Michigan Math. J. 62 (2013), 507–536

Coherence and Negative Sectional Curvature in Complexes of Groups

Eduardo Martínez-Pedroza & Daniel T. Wise

1. Introduction

A group G is *coherent* if finitely generated subgroups are finitely presented. A group G is *locally quasiconvex* if each finitely generated subgroup is quasiconvex. A subgroup H of G is *quasiconvex* if there is a constant L such that every geodesic in the Cayley graph of G that joins two elements of H lies in an L-neighborhood of H. While L depends upon the choice of Cayley graph, it is well known that the quasiconvexity of H is independent of the finite generating set when G is hyperbolic. As quasiconvex subgroups are finitely presented, it is clear local quasiconvexity implies coherence.

The class of coherent groups includes fundamental groups of compact 3-manifolds by a result of Scott [13], mapping tori of free group automorphisms by work of Feighn and Handel [3], and 1-relator groups with sufficient torsion by McCammond and Wise [11]. In contrast, the class of locally quasiconvex groups is substantially smaller. It includes fundamental groups of infinite-volume hyperbolic 3-manifolds by a result of Thurston ([12, Prop. 7.1] or [9, Thm. 3.11]), and there are criteria for local quasiconvexity for certain classes of small cancellation groups [8; 11].

Criteria for proving coherence and local quasiconvexity of groups acting freely on simply connected 2-complexes was introduced in [15] based on a notion of combinatorial sectional curvature. These methods do not apply on groups with torsion unless they are known to be virtually torsion free. It is an open question whether negatively curved groups are virtually torsion free [5].

In this paper we revisit the notion of combinatorial sectional curvature. We provide criteria for coherence and local quasiconvexity of groups acting properly and cocompactly on simply connected 2-complexes. This extends the methods in [15] to groups with torsion. Our extension of these results involves a generalization of the notion of sectional curvature and an extension of the combinatorial Gauss–Bonnet theorem to complexes of groups, and it surprisingly requires the use of ℓ^2-Betti numbers.

We revisit the following notion of sectional curvature in Section 3.

Definition 1.1 (Sectional Curvature $\leq \alpha$). An *angled 2-complex X* is a combinatorial 2-complex with an assignment of a real number to each corner of each

Received May 14, 2012. Revision received February 18, 2013.
Both authors are supported by NSERC.

2-cell of X. A locally finite angled 2-complex X has *sectional curvature at most* α if the following two conditions hold.

(i) For each 0-cell x and each finite subgraph Δ of $\operatorname{link}(x)$ containing a cycle but no valence-1 vertex, we have $\operatorname{Curvature}(\Delta) \leq \alpha$, where

$$\operatorname{Curvature}(\Delta) = 2\pi - \pi \cdot \chi(\Delta) - \sum_{e \in \operatorname{Edges}(\Delta)} \measuredangle(e)$$

and $\measuredangle(e)$ is the angle assigned to the corner $e \in \operatorname{Edges}(\Delta)$; each edge of the link of a 0-cell $x \in X$ corresponds to a corner of a 2-cell whose attaching map contains x.

(ii) For each 2-cell f of X, we have $\operatorname{Curvature}(f) \leq 0$, where

$$\operatorname{Curvature}(f) = \left(\sum_{c \in \operatorname{Corners}(f)} \measuredangle(c) \right) - \pi(|\partial f| - 2),$$

and $\operatorname{Corners}(f)$ denotes the set of corners of the 2-cell f.

DEFINITION 1.2 (Angled G-Complex). Let G be a group. A complex X equipped with a cellular G-action without inversions is a *G-complex*. A G-complex X is *proper* (resp., *cocompact*) if the G-action is proper (resp., cocompact). An *angled G-complex* is a 2-dimensional G-complex equipped with a G-equivariant angle assignment. A G-complex is trivial if it is empty or a single point.

THEOREM 1.3 (Cocompact Core). *Let X be a simply connected, proper, and cocompact angled G-complex with negative sectional curvature. If H is a subgroup of G and $Y \subseteq X$ is a connected H-cocompact subcomplex of X, then there is a simply connected H-cocompact subcomplex Z such that $Y \subseteq Z \subseteq X$.*

The strategy of the proof is as follows. A sequence of H-equivariant immersions $Y_n \to X$ is constructed inductively from the inclusion map $Y_0 = Y \to X$. The complex Y_{i+1} is obtained from Y_i by either "killing a loop" or correcting a "failure of injectivity". From the construction, a computation shows that the orbifold Euler characteristic of $H \backslash Y_n$ is bounded from below by the first ℓ^2-Betti number of $H \backslash Y_0$. Using that X has negative sectional curvature, an analysis of the structure of Y_n shows that the number of orbits of 0-cells with nonnegative curvature does not increase with n. Then, using a version of the combinatorial Gauss–Bonnet theorem for orbihedra and the previous two upper bounds, we obtain that the number of orbits of 0-cells of Y_n with negative curvature is uniformly bounded, and hence the total number of orbits of 0-cells of Y_n is uniformly bounded. Then a counting argument shows that there are finitely many possibilities for the immersions $Y_i \to X$ up to G-equivalence, and therefore the sequence $Y_n \to X$ stabilizes in an embedding $Z \to X$ of a simply connected complex.

COROLLARY 1.4 (Coherence Criterion). *Let G be a group admitting a proper cocompact action on a simply connected 2-complex with negative sectional curvature. Then each finitely generated subgroup of G is finitely presented.*

Proof. Let X a simply connected proper and cocompact G-complex with negative sectional curvature. Since X is connected, for any finitely generated subgroup $H \leq G$ there is a connected and cocompact H-subcomplex of X; by Theorem 1.3 this subcomplex can be assumed to be simply connected. Then the corollary follows by the well-known fact that a group is finitely presented if and only if it acts properly and cocompactly on a simply connected 2-complex [7]. □

CONJECTURE 1.5. *Let X be a simply connected, proper, and cocompact angled G-complex with sectional curvature ≤ 0. If H is a subgroup of G and $Y \subseteq X$ is a connected H-cocompact subcomplex of X, then there is a simply connected H-cocompact subcomplex Z such that $Y \subseteq Z \subseteq X$.*

The main result of the paper is the following criterion for local quasiconvexity. A subspace Y of a geodesic space X is *quasiconvex* if there is a constant L such that every geodesic in X that joins two elements of Y lies in the L-neighborhood of Y.

THEOREM 1.6 (Quasiconvex Cores). *Let X be a 2-dimensional proper cocompact* CAT(0) *G-complex whose cells are convex. Assign angles as they arise from the* CAT(0)*-metric. Suppose X has negative sectional curvature. If $H < G$ and Y is a simply connected cocompact H-subcomplex, then Y is a quasiconvex subspace of X.*

COROLLARY 1.7 (Local Quasiconvexity Criterion). *Let G be a group admitting a proper cocompact action on a 2-dimensional* CAT(0)*-complex with convex cells and negative sectional curvature. Then G is a locally quasiconvex hyperbolic group.*

Proof. Let X be a G-complex as in the statement. Since X has negative sectional curvature and angles are positive, X satisfies Gersten's negative weight test [15, Lemma 2.11]. It follows that X satisfies a linear isoperimetric inequality and hence X is a δ-hyperbolic space and G is a hyperbolic group [4, Thm. A6]. By Theorem 1.3, every finitely generated subgroup H of G admits a simply connected and H-cocompact subcomplex of X; then Theorem 1.6 implies that these subcomplexes are quasiconvex in X. □

The strategy of the proof of Theorem 1.6 is the following. Let ℓ be a geodesic with respect to the CAT(0)-metric in X with endpoints in the 0-skeleton of Y. Since ℓ is not a combinatorial path in the cell structure of X, we approximate ℓ with a combinatorial path $P_L \to X$ that has the same endpoints and is uniformly close. We also take a path $P_Y \to Y$ between the endpoints of ℓ. The choices of P_Y and P_L are made under some technical assumptions; in particular, they minimize the area of the disk diagram D with boundary cycle $P_L P_Y^{-1} \to X$. Let u be a 0-cell of $P_L \to X$. Analyzing the cell structure of D, we show that there exists a *good path* $Q \to D$ with initial point u and terminal point in $P_Y \to X$; by "good" we mean that if a 0-cell x of $Q \to D \to X$ intersects ℓ then x is in the interior of D. Using that X has a CAT(0)-structure, we construct an H-equivariant immersion

$Z' \to X'$, where X' is a subdivision of X, any good path $Q \to D \to X$ lifts as an internal path of $Q' \to Z'$ (after subdividing), and the number of H-orbits of 0-cells with negative curvature of Z' is bounded by a constant independent of ℓ. By the existence of good paths, we can take a path $Q \to X$ of minimal length from u to $P_Y \to X$ such that there is a lifting $Q' \to Z'$ that is internal. Then the proof concludes in the following way. Since X' has nonpositive sectional curvature, if a 0-cell x of Z' is internal (i.e., if $\mathrm{link}(x)$ has a cycle) then its curvature is nonpositive. Since X has negative sectional curvature, if x has zero curvature in Z' and is internal then its image in X is not a 0-cell. Therefore the length of $Q \to X$ equals the number of 0-cells of $Q' \to Z'$ with negative curvature plus 1. By minimality, no two 0-cells of $Q \to X$ are in the same H-orbit; therefore $|Q|$ is bounded by the number of orbits of 0-cells of Z' with negative curvature plus 1. By construction of Z', this number is uniformly bounded independently of ℓ.

A (p, q, r)-complex is a combinatorial 2-complex X such that the attaching map of each 2-cell has length $\geq p$, for each $x \in X^0$ the $\mathrm{link}(x)$ has girth $\geq q$, and each 1-cell e of X appears no more than r times among the attaching maps of 2-cells. The second author provides a criterion for negative sectional curvature of (p, q, r)-complexes in [16], from which the following application follows.

COROLLARY 1.8. *Let X be a 2-complex, where*:

(i) *X is a $(p, 3, p-3)$-complex for $p \geq 7$;*
(ii) *X is a $(p, 4, p-2)$-complex for $p \geq 5$; or*
(iii) *X is a $(p, 5, p-1)$-complex for $p \geq 4$.*

Then any group G acting properly and cocompactly on X is a locally quasiconvex hyperbolic group.

OUTLINE OF THE PAPER. Section 2 discusses preliminaries. Section 3 contains definitions of sectional curvature and generalized sectional curvature that are used in the rest of the paper. Section 4 discusses the combinatorial Gauss–Bonnet theorem for angled G-complexes. Section 5 recalls some results in the literature on ℓ^2-Betti numbers. Section 6 discusses equivariant immersions as a preliminary of the proof of the main results of the paper. The proof of the simply connected core theorem is the content of Section 7. Section 8 contains the proof of Theorem 1.6. The last section discusses a criterion establishing that certain quotients of locally quasiconvex groups are locally quasiconvex.

ACKNOWLEDGMENT. We thank the anonymous referee for a very careful reading, many suggestions, and for pointing out a serious gap in an earlier version of the paper.

2. Preliminaries

2.1. Complexes and Disk Diagrams

This paper follows the notation used in [10], and in this section we quote much of the relevant notation for the convenience of the reader. All complexes considered in this paper are combinatorial 2-dimensional complexes, and all maps are combinatorial.

Definition 2.1 (Path and Cycle) [10, Def. 2.5]. A *path* is a map $P \to X$, where P is a subdivided interval or a single 0-cell. A *cycle* is a map $C \to X$, where C is a subdivided circle. Given two paths $P \to X$ and $Q \to X$ such that the terminal point of P and the initial point of Q map to the same 0-cell of X, their concatenation $PQ \to X$ is the obvious path whose domain is the union of P and Q along these points. The path $P \to X$ is *closed* if the endpoints of P map to the same 0-cell of X. A path or cycle is *simple* if the map is injective on 0-cells. The *length* of the path P or cycle C is the number of 1-cells in the domain and is denoted by $|P|$ or $|C|$, respectively. The *interior of a path* is the path minus its endpoints. A *subpath* Q of a path P is given by a path $Q \to P \to X$ in which distinct 1-cells of Q are sent to distinct 1-cells of P.

