GROUPS QUASI-ISOMETRIC TO RIGHT-ANGLED ARTIN GROUPS

JINGYIN HUANG and BRUCE KLEINER

Abstract

We characterize groups quasi-isometric to a right-angled Artin group (RAAG) G with finite outer automorphism group. In particular, all such groups admit a geometric action on a CAT(0) *cube complex that has an equivariant "fibering" over the Davis building of G. This characterization will be used in forthcoming work of the first author to give a commensurability classification of the groups quasi-isometric to certain RAAGs.*

Contents

1. Introduction

Overview

In this article we will study right-angled Artin groups (RAAGs). Like other authors, our motivation for considering these groups stems from the fact that they are an easily defined yet remarkably rich class of objects, exhibiting interesting features from many different vantage points: algebraic structure (subgroup structure, automorphism groups; see [17], [24], [47], [56]), finiteness properties (see [8], [11]), representation varieties (see [43]), CAT(0) geometry (see [21]), cube complex geometry (see

DUKE MATHEMATICAL JOURNAL
Vol. 167, No. 3, © 2018 DOI 10.1215/00127094-2017-0042
Received 14 March 2016. Revision received 31 July 2017.
First published online 9 January 2018.
2010 *Mathematics Subject Classification*. Primary 20F65; Secondary 20F69.

[34], [62]), and coarse geometry (see [5], [6], [9], [13], [36], [37], [61]). Further impetus for studying RAAGs comes from their role in the theory of special cube complexes, which was a key ingredient in Agol's spectacular solution of Thurston's virtual Haken and virtual fibered conjectures (see [2], [34], [40], [53], [62]).

Our focus here is on quasi-isometric rigidity, which is part of Gromov's program for quasi-isometric classification of groups and metric spaces. In this article we build on [9], [10], [37], and [36], which analyzed the structure of individual quasi-isometries $G \to G$, where G is a RAAG with finite outer automorphism group. Our main results are a structure theorem for groups of quasi-isometries (more precisely quasiactions) and a characterization of finitely generated groups quasi-isometric to such RAAGs. Both are formulated using a geometric description in terms of Caprace–Sageev [14] restriction quotients and the Davis [22] building.

Background

Prior results on quasi-isometric classification of RAAGs may be loosely divided into two types: internal quasi-isometry classification among (families of) RAAGs and quasi-isometry rigidity results characterizing arbitrary finitely generated groups quasi-isometric to a given RAAG. In the former category, it is known that, to classify RAAGs up to quasi-isometry, it suffices to consider the case when the groups are 1-ended and do not admit any nontrivial direct product decomposition or, equivalently, when their defining graphs are connected, contain more than one vertex, and do not admit a nontrivial join decomposition (see [36, Theorem 2.9], [41], [51]). This covers, for instance, the classification up to quasi-isometry of RAAGs that may be formed inductively by taking products or free products, starting from copies of $\mathbb{Z}$. Beyond this, internal classification is known for RAAGs whose defining graph is a tree (see [6]) or a higher dimension analogue (see [5]) or when the outer automorphism group is finite (see [9], [36]). General quasi-isometric classification results in the literature are much more limited; if H is a finitely generated group quasi-isometric to a RAAG G, then we have the following.

(i) If G is free or free Abelian, then H is virtually free or free Abelian, respectively (see [4], [25], [27], [57]).

(ii) If $G = F_k \times \mathbb{Z}^\ell$, then H is virtually $F_k \times \mathbb{Z}^\ell$ (see [60]).

(iii) If the defining graph of G is a tree, then H is virtually the fundamental group of a nongeometric graph manifold that has nonempty boundary in every Seifert fiber space component, and consequently, H is virtually cocompactly cubulated (see [6], [31], [42]).

(iv) If G is a product of free groups, then H acts geometrically on a product of trees (see [3], [41], [49]).

Unlike (i)–(iii), which give characterizations up to commensurability, the characterization in (iv) only asserts the existence of an action on a good geometric model; the stronger commensurability assertion is false, in view of examples of Wise [61] and Burger–Mozes [13].

The setup

We now recall some terminology and notation; see Section 3 for more details. If Γ is a finite simplicial graph with vertex set $V(\Gamma)$, we denote the associated RAAG by $G(\Gamma)$. This is the fundamental group of the Salvetti complex $S(\Gamma)$, a nonpositively curved cube complex that may be constructed by choosing a pointed unit-length circle $(S^1_v, \star_v)$ for every vertex $v \in V(\Gamma)$, forming the pointed product torus $\prod_v (S^1_v, \star_v)$, and passing to the union of the product subtori corresponding to the cliques (complete subgraphs) in Γ. The clique subtori are the *standard tori* in $S(\Gamma)$.

We denote the universal covering by $X(\Gamma) \to S(\Gamma)$; here $X(\Gamma)$ is a CAT(0) cube complex on which $G(\Gamma)$ acts geometrically by deck transformations. The inverse image of a standard torus in $S(\Gamma)$ under the universal covering $X(\Gamma) \to S(\Gamma)$ breaks up into connected components; these are the *standard flats* in $X(\Gamma)$ which we partially order by inclusion. Note that we include standard tori and standard flats of dimension 0.

The poset of standard flats in $X(\Gamma)$ turns out to be crucial to our story. Using it one may define a locally infinite CAT(0) cube complex $|\mathcal{B}|(\Gamma)$ whose cubes of dimension $k \geq 0$ are indexed by inclusions $F_1 \subset F_2$, and F_1, F_2 are standard flats where $\dim F_2 = \dim F_1 + k$. Elements of the 0-skeleton $|\mathcal{B}|^{(0)}(\Gamma)$ correspond to the trivial inclusions $F \subset F$, where F is a standard flat, so we will identify $|\mathcal{B}|^{(0)}(\Gamma)$ with the collection of standard flats and define the rank of a vertex of $|\mathcal{B}|(\Gamma)$ to be the dimension of the corresponding standard flat; in particular, we may identify the 0-skeleton $X^{(0)}(\Gamma)$ with the set of rank 0 vertices of $|\mathcal{B}|^{(0)}$. Since $G(\Gamma) \curvearrowright X(\Gamma)$ preserves the collection of standard flats, there is an induced action $G(\Gamma) \curvearrowright |\mathcal{B}|(\Gamma)$ by cubical isomorphisms. The above description is a slight variation on the original construction of the same object given by Davis, in which one views $|\mathcal{B}|(\Gamma)$ as the Davis realization of a certain right-angled building $\mathcal{B}(\Gamma)$ associated with $G(\Gamma)$, where the apartments of $\mathcal{B}(\Gamma)$ are modelled on the right-angled Coxeter group $W(\Gamma)$ with defining graph Γ (see [22] and Section 3). By abuse of terminology we will refer to this cube complex as the *Davis building associated with* $G(\Gamma)$; it has been called the modified Deligne complex in [19] and flat space in [10].

The following lemma is not difficult to prove.

LEMMA 1.1

- *Every isomorphism* $|\mathcal{B}|^{(0)}(\Gamma) \to |\mathcal{B}|^{(0)}(\Gamma)$ *of the poset of standard flats extends to a unique cubical isomorphism* $|\mathcal{B}|(\Gamma) \to |\mathcal{B}|(\Gamma)$ *(Section 3.4).*
- *Every cubical isomorphism of* $|\mathcal{B}| \to |\mathcal{B}|$ *induces a poset isomorphism* $|\mathcal{B}|^{(0)} \to |\mathcal{B}|^{(0)}$ *(Lemma 3.15).*
- *A bijection* $\phi^{(0)} : |\mathcal{B}|^{(0)}(\Gamma) \supset X^{(0)}(\Gamma) \to X^{(0)}(\Gamma) \subset |\mathcal{B}|^{(0)}(\Gamma)$ *induces/extends to a poset isomorphism* $|\mathcal{B}|^{(0)}(\Gamma) \to |\mathcal{B}|^{(0)}(\Gamma)$ *if and only if it is flat-preserving in the sense that, for every standard flat* $F_1 \subset X(\Gamma)$*, the* 0*-skeleton* $F_1^{(0)}$ *is mapped bijectively by* $\phi^{(0)}$ *onto the* 0*-skeleton of some standard flat* F_2 *(Section 5.1).*

Remark 1.2

We caution the reader that a cubical isomorphism $|\mathcal{B}|(\Gamma) \to |\mathcal{B}|(\Gamma)$ need not arise from an isomorphism $\mathcal{B}(\Gamma) \to \mathcal{B}(\Gamma)$ of the right-angled building, because a cubical isomorphism need not preserve residue types.

Rigidity and flexibility

We now fix a finite graph Γ such that the outer automorphism group $\mathrm{Out}(G(\Gamma))$ is finite; by work of [20] and [23], one may view this as the generic case. The reader may find it helpful to keep in mind the case when Γ is a pentagon. Since there is no ambiguity in Γ we will often suppress it in the notation below.

It is known that, even if $\mathrm{Out}(G(\Gamma))$ is finite, $X = X(\Gamma)$ is not quasi-isometrically rigid: there are quasi-isometries that are not at a finite sup distance from isometries, and there are finitely generated groups H that are quasi-isometric to X, but do not admit geometric actions on X (Corollary 6.11). On the other hand, quasi-isometries exhibit a form of partial rigidity that is captured by the building $|\mathcal{B}|$. The following result is a consequence of [36, Theorem 4.18] (see also [9, Theorem 1.6]) and Lemma 1.1.

THEOREM 1.3

Suppose $\mathrm{Out}(G(\Gamma))$ *is finite and* $G(\Gamma) \not\simeq \mathbb{Z}$*. If* $\phi : X^{(0)} \to X^{(0)}$ *is an* (L, A)*-quasi-isometry, then there is a unique cubical isomorphism* $|\mathcal{B}| \to |\mathcal{B}|$ *such that the associated flat-preserving bijection* $\bar{\phi} : X^{(0)} \to X^{(0)}$ *is at finite sup distance from* ϕ*, and moreover,*

$$d(\bar{\phi}, \phi) = \sup\{v \in X^{(0)} \mid d(\bar{\phi}(v), \phi(v))\} < D = D(L, A).$$

By the uniqueness assertion, we obtain a cubical action $\mathrm{QI}(X) \curvearrowright |\mathcal{B}|$ *of the quasi-isometry group of* X *on* $|\mathcal{B}|$*.*

We point out that the partial rigidity statement of the theorem does not hold for general RAAGs; in fact, it holds only for the RAAGs covered by [36, Theorem 4.18], because when $\mathrm{Out}(G(\Gamma))$ is infinite, there are automorphisms of $G(\Gamma)$—either transvections or partial conjugations (see [47], [56])—which do not satisfy the conclusion of the above theorem.

The main results

We will produce good geometric models quasi-isometric to $X(\Gamma)$ that are simultaneously compatible with group actions, the underlying building $|\mathcal{B}|$, and cubical structure. The key idea for expressing this is the following.

Definition 1.4

A cubical map $q : Y \to Z$ between CAT(0) cube complexes (see Definition 3.4) is a *restriction quotient* if it is surjective and the point inverse $q^{-1}(z)$ is a convex subset of Y for every $z \in Z$.

It turns out that restriction quotients as defined above are essentially equivalent to the class of mappings introduced by Caprace–Sageev [14] with a different definition (see Section 4 for the proof that the definitions are equivalent). Restriction quotients $Y \to |\mathcal{B}|$ provide a means to "resolve" or "blow up" the locally infinite building $|\mathcal{B}|$ to a locally finite CAT(0) cube complex.

THEOREM 1.5 (see Section 3 for definitions)

Let $H \curvearrowright X$ be a quasiaction of an arbitrary group on $X = X(\Gamma)$, where $\mathrm{Out}(G(\Gamma))$ *is finite and $G(\Gamma) \not\simeq \mathbb{Z}$. Then there is an H-equivariant restriction quotient $H \curvearrowright Y \xrightarrow{q} H \curvearrowright |\mathcal{B}|$ where the following hold.*

(a) *$H \curvearrowright |\mathcal{B}|$ is the cubical action arising from the quasiaction $H \curvearrowright X$ using Theorem 1.3, and $H \curvearrowright Y$ is a cubical action.*

(b) *The point inverse $q^{-1}(v)$ of every rank k vertex $v \in |\mathcal{B}|^{(0)}$ is a copy of $\mathbb{R}^k$ with the usual cubulation.*

(c) *$H \curvearrowright X$ is quasiconjugate to the cubical action $H \curvearrowright Y$.*

THEOREM 1.6

If $|\mathrm{Out}(G(\Gamma))| < \infty$ and $G(\Gamma) \not\simeq \mathbb{Z}$, then a finitely generated group H is quasi-isometric to $G(\Gamma)$ if and only if there is an H-equivariant restriction quotient $H \curvearrowright Y \xrightarrow{q} H \curvearrowright |\mathcal{B}|$, where the following hold.

(a) *$H \curvearrowright Y$ is a geometric cubical action.*

(b) *$H \curvearrowright |\mathcal{B}|$ is cubical.*

(c) *The point inverse* $q^{-1}(v)$ *of every rank* k *vertex* $v \in |\mathcal{B}|^{(0)}$ *is a copy of* $\mathbb{R}^k$ *with the usual cubulation.*

Remark 1.7
In fact, the restriction quotient $Y \to |\mathcal{B}|$ in Theorems 1.5 and 1.6 has slightly more structure (see Theorem 6.2).

In particular, we have the following result.

COROLLARY 1.8
Any group quasi-isometric to G *is cocompactly cubulated, that is, it has a geometric cubical action on a* CAT(0) *cube complex.*

One may compare Theorem 1.6 with rigidity theorems for symmetric spaces or products of trees, which characterize a quasi-isometry class of groups by the existence of a geometric action on a model space of a specified type (see [3], [25], [28], [41], [45], [49], [50], [57]–[59]). As in the case of products of trees—and unlike the case of symmetric spaces—there are finitely generated groups H as in Theorem 1.6 which do not admit a geometric action on the original model space X, so one is forced to pass to a different space Y (see [9], [36]). Also, Theorems 1.5 and 1.6 fail for general RAAGs, for instance, for free Abelian groups of rank at least 2 and for products of non-Abelian free groups $\prod_{1 \leq j \leq k} G_j$ for $k \geq 1$.

The quasi-isometry invariance of the existence of a cocompact cubulation as asserted in Corollary 1.8 is false in general. Some groups quasi-isometric to $\mathbb{H}^2 \times \mathbb{R}$ admit a cocompact cubulation, while others are not virtually CAT(0) (see [12]). Combining [6], [48], and [31], it follows that there is a pair of quasi-isometric CAT(0) graph manifold groups, one of which is the fundamental group of a compact special cube complex, while the other is not virtually cocompactly cubulated. The quasi-isometry invariance of cocompact cubulations fails to hold even among RAAGs: for $n > 1$ there are groups commensurable to $\mathbb{Z}^n$ that are not cocompactly cubulated (see [30]).

Earlier cocompact cubulation theorems in the spirit of Corollary 1.8 include the cases of groups quasi-isometric to trees, products of trees, and hyperbolic k-space $\mathbb{H}^k$ for $k \in \{2, 3\}$ (see [3], [7], [25], [26], [40], [41], [49], [57]). It is worth noting that each case requires different ingredients that are specific to the spaces in question.

Further results
We briefly discuss some further results here, referring the reader to the body of the article for details. One portion of the proof of Theorem 1.5 has to do with the geometry

of restriction quotients and, more specifically, restriction quotients with a right-angled building as the target. We view this as a contribution to cube complex geometry and to the geometric theory of graph products; beyond the references mentioned already, our treatment has been influenced by the papers of Januszkiewicz–Świątkowski [39] and Haglund [32]. The main results on this are as follows.

(a) We show in Section 4 that restriction quotients may be characterized in several different ways.

(b) We show that having a restriction quotient $q : Y \to Z$ is equivalent to knowing certain "fiber data" living on the target complex Z.

(c) When $|\mathcal{B}|$ is the Davis realization of a right-angled building $\mathcal{B}$ and $Y \to |\mathcal{B}|$ is a restriction quotient whose fibers are copies of $\mathbb{R}^k$ with dimension specified as in Theorems 1.5 and 1.6, the fiber data in (b) may be distilled even more, leading to what we call "blow-up data."

As by-products of (a)–(c), we obtain the following.

- A characterization of the quasiactions $H \curvearrowright X(\Gamma)$ that are quasiconjugate to isometric actions $H \curvearrowright X(\Gamma)$ (see Section 6.2).
- A characterization of the restriction quotients $Y \to |\mathcal{B}|$ satisfying Theorem 1.5(b) for which Y is quasi-isometric to X (see Corollary 6.4, Theorem 6.5).
- A proof of uniqueness of the right-angled building modelled on the right-angled Coxeter group $W(\Gamma)$ with defining graph Γ, with countably infinite rank 1 residues (Corollary 5.21). This was previously established in [33].

It follows from [46] that a finitely generated group H quasi-isometric to a symmetric space of noncompact type X admits an epimorphism $H \to \Lambda$ with finite kernel, where Λ is a cocompact lattice in the isometry group $\mathrm{Isom}(X)$. In contrast to this, we have the following result, which is inspired by [49, Theorem 9 and Corollary 10].

THEOREM 1.9 (see Theorem 6.10)
Suppose G is a RAAG with $G(\Gamma) \not\simeq \mathbb{Z}$ and $|\mathrm{Out}(G)| < \infty$. Then there are finitely generated groups H and H' quasi-isometric to G that do not admit discrete, virtually faithful cocompact representations into the same locally compact topological group.

Open questions
As mentioned above, Corollary 1.8 may be considered part of the quasi-isometry classification program for finitely generated groups. This leads to the following.

Question 1.10
If $\mathrm{Out}(G(\Gamma))$ is finite, what is the commensurability classification of groups quasi-

isometric to $G(\Gamma)$? Are they all commensurable to $G(\Gamma)$? What about cocompact lattices in the automorphism group of $X(\Gamma)$?

For comparison, we recall that any group quasi-isometric to a tree is commensurable to a free group, but there are groups quasi-isometric to a product of trees that contain no nontrivial finite-index subgroups and are therefore not commensurable to a product of free groups (see [13], [61]). We mention that Theorem 1.6 will be used in [38] to answer Question 1.10 in certain cases.

Another question motivated by Corollary 1.8 is the following.

Question 1.11
Under what conditions on a RAAG G must every group quasi-isometric to G be virtually cocompactly cubulated?

Discussion of the proofs
Before sketching the arguments for Theorems 1.5 and 1.6, we first illustrate them in the tautological case when $H = G$ and the quasiaction is the deck group action $G \curvearrowright X$. In this case, we cannot take $Y = X$, as there is no H-equivariant restriction quotient $H \curvearrowright X \to H \curvearrowright |\mathcal{B}|$ satisfying Theorem 1.6(c). Instead, we use a different geometric model.

Definition 1.12 (Graph products of spaces [32])
For every vertex $v \in V(\Gamma)$, choose a pointed geodesic metric space $(Z_v, \star_v)$. The Γ-*graph product of* $\{(Z_v, \star_v)\}_{v \in V(\Gamma)}$ is obtained by forming the product $\prod_v (Z_v, \star_v)$ and passing to the union of the subproducts corresponding to the cliques in Γ. We denote this by $\prod_\Gamma (Z_v, \star_v)$. When the Z_v's are nonpositively curved, then so is the graph product (see [32, Corollary 4.6]).

There are three graph products that are useful here.

(1) The Salvetti complex $S(\Gamma)$ is the graph product $\prod_\Gamma (S^1_v, \star_v)$, where $(S^1_v, \star_v)$ is a pointed unit circle.
(2) For every $v \in V(\Gamma)$, let $(L_v, \star_v)$ be a pointed lollipop, that is, L_v is the wedge of the unit circle S^1_v and a unit interval I_v, and the base point $\star_v \in L_v$ is the vertex of valence 1. Then the graph product $\prod_\Gamma (L_v, \star_v)$ is the *exploded Salvetti complex* $S_e = S_e(\Gamma)$. We denote its universal covering by $X_e \to S_e$.
(3) If $(Z_v, \star_v)$ is a unit interval and $\star_v \in Z_v$ is an endpoint for every $v \in V(\Gamma)$, then the graph product $\prod_\Gamma (Z_v, \star_v)$ is the *Davis chamber*, that is, it is a copy of the Davis realization $|c|$ of a chamber c in $|\mathcal{B}|(\Gamma)$; for this reason we will denote it by $|c|_\Gamma$.

By collapsing the interval I_v in each lollipop L_v to a point, we obtain a cubical map $S_e \to S$; this has contractible point inverses, and is therefore a homotopy equivalence. If we collapse the circles $S_v^1 \subset L_v$ to points instead, we get a map $S_e \to |c|_\Gamma$ to the Davis chamber whose point inverses are closed, locally convex tori. The point inverses of the composition $X_e \to S_e \to |c|_\Gamma$ cover the torus point inverses of $S_e \to |c|_\Gamma$, and their connected components form a "foliation" of X_e by flat convex subspaces. It turns out that by collapsing X_e along these flat subspaces, we obtain a copy of $|\mathcal{B}|$, and the quotient map $X_e \to |\mathcal{B}|$ is a restriction quotient $X_e \to |\mathcal{B}|$. Alternately, one may take the collection $\mathcal{K}$ of hyperplanes of X_e dual to edges $\sigma \subset X_e$ whose projection under $X_e \to |c|_\Gamma$ is an edge, and then one may form the restriction quotient using the Caprace–Sageev construction.

Remark 1.13
The exploded Salvetti complex and the restriction quotient $X_e \to |\mathcal{B}|$ were discussed in [9] in the 2-dimensional case, using an ad hoc construction that was initially invented for "ease of visualization." However, the authors were unaware of the general description above, and the notion of restriction quotient had not yet appeared.

We now discuss the proofs of Theorem 1.5 and the forward direction of Theorem 1.6.

The forward direction of Theorem 1.6 reduces to Theorem 1.5, by the standard observation that a quasi-isometry $H \to G \overset{qi}{\simeq} X$ allows us to quasiconjugate the left translation action $H \curvearrowright H$ to a quasiaction $H \curvearrowright X$. Therefore, we focus on Theorem 1.5.

Let $H \curvearrowright X$ be as in Theorem 1.5. By a bounded perturbation, we may assume that this quasiaction preserves the 0-skeleton $X^{(0)} \subset X$. Applying Theorem 1.3, we may further assume that we have an action $H \curvearrowright X^{(0)}$ by flat-preserving quasi-isometries. The fact the we have an action, rather than just a quasiaction, comes from the uniqueness in Theorem 1.3; this turns out to be a crucial point in the rest of this article.

Given a standard geodesic $\ell \subset X$, the parallel set $P_\ell \subset X$ decomposes as a product $\mathbb{R}_\ell \times Q_\ell$, where $\mathbb{R}_\ell$ is a copy of $\mathbb{R}$; likewise there is a product decomposition of 0-skeleta $P_\ell^{(0)} \simeq \mathbb{Z}_\ell \times Q_\ell^{(0)}$. One argues that the action $H \curvearrowright X^{(0)}$ permutes the collection of 0-skeleta $\{P_\ell^{(0)}\}_\ell$ and that, for any ℓ, the stabilizer $\mathrm{Stab}(P_\ell^{(0)}, H)$ of $P_\ell^{(0)}$ in H acts on $P_\ell^{(0)} \simeq \mathbb{Z}_\ell \times Q_\ell^{(0)}$, preserving the product structure. We call the action $H_\ell = \mathrm{Stab}(P_\ell^{(0)}, H) \curvearrowright \mathbb{Z}_\ell$ a *factor action*. The factor actions are by bijections with quasi-isometry constants bounded uniformly independent of ℓ. See Definition 5.32 for a rigorous definition of factor action and Section 6 for more details on the argument in this paragraph.

It turns out that factor actions play a central role in the story. For instance, when the action $H \curvearrowright X^{(0)}$ is the restriction of an action $H \curvearrowright X$ by cubical isometries, the factor action $H_\ell \curvearrowright \mathbb{Z}_\ell$ is also an action by isometries for every standard geodesic ℓ. In general, the factor actions can be arbitrary: up to isometric conjugacy, any action $A \curvearrowright \mathbb{Z}$ by quasi-isometries with uniform constants can arise as a factor action for some action as in Theorem 1.5. A key step in the proof is to show that such actions have a relatively simple structure.

PROPOSITION 1.14 (Semiconjugacy)
Let $U \overset{\rho_0}{\curvearrowright} \mathbb{Z}$ be an action of an arbitrary group by (L, A)-quasi-isometries. Then there is an isometric action $U \overset{\rho_1}{\curvearrowright} \mathbb{Z}$ and surjective equivariant (L', A')-quasi-isometry

$$U \overset{\rho_0}{\curvearrowright} \mathbb{Z} \longrightarrow U \overset{\rho_1}{\curvearrowright} \mathbb{Z},$$

where L' and A' depend only on L and A.

This is proved in Proposition 6.3.

Remark 1.15
The assumption that ρ_0 is an action, as opposed to a quasiaction, is crucial: if a group U has a nontrivial quasihomomorphism $\alpha : U \to \mathbb{R}$, then the translation quasiaction $U \overset{\hat{\alpha}}{\curvearrowright} \mathbb{R}$ defined by $\hat{\alpha}(u)(x) = x + \alpha(u)$ is quasiconjugate to a quasiaction on $\mathbb{Z}$, but not to an isometric action on $\mathbb{Z}$.

It follows immediately from Proposition 1.14 that $U \overset{\rho_0}{\curvearrowright} \mathbb{Z}$ is quasiconjugate to an isometric action on the tree $\mathbb{R}$. In that respect Proposition 1.14 is similar to the theorem of Mosher–Sageev–Whyte [49, Theorem 1] about promoting quasiactions on bushy trees to isometric actions on trees. Since $\mathbb{R}$ is not bushy, [49, Theorem 1] does not apply, and indeed, the example above shows that the assumption of bushiness is essential in that theorem.

Continuing with the proof of Theorem 1.5, we see that Proposition 1.14 gives a good geometric model for the factor action $\mathrm{Stab}(P_\ell^{(0)}, H) \curvearrowright \mathbb{Z}_\ell$: we simply extend each isometry $\mathbb{Z}_\ell \to \mathbb{Z}_\ell$ to an isometry $\mathbb{R}_\ell \to \mathbb{R}_\ell$, thereby obtaining a cubical action $\mathrm{Stab}(P_\ell^{(0)}, H) \curvearrowright \mathbb{R}_\ell$. In vague terms, the remainder of the proof is concerned with combining these cubical models into models for the fibers of a restriction quotient $Z \to |\mathcal{B}|$ in an H-equivariant way. This portion of the proof is covered by more general results about restriction quotients (see (b) and (c) in the "Further results" section).

Organization of the article

A summary of notation can be found in Section 2. Section 3 contains some background material on quasiactions, CAT(0) cube complexes, RAAGs, and buildings. One can proceed directly to later sections with Sections 2 and 3 as references.

