set of $\mathcal{P}(\Gamma)$ to the vertex set of Γ^e. Moreover, if v_1 and v_2 are adjacent, then α_{v_1} and α_{v_2} commute.

Now we define an inverse map. Pick $\alpha = gsg^{-1} \in \Gamma^e$ ($s \in S$). Then all standard geodesics which are stabilized by α are in the same parallel class. Let v_α be the vertex in $\mathcal{P}(\Gamma)$ associated with this parallel class. We map the vertex α of Γ^e to the vertex v_α. Now we show this map extends to the 1–skeleton. For $i = 1, 2$, let $\alpha_i = g_i s_i g_i^{-1} \in \Gamma^e$. By the centralizer theorem of [58], α_1 and α_2 commute if and only if $[s_1, s_2] = 1$ and there exists $g \in G(\Gamma)$ such that $\alpha_i = gs_i g^{-1}$. Thus v_{α_1} and v_{α_2} are adjacent in $\mathcal{P}(\Gamma)$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Since every edge in the standard geodesics of the same parallel class has the same label, the labeling of the edges of $X(\Gamma)$ induces a labeling of the vertices of $\mathcal{P}(\Gamma)$. Moreover, since $G(\Gamma) \curvearrowright X(\Gamma)$ by label-preserving cubical isomorphisms, we obtain an induced action $G(\Gamma) \curvearrowright \mathcal{P}(\Gamma)$ by label-preserving simplicial isomorphisms. Moreover, the unique label-preserving map from the vertices of $\mathcal{P}(\Gamma)$ to the vertices of $F(\Gamma)$ extends to a simplicial map

$$(4\text{-}1) \qquad\qquad\qquad \pi\colon \mathcal{P}(\Gamma) \to F(\Gamma).$$

Pick an arbitrary vertex $p \in X(\Gamma)$; one can obtain a simplicial embedding i_p from the flag complex $F(\Gamma)$ of Γ to $\mathcal{P}(\Gamma)$ by considering the collection of standard geodesics passing through p. We will denote the image of i_p by $(F(\Gamma))_p$. Note that for each vertex $p \in X(\Gamma)$, the composition $\pi \circ i_p\colon F(\Gamma) \to F(\Gamma)$ is the identity map.

Pick a $(k-1)$–simplex in $\mathcal{P}(\Gamma)$ with vertex set $\{v_i\}_{i=1}^k$, and pick a standard geodesic l_i in the parallel class associated with v_i for each $1 \le i \le k$. Since $P_{l_i} \cap P_{l_j} \ne \varnothing$ for $1 \le i \ne j \le k$, by Lemma 2.2, $\bigcap_{i=1}^k P_{l_i} \ne \varnothing$. By Corollary 3.2 and Lemma 3.4, there exist standard geodesics $\{l_i'\}_{i=1}^k$ satisfying:

(1) l_i' is parallel to l_i for each i.

(2) The convex hull of $\{l_i'\}_{i=1}^k$ is a standard k–flat denoted by F_k.

(3) $\bigcap_{i=1}^k P_{l_i} = P_{F_k}$.

Thus we have a one-to-one correspondence between the $(k-1)$–simplexes of $\mathcal{P}(\Gamma)$ and parallel classes of standard k–flats in $X(\Gamma)$. In particular, maximal simplexes in $\mathcal{P}(\Gamma)$, namely those simplexes which are not properly contained in some larger simplexes of $\mathcal{P}(\Gamma)$, are in one-to-one correspondence with maximal standard flats in $X(\Gamma)$. For standard flat $F \subset X(\Gamma)$, we denote the simplex in $\mathcal{P}(\Gamma)$ associated with the parallel class containing F by $\Delta(F)$.

Observation 4.3 *Let Δ_1, Δ_2 be two simplexes in $\mathcal{P}(\Gamma)$ such that $\Delta = \Delta_1 \cap \Delta_2 \ne \varnothing$. For $i = 1, 2$, let $F_i \subset X(\Gamma)$ be a standard flat such that $\Delta(F_i) = \Delta_i$. Set $(F_1', F_2') = \mathcal{I}(F_1, F_2)$. Then $\Delta(F_1') = \Delta(F_2') = \Delta$.*

We define the *reduced Tits boundary*, denoted $\overline{\partial}_T(X(\Gamma))$, to be the subset of $\partial_T(X(\Gamma))$ which is the union of Tits boundaries of standard flats in $X(\Gamma)$. For a standard flat $F \subset X(\Gamma)$, we triangulate $\partial_T F$ into all-right spherical simplexes which are the Tits boundaries of orthant subcomplexes in F. Pick another standard flat $F' \subset X(\Gamma)$; then $\partial_T F \cap \partial_T F'$ is a subcomplex in both $\partial_T F$ and $\partial_T F'$ by Lemma 3.1 and Remark 2.5. Thus we can endow $\overline{\partial}_T(X(\Gamma))$ with the structure of an all-right spherical complex.

Now we look at the relation between $\overline{\partial}_T(X(\Gamma))$ and $\mathcal{P}(\Gamma)$. For each standard flat $F \subset X(\Gamma)$, we can associate $\partial_T F$ with $\Delta(F) \subset \mathcal{P}(\Gamma)$. This induces a surjective simplicial map $s \colon \overline{\partial}_T(X(\Gamma)) \to \mathcal{P}(\Gamma)$ (s can be defined by induction on dimension). Note that the inverse image of each simplex in $\mathcal{P}(\Gamma)$ under s is a sphere in $\overline{\partial}_T(X(\Gamma))$. Then one can construct $\overline{\partial}_T(X(\Gamma))$ from $\mathcal{P}(\Gamma)$ as follows. We start with a collection of the $\mathbb{S}^0$ which are in one-to-one correspondence to vertices of $\mathcal{P}(\Gamma)$ and form a join of n copies of the $\mathbb{S}^0$ if and only if the corresponding n vertices in $\mathcal{P}(\Gamma)$ span an $(n-1)$–simplex. In other words, $\overline{\partial}_T(X(\Gamma))$ is obtained by applying the spherical complex construction in the sense of [12, Definition 2.1.22] to $\mathcal{P}(\Gamma)$.

Let $K_1 \subset X(\Gamma)$ be a standard subcomplex. We define $\overline{\partial}_T(K_1)$ to be the union of Tits boundaries of standard flats in K_1. Note that $\overline{\partial}_T(K_1) = \overline{\partial}_T(X(\Gamma)) \cap \partial_T K_1$, and it descends to a subcomplex in $\mathcal{P}(\Gamma)$, which will be denoted by $\Delta(K_1)$.

Lemma 4.4 *Let K_1 and K_2 be two standard subcomplexes of $X(\Gamma)$. Put $(K_1', K_2') = \mathcal{I}(K_1, K_2)$. Then $\Delta(K_1') = \Delta(K_2') = \Delta(K_1) \cap \Delta(K_2)$.*

Proof By Remark 2.5, we have $\partial_T K_1' = \partial_T K_2' = \partial_T K_1 \cap \partial_T K_2$; hence $\overline{\partial}_T K_1' = \overline{\partial}_T K_2' = \overline{\partial}_T K_1 \cap \overline{\partial}_T K_2$ and $\Delta(K_1') = \Delta(K_2') = \Delta(K_1) \cap \Delta(K_2)$. $\qquad\square$

Now we study how the extension complexes behave under quasi-isometries.

Lemma 4.5 *Pick Γ_1 and Γ_2 such that $\mathrm{Out}(G(\Gamma_i))$ is transvection-free for $i = 1, 2$. Then any quasi-isometry $q \colon X(\Gamma_1) \to X(\Gamma_2)$ induces a simplicial isomorphism $q_* \colon \mathcal{P}(\Gamma_1) \to \mathcal{P}(\Gamma_2)$. If only $\mathrm{Out}(G(\Gamma_1))$ is assumed to be transvection-free, we still have a simplicial embedding $q_* \colon \mathcal{P}(\Gamma_1) \to \mathcal{P}(\Gamma_2)$.*

Proof We prove the case when both $\mathrm{Out}(G(\Gamma_1))$ and $\mathrm{Out}(G(\Gamma_2))$ are transvection-free. The other case is similar. By Theorem 3.20, every vertex in Γ_1 is stable; thus q sends any parallel class of standard geodesics in $X(\Gamma_1)$ to another parallel class of standard geodesics in $X(\Gamma_2)$ up to finite Hausdorff distance. This induces a well-defined map q_* from the 0–skeleton of $\mathcal{P}(\Gamma_1)$ to the 0–skeleton of $\mathcal{P}(\Gamma_2)$. The map q_* is a bijection by considering the quasi-isometry inverse. Moreover, Observation 4.1

implies two vertices in $\mathcal{P}(\Gamma_1)$ are adjacent if and only if their images under q_* are adjacent. So we can extend q_* to be a graph isomorphism between the 1–skeleton of $\mathcal{P}(\Gamma_1)$ and the 1–skeleton of $\mathcal{P}(\Gamma_2)$. Since both $\mathcal{P}(\Gamma_1)$ and $\mathcal{P}(\Gamma_2)$ are flag complexes, we can extend the isomorphism to the whole complex. □

4.1.2 Extension complexes and their relatives Now we discuss the relation between $\mathcal{P}(\Gamma)$ with several other objects in the literature. The material in this subsection will not be used later.

We can endow $F(\Gamma)$ with the structure of complex of groups, which gives us an alternative definition of $\mathcal{P}(\Gamma)$. More specifically, $\mathcal{P}(\Gamma) = F(\Gamma) \times G(\Gamma)/\sim$; here $\mathrm{St}(v) \times g_1$ and $\mathrm{St}(v) \times g_2$ ($v \in F(\Gamma)$ is a vertex) are identified if and only if there exists an integer m such that $g_1^{-1} g_2 = v^m$ (we also view v as one of the generators of $G(\Gamma)$). Hence for k–simplex $\Delta^k \subset F(\Gamma)$ with vertex set $\{v_i\}_{i=1}^k$, we have that $\mathrm{St}(\Delta^k) \times g_1$ and $\mathrm{St}(\Delta^k) \times g_2$ are identified if and only if $g_1^{-1} g_2$ belongs to the $\mathbb{Z}^k$ subgroup of $G(\Gamma)$ generated by $\{v_i\}_{i=1}^k$. One can compare this with a similar construction for a Coxeter group in [21].

There is another important object associated with a right-angled Artin group, called the modified Deligne complex in [17] and the *flat space* in [9].

Definition 4.6 Let $\mathbb{P}(\Gamma)$ be poset of left cosets of standard abelian subgroups of $G(\Gamma)$ (including the trivial subgroup) such that the partial order is induced by inclusion of sets. Then the *modified Deligne complex* is defined to be the geometric realization of the derived poset of $\mathbb{P}(\Gamma)$.

Recall that elements in the *derived poset* of a poset $\mathbb{P}$ are totally ordered finite chains in $\mathbb{P}$. It can be viewed as an abstract simplex.

The extension complex $\mathcal{P}(\Gamma)$ can be viewed as a coarse version of the modified Deligne complex. Let A and B be two subsets of a metric space. We say A and B are *coarsely equivalent* if $A \stackrel{\infty}{=} B$, and A is *coarsely contained* in B if $A \subset_\infty B$. Let $\mathbb{P}'(\Gamma)$ be the poset whose elements are coarsely equivalent classes of left cosets of nontrivial standard abelian subgroups of $G(\Gamma)$, and the partial order is induced by coarse inclusion of sets.

Observation 4.7 *The poset $\mathbb{P}'(\Gamma)$ is an abstract simplicial complex, and it is isomorphic to $\mathcal{P}(\Gamma)$.*

Roughly speaking, $\mathbb{P}(\Gamma)$ captures the combinatorial pattern of how standard flats in $X(\Gamma)$ intersect with each other, and $\mathcal{P}(\Gamma)$ is about how they coarsely intersect with each other; thus $\mathbb{P}(\Gamma)$ contains more information than $\mathcal{P}(\Gamma)$. However, in certain cases, it is possible to recover information about $\mathbb{P}(\Gamma)$ from $\mathcal{P}(\Gamma)$, and this enable us to prove quasi-isometric classification/rigidity results for RAAGs.

We can define the poset $\mathbb{P}'(\Gamma)$ for an arbitrary Artin group by considering the collection of coarse equivalent classes of spherical subgroups in an Artin group under coarse inclusion. Then the geometric realization of the derived poset of $\mathbb{P}'(\Gamma)$ would be a natural candidate for the extension complex of an Artin group. It is interesting to ask how much of the results in [43] can be generalized to this context.

There is also a link between $\mathcal{P}(\Gamma)$ and the structure of hyperplanes in $X(\Gamma)$. Recall that for every CAT(0) cube complex X, the *crossing graph* of X, denoted by $C(X)$, is a graph whose vertices are in one-to-one correspondence to the hyperplanes in X, and two vertices are adjacent if and only if the corresponding hyperplanes intersect. The *contact graph*, introduced in [30] and denoted by $\mathcal{C}(X)$, has the same vertex set as $C(X)$, and two vertices are adjacent if and only if the carriers of the corresponding hyperplanes intersect.

There is a natural surjective simplicial map $p: C(X(\Gamma)) \to \Gamma^e$ defined as follows. Pick a vertex $v \in C(X(\Gamma))$ and let h be the corresponding hyperplane. Since all standard geodesics which intersect h at one point are in the same parallel class, we define $p(v)$ to be the vertex in Γ^e associated with this parallel class; see Lemma 4.2. It is clear that if $v_1, v_2 \in C(X(\Gamma))$ are adjacent vertices, then $p(v_1)$ and $p(v_2)$ are adjacent, so p extends to a simplicial map. Pick a vertex $w \in \Gamma^e$; then $p^{-1}(e)$ is the collection of hyperplanes dual to a standard geodesic.

Theorem 4.8 [42; 30] *If Γ is connected, then $C(X(\Gamma))$, $\mathcal{C}(X(\Gamma))$ and $\mathcal{P}(\Gamma)$ are quasi-isometric to each other; moreover, they are quasi-isometric to a tree.*

From this viewpoint, $\mathcal{P}(\Gamma)$ captures both the geometric information of $X(\Gamma)$ (the standard flats) and the combinatorial information (the hyperplanes).

4.2 Reconstruction of quasi-isometries

We show the boundary map $q_*: \mathcal{P}(\Gamma) \to \mathcal{P}(\Gamma')$ in Lemma 4.5 induces a well-defined map from $G(\Gamma)$ to $G(\Gamma')$.

Lemma 4.9 *Let F_1 and F_2 be two maximal standard flats in $X(\Gamma)$ and let Δ_1 and Δ_2 be the corresponding maximal simplexes in $\mathcal{P}(\Gamma)$. If F_1 and F_2 are separated by a hyperplane h, then there exist vertices $v_i \in \Delta_i$ for $i = 1, 2$ and $v \in \mathcal{P}(\Gamma)$ such that v_1 and v_2 are in different connected components of $\mathcal{P}(\Gamma) \setminus \mathrm{St}(v)$.*

Proof Let e be an edge dual to h and let l be the standard geodesic that contains e. Set $v = \Delta(l) \in \mathcal{P}(\Gamma)$. By Lemma 3.4, the parallel set P_l of l is isometric to $h \times \mathbb{E}^1$. Thus every standard geodesic parallel to l must have nontrivial intersection with h. Since $F_1 \cap h = \varnothing$, we see that F_1 cannot contain any standard geodesic parallel to l,

which means $v \notin \Delta_1$. Moreover, $\Delta_1 \nsubseteq \mathrm{St}(v)$ since Δ_1 is a maximal simplex. Similarly, $\Delta_2 \nsubseteq \mathrm{St}(v)$; thus we can find vertices $v_i \in \Delta_i \setminus \mathrm{St}(v)$ for $i = 1, 2$. We claim v_1, v_2 and v are the vertices we are looking for.

If there is a path $\omega \subset \mathcal{P}(\Gamma) \setminus \mathrm{St}(v)$ connecting v_1 and v_2, we can assume ω consists of a sequence of edges $\{e_i\}_{i=1}^k$ with $v_1 \in e_1$ and $v_2 \in e_k$. For each e_i, pick a maximal simplex Δ_i' that contains e_i, and let F_i' be the maximal standard flat such that $\Delta(F_i') = \Delta_i'$. Then $v \notin \Delta_i'$ for $1 \le i \le k$; hence $F_i' \cap h = \varnothing$.

Set $\Delta_0' = \Delta_1$, $\Delta_{k+1}' = \Delta_2$, $F_0' = F_1$ and $F_{k+1}' = F_2$. Since $\Delta_i' \cap \Delta_{i+1}'$ contains a vertex in ω, we have

(4-2)
$$(\Delta_i' \cap \Delta_{i+1}') \setminus \mathrm{St}(v) \neq \varnothing$$

for $0 \le i \le k$. Since F_0' and F_{k+1}' are in different sides of h, there exists i_0 such that h separates F_{i_0}' and F_{i_0+1}'. Let $(F_{i_0}'', F_{i_0+1}'') = \mathcal{I}(F_{i_0}', F_{i_0+1}')$. By Observation 4.3, $\Delta(F_{i_0}'') = \Delta(F_{i_0+1}'') = \Delta_i' \cap \Delta_{i+1}'$. However, by Lemma 2.7, there exists a convex subset of h parallel to F_{i_0}''; thus $F_{i_0}'' \subset_\infty h \subset P_l$. It follows from Observation 4.1 that $\Delta_i' \cap \Delta_{i+1}' \subset \mathrm{St}(v)$, which contradicts (4-2). □

Denote the Cayley graph of $G(\Gamma)$ with respect to the standard generating set S by $C(\Gamma)$. We pick an identification between $C(\Gamma)$ and the 1–skeleton of $X(\Gamma)$. Thus $G(\Gamma)$ is identified with the vertex set of $X(\Gamma)$.

Lemma 4.10 *Let Γ_1 be a simple graph such that:*

(1) *There is no separating closed star in $F(\Gamma_1)$.*

(2) *$F(\Gamma_1)$ is not contained in a union of two closed stars.*

Then any simplicial isomorphism $s : \mathcal{P}(\Gamma_1) \to \mathcal{P}(\Gamma_2)$ induces a unique map $s' : G(\Gamma_1) \to G(\Gamma_2)$ such that for any maximal standard flat $F_1 \subset X(\Gamma_1)$, vertices in F_1 are mapped by s' to vertices lying in a maximal standard flat $F_2 \subset X(\Gamma_2)$ with $\Delta(F_2) = s'(\Delta(F_1))$.

Proof Pick a vertex $p \in G(\Gamma_1)$. Let $\{F_i\}_{i=1}^k$ be the collection of maximal standard flats containing p. For $1 \le i \le k$, define $\Delta_i = \Delta(F_i)$ and $\Delta_i' = s(\Delta_i)$. Let $F_i' \subset X(\Gamma_2)$ be the maximal standard flat such that $\Delta(F_i') = \Delta_i'$. Let $K_p = (F(\Gamma_1))_p = \bigcup_{i=1}^k \Delta_i$ (recall that $K_p \cong F(\Gamma_1)$). We claim

(4-3)
$$\bigcap_{i=1}^k F_i' \neq \varnothing.$$

The lemma will then follow from (4-3). To see this, we deduce from condition (2) that $\bigcap_{i=1}^k \Delta_i = \varnothing$. Hence $\bigcap_{i=1}^k F_i = \{p\}$. It follows that $\bigcap_{i=1}^k \Delta_i' = \varnothing$. This together

with (4-3) imply that $\bigcap_{i=1}^{k} F_i'$ is exactly one point. We define s' by sending p to this point. One readily verifies that s' has the required properties.

