Geometry & Topology

Volume 27 (2023)

Hyperbolic groups acting improperly

Daniel Groves
Jason Fox Manning

msp

Geometry & Topology 27:9 (2023) 3387–3460
DOI: 10.2140/gt.2023.27.3387
Published: 5 December 2023

Hyperbolic groups acting improperly

DANIEL GROVES
JASON FOX MANNING

We study hyperbolic groups acting on CAT(0) cube complexes. The first main result is a structural result about the Sageev construction, in which we relate quasiconvexity of hyperplane stabilizers with quasiconvexity of cell stabilizers. The second main result generalizes both Agol's Theorem on cubulated hyperbolic groups and Wise's Quasiconvex Hierarchy Theorem.

20F65, 57M05

1 Introduction

In recent years, CAT(0) cube complexes have played a central role in many spectacular advances, most notably in Agol's proof of the Virtual Haken and Virtual Fibering Theorems in [1]. The main result of [1] is that a hyperbolic group which acts properly and cocompactly on a CAT(0) cube complex is virtually special. A key ingredient in Agol's proof was the work of Wise from [36], particularly Wise's Quasiconvex Hierarchy Theorem [36, Theorem 13.3]. One of the two main results of the current paper is Theorem D, which provides a simultaneous generalization of Agol's Theorem and Wise's Theorem. So far this generalization has been applied by Duong [12] and Einstein and Groves [13]. At the end of the introduction in Section 1.2, we explain how Theorem D (together with Theorem A) simplifies the proof of the Virtual Haken and Virtual Fibering Theorems for hyperbolic 3–manifolds, requiring only a single immersed quasi-Fuchsian surface instead of a ubiquitous family.

Cube complexes in group theory arise via the construction of Sageev [32] which takes as input a group G and a collection of *codimension-one* subgroups of G and produces a CAT(0) cube complex X, equipped with an isometric G–action on X with no global fixed point. The other main result of the current paper is Theorem A, which establishes some fundamental properties about the Sageev construction.

Sageev's construction works in great generality. However, in order to get more information from the G–action on X, it is useful to add geometric hypotheses. For example, if G is a hyperbolic group and the codimension-one subgroups are quasiconvex, Sageev proved that the associated cube complex is G–cocompact [33, Theorem 3.1]. Achieving a *proper* action is harder (see Bergeron and Wise [5] and Hruska and Wise [24] for conditions which ensure properness).

Even an improper action $G \curvearrowright X$ gives a description of G as the fundamental group of a *complex of groups* in the sense of Bridson and Haefliger (see [8, III.$\mathcal{C}$] or Section 2 below). In this description, the underlying space is the quotient $G\backslash X$ and the local groups can be identified with cell stabilizers for the action.

Our first main result links the geometry of the hyperplane stabilizers with that of the cell stabilizers.

Theorem A *Let G be hyperbolic. The following conditions on a cocompact G–action on a CAT(0) cube complex are equivalent*:

1. *All hyperplane stabilizers are quasiconvex.*
2. *All vertex stabilizers are quasiconvex.*
3. *All cell stabilizers are quasiconvex.*

Intersections of quasiconvex subgroups are quasiconvex, and cell stabilizers are intersections of vertex stabilizers. Therefore, the equivalence of (2) and (3) is trivial. We prove the equivalence of (1) and (2).

We remark that we actually prove the direction (1) $\Longrightarrow$ (2) in the more general setting of arbitrary finitely generated groups where we assume the relevant subgroups are *strongly* *quasiconvex* in the sense of Tran [35]. Note that in this more general setting, (2) and (3) are still equivalent. See Section 3 for more details. In Section 3.6 we explain how Theorem A implies the following result.

Corollary B *Suppose that G is a hyperbolic group acting cocompactly on a CAT(0) cube complex X with quasiconvex hyperplane stabilizers. Then*

(1) *X is δ–hyperbolic for some δ*;
(2) *there exists a $k \geq 0$ such that the fixed-point set of any infinite subgroup of G intersects at most k distinct cells*; *and*
(3) *the action of G on X is acylindrical* (in the sense of Bowditch [6, page 284]).

Anthony Genevois explained to us how conclusion (2) implies acylindricity for actions on hyperbolic CAT(0) cube complexes (see Section 3.6). The condition in (2) is not implied by acylindricity since X is not assumed to be locally compact.

Without the conclusion of δ–hyperbolicity, a more general version of Corollary B holds just as for Theorem A. See Remark 3.31 for more details.

In Sageev's construction, the stabilizers in G of hyperplanes in the resulting cube complex are commensurable with the chosen codimension-one subgroups of G. Therefore, we have the following result.

Corollary C *Let G be a hyperbolic group and let $\mathcal{H} = \{H_1, \dots, H_k\}$ be a collection of quasiconvex codimension-one subgroups. Let X be a CAT(0) cube complex obtained by applying the Sageev construction to $\mathcal{H}$.*

(1) *The stabilizers of cells in X are quasiconvex in G. In particular*, *they are finitely presented.*
(2) *X is δ–hyperbolic for some δ.*
(3) *There exists a $k \geq 0$ such that the fixed-point set of any infinite subgroup of G intersects at most k distinct cells.*
(4) *The action of G on X is acylindrical.*

As far as we are aware, even the corollary of item (1) that the cell stabilizers are finitely generated in the above result is new. We remark that the fact that cell stabilizers are finitely presented implies that the description of G as the fundamental group of the complex of groups associated to $G \backslash X$ is a finite description.

Some of the most dramatic uses of CAT(0) cube complexes have come from Haglund and Wise's theory of *special* cube complexes [20]. A cube complex is *special* if it admits a locally isometric immersion into the Salvetti complex of a right-angled Artin group. A group G is *virtually special* if there is a finite-index subgroup $G_0 \leq G$ and a CAT(0) cube complex X such that G_0 acts freely and cubically on X and $G_0 \backslash X$ is a compact special cube complex. (For some authors the quotient is allowed to be noncompact but have finitely many hyperplanes.)

As shown in [20], virtually special hyperbolic groups have many remarkable properties, such as being residually finite, linear over $\mathbb{Z}$ and possessing very strong subgroup separability properties.

Agol [1] proved that if a hyperbolic group G acts properly and cocompactly on a CAT(0) cube complex then G is virtually special. It is this result that implies the virtual Haken conjecture, as well as the virtual fibering conjecture (in the compact case), and many other results.

One of the key ingredients of the proof of Agol's Theorem, and another of the most important theorems in the area is Wise's Quasiconvex Hierarchy Theorem [36, Theorem 13.3] — see also [3, Theorem 10.2] — which states that if a hyperbolic group G can be expressed as $A*_C$ (resp. $A *_C B$), where C is quasiconvex in G and A is (resp. A and B are) virtually special then G is virtually special. This theorem can be rephrased as saying that if a hyperbolic group acts cocompactly on a *one-dimensional CAT*(0) *cube complex* (otherwise known as a "tree") with virtually special and quasiconvex cell stabilizers, then G is virtually special.

Our second main result is a common generalization of Agol's Theorem and Wise's Quasiconvex Hierarchy Theorem.

Theorem D *Suppose that G is a hyperbolic group acting cocompactly on a CAT(0) cube complex X with quasiconvex and virtually special cell stabilizers. Then G is virtually special.*

By Corollary C, Theorem D has the following immediate consequence.

Corollary E *Suppose that G is a hyperbolic group and that $\mathcal{H} = \{H_1, \dots, H_k\}$ is a collection of quasiconvex codimension-one subgroups. If the vertex stabilizers of the G–action on a cube complex obtained by applying the Sageev construction to $\mathcal{H}$ are virtually special, then G is virtually special.*

Since finding proper actions of hyperbolic groups on CAT(0) cube complexes is much harder than finding cocompact actions, Theorem D is expected to be a powerful new tool for proving that hyperbolic groups are virtually special. As mentioned above, Theorem D was used in [12] to show that random groups in the square model at density $< \frac{1}{3}$ are virtually special. Theorem D (as well as Corollary 6.6 below) are also applied in [13] to provide a characterization of relatively hyperbolic groups with 2–sphere boundary in terms of actions on cube complexes.

Theorem A is one of the key ingredients of the proof of Theorem D. We now explain how Theorem D is a consequence of the above-mentioned results of Agol and Wise, along with Theorem A and the following result (proved in Section 6).

Theorem F *Suppose that the hyperbolic group G acts cocompactly on a CAT(0) cube complex X and that cell stabilizers are virtually special and quasiconvex. There exists a quotient $\overline{G} = G/K$ such that*

1. *the quotient $K\backslash X$ is a CAT(0) cube complex;*
2. *the group $\overline{G}$ is hyperbolic; and*
3. *the action of $\overline{G}$ on $K\backslash X$ is proper (and cocompact).*

Proof of Theorem D Consider the hyperbolic group G, acting on a CAT(0) cube complex X as in the statement of Theorem D. By Theorem F there exists a hyperbolic quotient $\overline{G} = G/K$ of G such that $K\backslash X$ is a CAT(0) cube complex, and the $\overline{G}$–action on $K\backslash X$ is proper and cocompact. Let $Z = K\backslash X$.

By Agol's Theorem [1, Theorem 1.1], there is a finite-index subgroup $\overline{G}_0$ of $\overline{G}$ such that $\overline{G}_0\backslash Z$ is special. Let G_0 be the preimage in G of $\overline{G}_0$. Clearly, the underlying space of $G_0\backslash X$ is the same as that of $\overline{G}_0\backslash Z$, and in particular all of the hyperplanes are two-sided and embedded.

We cut successively along these hyperplanes, applying the complex of groups version of the Seifert–Van Kampen Theorem [8, Example III.$\mathcal{C}$.3.11(5) and Exercise III.$\mathcal{C}$.3.12]. In this way, we obtain a hierarchy of G_0 with the following properties:

1. The edge groups are quasiconvex (since they are stabilizers of hyperplanes, which are quasiconvex by Theorem A).
2. The terminal groups are virtually special (since they are finite-index subgroups of the vertex stabilizers in G).

Therefore, G_0 admits a quasiconvex hierarchy terminating in virtually special groups, so G_0 is virtually special by Wise's Quasiconvex Hierarchy Theorem [36, Theorem 13.3] (see [3, Theorem 10.3] for a somewhat different account). Since G_0 is finite index in G, the group G is virtually special, as required. □

Remark 1.1 We thank one of the referees for pointing out that one can replace the use of the complex of groups Seifert–Van Kampen Theorem in the previous proof with the following argument: Once all of the hyperplanes in $G_0\backslash X$ are two-sided and embedded, lift to X and consider the trees dual to the hyperplanes. This gives a collection of G_0–trees. Order them in some way. If V stabilizes a vertex in the first tree, consider the V–action on the second tree. The stabilizers in V for this second action act on the third tree, and so on. In this way, a quasiconvex hierarchy for G_0 is obtained, and the proof finishes as above.

We now briefly outline the contents of this paper. In Section 2 we recall those parts of the theory of complexes of groups from [8] which we need. In Section 3, we prove Theorem A and Corollary B. The proof of Theorem A depends on a quasiconvexity criterion (Theorem A.3) which is proved separately in the appendix. We separate out Theorem A.3 and its proof both because it may be of independent interest and because the methods, unlike in the rest of the paper, are pure δ–hyperbolic geometry. In Section 4 we investigate conditions on a group G acting on a CAT(0) cube complex X and a normal subgroup $K \trianglelefteq G$ which imply that the quotient $K\backslash X$ is a CAT(0) cube complex. In Section 5 we translate these conditions into group-theoretic statements. In Section 6 we prove various results about Dehn filling (in particular, Theorem 6.5 and Corollary 6.6, which may be of independent interest) to see that the conditions from Section 5 are satisfied for certain subgroups K which arise as kernels of long Dehn filling maps. We use this to deduce Theorem F.

1.1 Notation and conventions

The notation $A \dot{<} B$ indicates that A is a finite-index subgroup of B; similarly, $A \dot{\triangleleft} B$ indicates A is a finite-index normal subgroup. We write conjugation as $a^x = xax^{-1}$, or sometimes as $\mathrm{Ad}(x)(a)$. For p an element of a G–set, we denote the G–orbit by $[\![p]\!]$.

1.2 Virtual Haken and fibering with a single surface

Let M be a closed hyperbolic 3–manifold, and let $\Gamma = \pi_1(M)$. Agol's proof that M is virtually Haken and virtually fibered in [1] proceeds via proving that Γ is virtually special. This relies on Bergeron and Wise's Theorem that Γ acts properly and cocompactly on a CAT(0) cube complex [5]. In turn, Bergeron and Wise rely on work of Kahn and Markovic [26], which provides a "ubiquitous"[1] family of immersed

[1]This terminology is from Cooper and Futer [10].

quasi-Fuchsian surfaces in M. That there is such an abundance of surfaces follows from the proofs in [26], but is not explicitly stated there.

Here we point out that the results in this paper show that the fact that Γ is virtually special follows from the existence of a *single* immersed quasi-Fuchsian surface in M. It is explained in [36] how virtual Haken and virtual fibering follow.

Theorem 1.2 *Suppose that M is a closed hyperbolic 3–manifold and that M contains an immersed quasi-Fuchsian surface. Then $\pi_1(M)$ is virtually special.*

Proof If M is nonorientable, we replace it by its orientation double cover. Let $\Gamma \cong \pi_1 M$ be a lattice in $\mathrm{Isom}^+(\mathbb{H}^3)$, so that $M \cong \Gamma\backslash\mathbb{H}^3$. We note that in this setting a subgroup $W < \Gamma$ is geometrically finite as a Kleinian group if and only if it is quasiconvex in Γ; see [27, Theorem 2] or [34, Theorem 1.1 and Proposition 1.3].

Let $H < \Gamma$ be the subgroup corresponding to the immersed quasi-Fuchsian surface. Since H is quasiconvex and codimension-one in Γ, we can apply the Sageev construction to obtain a cocompact action of Γ on a CAT(0) cube complex X with no global fixed point, and with hyperplane stabilizers conjugate to H. Theorem A implies that the vertex stabilizers for this action are quasiconvex in Γ. To apply Theorem D, we will show that the vertex stabilizers admit quasiconvex hierarchies and hence are virtually special.

Let $V < \Gamma$ be a vertex stabilizer. Since V is quasiconvex in Γ it is a geometrically finite subgroup of $\mathrm{Isom}^+(\mathbb{H}^3)$. As V has infinite index in Γ, it acts with infinite covolume on $\mathbb{H}^3$. An argument of Thurston shows that every finitely generated subgroup of V is also geometrically finite [30, Proposition 7.1].

Since Γ contains no parabolics, neither does V. Thus a small closed neighborhood N of the convex core of $H\backslash\mathbb{H}^3$ is a compact 3–manifold with nonempty boundary, and hence is irreducible in the sense that every embedded 2–sphere bounds a ball [29, Propositions 2.36 and 3.1]. A compact irreducible 3–manifold with nonempty boundary is Haken; see [22, Chapter 6; 25, Chapter III]. In particular it has a Haken hierarchy [25, IV.12]. This topological hierarchy of N gives a group-theoretic hierarchy of V. The edge groups in the hierarchy are finitely generated. The previously mentioned argument of Thurston then implies that the edge groups are geometrically finite and hence quasiconvex in Γ. In particular, this is a quasiconvex hierarchy, and we may apply Wise's Quasiconvex Hierarchy Theorem to conclude that V is virtually special.

Since all vertex stabilizers of the action $\Gamma \curvearrowright X$ are quasiconvex and virtually special, we may apply Theorem D to conclude that Γ is itself virtually special. $\square$

Acknowledgements

We thank Richard Webb for suggesting that the direction (1) $\Longrightarrow$ (2) of Theorem A might hold in a more general setting than that of hyperbolic groups.

Thanks to Anthony Genevois for pointing out that his work allows us to deduce acylindricity (Corollary B(3)) from Corollary B(2).

We also thank Alessandro Sisto for pointing out the applicability of a "Greendlinger Lemma" type result in our joint work; see [19, Lemma 2.26]. This lemma inspired Theorem 6.7 in the current paper, which simplified the proof of Theorem 6.5.

Finally, we thank the referees whose careful readings and comments have led to substantial improvements in the exposition of this paper.

Groves was partially supported by the Simons Foundation, grant 342049, and the National Science Foundation, grants DMS-1507067 and DMS-1904913. Manning was also partially supported by the Simons Foundation, grant 524176, and the National Science Foundation, grant DMS-1462263.

2 Complexes of groups

In this section we give a brief account of those parts of the theory of complexes of groups which we need. Much more detail can be found in Bridson and Haefliger [8, III.$\mathcal{C}$].

2.1 Paths and homotopies in a category

The definitions here are mainly taken from [8, III.$\mathcal{C}$.A], though our notation is slightly different.

Let $\mathcal{C}$ be a category. For an arrow a of $\mathcal{C}$, we denote its source by $i(a)$ and its target by $t(a)$. An *oriented edge* of $\mathcal{C}$ is a symbol a^+ or a^-, where a is an arrow of $\mathcal{C}$. The source and target of an oriented edge are defined by

$$i(a^-) = i(a), \quad t(a^-) = t(a) \qquad \text{and} \qquad i(a^+) = t(a), \quad t(a^+) = i(a).$$

(We caution readers that this may be the opposite of what they expect. The signs are chosen so that concatenation of $+$ edges is homotopic to composition of the corresponding arrows; see Definition 2.1.)

We now define $\mathcal{C}$–paths. A *$\mathcal{C}$–path p of length* 0 is an object v of $\mathcal{C}$ with $i(p) = t(p) = v$. For $j > 0$, a *$\mathcal{C}$–path of length j* is a list $p = e_1 \cdot e_2 \cdots e_j$ where for each i we have $t(e_i) = i(e_{i+1})$. We have $i(p) = i(e_1)$ and $t(p) = t(e_j)$.

If p is a $\mathcal{C}$–path of length j, q is a $\mathcal{C}$–path of length k, and $t(p) = i(q)$, then the concatenation $p \cdot q$ is a $\mathcal{C}$–path of length $j + k$ with $i(p \cdot q) = i(p)$ and $t(p \cdot q) = t(q)$.[2]

The category $\mathcal{C}$ is *connected* if for any two objects v_0, v_1 in $\mathcal{C}$ there is a $\mathcal{C}$–path p with $i(p) = v_0$ and $t(p) = v_1$.

If p is a $\mathcal{C}$–path, then p is *nonbacktracking* if it contains no subpath of the form $a^+ \cdot a^-$ or $a^- \cdot a^+$.

Definition 2.1 *Homotopies* of $\mathcal{C}$–paths (see [8, III.$\mathcal{C}$.A.11]) are generated by the following *elementary homotopies*, valid whenever both sides are paths:

(1) $p \cdot a^+ \cdot a^- \cdot q \simeq p \cdot q$ or $p \cdot a^- \cdot a^+ \cdot q \simeq p \cdot q$;

(2) $p \cdot a^+ \cdot b^+ \cdot q \simeq p \cdot (ab)^+ \cdot q$ or $p \cdot b^- \cdot a^- \cdot q \simeq p \cdot (ab)^- \cdot q$ (here and below we write ab for the composition $a \circ b$); and

(3) $p \cdot \mathbb{1}_v^{\pm} \cdot q \simeq p \cdot q$ (where $\mathbb{1}_v$ is an identity arrow).

Any category has a *nerve* which is a simplicial complex whose 0–cells are the objects of $\mathcal{C}$, with 1–cells corresponding to arrows, 2–cells to composable pairs of arrows, and so on. The $\mathcal{C}$–paths we have just defined give edge-paths and the elementary homotopies correspond to simplicial homotopies in this complex.

2.2 Small categories without loops (scwols)

By a *scwol* (small category without loops) we mean a small category in which for every object v, the set of arrows from v to itself contains only the unit $\mathbb{1}_v$, and this unit $\mathbb{1}_v$ cannot be written as a composition of other arrows. An arrow is *trivial* if it is equal to $\mathbb{1}_v$ for some object v. We sometimes conflate v and $\mathbb{1}_v$. A (*nondegenerate*) *morphism* of scwols $f : \mathcal{A} \to \mathcal{B}$ is a functor which induces, for each object v of $\mathcal{A}$, a bijection between the arrows $\{a \mid i(a) = v\}$ and the arrows $\{a \mid i(a) = f(v)\}$.

[2] For purposes of concatenation, a $\mathcal{C}$–path of length 0 is regarded as an empty list.

Definition 2.2 (simple scwol, scwolification) A scwol in which there is at most one arrow with a given source and target will be called a *simple scwol*. Any small category $\mathcal{C}$ has a canonical quotient category $\mathrm{scwol}(\mathcal{C})$, which is a simple scwol. The objects of $\mathrm{scwol}(\mathcal{C})$ are equivalence classes of objects of $\mathcal{C}$, where $v \sim w$ if there are arrows a and b such that $i(a) = t(b) = v$ and $i(b) = t(a) = w$. Similarly, the arrows of $\mathrm{scwol}(\mathcal{C})$ are equivalence classes of arrows of $\mathcal{C}$, where $a \sim b$ whenever $i(a) \sim i(b)$ and $t(a) \sim t(b)$. We may refer to $\mathrm{scwol}(\mathcal{C})$ as the *scwolification of* $\mathcal{C}$. The map $\mathcal{C} \to \mathrm{scwol}(\mathcal{C})$ taking each object and arrow to its equivalence class will be called the *scwolification functor*.

Remark 2.3 The procedure of scwolification is natural. In particular, a group action on a small category $\mathcal{C}$ descends to an action on $\mathrm{scwol}(\mathcal{C})$.

A key example of a scwol is the (opposite) poset of cells of a simplicial or cubical complex, with arrows from each cell to all its faces. If two faces of some cell are glued together one obtains a nonsimple scwol.

Definition 2.4 [8, III.$\mathcal{C}$.1.3] A scwol $\mathcal{A}$ has a (*geometric*) *realization* which is a simplicial complex whose 0–cells are the objects of $\mathcal{A}$, with 1–cells corresponding to nontrivial arrows, 2–cells to composable pairs of such arrows, and so on.

The realization of $\mathcal{A}$ is naturally a subcomplex of the nerve of $\mathcal{A}$. Although the nerve is necessarily infinite-dimensional, the realization of a scwol has dimension equal to the length of the longest chain of nontrivial composable arrows. Every scwol which appears in the current paper has finite-dimensional realization.

Definition 2.5 If A is the realization of a scwol $\mathcal{A}$, then there is a canonical correspondence between combinatorial paths in the 1–skeleton of A and $\mathcal{A}$–paths without trivial arrows. If p is a combinatorial path in $A^{(1)}$, and q the corresponding $\mathcal{A}$–path, we say that p is the *realization* of q, and q is the *idealization* of p.

2.3 Complexes of groups

Definition 2.6 [8, III.$\mathcal{C}$.2.1] Let $\mathcal{A}$ be a scwol. A *complex of groups* $H(\mathcal{A})$ consists of

(1) for each object σ of $\mathcal{A}$, a *local group* (also called a *cell group*) H_σ;

(2) for each arrow a of $\mathcal{A}$, an injective group homomorphism $\psi_a \colon H_{i(a)} \to H_{t(a)}$ (if a is a trivial arrow, we require ψ_a to be the identity map);

(3) for each pair of composable arrows a and b with composition ab, a *twisting element* $z(a,b) \in H_{t(a)}$ (if either a or b is trivial, $z(a,b) = 1$).[3]

These data satisfy the following conditions (continuing to write ab for $a \circ b$) whenever all written compositions of arrows are defined:

(1) **Compatibility** $\mathrm{Ad}(z(a,b))\psi_{ab} = \psi_a \psi_b$.[4]

(2) **Cocycle** $\psi_a(z(b,c))z(a,bc) = z(a,b)z(ab,c)$.

The cocycle condition above applies to any arrangement of arrows of the form

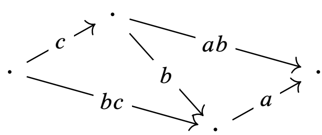

Definition 2.7 (the complex of groups coming from an action) Suppose G acts on a scwol $\mathcal{X}$ in such a way that any $g \in G$ fixing an object fixes every arrow from that object. Suppose further that $\mathcal{Y} = G \backslash \mathcal{X}$ is a scwol. We obtain a complex of groups $G(\mathcal{Y})$ once we have [8, III.$\mathcal{C}$.2.9]:

(1) For each object v of $\mathcal{Y}$, a choice of a lift $\tilde{v}$ to $\mathcal{X}$; this lift also determines lifts $\tilde{a}$ of all arrows a with $i(a) = v$.

(2) For each arrow a, a choice of $h_a \in G$ such that $t(h_a(\tilde{a})) = \widetilde{t(a)}$. (When a is a trivial arrow, we always take $h_a = 1$.)

Given these choices, one defines

(1) G_v as the stabilizer of $\tilde{v}$,

(2) $\psi_a = \mathrm{Ad}(h_a)|_{G_{i(a)}}$,

(3) $z(a,b) = h_a h_b h_{ab}^{-1}$.

The complex of groups $G(\mathcal{Y})$ can be used to recover the group G. There are two different ways of doing this. The first is explained in [8, III.$\mathcal{C}$.3.7], and involves $G(\mathcal{Y})$–*paths*. The second way is from [8, III.$\mathcal{C}$.A], and is the way that we proceed. The advantage to this second way, which uses categories and coverings of categories, is that lifting paths to covers is a canonical procedure (as with usual covering theory).

[3] In [8] the notation $g_{a,b}$ is used instead of $z(a,b)$.

[4] Recall $\mathrm{Ad}(z)(x) = zxz^{-1}$.

2.4 Fundamental groups and coverings of categories

In Definition 2.1 we defined homotopy of $\mathcal{C}$–paths, where $\mathcal{C}$ is a category.

Definition 2.8 Given a category $\mathcal{C}$ and an object v_0 of $\mathcal{C}$, the *fundamental group of $\mathcal{C}$ based at* v_0, denoted by $\pi_1(\mathcal{C}, v_0)$, is the set of homotopy classes of $\mathcal{C}$–loops based at v_0, with operation induced by concatenation of $\mathcal{C}$–paths.

Definition 2.9 [8, III.$\mathcal{C}$.A.15] Let $\mathcal{C}$ be a connected category. A functor $f\colon \mathcal{C}' \to \mathcal{C}$ is a *covering* if for each object σ' of $\mathcal{C}'$, the restriction of f to the collection of arrows that have σ' as their initial (resp. terminal) object is a bijection onto the set of arrows which have $f(\sigma')$ as their initial (resp. terminal) object.

The *universal cover* $\widetilde{\mathcal{C}}$ of a connected category $\mathcal{C}$ is described in [8, III.$\mathcal{C}$.A.19]: Fix a base vertex v_0 of $\mathcal{C}$, and define $\mathrm{Obj}(\widetilde{\mathcal{C}})$ to be the set of homotopy classes of $\mathcal{C}$–paths starting at v_0. If $[c]$ is a homotopy class of path, and α is an arrow from $t(c)$, then there is an arrow $\tilde{\alpha}$ of $\widetilde{\mathcal{C}}$ from $[c]$ to $[c \cdot \alpha^-]$. The projection $\pi\colon \widetilde{\mathcal{C}} \to \mathcal{C}$ sets $\pi([p]) = t(p)$ and if $\tilde{\alpha}$ is the arrow described above then $\pi(\tilde{\alpha}) = \alpha$. The fundamental group $\pi_1(\mathcal{C}, v_0)$ acts on $\widetilde{\mathcal{C}}$ by preconcatenation.

