HYPERBOLIC GROUPS ACTING IMPROPERLY

DANIEL GROVES AND JASON FOX MANNING

Abstract. In this paper we study hyperbolic groups acting on CAT(0) cube complexes. The first main result (Theorem A) is a structural result about the Sageev construction, in which we relate quasi-convexity of hyperplane stabilizers with quasi-convexity of cell stabilizers. The second main result (Theorem D) generalizes both Agol's theorem on cubulated hyperbolic groups and Wise's Quasi-convex Hierarchy Theorem.

Contents

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 Quasi-convex 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 in [13] and [12]. At the end of the introduction in Subsection 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*–1 subgroups of G and produces a CAT(0) cube complex X, equipped with an isometric G–action on X

The first author was partially supported by the Simons Foundation, #342049 to Daniel Groves and the National Science Foundation, DMS-1507067. The second author was also partially supported by the Simons Foundation #524176 to Jason Manning and the National Science Foundation, DMS-1462263.

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–1 subgroups are quasi-convex, Sageev proved that the associated cube complex is G–cocompact [33, Theorem 3.1]. Achieving a *proper* action is harder (see [5, 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–Haefliger (see [8, III.$\mathcal{C}$] or Section 2 below). In this description, the underlying space is ${}_{G}\backslash^{X}$ 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 quasi-convex.
(2) All vertex stabilizers are quasi-convex.
(3) All cell stabilizers are quasi-convex.

Intersections of quasi-convex subgroups are quasi-convex, 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) $\implies$ (2) in the more general setting of arbitrary finitely generated groups where we assume the relevant subgroups are *strongly quasi-convex* in the sense of [35]. Note that in this more general setting, (2) and (3) are still equivalent. See Section 3 for more details. In Subsection 3.7 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 quasi-convex hyperplane stabilizers. Then*

(1) X is δ–hyperbolic for some δ;
(2) there exists a $k \geq 0$ so 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, p. 284]*).*

Anthony Genevois explained to us how conclusion (2) implies acylindricity for actions on hyperbolic CAT(0) cube complexes (see Subsection 3.7). 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 3.44 for more details.

In Sageev's construction, the stabilizers in G of hyperplanes in the resulting cube complex are commensurable with the chosen codimension–1 subgroups of G. Therefore, we have the following result.

Corollary C. *Let G be a hyperbolic group and let $\mathcal{H} = \{H_1, \ldots, H_k\}$ be a collection of quasi-convex codimension–1 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 quasi-convex in G. In particular, they are finitely presented.

(2) *X is δ–hyperbolic for some δ.*
(3) *There exists a $k \geq 0$ so that the fixed 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 so 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 non-compact 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 Quasi-convex Hierarchy Theorem [36, Theorem 13.3] (see also [3, Theorem 10.2]) which states that if a hyperbolic group G can be expressed as $A*_C$ (respectively $A *_C B$) where C is quasi-convex in G and A is (respectively 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 1-*dimensional CAT*(0) *cube complex* (otherwise known as a 'tree') with virtually special and quasi-convex cell stabilizers, then G is virtually special.

Our second main result is a common generalization of Agol's theorem and Wise's Quasi-convex Hierarchy Theorem.

Theorem D. *Suppose that G is a hyperbolic group acting cocompactly on a CAT(0) cube complex X so that cell stabilizers are quasi-convex and virtually special. 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 $\mathcal{H} = \{H_1, \ldots, H_k\}$ is a collection of quasi-convex codimension–1 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, already Yen Duong [12] has used Theorem D to show that random groups in the square model at density $< 1/3$ are virtually special. Theorem D (as well as Corollary 6.5 below) are also applied in the paper [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 quasi-convex. There exists a quotient $\overline{G} = G/K$ so 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 so 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 $\overline{G}$ so that ${}_{\overline{G}_0}\backslash^{Z}$ is special. Let G_0 be the pre-image 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 quasi-convex (since they are stabilizers of hyperplanes, which are quasi-convex by Theorem A); and
(2) The terminal groups are virtually special (since they are finite-index subgroups of the vertex stabilizers in G).

Therefore, G_0 admits a quasi-convex hierarchy terminating in virtually special groups, so G_0 is virtually special by Wise's Quasi-convex 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. This completes the proof of Theorem D. □

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 quasi-convexity criterion (Theorem A.3) which is proved separately in Appendix A. 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 $K \trianglelefteq G$ so 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 6.4 and Corollary 6.5 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] relies on Bergeron–Wise's theorem that Γ acts properly and cocompactly on a CAT(0) cube complex [5]. In turn, Bergeron–Wise rely on work of Kahn–Markovic [26], which provides a "ubiquitous"[1] family of immersed 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 1.1. *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 non-orientable, 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 quasi-convex 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 quasi-convex and codimension 1 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 quasi-convex in Γ. To apply Theorem D, we will show that the vertex stabilizers admit quasi-convex hierarchies and hence are virtually special.

Let $V < \Gamma$ be a vertex stabilizer. Since V is quasi-convex 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, 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 quasi-convex in Γ. In particular, this is a quasi-convex hierarchy, and we may apply Wise's Quasi-convex Hierarchy Theorem to conclude that V is virtually special.

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

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

1.3. **Acknowledgments.** 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 the suggestion of using a result like Theorem 6.6 in our joint work [17]. This result simplifies the proof of Theorem 6.4.

2. The complex of groups coming from an action on a cube complex

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–Haefliger [8, III.$\mathcal{C}$].

2.1. **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 1_v. For an arrow a of a scwol, we denote its source by $i(a)$ and its target by $t(a)$. If $i(a) = t(a)$, we say a is a *trivial arrow*; it follows that $a = 1_v$ for some v. We sometimes conflate v and 1_v. A *(non-degenerate) morphism* of scwols $f\colon \mathcal{A} \to \mathcal{B}$ 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 $\{a \mid i(a) = f(v)\}$.

Notation 2.1. *Given a scwol $\mathcal{X}$, we denote the set of objects of $\mathcal{X}$ by $V(\mathcal{X})$ and the set of non-unit morphisms in $\mathcal{X}$ by $E(\mathcal{X})$. The set $E(\mathcal{X})$ comes equipped with two maps*

$$i\colon E(\mathcal{X}) \to V(\mathcal{X}), t\colon E(\mathcal{X}) \to V(\mathcal{X}),$$

Where a is a morphism from $i(a)$ to $t(a)$.

Let $E^{\pm}(\mathcal{X})$ be the set of symbols a^+ and a^- as a ranges over $E(\mathcal{X})$. We refer to elements e of $E^{\pm}(\mathcal{X})$ as oriented edges *of $\mathcal{X}$. If $e = a^+$ then $i(e) = t(a)$ and $t(e) = i(a)$, while if $e = a^-$ then $i(e) = i(a)$ and $t(e) = t(a)$.*

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.

Let X be a CAT(0) cube complex, and suppose that G acts on X combinatorially. 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 $\{g \mid gx = x, \forall x \in \sigma\}$ 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 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 2.2. If W is a cube complex, the *idealization* of W is a scwol $\mathcal{W}$ which has objects $V(\mathcal{W})$ in one-to-one correspondence with non-empty nested chains of cubes of W. There is at most one morphism in $\mathcal{W}$ between two objects: If c_1 contains c_2 as a sub-chain, there is an arrow from c_1 to c_2.

For example, if X is a single 1–dimensional cube e with endpoints a and b, the nontrivial arrows of the idealization $\mathcal{X}$ are as follows:

$$(a) \longleftarrow (a \subset e) \longrightarrow (e) \longleftarrow (b \subset e) \longrightarrow (b)$$

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

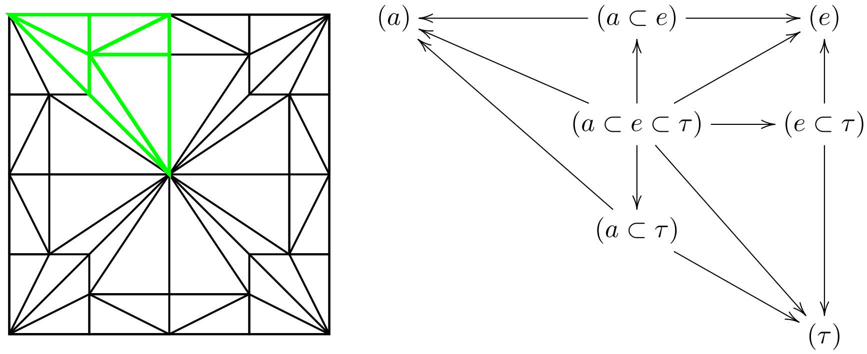

Any automorphism of a cube complex gives an automorphism of its idealization. Moreover if ϕ maps a chain of cubes to itself, then it also preserves all subchains. It follows that the quotient $\mathcal{Y} = {}_G\backslash^{\mathcal{X}}$ is also a scwol, and that the quotient map $\mathcal{X} \to \mathcal{Y}$ is non-degenerate morphism of scwols.

Remark 2.3. A small category $\mathcal{C}$ always has a *(geometric) realization* which is a simplicial complex whose 0–cells are the objects of $\mathcal{C}$, with 1–cells corresponding to morphisms, 2–cells to composable pairs of morphisms, and so on. The realization of the idealization of W is the second barycentric subdivision of W, so it is naturally homeomorphic to W.

2.2. **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. We define $\mathcal{C}$–paths to be lists of letters e, where $e = a^{\pm}$ for some arrow a of $\mathcal{C}$. For $e = a^{\pm}$ we have $i(a^+) = t(a) = t(a^-)$ and $t(a^+) = i(a) = i(a^-)$.

A *$\mathcal{C}$–path p of length* 0 is an object v of $\mathcal{C}$, with $i(p) = t(p) = v$. We also consider the path of length 0 at v to be an empty list (though it is an empty list based at v). For $j > 0$, a *$\mathcal{C}$–path of length* j is a list $p = (e_1, \ldots, e_k)$ where for each i we have $t(e_i) = i(e_{i+1})$. We have $i(p) = i(e_1)$ and $t(p) = t(e_k)$.

If p is a $\mathcal{C}$–path of length $j > 0$, q is a $\mathcal{C}$–path of length $k > 0$, 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 q is any $\mathcal{C}$–path, then $i(q)$ and $t(q)$ can be regarded as paths of length 0. We use the convention that $i(q) \cdot q = q \cdot t(q) = q$.

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

Definition 2.4. *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]By 'concatenation', we just mean concatenation of lists.

(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 $ab = a \circ b$); and
(3) $p \cdot 1_v^\pm \cdot q \simeq p \cdot q$ (where 1_v is an identity arrow).

Definition 2.5. If C is the realization of a category $\mathcal{C}$, then there is a canonical correspondence between combinatorial paths in the 1–skeleton of C and $\mathcal{C}$–paths. If p is a combinatorial path in $C^{(1)}$, and q the corresponding $\mathcal{C}$–path, we say that p is the *realization* of q, and q is the *idealization* of p.

Remark 2.6. 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.

2.3. Complexes of groups.

Definition 2.7. [8, III.$\mathcal{C}$.2.1] Let $\mathcal{A}$ be a scwol. A *complex of groups* $H(\mathcal{A})$ consists of the following data:

(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); and
(3) For each pair of composable arrows a, b with composition $a \circ b$, 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 (writing ab for $a \circ b$) whenever all written compositions of arrows are defined:

(1) (compatibility) $\mathrm{Ad}(z(a, b))\psi_{ab} = \psi_a\psi_b$; and
(2) (cocycle) $\psi_a(z(b, c))z(a, bc) = z(a, b)z(ab, c)$.

Definition 2.8 (The complex of groups coming from an action)**.** Suppose G acts on a scwol $\mathcal{X}$ so that any $g \in G$ fixing an object fixes every arrow from that object. Let $\mathcal{Y}$ be the quotient scwol. We obtain a complex of groups $G(\mathcal{Y})$ once we have made the following choices [8, III.$\mathcal{C}$.2.9]:

(1) For each object v of $\mathcal{Y}$, a lift $\widetilde{v}$ to $\mathcal{X}$; this lift also determines lifts $\widetilde{a}$ of all arrows a with $i(a) = v$.
(2) For each nontrivial arrow a, a choice of element h_a so that $t(h_a(\widetilde{a})) = \widetilde{t(a)}$.

Given these choices, one defines:

(1) G_v is the stabilizer of $\widetilde{v}$,
(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.

---

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

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).

2.4. **Fundamental groups and coverings of categories.** In Definition 2.4 we defined homotopy of $\mathcal{C}$–paths, where $\mathcal{C}$ is a category.

Definition 2.9. Given a category $\mathcal{C}$ and an object v_0 of $\mathcal{C}$, the *fundamental group of* $\mathcal{C}$ *based at* v_0, denoted $\pi_1(\mathcal{C}, v_0)$, is the set of homotopy classes of $\mathcal{C}$–loops based at v_0, with operation induced by concatenation of $\mathcal{C}$–paths.

Definition 2.10. [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 morphisms that have σ' as their initial (respectively, terminal) object is a bijection onto the set of morphisms 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 $\mathcal{Y}$, 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 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 $\widetilde{\alpha}$ is the arrow described above then $\pi(\widetilde{\alpha}) = \alpha$.

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

We record the following observation.

Lemma 2.11. *Let $\phi\colon \widetilde{\mathcal{C}} \to \mathcal{C}$ be a covering of categories, and suppose $\phi(\widetilde{v}) = v$, for objects v of $\mathcal{C}$ and $\widetilde{v}$ of $\widetilde{\mathcal{C}}$. Any $\mathcal{C}$–path p with $i(p) = v$ has a unique lift to a $\widetilde{\mathcal{C}}$–path $\widetilde{p}$ with $i(\widetilde{p}) = \widetilde{v}$. Moreover any elementary homotopy from p to a path p' gives a unique elementary homotopy of $\widetilde{p}$ to a lift $\widetilde{p}'$ of p' with the same endpoints as $\widetilde{p}$.*

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 2.12. [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) so 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 $z(a, x)$, $z(x, a)$ are always trivial.

Remark 2.13. The map $CG(\mathcal{Y}) \to \mathcal{Y}$ given by $(g, a) \to a$ is a functor. This map 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 2.20, we explain how to go back and forth between paths in covers of $CG(\mathcal{Y})$ and their associated scwols.

Theorem 2.14. [8, III.$\mathcal{C}$.3.15 and 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 $\mathcal{X}$, 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 $G \cong \pi_1(CG(\mathcal{Y}), v_0)$.*

Definition 2.15. 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 $(g, 1_v)$ is called a *group arrow.*

In later sections we abuse notation and refer to the edge $(g,1_v)^+$ (for a group arrow $(g,1_v)$) 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 $a^\pm$.

Lemma 2.16. *Every $CG(\mathcal{Y})$–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: $(g,a) = (g,t(a)) \circ (1,a)$. □

As described at the end of the last subsection, a choice of base vertex v_0 determines a universal covering map $\phi\colon \widetilde{CG(\mathcal{Y})} \to CG(\mathcal{Y})$ 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 $CG(\mathcal{Y})$.

What is important for us is that the group $\pi_1(CG(\mathcal{Y}),v_0)$ acts on the universal cover $\widetilde{CG(\mathcal{Y})}$ with quotient $CG(\mathcal{Y})$, and if $H < \pi_1(G(\mathcal{Y}),v_0)$ is any subgroup, then $H\backslash\widetilde{CG(\mathcal{Y})}$ is an intermediate cover of categories. We only consider covers of this form.

Definition 2.17. 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 2.18. *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. Lemma 2.16 gives a homotopy in $CG(\mathcal{Y})$ to a path of the desired form. Lemma 2.11 says that the homotopy lifts. □

We are particularly interested in covers of $CG(\mathcal{Y})$ corresponding to normal subgroups of $G \cong \pi_1(G(\mathcal{Y}),v_0)$. Fix some such $K \lhd G$, let $\widetilde{CG(\mathcal{Y})}$ be the universal cover of $CG(\mathcal{Y})$, and let $\mathcal{C}_K := K\backslash\widetilde{CG(\mathcal{Y})}$ be the corresponding cover. The group K also acts on the scwol $\mathcal{X}$, with quotient scwol $\mathcal{Z} := K\backslash\mathcal{X}$. We observe that (just as with ordinary covers) a $CG(\mathcal{Y})$–loop p represents an element of K if and only if it lifts to a loop in $\mathcal{C}_K$. (Since $K \lhd G$, the basepoints do not matter.)