The main part of the article is Sections 4–7, where we prove Theorem 1.6. In Section 4 we discuss restriction quotients, showing how to construct a restriction quotient $Y \to Z$ starting from the target Z and an admissible assignment of fibers to the cubes of Z. Then we discuss equivariance properties and the coarse geometry of restriction quotients.

In Section 5, we introduce blowups of buildings based on Section 4. These are restriction quotients $Y \to |\mathcal{B}|$, where the target is a right-angled building and the fibers are Euclidean spaces of varying dimension. We motivate our construction in Sections 5.1 and 5.2. Blowups of buildings are constructed in Section 5.3. Several properties of them are discussed in Sections 5.4 and 5.5. We incorporate a group action into our construction in Section 5.6.

In Section 6.1, we apply the construction in Section 5.6 to RAAGs and prove Theorem 1.6 modulo Theorem 1.14, which is postponed until Section 7. In Section 6.2 we answer several natural questions motivated by Theorem 1.6 and prove Theorem 1.9.

2. Index of notation

- $\mathcal{B}$: a combinatorial building (Section 3.4).
- $|\mathcal{B}|$: the Davis realization of a building (Section 3.4).
- Chambers in the combinatorial building $\mathcal{B}$ are c, c', d.
- $|c|_\Gamma$: the Davis chamber (the discussion after Definition 1.12, Section 3.4).
- S^r: the collection of all spherical residues in the building $\mathcal{B}$.
- $\mathrm{proj}_{\mathcal{R}} : \mathcal{B} \to \mathcal{R}$: the nearest point projection from $\mathcal{B}$ to a residue $\mathcal{R}$ (Section 3.4).
- $\Lambda_{\mathcal{B}}$: the collection of parallel classes of rank 1 residues in the combinatorial building $\mathcal{B}$. We also write Λ when the building $\mathcal{B}$ is clear (Section 5.3).
- T: a type map which assigns each residue of $\mathcal{B}$ a subset of $\Lambda_{\mathcal{B}}$ (Section 5.3).
- CCC: the category of nonempty CAT(0) cube complexes with morphisms given by convex cubical embeddings.
- P_C: the parallel set of a closed convex subset of a CAT(0) space (Section 3.2).
- $W(\Gamma)$: the right-angled Coxeter group with defining graph Γ (Section 3.4).
- $G(\Gamma)$: the RAAG with defining graph Γ.
- $X(\Gamma) \to S(\Gamma)$: the universal covering of the Salvetti complex (Section 3.3).
- $X_e(\Gamma) \to S_e(\Gamma)$: the universal covering of the exploded Salvetti complex (after Definition 1.12 and Section 5.1). We also write $X_e \to S_e$ when the graph Γ is clear.

- $\mathcal{P}(\Gamma)$: the extension complex (Definition 3.5).
- $X \to X(\mathcal{K})$: the restriction quotient arising from a set $\mathcal{K}$ of hyperplanes in a CAT(0) cube complex (Definition 1.4).
- $Lk(x, X)$ or $Lk(c, X)$: the link of a vertex x or a cell c in a polyhedral complex X.
- $\Gamma_1 \circ \Gamma_2$: the join of two graphs.
- $K_1 * K_2$: the join of two simplicial complexes.

3. Preliminaries

3.1. Quasiactions

We recall several definitions from coarse geometry.

Definition 3.1

An (L, A)-quasiaction of a group G on a metric space Z is a map $\rho : G \times Z \to Z$ such that $\rho(\gamma, \cdot) : Z \to Z$ is an (L, A) quasi-isometry for every $\gamma \in G$, $d(\rho(\gamma_1, \rho(\gamma_2, z)), \rho(\gamma_1\gamma_2, z)) < A$ for every $\gamma_1, \gamma_2 \in G$, $z \in Z$, and $d(\rho(e, z), z) < A$ for every $z \in Z$.

The action ρ is *discrete* if, for any point $z \in Z$ and any $R > 0$, the set of all $\gamma \in G$ such that $\rho(\gamma, z)$ is contained in the ball $B_R(z)$ is finite; ρ is *cobounded* if Z coincides with a finite tubular neighborhood of the "orbit" $\rho(G, z)$. If ρ is a discrete and cobounded quasiaction of G on Z, then the orbit map $\gamma \in G \to \rho(\gamma, z)$ is a quasi-isometry. Conversely, given a quasi-isometry between G and Z, it induces a discrete and cobounded quasiaction of G on Z.

Two quasiactions ρ and ρ' are *equivalent* if there exists a constant D such that

$$\sup_{\gamma \in G} \sup_{z \in Z} d\big(\rho(\gamma, z), \rho'(\gamma, z)\big) < D.$$

Definition 3.2

Let ρ and ρ' be quasiactions of G on Z and Z', respectively, and let $\phi : Z \to Z'$ be a quasi-isometry. Then ρ is *quasiconjugate* to ρ' via ϕ if there is a D such that

$$\sup_{\gamma \in G} \sup_{z \in Z} d(\phi \circ \rho(\gamma, z), \rho'\big(\gamma, \phi(z)\big) < D.$$

3.2. CAT(0) *cube complexes*

We refer to [12] for background about CAT(0) spaces ([12, Chapter II.1]) and cube complexes ([12, Chapter II.5]). We refer to [53] and [54] for CAT(0) cube complexes and hyperplanes.

A unit Euclidean n-cube is $[0, 1]^n$ with the standard metric. A *midcube* is the set of fixed points of a reflection with respect to some $[0, 1]$-factor of $[0, 1]^n$. A cube complex Y is obtained by taking a collection of unit Euclidean cubes and gluing them along isometric faces. The gluing metric on Y is CAT(0) if and only if Y is simply connected and the link of each vertex in Y is a flag simplicial complex (see [29]); in this case, Y is called a CAT(0) *cube complex*.

Let X be a complete CAT(0) space, and let $C \subset X$ be a closed convex subset. Then there is a well-defined nearest point projection from X to C, which we denote by $\pi_C : X \to C$. Two convex subsets C_1 and C_2 are *parallel* if $d(\cdot, C_2)|_{C_1}$ and $d(\cdot, C_1)|_{C_2}$ are constant functions. In this case, the convex hull of C_1 and C_2 is isometric to $C_1 \times [0, d(C_1, C_2)]$.

For a closed convex subset $C \subset X$, we define P_C, the *parallel set* of C, to be the union of all convex subsets of X which are parallel to C. If C has the geodesic extension property, then P_C is also a closed convex subset and admits a canonical splitting $P_C \cong C \times C^{\perp}$ (see [12, Chapter II.2.12]).

Suppose Y is a CAT(0) cube complex. Then two edges e and e' are parallel if and only if there exist sequences of edges $\{e_i\}_{i=1}^n$ such that $e_1 = e$, $e_n = e'$, and e_i, e_{i+1} are the opposite sides of a 2-cube in Y. For each edge $e \subset Y$, let N_e be the union of cubes in Y which contain an edge parallel to e. Then N_e is a convex subcomplex of Y; moreover, N_e has a natural splitting $N_e \cong h_e \times [0,1]$, where $[0,1]$ corresponds to the e-direction. The subset $h_e \times \{1/2\}$ is called the *hyperplane* dual to e, and N_e is called the *carrier* of this hyperplane. Each hyperplane is a union of midcubes and, hence, has a natural cube complex structure, which makes it a CAT(0) cube complex. The following statements are true for hyperplanes.

(1) Each hyperplane h is a convex subset of Y. Moreover, $Y \setminus h$ has exactly two connected components. The closure of each connected component is called a *half-space*. Each half-space is also a convex subset.

(2) Pick an edge $e \subset Y$. We identify e with $[0,1]$ and consider the CAT(0) projection $\pi_e : Y \to e \cong [0,1]$. Then $h = \pi_e^{-1}(1/2)$ is the hyperplane dual to e, and $\pi_e^{-1}([0,1/2]), \pi_e^{-1}([1/2,1])$ are two half-spaces associated with h. The closure of $\pi_e^{-1}((0,1))$ is the carrier of h.

Let Y be a CAT(0) cube complex, and let $l \in Y$ be a geodesic line (with respect to the CAT(0)-metric) in the 1-skeleton of Y. Let $e \subset l$ be an edge, and pick $x \in e$. We claim that if x is in the interior of e, then $\pi_l^{-1}(x) = \pi_e^{-1}(x)$. It is clear that $\pi_l^{-1}(x) \subset \pi_e^{-1}(x)$. Suppose $y \in \pi_e^{-1}(x)$. Recall that $\pi_e^{-1}(x) \subset N_e$. It follows from the splitting $N_e \cong h_e \times [0,1]$ as above that the geodesic segment $\overline{xy}$ is orthogonal to l, that is, $\angle_x(y, y') = \pi/2$ for any $y' \in l \setminus \{x\}$; thus, $y \in \pi_l^{-1}(x)$.

The above claim implies $\pi_l^{-1}(x)$ is a convex subset for any $x \in l$. Moreover, the following lemma is true.

LEMMA 3.3
Let Y and l be as before. Pick an edge $e \subset Y$. If e is parallel to some edge $e' \subset l$, then $\pi_l(e) = e'$; otherwise, $\pi_l(e)$ is a vertex of l.

Now we define an alternative metric on the CAT(0) cube complex Y, which is called the l^1-metric. One can view the 1-skeleton of Y as a metric graph with edge length equal to 1, and this metric extends naturally to a metric on Y. The distance between two vertices in Y with respect to this metric is equal to the number of hyperplanes separating these two vertices.

A *combinatorial geodesic* in Y is an edge path in $Y^{(1)}$ which is a geodesic with respect to the l^1-metric. However, we always refer to the CAT(0)-metric when we talk about a geodesic.

If Y is finite-dimensional, then the l^1-metric and the CAT(0)-metric on Y are quasi-isometric (see [14, Lemma 2.2]). In this article, we will use the CAT(0)-metric unless otherwise specified.

Definition 3.4 ([14, Section 2.1])
A cellular map between 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.

3.3. RAAGs

Pick a finite simplicial graph Γ, and recall that $G(\Gamma)$ is the RAAG with defining graph Γ. Let S be a standard generating set for $G(\Gamma)$, and label the vertices of Γ by elements in S. $G(\Gamma)$ has a nice Eilenberg–MacLane space $S(\Gamma)$, called the *Salvetti complex* (see [15], [18]). Recall that $S(\Gamma)$ is the graph product $\prod_\Gamma(S^1_v, \star_v)$, where $(S^1_v, \star_v)$ is a pointed unit circle (see Definition 1.12).

The 2-skeleton of $S(\Gamma)$ is the usual presentation complex of $G(\Gamma)$, so $\pi_1(S(\Gamma)) \cong G(\Gamma)$. The 0-skeleton of $S(\Gamma)$ consists of one point whose link is a flag complex, so $S(\Gamma)$ is nonpositively curved and $S(\Gamma)$ is an Eilenberg–MacLane space for $G(\Gamma)$ by the Cartan–Hadamard theorem (see [12, Theorem II.4.1]).

The closure of each k-cell in $S(\Gamma)$ is a k-torus. Tori of this kind are called *standard tori*. There is a one-to-one correspondence between the k-cells (or standard torus of dimension k) in $S(\Gamma)$ and k-cliques in Γ. 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 $X(\Gamma)$. A *standard k-flat* in $X(\Gamma)$ is a connected component of the inverse image of a standard k-torus under the covering map $X(\Gamma) \to S(\Gamma)$. When $k = 1$, we also call it a *standard geodesic*.

For each simplicial graph Γ, there is a simplicial complex $\mathcal{P}(\Gamma)$ called the *extension complex*, which captures the combinatorial pattern of how standard flats intersect each other in $X(\Gamma)$. This object was first introduced in [44]. We will define it in a slightly different way (see [36, Section 4.1] for more discussion).

Definition 3.5 (Extension complex)
The vertices of $\mathcal{P}(\Gamma)$ are in one-to-one correspondence with the parallel classes of standard geodesics in $X(\Gamma)$. Two distinct vertices $v_1, v_2 \in \mathcal{P}(\Gamma)$ are connected by an edge if and only if there is a standard geodesic l_i in the parallel class associated with v_i $(i = 1, 2)$ such that l_1 and l_2 span a standard 2-flat. Then $\mathcal{P}(\Gamma)$ is defined to be the flag complex of its 1-skeleton; namely, we build $\mathcal{P}(\Gamma)$ inductively from its 1-skeleton by filling a k-simplex whenever we see the $(k-1)$-skeleton of a k-simplex.

Since each complete subgraph in the 1-skeleton of $\mathcal{P}(\Gamma)$ gives rise to a collection of mutually orthogonal standard geodesic lines, there is a one-to-one correspondence between k-simplices in $\mathcal{P}(\Gamma)$ and parallel classes of standard $(k+1)$-flats in $X(\Gamma)$. In particular, there is a one-to-one correspondence between maximal simplices in $\mathcal{P}(\Gamma)$ and maximal standard flats in $X(\Gamma)$. Given a standard flat $F \subset X(\Gamma)$, we denote the simplex in $\mathcal{P}(\Gamma)$ associated with the parallel class containing F by $\Delta(F)$.

3.4. Right-angled buildings

We will follow the treatment in [1], [22], and [52]. In particular, we refer to [22, Sections I.1–I.3] for the definitions of chamber systems, galleries, residues (which are particular subsets of chambers), Coxeter groups, and buildings. We will focus on right-angled buildings, that is, the associated Coxeter group is right-angled, though most of the discussion below is valid for general buildings.

Let $W = W(\Gamma)$ be a right-angled Coxeter group with (finite) defining graph Γ. Let $\mathcal{B} = \mathcal{B}(\Gamma)$ be a right-angled building with the associated W-distance function denoted by $\delta : \mathcal{B} \times \mathcal{B} \to W$. We will also call $\mathcal{B}(\Gamma)$ a *right-angled* Γ*-building* for simplicity.

Let I be the vertex set of Γ. Recall that a subset $J \subset I$ is *spherical* if the subgroup of W generated by J is finite. Let S be the poset of spherical subsets of I (including the empty set), and let $|\mathsf{S}|_\Delta$ be the geometric realization of S, that is, $|\mathsf{S}|_\Delta$ is a simplicial complex such that its vertices are in one-to-one correspondence to elements in S and its n-simplices are in one-to-one correspondence to $(n+1)$-chains in S. Note that $|\mathsf{S}|_\Delta$ is isomorphic to the simplicial cone over the barycentric subdivision of the flag complex of Γ.

Recall that, for elements $x \leq y$ in S, the *interval* I_{xy} between x and y is a poset consisting of elements $z \in \mathsf{S}$ such that $x \leq z \leq y$ with the induced order from S. There is a natural simplicial embedding $|I_{xy}|_\Delta \hookrightarrow |\mathsf{S}|_\Delta$. Each $|I_{xy}|_\Delta$ is a simplicial cone over the barycentric subdivision of a simplex and, thus, can be viewed as a subdivision of a cube into simplices. (The Hasse diagram of I_{xy} can be identified as the 1-skeleton of this cube.) It is not hard to check that the collection of all intervals in S gives rise to a structure of a cube complex on $|\mathsf{S}|_\Delta$. Let $|\mathsf{S}|$ be the resulting cube complex. Then $|\mathsf{S}|$ is CAT(0).

A residue is *spherical* if it is a J-residue with $J \in \mathsf{S}$. The *rank* of this residue is the cardinality of J. Let S^r be the poset of all spherical residues in $\mathcal{B}$. For $x \in \mathsf{S}^r$ which comes from a J-residue, we define the *rank* of x to be the rank of the associated residue and define a *type map* $t : \mathsf{S}^r \to \mathsf{S}$ which maps x to $J \in \mathsf{S}$. Let $|\mathsf{S}^r|_\Delta$ be the geometric realization of S^r. Then the type map induces a simplicial map $t : |\mathsf{S}^r|_\Delta \to |\mathsf{S}|_\Delta$. For $x \in \mathsf{S}^r$, let S^r_x be the sub-poset made of elements in S^r which are $\geq x$. If x is of rank 0, then S^r_x is isomorphic to S; moreover, there is a natural simplicial embedding $|\mathsf{S}^r_x|_\Delta \to |\mathsf{S}^r|_\Delta$, and t maps the image of $|\mathsf{S}^r_x|_\Delta$ isomorphically onto $|\mathsf{S}|_\Delta$.

As before, the geometric realization of each interval in S^r is a subdivision of a cube into simplices. Moreover, the intersection of two intervals in S^r is also an interval. Thus, one gets a cube complex $|\mathcal{B}|$ whose cubes are in one-to-one correspondence with intervals in S^r. $|\mathcal{B}|$ is called the *Davis realization* of the building $\mathcal{B}$, and $|\mathcal{B}|$ is a CAT(0) cube complex by [22]. Moreover, the above type map induces a cubical map $t : |\mathcal{B}| \to |\mathsf{S}|$. Let $\mathcal{R} \subset \mathcal{B}$ be a residue. Since $\mathcal{R}$ also has the structure of a building, there is an isometric embedding $|\mathcal{R}| \to |\mathcal{B}|$ between their Davis realizations. $|\mathcal{R}|$ is called a *residue* in $|\mathcal{B}|$.

In the special case when $\mathcal{B}$ is equal to the associated Coxeter group W, there is a natural embedding from the Cayley graph of W to $|\mathcal{B}|$ such that vertices of the Cayley graph are mapped to vertices of rank 0 in $|\mathcal{B}|$. And $|\mathcal{B}|$ can be viewed as the first cubical subdivision of the cubical completion of the Cayley graph of W. (The cubical completion means we attach an n-cube to the graph whenever there is a copy of the 1-skeleton of an n-cube inside the graph.)

Each vertex of $|\mathcal{B}|$ corresponds to a J-residue in $\mathcal{B}$ and, thus, has a well-defined rank. For a vertex x of rank 0, the space $|\mathsf{S}^r_x|_\Delta$ discussed in the previous paragraph induces a subcomplex $|\mathcal{B}_x| \subset |\mathcal{B}|$. Note that t maps $|\mathcal{B}_x|$ isomorphically onto $|\mathsf{S}|$, while $|\mathcal{B}_x|$ is called a *chamber* in $|\mathcal{B}|$ (there is a one-to-one correspondence between chambers in $|\mathcal{B}|$ and chambers in $\mathcal{B}$). Let $|\mathcal{B}_x|$ and $|\mathcal{B}_y|$ be two chambers in $|\mathcal{B}|$. Since there is an apartment $\mathcal{A} \subset \mathcal{B}$ which contains both x and y, this induces an isometric embedding $|\mathcal{A}| \to |\mathcal{B}|$ whose image contains $|\mathcal{B}_x|$ and $|\mathcal{B}_y|$. Here, $|\mathcal{A}|$ is isomorphic to the Davis realization of the Coxeter group W and is called an *apartment* in $|\mathcal{B}|$.

Definition 3.6
For $c_1, c_2 \in \mathcal{B}$, define $d(c_1, c_2)$ to be the minimal length of word in W (with respect to the generating set I) that represents $\delta(c_1, c_2)$. For any two residues $\mathcal{R}_1, \mathcal{R}_2 \in \mathcal{B}$, we define $d(\mathcal{R}_1, \mathcal{R}_2) = \min\{d(c_1, c_2) \mid c_1 \in \mathcal{R}_1, c_2 \in \mathcal{R}_2\}$. It turns out that, for any $c_1 \in \mathcal{R}_1$ and $c_2 \in \mathcal{R}_2$ with $d(c_1, c_2) = d(\mathcal{R}_1, \mathcal{R}_2)$, $\delta(c_1, c_2)$ gives rise to the same element in W (see [1, Chapter 5.3.2]). This element is defined to be $\delta(\mathcal{R}_1, \mathcal{R}_2)$.

LEMMA 3.7
We have that $d(c_1, c_2) = 2d_{l^1}(c_1, c_2)$, where d_{l^1} means the l^1-distance in $|\mathcal{B}|$. Since c_1 and c_2 can also be viewed as vertices of rank 0 in $|\mathcal{B}|$, $d_{l^1}(c_1, c_2)$ makes sense.

Proof
If $\mathcal{B} = W$, then this lemma follows from the above description of the Davis realization of a Coxeter group. The general case can be reduced to this case by considering an apartment $|\mathcal{A}| \subset |\mathcal{B}|$ which contains c_1 and c_2. Note that $|\mathcal{A}|$ is convex in $|\mathcal{B}|$. $\square$

Given a residue $\mathcal{R} \subset \mathcal{B}$, there is a well-defined nearest point projection map as follows.

THEOREM 3.8 ([1, Proposition 5.34])
Let $\mathcal{R}$ be a residue, and let c be a chamber. Then there exists a unique $c' \in \mathcal{R}$ such that $d(c, c') = d(\mathcal{R}, c)$.

This projection is compatible with several other projections in the following sense. Let $|\mathcal{R}| \subset |\mathcal{B}|$ be the convex subcomplex corresponding to $\mathcal{R}$. Let c and c' be as above. We also view them as vertices of rank 0 in $|\mathcal{B}|$. Let c_1 be the combinatorial projection of c onto $|\mathcal{R}|$ (see [34, Lemma 13.8]), and let c_2 be the CAT(0) projection of c onto $|\mathcal{R}|$.

LEMMA 3.9
We have that $c' = c_1 = c_2$.

Proof
$c_1 = c_2$ follows from Lemma 3.10 below. To see $c' = c_1$, by Lemma 3.7, it suffices to prove that c_1 is of rank 0. When $\mathcal{B} = W$, this follows from $c_1 = c_2$, since we can work with the cubical completion of the Cayley graph of W instead of $|W|$ (the latter

is the cubical subdivision of the former) and apply [36, Lemma 2.3]. The general case follows by considering an apartment $|\mathcal{A}| \subset |\mathcal{B}|$ which contains c_1 and c. Note that in this case $|\mathcal{A}| \cap |\mathcal{R}|$ can be viewed as a residue in $|\mathcal{A}|$. ☐

LEMMA 3.10
Let C_1 be a convex subcomplex in a CAT(0) *cube complex C. Pick a vertex $x \in C$. Let x_1 (resp., x_2) be the combinatorial projection (resp.,* CAT(0) *projection) of x onto C_1. Then $x_1 = x_2$.*

Proof
By [36, Lemma 2.3], x_2 is a vertex. If $x_2 \neq x_1$, then by [34, Lemma 13.8], the concatenation of the combinatorial geodesic ω_1 which connects x_2 and x_1 and the combinatorial geodesic ω_2 which connects x_1 and x is a combinatorial geodesic connecting x and x_2. By [34, Proposition 13.7], $\omega_1 \subset C_1$. Let $e \subset \omega_1$ be the edge that contains x_2, and let y be the other endpoint of e. Then y and x are on the same side of the hyperplane dual to e (see [34, Lemma 13.1]). It is not hard to see $d(y,x) < d(x_2,x)$ (here d denotes the CAT(0) distance), which yields a contradiction. ☐

Definition 3.11
Let $\mathrm{proj}_{\mathcal{R}}$ be the map defined in Theorem 3.8. Two residues $\mathcal{R}_1$ and $\mathcal{R}_2$ are *parallel* if $\mathrm{proj}_{\mathcal{R}_1}(\mathcal{R}_2) = \mathcal{R}_1$ and $\mathrm{proj}_{\mathcal{R}_2}(\mathcal{R}_1) = \mathcal{R}_2$. In this case, $\mathrm{proj}_{\mathcal{R}_1}$ and $\mathrm{proj}_{\mathcal{R}_2}$ induce mutually inverse bijections between $\mathcal{R}_1$ and $\mathcal{R}_2$. These bijections are called *parallelism maps* between $\mathcal{R}_1$ and $\mathcal{R}_2$. They are also isomorphisms of chamber systems; that is, they map residues to residues (see [1, Proposition 5.37]).

It follows from the uniqueness of the projection map that if $f : \mathcal{R} \to \mathcal{R}'$ is the parallelism map between two parallel residues and $\mathcal{R}_1 \subset \mathcal{R}$ is a residue, then $\mathcal{R}_1$ and $f(\mathcal{R}_1)$ are parallel, and the parallelism map between $\mathcal{R}_1$ and $f(\mathcal{R}_1)$ is induced by f.

LEMMA 3.12
If $\mathcal{R}_1$ and $\mathcal{R}_2$ are parallel, then $|\mathcal{R}_1|$ and $|\mathcal{R}_2|$ are parallel with respect to the CAT(0)*-metric on $|\mathcal{B}|$. Moreover, the parallelism map between $\mathcal{R}_1$ and $\mathcal{R}_2$ induced by $\mathrm{proj}_{\mathcal{R}_1}$ and $\mathrm{proj}_{\mathcal{R}_2}$ is compatible with the* CAT(0) *parallelism between $|\mathcal{R}_1|$ and $|\mathcal{R}_2|$ induced by* CAT(0) *projections.*

Proof
By Lemma 3.9, it suffices to show that, for any residue $\mathcal{R} \in \mathcal{B}$, $|\mathcal{R}|$ is the convex hull of the vertices of rank 0 inside $|\mathcal{R}|$. This is clear when $\mathcal{B} = W$ if one considers the

cubical completion of the Cayley graph of W. The general case also follows, since $|\mathcal{R}|$ is a union of apartments in $|\mathcal{R}|$, and $|\mathcal{R}|$ is convex in $|\mathcal{B}|$. $\square$

It follows that if $\mathcal{R}_1$ and $\mathcal{R}_2$ are parallel residues and if $\mathcal{R}_2$ and $\mathcal{R}_3$ are parallel residues, then $\mathcal{R}_1$ is parallel to $\mathcal{R}_3$. Moreover, let f_{ij} be the parallelism map from $\mathcal{R}_i$ to $\mathcal{R}_j$ induced by the projection map. Then $f_{13} = f_{23} \circ f_{12}$.

Given chamber systems $C_1, \dots, C_k$ over $I_1, \dots, I_k$, their direct product $C_1 \times \cdots \times C_k$ is a chamber system over the disjoint union $I_1 \sqcup \cdots \sqcup I_k$. Its chambers are k-tuples $(c_1, \dots, c_k)$ with $c_t \in C_t$. For $i \in I_t$, $(c_1, \dots, c_k)$ is i-adjacent to $(d_1, \dots, d_k)$ if $c_j = d_j$ for $j \neq t$ and c_t and d_t are i-adjacent.

Suppose the defining graph Γ of the right-angled Coxeter group W admits a join decomposition $\Gamma = \Gamma_1 \circ \Gamma_2 \circ \cdots \circ \Gamma_k$. Let $I = \bigcup_{i=1}^{k} I_i$ be the corresponding decomposition of the vertex set of Γ, and let $W = \prod_{i=1}^{k} W_i$ be the induced product decomposition of W. Pick a chamber $c \in \mathcal{B}$, and let $\mathcal{B}_i$ be the I_i-residue that contains c. Define a map $\phi : \mathcal{B} \to \mathcal{B}_1 \times \mathcal{B}_2 \times \cdots \times \mathcal{B}_k$ by $\phi(d) = (\mathrm{proj}_{\mathcal{B}_1}(d), \mathrm{proj}_{\mathcal{B}_2}(d), \dots, \mathrm{proj}_{\mathcal{B}_k}(d))$ for any chamber $d \in \mathcal{B}$.