It remains to prove (4-3).

Suppose that (4-3) is not true. Then by Lemma 2.2, there exist i_1 and i_2 such that $F_{i_1}' \cap F_{i_2}' = \varnothing$. Thus F_{i_1}' and F_{i_2}' are separated by a hyperplane. It follows from Lemma 4.9 that there exist vertices $v' \in \mathcal{P}(\Gamma_2)$, $v_1' \in \Delta_{i_1}'$ and $v_2' \in \Delta_{i_2}'$ such that v_1' and v_2' are in different connected components of $\mathcal{P}(\Gamma_2) \setminus \mathrm{St}(v')$. Let $v = s^{-1}(v')$, $v_1 = s^{-1}(v_1')$ and $v_2 = s^{-1}(v_2')$. Then $K_p \setminus (K_p \cap \mathrm{St}(v))$ is disconnected (since $v_1, v_2 \in K_p$ and they are separated by $\mathrm{St}(v)$).

If $v \in K_p$, then K_p would contain a separating closed star, which yields a contradiction; thus (4-3) is true in this case.

Suppose $v \notin K_p$. Pick a standard geodesic l such that $\Delta(l) = v$ and let $\{h_i\}_{i=1}^{n}$ be the collection of hyperplanes in $X(\Gamma)$ such that each h_i separates p from the parallel set P_l of l (note that $p \notin P_l$). For $1 \le i \le n$, pick an edge e_i dual to h_i and let w_i be the unique vertex in K_p that has the same label as e_i. Let $w_0 \in K_p$ be the unique vertex which has the same label as v. We claim

$$\text{(4-4)} \qquad \mathrm{St}(v) \cap K_p = \bigcap_{i=0}^{n} (\mathrm{St}(w_i) \cap K_p).$$

For every $u \in K_p$, let l_u be the unique standard geodesic such that $\Delta(l_u) = u$ and $p \in l_u$.

Pick $u \in \mathrm{St}(v) \cap K_p$. Observation 4.1 implies $\mathcal{I}(l_u, P_l) = (l_u, l_u')$, where l_u' is some standard geodesic in P_l. Then for $1 \le i \le n$, the hyperplane h_i separates l_u from P_l, otherwise $h_i \cap l_u \ne \varnothing$ and Lemma 2.6 implies $h_i \cap P_l \ne \varnothing$, which is a contradiction. It follows from Corollary 3.2 that u and w_i are adjacent for $0 \le i \le n$; thus $u \in \bigcap_{i=0}^{n} (\mathrm{St}(w_i) \cap K_p)$. Therefore, $\mathrm{St}(v) \cap K_p \subset \bigcap_{i=0}^{n} (\mathrm{St}(w_i) \cap K_p)$.

Pick $u \in \bigcap_{i=0}^{n} (\mathrm{St}(w_i) \cap K_p)$. First we show $l_u \cap P_l = \varnothing$. Suppose there is a vertex z in $l_u \cap P_l$. Since v and w_0 have the same label and $u \in \mathrm{St}(w_0)$, it follows that the edge in l_u which contains z belongs to the parallel set P_l. Then $l_u \subset P_l$, contradicting the fact that $p \notin P_l$. Therefore, $l_u \cap P_l = \varnothing$.

Now we pick an edge path ω of shortest combinatorial length that travels from l_u to P_l. Let $\{f_j\}_{j=1}^{m}$ be the consecutive edges in ω such that $f_1 \cap l_u \ne \varnothing$. For each f_j, let $\bar{h}_j$ be the hyperplane dual to f_j. Then $\bar{h}_j$ separates l_u from P_l (otherwise ω would not be the shortest edge path), hence separates p from P_l. This and $u \in \bigcap_{i=0}^{n} (\mathrm{St}(w_i) \cap K_p)$ imply that $d(\pi(u), V_{f_j}) \le 1$ for each j, where π is the map in (4-1) and V_{f_j} is the

label of the edge f_j. It follows that ω is contained in the parallel set P_{l_u}, and hence the intersection $P_{l_u} \cap P_l$ contains some vertex z. Again, since $u \in \mathrm{St}(w_0)$, and since w_0 has the same label as v, we find that the standard geodesic $l'_u \subset P_{l_u}$ that is parallel to l_u and passes through z is contained in P_l. Therefore, $u \in \mathrm{St}(v) \cap K_p$, and (4-4) follows.

By condition (2) of Lemma 4.10, we have

$$(4\text{-}5) \qquad (\mathrm{St}(w_0) \cap K_p) \cup \left(\bigcap_{i=1}^{n} (\mathrm{St}(w_i) \cap K_p) \right) \subsetneqq K_p.$$

Let $A = K_p \setminus (\mathrm{St}(w_0) \cap K_p)$, and let $B = K_p \setminus (\bigcap_{i=1}^{n} (\mathrm{St}(w_i) \cap K_p))$. Then (4-5) implies $A \cap B \neq \varnothing$. Thus we have the following Mayer–Vietoris sequence for reduced homology:

$$\cdots \to \widetilde{H}_0(A \cap B) \to \widetilde{H}_0(A) \oplus \widetilde{H}_0(B) \to \widetilde{H}_0(A \cup B) \to 0.$$

Recall that $K_p \setminus (K_p \cap \mathrm{St}(v))$ is disconnected, we deduce that $\widetilde{H}_0(A \cup B)$ is nontrivial from (4-4). Thus $\widetilde{H}_0(A) \oplus \widetilde{H}_0(B)$ is nontrivial, which implies that either $\bigcap_{i=1}^{n} (\mathrm{St}(w_i) \cap K_p)$ or $\mathrm{St}(w_0) \cap K_p$ would separate K_p. Thus we can induct on n to deduce that there exists i_0 such that $\mathrm{St}(w_{i_0}) \cap K_p$ separates K_p. This yields a contradiction to condition (1) of Lemma 4.10. $\qquad \square$

There are counterexamples if we only assume (1) in Lemma 4.10. For example, let Γ_1 and Γ_2 be discrete graphs made of two points. Then $\mathcal{P}(\Gamma_1)$ and $\mathcal{P}(\Gamma_2)$ are discrete sets. Now it is not hard to construct a permutation of a discrete set to itself which does not satisfy the conclusion of Lemma 4.10. If we go back to the proof of Lemma 4.10, then the step using the Mayer–Vietoris sequence will fail, since we need $A \cap B \neq \varnothing$ in order to use the reduced version of Mayer–Vietoris sequence.

Corollary 4.11 *Suppose that $G(\Gamma_1)$ and $G(\Gamma_2)$ both satisfy the assumption of Lemma 4.10. Then they are isomorphic if and only if $\mathcal{P}(\Gamma_1)$ and $\mathcal{P}(\Gamma_2)$ are isomorphic as simplicial complexes.*

Proof The "only if" direction follows from the fact that $G(\Gamma_1)$ and $G(\Gamma_2)$ are isomorphic if and only if Γ_1 and Γ_2 are isomorphic; see [23]. It remains to prove the "if" direction. Pick an isomorphism $s\colon \mathcal{P}(\Gamma_1) \to \mathcal{P}(\Gamma_2)$, and let $s'\colon G(\Gamma_1) \to G(\Gamma_2)$ be the map in Lemma 4.10. Pick a vertex $p \in G(\Gamma_1)$ and let $q = s(p)$. We define $(F(\Gamma_1))_p \subset \mathcal{P}(\Gamma_1)$ and $(F(\Gamma_2))_q \subset \mathcal{P}(\Gamma_2)$ as in the first paragraph of the proof of Lemma 4.10. Then (4-3) implies $s((F(\Gamma_1))_p) \subset (F(\Gamma_2))_q$. This induces a graph embedding $\Gamma_1 \hookrightarrow \Gamma_2$. By repeating the previous discussion for s^{-1}, we obtain another graph embedding $\Gamma_2 \hookrightarrow \Gamma_1$. Since both Γ_1 and Γ_2 are finite simplicial graphs, they are isomorphic. Hence $G(\Gamma_1) \cong G(\Gamma_2)$. $\qquad \square$

Lemma 4.12 *Let $G(\Gamma)$ be a RAAG such that $\text{Out}(G(\Gamma))$ is finite and $G(\Gamma) \not\cong \mathbb{Z}$. Then $F(\Gamma)$ satisfies the assumption of Lemma 4.10.*

Proof It is clear that $F(\Gamma)$ should satisfy condition (1) of Lemma 4.10 since no nontrivial partial conjugation is allowed. If $F(\Gamma)$ is contained in a closed star, then Γ is a point. So if (2) is not true, then $F(\Gamma) = \text{St}(v) \cup \text{St}(w)$ for distinct vertices $v, w \in \Gamma$. Since the orthogonal complement $v^\perp$ satisfies $v^\perp \not\subseteq \text{St}(w)$, there exists $u \in v^\perp$ such that $d(u, w) \geq 2$. Pick any edge e such that $u \in e$; then $e \not\subseteq \text{St}(w)$, and so $e \subset \text{St}(v)$. This implies $u^\perp \subset \text{St}(v)$; hence $\text{Out}(G(\Gamma))$ is infinite, which yields a contradiction. $\square$

By Lemma 4.5, Lemma 4.12 and Corollary 4.11, we have following result, which in particular establishes Theorem 1.1 of the introduction.

Theorem 4.13 *Let Γ_1 and Γ_2 be two finite simplicial graphs such that $\text{Out}(G(\Gamma_i))$ is finite for $i = 1, 2$. Then $G(\Gamma_1)$ and $G(\Gamma_2)$ are quasi-isometric if and only if they are isomorphic. Moreover, for any (L, A)–quasi-isometry $q\colon X(\Gamma_1) \to X(\Gamma_2)$, there exist a bijection $q'\colon G(\Gamma_1) \to G(\Gamma_2)$ and a constant $D = D(L, A, \Gamma_1)$ such that:*

(1) *$d(q(v), q'(v)) < D$ for any $v \in G(\Gamma_1)$.*

(2) *For any standard flat $F_1 \subset X(\Gamma_1)$, there exists a standard flat $F_2 \subset X(\Gamma_2)$ such that q' induces a bijection between $F_1 \cap G(\Gamma_1)$ and $F_2 \cap G(\Gamma_2)$.*

If $G(\Gamma_1) \neq \mathbb{Z}$, then such a q' is unique.

Proof It suffices to look at the case where $G(\Gamma_1) \neq \mathbb{Z}$. Then $G(\Gamma_2) \neq \mathbb{Z}$. In this case, every vertex v in Γ_1 or Γ_2 is the intersection of maximal cliques that contain v (otherwise there exists a vertex w such that $w \neq v$ and $v^\perp \subset \text{St}(w)$). It follows that every standard geodesic in $X(\Gamma_1)$ or $X(\Gamma_2)$ is the intersection of finitely many maximal standard flats, and so is every standard flat. Let $q_*\colon \mathcal{P}(\Gamma_1) \to \mathcal{P}(\Gamma_2)$ be the map in Lemma 4.5. We apply Lemma 4.10 to q_* and q_*^{-1} to obtain q' with the required properties. Note that each vertex of $X(\Gamma)$ is the intersection of maximal standard flats that contain it; thus q' is unique. $\square$

4.3 The automorphism groups of extension complexes

Suppose $\text{Out}(G(\Gamma))$ is finite; by Theorem 4.13, each element in the simplicial automorphism group $\text{Aut}(\mathcal{P}(\Gamma))$ of $\mathcal{P}(\Gamma)$ induces a bijection $G(\Gamma) \to G(\Gamma)$. However, this bijection does not extend to an isomorphism from $X(\Gamma)$ to itself in general. We start by looking at the following example which was first pointed out in [9, Section 11] in a slightly different form.

Example 4.14 Let $l \subset X(\Gamma)$ be a standard geodesic, and let $\pi_l\colon X(\Gamma) \to l$ be the CAT(0) projection. We identify the vertex set of l with $\mathbb{Z}$. Let $X^{(0)}(\Gamma)$ be the vertex set of $X(\Gamma)$. Then the above projection induces a map $\pi_l\colon X^{(0)}(\Gamma) \to \mathbb{Z}$.

Recall that each edge of $X(\Gamma)$ is oriented and labeled, and $G(\Gamma)$ acts on $X(\Gamma)$ by transformations that preserve labels and orientations. There is a unique element $\alpha \in G(\Gamma)$ such that α translates l one unit in the positive direction.

We want to define a bijection $q\colon X^{(0)}(\Gamma) \to X^{(0)}(\Gamma)$ which basically flips $\pi_l^{-1}(0)$ and $\pi_l^{-1}(1)$. More precisely,

$$q(x) = \begin{cases} x & \text{if } \pi_l(x) \neq 0, 1, \\ \alpha(x) & \text{if } x \in \pi_l^{-1}(0), \\ \alpha^{-1}(x) & \text{if } x \in \pi_l^{-1}(1). \end{cases}$$

One can check the following:

(1) q is a quasi-isometry.

(2) q does not respect the word metric.

(3) q maps vertices in a standard flat to vertices in another standard flat. Thus q induces an element in $\text{Aut}(\mathcal{P}(\Gamma))$.

The above example implies that, in general, elements in $\text{Aut}(\mathcal{P}(\Gamma))$ do not respect the order along the standard geodesics of $X(\Gamma)$. There is another metric on $G(\Gamma)$ which "forgets about" the ordering. Following [43], we define the *syllable length* of a word ω to be the minimal l such that ω can be written as a product of l elements of the form $v_i^{k_i}$, where v_i is a standard generator and k_i is an integer.

An alternative definition is the following. Let $\{h_i\}_{i=1}^k$ be the collection of hyperplanes separating $\omega \in G(\Gamma)$ and the identity element (recall that we have identified $G(\Gamma)$ with the 0–skeleton of $X(\Gamma)$). For each i, pick a standard geodesic l_i dual to h_i. Then the syllable length of ω is the number of elements in $\{\Delta(l_i)\}_{i=1}^k$. The syllable length induces a left invariant metric on $G(\Gamma)$, which will be denoted by d_r. Note that the map in Example 4.14 is an isometry with respect to d_r.

Denote the word metric on $G(\Gamma)$ with respect to the standard generators by d_w.

Corollary 4.15 *Let Γ be a graph such that $\text{Out}(G(\Gamma))$ is finite, and denote the simplicial automorphism group of $\mathcal{P}(\Gamma)$ by $\text{Aut}(\mathcal{P}(\Gamma))$. Then*

$$\text{Aut}(\mathcal{P}(\Gamma)) \cong \text{Isom}(G(\Gamma), d_r).$$

Proof Let $\text{Perm}(G(\Gamma))$ be the permutation group of elements in $G(\Gamma)$. We have a group homomorphism $h_1\colon \text{Aut}(\mathcal{P}(\Gamma)) \to \text{Perm}(G(\Gamma))$ by Lemma 4.10. Take $\phi \in \text{Aut}(\mathcal{P}(\Gamma))$; by Lemma 4.12, $\varphi = h_1(\phi)$ and $\varphi^{-1} = h_1(\phi^{-1})$ satisfy the conclusion

of Lemma 4.10. Since every standard geodesic is the intersection of finitely many maximal standard flats, points in a standard geodesic are mapped to points in a standard geodesic by ϕ, which implies $d_r(\varphi(v_1), \varphi(v_2)) \le d_r(v_1, v_2)$ if $d_r(v_1, v_2) \le 1$. By the triangle inequality, we have $d_r(\varphi(v_1), \varphi(v_2)) \le d_r(v_1, v_2)$ for any $v_1, v_2 \in G(\Gamma)$. Similarly, $d_r(\varphi^{-1}(v_1), \varphi^{-1}(v_2)) \le d_r(v_1, v_2)$. Thus $\varphi \in \mathrm{Isom}(G(\Gamma), d_r)$, and we have a homomorphism $h_1 \colon \mathrm{Aut}(\mathcal{P}(\Gamma)) \to \mathrm{Isom}(G(\Gamma), d_r)$.

Now pick $\varphi \in \mathrm{Isom}(G(\Gamma), d_r)$. Let $v_1, v_2, v_3 \in G(\Gamma)$ such that $d_r(v_1, v_i) = 1$ for $i = 2, 3$. We claim

$$(4\text{-}6) \qquad \angle_{v_1}(v_2, v_3) = \tfrac{\pi}{2} \quad \Longleftrightarrow \quad \angle_{\varphi(v_1)}(\varphi(v_2), \varphi(v_3)) = \tfrac{\pi}{2}.$$

If $\angle_{v_1}(v_2, v_3) = \tfrac{\pi}{2}$, then we can find $v_4 \in G(\Gamma)$ such that $\{v_i\}_{i=1}^4$ are the vertices of a flat rectangle in $X(\Gamma)$. Note that

$$d_r(v_1, v_4) = d_r(v_2, v_3) = 2 \quad \text{and} \quad d_r(v_4, v_2) = d_r(v_4, v_3) = 1,$$

so

$$d_r(\varphi(v_1), \varphi(v_4)) = d_r(\varphi(v_2), \varphi(v_3)) = 2 \quad \text{and} \quad d_r(\varphi(v_4), \varphi(v_2)) = d_r(\varphi(v_4), \varphi(v_3)) = 1.$$

Now we consider the 4–gon formed by $\overline{\varphi(v_1)\varphi(v_2)}$, $\overline{\varphi(v_2)\varphi(v_4)}$, $\overline{\varphi(v_4)\varphi(v_3)}$ and $\overline{\varphi(v_3)\varphi(v_1)}$. Then the angles at the four vertices of this 4–gon are bigger or equal to $\tfrac{\pi}{2}$. It follows from CAT(0) geometry that the angles are exactly $\tfrac{\pi}{2}$ and the 4–gon actually bounds a flat rectangle. Thus one direction of (4-6) is proved; the other direction is similar.

We need another observation as follows. If three points $v_1, v_2, v_3 \in G(\Gamma)$ satisfy $d_r(v_i, v_j) = 1$ for $1 \le i \ne j \le 3$, then the angle at each vertex of the triangle $\Delta(v_1, v_2, v_3)$ could only be 0 or π; thus $\{v_i\}_{i=1}^3$ are inside a standard geodesic. It follows from this observation that points in a standard geodesic are mapped by φ to points in a standard geodesic.

We define $\phi \colon \mathcal{P}(\Gamma) \to \mathcal{P}(\Gamma)$ as follows. For vertex $w \in \mathcal{P}(\Gamma)$, let l be a standard geodesic such that $\Delta(l) = w$. Suppose $l' \subset X(\Gamma)$ is the standard geodesic such that $\phi(v(l)) \subset l'$ ($v(l)$ denotes the vertex set of l). Suppose $w' = \Delta(l')$. We define $w' = \phi(w)$; (4-6) implies w' does not depend on the choice of l, and $\phi(w_1)$ and $\phi(w_2)$ are adjacent if vertices $w_1, w_2 \in \mathcal{P}(\Gamma)$ are adjacent. Thus ϕ is a well-defined simplicial map. Note that φ^{-1} also induces a simplicial map from $\mathcal{P}(\Gamma)$ to itself in a similar way, so $\phi \in \mathrm{Aut}(\mathcal{P}(\Gamma))$. We define $\phi = h_2(\varphi)$. One readily verifies that $h_2 \colon \mathrm{Isom}(G(\Gamma), d_r) \to \mathrm{Aut}(\mathcal{P}(\Gamma))$ is a group homomorphism, and $h_2 \circ h_1 = h_1 \circ h_2 = \mathrm{Id}$. Thus the corollary follows. $\qquad\square$

Remark 4.16 If we drop the assumption in the above corollary about Γ, then there is still a monomorphism $h\colon \mathrm{Isom}(G(\Gamma), d_r) \to \mathrm{Aut}(\mathcal{P}(\Gamma))$; moreover, any $\varphi \in \mathrm{Isom}(G(\Gamma), d_r)$ maps vertices in a standard flat to vertices in a standard flat of the same dimension. The homomorphism h is surjective if $\mathrm{Out}(G(\Gamma))$ is finite.