The theory of coverings of categories is entirely analogous to ordinary covering theory. In fact it is a special case, as the coverings of a connected category $\mathcal{C}$ correspond bijectively to the covering spaces of its nerve.

We record the following observation.

Lemma 2.10 *Let $\phi\colon \widetilde{\mathcal{C}} \to \mathcal{C}$ be a covering of categories, and suppose $\phi(\tilde{v}) = v$ for objects v of $\mathcal{C}$ and $\tilde{v}$ of $\widetilde{\mathcal{C}}$. Any $\mathcal{C}$–path p with $i(p) = v$ has a unique lift to a $\widetilde{\mathcal{C}}$–path $\tilde{p}$ with $i(\tilde{p}) = \tilde{v}$. Moreover any elementary homotopy from p to a path p' gives a unique elementary homotopy of $\tilde{p}$ to a lift $\tilde{p}'$ of p' with the same endpoints as $\tilde{p}$.*

2.5 The category associated to a complex of groups

Any complex of groups $G(\mathcal{Y})$ has an associated category $CG(\mathcal{Y})$.

Definition 2.11 [8, III.$\mathcal{C}$.2.8] The objects of $CG(\mathcal{Y})$ are the objects of the scwol $\mathcal{Y}$. Arrows of $CG(\mathcal{Y})$ are pairs (g, a) such that a is an arrow of $\mathcal{Y}$ and $g \in G_{t(a)}$. Composition is defined by $(g, a) \circ (h, b) = (g\psi_a(h) z(a, b), ab)$.

Recall that if a is a trivial arrow then ψ_a is the identity homomorphism and $z(a, x)$ and $z(x, a)$ are always trivial.

Remark 2.12 The map $CG(\mathcal{Y}) \to \mathcal{Y}$ given by $(g,a) \to a$ is the scwolification functor (see Definition 2.2), and is a bijection on objects. This functor has an obvious section $a \mapsto (1,a)$. If there are nontrivial twisting elements, this is not a functor, but it does allow $\mathcal{Y}$–paths to be "unscwolified" to $CG(\mathcal{Y})$–paths. In Definition 2.22, we explain how to go back and forth between paths in covers of $CG(\mathcal{Y})$ and their associated scwols.

Theorem 2.13 [8, III.$\mathcal{C}$.3.15 and text before III.$\mathcal{C}$.A.13] *Suppose that the group G acts on the simply connected complex X, giving rise to an action of G on the scwol $\mathcal{X}$ as in Definition 2.7, and that v_0 is an object in $\mathcal{Y} = G\backslash\mathcal{X}$. Let $CG(\mathcal{Y})$ be the category associated to $G(\mathcal{Y})$. Then there is an isomorphism from $\pi_1(CG(\mathcal{Y}), v_0)$ to G taking any loop of the form $(g_1,a_1)^{\epsilon_1}\cdots(g_n,a_1)^{\epsilon_1}$ to the product $(g_1h_{a_1})^{\epsilon_1}\cdots(g_nh_{a_1})^{\epsilon_n}$.*

The exponents ϵ_i in the statement are taken from the set $\{+,-\}$. We use the mild abuse of notation that if g is a group element then $g^+ = g$ and $g^- = g^{-1}$.

Definition 2.14 Let a be a nontrivial arrow of $\mathcal{Y}$. The arrow $(1,a)$ of $CG(\mathcal{Y})$ is called a *scwol arrow*. Let $g \in G_v$ where v is a vertex of the scwol $\mathcal{Y}$. The arrow $(g,\mathbb{1}_v)$ is called a *group arrow at v*, or just a *group arrow* if v is unimportant.

In later sections we abuse notation and refer to the edge $(g,\mathbb{1}_v)^+$ (for a group arrow $(g,\mathbb{1}_v)$) as "(g,v)" or even just "g". We also blur the difference between the scwol arrow $(1,a)$ and the $\mathcal{Y}$–arrow a, and often refer to the scwol arrow by "a". We also blur the distinction between the $CG(\mathcal{Y})$–edge $(1,a)^{\pm}$ and the $\mathcal{Y}$–edge $a^{\pm}$.

Definition 2.15 Let $\mathcal{C} \to CG(\mathcal{Y})$ be a covering of categories. We say that an arrow is *labeled by* (g,a) if its image in $CG(\mathcal{Y})$ is (g,a). An arrow of $\mathcal{C}$ is said to be a *scwol* (*resp. group*) *arrow* if its label is a scwol (resp. group) arrow of $CG(\mathcal{Y})$.

Lemma 2.16 *If $\mathcal{C} \to CG(\mathcal{Y})$ is any cover, then every $\mathcal{C}$–path is homotopic to a concatenation of group and scwol arrows.*

Proof Observe that any $CG(\mathcal{Y})$–arrow (g,a) is a composition of a group arrow and a scwol arrow; $(g,a) = (g,t(a)) \circ (1,a)$. This gives a homotopy in $CG(\mathcal{Y})$ to a path of the desired form. Lemma 2.10 says that the homotopy lifts. □

As described at the end of the last subsection, a choice of base vertex v_0 determines a universal covering map $\widetilde{CG(\mathcal{Y})} \to CG(\mathcal{Y})$ sending a homotopy class of path $[p]$ to its terminal vertex $t(p)$, and the arrow from $[c]$ to $[c\cdot(g,a)^-]$ to the arrow (g,a) of $CG(\mathcal{Y})$.

The group $\pi_1(CG(\mathcal{Y}), v_0) \cong G$ acts on the universal cover $\widetilde{CG(\mathcal{Y})}$ by preconcatenation of paths. The quotient by this action is $CG(\mathcal{Y})$. If $H < \pi_1(CG(\mathcal{Y}), v_0)$ is any subgroup, then $H\backslash\widetilde{CG(\mathcal{Y})}$ is an intermediate cover of categories. Every connected cover of $CG(\mathcal{Y})$ is of this form. Indeed, covers of $CG(\mathcal{Y})$ correspond to coverings of the nerve N of $CG(\mathcal{Y})$. Since $\pi_1(CG(\mathcal{Y}, v_0))$ is canonically isomorphic to $\pi_1(N, v_0)$, connected covers of $CG(\mathcal{Y})$ are all of the form $H\backslash\widetilde{CG(\mathcal{Y})}$ for $H < \pi_1(CG(\mathcal{Y}), v_0)$. The isomorphism $\pi_1(CG(\mathcal{Y}), v_0) \cong G$ from Theorem 2.13 allows us to identify such H as subgroups of G.

Proposition 2.17 *Suppose that $CG(\mathcal{Y})$ arises from an action of G on a simply connected scwol $\mathcal{X}$ via Definitions 2.7 and 2.11. Let $\phi\colon \pi_1(CG(\mathcal{Y}), v_0) \to G$ be the isomorphism from Theorem 2.13. There is a ϕ–equivariant functor $\Theta\colon \widetilde{CG(\mathcal{Y})} \to \mathcal{X}$ which factors through an isomorphism of categories $\widehat{\Theta}\colon \mathrm{scwol}(\widetilde{CG(\mathcal{Y})}) \to \mathcal{X}$.*

Proof sketch Consider an arrow x labeled by (g,a) from $[\sigma_1]$ to $[\sigma_2]$ in $\widetilde{CG(\mathcal{Y})}$. We may suppose

$$\sigma_1 = (g_1,a_1)^{\epsilon_1}\cdots(g_n,a_n)^{\epsilon_n} \quad\text{and}\quad \sigma_2 = (g_1,a_1)^{\epsilon_1}\cdots(g_n,a_n)^{\epsilon_n}\cdot(g,a)^{-}.$$

We define

$$\Theta(x) = \prod_{i=1}^{n}(g_i h_{a_i})^{\epsilon_i}\tilde{a}$$

Examining the elementary homotopies, it is not hard to see that $\Theta(x)$ is well defined.

The map ϕ sends the homotopy class of the loop $(h_1,b_1)^{\delta_1}\cdots(h_k,b_k)^{\delta_k}$ to the group element

$$(h_1h_{b_1})^{\delta_1}\cdots(h_kh_{b_k})^{\delta_k}.$$

Since $\pi_1(CG(\mathcal{Y}), v_0)$ acts by preconcatenation of paths, Θ is clearly ϕ–equivariant.

To see that Θ is a functor, suppose that $x = yz$ in $\widetilde{CG(\mathcal{Y})}$, where x, y and z are labeled by (g,a), (h,b) and (k,c), respectively, and $i(x) = i(z) = [(g_1,a_1)^{\epsilon_1}\cdots(g_n,a_n)^{\epsilon_n}]$. Letting $p = \prod_{i=1}^{n}(g_ih_{a_i})^{\epsilon_i}$, we have $\Theta(x) = p\tilde{a}$, $\Theta(y) = ph_c^{-1}k^{-1}\tilde{b}$ and $\Theta(z) = p\tilde{c}$, and it is easily checked that $\Theta(y)\Theta(z) = \Theta(x)$.

Any two objects of $\widetilde{CG(\mathcal{Y})}$ identified under the scwolification map are separated by a group arrow. If x is a group arrow then $\Theta(x)$ is a trivial arrow, so Θ factors through a functor

$$\widehat{\Theta}\colon \mathrm{scwol}(\widetilde{CG(\mathcal{Y})}) \to \mathcal{X}.$$

It remains to show that $\widehat{\Theta}$ is an isomorphism of scwols. The fact that the homomorphism from Theorem 2.13 is surjective for any $v_0 \in \mathcal{Y}$ implies that $\widehat{\Theta}$ is also surjective, so the only difficult point is to see injectivity.

Assuming that $\Theta([\sigma]) = \Theta([\tau])$, we consider the images $\hat{\sigma}$ and $\hat{\tau}$ of the *paths* σ and τ, respectively. Since $\mathcal{X}$ is assumed to be simply connected, there is a sequence of elementary homotopies in $\mathcal{X}$ taking $\hat{\sigma}$ to $\hat{\tau}$. It can be shown that these homotopies all lift (nonuniquely) to homotopies in $\widetilde{CG(\mathcal{Y})}$, so we may assume that $\hat{\sigma} = \hat{\tau}$. If $\hat{\sigma} = \hat{\tau}$ is a degenerate path at $\tilde{v}_0$, then it is clear that σ and τ are separated by a group arrow. Otherwise, after a further homotopy in $\widetilde{CG(\mathcal{Y})}$, we can assume that

$$\sigma = (g_1, a_1)^{\epsilon_1} \cdots (g_n, a_n)^{\epsilon_n}, \quad \tau = (h_1, a_1)^{\epsilon_1} \cdots (h_n, a_n)^{\epsilon_n},$$

where the signs of the ϵ_i alternate and each arrow a_i is nontrivial. If k is the smallest index for which $g_k \neq h_k$ and $k < n$, one can find a short sequence of elementary homotopies taking σ to another path σ' which agrees with τ for the first k edges. If $k = n$, then there is a slightly shorter sequence of elementary homotopies taking σ to $\tau \cdot x$ for some group arrow x.

To sum up, if $\Theta([\sigma]) = \Theta([\tau])$, then there is a group arrow from $[\sigma]$ to $[\tau]$. It follows that the map $\widehat{\Theta}\colon \mathrm{scwol}(\widetilde{CG(\mathcal{Y})}) \to \mathcal{X}$ is injective, and hence an isomorphism. □

The map Θ also passes to quotients by subgroups of G:

Corollary 2.18 *Suppose that $CG(\mathcal{Y})$ arises from an action of G on a simply connected scwol $\mathcal{X}$ via Definitions 2.7 and 2.11, and that $\mathcal{Y} = G\backslash\mathcal{X}$ is simple. If $H < G$, the scwolification of $H\backslash\widetilde{CG(\mathcal{Y})}$ is canonically isomorphic to $H\backslash\mathcal{X}$.*

Proof For any small category $\mathcal{C}$, acted on by a group H, there is a canonical surjective functor $H\backslash \mathrm{scwol}\,\mathcal{C} \to \mathrm{scwol}(H\backslash\mathcal{C})$. This is an isomorphism if and only if $H\backslash \mathrm{scwol}\,\mathcal{C}$ is already a simple scwol. Since $\mathcal{Y} = G\backslash\mathcal{X}$ is a simple scwol, the intermediate quotient $H\backslash\mathcal{X}$ is also a simple scwol.

The naturality of scwolification (Remark 2.3) and the equivariance of Θ together mean that the map $\widehat{\Theta}$ from Proposition 2.17 is also equivariant, and so we get an isomorphism

$$\widehat{\Theta}_H : H\backslash \mathrm{scwol}(\widetilde{CG(\mathcal{Y})}) \to H\backslash\mathcal{X}.$$

Thus $\mathrm{scwol}(H\backslash\widetilde{CG(\mathcal{Y})}) = H\backslash \mathrm{scwol}(\widetilde{CG(\mathcal{Y})})$ is isomorphic to $H\backslash\mathcal{X}$. □

Remark 2.19 Although Corollary 2.18 does not explicitly appear in Bridson and Haefliger [8], it can be derived from results there, as we outline briefly in this remark.

Suppose $\mathcal{C}' = H\backslash\widetilde{CG(\mathcal{Y})} \to CG(\mathcal{Y})$ is a cover. According to [8, Proposition III.$\mathcal{C}$.A.24] there is an associated (category of a) complex of groups $CG(\mathcal{Y}')$, which is a subcategory of $\mathcal{C}'$ such that the inclusion is an equivalence $CG(\mathcal{Y}') \to \mathcal{C}'$. The construction of $CG(\mathcal{Y}')$ from $\mathcal{C}'$ is as in [8, Proposition III.$\mathcal{C}$.A.4], which uses the same equivalence relation as the definition of the scwolification functor as in Definition 2.2. From the construction of $CG(\mathcal{Y}')$, and the assumption that $\mathcal{Y}$ is a simple scwol, it follows that $\mathcal{Y}'$ is isomorphic to $\mathrm{scwol}(\mathcal{C}')$. Further, from the correspondence between coverings and subgroups (both for categories and also for complexes of groups), and the identification of G with $\pi_1(CG(\mathcal{Y}))$ as in Theorem 2.13, it follows that $\mathcal{Y}'$ is isomorphic to $H\backslash\mathcal{X}$, as required.

Rather than fully develop this approach, we chose to sketch a more direct approach using Proposition 2.17 and the naturality of the scwolification functor.

Definition 2.20 We denote the scwolification functor from $H\backslash\widetilde{CG(\mathcal{Y})}$ to $H\backslash\mathcal{X}$ by Θ_H. If $H = \{1\}$, we just write Θ, as in Proposition 2.17.

We observe the following.

Lemma 2.21 *Suppose that $CG(\mathcal{Y})$ arises from an action of G on a simply connected scwol $\mathcal{X}$ via Definitions 2.7 and 2.11, and that $\mathcal{Y} = G\backslash\mathcal{X}$ is simple.*

Let $\mathfrak{o}$ be an object of $\mathcal{X}$. Then $\mathrm{Stab}_G(\mathfrak{o})$ acts freely and transitively on $\Theta^{-1}(\mathfrak{o})$.

Definition 2.22 Let $K < G \cong \pi_1(CG(\mathcal{Y}), v_0)$, let $\mathcal{C}_K = K\backslash\widetilde{CG(\mathcal{Y})}$ and let $\mathcal{Z} = K\backslash\mathcal{X}$. Since $\Theta_K : \mathcal{C}_K \to \mathrm{scwol}(\mathcal{C}_K) = \mathcal{Z}$ is a functor, it gives a way to turn a $\mathcal{C}_K$–path p into a $\mathcal{Z}$–path p'. Deleting all the trivial arrows from p' produces a $\mathcal{Z}$–path, which we call the *scwolification* of p. Abusing the notation slightly, we denote the scwolification of p by $\Theta_K(p)$.

Conversely, if σ is a $\mathcal{Z}$–path, then any $\mathcal{C}_K$–path $\hat{\sigma}$ such that $\Theta_K(\hat{\sigma}) = \sigma$ is called an *unscwolification* of σ. The unscwolification is highly nonunique, but always exists.

The following can be deduced by examining the elementary homotopies.

Lemma 2.23 *Scwolifications of homotopic paths are homotopic.*

Given a $CG(\mathcal{Y})$–path p we can lift it to a $\mathcal{C}_K$–path $\hat{p}$, and then scwolify $\hat{p}$ to the $\mathcal{Z}$–path $\Theta_K(\hat{p})$.

Lemma 2.24 *Let p be a $CG(\mathcal{Y})$–loop at v and let $\hat{p}$ be a lift to $\mathcal{C}_K$. If $\Theta_K(\hat{p})$ is a loop, then there is a group arrow labeled by an element of G_v joining the endpoints of $\hat{p}$.*

2.6 The complex of groups coming from an action on a cube complex

Let X be a CAT(0) cube complex, and suppose that G acts on X by cubical automorphisms. The quotient $G\backslash X$ may or may not be a cube complex, depending on whether the groups $G_\sigma = \{g \mid g\sigma = \sigma\}$ and $\{g \mid gx = x \text{ for all } x \in \sigma\}$ agree for all cells σ.

Another way to phrase this issue is to note that, if $\mathcal{X}_0$ is the scwol of cells of X, then G acts by morphisms on $\mathcal{X}_0$, but the quotient map $\mathcal{X}_0 \to G\backslash\mathcal{X}_0$ may not be a nondegenerate morphism of scwols, since some isometry of X may fix the center of some cube, but permute faces of that cube. In order to obtain a complex of groups structure on G from the action $G \curvearrowright X$, we need a scwol quotient, so we replace $\mathcal{X}_0$ with $\mathcal{X}$, the scwol of cells of the first barycentric subdivision of X:

Definition 2.25 If W is a cube complex, the *idealization of* W is the scwol of cells of the first barycentric subdivision of W.

If every cube of the cube complex W embeds in W, there is another way of thinking of the idealization $\mathcal{W}$. Namely, the objects of $\mathcal{W}$ are in one-to-one correspondence with nonempty nested chains of cubes of W and there is at most one arrow in $\mathcal{W}$ between two objects: if c_1 contains c_2 as a subchain, there is an arrow from c_1 to c_2.

For example, if X is a single one-dimensional cube e with endpoints a and b, the nontrivial arrows of the idealization $\mathcal{X}$ are

$$(a) \leftarrow (a \subset e) \to (e) \leftarrow (b \subset e) \to (b). \tag{1}$$

Already a square τ with e as a face is much more complicated. The idealization is shown on the left of Figure 1 as a graph, with detail shown on the right for the highlighted portion.

Let X be a CAT(0) cube complex, and let $\mathcal{X}$ be its idealization. Any cubical automorphism of X gives a nondegenerate automorphism of $\mathcal{X}$. Moreover if such an automorphism maps a chain of cubes to itself, then it also preserves all subchains. In terms of the scwol structure this means that the stabilizer of an object also stabilizes every arrow from that object. It follows that the quotient $\mathcal{Y} = G\backslash\mathcal{X}$ is a simple scwol, and that the quotient map $\mathcal{X} \to \mathcal{Y}$ is nondegenerate morphism of scwols. Similarly, if

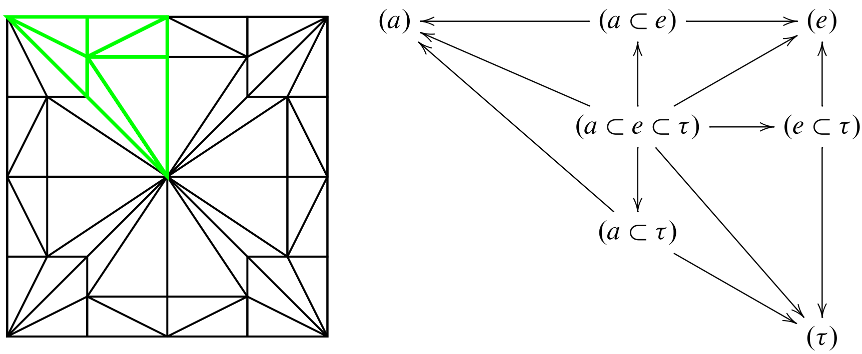

Figure 1: Idealization of a square.

$K < G$ is any subgroup, then $K\backslash\mathcal{X}$ is a simple scwol and the maps $\mathcal{X} \to K\backslash\mathcal{X} \to \mathcal{Y}$ are nondegenerate morphisms of scwols. In particular Corollary 2.18 applies to the covers of $CG(\mathcal{Y})$.

Example 2.26 Let X be the regular 3–valent tree dual to the Farey tessellation, and let

$$G = \mathrm{PSL}(2,\mathbb{Z}) \cong \langle x \mid x^2 \rangle * \langle y \mid y^3 \rangle$$

act on X in the standard way, with y rotating around a vertex v, and x rotating around the center of an adjacent edge e. Then $CG(\mathcal{Y})$ has the following form, omitting identity arrows and most arrow labels:

(2)
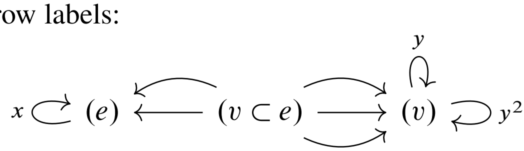

The two-fold cover corresponding to the subgroup generated by $\{y, xyx\}$ is of the following form, omitting identity arrows and all labels:

(3)
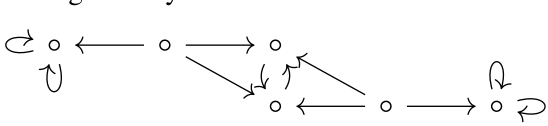

The scwolification of this cover is isomorphic to the scwol shown in (1). In both (2) and (3), the scwol arrows are exactly the horizontal ones.

This example is one-dimensional, so there are no nontrivial compositions of scwol arrows. The reader is invited to explore a simple two-dimensional example, for example the category of the complex of groups associated to a dihedral group acting on a square with quotient scwol equal to the right-hand part of Figure 1.

Remark 2.27 Suppose that $\mathcal{C}$ is the idealization of a cube complex C, so that the realization of $\mathcal{C}$ is the second barycentric subdivision of C. In later sections, we make use of the fact that the following types of paths have canonical idealizations in $\mathcal{C}$:

(1) combinatorial paths in the 1–skeleton of the first *cubical* subdivision C^b of C (Section 3);

(2) combinatorial paths in links of cells of C (Section 4).

In both cases, this follows from the fact that subdivisions of these graphs embed naturally in the 1–skeleton of the second barycentric subdivision.

3 Quasiconvexity in the Sageev construction

In this section, we prove Theorem A. Recall that we have a hyperbolic group G acting cocompactly on a CAT(0) cube complex X, and we are required to prove that the vertex stabilizers are quasiconvex if and only if the hyperplane stabilizers are quasiconvex.

To prepare for this proof it may be useful to think about the case that X is a tree. In that case, hyperplanes are midpoints of edges, and so the statement is that edge stabilizers are quasiconvex if and only if vertex stabilizers are. Edge stabilizers are intersections of vertex stabilizers, and intersections of quasiconvex subgroups are quasiconvex, so one direction is clear. The other direction is not much harder: Consider a geodesic joining two vertices of a vertex stabilizer. The vertex stabilizer is coarsely separated from the rest of the Cayley graph by appropriate cosets of edge stabilizers. The quasiconvexity of these cosets "traps" the geodesic close to the vertex stabilizer.

Now remove the assumption that X is a tree, and suppose that vertex stabilizers are quasiconvex. It still follows that edge stabilizers are quasiconvex, but a hyperplane stabilizer is much bigger than an edge stabilizer. We will express a hyperplane stabilizer as a union of cosets of edge stabilizers, intersecting in a controlled way, and use a quasiconvexity criterion proved in the appendix to conclude that the hyperplane stabilizer is quasiconvex.

If on the other hand we assume that hyperplane stabilizers are quasiconvex, we will use them as in the tree case to control geodesics joining points in a vertex stabilizer. We inductively use more and more hyperplanes to corral points on a geodesic in an argument which terminates because of the finite-dimensionality of the cube complex.

Throughout this section we suppose that X is a CAT(0) cube complex and that $\mathcal{X}$ is its idealization (see Definition 2.25). We suppose further that G is a group acting cocompactly on this cube complex. The quotient $G\backslash\mathcal{X}$ is a scwol $\mathcal{Y}$. Making choices as in Definition 2.7, we obtain a complex of groups $G(\mathcal{Y})$, with associated category $CG(\mathcal{Y})$. Choosing a vertex $v_0 \in \mathcal{Y}$, Theorem 2.13 gives an identification of G with $\pi_1(CG(\mathcal{Y}), v_0)$. It is helpful to assume (as we may do without loss of generality) that v_0 is the orbit of some 0–cube of X.

In Section 2.4, we defined $\widetilde{CG(\mathcal{Y})}$ to be the universal covering of the category $CG(\mathcal{Y})$. Recall that the objects of $\widetilde{CG(\mathcal{Y})}$ are homotopy classes of $CG(\mathcal{Y})$–paths, starting at the basepoint $v_0 \in \mathcal{Y}$, and arrows are labeled by arrows of $CG(\mathcal{Y})$ (Definition 2.15). The basepoint of $\widetilde{CG(\mathcal{Y})}$ is $\tilde{v}_0$, the homotopy class of the constant path at v_0. As in Section 2.4 we denote the universal covering map by $\phi\colon \widetilde{CG(\mathcal{Y})} \to CG(\mathcal{Y})$. Recall from Proposition 2.17 and Definition 2.20 that $\Theta\colon \widetilde{CG(\mathcal{Y})} \to \mathcal{X}$ is the scwolification functor.

We briefly describe the contents of the remainder of this section. In Section 3.1 we explain how we consider subsets of small categories as graphs. In Section 3.2 we identify certain subsets of $\widetilde{CG(\mathcal{Y})}$ which are tuned to the cubical geometry of X^b and associate graphs with these subsets. In Section 3.3 we prove the direction (1) $\Longrightarrow$ (2) of Theorem A. In fact, we prove the more general Theorem 3.19. In Section 3.4 we prove the direction (2) $\Longrightarrow$ (1) of Theorem A. In Section 3.5 we consider various possible generalizations of Theorem A. Finally, in Section 3.6 we prove Corollary B.

3.1 Graphs from subsets of small categories

Let $\mathcal{C}$ be a (small) category, and let S be a subset of the set of arrows of $\mathcal{C}$. There is an associated graph (really a 1–complex), which we denote by $\mathrm{Gr}(S)$, with vertex set the set of objects which are either the source or target of some arrow in S, and with edges in correspondence with the arrows S.

Example 3.1 Let S be the set of *all* arrows in $\mathcal{C}$. Then $\mathrm{Gr}(\mathcal{C}) := \mathrm{Gr}(S)$ is the 1–skeleton of the nerve of $\mathcal{C}$.