THEOREM 3.13 ([52, Theorem 3.10])
The definition of ϕ does not depend on the choice of c, and ϕ is an isomorphism of buildings.

It follows from the definition of the Davis realization that there is a natural isomorphism $|\mathcal{B}_1 \times \mathcal{B}_2 \times \cdots \times \mathcal{B}_k| \cong |\mathcal{B}_1| \times |\mathcal{B}_2| \times \cdots \times |\mathcal{B}_k|$. Thus, we have a product decomposition $|\mathcal{B}| \cong |\mathcal{B}_1| \times |\mathcal{B}_2| \times \cdots \times |\mathcal{B}_k|$, where the isomorphism is induced by CAT(0) projections from $|\mathcal{B}|$ to $|\mathcal{B}_i|$'s. (This is a consequence of Lemma 3.12.)

We define the *parallel set* of a residue $\mathcal{R} \subset \mathcal{B}$ to be the union of all residues in $\mathcal{B}$ that are parallel to $\mathcal{R}$. Now we show that parallelism preserves types of residues in right-angled buildings. (This may not be true for buildings which are not right-angled.)

LEMMA 3.14
Suppose $\mathcal{R}$ is a J-residue. Let $J^{\perp} \subset I$ be the collection of vertices in Γ which are adjacent to every vertex in J. Then the following statements hold.
(1) *If another residue $\mathcal{R}'$ is parallel to $\mathcal{R}$, then $\mathcal{R}'$ is a J-residue.*
(2) *The parallel set of $\mathcal{R}$ is the $J \cup J^{\perp}$-residue that contains $\mathcal{R}$.*

Proof
Suppose that $\mathcal{R}'$ is a J_1-residue. Let $w = \delta(\mathcal{R}, \mathcal{R}')$ (see Definition 3.6). It follows from [1, Lemma 5.36(2)] that $\mathcal{R}'$ is a $(J \cap wJ_1w^{-1})$-residue. Since $\mathcal{R}$ and $\mathcal{R}'$ are

parallel, they have the same rank; thus, $J = wJ_1w^{-1}$. By considering the Abelianization of the right-angled Coxeter group W, we deduce that $J = J_1$ (this proves the first assertion of the lemma) and w commutes with each element in J. Thus, w belongs to the subgroup generated by $J^{\perp}$, and $\mathcal{R}'$ is in the $J \cup J^{\perp}$-residue $\mathcal{S}$ that contains $\mathcal{R}$. Then the parallel set of $\mathcal{R}$ is contained in $\mathcal{S}$. It remains to prove that every J-residue in $\mathcal{S}$ is parallel to $\mathcal{R}$, but this follows from Theorem 3.13. □

Pick a vertex $v \in |\mathcal{B}|$ of rank k, and let $\mathcal{R} = \prod_{i=1}^{k} \mathcal{R}_i$ be the associated residue with its product decomposition. Let $\{v_\lambda\}_{\lambda\in\Lambda}$ be the collection of vertices that are adjacent to v. Then there is a decomposition $\{v_\lambda\}_{\lambda\in\Lambda} = \{v_\lambda \leq v\} \sqcup \{v_\lambda > v\}$, where $\{v_\lambda > v\}$ denotes the collection of vertices whose associated residues contain $\mathcal{R}$. This induces a decomposition $Lk(v, |\mathcal{B}|) = K_1 * K_2$ of the link of v in $|\mathcal{B}|$ (see [12, Definition I.7.15]) into a spherical join of two CAT(1) all-right spherical complexes. Note that K_2 is finite, since $\{v_\lambda > v\}$ is finite. Moreover, $K_1 \cong Lk(v, |\mathcal{R}|)$. However, $|\mathcal{R}| \cong \prod_{i=1}^{k} |\mathcal{R}_i|$; thus, K_1 is the spherical join of k discrete sets such that elements in each of these discrete sets are in one-to-one correspondence to elements in some $\mathcal{R}_i$. Now we can deduce from this the following result.

LEMMA 3.15
Suppose $\mathcal{B}$ is a right-angled building such that each of its residues of rank 1 *contains infinitely many elements. If $\alpha : |\mathcal{B}| \to |\mathcal{B}|$ is a cubical isomorphism, then α preserves the rank of vertices in $|\mathcal{B}|$.*

4. Restriction quotients

In this section we study restriction quotients, a certain type of mapping between CAT(0) cube complexes introduced by Caprace and Sageev [14]. These play a central role in our story.

We first show in Section 4.1 that restriction quotients can be characterized in several different ways (see Theorem 4.4). We then show in Section 4.2 that a restriction quotient $f : Y \to Z$ determines fiber data that satisfies certain conditions; conversely, given such fiber data, one may construct a restriction quotient inducing the given data, which is unique up to equivalence. This correspondence will later be applied to construct restriction quotients over right-angled buildings. Sections 4.3 and 4.4 deal with the behavior of restriction quotients under group actions and quasi-isometries.

4.1. Quotient maps between CAT(0) *cube complexes*

We recall the notion of a restriction quotient from [14, Section 2.3] (see [35] for the background on wallspaces).

Definition 4.1
Let Y be a CAT(0) cube complex, and let $\mathcal{H}$ be the collection of walls in the 0-skeleton $Y^{(0)}$ corresponding to the hyperplanes in Y. Pick a subset $\mathcal{K} \subset \mathcal{H}$, and let $Y(\mathcal{K})$ be the CAT(0) cube complex associated with the wallspace $(Y^{(0)}, \mathcal{K})$. Then every 0-cube of the wallspace $(Y^{(0)}, \mathcal{H})$ gives rise to a 0-cube of $(Y^{(0)}, \mathcal{K})$ by restriction. This can be extended to a surjective cubical map $q : Y \to Y(\mathcal{K})$, which is called the *restriction quotient* arising from the subset $\mathcal{K} \subset \mathcal{H}$.

The following example motivates many of the constructions in this article.

Example 4.2 (The canonical restriction quotient of $X_e(\Gamma)$)
For a fixed graph Γ, let $S_e \to |c|_\Gamma$ and $X_e \to S_e$ be the mappings associated with the exploded Salvetti complex, as defined after Definition 1.12. Let $\mathcal{K}$ be the collection of hyperplanes in $X_e(\Gamma)$ dual to edges $e \subset X_e$ that project to edges under the composition $X_e \to S_e \to |c|_\Gamma$. Then the *canonical restriction quotient of* $G = G(\Gamma)$ is the restriction quotient arising from $\mathcal{K}$.

Let $q : Y \to Y(\mathcal{K})$ be a restriction quotient. Pick an edge $e \subset Y$. If e is dual to some element in $\mathcal{K}$, then $q(e)$ is an edge; otherwise, $q(e)$ is a point. The edge e is called *horizontal* in the former case and *vertical* in the latter case. We use the words vertical and horizontal, since we would like to think of q as a kind of “fibration.”

We record the following simple observation.

LEMMA 4.3
Let $\alpha : Y \to Y$ be a cubical CAT(0) *automorphism of Y that maps vertical edges to vertical edges and horizontal edges to horizontal edges. Then α descends to an automorphism $Y(\mathcal{K}) \to Y(\mathcal{K})$.*

The following result shows that restriction quotients may be characterized in several different ways.

THEOREM 4.4
If $f : Y \to Z$ is a surjective cubical map between two CAT(0) *cube complexes, then the following conditions are equivalent.*

(1) *The inverse image of each vertex of Z is convex.*
(2) *The inverse image of every point in Z is convex.*
(3) *The inverse image of every convex subcomplex of Z is convex.*
(4) *The inverse image of every hyperplane in Z is a hyperplane.*

(5) *f is equivalent to a restriction quotient, that is, for some set of walls $\mathcal{K}$ in Y, there is a cubical isomorphism $\phi : Z \to Y(\mathcal{K})$ making the following diagram commute:*

$$\begin{tikzcd} Y \arrow[d, "f"'] \arrow[dr, "q"] & \\ Z \arrow[r, "\phi"'] & Y(\mathcal{K}) \end{tikzcd}$$

The proof of Theorem 4.4 will take several lemmas. For the remainder of Section 4.1, we fix CAT(0) cube complexes Y and Z and a (not necessarily surjective) cubical map $f : Y \to Z$.

LEMMA 4.5
Let $\sigma \subset Z$ be a cube, and let Y_σ be the union of cubes in Y whose image under f is exactly σ. Then the following statements hold.

(1) *If $y \in \sigma$ is an interior point, then $f^{-1}(y) \subset Y_\sigma$.*

(2) *$f^{-1}(y)$ has a natural induced structure as a cube complex; moreover, there is a natural isomorphism of cube complexes $Y_\sigma \cong f^{-1}(y) \times \sigma$.*

(3) *If $\sigma_1 \subset \sigma_2$ are cubes of Z and $y_i \in \sigma_i$ are interior points, then there is a canonical embedding $f^{-1}(y_2) \hookrightarrow f^{-1}(y_1)$. Moreover, these embeddings are compatible with the composition of inclusions.*

LEMMA 4.6

(1) *For every $y \in Z$, every connected component of $f^{-1}(y)$ is a convex subset of Y.*

(2) *For every convex subcomplex $A \subset Z$, every connected component of $f^{-1}(A)$ is a convex subcomplex of Y.*

Proof
First we prove (1). Let σ be the support of y, and let $Y_\sigma \cong f^{-1}(y) \times \sigma$ be the subcomplex defined as above. It suffices to show that Y_σ is locally convex. Pick a vertex $x \in Y_\sigma$, and let $\{e_i\}_{i=1}^n$ be a collection of edges in Y_σ that contains x. It suffices to show that if these edges span an n-cube $\eta \subset Y$, then $\eta \subset Y_\sigma$. It suffices to consider the case when all e_i's are orthogonal to σ, in which case it follows from Definition 3.4 that $\eta \times \sigma \subset Y_\sigma$.

To see (2), pick an n-cube $\eta \subset Y$, and let $\{e_i\}_{i=1}^n$ be the edges of η at one corner $c \subset \eta$. It suffices to show that if $f(e_i) \subset A$, then $f(\eta) \subset A$. Note that $f(\eta)$ is a cube, and every edge of this cube which emanates from the corner $f(c)$ is contained in A. Thus, $f(\eta) \subset A$ by the convexity of A. ☐

LEMMA 4.7
Let $f : Y \to Z$ be a cubical map as above. Then the following statements hold.
(1) *The inverse image of each hyperplane of Z is a disjoint union of hyperplanes in Y.*
(2) *If the inverse image of each hyperplane of Z is a single hyperplane, then for each point $y \in Z$, the point inverse $f^{-1}(y)$ is connected and, hence, is convex by Lemma 4.6.*

Proof
It follows from Definition 3.4 that the inverse image of each hyperplane of Z is a union of hyperplanes. If two of them were to intersect, then there would be a 2-cube in Y with two consecutive edges mapped to the same edge in Z, which is impossible.

Now we prove (2). It suffices to consider the case in which y is the center of some cube in Z. In this case, y is a vertex in the first cubical subdivision of Z, and f can be viewed as a cubical map from the first cubical subdivision of Y to the first cubical subdivision of Z such that the inverse image of each hyperplane is a single hyperplane; thus, it suffices to consider the case in which y is a vertex of Z.

Suppose that $f^{-1}(y)$ contains two connected components A and B. Pick a combinatorial geodesic ω of shortest distance that connects vertices in A and vertices in B. Note that $f(\omega)$ is a nontrivial edge-loop in Z; otherwise, we will have $\omega \subset f^{-1}(y)$. It follows that there exist two different edges e_1 and e_2 of ω mapping to parallel edges in Y. The hyperplanes dual to e_1 and e_2 are different, yet they are mapped to the same hyperplane in Y, which is a contradiction. $\square$

LEMMA 4.8
If f is surjective and, for any vertex $v \in Z$, $f^{-1}(v)$ is connected, then the inverse image of each hyperplane of Z is a single hyperplane.

Proof
Let $h \subset Z$ be a hyperplane. By Lemma 4.7, $f^{-1}(h) = \bigsqcup_{\lambda \in \Lambda} h_\lambda$, where each h_λ is a hyperplane in Y. Since f is surjective, $\{f(h_\lambda)\}_{\lambda \in \Lambda}$ is a collection of subcomplexes of h that cover h. Thus, there exist $h_1, h_2 \in \{h_\lambda\}_{\lambda \in \Lambda}$ ($h_1 \neq h_2$) and a vertex $u \in h$ such that $u \subset f(h_1) \cap f(h_2)$. Let $e \subset Z$ be the edge such that $u = e \cap h$. Then there exist edges $e_1, e_2 \subset Y$ such that $e_i \cap h_i \neq \emptyset$ and $f(e_i) = e$ for $i = 1, 2$. Since $h_1 \cap h_2 = \emptyset$, a case study implies that there exist x_1 and x_2 which are endpoints of e_1 and e_2, respectively, such that
(1) these two points are separated by at least one of h_1 and h_2;
(2) they are mapped to the same endpoint $y \in e$.
It follows that $f^{-1}(y)$ is disconnected, which is a contradiction. $\square$

Remark 4.9
If f is not surjective, then the above conclusion is not necessarily true. Consider the map from $A = [0,3] \times [0,1]$ to the unit square which collapses the $[0,1]$-factor in A and maps $[0,3]$ to three consecutive edges on the boundary of the unit square.

LEMMA 4.10
If $q : Y \to Y(\mathcal{K})$ is the restriction quotient as in Definition 4.1, then the inverse image of each hyperplane in $Y(\mathcal{K})$ is a single hyperplane in Y. Conversely, suppose that $f : Y \to Z$ is a surjective cubical map between CAT(0) *cube complexes such that the inverse image of each hyperplane is a hyperplane. Let $\mathcal{K}$ be the collection of walls arising from inverse images of hyperplanes in Z. Then there is a natural isomorphism $i : Z \cong Y(\mathcal{K})$ which fits into the following commutative diagram:*

$$\begin{tikzcd} Y \arrow[r, "f"] \arrow[dr, "q"'] & Z \arrow[d, "i"] \\ & Y(\mathcal{K}) \end{tikzcd}$$

Proof
Define two vertices of Y to be $\mathcal{K}$-equivalent if and only if they are not separated by any wall in $\mathcal{K}$. This defines an equivalence relation on vertices of Y, and the corresponding equivalence classes are called $\mathcal{K}$-classes. For each $\mathcal{K}$-class C and every wall in $\mathcal{K}$, we may choose the half-space that contains C; it follows that the points in C are exactly the set of vertices contained in the intersection of such half-spaces, and thus, C is the vertex set of a convex subcomplex of Y. Note that each $\mathcal{K}$-class determines a 0-cube of $(Y^0, \mathcal{K})$ and, hence, is mapped to this 0-cube under q. It follows that the inverse image of every vertex in $Y(\mathcal{K})$ is convex; thus, by Lemma 4.8, the inverse image of a hyperplane is a hyperplane.

It remains to prove the converse. Note that the inverse image of each half-space in Z under f is a half-space of Y. Moreover, the surjectivity of f implies that f maps hyperplanes to hyperplanes and half-spaces to half-spaces. Pick a vertex $y \in Z$, and let $\{H_\lambda\}_{\lambda \in \lambda}$ be the collection of half-spaces in Z that contain y. Then $f^{-1}(y) \subset \bigcap_{\lambda \in \Lambda} f^{-1}(H_\lambda)$, every vertex of $\bigcap_{\lambda \in \Lambda} f^{-1}(H_\lambda)$ is mapped to y by f, and thus, the vertex set of $f^{-1}(y)$ is a $\mathcal{K}$-class. This induces a bijective map from $Z^{(0)}$ to the vertex set of $Y(\mathcal{K})$, which extends to an isomorphism. The above diagram commutes, since it commutes when restricted to the 0-skeleton. $\square$

Proof of Theorem 4.4
The equivalence of (4) and (5) follows from Lemma 4.10. (1) $\Rightarrow$ (4) follows from Lemma 4.8, (4) $\Rightarrow$ (2) follows from Lemma 4.7, and (3) $\Rightarrow$ (1) is obvious. It suf-

fices to show (2) $\Rightarrow$ (3). Pick a convex subcomplex $K \subset Z$, and let $\{R_\lambda\}_{\lambda\in\Lambda}$ be the collection of cubes in K. For each R_λ, let Y_{R_λ} be the subcomplex defined in Lemma 4.5. $Y_{R_\lambda} \neq \emptyset$ since f is surjective and Y_{R_λ} is connected by (2). If $R_\lambda \subset R_{\lambda'}$, then $Y_{R_\lambda} \cap Y_{R_{\lambda'}} \neq \emptyset$. Thus, $f^{-1}(K) = \bigcup_{\lambda\in\Lambda} Y_{R_\lambda}$ is connected and, hence, convex. □

4.2. *Restriction maps versus fiber functors*

If $q : Y \to Z$ is a restriction quotient between CAT(0) cube complexes, then we may express the fiber structure in categorical language as follows. Let $\mathrm{Face}(Z)$ denote the face poset of Z, viewed as a category, that is, the objects are cubes and morphisms are inclusions. Let CCC denote the category whose objects are nonempty CAT(0) cube complexes and whose morphisms are convex cubical embeddings. By Lemma 4.5, we obtain a contravariant functor $\Psi_q : \mathrm{Face}(Z) \to \mathrm{CCC}$, which takes a face σ to $q^{-1}(y)$ for an interior point $y \in \sigma$.

Definition 4.11

The contravariant functor Ψ_q is the *fiber functor* of the restriction quotient $q : Y \to Z$.

Let Ψ be a contravariant functor from $\mathrm{Face}(Z)$ to CCC. For notational brevity, for any inclusion $i : \sigma_1 \to \sigma_2$, we will often denote the map $\Psi(i) : \Psi(\sigma_2) \to \Psi(\sigma_1)$ simply by $\Psi(\sigma_2) \to \Psi(\sigma_1)$, suppressing the name of the map.

Note that if $\sigma_1 \subset \sigma_2 \subset \sigma_3$, then the functor property implies that the image of $\Psi(\sigma_3) \to \Psi(\sigma_1)$ is a convex subcomplex of the image of $\Psi(\sigma_2) \to \Psi(\sigma_1)$. In particular, if v is a vertex of a cube σ, then the image of $\Phi(\sigma) \to \Psi(v)$ is a convex subcomplex of the intersection

$$\bigcap_{v \subsetneq e \subset \sigma^{(1)}} \mathrm{Im}\big(\Psi(e) \to \Psi(v)\big).$$

Definition 4.12

Let Z be a cube complex. A contravariant functor $\Psi : \mathrm{Face}(Z) \to \mathrm{CCC}$ is 1-*determined* if, for every cube $\sigma \in \mathrm{Face}(Z)$ and every vertex $v \in \sigma^{(0)}$,

$$\mathrm{Im}\big(\Psi(\sigma) \longrightarrow \Psi(v)\big) = \bigcap_{v \subsetneq e \subset \sigma^{(1)}} \mathrm{Im}\big(\Psi(e) \to \Psi(v)\big). \tag{4.13}$$

LEMMA 4.14

If $q : Y \to Z$ is a restriction quotient, then the fiber functor $\Psi : \mathrm{Face}(Z) \to \mathrm{CCC}$ is 1-determined.

Proof
Pick $\sigma \in \mathrm{Face}(Z)$ and $v \in \sigma^{(0)}$. We know that $\mathrm{Im}(\Psi(\sigma) \to \Psi(v))$ is a nonempty convex subcomplex of $\bigcap_{v \subsetneq e \subset \sigma^{(1)}} \mathrm{Im}(\Psi(e) \to \Psi(v))$, so to establish (4.13) we need only show that the two convex subcomplexes have the same 0-skeleton.

Pick a vertex $w \in \mathrm{Im}(\Psi(\sigma) \to \Psi(v))$, and let $w' \in \bigcap_{v \subsetneq e \subset \sigma^{(1)}} \mathrm{Im}(\Psi(e) \to \Psi(v))$ be a vertex adjacent to w. We let $\tau \in \mathrm{Face}(Y)$ denote the edge spanned by w, w'. For every edge e of Z with $v \subsetneq e \subset \sigma^{(1)}$, let $\hat{e} \subset Y^{(1)}$ denote the edge with $q(\hat{e}) = e$ that contains w. By assumption, the collection of edges $\{\tau\} \cup \{\hat{e}\}_{v \subsetneq e \subset \sigma^{(1)}}$ determines a complete graph in the link of w and, therefore, is contained in a cube $\hat{\sigma}$ of dimension $1 + \dim \sigma$. Then $q(\hat{\sigma}) = \sigma$ and $\tau \subset \hat{\sigma}$; this implies that $\tau \subset \mathrm{Im}(\Psi(\sigma) \to \Psi(v))$.

Since the 1-skeleton of $\bigcap_{v \subsetneq e \subset \sigma^{(1)}} \mathrm{Im}(\Psi(e) \to \Psi(v))$ is connected, we conclude that it coincides with the 1-skeleton of $\mathrm{Im}(\Psi(\sigma) \to \Psi(v))$. By convexity, we get (4.13). □

THEOREM 4.15
Let Z be a CAT(0) *cube complex, and let $\Psi : \mathrm{Face}(Z) \to \mathrm{CCC}$ be a 1-determined contravariant functor. Then there is a restriction quotient $q : Y \to Z$ such that the associated fiber functor $\Psi_q : \mathrm{Face}(Z) \to \mathrm{CCC}$ is equivalent by a natural transformation to Ψ.*

Proof
We first construct the cube complex Y and then verify that it has the desired properties. We begin with the disjoint union $\bigsqcup_{\sigma \in \mathrm{Face}(Z)} (\sigma \times \Psi(\sigma))$, and for every inclusion $\sigma \subset \tau$, we glue the subset $\sigma \times \Psi(\tau) \subset \tau \times \Psi(\tau)$ to $\sigma \times \Psi(\sigma)$ by using the map

$$\sigma \times \Psi(\tau) \xrightarrow{\mathrm{id}_\sigma \times \Psi(\sigma \subset \tau)} \sigma \times \Psi(\sigma).$$

One checks that the cubical structure on $\bigsqcup_{\sigma \in \mathrm{Face}(Z)} (\sigma \times \Psi(\sigma))$ descends to the quotient Y, the projection maps $\sigma \times \Psi(\sigma) \to \sigma$ descend to a cubical map $q : Y \to Z$, and for every $\sigma \in \mathrm{Face}(Z)$, the union of the cubes $\hat{\sigma} \subset Y$ such that $f(\hat{\sigma}) = \sigma$ is a copy of $\sigma \times \Psi(\sigma)$.

We now verify that links in Y are flag complexes. Let v be a 0-cube in Y, and suppose $\sigma_1, \ldots, \sigma_k$ are 1-cubes containing v such that for all $1 \le i \ne j \le k$ the 1-cubes σ_i, σ_j span a 2-cube σ_{ij} in the link of v. We may assume after reindexing that for some $h \ge 0$ the image $q(\sigma_i)$ is a 1-cube in Z if $i \le h$ and a 0-cube if $i > h$.

Since $\Psi(v)$ is a CAT(0) cube complex, the edges $\{\sigma_i\}_{i > h}$ span a cube $\sigma_{\mathrm{vert}} \subset q^{-1}(v)$.

For $1 \le i \ne j \le h$, the 2-cube σ_{ij} projects to a 2-cube $q(\sigma_{ij})$ spanned by the two edges $q(\sigma_i), q(\sigma_j)$. Since Z is a CAT(0) cube complex, the edges $\{q(\sigma_i)\}_{i \le h}$ span an h-cube $\bar{\sigma}_{\mathrm{hor}} \subset Z$. By the 1-determined property, we get that $\mathrm{Im}(\Psi(\bar{\sigma}_{\mathrm{hor}}) \to \Psi(v))$ contains v, and so there is an h-cube $\sigma_{\mathrm{hor}} \subset Y$ containing v such that $q(\sigma_{\mathrm{hor}}) = \bar{\sigma}_{\mathrm{hor}}$.

Fix $1 \le i \le h$. Then for $j > h$, the 2-cube σ_{ij} projects to $q(\sigma_i)$, and hence, σ_j belongs to $\mathrm{Im}(\Psi(\sigma_i) \to \Psi(v))$. If $j,k > h$, then σ_j, σ_k both belong to $\mathrm{Im}(\Psi(\sigma_i) \to \Psi(v))$, and by the convexity of $\mathrm{Im}(\Psi(\sigma_i) \to \Psi(v))$ in $\Psi(v)$, we get that σ_{jk} also belongs to $\mathrm{Im}(\Psi(\sigma_i) \to \Psi(v))$. Applying convexity again, we get that $\sigma_{\mathrm{vert}} \subset \mathrm{Im}(\Psi(\sigma_i) \to \Psi(v))$. By the 1-determined property, it follows that $\sigma_{\mathrm{vert}} \subset \mathrm{Im}(\Psi(\bar{\sigma}_{\mathrm{hor}}) \to \Psi(v))$. This yields a k-cube $\sigma \subset Y$ containing $\sigma_{\mathrm{hor}} \cup \sigma_{\mathrm{vert}}$, which is spanned by $\sigma_1, \ldots, \sigma_k$.

Thus, we have shown that links in Y are flag complexes. The fact that the fibers of $f : Y \to Z$ are contractible implies that Y is contractible (in particular, simply connected), so Y is CAT(0). ☐

We now observe that the construction of restriction quotients is compatible with product structure.