Remark 4.17 For any finite simplicial graphs Γ_1 and Γ_2, we have $G(\Gamma_1) \cong G(\Gamma_2)$ if and only if $(G(\Gamma_1), d_r)$ and $(G(\Gamma_2), d_r)$ are isometric as metric spaces. The "only if" direction follows from [23; 49]. For the other direction, let $\varphi\colon (G(\Gamma_1), d_r) \to (G(\Gamma_2), d_r)$ be an isometry. Pick $v \in G(\Gamma_1)$, and let $\{l_i\}_{i=1}^k$ be the collection of standard geodesics passing through v. Pick $v_i \in G(\Gamma_1)$ such that $v_i \in l_i \setminus \{v\}$. Then $d_r(v, v_i) = 1$ for $1 \le i \le k$, and $d_r(v_i, v_j) = 2$ for $1 \le i \ne j \le k$. So $d_r(\varphi(v), \varphi(v_i)) = 1$ for $1 \le i \le k$, and $d_r(\varphi(v_i), \varphi(v_j)) = 2$ for $1 \le i \ne j \le k$, and $\angle_v(v_i, v_j) = \frac{\pi}{2}$ if and only if $\angle_{\varphi(v)}(\varphi(v_i), \varphi(v_j)) = \frac{\pi}{2}$ by (4-6). This induces a graph embedding $\Gamma_1 \to \Gamma_2$. By considering φ^{-1}, we obtain another graph embedding $\Gamma_2 \to \Gamma_1$. Hence Γ_1 and Γ_2 are isomorphic.

Corollary 4.18 *If* $\mathrm{Out}(G(\Gamma))$ *is finite and* $\mathrm{QI}(G(\Gamma))$ *is the quasi-isometry group of* $G(\Gamma)$*, then we have the following commutative diagram, where* i_1*,* i_2 *and* i_3 *are injective homomorphisms:*

$$\mathrm{Isom}(G(\Gamma), d_w) \xrightarrow{\;i_1\;} \mathrm{QI}(G(\Gamma)) \xrightarrow{\;i_2\;} \mathrm{Isom}(G(\Gamma), d_r).$$
$$\underbrace{\phantom{\mathrm{Isom}(G(\Gamma), d_w) \xrightarrow{\;i_1\;} \mathrm{QI}(G(\Gamma)) \xrightarrow{\;i_2\;} \mathrm{Isom}(G(\Gamma), d_r)}}_{i_3}$$

Proof The homomorphisms i_1 and i_3 are obvious, and i_2 is given by Lemma 4.5 and Corollary 4.15. It is clear that i_2 is a group homomorphism and $i_3 = i_2 \circ i_1$. Note that i_3 is injective, so i_1 is injective. Pick $\alpha \in \mathrm{QI}(G(\Gamma))$; by Corollary 4.15, we know $i_2(\alpha) = \mathrm{Id}$ implies the image of every standard flat under α is uniformly Hausdorff close to itself; thus α is of bounded distance from the identity map. $\qquad\square$

5 Quasi-isometries and special subgroups

Let $G(\Gamma)$ be a RAAG with finite outer automorphism group. In this section, we characterize all other RAAGs quasi-isometric to $G(\Gamma)$.

5.1 Preservation of extension complex

Lemma 5.1 *Let* Γ *be a finite simplicial graph. Pick a vertex* $w \in \Gamma$*, and let* Γ_w *be the minimal stable subgraph containing* w*. Denote* $\Gamma_1 = \mathrm{lk}(w)$ *and* $\Gamma_2 = \mathrm{lk}(\Gamma_1)$ (*see Section 2.1 for the definition of links*)*. Then exactly one of the following is true:*

(1) Γ_w *is a clique. In this case,* $\mathrm{St}(w)$ *is a stable subgraph.*

(2) *Both* Γ_1 *and* $\Gamma_1 \circ \Gamma_2$ *are stable subgraphs of* Γ*. Moreover,* Γ_2 *is disconnected.*

Recall that we use $(\Gamma')^{\perp}$ to denote the orthogonal complement of the subgraph $\Gamma' \subset \Gamma$ (see Section 2.1), and we assume $(\varnothing)^{\perp} = \Gamma$.

Proof If $\Gamma_w \subset \mathrm{St}(w)$, then Γ_w is a clique by Lemma 3.23. We also deduce from Lemma 3.23 that each vertex of $\mathrm{St}(w) \setminus \Gamma_w$ is in $\Gamma_w^{\perp}$. Moreover, $\Gamma_w^{\perp} \subset w^{\perp}$ since $w \in \Gamma_w$. Thus $\mathrm{St}(w)$ is the full subgraph spanned by vertices in Γ_w and $\Gamma_w^{\perp}$. So $\mathrm{St}(w)$ is stable by Lemma 3.12.

If $\Gamma_w \not\subset \mathrm{St}(w)$, let Γ_{11} be the full subgraph spanned by vertices in $\Gamma_w \cap \mathrm{lk}(w)$, and let Γ_2' be the full subgraph spanned by vertices in $\Gamma_w \setminus \Gamma_{11}$. By Lemma 3.23, $\Gamma_w = \Gamma_{11} \circ \Gamma_2'$ and $\Gamma_2' = \Gamma_2$. Note that Γ_2 is disconnected with isolated point $w \in \Gamma_2$, and Γ_{11} may be empty.

Let $V_w = v(\Gamma_w)$ be the vertex set of Γ_w and let Γ_{12} be the full subgraph spanned by $V_w^{\perp}$. Then $\Gamma_w \circ \Gamma_{12} = \Gamma_{11} \circ \Gamma_2 \circ \Gamma_{12}$ is stable by Lemma 3.12. Pick a vertex $v \in \Gamma_1 \setminus \Gamma_{11}$; then $v \in w^{\perp} \subset \mathrm{St}(u)$ for any vertex $u \in \Gamma_w$ by Lemma 3.23. Thus $v \in \Gamma_{12}$ and $\Gamma_1 \subset \Gamma_{11} \circ \Gamma_{12}$. On the other hand, $w \in \Gamma_2$, so $\Gamma_{11} \circ \Gamma_{12} \subset \Gamma_1$ and $\Gamma_1 = \Gamma_{11} \circ \Gamma_{12}$. Since Γ_2 does not contain any clique factor and $\Gamma_{11} \circ \Gamma_2 \circ \Gamma_{12} = \Gamma_1 \circ \Gamma_2$ is stable, we know Γ_1 is stable in Γ by Theorem 2.9. $\qquad\square$

Remark 5.2 In the above proof, Γ_{12} may be empty. But if $\Gamma_{12} \neq \varnothing$, then it does not contain any clique join factor. Thus Γ_{11} is the maximal clique join factor of $\Gamma_{11} \circ \Gamma_2 \circ \Gamma_{12}$.

The next result answers the question at the end of Example 3.22.

Theorem 5.3 *Suppose* $\mathrm{Out}(G(\Gamma))$ *is finite and let* $q \colon X(\Gamma) \to X(\Gamma')$ *be a quasi-isometry. Then* q *induces a simplicial isomorphism* $q_* \colon \mathcal{P}(\Gamma) \to \mathcal{P}(\Gamma')$; *in particular,* $\mathrm{Out}(G(\Gamma'))$ *is transvection-free.*

In the following proof, we identify Γ with the one-skeleton of $F(\Gamma)$, which is the flag complex of Γ. Also recall that there are label-preserving projections $\pi \colon \mathcal{P}(\Gamma) \to F(\Gamma)$ and $\pi \colon \mathcal{P}(\Gamma') \to F(\Gamma')$.

Proof By Lemma 4.5, there is a simplicial embedding $q_* \colon \mathcal{P}(\Gamma) \to \mathcal{P}(\Gamma')$. Note that $q_*(\mathcal{P}(\Gamma))$ is a full subcomplex in $\mathcal{P}(\Gamma')$. To see this, pick a simplex $\Delta' \subset \mathcal{P}(\Gamma')$ with its vertices in $q_*(\mathcal{P}(\Gamma))$. Then each vertex of Δ' comes from a stable standard geodesic line in $X(\Gamma')$. Thus there exists a stable standard flat $F' \subset X(\Gamma')$ such that $\Delta(F') = \Delta'$ by Lemma 3.18. By considering the quasi-inverse of q, we know F' is Hausdorff close to the q–image of a stable standard flat in $X(\Gamma)$. Thus $\Delta(F') = \Delta' \subset q_*(\mathcal{P}(\Gamma))$.

Pick a vertex $p \in X(\Gamma)$, and let $\{\Delta_i\}_{i=1}^k$, $\{F_i\}_{i=1}^k$, $\{\Delta_i'\}_{i=1}^k$ and $\{F_i'\}_{i=1}^k$ be as in the proof of Lemma 4.10. We claim

$$(5\text{-}1) \qquad\qquad \bigcap_{i=1}^k F_i' \neq \varnothing.$$

Suppose (5-1) is not true. Then there exist $1 \leq i_1 \neq i_2 \leq k$ and hyperplane $h' \subset X(\Gamma)$ such that h' separates F_{i_1}' and F_{i_2}'. Let l' be a standard geodesic that intersects h' transversely, and let $v' = \Delta(l')$. By the discussion in Lemma 4.9, we can find vertices $v_1' \in \Delta_{i_1}'$ and $v_2' \in \Delta_{i_2}'$ such that v_1' and v_2' are separated by $\mathrm{St}(v')$. If there exists i_0 such that $F_{i_0}' \cap h \neq \varnothing$, then $v' \in q_*(\mathcal{P}(\Gamma'))$, and we can prove (5-1) as in Lemma 4.10. Now we assume $F_i' \cap h' = \varnothing$ for any i. Let $w' = \pi(v') \in \Gamma'$, and let $\Gamma_{w'}$ be the minimal stable subgraph of Γ' that contains w'.

We apply Lemma 5.1 to $w' \in \Gamma'$; if case (1) is true, let F' be the standard flat in $X(\Gamma')$ such that $l' \subset F'$ and $\Gamma_{F'} = \Gamma_{w'}$. Since $\Gamma_{w'}$ is stable, $\Delta(F') \subset q_*(\mathcal{P}(\Gamma'))$; in particular, $v' \in q_*(\mathcal{P}(\Gamma'))$, and we can prove (5-1) as in Lemma 4.10.

If case (2) is true, let $\Gamma_1' = \mathrm{lk}(w')$ and let $\Gamma_2' = \mathrm{lk}(\Gamma_1')$. Take K_1' and K' to be the standard subcomplexes in $X(\Gamma')$ such that: (a) the defining graphs $\Gamma_{K_1'}$ and $\Gamma_{K'}$ of K_1' and K' satisfy $\Gamma_{K_1'} = \Gamma_1'$ and $\Gamma_{K'} = \Gamma_1' \circ \Gamma_2'$; (b) $l' \subset K'$ and $K_1' \subset K'$. Set $M_1' = \Delta(K_1')$ and $M' = \Delta(K')$. Let K_2' be an orthogonal complement of K_1' in K'; ie K_2' is a standard subcomplex such that $\Gamma_{K_2'} = \Gamma_2'$ and $K' = K_1' \times K_2'$. It follows that $M' = M_1' * M_2'$ for $M_2' = \Delta(K_2')$. By construction, $v' \in M'$ and $\mathrm{lk}(v') = M_1'$.

Since K' and K_1' are stable, there exist stable standard subcomplexes K and K_1 in $X(\Gamma)$ such that $q(K) \overset{\infty}{=} K'$ and $q(K_1) \overset{\infty}{=} K_1'$. Moreover, by applying Theorem 2.9 to the quasi-isometry between K and K', there exists a standard subcomplex $K_2 \subset K$ such that $K = K_1 \times K_2$, and K_2 is quasi-isometric to K_2'. Thus Γ_{K_2} is also disconnected. Let $M_i = \Delta(K_i) \subset \mathcal{P}(\Gamma)$ for $i = 1, 2$, and let $M = M_1 * M_2 = \Delta(K)$. Then $q_*(M_1) \subset M_1'$ (at this stage we may not know $q_*(M_1) = M_1'$), and

$$(5\text{-}2) \qquad\qquad q_*^{-1}(M_1') = M_1.$$

To see this, pick a simplex $\Delta \subset \mathcal{P}(\Gamma)$ with $q_*(\Delta) \subset M_1'$. Suppose $\Delta = \Delta(F)$ for a stable standard flat $F \subset X(\Gamma)$. Then $q(F) \subset_\infty K_1'$; hence $F \subset_\infty K_1$ and $\Delta \subset M_1$.

Let $L = \bigcup_{i=1}^k \Delta_i$ and $L' = \bigcup_{i=1}^k \Delta_i'$. By the proof of Lemma 4.10, $L' \setminus (\mathrm{St}(v') \cap L')$ is disconnected; thus $L \setminus q_*^{-1}(\mathrm{St}(v') \cap L')$ is disconnected. Recall that $\mathrm{lk}(v') = M_1'$, and we are assuming $v' \notin L'$. Thus $(\mathrm{St}(v') \cap L') \subset M_1'$. Then $q_*^{-1}(\mathrm{St}(v') \cap L') \subset q_*^{-1}(M_1')$; hence $q_*^{-1}(\mathrm{St}(v') \cap L') \subset M_1$ by (5-2).

Let $N = \pi(q_*^{-1}(\mathrm{St}(v') \cap L'))$, and let $N_i = \pi(M_i)$ for $i = 1, 2$. Then N separates $F(\Gamma)$, $N \subset N_1$ and N_2 is disconnected. Pick vertices u_1, u_2 in different connected

components of N_2; then $d(u_1, u_2) \geq 2$ (since N_2 is the full subcomplex spanned by Γ_{K_2}). Since $\pi(M) = N_1 * N_2 \subset F(\Gamma)$, we have $N \subset \mathrm{St}(u_i) \setminus \{u_i\}$ for $i = 1, 2$. Let $\{C_j\}_{j=1}^d$ be the connected components of $F(\Gamma) \setminus N$. Then at most one of C_j is contained in $\mathrm{St}(u_1)$. If $d \geq 3$, then $\mathrm{St}(u_1)$ would separate $F(\Gamma)$, which is a contradiction. Now we suppose $d = 2$. Note that for $i = 1, 2$, there must exist j such that $C_j \subset \mathrm{St}(u_i)$: otherwise, $\mathrm{St}(u_i)$ would separate $F(\Gamma)$. Moreover, if $C_j \subset \mathrm{St}(u_i)$, then $u_i \in C_j$. So we can assume without loss of generality that $C_1 \subset \mathrm{St}(u_1)$ and $C_2 \subset \mathrm{St}(u_2)$, which implies $F(\Gamma) = \mathrm{St}(u_1) \cup \mathrm{St}(u_2)$, and again we have a contradiction by Lemma 4.12. Thus case (2) is impossible, and (5-1) is true.

Let $\{F_\lambda\}_{\lambda \in \Lambda}$ be the collection of maximal standard flats in $X(\Gamma)$. Then $X(\Gamma) = \bigcup_{\lambda \in \Lambda} F_\lambda$. For each λ, let F'_λ be the unique maximal standard flat in $X(\Gamma')$ such that $q(F_\lambda) \overset{\infty}{\cong} F'_\lambda$. Then

$$\tag{5-3} X(\Gamma') \overset{\infty}{\cong} \bigcup_{\lambda \in \Lambda} F'_\lambda.$$

Let $h \subset X(\Gamma')$ be an arbitrary hyperplane. Then $h \cap \left(\bigcup_{\lambda \in \Lambda} F'_\lambda \right) \neq \varnothing$: otherwise, $\bigcup_{\lambda \in \Lambda} F'_\lambda$ would stay on one side of the hyperplane since it is a connected set by (5-1), and this contradicts (5-3). Pick any standard geodesic $r \subset X(\Gamma')$, and let h_r be a hyperplane dual to r. Then there exists $\lambda \in \Lambda$ such that $F'_\lambda \cap h_r \neq \varnothing$. It follows that $r \subset_\infty F'_\lambda$. So $\Delta(r) \in \Delta(F'_\lambda) \subset q_*(\mathcal{P}(\Gamma))$, which implies q_* is surjective on the vertices. However, $q_*(\mathcal{P}(\Gamma))$ is a full subcomplex in $\mathcal{P}(\Gamma')$, so q_* is surjective. $\qquad\square$

5.2 Coherent ordering and coherent labeling

Throughout this section, we assume that $\mathrm{Out}(G(\Gamma))$ is finite and $G(\Gamma) \ncong \mathbb{Z}$. If $q\colon G(\Gamma) \to G(\Gamma')$ is a quasi-isometry, then $G(\Gamma')$ has a quasi-action (see [46, Definition 2.2]) on $G(\Gamma)$, which induces a group homomorphism

$$H\colon G(\Gamma') \to \mathrm{QI}(G(\Gamma)).$$

On the other hand, since $G(\Gamma)$ acts by isometries on $X(\Gamma)$, we can identify $G(\Gamma)$ as a subgroup of $\mathrm{QI}(G(\Gamma))$ (more precisely, we embed $G(\Gamma)$ into $\mathrm{Isom}(G(\Gamma), d_w)$ and embed $\mathrm{Isom}(G(\Gamma), d_w)$ into $\mathrm{QI}(G(\Gamma))$ by Corollary 4.18). In this subsection, we will seek to answer the following question:

Does there exist $g \in \mathrm{QI}(G(\Gamma))$ such that $g \cdot H(G(\Gamma')) \cdot g^{-1} \subset G(\Gamma)$?

Recall that we have picked an identification between $G(\Gamma)$ and the 0–skeleton of $X(\Gamma)$. Each circle in the 1–skeleton of the Salvetti complex of $G(\Gamma)$ is labeled by an element in the standard generating set S of $G(\Gamma)$. Moreover, we have chosen an orientation

for each such circle. By pulling back the labeling and orientation of edges to the universal cover $X(\Gamma)$, we obtain a $G(\Gamma)$–invariant directed labeling of edges in $X(\Gamma)$. Moreover, both the labeling and orientation of edges in $X(\Gamma)$ are compatible with parallelism between edges. This also induces an associated $G(\Gamma)$–invariant labeling of vertices in $\mathcal{P}(\Gamma)$.

Let $\{l_\lambda\}_{\lambda \in \Lambda}$ be the collection of standard geodesics in $X(\Gamma)$, and let $V_\lambda = v(l_\lambda)$ be the vertex set of l_λ. A *coherent ordering* of $G(\Gamma)$ is obtained by assigning a collection of bijections $f_\lambda \colon V_\lambda \to \mathbb{Z}$ for each $\lambda \in \Lambda$ such that if l_{λ_1} and l_{λ_2} are parallel, then the $f_{\lambda_2} \circ p \circ f_{\lambda_1}^{-1} \colon \mathbb{Z} \to \mathbb{Z}$ is a translation, where $p \colon V_{\lambda_1} \to V_{\lambda_2}$ is the map induced by parallelism. The map f_λ pulls back the total order on $\mathbb{Z}$ to V_λ, which we denote by $\leq_\lambda$. Then $p \colon V_{\lambda_1} \to V_{\lambda_2}$ is order preserving.