Definition 2.2 (Disc Diagram) [10, Def. 2.6]. A *disc diagram* D is a compact contractible 2-complex with a fixed embedding in the plane. A *boundary cycle* P of D is a closed path in ∂D that travels entirely around D (in a manner respecting the planar embedding of D). For a precise definition we refer the reader to [11].

Let $P \to X$ be a closed null-homotopic path. A *disc diagram in X for P* is a disc diagram D together with a map $D \to X$ such that the closed path $P \to X$ factors as $P \to D \to X$, where $P \to D$ is the boundary cycle of D. Define $\text{area}(D)$ as the number of 2-cells in D.

Definition 2.3 (Arc) [10, Def. 5.4]. An *arc* in a diagram D is an embedded path $P \to D$ such that each of its internal 0-cells is mapped to a 0-cell with valence 2 in D. The arc is *internal* if its interior lies in the interior of D, and it is a *boundary arc* if it lies entirely in ∂D.

Definition 2.4 (Internal Path). A path $P \to X$ is *internal* if each 0-cell in the interior of P is mapped to a 0-cell of X whose link contains an embedded cycle.

Definition 2.5 (Links) [10, Def. 4.1]. Let X be a locally finite complex and let x be a 0-cell of X. The cells of X each have a natural partial metric obtained by making every 1-cell isometric to the unit interval and every n-sided 2-cell isometric to a Euclidean disc of circumference n whose boundary has been subdivided into n curves of length 1. In this metric, the set of points that are a distance equal to ε from x will form a finite graph. If ε is sufficiently small, then the graph obtained is independent of the choice of ε. This well-defined graph is the link of x in X and is denoted by $\text{link}(x, X)$.

Definition 2.6 (Immersions, Near-Immersions) [10, Def. 2.13]. The map $Y \to X$ is an *immersion* if it is locally injective. The map $Y \to X$ is a *near-immersion* if $Y \backslash Y^{(0)} \to X$ is locally injective. Equivalently, a map is an immersion if the induced maps on links of 0-cells are embeddings and a map is a near-immersion if the induced maps on links of 0-cells are immersions.

The following lemma is essential for the rest of the paper. It is an immediate consequence of the fact that immersions of graphs are π_1-injective.

LEMMA 2.7 (Near-Immersions Map Internal 0-Cells to Internal 0-Cells). *Let $X \to Y$ be a near-immersion mapping the 0-cell x to y. If* $\operatorname{link}(x)$ *has an embedded cycle then* $\operatorname{link}(y)$ *has an embedded cycle.*

DEFINITION 2.8 (Corners) [10, Def. 4.2]. Let X be a 2-complex, let x be a 0-cell of X, and let $R \to X$ be a 2-cell of X. Regard the 2-cells of X as polygons. Then the edges of $\operatorname{link}(x)$ correspond to the corners of these polygons attached to x. We will refer to a particular edge in $\operatorname{link}(x)$ as a *corner* of R at x.

2.2. *G-Complexes*

All group actions on complexes are without inversions; that is, a setwise stabilizer of a cell is a pointwise stabilizer. Under this assumption, quotients of complexes by group actions have an induced cell structure.

DEFINITION 2.9 ($I^p(G, X)$ and $|G_\sigma|^{-1}$). Let X be a G-complex X. Let $I^p(G, X)$ be the set of orbits of p-dimensional cells. For $\sigma \in I^p(G, X)$, let $|G_\sigma|^{-1}$ denote the reciprocal of the order of the G-stabilizer of a representative of σ in X, where $|G_\sigma|^{-1}$ is understood to be zero if the order is infinite.

DEFINITION 2.10 (Angled K-Graph). Let K be a group. An *angled K-graph* is a graph Γ equipped with a K-action and a K-map $\measuredangle\colon \operatorname{Edges}(\Gamma) \to \mathbb{R}$. For an edge e of Γ, the number $\measuredangle(e)$ is called the *angle at e*.

REMARK 2.11 (Connection with Angled G-Complexes). If X is an angled G-complex and x is a 0-cell of X, then each edge e of $\operatorname{link}(x)$ corresponds to a corner of a 2-cell of X and is thus associated to a real number $\measuredangle(e)$. This assignment of angles to the edges of $\operatorname{link}(x)$ is preserved under the G_x-action. In particular, $\operatorname{link}(x, X)$ is an angled G_x-graph.

3. Curvature

DEFINITION 3.1 (Curvature of K-Graphs). Let K be a group. For a cocompact angled K-graph Γ, define

$$\operatorname{Curvature}(K, \Gamma) = 2\pi \cdot |K|^{-1} - \sum_{v \in I^0(K,\Gamma)} \pi \cdot |K_v|^{-1} + \sum_{e \in I^1(K,\Gamma)} (\pi - \measuredangle(e)) \cdot |K_e|^{-1}.$$

DEFINITION 3.2 (Regular Section). Suppose that Γ is an angled K-graph and H is a subgroup of K. An *H-subgraph* Δ is an H-invariant subgraph of Γ, and an *H-section* is an H-subgraph that is H-cocompact. An edge having a vertex with valence 1 is called a *spur,* and a graph with no vertices of valence 1 is called *spurless.* A spurless, connected, and not edgeless H-section is called *regular.* An edgeless H-section is called *trivial.*

DEFINITION 3.3 (Sectional Curvature $\leq \alpha$). An angled complex X has *sectional curvature $\leq \alpha$* if the following two conditions hold.

(i) For each 0-cell x, each regular section of the angled $\mathbf{1}$-graph $\mathrm{link}(x)$ has curvature $\leq \alpha$, where $\mathbf{1}$ denotes the trivial group.
(ii) For each 2-cell f of X we have $\mathrm{Curvature}(f) \leq 0$, where

$$\mathrm{Curvature}(f) = \left(\sum_{c \in \mathrm{Corners}(f)} \measuredangle(c) \right) - \pi(|\partial f| - 2).$$

If $\alpha \leq 0$ then we say that X has *nonpositive sectional curvature*.

DEFINITION 3.4 (Generalized Sectional Curvature $\leq \alpha$). An angled G-complex X has *generalized sectional curvature* $\leq \alpha$ if:

(i) for each 0-cell x and each $H \leq G_x$, each regular H-section of the angled G_x-graph $\mathrm{link}(x)$ has curvature $\leq \alpha$; and
(ii) for each 2-cell f of X, we have $\mathrm{Curvature}(f) \leq 0$.

When X is a proper G-complex, these two notions are equivalent.

PROPOSITION 3.5. *Let X be a proper, cocompact, and angled G-complex with sectional curvature $\leq \alpha \leq 0$. Let $K = K(G, X)$ be an upper bound for the cardinality of 0-cell stabilizers. Then X has generalized sectional curvature $\leq \alpha/K$.*

Proof. Let x be a 0-cell of X, let H be a subgroup of G_x, and let Δ be an H-invariant regular section of $\mathrm{link}(x)$. The result follows by observing that

$$|H| \cdot \mathrm{Curvature}(H, \Delta) = \mathrm{Curvature}(\mathbf{1}, \Delta). \qquad \square$$

DEFINITION 3.6 (Corners, Sides). Let X be a G-complex and suppose that G acts without inversions on X. For $v \in I^0(G, X)$, let $\mathrm{Corners}(v)$ and $\mathrm{Sides}(v)$ denote the sets of edges and vertices of the link of v in $G \backslash X$. Let $\mathrm{Corners}(G, X)$ denote the disjoint union $\bigsqcup_{v \in I^0} \mathrm{Corners}(v)$, and analogously for $f \in I^2(G, X)$ let $\mathrm{Corners}(f)$ denote the subset of $\mathrm{Corners}(G, X)$ determined by f. For $e \in \mathrm{Corners}(v)$, let $|G_e|^{-1}$ denote $|G_\sigma|^{-1}$, where σ is the 2-cell of $G \backslash X$ determined by e. For $a \in \mathrm{Sides}(v)$ define $|G_a|^{-1}$ analogously.

REMARK 3.7. Each element of $\mathrm{Sides}(v)$ is determined by a 1-cell in $G \backslash X$. In particular, there is a natural two-to-one surjection

$$\bigsqcup_{v \in I^0} \mathrm{Sides}(v) \longrightarrow I^1(G, X).$$

DEFINITION 3.8. Let X be a cocompact G-complex. For $v \in I^0(G, X)$, the curvature $\kappa(v)$ is defined by

$$\kappa(v) = 2\pi \cdot |G_v|^{-1} - \sum_{e \in \mathrm{Sides}(v)} \pi \cdot |G_e|^{-1} + \sum_{c \in \mathrm{Corners}(v)} (\pi - \measuredangle(c)) \cdot |G_c|^{-1}.$$

The curvature of $f \in I^2(G, X)$ is defined by

$$\kappa(f) = \left[\left(\sum_{c \in \mathrm{Corners}(f)} \measuredangle(c) \right) - \pi(|\partial f| - 2) \right] \cdot |G_f|^{-1}.$$

REMARK 3.9. Let X be a cocompact angled G-complex, let $v \in I^0(G,X)$, and let $f \in I^2(G,X)$. Observe that $\kappa(v)$ and $\kappa(f)$ are finite real numbers. Moreover,

$$\kappa(v) = \text{Curvature}(G_x, \text{link}(x)),$$

where x is a representative of v in X, and analogously

$$\kappa(f) = \text{Curvature}(\sigma) \cdot |G_\sigma|^{-1},$$

where σ is a representative of f in X.

4. The Combinatorial Gauss–Bonnet Formula

DEFINITION 4.1. (Euler Characteristic). Let X be a 2-dimensional cocompact G-complex. The *Euler characteristic* $\chi(G,X)$ is defined by

$$\chi(G,X) = \sum_{\sigma \in I^0(G,X)} |G_\sigma|^{-1} - \sum_{\sigma \in I^1(G,X)} |G_\sigma|^{-1} + \sum_{\sigma \in I^2(G,X)} |G_\sigma|^{-1}.$$

THEOREM 4.2 (Combinatorial Gauss–Bonnet). *If X is an angled and cocompact G-complex, then*

$$2\pi \cdot \chi(G,X) = \sum_{v \in I^0(G,X)} \kappa(v) + \sum_{f \in I^2(G,X)} \kappa(f). \tag{1}$$

Proof. From the definition of $\kappa(v)$ and the natural two-to-one surjection of Remark 3.7, we have

$$\sum_{v \in I^0} \kappa(v) = \sum_{v \in I^0} 2\pi \cdot |G_v|^{-1} - \sum_{e \in I^1} 2\pi \cdot |G_e|^{-1} + \sum_{v \in I^0} \sum_{c \in \text{Corners}(v)} (\pi - \measuredangle(c)) \cdot |G_c|^{-1}. \tag{2}$$

Observe that

$$\begin{aligned} \sum_{f \in I^2} \kappa(f) &= \sum_{f \in I^2} 2\pi \cdot |G_f|^{-1} + \sum_{f \in I^2} \left[\left(\sum_{c \in \text{Corners}(f)} \measuredangle(c) \right) - \pi \cdot |\partial f| \right] \cdot |G_f|^{-1} \\ &= \sum_{f \in I^2} 2\pi \cdot |G_f|^{-1} - \sum_{f \in I^2} \sum_{c \in \text{Corners}(f)} (\pi - \measuredangle(c)) \cdot |G_f|^{-1}, \end{aligned} \tag{3}$$

where the first equality follows from the definition of $\kappa(f)$ and the second equality holds since $|\partial f| = |\text{Corners}(f)|$ for each 2-cell f.