LEMMA 4.16 (Behavior under products)
For $i \in \{1,2\}$ let $q_i : Y_i \to Z_i$ be a restriction quotient with fiber functor $\Psi_i : \mathrm{Face}(Z_i) \to \mathrm{CCC}$. Then the product $q_1 \times q_2 : Y_1 \times Y_2 \to Z_1 \times Z_2$ is a restriction quotient with fiber functor given by the product

$$\mathrm{Face}(Z_1 \times Z_2) \simeq \mathrm{Face}(Z_1) \times \mathrm{Face}(Z_2) \xrightarrow{\Psi_1 \times \Psi_2} \mathrm{CCC} \times \mathrm{CCC} \xrightarrow{\times} \mathrm{CCC}.$$

In particular, if one starts with CAT(0) *cube complexes Z_i and fiber functors $\Psi_i : Z_i \to \mathrm{CCC}$ for $i \in \{1,2\}$, then the product fiber functor defined as above is the fiber functor of the product of the restriction quotients associated to the Ψ_i's.*

4.3. Equivariance properties
We now discuss isomorphisms between restriction quotients and the naturality properties of the restriction quotient associated with a fiber functor. Suppose that we have a commutative diagram

$$\begin{tikzcd} Y_1 \arrow[r, "\hat{\alpha}"] \arrow[d, "q_1"'] & Y_2 \arrow[d, "q_2"'] \\ Z_1 \arrow[r, "\alpha"] & Z_2 \end{tikzcd}$$

where the q_i's are restriction quotients and $\alpha, \hat{\alpha}$ are cubical isomorphisms. Let $\Psi_i : \mathrm{Face}(Z_i) \to \mathrm{CCC}$ be the fiber functor associated with q_i. Note that the pair $\alpha, \hat{\alpha}$ allows us to compare the two fiber functors, since for every $\sigma \in \mathrm{Face}(Z_1)$, the map

$\hat{\alpha}$ induces a cubical isomorphism between $\Psi_1(\sigma)$ and $\Psi_2(\alpha(\sigma))$, and this is compatible with maps induced with inclusions of faces. This may be stated more compactly by saying that $\hat{\alpha}$ induces a natural isomorphism between the fiber functors Ψ_1 and $\Psi_2 \circ \mathrm{Face}(\alpha)$, where $\mathrm{Face}(\alpha) : \mathrm{Face}(Z_1) \to \mathrm{Face}(Z_2)$ is the poset isomorphism induced by α. Here the term *natural isomorphism* is being used in the sense of category theory, that is, a natural transformation that has an inverse that is also a natural transformation.

Now suppose that for $i \in \{1,2\}$ we have a CAT(0) cube complex Z_i and a 1-determined fiber functor $\Psi_i : \mathrm{Face}(Z_i) \to \mathrm{CCC}$. Let $f_i : Y_i \to Z_i$ be the associated restriction quotients. If we have a pair α, β, where $\alpha : Z_1 \to Z_2$ is a cubical isomorphism, and β is a natural isomorphism between the fiber functors Ψ_1 and $\Psi_2 \circ \mathrm{Face}(\alpha)$, then we get an induced map $\hat{\alpha} : Y_1 \to Y_2$, which may be defined by using the description of Y_i as the quotient of the disjoint collection $\{\sigma \times \Psi_i(\sigma)\}_{\sigma \in \mathrm{Face}(Z_i)}$. As a consequence of the above, having an action of a group G on a restriction quotient $f : Y \to Z$ is equivalent to having an action $G \curvearrowright Z$ together with a compatible "action" on the fiber functor Ψ_f, that is, a family $\{(\alpha(g), \beta(g))\}_{g \in G}$ as above that also satisfies an appropriate composition rule.

4.4. Quasi-isometric properties

We now consider the coarse geometry of restriction quotients; this amounts to a "coarsification" of the discussion in the preceding section. The relevant definition is a coarsification of the natural isomorphisms between fiber functors.

Definition 4.17

Let Z be a CAT(0) cube complex, and let $\Psi_i : \mathrm{Face}(Z) \to \mathrm{CCC}$ be fiber functors for $i \in \{1,2\}$. An (L, A)-*quasinatural isomorphism from* Ψ_1 *to* Ψ_2 is a collection $\{\phi(\sigma) : \Psi_1(\sigma) \to \Psi_2(\sigma)\}_{\sigma \in \mathrm{Face}(Z)}$ such that $\phi(\sigma)$ is an (L, A)-quasi-isometry for all $\sigma \in \mathrm{Face}(Z)$, and for every inclusion $\sigma \subset \tau$, the diagram

$$\begin{array}{ccc} \Psi_1(\tau) & \xrightarrow{\phi(\tau)} & \Psi_2(\tau) \\ \downarrow & & \downarrow \\ \Psi_1(\sigma) & \xrightarrow{\phi(\sigma)} & \Psi_2(\sigma) \end{array}$$

commutes up to error L.

Now for $i \in \{1,2\}$ let $f_i : Y_i \to Z$ be a finite-dimensional restriction quotient, with respective fiber functor $\Psi_i : \mathrm{Face}(Z) \to \mathrm{CCC}$. For any $\sigma \in \mathrm{Face}(Z)$, we identify $\Psi_i(\sigma)$ with $f_i^{-1}(b_\sigma)$, where $b_\sigma \in \sigma$ is the barycenter.

LEMMA 4.18
Suppose that we have a commutative diagram

$$\begin{tikzcd} Y_1 \arrow[rr, shift left, "\phi"] \arrow[dr, "f_1"'] & & Y_2 \arrow[ll, shift left, "\phi'"] \arrow[dl, "f_2"] \\ & Z & \end{tikzcd}$$

where ϕ, ϕ' are (L, A)-quasi-isometries that are A-quasi-inverses, that is, the compositions $\phi \circ \phi'$, $\phi' \circ \phi$ are at a distance less than A from the identity maps. Then the collection

$$\Big\{\Psi_1(\sigma) = f_1^{-1}(b_\sigma) \xrightarrow{\phi\big|_{f_1^{-1}(b_\sigma)}} f_2^{-1}(b_\sigma) = \Psi_2(\sigma)\Big\}_{\sigma\in\mathrm{Face}(Z)}$$

is an (L', A')-quasinatural isomorphism, where $L' = L'(L, A, \dim Y_i)$, $A' = A'(L, A, \dim Y_i)$.

Proof
By Theorem 4.4, the fiber $f_i^{-1}(b_\sigma)$ is a convex subset of Y_i and, hence, is isometrically embedded. Therefore, ϕ and ϕ' induce (L, A)-quasi-isometric embeddings $f_1^{-1}(b_\sigma) \to f_2^{-1}(b_\sigma)$, $f_2^{-1}(\sigma_b) \to f_1^{-1}(b_\sigma)$. If $\sigma \subset \tau$, then any point $x \in f_i^{-1}(b_\tau)$ lies at a distance less than $C = C(\dim Y_i)$ from a point in $f_i^{-1}(b_\sigma)$, and this implies that the collection of maps $\{\Psi_1(\sigma) \to \Psi_2(\sigma)\}_{\sigma\in\mathrm{Face}(Z)}$ is an (L', A')-quasinatural isomorphism as claimed. □

LEMMA 4.19
If $\{\phi(\sigma) : \Psi_1(\sigma) \to \Psi_2(\sigma)\}_{\sigma\in\mathrm{Face}(Z)}$ is an (L, A)-quasinatural isomorphism from Ψ_1 to Ψ_2, then it arises from a commutative diagram as in the previous lemma, where ϕ, ϕ' are (L', A')-quasi-isometries that are A'-quasi-inverses, and L', A' depend only on L, A, and $\dim Y_i$.

Proof
For every $\sigma \in \mathrm{Face}(Z)$, we may choose a quasi-inverse $\phi'(\sigma) : \Psi_2(\sigma) \to \Psi_1(\sigma)$ with uniform constants; this is also a quasinatural isomorphism. Identifying $f_i^{-1}(\mathrm{Int}(\sigma))$ with the product $\mathrm{Int}(\sigma) \times \Psi_i(\sigma)$, we define $\phi\big|_{f_1^{-1}(\mathrm{Int}(\sigma))}$ by

$$f_1^{-1}\big(\mathrm{Int}(\sigma)\big) = \mathrm{Int}(\sigma) \times \Psi_1(\sigma) \xrightarrow{\mathrm{id}_{\mathrm{Int}(\sigma)} \times \phi(\sigma)} \mathrm{Int}(\sigma) \times \Psi_2(\sigma) = f_2^{-1}(\sigma),$$

and we define ϕ' similarly using $\{\phi'(\sigma)\}_{\sigma\in\mathrm{Face}(Z)}$. One readily checks that ϕ, ϕ' are quasi-isometric embeddings that are also quasi-inverses, where the constants depend on L, A, and $\dim Y_i$. □

5. The $\mathbb{Z}$-blowup of a right-angled building

In this section Γ will be an arbitrary finite simplicial graph, and all buildings will be right-angled buildings modelled on the right-angled Coxeter group $W(\Gamma)$ with defining graph Γ. The reader may wish to review Section 3.4 for terminology and notation regarding buildings before proceeding.

The goal of this section is to examine restriction quotients $q : Y \to |\mathcal{B}|$, where the fibers are Euclidean spaces satisfying a dimension condition as in Theorems 1.5 and 1.6. For such restriction quotients, the fiber functor may be distilled down to a simpler set of information, called 1-data (see Definition 5.3); this is discussed in Section 5.2. Conversely, given a building $\mathcal{B}$ and certain blow-up data (Definition 5.6), one can construct a corresponding 1-determined fiber functor as in Section 4.2 (see Section 5.3). On the one hand, this establishes item (c) in the "Further results" section of Section 1 and, on the other hand, provides us a simpler set of information to work with which is equivalent to the fiber functor. Most of the material in Section 5.2 is not used in the later part of the article, so a hurried reader can skip this section and come back when needed. In the later parts of this section, we study several key properties of this blow-up construction and formulate a version which takes group actions into account.

5.1. *The canonical restriction quotient for a RAAG*

Let $G(\Gamma)$ be the RAAG with defining graph Γ, and let $\mathcal{B}_0(\Gamma)$ be the building associated with $G(\Gamma)$ (see [22, Section 5]). Then $G(\Gamma)$ can be identified with the set of chambers of $\mathcal{B}_0(\Gamma)$. Under this identification, the J-residues of $\mathcal{B}_0$, for J a collection of vertices in Γ, are the left cosets of the standard subgroups of $G(\Gamma)$ generated by J. Thus, the poset of spherical residues is exactly the poset of left cosets of standard Abelian subgroups of $G(\Gamma)$, which is also isomorphic to the poset of standard flats in $X(\Gamma)$.

We now revisit the discussion after Definition 1.12 and Example 4.2 in more detail and relate them to buildings. To simplify notation, we will write $G = G(\Gamma)$, $\mathcal{B}_0 = \mathcal{B}_0(\Gamma)$, and $X = X(\Gamma)$.

Let $|\mathcal{B}_0|$ be the Davis realization of the building $\mathcal{B}_0$. Then we have an induced isometric action $G \curvearrowright |\mathcal{B}_0|$, which is cocompact but not proper. It turns out that there is a natural way to blow up $|\mathcal{B}_0|$ to obtain a space $X_e = X_e(\Gamma)$ such that there is a geometric action $G \curvearrowright X_e$ and a G-equivariant restriction quotient map $X_e \to |\mathcal{B}_0|$.

X_e can be constructed as follows. First we constructed the *exploded Salvetti complex* $S_e = S_e(\Gamma)$, which was introduced in [9] (see also the discussion after Definition 1.12). For each vertex v in the vertex set $V(\Gamma)$ of Γ, we associate a copy of the "lollipop" $L_v = S_v \cup I_v$, which is the union of a unit circle S_v and a unit interval I_v along one point. Let $\star_v \in L_v$ be the free end of I_v. Let $T = \prod_{v \in V(\Gamma)} L_v$. Each clique $\Delta \subset \Gamma$ gives rise to a subcomplex $T_\Delta = \prod_{v\in\Delta} L_v \times \prod_{v\notin\Delta}\{\star_v\}$. Then S_e is the subcomplex of T which is the union of all such T_Δ's; here Δ is allowed to be empty. It is easy to check that S_e is a nonpositively curved cube complex. A *standard torus* in S_e is a subcomplex of form $\prod_{v\in\Delta} S_v \times \prod_{v\notin\Delta}\{\star_v\}$, where $\Delta \subset \Gamma$ is a clique. Note that there is a unique standard torus of dimension 0, which corresponds to the empty clique. There is a natural map $S_e = S_e(\Gamma) \to S(\Gamma)$ by collapsing the I_v-edge in each L_v-factor. This maps induces a one-to-one correspondence between standard tori in S_e and standard tori in $S(\Gamma)$. Note that there is also a one-to-one correspondence between vertices in S_e and standard tori in S_e.

Let X_e be the universal cover of S_e. Then X_e is a CAT(0) cube complex, and the action $G \curvearrowright X_e$ is geometric. The inverse images of standard tori in S_e are called *standard flats*. Note that each vertex in X_e is contained in a unique standard flat. We define a map between the 0-skeletons $p : X_e^{(0)}(\Gamma) \to |\mathcal{B}_0|^{(0)}$ as follows. Pick a G-equivariant identification between 0-dimensional standard flats in X and elements in G, and pick a G-equivariant map $\phi : X_e \to X$ induced by $S_e = S_e(\Gamma) \to S(\Gamma)$ described as above. Note that ϕ induces a one-to-one correspondence between standard flats in X_e and standard flats in X. This gives rise to a one-to-one correspondence between standard flats in X_e and left cosets of standard Abelian subgroups of G. For each $x \in X_e^{(0)}(\Gamma)$, we define $p(x)$ to be the vertex in $|\mathcal{B}_0|^{(0)}$ that represents the left coset of the standard Abelian subgroup of G which corresponds to the unique standard flat that contains x.

A *vertical edge* of X_e is an edge which covers some S_v-circle in S_e. A *horizontal edge* of X_e is an edge which covers some I_v-interval in S_e. Two endpoints of every vertical edge are in the same standard flat; thus, they are mapped by p to the same point in $|\mathcal{B}_0|^{(0)}$. More generally, for any given *vertical cube*, that is, every edge in this cube is a vertical edge, its vertex set is mapped by p to one point in $|\mathcal{B}_0|^{(0)}$. Pick a horizontal edge, and let $F_1, F_2 \subset X_e$ be standard flats which contain the two endpoints of this edge, respectively. Then $\phi(F_1)$ and $\phi(F_2)$ are two standard flats in X such that one is contained as a codimension 1 flat inside another. More generally, if σ is a horizontal cube, that is, each edge of σ is a horizontal edge, then by looking at the image of σ under the covering map $X_e \to S_e$, we know that the vertex set of σ corresponds to an interval in the poset of standard flats of X. Every cube in X_e splits as a product of a vertical cube and a horizontal cube. (Again this is clear by looking at cells in S_e.) Thus, we can extend p to a cubical map $p : X_e \to |\mathcal{B}_0|$.

By construction, for a vertex $v \in |\mathcal{B}_0|$ of rank n, $p^{-1}(v)$ is isometric to $\mathbb{E}^n$. It follows from Theorem 4.4 that p arises from a restriction quotient, and this is called the *canonical restriction quotient* for the RAAG G. This restriction quotient is exactly the one described in Example 4.2, since the hyperplanes in $\mathcal{K}$ of Example 4.2

are those which are dual to horizontal edges. We record the following immediate consequence of this construction.

LEMMA 5.1
Let $\sigma \subset |\mathcal{B}_0|$ be a cube, and let $v \in \sigma$ be the vertex of minimal rank in σ. Then for any interior point $x \in \sigma$, $p^{-1}(x)$ is isometric to $\mathbb{E}^{\mathrm{rank}(v)}$.

Remark 5.2
In the literature, there is a related cubical map $X \to |\mathcal{B}_0|$ defined as follows. First, we recall an alternative description of X. Actually, similar spaces can be defined for all Artin groups (not necessarily right-angled) and were introduced by Salvetti. We will follow the description in [16]. Let $G \to W(\Gamma)$ be the natural projection map. This map has a set-theoretic section defined by representing an element $w \in W$ by a minimal length positive word with respect to the standard generating set and setting $\sigma(w)$ to be the image of this word G. It follows from fundamental facts about Coxeter groups that σ is well defined. Let I be the vertex set of Γ, and for any $J \subset I$, let $W(J)$ be the subgroup of $W(\Gamma)$ generated by J. Let K be the geometric realization of the following poset:

$$\{g\sigma(W(J)) \mid g \in G, J \subset I, W(J) \text{ is finite}\}.$$

It turns out that K is isomorphic to the first barycentric subdivision of X. Let $G(J) \leq G$ be the subgroup generated by J. We associate each $g\sigma(W(J))$ with the left coset $gG(J)$, and this induces a cubical map from the first cubical subdivision of X to $|\mathcal{B}_0|$. However, this map is not a restriction quotient, since it has a lot of foldings (think of the special case when $G \cong \mathbb{Z}$).

5.2. Restriction quotients with Euclidean fibers
We remind the reader that, in this section, $W = W(\Gamma)$ will be the right-angled Coxeter group with defining graph Γ and standard generating set I. Let $\mathcal{B}$ be an arbitrary right-angled building modelled on W. Let S be the poset of spherical subsets of I, and let $|\mathcal{B}|$ be the Davis realization of $\mathcal{B}$. Let $q : Y' \to |\mathcal{B}|$ be an arbitrary restriction quotient satisfying the conclusion of Lemma 5.1. ($\mathcal{B}$ does not have to be the building associated with a RAAG and q does not have to be the canonical restriction quotient.)

Let Φ be the fiber functor associated with q (see Section 4.2). For any vertices $v, w \in |\mathcal{B}|$, we will write $v \leq w$ if and only if the residue associated with v is contained in the residue associated with w.

Let S^r be the poset of spherical residues in $\mathcal{B}$. Then Φ induces a functor Φ' from S^r to CCC (Section 4.2) as follows. Each element in S^r is associated with the fiber of the corresponding vertex in $|\mathcal{B}|$. If $s, t \in \mathsf{S}^r$ are two elements such that

$\mathrm{rank}(t) = \mathrm{rank}(s) + 1$ and $s < t$, then the associated vertices in $v_s, v_t \in |\mathcal{B}|$ are joined by an edge e_{st}. In this case $\Phi(e_{st}) \to \Phi(v_s)$ is an isomorphism, so we define the morphism $\Phi'(s) \to \Phi'(t)$ to be the map induced by $\Phi(e_{st}) \to \Phi(v_t)$. If $s, t \in \mathsf{S}^r$ are two arbitrary elements with $s \leq t$, then we find an ascending chain from s to t such that the difference between the ranks of adjacent elements in the chain is 1, and we define $\Phi'(s) \to \Phi'(t)$ to be the composition of those maps induced by the chain. It follows from the functor property of Φ that $\Phi'(s) \to \Phi'(t)$ does not depend on the choice of the chain, and Φ' is a functor. Recall that there is a one-to-one correspondence between elements in S^r and vertices of $|\mathcal{B}|$, so we will also view Φ' as a functor from the vertex set of $|\mathcal{B}|$ to CCC. Let $\sigma_1 \subset \sigma_2$ be faces in $|\mathcal{B}|$, and let v_i be the vertex of minimal rank in σ_i for $i = 1, 2$. Then by our construction, the morphism $\Phi(\sigma_2) \to \Phi(\sigma_1)$ is the same as $\Phi'(v_2) \to \Phi'(v_1)$.

Definition 5.3 (1-data)
Pick a vertex $v \in |\mathcal{B}|$ of rank 1, and let $\mathcal{R}_v$ be the associated residue. Let $\{v_\lambda\}_{\lambda \in \Lambda}$ be the collection of vertices in $|\mathcal{B}|$ which are $< v$, and let e_λ be the edge joining v and v_λ. Then there is a one-to-one correspondence between elements in $\mathcal{R}_v$ and the v_λ's. Each v_λ determines a point in $\Phi(v)$ by considering the image of $\Phi(e_\lambda) \to \Phi(v)$. This induces a map $f_{\mathcal{R}_v} : \mathcal{R}_v \to \Phi(v)$. The collection of all such $f_{\mathcal{R}_v}$'s with v ranging over all rank 1 vertices of $|\mathcal{B}|$ is called the 1-*data* associated with the restriction quotient $q : Y' \to |\mathcal{B}|$.

LEMMA 5.4
Pick two vertices $v, u \in |\mathcal{B}|$ of rank 1, and let $\mathcal{R}_v, \mathcal{R}_u$ be the corresponding residues. Suppose these two residues are parallel with the parallelism map given by $p : \mathcal{R}_v \to \mathcal{R}_u$. Then the following statements hold.
(1) *$\Phi(v)$ and $\Phi(u)$, considered as convex subcomplexes of Y', are parallel.*
(2) *If $p' : \Phi(v) \to \Phi(u)$ is the parallelism map, then the following diagram commutes:*

$$\begin{tikzcd}
\mathcal{R}_v \arrow[r, "p"] \arrow[d, "f_{\mathcal{R}_v}"'] & \mathcal{R}_u \arrow[d, "f_{\mathcal{R}_u}"'] \\
\Phi(v) \arrow[r, "p'"] & \Phi(u)
\end{tikzcd}$$

Proof
It follows from Lemma 3.14 that there is a finite chain of residues, starting at $\mathcal{R}_v$ and ending at $\mathcal{R}_u$, such that adjacent elements in the chain are parallel residues in a spherical residue of rank 2. Thus, we can assume without loss of generality that $\mathcal{R}_v$

and $\mathcal{R}_u$ are contained in a spherical residue $\mathcal{S}$ of type $J = \{j, j'\}$, and we assume that both $\mathcal{R}_v$ and $\mathcal{R}_u$ are j-residues.

Pick $x \in \mathcal{R}_v$. By Theorem 3.13, there is a j'-residue $\mathcal{W}$ which contains both x and $p(x)$. Let $s, w \in |\mathcal{B}|$ be the vertex corresponding to $\mathcal{S}$ and $\mathcal{W}$. Note that there is a 2-cube in $|\mathcal{B}|$ such that v, w, s are its vertices. Since Φ is 1-determined, $\mathrm{Im}(\Phi'(v) \to \Phi'(s))$ and $\mathrm{Im}(\Phi'(w) \to \Phi'(s))$ are orthogonal lines in the 2-flat $\Phi'(s)$. Moreover, the intersection of these two lines is the image of $f_{\mathcal{R}_v}(x)$ under the morphism $\Phi'(v) \to \Phi'(s)$. Similarly, the images of $\Phi'(u) \to \Phi'(s)$ and $\Phi'(w) \to \Phi'(s)$ are orthogonal lines $\Phi'(s)$, and their intersection is the image of $f_{\mathcal{R}_u}(p(x))$ under $\Phi'(v) \to \Phi'(s)$. It follows that $\mathrm{Im}(\Phi'(v) \to \Phi'(s))$ and $\mathrm{Im}(\Phi'(u) \to \Phi'(s))$ are parallel; hence, $\Phi(v)$ and $\Phi(u)$, considered as convex subcomplexes of Y', are parallel. Moreover, since the image of $f_{\mathcal{R}_v}(x)$ under $\Phi'(v) \to \Phi'(s)$ and the image of $f_{\mathcal{R}_u}(p(x))$ under $\Phi'(v) \to \Phi'(s)$ are in the line $\mathrm{Im}(\Phi'(w) \to \Phi'(s))$, the diagram in (2) commutes. □

5.3. Construction of the $\mathbb{Z}$-blowup

In the previous section, we started from a restriction quotient $q : Y' \to |\mathcal{B}|$ and produced associated 1-data (Definition 5.3), which is compatible with parallelism in the sense of Lemma 5.4. In this section, we will consider the inverse, namely, we want to construct a restriction quotient from this data.

Let $\Lambda_{\mathcal{B}}$ be the collection of parallel classes of i-residues in $\mathcal{B}$ (i could be any element in I). There is another type map T which maps a spherical J-residue $\mathcal{R}$ to $\{\lambda \in \Lambda_{\mathcal{B}} \mid \lambda \text{ contains a representative in } \mathcal{R}\}$. In other words, let $\mathcal{R} \cong \prod_{i\in I} \mathcal{R}_i$ be the product decomposition as in Theorem 3.13, where each $\mathcal{R}_i$ is an i-residue in $\mathcal{R}$ ($i \in J$). Then $T(\mathcal{R})$ is the collection of parallel classes represented by those $\mathcal{R}_i$'s. Let $\mathbb{Z}^{T(\mathcal{R})}$ be the collection of maps from $T(\mathcal{R})$ to $\mathbb{Z}$, and let $\mathbb{Z}^{\emptyset}$ be a single point.

Remark 5.5 (Remark on the notation)

The cardinality of $T(\mathcal{R})$ is equal to the rank of $\mathcal{R}$. If $\mathcal{R}$ is of rank 1, then $T(\mathcal{R}) = \lambda$ is a singleton and $\mathbb{Z}^{T(\mathcal{R})} \cong \mathbb{Z}$. However, in this case, we will still write $\mathbb{Z}^{T(\mathcal{R})}$ or $\mathbb{Z}^{\{\lambda\}}$ to emphasize that this is the copy of $\mathbb{Z}$ associated with the parallel class of $\mathcal{R}$. To simplify notation, we will also use $\mathbb{Z}^{\lambda}$ instead of $\mathbb{Z}^{\{\lambda\}}$ to denote the copy of $\mathbb{Z}$ associated with $\lambda \in \Lambda_{\mathcal{B}}$.

Our goal in this section is to construct a restriction quotient from the following data.

Definition 5.6 (Blow-up data)

A *blow-up data* is a collection of maps $\{h_{\mathcal{R}} : \mathcal{R} \to \mathbb{Z}^{T(\mathcal{R})}\}$, one for each residue of rank 1, such that if two rank 1 residues $\mathcal{R}_1$ and $\mathcal{R}_2$ are parallel with the parallelism map given by $h_{12} : \mathcal{R}_1 \to \mathcal{R}_2$, then $h_{\mathcal{R}_1} = h_{\mathcal{R}_2} \circ h_{12}$.

If $\mathcal{R}$ is a spherical residue with product decomposition given by $\mathcal{R} \cong \prod_{i \in I} \mathcal{R}_i$, then we define the map $h_{\mathcal{R}} : \mathcal{R} \to \mathbb{Z}^{T(\mathcal{R})}$ to be the product of the maps $\{h_{\mathcal{R}_i} : \mathcal{R}_i \to \mathbb{Z}\}$. It follows from the definition of $h_{\mathcal{R}}$ and the discussion after Definition 3.11 that if two spherical residues $\mathcal{R}$ and $\mathcal{R}'$ are parallel with the parallelism map given by $h : \mathcal{R} \to \mathcal{R}'$, then $h_{\mathcal{R}} = h_{\mathcal{R}'} \circ h$.

The following result is a consequence of Theorem 3.13.

LEMMA 5.7
Let $\mathcal{R}$ be a spherical J-residue. Let $g : \mathcal{R} \cong \prod_{i=1}^{n} \mathcal{R}_i$ be the product decomposition induced by $J = \bigsqcup_{i=1}^{n} J_i$ (see Theorem 3.13). Then $h_{\mathcal{R}} = (\prod_{i=1}^{n} h_{\mathcal{R}_i}) \circ g$.