Two coherent orderings Ω_1 and Ω_2 are *equivalent*, denoted by $\Omega_1 = \Omega_2$, if their collections of bijections agree up to a translation of $\mathbb{Z}$. Recall that we have a $G(\Gamma)$–invariant orientation of edges in $X(\Gamma)$ which is compatible with parallelism between edges. This induces a unique coherent ordering Ω of $G(\Gamma)$ up to the equivalence relation defined before. Moreover, for any element $g \in G(\Gamma)$, the pull-back $g^*(\Omega)$ is also a coherent ordering; additionally, $g^*(\Omega) = \Omega$.

Recall that for any vertex $v \in X(\Gamma)$, there is a label-preserving simplicial embedding $i_v \colon F(\Gamma) \to \mathcal{P}(\Gamma)$ by considering the standard geodesics passing through v. A *coherent labeling* of $G(\Gamma)$ is a simplicial map $a \colon \mathcal{P}(\Gamma) \to F(\Gamma)$ such that $a \circ i_v \colon F(\Gamma) \to F(\Gamma)$ is a simplicial isomorphism for every vertex $v \in X(\Gamma)$.

The label-preserving projection $L \colon \mathcal{P}(\Gamma) \to F(\Gamma)$ gives rise to a coherent labeling of $G(\Gamma)$. Recall that $G(\Gamma)$ acts on $\mathcal{P}(\Gamma)$ by simplicial automorphisms, and the labeling of vertices in $\mathcal{P}(\Gamma)$ is $G(\Gamma)$–invariant. Thus for any element $g \in G(\Gamma)$, the pull-back $g^*(L)$ is also a coherent labeling and $g^*(L) = L$.

We have the following alternative characterization of elements in $\mathrm{Isom}(G(\Gamma), d_r)$.

Lemma 5.4 *There is a one-to-one correspondence which associates each element of* $\mathrm{Isom}(G(\Gamma), d_r)$ *to a triple consisting of*

(1) *a point* $v \in G(\Gamma)$,

(2) *a coherent ordering of* $G(\Gamma)$ *(up to the equivalence relation defined above),*

(3) *a coherent labeling of* $G(\Gamma)$.

Proof Pick $\phi \in \mathrm{Isom}(G(\Gamma), d_r)$ and let $\varphi = h(\phi) \colon \mathcal{P}(\Gamma) \to \mathcal{P}(\Gamma)$, where h is the monomorphism in Remark 4.16. Then $\varphi^* L = L \circ \varphi \colon \mathcal{P}(\Gamma) \to F(\Gamma)$ is a coherent labeling of $G(\Gamma)$. Pick a standard geodesic $l_1 \subset X(\Gamma)$. Then the parallel set P_{l_1}

admits a splitting $P_{l_1} = l_1 \times l_1^{\perp}$. Since ϕ maps vertices in a standard flat bijectively to vertices in a standard flat, there exists a standard geodesic $l_2 \subset X(\Gamma)$ such that $\phi(v(l_1)) = v(l_2)$ and $\phi(v(P_{l_1})) = P_{l_2}$; moreover, ϕ respects the product structure on P_{l_1}. Thus the pull-back $\phi^*\Omega$ is a coherent ordering of $G(\Gamma)$. Now we can set up the correspondence in one direction:

$$\phi \rightsquigarrow (\phi(\mathrm{id}), \phi^*\Omega, \varphi^*L).$$

Here, id denotes the identity element of $G(\Gamma)$.

Conversely, given a point $v \in G(\Gamma)$, a coherent ordering Ω' and a coherent labeling L', we can construct a map ϕ as follows. Set $\phi(\mathrm{id}) = v$. For $u \in G(\Gamma)$, pick a word $w_u = a_1 a_2 \cdots a_n$ representing u. Let u_i be the point in $G(\Gamma)$ represented by the word $a_1 a_2 \cdots a_i$ for $1 \leq i \leq n$, and let $u_0 = \mathrm{id}$. We define $q_i = \phi(a_1 a_2 \cdots a_i) \in G(\Gamma)$ inductively as follows. Set $q_0 = v$, and suppose q_{i-1} is already defined. Denote the standard geodesic containing u_{i-1} and u_i by l_i. Let $v_i = L'(\Delta(l_i))$, which is a vertex of Γ, and let l_i' be the standard line that contains q_{i-1} and is labeled by v_i. Denote the vertex set of l_i with the order from Ω' by $(v(l_i), \leq_{\Omega'})$. Suppose that $k \colon (v(l_i), \leq_{\Omega'}) \to (v(l_i'), \leq_{\Omega})$ is the unique order-preserving bijection such that $k(u_{i-1}) = q_{i-1}$. Then we define $q_i = k(u_i)$.

We claim that for any other word w_u' representing u, we have $\phi(w_u) = \phi(w_u')$, and hence there is a well-defined map $\phi \colon G(\Gamma) \to G(\Gamma)$. To see this, recall that one can obtain w_u from w_u' by performing the following two basic moves:

(1) $w_1 a a^{-1} w_2 \to w_1 w_2$,

(2) $w_1 a b w_2 \to w_1 b a w_2$ when a and b commute.

It is clear that $\phi(w_1 a a^{-1} w_2) = \phi(w_1 w_2)$. For the second move, let u_{i-1}, u_i, u_i' and u_{i+1} be points in $G(\Gamma)$ represented by $w_1, w_1 a, w_1 b$ and $w_1 ab = w_1 ba$, respectively. Define $q_{i-1} = \phi(w_1)$, $q_i = \phi(w_1 a)$, $q_i' = \phi(w_1 b)$, $q_{i+1} = \phi(w_1 ab)$ and $q_{i+1}' = \phi(w_1 ba)$. Since L' is a coherent labeling, $\angle_{q_i}(q_{i+1}, q_{i-1}) = \angle_{q_{i-1}}(q_i, q_i') = \angle_{q_i'}(q_{i-1}, q_{i+1}') = \frac{\pi}{2}$; moreover, the standard geodesic containing q_i and q_{i+1} is parallel to the standard geodesic containing q_{i-1} and q_i'. Since Ω' is a coherent ordering, $d(q_i, q_{i+1}) = d(q_{i-1}, q_i')$; thus $\overline{q_i q_{i+1}}$ and $\overline{q_{i-1} q_i'}$ are parallel. Similarly, $\overline{q_{i-1} q_i}$ and $\overline{q_i' q_{i+1}'}$ are parallel; thus $q_{i+1} = q_{i+1}'$.

Now we define another map $\phi' \colon G(\Gamma) \to G(\Gamma)$, which serves as the inverse of ϕ. Set $\phi'(v) = \mathrm{id}$ and pick a word $w = a_1 a_2 \cdots a_n$. Let r_i be the point in $G(\Gamma)$ represented by $v a_1 a_2 \cdots a_i$ for $1 \leq i \leq n$, and let $r_0 = v$. We define $p_i = \phi'(v a_1 a_2 \cdots a_i)$ inductively as follows. Put $p_0 = \mathrm{id}$, and suppose p_{i-1} is already defined. Since L' is a coherent labeling, there exists a unique standard geodesic l_i containing p_{i-1} such

that $L'(\Delta(l_i))$ and the edge $\overline{r_{i-1}r_i}$ share the same label. Let l_i' be the unique standard geodesic containing r_{i-1} and r_i, and let $k'\colon (v(l_i'), \leq_\Omega) \to (v(l_i), \leq_{\Omega'})$ be the unique order-preserving bijection such that $k'(r_{i-1}) = p_{i-1}$. Put $p_i = k'(r_i)$. By a similar argument as above, $\phi'\colon G(\Gamma) \to G(\Gamma)$ is well defined. It is not hard to deduce the following properties from our construction:

(1) $\phi' \circ \phi = \phi \circ \phi' = \mathrm{Id}$.

(2) $d_r(\phi(v_1), \phi(v_2)) \leq d_r(v_1, v_2)$ and $d_r(\phi'(v_1), \phi'(v_2)) \leq d_r(v_1, v_2)$ for any vertices $v_1, v_2 \in G(\Gamma)$.

(3) If $L' = L$ and $\Omega' = \Omega$, then ϕ is a left translation. If, in addition, $v = \mathrm{id}$, then $\phi = \mathrm{Id}$.

It follows from (1) and (2) that $\phi \in \mathrm{Isom}(G(\Gamma), d_r)$. Moreover, $v = \phi(\mathrm{id})$, $L' = \varphi^* L$ ($\varphi = h(\phi)$, where h is the monomorphism in Remark 4.16) and $\Omega' = \phi^* \Omega$; thus we have established the required one-to-one correspondence. □

Pick finite simplicial graphs Γ and Γ' such that: (1) $\mathrm{Out}(G(\Gamma))$ is finite; (2) there exists a simplicial isomorphism $s\colon \mathcal{P}(\Gamma) \to \mathcal{P}(\Gamma')$. By Lemma 4.10, s induces a map $\phi\colon G(\Gamma) \to G(\Gamma')$. For every $g' \in G(\Gamma')$, there is a left translation

$$\overline{\phi}_{g'}\colon G(\Gamma') \to G(\Gamma'),$$

which gives rise to a simplicial isomorphism $\overline{s}_{g'}\colon \mathcal{P}(\Gamma') \to \mathcal{P}(\Gamma')$. Let $s_{g'} = s^{-1} \circ \overline{s}_{g'} \circ s$. Then $s_{g'}$ gives rise to a map $\phi_{g'} \in \mathrm{Isom}(G(\Gamma), d_r)$ by Corollary 4.15; moreover, by Lemma 4.10,

$$(5\text{-}4) \qquad\qquad \overline{\phi}_{g'} \circ \phi = \phi \circ \phi_{g'}$$

for any $g' \in G(\Gamma')$. So $G(\Gamma')$ acts on $G(\Gamma)$, and we can define a homomorphism $\Phi\colon G(\Gamma') \to \mathrm{Isom}(G(\Gamma), d_r)$ by sending g' to $\phi_{g'}$. Φ is injective since each step in defining Φ is injective.

Lemma 5.5 *In the above setting, there exists an element $\phi_1 \in \mathrm{Isom}(G(\Gamma), d_r)$ such that it conjugates the image of Φ to a finite-index subgroup of $G(\Gamma)$.*

We identify $G(\Gamma)$ as a subgroup of $\mathrm{Isom}(G(\Gamma), d_r)$ via the left action of $G(\Gamma)$ on itself.

Proof Pick a reference point $q \in \mathrm{Im}\,\phi$, and let $K_q = (F(\Gamma'))_q$. Denote the points in $\phi^{-1}(q)$ by $\{p_\lambda\}_{\lambda \in \Lambda}$, and let $K_{p_\lambda} = (F(\Gamma))_{p_\lambda}$. Since the $\{\phi(K_{p_\lambda})\}_{\lambda \in \Lambda}$ are distinct subcomplexes of K_q, the set Λ must be finite.

Let $L\colon \mathcal{P}(\Gamma) \to F(\Gamma)$ and Ω be the coherent labeling and coherent ordering induced by the $G(\Gamma)$–invariant labeling of $X(\Gamma)$ and $\mathcal{P}(\Gamma)$. In a similar fashion, we can obtain

a coherent labeling $L' \colon \mathcal{P}(\Gamma') \to F(\Gamma')$ and a coherent ordering Ω' for $G(\Gamma')$ which are invariant under the $G(\Gamma')$–action, ie

$$(5\text{-}5) \qquad\qquad (\bar{s}_{g'})^* L' = L' \quad \text{and} \quad (\bar{\phi}_{g'})^* \Omega' = \Omega'.$$

Our goal is to find a coherent labeling L_1 and a coherent ordering Ω_1 of $G(\Gamma)$ such that $(s_{g'})^* L_1 = L_1$ and $(\phi_{g'})^* \Omega_1 = \Omega_1$ for any $g' \in G(\Gamma')$.

Let $i_q \colon F(\Gamma') \to \mathcal{P}(\Gamma')$ be the canonical embedding, and let

$$L_1 = L \circ s^{-1} \circ i_q \circ L' \circ s$$

be the simplicial map from $\mathcal{P}(\Gamma)$ to $F(\Gamma)$. Pick an arbitrary $p \in G(\Gamma)$, and let $i_p \colon F(\Gamma) \to \mathcal{P}(\Gamma)$ be the canonical embedding. We need to show $L_1 \circ i_p$ is a simplicial isomorphism. Let $K_p = i_p(F(\Gamma))$, and let $g'_1 \in G(\Gamma')$ such that $g'_1 \cdot \phi(p) = q$. Then $i_q \circ L'|_{s(K_p)} = \bar{s}_{g'_1}|_{s(K_p)}$. Thus

$$L_1 \circ i_p = L \circ s^{-1} \circ i_q \circ L' \circ s \circ i_p = L \circ s^{-1} \circ \bar{s}_{g'_1} \circ s \circ i_p = L \circ s_{g'_1} \circ i_p,$$

which is a simplicial isomorphism by [Lemma 4.10](#). It follows that L_1 is a coherent labeling; moreover,

$$(s_{g'})^* L_1 = (L \circ s^{-1} \circ i_q \circ L' \circ s) \circ (s^{-1} \circ \bar{s}_{g'} \circ s) = L \circ s^{-1} \circ i_q \circ L' \circ \bar{s}_{g'} \circ s$$
$$= L \circ s^{-1} \circ i_q \circ L' \circ s = L_1$$

for any $g' \in G(\Gamma')$, where the third equality follows from [(5-5)](#). So L_1 is the required coherent labeling.

To simplify notation, we will write $x <_{\Omega} y$ if $x < y$ under the ordering Ω. We define Ω_1 as follows. Let $p_1, p_2 \in G(\Gamma)$ be two distinct points in a standard geodesic line. If $\phi(p_1) \neq \phi(p_2)$, then we set $p_1 <_{\Omega_1} p_2$ if and only if $\phi(p_1) <_{\Omega'} \phi(p_2)$. If $\phi(p_1) = \phi(p_2)$, then by [(5-4)](#), there exists a unique $g' \in G(\Gamma')$ such that $\phi_{g'}(p_i) \in \phi^{-1}(q)$ for $i = 1, 2$, and we set $p_1 <_{\Omega_1} p_2$ if and only if $\phi_{g'}(p_1) <_{\Omega} \phi_{g'}(p_2)$. It follows from [(5-5)](#), [(5-4)](#) and our construction that $p_1 <_{\Omega_1} p_2$ if and only if $\phi_{g'}(p_1) <_{\Omega_1} \phi_{g'}(p_2)$ for any p_1, p_2 in the same standard geodesic line and any $g' \in G(\Gamma')$; thus $(\phi_{g'})^* \Omega_1 = \Omega_1$.

To verify Ω_1 is coherent, pick parallel standard geodesics l_1 and l_2 in $X(\Gamma)$, and pick distinct vertices $p_{11}, p_{12} \in l$. Let p_{21}, p_{22} be the corresponding vertices in l_2 via parallelism. We assume $p_{11} <_{\Omega_1} p_{12}$; it suffices to prove $p_{21} <_{\Omega_1} p_{22}$.

Case 1 We assume $\phi(p_{11}) \neq \phi(p_{12})$. Recall that l_1 can be realized as an intersection of finitely many maximal standard flats, so by [Lemma 4.10](#), there exists a standard geodesic line $l'_1 \subset X(\Gamma')$ such that $\phi(v(l_1)) \subset v(l'_1)$ and $\phi(v(P_{l_1})) \subset v(P_{l'_1})$; moreover, ϕ respects the product structures of P_{l_1} and $P_{l'_1}$. Thus $\overline{\phi(p_{11})\phi(p_{21})}$ and $\overline{\phi(p_{21})\phi(p_{22})}$

are the opposite sides of a flat rectangle in $X(\Gamma')$. Now $p_{21} <_{\Omega_1} p_{22}$ follows since Ω' is coherent.

Case 2 We assume $\phi(p_{11}) = \phi(p_{12}) \neq \phi(p_{21})$. In this case, we can assume without loss of generality that $\phi(p_{11}) = \phi(p_{12}) = q$ (since $(\phi_{g'})^*\Omega_1 = \Omega_1$), and the points p_{11} and p_{21} stay in the same standard geodesic. For $i = 1, 2$, let r_i be the standard geodesic passing p_{1i} and p_{2i}. Take $r_i' \subset X(\Gamma')$ and $l_i' \subset X(\Gamma')$ to be the standard geodesics such that $\phi(v(r_i)) \subset v(r_i')$ and $\phi(v(l_i)) \subset v(l_i')$, respectively. Let $q' = \phi(p_{21})$. Since ϕ restricted to $v(P_{l_1})$ respects the product structure, $\phi(p_{21}) = \phi(p_{22}) = q'$ and $r_1' = r_2'$.

Let $\bar{\phi}_{g'}$ be the left translation such that $\bar{\phi}_{g'}(q') = q$. Since $q' \in r_1'$ and $q \in r_1'$, we have that $\bar{\phi}_{g'}$ is a translation along r_1', and $\bar{s}_{g'}$ fixes every point in $\mathrm{St}(\Delta(r_1'))$; hence $s_{g'}$ fixes every point in $s^{-1}(\mathrm{St}(\Delta(r_1'))) = \mathrm{St}(\Delta(r_1))$, and

$$(5\text{-}6) \qquad\qquad\qquad\qquad \phi_{g'}(r_i) = r_i$$

for $i = 1, 2$. Let $l_3 = \phi_{g'}(l_2)$. Then l_3 is parallel to l_1 (or l_2). To see this, note that $\Delta(l_1) \in \mathrm{St}(\Delta(r_1))$; hence $\Delta(l_1)$ is fixed by $s_{g'}$. Put $p_{3i} = \phi_{g'}(p_{2i})$ for $i = 1, 2$. Then $p_{3i} \in r_i$ by (5-6); hence $\overline{p_{11}p_{12}}$ and $\overline{p_{31}p_{32}}$ are the opposite sides of a flat rectangle. Moreover, $p_{3i} \in \phi^{-1}(q)$ for $i = 1, 2$ by (5-4), so $p_{31} <_{\Omega_1} p_{32}$ since Ω is coherent, and $\Omega = \Omega_1$ while restricted on $\phi^{-1}(q)$. Now the $G(\Gamma')$–invariance of Ω_1 implies $p_{21} <_{\Omega_1} p_{22}$.

Case 3 If $\phi(p_{11}) = \phi(p_{12}) = \phi(p_{21})$, then we can assume without loss of generality that they all equal to q. It follows that $\phi(p_{22}) = q$ since ϕ respects the product structure while restricted to $v(P_{l_1})$. Thus $p_{21} <_{\Omega_1} p_{22}$ by definition.

By Lemma 5.4, there exists $\phi_1 \in \mathrm{Isom}(G(\Gamma), d_r)$ such that $\phi_1^*\Omega = \Omega_1$ and $s_1^*L = L_1$ ($s_1 = h(\phi_1)$ where h is the monomorphism in Remark 4.16). Thus

$$(\phi_1 \circ \phi_{g'} \circ \phi_1^{-1})^*\Omega = (\phi_1^{-1})^* \circ (\phi_{g'})^* \circ (\phi_1^*\Omega) = (\phi_1^{-1})^* \circ (\phi_{g'})^*\Omega_1$$
$$= (\phi_1^{-1})^*\Omega_1 = \Omega$$

for any $g' \in G(\Gamma')$. Similarly, $(s_1 \circ s_{g'} \circ s_1^{-1})^*L = L$ for any $g' \in G(\Gamma')$. Note that $s_1 \circ s_{g'} \circ s_1^{-1} = h(\phi_1 \circ \phi_{g'} \circ \phi_1^{-1})$; thus by Lemma 5.4, $G(\Gamma')$ acts on $G(\Gamma)$ by left translations via $g' \to \phi_1 \circ \phi_{g'} \circ \phi_1^{-1}$. This induces a monomorphism $G(\Gamma') \to G(\Gamma)$. Moreover, by (5-4) and the fact that $\phi^{-1}(q)$ is finite, this action has finite quotient; thus we can realize $G(\Gamma')$ as a finite-index subgroup of $G(\Gamma)$. $\square$

The next result basically says under suitable conditions, if there exists a quasi-isometry $q\colon G(\Gamma) \to G(\Gamma')$, then there exists a very "nice" quasi-isometry $q'\colon G(\Gamma) \to G(\Gamma')$. However, we do not insist that q' is of bounded distance away from q (compared to Theorem 4.13).