Moreover,

$$\begin{aligned} \sum_{f \in I^2} \sum_{c \in \text{Corners}(f)} (\pi - \measuredangle(c)) \cdot |G_f|^{-1} &= \sum_{f \in I^2} \sum_{c \in \text{Corners}(f)} (\pi - \measuredangle(c)) \cdot |G_c|^{-1} \\ &= \sum_{v \in I^0} \sum_{c \in \text{Corners}(v)} (\pi - \measuredangle(c)) \cdot |G_c|^{-1}, \end{aligned} \tag{4}$$

where the first equality follows from $G_c = G_f$ for each $c \in \text{Corners}(f)$ and the second equality holds since $\{\text{Corners}(v)\}_{v \in I^0}$ and $\{\text{Corners}(f)\}_{f \in I^2}$ are partitions of the set $\text{Corners}(G,X)$.

The chains of equalities (3) and (4) imply that

$$\sum_{f\in I^2}\kappa(f)=\sum_{f\in I^2}2\pi\cdot|G_f|^{-1}-\sum_{v\in I^0}\sum_{c\in\mathrm{Corners}(v)}(\pi-\measuredangle(c))\cdot|G_c|^{-1}. \tag{5}$$

The Gauss–Bonnet formula (1) now follows by adding equations (2) and (5). □

5. Euler Characteristic and ℓ^2-Betti Numbers

For a G-complex X, the pth ℓ^2-Betti number $b_p^{(2)}(G,X)$ of X is a element of the extended interval $[0,\infty]$. We follow the approach of Lück and refer the reader to [6] for definitions and a general exposition on the subject. The approach of Lück to ℓ^2-Betti numbers fits the work of this paper because there are no assumptions on the G-action on X; in particular, the G-action is not required to be free.

THEOREM 5.1 (Atiyah's Formula) [6, Thm. 6.80]. *For a cocompact G-complex X,*

$$\chi(G,X)=\sum_{p\geq 0}(-1)^p\cdot b_p^{(2)}(G,X).$$

DEFINITION 5.2. For a G-space X and $H\leq G$, let X^H denote the subspace of X consisting of points fixed by all elements of H. For a G-map $X\to Y$ and $H<G$, we have $f(X^H)\subseteq Y^H$; denote by f^H the restriction of f to $X^H\to Y^H$.

THEOREM 5.3 [6, Thm. 6.54(1)]. *Let X and Y be G-complexes, and let $f\colon X\to Y$ be a G-map. Suppose for $n\geq 1$ that, for each subgroup $H\leq G$, the induced map $f^H\colon X^H\to Y^H$ is $\mathbb{C}$-homologically n-connected; that is, the map*

$$H_p^{\mathrm{sing}}(f^H;\mathbb{C})\colon H_p^{\mathrm{sing}}(X^H;\mathbb{C})\longrightarrow H_p^{\mathrm{sing}}(Y^H;\mathbb{C})$$

induced by f^H on singular homology with complex coefficients is bijective for $p<n$ and surjective for $p=n$. Then

$$b_p^{(2)}(G,X)=b_p^{(2)}(G,Y)\quad \textit{for}\ \ p<n,$$
$$b_n^{(2)}(G,X)\geq b_n^{(2)}(G,Y)\quad \textit{for}\ \ p=n.$$

THEOREM 5.4 [6, Thm. 6.54(3)]. *Let X be a G-complex. Suppose that for all $x\in X$ the stabilizer G_x is finite or satisfies $b_p^{(2)}(G_x)=0$ for all $p\geq 0$. Then $b_p^{(2)}(G,X)=b_p^{(2)}(G,EG\times X)$ for $p\geq 0$, where EG is a classifying space for G.*

COROLLARY 5.5. *Let X and Y be connected G-complexes, and let $f\colon X\to Y$ be a G-map such that the induced map on singular homology, $f_*\colon H_1(X,\mathbb{C})\to H_1(Y,\mathbb{C})$, is surjective. Suppose that for all $x\in X$ the stabilizer G_x is finite or satisfies $b_p^{(2)}(G_x)=0$ for all $p\geq 0$, and analogously for all $y\in Y$. Then $b_1^{(2)}(G,X)\geq b_1^{(2)}(G,Y)$.*

Proof. Since $EG\times X$ and $EG\times Y$ are connected free G-complexes, the map $\mathbf{1}\times f\colon EG\times X\to EG\times Y$ is $\mathbb{C}$-homologically 1-connected. By Theorem 5.3, it follows that $b_1^{(2)}(G, EG \times X) \geq b_1^{(2)}(G, EG \times Y)$. Since isotropy groups are finite or satisfy $b_p^{(2)}(G_x) = 0$ for $p \geq 0$, Theorem 5.4 implies that $b_1^{(2)}(G, X) \geq b_1^{(2)}(G, Y)$. $\square$

COROLLARY 5.6. *Let X and Y be contractible G-complexes, let $f : X \to Y$ be a G-map, and suppose that for all $x \in X$ the stabilizer G_x is finite or satisfies $b_p^{(2)}(G_x) = 0$ for all $p \geq 0$, and analogously for all $y \in Y$. Then $b_p^{(2)}(G, X) = b_p^{(2)}(G, Y)$ for $p \geq 0$. In particular, $\chi(G, X) = \chi(G, Y)$.*

Proof. Since $EG \times X$ and $EG \times Y$ are free G-complexes, the Künneth formula implies that the map $\mathbf{1} \times f : EG \times X \to EG \times Y$ is $\mathbb{C}$-homologically n-connected for all n. Theorem 5.3 implies $b_p^{(2)}(G, EG \times X) = b_p^{(2)}(G, EG \times Y)$ for $p \geq 0$. Since isotropy groups are finite or satisfy $b_p^{(2)}(G_x) = 0$ for $p \geq 0$, Theorem 5.4 implies that $b_p^{(2)}(G, X) = b_p^{(2)}(G, Y)$ for $p \geq 0$. $\square$

6. Equivariant Immersions

DEFINITION 6.1. Let X be a G-complex. Define

- $\mathrm{Bnd}(G, X)$ as the subset of $v \in I^0(G, X)$ such that $\mathrm{link}(x, X)$ is a single vertex or has spurs for a representative $x \in X$ of v, and
- $\mathrm{Isd}(G, X)$ as the subset of $v \in I^0(G, X)$ such that $\mathrm{link}(x, X)$ has at least one vertex of valence 0 for a representative $x \in X$ of v.

DEFINITION 6.2 (Essential Path). A path $P \to Y$ is *essential* if the lift to the universal cover $P \to \tilde{Y}$ is not closed.

THEOREM 6.3 (Collapsing Essential Paths). *Let X be a simply connected H-complex, let $Y \to X$ be an H-equivariant immersion, and let $P \to Y$ be an essential path. Suppose $P \to Y \to X$ is a closed path that is simple in the sense that it embeds except at its endpoints. Then $Y \to X$ factors as a composition $Y \to Z \to X$ of H-equivariant maps, where $Y \to Z$ is π_1-surjective, $Z \to X$ is an immersion, and the path $P \to Y \to Z$ is closed and null-homotopic. Moreover, we can choose Z such that the following hold*:

(i) *if Y is connected then Z is connected*;
(ii) *if Y is H-cocompact then Z is H-cocompact*;
(iii) $|\mathrm{Bnd}(H, Z) \cup \mathrm{Isd}(H, Z)| \leq |\mathrm{Bnd}(H, Y) \cup \mathrm{Isd}(H, Y)|$.

The strategy of the proof is as follows. A disk diagram $D \to X$ with boundary path $P \to X$ is equivariantly attached to Y to obtain an H-complex Z' and H-maps $Y \to Z' \to X$; this is performed using pushouts of equivariant immersions. Then the complex Z' is equivariantly folded to obtain H-maps $Z' \to Z \to X$ such that $Z \to X$ is an immersion. The proof of the theorem requires some preliminary results.

6.1. Equivariant Folding and Pushouts

LEMMA 6.4 (Equivariant Folding). *Let $W \to X$ be a G-map with W locally finite. Then $W \to X$ factors as $W \to Z \to X$, where $Z \to X$ is an immersion, $W \to Z$ is both a surjection and a π_1-surjection, and all maps are G-invariant.*

Proof. The statement is clear when W is compact and G is trivial. In general, let $W = \bigcup_i W_i$ be a filtration by compact sets. For each i, let $W_i \to Z_i \to X$ be a factorization such that $Z_i \to X$ is an immersion and $W_i \to Z_i$ is surjective and π_1-surjective. Observe that for $i < j$ there is a commutative diagram on the left below.

$$\begin{array}{ccc} W_i & \subseteq & W_j \\ \downarrow & & \downarrow \\ Z_i & \longrightarrow & Z_j \end{array} \qquad \begin{array}{ccccc} U_1 & \longrightarrow & U_2 & \longrightarrow & U_3 \\ \downarrow & & \downarrow & & \downarrow \;\cdots \\ V_1 & \longrightarrow & V_2 & \longrightarrow & V_3 \end{array}$$

For a sequence $U_1 \to U_2 \to U_3 \to \cdots$, we define U_∞ to be the direct limit. Specifically, U_∞ is the combinatorial complex whose p-cells are tales of p-cells $c_i \to c_{i+1} \to c_{i+2} \to \cdots$ for some $i \geq 1$, where two tales are equivalent if they are eventually the same. If $\{U_i\}$ is a filtration of U, then U_∞ equals U. A morphism $\{U_i\} \to \{V_i\}$ between two such sequences is a commutative diagram as on the right above. Any such morphism induces a map $U_\infty \to V_\infty$. Observe that if the vertical arrows of the morphism are surjective, then the map $U_\infty \to V_\infty$ is surjective.

By surjectivity, the morphism $\{W_i\} \to \{Z_i\}$ induces a surjective map $W_\infty \to Z_\infty$. As before, W_∞ equals W. Let $X_i = X$ and let $X_1 \to X_2 \to X_3 \to \cdots$ be the identity sequence, and note that X_∞ equals X. It follows that we have a map $W \to Z \to X$, where $Z = Z_\infty$.

Let us verify that $Z \to X$ is an immersion. Since each $Z_i \to X$ is an immersion and X is locally finite, Z is locally finite. In particular, for each 0-cell $z \in Z$ there is an index i for which $Z_i \to Z$ maps $z_i \mapsto z$ and $\text{link}(z_i)$ maps isomorphically onto $\text{link}(z)$; the factorization $Z_i \to Z \to X$ shows that $Z \to X$ is locally injective at z.

To see that Z is an H-complex and $W \to Z$ is an H-map, we observe that for $h \in H$ there is another filtration $W = \bigcup_i hW_i$. For each i there is an i_* such that $hW_i \subseteq W_{i_*}$, and this induces a map $h_i \colon Z_i \to Z_{i_*}$; it follows that $\{h_i\}$ induces a map $h\colon Z \to Z$ defining an action of H onto Z. By construction, the action commutes with the map $W \to Z$.

To see that $W \to Z$ is π_1-surjective, observe that each closed path σ in Z occurs in some Z_i. Since $W_i \to Z_i$ is π_1-surjective, $\sigma_i \to Z_i$ is homotopic to (the image of) a closed path $\sigma_i \to W_i \to W$. $\square$

LEMMA 6.5 (Pushouts of Equivariant Maps). *Let $\phi\colon C \to A$ and $\psi\colon C \to B$ be G-maps of complexes. Then there is a G-complex Z and there are G-maps $\imath\colon A \to Z$ and $\jmath\colon B \to Z$ such that $\imath \circ \phi = \jmath \circ \psi$. The pushout $(Z, \imath, \jmath)$ is universal in the sense that the equivariant maps $A \to X$ and $B \to X$ for which* *the diagram below commutes induce a unique equivariant map $Z \to X$ that also makes the diagram commute.*

$$\begin{tikzcd} & A \arrow[dr, "\iota"] \arrow[drr] & & \\ C \arrow[ur, "\phi"] \arrow[dr, "\psi"'] & & Z \arrow[r, dashed] & X \\ & B \arrow[ur, "\jmath"'] \arrow[urr] & & \end{tikzcd}$$

Moreover, if A and B are G-cocompact (*resp., proper*) *then Z is G-cocompact* (*resp., proper*). *If A is connected and $\psi(C)$ intersects all connected components of B, then Z is connected.*

Proof. The construction of the pushout of ϕ and ψ is standard and is briefly described. Let Z be the combinatorial complex obtained by taking the quotient of the disjoint union of A and B by the relation $\phi(\sigma) = \psi(\sigma)$ for σ in C. The resulting complex admits a natural G-action. The statements on inheritance of cocompactness, properness, and connectedness are routine; the details are left to the reader. □

6.2. Proof of Theorem 6.3

Since X is simply connected, there is a near-immersion $D \to X$ of a disk diagram with boundary path $P \to X$. Since $P \to X$ is a simple closed path, D is homeomorphic to an Euclidean disk.