Example 3.2 Suppose $\mathcal{C}$ is a group (ie $\mathcal{C}$ has a single object and each arrow of $\mathcal{C}$ is invertible), and $S_0 \subset \mathcal{C}$ is a generating set. Let S be the set of arrows in the universal cover $\widetilde{\mathcal{C}}$ with label in S_0. Then $\mathrm{Gr}(S)$ is the Cayley graph of G with respect to S_0.

3.2 Cubical paths

The group $\pi_1(CG(\mathcal{Y}), v_0)$ acts on $\widetilde{CG(\mathcal{Y})}$, with quotient the category $CG(\mathcal{Y})$. As in Theorem 2.13, we can identify G with $\pi_1(CG(\mathcal{Y}), v_0)$. The proof of each direction of Theorem A begins with choosing a certain connected G–cocompact subgraph Γ of $\mathrm{Gr}(\widetilde{CG(\mathcal{Y})})$ (see Lemma 3.9). The graph Γ admits a G–equivariant map to the 1–skeleton of the cubical subdivision of X, which we now recall.

Definition 3.3 Suppose that X is a cube complex. The (*first*) *cubical subdivision* of X, denoted by X^b, is the cube complex obtained by replacing each n–cube in X by 2^n n–cubes, found by subdividing each coordinate interval into two equal halves, and then gluing in the obvious way induced from the structure of X.

Of course, X^b is canonically homothetic to X, and X^b is NPC (respectively, CAT(0)) if and only if X is. We suppose that X is CAT(0), and therefore X^b is also.

Observation 3.4 The vertices of X^b are in bijection with the cubes of X.

The cells of X^b are in bijection with pairs $(\tilde{\sigma}_1, \tilde{\sigma}_2)$ of cubes in X such that $\tilde{\sigma}_1 \subseteq \tilde{\sigma}_2$. The dimension of the cube corresponding to $(\tilde{\sigma}_1, \tilde{\sigma}_2)$ is $\dim(\tilde{\sigma}_2) - \dim(\tilde{\sigma}_1)$.

Thus, a 1–cell in X^b corresponds to a pair of cubes $(\tilde{\sigma}_1, \tilde{\sigma}_2)$ where $\tilde{\sigma}_1$ is a codimension-one face of $\tilde{\sigma}_2$. Moreover, each cell of X^b can be naturally identified with an object of $\mathcal{X}$.

As noted in Remark 2.27 any path in the 1–skeleton of X^b has a canonical idealization in $\mathcal{X}$. Each 1–cell e of X^b corresponds to some pair of cells $(\tilde{\sigma}_1 \subseteq \tilde{\sigma}_2)$ with $\tilde{\sigma}_1$ of codimension one in $\tilde{\sigma}_2$. If the path p traverses the 1–cell e, its idealization $\hat{p}$ contains consecutive arrows labeled $(\tilde{\sigma}_1 \subseteq \tilde{\sigma}_2) \to \tilde{\sigma}_1$ and $(\tilde{\sigma}_1 \subseteq \tilde{\sigma}_2) \to \tilde{\sigma}_2$, and every arrow of $\hat{p}$ has such a label. By Lemma 2.16, every $\widetilde{CG(\mathcal{Y})}$–path is homotopic to a concatenation of group arrows and scwol arrows. The graph that we use to prove Theorem A uses only scwol arrows that occur in pairs corresponding to the above description. Thus we make the following definition.

Definition 3.5 A *pair of opposable scwol arrows in* $CG(\mathcal{Y})$ is a pair of scwol arrows $(\gamma^{\downarrow}, \gamma^{\uparrow})$ such that

(1) $c = i(\gamma^{\downarrow}) = i(\gamma^{\uparrow})$ is an orbit of chain $(\tilde{\sigma}_1 \subset \tilde{\sigma}_2)$, where $\tilde{\sigma}_1$ has codimension one in $\tilde{\sigma}_2$;

(2) $\gamma^{\downarrow} = (1, a)$, where a is the arrow in $\mathcal{Y}$ corresponding to the G–orbit of the arrow $(\tilde{\sigma}_1 \subset \tilde{\sigma}_2) \to \tilde{\sigma}_1$ in $\mathcal{X}$; and

(3) $\gamma^{\uparrow} = (1, b)$, where b is the arrow in $\mathcal{Y}$ corresponding to the G–orbit of the arrow $(\tilde{\sigma}_1 \subset \tilde{\sigma}_2) \to \tilde{\sigma}_2$ in $\mathcal{X}$.

Suppose $c = i(\gamma^{\downarrow}) = i(\gamma^{\uparrow})$ for a pair of opposable scwol arrows $(\gamma^{\downarrow}, \gamma^{\uparrow})$. If $\tilde{c}$ is a lift of c to $\widetilde{CG(\mathcal{Y})}$, there are unique lifts $\tilde{\gamma}^{\downarrow}$ and $\tilde{\gamma}^{\uparrow}$ with source $\tilde{c}$. The pair $(\tilde{\gamma}^{\downarrow}, \tilde{\gamma}^{\uparrow})$ is a *pair of opposable scwol arrows in* $\widetilde{CG(\mathcal{Y})}$.

We remark that the image under Θ of a pair of opposable scwol arrows is quite restricted. In the right-hand part of Figure 1, for example, the possible images are the horizontal pair of arrows at the top of the diagram or the vertical pair at the right.

Definition 3.6 An object in $CG(\mathcal{Y})$ (equivalently, in $\mathcal{Y}$, since the objects of these two categories are the same) is *cubical* if it is an orbit of cubes in X (rather than an orbit of chains of cubes of length greater than 1). An object in $\widetilde{CG(\mathcal{Y})}$ is *cubical* if its projection to $CG(\mathcal{Y})$ is cubical.

A path p in $CG(\mathcal{Y})$ is *cubical* if

(1) the initial and terminal objects of p are cubical;

(2) p is a concatenation of group arrows and scwol arrows; and

(3) the scwol arrows occur in consecutive pairs, as pairs of opposable scwol arrows.

A path in $\widetilde{CG(\mathcal{Y})}$ is *cubical* if its projection to $CG(\mathcal{Y})$ is cubical.

It follows from the definition that all group arrows for a cubical path occur at cubical objects.

Proposition 3.7 *Suppose that v and w are cubical vertices of $CG(\mathcal{Y})$ and σ is a $CG(\mathcal{Y})$–path between v and w. Then σ is homotopic to a cubical path.*

In particular, every $g \in G = \pi_1(CG(\mathcal{Y}), v_0)$ is represented by a cubical $CG(\mathcal{Y})$–loop starting and ending at v_0.

Proof By path lifting, it suffices to prove the analogous statement for $\widetilde{CG(\mathcal{Y})}$–paths. We already observed in the proof of Proposition 2.17 that homotopies in $\mathcal{X}$ can be lifted to homotopies of $\widetilde{CG(\mathcal{Y})}$–paths, thus a given path can be homotoped to a path whose image under Θ stays in the idealization of $(X^b)^{(1)}$. Lemma 2.16 turns this into a concatenation of group arrows and scwol arrows (without changing its image under Θ). Any subpath $(1,a)^+ \cdot (g, \mathbb{1}_{i(a)})^{\pm}$ is homotopic to $(\psi(a)(g), \mathbb{1}_{t(a)})^{\pm} \cdot (1,a)^+$. In particular, after a homotopy, we may assume our path contains no such subpath, and whenever a scwol arrow occurs it occurs in a pair of opposable scwol arrows. $\square$

Definition 3.8 Suppose that for each cubical object $\mathfrak{o}$ of $\mathcal{Y}$ we choose a set $\mathbb{A}_{\mathfrak{o}} \subset G_{\mathfrak{o}}$. These determine a subset $S(\mathbb{A})$ of the arrows of $\widetilde{CG(\mathcal{Y})}$ which is the union of

(1) the set of (group) arrows with label $(g, \mathbb{1}_{\mathfrak{o}})$ for some $\mathfrak{o}$ and some $g \in \mathbb{A}_{\mathfrak{o}}$, and

(2) the set of scwol arrows occurring in some pair of opposable scwol arrows.

As discussed in Section 3.1, there is an associated graph $\mathrm{Gr}(S(\mathbb{A}))$ which we denote by $\Gamma(\mathbb{A})$. A vertex of this graph is called *cubical* if it comes from a cubical object, and otherwise it is called *central*.

Note that any central vertex of $\Gamma(\mathbb{A})$ only meets opposable scwol arrows and thus has valence exactly two, and each of its neighbors is a cubical vertex of $\Gamma(\mathbb{A})$. The valence of a cubical vertex coming from the object $\mathfrak{o}$ is equal to the number of opposable scwol arrows with terminus $\mathfrak{o}$ plus twice the cardinality of $\mathbb{A}_{\mathfrak{o}}$. Since $\mathcal{Y}$ is finite, the graph $\Gamma(\mathbb{A})$ is locally finite if and only if every $A_{\mathfrak{o}}$ is finite.

Lemma 3.9 *Suppose that $G = \pi_1(CG(\mathcal{Y}), v_0)$ is finitely generated. Then we can choose $\mathbb{A}$ so that $\Gamma(\mathbb{A})$ is locally finite, connected, and G–cocompact.*

Proof Let S be a finite generating set for G. For each $s \in S$ choose a cubical loop $p(s)$ in $CG(\mathcal{Y})$ based at v_0 representing s. For a cubical object $\mathfrak{o}$ of $\mathcal{Y}$, let $\mathbb{A}_{\mathfrak{o}}$ consist of those group arrows at $\mathfrak{o}$ which occur in some $p(s)$. There are only finitely many such, so the graph $\Gamma(\mathbb{A})$ defined in Definition 3.8 is locally finite.

To see that $\Gamma(\mathbb{A})$ is connected, let w be any vertex of $\Gamma(\mathbb{A})$. There is a path composed of opposable scwol arrows joining w to a vertex v in the G–orbit of $\tilde{v}_0$. Since S generates G, there is a concatenation of the cubical loops $p(s)$ which lifts to a path in $\Gamma(\mathbb{A})$ joining $\tilde{v}_0$ to v.

The set of G–orbits of cubical vertices of $\Gamma(\mathbb{A})$ injects into the set of objects of $\mathcal{Y}$, so it is finite and $\Gamma(\mathbb{A})$ is G–cocompact. $\square$

Convention 3.10 For the rest of this section, we fix $\Gamma = \Gamma(\mathbb{A})$ as in the conclusion of Lemma 3.9.

The functor $\Theta\colon \widetilde{CG(\mathcal{Y})} \to \mathcal{X}$ induces a map

$$\Psi\colon \Gamma \to (X^b)^{(1)}.$$

This map is simplicial after barycentrically subdividing the target. Each pair of edges of Γ coming from a pair of opposable scwol arrows maps to a single 1–cell of $(X^b)^{(1)}$. Any central vertex of Γ maps under Ψ to an intersection of an edge of X^b with a hyperplane of X^b.

Note that the map Ψ is continuous, G–equivariant and Lipschitz.

Definition 3.11 (cubical neighborhood) Let v be a vertex of X. The *cubical neighborhood of* v is the union of those cubes of X^b which contain v. It will be denoted below by $N(v)$.

Proposition 3.12 *Let v be a vertex of X. Then $\Psi^{-1}(N(v))$ is finite Hausdorff distance from $\Psi^{-1}(v)$ in Γ.*

Proof Since G acts cocompactly on X, there are finitely many $\mathrm{Stab}(v)$–orbits of pairs of cubes $(\sigma \subset \tau)$, so σ contains v and is a codimension-one face of τ.

Let $(\sigma \subset \tau)$ be one such pair of cubes. By Lemma 2.21, $\mathrm{Stab}(\tau)$ acts freely and transitively on $\Theta^{-1}((\tau))$. The subgroup $\mathrm{Stab}(\sigma) \cap \mathrm{Stab}(\tau) = \mathrm{Stab}((\sigma \subset \tau))$ preserves the collection of pairs of opposable scwol arrows joining $\Theta^{-1}((\tau))$ to $\Theta^{-1}((\sigma))$. Moreover, $\mathrm{Stab}((\sigma \subset \tau))$ is finite index in $\mathrm{Stab}(\sigma)$. It follows that there is some $c(\tau,\sigma) > 0$ such that every vertex of $\Psi^{-1}((\tau))$ is distance at most $c(\tau,\sigma)$ from a vertex of $\Psi^{-1}((\sigma))$. As there are only finitely many $\mathrm{Stab}(v)$–orbits of pairs $(\sigma \subset \tau)$ of such faces, there is some $c > 0$ which works for every such pair.

If $x \in \Psi^{-1}(N(v))$ is a cubical vertex, it is therefore distance at most $c \cdot \dim(X)$ from a vertex of $\Psi^{-1}(v)$. If $x \in \Psi^{-1}(N(v))$ is central, then it is distance 1 from a cubical vertex of $\Psi^{-1}(N(v))$. We have shown that $\Psi^{-1}(N(v))$ is contained in the $(c \cdot \dim(X)+1)$–neighborhood of $\Psi^{-1}(v)$. Since $\Psi^{-1}(v) \subset \Psi^{-1}(N(v))$, we are finished. □

Part of our reason for working in X^b is that Proposition 3.12 would fail if we defined $N(v)$ to be the union of those cubes of X meeting v.

In the following statements we use the convention that the empty intersection of hyperplanes of X is X^b and the empty intersection of subgroups is G.

Lemma 3.13 *Suppose that $W_1, \dots, W_k$ are hyperplanes in X and that $I = \bigcap_{i=1}^k W_i$ is nonempty. For any cell τ intersecting I, the subgroup $\mathrm{Stab}(I) \cap \mathrm{Stab}(\tau)$ is finite index in $\mathrm{Stab}(\tau)$.*

Proof Finitely many hyperplanes intersect τ and these are permuted by any element of $\mathrm{Stab}(\tau)$. Thus, a finite-index subgroup of $\mathrm{Stab}(\tau)$ fixes all the hyperplanes in I. □

We observe the following consequence of the cocompactness of $G \curvearrowright X$.

Lemma 3.14 *There are finitely many G–orbits of finite sets*

$$\{W_1, \dots, W_k\}$$

of distinct hyperplanes of X with nonempty intersection $\bigcap_{i=1}^k W_i$. For each such set, the intersection $\bigcap_{i=1}^k \mathrm{Stab}(W_i)$ is finite index in $\mathrm{Stab}\bigl(\bigcap_{i=1}^k W_i\bigr)$.

Definition 3.15 Let $D > 0$. An action of a group on a metric space is *D–cobounded* if there is a set of diameter D which meets every orbit.

Proposition 3.16 *There is a constant $D > 0$ such that for any nonempty intersection I of hyperplanes of X, the subgroup $\mathrm{Stab}(I)$ acts D–coboundedly on $\Psi^{-1}(I)$.*

Proof Since there are finitely many orbits of nonempty intersections of hyperplanes, it suffices to consider a single such intersection.

Since the action of G on X is cocompact, so is the action of $\mathrm{Stab}(I)$ on I. In particular, there are finitely many $\mathrm{Stab}(I)$ orbits of cubical or central vertices in the idealization of I. Let $\mathfrak{o}$ be the object of $\mathcal{X}$ corresponding to one of these vertices. By Lemma 2.21, $\mathrm{Stab}_G(\mathfrak{o})$ acts transitively on $\Theta^{-1}(\mathfrak{o})$. By Lemma 3.13, there is a finite-index subgroup of $\mathrm{Stab}_G(\mathfrak{o})$ in $\mathrm{Stab}(I)$. Thus we see that $\Psi^{-1}(I)$ contains finitely many $\mathrm{Stab}(I)$–orbits of vertices. □

3.3 If hyperplane stabilizers are QC then cell stabilizers are QC

In this section we prove the direction (1) $\Longrightarrow$ (2) of Theorem A. As mentioned in the introduction, we prove this in greater generality than that of a hyperbolic group acting cocompactly on a CAT(0) cube complex with quasiconvex hyperplane stabilizers. The right general setting for this proof is that of *strongly quasiconvex subgroups* of finitely generated groups, as defined by Tran in [35]. (Such subgroups were also studied by Genevois [14] under the name *Morse subgroups*.)

Definition 3.17 [35, Definition 1.1] Let X be a geodesic metric space. A subset $Q \subseteq X$ is *strongly quasiconvex* if for every $K \geq 1$ and $C \geq 0$ there is some $M = M(K, C)$ such that every (K, C)–quasigeodesic in X with endpoints in Q is contained in the M–neighborhood of Q. The function $M(K, C)$ is called a *Morse gauge*.

Strong quasiconvexity persists under quasi-isometries of pairs. This is presumably known to the experts, and is closely related to [35, Proposition 4.2], but we do not see it in the literature so we provide a proof sketch.

Theorem 3.18 *Suppose X and Y are geodesic metric spaces, that $A \subset X$ is strongly quasiconvex, that $\phi\colon X \to Y$ is a quasi-isometry and that $B \subset Y$ is finite Hausdorff distance from $\phi(A)$. Then B is a strongly quasiconvex subset of Y.*

Proof sketch This is proved essentially in the same way as the corresponding fact about quasiconvex subsets of hyperbolic spaces. The difference is that instead of a single constant of quasiconvexity, we must produce a Morse gauge.

Suppose that $\phi\colon X \to Y$ and $\psi\colon Y \to X$ are (λ, ϵ)–quasi-isometries which are ϵ–quasi-inverses, and that $d_{\mathrm{Haus}}(B, \phi(A)) \le \epsilon$.

Any quasigeodesic γ joining points in B can be extended by a pair of geodesic segments of length $\le \epsilon$ to make a quasigeodesic γ' joining points in $\phi(A)$. The image of γ' under ψ can likewise be extended to a quasigeodesic γ'' between points of A. If γ was a (K, C)–quasigeodesic, then γ'' is a (K', C')–quasigeodesic where K' and C' depend only on K, C, λ and ϵ. If M is the Morse gauge for A in X, then let $M_1 = M(K', C')$. For any point p on γ, the point $\psi(p)$ is on γ'' so it is within M_1 of some point in A. Using ϕ to move back to X, we see that p is within $\lambda M_1 + 3\epsilon$ of some point of B. We can thus define a Morse gauge M' for B in Y by $M'(K, C) = \lambda M(K', C') + 3\epsilon$. □

In particular, the notion of strong quasiconvexity makes sense for subgroups of finitely generated groups.

In this subsection, we prove the following theorem.

Theorem 3.19 *Suppose that a finitely generated group G acts cocompactly on a CAT(0) cube complex X and that the hyperplane stabilizers are strongly quasiconvex. Then the cell stabilizers are strongly quasiconvex.*

Since quasiconvexity is equivalent to strong quasiconvexity for subgroups of hyperbolic groups, Theorem 3.19 immediately implies the direction (1) $\Longrightarrow$ (2) of Theorem A.

Note that each cell stabilizer is a finite intersection of vertex stabilizers. Tran shows that a finite intersection of strongly quasiconvex subgroups is strongly quasiconvex [35, Theorem 1.2(2)], so we only need to show that vertex stabilizers are strongly quasiconvex whenever hyperplane stabilizers are.

We will use the following general statement about intersections of strongly quasiconvex sets, analogous to [8, III.Γ.4.13].

Proposition 3.20 *For any Morse gauge M and any D and N, there is a function $R\colon [0,\infty) \to [0,\infty)$ such that the following holds: Let X be a graph of valence at most N with a free G–action, let $A, B < G$ be subgroups acting D–coboundedly on M–strongly quasiconvex subgraphs Y_A and Y_B, respectively. If $Y_C = Y_A \cap Y_B$ is nonempty, then, for any $p \in X$,*

$$d(p, Y_A \cap Y_B) \le R\big(\max\{d(p,Y_A), d(p,Y_B)\}\big).$$

Proof Let $r > 0$. We must describe $R(r)$.

Note that a concatenation of a geodesic of length r with a geodesic of any length is a $(1,2r)$–quasigeodesic. Let $M_0 = M(1,2r)$. Let $R(r)$ be a bound for the number of pointed oriented simplicial paths in X of length $\le 2(M_0 + D)$, up to the G–action. (A bound can be chosen depending only on N, M, D and r.)

Let $p \in X$ be chosen such that $\max\{d(p,Y_A), d(p,Y_B)\} \le r$. Let q be a closest point in Y_C to p. Suppose $d(p,q) > R(r)$, and let γ be a geodesic from p to q. Every vertex on γ lies within M_0 of both Y_A and Y_B. It follows from D–coboundedness that every vertex v on γ lies within $M_0 + D$ of some aq and some bq for $a \in A$ and $b \in B$. By our choice of $R(r)$, there must be a pair of distinct vertices v_1 and v_2 on γ and paths σ_i joining v_i to $a_i q$, and τ_i joining v_i to $b_i q$ of length at most $M_0 + D$, and an element $h \in G$ such that $hv_1 = v_2$, $h\sigma_1 = \sigma_2$ and $h\tau_1 = \tau_2$. We may assume that v_1 is closer to q than v_2 is.

Note that we have $hAq \cap Aq \neq \varnothing$. Since the action of G on X is free, this implies that $h \in A$. By the same argument $h \in B$, so $h \in C$, and thus $hq \in Y_C$. But hq is closer to p than q is, contradicting our choice of q. □

Remark 3.21 It is straightforward to see, using the above proof, that the set Y_C is $(R \circ M)$–strongly quasiconvex, and also that C acts coboundedly on Y_C (with constants depending only on D, M and R).

Towards proving Theorem 3.19, suppose that G is a finitely generated group acting cocompactly on a CAT(0) cube complex X, and suppose that hyperplane stabilizers are strongly quasiconvex in G. Recall the graph Γ from Lemma 3.9, and the continuous, G–equivariant, Lipschitz map $\Psi\colon \Gamma \to (X^b)^{(1)}$.

Lemma 3.22 *Suppose that $W_1, \dots, W_k$ are hyperplanes in X and that $I = \bigcap_{i=1}^k W_i$ is nonempty. Then* $\mathrm{Stab}(I)$ *is strongly quasiconvex in G.*

Proof By Lemma 3.14, the stabilizer of I is a finite-index supergroup of the intersection $\bigcap_{i=1}^k \mathrm{Stab}(W_i)$. The intersection of strongly quasiconvex subgroups is strongly quasiconvex by [35, Theorem 1.2(2)]. □

Proposition 3.23 *There exists a Morse gauge M such that for any collection of hyperplanes in X with nonempty intersection I, the set $\Psi^{-1}(I)$ is strongly quasiconvex in Γ with Morse gauge M.*

Proof The G–action on X is cocompact, so there are finitely many G–orbits of sets $\{W_1, \dots, W_k\}$ of hyperplanes of X such that $\bigcap_{i=1}^k W_i \neq \varnothing$. If I is such a set and $g \in G$ then $\Psi^{-1}(g \cdot I) = g \cdot \Psi^{-1}(I)$. Therefore, it suffices to consider a (finite) collection of representatives of G–orbits of intersections and prove that each is individually strongly quasiconvex, and then take a maximum over the finitely many Morse gauges for these subsets.

By Proposition 3.16, $\mathrm{Stab}(I)$ acts cocompactly on $\Psi^{-1}(I)$ and, by Lemma 3.22, $\mathrm{Stab}(I)$ is strongly quasiconvex in G. Therefore, by considering an orbit map $G \to \Gamma$ and applying Theorem 3.18, we see that each $\Psi^{-1}(I)$ is a strongly quasiconvex subset of Γ. □

We now give the main part of the argument of the proof of Theorem 3.19, namely that if hyperplane stabilizers are strongly quasiconvex, vertex stabilizers are also strongly quasiconvex. We therefore fix a vertex v of X.

Note that $\Psi^{-1}(v)$ is a nonempty and $\mathrm{Stab}(v)$–invariant set of vertices of Γ consisting of finitely many $\mathrm{Stab}(v)$–orbits. Thus in order to show $\mathrm{Stab}(v)$ is strongly quasiconvex in G, it suffices (by Theorem 3.18) to show that the preimage $\Psi^{-1}(v)$ is a strongly quasiconvex subset of Γ.

We fix constants $K \geq 1$ and $C \geq 0$, suppose that a and b are vertices in $\Psi^{-1}(v)$ and let γ be a (K, C)–quasigeodesic in Γ between a and b. Let y be an arbitrary vertex on γ. We have to show $d(y, \Psi^{-1}(v))$ is bounded independent of a and b. By Proposition 3.12, $\Psi^{-1}(v)$ is finite Hausdorff distance from $\Psi^{-1}(N(v))$ (recall $N(v)$ is the cubical neighborhood of v), so it is enough to show $d\big(y, \Psi^{-1}(N(v))\big)$ is bounded independent of a and b.

Here is a description of our bound: Let N be a bound for the valence of Γ and D the bound from Proposition 3.16. Let $R_0 = M(K, C)$. Assuming R_i has been defined, we let R_{i+1} be the maximum of R_i and the number $R(R_i)$, where R is the function from the conclusion of Proposition 3.20, with the Morse gauge M and the above D and N. We will prove that $d(y, \Psi^{-1}(v)) \le R_{\dim X}$.

We build a sequence of points $y = z_0, z_1, \dots, z_t$ in Γ and hyperplanes $W_1, \dots, W_t$ in X for some $t \le \dim X$ such that for each i,

$$(*_i) \qquad d(y, z_i) \le R_i, \quad \Psi(z_i) \in \bigcap_{s=1}^{i} W_s, \quad \text{and} \quad \bigcap_{s=1}^{i} W_s \cap N(v) \neq \varnothing.$$

Notice that $(*_0)$ holds, since $d(y, z_0) = 0 \le R_0$ and the empty intersection of hyperplanes is X^b.

Proposition 3.24 *Suppose $i \ge 0$ and that $z_0, \dots, z_i$ and $W_1, \dots, W_i$ have been defined and satisfy $(*_i)$. Then either $\Psi(z_i) \in N(v)$ or there is some z_{i+1} and W_{i+1} such that z_{i+1} and $W_1, \dots, W_{i+1}$ satisfy $(*_{i+1})$.*

Proof Suppose that $\Psi(z_i) \notin N(v)$. We are given hyperplanes $W_1, \dots, W_i$ such that $I = \bigcap_{s=1}^{i} W_s$ is nonempty and $I \cap N(v) \neq \varnothing$. Notice that I is a combinatorially convex subcomplex of X^b. Choose a geodesic p in the 1–skeleton of I from $I \cap N(v)$ to $\Psi(z_i)$. The first edge of p joins a vertex u_1 in $I \cap N(v)$ to a vertex u_2 which is not in $I \cap N(v)$. Thus, the vertex u_1 corresponds to a cube τ of X which contains v, and u_2 corresponds to a cube σ which is a codimension one face of τ such that σ does not contain v. Let W_{i+1} be the hyperplane of X meeting τ in a mid-cube parallel to σ. Since $u_1 \in I \cap N(v)$ we see that $I \cap W_{i+1} \cap N(v) \neq \varnothing$, since it contains the vertex u_1 of X^b. Also, because p is a geodesic in $(X^b)^{(1)}$, W_{i+1} separates $N(v)$ from $\Psi_\Gamma(z_i)$. It is clear that

$$\bigcap_{s=1}^{i+1} W_s \cap N(v) = W_{i+1} \cap I \cap N(v) \neq \varnothing,$$

so it remains to find z_{i+1} satisfying the first two conditions of $(*_{i+1})$.