Theorem 5.6 Let Γ and Γ' be finite simplicial graphs such that $\mathrm{Out}(G(\Gamma))$ is finite and $G(\Gamma')$ is quasi-isometric to $G(\Gamma)$. Then there exists a cubical map (see Definition 2.1) $\varphi\colon X(\Gamma) \to X(\Gamma')$ such that:

(1) The map φ is onto, and φ maps any standard flat in $X(\Gamma)$ onto a standard flat in $X(\Gamma')$ of the same dimension.

(2) The map φ maps combinatorial geodesics in the 1–skeleton of $X(\Gamma)$ to combinatorial geodesics in the 1–skeleton of $X(\Gamma')$.

(3) The map φ is a quasi-isometry.

Proof Let $f\colon G(\Gamma) \to G(\Gamma')$ be a quasi-isometry. By Theorem 5.3, f induces a simplicial isomorphism $s\colon \mathcal{P}(\Gamma) \to \mathcal{P}(\Gamma')$. By Lemma 4.10, s induces a map $\phi\colon G(\Gamma) \to G(\Gamma')$ such that $d_w(f(x), \phi(x)) < D$ for any $x \in G(\Gamma)$. Let ϕ_1 be the map in Lemma 5.5 and let $\varphi = \phi \circ \phi_1^{-1}$. We will use the same notation as in the proof of Lemma 5.5.

We claim that if $F = \bigcap_{i=1}^h F_i$, where each F_i is a maximal standard flat, then there exists a unique standard flat $F' \subset G(\Gamma')$ such that $\phi(v(F)) = v(F')$. To see this, let F_i' be the maximal standard flat in $X(\Gamma')$ such that $\Delta(F_i') = s(\Delta(F_i))$ for $1 \le i \le h$, and let $F' = \bigcap_{i=1}^h F_i'$. Then it follows from Lemma 4.10 that $\phi(v(F)) \subset v(F')$. Recall that $G(\Gamma')$ acts on $G(\Gamma')$, $\mathcal{P}(\Gamma')$, $G(\Gamma)$ and $\mathcal{P}(\Gamma)$. The stabilizer $\mathrm{Stab}(v(F'))$ fixes $\Delta(F_i')$ for all i; hence it fixes Δ_i for all i, and $\mathrm{Stab}(v(F')) \subset \mathrm{Stab}(v(F))$. Since $\mathrm{Stab}(v(F'))$ acts on $v(F')$ transitively, (5-4) implies $\phi(v(F)) = v(F')$ and $|\phi^{-1}(y) \cap F| = |\phi^{-1}(y') \cap F|$ for any $y, y' \in v(F')$. It also follows that $\mathrm{Stab}(v(F)) \subset \mathrm{Stab}(v(F'))$; thus $\mathrm{Stab}(v(F')) = \mathrm{Stab}(v(F))$.

Note that the above claim is also true for φ, and any standard geodesic satisfies the assumption of the claim. Moreover, φ is surjective since ϕ_1 is surjective by (5-4). Pick standard geodesics $l \subset X(\Gamma)$ and $l' \subset X(\Gamma')$ such that $v(l') = \varphi(v(l))$, and we identify $v(l)$ and $v(l')$ with $\mathbb{Z}$ in an order-preserving way. Then the above claim and the construction of ϕ_1 imply that $\varphi|_{v(l)}$ is of the form

(5-7) $$\varphi(a) = \lfloor a/d \rfloor + r$$

for some integers r and d (with $d \ge 1$). In particular, φ can be extended to a simplicial map from the Cayley graph $C(\Gamma)$ of $G(\Gamma)$ to $C(\Gamma')$.

Pick a combinatorial geodesic $\omega \subset C(\Gamma)$ connecting vertices x and y; we claim that $\omega' = \phi(\omega)$ is also a geodesic in $C(\Gamma')$ (it could be a point). Let $\{v_i\}_{i=0}^n$ be vertices in ω such that for $0 \le i \le n-1$, we have that $[v_i, v_{i+1}]$ is a maximal subsegment of ω that is contained in a standard geodesic ($v_0 = x$ and $v_n = y$). Denote the corresponding standard geodesic by l_i. For $0 \le i \le n-1$, let $l_i' \subset X(\Gamma')$ be the

standard geodesic such that $v(l_i') = \varphi(v(l_i))$, and let $\omega_i' = \phi([v_i, v_{i+1}])$. Then ω_i' is a (possibly degenerate) segment in l_i' by (5-7). Since ω is a geodesic, no two geodesics in $\{l_i\}_{i=0}^{n-1}$ are parallel. Note that φ is induced by a simplicial isomorphism between $\mathcal{P}(\Gamma)$ and $\mathcal{P}(\Gamma')$; thus the same property is true for the collection of geodesics $\{l_i'\}_{i=0}^{n-1}$. It follows that no hyperplane in $X(\Gamma')$ could intersect ω' at more than one point; hence ω' is a combinatorial geodesic.

Let $u_i = \varphi(v_i)$. Then $d_w(u_i, u_{i+1}) \leq d_w(v_i, v_{i+1})$ by (5-7) (recall that d_w denotes the word metric on the corresponding group). Thus

$$(5\text{-}8) \qquad d_w(\varphi(x), \varphi(y)) = \sum_{i=0}^{n-1} d_w(u_i, u_{i+1}) \leq \sum_{i=0}^{n-1} d_w(v_i, v_{i+1}) = d_w(x, y)$$

for any $x, y \in G(\Gamma)$.

Pick $p \in G(\Gamma')$ and let $k = |\varphi^{-1}(p)|$. Then k does not depend on p by (5-4). It follows that $d_w(\varphi(x), \varphi(y)) \geq 1$ whenever $d_w(x, y) \geq k+1$. Now we can cut ω into pieces of length $k+1$. Since $\varphi(\omega)$ is a combinatorial geodesic,

$$d_w(\varphi(x), \varphi(y)) \geq \frac{d_w(x, y)}{k+1} - 1.$$

Note that φ naturally extends to a cubical map from $X(\Gamma)$ to $X(\Gamma')$, which satisfies all the required properties. □

Theorem 5.7 *If Γ and Γ' are finite simplicial graphs such that $\mathrm{Out}(G(\Gamma))$ is finite, then the following are equivalent:*

(1) *$G(\Gamma')$ is quasi-isometric to $G(\Gamma)$.*

(2) *$\mathcal{P}(\Gamma')$ is isomorphic to $\mathcal{P}(\Gamma)$ as simplicial complexes.*

(3) *$G(\Gamma')$ is isomorphic to a subgroup of finite index in $G(\Gamma)$.*

Proof (1) $\implies$ (2) follows from Theorem 5.3. (2) $\implies$ (3) follows from Lemma 5.5. (3) $\implies$ (1) is trivial. □

This establishes Theorem 1.2 in the introduction.

6 The geometry of finite-index RAAG subgroups

Throughout this section, we assume $G(\Gamma) \not\cong \mathbb{Z}$, since the main results of this section (Theorems 6.13 and 6.19) are trivial when $G(\Gamma) \cong \mathbb{Z}$.

6.1 Constructing finite-index RAAG subgroups

A *right-angled Artin subgroup* is a subgroup which is also a right-angled Artin group. In this section, we introduce a process to obtain finite-index RAAG subgroups of an arbitrary RAAG.

Lemma 6.1 *Let X be a CAT(0) cube complex, let $l \subset X$ be a geodesic in the 1–skeleton and let $\{h_i\}_{i \in \mathbb{Z}}$ be consecutive hyperplanes dual to l. Let $\pi_l \colon X \to l$ be the CAT(0) projection. Then:*

(1) *For every edge $e \subset X$, if $e \cap h_i = \varnothing$ for all i, then $\pi_l(e)$ is a vertex in l, and if $e \cap h_i \neq \varnothing$ for some i, then $\pi_l(e)$ is an edge in l.*

(2) *If K is any connected subcomplex such that $e \cap h_i = \varnothing$ for all i, then $\pi_l(K)$ is a vertex in l; moreover, if K stays between h_i and h_{i+1}, then $\pi_l(K)$ is the vertex in l that stays between h_i and h_{i+1}.*

(3) *For every interval $[a, b] \subset l$, we have that $\pi_l^{-1}([a, b])$ is a convex set in X. In particular, if $x \in l$ is a vertex, then $\pi_l^{-1}(x)$ is a convex subcomplex of X.*

(4) *If K is a convex subcomplex such that $K \cap l \neq \varnothing$, then $\pi_l(K) = K \cap l$.*

Proof Here (1) and (3) follow from the fact the every hyperplane has a carrier, and (2) follows from (1). To see (4), it suffices to show that for every i such that $h_i \cap l \neq \varnothing$ and $h_i \cap K \neq \varnothing$, we have $e_i \subset K$ (e_i is the edge in l dual to h_i). Let N_{h_i} be the carrier of h_i. By [Lemma 2.3](), $d(x, N_{h_i} \cap K) \equiv c$ for any $x \in e_i$. Moreover, $d(x, N_{h_i} \cap K) = d(x, K)$ for x in the interior of e_i, so we must have $c = 0$: otherwise, the convexity of $d(\cdot, K)$ would imply $K \cap l = \varnothing$. $\qquad\square$

Recall that $v(\mathcal{P}(\Gamma) \setminus \mathrm{St}(\Delta(l)))$ is the collection of vertices in $\mathcal{P}(\Gamma) \setminus \mathrm{St}(\Delta(l))$.

Lemma 6.2 *Let $l \subset X(\Gamma)$ be a standard geodesic. Then there is a map*

$$\pi_{\Delta(l)} \colon v(\mathcal{P}(\Gamma) \setminus \mathrm{St}(\Delta(l))) \to v(l)$$

such that if v_1 and v_2 are in the same connected component of $\mathcal{P}(\Gamma) \setminus \mathrm{St}(\Delta(l))$, then $\pi_{\Delta(l)}(v_1) = \pi_{\Delta(l)}(v_2)$.

Proof Let $\pi_l \colon X(\Gamma) \to l$ be the CAT(0) projection and let $l_1 \subset X(\Gamma)$ be a standard geodesic such that $d(\Delta(l_1), \Delta(l)) \geq 2$. Then $\pi_l(l_1)$ is a vertex in l by [Lemma 3.1]() and [Corollary 3.2](). Moreover, we claim $\pi_l(l_1) = \pi_l(l_2)$ if l_2 is a standard geodesic parallel to l_1. It suffices to prove the case when there is a unique hyperplane h separating l_1 from l_2. Note that $d(\Delta(l_1), \Delta(l)) \geq 2$ yields $h \cap l = \varnothing$, so l_1 and l_2 are pinched by two hyperplanes dual to l; then the claim follows from [Lemma 6.1](). Thus π_l

induces a well-defined map $\pi_{\Delta(l)}\colon v(\mathcal{P}(\Gamma)\setminus \mathrm{St}(\Delta(l))) \to v(l)$. If $\Delta(l_1)$ and $\Delta(l_2)$ are connected by an edge, then there exist standard geodesics l_1' and l_2' such that $l_1' \cap l_2' \neq \varnothing$ and l_i' is parallel to l_i for $i = 1, 2$. Thus $\pi_l(l_1) = \pi_l(l_1') = \pi_l(l_2') = \pi_l(l_2)$, and $\pi_{\Delta(l)}(\Delta(l_1)) = \pi_{\Delta(l)}(\Delta(l_2))$. $\qquad\qquad\square$

Pick a standard generating set S of $G(\Gamma)$, and let $C(\Gamma, S)$ be the Cayley graph. We identify $G(\Gamma)$ as a subset of $C(\Gamma, S)$ and attach higher-dimensional cubes to $C(\Gamma, S)$ to obtain a CAT(0) cube complex $X(\Gamma, S)$, which is basically the universal cover of the Salvetti complex. Here we would like to think of $G(\Gamma)$ as a fixed set and of $C(\Gamma, S)$ and $X(\Gamma, S)$ as objects formed by adding edges and cubes to $G(\Gamma)$ in a particular way determined by S, so we write S explicitly. We will choose a $G(\Gamma)$–equivariant orientation for edges in $X(\Gamma, S)$ as before.

An S–*flat* (or an S–*geodesic*) in $G(\Gamma)$ is defined to be the vertex set of a standard flat (or geodesic) in $X(\Gamma, S)$. We define $\mathcal{P}(\Gamma, S)$ as before such that its vertices correspond to coarse equivalence classes of S–geodesics.

We define an isometric embedding $I\colon G(\Gamma) \to \ell^1(v(\mathcal{P}(\Gamma, S)))$ which depends on S and the orientation of edges in $X(\Gamma, S)$. Pick a standard geodesic $l \subset X(\Gamma, S)$, and let $\pi_l\colon X(\Gamma, S) \to l$ be the CAT(0) projection. We identify $v(l)$ with $\mathbb{Z}^{\Delta(l)}$ in an orientation-preserving way such that $\pi_l(\mathrm{id}) = 0$ (id is the identity element in $G(\Gamma)$). Then π_l induces a coordinate function $I_{\Delta(l)}\colon G(\Gamma) \to \mathbb{Z}^{\Delta(l)}$. If we change l to a standard geodesic l_1 parallel to l, then $I_{\Delta(l)}$ and $I_{\Delta(l_1)}$ are identical by Lemma 6.1. Thus for every vertex $v \in \mathcal{P}(\Gamma)$, there is a well-defined coordinate function $I_v\colon G(\Gamma) \to \mathbb{Z}^v$. These coordinate functions induce a map $I\colon G(\Gamma) \to \mathbb{Z}^{(v(\mathcal{P}(\Gamma)))}$.

The map I is an embedding since every two points in $G(\Gamma)$ are separated by some hyperplane. $I(G(\Gamma)) \subset \ell^1(v(\mathcal{P}(\Gamma)))$ since for any $g \in G(\Gamma)$, there are only finitely many hyperplanes separating id and g. I naturally extends to a map $I\colon X(\Gamma, S) \to \ell^1(v(\mathcal{P}(\Gamma)))$, and it maps combinatorial geodesics to geodesics by the argument in Theorem 5.6. Thus I is an isometric embedding with respect to the ℓ^1 metric on $X(\Gamma, S)$. We say a convex subcomplex $K \subset X(\Gamma, S)$ is *nonnegative* if each point in $I(K)$ has nonnegative coordinates (this notion depends on the orientation of edges in $X(\Gamma, S)$). Let $\mathrm{CN}(\Gamma, S)$ be the collection of compact, convex, nonnegative subcomplexes of $X(\Gamma, S)$ that contain the identity.

For any $K \in \mathrm{CN}(\Gamma, S)$, we find a maximal collection of standard geodesics $\{c_i\}_{i=1}^s$ such that $c_i \cap K \neq \varnothing$ for all i and $\Delta(c_i) \neq \Delta(c_j)$ for any $i \neq j$. Let $g_i \in S$ be the label of edges in c_i and let $\alpha_i = \pi_{c_i}(\mathrm{id})$. Put $n_i = |v(K \cap c_i)|$ and $v_i = \alpha_i g_i^{n_i} \alpha_i^{-1}$. Let G' be the subgroup generated by $\{v_i\}_{i=1}^s$. It follows from the convexity of K that if a standard geodesic c is parallel to c_i and $c \cap K \neq \varnothing$, then $|v(K \cap c_i)| = |v(K \cap c)|$. Thus $\{v_i\}_{i=1}^s$ and G' do not depend on the choice of the c_i.

Lemma 6.3 G' *is a finite-index subgroup of* $G(\Gamma)$.

Proof We prove this by showing $G' \cdot v(K) = G(\Gamma)$. Let d_r be the syllable metric on $G(\Gamma)$ defined in Section 4.3. Pick a word $\alpha \in G(\Gamma)$ and assume $\alpha \in G' \cdot v(K)$ when $d_r(\alpha, \mathrm{id}) \leq k - 1$. If $d_r(\alpha, \mathrm{id}) = k$, then there exists $\beta \in G(\Gamma)$ such that $d_r(\mathrm{id}, \beta) = k - 1$ and $d_r(\beta, \alpha) = 1$. Let $\beta = \beta_1 \beta_2$ for $\beta_1 \in G'$ and $\beta_2 \in v(K)$. Then $d_r(\beta_2, \beta_1^{-1}\alpha) = 1$. Suppose c is the standard geodesic containing β_2 and $\beta_1^{-1}\alpha$. Then there exists i such that c_i and c are parallel. Note that $P_c \cap K$ is a convex set in the parallel set P_c, hence respects the natural splitting $P_c = c \times c^{\perp}$; moreover, the left action of v_i translates the c factor by n_i units and fixes the other factor. Thus there exists $d \in \mathbb{Z}$ and $\beta_2' \in K \cap c$ such that $v_i^d \beta_2' = \beta_1^{-1}\alpha$, which implies $\alpha = \beta_1 v_i^d \beta_2' \in G' \cdot v(K)$. $\qquad\square$

Let Γ' be the full subgraph of $\mathcal{P}(\Gamma)$ spanned by points $\{\Delta(c_i)\}_{i=1}^s$. Then there is a natural homomorphism $G(\Gamma') \to G'$.

Lemma 6.4 *The homomorphism* $G(\Gamma') \to G'$ *is actually an isomorphism. Hence* G' *is a finite-index RAAG subgroup of* $G(\Gamma)$.

We will follow the strategy in [47], where the following version of the ping-pong lemma for right-angled Artin groups was used.

Theorem 6.5 [47, Theorem 4.1] *Let* $G = G(\Gamma)$ *and let* X *be a set with a* G*–action. Suppose the following hold:*

(1) *For each vertex* v_i *of* Γ, *there exists a subset* $X_i \subset X$ *such that the union of all the* X_i *is properly contained in* X.

(2) *For each nonzero* $k \in \mathbb{Z}$ *and vertices* v_i, v_j *joined by en edge,* $v_i^k(X_j) \subset X_j$.

(3) *For each nonzero* $k \in \mathbb{Z}$ *and vertices* v_i, v_j *not joined by en edge,* $v_i^k(X_j) \subset X_i$.

(4) *There exists* $x_0 \in X \setminus \bigcup_{i \in V} X_i$ (V *is the vertex set of* Γ) *such that* $v_i^k(x_0) \in X_i$ *for each nonzero* $k \in \mathbb{Z}$.

Then the G*–action is faithful.*

Proof of Lemma 6.4 We will apply Theorem 6.5 with $X = X(\Gamma, S)$ and $G = G(\Gamma')$. For $1 \leq i \leq s$, we identify c_i and $\mathbb{R}$ in an orientation-preserving way such that $\pi_{c_i}(\mathrm{id})$ corresponds to $0 \in \mathbb{R}$. Define $X_i^+ = \pi_{c_i}^{-1}([n_i - \frac{1}{2}, \infty))$, $X_i^- = \pi_{c_i}^{-1}((-\infty, -\frac{1}{2}])$ and $X_i = X_i^+ \cup X_i^-$. It clear that the identity element id does not lie in X_i for any i, so Theorem 6.5(1) is true. Each $v_i = \alpha_i g_i^{n_i} \alpha_i^{-1}$ translates c_i by n_i units, so (4) is also true with $x_0 = \mathrm{id}$.