Given a regular cover $\mathcal{C}_K$ of $CG(\mathcal{Y})$, corresponding to the normal subgroup $K \unlhd G$, there is a natural quotient $\overline{\mathcal{C}}_K$ of $\mathcal{C}_K$, defined as follows:

Define an equivalence relation on the objects of $\mathcal{C}_K$ where two objects are equivalent if they differ by an invertible arrow. Define an equivalence relation on arrows of $\mathcal{C}_K$ by setting $\gamma_1 \sim \gamma_2$ if $i(\gamma_1) \sim i(\gamma_2)$, $t(\gamma_1) \sim t(\gamma_2)$ and if there are invertible arrows: $\rho_1\colon i(\gamma_1) \to i(\gamma_2)$ and $\rho_2\colon t(\gamma_1) \to t(\gamma_2)$ so that $\rho_2\gamma_1 = \gamma_2\rho_1$.

Objects and arrows of $\mathcal{C}_K$ up to the above equivalences form a quotient category $\overline{\mathcal{C}}_K$. Since $\mathrm{Mor}(v,v)$ is a group for any object v of $\mathcal{C}_K$, the following lemma is straightforward.

Lemma 2.19. *The category $\overline{\mathcal{C}}_K$ is a scwol, and the quotient map $\mathcal{C}_K \to \overline{\mathcal{C}}_K$ is a functor. The category $\overline{\mathcal{C}}_K$ is G/K–equivariantly isomorphic to the scwol $\mathcal{Z} = K\backslash\mathcal{X}$.*

Definition 2.20. We denote the functor from $\mathcal{C}_K$ to $\mathcal{Z}$ given by Lemma 2.19 by $\Theta_K\colon \mathcal{C}_K \to \mathcal{Z}$. Since Θ_K is a functor, it also gives a way to turn a $\mathcal{C}_K$–path p into a $\mathcal{Z}$–path $\bar{p}$. Deleting all the trivial arrows from $\bar{p}$ 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 $\widehat{\sigma}$ so that $\Theta_K(\widehat{\sigma}) = \sigma$ is called an *unscwolification* of σ. The unscwolification is highly non-unique, but always exists.

The following can be deduced by examining the elementary homotopies.

Lemma 2.21. *Scwolifications of homotopic paths are homotopic.*

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

Lemma 2.22. *Let p be a $CG(\mathcal{Y})$–loop at v and let $\widehat{p}$ be a lift to $\mathcal{C}_K$. If $\Theta_K(\widehat{p})$ is a loop, then there is a group arrow labeled by an element of G_v joining the endpoints of $\widehat{p}$.*

In case $K = \{1\}$ we have $\mathcal{Z} = \mathcal{X}$ and $\mathcal{C}_K = \widetilde{CG(\mathcal{Y})}$. In this case we just write $\Theta\colon \widetilde{CG(\mathcal{Y})} \to \mathcal{X}$.

3. Quasi-convexity 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 quasi-convex if and only if the hyperplane stabilizers are quasi-convex.

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 quasi-convex if and only if vertex stabilizers are. Edge stabilizers are intersections of vertex stabilizers, and intersections of quasi-convex subgroups are quasi-convex, 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 quasi-convexity 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 quasi-convex. It still follows that edge stabilizers are quasi-convex, 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 quasi-convexity criterion proved in the Appendix to conclude that the hyperplane stabilizer is quasi-convex.

If on the other hand we assume that hyperplane stabilizers are quasi-convex, 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.

As mentioned in the introduction, we prove the direction (1) $\implies$ (2) in a more general setting. Therefore, for the beginning of this section we do not assume that G is hyperbolic, merely that it acts cocompactly on a CAT(0) cube complex X.

We briefly describe the contents of the remainder of this section. In Subsection 3.1 we explain how we consider subset of small categories as graphs. In Subsection 3.2 we identify certain subsets of $CG(\mathcal{Y})$ which are tuned to the cubical geometry of X^b and associate graphs with these subsets. In Subsection 3.3 we build graphs upon which intersections of stabilizers of hyperplanes act. In Subsection 3.4 we prove the direction (1) $\implies$ (2) of Theorem A. In fact, we prove the more general Theorem 3.26. In Subsection 3.5 we prove the direction (2) $\implies$ (1) of Theorem A. In Subsection 3.6 we consider various possible generalizations of Theorem A. Finally, in Subsection 3.7 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 $\mathrm{Gr}(S)$, 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. For $\mathrm{Gr}(S)$ a graph constructed this way, we denote the original set of arrows as $\mathrm{Ar}(\mathrm{Gr}(S)) = S$.

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 geometric realization of $\mathcal{C}$.*

Example 3.2. *Suppose $\mathcal{C}$ is a group (i.e. $\mathcal{C}$ has a single object and each morphism 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.** It is convenient for us to work in the *cubical subdivision* of X, which we now describe.

Definition 3.3. Suppose that X is a cube complex. The *cubical subdivision* of X, denoted 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.

If a group G acts by cubical automorphisms on X, then it clearly does so on X^b. Moreover, the carrier of a hyperplane W in X^b is homeomorphic to $W \times [0,1]$. With appropriate choice of orientation, $W \times \{1\}$ is a hyperplane $W^\uparrow$ of X, and $\mathrm{Stab}(W) \subset \mathrm{Stab}(W^\uparrow)$ is an inclusion with index 1 or 2 (depending on whether or not there is an element of G which fixes $W^\uparrow$ but exchanges the two sides of $W^\uparrow$). We denote $W \times \{0\}$ by $W^\downarrow$. Note that $W^\downarrow$ naturally corresponds to the cubical subdivision of a sub-complex of X.

We observe the following.

Lemma 3.4. *Suppose that a hyperbolic group G acts on a CAT(0) cube complex X. Then hyperplane stabilizers for the G–action on X are quasi-convex if and only if hyperplane stabilizers for the G–action on X^b are.*

Thus, in order to prove Theorem A, we can consider either hyperplane stabilizers for hyperplanes in X or for hyperplanes in X^b. A similar result holds in the setting of strongly quasi-convex subgroups of a finitely generated group, as discussed in Subsection 3.4 below.

Observation 3.5. *The vertices of X^b are in bijection with the cubes of X.*

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

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

In Section 2.4, we defined $\widetilde{CG(\mathcal{Y})}$ to be the universal covering of the category $CG(\mathcal{Y})$ associated to the complex of groups $G(\mathcal{Y})$. Recall that the objects of $\widetilde{CG(\mathcal{Y})}$ are homotopy classes of $CG(\mathcal{Y})$–paths, starting at a (fixed) basepoint $v_0 \in \mathcal{Y}$, and arrows are labeled by arrows of $CG(\mathcal{Y})$ (Definition 2.17). It is helpful to assume (as we may do without loss of generality) that v_0 comes from a 0–cube of X. The basepoint of $\widetilde{CG(\mathcal{Y})}$ is $\widetilde{v}_0$, the homotopy class of the constant path at v_0.

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 2.14, 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 $\mathrm{Gr}(\widetilde{CG(\mathcal{Y})})$. The G-cocompact graph is different in the two directions of the proof, primarily because when we assume that hyperplane stabilizers are quasi-convex, we do not a priori know that vertex stabilizers are finitely generated (in fact, this is part of the desired conclusion of Theorem A). In both directions, the graph we choose is chosen to reflect the cubical geometry of X, or rather that of X^b.

As noted in Remark 2.6 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 $(\widetilde{\sigma}_1 \subseteq \widetilde{\sigma}_2)$ with $\widetilde{\sigma}_1$ of codimension 1 in $\widetilde{\sigma}_2$. If the path p passes over the edge e, its idealization $\widehat{p}$ contains consecutive arrows labelled $(\widetilde{\sigma}_1 \subseteq \widetilde{\sigma}_2) \to \widetilde{\sigma}_1$ and $(\widetilde{\sigma}_1 \subseteq \widetilde{\sigma}_2) \to \widetilde{\sigma}_2$, and every arrow of $\widehat{p}$ has such a label.

Definition 3.6. A *pair of opposable scwol arrows* in $\widetilde{CG(\mathcal{Y})}$ is a pair of scwol arrows γ_1, γ_2 so that

1. $c = i(\gamma_1) = i(\gamma_2)$ has the property that $\Theta(c)$ is a chain $(\widetilde{\sigma}_1 \subset \widetilde{\sigma}_2)$ where $\widetilde{\sigma}_1$ has codimension one in $\widetilde{\sigma}_2$;
2. the label of γ_1 is $(1, a)$ where a is the arrow in $\mathcal{Y}$ corresponding to the G–orbit of the arrow $(\widetilde{\sigma}_1 \subset \widetilde{\sigma}_2) \to \widetilde{\sigma}_1$ in $\mathcal{X}$; and
3. the label of γ_2 is $(1, b)$, where b is the arrow in $\mathcal{Y}$ corresponding to the G–orbit of the arrow $(\widetilde{\sigma}_1 \subset \widetilde{\sigma}_2) \to \widetilde{\sigma}_2$ in $\mathcal{X}$.

The *center* of the pair of opposable arrows (γ_1, γ_2) is the object $c = i(\gamma_1) = i(\gamma_2)$. The image in $CG(\mathcal{Y})$ of a pair of opposable scwol arrows is also referred to as a pair of opposable scwol arrows.

Definition 3.7. 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.

The following result is straightforward to prove, starting with an arbitrary $CG(\mathcal{Y})$–path and applying relations until it is of the desired form.

Proposition 3.8. *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 $CG(\mathcal{Y})$–path starting and ending at v_0.

Definition 3.9. 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 the following two sets:

1. $S_1(\mathbb{A})$ is the set of (group) arrows with label in some $\mathbb{A}_{\mathfrak{o}}$.
2. $S_2(\mathbb{A})$ is 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 $\Gamma(\mathbb{A})$.

Example 3.10. *For each cubical object $\mathfrak{o}$ of $\mathcal{Y}$, let $\mathbb{U}_{\mathfrak{o}} = G_{\mathfrak{o}}$. It follows from Proposition 3.8 that $\Gamma(\mathbb{U})$ is connected (see also Proposition 3.23 below).*

The functor $\Theta\colon \widetilde{CG(\mathcal{Y})} \to \mathcal{X}$ induces a simplicial map
$$\Psi\colon \Gamma(\mathbb{U}) \to \left(X^b\right)^{(1)}.$$

All the graphs we construct are subgraphs of $\Gamma(\mathbb{U})$, and we keep the terminology of cubical vertices and central vertices for these subgraphs.

3.3. **Graphs associated to tuples of intersecting hyperplanes.** In this subsection we consider tuples $(W_1, \ldots, W_k)$ of hyperplanes of X^b so that $\bigcap_i W_i$ is nonempty. We include the possibility of the empty list $()$, in which case we use the convention that the intersection is all of X^b. We observe the following consequence of the cocompactness of $G \curvearrowright X$:

Lemma 3.11. *There are finitely many G–orbits of finite ordered lists*
$$(W_1, W_2, \ldots, W_k)$$
of distinct hyperplanes so that $\cap_{i=1}^k W_i \neq \emptyset$. For each such list,
$$\mathrm{Stab}(\mathfrak{C}) = \mathrm{Stab}(W_1) \cap \cdots \cap \mathrm{Stab}(W_k),$$
using the convention that the empty intersection of subgroups is G.

Definition 3.12 (Choosing representatives)**.** We choose representatives of these ordered lists to preserve inclusion. Namely, first choose a collection of representatives of hyperplanes. If W is a hyperplane, denote its representative by $\overline{W}$. Next, choose representatives of pairs (W_1, W_2) with $W_1 \cap W_2 \neq \emptyset$ so that the representative pair is $(\overline{W_1}, W_2')$, where W_2 is some hyperplane in the G–orbit of W_2, though not necessarily $\overline{W_2}$. What we require is that $(\overline{W_1}, W_2')$ is in the G–orbit of (W_1, W_2), in the sense that there is some $g \in G$ so that $gW_1 = \overline{W_1}$ and $gW_2 = W_2'$. More generally, the representative of a list $(W_1, \ldots, W_k)$ starts with the representative of the list $(W_1, \ldots, W_{k-1})$ and then appends an appropriate element of the orbit of W_k (though not necessarily the element $\overline{W_k}$). Let $\mathfrak{W}$ denote the chosen finite collection of representative ordered lists of hyperplanes. For $0 \le i \le \dim X$, let $\mathfrak{W}_i$ denote the subcollection consisting of ordered lists of length i.

Definition 3.13 (Pushing down)**.** Suppose that W is a hyperplane of X^b. As in Subsection 3.2, we identify the carrier of W in X^b with a product $W \times [0,1]$, where $W \times \{1\}$ is a hyperplane W' of X, and $W \times \{0\}$ is one of the connected components of the frontier of the carrier of W'. As sub-complexes of X^b, we write these as $W^\uparrow = W \times \{1\}$ and $W^\downarrow = W \times \{0\}$. If U is a sub-complex of W, then define $U^\uparrow \subseteq W^\uparrow$ and $U^\downarrow \subseteq W^\downarrow$ in the analogous way. It is possible that such a U is a sub-complex of more than one hyperplane, in which case we need to specify *which* hyperplane we are focusing on. We write $U^{\uparrow(W)}$ and $U^{\downarrow(W)}$ when there is some ambiguity.

Definition 3.14 (The graph $\widehat{\Gamma}(\mathfrak{C})$)**.** Suppose that $\mathfrak{C} = (W_1, \ldots, W_k) \in \mathfrak{W}$.

If $\mathfrak{C}$ is the empty list of hyperplanes, let $I_{\mathfrak{C}} = X^b$, let $V_\square(\mathfrak{C})$ be the set of cubical objects of X^b, and let $\widehat{\Gamma}(\mathfrak{C}) = \Gamma(\mathbb{U})$.

Now suppose that $k \ge 1$. Let $\mathfrak{C}^- = (W_1, \ldots, W_{k-1})$ and suppose by induction that $I_{\mathfrak{C}^-}$, $V_\square(\mathfrak{C}^-)$ and $\widehat{\Gamma}(\mathfrak{C}^-)$ have been defined. By the way the set $\mathfrak{W}$ was chosen, $\mathfrak{C}^- \in \mathfrak{W}$.

Define

$$I_{\mathfrak{C}} = (I_{\mathfrak{C}^-} \cap W_k)^{\downarrow(W_k)} .$$

This is obtained by pushing the first hyperplane into a sub-complex of X^b, intersecting with the next hyperplane and then pushing this intersection into a sub-complex of X^b, and repeating. In particular $I_{\mathfrak{C}}$ is a subcomplex of $I_{\mathfrak{C}^-}$. The intersection $I_{\mathfrak{C}^-} \cap W_k \cap (X^b)^{(1)}$ is a collection of midpoints of edges in X^b dual to W_k. Let $\mathcal{O}(\mathfrak{C}) \subset \mathcal{X}$ be the set of idealizations of these midpoints. Each such idealization is a chain $c = (\widetilde{\sigma} \subset \widetilde{\nu})$ where $\widetilde{\nu}$ is a cube of X contained in $I_{\mathfrak{C}^-}$ which is cut in half by $W_k^\uparrow$. For $v \in \Theta^{-1}(c)$ there is a unique scwol arrow γ_v in $\widetilde{CG(\mathcal{Y})}$ with $i(\gamma_v) = v$ so that $\Theta(\gamma_v) = (\widetilde{\sigma} \subset \widetilde{\nu}) \to \widetilde{\sigma}$.