Claim $$d(y, \Psi^{-1}(W_{i+1})) \le R_i.$$

Proof If $\Psi(y) \in W_{i+1}$ then $d(y, \Psi^{-1}(W_{i+1})) = 0$, so we suppose $\Psi(y) \notin W_{i+1}$. There are two cases.

Suppose first W_{i+1} separates v from $\Psi(y)$. In this case, we know that γ must cross $\Psi^{-1}(W_{i+1})$ between a and y. However, $\Psi_\Gamma(\gamma)$ is a loop, so γ must also cross $\Psi^{-1}(W_{i+1})$ in the segment of γ between y and b. Thus, there is a (quasigeodesic) subsegment γ_1 of γ which contains y and which starts and finishes on $\Psi^{-1}(W_{i+1})$. Since M is a Morse gauge for $\Psi^{-1}(W_{i+1})$, and $M(K,C) \leq R_i$, the claim follows in this case.

Now suppose W_{i+1} does not separate v from $\Psi(y)$. In this case, since W_{i+1} *does* separate v from $\Psi(z_i)$ we know that W_{i+1} must separate $\Psi(y)$ from $\Psi(z_i)$. (In particular $i > 0$ in this case.) Since $d(y, z_i) \leq R_i$, and any path from y to z_i must intersect $\Psi^{-1}(W_{i+1})$, we must have $d(y, \Psi^{-1}(W_{i+1})) \leq R_i$, as required. □

Let $I = \bigcap_{s=1}^{i} W_s$. It follows immediately from the first two conditions of $(*_i)$ that $d(y, \Psi^{-1}(I)) \leq R_i$. Moreover, by the claim, $d(y, \Psi_\Gamma^{-1}(W_{i+1})) \leq R_i$. Furthermore,

$$\Psi^{-1}(I \cap W_{i+1}) = \Psi^{-1}(I) \cap \Psi^{-1}(W_{i+1}).$$

By Proposition 3.16 each of the sets $\Psi^{-1}(I)$, $\Psi^{-1}(W_{i+1})$ and $\Psi^{-1}(I \cap W_{i+1})$ is D–cobounded under the action of its stabilizer. Proposition 3.20 thus gives

$$d(y, \Psi^{-1}(I \cap W_{i+1})) \leq R(R_i) \leq R_{i+1}.$$

We choose z_{i+1} to be any point of $\Psi^{-1}(I \cap W_{i+1})$ which is closest to y. The point z_{i+1} satisfies the first two conditions of $(*_{i+1})$ so the proof is complete. □

For $j > \dim X$, there cannot exist a point z_j satisfying $(*_j)$, since there are no j–tuples of hyperplanes with nonempty intersection. Therefore, Proposition 3.24 asserts that for some $i \leq \dim X$, $\Psi(z_i) = v$. We conclude that $d(y, \Psi^{-1}(v)) \leq R_i \leq R_{\dim X}$, as desired.

This completes the proof of Theorem 3.19.

3.4 If cell stabilizers are QC then hyperplane stabilizers are QC

In this section we prove the direction (2) $\Longrightarrow$ (1) of Theorem A. Therefore, suppose that G is a hyperbolic group acting cocompactly on a CAT(0) cube complex X, and suppose that the vertex stabilizers are quasiconvex in G.

Let W be a hyperplane in X. Then as we have noted W is a subcomplex of X^b. Given w in $W \cap (X^b)^{(0)}$, let $Y(w)$ be the (closed) 1–neighborhood in $\Psi^{-1}(W)$ of $\Psi^{-1}(w)$.

Lemma 3.25 *The sets $Y(w)$ are quasiconvex subsets of Γ with constants which do not depend on w.*

Proof Since $\mathrm{Stab}(W)$ acts cocompactly on W, there are finitely many $\mathrm{Stab}(W)$–orbits of sets $Y(w)$, so the uniformity of constants will follow immediately if we can prove each $Y(w)$ is a quasiconvex subset of Γ. We therefore fix such a w.

The stabilizer $\mathrm{Stab}(w)$ is equal to the stabilizer of some cube σ of X. Thus $\mathrm{Stab}(w)$ is virtually the intersection of the stabilizers of the vertices of σ. The vertex stabilizers are assumed to be quasiconvex, so $\mathrm{Stab}(w)$ is also quasiconvex.

Since G acts freely and cocompactly on Γ and Ψ is equivariant, $\mathrm{Stab}(w)$ acts freely and cocompactly on $\Psi^{-1}(w)$. By Lemma 3.13, $\mathrm{Stab}(W) \cap \mathrm{Stab}(w)$ is a finite-index subgroup of $\mathrm{Stab}(w)$, so it is also quasiconvex and acts freely and cocompactly on $\Psi^{-1}(w)$. It moreover acts cocompactly on $Y(w)$.

The result follows (for example, by Theorem 3.18). □

We are ready to prove the direction (2) $\Longrightarrow$ (1) of Theorem A, which is the content of the following theorem. For this result, we assume Theorem A.3, which is proved in the appendix.

Theorem 3.26 *Suppose that the hyperbolic group G acts cocompactly on the cube complex X, and that for every vertex v of X, the stabilizer $\mathrm{Stab}(v)$ is quasiconvex. Then, for every hyperplane $W \subset X$, the stabilizer $\mathrm{Stab}(W)$ is a quasiconvex subgroup of G.*

Proof As we have already remarked, quasiconvexity of vertex stabilizers implies quasiconvexity of all cell stabilizers.

Let Γ, $\Psi^{-1}(W)$ and the $Y(w)$ be as discussed above. Since G acts freely and cocompactly on Γ, we know that Γ is δ–hyperbolic for some δ. Let ϵ be a constant such that $Y(w)$ is ϵ–quasiconvex for every w (Lemma 3.25).

Since $\mathrm{Stab}(W)$ acts freely and cocompactly on $\Psi^{-1}(W)$, in order to prove the theorem it suffices to prove that $\Psi^{-1}(W)$ is quasiconvex in Γ, so let $p, q \in \Psi^{-1}(W)$.

Consider a geodesic γ in $(X^b)^{(1)}$ between $\Psi(p)$ and $\Psi(q)$. Both $\Psi(p)$ and $\Psi(q)$ lie in W. Since W is combinatorially convex in X^b, the geodesic γ is entirely contained in the 1–skeleton of W (considered as a subcomplex of X^b). The vertices $w_1, \dots, w_n$ on γ

correspond to cells of X contained in W. The sets $Y(w_i)$ corresponding to these cells satisfy the hypotheses of Theorem A.3 with $m = 2$, $c = 1$, and ϵ the quasiconvexity constant chosen above. Theorem A.3 then implies that $Y(w_1) \cup \cdots \cup Y(w_n)$ is ϵ'–quasiconvex, for a constant ϵ' depending only on ϵ and δ.

In particular, a Γ–geodesic between p and q lies within ϵ' of $Y(w_1) \cup \cdots \cup Y(w_n)$. Since each of these $Y(w_i)$ is contained in $\Psi^{-1}(W)$, the Γ–geodesic between p and q stays uniformly close to $\Psi^{-1}(W)$, as required. □

Together with Theorem 3.19, this completes the proof of Theorem A.

3.5 On generalizations of Theorem A

For a subgroup H of a hyperbolic group G, the following three conditions are equivalent:

(a) H is strongly quasiconvex in G.

(b) H is quasiconvex in G.

(c) H is undistorted in G.

Dropping the condition that G is hyperbolic, condition (b) ceases to be a useful notion, but conditions (a) and (c) still make sense.

One can ask for versions of Theorem A where the hypothesis of hyperbolicity is removed and condition (b) is replaced by either condition (a) or (c).

3.5.1 Strong quasiconvexity Replacing quasiconvexity with strong quasiconvexity we can ask about the following conditions for a finitely generated group G acting cocompactly on a CAT(0) cube complex:

(1S) Hyperplane stabilizers are strongly quasiconvex.

(2S) Vertex stabilizers are strongly quasiconvex.

(3S) All cell stabilizers are strongly quasiconvex.

As remarked earlier, (2S) $\iff$ (3S) follows from [35, Theorem 1.2(2)]. Theorem 3.19 states that (1S) $\implies$ (2S).

The remaining implication (3S) $\implies$ (1S) is *false*, as shown for example by $\mathbb{Z}^2$ acting freely on a cubulated $\mathbb{R}^2$.

3.5.2 Undistortedness The situation when replacing quasiconvexity with quasi-isometric embeddedness is murkier. We consider the following conditions, for a finitely generated group G acting cocompactly on a CAT(0) cube complex X:

(1U) Hyperplane stabilizers are undistorted.

(2U) Vertex stabilizers are undistorted.

(3U) All cell stabilizers are undistorted.

If X is a tree, (1U) and (3U) each imply (2U), but not conversely. For example, the double of a finitely generated group over a distorted group acts on a tree with undistorted vertex stabilizers but distorted edge/hyperplane stabilizers.

We do not know the relationship between (1U) and (3U) in general, so we ask the question.

Question 3.27 *For finitely generated groups acting cocompactly on CAT(0) cube complexes does (1U) $\Longrightarrow$ (3U)? Does (3U) $\Longrightarrow$ (1U)?*

3.6 Height of families and the proof of Corollary B

The *height* of a subgroup was introduced in [16]. We need a generalization of this notion to families of subgroups.

Definition 3.28 (height of a family) Suppose that G is a group and $\mathcal{H}$ is a collection of subgroups. The *height* of $\mathcal{H}$ is the minimum number n such that for every tuple of distinct cosets $(g_0H_0, g_1H_1, \dots, g_nH_n)$ with $H_i \in \mathcal{H}$ (and $g_i \in G$), the intersection $\bigcap_{i=0}^n H_i^{g_i}$ is finite. If there is no such n then we say the height of $\mathcal{H}$ is infinite.

When $\mathcal{H} = \{H\}$ is a single subgroup, we recover the familiar notion of the height of a subgroup from [16].

The following result for a single subgroup is part of [16, Main Theorem]. The proof of that result from [1, Corollary A.40] can be adapted in the obvious way to prove the result for finite families. This result was proved in the more general setting of strongly quasiconvex subgroups by Tran [35, Theorem 1.2(3)].

Proposition 3.29 *Let G be a hyperbolic group and $\mathcal{H}$ a finite collection of quasiconvex subgroups of G. Then the height of $\mathcal{H}$ is finite.*

We also use a special case of a theorem of Charney and Crisp [9, Theorem 5.1]:

Theorem 3.30 *Suppose that G acts cocompactly on a cube complex X. Then X is quasi-isometric to the space obtained from the Cayley graph of G by coning cosets of stabilizers of vertices to points.*

We now prove Corollary B. For convenience, we recall the statement.

Corollary B *Suppose that G is a hyperbolic group acting cocompactly on a CAT(0) cube complex X with quasiconvex hyperplane stabilizers. Then*

(1) *X is δ–hyperbolic for some δ;*

(2) *there exists a $k \geq 0$ such that the fixed-point set of any infinite subgroup of G intersects at most k distinct cells; and*

(3) *the action of G on X is acylindrical* (in the sense of Bowditch [6, page 284]).

Proof If G is a hyperbolic group acting cocompactly on a CAT(0) cube complex, and if the stabilizers in G of vertices in X are quasiconvex, then [7, Theorem 7.11], due to Bowditch, implies that this coned graph is δ–hyperbolic for some δ. Theorem 3.30 then implies that the cube complex X is δ–hyperbolic for some (possibly different) δ. Thus, we have the first statement from Corollary B.

Now we prove the statement about fixed-point sets of infinite subgroups. Let I be a collection of orbit representatives of cells in X. For $i \in I$, let $Q_i = \{g \in G \mid gi = i\}$, and let $\mathcal{Q} = \{Q_i\}_{i \in I}$. Then $\mathcal{Q}$ is a finite collection of quasiconvex subgroups of G, so it has some finite height k by Proposition 3.29. If $H < G$ is infinite with nonempty fixed-point set and σ is a cell meeting the fixed-point set of H, then $H < Q_i^g$, where $\sigma = gi$. Since the height of $\mathcal{Q}$ is k, at most k such cells appear.

In [15], Genevois studies actions of groups on hyperbolic CAT(0) cube complexes and shows in Theorem 8.33 that, in this setting, acylindricity is equivalent to the condition:

(G) $$\exists L, R \ \forall x, y \in X^{(0)} \ d(x,y) \geq L \implies \#(\mathrm{Stab}(x) \cap \mathrm{Stab}(y)) \leq R.$$

We take R to be the maximum size of a finite subgroup and $L = k$. Suppose $d(x,y) \geq L$. Then the union of the combinatorial geodesics joining x to y contains finitely many (but at least $k+1$) vertices. There is a finite-index subgroup of $\mathrm{Stab}(x) \cap \mathrm{Stab}(y)$ which fixes all of these vertices. This finite-index subgroup fixes more than k cells, so it is finite. This implies $\mathrm{Stab}(x) \cap \mathrm{Stab}(y)$ is finite, as desired. □

Remark 3.31 In the context where G is a finitely generated group acting cocompactly on a cube complex X with strongly quasiconvex hyperplane stabilizers, the same proof of conclusion (2) works as written, replacing the reference to Proposition 3.29 with a reference to [35, Theorem 1.2(3)].

4 Conditions for quotients to be CAT(0)

As noted in the introduction, Theorem D follows quickly from Theorems A and F, Agol's Theorem [1, Theorem 1.1] and Wise's Quasiconvex Hierarchy Theorem [36, Theorem 13.3]. Thus, other than Theorem A.3 in the appendix (which is independent of everything else in this paper), it remains to prove Theorem F. Therefore, we are interested in conditions on a group G acting on a CAT(0) cube complex X and a normal subgroup $K \trianglelefteq G$ which ensure that the quotient $K \backslash X$ is a CAT(0) cube complex. In this section we develop criteria in terms of complexes of groups to ensure this. In the next section, we translate these conditions into algebraic conditions on $K \trianglelefteq G$.

Three conditions need to be ensured in order for the complex $Z = K \backslash X$ to be a CAT(0) cube complex:

(1) Z must be simply connected;

(2) Z must be a cube complex (rather than a complex made out of cells which are quotients of cubes); and

(3) Z must be nonpositively curved.

We investigate these three properties in turn.

4.1 Ensuring the quotient is simply connected

First, we give a sufficient condition for $K \backslash X$ to be simply connected.

Since X is a finite-dimensional cube complex, it has finitely many shapes, and we can use the following application of a theorem of Armstrong:

Theorem 4.1 *Let X be a simply connected metric polyhedral complex with finitely many shapes, and let K be a group of isometries of X respecting the polyhedral structure, generated by elements with fixed points. Then $K \backslash X$ is simply connected.*

Proof sketch A theorem of Armstrong [4, Theorem 3] shows that $K \backslash X$ is simply connected with the CW topology. We have to show it is still simply connected with the metric topology.

Because X has finitely many shapes, there is an equivariant triangulation $\mathcal{T}$ and an $\epsilon > 0$ such that for every finite subcomplex Y, the ϵ–neighborhood of Y deformation retracts to Y. If $f : S^1 \to X$ is any loop, then a compactness argument shows that it lies in an ϵ–neighborhood of some such finite complex. We can then homotope f to have image in Y and apply the simple connectedness of $K\backslash X$ with the CW topology. □

We remark that the hypothesis of finitely many shapes is necessary even when X is CAT(0) as the following example shows:

Example 4.2 For $n \in \{2, 3, \dots\}$, let D_n be the Euclidean cone of radius 1 on a loop σ_n of length $2\pi/n^2$. For each n mark a point on σ_n. Let Y be obtained from $\bigcup_n D_n$ by identifying the marked points. Unwrapping all the cones to Euclidean discs gives a tree of Euclidean discs of radius 1. We call this CAT(0) space $\widetilde{Y}$. There is a discrete group of isometries $\Gamma = \langle \gamma_2, \gamma_3, \dots \rangle$ acting on $\widetilde{Y}$ with quotient Y such that each γ_n fixes the center of some disc and rotates it by an angle of $2\pi/n^2$. Nonetheless Y is not simply connected, as the infinite concatenation of the loops σ_n has finite length, but cannot be contracted to a point.

4.2 Ensuring the quotient is a cube complex

We now turn to the question of when $K\backslash X$ is a cube complex.

In order that the quotient $Z = K\backslash X$ be a cube complex, there needs to be no element of K which fixes a cell of X setwise but not pointwise.

Suppose that σ is a cube of X. The stabilizer G_σ has a finite-index subgroup Q_σ consisting of those elements which fix σ pointwise. Let $\{\sigma_1, \dots, \sigma_k\}$ be a set of representatives of G–orbits of cubes in X. The following result is straightforward:

Proposition 4.3 *Suppose that G acts cocompactly on the cube complex X and that K is a normal subgroup of G such that for each i we have $G_{\sigma_i} \cap K \le Q_{\sigma_i}$. Then the quotient $K\backslash X$ is a cube complex and the links of vertices in $K\backslash X$ inherit a cellular structure from the simplicial structure of cells in X.*

4.3 Ensuring the quotient is nonpositively curved

The most complicated condition to ensure is that $K\backslash X$ is nonpositively curved.

Throughout this subsection we suppose that X is a CAT(0) cube complex and that $\mathcal{X}$ is its idealization (see Definition 2.25). We suppose further that G is a group acting cocompactly on this cube complex. The induced action of G on $\mathcal{X}$ has quotient a scwol $\mathcal{Y}$. Making choices as in Definition 2.7, we obtain a complex of groups $G(\mathcal{Y})$, with associated category $CG(\mathcal{Y})$. Choosing a vertex $v_0 \in \mathcal{Y}$, there is then an identification of G with $\pi_1(CG(\mathcal{Y}), v_0)$. Moreover, we choose a normal subgroup $K \trianglelefteq G$ such that $K \backslash X$ is a cube complex. Throughout this section we let $\mathcal{Z} = K \backslash \mathcal{X}$.[5]

In Section 4.2 we discussed how to find subgroups K such that $K \backslash X$ is a cube complex, but for this section we just assume that this is the case.

Let $\mathcal{C}_K$ be the cover of the category $CG(\mathcal{Y})$ corresponding to the subgroup K. Observe that $CG(\mathcal{Y})$–loops lift to $\mathcal{C}_K$ if and only if they represent elements of K. (Basepoints are mostly omitted in this section, since we deal with a normal subgroup K.)

Standing Assumption 4.4 Lemma 2.16 tells us that any $CG(\mathcal{Y})$–path is homotopic to a concatenation of group and scwol arrows. Throughout this section we assume all $CG(\mathcal{Y})$–paths have been homotoped to this form, though we do not assume that the arrows *alternate* between group and scwol arrows. We blur the distinction between the scwol arrow $(1, a_i)$ in $CG(\mathcal{Y})$ and the $\mathcal{Y}$–arrow a_i, and also between the group arrow $(g, \mathbb{1}_v)$ and the element $g \in G_v$.

Thus we may write a $CG(\mathcal{Y})$ path for example as a list $g_1 \cdot e_1 \cdot e_2 \cdot g_2 \cdots$, where each g_i is an element of a local group G_v and represents the edge $(g, \mathbb{1}_v)^+$, corresponding to the group arrow $(g, \mathbb{1}_v)$, and each e_i is equal to some $a_i^{\pm}$ for a scwol arrow $(1, a_i)$. We implicitly assume that each concatenation we write defines a path, which forces the group arrows labeled g_i to be elements of particular local groups. Whenever we consider a $CG(\mathcal{Y})$–path of length 1 consisting of a single group arrow we are either explicit about the local group or else it is clear from the context.

If we have an edge of the form $(1, \mathbb{1}_v)^{\pm}$, we often implicitly (or explicitly) omit this arrow from our path. Again this only changes the path by an elementary homotopy.

Definition 4.5 (link of a cube) Given an n–cube σ in a cube complex Z, let b_σ be the barycenter. For sufficiently small $\epsilon > 0$, the sphere $\{x \in Z \mid d_Z(b_\sigma, x) = \epsilon\}$ is the join of an $(n-1)$–sphere with some simplicial complex (here we take the join of a (-1)–sphere with K to be equal to K). This simplicial complex is what we refer to as the *link* of σ, or $\mathrm{link}(\sigma)$. It is naturally triangulated by simplices corresponding to inclusions of σ as a face of some higher-dimensional cube. In particular $\mathrm{link}(\sigma)$ is a Δ*–complex* [21, Chapter 2.1] (though it may not be simplicial).

[5]We do not make any further assumptions than these about G and X in this subsection. This may be an important observation for future applications.

Though the link of a cube naturally has the structure of a piecewise spherical complex, we will think of it as just a combinatorial object. When we refer to paths or loops in a link, these will always be concatenations of 1–cells. The *length* of such a path is the number of 1–cells traversed.

Theorem 4.6 (Gromov's cubical link condition [8, Theorem II.5.20]) *The cube complex Z is a nonpositively curved cube complex if and only if the link of each vertex in Z is flag.*

In this section, we provide a set of conditions on the subgroup K which imply the link condition for $Z = K\backslash X$.

We record two elementary observations:

Lemma 4.7 *Let v be a vertex of a cube complex, and let $L = \mathrm{link}(v)$. Then L is simplicial if and only if for every cell σ containing v, the 1–skeleton of $\mathrm{link}(\sigma)$ contains no immersed loop of length 1 or 2.*

Proof If L fails to be simplicial, there is either a nonembedded simplex, or a pair of simplices which intersect in a set which is not a face of both. If a simplex is nonembedded, we obtain a loop of length 1 in L. If two embedded simplices τ_1 and τ_2 of L intersect in a set which is not a single face, let F_1 and F_2 be different maximal faces in the intersection, and let $f = F_1 \cap F_2$. For $i \in \{1, 2\}$, let v_i be a vertex in $F_i - f$. Then the simplices spanned by $v_1 \cup f$ and $v_2 \cup f$ correspond to points in $\mathrm{link}(f)$ which lie on an immersed loop of length 2. But f corresponds to some cube containing σ, and $\mathrm{link}(f) \subset L$ is isomorphic to the link of that cube. □

Lemma 4.8 *Let v be a vertex of a cube complex, and suppose $L = \mathrm{link}(v)$ is simplicial. Then L is a flag complex if and only if for every cell σ containing v, every loop of length 3 in $\mathrm{link}(\sigma)$ is filled by a 2–cell.*

Proof If σ is a cube and ϕ is a cube with σ as a face, of dimension one higher, then ϕ corresponds to a vertex f of the link L of σ. The link of ϕ is isomorphic to the link in L of f. The result now follows from [8, Remark II.5.16(4)]. □

Therefore, in order to ensure Z is nonpositively curved, for each cell σ in Z we must rule out loops of length 1 and 2 in $\mathrm{link}(\sigma)$ and also ensure that any loop of length 3 in $\mathrm{link}(\sigma)$ is filled by a 2–cell. We first explain how we translate between 1–cells in links in Z and $CG(\mathcal{Y})$–paths. Then we develop the required conditions to rule out loops, finally dealing with loops of length 3 which must be filled by 2–cells.

4.3.1 $CG(\mathcal{Y})$–paths associated to 1–cells in link(σ) Below we choose, for each cube σ in X and each oriented 1–cell α in the link of σ, a $CG(\mathcal{Y})$–path $p_{[\![\alpha]\!]}$ which is the label of an unscwolification of the idealization of α. As indicated by the notation, this label is the same for two such 1–cells in the same G–orbit. (In fact we will choose these paths for slightly more general objects than 1–cells in links of cubes.) If α is an oriented 1–cell, we write $\overleftarrow{\alpha}$ for the same 1–cell with the opposite orientation.

We fix a cube σ of X. The second barycentric subdivision of the link of σ embeds naturally in the geometric realization of $\mathcal{X}$. The vertices of the image of link(σ) are precisely the length ≥ 2 chains of cubes whose minimal element is σ.

We first consider an oriented 1–cell β of link(σ). The 1–cell β corresponds to a triple of cubes $\epsilon_1, \epsilon_2, \phi$ in X such that

$$\dim(\sigma) = \dim(\epsilon_i) - 1 = \dim(\phi) - 2,$$

with $\epsilon_1, \epsilon_2 \subset \phi$ and $\epsilon_1 \cap \epsilon_2 = \sigma$.

In particular, β has idealization an $\mathcal{X}$–path of length 4, made up of arrows

$$(\sigma \subset \epsilon_1) \leftarrow (\sigma \subset \epsilon_1 \subset \phi) \rightarrow (\sigma \subset \phi) \leftarrow (\sigma \subset \epsilon_2 \subset \phi) \rightarrow (\sigma \subset \epsilon_2). \tag{4}$$

In order to formulate Definition 4.9 and Lemma 4.10 below we must consider a slightly more general situation: we suppose ϵ_1, ϵ_2 and ϕ are cubes in X with ϵ_1 and ϵ_2 codimension-one subcubes of ϕ, and γ is a chain of cubes in X such that each element of γ is contained in each of ϵ_1, ϵ_2 and ϕ. We can naturally extend γ to chains which we denote by $(\gamma \subset \epsilon_1)$, $(\gamma \subset \epsilon_2)$, $(\gamma \subset \phi)$, and $(\gamma \subset \epsilon_i \subset \phi)$. This sequence of chains corresponds to a 1–cell α in an "iterated link" (a link of a cell in a link, etc), which we fix from now through Definition 4.9. The 1–cell α has idealization an $\mathcal{X}$–path of length 4,

$$(\gamma \subset \epsilon_1) \leftarrow (\gamma \subset \epsilon_1 \subset \phi) \rightarrow (\gamma \subset \phi) \leftarrow (\gamma \subset \epsilon_2 \subset \phi) \rightarrow (\gamma \subset \epsilon_2). \tag{5}$$

The $\mathcal{X}$–path (5) may not embed in $\mathcal{Y}$. There are two ways this could happen. The first is that there is an element of Stab(γ) which sends ϵ_1 to ϵ_2, but no such element fixes ϕ. In this case, the image in $\mathcal{Y}$ is a nonbacktracking loop. The second possibility is that there is an element $g \in G$ sending each of γ and ϕ to itself, but exchanging ϵ_1 and ϵ_2. If there is such a g, the idealization of the 1–cell α backtracks in $\mathcal{Y}$, forming a "half 1–cell". Since we are assuming Z is a cube complex, the second possibility does not arise for the image of (5) in $\mathcal{Z}$, though the first possibility may.