If $\Delta(c_i)$ and $\Delta(c_j)$ are connected by an edge in $\mathcal{P}(\Gamma)$, then v_i stabilizes every hyperplane dual to v_j; thus $v_i^k(X_j) = X_j$, and (2) is true. If

$$(6\text{-}1) \qquad\qquad\qquad d(\Delta(c_i), \Delta(c_j)) \geq 2,$$

then $\pi_{c_j}(c_i)$ is a point. Lemma 6.1 and $c_i \cap K \neq \varnothing$ yield that $\pi_{c_j}(c_i) \subset \pi_{c_j}(K) = c_j \cap K = [0, n_j - 1]$; thus

$$(6\text{-}2) \qquad\qquad\qquad c_i \cap X_j = \varnothing.$$

Similarly, $c_i \cap X_j = \varnothing$. Let $h = \pi_{c_j}^{-1}\left(-\frac{1}{2}\right)$ be the boundary of X_j^-, and let N_h be the carrier of h. Then (6-1) implies that h has empty intersection with any hyperplane dual to c_i, and so does N_h. It follows from Lemma 6.1 that $\pi_{c_i}(h) = \pi_{c_i}(N_h) = p$ is a vertex in c_i. If $h_1 = \pi_{c_i}^{-1}\left(p - \frac{1}{2}\right)$ and $h_2 = \pi_{c_i}^{-1}\left(p + \frac{1}{2}\right)$ are two hyperplanes that pinch p, then $h \cap h_k = \varnothing$ for $k = 1, 2$. This and (6-2) yield $X_j^- \cap h_k = \varnothing$; hence $\pi_{c_i}(X_j^-) = p$ by Lemma 6.1. Similarly, $\pi_{c_i}(X_j^+) = p$, so

$$p = \pi_{c_i}(X_j) = \pi_{c_i}(c_j) \subset \pi_{c_i}(K) = c_i \cap K = [0, n_i - 1].$$

Note $(\pi_{c_i} \circ v_i^k)(X_j) = (v_i^k \circ \pi_{c_i})(X_j) = v_i^k(p) = p + k\,n_i$, so $v_i^k(X_j) \subset X_i$ for $k \neq 0$. $\square$

The discussion in this subsection yields a well-defined map

$$\Theta_S \colon \mathrm{CN}(\Gamma, S) \to \{\text{finite-index RAAG subgroups of } G(\Gamma)\}.$$

The images of Θ_S are called S–*special subgroups* of $G(\Gamma)$. A subgroup of $G(\Gamma)$ is *special* if it is S–special for some standard generating set S of $G(\Gamma)$.

6.2 Rigidity of RAAG subgroups

In this subsection, we will assume $G(\Gamma')$ is a finite-index RAAG subgroup in $G(\Gamma)$ and $\mathrm{Out}(G(\Gamma))$ is finite. We will show that under such conditions, $G(\Gamma')$ must arise from the process described in the previous subsection. We will prove this in three steps. First we produce a convex subcomplex of $X(\Gamma, S)$ from $G(\Gamma')$. Then we will modify this convex subcomplex such that it is an element in $\mathrm{CN}(\Gamma, S)$. Thus we have defined a map from finite-index RAAG subgroups of $G(\Gamma)$ to elements in $\mathrm{CN}(\Gamma, S)$. In the last step, we show the map defined in Step 2 is an inverse to the map Θ_S defined in Section 6.1.

Also near the end of this subsection, we will leave several relatively long remarks which discuss relevant material in the literature. The reader can skip these remarks at first reading.

Recall that $\mathrm{Out}(G(\Gamma))$ is finite and $\mathrm{Out}(G(\Gamma'))$ is transvection-free (Theorem 5.3), so any two standard generating sets of $G(\Gamma)$ (or $G(\Gamma')$) differ by a sequence of

conjugations or partial conjugations. Then given any two standard generating sets S and S_1 for $G(\Gamma)$, there is a canonical way to identify $\mathcal{P}(\Gamma, S)$ and $\mathcal{P}(\Gamma, S_1)$ (every S–geodesic is Hausdorff close to an S_1–geodesic). Thus we will write $\mathcal{P}(\Gamma)$ and $\mathcal{P}(\Gamma')$ and omit the generating set.

Lemma 6.6 *Let ϕ and s be as in the discussion before Lemma 5.5. Let $l \subset X(\Gamma)$ and $l' \subset X(\Gamma')$ be standard geodesics such that $\phi(v(l)) = v(l')$. Then $\phi \circ \pi_{\Delta(l)} = \pi_{\Delta(l')} \circ s$.*

Proof Pick standard geodesics $r \subset X(\Gamma)$ and $r' \subset X(\Gamma')$ such that $\phi(v(r)) = v(r')$; then $s(\Delta(r)) = \Delta(r')$ by Lemma 4.10 (recall that r is the intersection of maximal standard flats). Therefore, by the definition of $\pi_{\Delta(l)}$, it suffices to show $\phi \circ \pi_l(x) = \pi_{l'} \circ \phi(x)$ for any vertex $x \in X(\Gamma)$. Let y be a vertex such that $y \notin l$, and let $x = \pi_l(y)$. By Lemma 6.1, we can approximate $\overline{xy}$ by a combinatorial geodesic ω in the 1–skeleton of $\pi_l^{-1}(y)$; then no hyperplane could intersect both l and ω. Let $\{v_i\}_{i=0}^n$ be vertices in ω such that for $0 \leq i \leq n-1$, we have that each $[v_i, v_{i+1}]$ is a maximal subsegment of ω that is contained in a standard geodesic ($v_0 = x$ and $v_n = y$). Denote the corresponding standard geodesic by l_i. Then $\Delta(l) \neq \Delta(l_i)$ for all i. Let $u_i = \phi(v_i)$ and let l_i' be the standard geodesic such that $\phi(v(l_i)) = v(l_i')$. Then $\overline{u_i u_{i+1}} \subset l_i'$ and $\Delta(l') \neq \Delta(l_i')$ for all i; thus $\pi_{l'}(l_i')$ is a point by Corollary 3.2, and $\pi_{l'}(u_i) = \pi_{l'}(u_j)$ for all $1 \leq i, j \leq n$. $\square$

Step 1 We produce a convex subcomplex of $X(\Gamma, S)$ from $G(\Gamma')$.

The left action $G(\Gamma) \curvearrowright G(\Gamma)$ induces $G(\Gamma') \curvearrowright G(\Gamma)$ and $G(\Gamma') \curvearrowright X(\Gamma, S)$. By choosing a standard generating set S' of $G(\Gamma')$, we have left action $G(\Gamma') \curvearrowright X(\Gamma', S')$. For $h \in G(\Gamma')$, we use ϕ_h, $\overline{\phi}_h$, s_h and $\overline{s}_h$ to denote the action of h on $G(\Gamma)$, $G(\Gamma')$, $\mathcal{P}(\Gamma)$ and $\mathcal{P}(\Gamma')$ respectively. Pick a $G(\Gamma')$–equivariant quasi-isometry $q: X(\Gamma, S) \to X(\Gamma', S')$ such that $q|_{G(\Gamma')} = \mathrm{Id}$. By Theorem 5.3 and Lemma 4.10, q induces surjective $G(\Gamma')$–equivariant maps $\phi: G(\Gamma) \to G(\Gamma')$ and $s: \mathcal{P}(\Gamma) \to \mathcal{P}(\Gamma')$. Note that ϕ depends on the choice of generating set S and S', and this flexibility comes from the automorphism groups of $G(\Gamma)$ and $G(\Gamma')$.

The key of Step 1 is to choose a "nice" standard generating set S' of $G(\Gamma')$ such that ϕ behaves like φ in Theorem 5.6.

Lemma 6.7 *By choosing a possibly different standard generating set S' for $G(\Gamma')$, we can assume the map ϕ satisfies $\phi(\mathrm{id}) = \mathrm{id}$, where id denotes the identity element in the corresponding group.*

Proof Assume $\phi(\mathrm{id}) = a \neq \mathrm{id}$; we claim if we change the generating set from S' to $aS'a^{-1}$, then the resulting ϕ will satisfy our requirement. By the construction

of ϕ, it suffices to show for any maximal S'–flat F_1' such that $a \in F_1'$, there exists a maximal $aS'a^{-1}$–flat F_2' such that id $\in F_2'$ and $d_H(F_1', F_2') < \infty$. Let us assume $F_1' = \{ag^k\}_{k \in \mathbb{Z}}$ for some $g \in S'$. Then $F_2' = \{(aga^{-1})^k\}_{k \in \mathbb{Z}}$ would satisfy the required condition. We can prove the general case in a similar way.   □

Pick a standard geodesic $l \subset X(\Gamma, S)$; we want to flip the order of points of l in a $G(\Gamma')$–equivariant way such that (5-7) is true. We choose an order-preserving identification of $v(l)$ and $\mathbb{Z}$. Let $d = |\phi^{-1}(\phi(p)) \cap v(l)|$ where p is a vertex in $v(l)$. Let Stab$(v(l))$ be the stabilizer of $v(l)$ under the action $G(\Gamma') \curvearrowright G(\Gamma)$. By the second paragraph of the proof of Theorem 5.6, d does not depend on the choice of p in $v(l)$, and Stab$(v(l))$ acts on $v(l)$ in the same way as $d\mathbb{Z}$ acts on $\mathbb{Z}$ (recall that ϕ is $G(\Gamma')$–equivariant and the action of $G(\Gamma')$ on $G(\Gamma)$ is induced from the left action of $G(\Gamma)$ on itself).

We will write $\chi(l) = d$. If $\bar{l}$ and l are parallel, then $\chi(l) = \chi(\bar{l})$. Thus $\chi: \mathcal{P}(\Gamma) \to \mathbb{Z}$ is well defined. Since $\chi(l)$ only depends on how Stab$(v(l))$ acts on $v(l)$, we see that χ does not depend on the standard generating set S'. However, χ descends to $\chi: S' \to \mathbb{Z}$ for any choice of S' by the $G(\Gamma')$–equivariance of ϕ.

Let $\phi(0) = a$. Then Stab$(v(l))$ is generated by aha^{-1} for some $h \in S'$. By the same reasoning as in Lemma 6.7, we can assume $a =$ id. Let $S' = \{h_\lambda\}_{\lambda \in \Lambda}$. For each $h_\lambda \in S'$, we associated an integer n_λ as follows. If $h_\lambda h = h h_\lambda$, we set $n_\lambda = 0$. Now we consider the case where $h_\lambda h \neq h h_\lambda$. Let $l_\lambda' \subset X(\Gamma', S')$ be the standard geodesic that contains all powers of h_λ, and let $b_\lambda = \pi_{\Delta(l)} \circ s^{-1}(\Delta(l_\lambda'))$ ($\pi_{\Delta(l)}$ is the map in Lemma 6.2). Then n_λ is defined to be the unique integer such that $b_\lambda + n_\lambda d \in [0, d-1]$ (recall that $d = \chi(l)$). Define $f: S' \to G(\Gamma')$ by sending h_λ to $h^{n_\lambda} h_\lambda h^{-n_\lambda}$; then f extends to an automorphism of $G(\Gamma')$, and $S'' = \{f(h_\lambda)\}_{\lambda \in \Lambda}$ is also a standard generating set. Indeed, if $\Delta(l_{\lambda_1}')$ and $\Delta(l_{\lambda_2}')$ stay in the same connected component of $\mathcal{P}(\Gamma') \setminus \mathrm{St}(\Delta(l))$, then $b_{\lambda_1} = b_{\lambda_2}$ by Lemma 6.2; hence $n_{\lambda_1} = n_{\lambda_2}$. It follows that f can be realized as a composition of partial conjugations.

Lemma 6.8 *Define* ϕ_1 *by replacing* S' *by* S'' *in the definition of* ϕ. *Then* $\phi_1|_{v(l_1)}$ *satisfies* (5-7) *for any standard geodesic* $l_1 \subset X(\Gamma, S)$ *with* $\Delta(l_1) \in \{s_h(\Delta(l))\}_{h \in G(\Gamma')}$.

Recall that for any $h \in G(\Gamma')$, we use s_h to denote the action of h on $\mathcal{P}(\Gamma)$.

Proof It suffices to show $\phi_1|_{v(l)}$ satisfies (5-7), and the rest follows from the $G(\Gamma')$–equivariance of ϕ_1. To show this, we only need to prove $\phi_1(i) =$ id for any $i \in [0, d-1]$. Let Λ, b_λ and n_λ be as above.

We pick $i \in [0, d-1]$. Then there exists $\lambda \in \Lambda$ such that $b_\lambda + n_\lambda d = i$. By Lemma 6.6, $\phi(b_\lambda) =$ id; hence $\phi(i) = h^{n_\lambda}$. Let l_i be a standard geodesic such

that $b_\lambda \in l_i$ and $d(\Delta(l_i), \Delta(l)) \geq 2$. Then there exists $h_{\lambda'} \in S'$ with $b_{\lambda'} = b_\lambda$ such that $\phi(v(l_i)) = \{h_{\lambda'}^k\}_{k\in\mathbb{Z}}$. Then $(\phi_h)^{n_\lambda}(v(l_i))$ is an S–geodesic passing through i, and $(\phi \circ (\phi_h)^{n_\lambda})(v(l_i)) = ((\bar{\phi}_h)^{n_\lambda} \circ \phi)(v(l_i)) = \{h^{n_\lambda} h_{\lambda'}^k\}_{k\in\mathbb{Z}}$. Note that

$$(6\text{-}3) \qquad d_H\big(\{h^{n_\lambda} h_{\lambda'}^k\}_{k\in\mathbb{Z}}, \{(f(h_{\lambda'}))^k\}_{k\in\mathbb{Z}}\big) < \infty.$$

Now we look at the new map ϕ_1. Note that $\phi_1(0) = \mathrm{id}$ is still true. Moreover, (6-3) and Lemma 6.6 imply $\phi_1(i) = \mathrm{id}$. Thus the lemma follows. $\qquad\square$

The next lemma basically says the above change-of-basis process does not significantly affect other geodesics.

Lemma 6.9 *Let r be a standard geodesic in $X(\Gamma, S)$ which satisfies the condition that $\Delta(r) \notin \{s_h(\Delta(l))\}_{h\in G(\Gamma')}$. Pick two different vertices $x_1, x_2 \in r$. If $\phi(x) = \phi(y)$, then $\phi_1(x) = \phi_1(y)$.*

Proof For $i = 1, 2$, let $r_i \subset X(\Gamma, S)$ be a standard geodesic containing x_i such that $d(\Delta(r_i), \Delta(r)) \geq 2$ for $i = 1, 2$. Let r' (resp. r'') be an S'–geodesic (resp. S''–geodesic) such that $\phi(v(r)) = v(r')$ (resp. $\phi_1(v(r)) = v(r'')$). Let $\alpha = \phi(x) = \phi(y)$. Then there exist elements h_λ, h_{λ_1} and h_{λ_2} in S' such that $\phi(v(r_i)) = \{\alpha h_{\lambda_i}^k\}_{k\in\mathbb{Z}}$ for $i = 1, 2$, and $r' = \{\alpha h_\lambda^k\}_{k\in\mathbb{Z}}$. Note that

$$(6\text{-}4) \qquad h \neq h_\lambda, \quad h_{\lambda_1} \neq h_\lambda \quad \text{and} \quad h_{\lambda_2} \neq h_\lambda.$$

Recall that h is the generator of $\mathrm{Stab}(v(l))$. The first inequality of (6-4) follows from $\Delta(r) \notin \{s_h(\Delta(l))\}_{h\in G(\Gamma')}$.

It suffices to show there exist S''–geodesics r_1'' and r_2'' such that

$$(6\text{-}5) \qquad d_H\big(\phi(v(r_i)), r_i''\big) < \infty$$

for $i = 1, 2$, and

$$(6\text{-}6) \qquad \pi_{\Delta(r'')}(\Delta(r_1'')) = \pi_{\Delta(r'')}(\Delta(r_2'')),$$

then $\phi_1(x) = \phi_1(y)$ follows from Lemma 6.6. Define $r_i'' = \{\alpha h^{-n_{\lambda_i}} (f(h_{\lambda_i}))^k\}_{k\in\mathbb{Z}}$; then (6-5) is immediate. Note that for any $a \in r_1'$ and $b \in r_2'$, we have

$$b = a \cdot (f(h_{\lambda_1}))^{k_1} \cdot h^{n_{\lambda_1} - n_{\lambda_2}} \cdot (f(h_{\lambda_2}))^{k_2}$$

for some $k_1, k_2 \in \mathbb{Z}$; then (6-6) follows from (6-4) and the definition of $\pi_{\Delta(r'')}$. $\qquad\square$

Similarly, we can prove that if we change ϕ with respect to the conjugation $S' \to aS'a^{-1}$, then Lemma 6.9 is still true with r being an arbitrary standard geodesic.

By Lemma 6.8 and Lemma 6.9, we can apply the above change-of-basis procedure finitely many times to find an appropriate standard generating set S' of $G(\Gamma')$ such that the corresponding map ϕ satisfies (5-7) when restricted to any standard geodesic in $X(\Gamma, S)$. By the proof of Theorem 5.6, we can extend ϕ to a cubical map $\phi \colon X(\Gamma, S) \to X(\Gamma', S')$ such that combinatorial geodesics in $C(\Gamma, S)$ are mapped to combinatorial geodesics in $C(\Gamma', S')$. Thus $\phi^{-1}(\mathrm{id})$ is a combinatorially convex subcomplex. The subcomplex $\phi^{-1}(\mathrm{id})$ is also compact since $\phi^{-1}(\mathrm{id})$ contains finitely many vertices. Recall that combinatorial convexity in ℓ^1 metric and convexity in CAT(0) metric are the same for subcomplexes of CAT(0) cube complexes [33], so we have constructed a compact convex subcomplex $\phi^{-1}(\mathrm{id}) \subset X(\Gamma, S)$ from a given finite-index RAAG subgroup $G(\Gamma') \leq G(\Gamma)$.

Step 2 We show $\phi^{-1}(\mathrm{id})$ can be assumed to be an element in $\mathrm{CN}(\Gamma, S)$.

For $K \subset G(\Gamma)$, denote the union of all standard geodesics in $X(\Gamma, S)$ that have nontrivial intersection with K by K^*. K is S–*convex* if and only if K is the vertex set of some convex subcomplex in $X(\Gamma, S)$. Now we return to ϕ. By Step 1, we can assume $\phi(\mathrm{id}) = \mathrm{id}$, and $\phi^{-1}(y)$ is S–convex for any $y \in G(\Gamma')$.

Step 2.1 Let $\{l_i\}_{i=1}^q$ be the collection of standard geodesics passing through id, and let $\Lambda_1 = \{\mathrm{id}\}$. Let $I \colon G(\Gamma) \to \ell^1(v(\mathcal{P}(\Gamma, S)))$ and $I_{\Delta(l)} \colon G(\Gamma) \to \mathbb{Z}^{\Delta(l)}$ be the maps defined in Section 6.1. Since $v(l_i)$ and $v(l_j)$ are in different $G(\Gamma')$–orbits for $i \neq j$, by Lemma 6.8 and Lemma 6.9, we can apply the change-of-basis procedure in Step 1 to find a standard generating set S' for $G(\Gamma')$ such that for each $1 \leq i \leq q$,

$$(6\text{-}7) \qquad\qquad I_{\Delta(l_i)}^{-1}\big([0, \chi(l_i) - 1]\big) \cap v(l_i) \subset \phi^{-1}(\mathrm{id}).$$

Step 2.2 Let $\Lambda_2 = \Lambda_1^* \cap \phi^{-1}(\mathrm{id})$. Pick a vertex $x \in \Lambda_2 \setminus \Lambda_1$ (if such x does not exist, then our process terminates). Let l be a standard geodesic such that $x \in l$. If l is parallel to some l_i in Step 2.1, then (6-7) with l_i replaced by l is automatically true without any modification on S', because both I and ϕ respect the product structure of P_{l_i}. If l is not parallel to any l_i, then $I_{\Delta(l)}(x) = 0$. Moreover, $\Delta(l)$ is not in the $G(\Gamma')$–orbits of the $\Delta(l_i)$, so we can modify S' as before such that both (6-7) and $I_{\Delta(l)}^{-1}([0, \chi(l) - 1]) \cap v(l) \subset \phi^{-1}(\mathrm{id})$ are true. We deal with other standard geodesics passing through x and other points in $\Lambda_2 \setminus \Lambda_1$ in a similar way.