Define the set

$$(\dagger) \qquad V_\square(\mathfrak{C}) = \left\{ t(\gamma_v) \mid v \in \Theta^{-1}\left(\mathcal{O}(\mathfrak{C})\right) \right\} \cap V_\square(\mathfrak{C}^-).$$

The graph $\widehat{\Gamma}(\mathfrak{C})$ is now defined to be $\mathrm{Gr}\left(S(\mathfrak{C})\right)$, where $S(\mathfrak{C})$ is the union of the following two sets of arrows:

(1) $S_1(\mathfrak{C})$ is the collection of group arrows between objects of $V_\square(\mathfrak{C})$.
(2) $S_2(\mathfrak{C})$ is the set of scwol arrows occurring in a pair (γ_1, γ_2) of opposable scwol arrows with both $t(\gamma_1)$ and $t(\gamma_2)$ in $V_\square(\mathfrak{C})$.

The following lemmas are clear from the definition.

Lemma 3.15.

$$\widehat{\Gamma}(\mathfrak{C}) \subseteq \widehat{\Gamma}(\mathfrak{C}^-).$$

Lemma 3.16. *For any $\mathfrak{C} \in \mathcal{W}$, $\Theta(V_\square(\mathfrak{C}))$ lies in the idealization of $I_{\mathfrak{C}}$.*

The next proposition shows that $\widehat{\Gamma}(\mathfrak{C})$ is $\mathrm{Stab}(\mathfrak{C})$–equivariant in a strong sense.

Proposition 3.17. *Let $\mathfrak{C} = (W_1, \ldots, W_k)$ be as above. The graph $\widehat{\Gamma}(\mathfrak{C})$ is $\mathrm{Stab}(\mathfrak{C})$–invariant. Moreover, if $g\widehat{\Gamma}(\mathfrak{C}) \cap \widehat{\Gamma}(\mathfrak{C}) \neq \emptyset$, then $g \in \mathrm{Stab}(\mathfrak{C})$.*

Proof. We continue to use the notation $\mathfrak{C}^- = (W_1, \ldots, W_{k-1})$. (Note that for $k = 1$, $\mathrm{Stab}(\mathfrak{C}^-) = G$.)

First observe that g preserves $\widehat{\Gamma}(\mathfrak{C})$ if and only if $g \in \mathrm{Stab}(V_\Box(\mathfrak{C}))$, defined above in equation (†).

Suppose that $g \in \mathrm{Stab}(\mathfrak{C})$. Then $g \in \mathrm{Stab}(\mathfrak{C}^-)$, so by induction, g preserves $V_\Box(\mathfrak{C}^-)$. To establish that $g \in \mathrm{Stab}(V_\Box(\mathfrak{C}))$, it suffices to show that g preserves $\mathcal{O}(\mathfrak{C})$, the set of idealizations of points in $I_{\mathfrak{C}^-} \cap W_k \cap (X^b)^{(1)}$. This holds because $gW_i = W_i$ for each $i \in \{1, \ldots, k\}$. The first conclusion is proved.

We now assume $g\widehat{\Gamma}(\mathfrak{C}) \cap \widehat{\Gamma}(\mathfrak{C}) \neq \emptyset$ and show that $g \in \mathrm{Stab}(\mathfrak{C})$. Suppose that $z \in g\widehat{\Gamma}(\mathfrak{C}) \cap \widehat{\Gamma}(\mathfrak{C})$ is a vertex. If z is central, then both its (cubical) neighbors are also in $g\widehat{\Gamma}(\mathfrak{C}) \cap \widehat{\Gamma}(\mathfrak{C})$, so we may assume z is cubical.

In this case $z \in V_\Box(\mathfrak{C})$, so $z = t(\gamma_v)$ where $v \in \Theta^{-1}(\mathcal{O}(\mathfrak{C}))$ is the idealization of some point of $I_{\mathfrak{C}^-} \cap W_k \cap (X^b)^{(1)}$. Similarly $z' = g^{-1}(z) = t(\gamma_w)$ for some $w \in \Theta^{-1}(\mathcal{O}(\mathfrak{C}))$. Because the orbit map from $\widetilde{CG(\mathcal{Y})}$ to $CG(\mathcal{Y})$ is a covering map, we must have $g\gamma_w = \gamma_v$. But the realization of $\Theta(i(\gamma_v))$ is contained in a unique hyperplane of X^b, namely W_k, and similarly for $\Theta(i(\gamma_w))$. We must therefore have $gW_k = W_k$, so $g \in \mathrm{Stab}(W_k)$. Induction shows that $g \in \mathrm{Stab}(\mathfrak{C}^-)$. Since $\mathrm{Stab}(\mathfrak{C}^-) \cap \mathrm{Stab}(W_k) = \mathrm{Stab}(\mathfrak{C})$, we are finished. □

Lemma 3.18. *Suppose that $v \in V_\Box(\mathfrak{C})$. Let $\mathfrak{o} = [\![\Theta(v)]\!]$ be the corresponding object in $\mathcal{Y}$. The set of labels of group arrows in $\mathrm{Ar}(\widehat{\Gamma}(\mathfrak{C}))$ adjacent to v forms a subgroup of $G_\mathfrak{o}$ isomorphic to $\mathrm{Stab}(\mathfrak{C}) \cap \mathrm{Stab}(\Theta(v))$. If v, w are in the same $\mathrm{Stab}(\mathfrak{C})$–orbit the corresponding subgroups of $G_\mathfrak{o}$ are equal.*

Proof. Fix $v \in V_\Box(\mathfrak{C})$. By the definition of $V_\Box(\mathfrak{C})$, there is a unique scwol arrow γ with $t(\gamma) = v$ and $i(\gamma) \in \Theta^{-1}(\mathcal{O}(\mathfrak{C}))$. Let a be a group arrow from v to w for some $w \in V_\Box(\mathfrak{C})$. There must similarly be a scwol arrow γ' with the same label as γ satisfying $t(\gamma') = w$ and $i(\gamma') \in \Theta^{-1}(\mathcal{O}(\mathfrak{C}))$. In fact we have $\Theta(\gamma) = \Theta(\gamma')$, so $i(\gamma)$ is connected to $i(\gamma')$ by a group arrow.

Conversely, if v_0 and w_0 are two objects in $\mathcal{O}(\mathfrak{C})$ which are joined by a group arrow then there is a square with two scwol arrows γ_{v_0} and γ_{w_0} and a group arrow between the corresponding objects in $V_\Box(\mathfrak{C})$. It is clear that the set of all such group arrows adjacent to v form a group which is isomorphic to $G_{[\![\Theta(i(\gamma))]\!]}$. This group is isomorphic to $\mathrm{Stab}(\mathfrak{C}) \cap \mathrm{Stab}(\Theta(v))$. The final assertion is clear. □

Definition 3.19. Let $v \in V_\Box(\mathfrak{C})$, and let $\mathfrak{o} = [\![\Theta(v)]\!]$ be the corresponding G–orbit (object in $\mathcal{Y}$). Let $H_{v,\mathfrak{C}}$ be the subgroup of $G_\mathfrak{o}$ given by Lemma 3.18.

Note that by Lemma 3.18, $H_{v,\mathfrak{C}}$ depends only on $\mathfrak{C}$ and the $\mathrm{Stab}(\mathfrak{C})$–orbit of v, and not on the particular choice of v.

Lemma 3.20. *There are finitely many $\mathrm{Stab}(\mathfrak{C})$–orbits of cubical vertices in $\widehat{\Gamma}(\mathfrak{C})$.*

Proof. Suppose $\mathfrak{C} = (W_1, \ldots, W_k)$, $\mathfrak{C}^- = (W_1, \ldots, W_{k-1})$, as above.

The cubical vertices in $\widehat{\Gamma}(\mathfrak{C})$ are exactly $V_\Box(\mathfrak{C})$, so we need to show $\mathrm{Stab}(\mathfrak{C})$ acts cofinitely on $V_\Box(\mathfrak{C})$.

We claim first that there is a $\mathrm{Stab}(\mathfrak{C})$–equivariant bijection $\alpha\colon V_\square(\mathfrak{C}) \to \Theta^{-1}(\Xi)$ where Ξ is the set of objects

$$\tau_{\mathfrak{C}} := (I_{\mathfrak{C}} \cap \tau \subset I_{\mathfrak{C}^-} \cap \tau \subset \cdots \subset \tau),$$

where τ ranges over those cubes of X meeting $\bigcap_{i=1}^k W_i$. Indeed this can be seen by induction on k. In case $k = 0$, these sets are equal. Suppose $k > 0$. If $v \in V_\square(\mathfrak{C})$, then $v = t(\gamma_1)$ where γ_1 has scwolification

$$(I_{\mathfrak{C}} \cap \tau \subset I_{\mathfrak{C}^-} \cap \tau) \to (I_{\mathfrak{C}} \cap \tau)$$

for some cube τ of X. If $k = 1$, then $I_{\mathfrak{C}^-} \cap \tau = \tau$, and we define $\alpha(v) = i(\gamma_1)$. Otherwise by induction there is a unique scwol arrow γ_2 with $t(\gamma_2) = i(\gamma_1)$ and scwolification

$$\tau_{\mathfrak{C}} \to (I_{\mathfrak{C}} \cap \tau \subset I_{\mathfrak{C}^-} \cap \tau)$$

and we define $\alpha(v) = i(\gamma_2)$. We note that in either case $\alpha\colon V_\square(\mathfrak{C}) \to \Theta^{-1}(\Xi)$ is a bijection.

For any cube τ meeting $\bigcap_{i=1}^k W_i$ we have $\mathrm{Stab}(\tau) \cap \mathrm{Stab}(\mathfrak{C}) = \mathrm{Stab}(\tau_{\mathfrak{C}})$, so $\Theta^{-1}(\tau_{\mathfrak{C}})$ meets exactly one $\mathrm{Stab}(\mathfrak{C})$–orbit in $\Theta^{-1}(\Xi)$. Since there are finitely many $\mathrm{Stab}(\mathfrak{C})$–orbits of such cubes, the set $\Theta^{-1}(\Xi)$ is $\mathrm{Stab}(\mathfrak{C})$–cofinite. □

Let $\widetilde{v_{\mathfrak{C}}}$ be a basepoint of $\widehat{\Gamma}(\mathfrak{C})$. Let $\widetilde{d_{\mathfrak{C}}}$ be a cubical $\widetilde{CG(\mathcal{Y})}$–path from $\widetilde{v_0}$ to $\widetilde{v_{\mathfrak{C}}}$. The paths without '$\,\widetilde{}\,$' on them denote the projected $CG(\mathcal{Y})$–paths.

Lemma 3.21. *Every element of* $\mathrm{Stab}(\mathfrak{C})$ *can be represented by a* $CG(\mathcal{Y})$*–path of the form*

$$d_{\mathfrak{C}} \cdot p \cdot \overline{d_{\mathfrak{C}}},$$

where p *is a cubical* $CG(\mathcal{Y})$*–path so that the lift* $\widetilde{p}$ *of* p *to* $\widetilde{CG(\mathcal{Y})}$ *starting at* $\widetilde{v_{\mathfrak{C}}}$ *is a concatenation of arrows in* $\mathrm{Ar}(\widehat{\Gamma}(\mathfrak{C}))$.

Definition 3.22. Choose a collection of $\mathrm{Stab}(\mathfrak{C})$–orbit representatives of objects of $V_\square(W)$. For each such representative v, choose a subset $\mathbb{B}_v^{\mathfrak{C}}$ of $H_{v,\mathfrak{C}}$. Let $\widetilde{\mathbb{B}_v^{\mathfrak{C}}}$ be the set of arrows from v which are labelled by elements of $\mathbb{B}_v^{\mathfrak{C}}$, and extend this definition equivariantly across the $\mathrm{Stab}(\mathfrak{C})$–orbit of v. Let $\mathbb{B}^{\mathfrak{C}}$ be the union of the $\widetilde{\mathbb{B}_v^{\mathfrak{C}}}$ and let $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}}) = \mathrm{Gr}\left(\mathbb{B}^{\mathfrak{C}} \cup S_2(\mathfrak{C})\right)$.

Proposition 3.23. *The graph* $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ *admits a free* $\mathrm{Stab}(\mathfrak{C})$*–action. If each* $\mathbb{B}_v^{\mathfrak{C}}$ *is finite then the* $\mathrm{Stab}(\mathfrak{C})$*–action is cocompact.*

Suppose that $\mathrm{Stab}(\mathfrak{C})$ *is generated by a finite set* F*, and that each* $g \in F$ *is represented by a* $CG(\mathcal{Y})$*–path* $d_{\mathfrak{C}} \cdot p_g \cdot \overline{d_{\mathfrak{C}}}$ *as in the conclusion of Lemma 3.21. If each group arrow occurring in each* p_g *has label in* $\langle \mathbb{B}_v^{\mathfrak{C}} \rangle$ *for the appropriate* $\mathbb{B}_v^{\mathfrak{C}}$ *then* $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ *is connected.*

Proof. Proposition 3.17 implies that $\widehat{\Gamma}(\mathfrak{C})$ is $\mathrm{Stab}(\mathfrak{C})$–invariant. It then follows from the construction that $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ is also $\mathrm{Stab}(\mathfrak{C})$–invariant. The free G–action on $\widetilde{CG(\mathcal{Y})}$ by deck transformations thus restricts to a free action of $\mathrm{Stab}(\mathfrak{C})$ on $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$.

Now suppose each $\mathbb{B}_v^{\mathfrak{C}}$ is finite. Since $\mathcal{Y}$ is finite, at each object v of $\widetilde{CG(\mathcal{Y})}$ there are only finitely many scwol arrows which begin or end at v. Since $\mathbb{B}_v^{\mathfrak{C}}$ is finite, this implies that the graph $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ is locally finite. It now follows from Lemma 3.20 that the $\mathrm{Stab}(\mathfrak{C})$–action is cocompact.

Now suppose that $\mathrm{Stab}(\mathfrak{C}) = \langle F \rangle$ for a finite set F, and that the hypotheses of the final assertion of the proposition is satisfied. We first note that from any vertex v of $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ there is a $g \in \mathrm{Stab}(\mathfrak{C})$ and a path consisting entirely of scwol arrows between v and $g.\widetilde{v_{\mathfrak{C}}}$. Therefore, it suffices to find a path in $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ between $\widetilde{v_{\mathfrak{C}}}$ an $g.\widetilde{v_{\mathfrak{C}}}$, for an arbitrary $g \in \mathrm{Stab}(\mathfrak{C})$. We can represent g as a product of elements of F and their inverses. Each of these elements of F is represented by a path based at v_0 of the form $d_{\mathfrak{C}} \cdot p \cdot \overline{d_{\mathfrak{C}}}$ as in Lemma 3.21. Lifting to a path starting at $\widetilde{v_0}$, these paths determine a path between $\widetilde{v_0}$ and $g.\widetilde{v_0}$ which is homotopic to a path $d_{\mathfrak{C}} \cdot \widetilde{q} \cdot \overline{d_{\mathfrak{C}}}$ where $\widetilde{q}$ is a cubical path starting at $\widetilde{v_{\mathfrak{C}}}$ which is a concatenation of arrows whose labels all lie in the appropriate $\langle \mathbb{B}^{\mathfrak{C}}_v \rangle$, by the choice of F. Therefore, $\widetilde{q}$ lies in $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$, which proves that $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ is connected, as required. □

We denote the restriction of $\Psi\colon \Gamma(\mathbb{U}) \to \left(X^b\right)^{(1)}$ to $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ by

$$\Psi_{\mathfrak{C}}\colon \Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}}) \to (X^b)^{(1)}.$$

3.4. **If hyperplane stabilizers are QC then cell stabilizers are QC.** In this section we prove the direction (1) $\implies$ (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 quasi-convex hyperplane stabilizers. The right general setting for this proof is that of *strongly quasi-convex subgroups* of finitely generated groups, as defined by Tran in [35]. (Such subgroups were also studied by Genevois [15] under the name *Morse subgroups*.)

Definition 3.24. [35, Definition 1.1] Let X be a geodesic metric space. A subset $Q \subseteq X$ is *strongly quasi-convex* if for every $K \geq 1$, $C \geq 0$ there is some $M = M(K,C)$ so that every (K,C)–quasi-geodesic 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 quasi-convexity 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 3.25. *Suppose that X and Y are geodesic metric spaces, that $A \subset X$ is strongly quasi-convex, that $\phi\colon X \to Y$ is a quasi-isometry and that $B \subset Y$ lies at finite Hausdorff distance from $\phi(A)$. Then B is a strongly quasi-convex subset of Y.*

Proof sketch. This is proved essentially in the same way as the corresponding fact about quasi-convex subsets of hyperbolic spaces. The difference is that instead of a single constant of quasi-convexity, 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 $d_{\mathrm{Haus}}(B, \phi(A)) \leq \epsilon$.