Let $y_{[\![\alpha]\!]} = a_1^+ \cdot a_2^- \cdot a_3^+ \cdot a_4^-$ be the $\mathcal{Y}$–path which is the image of the $\mathcal{X}$–path (5). For more compact notation, we define

$$\nu = [\![(\gamma \subset \phi)]\!], \quad \mu_i = [\![(\gamma \subset \epsilon_i \subset \phi)]\!], \quad \xi_i = [\![(\gamma \subset \epsilon_i)]\!] \quad \text{for } i \in \{1,2\}.$$

Then we have the injective homomorphisms

$$\psi_{a_2}\colon G_{\mu_1} \to G_\nu \quad \text{and} \quad \psi_{a_3}\colon G_{\mu_2} \to G_\nu.$$

The images of ψ_{a_2} and ψ_{a_3} are equal. The projection to $\mathcal{Y}$ of the path (5) associated to α depends only on $[\![\alpha]\!]$. We denote the common image of ψ_{a_2} and ψ_{a_3} in G_ν by G_ν^+. Note that G_ν^+ either has index 2 in G_ν (if there is a g fixing γ and ϕ and exchanging ϵ_1 with ϵ_2) or else $G_\nu^+ = G_\nu$ (if there is no such g). If G_ν^+ has index 2 in G_ν, we fix a choice of $g_\nu \in G_\nu - G_\nu^+$. We make this choice once and for all for each orbit of $(\gamma, \epsilon_1, \epsilon_2, \phi)$, so that the choice depends only on the orbit and not on the representative.

In the sequel, we refer to the vertex groups by $G_{i([\![\alpha]\!])}$ (for G_{ξ_1}) and $G_{t([\![\alpha]\!])}$ (for G_{ξ_2}). We further define "edge-inclusions" $\psi_{[\![\alpha]\!]}\colon G_\nu^+ \to G_{t([\![\alpha]\!])}$ and $\psi_{[\![\bar{\alpha}]\!]}\colon G_\nu^+ \to G_{i([\![\alpha]\!])}$ by

$$\psi_{[\![\alpha]\!]} = \psi_{a_4} \circ \psi_{a_3}^{-1} \quad \text{and} \quad \psi_{[\![\bar{\alpha}]\!]} = \psi_{a_1} \circ \psi_{a_2}^{-1}.$$

Let $E_{[\![\alpha]\!]}$ denote the image of $\psi_{[\![\alpha]\!]}$ in $G_{t([\![\alpha]\!])}$, and $E_{[\![\bar{\alpha}]\!]}$ denote the image of $\psi_{[\![\bar{\alpha}]\!]}$ in $G_{i([\![\alpha]\!])}$.

Definition 4.9 If $G_\nu^+ \neq G_\nu$, associate to α the $CG(\mathcal{Y})$–path

$$p_\alpha = a_1^+ \cdot a_2^- \cdot g_\nu \cdot a_3^+ \cdot a_4^-,$$

where $g_\nu \in G_\nu - G_\nu^+$ is the element fixed above. Note that in this case $a_2 = a_3$ and $a_1 = a_4$. If $G_\nu^+ = G_\nu$, let

$$p_\alpha = a_1^+ \cdot a_2^- \cdot a_3^+ \cdot a_4^-.$$

In either case some lift of p_α to $\widetilde{CG(\mathcal{Y})}$ is an unscwolification of the idealization of α. The $CG(\mathcal{Y})$–path p_α depends only on $[\![\alpha]\!]$; if there is a g with $g\sigma = \sigma$ and $g\alpha' = \alpha$, then $p_{\alpha'} = p_\alpha$. Therefore, we have a well-defined $CG(\mathcal{Y})$–path $p_{[\![\alpha]\!]}$.

Next we consider $CG(\mathcal{Y})$–paths whose scwolifications traverse $y_{[\![\alpha]\!]}$ and whose lifts to $\widetilde{CG(\mathcal{Y})}$ have nonbacktracking scwolifications. (Since $Z = K\backslash X$ is assumed to be a cube complex, these paths also have lifts in $\mathcal{C}_K$ whose scwolifications are nonbacktracking in $\mathcal{Z}$.) We observe in the following lemma that such paths can be converted by a sequence of elementary homotopies to paths which consist of a copy of $p_{[\![\alpha]\!]}$ bookended by group arrows.

Lemma 4.10 *Suppose that $G_\nu^+ = G_\nu$. Then any $CG(\mathcal{Y})$–path*

$$g_0 \cdot a_1^+ \cdot g_1 \cdot a_2^- \cdot g_2 \cdot a_3^+ \cdot g_3 \cdot a_4^- \cdot g_4$$

is homotopic to a $CG(\mathcal{Y})$–path of the form

$$g_0' \cdot p_{[\![\alpha]\!]} \cdot g_1'.$$

Suppose that $G_\nu^+ \neq G_\nu$. Then any $CG(\mathcal{Y})$–path

$$g_0 \cdot a_1^+ \cdot g_1 \cdot a_2^- \cdot g_2 \cdot a_3^+ \cdot g_3 \cdot a_4^- \cdot g_4$$

such that $g_2 \notin E_{[\![\alpha]\!]}$ is homotopic to a $CG(\mathcal{Y})$–path of the form

$$g_0' \cdot p_{[\![\alpha]\!]} \cdot g_1'.$$

In both cases, the scwolification of the path is fixed during the homotopy. Moreover any lift of the homotopy to a cover of $CG(\mathcal{Y})$ gives a sequence of paths with constant scwolification.

Notation 4.11 We fix some notation in order to study paths in Z and also $\mathcal{Y}$–paths and $CG(\mathcal{Y})$–paths. As above, we use $[\![\cdot]\!]$ to denote a G–orbit in $\mathcal{Z}$, which corresponds to its image in $\mathcal{Y}$ under the projection $\pi : \mathcal{Z} \to \mathcal{Y}$.

Let $p_{[\![\alpha]\!]}$ be one of the $CG(\mathcal{Y})$–paths fixed in Definition 4.9, corresponding to a 1–cell α in some link (or iterated link) of a cube of Z. The $CG(\mathcal{Y})$–path $p_{[\![\alpha]\!]}$ has an underlying $\mathcal{Y}$–path, which we denote by $y_{[\![\alpha]\!]}$. Define $t([\![\alpha]\!]) = t(y_{[\![\alpha]\!]})$ and $i([\![\alpha]\!]) = i(y_{[\![\alpha]\!]})$. This is so we can denote the corresponding local groups by $G_{i([\![\alpha]\!])}$ and $G_{t([\![\alpha]\!])}$.

We will also need to refer to the subgroups $E_{[\![\alpha]\!]} < G_{t([\![\alpha]\!])}$ and $E_{[\![\bar{\alpha}]\!]} < G_{i([\![\alpha]\!])}$ defined just before Definition 4.9. Each of these subgroups can be thought of as the pointwise stabilizer of some translate of a lift of α to X.

4.3.2 Loops in link($\breve{\sigma}$) We are now ready to formulate the conditions on K which characterize whether or not $K\backslash X$ is nonpositively curved. We use Lemmas 4.7 and 4.8 repeatedly.

Recall that we have fixed a $K \lhd G$ such that $Z = K\backslash X$ is a cube complex. We also fix a cube $\breve{\sigma}$ of Z, and a lift σ of $\breve{\sigma}$ to X. If α is a 1–cell in link(σ), then as described above there is a corresponding $\mathcal{X}$–path of length 4 (its idealization). Similarly a 1–cell in the link of $\breve{\sigma}$ has a corresponding $\mathcal{Z}$–path of length 4. We sometimes conflate a concatenation of 1–cells with a concatenation of these idealizations.

We next give an algebraic characterization of 1–cells in $\mathrm{link}(\sigma)$ which project to loops of length 1 in $\mathrm{link}(\breve{\sigma})$. Recall $\mathcal{C}_K = K\backslash\widetilde{CG(\mathcal{Y})}$.

Lemma 4.12 *Let α be a 1–cell in $\mathrm{link}(\sigma)$, and $\breve{\alpha}$ the projection of α to $\mathrm{link}(\breve{\sigma})$. The endpoints of $\breve{\alpha}$ are equal if and only if there is a $CG(\mathcal{Y})$–loop of the form $p_{[\![\alpha]\!]} \cdot g$ that represents a conjugacy class in K.*

Proof Thinking of X as the geometric realization of $\mathcal{X}$, the 1–cell α is the realization of an $\mathcal{X}$–path q_α of length 4, which projects to a $\mathcal{Y}$–path $a_1^+ \cdot a_2^- \cdot a_3^+ \cdot a_4^-$. Let $\hat{q}_\alpha$ be an unscwolification of q_α in $\widetilde{CG(\mathcal{Y})}$, which we may choose to have label

$$a_1^+ \cdot a_2^- \cdot g_1 \cdot a_3^+ \cdot a_4^- \tag{6}$$

for some group arrow g_1.

Suppose first that the endpoints of $\breve{\alpha}$ coincide. Then the path (6) projects in $\mathcal{C}_K$ to a path with endpoints separated by a group arrow, and so there is a $\mathcal{C}_K$–loop with label

$$a_1^+ \cdot a_2^- \cdot g_1 \cdot a_3^+ \cdot a_4^- \cdot g_2.$$

Lemma 4.10 implies that this loop is homotopic to a loop with label $g_0 \cdot p_{[\![\alpha]\!]} \cdot g_1$ for some g_0 and g_1, and starting this loop at a different place gives the required $CG(\mathcal{Y})$–loop $p_{[\![\alpha]\!]} \cdot g$ representing a conjugacy class of K.

Conversely, suppose a conjugacy class in K is represented by a $CG(\mathcal{Y})$–loop of the form $p_{[\![\alpha]\!]} \cdot g$. Then $p_{[\![\alpha]\!]} \cdot g$ lifts to a loop in $\mathcal{C}_K$ whose scwolification is a projection of a translate of q_α by some element of G. In particular, q_α must project to a loop, and so the endpoints of $\breve{\alpha}$ coincide. □

Definition 4.13 If $\mathfrak{o}$ is an object of $\mathcal{Y}$, let $K_{\mathfrak{o}} \lhd G_{\mathfrak{o}}$ be $K \cap G_{\mathfrak{o}}$.

Definition 4.14 A $CG(\mathcal{Y})$–path p is *K–nonbacktracking* if for some (equivalently any) lift $\hat{p}$ to $\mathcal{C}_K$, the scwolification $\Theta_K(\hat{p})$ is nonbacktracking. A $CG(\mathcal{Y})$–loop can be thought of as a path starting at any of its vertices. The loop is K–nonbacktracking if all these paths are K–nonbacktracking.

Lemma 4.15 *A $CG(\mathcal{Y})$–path $g_0 \cdot p_{[\![\alpha_1]\!]} \cdot g_1 \cdot p_{[\![\alpha_2]\!]} \cdots g_{k-1} \cdot p_{[\![\alpha_k]\!]} \cdot g_k$ is K–nonbacktracking if and only if*

(1) *for $i \in \{1, \dots, k-1\}$, if $[\![\alpha_{i+1}]\!] = [\![\overleftarrow{\alpha}_i]\!]$ then $g_i \notin E_{[\![\alpha_i]\!]} K_{t([\![\alpha_i]\!])} \subset G_{t([\![\alpha_i]\!])}$.*

Furthermore, a $CG(\mathcal{Y})$–loop with such a label is K–nonbacktracking if and only if (1) *and*

(2) *if $[\![\alpha_1]\!] = [\![\overleftarrow{\alpha}_k]\!]$, then $g_k g_0 \notin E_{[\![\alpha_k]\!]} K_{t([\![\alpha_k]\!])} \subset G_{t([\![\alpha_k]\!])}$.*

The following result algebraically characterizes immersed loops of length 2 in $\mathrm{link}(\sigma)$.

Lemma 4.16 *Let p be a path in $\mathrm{link}(\breve{\sigma})$ which is a concatenation of two 1–cells, $\breve{\alpha}$ and $\breve{\beta}$, which lift respectively to 1–cells α and β in X. The following are equivalent:*

(1) *There is a path $p' = \breve{\alpha}' \cdot \breve{\beta}'$ in $\mathrm{link}(\breve{\sigma})$ with $[\![\breve{\alpha}']\!] = [\![\breve{\alpha}]\!]$ and $[\![\breve{\beta}']\!] = [\![\breve{\beta}]\!]$ such that p' is an immersed loop.*

(2) *There is a K–nonbacktracking $CG(\mathcal{Y})$–loop $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2$ that represents a conjugacy class in K.*

Moreover, when these conditions hold, the path p' can be chosen to be the scwolification of a lift of $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2$ (and conversely $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2$ is the $CG(\mathcal{Y})$–path which labels the unscwolification of p').

Proof Suppose that there is an immersed loop $p' = \breve{\alpha}' \cdot \breve{\beta}'$ as in (1). The idealization of p' is a $\mathcal{Z}$–path $q_{p'}$ of length 8, labeled by a $\mathcal{Y}$–path $a_1^+ \cdot a_2^- \cdot a_3^+ \cdot a_4^- \cdot b_1^+ \cdot b_2^- \cdot b_3^+ \cdot b_4^-$ as discussed above. Using Lemma 4.10, we can choose an unscwolification $\hat{q}_{p'}$ of $q_{p'}$ in $\mathcal{C}_K$ with label

$$p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]},$$

where g_1 is a group arrow. But the unscwolification $\hat{q}_{p'}$ has endpoints separated by a group arrow g_2, so there is a loop labeled $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2$ as desired. It is K–nonbacktracking since its scwolification is the path $q_{p'}$.

Conversely, suppose that there is a K–nonbacktracking $CG(\mathcal{Y})$–loop

$$p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2,$$

which represents an element of K. Then $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2$ lifts to a loop in $\mathcal{C}_K$. The scwolification of this loop gives a path p' as in condition (1). □

The following is elementary.

Lemma 4.17 *Let Q be a complex such that there are no loops of length 1 or 2 in its 1–skeleton. Any loop in Q of length 3 is nonbacktracking.*

The utility of Lemma 4.17 is that once we have found conditions to ensure that links in $K\backslash X$ have no loops of length 1 or 2 then loops of length 3 are automatically nonbacktracking.

Given Lemma 4.17, the following is proved in the same way as Lemma 4.16.

Lemma 4.18 *Suppose that* $\operatorname{link}(\breve{\sigma})$ *is simplicial, and suppose that* p *is a path in* $\operatorname{link}(\breve{\sigma})$ *which is a concatenation of three* 1*–cells,* $\breve{\alpha}$, $\breve{\beta}$ *and* $\breve{\gamma}$, *with lifts* α, β *and* γ *to* X. *The following are equivalent:*

(1) *There is a path* $p' = \breve{\alpha}' \cdot \breve{\beta}' \cdot \breve{\gamma}'$ *in* $\operatorname{link}(\breve{\sigma})$ *such that* $[\![\breve{\alpha}']\!] = [\![\breve{\alpha}]\!]$, $[\![\breve{\beta}']\!] = [\![\breve{\beta}]\!]$ *and* $[\![\breve{\gamma}']\!] = [\![\breve{\gamma}]\!]$, *and* p' *is an immersed loop.*

(2) *There is a* $CG(\mathcal{Y})$*–loop of the form* $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2 \cdot p_{[\![\gamma]\!]} \cdot g_3$ *that represents a conjugacy class in* K.

Moreover, when these conditions hold, the path p' *can be chosen to be the scwolification of a lift of* $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2 \cdot p_{[\![\gamma]\!]} \cdot g_3$ *(and conversely* $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2 \cdot p_{[\![\gamma]\!]} \cdot g_3$ *is the* $CG(\mathcal{Y})$*–path which labels the unscwolification of* p'*).*

If X has dimension greater than 2, there are certainly some σ such that there are loops of length 3 in $\operatorname{link}(\sigma)$. This introduces some subtleties, which we discuss in the next subsection.

4.3.3 Loops of length 3 filled by 2–cells We assume for the rest of the section that cubes of Z are embedded, so that objects of $\mathcal{Z}$ can be unambiguously described by chains of cubes of Z. The phenomenon we are concerned with in this section is illustrated by the following example.

Example 4.19 Let $\mathcal{Y}$ be a single 2–simplex, and consider the complex of groups $G(\mathcal{Y})$ such that $G_v \cong \mathbb{Z}$ for each vertex v, and all the other local groups are trivial. Let $x, y, z \in \pi_1(G(\mathcal{Y}))$ generate the three vertex groups. The universal cover X of $G(\mathcal{Y})$ is an infinite-valence "tree of triangles". Let $K = \langle\!\langle x^3, y^3, z^3, xyz \rangle\!\rangle$. Then $K \backslash X$ can be realized as a subset of the Euclidean plane, consisting of every other triangle of a tessellation by equilateral triangles. Moreover, if $\alpha\beta\gamma$ is the path in the 1–skeleton of $\mathcal{Y}$ labeling the boundary of $\mathcal{Y}$, there are paths in $K \backslash X$ projecting to $\alpha\beta\gamma$, but which are not filled by a 2–cell in $K \backslash X$. The issue here, as we will see, is that $xyz \in K$ is not an element of $K_x K_y K_z$, where $K_x = K \cap \langle x \rangle$, and so on.

Of course X is not a cube complex, but it can be realized as the link of a vertex of a cube complex, covering a complex of groups in which $G(\mathcal{Y})$ is embedded.

Definition 4.20 Let $\breve{\sigma}$ be a cube of Z, and let $\breve{\tau}$ be a 2–cell in $\operatorname{link}(\breve{\sigma})$. Then $\partial\breve{\tau}$ is a loop, composed of three oriented 1–cells, $\breve{\alpha} \cdot \breve{\beta} \cdot \breve{\gamma}$, which we may lift to 1–cells α, β and γ, forming the boundary of a lift τ of $\breve{\tau}$ in X. These 1–cells are associated to $CG(\mathcal{Y})$–paths $p_{[\![\alpha]\!]}$, $p_{[\![\beta]\!]}$ and $p_{[\![\gamma]\!]}$ as in Definition 4.9. Consider a $CG(\mathcal{Y})$–path of

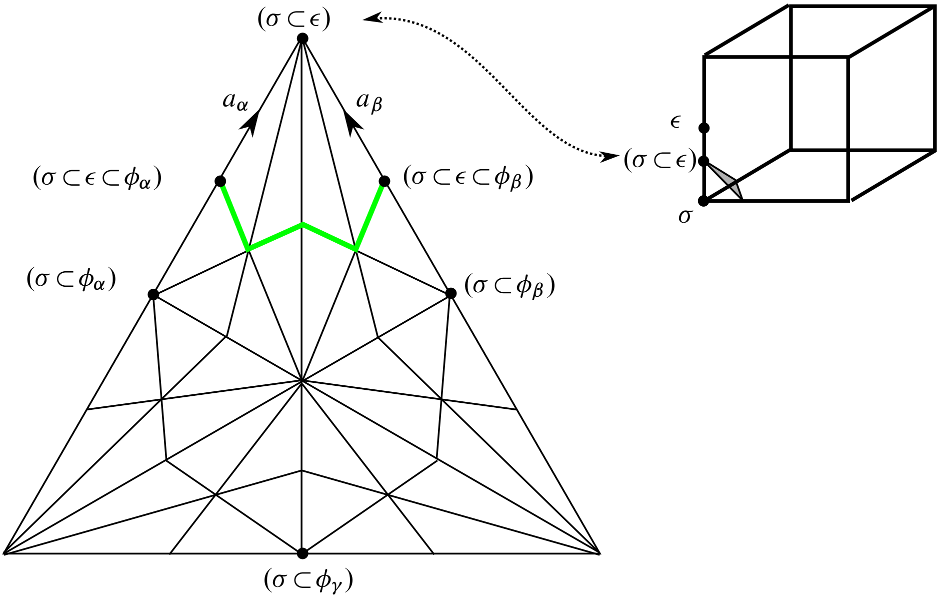

Figure 2: A part of $\mathcal{X}$ representing part of the link of σ, containing the idealization of the 1–cell ζ in green. Directions of most arrows have been omitted.

the form $q = p_{[\![\alpha]\!]} \cdot g \cdot p_{[\![\beta]\!]}$. Let $\hat{q}$ be a lift to $\mathcal{C}_K$. The realization of $\Theta_K(\hat{q})$ is a concatenation of two 1–cells $\breve{\alpha}' \cdot \breve{\beta}'$. We say that q *K–bounds a (τ, α)–corner* if there is a cube $\breve{\sigma}'$, a 2–cell $\breve{\tau}'$ in link$(\breve{\sigma}')$, and an $h \in G$ such that $\breve{\sigma}' = h\breve{\sigma}$, $\breve{\tau}' = h\breve{\tau}$, $\breve{\alpha}' = h\breve{\alpha}$ and $\breve{\beta}' = h\breve{\beta}$. If there is some (τ, α) for which the path q K–bounds a (τ, α)–corner, we may just say q *K–bounds a corner.*

If there exists a path q as above which K–bounds a (τ, α)–corner, there are X–cubes ϵ, ϕ_α, ϕ_β and ψ, all containing σ, such that $\epsilon \subset \phi_\alpha, \phi_\beta \subset \psi$, and

$$\dim(\psi) = \dim(\phi_\alpha) + 1 = \dim(\phi_\beta) + 1 = \dim(\epsilon) + 2 = \dim(\sigma) + 3.$$

There is a copy of the link of ϵ contained in the link of σ. The cubes ϵ, ϕ_α, ϕ_β and ψ determine an oriented 1–cell ζ in this copy of the link of ϵ. Its idealization is shown in Figure 2. The idealization of ζ begins at the object $(\sigma \subset \epsilon \subset \phi_\alpha)$ and ends at $(\sigma \subset \epsilon \subset \phi_\alpha)$. Let a_α be the arrow pointing from $(\sigma \subset \epsilon \subset \phi_\alpha)$ to $(\sigma \subset \epsilon)$, and let a_β be the arrow pointing from $(\sigma \subset \epsilon \subset \phi_\beta)$ to $(\sigma \subset \epsilon)$. These arrows project to arrows $[\![a_\alpha]\!]$ and $[\![a_\beta]\!]$ in $\mathcal{Y}$, and the path $p_{[\![\zeta]\!]}$ (defined as in Definition 4.9) travels from $i([\![a_\alpha]\!])$ to $i([\![a_\beta]\!])$.

Lemma 4.21 *The $CG(\mathcal{Y})$–loop*

$$[\![a_\alpha]\!]^+ \cdot p_{[\![\zeta]\!]} \cdot [\![a_\beta]\!]^-,$$

which is based at $[\![\sigma \subset \epsilon]\!]$, represents an element of $G_{[\![\sigma\subset\epsilon]\!]}$.

Proof All the chains which occur in this proof have the same minimal element σ, so we omit the prefix "$\sigma \subset$" from all chains until the end of the proof of the lemma. We therefore have a diagram in $\mathrm{link}(\sigma)$ in the scwol $\mathcal{X}$ as follows:

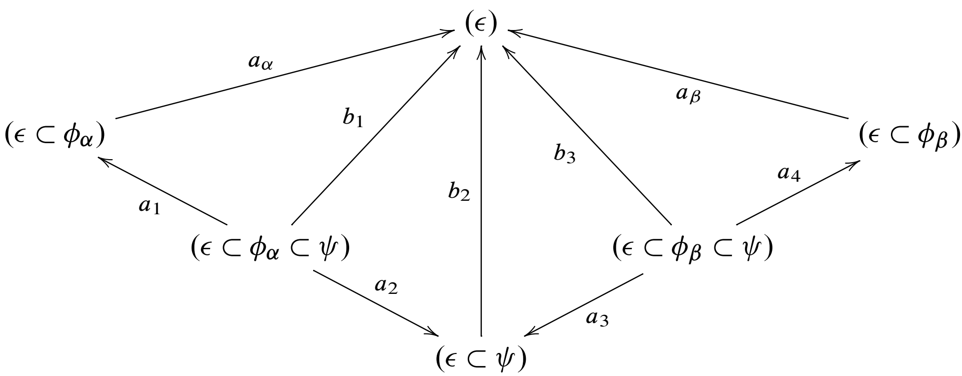

We have the identities $a_\alpha a_1 = b_1 = b_2 a_2$ and $a_\beta a_4 = b_3 = b_2 a_3$ in the category $\mathcal{X}$. The path in the statement of the lemma is equal to

$$[\![a_\alpha]\!]^+ \cdot [\![a_1]\!]^+ \cdot [\![a_2]\!]^- \cdot g_{(\epsilon\subset\psi)} \cdot [\![a_3]\!]^+ \cdot [\![a_4]\!]^- \cdot [\![a_\beta]\!]^-,$$

where $g_{(\epsilon\subset\psi)}$ is the element of $G_{[\![(\epsilon\subset\psi)]\!]}$ chosen for the path p_ζ as in Definition 4.9.

Define the elements of $G_{[\![\epsilon]\!]}$

$$\begin{aligned} h_1 &= z([\![a_\alpha]\!], [\![a_1]\!]) z([\![b_2]\!], [\![a_2]\!])^{-1}, \\ h_2 &= h_1 \psi_{[\![b_2]\!]}(g_{(\epsilon\subset\psi)}), \\ h_3 &= h_2 z([\![b_2]\!], [\![a_3]\!]) z([\![a_\beta]\!], [\![a_4]\!])^{-1}, \end{aligned}$$

where the $z([\![a]\!], [\![b]\!])$ are the twisting elements determined by the complex of groups structure on $G(\mathcal{Y})$.

We now have the sequence of elementary homotopies of $CG(\mathcal{Y})$–paths (all of which consist of applying the moves in Definition 2.1, and the rule of arrow composition in $CG(\mathcal{Y})$ from Definition 2.11),

$$[\![a_\alpha]\!]^+ \cdot p_{[\![\zeta]\!]} \cdot [\![a_\beta]\!]^- \simeq [\![a_\alpha]\!]^+ \cdot [\![a_1]\!]^+ \cdot [\![a_2]\!]^- \cdot g_{(\epsilon\subset\psi)} \cdot [\![a_3]\!]^+ \cdot [\![a_4]\!]^- \cdot [\![a_\beta]\!]^-$$

$$\begin{aligned}
&\simeq z([\![a_\alpha]\!],[\![a_1]\!])\cdot[\![b_1]\!]^+\cdot[\![a_2]\!]^-\cdot g_{(\epsilon\subset\psi)}\cdot[\![a_3]\!]^+\cdot[\![a_4]\!]^-\cdot[\![a_\beta]\!]^-\\
&\simeq z([\![a_\alpha]\!],[\![a_1]\!])z([\![b_2]\!],[\![a_2]\!])^{-1}\cdot[\![b_2]\!]^+\cdot[\![a_2]\!]^+\cdot[\![a_2]\!]^-\\
&\qquad\qquad\cdot g_{(\epsilon\subset\psi)}\cdot[\![a_3]\!]^+\cdot[\![a_4]\!]^-\cdot[\![a_\beta]\!]^-\\
&\simeq h_1\cdot[\![b_2]\!]^+\cdot g_{(\epsilon\subset\psi)}\cdot[\![a_3]\!]^+\cdot[\![a_4]\!]^-\cdot[\![a_\beta]\!]^-\\
&\simeq h_1\cdot\psi_{[\![b_2]\!]}(g_{(\epsilon\subset\psi)})\cdot[\![b_2]\!]^+\cdot[\![a_3]\!]^+\cdot[\![a_4]\!]^-\cdot[\![a_\beta]\!]^-\\
&\simeq h_2 z([\![b_2]\!],[\![a_3]\!])\cdot[\![b_3]\!]^+\cdot[\![a_4]\!]^-\cdot[\![a_\beta]\!]^-\\
&\simeq h_2 z([\![b_2]\!],[\![a_3]\!])z([\![a_\beta]\!],[\![a_4]\!])^{-1}\cdot[\![a_\beta]\!]^+\cdot[\![a_4]\!]^+\cdot[\![a_4]\!]^-\cdot[\![a_\beta]\!]^-\\
&\simeq h_3\cdot[\![a_\beta]\!]^+\cdot[\![a_\beta]\!]^-\\
&\simeq h_3. \qquad\square
\end{aligned}$$

Notation 4.22 The element of $G_{[\![\epsilon]\!]}$ represented by $[\![a_\alpha]\!]^+\cdot p_{[\![\zeta]\!]}\cdot[\![a_\beta]\!]^-$ is denoted by $g_{\tau,\alpha}$.