To simplify notation, we will write $h_{\mathcal{T}} = \prod_{i=1}^{n} h_{\mathcal{T}_i}$ instead of $h_{\mathcal{T}} = (\prod_{i=1}^{n} h_{\mathcal{T}_i}) \circ g$. Let J and $\mathcal{R} = \prod_{i \in J} \mathcal{R}_i$ be as before. A J'-residue $\mathcal{R}' \subset \mathcal{R}$ can be expressed as $(\prod_{i \in J'} \mathcal{R}_i) \times (\prod_{i \in J \setminus J'} \{c_i\})$, where c_i is a chamber in $\mathcal{R}_i$. We define an inclusion $h_{\mathcal{R}'\mathcal{R}} : \mathbb{Z}^{T(\mathcal{R}')} \to \mathbb{Z}^{T(\mathcal{R})}$ by $h_{\mathcal{R}'\mathcal{R}}(a) = \{a\} \times \prod_{i \in J \setminus J'} \{h_{\mathcal{R}_i}(c_i)\}$. Since $h_{\mathcal{R}} = h_{\mathcal{R}'} \times (\prod_{i \in J \setminus J'} h_{\mathcal{R}_i})$, $h_{\mathcal{R}'\mathcal{R}}$ fits into the following commutative diagram:

$$\begin{array}{ccc} \mathcal{R}' & \longrightarrow & \mathcal{R} \\ {\scriptstyle h_{\mathcal{R}'}}\big\downarrow & & {\scriptstyle h_{\mathcal{R}}}\big\downarrow \\ \mathbb{Z}^{T(\mathcal{R}')} & \xrightarrow{h_{\mathcal{R}'\mathcal{R}}} & \mathbb{Z}^{T(\mathcal{R})} \end{array}$$

Suppose $\mathcal{R}''$ is a J''-residue such that $\mathcal{R}'' \subset \mathcal{R}' \subset \mathcal{R}$. Since $h_{\mathcal{R}} = h_{\mathcal{R}'} \times (\prod_{i \in J \setminus J'} h_{\mathcal{R}_i}) = h_{\mathcal{R}''} \times (\prod_{i \in J' \setminus J''} h_{\mathcal{R}_i}) \times (\prod_{i \in J \setminus J'} h_{\mathcal{R}_i})$, we have

$$h_{\mathcal{R}''\mathcal{R}} = h_{\mathcal{R}'\mathcal{R}} \circ h_{\mathcal{R}''\mathcal{R}'}. \tag{5.8}$$

Now we define a contravariant functor $\Psi : \mathrm{Face}(|\mathcal{B}|) \to \mathrm{CCC}$ as follows. Let f be a face of $|\mathcal{B}|$, and let $v_f \in f$ be the unique vertex which has minimal rank among the vertices of f. Let $\mathcal{R}_f \subset \mathcal{B}$ be the residue associated with v_f. We define $\Psi(f) = \mathbb{R}^{T(\mathcal{R}_f)}$ ($\mathbb{R}^{\emptyset}$ is a single point), where $\mathbb{R}^{T(\mathcal{R}_f)}$ is endowed with the standard cubical structure, and we identify $\mathbb{Z}^{T(\mathcal{R}_f)}$ with the 0-skeleton of $\mathbb{R}^{T(\mathcal{R}_f)}$.

An inclusion of faces $f \to f'$ induces an inclusion $\mathcal{R}_{f'} \to \mathcal{R}_f$. We define the morphism $\Psi(f') \to \Psi(f)$ to be the embedding induced by $h_{\mathcal{R}_{f'}\mathcal{R}_f} : \mathbb{Z}^{T(\mathcal{R}_{f'})} \to \mathbb{Z}^{T(\mathcal{R}_f)}$.

LEMMA 5.9
We have that Ψ is a contravariant functor.

Proof
It is easy to check that passing from an inclusion of faces $f \to f'$ to $\mathcal{R}_{f'} \to \mathcal{R}_f$ is a functor. And it follows from (5.8) that passing from $\mathcal{R}_{f'} \to \mathcal{R}_f$ to $h_{\mathcal{R}_{f'}\mathcal{R}_f} : \mathbb{Z}^{T(\mathcal{R}_{f'})} \to \mathbb{Z}^{T(\mathcal{R}_f)}$ is a functor. $\square$

LEMMA 5.10
We have that Ψ is 1-determined.

Proof
Let $\sigma \subset |\mathcal{B}|$ be a face, and pick a vertex $v \in \sigma$. Let $\{v_i\}_{i=1}^k$ be the vertices in σ that are adjacent to v along edges $\{e_i\}_{i=1}^k$. Let $\sigma_{<v}$ be the subcube of σ that is spanned by the e_i's such that $v_i < v$. (If v is a minimal vertex in σ, then we define $\sigma_{<v} = \{v\}$.) Note that $v \in \sigma_{<v}$, since v is inside each e_i. We define $\sigma_{>v}$ similarly. Then $\sigma = \sigma_{<v} \times \sigma_{>v}$. Moreover, v is the maximal vertex in $\sigma_{<v}$ and the minimum vertex in $\sigma_{>v}$. Note that $\Psi(e_i) \to \Psi(v)$ is an isometry if $v_i > v$. Thus, it suffices to consider the case where v is the maximal vertex of σ.

Let v_m be the minimal vertex of σ. Note that $\mathrm{Im}(\Psi(\sigma) \to \Psi(v)) \subset \bigcap_{i=1}^k \mathrm{Im}(\Psi(e) \to \Psi(v))$ is a cubical convex embedding of Euclidean subspaces; it suffices to show they have the same dimension. Let $\mathcal{R}(v)$ be the spherical residue corresponding to the vertex v. Note that $T(\mathcal{R}(v_m)) = \bigcap_{i=1}^k T(\mathcal{R}(v_i))$. ($T$ is the type map defined at the beginning of Section 5.3.) Thus, the dimension of $\bigcap_{i=1}^k \mathrm{Im}(\Psi(e) \to \Psi(v))$ equals the cardinality of $T(\mathcal{R}(v_m))$, which is the dimension of $\mathrm{Im}(\Psi(\sigma) \to \Psi(v))$. $\square$

Ψ is called the *fiber functor associated with the blow-up data* $\{h_{\mathcal{R}}\}$, and the restriction quotient $q : Y \to |\mathcal{B}|$ which arises from the fiber functor Ψ (see Theorem 4.15) is called the *restriction quotient associated with the blow-up data* $\{h_{\mathcal{R}}\}$. It is clear from the construction that the 1-data of q (see Definition 5.3) is the blow-up data $\{h_{\mathcal{R}}\}$. (We naturally identify the $\mathbb{Z}^{T(\mathcal{R})}$'s in the blow-up data with the 0-skeleton of the q-fibers of rank 1 vertices in $|\mathcal{B}|$.) We summarize the above discussion in the following theorem.

THEOREM 5.11
Given the blow-up data $\{h_{\mathcal{R}}\}$ as in Definition 5.6, there exists a restriction quotient $q : Y \to |\mathcal{B}|$ whose 1-data is the blow-up data we start with.

Remark 5.12
Here we blow up the building $\mathcal{B}$ with respect to a collection of $\mathbb{Z}$'s, since we want to apply the construction for RAAGs. However, in other cases (e.g., for more general

graph products of groups), it may be natural to blow it up with respect to other objects. Here is a variation. To each parallel class of rank 1 residues $\lambda \in \Lambda_{\mathcal{B}}$, we associate a CAT(0) cube complex Z_λ. For each rank 1 residue $\mathcal{R}$ in the class λ, we define a map $h_{\mathcal{R}}$ which assigns each element of $\mathcal{R}$ a convex subcomplex of Z_λ. We require these $\{h_{\mathcal{R}}\}$ to be compatible with parallelism between rank 1 residues. Given this set of blow-up data, we can repeat the previous construction to obtain a restriction quotient over $|\mathcal{B}|$, whose vertex fibers are products of some subcomplexes of Z_λ's.

Now we show that the construction in this section is indeed a converse to Section 5.2 in the following sense. Let $q : Y' \to |\mathcal{B}|$ be a restriction quotient as in Section 5.2, and let Φ and Φ' be the functors introduced there. For each vertex $v \in |\mathcal{B}|$ of rank 1 and its associated residue $\mathcal{R}_v$, we pick an isometric embedding $\eta_v : \mathbb{Z}^{T(\mathcal{R}_v)} \to \Phi(v)$ such that its image is the vertex set of $\Phi(v)$. We also require that these η_v's respect parallelism. More precisely, let $u \in |\mathcal{B}|$ be a vertex of rank 1 such that $\Phi(v)$ and $\Phi(u)$ (understood as subcomplexes of Y') are parallel with the parallelism map given by $p : \Phi(v) \to \Phi(u)$. Then $p \circ \eta_v = \eta_u$. (Note that $T(\mathcal{R}_v) = T(\mathcal{R}_u)$ by Lemma 5.4.)

Let Ψ be the functor constructed in this section from the blow-up data $\{h_{\mathcal{R}_v} = \eta_v^{-1} \circ f_{\mathcal{R}_v} : \mathcal{R}_v \to \mathbb{Z}^{T(\mathcal{R}_v)}\}_{v \in |\mathcal{B}|}$, where v ranges over all vertices of rank 1 in $|\mathcal{B}|$, $\mathcal{R}_v$ is the residue associated with v, and $f_{\mathcal{R}_v}$ is the map in Definition 5.3. Pick a face $\sigma \in |\mathcal{B}|$, and let $u \in \sigma$ be the vertex of minimal rank. Let $\mathcal{R}_u$ be the associated J-residue with its product decomposition given by $\mathcal{R}_u = \prod_{j \in J} \mathcal{R}_{v_j}$. (The v_j's are rank 1 vertices $\leq u$.) Let $\xi_\sigma : \Psi(\sigma) \to \Phi(\sigma)$ be the isometry induced by

$$\prod_{j \in J} \eta_{v_j} : \mathbb{Z}^{T(\mathcal{R}_u)} \to \prod_{j \in J} \Phi(v_j) \cong \Phi(u) \cong \Phi(\sigma).$$

(To see the product decomposition $\prod_{j \in J} \Phi(v_j) \cong \Phi(u)$, recall that Φ is 1-determined; thus, $\{\mathrm{Im}(\Phi'(v_j) \to \Phi'(u))\}_{j \in J}$ are mutually orthogonal lines in $\Phi'(u)$, and then the required product decomposition follows from this together with Lemma 5.4.) It is not hard to check that the collection of ξ_σ's gives a natural isomorphism from Ψ to Φ.

COROLLARY 5.13
The maps $\{\xi_\sigma\}_{\sigma \in \mathrm{Face}(|\mathcal{B}|)}$ induce a natural isomorphism between Φ and Ψ. Thus for any restriction quotient $q : Y' \to |\mathcal{B}|$ which satisfies the conclusion of Lemma 5.1, if q' is the restriction quotient whose blow-up data is the 1*-data of q, then q' is equivalent to q up to a natural isomorphism between their fiber functors.*

COROLLARY 5.14
Let $q : Y \to |\mathcal{B}|$ be a restriction quotient which satisfies the conclusion of Lemma 5.1. Let $\mathcal{B} \cong \mathcal{B}_1 \times \mathcal{B}_2$ be a product decomposition of the building $\mathcal{B}$ induced by the

join decomposition $\Gamma = \Gamma_1 \circ \Gamma_2$ of the defining graph of the associated right-angled Coxeter group. Then there are two restriction quotients $q_1 : Y_1 \to |\mathcal{B}_1|$ and $q_2 : Y_2 \to |\mathcal{B}_2|$ such that $Y = Y_1 \times Y_2$ and $q = q_1 \times q_2$. Moreover, q_1 and q_2 also satisfy the conclusion of Lemma 5.1.

Proof
By Corollary 5.13, we can assume that q is the restriction quotient associated with a set of blow-up data $\{h_{\mathcal{R}}\}$. For every $\mathcal{B}_1$-slice in $\mathcal{B}$, we can restrict $\{h_{\mathcal{R}}\}$ to $\mathcal{B}_1$ to obtain a blow-up data for $\mathcal{B}_1$. This does not depend on our choice of the $\mathcal{B}_1$-slice, since the blow-up data respects parallelism. We obtain a blow-up data for $\mathcal{B}_2$ in a similar way. It follows from the above construction that the fiber functor associated with $\{h_{\mathcal{R}}\}$ is the product of the fiber functors associated with the blow-up data on $\mathcal{B}_1$ and $\mathcal{B}_2$. Thus, this corollary is a consequence of Lemma 4.16. $\square$

5.4. More properties of the blow-up buildings
In this section, we look at the restriction quotient $q : Y \to |\mathcal{B}|$ associated with the blow-up data $\{h_{\mathcal{R}}\}$ as in Definition 5.6 (or, equivalently, a restriction quotient $q : Y \to |\mathcal{B}|$ which satisfies the conclusion of Lemma 5.1) in more detail, and we record several basic properties of Y. A hurried reader can go through Definition 5.15, then proceed directly to Section 5.5, and come back to this part later.

Definition 5.15
A vertex $y \in Y$ is of *rank* k if $q(y)$ is vertex of rank k. Thus, q induces a bijection between rank 0 vertices in Y and rank 0 vertices in $|\mathcal{B}|$. Since rank 0 vertices in $|\mathcal{B}|$ can be identified with chambers in $\mathcal{B}$, q^{-1} induces a well-defined map $q^{-1} : \mathcal{B} \to Y$ from the set of chambers of $\mathcal{B}$ (or rank 0 vertices of $|\mathcal{B}|$) to rank 0 vertices in Y.

LEMMA 5.16
For any residue $\mathcal{R} \subset \mathcal{B}$, we view $\mathcal{R}$ as a building and restrict the blow-up data over $\mathcal{B}$ to the blow-up data over $\mathcal{R}$. Let $q_{\mathcal{R}} : Y_{\mathcal{R}} \to |\mathcal{R}|$ be the associated restriction quotient. Then there exists an isometric embedding $i : Y_{\mathcal{R}} \to Y$ which fits into the following commutative diagram:

$$\begin{array}{ccc} Y_{\mathcal{R}} & \xrightarrow{i} & Y \\ {\scriptstyle q_{\mathcal{R}}}\downarrow & & {\scriptstyle q}\downarrow \\ |\mathcal{R}| & \xrightarrow{i'} & |\mathcal{B}| \end{array}$$

Moreover, $i(Y_{\mathcal{R}}) = q^{-1}(i'(|\mathcal{R}|))$.

The lemma is a direct consequence of the construction in Section 5.3.

Pick a vertex $v \in |\mathcal{B}|$. The *downward complex of* v is the smallest convex subcomplex of $|\mathcal{B}|$ which contains all vertices which are $\le v$. If $\mathcal{R}_v$ is the residue associated with v, then the downward complex is the image of the embedding $|\mathcal{R}_v| \hookrightarrow |\mathcal{B}|$. The next result follows from Lemma 5.16 and Corollary 5.14.

LEMMA 5.17
Let D_v be the downward complex of a vertex $v \in \mathcal{B}$, and let $\mathcal{R}_v = \prod_{i=1}^{k} \mathcal{R}_i$ be the product decomposition of the residue associated with v. Then $q^{-1}(D_v)$ is isomorphic to the product of the mapping cylinders of $f_i : \mathcal{R}_i \to \mathbb{R}^{T(\mathcal{R}_i)}$ for $1 \le i \le k$, where each f_i is the composition $\mathcal{R}_i \xrightarrow{h_{\mathcal{R}_i}} \mathbb{Z}^{T(\mathcal{R}_i)} \to \mathbb{R}^{T(\mathcal{R}_i)}$.

LEMMA 5.18
(1) *If $h_{\mathcal{R}}^{-1}(x)$ is finite for any rank 1 residue $\mathcal{R}$ and $x \in \mathbb{Z}^{T(\mathcal{R})}$, then Y is locally finite. If there is a uniform upper bound for the cardinality of $h_{\mathcal{R}}^{-1}(x)$, then Y is uniformly locally finite.*
(2) *If there exists $D > 0$ such that the image of each $h_{\mathcal{R}}$ is D-dense in $\mathbb{Z}^{T(\mathcal{R})}$, then there exists D' which depends on D and the dimension of $|\mathcal{B}|$ such that the collection of inverse images of rank 0 vertices in $|\mathcal{B}|$ is D'-dense in Y.*

Proof
We prove (1) first. Pick a vertex $y \in Y$. Let $v = q(y)$. It suffices to show that the set of edges in $|\mathcal{B}|$ which contains v and which can be lifted to an edge in Y that contains y is finite. Since there are only finitely many vertices in $|\mathcal{B}|$ which are $\ge v$, it suffices to consider the edges of the form $\overline{v_\lambda v}$ with $v_\lambda < v$. It follows from our assumption and Lemma 5.17 that there are only finitely many such edges which have the required lift. The proof of uniform local finiteness is similar.

To see (2), notice that $\bigcup_{v \in |\mathcal{B}|} \Psi(v)$ is 1-dense in Y, where v ranges over all vertices of $|\mathcal{B}|$. It follows from Lemma 5.17 that every point in $\Psi(v)$ can be approximated by the inverse image of some rank 0 vertex up to distance D'. □

Next we discuss the relation between Y and the exploded Salvetti complex $S_e = S_e(\Gamma)$ introduced in Section 5.1. Let Ψ be the fiber functor associated with $q : Y \to |\mathcal{B}|$.

First we label each vertex $v \in Y$ by a clique in Γ as follows. Recall that $q(v)$ is associated with a J-residue $\mathcal{R} \subset \mathcal{B}$, where J is the vertex set of a clique in Γ. Thus, we label v by this clique. We also label each vertex of S_e by a clique. Any vertex $v \in S_e$ is contained in a unique standard torus. Recall that a standard torus arises from a clique in Γ; thus, we label v by this clique. Note that some vertices of Y and S_e are labeled by the empty set. There is a unique label-preserving map $p : Y^{(0)} \to S_e^{(0)}(\Gamma)$.

An edge in Y or S_e is *horizontal* if the labels on its two endpoints are different; otherwise, this edge is *vertical*. When $Y = X_e$, this definition coincides with the one in Section 5.1. Moreover, horizontal (or vertical) edges in X_e are lifts of horizontal (or vertical) edges in S_e.

Horizontal edges in Y are exactly those ones whose dual hyperplanes are mapped by q to hyperplanes in $|\mathcal{B}|$, and the q-image of any vertical edge is a point. Now we label each vertical edge of Y by vertices in Γ as follows. Pick a vertical edge $e \subset Y$, and let $v = q(e)$. Let $\mathcal{R} = \prod_{i=1}^{k} \mathcal{R}_i$ be the product decomposition of the residue associated with v. There is a corresponding product decomposition $\Psi(v) = \prod_{i=1}^{k} \ell_i$, where ℓ_i is a line which is parallel to $\Psi(v_i)$ and $v_i \in |\mathcal{B}|$ is the vertex associated with $\mathcal{R}_i$, and we view $\Psi(v_i)$ and $\Psi(v)$ as subcomplexes of Y. If e is in the ℓ_i-direction, then we label e by the type of $\mathcal{R}_i$, which is a vertex in Γ. A case study implies that if two vertical edges are the opposite sides of a 2-cube, then they have the same label. Hence, all parallel vertical edges have the same label. Now we label vertical edges in S_e. Recall that the map $S_e \to S(\Gamma)$ induces a one-to-one correspondence between vertical edges in S_e and edges in $S(\Gamma)$, and edges in $S(\Gamma)$ are labeled by vertices of Γ. This induces a labeling of vertical edges in S_e.

We pick an orientation for each vertical edge in S_e and orient every vertical edge in Y in the following way. A *vertical line* is a geodesic line made of vertical edges. It is easy to see that every vertical edge is contained in a vertical line. For two vertical lines ℓ_1 and ℓ_2, if there exist edges $e_i \in \ell_i$ for $i = 1, 2$ such that they are parallel, then ℓ_1 and ℓ_2 are parallel. To see this, it suffices to consider the case where e_1 and e_2 are the opposite sides of a 2-cube, and this follows from a similar case study as before. Now we pick an orientation for each parallel class of vertical lines, and this induces a well-defined orientation on each vertical edge of Y; moreover, this orientation respects parallelism of edges.

There is a unique way to extend $p : Y^{(0)} \to S_e^{(0)}(\Gamma)$ to $p : Y^{(1)} \to S_e^{(1)}(\Gamma)$ such that p preserves the orientation and labeling of vertical edges. One can further extend p to higher-dimensional cells as follows. A cube $\sigma \subset Y$ is of *type* (m, n) if σ is the product of m vertical edges and n horizontal edges. We extend p according to the type.

(1) If σ is of type $(m, 0)$, then we can define p on σ, since the orientation of vertical edges in Y respects parallelism, and p preserves the labeling and orientation of vertical edges. In this case, $p(\sigma)$ is an m-dimensional standard torus.

(2) If σ is of type $(0, n)$, then we can define p on σ, since p preserves the labeling of vertices. In this case, $p(\sigma) \cong [0, 1]^n$.

(3) If σ is of type (m, n), then we can define p on σ for similar reasons as before. In this case, $p(\sigma) \cong \mathbb{T}^m \times [0,1]^n$.

If $y \in Y$ is a vertex, then p induces a simplicial map between the vertex links $p_y : Lk(y, Y) \to Lk(p(y), S_e)$. The above case study implies that p_y is a *combinatorial* map; that is, p_y maps each simplex isomorphically onto its image.

THEOREM 5.19
If each map $h_{\mathfrak{R}}$ in the blow-up data is a bijection, then Y is isomorphic to $X_e = X_e(\Gamma)$, which is the universal cover of the exploded Salvetti complex $S_e = S_e(\Gamma)$.

Proof
We prove the theorem by showing that $p : Y \to S_e$ is a covering map. It suffices to show, for each vertex $y \in Y$, that the above map p_y is an isomorphism. Suppose that y is labeled by a clique $\Delta \subset \Gamma$. We look at edges which contain y and which fall into three classes:
(1) vertical edges;
(2) horizontal edges whose other endpoints are labeled by cliques in Δ; and
(3) horizontal edges whose other endpoints are labeled by cliques that contain Δ.
Note that there is a one-to-one correspondence between edges in (3) and cliques which contain Δ and have exactly one vertex not in Δ. For any clique $\Delta' \subset \Delta$ which contains all but one vertex of Δ, there exists a unique edge in (2) such that its other endpoint is labeled by Δ', since if this edge does not exist, then some $h_{\mathfrak{R}}$ will not be surjective; if there exists more than one such edge, then some $h_{\mathfrak{R}}$ will not be injective. Thus, there is a one-to-one correspondence between horizontal edges which contain y and horizontal edges which contain $p(y)$. Hence, p_y induces a bijection between the 0-skeletons. Moreover, edges in (3) are orthogonal to edges in (1) and (2), so a case study implies that if two edges at $p(y)$ form the corner of a 2-cube, then their lifts at y (if any exist) also form the corner of a 2-cube. It follows that p_y induces an isomorphism between the 1-skeletons. Since both $Lk(y, Y)$ and $Lk(p(y), S_e)$ are flag complexes, p_y is an isomorphism. □

Remark 5.20
If each map $h_{\mathfrak{R}}$ is injective (or surjective), then p is locally injective (or locally surjective).

COROLLARY 5.21
Let $\mathcal{B}_1 = \mathcal{B}_1(\Gamma)$ and $\mathcal{B}_2 = \mathcal{B}_2(\Gamma)$ be two right-angled Γ-buildings with countably infinite rank 1 residues. Then they are isomorphic as buildings.

Proof
We pick a blowup for $\mathcal{B}_1$ such that each map in the blow-up data is a bijection. Let $Y \to |\mathcal{B}_1|$ be the associated restriction quotient, and let $p : Y \to S_e$ be the covering map as in Theorem 5.19. Note that p sends vertical edges to vertical edges and horizontal edges to horizontal edges, and p preserves the labeling of vertices and edges. So does the lift $\tilde{p} : Y \to X_e$ of p. Lemma 4.3 implies $\tilde{p}$ descends to a cubical isomorphism $|\mathcal{B}_1| \to |\mathcal{B}|$, where $|\mathcal{B}|$ is the building associated with $G(\Gamma)$. Since $\tilde{p}$ is label-preserving, this cubical isomorphism induces a building isomorphism $\mathcal{B}_1 \to \mathcal{B}$. Similarly, we can obtain a building isomorphism $\mathcal{B}_2 \to \mathcal{B}$. Hence, the corollary follows. □

THEOREM 5.22
Suppose that Γ does not admit a join decomposition $\Gamma = \Gamma_1 \circ \Gamma_2$, where Γ_1 is a discrete graph with more than one vertex. If $\mathcal{B}$ is a Γ-building and $q : Y \to |\mathcal{B}|$ is a restriction quotient with blow-up data $\{h_{\mathcal{R}}\}$, then any automorphism $\alpha : Y \to Y$ descends to an automorphism $\alpha' : |\mathcal{B}| \to |\mathcal{B}|$.

Proof
By Lemma 4.3, it suffices to show that α preserves the rank (see Definition 5.15) of vertices of Y. Let $F(\Gamma)$ be the flag complex of Γ. Here we change the label of each vertex in Y from some clique in Γ to the associated simplex in $F(\Gamma)$. Suppose that $y \in Y$ is a vertex of rank k labeled by Δ. Then Lemma 5.17 and the proof of Theorem 5.19 imply $Lk(y, Y) \cong K_1 * K_2 * \cdots * K_k * Lk(\Delta, F(\Gamma))$, where each K_i is discrete with cardinality at least 2, and $Lk(\Delta, F(\Gamma))$ is understood to be $F(\Gamma)$ when $\Delta = \emptyset$. Note that $\{K_i\}_{i=1}^k$ comes from vertices adjacent to y of rank at most k, and $Lk(\Delta, F(\Gamma))$ comes from vertices adjacent to y of rank greater than k. Thus, α preserves the collection of rank 0 vertices.

Now we assume that α preserves the collection of rank i vertices for $i \le k-1$. A rank k vertex in Y is of *type I* if it is adjacent to a vertex of rank $k-1$; otherwise, it is a vertex of *type II*. It is clear that α preserves the collection of rank k vertices of type I. Before we deal with type II vertices, we need the following claim. Suppose that $w \in Y$ is a vertex of rank k such that $\alpha(w)$ is also of rank k. If there exist k vertices $\{z_i\}_{i=1}^k$ adjacent to w such that

(1) rank$(z_i) \le k$ and rank$(\alpha(z_i)) \le k$;
(2) the edges $\{\overline{z_i w}\}_{i=1}^k$ are mutually orthogonal,

then rank$(\alpha(z)) \le k$ for any z adjacent to w with rank$(z) \le k$.

Now we verify the claim. Let $w' = \alpha(w)$. Suppose that w and w' are labeled by Δ and Δ'. Then α induces an isomorphism between the links of w and w' in Y:

$$\alpha_* : K_1 * \cdots * K_k * Lk\big(\Delta, F(\Gamma)\big) \to K_1' * \cdots * K_k' * Lk\big(\Delta', F(\Gamma)\big).$$

Each edge $\overline{z_i w}$ gives rise to a vertex in K_i, and each edge $\overline{\alpha(z_i)w'}$ gives rise to a vertex in K'_i. Thus, $\alpha_*(K_1 * \cdots * K_k) = K'_1 * \cdots * K'_k$. Since the edge $\overline{zw}$ gives rise to a vertex in $K_1 * \cdots * K_k$, the edge $\overline{\alpha(z)w'}$ gives rise to a vertex in $K'_1 * \cdots * K'_k$. Then $\alpha(z)$ is of rank at most k.