Any quasi-geodesic γ joining points in B can be extended by a pair of geodesic segments of length $\leq \epsilon$ to make a quasi-geodesic γ' joining points in $\phi(A)$. The image of γ' under ψ can likewise be extended to a quasi-geodesic γ'' between points of A. If γ was a (K,C)–quasi-geodesic, then γ'' is a (K',C')–quasi-geodesic where K', 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 therefore 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 quasi-convexity makes sense for subgroups of finitely generated groups.

In this subsection, we prove the following theorem.

Theorem 3.26. *Suppose that a finitely generated group G acts cocompactly on a $CAT(0)$ cube complex X and that the hyperplane stabilizers are strongly quasi-convex. Then the cell stabilizers are strongly quasi-convex.*

Since quasi-convexity is equivalent to strong quasi-convexity for subgroups of hyperbolic groups, Theorem 3.26 immediately implies the direction (1) $\implies$ (2) of Theorem A.

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

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

Proposition 3.27. *For any Morse gauge M and any $D, N, r > 0$ there is an $R > 0$ so that the following holds. Let X be a graph of valence $\leq D$ with a group G acting on X with at most N orbits of vertices. Let A, B be M–strongly quasi-convex subsets of X satisfying:*

(1) If $g \in G$ satisfies $gA \cap A \neq \emptyset$ then $gA = A$.
(2) If $g \in G$ satisfies $gB \cap B \neq \emptyset$ then $gB = B$.
(3) $A \cap B$ is nonempty.

If $\max\{d(p, A), d(p, B)\} \leq r$, then $d(p, A \cap B) \leq R$.

Proof. Note that a concatenation of a geodesic of length r with a geodesic of any length is a $(1, 2r)$–quasi-geodesic. Let $M_0 = M(1, 2r)$. Let R be the number of pointed oriented simplicial paths in X of length $\leq 2M_0$, up to the G–action.

Let q be the closest point in $A \cap B$ to p. Suppose $d(p, q) > R$, and let γ be a geodesic from p to q. Every vertex on γ lies within M_0 of both A and B. By our choice of R, there must be a pair of distinct vertices a_1, a_2 on γ and paths σ_i joining a_i to A, and τ_i joining a_i to B of length at most M_0, and an element $h \in G$, so that $ha_1 = a_2$, $h\sigma_1 = \sigma_2$ and $h\tau_1 = \tau_2$. We may assume that a_1 is closer to q than a_2 is.

Since $hA \cap A$ and $hB \cap B$ are nonempty, h must stabilize both A and B. Thus $hq \in A \cap B$. But hq is closer to p than q is, contradicting our choice of q. □

Towards proving Theorem 3.26, suppose that G is a finitely generated group acting cocompactly on a CAT(0) cube complex X, and suppose that hyperplane stabilizers are strongly quasi-convex in G. An index 2 subgroup of a strongly quasi-convex subgroup is strongly quasi-convex, so the stabilizers of hyperplanes in X^b are strongly quasi-convex in G. We build an appropriate G–cocompact subgraph of $\mathrm{Gr}\left(\widetilde{CG(\mathcal{Y})}\right)$, using the structure of intersections of hyperplane stabilizers.

This graph will be $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$, for some coherent choices of $\mathbb{B}^{\mathfrak{C}}$ over all of the representative lists $\mathfrak{C} \in \mathcal{W}$ (see Definitions 3.12 and 3.22).

Definition 3.28. Suppose that $\mathfrak{C} = (W_1, W_2, \ldots, W_k) \in \mathfrak{W}$ is a representative ordered list of hyperplanes as above. An *initial segment of* $\mathfrak{C}$ is a list $(W_1, \ldots, W_i)$ for some $i < k$.

We remark that by the way the representative lists were chosen, if $\mathfrak{C} \in \mathfrak{W}$ then any initial segment of $\mathfrak{C}$ is also in $\mathfrak{W}$.

Definition 3.29. Suppose that $\mathfrak{C} \in \mathfrak{W}$. We define

$$\mathcal{I}(\mathfrak{C}) = \{\mathfrak{C}' \in \mathfrak{W} \mid \mathfrak{C} \text{ is an initial segment of } \mathfrak{C}'\}.$$

The following holds because the stabilizer of each $\mathfrak{C}$ is an intersection of hyperplane stabilizers, and because intersections of strongly quasi-convex subgroups are strongly quasi-convex by [35, Theorem 1.2.(2)]. Tran also proves in [35, Theorem 1.2.(1)] that strongly quasi-convex subgroups are finitely generated.

Lemma 3.30. *Each* $\mathrm{Stab}(\mathfrak{C})$ *is strongly quasi-convex and hence finitely generated.*

Now choose finite generating sets $\mathcal{A}_{\mathfrak{C}}$ for each $\mathrm{Stab}(\mathfrak{C}) \in \mathfrak{W}$. By Lemma 3.21, each element g of $\mathcal{A}_{\mathfrak{C}}$ can be represented by a $CG(\mathcal{Y})$–path $d_{\mathfrak{C}} \cdot p_g \cdot \overline{d_{\mathfrak{C}}}$ where p_g is a cubical $CG(\mathcal{Y})$–path so that the lift of p_g to $\widetilde{CG(\mathcal{Y})}$ starting at $\widetilde{v_{\mathfrak{C}}}$ is a concatenation of arrows in $\mathrm{Ar}(\widehat{\Gamma}(\mathfrak{C}))$.

For an object v, choose $\mathbb{B}_v^{\mathfrak{C}}$ to consist of the following collections of group arrows:

1. all of the group arrows at v that occur in the paths p_g for $g \in \mathcal{A}_{\mathfrak{C}}$;
2. all of the group arrows at v in the p_g for $g \in \mathcal{A}_{\mathfrak{C}'}$ for any $\mathfrak{C}' \in \mathcal{I}(\mathfrak{C})$; and

Use these choices to define a set $\mathbb{B}^{\mathfrak{C}}$ and a graph $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ as in Definition 3.22. The choices made above give us such a graph $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ for each $\mathfrak{C} \in \mathfrak{W}$.

Proposition 3.31. *For each* $\mathfrak{C} \in \mathfrak{W}$ *the graph* $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ *is connected, and* $\mathrm{Stab}(\mathfrak{C})$*–invariant. Moreover* $\mathrm{Stab}(\mathfrak{C})$ *acts freely and cocompactly on* $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$*. For any* $\mathfrak{C}' \in \mathcal{I}(\mathfrak{C})$ *we have* $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}}) \subset \Gamma_{\mathfrak{C}'}(\mathbb{B}^{\mathfrak{C}'})$.

Proof. Other than the final statement, the result follows immediately from the construction and Proposition 3.23. The final statement follows immediately from Lemma 3.15 and the second condition in the choice of $\mathbb{B}_v^{\mathfrak{C}}$. □

Finally, choose a finite generating set $\mathcal{A}$ for G, and represent each element of $\mathcal{A}$ as a cubical path as in Proposition 3.8. Let $\mathbb{A}$ consist of all of the group arrows appearing in these paths, together with all of the group arrows $\mathbb{B}^{\mathfrak{C}}$ for $\mathfrak{C} \in \mathfrak{W}$, and use this set to build the graph $\Gamma(\mathbb{A})$. The following is entirely analogous to Proposition 3.31.

Proposition 3.32. *The graph* $\Gamma(\mathbb{A})$ *is connected and* G *acts freely and cocompactly on* $\Gamma(\mathbb{A})$*. For any* $\mathfrak{C} \in \mathfrak{W}$ *we have* $\Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}}) \subset \Gamma(\mathbb{A})$.

For the remainder of this subsection, we write $\Gamma = \Gamma(\mathbb{A})$ and for $\mathfrak{C} \in \mathfrak{W}$ we write $\Gamma_{\mathfrak{C}} = \Gamma_{\mathfrak{C}}(\mathbb{B}^{\mathfrak{C}})$ for the choices of $\mathbb{A}$ and $\mathbb{B}^{\mathfrak{C}}$ as made above.Denote the restriction of the map $\Psi\colon \Gamma(\mathbb{U}) \to (X^b)^{(1)}$ to Γ by Ψ_Γ. Note that Ψ_Γ is continuous, Lipschitz and G–equivariant.

We will need the following lemma in order to apply Proposition 3.27.

Lemma 3.33. *Let* $\mathfrak{C} = (W_1, \ldots, W_k) \in \mathfrak{W}$*, and let* $\mathfrak{C}^- = (W_1, \ldots, W_{k-1})$*. Then* $\Gamma_{\mathfrak{C}^-} \cap \Psi_\Gamma^{-1}(W_k) \neq \emptyset$.

Proof. This intersection is nonempty since W_k intersects $I_{\mathfrak{C}^-}$ nontrivially and Ψ_Γ surjects the 1–skeleton of X^b. □

The group G acts properly and cocompactly on Γ and the strongly quasi-convex subgroup $\mathrm{Stab}(\mathfrak{C})$ acts properly and cocompactly on $\Gamma_{\mathfrak{C}} \subset \Gamma$. Therefore, by considering an orbit map $G \to \Gamma_{\mathfrak{C}}$ and applying Theorem 3.25, we see that each $\Gamma_{\mathfrak{C}}$ is a strongly quasi-convex subset of Γ. Also, for each hyperplane W of X^b the set $\Psi_\Gamma^{-1}(W)$ is a strongly quasi-convex subset of Γ. For $\mathfrak{C} \in \mathfrak{W}$ let $M_{\mathfrak{C}}$ be a Morse gauge for $\Gamma_{\mathfrak{C}}$, and for W a hyperplane of X^b, let M_W be a Morse gauge for $\Psi_\Gamma^{-1}(W)$. Note that there are finitely many distinct gauges M_W as W ranges over all the hyperplanes of X^b. Define the Morse gauge M to be the maximum of the $M_{\mathfrak{C}}$ and the M_W.

We now give the main part of the argument of the proof of Theorem 3.26, namely that if hyperplane stabilizers are strongly quasi-convex, vertex stabilizers are also strongly quasi-convex. We therefore fix a vertex v of X.

Note that $\Psi_\Gamma^{-1}(v)$ is a non-empty 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 quasi-convex in G, it suffices (by Theorem 3.25) to show that the pre-image $\Psi_\Gamma^{-1}(v)$ is a strongly quasi-convex subset of Γ.

We fix constants $K \geq 1$ and $C \geq 0$, suppose that a and b are vertices in $\Psi_\Gamma^{-1}(v)$ and let γ be a (K, C)–quasi-geodesic in Γ between a and b. Let y be an arbitrary vertex on γ. We have to show $d(y, \Psi_\Gamma^{-1}(v))$ is bounded independent of a and b.

Here is a description of our bound: Let D be a bound for the valence of Γ, N a bound for the number of G–orbits of vertices in Γ. Let $R_1 = M(K, C)$. Assuming R_i has been defined, we let R_{i+1} be the maximum of R_i and the constant R in the conclusion of Proposition 3.27, with the above D, N and with $r = R_i + 1$. We will prove that $d(y, \Psi_\Gamma^{-1}(v))$ is bounded above by $R_{\dim X} + 1$.

If $\Psi_\Gamma(y) = v$, there is nothing to prove, so we assume that $\Psi_\Gamma(y) \neq v$.

We will build a sequence of points $z_1, \ldots, z_t$ for some $t \leq \dim X$ so that for each i, the following conditions are satisfied:

$$(*_i) \qquad \begin{cases} d(y, z_i) \leq R_i + 1; \text{ and} \\ \exists g_i \in G,\ \mathfrak{C}_i \in \mathfrak{W}_i,\ \text{so that } v \in g_i I_{\mathfrak{C}_i},\ z_i \in g_i \Gamma_{\mathfrak{C}_i}. \end{cases}$$

We first find z_1. Let $[v, \Psi_\Gamma(y)]$ be a geodesic in the 1–skeleton of X^b from v to $\Psi_\Gamma(y)$. The first edge of $[v, \Psi_\Gamma(y)]$ is dual to some hyperplane W_1 of X^b.

Lemma 3.34.

$$d(y, \Psi_\Gamma^{-1}(W_1)) \leq M(K, C)$$

Proof. Since W_1 separates v from $\Psi_\Gamma(y)$, we know that γ must cross $\Psi_\Gamma^{-1}(W_1)$ between a and y. However, $\Psi_\Gamma(\gamma)$ is a loop, so γ must also cross $\Psi_\Gamma^{-1}(W_1)$ in the segment of γ between y and b. Thus, there is a (quasi-geodesic) subsegment γ_1 of γ which contains y and which starts and finishes on $\Psi_\Gamma^{-1}(W_1)$. Since M is a Morse gauge for $\Psi_\Gamma^{-1}(W_1)$, the lemma follows. □

Suppose that $\overline{W_1}$ is the chosen representative of W_1, so that $\mathfrak{C}_1 = (\overline{W}_1) \in \mathfrak{W}$. Suppose further that $W_1 = g_1 \overline{W_1}$. By the definition of $\Gamma_{\mathfrak{C}_1}$ in terms of the objects $V_\Box(\mathfrak{C})$, it is straightforward to see that $\Psi_\Gamma^{-1}(W_1)$ is contained in the 1–neighborhood of $g_1 \Gamma_{\mathfrak{C}_1}$ in Γ. Thus, there is a vertex z_1 in $g_1 \Gamma_{\mathfrak{C}_1}$ which is distance at most $M(K, C) + 1$ from y.

We claim that z_1 satisfies $(*_1)$ for this choice of g_1 and $\mathfrak{C}_1$. The only thing we have left to prove is that $v \in g_1 I_{\mathfrak{C}_1}$. This follows immediately from the choice of W_1.

Proposition 3.35. *Suppose z_i satisfies $(*_i)$. Then either $\Psi_\Gamma(z_i) = v$ or there is some z_{i+1} satisfying $(*_{i+1})$.*

Proof. Suppose that $\Psi_\Gamma(z_i) \neq v$. Choose a geodesic $[v, \Psi_\Gamma(z_i)]$ in the 1–skeleton of the combinatorially convex subcomplex $g_i I_{\mathfrak{C}_1} \subset X^b$. Let W_{i+1} be the hyperplane of X^b which is dual to the first edge of $[v, \Psi_\Gamma(z_i)]$. Then W_{i+1} separates v from $\Psi_\Gamma(z_i)$.

Claim.

$$d(y, \Psi_\Gamma^{-1}(W_{i+1})) \leq R_i + 1.$$

Proof. There are two cases.

Suppose first W_{i+1} separates v from $\Psi_\Gamma(y)$. In this case, we argue as in Lemma 3.34 that $d(y, \Psi_\Gamma^{-1}(W_{i+1})) \leq M(K, C) \leq R_i + 1$.

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

Suppose that $\mathfrak{C}_i = (\overline{W}_1, \ldots, \overline{W}_i)$. We choose g_{i+1} and $\overline{W}_{i+1}$ so that $\mathfrak{C}_{i+1} = (\overline{W}_1, \ldots, \overline{W}_i, \overline{W}_{i+1}) \in \mathfrak{W}$ and $g_{i+1}(\mathfrak{C}_i) = g_i \mathfrak{C}_i$ and $g_{i+1}(\overline{W}_{i+1}) = W_{i+1}$. The hyperplane $\overline{W}_{i+1}$ is uniquely determined, but g_{i+1} is only determined up to multiplication by something in $\mathrm{Stab}(\mathfrak{C}_{i+1})$.

We want to apply Proposition 3.27. Take $p = y$, $A = g_{i+1}\Gamma_{\mathfrak{C}_i}$ and $B = \Psi_\Gamma^{-1}(W_{i+1})$, and $r = R_i + 1$. The parameters D and N are the maximum valence and number of G–orbits in Γ, respectively.