Lemma 4.23 *A path $p_{[\![\alpha]\!]}\cdot g\cdot p_{[\![\beta]\!]}$ K–bounds a (τ,α)–corner if and only if there exists a $CG(\mathcal{Y})$–loop*

$$[\![a_\alpha]\!]^+\cdot g_1\cdot p_{[\![\zeta]\!]}\cdot g_2\cdot[\![a_\beta]\!]^-\cdot g^{-1} \tag{7}$$

which represents an element of K.

Proof First suppose that there is a $CG(\mathcal{Y})$–loop of the form (7) representing an element of K.

Then the $CG(\mathcal{Y})$–paths

$$p_{[\![\alpha]\!]}\cdot g\cdot p_{[\![\beta]\!]},\quad p_{[\![\alpha]\!]}\cdot[\![a_\alpha]\!]^+\cdot g_1\cdot p_{[\![\zeta]\!]}\cdot g_2\cdot[\![a_\beta]\!]^-\cdot g^{-1}\cdot g\cdot p_{[\![\beta]\!]}$$

differ by an element of K. Thus they together form a loop which lifts to $\mathcal{C}_K$.

This second path is homotopic to a $CG(\mathcal{Y})$–path whose scwolification avoids the vertex $t([\![\alpha]\!])$ after $p_{[\![\alpha]\!]}$ but instead travels across the first three edges of $p_{[\![\alpha]\!]}$, traverses $p_{[\![\zeta]\!]}$, and then travels across the final three edges of $p_{[\![\beta]\!]}$. The homotopy lifts to $\mathcal{C}_K$, and the image in $\mathcal{Z}$ of this homotopy under the scwolification Θ_K shows that there is a 2–cell $\breve{\tau}'$ between the 1–cells $\breve{\alpha}'$ and $\breve{\beta}'$ whose idealizations are the scwolifications of the lifts of $p_{[\![\alpha]\!]}$ and $p_{[\![\beta]\!]}$, respectively. This shows that the path $p_{[\![\alpha]\!]}\cdot g\cdot p_{[\![\beta]\!]}$ K–bounds a (τ,α)–corner.

Conversely, suppose that the $CG(\mathcal{Y})$–path $q = p_{[\![\alpha]\!]} \cdot g \cdot p_{[\![\beta]\!]}$ K–bounds a (τ,α)–corner. Lift to a $\mathcal{C}_K$–path $\hat{q}$ and consider the scwolification $\Theta_K(\hat{q})$ in $\mathcal{Z}$. As in Definition 4.20, the realization of $\hat{q}$ is the concatenation of two 1–cells $\breve{\alpha}'$ and $\breve{\beta}'$ in link$(\breve{\sigma}')$ for some cube $\breve{\sigma}'$ in the orbit of $\breve{\sigma}$. Moreover, there is a 2–cell $\breve{\tau}'$ with $\breve{\alpha}'$ and $\breve{\beta}'$ in the boundary of $\breve{\tau}'$ and an element h of G such that $\breve{\sigma}' = h\breve{\sigma}$, $\breve{\tau}' = h\breve{\tau}$, $\breve{\alpha}' = h\breve{\alpha}$ and $\breve{\beta}' = h\breve{\beta}$. Let v' be the vertex of link$(\breve{\sigma}')$ where $\breve{\alpha}'$ and $\breve{\beta}'$ meet.

Consider the loop $q_0 = [\![a_\alpha]\!]^+ \cdot p_{[\![\zeta]\!]} \cdot [\![a_\beta]\!]^-$ as in Lemma 4.21. This represents an element of $G_{t([\![\alpha]\!])}$, and there is a lift $\hat{q}_0$ of q_0 to $\mathcal{C}_K$ such that $\Theta_K(\hat{q}_0)$ is a loop based at v' and traveling across the corner of $\breve{\tau}'$ from $\breve{\alpha}'$ to $\breve{\beta}'$. The paths q_0 and q have lifts to $\mathcal{C}_K$ forming a subdiagram

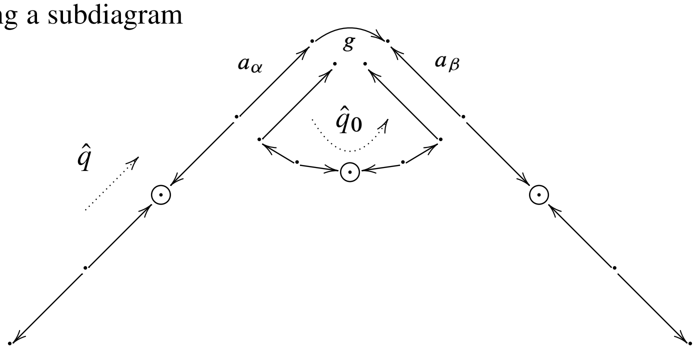

The circled dots represent either single objects or pairs of objects separated by a group arrow, depending on whether the paths $p_{[\![x]\!]}$ have length four or five for $x \in \{\alpha,\beta,\zeta\}$. The scwolification of this diagram in $\mathcal{Z}$ looks like

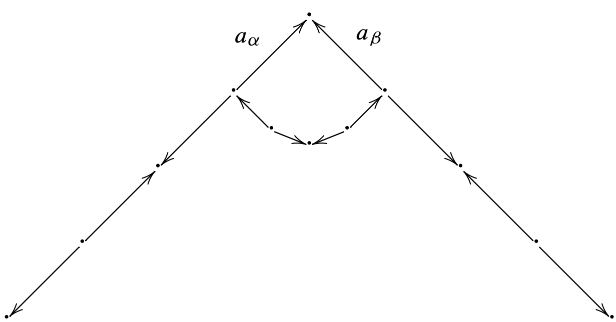

The edges which scwolify to a_α in $\hat{q}$ and $\hat{q}_0$ have sources connected by a group arrow labeled by some g_1. Similarly the edges which scwolify to a_β have sources connected by group arrow with some label g_2. We thus obtain a loop in $\mathcal{C}_K$ of the form (7). □

Given the criterion from Lemma 4.23, the following result is straightforward. Recall the definition of the element $g_{\tau,\alpha}$ from Notation 4.22.

Proposition 4.24 *Suppose that τ is a 2–cell in* $\mathrm{link}(\sigma)$ *and that the boundary of τ is $\alpha.\beta.\gamma$. For any $g \in G_{t[\![\alpha]\!]}$ the $CG(\mathcal{Y})$–path $p_{[\![\alpha]\!]} \cdot g \cdot p_{[\![\beta']\!]}$ K–bounds a (τ,α)–corner if and only if*

(1) $[\![\beta']\!] = [\![\beta]\!]$, *and*

(2) $g \in E_{[\![\alpha]\!]} g_{\tau,\alpha} E_{[\![\bar{\beta}]\!]} . K_{t([\![\alpha]\!])}$

Proof Recall from Notation 4.22 that $g_{\tau,\alpha}$ is the element of $G_{t[\![\alpha]\!]}$ represented by the $CG(\mathcal{Y})$–loop

$$[\![a_\alpha]\!]^+ \cdot p_{[\![\gamma]\!]} \cdot [\![a_\beta]\!]^-.$$

Suppose that $p_{[\![\alpha]\!]} \cdot g \cdot p_{[\![\beta]\!]}$ K–bounds a (τ,α)–corner. Then consider the path

$$[\![a_\alpha]\!]^+ \cdot g_1 \cdot p_{[\![\gamma]\!]} \cdot g_2 \cdot [\![a_\beta]\!]^- \cdot g^{-1}$$

from Lemma 4.23 which represents an element of K.

We have homotopies

$$\begin{aligned}[\![a_\alpha]\!]^+ \cdot g_1 \cdot p_{[\![\gamma]\!]} \cdot g_2 \cdot [\![a_\beta]\!]^- \cdot g^{-1} &\simeq \psi_{[\![a_\alpha]\!]}(g_1) \cdot [\![a_\alpha]\!]^+ \cdot p_{[\![\gamma]\!]} \cdot [\![a_\beta]\!]^- \cdot (\psi_{[\![a_\beta]\!]}(g_2) g^{-1}) \\ &\simeq \psi_{[\![a_\alpha]\!]}(g_1) \cdot g_{\tau,\alpha} \cdot (\psi_{[\![a_\beta]\!]}(g_2) g^{-1}) \\ &\simeq \psi_{[\![a_\alpha]\!]}(g_1) g_{\tau,\alpha} \psi_{[\![a_\beta]\!]}(g_2) g^{-1}.\end{aligned}$$

Since $\psi_{[\![a_\alpha]\!]}(g_1) \in E_{[\![\alpha]\!]}$, $\psi_{[\![a_\beta]\!]}(g_2) \in E_{[\![\beta]\!]}$ and the whole expression above is an element of $K \cap G_{[\![v]\!]} = K_{t([\![\alpha]\!])}$,

$$g \in E_{[\![\alpha]\!]} g_{\tau,\alpha} E_{[\![\beta]\!]} K_{t([\![\alpha]\!])},$$

as required.

In order to prove the other direction, this computation may be performed in reverse. □

Lemma 4.25 *Suppose that* $\mathrm{link}(\breve{\sigma})$ *is simplicial and contains 1–cells $\breve{\alpha}$, $\breve{\beta}$ and $\breve{\gamma}$ which lift respectively to 1–cells α, β and γ in* $\mathrm{link}(\sigma)$ *in X. Let*

$$q = p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2 \cdot p_{[\![\gamma]\!]} \cdot g_3$$

be a $CG(\mathcal{Y})$–loop which represents an element of K. Suppose $\breve{\eta} \subset$ $\mathrm{link}(\breve{\sigma})$ *is the realization of the scwolification of some lift of q to $\mathcal{C}_K$.*

If any one of $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]}$, $p_{[\![\beta]\!]} \cdot g_2 \cdot p_{[\![\gamma]\!]}$ or $p_{[\![\gamma]\!]} \cdot g_3 \cdot p_{[\![\alpha]\!]}$ K–bounds a corner, then $\breve{\eta}$ bounds a 2–cell in $\mathrm{link}(\breve{\sigma})$.

Proof Note that since q represents an element of K, any lift to $\mathcal{C}_K$ is a loop, and so the realization $\breve{\eta}$ is also a loop. Since $\mathrm{link}(\breve{\sigma})$ is simplicial, this loop is embedded of length 3 in $\mathrm{link}(\breve{\sigma})$ by Lemma 4.17.

Think of q as given by a cyclic word in the arrows of $CG(\mathcal{Y})$, and suppose that one of the three given subpaths of q K–bounds a corner. By relabeling and cyclically rotating we can assume it is the subpath $p = p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]}$, so there is some 2–cell $\breve{\tau}$ in $\mathrm{link}(\breve{\sigma})$ and lift τ to $\mathrm{link}(\sigma)$ and p K–bounds a (τ,α)–corner. It follows that some translate $\breve{\tau}'$ of $\breve{\tau}$ in $\mathrm{link}(\breve{\sigma})$ has boundary given by a path $\breve{\alpha}' \cdot \breve{\beta}' \cdot \breve{\gamma}'$, where $\alpha' \cdot \beta'$ are the first two 1–cells of the path $\breve{\eta}$. If the third 1–cell of $\partial\breve{\tau}'$ is not the third 1–cell of $\breve{\eta}$, we obtain 1–cells in $\mathrm{link}(\breve{\sigma})$ with the same endpoints, contradicting the assumption that $\mathrm{link}(\breve{\sigma})$ is simplicial. So η bounds the 2–cell $\breve{\tau}'$. □

Since there are finitely many $\mathrm{Stab}(\breve{\sigma})$–orbits of 2–cells in $\mathrm{link}(\breve{\sigma})$, we obtain the following.

Proposition 4.26 *Suppose that* $\mathrm{link}(\breve{\sigma})$ *is simplicial. There are finitely many* 2*–cells* $\breve{\tau}_i$ *in* $\mathrm{link}(\breve{\sigma})$ (*with boundary* $\breve{\alpha}_i \cdot \breve{\beta}_i \cdot \breve{\gamma}_i$*, and lifts* α_i*,* β_i *and* γ_i *to* $\mathrm{link}(\sigma)$) *such that the following holds*:

Every loop of length 3 *in* $\mathrm{link}(\breve{\sigma})$ *is filled by a* 2*–cell if and only if, for every* $CG(\mathcal{Y})$*–path*

$$p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2 \cdot p_{[\![\gamma]\!]} \cdot g_3 \tag{$*$}$$

which represents an element of K*, there exists an* i *such that*

(1) $[\![\alpha]\!] = [\![\alpha_i]\!]$, $[\![\beta]\!] = [\![\beta_i]\!]$ *and* $[\![\gamma]\!] = [\![\gamma_i]\!]$,
(2) $g_1 \in E_{[\![\alpha_i]\!]} g_{\tau_i,\alpha_i} E_{[\![\overleftarrow{\beta}_i]\!]} K_{t([\![\alpha_i]\!])}$,
(3) $g_2 \in E_{[\![\beta_i]\!]} g_{\tau_i,\beta_i} E_{[\![\overleftarrow{\gamma}_i]\!]} K_{t([\![\beta_i]\!])}$, *and*
(4) $g_3 \in E_{[\![\gamma_i]\!]} g_{\tau_i,\gamma_i} E_{[\![\overleftarrow{\alpha}_i]\!]} K_{t([\![\gamma_i]\!])}$.

Proof Choose the 2–cells $\breve{\tau}_i$ to be representatives of the $\mathrm{Stab}(\breve{\sigma})$–orbits of 2–cells (together with a fixed vertex to label the boundary — so that a single orbit may appear up to three times in the list).

Suppose first that the condition about paths of the form $(*)$ representing elements of K is satisfied, and suppose that p is a loop of length 3 in $\mathrm{link}(\breve{\sigma})$ which is labeled by 1–cells $\breve{\alpha}'$, $\breve{\beta}'$ and $\breve{\gamma}'$, in order. By Lemma 4.18 there exists a $CG(\mathcal{Y})$–path λ of the form $(*)$ which is the label of an unscwolification of p. Because of our hypothesis, there exists an i such that conditions (1)–(4) are satisfied. By Proposition 4.24, the $CG(\mathcal{Y})$–path λ K–bounds a corner at each of its three corners, and so by Lemma 4.25 the path p bounds a 2–cell, as required.

Conversely, suppose that every loop of length 3 in $\mathrm{link}(\breve{\sigma})$ bounds a 2–cell, and consider a $CG(\mathcal{Y})$–path λ of the form $(*)$ which represents an element of K. By Lemma 4.18 the scwolification of a lift of λ is (the idealization of) an immersed path of length 3. This immersed path must then bound a 2–cell $\check{\tau}$. Suppose that $\check{\tau}_i$ is the representative in the $\mathrm{Stab}(\breve{\sigma})$–orbit of the 2–cell $\check{\tau}$, so condition (1) is satisfied. According to Lemma 4.23, applied to all three corners of this 2–cell, the path λ satisfies conditions (2)–(4). □

To summarize, given Lemma 4.8, Lemmas 4.12, 4.16 and 4.18 and Proposition 4.26 give descriptions of various types of $CG(\mathcal{Y})$–paths such that the cube complex $Z = K\backslash X$ is nonpositively curved if and only if no such path lifts to $\mathcal{C}_K$.

5 Algebraic translation

In this section, we continue to work in the context of a group G acting cocompactly on a CAT(0) cube complex X. The induced action on the associated scwol $\mathcal{X}$ has quotient scwol $\mathcal{Y}$, the underlying scwol for a complex of groups structure $G(\mathcal{Y})$ on G. We let $\mathcal{Q}(G)$ be the set of cube stabilizers for $G \curvearrowright X$; equivalently $\mathcal{Q}(G)$ is the set of conjugates of the local groups for the complex of groups $G(\mathcal{Y})$.

We translate the conditions from the previous section into algebraic statements about elements of G and of $\mathcal{Q}(G)$, with an eye toward finding conditions on $K \lhd G$ implying that $K\backslash X$ is nonpositively curved. In Section 6 we use hyperbolic Dehn filling to find K which satisfy the conditions, under certain hyperbolicity assumptions on G and $\mathcal{Q}(G)$.

We fix a basepoint v_0 for $\mathcal{Y}$ and an isomorphism $\pi_1(CG(\mathcal{Y}), v_0) \cong G$ as in Section 2. The scwolification functor

$$\Theta\colon \widetilde{CG(\mathcal{Y})} \to \mathcal{X}$$

is G–equivariant. Recall also that the objects of $\widetilde{CG(\mathcal{Y})}$ are homotopy classes of paths starting at v_0.

Fix also a maximal (undirected) tree T in $\mathcal{Y}$. For each object v of $\mathcal{Y}$ which represents an orbit of cubes in X, let c_v be the unique $\mathcal{Y}$–path in T from v_0 to v. By using scwol arrows, we also consider c_v to be a $CG(\mathcal{Y})$–path in the natural way. For an object v of $\mathcal{Y}$ which represents a chain of cubes of length longer than 1, we define a $\mathcal{Y}$–path c_v

from v_0 to v as follows: if v is represented by $(\sigma_1 \subset \sigma_2 \subset \cdots \subset \sigma_k)$ (a nested chain of cubes in X) then define c_v to be the concatenation of $c_{[\![\sigma_1]\!]}$ with the path consisting of the arrows $(\sigma_1 \subset \cdots \subset \sigma_i) \to (\sigma_1 \subset \cdots \subset \sigma_{i+1})$, for $i = 1, 2, \ldots, k-1$.

We use the paths c_v to define a map from (homotopy classes rel endpoints of) $CG(\mathcal{Y})$–paths to (homotopy classes of) $CG(\mathcal{Y})$–loops based at v_0 by

$$p \mapsto c_{i(p)} \cdot p \cdot \overleftarrow{c}_{t(p)}.$$

(Here and below, $\overleftarrow{c}$ denotes the reverse of the $CG(\mathcal{Y})$–path c.)

Given a path p, let $\ell_p = [c_{i(p)} \cdot p \cdot \overleftarrow{c}_{t(p)}] \in \pi_1(CG(\mathcal{Y}), v_0)$.

The following results are straightforward:

Lemma 5.1 *For any $CG(\mathcal{Y})$–paths p and p' such that $t(p) = i(p')$,*

$$\ell_{\overleftarrow{p}} = \ell_p^{-1}, \quad \ell_{p \cdot p'} = \ell_p \ell_{p'}.$$

Lemma 5.2 *Suppose that p is a $CG(\mathcal{Y})$–path starting at v_0. Let $[p]$ be the equivalence class of p in $\widetilde{CG(\mathcal{Y})}$, and let $x = \Theta([p])$. Then*

$$\mathrm{Stab}_G(x) = \{[p \cdot g \cdot \overleftarrow{p}] \mid g \in G_{[\![x]\!]}\}.$$

Definition 5.3 Given an object v of $\mathcal{Y}$, define

$$Q_v = \{[c_v \cdot g \cdot \overleftarrow{c}_v] \mid g \in G_v\}.$$

Definition 5.3 gives an explicit identification of the local groups of the complex of groups $G(\mathcal{Y})$ with finitely many elements of $\mathcal{Q}(G)$.

5.1 Algebraic formulation of the link conditions

Suppose that $K \trianglelefteq G$. In order for $Z = K \backslash X$ to be nonpositively curved, there are five conditions that need to be ensured on links in Z. Roughly speaking, they are

(1) no loop of length 1,

(2) no loop of length 2 consisting of 1–cells in different G–orbits,

(3) no loop of length 2 consisting of 1–cells in the same G–orbit,

(4) no loop of length 3 whose image in $\mathcal{Y}$ does not bound a 2–cell, and

(5) no loop of length 3 which does not bound a 2–cell but whose image in $\mathcal{Y}$ does bound a 2–cell.

More precisely, the "image in $\mathcal{Y}$" means the image in $\mathcal{Y}$ of the idealization. And we say this image p "bounds a 2–cell" if there is an unscwolification $\hat{p}$ and a lift $\tilde{p}$ of $\hat{p}$ to $\widetilde{CG(\mathcal{Y})}$ such that the realization of the scwolification of $\tilde{p}$ bounds a 2–cell in some link of a cube in X.

If $K\backslash X$ is a simply connected cube complex and we ensure each of these conditions, then Lemmas 4.7, 4.8 and 4.17 imply that $K\backslash X$ is CAT(0).

In this subsection, we formulate five results which give algebraic conditions to enforce each of these five conditions in turn. These results follow quickly from the results in Section 4 using the translation from the beginning of this section. In each case, since G acts cocompactly on a CAT(0) cube complex, there are finitely many G–orbits of links and in each link finitely many G–orbits of each of the five kinds of paths in the above list, and we can rule out each orbit behaving badly in $K\backslash X$ in turn.

Assumption 5.4 The group G acts cocompactly on the CAT(0) cube complex X, and $\mathcal{Q}(G)$ is the collection of cell stabilizers of the action.

Terminology 5.5 Under Assumption 5.4, a normal subgroup $K \trianglelefteq G$ is *cocubical* if $K\backslash X$ is a cube complex.

The following is a straightforward translation of Lemma 4.12. We spell out the proof since we use similar techniques for other more complicated results later in the section.

Theorem 5.6 *Under Assumption 5.4, there exists a finite set $F_1 \subset \mathcal{Q}(G) \times G$ such that for each $(Q, p) \in F_1$ we have $p \notin Q$ and such that, if*

(i) *$K \trianglelefteq G$ is cocubical, and*

(ii) *for each $(Q, p) \in F_1$ we have $p \notin Q.K$,*

then no link in $K\backslash X$ contains a loop of length 1.

Proof Up to the action of G, there are finitely many pairs $(\tilde{\sigma}, \tilde{\alpha})$, where $\tilde{\sigma}$ is a cube of X and $\tilde{\alpha}$ is a 1–cell in $\operatorname{link}(\tilde{\sigma})$ whose endpoints are identified by some element of G. For each such pair we will give a pair (Q, p) as in the statement of the theorem.

For such a pair, let (σ, α) be the image in $K\backslash X$. Since K is assumed to act cocubically, α is embedded in $\operatorname{link}(\sigma)$, except that its endpoints may have been identified, making it a loop. According to Lemma 4.12, α is a loop if and only if there is a $CG(\mathcal{Y})$–loop of the form $p_{[\![\alpha]\!]}.g$ that represents a conjugacy class in K. In particular, this condition only depends on the orbit $[\![\alpha]\!]$ and not on α itself. We associate to α the element $p = \ell_{p_{[\![\alpha]\!]}}$ and the subgroup $Q = Q_{t([\![\alpha]\!])}$, as described in the preamble to this section.

Since X itself is a CAT(0) cube complex, the 1–cell $\tilde{\alpha}$ is not a loop. Applying Lemma 4.12 in the case $K = \{1\}$, we see that $p \notin Q$. On the other hand, to say that $p \notin Q.K$ is the same as saying there is no $CG(\mathcal{Y})$–loop of the form $p_{[\![\alpha]\!]}.g$ which represents an element of K (since in such a $CG(\mathcal{Y})$–loop the element g must be in the local group $G_{t([\![\alpha]\!])}$). □

The next result is an application of Lemma 4.16 to paths of length 2 consisting of 1–cells in different G–orbits (since then the K–nonbacktracking condition is vacuous).

Theorem 5.7 *Under Assumption 5.4 there exists a finite set $F_2 \subset \mathcal{Q}(G)^2 \times G^2$ such that for each $(Q_1, Q_2, p_1, p_2) \in F_2$,*

$$1 \notin p_1 Q_1 p_2 Q_2,$$

and such that, if

(i) *$K \trianglelefteq G$ is cocubical, and*

(ii) *for each $(Q_1, Q_2, p_1, p_2) \in F_2$,*

$$K \cap p_1 Q_2 p_2 Q_2 = \varnothing,$$

then every loop of length 2 in a link in $K\backslash X$ consists of 1–cells in the same G–orbit.

Proof The proof is similar to the proof of Theorem 5.6 above. Lemma 4.16 implies that it is enough to verify that no link in a cube of $K\backslash X$ contains a pair of 1–cells α and β in distinct G–orbits $[\![\alpha]\!]$ and $[\![\beta]\!]$ such that there is a $CG(\mathcal{Y})$–loop $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2$ representing an element of K.

There are finitely many pairs of such orbits, and to each such pair we can associate the elements $p_1 = \ell_{p_{[\![\alpha]\!]}}$, $p_2 = \ell_{p_{[\![\beta]\!]}}$, $Q_1 = Q_{t([\![\alpha]\!])}$ and $Q_2 = Q_{t([\![\beta]\!])}$.

Since X is a CAT(0) cube complex, there are no nonbacktracking loops of length 2 in any links in X, so applying Lemma 4.16 with $K = \{1\}$ we see that $1 \notin p_1 Q_1 p_2 Q_2$. The result now follows from Lemma 4.16 with our choice of K. □

For paths of length 2 consisting of 1–cells in the same G–orbit, the condition is slightly more complicated, as K–backtracking paths are possible.

Theorem 5.8 *Under Assumption 5.4 there exists a finite set $F_3 \subset \mathcal{Q}(G)^2 \times G^2$ such that, for each $(Q_1, Q_2, p_1, p_2) \in F_3$,*

$$1 \notin p_1(Q_1 - Q_2^{p_2}) p_2(Q_2 - Q_1^{p_1}), \tag{8}$$

and such that, if

(i) *$K \trianglelefteq G$ is cocubical,*

(ii) *no link in $K\backslash X$ contains a loop of length 1, and*

(iii) *for every $(Q_1, Q_2, p_1, p_2) \in F_3$,*

$$K \cap p_1\big(Q_1 - (Q_2^{p_2}(K \cap Q_1))\big) p_2\big(Q_2 - (Q_1^{p_1}(K \cap Q_2))\big) = \varnothing, \tag{9}$$

then no link in $K\backslash X$ contains an immersed loop of length 2 consisting of 1–cells in the same orbit.