Let $y \in Y$ be a rank k vertex of type II. Then there exists an edge path ω from y to a type I vertex y_1 such that every vertex in ω is of rank k. Let $\{y_i\}_{i=1}^m$ be consecutive vertices in ω such that $y_m = y$. Note that there are k vertices of rank $k-1$ adjacent to y_1. By the induction assumption, they are sent to vertices of rank $k-1$ by α. Moreover, $\operatorname{rank}(\alpha(y_1)) = k$, since y_1 is of type I. Thus, the assumption of the claim is satisfied for y_1. Then $\operatorname{rank}(\alpha(y_2)) \le k$; hence, $\operatorname{rank}(\alpha(y_2)) = k$ by the induction assumption. Next we show that y_2 satisfies the assumption of the claim. Let $\{z_i\}_{i=1}^k$ be vertices of rank k such that they are adjacent to y_1 and $\{\overline{z_i y_1}\}_{i=1}^k$ are mutually orthogonal. We also assume that $y_2 = z_1$. Then $\operatorname{rank}(\alpha(z_i)) = k$ for all i. Hence, all the $\alpha(\overline{z_i y_1})$'s are vertical edges. For $i \ge 2$, let z'_i be the vertex adjacent to y_2 such that $\overline{z'_i y_2}$ and $\overline{z_i y_1}$ are parallel. Then $\alpha(\overline{z'_i y_2})$ is a vertical edge for $i \ge 2$. Thus, $\operatorname{rank}(\alpha(z'_i)) = k$ and the assumption of the claim is satisfied for y_2. We can repeat this argument finitely many times to deduce that $\operatorname{rank}(\alpha(y)) = k$. $\square$

Remark 5.23

If the assumption on Γ in Theorem 5.22 is not satisfied, then there exist a blowup $Y \to |\mathcal{B}|$ and an automorphism of Y that does not descend to an automorphism of $|\mathcal{B}|$. By Corollary 5.14, it suffices to construct an example in the case when Γ is a discrete graph with n vertices with $n \ge 2$. If $n \ge 3$, then we define each $h_{\mathcal{R}}$ to be a surjective map such that the inverse image of each point has $n-2$ points. Then Y is a tree with valence $= n$. If $n = 2$, then we define $h_{\mathcal{R}}$ to be an injective map whose image is the set of even integers. Then Y is isomorphic to the first subdivision of a tree of valence 3. In both cases, it is not hard to find an automorphism of Y which maps some vertex of rank 0 to a vertex of rank 1.

5.5. *Morphisms between blow-up data*

Let $\mathcal{B}$ and $\mathcal{B}'$ be two buildings modeled on the same right-angled Coxeter group $W(\Gamma)$. A cubical isomorphism $\eta : |\mathcal{B}| \to |\mathcal{B}'|$ is *rank-preserving* if, for each vertex $v \in |\mathcal{B}|$, v and $\eta(v)$ have the same rank. (In general, cubical isomorphisms may not preserve the rank of vertices (see Remark 5.23).) Note that such an η induces a bijection $\eta' : \mathcal{B} \to \mathcal{B}'$ which preserves the spherical residues. Conversely, every bijection $\mathcal{B} \to \mathcal{B}'$ which preserves the spherical residues induces a rank-preserving isomorphism $|\mathcal{B}| \to |\mathcal{B}'|$. Note that η' maps parallel residues of rank 1 to parallel residues of rank 1; thus, η' induces a bijection $\bar{\eta} : \Lambda_{\mathcal{B}} \to \Lambda_{\mathcal{B}'}$, where $\Lambda_{\mathcal{B}}$ and $\Lambda_{\mathcal{B}'}$ denote the collection of parallel classes of residues of rank 1 in $\mathcal{B}$ and $\mathcal{B}'$, respectively (see Section 5.3). See Remark 5.5 for our use of notation in the following discussion.

Definition 5.24 (η-Isomorphism)
Suppose that the blow-up data (see Definition 5.6) of $|\mathcal{B}|$ and $|\mathcal{B}'|$ are given by $\{h_{\mathcal{R}}\}$ and $\{h'_{\mathcal{R}}\}$, respectively. An *η-isomorphism* between the blow-up data is defined to be a collection of isometries $\{f_\lambda : \mathbb{Z}^\lambda \to \mathbb{Z}^{\bar{\eta}(\lambda)}\}_{\lambda \in \Lambda_{\mathcal{B}}}$ such that the following diagram commutes for every rank 1 residue $\mathcal{R} \subset \mathcal{B}$:

$$\begin{array}{ccc} \mathcal{R} & \xrightarrow{h_{\mathcal{R}}} & \mathbb{Z}^{T(\mathcal{R})} \\ \eta' \big\downarrow & & f_{T(\mathcal{R})} \big\downarrow \\ \eta'(\mathcal{R}) & \xrightarrow{h'_{\eta'(\mathcal{R})}} & \mathbb{Z}^{\bar{\eta}(T(\mathcal{R}))} \end{array}$$

Here T is the type map defined at the beginning of Section 5.3. The map $h_{\mathcal{R}}$ is *nondegenerate* if its image contains more than one point. In this case, if $f_{T(\mathcal{R})}$ exists, then it is unique. If $h_{\mathcal{R}}$ is degenerate, then we have two choices for $f_{T(\mathcal{R})}$.

Let $\eta_1 : |\mathcal{B}_1| \to |\mathcal{B}_2|$, $\eta_2 : |\mathcal{B}_2| \to |\mathcal{B}_3|$, and $\eta : |\mathcal{B}_1| \to |\mathcal{B}_3|$ be rank-preserving isomorphisms such that $\eta = \eta_2 \circ \eta_1$. We fix a blow-up data for each $\mathcal{B}_i$. Let $\{f_\lambda : \mathbb{Z}^\lambda \to \mathbb{Z}^{\bar{\eta}_1(\lambda)}\}_{\lambda \in \Lambda_{\mathcal{B}_1}}$ and $\{g_\lambda : \mathbb{Z}^\lambda \to \mathbb{Z}^{\bar{\eta}_2(\lambda)}\}_{\lambda \in \Lambda_{\mathcal{B}_2}}$ be some η_1-isomorphism and η_2-isomorphism, respectively, between the corresponding blow-up data. We define the *composition* of them to be $\{g_{\bar{\eta}_1(\lambda)} \circ f_\lambda\}_{\lambda \in \Lambda}$, which turns out to be an η-isomorphism.

Let Ψ and Ψ' be the fiber functor associated with the blow-up data $\{h_{\mathcal{R}}\}$ and $\{h'_{\mathcal{R}}\}$, respectively, and let $Y \to |\mathcal{B}|$ and $Y' \to |\mathcal{B}'|$ be the associated restriction quotient.

LEMMA 5.25
Every η-isomorphism induces a natural isomorphism from Ψ to Ψ'; hence by Section 4.3, it induces an isomorphism $Y \to Y'$, which is a lift of $\eta : |\mathcal{B}| \to |\mathcal{B}'|$. Moreover, the composition of η-isomorphisms gives rise to the composition of natural transformations of the associated fiber functors.

Proof
For every spherical residue $\mathcal{R} \subset \mathcal{B}$, η' respects the product decomposition of $\mathcal{R}$. Thus, the following diagram commutes:

$$\begin{array}{ccc} \mathcal{R} & \xrightarrow{h_{\mathcal{R}}} & \mathbb{Z}^{T(\mathcal{R})} \\ \eta' \big\downarrow & & \prod_{\lambda \in T(\mathcal{R})} f_\lambda \big\downarrow \\ \eta'(\mathcal{R}) & \xrightarrow{h'_{\eta'(\mathcal{R})}} & \mathbb{Z}^{\bar{\eta}(T(\mathcal{R}))} \end{array}$$

Here $\prod_{\lambda \in T(\mathcal{R})} f_\lambda$ induces an isometry $\mathbb{R}^{T(\mathcal{R})} \to \mathbb{R}^{\bar{\eta}(T(\mathcal{R}))}$. This gives rise to a collection of isometries between objects of Ψ and Ψ'. It follows from the construction in Section 5.3 that these isometries give the required natural isomorphism between Ψ and Ψ'. The second assertion in the lemma is straightforward. $\square$

Remark 5.26
If we weaken the assumption of Definition 5.24 by assuming that each f_λ is a bijection, then we can obtain a bijection between the vertex sets of Y and Y'. This bijection preserves the fibers; however, we may not be able to extend it to a cubical map.

Definition 5.27 (η-Quasimorphism)
We follow the notation in Definition 5.24. An (η, L, A)-*quasimorphism* between the blow-up data $\{h_{\mathcal{R}}\}$ and $\{h'_{\mathcal{R}}\}$ is a collection of (L, A)-quasi-isometries $\{f_\lambda : \mathbb{Z}^\lambda \to \mathbb{Z}^{\bar{\eta}(\lambda)}\}_{\lambda \in \Lambda_{\mathcal{B}}}$ such that the diagram in Definition 5.24 commutes up to error A.

LEMMA 5.28
Each (η, L, A)-quasimorphism between $\{h_{\mathcal{R}}\}$ and $\{h'_{\mathcal{R}}\}$ induces an (L', A')-quasi-isometry $Y \to Y'$ with L', A' depending on L, A, and the dimension of $|\mathcal{B}|$.

Proof
By Lemma 4.19, it suffices to produce an (L', A')-quasinatural isomorphism from Ψ to Ψ'. This can be done by considering maps of the form $\prod_{\lambda \in T(\mathcal{R})} f_\lambda$ as in Lemma 5.25. $\square$

Remark 5.29 (A nice representative)
Let Y_0 be the collection of rank 0 vertices in Y (see Definition 5.15). We define Y'_0 similarly. If the assumption in Lemma 5.18(2) is satisfied, then Y_0 and Y'_0 are D-dense in Y and Y', respectively. In this case, the quasi-isometry $Y \to Y'$ in Lemma 5.28 can be represented by $\phi : Y_0 \to Y'_0$, where ϕ is the bijection induced by $\eta : |\mathcal{B}| \to |\mathcal{B}'|$. (Recall that we can identify Y_0 and Y'_0 with rank 0 vertices in $|\mathcal{B}|$ and $|\mathcal{B}'|$, respectively (see Definition 5.15).) The fact that ϕ is a quasi-isometry follows from the construction in the proof of Lemma 4.19.

COROLLARY 5.30
If there exists a constant $D > 0$ such that each map $h_{\mathcal{R}}$ in the blow-up data satisfies:
(1) *for any $x \in \mathbb{Z}^{T(\mathcal{R})}$, $|h_{\mathcal{R}}^{-1}(x)| \leq D$;*
(2) *the image of $h_{\mathcal{R}}$ is D-dense in $\mathbb{Z}^{T(\mathcal{R})}$;*
then Y is quasi-isometric to $G(\Gamma)$.

Proof
By the assumptions, there exists another set of blow-up data $\{h'_{\mathcal{R}}\}$ such that each $h'_{\mathcal{R}}$ is a bijection and there exists an (η, L, A)-quasi-isomorphism $\{f_\lambda\}_{\lambda \in \Lambda_{\mathcal{B}}}$ from $\{h'_{\mathcal{R}}\}$ to $\{h_{\mathcal{R}}\}$, where η is the identity map. It follows from Lemma 5.28 and Theorem 5.19 that Y is quasi-isometric to X_e, the universal cover of the exploded Salvetti complex; hence, Y is quasi-isometric to $G(\Gamma)$. □

In the rest of this section, we look at the special case when $\mathcal{B} = \mathcal{B}(\Gamma)$ is the Davis building of $G(\Gamma)$ (see the beginning of Section 5.1), and we record an observation for later use. In this case, we identify elements of $G(\Gamma)$ with chambers in $\mathcal{B}$.

We denote the word metric on $G(\Gamma)$ by d_w. If we identify $G(\Gamma)$ with chambers of the building $\mathcal{B} = \mathcal{B}(\Gamma)$, then there is another metric on $G(\Gamma)$ defined in Definition 3.6. We caution the reader that these two metrics are not the same. We pick a set of blow-up data $\{h_{\mathcal{R}}\}$ on $\mathcal{B}$, and let $q : Y \to |\mathcal{B}|$ be the associated restriction quotient. Recall that vertices of rank 0 in $|\mathcal{B}|$ can be identified with chambers in $\mathcal{B}$ and, hence, can be identified with $G(\Gamma)$. Thus, the map $q^{-1} : G(\Gamma) \to Y$ is well defined.

LEMMA 5.31
If there exist $L, A > 0$ such that all $\{h_{\mathcal{R}} : \mathcal{R} \to \mathbb{Z}^{T(\mathcal{R})}\}$ are (L, A)-quasi-isometries (here we identify chambers in $\mathcal{R}$ with a subset of $G(\Gamma)$; hence, $\mathcal{R}$ is endowed with an induced metric from d_w), then $q^{-1} : (G(\Gamma), d_w) \to Y$ is an (L', A')-quasi-isometry with its constants depending on L, A, and Γ.

Proof
Let $q' : X_e \to |\mathcal{B}|$ be the $G(\Gamma)$-equivariant canonical restriction quotient constructed in Section 5.1. In this case, $(q')^{-1} : G(\Gamma) \to X_e$ is a quasi-isometry whose constants depend on Γ. Let $h'_{\mathcal{R}}$ be the blow-up data which arises from the 1-data (see Definition 5.3) of q'. Then each $h'_{\mathcal{R}}$ is an isometry. It follows from the assumption that there exists an (η, L, A)-quasi-isomorphism from the blow-up data $\{h'_{\mathcal{R}}\}$ to $\{h_{\mathcal{R}}\}$ with η being the identity map. Thus, there exists a quasi-isometry $X_e \to Y$ which can be represented by a map ϕ of the form in Remark 5.29. Since $q^{-1} = \phi \circ (q')^{-1}$, the lemma follows. □

5.6. An equivariant construction

Let $\mathcal{B} = \mathcal{B}(\Gamma)$ be a right-angled building. Let K be a group which acts on $|\mathcal{B}|$ by automorphisms which preserve the rank of its vertices, and let $K \curvearrowright \mathcal{B}$ and $K \curvearrowright \Lambda_{\mathcal{B}}$ be the induced actions. ($\Lambda_{\mathcal{B}}$ is defined at the beginning of Section 5.3.)

Definition 5.32 (Factor actions)

Pick $\lambda \in \Lambda_{\mathcal{B}}$, and let $\mathcal{R}_\lambda \subset \mathcal{B}$ be a residue of rank 1 such that $T(\mathcal{R}_\lambda) = \lambda$. ($T$ is the type map defined in Section 5.3.) Let K_λ be the stabilizer of λ with respect to the action $K \curvearrowright \Lambda_{\mathcal{B}}$, and let $P(\mathcal{R}_\lambda) = \mathcal{R}_\lambda \times \mathcal{R}_\lambda^\perp$ be the parallel set of $\mathcal{R}_\lambda$ with its product decomposition (see Lemma 3.14, Theorem 3.13). Then $P(\mathcal{R}_\lambda)$ is K_λ-invariant, and K_λ respects the product decomposition of $P(\mathcal{R}_\lambda)$. (Note that, in general, K_λ is smaller than the stabilizer of $P(\mathcal{R}_\lambda)$.) Let $\rho_\lambda : K_\lambda \curvearrowright \mathcal{R}_\lambda$ be the action of K_λ on the $\mathcal{R}_\lambda$-factor. This action ρ_λ is called a *factor action.*

We construct equivariant blow-up data as follows. Pick one representative from each K-orbit of $K \curvearrowright \Lambda_{\mathcal{B}}$, and form the set $\{\lambda_u\}_{u\in U}$. Let K_u be the stabilizer of λ_u. Pick a residue $\mathcal{R}_u \subset \mathcal{B}$ of rank 1 such that $T(\mathcal{R}_u) = \lambda_u$, and let $\rho_u : K_u \curvearrowright \mathcal{R}_u$ be the factor action defined as above.

To obtain a K-equivariant blow-up data, we pick an isometric action $K_u \curvearrowright \mathbb{Z}^{\lambda_u}$ and a K_u-equivariant map $h_{\mathcal{R}_u} : \mathcal{R}_u \to \mathbb{Z}^{\lambda_u}$. (In general, there may not exist nontrivial action $K_u \curvearrowright \mathbb{Z}^{\lambda_u}$ and $h_{\mathcal{R}_u}$; however, in the next section, we will consider situations where there exist nontrivial actions and equivariant maps, and we will plug them into the construction.) If $\mathcal{R}$ is parallel to $\mathcal{R}_u$ with the parallelism map given by $p : \mathcal{R} \to \mathcal{R}_u$, we define $h_{\mathcal{R}} = h_{\mathcal{R}_u} \circ p$. By the previous discussion, there is a factor action $K_u \curvearrowright \mathcal{R}$, and $h_{\mathcal{R}}$ is K_u-equivariant. We run this process for each element in $\{\lambda_u\}_{u\in U}$. If $\lambda \notin \{\lambda_u\}_{u\in U}$, then we fix an element $g_\lambda \in K$ such that $g_\lambda(\lambda) \in \{\lambda_u\}_{u\in U}$. For a rank 1 element $\mathcal{R}$ with $T(\mathcal{R}) = \lambda$, we define

$$h_{\mathcal{R}} = \mathrm{Id} \circ h_{g_\lambda(\mathcal{R})} \circ g_\lambda, \tag{5.33}$$

where $\mathrm{Id} : \mathbb{Z}^{g_\lambda(\lambda)} \to \mathbb{Z}^\lambda$ is the identity map. Let $K_\lambda = g_\lambda^{-1} K_{g_\lambda(\lambda)} g_\lambda$ be the stabilizer of λ. We define the action $K_\lambda \curvearrowright \mathbb{Z}^\lambda$ by letting $g_\lambda^{-1} g g_\lambda$ act on $\mathbb{Z}^\lambda$ by $\mathrm{Id} \circ g \circ \mathrm{Id}^{-1}$ ($g \in K_{g_\lambda(\lambda)}$). Then $h_{\mathcal{R}}$ becomes K_λ-equivariant.

LEMMA 5.34

There exists a collection of maps $\{f_{g,\mathcal{R}} : \mathbb{Z}^{T(\mathcal{R})} \to \mathbb{Z}^{T(g(\mathcal{R}))}\}$ with g ranging over elements in K and $\mathcal{R}$ ranging over rank 1 *residues in $\mathcal{B}$ such that*

(1) *the following diagram commutes:*

$$\begin{array}{ccc} \mathcal{R} & \xrightarrow{h_{\mathcal{R}}} & \mathbb{Z}^{T(\mathcal{R})} \\ g\Big\downarrow & & f_{g,\mathcal{R}}\Big\downarrow \\ g(\mathcal{R}) & \xrightarrow{h_{g(\mathcal{R})}} & \mathbb{Z}^{T(g(\mathcal{R}))} \end{array}$$

(2) $\quad f_{g_1 g_2,\mathcal{R}} = f_{g_1,g_2(\mathcal{R})} \circ f_{g_2,\mathcal{R}}$ *for any* $g_1, g_2 \in K$.

Proof

First we define $f_{g,\mathcal{R}}$. Let $\mathcal{R}$ be a rank 1 residue in $\mathcal{B}$ such that $T(\mathcal{R}) = \lambda_1$. Suppose that $T(g(\mathcal{R})) = \lambda_2$. Then $\lambda_2 = g\lambda_1$. Let g_{λ_1} and g_{λ_2} be the elements we chose before such that $g_{\lambda_i}(\lambda_i) \in \{\lambda_u\}_{u\in U}$ for $i = 1, 2$. (If λ_i is already in $\{\lambda_u\}_{u\in U}$, then let g_{λ_i} be the identity element.) By our choice of $\{\lambda_u\}_{u\in U}$, $g_{\lambda_1}(\lambda_1) = g_{\lambda_2}(\lambda_2)$, and we assume without loss of generality that they are equal to λ_{u_1}. Note that $g_{\lambda_2} g g_{\lambda_1}^{-1}(\lambda_{u_1}) = \lambda_{u_1}$. Thus, $g_{\lambda_2} g g_{\lambda_1}^{-1} \in K_{u_1}$; in particular, this element also acts on $\mathbb{Z}^{\lambda_{u_1}}$. Define

$$f_{g,\mathcal{R}} = \mathbb{Z}^{T(\mathcal{R})} \xrightarrow{\mathrm{Id}} \mathbb{Z}^{\lambda_{u_1}} \xrightarrow{g_{\lambda_2} g g_{\lambda_1}^{-1}} \mathbb{Z}^{\lambda_{u_1}} \xrightarrow{\mathrm{Id}} \mathbb{Z}^{T(g(\mathcal{R}))}.$$

Let $\mathcal{R}_1$ be a rank 1 residue with $T(\mathcal{R}_1) = \lambda_{u_1}$. Now we verify the commutativity of the diagram. To simplify our notation, we identify all these copies of $\mathbb{Z}$'s with a single $\mathbb{Z}$, and we omit the identification maps between them. Then

$$\begin{aligned} f_{g,\mathcal{R}} \circ h_{\mathcal{R}} &= (g_{\lambda_2} g g_{\lambda_1}^{-1}) \circ (h_{\mathcal{R}_1} \circ g_{\lambda_1}) = (g_{\lambda_2} g g_{\lambda_1}^{-1} \circ h_{\mathcal{R}_1}) \circ g_{\lambda_1} \\ &= (h_{\mathcal{R}_1} \circ g_{\lambda_2} g g_{\lambda_1}^{-1}) \circ g_{\lambda_1} = (h_{\mathcal{R}_1} \circ g_{\lambda_2}) \circ g \\ &= h_{g(\mathcal{R})} \circ g, \end{aligned}$$

where the first and last equalities follow from (5.33), and the third equality follows from the K_{u_1}-equivariance of $h_{\mathcal{R}_1}$.

To see (2), let $T(\mathcal{R}) = \lambda_1$, $T(g_2(\mathcal{R})) = \lambda_2$, and $T(g_1 g_2(\mathcal{R})) = \lambda_3$. We define g_{λ_1}, g_{λ_2}, and g_{λ_3} in a similar way as before. Then (2) follows from the equality $g_{\lambda_3} g_1 g_2 g_{\lambda_1}^{-1} = g_{\lambda_3} g_1 g_{\lambda_2}^{-1} \cdot g_{\lambda_2} g_2 g_{\lambda_1}^{-1}$. □

Let Ψ be the fiber functor associated with the above blow-up data, and let $q : Y \to |\mathcal{B}|$ be the corresponding restriction quotient. Lemma 5.25 implies that K acts on Ψ by natural transformations; hence, there is an induced action $K \curvearrowright Y$ and q is K-equivariant.

Remark 5.35

The previous construction depends on several choices:

(1) the choice of the set $\{\lambda_u\}_{u\in U}$,
(2) the choice of the isometric action $K_u \curvearrowright \mathbb{Z}^{\lambda_u}$ and the K_u-equivariant map $h_{\mathcal{R}_u} : \mathcal{R}_u \to \mathbb{Z}^{\lambda_u}$, and
(3) the choice of the elements g_λ.

6. Quasiactions on RAAGs

In this section we will apply the construction in Section 5.6 to study quasiactions on RAAGs. We assume $G(\Gamma) \neq \mathbb{Z}$ throughout Section 6.

6.1. *The cubulation*

Recall that $G(\Gamma)$ acts on $X(\Gamma)$ by deck transformations, and this action is simply transitive on the vertex set of $X(\Gamma)$. By picking a basepoint in $X(\Gamma)$, we identify $G(\Gamma)$ with the 0-skeleton of $X(\Gamma)$.

Definition 6.1

A quasi-isometry $\phi : G(\Gamma) \to G(\Gamma)$ is *flat-preserving* if it is a bijection and for every standard flat $F \subset X(\Gamma)$ there is a standard flat $F' \subset X(\Gamma)$ such that ϕ maps the 0-skeleton of F bijectively onto the 0-skeleton of F'. The standard flat F' is uniquely determined, and we denote it by $\phi_*(F)$. Note that if ϕ is flat-preserving, then ϕ^{-1} is also flat-preserving.

Our main goal for this section is the following theorem, which establishes Theorem 1.5.

THEOREM 6.2

If the outer automorphism group $\mathrm{Out}(G(\Gamma))$ *is finite and* $G(\Gamma) \not\simeq \mathbb{Z}$, *then any* (L, A)-*quasiaction* $\rho : H \curvearrowright X(\Gamma)$ *is quasiconjugate to an action* $\hat{\rho}$ *of* H *by cubical isometries on a uniformly locally finite* $\mathrm{CAT}(0)$ *cube complex* Y. *Moreover, the following hold.*

(1) *If* ρ *is cobounded, then* $\hat{\rho}$ *is cocompact.*
(2) *If* ρ *is proper, then* $\hat{\rho}$ *is proper.*
(3) *Let* $|\mathcal{B}|$ *be the Davis realization of the right-angled building associated with* $G(\Gamma)$, *let* $H \curvearrowright |\mathcal{B}|$ *be the induced action, and let* $X_e = X_e(\Gamma)$ *be the universal cover of the exploded Salvetti complex for* $G(\Gamma)$. *Then* Y *fits into the following commutative diagram:*

$$\begin{tikzcd} X_e \arrow[rr, "\phi"] \arrow[dr, "q'"'] & & Y \arrow[dl, "q"] \\ & {|\mathcal{B}|} & \end{tikzcd}$$

Here q', q, and ϕ are restriction quotients. The map ϕ is a quasi-isometry whose constants depend on the constants of the quasiaction ρ, and q is H-equivariant.

Proof
By Theorem 1.3, without loss of generality we can assume that $\rho : H \curvearrowright G(\Gamma)$ is an action by flat-preserving bijections which are also (L, A)-quasi-isometries.

On the one hand, we want to think of $G(\Gamma)$ as a metric space with the word metric with respect to its standard generating set or, equivalently, with the induced l^1-metric from $X(\Gamma)$. On the other hand, we want to treat $G(\Gamma)$ as a right-angled building (see Section 5.1); more precisely, we want to identify points in $G(\Gamma)$ with chambers in the associated right-angled building of $G(\Gamma)$. Then ρ preserves the spherical residues in $G(\Gamma)$; thus, there is an induced $\rho_{|\mathcal{B}|} : H \curvearrowright |\mathcal{B}|$ on the Davis realization $|\mathcal{B}|$ of $G(\Gamma)$.