That $gA \cap A \neq \emptyset$ implies $gA = A$ is contained in Proposition 3.17.

It is easy to see that if $gB \cap B \neq \emptyset$, then $gB = B$, since a hyperplane is determined by any edge dual to it.

Lemma 3.33 asserts that $A \cap B \neq \emptyset$. By the definition of R_{i+1}, Proposition 3.27 implies

$$d(y, g_{i+1}\Gamma_{\mathfrak{C}_i} \cap \Psi_\Gamma^{-1}(W_{i+1})) \leq R_{i+1}$$

The intersection $g_{i+1}\Gamma_{\mathfrak{C}_i} \cap \Psi_\Gamma^{-1}(W_{i+1})$ lies in the 1–neighborhood of $g_{i+1}\Gamma_{\mathfrak{C}_{i+1}}$, so there is a point $z_{i+1} \in g_{i+1}\Gamma_{\mathfrak{C}_{i+1}}$ with $d(y, z_{i+1}) \leq R_{i+1} + 1$. By the choice of W_{i+1} $v \in g_{i+1} I_{\mathfrak{C}_{i+1}}$, so we have established $(*_{i+1})$ and finished the proof. □

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 3.35 asserts that for some $i \leq \dim X$, $\Psi_\Gamma(z_i) = v$. We conclude that $d(y, \Psi_\Gamma^{-1}(v)) \leq R_i + 1 \leq R_{\dim X} + 1$, as desired.

This completes the proof of Theorem 3.26.

3.5. **If cell stabilizers are QC then hyperplane stabilizers are QC.** In this section we prove the direction (2) $\implies$ (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 quasi-convex in G. In particular, these vertex stabilizers will be finitely generated. Note that stabilizers of other cells are intersections of vertex stabilizers, so they are also quasi-convex, and so finitely generated.

Let W be a hyperplane in X^b. We simplify the notation set up in Section 3.3 by writing 'W' instead of the 1–tuple '(W)'.

Lemma 3.36. *Let $v \in V_{\square}(W)$, and let $H_{v,W}$ be as in Definition 3.19. There is a cube σ of X so that $H_{v,W}$ is naturally isomorphic to a finite-index subgroup of $G_{[\![\sigma]\!]}$.*

Proof. The stabilizer in G of some $v \in V_{\square}(W)$ naturally corresponds to the stabilizer of a pair $(\sigma_0 \subset \sigma)$, a codimension–1 inclusion of cubes of X. But $\mathrm{Stab}(W) \cap \mathrm{Stab}(\Theta_W(v))$ is exactly the stabilizer of $(\sigma_0 \subset \sigma)$, which has finite-index in the stabilizer of σ, which in turn is isomorphic to $G_{[\![\sigma]\!]}$. □

By Lemma 3.36 the groups $H_{v,W}$ from Definition 3.19 are finitely generated. For each such $H_{v,W}$, let $\mathbb{B}_v$ be a finite generating set. For each object $\mathfrak{o}$ in $\mathcal{Y}$, choose a finite generating set $\mathbb{A}_{\mathfrak{o}}$ for $G_{\mathfrak{o}}$ so that if $H_{v,W} \le G_{\mathfrak{o}}$ for some v, W then $\mathbb{B}_v \subseteq \mathbb{A}_{\mathfrak{o}}$. Form the graphs $\Gamma_W = \Gamma_W(\mathbb{B}^W)$ and $\Gamma = \Gamma(\mathbb{A})$ as described in Definitions 3.9 and 3.22.

Proposition 3.37. *The graph Γ is connected and G–invariant and G acts freely and cocompactly on Γ. The graph Γ_W is connected and* $\mathrm{Stab}(W)$*–invariant, and* $\mathrm{Stab}(W)$ *acts freely and cocompactly on Γ_W. Moreover, $\Gamma_W \subset \Gamma$.*

Proof. The special cases that $\mathfrak{C} = ()$ and $\mathfrak{C} = (W)$ of Proposition 3.23 give the first two sentences of the proposition. The final assertion follows quickly from the condition that $\mathbb{B}_v \subseteq \mathbb{A}_{\mathfrak{o}}$. □

Given a vertex w of $W^{\downarrow}$, let $Y(w)$ be the (closed) 1-neighborhood in Γ_W of $\Psi_W^{-1}(w)$.

Lemma 3.38. *The sets $Y(w)$ are quasi-convex subsets of Γ with constants which do not depend on w.*

Proof. Since $\mathrm{Stab}(W)$ acts cocompactly on $W^{\downarrow}$, 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 quasi-convex subset of Γ.

The stabilizer in $\mathrm{Stab}(W)$ of $Y(w)$ is the same as the subgroup $H_{w,W}$ from Definition 3.19. By Lemma 3.36 this is a finite index subgroup of some cell group of the G–action on X. Thus, by hypothesis, $H_{w,W}$ is a quasi-convex subgroup of G. Since G acts properly and cocompactly on Γ, and $H_{w,W}$ acts properly and cocompactly on $Y(w) \subset \Gamma$, the result follows (for example, by Theorem 3.25). □

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

Theorem 3.39. *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 quasi-convex. Then, for every hyperplane $W \subset X$, the stabilizer* $\mathrm{Stab}(W)$ *is a quasi-convex subgroup of G.*

Proof. As we have already remarked, quasi-convexity of vertex stabilizers implies quasi-convexity of all cell stabilizers.

Let Γ, Γ_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 so that $Y(w)$ is ϵ–quasi-convex for every w (Lemma 3.38).

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

Consider a geodesic γ in $(X^b)^{(1)}$ between $\Psi(p)$ and $\Psi(q)$. Both $\Psi(p)$ and $\Psi(q)$ lie in $W^{\downarrow}$. Since $W^{\downarrow}$ is combinatorially convex in X^b, the geodesic γ is entirely contained in the 1-skeleton of $W^{\downarrow}$. The vertices $w_1, \ldots, w_n$ on γ correspond to cells of X contained in $W^{\downarrow}$. The sets $Y(w_i)$ corresponding to these cells satisfy the hypotheses of Theorem A.3 with $m = 2$, $c = 1$, and ϵ the quasi-convexity constant chosen above. Theorem A.3 then implies that $Y(w_1) \cup \cdots \cup Y(w_n)$ is ϵ'–quasi-convex, 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 Γ_W, the Γ–geodesic between p and q stays uniformly close to Γ_W, as required. □

Together with Theorem 3.26, this completes the proof of Theorem A.

3.6. **On generalizations of Theorem A.** For a subgroup H of a hyperbolic group G, the following three conditions are equivalent:

(a) H is strongly quasi-convex in G.
(b) H is quasi-convex in G.
(c) H is undistorted in G.

Dropping the condition that G is hyperbolic, condition (b) ceases to be well-defined, but the 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).

3.6.1. *Strong quasi-convexity.* Replacing quasi-convexity with strong quasi-convexity 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 quasi-convex.
(2S) Vertex stabilizers are strongly quasi-convex.
(3S) All cell stabilizers are strongly quasi-convex.

As remarked earlier (2S)$\Longleftrightarrow$(3S) follows from [35, Theorem 1.2.(2)]. Theorem 3.26 states that (1S) $\Longrightarrow$ (2S)

The remaining implication (3S) $\Longrightarrow$ (1S) is *false*, as the example of $\mathbb{Z}^2$ acting freely on a cubulated $\mathbb{R}^2$ shows.

3.6.2. *Undistortedness.* The situation when replacing quasi-convexity 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 implies (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 3.40. *For finitely generated groups acting cocompactly on CAT*(0) *cube complexes does (1U)* $\Longrightarrow$ *(3U)? Does (3U)* $\Longrightarrow$ *(1U)?*

3.7. **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 3.41. [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 so that for every tuple of distinct cosets $(g_0H_0, g_1H_1, \ldots, g_nH_n)$ with $H_i \in \mathcal{H}$ (and $g_i \in G$), the intersection $\cap_{i=0}^{n} H_i^{g_i}$ is finite. If there is no such n then we say the height of $\mathcal{H}$ is infinite.

In case $\mathcal{H} = \{H\}$ 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 in that paper) 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 quasi-convex subgroups by Tran [35, Theorem 1.2.(3)].

Proposition 3.42. *Let G be a hyperbolic group and $\mathcal{H}$ a finite collection of quasi-convex subgroups of G. Then the height of $\mathcal{H}$ is finite.*

We also use the following special case of a theorem of Charney–Crisp [9, Theorem 5.1].

Theorem 3.43. *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 quasi-convex hyperplane stabilizers. Then*

(1) X is δ–hyperbolic for some δ;
(2) there exists a $k \geq 0$ so 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, p. 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 quasi-convex, then [7, Theorem 7.11], due to Bowditch, implies that this coned graph is δ–hyperbolic for some δ. Theorem 3.43 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 quasi-convex subgroups of G, so it has some finite height k by Proposition 3.42. If $H < G$ is infinite with nonempty fixed 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 [14], 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:

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

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 3.44. In the context where G is a finitely generated group acting cocompactly on a cube complex X with strongly quasi-convex hyperplane stabilizers, the same proof of conclusion (2) works as written, replacing the reference to Proposition 3.42 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 Theorem A, Theorem F, Agol's Theorem [1, Theorem 1.1] and Wise's Quasi-convex Hierarchy Theorem [36, Theorem 13.3]. Thus, other than Theorem A.3 in Appendix A (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 $\overline{X} = {}_K\backslash^X$ to be a CAT(0) cube complex:

(1) $\overline{X}$ must be simply-connected;
(2) $\overline{X}$ must be a cube complex (rather than a complex made out of cells which are quotients of cubes); and
(3) $\overline{X}$ must be non-positively 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 $\epsilon > 0$ so that for every finite subcomplex K, the ϵ–neighborhood of K deformation retracts to K. If $f\colon S^1 \to X$ is any loop, then a compactness argument shows it lies in an ϵ–neighborhood of some such finite complex. We can then homotope f to have image in K 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 $n \in \{2, 3, \ldots\}$, let D_n be the Euclidean cone of radius 1 on a loop σ_n of length $\frac{2\pi}{n^2}$. For each n mark a point on σ_n. Let Y be obtained from $\bigcup D_n$ 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 $\Gamma = \langle \gamma_2, \gamma_3, \ldots \rangle$ acting on $\widetilde{Y}$ with quotient Y, so that each γ_n fixes the center of some disc and rotates it by an angle of $\frac{2\pi}{n^2}$. 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 set-wise but not point-wise.

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 $\{\sigma_1, \ldots, \sigma_k\}$ 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 so that for each i we have $G_{\sigma_i} \cap K \leq Q_{\sigma_i}$. 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 2.2). 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 2.8, 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 $K \trianglelefteq G$ so that ${}_K\backslash^X$ is a cube complex. [4]

In Subsection 4.2 we discuss how to find subgroups K so 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. *Through this section we write $CG(\mathcal{Y})$–paths as a concatenation of group arrows and scwol arrows (which need not alternate between group arrows and scwol arrows). Thus in a list of arrows such as $g_1 \cdot e_1 \cdot e_2 \cdot g_2 \cdot \ldots$, each g_i is an element of a local group G_v and represents the edge $(g, 1_v)^+$, corresponding to the group arrow $(g, 1_v)$. The e_i represent $a_i^\pm$ for a scwol arrow $(a_i, 1)$, and we blur the distinction between the scwol arrow $(a_i, 1)$ in $CG(\mathcal{Y})$ and the $\mathcal{Y}$-arrow a_i, and also between the $CG(\mathcal{Y})$–path $(1, a_i)^\pm$ and the $\mathcal{Y}$–edge $a_i^\pm$. We implicitly*

[4] We remark that 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.

assume that each concatenation we write defines a path, which often forces the group arrows labelled 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.*

In case we have a group arrow of the form $(1, 1_v)$*, we often implicitly (or explicitly) omit this arrow from our path.*

Theorem 4.5. *(Gromov's Cubical Link condition,* [8, Theorem II.5.20]*) The cube complex complex Z is a non-positively curved cube complex if and only if for each vertex $\overline{v} \in Z$ the link of $\overline{v}$ is flag.*

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

The link of a cube σ in a cube complex has a natural cellulation by spherical simplices, one coming from each inclusion of σ into a higher-dimensional cube. In particular $\mathrm{lk}(\sigma)$ is a Δ–*complex* [21, Chapter 2.1] (though it may not be simplicial).

We record two elementary observations:

Lemma 4.6. *Let v be a vertex of a cube complex, and let $L = \mathrm{lk}(v)$. Then L is simplicial if and only if for every cell σ containing v, the* 1*–skeleton of* $\mathrm{lk}(\sigma)$ *contains no immersed loop of length* 1 *or* 2*.*

Proof. If L fails to be simplicial, there is either a non-embedded simplex, or a pair of simplices which intersect in a set which is not a face of both. If a simplex is non-embedded, 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 $F_i \setminus f$. Then the simplices spanned by $v_1 \cup f$ and $v_2 \cup f$ correspond to points in $\mathrm{lk}(f)$ which lie on an immersed loop of length 2. But f corresponds to some cube containing σ, and $\mathrm{lk}(f) \subset L$ is isomorphic to the link of that cube. □

Lemma 4.7. *Let v be a vertex of a cube complex, and suppose $L = \mathrm{lk}(v)$ is simplicial. Then L is a flag complex if and only if for every cell σ containing v, every loop of length* 3 *in the* 1*–skeleton of* $\mathrm{lk}(\sigma)$ *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 $\mathrm{lk}(\sigma)$ and also ensure that any loop of length 3 in $\mathrm{lk}(\sigma)$ 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* $\mathrm{lk}(\sigma)$. Below we choose, for each cube σ in Z and each 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.)

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

A 1–cell α in $\mathrm{lk}(\sigma)$ corresponds to a triple of cubes $\epsilon_1, \epsilon_2, \phi$ in Z so 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, an oriented 1–cell α of $\mathrm{lk}(\sigma)$ has idealization a $\mathcal{Z}$–path of length 4, made up of arrows

$$(\sigma \subset \epsilon_1) \longleftarrow (\sigma \subset \epsilon_1 \subset \phi) \longrightarrow (\sigma \subset \phi) \longleftarrow (\sigma \subset \epsilon_2 \subset \phi) \longrightarrow (\sigma \subset \epsilon_2). \tag{1}$$

In what follows, we want a slightly more general situation, where $\epsilon_1, \epsilon_2, \phi$ are cubes in Z with ϵ_1, ϵ_2 codimension–1 sub-cubes of ϕ, and γ is a chain of cubes in Z so that each element of γ is contained in each of $\epsilon_1, \epsilon_2, \phi$. We can naturally extend γ to chains which we denote $(\gamma \subset \epsilon_1), (\gamma \subset \epsilon_2)$, $(\gamma \subset \phi)$, and $(\gamma \subset \epsilon_i \subset \phi)$, and this triple of cubes correspond to a 1–cell in an 'iterated link' (a link of a cell in a link, etc.), and also has idealization a $\mathcal{Z}$–path of length 4 as follows:

$$(\gamma \subset \epsilon_1) \longleftarrow (\gamma \subset \epsilon_1 \subset \phi) \longrightarrow (\gamma \subset \phi) \longleftarrow (\gamma \subset \epsilon_2 \subset \phi) \longrightarrow (\gamma \subset \epsilon_2). \tag{2}$$

The $\mathcal{Z}$–path (2) may not embed in $\mathcal{Y}$. There are two ways this could happen. The first is that there is an element of $\mathrm{Stab}(\gamma)$ which sends ϵ_1 to ϵ_2, but no such element fixes ϕ. In this case, the image in $\mathcal{Y}$ is a non-backtracking 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-edge'.