Proof Because of assumptions (i) and (ii) we only need to be concerned with the following situation: there is some cube $\tilde{\sigma}$ of X and some 1–cell $\tilde{\alpha}$ in its link such that

(1) there is some $g \in G$ such that g fixes $t(\tilde{\alpha})$ but not $\tilde{\alpha}$;

(2) there is some $h \in G$ such that $h(i(\tilde{\alpha})) = i(g\tilde{\alpha})$ but $h^{-1}g\tilde{\alpha} \neq \tilde{\alpha}$.

There are finitely many orbits of pairs $(\tilde{\sigma}, \tilde{\alpha})$ of this type. For each orbit we pick a representative, and describe an element of $\mathcal{Q}(G)^2 \times G^2$ as in the theorem. If (9) is satisfied for this element, then $K\backslash X$ will contain no immersed loop of length 2 consisting of 1–cells in the orbit of $\tilde{\alpha}$.

We apply Lemma 4.16 to a path of length 2 of the form $\alpha.\alpha'$ where α is the image of $\tilde{\alpha}$ in $K\backslash X$ and α' is the (oppositely oriented) image of a translate of $\tilde{\alpha}$ by an element of the stabilizer of $\tilde{\sigma}$. Any immersed loop of the type we are trying to rule out gives rise to a K–nonbacktracking $CG(\mathcal{Y})$–loop $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\bar{\alpha}]\!]} \cdot g_2$ representing a conjugacy class in K. We let $p_1 = l_{p_{[\![\alpha]\!]}}$, $p_2 = l_{p_{[\![\bar{\alpha}]\!]}}$, $Q_1 = Q_{t([\![\alpha]\!])}$ and $Q_2 = Q_{i([\![\alpha]\!])}$. Using Lemma 4.15, the loop $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\bar{\alpha}]\!]} \cdot g_2$ is K–nonbacktracking if and only if $g_1 \notin E_{[\![\alpha]\!]} K_{t([\![\alpha]\!])}$ and $g_2 \notin E_{[\![\bar{\alpha}]\!]} K_{i([\![\alpha]\!])}$. The subgroup of Q_1 corresponding to $E_{[\![\bar{\alpha}]\!]}$ is equal to $Q_1 \cap Q_2^{p_2}$, and the subgroup of Q_2 corresponding to $E_{[\![\bar{\alpha}]\!]}$ is $Q_2 \cap Q_1^{p_1}$. Thus an element $p_1 q_1 p_2 q_2$ of $p_1 Q_1 p_2 Q_2$ comes from a K–nonbacktracking $CG(\mathcal{Y})$–loop if and only if $q_1 \notin Q_2^{p_2}(K \cap Q_1)$ and $q_2 \notin Q_1^{p_1}(K \cap Q_2)$. Applying Lemmas 4.15 and 4.16 in the case $K = \{1\}$ and $K\backslash X = X$ is CAT(0), we see that our tuple satisfies (8). For an arbitrary K we see that when (9) is satisfied, there is no immersed loop of length 2 in a link in $K\backslash X$ consisting of images of translates of $\tilde{\alpha}$. □

In order to apply Lemma 4.17, in each of the following two results we make the extra assumption that K is such that no link in $K\backslash X$ contains a loop of length 1 or 2. The following result is a translation of Lemma 4.18.

Theorem 5.9 *Under Assumption 5.4 there exists a finite set $F_4 \subset \mathcal{Q}(G)^3 \times G^3$ such that, for each $(Q_1, Q_2, Q_3, p_1, p_2, p_3) \in F_4$,*

$$1 \notin p_1 Q_1 p_2 Q_2 p_3 Q_3$$

and such that if

(i) *$K \trianglelefteq G$ is cocubical,*

(ii) *no link in $K\backslash X$ contains a loop of length 1 or 2, and*

(iii) *for all $(Q_1, Q_2, Q_3, p_1, p_2, p_3) \in F_4$,*

$$K \cap p_1 Q_1 p_2 Q_2 p_3 Q_3 = \varnothing,$$

then every loop of length 3 in a link of $K\backslash X$ has image in $\mathcal{Y}$ which bounds a 2–cell.

Proof Condition (ii) and Lemma 4.17 imply that it suffices to consider immersed loops of length 3 in links in $K\backslash X$. For each choice of triple of G–orbits $[\![\alpha]\!], [\![\beta]\!], [\![\gamma]\!]$ of 1–cells in links in X whose image in $\mathcal{Y}$ forms a loop, but whose image does not bound a 2–cell in $\mathcal{Y}$ (in the sense described at the beginning of this subsection), we proceed as follows. We associate the elements $p_1 = l_{p_{[\![\alpha]\!]}}$, $p_2 = l_{p_{[\![\beta]\!]}}$, $p_3 = l_{p_{[\![\gamma]\!]}}$, $Q_1 = Q_{t([\![\alpha]\!])}$, $Q_2 = Q_{t([\![\beta]\!])}$ and $Q_3 = Q_{t([\![\gamma]\!])}$.

Since X is a CAT(0) cube complex, we can apply Lemma 4.18 to see that

$$1 \notin p_1 Q_1 p_2 Q_2 p_3 Q_3.$$

Now let $K \trianglelefteq G$ be cocubical, and satisfy conditions (i)–(iii) from the statement. Condition (iii) implies that condition (2) from Lemma 4.18 does not hold, and by that lemma there is no immersed loop of length 3 in a link in $K\backslash X$ whose image in $\mathcal{Y}$ is $[\![\alpha]\!], [\![\beta]\!], [\![\gamma]\!]$.

Since there are finitely many such triples $[\![\alpha]\!], [\![\beta]\!], [\![\gamma]\!]$, the theorem follows. □

Finally, we deal with loops of length 3 in links in $K\backslash X$ whose image in $\mathcal{Y}$ does bound a 2–cell.

Terminology 5.10 Suppose that

$$A = (Q_1, Q_2, Q_3, p_1, p_2, p_3, h_1, h_2, h_3) \in \mathcal{Q}(G)^3 \times G^6.$$

With indices read mod 3, let

$$A_i^- = Q_{i-1}^{p_{i-1}} \cap Q_i, \quad A_i^+ = Q_i \cap Q_{i+1}^{p_{i+1}}.$$

Furthermore, let

$$B_i = A_i^- h_i A_i^+.$$

Using this terminology, we have the following translation of Proposition 4.26.

Theorem 5.11 *Under Assumption 5.4 there exists a finite set $F_5 \subseteq \mathcal{Q}(G)^3 \times G^6$ such that, for each $A = (Q_1, Q_2, Q_3, p_1, p_2, p_3, h_1, h_2, h_3)$,*

$$1 \notin p_1(Q_1 - B_1)p_2(Q_2 - B_2)p_3(Q_3 - B_3)$$

and such that, if

(i) *$K \trianglelefteq G$ is cocubical,*

(ii) *no link in $K \backslash X$ contains a loop of length 1 or 2, and*

(iii) *for all $(Q_1, Q_2, Q_3, p_1, p_2, p_3, h_1, h_2, h_3) \in F_5$,*

$$K \cap p_1(Q_1 - B_1(K \cap Q_1))p_2(Q_2 - B_2(K \cap Q_2))p_3(Q_3 - B_3(K \cap Q_3)) = \varnothing,$$

then no link in $K \backslash X$ contains a loop of length 3 which does not bound a 2–cell but whose image in $\mathcal{Y}$ bounds a 2–cell.

Proof For each choice of triple of orbits $[\![\alpha]\!], [\![\beta]\!], [\![\gamma]\!]$ whose image in $\mathcal{Y}$ bounds a 2–cell (in the sense described at the beginning of this subsection), we proceed as follows. Without loss of generality we choose representatives α, β and γ of these orbits such that there is a 2–cell τ with boundary $\alpha \cdot \beta \cdot \gamma$. We associate the elements $p_1 = \ell_{p_{[\![\alpha]\!]}}$, $p_2 = \ell_{p_{[\![\beta]\!]}}$, $p_3 = \ell_{p_{[\![\gamma]\!]}}$, $Q_1 = Q_{t([\![\alpha]\!])}$, $Q_2 = Q_{t([\![\beta]\!])}$, $Q_3 = Q_{t([\![\gamma]\!])}$, $h_1 = [c_{t([\![\alpha]\!])} \cdot g_{\tau,\alpha} \cdot \overleftarrow{c}_{t([\![\alpha]\!])}]$, $h_2 = [c_{t([\![\beta]\!])} \cdot g_{\tau,\beta} \cdot \overleftarrow{c}_{t([\![\beta]\!])}]$ and $h_3 = [c_{t([\![\gamma]\!])} \cdot g_{\tau,\gamma} \cdot \overleftarrow{c}_{t([\![\gamma]\!])}]$. Once again, since X is a CAT(0) cube complex, we can apply Proposition 4.26 to see that

$$1 \notin p_1(Q_1 - B_1)p_2(Q_2 - B_2)p_3(Q_3 - B_3).$$

When conditions (2)–(4) from the statement of Proposition 4.26 are translated into statements about the group G, we get exactly $g_1 \in B_1(K \cap Q_1)$, etc, which gives the statement in the conclusion of the result.

Since there are finitely many such triples $[\![\alpha]\!], [\![\beta]\!], [\![\gamma]\!]$, the theorem follows. □

6 Dehn filling

In this section we prove some results about group-theoretic Dehn filling. Theorem 6.5 gives a "weak separability" of certain multicosets, and generalizations of multicosets, and is used to find subgroups K which satisfy the conditions from Theorems 5.6–5.11. Theorem 6.5 may be of independent interest, and we expect it to have applications beyond the scope of this paper. The second main result of this section is Theorem 6.9, from which Theorem F from the introduction follows quickly by induction.

6.1 Dehn fillings

Let $(G,\mathcal{P})$ be a group pair, and let $\mathcal{N}=\{N_P \lhd P \mid P\in\mathcal{P}\}$ be a choice of normal subgroups of the peripheral groups. The collection $\mathcal{N}$ determines a *(Dehn) filling* $(\overline{G},\overline{\mathcal{P}})$ *of* $(G,\mathcal{P})$, where $\overline{G}=G/K$ for K the normal closure of $\bigcup\mathcal{N}$, and $\overline{\mathcal{P}}$ is equal to the collection of images of elements of $\mathcal{P}$ in $\overline{G}$. The elements of $\mathcal{N}$ are called *filling kernels*. We sometimes write such a filling as

$$\pi\colon (G,\mathcal{P})\to(\overline{G},\overline{\mathcal{P}}),$$

omitting mention of the particular filling kernels.

If $N_P \mathrel{\dot{<}} P$ (ie N_P is finite index in P) for all $P\in\mathcal{P}$, we say that the filling is *peripherally finite*. If $H<G$ and, for all $g\in G$, $|H\cap P^g|=\infty$ implies $N_P^g\subseteq H$, then the filling is an *H–filling*. If $\mathcal{H}$ is a family of subgroups, the filling is an *$\mathcal{H}$–filling* whenever it is an H–filling for every $H\in\mathcal{H}$.

A property P holds *for all sufficiently long fillings* of $(G,\mathcal{P})$ if there is a finite set $S\subseteq\bigcup\mathcal{P}-\{1\}$ such that P holds whenever $\left(\bigcup\mathcal{N}\right)\cap S=\varnothing$. It is frequently useful to restrict attention to specific types of fillings (peripherally finite, H–fillings, etc). If A is a property of fillings we say that P holds for *all sufficiently long* A*–fillings* if, for all sufficiently long fillings, either P holds or A does not hold.

6.2 Relatively hyperbolic group pairs

We refer the reader to [17] for background on relatively hyperbolic groups. In that paper, given a group pair $(G,\mathcal{P})$ (consisting of finitely generated groups) a space called the *cusped space* is built, which is δ–hyperbolic (for some δ) if and only if $(G,\mathcal{P})$ is relatively hyperbolic. The cusped space is built by attaching *combinatorial horoballs* to a Cayley graph for G. Each combinatorial horoball H has vertex set $tP\times\mathbb{Z}_{\geq 0}$ for some coset tP of some $P \in \mathcal{P}$, and is hyperbolic. A vertex (g, n) of such a horoball is said to have *depth* n. The depth 0 vertices of the cusped space are exactly the vertices of the Cayley graph; if two vertices are connected by an edge, then their depths differ by at most one. See [17, Section 3] for more information about the construction and geometry of the cusped space. The following result is essentially contained in [7, Theorem 7.11].

Theorem 6.1 *Suppose that G is a hyperbolic group and that $\mathcal{P}$ is a finite collection of subgroups of G. Then $(G, \mathcal{P})$ is relatively hyperbolic if and only if $\mathcal{P}$ is an almost malnormal family of quasiconvex subgroups.*

Recall that $\mathcal{P} = \{P_1, \dots, P_n\}$ is *almost malnormal* if whenever $P_i \cap P_j^g$ is infinite, we have $i = j$ and $g \in P_i$.

We can use the notion of height (see Definition 3.28) to measure how far away a family of subgroups is from being almost malnormal.

We now define the *induced peripheral structure* on G associated to a finite collection of quasiconvex subgroups of a hyperbolic group, in analogy with the construction from [2, Section 3.1].

Definition 6.2 Suppose that G is a hyperbolic group and $\mathcal{H}$ is a finite collection of quasiconvex subgroups of G. The *peripheral structure on G induced by $\mathcal{H}$* is obtained as follows:

Start by taking the collection of minimal infinite subgroups of the form

$$H_1 \cap H_2^{g_2} \cap \cdots \cap H_k^{g_k}$$

where the H_i are in $\mathcal{H}$ and the cosets $\{H_1, g_2H_2, \dots, g_kH_k\}$ are all distinct. Replace each element in this collection by its commensurator in G, and then choose one from each G–conjugacy class. The resulting collection $\mathcal{P}$ is the induced peripheral structure.

If $H \in \mathcal{H}$ then the *induced peripheral structure on H with respect to $\mathcal{H}$* is a choice of H–conjugacy representatives of intersections with H of G–conjugates of elements of $\mathcal{P}$.

We remark that the fact that there is a bound on the number k of g_iH_i as above follows from Proposition 3.29.

To state the next lemma we need a definition from [2]:

Definition 6.3 Let $H < G$ and suppose $(H, \mathcal{D})$ and $(G, \mathcal{P})$ are relatively hyperbolic. Suppose furthermore that every $D \in \mathcal{D}$ is conjugate into some $P \in \mathcal{P}$. Then — see [2, Lemma 3.1] — there is an induced H–equivariant map from the cusped space of $(H, \mathcal{D})$ to the cusped space of $(G, \mathcal{P})$. The subgroup $(H, \mathcal{D})$ is *relatively quasiconvex* if this induced map has quasiconvex image.

In [28, Appendix A] it is proved that this is the same notion as the various notions of relative quasiconvexity discussed by Hruska in [23].

The following can be proved in the same way as [2, Proposition 3.12].

Lemma 6.4 *Suppose that G is hyperbolic and $\mathcal{H}$ is a finite collection of quasiconvex subgroups of G.*

(1) *The induced peripheral structure $\mathcal{P}$ is a finite collection of groups. The pair $(G, \mathcal{P})$ is relatively hyperbolic.*
(2) *If $H \in \mathcal{H}$ then the induced peripheral structure $\mathcal{D}$ of H with respect to $\mathcal{H}$ is finite. The pair $(H, \mathcal{D})$ is relatively hyperbolic.*
(3) *For any $H \in \mathcal{H}$, the pair $(H, \mathcal{D})$ is full relatively quasiconvex in $(G, \mathcal{P})$.*

A subgroup H is *full* if whenever P is a parabolic subgroup such that $H \cap P$ is infinite we have $H \cap P \dot{<} P$.

6.3 The appropriate metacondition

The goal of this subsection is to prove Theorem 6.5 below. The special case that $n = 1$ and $S_1 = \varnothing$ is [2, Proposition 4.5], which is about keeping elements out of full quasiconvex subgroups when performing long Dehn fillings. Here we generalize to multicosets of full quasiconvex subgroups, possibly with some elements deleted. Although the present result is more general, our proof is simpler, using the more appealing "Greendlinger Lemma"-type Theorem 6.7 below in place of the somewhat technical [2, Lemmas 4.1 and 4.2].[6]

Theorem 6.5 *Let $(G, \mathcal{P})$ be relatively hyperbolic, and let $\mathcal{Q}$ be a collection of full relatively quasiconvex subgroups. For $1 \le i \le n$, let $p_i \in G$, $Q_i \in \mathcal{Q}$ and $S_i \subseteq Q_i$ be chosen to satisfy*

$$1 \notin p_1(Q_1 - S_1) \cdots p_n(Q_n - S_n). \tag{10}$$

[6] Using such a Greendlinger Lemma in place of the results of [2] was suggested to us by Alessandro Sisto while we were collaborating on [19].

Then, for sufficiently long $\mathcal{Q}$–fillings $G \to G/K$, the kernel K contains no element of the form

$$p_1 t_1 \cdots p_n t_n \tag{11}$$

where $t_i \in Q_i - ((K \cap Q_i)S_i)$.

The five conditions in the conclusions of Theorems 5.6–5.11 each fall into the scheme of the conditions in Theorem 6.5. Therefore, we may apply Theorem 6.5 to obtain the following result. We remark that the following result is stated in the generality of relatively hyperbolic groups acting cocompactly on cube complexes with full relatively quasiconvex subgroups. This is greater generality than is strictly required for the proof of Theorem F. However, we believe that this extra generality will be of use in future work, and should be of independent interest.

Corollary 6.6 *Suppose that $(G, \mathcal{P})$ is relatively hyperbolic and that G acts cocompactly on the CAT(0) cube complex X. Suppose every parabolic element of G fixes some point of X, and that cell stabilizers are full relatively quasiconvex. Let $\sigma_1, \dots, \sigma_k$ be representatives of the G–orbits of cubes of X. For each i let Q_i be the finite-index subgroup of* $\mathrm{Stab}(\sigma_i)$ *consisting of elements which fix σ_i pointwise. Let $\mathcal{Q} = \{Q_1, \dots, Q_k\}$.*

For sufficiently long $\mathcal{Q}$–fillings

$$G \to \overline{G} = G(N_1, \dots, N_m)$$

of $(G, \mathcal{P})$, with kernel K, the quotient $K \backslash X$ is a CAT(0) cube complex.

Proof The kernels of Dehn fillings are always generated by parabolic elements, and the parabolic elements act elliptically by assumption. Thus the kernel of *any* Dehn filling is generated by elliptic elements, so $K \backslash X$ is simply connected by Theorem 4.1. For sufficiently long $\mathcal{Q}$–fillings the fact that $G_{\sigma_i} \cap K \le Q_i$ follows from [2, Proposition 4.4], so by Proposition 4.3 for such fillings $K \backslash X$ is a cube complex. Therefore, we may assume that the subgroup K is cocubical (in the sense of Terminology 5.5).

It remains to show that for sufficiently long $\mathcal{Q}$–fillings $K \backslash X$ is nonpositively curved. It follows from Theorems 5.6–5.8 and 6.5 that for sufficiently long $\mathcal{Q}$–fillings each link of each cell in $K \backslash X$ is simplicial. Thus it follows from Theorems 5.9, 5.11 and 6.5 that for sufficiently long $\mathcal{Q}$–fillings, each link of each cell in $K \backslash X$ is also flag, which means that $K \backslash X$ is nonpositively curved by Theorem 4.6. □

To prove Theorem 6.5, we use the following "Greendlinger Lemma" — cf [19, Lemma 2.26].

Theorem 6.7 *Let $C_1, C_2 > 0$. Suppose that $(G, \mathcal{P})$ is relatively hyperbolic, with cusped space X. For all sufficiently long fillings $G \to G/K$, and any geodesic γ in X joining 1 to $g \in K - \{1\}$, there is a horoball A such that*

1. *γ contains a depth C_1 vertex of A, and*
2. *there is an element k of K stabilizing A such that, for two points $a, b \in A$ and lying on γ at depth at least C_1, $d(a, kb) < d(a, b) - C_2$ (in particular, $d(1, kg) < d(1, g) - C_2$).*

Proof Let $\delta > 0$ be such that X is δ–hyperbolic, and so are the cusped spaces for sufficiently long fillings (that there exists such a δ is [2, Proposition 2.3]). We only consider such fillings, without further mention of this assumption.

Now choose L and ϵ such that every L–local $(1, C_2)$–quasigeodesic lies within an ϵ–neighborhood of any geodesic with the same endpoints. (Such L and ϵ only depend on δ and C_2; see [11, Chapter 3].)

Now choose a filling long enough that every $(2L+C_1+2\epsilon)$–ball centered on the Cayley graph embeds in the quotient cusped space. Let K be the kernel of the filling, and choose $g \in K - \{1\}$. Let γ be a geodesic from 1 to g, and let $\bar{\gamma}$ be the projection to the cusped space $K \backslash X$ for G/K. Within an $(L+C_1+2\epsilon)$–neighborhood of the Cayley graph, $\bar{\gamma}$ is an L–local geodesic. But $\bar{\gamma}$ cannot be an L–local $(1, C_2)$–quasigeodesic everywhere, since it is a loop with diameter larger than ϵ.

In particular, there is a subsegment σ of $\bar{\gamma}$ of length $l \le L$ such that the endpoints $\bar{a}$ and $\bar{b}$ of σ are less than $l - C_2$ apart. This subsegment σ must moreover lie in the image of a single horoball.

The corresponding points a and b on γ lie at depth at least C_1 in a horoball A of X. Since $d(\bar{a}, \bar{b}) < l - C_2$, there is some element $k \in K$ stabilizing A such that $d(a, kb) < l - C_2$, as desired. □

The following result follows immediately from [28, A.6].

Lemma 6.8 *Suppose that $(G, \mathcal{P})$ is relatively hyperbolic with cusped space X and that $(H, \mathcal{D}) \le (G, \mathcal{P})$ is a full relatively quasiconvex subgroup. There exists a constant κ satisfying the following:*

Suppose that $g \in G$ and that $x_1, x_2 \in gH$. Suppose that γ is a geodesic in X between x_1 and x_2. Further, suppose that aP (for $a \in G$ and $P \in \mathcal{P}$) is a coset such that γ intersects the horoball corresponding to aP to depth at least κ. Then P is infinite and $P^a \cap H^g$ has finite index in P^a.

Proof of Theorem 6.5 Let X be the cusped space associated to $(G, \mathcal{P})$ and suppose that X is δ–hyperbolic. Let C_2 be any positive number, and let

$$C_1 = \max\{|p_i|, \kappa\} + 2(n+100)\delta,$$

where κ is the constant from Lemma 6.8 above. Suppose that K is the kernel of a filling which is long enough to satisfy the conclusion of Theorem 6.7 with these constants.

In order to obtain a contradiction, suppose that there is an element $g \in K$ which is of the form

$$g = p_1 t_1 \cdots p_n t_n,$$

where $t_i \in Q_i - ((K \cap Q_i)S_i)$, and suppose that g is chosen such that $d_X(1, g)$ is minimal amongst all such choices.

Since for each i we have $Q_i - ((K \cap Q_i)S_i) \subseteq Q_i - S_i$, the assumption of the theorem implies that $g \neq 1$. We can represent the equation $g = p_1 t_1 \cdots t_n p_n$ by a geodesic $(2n{+}1)$–gon in X, joining the appropriate elements of the Cayley graph in turn by X–geodesics. Let γ be the geodesic for g, ρ_i the geodesic for p_i and τ_i the geodesic for t_i.

Since $g \in K - \{1\}$, by Theorem 6.7 there exist a horoball A in X, an element $k \in K$ stabilizing A, and points a and b on γ at depth at least C_1 such that k stabilizes A and $d(a, kb) < d(a, b) - C_2$. In particular, we have $d(x, kgx) < d(x, gx) - C_2$. The geodesic $(2n{+}1)$–gon is $(2n{-}1)\delta$–thin, so b lies within distance $(2n-1)\delta$ of some side other than γ. The paths ρ_i do not go deep enough into any horoballs to be this close to b, so b lies within $(2n-1)\delta$ of some point b' on some τ_i. By the choice of C_1, b' lies at depth at least κ in A.

Write $A = aP$ for some $P \in \mathcal{P}$. Note that τ_i is a geodesic between two points in the coset $p_1 t_1 \cdots p_i Q_i$. By Lemma 6.8, $P^a \cap Q_i^{p_1 t_1 \cdots p_i}$ has finite index in P^a. Since the filling is a $\mathcal{Q}$–filling, we have that $k \in Q_i^{p_1 t_1 \cdots p_i}$.

Let $k' = k^{(p_1 t_1 \cdots p_i)^{-1}}$, and let $t_i' = k' t_i$. Then $k' \in K \cap Q_i$.

Note that $kg = p_1 t_1 \cdots p_i (k' t_i) p_{i+1} \cdots p_n t_n$. Since $t_i \notin (K \cap Q_i)S_i$, we have that $t_i' \notin (K \cap Q_i)S_i$. Therefore, the element kg is another element of the required form, contradicting the choice of g as the shortest such. □

6.4 Dehn fillings which induce CAT(0) quotient cube complexes

Theorem 6.9 *Suppose that the hyperbolic group G acts cocompactly on the CAT(0) cube complex X, and that cell stabilizers are virtually special and quasiconvex. Let $\sigma_1, \dots, \sigma_k$ be representatives of the G–orbits of cubes of X, and for each i let Q_i be the finite-index subgroup of $\mathrm{Stab}(\sigma_i)$ consisting of elements which fix σ_i pointwise. Let $\mathcal{Q} = \{Q_1, \dots, Q_k\}$, and let $\mathcal{P}$ be the peripheral structure on G induced by $\mathcal{Q}$, as in Definition 6.2.*

If some element of $\mathcal{Q}$ is infinite, then there exists a Dehn filling

$$G \twoheadrightarrow \overline{G} = G(N_1, \dots, N_m)$$

of $(G, \mathcal{P})$, with kernel $K_{\mathcal{P}}$ such that

1. (1) *$\overline{G}$ is hyperbolic;*
2. (2) *$\overline{\mathcal{Q}}$ consists of virtually special quasiconvex subgroups of $\overline{G}$;*
3. (3) *$K_{\mathcal{P}}$ is generated by elements in cell stabilizers;*
4. (4) *for each i, we have $K_{\mathcal{P}} \cap \mathrm{Stab}(\sigma_i) \le Q_i$;*
5. (5) height$(\overline{\mathcal{Q}}) <$ height$(\mathcal{Q})$;
6. (6) *$K_{\mathcal{P}} \backslash X$ is a CAT(0) cube complex.*

Proof Let G, X, $\mathcal{Q}$ and $\mathcal{P}$ be as in the statement of the theorem. By Lemma 6.4, $(G, \mathcal{P})$ is relatively hyperbolic. Moreover, for each $Q \in \mathcal{Q}$, the induced structure $\mathcal{D}_Q$ on Q makes $(Q, \mathcal{D}_Q)$ relatively hyperbolic, and Q is full relatively quasiconvex in $(G, \mathcal{P})$. Note that the assumption that some element of $\mathcal{Q}$ is infinite implies (by the definition of $\mathcal{P}$) that some element of $\mathcal{P}$ is infinite.

Property (1) holds for sufficiently long peripherally finite fillings of $(G, \mathcal{P})$ by the basic result of relatively hyperbolic Dehn fillings [31, Theorem 1.1]. We always assume that we have taken a filling such that $\overline{G}$ is hyperbolic.