Let $\bigcup_H P$ denote the disjoint union of copies of P, one for each element of H. Define $\bigcup_H D$ analogously. By Lemma 6.5, let Z' be the pushout of the natural H-maps $\bigcup_H P \to \bigcup_H D$ and $\bigcup_H P \to Y$. By the universal property of the pushout, the immersion $Y \to X$ and the natural map $\bigcup_H D \to X$ induce an H-map $Z' \to X$. By Lemma 6.4, let $Z \to X$ be the H-equivariant immersion obtained after an equivariant folding of $Z' \to X$. We refer to the following commutative diagram.

$$\begin{tikzcd} & Y \arrow[dr] \arrow[drrr] & & & \\ \bigcup_H P \arrow[ur] \arrow[dr] & & Z' \arrow[r] & Z \arrow[r] & X \\ & \bigcup_H D \arrow[ur] \arrow[urrr] & & & \end{tikzcd}$$

The main conclusions of Theorem 6.3 are proved in Lemmas 6.6–6.9.

LEMMA 6.6. *$Y \to Z' \to Z$ is π_1-surjective.*

Proof. By Lemma 6.4, $Z' \to Z$ is π_1-surjective. Hence it remains to prove that $Y \to Z'$ is π_1-surjective. Let Y' denote the image of Y in Z', and note that $Y' \to Z'$ is π_1-surjective since the complement of Y' in Z' is a collection of disjoint open disks. Observe that $\pi_1 Y'$ is generated by the image of $\pi_1 Y \to \pi_1 Y'$ together with closed paths corresponding to H-translates of $P \to Y \to Z'$. Since these additional paths become null-homotopic by the addition of the H-translates of $D \to Z'$, the result follows. $\square$

LEMMA 6.7. *If Y is H-cocompact then Z is H-cocompact. If Y is connected then Z is connected.*

Proof. Suppose Y is H-cocompact. By Lemma 6.5, the pushout Z' is H-cocompact. Lemma 6.4 implies that $Z' \to Z$ is surjective. It follows that Z is H-cocompact. Suppose that Y is connected. Then Z' is connected by Lemma 6.5. Since $Z' \to Z$ is surjective, Z is connected. $\square$

For the next two lemmas, consider the natural map $I^0(H, Y) \to I^0(H, Z)$ induced by $Y \to Z$.

LEMMA 6.8. *The image of* $\operatorname{Isd}(H, Y)$ *contains* $\operatorname{Isd}(H, Z)$.

Proof. Suppose that $v \in \operatorname{Isd}(H, Z)$ and let $z \in Z$ be a representative. Let e be the 1-cell of Z giving rise to the isolated vertex of $\operatorname{link}(z, Z)$. Since $Z' \to Z$ is surjective, there is a 1-cell e' of Z' mapping to e and a corresponding 0-cell z' mapping to z. Since $\operatorname{link}(z, Z)$ has an isolated vertex induced by e, $\operatorname{link}(z', Z')$ has an isolated vertex s induced by e'. Suppose that e' has a preimage f in $\bigcup_H D$. Then f is on the boundary path of a component of $\bigcup_H D$. It follows that e' has a preimage in $\bigcup_H P$ and so e' has a preimage in Y. Hence there is a $y \in Y$ such that the image of $\operatorname{link}(y, Y)$ in $\operatorname{link}(z', Z')$ contains the isolated vertex s. It follows that $\operatorname{link}(y, Y)$ has an isolated vertex and thus there is a $u \in \operatorname{Isd}(H, Y)$ mapping to v. $\square$

LEMMA 6.9. *The image of* $\operatorname{Bnd}(H, Y) \cup \operatorname{Isd}(H, Y)$ *contains* $\operatorname{Bnd}(H, Z)$.

Proof. Suppose that $v \in \operatorname{Bnd}(H, Z)$. Let $z \in Z$ be a representative of v. If $\operatorname{link}(z, Z)$ is a single point, then $v \in \operatorname{Isd}(H, Y)$ and Lemma 6.8 shows that there is a $u \in \operatorname{Isd}(H, Y)$ that maps to v. Consider the case where $\operatorname{link}(z, Z)$ has a spur s with terminal vertex t.

Suppose that s corresponds to a 2-cell in the image of $Y \to Z$. Then there is a $y \in Y$ mapping to z such that the image of $\operatorname{link}(y, Y) \to \operatorname{link}(z, Z)$ contains s. Since $Y \to Z$ is an immersion, $\operatorname{link}(y, Y)$ is a subgraph of $\operatorname{link}(z, Z)$ and hence it has a spur. In particular, there is a $u \in \operatorname{Bnd}(H, Y)$ that maps to v.

Otherwise, there is a 0-cell w of $\bigcup_H D$ that maps to z, and the image of $\operatorname{link}\big(w, \bigcup_H D\big) \to \operatorname{link}(z, Z)$ contains s. In particular, $\operatorname{link}\big(w, \bigcup_H D\big)$ has a spur and therefore w is in the boundary of a connected component of $\bigcup_H D$. Hence there is also a $y \in Y$ that maps to z, and the image of $\operatorname{link}(y, Y) \to \operatorname{link}(z, Z)$ contains t. As $\bigcup_H D \to Z$ and $Y \to Z$ are immersions, $\operatorname{link}(y, Y)$

and $\mathrm{link}\big(w, \bigcup_H D\big)$ are subgraphs of $\mathrm{link}(z, Z)$. Both subgraphs contain the vertex t, but $\mathrm{link}(y, Y)$ does not contain s. It follows that t is an isolated vertex of $\mathrm{link}(y, Y)$. Therefore, there is a $u \in \mathrm{Isd}(H, Y)$ that maps to v. □

Lemmas 6.8 and 6.9 imply that the image of $\mathrm{Bnd}(H, Y) \cup \mathrm{Isd}(H, Y)$ contains $\mathrm{Bnd}(H, Z) \cup \mathrm{Isd}(H, Z)$, and this concludes the proof of Theorem 6.3.

6.3. No Self-Immersions

LEMMA 6.10 (No Self-Immersions). *Let X be a G-cocompact, proper, and connected complex. Any G-equivariant immersion $\phi\colon X \to X$ is an isomorphism.*

Proof. The quotient space X/G is a combinatorial complex, and ϕ induces a self-immersion $\psi\colon X/G \to X/G$. Since X/G is compact and connected, ψ is an isomorphism [15, Lemma 6.3] and hence ϕ is onto.

Let u be a 0-cell. Since ψ is an isomorphism, all elements of $\phi^{-1}(u)$ are G-equivalent. Therefore $|\phi^{-1}(u)|$ is a lower bound for the size of the G-stabilizer of u. Since X is G-cocompact and proper, there is an upper bound on the cardinality of cell stabilizers. Therefore $|\phi^{-1}(\phi^n(u))| = 1$ for some $n > 0$ depending on u. By cocompactness, there is an $m > 0$ such that ϕ restricted to $\phi^m(X) = X$ is injective. □

7. Existence of Cores

This section contains the proof of Theorem 1.3. The section is divided into five subsections. The first four subsections contain preliminary results, and the last subsection discusses the proof of the theorem.

7.1. Angled Graphs

LEMMA 7.1 (Curvature and Connected Components). *Let Δ be a cocompact angled H-graph. Let $\Delta_1, \dots, \Delta_\ell$ be a collection of representatives of H-orbits of connected components of Δ, and let H_i be the stabilizer of Δ_i. If* $\mathrm{Curvature}(H_i, \Delta_i) \le \pi \cdot |H_i|^{-1}$ *for* $1 \le i \le \ell$, *then* $\mathrm{Curvature}(H, \Delta) \le \mathrm{Curvature}(H_i, \Delta_i)$ *for* $1 \le i \le \ell$.

Proof. First notice that

$$\mathrm{Curvature}(H, \Delta) = \sum_{i=1}^{\ell} \mathrm{Curvature}(H_i, \Delta_i) + 2\pi \cdot \left(|H|^{-1} - \sum_{i=1}^{\ell} |H_i|^{-1} \right). \tag{6}$$

If $\ell = 1$ and $H = H_1$, then $\Delta = \Delta_1$ and obviously $\kappa(H, \Delta) = \kappa(H_1, \Delta_1)$. Otherwise, $2|H|^{-1} \le \sum_{i=1}^{\ell} |H_i|^{-1}$ and hence equation (6) implies that

$$\mathrm{Curvature}(H, \Delta) \le \sum_{i=1}^{\ell} (\mathrm{Curvature}(H_i, \Delta_i) - \pi \cdot |H_i|^{-1}).$$

Since $\mathrm{Curvature}(H_i, \Delta_i) \le \pi \cdot |H_i|^{-1}$, it follows that $\mathrm{Curvature}(H, \Delta) \le \mathrm{Curvature}(H_i, \Delta_i)$. □

PROPOSITION 7.2. *Let Γ be an angled G-graph such that, for each subgroup $K \leq G$, every regular K-section has curvature $\leq \alpha \leq 0$. Suppose H is a subgroup of G and Δ is a nonempty and spurless H-section of Γ. The following statements hold.*

(i) *If Δ contains an edge, then* $\mathrm{Curvature}(H, \Delta) \leq \alpha$.
(ii) $\mathrm{Curvature}(H, \Delta) > 0$ *if and only if Δ is a single vertex and H is a finite group.*

Proof. Let Δ_i, H_i, and ℓ be as in the statement of Lemma 7.1. Since Δ is spurless, for each i either Δ_i is a regular H_i-section of Δ or else Δ_i is a single vertex. It follows that $\mathrm{Curvature}(H_i, \Delta_i) \leq 0 \leq \pi|H_i|^{-1}$ for each i.

Suppose Δ contains an edge. Without loss of generality, assume that Δ_1 is a connected component with at least one edge. Then Δ_1 is a regular H_1-section and, by Lemma 7.1, $\mathrm{Curvature}(H, \Delta) \leq \mathrm{Curvature}(H_1, \Delta_1) \leq \alpha$.

Suppose that $\mathrm{Curvature}(H, \Delta) > 0$. Since $\alpha \leq 0$, Lemma 7.1 implies that each connected component of Δ is a single point. Therefore $0 < \mathrm{Curvature}(H, \Delta) = 2\pi|H|^{-1} - \pi \sum_{i=1}^{\ell} |H_i|^{-1}$. This implies that $\ell = 1$, H is a finite group, and $H_1 = H$. In particular, Δ is a single point and H is a finite group. The "if" part of statement (ii) is immediate. □

COROLLARY 7.3. *Let Γ be an angled G-graph such that each regular section has curvature $\leq \alpha < 0$. Suppose that Δ is an H-section of Γ such that Δ is nonempty and spurless. Then* $\mathrm{Curvature}(H, \Delta) = 0$ *if and only if Δ is an edgeless graph and either*

(i) *Δ consists of two vertices and H acts trivially on Δ or*
(ii) *the stabilizer of each vertex of Δ is infinite.*

Proof. Let Δ_i, H_i, and ℓ be as in the statement of Lemma 7.1. Suppose that $\mathrm{Curvature}(H, \Delta) = 0$. Since $\alpha < 0$, Proposition 7.2 implies that each Δ_i is a single vertex. Then

$$0 = \mathrm{Curvature}(H, \Delta) = 2\pi \cdot |H|^{-1} - \pi \cdot \sum_{i=1}^{\ell} |H_i|^{-1}.$$

If H is an infinite group then each H_i is infinite, and therefore the stabilizer of each point of Δ is an infinite subgroup of H. If H is a finite group, then $2 = \sum_{i=1}^{\ell} [H : H_i]$, where $[H : H_i]$ is the index of H_i in H. Hence $\ell = 2$ and $H_i = H$ for $i = 1, 2$; in particular, the action of H on Δ is trivial.