Step 2.3 Let $\Lambda_3 = \Lambda_2^* \cap \phi^{-1}(\mathrm{id})$. For each vertex in $\Lambda_3 \setminus \Lambda_2$, we repeat the procedure in Step 2.2. Then we can define $\Lambda_4, \Lambda_5, \dots$. Since $|\phi^{-1}(\mathrm{id})|$ is finite and this number does not change after adjusting S', our procedure must terminate after finitely many steps. Since $\phi^{-1}(\mathrm{id})$ remains connected in each step, once the procedure terminates, we must have already dealt with each point in $\phi^{-1}(\mathrm{id})$ and each standard geodesic passing

through each point in $\phi^{-1}(\mathrm{id})$. By construction, the resulting ϕ satisfies $\mathrm{id} \in \phi^{-1}(\mathrm{id})$ and $I_{\Delta(l)}^{-1}([0, \chi(l)-1]) \cap v(l) \subset \phi^{-1}(\mathrm{id})$ for each standard geodesic l which intersects $\phi^{-1}(\mathrm{id})$. Thus $\phi^{-1}(\mathrm{id})$ is nonnegative.

Note that the sets Λ_i actually do not depend on the map ϕ from step $i-1$. They only depend on the map $\chi \colon v(\mathcal{P}(\Gamma)) \to \mathbb{Z}$. Thus the nonnegative subset $\phi^{-1}(\mathrm{id}) \subset G(\Gamma)$ produced above depends only on S and the subgroup $G(\Gamma') \leq G(\Gamma)$. Then we have a well-defined map

$$\Xi_S \colon \{\text{Finite-index RAAG subgroups of } G(\Gamma)\} \to \mathrm{CN}(\Gamma, S).$$

Step 3 We show Ξ_S is an inverse to the map Θ_S defined in Section 5.2.

First we prove $\Theta_S \circ \Xi_S = \mathrm{Id}$. Let $K = \Xi_S(G(\Gamma'))$. Let S' be the corresponding standard generating set for $G(\Gamma')$ and let $\phi \colon G(\Gamma) \to G(\Gamma')$ be the corresponding map. We find a maximal collection of standard geodesics $\{c_i\}_{i=1}^s$ such that $c_i \cap K \neq \varnothing$ for all i and $\Delta(c_i) \neq \Delta(c_j)$ for any $i \neq j$. Let $n_i = \chi(c_i)$, and let $g_i \in S$ be the label of edges in c_i. Suppose $\alpha_i = \pi_{c_i}(\mathrm{id})$ where $\pi_{c_i} \colon X(\Gamma, S) \to c_i$ is the CAT(0) projection. Then it suffices to prove the following lemma.

Lemma 6.10 $$S' = \{\alpha_i g_i^{n_i} \alpha_i^{-1}\}_{i=1}^s.$$

Proof Pick $h \in S'$ and let $c_h \subset X(\Gamma', S')$ be the standard geodesic containing id and h. Then there exists a unique i such that $\phi(v(c_i)) = c_h$. To see this, let c be a standard geodesic in $X(\Gamma, S)$ such that $s(\Delta(c)) = \Delta(c_h)$. Then $\phi(v(c))$ and c_h are parallel and there exists $u \in G(\Gamma')$ which sends $\phi(v(c))$ to $v(c_h)$. Thus $\phi \circ \phi_u(v(c)) = v(c_h)$ by (5-4), where ϕ_u is defined in the beginning of Step 1. Note that $\phi_u(v(c))$ has nontrivial intersection with K. We choose c_i to be the geodesic parallel to $\phi_u(v(c))$. Then $\phi(v(c_i)) = v(c_h)$.

For any standard geodesic c_i' parallel to c_i, we have that $\phi(c_i')$ is parallel to c_h, so $h \in \mathrm{Stab}(v(\phi(c_i'))) = \mathrm{Stab}(v(c_i'))$. It follows that ϕ_h stabilizes the parallel set P_{c_i} and acts by translation along the c_i–direction. Note that $(I_{\Delta(c_i)} \circ \phi_h)(x) = I_{\Delta(c_i)}(x) + \chi(c_i)$ for any $x \in v(P_{c_i})$, so $h = \phi_h(\mathrm{id}) = \alpha_i g_i^{n_i} \alpha_i^{-1}$ and the claim follows. $\square$

It remains to show $\Xi_S \circ \Theta_S = \mathrm{Id}$. The following result implies this.

Lemma 6.11 *Let Γ be an arbitrary finite simplicial graph. Pick a standard generating set S for $G(\Gamma)$ and $K \in \mathrm{CN}(\Gamma, S)$. Let $G(\Gamma') = \Theta_S(K)$ and let S' be the corresponding generating set. Suppose $q \colon G(\Gamma) \to G(\Gamma')$ is a $G(\Gamma')$–equivariant quasi-isometry such that $q|_{G(\Gamma')}$ is the identity map. Then:*

(1) *q* induces a simplicial isomorphism $q_*: \mathcal{P}(\Gamma, S) \to \mathcal{P}(\Gamma', S')$.

(2) q_* induces a $G(\Gamma')$–equivariant retraction $r: G(\Gamma) \to G(\Gamma')$ such that r sends every S–flat to an S'–flat.

(3) r extends to a surjective cubical map $r: X(\Gamma, S) \to X(\Gamma', S')$ with $r^{-1}(\mathrm{id}) = K$. In particular, the vertex set of K is the strict fundamental domain for the left action $G(\Gamma') \curvearrowright G(\Gamma)$.

Proof It suffices to prove the case when Γ does not admit a nontrivial join decomposition and Γ is not a point.

By the construction of Θ_S, we know the q–image of any S–flat which intersects K is Hausdorff close to an S'–flat which contains the identity. Moreover, if the S–flat is maximal, then the corresponding S'–flat is unique. Since $G(\Gamma') \cdot v(K) = G(\Gamma)$, the equivariance of q implies the q–image of every S–flat is Hausdorff close to an S'–flat. Since q is a quasi-isometry, images of parallel S–geodesics are Hausdorff closed to each other. This induces $q_*: \mathcal{P}(\Gamma, S) \to \mathcal{P}(\Gamma', S')$, which is injective since q is a quasi-isometry, and surjective by the $G(\Gamma')$–equivariance.

Pick $x \in G(\Gamma)$, and let $\{F_i\}_{i \in I}$ be the collection of maximal S–flats containing x. For each i, let F_i' be the unique maximal S'–flat such that $d_H(q(F_i), F_i') < \infty$. Note that $\bigcap_{i \in I} F_i = x$ by our assumption on Γ. So $\bigcap_{i \in I} F_i'$ is either empty or one point. Note that if $x \in K$, then $\bigcap_{i \in I} F_i' = \mathrm{id}$. The equivariance of q_* implies that for every x, $\bigcap_{i \in I} F_i'$ is a point, which is defined to be $r(x)$. It is clear that $v(K) \subset r^{-1}(\mathrm{id})$, but $|G(\Gamma): G(\Gamma')| \le |v(K)|$, so $v(K) = r^{-1}(\mathrm{id})$. It follows that $v(K)$ is the strict fundamental domain for the left action of $G(\Gamma')$, and r is a $G(\Gamma')$–equivariant map which maps $v(K)$ to id.

Note that $r(\mathrm{id}) = \mathrm{id}$. Then the $G(\Gamma')$–equivariance of r implies $r(g) = g$ for any $g \in G(\Gamma') \subset G(\Gamma)$. Thus r is a retraction. Similarly, by using the $G(\Gamma')$–equivariance of r, we deduce that r sends every S–flat that intersects K to an S'–flat passing through the identity element of $G(\Gamma')$. Thus r sends every S–flat to an S'–flat by the equivariance of r. It is easy to see r extends to a cubical map $r: X(\Gamma, S) \to X(\Gamma', S')$ such that $r^{-1}(\mathrm{id}) = K$. □

Remark 6.12 We can generalize some of the results in Lemma 6.11 to infinite convex subcomplexes of $X(\Gamma, S)$. A convex subcomplex $K \subset X(\Gamma, S)$ is *admissible* if for any standard geodesic l, the CAT(0) projection $\pi_l(K)$ is either a finite interval or the whole of l (a ray is not allowed). Let $\{l_\lambda\}_{\lambda \in \Lambda}$ be a maximal collection of standard geodesics such that (1) $l_\lambda \cap K \ne \varnothing$; (2) l_λ and $l_{\lambda'}$ are not parallel for $\lambda \ne \lambda'$; (3) $\pi_{l_\lambda}(K)$ is a finite interval. For each l_λ, let $\alpha_\lambda \in G(\Gamma)$ be an element which translates along l_λ with translation length $= 1 + \mathrm{length}(\pi_{l_\lambda}(K))$. Let G_K be the subgroup generated by $S' = \{\alpha_\lambda\}_{\lambda \in \Lambda}$. If K is admissible, we can prove $G_K \cdot v(K) = G(\Gamma)$ as

before. Moreover, for any finite subset $S_1' \subset S'$, the subgroup G_1 generated by S_1' is a right-angled Artin group, and $G_1 \hookrightarrow G_K$ is an isometric embedding with respect to the word metric. We can define an S'–flat as before and view each vertex of G_K as a 0–dimensional S'–flat.

Now we show $v(K)$ is a strict fundamental domain for the action $G_K \curvearrowright G(\Gamma)$. It suffices to show $\alpha(K) \cap K = \varnothing$ for each nontrivial $\alpha \in G_K$. We can assume there is a right-angled Artin group G_1 such that $\alpha \in G_1 \subset G_K$. Let $\alpha = w_1 w_2 \cdots w_n$ be a canonical form of α; see [16, Section 2.3]. Then:

(1) Each w_i belongs to an abelian standard subgroup of G_1.

(2) For each i, let $w_i = r_{i,1}^{k_{i,1}} r_{i,2}^{k_{i,2}} \cdots r_{i,n_i}^{k_{i,n_i}}$ ($r_{i,j} \in S'$). Then for each $r_{i+1,j}$ ($1 \le j \le n_{i+1}$), there exists $r_{i,j'}$ which does not commute with $r_{i+1,j}$.

We associate each generator $r_{i,j}$ with a subset $X_{i,j} \subset X(\Gamma, S)$ as in the proof of Lemma 6.4, and claim there exists j with $1 \le j \le n_1$ such that $\alpha(K) \subset X_{1,j}$; then $\alpha(K) \cap K = \varnothing$ follows. We prove by induction on n and assume $w_2 w_3 \cdots w_n(K) \subset X_{2,j'}$. By (2), there is $r_{1,j}$ such that $r_{1,j}$ and $r_{2,j'}$ does not commute, so we have $r_{1,j}^{k_{1,j}}(X_{2,j'}) \subset X_{1,j}$. Moreover, by (1), $r_{1,h}^{k_{1,h}}(X_{1,j}) = X_{1,j}$ for $h \ne j$, so $\alpha(K) \subset w_1(X_{2,j'}) \subset X_{1,j}$.

Now we can define a G_K–equivariant map $r \colon G(\Gamma) \to G_K$ by sending $v(K)$ to the identity of G_K. We prove as before that r maps S–flats to (possibly lower-dimensional or 0–dimensional) S'–flats; thus r is 1–Lipschitz with respect to the word metric. Let $i \colon G_K \hookrightarrow G(\Gamma)$ be the inclusion. Then by the equivariance of r, the composition $r \circ i$ is a left translation of G_K. In particular, if K contains the identity, then r is a retraction. It follows that if S' is finite, then i is a quasi-isometric embedding.

Note that a related construction in the case of right-angled Coxeter groups has been discussed in [34]. By taking larger and larger convex compact subcomplexes of $X(\Gamma, S)$, we know $G(\Gamma)$ is residually finite. Moreover, pick $\beta \in \mathrm{Stab}(K) \subset G(\Gamma)$. By definition of S', we have $S' = \beta S' \beta^{-1}$, so $\mathrm{Stab}(K)$ normalizes G_K. Now we have obtained a direct proof of the fact that every word-quasiconvex subgroup of a finitely generated right-angled Artin group is separable (Theorem F of [34]) by using the above discussion together with the outline in Section 1.5 of [34].

The following result follows readily from the above discussion.

Theorem 6.13 *Let $G(\Gamma)$ be a RAAG with $\mathrm{Out}(G(\Gamma))$ finite. We pick a standard generating set S for $G(\Gamma)$. Then there is a one-to-one correspondence between nonnegative convex compact subcomplexes of $X(\Gamma, S)$ that contain the identity and finite-index RAAG subgroups of $G(\Gamma)$. In particular, these subgroups are generated by conjugates of powers of elements in S.*

In particular, Theorem 1.3 in the introduction follows from Theorem 6.13.

Remark 6.14 If we drop the finite automorphism group assumption in the above theorem, then there exist a RAAG $G(\Gamma_1)$ and its finite index RAAG subgroup $G(\Gamma_2)$ such that $G(\Gamma_2)$ is not isomorphic to any special subgroup of $G(\Gamma_1)$. To see this, let $G(\Gamma_1)$ be a right-angled Artin group such that $\mathrm{Out}(G(\Gamma_1))$ is transvection-free. Then Lemma 6.11 and Theorem 3.20 imply each special subgroup of $G(\Gamma_1)$ does not admit a nontrivial transvection in its outer automorphism group. Let Γ_1 and Γ_2 be the graphs in Example 3.22. Then $G(\Gamma_2)$ is a right-angled Artin subgroup of $G(\Gamma_1)$, and there are nontrivial transvections in $\mathrm{Out}(G(\Gamma_2))$. Thus $G(\Gamma_2)$ is not isometric to any special subgroup of $G(\Gamma_1)$.

Remark 6.15 Pick $G(\Gamma)$ such that $\mathrm{Out}(G(\Gamma))$ is finite; then Theorem 6.13 can be used to show a certain subgroup of $G(\Gamma)$ is not a RAAG. For example, let $\{v_i\}_{i=1}^k$ be a subset of some standard generating set for $G(\Gamma)$. We define a homomorphism $h\colon G(\Gamma) \to \mathbb{Z}/2$ by sending each v_i to the nontrivial element in $\mathbb{Z}/2$ and killing all other generators. Then $\ker(h)$ is a RAAG if and only if $k = 1$. One can compare this example to Example 3.22.

Remark 6.16 It is shown in [42, Theorem 2] that if $F(\Gamma')$ embeds into $\mathcal{P}(\Gamma)$ as a full subcomplex, then there exists a monomorphism $G(\Gamma') \hookrightarrow G(\Gamma)$. This result can be recovered by our previous discussion as follows. Let Γ be an arbitrary finite simplicial graph. Let S be a standard generating set for $G(\Gamma)$. For any vertex $w \in \mathcal{P}(\Gamma)$, let $\alpha_w \in G(\Gamma)$ be a conjugate of some element in S such that $\alpha_w(l) = l$ for every standard geodesic $l \subset X(\Gamma, S)$ with $\Delta(l) = w$.

Suppose $M \subset \mathcal{P}(\Gamma, S)$ is a compact full subcomplex and Γ' is the 1–skeleton of M. Denote the vertex set of M by $\{w_i\}_{i=1}^n$, and let l_i be a standard geodesic with $\Delta(l_i) = w_i$. We identify each l_i in an orientation-preserving way with $\mathbb{R}$ such that $0 \in \mathbb{R}$ is identified with $\pi_{l_i}(\mathrm{id}) \subset l_i$, where π_{l_i} is the CAT(0) projection to l_i and id is the identity element of $G(\Gamma)$.

For $1 \le i \le n$, define $\Lambda_i = \{1 \le j \le n \mid d(w_i, w_j) \ge 2\}$. For each i, we define a pair of integers a_i and k_i as follows. If $\Lambda_i \neq \varnothing$, then let $[a_i, a_i + k_i] \subset \mathbb{R}$ be the minimal interval such that $\bigcup_{j \in \Lambda_i} \pi_{l_i}(l_j) \subset [a_i, a_i + k_i]$ (recall that l_i is identified with $\mathbb{R}$). If $\Lambda_i = \varnothing$, then we pick an arbitrary a_i and set $k_i = 0$. Define $X_i = \pi_{c_i}^{-1}\big((-\infty, a_i - \tfrac{1}{2}]\big) \cup \pi_{c_i}^{-1}\big([a_i + k_i + \tfrac{1}{2}, \infty)\big)$. Then by construction, $X_i \cap X_j = \varnothing$ for i, j satisfying $d(w_i, w_j) \ge 2$. Using the argument in Section 6.1, we can show the subgroup generated by $S' = \{\alpha_{w_i}^{k_i + 1}\}_{i=1}^n$ is a RAAG with defining graph Γ'.

At this point it is natural to ask the following question.

Question 6.17 *Let S be a standard generating set of $G(\Gamma)$, and let S' be a finite collection of elements of the form $\alpha r^k \alpha^{-1}$, where $r \in S$, $k \in \mathbb{Z}$ and $\alpha \in G(\Gamma)$. Suppose G is the subgroup generated by S'. Is G a right-angled Artin group?*

6.3 Generalized star extension

Our goal in this subsection is to find an algorithm to determine whether $G(\Gamma)$ and $G(\Gamma')$ are quasi-isometric or not, given that $\text{Out}(G(\Gamma))$ is finite.

For a convex subcomplex $E \subset X(\Gamma)$, we denote the full subcomplex in $\mathcal{P}(\Gamma, S)$ spanned by $\{\Delta(l_\lambda)\}_{\lambda \in \Lambda}$ by $\widehat{E}$, where $\{l_\lambda\}_{\lambda \in \Lambda}$ is the collection of standard geodesics in $X(\Gamma)$ with $l_\lambda \cap E \neq \varnothing$.

Now we describe a process to construct a graph Γ' from Γ such that $G(\Gamma')$ is isomorphic to a special subgroup of $G(\Gamma)$. Let $\Gamma_1 = \Gamma$, and let K_1 be one point. We will construct a pair (Γ_i, K_i) inductively such that:

(1) K_i is a compact CAT(0) cube complex, and there is a cubical embedding $f \colon K_i \to X(\Gamma)$ such that $f(K_i)$ is convex in $X(\Gamma)$.

(2) Γ_i is a finite simplicial graph, and there is a simplicial isomorphism $g \colon F(\Gamma_i) \to \widehat{f(K_i)}$.

Note that these assumptions are true for $i = 1$.

We associate each edge $e \subset K_i$ with a vertex in Γ_i, denoted by v_e, as follows. Let l_e be the standard geodesic in $X(\Gamma)$ that contains $f(e)$. We define $v_e := g^{-1}(\Delta(l_e))$. Each vertex $x \in K_i$ can be associated with a full subcomplex $\Phi(x) \subset F(\Gamma_i)$ defined by $\Phi(x) = g^{-1}(\widehat{x})$.