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 $\mathcal{Z}$–path above. Let ν be the projection of $(\gamma \subset \phi)$ in $\mathcal{Y}$, and μ_i be the projection of $(\gamma \subset \epsilon_i \subset \phi)$. Let ξ_i be the projections of $(\gamma \subset \epsilon_i)$. Then we have the injective homomorphisms

$$\psi_{a_2} \colon G_{\mu_1} \to G_\nu,$$

and

$$\psi_{a_3} \colon G_{\mu_2} \to G_\nu.$$

The images of ψ_{a_2} and ψ_{a_3} are equal. The projections of all the data associated to α depend only on the orbit of α under the stabilizer of σ. We shall denote this orbit by $[\![\alpha]\!]$, and 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 (in case there is a g fixing ϕ and exchanging ϵ_1 with ϵ_2) or else $G_\nu^+ = G_\nu$ (if there is no such g). In case G_ν^+ has index 2 in G_ν, we fix a choice of $g_\nu \in G_\nu \smallsetminus G_\nu^+$. 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 $\psi_{\overline{[\![\alpha]\!]}} \colon G_\nu^+ \to G_{i([\![\alpha]\!])}$ by

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

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

Definition 4.8. In case $G_\nu^+ \neq G_\nu$, 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 $g_\nu \in G_\nu \smallsetminus G_\nu^+$ is the element fixed above. Note that in this case $a_2 = a_3$ and $a_1 = a_4$. In case $G_\nu^+ = G_\nu$ let

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

In either case some lift of p_α to $\mathcal{C}_K$ 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 $p_{[\![\alpha]\!]}$. By slight abuse of notation, we define $\overline{[\![\alpha]\!]} = [\![\overline{\alpha}]\!]$.

Lemma 4.9. *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

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

so 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.

Remark 4.10. We remark that in case $G_\nu^+ \neq G_\nu$ the paths considered in the second half of the above statement are exactly those $CG(\mathcal{Y})$–paths traversing $y_{[\![\alpha]\!]}$ which lift and scwolify (using Θ) to non-backtracking paths in $\mathcal{X}$ (see the discussion at the end of Section 2.5). As $Z = {}_K\backslash^X$ is a cube complex, these $CG(\mathcal{Y})$–paths also lift and scwolify (using Θ_K) to non-backtracking paths in $\mathcal{Z}$.

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 $[\![.]\!]$ to denote a G-orbit in $\mathcal{Z}$, which corresponds to its image in $\mathcal{Y}$ under the projection $\pi \colon \mathcal{Z} \to \mathcal{Y}$.*

Let $p_{[\![\alpha]\!]}$ be one of the $CG(\mathcal{Y})$–paths fixed in Definition 4.8, 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 as $G_{i([\![\alpha]\!])}$ and $G_{t([\![\alpha]\!])}$.

We will also need to refer to the subgroups $E_{[\![\alpha]\!]} < G_{t([\![\alpha]\!])}$ and $E_{\overline{[\![\alpha]\!]}} < G_{i([\![\alpha]\!])}$ defined just before Definition 4.8. 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* $\mathrm{lk}(\sigma)$. We are now ready to formulate the conditions on K which characterize whether or not ${}_K\backslash^X$ is non-positively curved. We use Lemmas 4.6 and 4.7 repeatedly.

Recall that we have fixed a $K \lhd G$ so that $Z = {}_K\backslash^X$ is a cube complex. We also fix a cube σ of Z. If α is a 1–cell in $\mathrm{lk}(\sigma)$, there is a corresponding $\mathcal{Z}$–path of length 4, and we sometimes conflate the two.

We next give an algebraic characterization of loops of length 1 in $\mathrm{lk}(\sigma)$.

Lemma 4.12. *Let α be a 1–cell in $\mathrm{lk}(\sigma)$. The endpoints of α 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 Z as the geometric realization of $\mathcal{Z}$, the 1–cell α is the realization of a $\mathcal{Z}$–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 $\widehat{q}_\alpha$ be an unscwolification of q_α in $\mathcal{C}_K$, which we may choose to have label

$$a_1^+ \cdot a_2^- \cdot g_1 \cdot a_3^+ \cdot a_4^-, \tag{3}$$

for some group arrow g_1.

Suppose first that the endpoints of α coincide. Then the path (3) has 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 4.9 implies that this loop is homotopic to a loop with label $p_{[\![\alpha]\!]} \cdot g$ for some g.

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 translate of q_α by some element of G. In particular, q_α must be a loop, and so the endpoints of α coincide. □

Definition 4.13. If v is a vertex of $\mathcal{Y}$, let $K_v \lhd G_v$ be $K \cap G_v$.

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

Lemma 4.15. *A $CG(\mathcal{Y})$–path $g_0 \cdot p_{[\![\alpha_1]\!]} \cdot g_1 \cdot p_{[\![\alpha_2]\!]} \cdot \ldots \cdot g_{k-1} \cdot p_{[\![\alpha_k]\!]} \cdot g_k$ is K–non-backtracking if and only if the first of the following two conditions holds. A $CG(\mathcal{Y})$–loop with such a label is K–non-backtracking if and only if both conditions hold.*

1. *For $i \in \{1, \ldots, k-1\}$, if $[\![\alpha_{i+1}]\!] = \overline{[\![\alpha_i]\!]}$ then $g_i \notin E_{[\![\alpha_i]\!]} K_{t([\![\alpha_i]\!])} \subset G_{t([\![\alpha_i]\!])}$; and*
2. *If $[\![\alpha_1]\!] = \overline{[\![\alpha_k]\!]}$ 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 $\mathrm{lk}(\sigma)$.

Lemma 4.16. *Let p be a path in $\mathrm{lk}(\sigma)$ which is a concatenation of two 1–cells, α and β. The following are equivalent:*

1. *There is a path $p' = \alpha'.\beta'$ in $\mathrm{lk}(\sigma)$ with $[\![\alpha']\!] = [\![\alpha]\!]$ and $[\![\beta']\!] = [\![\beta]\!]$ so that p' is an immersed loop.*
2. *There is a K–non-backtracking $CG(\mathcal{Y})$–loop $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{[\![\beta]\!]} \cdot g_2$ that represents a conjugacy class in K.*

Moreover, in case 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 a immersed loop $p' = \alpha'.\beta'$. The idealization of p' is a scwol-path $q_{p'}$ of length 8 in $\mathcal{Z}$, 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 4.9, we can choose an unscwolification $\widehat{q}_{p'}$ of $q_{p'}$ in $\mathcal{C}_Y$ with label

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

where g_1 is a group arrow. But the unscwolification $\widehat{q}_{p'}$ 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–non-backtracking since its scwolification is the path $q_{p'}$.

Conversely, suppose that there is a K–non-backtracking $CG(\mathcal{Y})$–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 so that there are no edge-loops of length 1 or 2. Any edge-loop in Q of length 3 is non-backtracking.*

The utility of Lemma 4.17 is that once we have found conditions to ensure that links in ${}_K\backslash^X$ have no edge-loops of length 1 or 2 then edge-loops of length 3 are automatically non-backtracking.

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

Lemma 4.18. *Suppose that $\mathrm{lk}(\sigma)$ is simplicial, and suppose that p is a path in $\mathrm{lk}(\sigma)$ which is a concatenation of three 1–cells, α, β and γ. The following are equivalent:*

1. *There is a path $p' = \alpha'.\beta'.\gamma'$ in $\mathrm{lk}(\sigma)$ so that $[\![\alpha']\!] = [\![\alpha]\!]$, $[\![\beta']\!] = [\![\beta]\!]$ and $[\![\gamma']\!] = [\![\gamma]\!]$, 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, in case 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 so that there are loops of length 3 in $\mathrm{lk}(\sigma)$. This introduces some subtleties, which we discuss in the next subsection.

4.3.3. *Loops of length 3 filled by 2–cells.* 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 $G(\mathcal{Y})$ so 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 $K_xK_yK_z$, 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 σ be a cube of Z, and let τ be a 2–cell in $\mathrm{lk}(\sigma)$. Then $\partial\tau$ is a loop composed of three oriented 1–cells $\alpha.\beta.\gamma$. These 1–cells are associated to $CG(\mathcal{Y})$–paths $p_{[\![\alpha]\!]}, p_{[\![\beta]\!]}, p_{[\![\gamma]\!]}$ as in Definition 4.8. Consider a $CG(\mathcal{Y})$–path of the form $q = p_{[\![\alpha]\!]} \cdot g \cdot p_{[\![\beta]\!]}$. Let $\widehat{q}$ be a lift to $\mathcal{C}_K$. The realization of $\Theta_K(\widehat{q})$ is a concatenation of two 1–cells $\alpha'.\beta'$. We say that *q K–bounds a (τ,α)–corner* if there is a cube σ', a 2–cell τ' in $\mathrm{lk}(\sigma')$, and an $h \in G$ so that $\sigma' = h\sigma$, $\tau' = h\tau$, $\alpha' = h\alpha$ and $\beta' = h\beta$. 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.*

In case there exists a path q as above which K–bounds a (τ,α)–corner, there are cubes $\epsilon, \phi_\alpha, \phi_\beta$ and ψ, all containing σ, so 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 $\epsilon, \phi_\alpha, \phi_\beta$ and ψ determine an oriented 1–cell ζ in this copy of the link of ϵ. Its idealization is shown in Figure 1. The

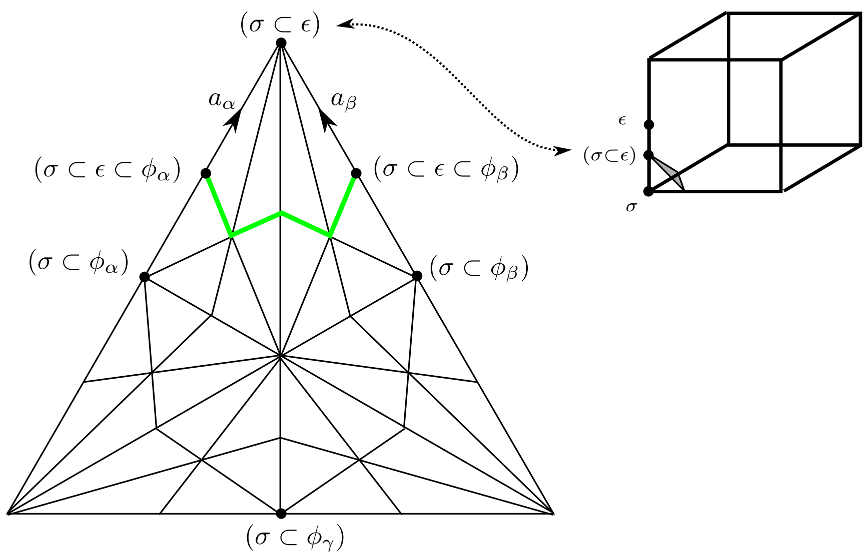

FIGURE 1. A part of $\mathcal{Z}$ representing part of the link of σ, containing the idealization of the 1–cell ζ in green. Directions of most arrows have been omitted.

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 4.8) 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 $\mathcal{Z}$ as follows:

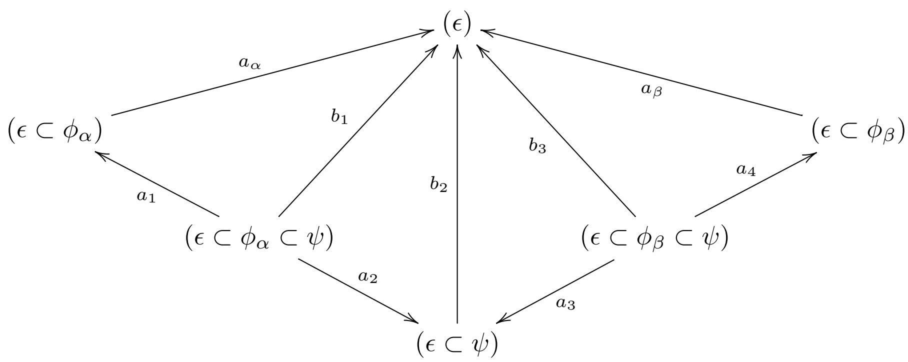

We have the following identities of morphisms in the category $\mathcal{Z}$: $a_\alpha a_1 = b_1 = b_2 a_2$ and $a_\beta a_4 = b_3 = b_2 a_3$. The path in the statement of the lemma is equal to:

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

where $g_{(\epsilon\subset\psi)}$ is the element of $G_{\llbracket(\epsilon\subset\psi)\rrbracket}$ chosen for the path p_ζ as in Definition 4.8.

Define the following elements of $G_{\llbracket\epsilon\rrbracket}$:

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

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

We now have the following sequence of elementary homotopies of $CG(\mathcal{Y})$–paths (all of which consist of applying the moves in Definition 2.4, and the rule of arrow composition in $CG(\mathcal{Y})$ from Definition 2.12).

$$\begin{aligned} \llbracket a_\alpha \rrbracket^+ \cdot p_{\llbracket\zeta\rrbracket} \cdot \llbracket a_\beta \rrbracket^- &\simeq \llbracket a_\alpha \rrbracket^+ \cdot \llbracket a_1 \rrbracket^+ \cdot \llbracket a_2 \rrbracket^- \cdot g_{(\epsilon\subset\psi)} \cdot \llbracket a_3 \rrbracket^+ \cdot \llbracket a_4 \rrbracket^- \cdot \llbracket a_\beta \rrbracket^- \\ &\simeq z(\llbracket a_\alpha \rrbracket, \llbracket a_1 \rrbracket) \cdot \llbracket b_1 \rrbracket^+ \cdot \llbracket a_2 \rrbracket^- \cdot g_{(\epsilon\subset\psi)} \cdot \llbracket a_3 \rrbracket^+ \cdot \llbracket a_4 \rrbracket^- \cdot \llbracket a_\beta \rrbracket^- \\ &\simeq z(\llbracket a_\alpha \rrbracket, \llbracket a_1 \rrbracket) z(\llbracket b_2 \rrbracket, \llbracket a_2 \rrbracket)^{-1} \cdot \llbracket b_2 \rrbracket^+ \cdot \llbracket a_2 \rrbracket^+ \cdot \llbracket a_2 \rrbracket^- \cdot g_{(\epsilon\subset\psi)} \cdot \llbracket a_3 \rrbracket^+ \cdot \llbracket a_4 \rrbracket^- \cdot \llbracket a_\beta \rrbracket^- \\ &\simeq h_1 \cdot \llbracket b_2 \rrbracket^+ \cdot g_{(\epsilon\subset\psi)} \cdot \llbracket a_3 \rrbracket^+ \cdot \llbracket a_4 \rrbracket^- \cdot \llbracket a_\beta \rrbracket^- \\ &\simeq h_1 \cdot \psi_{\llbracket b_2 \rrbracket}(g_{(\epsilon\subset\psi)}) \cdot \llbracket b_2 \rrbracket^+ \cdot \llbracket a_3 \rrbracket^+ \cdot \llbracket a_4 \rrbracket^- \cdot \llbracket a_\beta \rrbracket^- \\ &\simeq h_2 z(\llbracket b_2 \rrbracket, \llbracket a_3 \rrbracket) \cdot \llbracket b_3 \rrbracket^+ \cdot \llbracket a_4 \rrbracket^- \cdot \llbracket a_\beta \rrbracket^- \\ &\simeq h_2 z(\llbracket b_2 \rrbracket, \llbracket a_3 \rrbracket) z(\llbracket a_\beta \rrbracket, \llbracket a_4 \rrbracket)^{-1} \cdot \llbracket a_\beta \rrbracket^+ \cdot \llbracket a_4 \rrbracket^+ \cdot \llbracket a_4 \rrbracket^- \cdot \llbracket a_\beta \rrbracket^- \\ &\simeq h_3 \cdot \llbracket a_\beta \rrbracket^+ \cdot \llbracket a_\beta \rrbracket^- \\ &\simeq h_3 \end{aligned}$$

This proves the result. □

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

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

$$\llbracket a_\alpha \rrbracket^+ \cdot g_1 \cdot p_{\llbracket\zeta\rrbracket} \cdot g_2 \cdot \llbracket a_\beta \rrbracket^- \cdot g^{-1} \tag{4}$$

which represents an element of K.

Proof. First suppose that there is a $CG(\mathcal{Y})$–loop of the form (4) representing an element of K.

Then the following two $CG(\mathcal{Y})$–paths differ by an element of K:

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

and

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

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 τ' between the edges α' and β' which are the images 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.