The "if" part of the statement follows by a direct computation of $\mathrm{Curvature}(H, \Delta)$ and is left to the reader. □

7.2. Counting Immersions in Nonpositively Curved Complexes

DEFINITION 7.4 (G-Equivalent Maps). Let X be a G-complex, let H be a subgroup of G, and for $i = 1, 2$ let Y_i be an H-complex. A pair of H-equivariant immersions $\phi_1 \colon Y_1 \to X$ and $\phi_2 \colon Y_2 \to X$ are *G-equivalent* if there is an H-isomorphism of complexes $\psi \colon Y_1 \to Y_2$ and a $g \in G$ such that $g \circ \phi_1 = \phi_2 \circ \psi$.

DEFINITION 7.5. Let X be a nontrivial cocompact and proper angled G-complex. Define

- $\operatorname{Zero}(G, X)$ as the set of $v \in I^0(G, X)$ with $\kappa(v) = 0$,
- $\operatorname{Neg}(G, X)$ as the set of $v \in I^0(G, X)$ with $\kappa(v) < 0$, and
- $\operatorname{Pos}(G, X)$ as the set of $v \in I^0(G, X)$ with $\kappa(v) > 0$.

THEOREM 7.6 (Counting Immersions in Nonpositively Curved G-Complexes). *Let X be a nontrivial, cocompact, and proper angled G-complex with sectional curvature* ≤ 0. *Let H be a subgroup of G, and let r, s, t be fixed numbers. Up to G-equivalence, there are finitely many H-equivariant immersions $Y \to X$ with the following properties*:

(i) *Y is H-cocompact and connected*;
(ii) $\chi(H, Y) \geq r$;
(iii) $|\operatorname{Zero}(H, Y)| \leq s$;
(iv) $|\operatorname{Bnd}(H, Y)| \leq t$.

LEMMA 7.7 (Immersions Determined by a Compact Complex). *Let X be a proper G-complex, let H be a subgroup of G, let K be a finite connected complex, and let $\psi\colon K \to X$ be an immersion. There are finitely many G-equivalence classes of H-equivariant immersions $\phi\colon Y \to X$ such that there is an embedding $\iota\colon K \hookrightarrow Y$ satisfying $\phi \circ \iota = \psi$ and the H-translates of K cover Y.*

Proof. Observe that if $\phi\colon Y \to X$ is an H-equivariant immersion and K is a subcomplex with $\bigcup_{h \in H} K = Y$, then ϕ is completely determined by its restriction to K.

The H-proper complex Y is completely determined by the finite set of elements $\{g \in H : K \cap gK \neq \emptyset\}$ of H and the isomorphisms between the complexes $J_g = g^{-1}K \cap K$ and $J_g' = K \cap gK$. Indeed, one can recover Y by taking $H \times K$ and identifying the various $h \times J_g$ with $hg \times J_g'$ using the isomorphism.

To show that there are finitely many possibilities for the above data, we argue as follows. Let $g_1, \ldots, g_n$ be the set of elements of H such that $\psi K \cap g\psi K \neq \emptyset$, and note that this set is finite since G acts properly on X. For each i, there are finitely many choices of isomorphisms between subcomplexes $J_i \subset K$ and $J_i' \subset K$. □

LEMMA 7.8 (Counting Immersions). *Let X be a proper, cocompact, and connected G-complex, let H be a subgroup of G, and let M be a positive integer. Up to G-equivalence, there are finitely many H-equivariant immersions $Y \to X$ such that Y is connected and $|I^0(H, Y)| < M$.*

Proof. Observe that every H-cocompact complex Y with $|I^0(H, Y)| < M$ contains a connected subcomplex K with $|K^0| < M$. Since X is proper and G-cocompact, there are finitely many G-equivalent classes of immersions $K \to X$ where K is connected and $|K^0| < M$. The result then follows from Lemma 7.7. □

LEMMA 7.9. *Let X be a nontrivial, cocompact, and proper angled G-complex with generalized sectional curvature ≤ 0. Then*

$$I^0(G, X) = \mathrm{Zero}(G, X) \cup \mathrm{Neg}(G, X) \cup \mathrm{Bnd}(G, X).$$

Proof. Let $v \in I^0(G, X)$, and suppose that $v \notin \mathrm{Bnd}(G, X)$. Since X is nontrivial and connected, $\mathrm{link}(x, X)$ is nonempty. Since $v \notin \mathrm{Bnd}(G, X)$, $\mathrm{link}(x)$ is spurless and not a single point. Since X has generalized sectional curvature ≤ 0, Proposition 7.2 implies that $\mathrm{Curvature}(G_x, \mathrm{link}(x, X)) \leq 0$. Hence $v \in \mathrm{Zero}(G, X) \cup \mathrm{Neg}(G, X)$. □

DEFINITION 7.10 ($\mathrm{N}(G, X), \mathrm{P}(G, X)$). Let X be a cocompact and proper angled G-complex. There are finitely many 0-cells of X up to the action of G. Since X is proper and cocompact it follows that, for each 0-cell x of X, the angled G_x-graph $\mathrm{link}(x)$ has finitely many sections. Let $\mathrm{N}(G, X)$ be the maximum curvature among all possible such sections with negative curvature; if there are no sections with negative curvature let $\mathrm{N}(G, X) = -1$. Analogously, let $\mathrm{P}(G, X)$ be the maximum curvature among all sections with nonnegative curvature; if there are no sections with nonnegative curvature let $\mathrm{P}(G, X) = 0$.

REMARK 7.11. For a cocompact and proper angled G-complex X, $\mathrm{N}(G, X) < 0 \leq \mathrm{P}(G, X)$.

LEMMA 7.12. *Let X be a cocompact and proper angled G-complex. Suppose that $\kappa(f) \leq 0$ for each 2-cell of X. If H is a subgroup of G, Y is a cocompact H-complex, and $Y \to X$ is an H-equivariant immersion, then*

$$|\mathrm{Neg}(H, Y)| \leq \frac{2\pi \cdot \chi(H, Y) - \mathrm{P}(G, X) \cdot |\mathrm{Pos}(H, Y)|}{\mathrm{N}(G, X)}.$$

Proof. For $v \in I^0(H, Y)$ we have $\kappa(v) = \mathrm{Curvature}(H_y, \mathrm{link}(y, Y))$, where y is a representative of v. Let $x \in X$ be the image of y. Since $Y \to X$ is an immersion, $\mathrm{link}(y, Y)$ is an H_y-section of the G_x-graph $\mathrm{link}(y, X)$. In particular, $\kappa(v) \leq \mathrm{N}(G, X) < 0$ or $0 \leq \kappa(v) \leq \mathrm{P}(G, X)$.

By the Gauss–Bonnet Theorem 4.2 and the assumption that 2-cells have nonpositive curvature,

$$2\pi \cdot \chi(H, Y) \leq \mathrm{N}(G, X) \cdot |\mathrm{Neg}(H, Y)| + \mathrm{P}(G, X) \cdot |\mathrm{Pos}(H, Y)|.$$

Since $\mathrm{N}(G, X) < 0$, the conclusion of the lemma is immediate. □

Proof of Theorem 7.6. Let $Y \to X$ be an H-equivariant immersion satisfying the listed properties. By Proposition 3.5, X has generalized sectional curvature ≤ 0. Therefore Y has generalized sectional curvature ≤ 0 as well. By Lemma 7.9, $I^0(H, Y) = \mathrm{Zero}(H, Y) \cup \mathrm{Neg}(H, Y) \cup \mathrm{Bnd}(H, Y)$. By Lemmas 7.9 and 7.12,

$$\begin{aligned} |I^0(H, Y)| &\leq |\mathrm{Zero}(H, Y)| + |\mathrm{Bnd}(H, Y)| + |\mathrm{Neg}(H, Y)| \\ &\leq s + t + \frac{2\pi \cdot r - \mathrm{P}(G, X) \cdot t}{\mathrm{N}(G, X)}. \end{aligned}$$

Since there is an upper bound for the size of $I^0(H,Y)$ that is independent of H and Y, Lemma 7.8 implies that there are finitely many posibilities for $Y \to X$. □

7.3. Proof of the Simply Connected Core Theorem 1.3

Construct a sequence of H-equivariant immersions $\phi_n\colon Y_{n-1} \to Y_n$ for $n \geq 1$ and $\psi_n\colon Y_n \to X$ for $n \geq 0$ in the following way.

$$\begin{tikzcd} Y_0 \arrow[r, "\phi_1"] \arrow[d, "\psi_0"'] & Y_1 \arrow[r, "\phi_2"] \arrow[dl, "\psi_1"] & \cdots \arrow[r, "\phi_n"] & Y_n \arrow[r, "\phi_{n+1}"] \arrow[dlll, "\psi_n"] & \cdots \\ X & & & & \end{tikzcd}$$

Let $Y_0 = Y$ and let $\psi_0\colon Y_0 \to X$ be the inclusion map. Assume that $\psi_n\colon Y_n \to X$ has been defined. Suppose there is a nontrivial path $P \to Y_n$ such that either

- $P \to Y_n$ is a closed path that is not null-homotopic or
- $P \to Y_n$ is not a closed path but $P \to Y_n \to X$ is a closed path.

Such a path is called *essential*. Choose a path $P \to Y_n$ of minimal length with the above property. Observe that $P \to Y_n \to X$ is a simple closed path by minimality. Let $\phi_{n+1}\colon Y_n \to Y_{n+1}$ be the π_1-surjective map and let $\psi_{n+1}\colon Y_{n+1} \to X$ be the H-equivariant immersion provided by Theorem 6.3 applied to the immersion $\psi_n\colon Y_n \to X$ and the path $P \to Y_n$. If there is no path $P \to Y_n$ as above, then the sequence stabilizes in the sense that $Y_m = Y_n$, $\psi_m = \psi_n$, and ϕ_m is the identity map for all $m \geq n+1$. (We will show that the sequence always stabilizes.)

LEMMA 7.13. *Y_n is connected and H-cocompact.*

Proof. This follows by induction using Theorem 6.3 with base case that Y_0 is connected and H-cocompact. □

LEMMA 7.14. *ϕ_n is an immersion and is π_1-surjective.*

Proof. The property of π_1-surjectivity follows from the definition of ϕ_n in terms of Theorem 6.3. Since $\psi_{n+1} \circ \phi_{n+1} = \psi_n$ and ψ_n is an immersion, ϕ_{n+1} is an immersion. □

LEMMA 7.15. $|\mathrm{Zero}(H,Y_n) \cup \mathrm{Bnd}(H,Y_n)| \leq |\mathrm{Isd}(H,Y_0) \cup \mathrm{Bnd}(H,Y_0)|$.