We remark that, because each element of $\mathcal{Q}$ is finite-index in a cell stabilizer, each element of $\mathcal{Q}$ is hyperbolic and virtually special. Moreover, since each element of $\mathcal{P}$ has a finite-index subgroup which is a quasiconvex subgroup of some element of $\mathcal{Q}$ by construction, each element of $\mathcal{P}$ is also hyperbolic and virtually special. In particular, each element of $\mathcal{P}$ is residually finite. We choose particular fillings with $N_i \dot{<} P_i$, and residual finiteness guarantees the existence of the fillings that we seek.

We now explain how to ensure the properties of the conclusion of the result.

Suppose that $Q \in \mathcal{Q}$. Since $\mathcal{P}$ is the peripheral structure induced by $\mathcal{Q}$, we can choose finite-index subgroups of elements of $\mathcal{P}$ which induce $\mathcal{Q}$–fillings, and any such filling $\overline{G}$ of G naturally induces a filling $\overline{Q}$ of Q. By the Malnormal Special Quotient Theorem [36, Theorem 12.3] — see also [3, Corollary 2.8] — for each $P_i \in \mathcal{P}$ there is a subgroup $\dot{P}_i(Q) \dot{\lhd} P_i$ such that if each filling kernel N_i satisfies $N_i \leq \dot{P}_i(Q)$ then the induced filling $\overline{Q}$ is virtually special (and hyperbolic). Let $\dot{P}_i$ be the intersection of the $\dot{P}_i(Q)$ for all $Q \in \mathcal{Q}$. Thus, if we choose filling kernels $N_i \leq \dot{P}_i$ then each of the induced fillings of each element of $\mathcal{Q}$ is virtually special. By [18, Proposition 4.6], the natural map from $\overline{Q}$ to $\overline{G}$ is injective for all sufficiently long fillings.[7] If we choose a sufficiently long peripherally finite filling of $(G, \mathcal{P})$ with $N_i \leq \dot{P}_i$ then [18, Proposition 4.5] implies that each $\overline{Q}$ is quasiconvex in $\overline{G}$. This ensures property (2).

For the remaining properties, we show that they hold for sufficiently long peripherally finite $\mathcal{Q}$–fillings of $(G, \mathcal{P})$. Therefore, to ensure that all of the properties hold, it suffices to take a sufficiently long $\mathcal{Q}$–filling with each $N_i \dot{<} \dot{P}_i$.

Property (3) holds automatically for any $\mathcal{Q}$–filling, since $K_{\mathcal{P}}$ is generated by conjugates of elements in $\mathcal{Q}$, and each such conjugate lies in a cell stabilizer.

We now explain how to ensure each of the remaining properties in turn for sufficiently long $\mathcal{Q}$–fillings.

For property (4), suppose that $\mathcal{F}_i \sqcup \{1\}$ is a set of coset representatives for Q_i in $\mathrm{Stab}(\sigma_i)$. To ensure that (4) holds, it suffices to keep (the image of) each element of $\mathcal{F}_i$ out of the image of $\mathrm{Stab}(\sigma_i)$ in $\overline{G}$. This is true for sufficiently long $\mathcal{Q}$–fillings by [1, Theorem A.43.4], because Q_i has finite index in $\mathrm{Stab}(\sigma_i)$.

Property (5) holds for sufficiently long peripherally finite $\mathcal{Q}$–fillings of $(G, \mathcal{P})$ by an entirely analogous argument to that of [1, Theorem A.47].

Finally, property (6) holds for sufficiently long $\mathcal{Q}$–fillings by Corollary 6.6. □

The group $\overline{G}$ as above acts isometrically on $\overline{X} = K_{\mathcal{P}} \backslash X$ with quotient naturally isomorphic (as a topological space, but not as a complex of groups) to $G \backslash X$. Therefore, if the action of $\overline{G}$ on $\overline{X}$ is not proper, we can apply Theorem 6.9 to this action, to obtain a further quotient. By induction on height, we obtain the following result from the introduction.

[7] Lemma 3.7 of [18] ensures that sufficiently long $\mathcal{Q}$–fillings are sufficiently wide, in the terminology of that paper.

Theorem F *Suppose that the hyperbolic group G acts cocompactly on a CAT(0) cube complex X and that cell stabilizers are virtually special and quasiconvex. There exists a quotient $\overline{G} = G/K$ such that*

(1) *the quotient $K\backslash X$ is a CAT(0) cube complex;*

(2) *the group $\overline{G}$ is hyperbolic; and*

(3) *the action of $\overline{G}$ on $K\backslash X$ is proper (and cocompact).*

Appendix A quasiconvexity criterion

In this appendix, we give a criterion (Theorem A.3) for a possibly infinite union of quasiconvex sets in a hyperbolic space to be quasiconvex. This criterion is used in the forward direction of Theorem A: quasiconvex cell stabilizers imply quasiconvex hyperplane stabilizers. This criterion may be of independent interest.

Since any subset is a union of points, clearly some assumptions are needed.

We begin with a basic lemma about finite unions of quasiconvex subsets.

Lemma A.1 *Suppose that Y is δ–hyperbolic, and $P \subset Y$ is a union of k ϵ–quasiconvex subsets $P_1, \dots, P_k$ such that $P_i \cap P_{i+1} \neq \varnothing$ for each i. Then P is ρ–quasiconvex, where*

$$\rho = \delta(\log_2(k) + 1) + \epsilon.$$

Proof Consider a pair of points $x \in P_r$ and $y \in P_s$. Without loss of generality, assume that $r < s$ (the case $r = s$ being straightforward).

Now choose a sequence of points $p_i \in P_i \cap P_{i+1}$ for $r \le i < s$, let σ be a geodesic between x and y and let u be a point on σ. Our task is to bound the distance from u to P.

Consider the broken geodesic $\gamma = [x, p_r, p_{r+1}, \dots, p_{s-1}, y]$. Since the P_i are ϵ–quasiconvex, γ is contained in an ϵ–neighborhood of $P_r \cup \dots \cup P_s \subset P$.

Consider the geodesic polygon with one side the geodesic $\sigma = [x, y]$ and the other sides the geodesics forming γ. Let $r_0 = \left\lfloor \frac{1}{2}(r+s) \right\rfloor$, and consider the geodesic triangle $\sigma, [x, p_{r_0}], [p_{r_0}, y]$. By δ–hyperbolicity, u lies within δ of one of $[x, p_{r_0}]$ and $[p_{r_0}, y]$. Suppose it is $[x, p_{r_0}]$ (the other case being entirely similar), and suppose that $u_1 \in [x, p_{r_0}]$ is within δ of u.

Now let $r_1 = \lfloor \frac{1}{2}(r + r_0) \rfloor$ and consider the geodesic triangle $[x, p_{r_1}], [p_{r_1}, p_{r_0}], [p_{r_0}, x]$. By δ–hyperbolicity, u_1 is within δ of one of $[x, p_{r_1}]$ or $[p_{r_1}, p_{r_0}]$, so there is u_2 on one of these sides within δ of u_1 and within 2δ of u.

We proceed in this manner, in each case making the interval of indices half as long. After t steps of this argument we find a point u_t which is within $t\delta$ of u.

After at most $d = \log_2(k) + 1$ steps, we have a geodesic triangle where two sides are $[p_l, p_{l+1}], [p_{l+1}, p_{l+2}]$ (or maybe one endpoint x or y), and we have u_d within $d\delta$ of u, but also within ϵ of P. □

The following straightforward instance of "linear-beats-log" is tailored for use in the proof of Theorem A.3.

Lemma A.2 *Fix $\delta, \epsilon > 0$, and let $g(x) = \delta(\log_2(x+1)+1)+\epsilon$. For any $m > 0$ and $c \geq 0$ there exists a natural number $R_{m,\epsilon,\delta}$ such that for all $R_0 > R_{m,\epsilon,\delta}$,*

$$g(R_0) < \frac{1}{200} m \left(\frac{1}{4} R_0 - \frac{2g(R_0)+1}{m} - 3c \right).$$

The next result states that under appropriate hypotheses, the union of an arbitrary number of quasiconvex subsets is itself quasiconvex, with constant not depending on the number of such subsets.

Theorem A.3 *Suppose that Υ is a δ–hyperbolic space and that $m, \epsilon > 0$ and $c \geq 0$ are real numbers. There exists a constant ϵ' such that for any (finite or countably infinite) collection of subsets $\{X_i\}_{i=1}^{\Lambda}$ of Υ for which*

1. *each X_i is ϵ–quasiconvex,*
2. *for each i we have $X_i \cap X_{i+1} \neq \varnothing$, and*
3. *for any i, j, if $x \in X_i$ and $y \in X_j$, we have $d(x, y) \geq m(|i-j| - c)$,*

the set $\boldsymbol{X} = \bigcup_i X_i$ is ϵ'–quasiconvex.

Proof Let $g(x) = \delta(\log_2(x+1)+1)+\epsilon$, and let $R = R_{m,\epsilon,\delta}$ be the number from Lemma A.2. Without loss of generality we may assume that $R \geq 1$.

If $\Lambda \leq 100R$ then Lemma A.1 implies $\boldsymbol{X}$ is ρ–quasiconvex with

$$\rho = \delta(\log_2(100R)+1)+\epsilon = g(100R-1).$$

On the other hand, suppose that $\Lambda > 100R$ and fix $u, v \in \boldsymbol{X}$. Let j and k be such that $u \in X_j$ and $v \in X_k$, and without loss of generality suppose that $j \leq k$. It suffices to show that any geodesic $[u, v]$ stays uniformly close to $X_j \cup \dots \cup X_k$. If $|k - j| \leq 100R$ then this follows from Lemma A.1, so suppose that $|k - j| > 100R$. Let $Y = X_j \cup \dots \cup X_k$.

Our strategy is to build a path between u and v which is uniformly quasigeodesic and stays uniformly close to Y. The theorem then follows by quasigeodesic stability. Choose a sequence of indices $t_0 = j, t_1, \dots, t_{s-1}, t_s = k$ such that for each $0 \leq r \leq s-2$,

$$t_{r+1} - t_r = 100R,$$

and

$$t_s - t_{s-1} \in \mathbb{Z} \cap [100R, \dots, 200R].$$

Moreover, for each $0 \leq r \leq s$ choose some $u_r \in X_{t_r}$. We require $u_0 = u$ and $u_s = v$.

For $r \in \{0, \dots, s-1\}$, let γ_r be a geodesic between u_r and u_{r+1}. Let

$$K = g(200R) = \delta(\log_2(200R + 1) + 1) + \epsilon.$$

Since we assume $R \geq 1$, we know that $K > \delta$.

Since we know that for each $r \in \{0, \dots, s-1\}$ we have $t_{r+1} - t_r \leq 200R$, we know that the set

$$Y_r = \bigcup_{k=t_r}^{t_{r+1}} X_k$$

is K–quasiconvex, by Lemma A.1. In particular, the geodesic γ_r lies in a K–neighborhood of Y_r.

For each $r \in \{0, \dots, s-1\}$ and each $x \in \gamma_r$, let $\pi_r(x)$ denote the set of closest points on Y_r to x. Furthermore, let $I_r(x)$ be the set of indices l such that $\pi_r(x) \cap X_l \neq \varnothing$.

Claim A.3.1 *For any $v \in \{t_r, \dots, t_{r+1}\}$ there exists $x_v \in \gamma_r$ such that*

$$d_{\mathbb{N}}(v, I_r(x_v)) \leq \frac{1}{2}\left(\frac{2K+1}{m} + c\right).$$

Proof For any $y \in \pi_r(x)$ we have $d(x, y) \leq K$. Now, if x and x' are adjacent vertices and $y \in \pi_r(x)$ with $y \in X_k$ and $z \in \pi_r(x')$ with $z \in X_l$ then

$$m(|k - l| - c) \leq d(y, z) \leq d(y, x) + d(x, x') + d(x', z) \leq 2K + 1,$$

so $|k - l| \leq (2K + 1)/m + c$.

The claim now follows immediately from the fact that $t_r \in I_r(u_r)$ and $t_{r+1} \in I_r(u_{r+1})$, letting x and x' run over adjacent pairs of vertices in γ_r. □

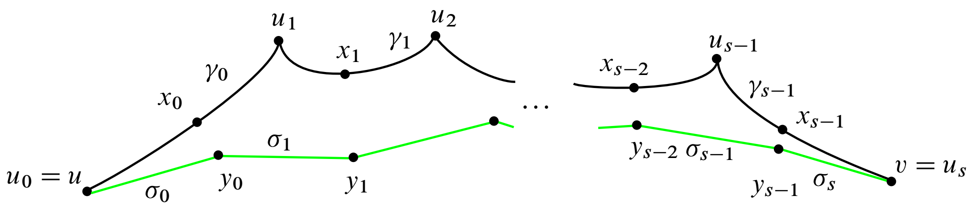

Figure 3: The σ_i forming a broken geodesic.

Suppose $0 \le r \le s-1$. Using Claim A.3.1, we can choose a point $x_r \in \gamma_r$ and a point $y_r \in \pi_r(x_r)$ such that $y_r \in X_{k_r}$ and

$$(\dagger) \qquad \left| k_r - \frac{t_r + t_{r+1}}{2} \right| \le \frac{2K+1}{2m} + c.$$

Now, for each $r \in \{1, \dots, s-1\}$, let σ_r be a geodesic between y_{r-1} and y_r. Further, let σ_0 be a geodesic from u to y_0 and let σ_s be a geodesic from y_{s-1} to v (note that there is no point y_s); see Figure 3. We bound the Gromov product between σ_t and σ_{t+1} for each t. (There is no reason to expect such a bound on the Gromov product between γ_r and γ_{r+1}.)

Though we have no control on the lengths of the segments σ_0 and σ_s, the lengths of the other segments can be bounded below:

Claim A.3.2 *Suppose $0 < r < s$. The length of σ_r is at least $200K$.*

Proof By the choice of the index k_r in $(\dagger)$,

$$k_r - k_{r-1} \ge \frac{t_{r+1} - t_{r-1}}{2} - \frac{2K+1}{m} - 2c = 100R - \frac{2K+1}{m} - 2c > 50R - \frac{2K+1}{m} - 2c$$

(the equality follows from the choice of t_r).

Below, we apply Lemma A.2 with $R_0 = 200R$, noting that $K = g(200R)$, where g is the function from that lemma. We have

$$|\sigma_r| = d_G(y_{r-1}, y_r) \ge m(k_r - k_{r-1} - c) > m\left(50R - \frac{2K+1}{m} - 3c\right) \ge 200K.$$

The second inequality above follows from the fact that $y_i \in X_{k_i}$ so such points are at least distance $m(k_r - k_{r-1} - c)$ apart. The final inequality follows from the promised use of Lemma A.2. □

Claim A.3.3 *Let $0 \le r \le s-1$. The Gromov product of σ_r and σ_{r+1} is at most $8K$.*

Proof We first handle the case that $0 < r < s-1$.

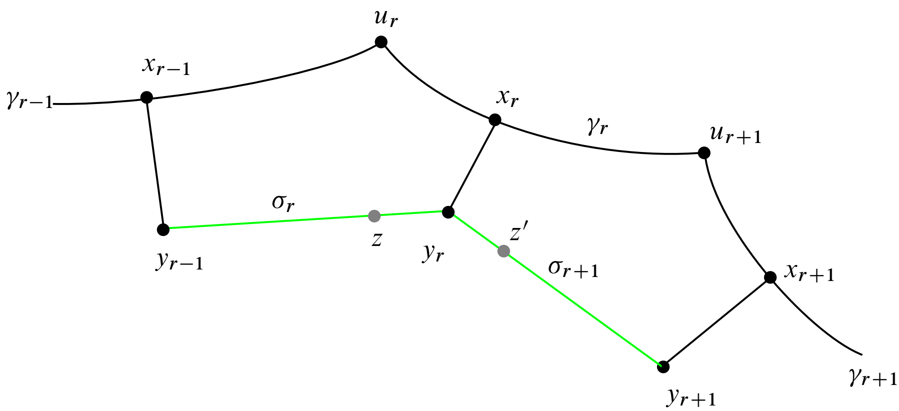

Figure 4: Computing the Gromov product of σ_r and σ_{r+1}.

For $i \in \{r, r+1\}$, the path σ_i is one side of a pentagon. The other sides are

(a) two sides of length at most K at either end of σ_i, and

(b) two "halves" of adjacent geodesics: the second "half" of γ_{i-1} and the first "half" of γ_i, joined at u_i.

See Figure 4.

By Claim A.3.2 the geodesics σ_r and σ_{r+1} have length at least $200K$. Let z be the point on σ_r at distance exactly $8K$ from y_r.

Since geodesic pentagons are 3δ–slim, we know that z must be distance at most 3δ from some point on one of the other four sides. However, it cannot be within distance 3δ of the geodesic between x_r and y_r since that geodesic has length at most K. Similarly, since $|\sigma_r| \geq 200K$, z cannot be within 3δ of the geodesic between x_{r-1} and y_{r-1}. We claim that z also cannot be within 3δ of the part of γ_{r-1} contained in the pentagon.

Indeed, suppose $w \in \gamma_{r-1}$, and choose $i_w \in I_{r-1}(w) \subset [t_{r-1}, t_r]$. There is a point w' of $\pi_{r-1}(w)$ in X_{i_w}; thus $d(w, w') \leq K$. The point x_r is likewise within K of some X_{k_r} where k_r satisfies the inequality (†). This implies that

$$|k_r - i_w| \geq \frac{t_{r+1} - t_r}{2} - \frac{2K+1}{2m} - c,$$

and so, using Lemma A.2 again,

$$d(x_r, w) \geq m\left(\frac{t_{r+1} - t_r}{2} - \frac{2K+1}{2m} - 2c\right) - 2K \geq m\left(50R - \frac{2K+1}{m} - 3c\right) - 2K \geq 198K.$$

But this contradicts $d(x_r, w) \leq d(x_r, z) + d(z, w) \leq 9K + 3\delta \leq 12K$.

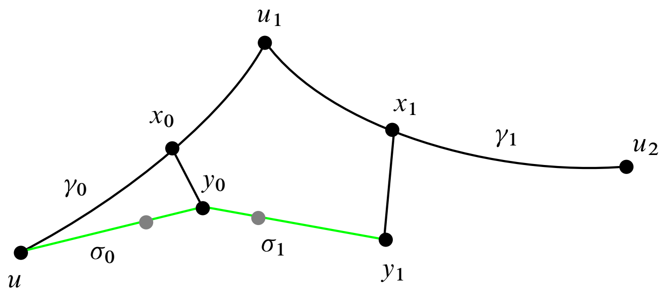

Figure 5: Computing the Gromov product of σ_0 and σ_1.

We have shown that there is some point w on γ_r between u_r and x_r within 3δ of z. Note that $d(x_r, w) \geq d(y_r, z) - K - 3\delta \geq 4K$, since $K \geq \delta$.

Now consider the pentagon formed with σ_{r+1} on one side, and the point z' on σ_{r+1} which is distance exactly $8K$ from y_r. An entirely analogous argument to the above shows that there is some w' between x_r and u_{r+1} on γ_r such that $d(z', w') \leq 3\delta$, and $d(x_r, w') \geq 4K$. Since γ_r is geodesic,

$$d(w, w') = d(w, x_r) + d(x_r, w) \geq 8K.$$

It follows that $d(z, z') \geq 8K - 6\delta \geq 2K > \delta$. Therefore, the Gromov product $(y_{r-1}, y_{r+1})_{y_r}$ is strictly less than $d(z, y_r) = d(z', y_r) = 8K$ whenever $0 < r < s$.

The cases $r = 0$ and $r = s-1$ are symmetric, so it suffices to handle the case $r = 0$; see Figure 5. We are trying to show that $(u, y_1)_{y_0} \leq 8K$, so we may suppose without loss of generality that $d(y_0, u) > 8K$. Thus there is a point z on σ_0 at distance exactly $8K$ from y_0. Since $d(x_0, y_0) \leq K$, this point is within δ of a point w on γ_0 between u and x_0.

For the point z' on σ_1 at distance $8K$ from y_0, we argue as before. We are again able to deduce that $d(z, z') > \delta$, and so $(u, y_1)_{y_0} \leq 8K$. □

Thus, we have a collection of arcs σ_i which form a broken geodesic between u and v with segments of length at least $200K$ (except possibly the first and last) and all Gromov product at most $8K$ at the corners. Thus the union of the σ_i forms a global quasigeodesic with uniformly bounded parameters. However, each σ_i lies within a $(3\delta+K)$–neighborhood of the union of the γ_i, which in turn lie in a K–neighborhood of the union of the X_i. As explained above, this suffices to prove that the union of the X_i is ϵ'–quasiconvex with the constant ϵ' depending on the quantities δ, m and ϵ, but not on the number of the X_i, as required. □

References

[1] **I Agol**, *The virtual Haken conjecture*, Doc. Math. 18 (2013) 1045–1087 MR Zbl With an appendix by I Agol, D Groves and J Manning

[2] **I Agol**, **D Groves**, **J F Manning**, *Residual finiteness, QCERF and fillings of hyperbolic groups*, Geom. Topol. 13 (2009) 1043–1073 MR Zbl

[3] **I Agol**, **D Groves**, **J F Manning**, *An alternate proof of Wise's malnormal special quotient theorem*, Forum Math. Pi 4 (2016) art. id. E1 MR Zbl

[4] **M A Armstrong**, *On the fundamental group of an orbit space*, Proc. Cambridge Philos. Soc. 61 (1965) 639–646 MR Zbl

[5] **N Bergeron**, **D T Wise**, *A boundary criterion for cubulation*, Amer. J. Math. 134 (2012) 843–859 MR Zbl

[6] **B H Bowditch**, *Tight geodesics in the curve complex*, Invent. Math. 171 (2008) 281–300 MR Zbl

[7] **B H Bowditch**, *Relatively hyperbolic groups*, Internat. J. Algebra Comput. 22 (2012) art. id. 1250016 MR Zbl

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

[9] **R Charney**, **J Crisp**, *Relative hyperbolicity and Artin groups*, Geom. Dedicata 129 (2007) 1–13 MR Zbl

[10] **D Cooper**, **D Futer**, *Ubiquitous quasi-Fuchsian surfaces in cusped hyperbolic 3–manifolds*, Geom. Topol. 23 (2019) 241–298 MR Zbl

[11] **M Coornaert**, **T Delzant**, **A Papadopoulos**, *Géométrie et théorie des groupes: les groupes hyperboliques de Gromov*, Lecture Notes in Math. 1441, Springer (1990) MR Zbl

[12] **Y Duong**, *On random groups: the square model at density $d < 1/3$ and as quotients of free nilpotent groups*, PhD thesis, University of Illinois at Chicago (2017) MR Available at `https://www.proquest.com/docview/2001345286`

[13] **E Einstein**, **D Groves**, *Relative cubulations and groups with a 2–sphere boundary*, Compos. Math. 156 (2020) 862–867 MR Zbl

[14] **A Genevois**, *Hyperbolicities in* CAT(0) *cube complexes*, Enseign. Math. 65 (2019) 33–100 MR Zbl

[15] **A Genevois**, *Coning-off* CAT(0) *cube complexes*, Ann. Inst. Fourier (Grenoble) 71 (2021) 1535–1599 MR Zbl

[16] **R Gitik**, **M Mitra**, **E Rips**, **M Sageev**, *Widths of subgroups*, Trans. Amer. Math. Soc. 350 (1998) 321–329 MR Zbl

[17] **D Groves**, **J F Manning**, *Dehn filling in relatively hyperbolic groups*, Israel J. Math. 168 (2008) 317–429 MR Zbl

[18] **D Groves**, **J F Manning**, *Quasiconvexity and Dehn filling*, Amer. J. Math. 143 (2021) 95–124 MR Zbl

[19] **D Groves**, **J F Manning**, **A Sisto**, *Boundaries of Dehn fillings*, Geom. Topol. 23 (2019) 2929–3002 MR Zbl

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

[21] **A Hatcher**, *Algebraic topology*, Cambridge Univ. Press (2002) MR Zbl

[22] **J Hempel**, *3–manifolds*, Annals of Mathematics Studies 86, Princeton Univ. Press (1976) MR Zbl

[23] **G C Hruska**, *Relative hyperbolicity and relative quasiconvexity for countable groups*, Algebr. Geom. Topol. 10 (2010) 1807–1856 MR Zbl

[24] **G C Hruska**, **D T Wise**, *Finiteness properties of cubulated groups*, Compos. Math. 150 (2014) 453–506 MR Zbl

[25] **W Jaco**, *Lectures on three-manifold topology*, CBMS Regional Conference Series in Mathematics 43, Amer. Math. Soc., Providence, RI (1980) MR Zbl

[26] **J Kahn**, **V Markovic**, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, Ann. of Math. 175 (2012) 1127–1190 MR Zbl

[27] **I Kapovich**, **H Short**, *Greenberg's theorem for quasiconvex subgroups of word hyperbolic groups*, Canad. J. Math. 48 (1996) 1224–1244 MR Zbl

[28] **J F Manning**, **E Martínez-Pedroza**, *Separation of relatively quasiconvex subgroups*, Pacific J. Math. 244 (2010) 309–334 MR Zbl

[29] **K Matsuzaki**, **M Taniguchi**, *Hyperbolic manifolds and Kleinian groups*, Oxford Univ. Press (1998) MR Zbl

[30] **J W Morgan**, *On Thurston's uniformization theorem for three-dimensional manifolds*, from "The Smith conjecture" (J W Morgan, H Bass, editors), Pure Appl. Math. 112, Academic, Orlando, FL (1984) 37–125 MR Zbl

[31] **D V Osin**, *Peripheral fillings of relatively hyperbolic groups*, Invent. Math. 167 (2007) 295–326 MR Zbl

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

[33] **M Sageev**, *Codimension-1 subgroups and splittings of groups*, J. Algebra 189 (1997) 377–389 MR Zbl

[34] **G A Swarup**, *Geometric finiteness and rationality*, J. Pure Appl. Algebra 86 (1993) 327–333 MR Zbl

[35] **H C Tran**, *On strongly quasiconvex subgroups*, Geom. Topol. 23 (2019) 1173–1235 MR Zbl

[36] **D T Wise**, *The structure of groups with a quasiconvex hierarchy*, Annals of Mathematics Studies 209, Princeton Univ. Press (2021) MR Zbl

Department of Mathematics, Statistics and Computer Science, University of Illinois at Chicago
Chicago, IL, United States

Department of Mathematics, Cornell University
Ithaca, NY, United States

dgroves@uic.edu, jfmanning@cornell.edu

Proposed: Martin R Bridson
Seconded: Benson Farb, Bruce Kleiner

Received: 29 April 2019
Revised: 14 March 2021

GEOMETRY & TOPOLOGY

msp.org/gt

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

The subscription price for 2023 is US $740/year for the electronic version, and $1030/year (+$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

GEOMETRY & TOPOLOGY

Volume 27 Issue 9 (pages 3387–3831) 2023