Let Λ be the collection of parallel classes of standard geodesic lines in $X(\Gamma)$ (i.e., Λ is the collection of parallel classes of rank 1 residues in $G(\Gamma)$), and let T be the type map defined at the beginning of Section 5.3; in other words, if $\mathcal{R}$ is a spherical residue which comes from a standard flat in $X(\Gamma)$, then $T(\mathcal{R})$ is the set of parallel classes of standard lines that have representatives contained in this standard flat. Since $\rho : H \curvearrowright G(\Gamma)$ is flat-preserving (see Theorem 1.3), it preserves parallel classes of standard flats. Thus, there is an induced action $\rho_\Lambda : H \curvearrowright \Lambda$. For each $\lambda \in \Lambda$, let H_λ be the stabilizer of λ. Pick a residue $\mathcal{R}$ in the parallel class λ, and let $\rho_\lambda : H_\lambda \curvearrowright \mathcal{R}$ be the factor action in Definition 5.32. Note that $\mathcal{R}$ is an isometrically embedded copy of $\mathbb{Z}$ with respect to the metric on $G(\Gamma)$; moreover, ρ_λ is an action such that its maps are (L', A')-quasi-isometries. Here we can choose L' and A' such that they depend only on the constant D of Theorem 1.3, so in particular they do not depend on λ and $\mathcal{R}$.

For the action $\rho_\Lambda : H \curvearrowright \Lambda$, we pick a representative from each H-orbit and form the set $\{\lambda_u\}_{u \in U}$. By the construction in Section 5.6, it remains to choose an isometric action $H_{\lambda_u} \curvearrowright \mathbb{Z}^{\lambda_u}$ and an H_{λ_u}-equivariant map $h_{\mathcal{R}_u} : \mathcal{R}_u \to \mathbb{Z}^{\lambda_u}$ for each $u \in U$. ($\mathcal{R}_u$ is a residue in the parallel class λ_u.) The choice is provided by the following result, whose proof is postponed to Section 7.

PROPOSITION 6.3
If a group U has an action on $\mathbb{Z}$ by (L, A)-quasi-isometries, then there exists another action $U \curvearrowright \mathbb{Z}$ by isometries which is related to the original action by a surjective equivariant (L', A')-quasi-isometry $f : \mathbb{Z} \to \mathbb{Z}$, where L' and A' depend on L and A.

We caution the reader that here we require $U \curvearrowright \mathbb{Z}$ to be an action, not a quasiaction (see Remark 1.15).

From the above data, we produce H-equivariant blow-up data $h_{\mathcal{R}} : \mathcal{R} \to \mathbb{Z}^{T(\mathcal{R})}$ for each rank 1 residue $\mathcal{R} \subset G(\Gamma)$ as in Section 5.6. Note that each $h_{\mathcal{R}}$ is an (L'', A'')-quasi-isometry with constants depending only on L and A.

Let $q : Y \to |\mathcal{B}|$ be the restriction quotient associated with the above blow-up data. Then there is an induced action $H \curvearrowright Y$ by isomorphisms, and q is H-equivariant. It follows from Lemma 5.18 that Y is uniformly locally finite.

Claim
There exists an (L_1, A_1)-quasi-isometry $G(\Gamma) \to Y$ with L_1 and A_1 depending only on L and A.

Proof of claim
Let $h'_{\mathcal{R}} : \mathcal{R} \to \mathbb{Z}^{T(\mathcal{R})}$ be another blow-up data such that each $h'_{\mathcal{R}}$ is an isometry (such blow-up data always exists), and let $q' : Y' \to |\mathcal{B}|$ be the associated restriction quotient. By Theorem 5.19, Y' is isomorphic to X_e, which is the universal cover of the exploded Salvetti complex introduced in Section 5.1. For any $\lambda \in \Lambda$, we define $f_\lambda = h_{\mathcal{R}} \circ (h'_{\mathcal{R}})^{-1}$, where $\mathcal{R}$ is a residue such that $T(\mathcal{R}) = \lambda$ and the definition of f_λ does not depend on $\mathcal{R}$. Each f_λ is an (L'', A'')-quasi-isometry, and the collection of all f_λ's induces a quasi-isomorphism between the blow-up data $\{h'_{\mathcal{R}}\}$ and $\{h_{\mathcal{R}}\}$. It follows from Lemma 5.28 that there exists a quasi-isometry between $\varphi : Y' \cong X_e \to Y$, and the claim follows. ☐

Let B_0 be the set of vertices of rank 0 in $|\mathcal{B}|$. There is a natural identification of B_0 with $G(\Gamma)$. Letting $Y_0 = q^{-1}(B_0)$, we get that q induces a bijection between Y_0 and B_0. Let $q' : Y' \to |\mathcal{B}|$ be as in the previous paragraph, and define $Y'_0 = (q')^{-1}(B_0)$. It follows from Lemma 5.18(2) that Y'_0 and Y_0 are D-dense in Y' and Y, respectively, for D depending on L and A. Note that $q^{-1} : G(\Gamma) \to Y_0$ is H-equivariant, and if the action $\rho : H \curvearrowright G(\Gamma)$ is cobounded, then $H \curvearrowright Y$ is cocompact.

The above quasi-isometry φ can be represented by $q^{-1} \circ q' : Y'_0 \to Y_0$ (see Remark 5.29). By Lemma 5.31, $(q')^{-1} : B_0 = G(\Gamma) \to Y'_0$ is also a quasi-isometry, and thus, $q^{-1} : G(\Gamma) \to Y_0$ is a quasi-isometry. This map is H-equivariant, so if $\rho : H \curvearrowright G(\Gamma)$ is proper, then $H \curvearrowright Y$ is also proper.

It remains to produce the map ϕ in Theorem 6.2(3). This relies on a refinement of the above discussion. Instead of requiring each $h'_{\mathcal{R}}$ in the proof of the above claim to be an isometry, it is possible to choose each $h'_{\mathcal{R}}$ such that:

(1) $h'_{\mathcal{R}}$ is a bijection;

(2) $h'_{\mathcal{R}}$ is an (L_2, A_2)-quasi-isometry; and

(3) $f_\lambda : \mathbb{Z}^\lambda \to \mathbb{Z}^\lambda$ is a surjective map which is either increasing or decreasing and, hence, can be extended to a surjective cubical map $\mathbb{R}^\lambda \to \mathbb{R}^\lambda$.

The surjectivity in (3) comes from our choice in Proposition 6.3. In this case, the space Y' is still isomorphic to X_e (Theorem 5.19). Let Ψ and Ψ' be the fiber functors associated with the blow-up data $\{h_{\mathcal{R}}\}$ and $\{h'_{\mathcal{R}}\}$. As in the proof of Lemma 5.25, the f_λ's induce a natural transformation from Ψ' to Ψ which is made up of a collection of surjective cubical maps from objects in Ψ' to objects in Ψ; moreover, these maps are quasi-isometries with uniform quasi-isometry constants. Recall that we can describe Y as the quotient of the disjoint collection $\{\sigma \times \Psi(\sigma)\}_{\sigma \in \mathrm{Face}(|\mathcal{B}|)}$ (see the proof of Theorem 4.15), and a similar description holds for Y'. Thus, there is a surjective cubical map $\phi : Y' \to Y$ induced by the natural transformation. Actually ϕ is a restriction quotient, since the inverse image of each hyperplane is a hyperplane. We also know that ϕ is a quasi-isometry by Lemma 4.19. ☐

COROLLARY 6.4

Suppose the outer automorphism group $\mathrm{Out}(G(\Gamma))$ *is finite. Then H is quasi-isometric to $G(\Gamma)$ if and only if there exists an H-equivariant restriction quotient map $q : Y \to |\mathcal{B}|$ such that:*

(1) *$|\mathcal{B}|$ is the Davis realization of some right-angled Γ-building;*

(2) *the action $H \curvearrowright Y$ is geometric; and*

(3) *if $v \in |\mathcal{B}|$ is a vertex of rank k, then $q^{-1}(v) = \mathbb{E}^k$.*

Proof

The only if direction follows from Theorem 6.2. For the if direction, it suffices to show that Y is quasi-isometric to $G(\Gamma)$. Let Φ be the fiber functor associated with q.

Pick a vertex $v \in |\mathcal{B}|$ of rank k, and let $F_v = q^{-1}(v)$. We claim $\mathrm{Stab}(v)$ acts cocompactly on F_v. By a standard argument, to prove this it suffices to show that $\{h(F_v)\}_{h \in H}$ is a locally finite family in Y. Suppose there exists an R-ball $N \subset Y$ such that there are infinitely many distinct elements in $\{h(F_v)\}_{h \in H}$ which have nontrivial intersection of N. Since Y admits a geometric action, it is locally finite, and thus, there exists a vertex $x \in |\mathcal{B}|$ which is contained in infinitely many distinct elements in $\{h(F_v)\}_{h \in H}$. This is impossible, since if $h(F_v) \neq h'(F_v)$, then $q(h(F_v))$ and $q(h'(F_v))$ are distinct vertices in $|\mathcal{B}|$ by the H-equivariance of q.

Consider a cube $\sigma \subset |\mathcal{B}|$, and let v be its vertex of minimal rank. We claim $\Phi(\sigma) \to \Phi(v)$ is surjective and, hence, is an isometry. By (3), the action $H \curvearrowright |\mathcal{B}|$ preserves the rank of the vertices; thus, $\mathrm{Stab}(\sigma) \subset \mathrm{Stab}(v)$. We know that $\mathrm{Stab}(v)$ acts cocompactly on $q^{-1}(v)$; since the poset $\{w \geq v\}$ is finite, $\mathrm{Stab}(\sigma)$ has finite index in $\mathrm{Stab}(v)$, and so $\mathrm{Stab}(\sigma)$ also acts cocompactly on $q^{-1}(v)$. Now the image of $\Phi(\sigma) \to \Phi(v)$ is a convex subcomplex of $q^{-1}(v)$ that is $\mathrm{Stab}(\sigma)$-invariant, so it coincides with $q^{-1}(v)$.

By Corollary 5.13, we can assume that q is the restriction quotient of a set of blow-up data $\{h_{\mathcal{R}}\}$. Pick a vertex $v \in |\mathcal{B}|$ of rank 1, and let $D(v)$ be the downward complex of v (see Section 5.4). Let $\mathcal{R}_v \subset \mathcal{B}$ be the associated residue, and let $\mathcal{R}_v \to \mathbb{R}^{T(\mathcal{R})}$ be the map induced by $h_{\mathcal{R}_v}$. Then $q^{-1}(D_v)$ is isomorphic to the mapping cylinder of this map. Since $\mathrm{Stab}(v)$ acts cocompactly on $q^{-1}(D_v)$, there are only finite many orbits of vertices of rank 1, and the assumptions of Corollary 5.30 are satisfied. It follows that Y is quasi-isometric to $G(\Gamma)$. $\square$

It is possible to drop the H-equivariant assumption on q under the following conditions. Here we do not put any assumption on Γ.

THEOREM 6.5
Let $\mathcal{B}$ be a right-angled Γ-building. Suppose that $q : Y \to |\mathcal{B}|$ is a restriction quotient such that, for every cube $\sigma \subset |\mathcal{B}|$ and every interior point $x \in \sigma$, the point inverse $q^{-1}(x)$ is a copy of $\mathbb{E}^{\mathrm{rank}(v)}$, where $v \in \sigma$ is the vertex of minimal rank. If H acts geometrically on Y by automorphisms, then H is quasi-isometric to $G(\Gamma)$.

Proof
First we assume Γ satisfies the assumption of Theorem 5.22. Then the above result is a consequence of Corollary 5.13, Theorem 5.22, and the argument in Corollary 6.4.

For arbitrary Γ, we make a join decomposition $\Gamma = \Gamma_1 \circ \Gamma_2 \circ \cdots \circ \Gamma_k \circ \Gamma'$, where Γ' satisfies the assumption of Theorem 5.22, and all the Γ_i's are discrete graphs with more than one vertex. By Corollary 5.14, there is an induced cubical product decomposition $Y = Y_1 \times Y_2 \times \cdots \times Y_k \times Y'$ and also restriction quotients $q_i : Y_i \to |\mathcal{B}_i|$ and $q' : Y' \to |\mathcal{B}'|$ which satisfy the assumption of the theorem. By [14, Proposition 2.6], we assume that H respects this product decomposition by passing to a finite-index subgroup. Since Y' is locally finite and cocompact, the same argument in Corollary 6.4 implies Y' is quasi-isometric to $G(\Gamma')$. Each Y_i is a locally finite and cocompact tree which is not quasi-isometric to a line. So Y_i is quasi-isometric to $G(\Gamma_i)$. Thus, Y is quasi-isometric to $G(\Gamma)$. $\square$

COROLLARY 6.6
Suppose $\mathrm{Out}(G(\Gamma))$ *is finite and* $G(\Gamma) \not\simeq \mathbb{Z}$. *Let $\mathcal{B}$ be the right-angled building of* $G(\Gamma)$. *Then H is quasi-isometric to* $G(\Gamma)$ *if and only if H acts geometrically on a blowup of $\mathcal{B}$ in the sense of Section 5.3 by automorphisms.*

6.2. Reduction to nicer actions

Though every action $\rho : H \curvearrowright G(\Gamma)$ by flat-preserving bijections which are also (L, A)-quasi-isometries is quasiconjugate to an isometric action $H \curvearrowright Y$ as in Theorem 6.2, it is in general impossible to take $Y = X(\Gamma)$, even if the action ρ is proper and cobounded.

Definition 6.7

Let $V = \mathbb{Z}/2 \oplus \mathbb{Z}$ with the generator of $\mathbb{Z}/2$ and $\mathbb{Z}$ denoted by a and b, respectively. Let $V \overset{\rho_0}{\curvearrowright} \mathbb{Z}$ be the action where $\rho_0(b)(n) = n + 2$ and $\rho_0(a)$ acts on $\mathbb{Z}$ by flipping $2n$ and $2n + 1$ for all $n \in \mathbb{Z}$. An action $U \curvearrowright \mathbb{Z}$ is 2-*flipping* if it factors through the action $V \overset{\rho_0}{\curvearrowright} \mathbb{Z}$ via an epimorphism $U \to V$.

LEMMA 6.8

Let $\rho_U : U \curvearrowright \mathbb{Z}$ be a 2-flipping action. Then ρ_U is not conjugate to an action by isometries on $\mathbb{Z}$ (with respect to the standard metric on $\mathbb{Z}$).

Proof

Suppose there exists a permutation $p : \mathbb{Z} \to \mathbb{Z}$ which conjugates ρ_U to an isometric action. Let $h : U \to V$ be the epimorphism in Definition 6.7. Pick $k_1, k_2 \in U$ such that $h(k_1)$ is of order 2 and $h(k_2)$ is of order infinity. Then pk_1p^{-1} is a reflection of $\mathbb{Z}$, and pk_2p^{-1} is a translation. However, this is impossible since $h(k_1)$ and $h(k_2)$ commute. $\square$

LEMMA 6.9

There does not exist an action $\rho_U : U \curvearrowright \mathbb{Z}$ by (L, A)-quasi-isometries with the following property. U has two subgroups U_1 and U_2 such that $\rho_U|_{U_1}$ is conjugate to a 2-flipping action and $\rho_U|_{U_2}$ is conjugate to a transitive action on $\mathbb{Z}$ by translations.

Proof

By Proposition 6.3, there exists an isometric action $\rho'_U : U \curvearrowright \mathbb{Z}$ and a U-equivariant surjective map $f : U \overset{\rho_U}{\curvearrowright} \mathbb{Z} \longrightarrow U \overset{\rho'_U}{\curvearrowright} \mathbb{Z}$. We claim f is also injective. Given this claim, we can deduce a contradiction to Lemma 6.8 by restricting the action to U_1. To see the claim, we restrict the action to U_2. Thus, we can assume without loss of generality that ρ_U is a transitive action by translations. Suppose that $f(a_1) = f(a_1 + k)$ for $a_1, k \in \mathbb{Z}$ and $k \neq 0$. Then the equivariance of f implies $f(a_1) = f(a_1 + nk)$ for any integer $n \in \mathbb{Z}$, which contradicts that f is a quasi-isometry. $\square$

THEOREM 6.10

Suppose that $G(\Gamma)$ is a RAAG with $|\mathrm{Out}(G(\Gamma))| < \infty$ and $G(\Gamma) \not\simeq \mathbb{Z}$. Then there is a *pair H, H' of finitely generated groups quasi-isometric to $G(\Gamma)$ that does not admit discrete, virtually faithful cocompact representations into the same locally compact topological group.*

Recall that a discrete, virtually faithful cocompact representation from H to a locally compact group $\hat{G}$ is a homomorphism $h : H \to \hat{G}$ such that its kernel is finite, and its image is a cocompact lattice.

Proof
Pick a vertex $u \in \Gamma$. Let Γ_1 and Γ_2 be the full subgraphs of Γ spanned by vertices in $\Gamma \setminus \{u\}$ and vertices adjacent to u, respectively. For $i = 1, 2$, let Y_i be the Salvetti complex $S(\Gamma_i)$. There is a natural embedding $Y_2 \hookrightarrow Y_1$.

Let $\{\mathcal{G}_i\}_{i=1}^4$ be the four graphs in the following figure, from left to right, and let $V_i \subset \mathcal{G}_i$ be the collection of dotted vertices indicated in the graph. For $1 \leq i \leq 4$, we define Z_i to be the space obtained by gluing $Y_1 \times V_i$ and $Y_2 \times \mathcal{G}_i$ along $Y_2 \times V_i$. Note that Z_1 is the Salvetti complex $S(\Gamma)$. Let $\tilde{Z}_i$ be the universal cover of Z_i.

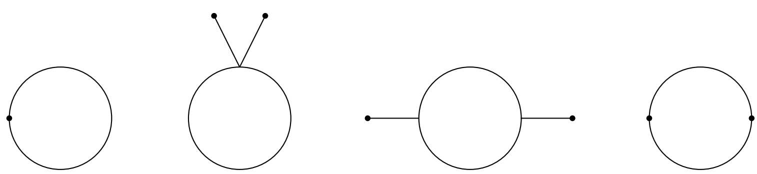

Let y_2 be the unique vertex in Y_2. Then $\{y_2\} \times V_i$ gives a collection of vertices in Z_i. Let $\tilde{V}_i$ be the collection of vertices in $\tilde{Z}_i$ which are mapped to vertices in $\{y_2\} \times V_i$ under the covering map $\tilde{Z}_i \to Z_i$.

We have a homotopy equivalence $\mathcal{G}_3 \to \mathcal{G}_2$ by collapsing the upper half of the circle and a homotopy equivalence $\mathcal{G}_3 \to \mathcal{G}_4$ by collapsing the two edges. These maps induce homotopy equivalences $Z_3 \to Z_2$ and $Z_3 \to Z_4$, which give us identifications between $\tilde{V}_3$, $\tilde{V}_2$, and $\tilde{V}_4$ up to choices of base points. Since there is a two-sheeted covering map $Z_4 \to Z_1$, we identify $\tilde{V}_4$ with $G(\Gamma)$.

The involution of $\mathcal{G}_2$ which fixes each point in the circle and flips the two points in V_2 induces an involution $\alpha : Z_2 \to Z_2$. This gives rise to an action of $H = \pi_1 Z_2 \rtimes \mathbb{Z}/2\mathbb{Z}$ on $\tilde{Z}_2$. Note that H preserves the set $\tilde{V}_2$. Since we have identified V_2 with $G(\Gamma)$ in the previous paragraph, we obtain an action $\rho_H : H \curvearrowright G(\Gamma)$ by quasi-isometries which is discrete and cobounded. One may readily verify from our construction that ρ_H is an action by flat-preserving quasi-isometries; moreover, the factor actions of H are either transitive actions on $\mathbb{Z}$ or 2-flipping actions.

We claim that $G(\Gamma)$ and H do not admit discrete, virtually faithful cocompact representations into the same locally compact topological group. Suppose such a topological group $\hat{G}$ exists. Then by [49, Chapter 6], $\hat{G}$ has a quasiaction on $G(\Gamma)$. We assume $\hat{G}$ acts on $G(\Gamma)$ by flat-preserving quasi-isometries as before. Then there are restriction actions $\rho'_{G(\Gamma)} : G(\Gamma) \curvearrowright G(\Gamma)$ and $\rho'_H : H \curvearrowright G(\Gamma)$ which are discrete and cobounded. Since any two discrete and cobounded quasiactions $H \curvearrowright G(\Gamma)$ are quasiconjugate, it follows from Theorem 1.3 that ρ_H and ρ'_H are conjugate by a flat-preserving quasi-isometry. Thus, factor actions of ρ'_H are conjugate to factor actions of ρ_H by bijective quasi-isometries. Similarly, we deduce that the factor actions of $\rho'_{G(\Gamma)}$ are conjugate to transitive actions by left translations on $\mathbb{Z}$ via bijective quasi-isometries. Note that the factor actions of $\rho'_{G(\Gamma)}$ and the factor actions of ρ'_H are both restrictions of factor actions of $\hat{G} \curvearrowright G(\Gamma)$; however, this is impossible by Lemma 6.9. $\square$

COROLLARY 6.11
The group $H = G(\Gamma') \rtimes \mathbb{Z}/2$ cannot act geometrically on $X(\Gamma)$.

We now give a criterion for when one can quasiconjugate a quasiaction on $X(\Gamma)$ to an isometric action $H \curvearrowright X(\Gamma)$.

THEOREM 6.12
Let $\rho : H \curvearrowright G(\Gamma)$ be an action by flat-preserving bijections. If, for each $\lambda \in \Lambda$, the factor action $\rho_\lambda : H_\lambda \curvearrowright \mathbb{Z}$ can be conjugate to an action by isometries with respect to the word metric of $\mathbb{Z}$, then there is a flat-preserving bijection $g : G(\Gamma) \to G(\Gamma)$ which conjugates $\rho : H \curvearrowright G(\Gamma)$ to an action $\rho' : H \curvearrowright X(\Gamma)$ by flat-preserving isometries. If ρ is also an action by (L, A)-quasi-isometries, then g can be taken to be a quasi-isometry.

Proof
We repeat the construction in Section 6.1 and assume that each $h_{\mathcal{R}} : \mathcal{R} \to \mathbb{Z}^{T(\mathcal{R})}$ is a bijection. Let $q : Y \to |\mathcal{B}|$, Y_0, $q^{-1} : G(\Gamma) \to Y_0$, and the action $\hat{\rho} : H \curvearrowright Y$ by automorphisms be as in Section 6.1. Recall that q^{-1} is H-equivariant. There is an isomorphism $i : Y \to X_e$ by Theorem 5.19; moreover, $i(Y_0)$ is exactly the collection of 0-dimensional standard flats X_0 in X_e. We deduce from the construction of i that the isometric action $H \curvearrowright X_e$ induced by $\hat{\rho}$ preserves standard flats in X_e. By the construction of X_e, there exists a natural identification $f : X_0 \to G(\Gamma)$ such that any automorphism of X_e which preserves its standard flats induces a flat-preserving isometry of $G(\Gamma)$ (with respect to the word metric) via f. It suffices to take $g = f \circ i \circ q^{-1}$. $\square$

Suppose that we have already conjugated the flat-preserving action $\rho : H \curvearrowright G(\Gamma)$ to an action $\rho' : H \curvearrowright X(\Gamma)$ (or $H \curvearrowright G(\Gamma)$) by flat-preserving isometries. We ask whether it is possible to further conjugate ρ' to an action by left translations.

We can orient each 1-cell in $S(\Gamma)$ and label it by the associated generator. This lifts to orientations and labels of edges of $X(\Gamma)$. If H preserves this orientation and labeling, then ρ' is already an action by left translations. In general, it suffices to require that H preserves a possibly different orientation and labeling which satisfy several compatibility conditions. Now we recall the following definitions from [36].

Definition 6.13 (Coherent ordering)
A *coherent ordering* for $G(\Gamma)$ is a blow-up data for $G(\Gamma)$ such that each map $h_{\mathcal{R}}$ is a bijection. Two coherent orderings are *equivalent* if their maps agree up to translations.

Let $\mathcal{P}(\Gamma)$ be the extension complex defined in Section 3.3. Note that we can identify $\Lambda_{G(\Gamma)}$ with the 0-skeleton of $\mathcal{P}(\Gamma)$. Any flat-preserving action $H \curvearrowright G(\Gamma)$ induces an action $H \curvearrowright \mathcal{P}(\Gamma)$ by simplicial isomorphisms. Let $F(\Gamma)$ be the flag complex of Γ.

Definition 6.14 (Coherent labeling)
Recall that, for each vertex $x \in X(\Gamma)$, there is a natural simplicial embedding $i_x : F(\Gamma) \to \mathcal{P}(\Gamma)$ by considering the standard flats passing through x. A *coherent labeling* of $G(\Gamma)$ is a simplicial map $L : \mathcal{P}(\Gamma) \to F(\Gamma)$ such that $L \circ i_x : F(\Gamma) \to F(\Gamma)$ is a simplicial isomorphism for every vertex $x \in X(\Gamma)$.

The next result follows from [36, Lemma 5.7].

LEMMA 6.15
Let $\rho' : H \curvearrowright G(\Gamma)$ be an action by flat-preserving bijections, and let $H \curvearrowright \mathcal{P}(\Gamma)$ be the induced action. If there exists an H-invariant coherent ordering and an H-invariant coherent labeling, then ρ' is conjugate to an action by left translations.

Since each vertex of $\mathcal{P}(\Gamma)$ corresponds to a parallel class of v-residues for a vertex $v \in \Gamma$, this gives a labeling of vertices of $\mathcal{P}(\Gamma)$ by vertices of Γ. We can extend this labeling map to a simplicial map $L : \mathcal{P}(\Gamma) \to F(\Gamma)$, which gives rise to a coherent labeling.

COROLLARY 6.16
Let $\rho : H \curvearrowright G(\Gamma)$ be an action by flat-preserving bijections. Suppose that the following statements hold.

(1) *The induced action* $H \curvearrowright \mathcal{P}(\Gamma)$ *preserves the vertex labeling of* $\mathcal{P}(\Gamma)$ *as above.*
(2) *For each vertex* $v \in \mathcal{P}(\Gamma)$*, the action* $\rho_v : H_v \curvearrowright \mathbb{Z}$ *is conjugate to an action by translations.*

Then ρ *is conjugate to an action* $H \curvearrowright G(\Gamma)$ *by left translations.*

Note that condition (2) is equivalent to the existence of an H-invariant coherent ordering.

7. Actions by quasi-isometries on $\mathbb{Z}$

In this section we prove Proposition 6.3.

7.1. Tracks

Tracks were introduced in [25]. They are hypersurface-like objects in 2-dimensional simplicial complexes.

Definition 7.1 (Tracks)
Let K be a 2-dimensional simplicial complex. A *track* $\tau \subset K$ is a connected embedded finite simplicial graph such that the following hold.
(1) For each 2-simplex $\Delta \subset K$, $\tau \cap \Delta$ is a finite disjoint union of arcs such that the endpoints of each arc are in the interior of edges of Δ.
(2) For each edge $e \in K$, $\tau \cap e$ is a discrete set in the interior of e. Let $\{\Delta_\lambda\}_{\lambda \in \Lambda}$ be the collection of 2-simplices that contains e. If $v \in \tau \cap e$, then for each λ, $\tau \cap \Delta_\lambda$ contains an arc that ends with v.