Proof. By Theorem 6.3(iii), an induction argument shows that

$$|\mathrm{Isd}(H,Y_n) \cup \mathrm{Bnd}(H,Y_n)| \leq |\mathrm{Isd}(H,Y_0) \cup \mathrm{Bnd}(H,Y_0)|.$$

Therefore it is enough to prove that $\mathrm{Zero}(H,Y_n) \cup \mathrm{Bnd}(H,Y_n)$ is a subset of $\mathrm{Isd}(H,Y_n) \cup \mathrm{Bnd}(H,Y_n)$. Let $v \in \mathrm{Zero}(H,Y_n) - \mathrm{Bnd}(H,Y_n)$ and let y be a representative of v. Since $Y_n \to X$ is an immersion, Y_n has negative sectional curvature. By Corollary 7.3, $\mathrm{link}(y, Y_n)$ consists of two vertices and no edge. Thus $v \in \mathrm{Isd}(H,Y_n)$. □

LEMMA 7.16. $\chi(H, Y_n) \geq -b_1^{(2)}(H, Y_0)$.

Proof. By Lemma 7.14 and an induction argument, $Y_0 \to Y_{n+1}$ is π_1-surjective. Then Corollary 5.5 implies that $b_1^{(2)}(H, Y_0) \geq b_1^{(2)}(H, Y_n)$. Then the conclusion follows from $\chi(H, Y_n) \geq -b_1^{(2)}(H, Y_n)$. □

For positive integers $m < n$, let $\phi_{m,n}$ denote the immersion $\phi_{n-1} \circ \cdots \circ \phi_{m+1}$ from Y_m to Y_n.

LEMMA 7.17. *There is an $m_0 > 0$ such that $\phi_{m_0,n}$ is an H-equivariant isomorphism for every $n > m_0$.*

Proof. By Theorem 7.6 and the previous lemmas, the sequence contains only finitely many non–G-equivalent immersions. If the statement is false, then there are positive integers $m < n$ such that $Y_m \to X$ and $Y_n \to X$ are G-equivalent but the H-equivariant immersion $\phi_{m,n}\colon Y_m \to Y_n$ is not an isomorphism. Since Y_m and Y_n are isomorphic as H-complexes, $\phi_{m,n}$ would be a self-immersion that is not an isomorphism—contradicting Lemma 6.10. □

Conclusion of the Proof of Theorem 1.3. By Lemma 7.17, it follows that there exists an $m_0 > 0$ such that Y_{m_0} has no essential paths as defined in the construction of the Y_i. In particular, ψ_{m_0} is an embedding and Y_{m_0} is simply connected. Therefore $\psi_{m_0}(Y_{m_0})$ is a simply connected H-cocompact subcomplex of X containing Y_0.

8. Quasiconvex Cores

This section contains the proof of Theorem 1.6. Let X be a proper cocompact CAT(0) G-complex whose cells are convex. This is a complete geodesic metric space [1]; for background on CAT(0) cell-complexes we refer the reader to [2].

The proof is split into several subsections. Assign angles as they arise from the CAT(0)-metric and suppose X has negative sectional curvature. Let H be a subgroup of G and let Y be a simply connected cocompact H-subcomplex. Let ℓ be a geodesic segment in X such that its endpoints are 0-cells of Y.

8.1. The Carrier L of ℓ, the Paths P_L, P_Y, and the Disk Diagram D

Let $R_1, \dots, R_n$ be the sequence of open cells of X that intersect ℓ in the order in which they are traversed by ℓ. Since cells of X are convex, there are no repetitions in the sequence and $\partial R_i \cap \partial R_j$ is connected. Let L be the complex constructed by taking the disjoint union of closures $\bar{R}_1 \sqcup \cdots \sqcup \bar{R}_n$ and identifying the two copies of $\partial R_i \cap \partial R_{i+1}$ in $\bar{R}_i$ and $\bar{R}_{i+1}$ for $1 \leq i < n$. Observe that $L \to X$ is a near-immersion, L is simply connected, and $\ell \hookrightarrow X$ factors as $\ell \to L \to X$. See Figure 1.

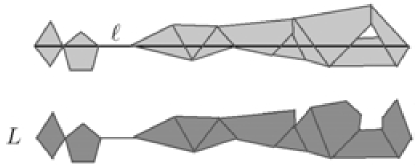

Figure 1 Sequence of open cells of X intersecting the geodesic ℓ (top); the resulting complex L (bottom)

Since Y and L are connected, there are edge paths $P_Y \to Y$ and $P_L \to L$ connecting the endpoints s and t of ℓ. Since X is simply connected, there is a disk diagram $D \to X$ between $P_Y \to X$ and $P_L \to X$. Choose $P_Y \to Y$ and $P_L \to L$ and $D \to X$ such that $(\text{area}(D), |\partial D|)$ is minimal in the lexicographical order.

LEMMA 8.1 (ℓ Is Uniformly Close to P_L). *If $R \to X$ is an open cell intersecting ℓ, then its closure $\bar{R} \to X$ intersects the image of $P_L \to X$.*

Proof. Since $P_L \to L$ connects the endpoints of $\ell \to L$, if a closed cell S disconnects L then $P_L \to L$ intersects S. By definition of L, if R is an open cell of X intersecting ℓ then the closure of R in L disconnects L. □

DEFINITION 8.2 (Cut 0-Cells and Cut Components). A 0-cell v is called a *cut 0-cell of* D provided that $D - \{v\}$ is not connected. Let V be the set of all cut 0-cells of D. Closures of connected components of $D - V$ are *cut components.* A cut component is *nonsingular* if it contains a 2-cell.

LEMMA 8.3. *The path $P_Y \to D$ is embedded.*

Proof. Suppose that $P_Y \to D$ is not embedded. Then P_Y is a concatenation $U_1 V U_2$ such that (a) the terminal point of $U_1 \to P_L \to D$ and also the initial point of $U_2 \to P_L \to D$ is a 0-cell u and (b) $V \to P_Y \to D$ is a nontrivial path with no internal 0-cells in common with $P_L \to D$. Observe that u is a cut 0-cell of D. Let $\dot{P}_Y \to D$ be the path $U_1 U_2 \to D$, and observe that $\dot{P}_Y \to D \to X$ factors through $Y \to X$ and $\dot{P}_Y \to X$ connects the endpoints of ℓ. The paths $\dot{P}_Y \to D$ and $P_L \to D$ bound a subdiagram $\dot{D}$ of D. Since $V \to D$ is nontrivial, the complexity of $\dot{D}$ is strictly smaller than the complexity of D. Then $\dot{P}_Y$ and $\dot{D}$ violate the minimality of D. □

LEMMA 8.4 (2-Cells of D Intersecting P_L). *If $R \to D$ is a 2-cell such that $\partial R \to D$ intersects $P_L \to D$, then the interior of $R \to X$ does not intersect ℓ.*

Proof. If the interior of $R \to X$ intersects ℓ, then $\partial R \to X$ factors through $L \to X$. Since $\partial R \to D$ intersects $P_L \to D$, it follows that P_L is a concatenation $U_1 c U_2$ where c is a cell mapped into $\partial R \to D$. In particular, ∂R is a concatenation cQ where $Q \to \partial R$ is a path; if c is a 0-cell then c is the initial point of the path Q. Let $\dot{D}$ be the subdiagram of D obtained by removing the interiors of c

and R; in the case that c is a 0-cell, remove the interior of R, remove c, and add two copies of c to obtain a simply connected diagram. Let $\dot{P}_L \to L$ be the path $U_1Q^{-1}U_2$ (or U_1QU_2) and observe that $\dot{P}_L$ and $\dot{D}$ violate the minimality of D. See Figure 2. □

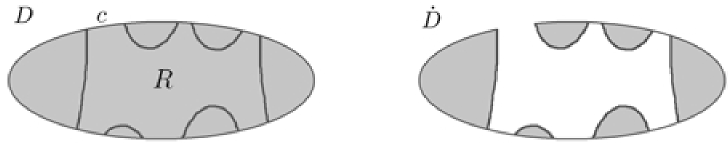

Figure 2 Reduction of complexity in the proof of Lemma 8.4: if a cut component of D has a 2-cell $R \to D$ with two different boundary arcs intersecting ℓ, then D has no minimal complexity

LEMMA 8.5 (Internal Paths in D). *If there is a nontrivial internal path $T \to D$, as described in Definition 2.4, such that $T \to D \to X$ factors through $L \to X$ and has endpoints in $P_L \to D$, then D does not have minimal complexity.*

Proof. Suppose that $T \to D$ is an internal path satisfying the hypothesis of the lemma. Without loss of generality, we can assume that $T \to D$ is an embedding; indeed, its image contains an internal and embedded path $T' \to D$ with the same endpoints as $T \to D$. The endpoints of T split the boundary path of D as a concatenation of paths $UV \to \partial D$ such that T and U have the same endpoints and $P_Y \to D$ is a subpath of U^{-1} (the path V may be trivial). Then $P_L \to D$ is a concatenation U_1VU_2, where each U_i is a subpath of U. Let $\dot{P}_L \to D$ be the path $U_1T^{-1}U_2$ and observe that $\dot{P}_L \to D \to X$ factors through $L \to X$. The paths $P_Y \to D$ and $\dot{P}_L \to D$ bound a subdiagram $\dot{D}$ of D. Since $T \to D$ is internal, the complexity of $\dot{D}$ is strictly smaller than the complexity of D. Then $\dot{P}_L$ and $\dot{D}$ violate the minimality of D. See Figure 3. □

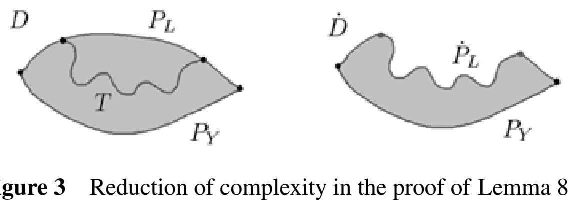

Figure 3 Reduction of complexity in the proof of Lemma 8.5

LEMMA 8.6 (2-Cells of D Intersecting ℓ). *Let $R \to D$ be a 2-cell. Suppose $\partial R \to D$ is a concatenation $S_1T_1 \dots S_mT_m$, where each $S_i \to D$ is a boundary path $S_i \to \partial D$ and each $T_i \to D$ is a nontrivial internal path in D. Then at most one subpath $S_i \to X$ intersects ℓ and factors through $P_L \to X$.*

Proof. Suppose that two different paths $S_i \to X$ intersect ℓ and factor through $P_L \to X$. By convexity of R, either the interior of $R \to X$ intersects ℓ or else there is a path $T_i \to D$ that maps into ℓ. The former case is impossible by Lemma 8.4 and the latter case is impossible by Lemma 8.5. □

LEMMA 8.7 (Cut 0-Cells of D). *Let u be a cut 0-cell of D.*

(i) *If u is in $P_Y \to D$, then u is also in $P_L \to D$.*
(ii) *If $u \to D \to X$ is contained in ℓ, then u is in $P_Y \to D$.*

Proof. Since $P_Y \to D$ is embedded by Lemma 8.3, the first statement is immediate. To prove the second statement, suppose that the image of u in X is contained in ℓ and that u is not in $P_Y \to D$. Hence u is in $P_L \to D$. Observe that all points of P_L that map to $u \in D$ map to the same point in L. This follows since $\ell \to X$ is injective, contains the cell $u \to X$, and factors through $L \to X$. Let $V \to D$ be the subpath of $P_L \to D$ starting and ending at u and traveling around the components of $D - U$ not containing $P_Y \to D$. Express P_L as a concatenation $U_1 V U_2$ and note that $V \to D$ and $P_Y \to D$ are disjoint. Let $\dot{P}_L$ be the path $U_1 U_2$. Observe that $\dot{P}_L \to X$ is a path that connects the endpoints of ℓ, factors through $L \to X$, and factors through $P_L \to D$. Moreover, the paths $\dot{P}_L \to D$ and $P_Y \to D$ bound a subdiagram $\dot{D}$ of D. Since $V \to D$ is nontrivial, the complexity of $\dot{D}$ is strictly smaller than the complexity of D. Then $\dot{P}_L$ and $\dot{D}$ violate the minimality of D. □

8.2. *Good Path in D*

The main goal of this subsection is to prove Proposition 8.9.