To define (Γ_{i+1}, K_{i+1}), pick a vertex $v \in \Gamma_i$, and let $\{x_j\}_{j=1}^m$ be the collection of vertices in K_i such that $v \in \Phi(x_j)$. Then $\{f(x_j)\}_{j=1}^k$ are exactly the vertices in $P_l \cap f(K_i)$, where l is a standard geodesic such that $\Delta(l) = g(v)$. Let L be the convex hull of $\{x_j\}_{j=1}^m$ in K_i. Then $e \subset L$ for any edge $e \subset K_i$ with $v_e = v$.

Since $f(L) = P_l \cap f(K_i)$, the natural product decomposition $P_l \cong l \times l^\perp$ induces a product decomposition of $L = h \times [0, a]$. Note that it is possible that $a = 0$, and $a > 0$ if and only if there exists an edge $e \subset K_i$ with $v_e = v$. If $a > 0$, then h is isomorphic to the hyperplane dual to e, and for any edge $e' \in K_i$ with $v_{e'} = v$, the projection of e' to the interval factor $[0, a]$ is an edge.

Let $L_i = h \times \{a\} \subset L$, and let $M_i = \bigcup_{x \in L_i} \Phi(x)$ (where x is a vertex). We define $F(\Gamma_{i+1})$ to be the simplicial complex obtained by gluing $F(\Gamma_i)$ and M_i along $\text{St}(v, M_i)$ (see Section 2.1 for the notation), and define K_{i+1} to be the CAT(0) cube complex obtained by gluing K_i and $L_i \times [0, 1]$ along L_i. One readily verifies that one

can extend f to a cubical embedding $f'\colon K_{i+1} \to X(\Gamma)$ such that $f'(K_{i+1})$ is convex. This also induces an isomorphism $g'\colon F(\Gamma_{i+1}) \to \hat{K}_{i+1}$ which is an extension of g.

By construction, each $G(\Gamma_i)$ is isomorphic to a special subgroup of $G(\Gamma)$; moreover, the associated convex subcomplex of this special subgroup is K_i. Also note that the above induction process actually does not depend on knowing what $X(\Gamma)$ is. Thus it also provides a way to construct convex subcomplexes of $X(\Gamma)$ by hand.

The above process of obtaining (Γ_{i+1}, K_{i+1}) from (Γ_i, K_i) is called a *generalized star extension* (GSE) at v. Note that the following are equivalent:

(1) $\Gamma_i \subsetneq \Gamma_{i+1}$.

(2) $P_l \subsetneq X(\Gamma)$, where l is the standard geodesic in $X(\Gamma)$ such that $\Delta(l) = g(v)$.

(3) $\mathrm{St}(\pi(g(v))) \subsetneq F(\Gamma)$, where $\pi\colon \mathcal{P}(\Gamma) \to F(\Gamma)$ is the natural label-preserving projection defined in (4-1).

A GSE is *nontrivial* if $\Gamma_i \subsetneq \Gamma_{i+1}$. If Γ is not a clique, then at each stage, there exists a vertex $v \in \Gamma_i$ such that the GSE at v is nontrivial.

Lemma 6.18 *Suppose* $G(\Gamma')$ *is isomorphic to a special subgroup of* $G(\Gamma)$. *Then we can construct* Γ' *from* Γ *by using finitely many GSEs.*

Proof   Let Θ_S and $\mathrm{CN}(\Gamma, S)$ be the objects defined in Section 6.1. Suppose $G(\Gamma')$ is isomorphic to $\Theta_S(K)$ for $K \in \mathrm{CN}(\Gamma, S)$. We define a sequence of convex subcomplexes in K by induction. Let K_1 be the identity element in $G(\Gamma)$. Suppose K_i is already defined. If $K_i = K$, then the induction terminates. If $K_i \subsetneq K$, pick an edge $e_i \subset K$ such that $e_i \cap K_i$ is a vertex and let K_{i+1} be the convex hull of $K_i \cup e_i$. Let $\{K_i\}_{i=1}^s$ be the resulting collection of convex subcomplexes. An alternative way of describing K_{i+1} is the following. If h_i is the hyperplane in K dual to e_i, and N_i is the carrier of h_i in K, then $h_i \cap K_i = \varnothing$ by the convexity of K_i. Thus $K_i \cap N_i$ is disjoint from h_i. Hence there is a copy of $(K_i \cap N_i) \times [0, 1]$ inside N_i, which is denoted by M_i. Then $K_{i+1} = K_i \cup M_i$. Now one readily verifies that one can obtain $(\hat{K}_{i+1}, K_{i+1})$ from $(\hat{K}_i, K_i)$ by a GSE. $\square$

The above construction gives rise to an algorithm to detect whether $G(\Gamma')$ is isomorphic to a special subgroup of $G(\Gamma)$. If there are n vertices in Γ', then Γ' can be obtained from Γ by at most n nontrivial GSEs. So we can start with Γ, enumerate all possible n–step nontrivial GSEs from Γ, and compare each resulting graph with Γ'. By Theorem 5.7 and Theorem 6.13, we have the following result.

Theorem 6.19   *If* $\mathrm{Out}(G(\Gamma))$ *is finite, then* $G(\Gamma')$ *is quasi-isometric to* $G(\Gamma)$ *if and only if* Γ' *can be obtained from* Γ *by finitely many GSEs. In particular, there is an algorithm to determine whether* $G(\Gamma')$ *and* $G(\Gamma)$ *are quasi-isometric.*

Note that a GSE gives rise to a pair (Γ_i, K_i). If one does not care about the associated convex subcomplex K_i, then there is a simpler description of GSE when $\mathrm{Out}(G(\Gamma))$ is finite. Suppose we have already obtained $F(\Gamma_i)$ together with a finite collection of full subcomplexes $\{G_\lambda\}_{\lambda \in \Lambda_i}$ such that:

(1) $\{G_\lambda\}_{\lambda \in \Lambda_i}$ is a covering of $F(\Gamma_i)$.

(2) Each G_λ is isomorphic to $F(\Gamma)$.

When $i = 1$, we pick the trivial cover of $F(\Gamma)$ by itself. To construct Γ_{i+1}, pick a vertex $v \in F(\Gamma_i)$, let $\Lambda_v = \{\lambda \in \Lambda_i \mid v \in G_\lambda\}$ and let $\Gamma_v = \bigcup_{\lambda \in \Lambda_v} G_\lambda$. Suppose $\{C_j\}_{j=1}^m$ is the collection of connected components of $\Gamma_v \setminus \mathrm{St}(v, \Gamma_v)$, and suppose $C_j' = C_j \cup \mathrm{St}(v, \Gamma_v)$. Then $F(\Gamma_{i+1})$ is defined by gluing C_1' and $F(\Gamma_i)$ along $\mathrm{St}(v, \Gamma_v)$, and Γ_{i+1} is the 1–skeleton of $F(\Gamma_{i+1})$.

Lemma 6.20 *Suppose* $\mathrm{Out}(G(\Gamma))$ *is finite. Then the above simplified process is consistent with* GSE.

Proof We assume inductively that there is a CAT(0) cube complex K_i such that the two induction assumptions for GSE are satisfied; moreover, $\{G_\lambda\}_{\lambda \in \Lambda_i}$ coincides with $\{\Phi(x)\}_{x \in K_i}$ (where x is a vertex). Let $L = h \times [0, a]$ be as before and let $L_j = h \times \{j\} \subset L$ for each integer $j \in [0, a]$. It suffices to show there is a one-to-one correspondence between $\{L_j\}_{j=0}^a$ and $\{C_j'\}_{j=1}^m$ such that for each j, there exists a unique j' with $\widehat{f(L_j)} = g(C_{j'}')$. Pick adjacent vertices $x_1, x_2 \in f(L_j)$ and let $\overline{w} \in \Gamma$ be the label of edge $\overline{x_1 x_2}$. Suppose $\overline{v} = \pi(g(v))$. Then $d(\overline{w}, \overline{v}) = 1$. Since $\mathrm{Out}(G(\Gamma))$ is finite, the orthogonal complement of $\overline{w}$ satisfies $\overline{w}^\perp \not\subseteq \mathrm{St}(\overline{v})$. Then there is a vertex $\overline{u} \in \overline{w}^\perp$ such that $d(\overline{u}, \overline{v}) = 2$. The lifts of $\overline{u}$ in $\hat{x}_1$ and $\hat{x}_2$ are the same point, so $(\hat{x}_1 \cap \hat{x}_2) \setminus \mathrm{St}(g(v))$ contains a vertex. Since $F(\Gamma)$ does not have separating closed stars, $\hat{x}_i \setminus \mathrm{St}(g(v))$ is connected for $i = 1, 2$. Thus $(\hat{x}_1 \cap \hat{x}_2) \setminus \mathrm{St}(g(v))$ is connected. It follows that $\widehat{f(L_j)} \setminus \mathrm{St}(g(v))$ is connected. Moreover, from Lemma 4.9, $\widehat{f(L_{j_1})} \setminus \mathrm{St}(g(v))$ and $\widehat{f(L_{j_2})} \setminus \mathrm{St}(g(v))$ are in different components of $\mathcal{P}(\Gamma) \setminus \mathrm{St}(g(v))$ when $j_1 \neq j_2$, so there exists a unique j' such that $\widehat{f(L_j)} = g(C_{j'}')$. $\square$

References

[1] **A Abrams**, **N Brady**, **P Dani**, **M Duchin**, **R Young**, *Pushing fillings in right-angled Artin groups*, J. Lond. Math. Soc. 87 (2013) 663–688 MR

[2] **I Agol**, *The virtual Haken conjecture*, Doc. Math. 18 (2013) 1045–1087 MR

[3] **H Baik**, **S-h Kim**, **T Koberda**, *Right-angled Artin subgroups of the C^∞ diffeomorphism group of the real line*, preprint (2014) arXiv

[4]   **J Behrstock**, **R Charney**, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann. 352 (2012) 339–356   MR

[5]   **J A Behrstock**, **T Januszkiewicz**, **W D Neumann**, *Quasi-isometric classification of some high dimensional right-angled Artin groups*, Groups Geom. Dyn. 4 (2010) 681–692   MR

[6]   **J Behrstock**, **B Kleiner**, **Y Minsky**, **L Mosher**, *Geometry and rigidity of mapping class groups*, Geom. Topol. 16 (2012) 781–888   MR

[7]   **J A Behrstock**, **W D Neumann**, *Quasi-isometric classification of graph manifold groups*, Duke Math. J. 141 (2008) 217–240   MR

[8]   **M Bestvina**, **N Brady**, *Morse theory and finiteness properties of groups*, Invent. Math. 129 (1997) 445–470   MR

[9]   **M Bestvina**, **B Kleiner**, **M Sageev**, *The asymptotic geometry of right-angled Artin groups, I*, Geom. Topol. 12 (2008) 1653–1699   MR

[10] **M Bestvina**, **B Kleiner**, **M Sageev**, *Quasiflats in* CAT(0) *2–complexes*, Algebr. Geom. Topol. 16 (2016) 2663–2676   MR

[11] **N Brady**, **J Meier**, *Connectivity at infinity for right angled Artin groups*, Trans. Amer. Math. Soc. 353 (2001) 117–132   MR

[12] **N Brady**, **T Riley**, **H Short**, *The geometry of the word problem for finitely generated groups*, Birkhäuser, Basel (2007)   MR

[13] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundl. Math. Wissen. 319, Springer (1999)   MR

[14] **P-E Caprace**, **N Monod**, *Isometry groups of non-positively curved spaces: structure theory*, J. Topol. 2 (2009) 661–700   MR

[15] **P-E Caprace**, **M Sageev**, *Rank rigidity for* CAT(0) *cube complexes*, Geom. Funct. Anal. 21 (2011) 851–891   MR

[16] **R Charney**, *An introduction to right-angled Artin groups*, Geom. Dedicata 125 (2007) 141–158   MR

[17] **R Charney**, **M W Davis**, *Finite $K(\pi, 1)s$ for Artin groups*, from "Prospects in topology" (F Quinn, editor), Ann. of Math. Stud. 138, Princeton Univ. Press (1995) 110–124   MR

[18] **R Charney**, **M Farber**, *Random groups arising as graph products*, Algebr. Geom. Topol. 12 (2012) 979–995   MR

[19] **M T Clay**, **C J Leininger**, **J Mangahas**, *The geometry of right-angled Artin subgroups of mapping class groups*, Groups Geom. Dyn. 6 (2012) 249–278   MR

[20] **C B Croke**, **B Kleiner**, *Spaces with nonpositive curvature and their ideal boundaries*, Topology 39 (2000) 549–556   MR

[21] **M W Davis**, *Groups generated by reflections and aspherical manifolds not covered by Euclidean space*, Ann. of Math. 117 (1983) 293–324   MR

[22] **M B Day**, *Finiteness of outer automorphism groups of random right-angled Artin groups*, Algebr. Geom. Topol. 12 (2012) 1553–1583 MR

[23] **C Droms**, *Isomorphisms of graph groups*, Proc. Amer. Math. Soc. 100 (1987) 407–408 MR

[24] **M J Dunwoody**, *The accessibility of finitely presented groups*, Invent. Math. 81 (1985) 449–457 MR

[25] **A Eskin**, *Quasi-isometric rigidity of nonuniform lattices in higher rank symmetric spaces*, J. Amer. Math. Soc. 11 (1998) 321–361 MR

[26] **A Eskin**, **B Farb**, *Quasi-flats and rigidity in higher rank symmetric spaces*, J. Amer. Math. Soc. 10 (1997) 653–692 MR

[27] **B Farb**, **M Handel**, *Commensurations of* $\mathrm{Out}(\mathrm{F}_n)$, Publ. Math. Inst. Hautes Études Sci. 105 (2007) 1–48 MR

[28] **V N Gerasimov**, *Semi-splittings of groups and actions on cubings*, from "Algebra, geometry, analysis and mathematical physics" (Y G Reshetnyak, L A Bokut', S K Vodop'yanov, I A Taĭmanov, editors), Izdat. Ross. Akad. Nauk Sib. Otd. Inst. Mat., Novosibirsk (1997) 91–109 MR In Russian; translated as *Fixed-point-free actions on cubings*, Siberian Adv. Math. 8 (1998) 36–58

[29] **M Gromov**, *Hyperbolic groups*, from "Essays in group theory" (S M Gersten, editor), Math. Sci. Res. Inst. Publ. 8, Springer (1987) 75–263 MR

[30] **M F Hagen**, *Weak hyperbolicity of cube complexes and quasi-arboreal groups*, J. Topol. 7 (2014) 385–418 MR

[31] **M F Hagen**, **D T Wise**, *Cubulating hyperbolic free-by-cyclic groups: the general case*, Geom. Funct. Anal. 25 (2015) 134–179 MR

[32] **M F Hagen**, **D T Wise**, *Cubulating hyperbolic free-by-cyclic groups: the irreducible case*, Duke Math. J. 165 (2016) 1753–1813 MR

[33] **F Haglund**, *Isometries of* CAT(0) *cube complexes are semi-simple*, preprint (2007) arXiv

[34] **F Haglund**, *Finite index subgroups of graph products*, Geom. Dedicata 135 (2008) 167–209 MR

[35] **F Haglund**, **D T Wise**, *Special cube complexes*, Geom. Funct. Anal. 17 (2008) 1551–1620 MR

[36] **F Haglund**, **D T Wise**, *Coxeter groups are virtually special*, Adv. Math. 224 (2010) 1890–1903 MR

[37] **U Hamenstädt**, *Geometry of the mapping class groups, III: Quasi-isometric rigidity*, preprint (2005) arXiv

[38] **J Huang**, *Quasi-isometry classification of right-angled Artin groups, II: Several infinite out cases*, preprint (2016) arXiv

[39] **J Huang**, *Top-dimensional quasiflats in* CAT(0) *cube complexes*, Geom. Topol. 21 (2017) 2281–2352 MR

[40] **J Huang**, **B Kleiner**, *Groups quasi-isometric to RAAG's*, preprint (2016) arXiv To appear in Duke Math. J.

[41] **M Kapovich**, **B Kleiner**, **B Leeb**, *Quasi-isometries and the de Rham decomposition*, Topology 37 (1998) 1193–1211 MR

[42] **S-h Kim**, **T Koberda**, *Embedability between right-angled Artin groups*, Geom. Topol. 17 (2013) 493–530 MR

[43] **S-H Kim**, **T Koberda**, *The geometry of the curve graph of a right-angled Artin group*, Internat. J. Algebra Comput. 24 (2014) 121–169 MR

[44] **S-h Kim**, **T Koberda**, *Anti-trees and right-angled Artin subgroups of braid groups*, Geom. Topol. 19 (2015) 3289–3306 MR

[45] **B Kleiner**, **B Leeb**, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, C. R. Acad. Sci. Paris Sér. I Math. 324 (1997) 639–643 MR

[46] **B Kleiner**, **B Leeb**, *Groups quasi-isometric to symmetric spaces*, Comm. Anal. Geom. 9 (2001) 239–260 MR

[47] **T Koberda**, *Ping-pong lemmas with applications to geometry and topology*, from "Geometry, topology and dynamics of character varieties" (W Goldman, C Series, S P Tan, editors), Lect. Notes Ser. Inst. Math. Sci. Natl. Univ. Singap. 23, World Scientific, Hackensack, NJ (2012) 139–158 MR

[48] **T Koberda**, *Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups*, Geom. Funct. Anal. 22 (2012) 1541–1590 MR

[49] **M R Laurence**, *A generating set for the automorphism group of a graph group*, J. London Math. Soc. 52 (1995) 318–334 MR

[50] **L Mosher**, **M Sageev**, **K Whyte**, *Quasi-actions on trees, II: Finite depth Bass–Serre trees*, Mem. Amer. Math. Soc. 1008, Amer. Math. Soc., Providence, RI (2011) MR

[51] **G D Mostow**, *Strong rigidity of locally symmetric spaces*, Annals of Mathematics Studies 78, Princeton Univ. Press (1973) MR

[52] **Y Ollivier**, **D T Wise**, *Cubulating random groups at density less than* 1/6, Trans. Amer. Math. Soc. 363 (2011) 4701–4733 MR

[53] **P Papasoglu**, **K Whyte**, *Quasi-isometries between groups with infinitely many ends*, Comment. Math. Helv. 77 (2002) 133–144 MR

[54] **P Przytycki**, **D T Wise**, *Mixed 3-manifolds are virtually special*, preprint (2012) arXiv

[55] **P Przytycki**, **D T Wise**, *Graph manifolds with boundary are virtually special*, J. Topol. 7 (2014) 419–435 MR

[56] **M Sageev**, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. 71 (1995) 585–617 MR

[57] **H Servatius**, *Automorphisms of graph groups*, J. Algebra 126 (1989) 34–60 MR

[58] **H Servatius**, **C Droms**, **B Servatius**, *Surface subgroups of graph groups*, Proc. Amer. Math. Soc. 106 (1989) 573–578 MR

[59] **S J Taylor**, *Right-angled Artin groups and* Out($\mathbb{F}_n$), *I: Quasi-isometric embeddings*, Groups Geom. Dyn. 9 (2015) 275–316 MR

[60] **D T Wise**, *Cubulating small cancellation groups*, Geom. Funct. Anal. 14 (2004) 150–214 MR

[61] **D T Wise**, *Research announcement: the structure of groups with a quasiconvex hierarchy*, Electron. Res. Announc. Math. Sci. 16 (2009) 44–55 MR

[62] **D T Wise**, *The structure of groups with a quasiconvex hierarchy*, preprint (2011) Available at http://www.math.mcgill.ca/wise/papers.html

The Department of Mathematics and Statistics, McGill University
Montreal, QC, Canada

jingyin.huang@mcgill.ca