Given a track $\tau \subset K$, we define the *support* of τ, denoted $\mathrm{Spt}(\tau)$, to be the minimal subcomplex of K which contains τ. We can view hyperplanes defined in Section 3.2 as analogues of tracks in the cubical setting. Each track $\tau \subset K$ has a regular neighborhood which fibers over τ. When K is simply connected, $K \setminus \tau$ has two connected components; moreover, the regular neighborhood of τ is homeomorphic to $\tau \times (-\epsilon, \epsilon)$.

Two tracks τ_1 and τ_2 are *parallel* if $\mathrm{Spt}(\tau_1) = \mathrm{Spt}(\tau_2)$ and there is a region homeomorphic to $\tau_1 \times (0, \epsilon)$ bounded by τ_1 and τ_2. A track $\tau \subset K$ is *essential* if the components of $K \setminus \tau$ are unbounded. The following result follows from [25, Proposition 3.1].

LEMMA 7.2
If K *is simply connected and has more than one end, then there exists an essential track* $\tau \subset K$.

Next we look at essential tracks which are "minimal"; these turn out to behave like minimal surfaces. First we metrize K as in [55].

Let $\Delta = \Delta(\xi_1\xi_2\xi_3)$ be an ideal triangle in the hyperbolic plane. We mark a point on each side of the triangle as follows. Let ϕ be the unique isometry which fixes ξ_3 and flips ξ_1 and ξ_2. We mark the unique point in $\overline{\xi_1\xi_2}$ which is fixed by ϕ. Other sides of Δ are marked similarly. This is called a *marked ideal triangle*.

We identify each 2-simplex of K with a marked ideal triangle in the hyperbolic plane and glue these triangles by isometries which identify the marked points. This gives a collection of complete metrics on each connected component of $K - K^{(0)}$ which is not homeomorphic to the interval $(0, 1)$. We denote this collection of metrics by $d_{\mathbb{H}}$. If a group G acts on K by simplicial isomorphisms, then G also acts by isometries on $(K, d_{\mathbb{H}})$. The original definition in [55] did not require these marked points (see Remark 7.4 for why we add them).

We assume each arc in the track is rectifiable. Thus, each track τ has a well-defined $d_{\mathbb{H}}$-length, which we denote by $l(\tau)$. We also define the *weight* of τ, denoted by $w(\tau)$, to be cardinality of $\tau \cap K^{(1)}$. The *complexity* $c(\tau)$ is defined to be the ordered pair $(w(\tau), l(\tau))$. We order the complexity lexicographically, namely, $c(\tau_1) < c(\tau_2)$ if and only if $w(\tau_1) < w(\tau_2)$ or $w(\tau_1) = w(\tau_2)$ and $l(\tau_1) < l(\tau_2)$.

The following result follows from [55, Lemmas 2.11, 2.14].

LEMMA 7.3
Suppose that K is a uniformly locally finite and simply connected simplicial 2-complex with at least two ends. Suppose that K does not contain separating vertices. Then there exists an essential track $\tau \subset K$ which has the least complexity with respect to $d_{\mathbb{H}}$ among all essential tracks in K.

Remark 7.4
Let $\{\tau_i\}_{i=1}^{\infty}$ be a minimizing sequence. Since K is uniformly locally finite, there are only finitely many combinatorial possibilities for $\{\tau_i\}_{i=1}^{\infty}$. Thus, we can assume all the τ_i's are inside a finite subcomplex L. Moreover, we can construct a hyperbolic metric $d_{\mathbb{H}}$ on L as above, and it suffices to work in the space $(L, d_{\mathbb{H}})$. However, if we do not use marked points in the construction of the hyperbolic metric on K, then each τ_i may sit inside a copy of L with different shears along the edges of L.

In [55], K is assumed to be cocompact, so one does not need to worry about the above issue.

Remark 7.5
If we metrize each simplex in K with the Euclidean metric, then Lemmas 7.3 and 7.6 may not be true. For example, one can take the following picture, where the dotted

line is part of some track τ. Once we shorten τ, it may hit the central vertex of the hexagon. However, this cannot happen if we have hyperbolic metrics on each simplex. Once τ gets too close to some vertex, then it takes a large amount of length for τ to escape that vertex, since $d_{\mathbb{H}}$ is complete. (Actually it does not matter if $d_{\mathbb{H}}$ is not complete, since we also have an upper bound on the weight of τ.)

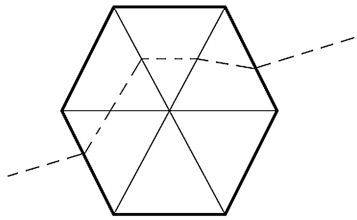

The next result can be proved in a similar fashion.

LEMMA 7.6
Let K be a simply connected simplicial 2-complex. Let $A \subset K$ be a uniformly locally finite subcomplex such that
(1) *A contains an essential track of K, and*
(2) *A does not contain any separating vertex of K.*
Then there exists an essential track τ of K which has the least complexity among all essential tracks of K with support in A.

LEMMA 7.7 ([55, Lemma 2.7])
Let τ_1 and τ_2 be two essential tracks of K which are minimal in the sense of Lemma 7.3 or Lemma 7.6. Then either $\tau_1 = \tau_2$ or $\tau_1 \cap \tau_2 = \emptyset$.

7.2. The proof of Proposition 6.3
First we briefly recall the notion of the Rips complex. See [12, Chapter III.Γ.3] for more details. Let (X, d) be a metric space, and pick $R > 0$. The *Rips complex* $P_R(X, d)$ is the geometric realization of the simplicial complex with vertex set X whose n-simplices are the $(n+1)$-element subsets $\{x_0, \ldots, x_n\} \subset X$ of diameter at most R.

Let d be the usual metric on $\mathbb{Z}$. Define a new metric $\bar{d}$ on $\mathbb{Z}$ by

$$\bar{d}(x, y) = \sup_{g \in U} d\big(g(x), g(y)\big).$$

Note that $(\mathbb{Z}, \bar{d})$ is quasi-isometric to $(\mathbb{Z}, d)$, and U acts on $(\mathbb{Z}, \bar{d})$ by isometries. Since $(\mathbb{Z}, \bar{d})$ is Gromov hyperbolic, the Rips complex $P_R(\mathbb{Z}, \bar{d})$ is contractible for some $R = R(L, A)$ (see [12, Proposition III.Γ.3.23]). Let K be the 2-skeleton of $P_R(\mathbb{Z}, \bar{d})$. Then K is simply connected, uniformly locally finite, and 2-ended.

We make K a piecewise Euclidean complex by identifying each 2-face with an equilateral triangle and identifying each edge with $[0, 1]$. Let $d_{\mathbb{E}}$ be the resulting length metric. There is an inclusion map $i : (\mathbb{Z}, d) \to (K, d_{\mathbb{E}})$ which is a quasi-isometry with quasi-isometry constants depending only on L and A.

Claim 7.8
There exist $D_1 = D_1(L, A)$ and a collection of disjoint essential tracks $\{\tau_i\}_{i\in I}$ of K such that the following hold.
(1) $\{\tau_i\}_{i\in I}$ is U-invariant.
(2) The diameter of each τ_i with respect to $d_{\mathbb{E}}$ is at most D_1.
(3) Each connected component of $K \setminus (\bigcup_{i\in I} \tau_i)$ has diameter at most D_1.

In the following proof, we denote the ball of radius D centered at x in K with respect to $d_{\mathbb{E}}$ by $B_{\mathbb{E}}(x, D)$. Let $\mathrm{diam}_{\mathbb{E}}$ be the diameter with respect to $d_{\mathbb{E}}$.

Proof of Claim 7.8
First we assume K does not have separating vertices. Since K is quasi-isometric to $\mathbb{Z}$, there exists $D = D(L, A)$ such that $K \setminus B_{\mathbb{E}}(x, D)$ has at least two unbounded components for each $x \in K$. Thus, every $(D + 1)$-ball contains an essential track with weight bounded above by $D' = D'(L, A)$. We put a U-invariant hyperbolic metric $d_{\mathbb{H}}$ on K as in Section 7.1. By Lemma 7.3, there exists an essential track $\tau \subset K$ of least complexity. Note that $\mathrm{diam}_{\mathbb{E}}(\tau) \le D'$, since the weight $w(\tau) \le D'$. Lemma 7.7 implies the U-orbits of τ give rise to a collection of disjoint essential tracks in K.

A collection of tracks $\{\tau_i\}_{i\in I}$ of K is *admissible* if the following hold.
(1) Each track in $\{\tau_i\}_{i\in I}$ is essential, and different tracks have empty intersection.
(2) No two tracks in $\{\tau_i\}_{i\in I}$ are parallel.
(3) The collection $\{\tau_i\}_{i\in I}$ is U-invariant.
(4) $\mathrm{diam}_{\mathbb{E}}(\tau_i) \le D'$ for each $i \in I$.
There exists a nonempty admissible collection of tracks by previous discussion. (In particular, (2) follows from the hyperbolicity of the metrics on the faces of K.)

Let $\{\tau_i\}_{i\in I}$ be a maximal admissible collection of tracks. Then this collection satisfies the above claim with $D_1 = 2D' + 5D$. To see this, let C be one connected component of $K \setminus (\bigcup_{i\in I} \tau_i)$. Since each track is essential and K is 2-ended, either $\mathrm{diam}_{\mathbb{E}}(C) < \infty$ and $\bar{C} \setminus C$ ($\bar{C}$ is the closure of C) is made of two tracks τ_1 and τ_2, or $\mathrm{diam}_{\mathbb{E}}(C) = \infty$ and $\bar{C} \setminus C$ is made of one track. Let us assume the former case is true. The latter case can be dealt with in a similar way. Let A be the maximal subcomplex of K which is contained in C. Then A is uniformly locally finite, and $C \setminus A$ is contained in the 1-neighborhood of $\tau_1 \cup \tau_2$.

Suppose that $\operatorname{diam}_{\mathbb{E}}(C) \geq 2D' + 5D$. Since $\operatorname{diam}_{\mathbb{E}}(\tau_i) \leq D'$ for $i = 1, 2$, there exists $x \in A$ such that $B_{\mathbb{E}}(x, 2D) \subset A$. Thus, A contains an essential track of X with its weight bounded above by D'. Let $\eta \subset A$ be an essential track of K which has the least complexity in the sense of Lemma 7.6. Then $w(\eta) \leq D'$; hence, $\operatorname{diam}_{\mathbb{E}}(\eta) \leq D'$. Moreover, by Lemma 7.7, for each $g \in \operatorname{Stab}(A) = \operatorname{Stab}(C)$, either $g \cdot \eta = \eta$ or $g \cdot \eta \cap \eta = \emptyset$. Thus, we can enlarge the original admissible collection of tracks by adding the U-orbits of η, which yields a contradiction.

The case when K has separating vertices is actually easier, since one can find essential tracks on the ϵ-sphere of each separating vertex. The rest of the proof is identical. $\square$

We now continue with the proof of Proposition 6.3.

Pick a regular neighborhood $N(\tau_i)$ for each τ_i such that the collection $\{N(\tau_i)\}_{i \in I}$ is disjoint and U-invariant. Then we define a map ϕ from K to a tree T by collapsing each component of $Y \setminus \bigcup_{i \in I} N(\tau_i)$ to a vertex and collapsing each $N(\tau_i)$, which is homeomorphic to $\tau_i \times (0,1)$, to the $(0,1)$-factor. It is easy to make ϕ equivariant under U, and by the above claim, ϕ is a quasi-isometry with quasi-isometry constants depending only on L and A. Note that T is actually a line, since τ is essential and K is 2-ended. Then Proposition 6.3 follows by considering the U-equivariant map $\phi \circ i : (\mathbb{Z}, d) \to T$.

Remark 7.9

If the action $U \curvearrowright \mathbb{Z}$ by (L, A)-quasi-isometries in Proposition 6.3 is not cobounded, then the resulting isometric action $\Lambda : U \curvearrowright \mathbb{Z}$ is also not cobounded; hence there are two possibilities:

(1) if U coarsely preserves the orientation of $\mathbb{Z}$, then Λ is trivial;

(2) otherwise Λ factors through a $\mathbb{Z}/2$-action by reflection.

Acknowledgments. The second author would like to thank Mladen Bestvina and Michah Sageev for many discussions related to atomic RAAGs, which planted the seeds for the cubulation result in this article. In particular, Proposition 6.3 originated in these discussions. We also thank the anonymous referees for carefully reading the article and offering many helpful comments.

Kleiner's work was partially supported by National Science Foundation grant DMS-1405899 and a Simons Fellowship.

References

[1] P. ABRAMENKO and K. S. BROWN, *Buildings: Theory and Applications*, Grad. Texts in Math. **248**, Springer, New York, 2008. MR 2439729. DOI 10.1007/978-0-387-78835-7. *(551, 553, 554, 555)*

[2] I. AGOL, *The virtual Haken conjecture*, Doc. Math. **18** (2013), 1045–1087. MR 3104553. *(538)*

[3] A. R. AHLIN, *The large scale geometry of products of trees*, Geom. Dedicata **92** (2002), 179–184. MR 1934017. DOI 10.1023/A:1019630514124. *(538, 542)*

[4] H. BASS, *The degree of polynomial growth of finitely generated nilpotent groups*, Proc. Lond. Math. Soc. (3) **25** (1972), 603–614. MR 0379672. DOI 10.1112/plms/s3-25.4.603. *(538)*

[5] J. A. BEHRSTOCK, T. JANUSZKIEWICZ, and W. D. NEUMANN, *Quasi-isometric classification of some high dimensional right-angled Artin groups*, Groups Geom. Dyn. **4** (2010), 681–692. MR 2727658. DOI 10.4171/GGD/100. *(538)*

[6] J. A. BEHRSTOCK and W. D. NEUMANN, *Quasi-isometric classification of graph manifold groups*, Duke Math. J. **141** (2008), 217–240. MR 2376814. DOI 10.1215/S0012-7094-08-14121-3. *(538, 542)*

[7] N. BERGERON and D. T. WISE, *A boundary criterion for cubulation*, Amer. J. Math. **134** (2012), 843–859. MR 2931226. DOI 10.1353/ajm.2012.0020. *(542)*

[8] M. BESTVINA and N. BRADY, *Morse theory and finiteness properties of groups*, Invent. Math. **129** (1997), 445–470. MR 1465330. DOI 10.1007/s002220050168. *(537)*

[9] M. BESTVINA, B. KLEINER, and M. SAGEEV, *The asymptotic geometry of right-angled Artin groups, I*, Geom. Topol. **12** (2008), 1653–1699. MR 2421136. DOI 10.2140/gt.2008.12.1653. *(538, 540, 542, 545, 566)*

[10] ———, *Quasiflats in* CAT(0) *2-complexes*, Algebr. Geom. Topol. **16** (2016), 2663–2676. MR 3572343. DOI 10.2140/agt.2016.16.2663. *(538, 539)*

[11] N. BRADY and J. MEIER, *Connectivity at infinity for right angled Artin groups*, Trans. Amer. Math. Soc. **353**, (2001), 117–132. MR 1675166. DOI 10.1090/S0002-9947-00-02506-X. *(537)*

[12] M. R. BRIDSON and A. HAEFLIGER, *Metric Spaces of Non-Positive Curvature*, Grundlehren Math. Wissen. **319**, Springer, Berlin, 1999. MR 1744486. DOI 10.1007/978-3-662-12494-9. *(542, 548, 549, 550, 556, 596, 597)*

[13] M. BURGER and S. MOZES, *Lattices in product of trees*, Inst. Hautes Études Sci. Publ. Math. **92** (2000), 151–194. MR 1839489. *(538, 539, 544)*

[14] P.-E. CAPRACE and M. SAGEEV, *Rank rigidity for* CAT(0) *cube complexes*, Geom. Funct. Anal. **21** (2011), 851–891. MR 2827012. DOI 10.1007/s00039-011-0126-7. *(538, 541, 550, 556, 589)*

[15] R. CHARNEY, *An introduction to right-angled Artin groups*, Geom. Dedicata **125** (2007), 141–158. MR 2322545. DOI 10.1007/s10711-007-9148-6. *(550)*

[16] ———, *Problems related to Artin groups*, preprint, http://people.brandeis.edu/~charney/papers/Artin_probs.pdf. *(568)*

[17] R. CHARNEY, J. CRISP, and K. VOGTMANN, *Automorphisms of 2-dimensional right-angled Artin groups*, Geom. Topol. **11** (2007), 2227–2264. MR 2372847. DOI 10.2140/gt.2007.11.2227. *(537)*

[18] R. CHARNEY and M. W. DAVIS, "Finite $K(\pi,1)$s for Artin groups" in *Prospects in Topology (Princeton, NJ, 1994)*, Ann. of Math. Stud. **138**, Princeton Univ. Press, Princeton, 1995, 110–124. MR 1368655. *(550)*

[19] ———, *The $K(\pi,1)$-problem for hyperplane complements associated to infinite reflection groups*, J. Amer. Math. Soc. **8** (1995), 597–627. MR 1303028. DOI 10.2307/2152924. *(539)*

[20] R. CHARNEY and M. FARBER, *Random groups arising as graph products*, Algebr. Geom. Topol **12** (2012), 979–995. MR 2928902. DOI 10.2140/agt.2012.12.979. *(540)*

[21] C. B. CROKE and B. KLEINER, *Spaces with nonpositive curvature and their ideal boundaries*, Topology **39** (2000), 549–556. MR 1746908. DOI 10.1016/S0040-9383(99)00016-6. *(537)*

[22] M. W. DAVIS, "Buildings are CAT(0)" in *Geometry and Cohomology in Group Theory (Durham, England, 1994)*, London Math. Soc. Lecture Note Ser. **252**, Cambridge Univ. Press, Cambridge, 1998, 108–123. MR 1709955. DOI 10.1017/CBO9780511666131.009. *(538, 539, 551, 552, 566)*

[23] M. B. DAY, *Finiteness of outer automorphism groups of random right-angled Artin groups*, Algebr. Geom. Topol. **12** (2012), 1553–1583. MR 2966695. DOI 10.2140/agt.2012.12.1553. *(540)*

[24] C. DROMS, *Isomorphisms of graph groups*, Proc. Amer. Math. Soc. **100** (1987), 407–408. MR 0891135. DOI 10.2307/2046419. *(537)*

[25] M. J. DUNWOODY, *The accessibility of finitely presented groups*, Invent. Math. **81** (1985), 449–457. MR 0807066. DOI 10.1007/BF01388581. *(538, 542, 594)*

[26] R. GITIK, M. MITRA, E. RIPS, and M. SAGEEV, *Widths of subgroups*, Trans. Amer. Math. Soc. **350** (1998), 321–329. MR 1389776. DOI 10.1090/S0002-9947-98-01792-9. *(542)*

[27] M. GROMOV, *Groups of polynomial growth and expanding maps*, Inst. Hautes Études Sci. Publ. Math. **53** (1981), 53–78. MR 0623534. *(538)*

[28] ———, "Hyperbolic manifolds, groups and actions" in *Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference (Stony Brook, N.Y., 1978)*, Ann. of Math. Stud. **97**, Princeton Univ. Press, Princeton, 1981, 183–213. MR 0624814. *(542)*

[29] ———, "Hyperbolic groups" in *Essays in Group Theory*, Math. Sci. Res. Inst. Publ. **8**, Springer, New York, 1987, 75–263. MR 0919829. DOI 10.1007/978-1-4613-9586-7_3. *(549)*

[30] M. F. HAGEN, *Cocompactly cubulated crystallographic groups*, J. Lond. Math. Soc. (2) **90** (2014), 140–166. MR 3245140. DOI 10.1112/jlms/jdu017. *(542)*

[31] M. F. HAGEN and P. PRZYTYCKI, *Cocompactly cubulated graph manifolds*, Israel J. Math. **207** (2015), 377–394. MR 3358051. DOI 10.1007/s11856-015-1177-5. *(538, 542)*

[32] F. HAGLUND, *Finite index subgroups of graph products*, Geom. Dedicata **135** (2008), 167–209. MR 2413337. DOI 10.1007/s10711-008-9270-0. *(543, 544)*

[33] F. HAGLUND and F. PAULIN, *Constructions arborescentes d'immeubles*, Math. Ann. **325** (2003), 137–164. MR 1957268. DOI 10.1007/s00208-002-0373-x. *(543)*

[34] F. HAGLUND and D. T. WISE, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), 1551–1620. MR 2377497. DOI 10.1007/s00039-007-0629-4. *(538, 553, 554)*

[35] G. C. HRUSKA and D. T. WISE, *Finiteness properties of cubulated groups*, Compos. Math. **150** (2014), 453–506. MR 3187627. DOI 10.1112/S0010437X13007112. *(556)*

[36] J. HUANG, *Quasi-isometric classification of right-angled Artin groups, I: The finite out case*, Geom. Topol. **21** (2017), 3467–3537. MR 3692971. DOI 10.2140/gt.2017.21.3467. *(538, 540, 541, 542, 551, 554, 593)*

[37] ———, *Top-dimensional quasiflats in* CAT(0) *cube complexes*, Geom. Topol. **21** (2017), 2281–2352. MR 3654109. DOI 10.2140/gt.2017.21.2281. *(538)*

[38] ———, *Commensurability of groups quasi-isometric to RAAGs*, preprint, arXiv:1603.08586v2 [math.GT]. *(544)*

[39] T. JANUSZKIEWICZ and J. ŚWIĄTKOWSKI, *Commensurability of graph products*, Algebr. Geom. Topol. **1** (2001), 587–603. MR 1875609. DOI 10.2140/agt.2001.1.587. *(543)*

[40] J. KAHN and V. MARKOVIC, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, Ann. of Math. (2) **175** (2012), 1127–1190. MR 2912704. DOI 10.4007/annals.2012.175.3.4. *(538, 542)*

[41] M. KAPOVICH, B. KLEINER, and B. LEEB, *Quasi-isometries and the de Rham decomposition*, Topology **37** (1998), 1193–1211. MR 1632904. DOI 10.1016/S0040-9383(97)00091-8. *(538, 542)*

[42] M. KAPOVICH and B. LEEB, *Quasi-isometries preserve the geometric decomposition of Haken manifolds*, Invent. Math. **128** (1997), 393–416. MR 1440310. DOI 10.1007/s002220050145. *(538)*

[43] M. KAPOVICH and J. J. MILLSON, *On representation varieties of Artin groups, projective arrangements and the fundamental groups of smooth complex algebraic varieties*, Inst. Hautes Études Sci. Publ. Math. **88** (1998), 5–95. MR 1733326. *(537)*

[44] S.-H. KIM and T. KOBERDA, *Embedability between right-angled Artin groups*, Geom. Topol. **17** (2013), 493–530. MR 3039768. DOI 10.2140/gt.2013.17.493. *(551)*

[45] B. KLEINER and 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 1447034. DOI 10.1016/S0764-4442(97)86981-9. *(542)*

[46] ———, *Groups quasi-isometric to symmetric spaces*, Comm. Anal. Geom. **9** (2001), 239–260. MR 1846203. DOI 10.4310/CAG.2001.v9.n2.a1. *(543)*

[47] M. R. LAURENCE, *A generating set for the automorphism group of a graph group*, J. Lond. Math. Soc. (2) **52** (1995), 318–334. MR 1356145. DOI 10.1112/jlms/52.2.318. *(537, 541)*

[48] B. LEEB, 3-*manifolds with(out) metrics of nonpositive curvature*, Invent. Math. **122** (1995), 277–289. MR 1358977. DOI 10.1007/BF01231445. *(542)*

[49] L. MOSHER, M. SAGEEV, and K. WHYTE, *Quasi-actions on trees, I: Bounded valence*, Ann. of Math. (2) **158** (2003), 115–164. MR 1998479. DOI 10.4007/annals.2003.158.115. *(538, 542, 543, 546, 591)*

[50] P. PANSU, *Métriques de Carnot-Carathéodory et quasiisométries des espaces symétriques de rang un*, Ann. of Math. (2) **129** (1989), 1–60. MR 0979599. DOI 10.2307/1971484. *(542)*

[51] P. PAPASOGLU and K. WHYTE, *Quasi-isometries between groups with infinitely many ends*, Comment. Math. Helv. **77** (2002), 133–144. MR 1898396. DOI 10.1007/s00014-002-8334-2. *(538)*

[52] M. RONAN, *Lectures on Buildings: Updated and Revised*, Univ. Chicago Press, Chicago, 2009. MR 2560094. *(551, 555)*

[53] M. SAGEEV, *Ends of group pairs and non-positively curved cube complexes*, Proc. Lond. Math. Soc. (3) **71** (1995), 585–617. MR 1347406. DOI 10.1112/plms/s3-71.3.585. *(538, 548)*

[54] ———, "CAT(0) cube complexes and groups" in *Geometric Group Theory*, IAS/Park City Math. Ser. **21**, Amer. Math. Soc., Providence, 2014, 7–54. MR 3329724. *(548)*

[55] G. P. SCOTT and G. A. SWARUP, *An algebraic annulus theorem*, Pacific J. Math. **196** (2000), 461–506. MR 1800588. *(595, 596)*

[56] H. SERVATIUS, *Automorphisms of graph groups*, J. Algebra **126** (1989), 34–60. MR 1023285. DOI 10.1016/0021-8693(89)90319-0. *(537, 541)*

[57] J. R. STALLINGS, *On torsion-free groups with infinitely many ends*, Ann. of Math. (2) **88** (1968), 312–334. MR 0228573. DOI 10.2307/1970577. *(538, 542)*

[58] D. SULLIVAN, "On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions" in *Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference (Stony Brook, N.Y., 1978)*, Ann. of Math. Stud. **97**, Princeton Univ. Press, Princeton, 1981, 465–496. MR 0624833. *(542)*

[59] P. TUKIA, *On quasiconformal groups*, J. Analyse Math. **46** (1986), 318–346. MR 0861709. DOI 10.1007/BF02796595. *(542)*

[60] K. WHYTE, *Coarse bundles*, preprint, arXiv:1006.3347v1 [math.GT]. *(538)*

[61] D. T. WISE, *Non-positively curved squared complexes: Aperiodic tilings and non-residually finite groups*, Ph.D. dissertation, Princeton University, Princeton, N.J., 1996. MR 2694733. *(538, 539, 544)*

[62] ———, *The structure of groups with a quasiconvex hierarchy*, preprint, 2011. *(538)*

Huang
McGill University, Department of Mathematics and Statistics, Montreal, Quebec, Canada; huangjingyin@gmail.com

Kleiner
Courant Institute, New York University, New York, New York, USA; bkleiner@cims.nyu.edu