DEFINITION 8.8 (Good Paths in D). A path $Q \to D$ is a *good path* if, for every 0-cell c in the interior of the path $Q \to D$, either c is an internal cell of the diagram D or the image of $c \to D \to X$ does not intersect ℓ.

PROPOSITION 8.9 (Good Paths with Terminal Point in P_Y). *Let u be a 0-cell of ∂D. Suppose u is not an internal 0-cell of a boundary arc of D. Then there is a good path $Q \to D$ from u to a 0-cell of $P_Y \to D$.*

The following lemma is immediate.

LEMMA 8.10 (Concatenation of Good Paths). *Let $P_1 \to D$ and $P_2 \to D$ be good paths such that the terminal point w of P_1 equals the initial point of P_2. Suppose the image of w in $D \to X$ does not intersect ℓ. Then the concatenation $P_1 P_2 \to D$ is a good path.*

LEMMA 8.11 (Good Paths along a Boundary of a 2-Cell). *Let $R \to D$ be a 2-cell such that ∂R is a concatenation $S_1 T_1 \dots S_m T_m$, where each $S_i \to D$ is a boundary arc $S_i \to \partial D$ and each $T_i \to D$ is a nontrivial internal path in D.*

(i) *Suppose each $S_i \to D$ factors through $P_L \to D$. If u and v are distinct 0-cells in the intersection of $\partial R \to D$ and ∂D and if u and v are not internal 0-cells of a boundary arc of D, then there is a good path between them.*

(ii) *Suppose $\partial R \to D$ intersects $P_L \to D$ but some $S_i \to D$ does not factor through $P_L \to D$. If u is a 0-cell in the intersection of $\partial R \to D$ and ∂D and if u is not an internal cell of a boundary arc of D, then there is a good path from u to a 0-cell v in the intersection of $\partial R \to D$ and $P_Y \to D$.*

Proof. Suppose that all $S_i \to D$ factor through $P_L \to D$. By Lemma 8.6, we can assume without loss of generality that $S_i \to D \to X$ does not intersect ℓ for $i \geq 2$. Then the path $P = T_1 S_2 \dots S_n T_n \to D$ is a good path since if a 0-cell is mapped into ℓ then it must be in the interior of D. Since u and v are not internal cells of a boundary arc of ∂D, u and v are in $P \to D$. The first statement follows.

Suppose $\partial R \to D$ intersects $P_L \to D$ but some $S_i \to D$ does not factor through $P_L \to D$. Since u is not an internal cell of a boundary arc, it follows that u is an endpoint of $S_j \to D$ for some j. If $S_i \to D$ does not factor through $P_L \to D$, then it factors through $P_Y \to D$. Therefore $\partial R \to D$ has 0-cells of $P_Y \to D$, and in particular there is a 0-cell v that is an endpoint of $S_i \to D$ for some i and is in $P_Y \to D$. Lemma 8.6 implies that at most one of the paths $S_i \to D \to X$ intersects ℓ and factors through $P_L \to D$. Without loss of generality, assume that if there is such a path then it is $S_1 \to D$. It follows that $T_1 S_2 \dots S_n T_n$ is a good path containing u and v. □

Proof of Proposition 8.9. There is a sequence of cells $c_1, c_2, \dots, c_n$ in D such that:

(i) each c_i is either an open 2-cell or an open 1-cell disconnecting D;
(ii) $\bar{c}_i \cap \bar{c}_{i+1}$ either is a cut 0-cell or contains a 1-cell;
(iii) $u \in \bar{c}_1$, and $\bar{c}_n$ intersects $P_Y \to D$;
(iv) c_i is not equal to c_{i+1}; and
(v) $\bar{c}_i$ does not intersect $P_Y \to D$ for $i < n$.

Indeed, consider a path $S \to D$ from u to $P_Y \to D$. Each open 1-cell of $S \to D$ either disconnects D or lies in the closure of a 2-cell. For consecutive edges e_1, e_2 that lie in 2-cells, either the 0-cell a that lies between e_1 and e_2 is a cut 0-cell of D or we add a sequence of 2-cells corresponding to a path in $\text{link}(a, D)$ between the vertices associated to e_1 and e_2. The last two properties are guaranteed by possibly passing to a subsequence.

If $n = 1$ and c_1 is a 2-cell, the proof concludes by letting $P \to D$ be a good path from u to $P_Y \to D$ as provided by Lemma 8.11(ii). If $n = 1$ and c_1 is a 1-cell, then the conclusion is immediate.

Suppose $n > 1$. A good path $Q \to D$ from u to $P_Y \to D$ is constructed as a concatenation $Q = P_1 \dots P_n$ of good paths $P_i \to D$ as follows. For each $i > 0$, either let b_i be a 1-cell in $\bar{c}_i \cap \bar{c}_{i+1}$ or let b_i be the cut 0-cell of D equal to $\bar{c}_i \cap \bar{c}_{i+1}$. Let $b_0 = u$. If c_i is an isolated 1-cell of D, then let $P_i \to D$ equal c_i.

Suppose that c_i is a 2-cell and $i < n$. Observe that if b_i is a 0-cell then b_i is a cut 0-cell of D and, by Lemma 8.7, does not map into ℓ. Suppose that b_i is a 1-cell. By Lemma 8.6 it is impossible for both endpoints of b_i to lie in ∂D and also to lie on $P_L \to D$ and map into ℓ. If both endpoints of b_i lie on ∂D, then property (v) of the sequence $\{c_n\}$ implies that the endpoints of b_i lie on $P_L \to Y$. Therefore either b_i is a cut 0-cell of D, or some (chosen) endpoint of b_i is either internal or does not map into ℓ. Since $\bar{c}_i$ is disjoint from $P_Y \to D$, Lemma 8.11(i) provides a good path $P_i \to \partial D$ from (our chosen points in) b_{i-1} to b_i. Note that the hypotheses are satisfied. Indeed, if b_i is a 0-cell then it is a cut 0-cell of D and therefore is not an internal 0-cell of a boundary arc of D. If b_i is a 1-cell, then either the chosen endpoint is internal, or the chosen endpoint is on ∂D and does not map into ℓ and is not an internal 0-cell of a boundary arc of D.

If c_n is a 2-cell, Lemma 8.11(ii) implies that there is a good path $P_i \to \partial D$ from the chosen endpoint of b_{n-1} to a 0-cell of $P_Y \to D$.

Finally, the good path $P \to D$ is the concatenation $P_1 \dots P_n$, which is a good path by Lemma 8.10. □

8.3. Subdivisions, Good Paths, and Internal Paths

LEMMA 8.12 (Subdividing along ℓ). *There is an H-equivariant subdivision X' of X that satisfies the following properties*:

(i) *ℓ is contained in the 1-skeleton of X'*;
(ii) *each cell of X' is convex*;
(iii) *each 0-cell of X' with nonzero curvature is a 0-cell of X*; *and*
(iv) *if X has nonpositive sectional curvature then so does X'.*

Proof. Since ℓ intersects finitely many cells and H acts properly on X, we see that each cell intersects finitely many H-translates of ℓ. The 1-skeleton of X' equals $X^1 \cup H\ell$. In particular, the 0-skeleton of X' consists of three types of cells: 0-cells of X, intersections of H-translates of ℓ with open 1-cells of X, and self-intersections of distinct H-translates ℓ within open 2-cells of X. Since each cell of X is convex and ℓ is a geodesic segment, it follows that each cell of X' is convex.

We now verify the third statement. Observe that each new 0-cell in X' is in the interior of either a 2-cell or a 1-cell of X. In the former case, the link is a circle with 2π-angle sum. In the latter case, the link is a finite subdivision of a bipartite graph Θ with two vertices, and each edge of Θ has angle π.

Finally, $\text{link}(x, X')$ is a subdivision of $\text{link}(x, X)$ when $x \in X^0$. Thus the last statement follows. □

The subdivision X' of X induces a subdivision of any complex that is immersed in X. Induced subdivisions of P_L, P_Y, D, Y, and L are denoted by P_L', P_Y', D', Y', and L', respectively. The geodesic ℓ is an edge path in L'. Divide the path $P_L' \to L'$ into paths $P_1', \dots, P_k'$ such that each P_i' either is a subpath of ℓ or else intersects ℓ only at its endpoints. Let ℓ_i be the subpath of ℓ between the endpoints of P_i'. (Note that the concatenation $\ell_1 \dots \ell_k$ may not equal ℓ, although it is a path within ℓ.) Let $K_i' \subseteq L'$ be the subdiagram between ℓ_i and P_i'. Let K' be the complex obtained by attaching to P_L' a copy of K_i' along the subpath P_i' for each i. Observe that $K' \to L'$ is a near-immersion and that $P_L' \to K'$ and $\ell_1 \dots \ell_k \to K'$ are embeddings. Since K' is contractible, $K' \to L'$ is a disk diagram between $P_L' \to L'$ and $\ell_1 \dots \ell_k \to L'$. See Figure 4.

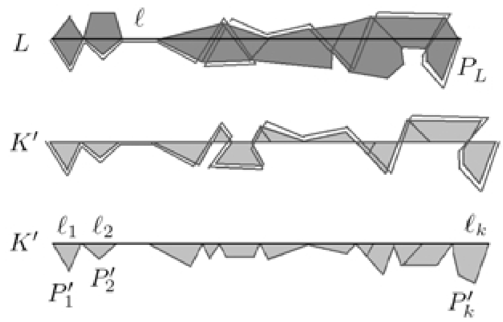

Figure 4 The paths $\ell \to L$ and $P_L \to L$ (top); the embedding $P'_L \to K'$ (bottom)

Let E' be the complex obtained by identifying D' and K' along the images of $P'_L \to D'$ and $P'_L \to K'$. Since D' and K' are disk diagrams and $P'_L \to K'$ is an embedding, it follows that E' is a disk diagram. The minimality assumption on D implies that $E' \to X'$ is a near-immersion; see Lemma 8.13. Therefore, $E' \to X'$ is a disk diagram between $P'_Y \to X'$ and $\ell_1 \dots \ell_k \to X'$. See Figure 5.

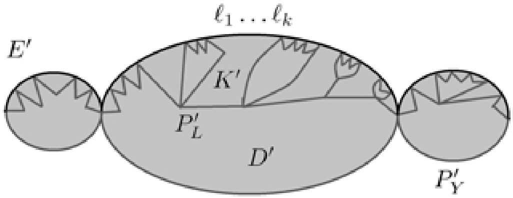

Figure 5 The disk diagram $E' \to X'$ between the paths $P'_Y \to X'$ and $\ell_1 \dots \ell_k \to X'$

LEMMA 8.13. *The map $E' \to X'$ is a near-immersion.*

Proof. As $D \to X$ and $K' \to X'$ are near-immersions, it suffices to examine the 1-cells of E' along $P'_L \to E'$. Let e' be a 1-cell of P'_L. Suppose R'_1 and R'_2 are distinct 2-cells of E' at (the image of) e' that map to the same 2-cell in X'. Since both $P'_L \to K'$ and $D' \to E'$ are embeddings, assume without loss of generality that R'_1 is in K' and that R'_2 is in D'.

We will show that the minimality of D is violated. Let e be the 1-cell of P_L containing e', let R_1 be the 2-cell of L containing the image of R'_1, and let R_2 be the 2-cell of D containing the image of R'_2. Let $\partial R_1 = Qe^{-1}$ and let $\dot{P}_L$ be formed from P_L by replacing e with Q. Similarly, let $\dot{D}$ be the subdiagram of D obtained by removing the open cells e and R_2. Observe that $\dot{D} \to X$ is a disk diagram between $\dot{P}_L \to X$ and $P_Y \to X$ that violates the minimality of D.