Now 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 $\widehat{q}$ and consider the scwolification $\Theta_K(\widehat{q})$ in $\mathcal{Z}$. As in Definition 4.20, the realization of $\widehat{q}$ is the concatenation of two 1–cells α', β' in $\mathrm{lk}(\sigma')$ for some cube σ' in the orbit of σ. Moreover, there is a 2–cell τ' with α', β' in the boundary of τ' and an element h of G so that $\sigma' = h\sigma$, $\tau' = h\tau$, $\alpha' = h\alpha$ and $\beta' = h\beta$. Let v' be the vertex of $\mathrm{lk}(\sigma')$ where α' and β' 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 $\widehat{q_0}$ of q_0 to $\mathcal{C}_K$ so that $\overline{q}_0 = \Theta_K(\widehat{q_0})$ is a loop based at v' and traveling across the corner of τ' from α' to β'. The paths q_0 and q have lifts to $\mathcal{C}_K$ forming a sub-diagram:

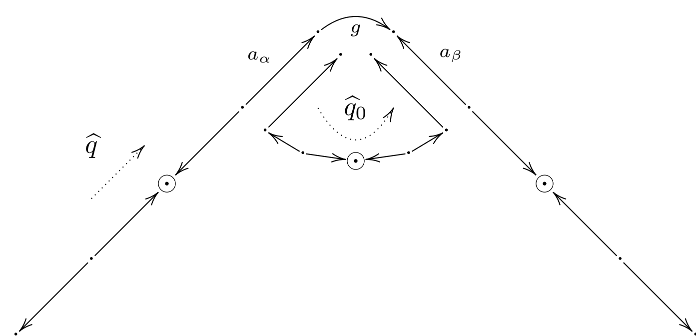

The circled dots represent either single objects or pairs of objects separated by a group arrow, depending on whether the paths $p_{[\![x]\!]}$ have length four or five for

$x \in \{\alpha, \beta, \zeta\}$. The scwolification of this diagram in $\mathcal{Z}$ looks like this:

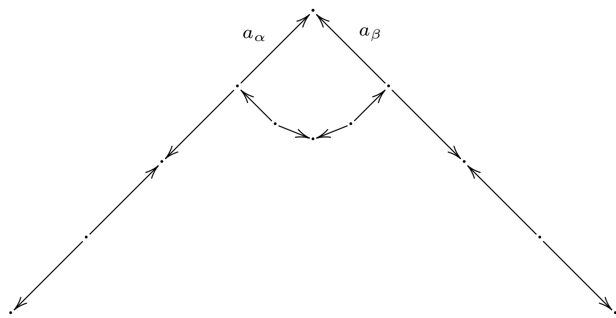

The edges which scwolify to a_α in $\widehat{q}$ and $\widehat{q}_0$ have sources connected by a group arrow labeled by some g_1. Similarly the edges which scwolify to a_β have sources connected by group arrow with some label g_2. We thus obtain a loop in $\mathcal{C}_K$ of the form (4). □

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 $\mathrm{lk}(\sigma)$ 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) $g \in E_{[\![\alpha]\!]} g_{\tau,\alpha} E_{\overline{[\![\beta]\!]}}.K_{t([\![\alpha]\!])}$

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

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

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 $K \cap G_{[\![v]\!]} = K_{t([\![\alpha]\!])}$, we have

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

as required.

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

Lemma 4.25. *Suppose that $\mathrm{lk}(\sigma)$ is simplicial and contains 1–cells α, β, and γ. 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 $\overline{q} \subset \mathrm{lk}(\sigma)$ 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 $\overline{q}$ bounds a 2–cell in $\mathrm{lk}(\sigma)$.

Proof. Note that since q represents an element of K, any lift to $\mathcal{C}_K$ is a loop, and so the realization $\overline{q}$ is also a loop. Since $\mathrm{lk}(\sigma)$ is simplicial, this loop is embedded of length 3 in $\mathrm{lk}(\sigma)$ 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 relabelling 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 τ and p K–bounds a (τ, α)–corner. It follows that some translate τ' of τ in $\mathrm{lk}(\sigma)$ has boundary given by a path $\alpha'.\beta'.\gamma'$, where $\alpha'.\beta'$ are the first two 1–cells of the path $\overline{q}$. If the third 1–cell of $\partial\tau'$ is not the third 1–cell of $\overline{q}$, we obtain 1–cells in $\mathrm{lk}(\sigma)$ with the same endpoints, contradicting the assumption that $\mathrm{lk}(\sigma)$ is simplicial. So $\overline{q}$ bounds the 2–cell τ'. □

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

Proposition 4.26. *Suppose that $\mathrm{lk}(\sigma)$ is simplicial. There are finitely many 2–cells τ_i in $\mathrm{lk}(\sigma)$ (with boundary $\alpha_i.\beta_i.\gamma_i$) so that $\mathrm{lk}(\sigma)$ is flag 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 so that

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

Proof. Choose the 2–cells τ_i to be representatives of the $\mathrm{Stab}(\sigma)$–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 form $(*)$ representing elements of K is satisfied, and suppose that p is an edge-loop of length 2 in $\mathrm{lk}(\sigma)$ which is labelled by 1–cells α', β', γ', 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 so 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 edge-loop of length 3 in $\mathrm{lk}(\sigma)$ 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 p of λ is an immersed edge-path of length 3, which hence must bound a 2–cell, τ say. Suppose that τ_i is the representative in the $\mathrm{Stab}(\sigma)$–orbit of the 2–cell τ, 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). This finishes the proof. □

To summarize, Lemmas 4.12, 4.16, 4.18, and Proposition 4.26 give descriptions of various types of $CG(\mathcal{Y})$–paths so 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 $K \lhd G$ so that ${}_K\backslash^X$ is non-positively 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 consider c_v also 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

$$p \mapsto c_{i(p)} \cdot p \cdot \overline{c_{t(p)}}.$$

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

The following results are all straightforward.

Lemma 5.1. *For any $CG(\mathcal{Y})$–paths p, p' so that $t(p) = i(p')$, we have*

$$\ell_{\overline{p}} = \ell_p^{-1}$$
$$\ell_{p \cdot p'} = \ell_p \ell_{p'}.$$

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*

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

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

$$Q_v = \{ [c_v \cdot g \cdot \overline{c_v}] \mid g \in G_v \}.$$

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 non-positively curved, there are five conditions that need to be ensured on links in Z. Roughly speaking, they are:

(1) No edge-loop of length 1,
(2) No edge-loop of length 2 consisting of 1–cells in different G–orbits,
(3) No edge-loop of length 2 consisting of 1–cells in the same G–orbit,
(4) No edge-loop of length 3 whose image in $\mathcal{Y}$ does not bound a 2–cell, and
(5) No edge-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 $\widehat{p}$ and a lift $\widetilde{p}$ of $\widehat{p}$ to $\widetilde{CG(\mathcal{Y})}$ so that the realization of the scwolification of $\widetilde{p}$ 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 4.6, 4.7 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* co-cubical *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 $F_1 \subset \mathcal{Q}(G) \times G$ so that for each $(Q, p) \in F_1$ we have $p \notin Q$ and so that if (i) $K \trianglelefteq G$ is co-cubical; and (ii) for each $(Q, p) \in F_1$ we have $p \notin Q.K$, then no link in ${}_K\backslash^X$ contains an edge-loop of length* 1*.*

Proof. Up to the action of G, there are finitely many pairs $(\widetilde{\sigma}, \widetilde{\alpha})$, where $\widetilde{\sigma}$ is a cube of X and $\widetilde{\alpha}$ is a 1–cell in link$(\widetilde{\sigma})$ 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 co-cubically, α is embedded in link(σ), 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 $\widetilde{\alpha}$ is not a loop. Applying Lemma 4.12 in 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 $G_{t([\![\alpha]\!])}$). This proves the result. $\square$

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

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

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

and so that if (i) $K \trianglelefteq G$ *is co-cubical; and (ii) for each* $(Q_1, Q_2, p_1, p_2) \in F_2$ *we have*

$$K \cap p_1 Q_2 p_2 Q_2 = \emptyset$$

then every edge-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 $[\![\alpha]\!], [\![\beta]\!]$ so 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 $p_1 = \ell_{p_{[\![\alpha]\!]}}, p_2 = \ell_{p_{[\![\beta]\!]}}, Q_1 = Q_{t([\![\alpha]\!])}, Q_2 = Q_{t([\![\beta]\!])}$.

Since X is a CAT(0) cube complex, there are no non-backtracking edge-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. $\square$

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

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

$$1 \notin p_1 \left(Q_1 \smallsetminus Q_2^{p_2}\right) p_2 \left(Q_2 \smallsetminus Q_1^{p_1}\right), \tag{5}$$

and so that if (i) $K \trianglelefteq G$ *is co-cubical; (ii) no link in* ${}_K\backslash^X$ *contains an edge-loop of length 1; and (iii) for every* $(Q_1, Q_2, p_1, p_2) \in F_3$ *we have*

$$K \cap p_1 \left(Q_1 \smallsetminus \left(Q_2^{p_2} \left(K \cap Q_1\right)\right)\right) p_2 \left(Q_2 \smallsetminus \left(Q_1^{p_1} \left(K \cap Q_2\right)\right)\right) = \emptyset \tag{6}$$

then no link in ${}_K\backslash^X$ *contains an immersed edge-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 so that the following hold.

(1) There is some $g \in G$ so that g fixes $t(\tilde{\alpha})$ but not $\tilde{\alpha}$.
(2) There is some $h \in G$ so 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 Equation (6) is satisfied for this element, then ${}_K\backslash^X$ will contain no immersed edge-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 $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{\overline{[\![\alpha]\!]}} \cdot g_2$ representing a conjugacy class in K. We let $p_1 = l_{p_{[\![\alpha]\!]}}$, $p_2 = l_{p_{\overline{[\![\alpha]\!]}}}$, $Q_1 = Q_{t([\![\alpha]\!])}$ and $Q_2 = Q_{i([\![\alpha]\!])}$. Using Lemma 4.15, the loop $p_{[\![\alpha]\!]} \cdot g_1 \cdot p_{\overline{[\![\alpha]\!]}} \cdot g_2$ is K–nonbacktracking if and only if $g_1 \notin E_{[\![\alpha]\!]} K_{t([\![\alpha]\!])}$ and $g_2 \notin E_{\overline{[\![\alpha]\!]}} K_{i([\![\alpha]\!])}$. The subgroup of Q_1 corresponding to $E_{\overline{[\![\alpha]\!]}}$ is equal to $Q_1 \cap Q_2^{p_2}$, and the subgroup of Q_2 corresponding to $E_{\overline{[\![\alpha]\!]}}$ 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 case $K = \{1\}$ and $K\backslash X = X$ is CAT(0), we see that our tuple satisfies Equation (5). For an arbitrary K we see that when Equation (6) is satisfied, there is no immersed edge-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 so that no link in $K\backslash X$ contains an edge-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 $F_4 \subset \mathcal{Q}(G)^3 \times G^3$ so that for each $(Q_1, Q_2, Q_3, p_1, p_2, p_3) \in F_4$ we have*

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

and so that if (i) $K \trianglelefteq G$ is co-cubical; (ii) no link in $K\backslash X$ contains an edge-loop of length 1 or 2; and (iii) for all $(Q_1, Q_2, Q_3, p_1, p_2, p_3) \in F_4$ we have

$$K \cap p_1 Q_1 p_2 Q_2 p_3 Q_3 = \emptyset$$

then every edge-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 co-cubical, 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 edge-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*

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

and let

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

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 $F_5 \subseteq \mathcal{Q}(G)^3 \times G^6$ so that for each $A = (Q_1, Q_2, Q_3, p_1, p_2, p_3, h_1, h_2, h_3)$ we have*

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

and so that if (i) $K \trianglelefteq G$ is co-cubical; (ii) no link in $K\backslash X$ contains an edge-loop of length 1 or 2; and (iii) for all $(Q_1, Q_2, Q_3, p_1, p_2, p_3, h_1, h_2, h_3) \in F_5$ we have

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

then no link in $K\backslash X$ contains an edge-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, γ of these orbits so that there is a 2–cell τ with boundary $\alpha \cdot \beta \cdot \gamma$. We associate the elements $p_1 = \ell_{p_{[\![\alpha]\!]}}, p_2 = \ell_{p_{[\![\beta]\!]}}, p_3 = \ell_{p_{[\![\gamma]\!]}}$, $Q_1 = Q_{t([\![\alpha]\!])}, Q_2 = Q_{t([\![\beta]\!])}$, $Q_3 = Q_{t([\![\gamma]\!])}$, and $h_1 = [c_{t([\![\alpha]\!])} \cdot g_{\tau,\alpha} \cdot \overline{c_{t([\![\alpha]\!])}}]$, $h_2 = [c_{t([\![\beta]\!])} \cdot g_{\tau,\beta} \cdot \overline{c_{t([\![\beta]\!])}}]$, $h_3 = [c_{t([\![\gamma]\!])} \cdot g_{\tau,\gamma} \cdot \overline{c_{t([\![\gamma]\!])}}]$.

Once again, since X is a CAT(0) cube complex, we can apply Proposition 4.26 to see that

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

When the conditions $g_1 \in E_{[\![\alpha_i]\!]} g_{\tau_i,\alpha_i} E_{\overline{[\![\beta]\!]}} K_{t([\![\alpha_i]\!])}$, etc. 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 6.4 gives a 'weak separability' of certain multi-cosets, and generalizations of multi-cosets, and is used to find subgroups K which satisfy the conditions from Theorems 5.6–5.11. Theorem 6.4 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 6.8, 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 $\overline{\mathcal{P}}$ 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 using the notation

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

omitting mention of the particular filling kernels.

If $N_P \dot{<} P$ (i.e. 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 $S \subseteq \bigcup \mathcal{P} \setminus \{1\}$ so that P holds whenever $(\bigcup \mathcal{N}) \cap S = \emptyset$. 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 [18] for a 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. See [18, Section 3] for the construction and basic 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 quasi-convex subgroups.*

Recall that $\mathcal{P} = \{P_1, \ldots, P_n\}$ 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 3.41) 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 quasi-convex 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 quasi-convex 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 $H_1 \cap H_2^{g_2} \cap \ldots \cap H^{g_k}$ where the H_i are in $\mathcal{H}$ and the cosets $\{H_1, g_2H_2, \ldots, g_kH_k\}$ 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 3.42.

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

Lemma 6.3. *Suppose that G is hyperbolic and $\mathcal{H}$ is a finite collection of quasi-convex 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 quasi-convex in $(G, \mathcal{P})$.*

The definition we use for relatively quasi-convex is that from [2]. In [28, Appendix A] it is proved that this is the same notion as the various notions defined in [23]. A subgroup H is *full* if whenever P is a parabolic subgroup so that $H \cap P$ is infinite we have $H \cap P \dot{<} P$.

6.3. **The appropriate meta-condition.** The goal of this subsection is to prove Theorem 6.4 below. The special case that $n = 1$ and $S_1 = \emptyset$ is [2, Proposition 4.5], which is about keeping elements out of full quasi-convex subgroups when performing long Dehn fillings. Here we generalize to multi-cosets of full quasi-convex 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 6.6 below in place of the somewhat technical [2, Lemmas 4.1 and 4.2].[5]

Theorem 6.4. *Let $(G, \mathcal{P})$ be relatively hyperbolic, and let $\mathcal{Q}$ be a collection of full relatively quasi-convex 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:*

$$1 \notin p_1(Q_1 \setminus S_1) \cdots p_n(Q_n \setminus S_n) \tag{7}$$

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

$$p_1 t_1 \cdots p_n t_n \tag{8}$$

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

The five conditions in the conclusions of Theorems 5.6 – 5.11 each fall into the scheme of the conditions in Theorem 6.4. Therefore, we may apply Theorem 6.4 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 quasi-convex 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 6.5. *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 quasi-convex. Let $\sigma_1, \ldots, \sigma_k$ 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 $\mathcal{Q} = \{Q_1, \ldots, Q_k\}$.*

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

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

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

[5]Using such a Greendlinger Lemma in place of the results of [2] was suggested to us by Alessandro Sisto while we were collaborating on [17].

complex. Therefore, we may assume that the subgroup K is co-cubical (in the sense of Terminology 5.5).

It remains to show that for sufficiently long $\mathcal{Q}$–fillings ${}_K\backslash^X$ is non-positively curved. It follows from Theorems 5.6–5.8 and 6.4 that for sufficiently long $\mathcal{Q}$–fillings each link of each cell in ${}_K\backslash^X$ is simplicial. Thus it follows from Theorem 5.9, 5.11 and 6.4 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 non-positively curved by Theorem 4.5. $\square$

To prove Theorem 6.4, we use the following "Greendlinger Lemma" (cf. [17, Lemma 2.41]):

Theorem 6.6. *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 $g \in K \setminus \{1\}$, there is a horoball A so that*

1. *γ penetrates A to depth at least C_1, and*
2. *there is an element k of K stabilizing A, so 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, ϵ so that every L–local $(1, C_2)$–quasi-geodesic lies within an ϵ–neighborhood of any geodesic with the same endpoints. (Such L, ϵ only depend on δ and C_2. See [11, Ch. 3].)

Now choose a filling long enough so 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 $g \in K \setminus \{1\}$. Let γ be a geodesic from 1 to g; let $\overline{\gamma}$ be the projection to the cusped space X/K for G/K. Within an $(L + C_1 + 2\epsilon)$–neighborhood of the Cayley graph, $\overline{\gamma}$ is an L–local geodesic. But $\overline{\gamma}$ cannot be an L–local $(1, C_2)$–quasi-geodesic everywhere, since it is a loop with diameter larger than ϵ.

In particular, there is a subsegment σ of $\overline{\gamma}$ of length $l \le L$ so that the endpoints $\overline{a}$ and $\overline{b}$ 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 $d(\overline{a}, \overline{b}) < l - C_2$, there is some element $k \in K$ stabilizing A so that $d(a, kb) < l - C_2$, as desired. $\square$

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

Lemma 6.7. *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 quasi-convex 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 so 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 6.4. 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 6.7 above. Suppose that K is the kernel of a filling which is long enough to satisfy the conclusion of Theorem 6.6 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 $t_i \in Q_i \setminus ((K \cap Q_i)S_i)$, and suppose that g is chosen so that $d_X(1,g)$ is minimal amongst all such choices.

Since for each i we have $Q_i \setminus ((K \cap Q_i)S_i) \subseteq Q_i \setminus S_i$, 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 $g \in K \setminus \{1\}$, by Theorem 6.6 there exist a horoball A in X, an element $k \in K$ stabilizing A, and points a, b on γ at depth at least C_1 so 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 deeply 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 6.7, $P^a \cap Q_i^{p_1 t_1 \cdots p_i}$ 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. This completes the proof of Theorem 6.4. $\square$

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

Theorem 6.8. *Suppose that the hyperbolic group G acts cocompactly on the CAT(0) cube complex X, and that cell stabilizers are virtually special and quasi-convex. Let $\sigma_1, \ldots, \sigma_k$ 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 $\mathcal{Q} = \{Q_1, \ldots, Q_k\}$, 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

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

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

1. *$\overline{G}$ is hyperbolic;*
2. *$\overline{\mathcal{Q}}$ consists of virtually special quasi-convex subgroups of $\overline{G}$.*
3. *$K_{\mathcal{P}}$ is generated by elements in cell stabilizers.*
4. *For each i, we have $K_{\mathcal{P}} \cap \mathrm{Stab}(\sigma_i) \le Q_i$;*
5. *$\mathrm{height}(\overline{\mathcal{Q}}) < \mathrm{height}(\mathcal{Q})$.*
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 6.3, $(G, \mathcal{P})$ 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 quasi-convex 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 so 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 quasi-convex 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 $\dot{P}_i(Q) \dot{\lhd} P_i$ so that if each filling kernel N_i satisfies $N_i \le \dot{P}_i(Q)$ 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 $N_i \le \dot{P}_i$ then each of the induced fillings of each element of $\mathcal{Q}$ is virtually special. By [19, Proposition 4.6], the natural map from $\overline{Q}$ to $\overline{G}$ is injective for all sufficiently long fillings.[6] If we choose a sufficiently long peripherally finite filling of $(G, \mathcal{P})$ with $N_i \le \dot{P}_i$ then [19, Proposition 4.5] implies that each $\overline{Q}$ is quasi-convex in $\overline{G}$. This ensures Property (2).

[6] [19, Lemma 3.5] ensures that sufficiently long $\mathcal{Q}$–fillings are sufficiently wide, in the terminology of that paper.

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 6.5. □

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 6.8 to this action, to obtain a further quotient. By induction on height, we obtain the following result from the introduction.

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 quasi-convex. There exists a quotient $\overline{G} = G/K$ so 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. A quasi-convexity criterion

In this appendix, we give a criterion (Theorem A.3) for a possibly infinite union of quasi-convex sets in a hyperbolic space to be quasi-convex. This criterion is used in the forward direction of Theorem A: quasi-convex cell stabilizers imply quasi-convex 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 quasi-convex subsets.

Lemma A.1. *Suppose that Y is δ–hyperbolic, and $P \subset Y$ is a union of k ϵ–quasi-convex subsets $P_1, \ldots, P_k$ so that $P_i \cap P_{i+1} \neq \emptyset$ for each i. Then P is ρ–quasi-convex where*

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

Proof. Consider a pair of points $x \in P_r$, $y \in P_s$. 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 $r \leq i < s$, 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 $\gamma = [x, p_r, p_{r+1}, \ldots, p_{s-1}, y]$. Since the P_i are ϵ–quasi-convex, γ is contained in an ϵ–neighborhood of $P_r \cup \ldots \cup P_s \subset P$.

Consider the geodesic polygon with one side the geodesic $\sigma = [x, y]$ and the other sides the geodesics forming γ. Let $r_0 = \lfloor \frac{r+s}{2} \rfloor$, and consider the geodesic triangle σ, $[x, p_{r_0}], [p_{r_0}, y]$. 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 $r_1 = \lfloor \frac{r+r_0}{2} \rfloor$ 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 $[x, p_{r_1}], [p_{r_1}, p_{r_0}]$, 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 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. This proves the lemma. □

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 $R_{m,\epsilon,\delta}$ so that for all $R_0 > R_{m,\epsilon,\delta}$, we have*

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

The next result states that under appropriate hypotheses, the union of an arbitrary number of quasi-convex subsets is itself quasi-convex, 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' so 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–quasi-convex;*
2. *For each i we have $X_i \cap X_{i+1} \neq \emptyset$; 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 $\mathbf{X} = \cup X_i$ is ϵ'–quasi-convex.

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 $\mathbf{X}$ is ρ-quasi-convex with

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

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

Our strategy is to build a path between u and v which is (i) uniformly quasi-geodesic; and (ii) stays uniformly close to Y. The theorem then follows by quasi-geodesic stability. Choose a sequence of indices $t_0 = j, t_1, \ldots, t_{s-1}, t_s = k$ so that for each $0 \leq r \leq s-2$ we have

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

and

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

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–quasi-convex, 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 so that $\pi_r(x) \cap X_l \neq \emptyset$.

Claim A.3.1. *For any $v \in \{t_r, \dots, t_{r+1}\}$ there exists $x_v \in \gamma_r$ so that*
$$d_{\mathbb{N}}(v, I_r(x_v)) \leq \frac{1}{2}\left(\frac{2K+1}{m} + c\right).$$

Proof of Claim A.3.1. 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 $|k-l| \leq \frac{2K+1}{m} + c$.

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. This finishes the proof of Claim A.3.1. □

Suppose $0 \leq r \leq s-1$. Using Claim A.3.1, we can choose a point $x_r \in \gamma_r$ and a point $y_r \in \pi_r(x_r)$ so that $y_r \in X_{k_r}$ and
$$(\dagger) \qquad \left| k_r - \frac{t_r + t_{r+1}}{2} \right| \leq \frac{2K+1}{2m} + c.$$

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 2. We bound the Gromov product

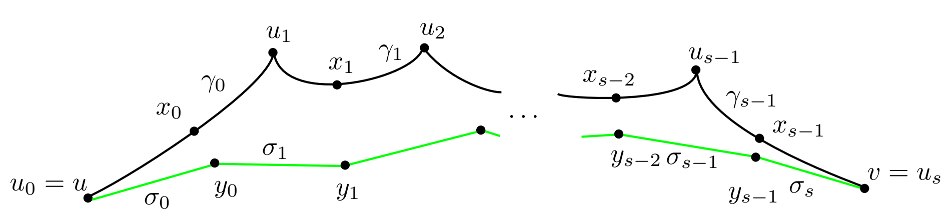

FIGURE 2. The σ_i forming a broken geodesic.

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 Equation (†) we have

$$\begin{aligned} k_r - k_{r-1} &\geq \frac{t_{r+1} - t_{r-1}}{2} - \frac{2K+1}{m} - 2c \\ &= 100R - \frac{2K+1}{m} - 2c \\ &> 50R - \frac{2K+1}{m} - 2c \end{aligned}$$

(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

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

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. This completes the proof of Claim A.3.2. □

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

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

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 3.

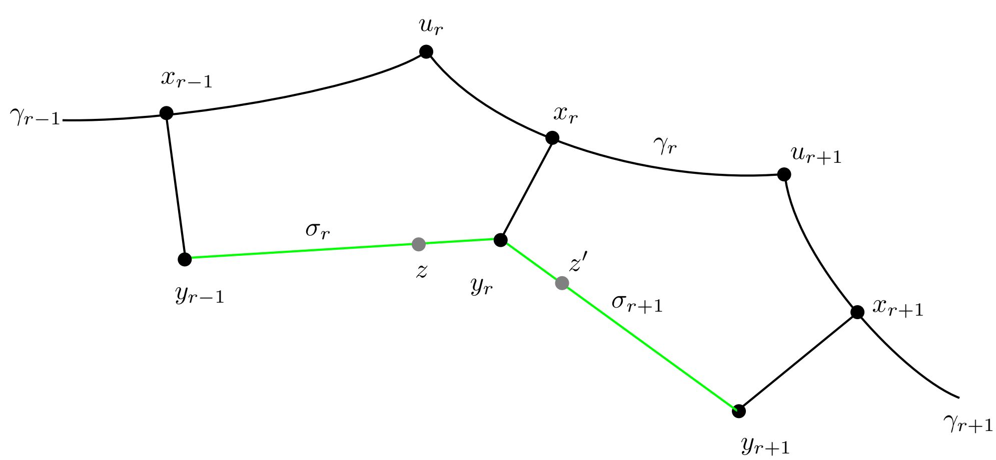

FIGURE 3. Computing the Gromov product of σ_r and σ_{r+1}.

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

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

using Lemma A.2 again. But this contradicts $d(x_r, w) \leq d(x_r, z) + d(z, w) \leq 9K + 3\delta \leq 12K$.

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 so that $d(z', w') \leq 3\delta$, and $d(x_r, w') \geq 4K$. Since γ_r is geodesic, we have

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

It follows that $d(z, z') \geq 8K - 6\delta \geq 2K > \delta$. It follows that 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 4. We are trying to show that $(u, y_1)_{y_0} \leq 8K$, so we may suppose

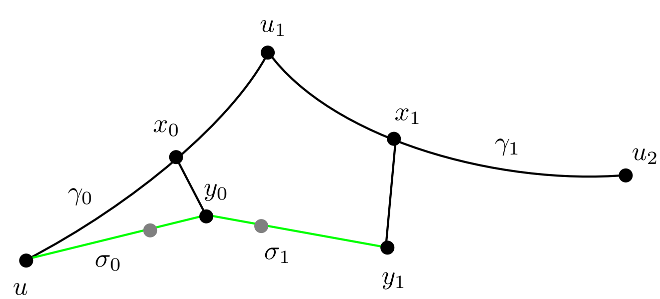

FIGURE 4. Computing the Gromov product of σ_0 and σ_1.

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 quasi-geodesic 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'–quasi-convex with the constant ϵ' depending on the quantities δ, m, and ϵ, but not on the number of the X_i, as required. This completes the proof of Theorem A.3. □

References

[1] I. Agol. The Virtual Haken Conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.

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

[3] I. Agol, D. Groves, and J. F. Manning. An alternate proof of Wise's malnormal special quotient theorem. *Forum Math. Pi*, 4:e1, 54, 2016.

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

[5] N. Bergeron and D. T. Wise. A boundary criterion for cubulation. *Amer. J. Math.*, 134(3):843–859, 2012.

[6] B. H. Bowditch. Tight geodesics in the curve complex. *Invent. Math.*, 171(2):281–300, 2008.

[7] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.

[8] M. R. Bridson and A. Haefliger. *Metric Spaces of Non–Positive Curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer–Verlag, Berlin, 1999.

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

[10] D. Cooper and D. Futer. Ubiquitous quasi-Fuchsian surfaces in cusped hyperbolic 3–manifolds. *Geom. Topol.*, 23(1):241–298, 2019.

[11] M. Coornaert, T. Delzant, and A. Papadopoulos. *Géométrie et théorie des groupes: Les groupes hyperboliques de Gromov*, volume 1441 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1990.

[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.

[13] E. Einstein and D. Groves. Relative cubulations and groups with a 2-sphere boundary, 2019.

[14] A. Genevois. Coning-off CAT(0) cube complexes. Preprint, arXiv:1603.06513, 2016.

[15] A. Genevois. Hyperbolicities in CAT(0) cube complexes. Preprint, arXiv:1709.08843, 2017.

[16] R. Gitik, M. Mitra, E. Rips, and M. Sageev. Widths of subgroups. *Trans. Amer. Math. Soc.*, 350(1):321–329, 1998.

[17] D. Groves, J. Manning, and A. Sisto. Boundaries of Dehn fillings. Preprint, arXiv:1612.03497, 2016.

[18] D. Groves and J. F. Manning. Dehn filling in relatively hyperbolic groups. *Israel Journal of Mathematics*, 168:317–429, 2008.

[19] D. Groves and J. F. Manning. Quasiconvexity and Dehn filling. Preprint, arXiv:1708.07968, 2017.

[20] F. Haglund and D. T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.

[21] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

[22] J. Hempel. 3-*Manifolds*. Princeton University Press, Princeton, N. J.; University of Tokyo Press, Tokyo, 1976. Ann. of Math. Studies, No. 86.

[23] G. C. Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.

[24] G. C. Hruska and D. T. Wise. Finiteness properties of cubulated groups. *Compos. Math.*, 150(3):453–506, 2014.
[25] W. Jaco. *Lectures on three-manifold topology*, volume 43 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, R.I., 1980.
[26] J. Kahn and V. Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)*, 175(3):1127–1190, 2012.
[27] I. Kapovich and H. Short. Greenberg's theorem for quasiconvex subgroups of word hyperbolic groups. *Canad. J. Math.*, 48(6):1224–1244, 1996.
[28] J. F. Manning and E. Martínez-Pedroza. Separation of relatively quasiconvex subgroups. *Pacific J. Math.*, 244(2):309–334, 2010.
[29] K. Matsuzaki and M. Taniguchi. *Hyperbolic Manifolds and Kleinian Groups*. Oxford University Press, Oxford, 1998.
[30] J. W. Morgan. On Thurston's uniformization theorem for three-dimensional manifolds. In *The Smith conjecture (New York, 1979)*, volume 112 of *Pure Appl. Math.*, pages 37–125. Academic Press, Orlando, FL, 1984.
[31] D. V. Osin. Peripheral fillings of relatively hyperbolic groups. *Invent. Math.*, 167(2):295–326, 2007.
[32] M. Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.
[33] M. Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.
[34] G. A. Swarup. Geometric finiteness and rationality. *J. Pure Appl. Algebra*, 86(3):327–333, 1993.
[35] H. C. Tran. On strongly quasiconvex subgroups. Preprint, arXiv:1707.05581, 2017.
[36] D. T. Wise. The structure of groups with a quasiconvex hierarchy. Unpublished manuscript.

Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago, 322 Science and Engineering Offices (M/C 249), 851 S. Morgan St., Chicago, IL 60607-7045

E-mail address: groves@math.uic.edu

Department of Mathematics, 310 Malott Hall, Cornell University, Ithaca, NY 14853

E-mail address: jfmanning@math.cornell.edu