We now consider the case where both R'_1 and R'_2 are in K'. Observe that R'_1, R'_2 meet each other along the 1-cell c' of D' (the image of e'). Again, let e be the 1-cell of P_L containing e' and let c be the 1-cell of D containing c'. Then P_L is a concatenation $S_1eS_2fS_3$, where e and f travel through the 1-cell c. Let $\dot{P}_L$ equal S_1S_3 and note that $P_Y^{-1}\dot{P}_L \to D$ bounds a subdiagram $\dot{D}$ of D—namely, the part of D that remains after removing the subdiagram bounded by eS_2f. Observe that $\dot{P}_L$ is a path in L since e and f map to the same 1-cell of L. As before, $\dot{D}$ violates the minimality of D. □

8.4. The Immersion $Z' \to X'$

Let $W' = Y' \sqcup_{HP_Y} HE'$ denote the union of Y' and copies of E' attached along the distinct H-translates of P_Y. Since E' is a finite complex, W' is a cocompact H-complex. Since Y' and E' are simply connected, W' is simply connected. In view of Lemma 6.4, the H-map $W' \to X'$ factors as the composition of a surjection $W' \to Z'$ and an immersion $Z' \to X'$, where Z' is a simply connected cocompact H-complex.

LEMMA 8.14 (ℓ^2-Euler Characteristic). $\chi(H, Z') = \chi(H, Y)$.

Proof. Since $Z' \to X'$ is a immersion, Z' is a locally CAT(0)-space. Since Z' is simply connected, it is contractible. Analogously, Y' is contractible. Since Y' and Z' are both proper H-complexes and the embedding $Y' \to Z'$ is an H-map, Corollary 5.6 implies that $\chi(H, Z') = \chi(H, Y')$. Moreover, $\chi(H, Y') = \chi(H, Y)$ by the Gauss–Bonnet theorem or because the definition of ℓ^2-Betti numbers is independent of the cell structure of the space. □

LEMMA 8.15 (0-Cells with Positive Curvature). $|\mathrm{Pos}(H, Z')| \leq |\mathrm{Pos}(H, Y)|$.

Proof. It is enough to show that if z is a 0-cell of Z' such that

$$\mathrm{Curvature}(\mathrm{link}(z), H_z) > 0$$

then z is in the image of $Y' \to Z'$. Indeed, this statement implies that there is an injective map $\mathrm{Pos}(H, Z') \to \mathrm{Pos}(H, Y')$; moreover, $|\mathrm{Pos}(H, Y')| = |\mathrm{Pos}(H, Y)|$ since 0-cells of Y' arising as a result of the subdivision have zero curvature by Lemma 8.12.

Let z be a 0-cell of Z' and suppose that z is not in the image of $Y' \to Z'$. We show that $\mathrm{Curvature}(\mathrm{link}(z), H_z) \leq 0$. Since Z' has nonpositive sectional curvature and is positively angled, if $\mathrm{link}(z, Z')$ has a cycle or is disconnected then it is easy to see that $\mathrm{Curvature}(\mathrm{link}(z), H_z) \leq 0$. Suppose that $\mathrm{link}(z, Z')$ is a tree.

Let w be a preimage of z by $W' \to Z'$. Since z is not in the image of $Y' \to Z'$, it follows that w is a 0-cell in the image of $hE' \to W'$ for some $h \in H$. Without loss of generality, assume that $h = 1$. Since $E' \to Z' \to X'$ is a near-immersion and $\mathrm{link}(z)$ is a tree, it follows that $\mathrm{link}(w)$ is a tree. Since E' is a disk diagram, it follows that w is in the image of $\partial E \to W'$. Since the image of w by $W' \to Z'$ is not contained in the image of $Y' \to Z'$, it follows that w is not in the image of $P'_Y \to E'$. Therefore w is a 0-cell in the interior of $\ell_1 \dots \ell_k \to W'$, and hence z is in the interior of $\ell_1 \dots \ell_k \to Z'$.

Suppose that $\ell_1 \dots \ell_k \to Z'$ is locally a geodesic at z. By the construction of K', if e_1 and e_2 are 1-cells of $\ell_1 \dots \ell_k$ with a common endpoint z then $\mathrm{link}(z, Z')$

has a path between the vertices induced by e_1 and e_2 with angle sum at least π; hence $\text{Curvature}(\text{link}(z), H_z) \leq 0$.

Suppose $\ell_1 \ldots \ell_k \to Z'$ is not locally an embedding around z (i.e., suppose there is a backtrack). By the construction of K', z is the terminal point of ℓ_i for some $i < k$ and the path $\ell_i \ell_{i+1} \to X'$ has a backtrack. It follows that $\text{link}(z, K')$ consists of two nonedgeless components and that the angle sum of $\text{link}(z, K')$ is π. Observe that the two components of $\text{link}(z, K')$ are mapped into $\text{link}(z, X')$ to a path with angle sum equal to π; in particular, the angle sum of $\text{link}(z, Z')$ is no less than π and hence $\text{Curvature}(\text{link}(z), H_z) \leq 0$. □

LEMMA 8.16 (0-Cells with Negative Curvature).

$$|\text{Neg}(H, Z')| \leq \frac{2\pi \cdot \chi(H, Y) - \text{P}(G, X) \cdot |\text{Pos}(H, Y)|}{\text{N}(G, X)}.$$

Proof. Since angles are positive, observe that the constants of Definition 7.10 satisty $\text{N}(G, X') = \text{N}(G, X)$ and $\text{P}(G, X) = \text{P}(G, X')$. By Lemma 7.12,

$$|\text{Neg}(H, Z')| \leq \frac{2\pi \cdot \chi(H, Z') - \text{P}(G, X) \cdot |\text{Pos}(H, Z')|}{\text{N}(G, X)}.$$

The conclusion follows from the previous two lemmas and the above inequality. □

8.5. *Conclusion of the Proof of the Quasiconvex Core Theorem*

LEMMA 8.17 (From Good to Internal). *Suppose that $Q \to D$ is a good path whose interior does not intersect $P_Y \to D$. Then all 0-cells of $Q \to D$ are mapped to internal 0-cells of Z' by $Q' \to Z'$.*

Proof. Since $D' \to E'$ is an embedding, each 0-cell of the interior of D' is mapped to the interior of E'. Suppose u is a 0-cell of $Q \to D$ that is not in the interior of D. Since $Q \to D$ is good path, the image of u in X does not intersect ℓ; and by assumption, u is not in $P_Y \to D$. Therefore u is mapped into the interior of E'. Since $E' \to Z'$ is a near-immersion, the conclusion follows. □

Since X is G-cocompact, there is an upper bound $C = C(X)$ on the length of boundary paths of 2-cells of X.

LEMMA 8.18 (P_L is uniformly close to P_Y). *Let u be a 0-cell of P_L. Then the combinatorial distance from u to the subcomplex Y is bounded by the constant*

$$1 + C(X) + \frac{2\pi \cdot \chi(H, Y) - \text{P}(G, X) \cdot |\text{Pos}(H, Y)|}{\text{N}(G, X)}.$$

Proof. If u is in the $C(X)$-neighborhood of Y in X then the statement is clear. Otherwise, u is not in $P_Y \to D$ and so Lemma 8.7 implies that u is not a cut 0-cell of D. It follows that u is in the closure of a boundary arc of D and hence there is a v in P_L that is not an internal cell of a boundary arc of D and the distance between u and v is bounded by $C(X)$.

By Proposition 8.9, there is a good path $Q \to D$ from v to a 0-cell of $P_Y \to X$. Assume that $Q \to X$ has minimal combinatorial length. Then the interior of $Q \to X$ does not intersect $P_Y \to D$. By Lemma 8.17, $Q \to X$ factors as $Q' \to Z' \to X'$ and each 0-cell z of $Q \to D$ is mapped to an internal cell of Z'.

Since $Z' \to X'$ is an immersion, $\mathrm{link}(z, Z')$ has a cycle for each 0-cell z of $Q \to D$. Since X has sectional curvature $\leq \alpha < 0$, it follows that

$$\mathrm{Curvature}(H_z, \mathrm{link}(z, Z')) < 0.$$

By minimality, $Q \to X$ is embedded and, moreover, no pair of distinct 0-cells of $Q' \to Z'$ are in the same H-orbit. Therefore $|Q| \leq 1 + |\mathrm{Neg}(H, Z')|$.

An upper bound for the combinatorial distance between v and $P_Y \to X$ follows from the previous inequality and Lemma 8.16. The proof concludes by adding the upper bound $C(X)$ on the distance between u and v. □

By Lemmas 8.1 and 8.18, there is a uniform upper bound for the distance between ℓ and Y that is independent of ℓ. This concludes the proof of Theorem 1.6.

9. Large Quotients

THEOREM 9.1. *Let X be a* CAT(0) *cocompact and proper G-complex with sectional curvature $\leq \alpha < 0$. Let $g \in G$ be an infinite-order element. Let γ be an axis for g and suppose $\gamma \cap f\gamma$ is a discrete set for any $f \in G - \langle g \rangle$.*

There exists an $N > 0$ such that, for any $n \geq N$, the group $\bar{G} = G/\langle\langle g^n \rangle\rangle$ has a CAT(0) *cocompact $\bar{G}$-complex with sectional curvature $\leq \bar{\alpha} < 0$.*

LEMMA 9.2. *Let $\bar{\Gamma}$ be a graph, let e be an edge, and let Γ equal $\bar{\Gamma} - e$. Suppose Γ is an angle graph with nonnegative angles and sectional curvature $\leq \alpha \leq 0$, and suppose the angle distance $\measuredangle(\iota e, \tau e)$ in Γ is no less than π. Then $\bar{\Gamma}$ has sectional curvature ≤ 0 by assigning $\measuredangle(e) = \pi$.*

Suppose that Γ has sectional curvature $\leq \alpha < 0$, $\measuredangle(\iota e, \tau e) = \theta > \pi$, and $\measuredangle(e) > \pi + \alpha$. Then $\bar{\Gamma}$ has sectional curvature < 0.

Proof. Let $\bar{\Lambda}$ be a connected, spurless, and not edgeless subgraph of $\bar{\Gamma}$, and let Λ be $\bar{\Lambda} \cap \Gamma$. If $\Lambda = \bar{\Lambda}$ then $\mathrm{Curvature}(\bar{\Lambda}) = \mathrm{Curvature}(\Lambda) \leq \alpha$ by hypothesis. If e is an edge of $\bar{\Lambda}$ then $\mathrm{Curvature}(\bar{\Lambda}) = \mathrm{Curvature}(\Lambda) + \pi - \measuredangle(e)$. Therefore, $\mathrm{Curvature}(\bar{\Lambda}) \leq \alpha + \pi - \measuredangle(e) \leq 0$ if Λ was not a tree because removing spurs gives a section. Observe that the last inequality is strict if $\measuredangle(e) > \pi + \alpha$. Otherwise, removing some spurs gives rise to a subdivided interval and hence $\mathrm{Curvature}(\Lambda) \leq \pi - \measuredangle(\iota e, \tau e)$. In this case, $\mathrm{Curvature}(\bar{\Lambda}) \leq 0$ with strict inequality if $\measuredangle(\iota e, \tau e) > \pi$. □

Proof of Theorem 9.1. Slightly decreasing all the angles yields a $\mathrm{CAT}(-\varepsilon)$ structure on X. Let $\hat{X}$ be the quotient of X by $\langle\langle g^n \rangle\rangle$.

Let X' be a G-invariant subdivision of X such that γ lies in the 1-skeleton of X'. If we slightly increase the curvature of all 2-cells (increasing all the angles), then

X' has negative sectional curvature at all new 0-cells corresponding to intersections of γ with 1-cells of X and translates of γ.

Let σ_n be the subpath of γ from p to $g^n p$. Form $\bar{X}$ from $\hat{X}$ by attaching a 2-cell $\bar{h}R$ along the cycle $\bar{h}\sigma^n$ of $\hat{X}$; attach a 2-cell for each left coset $\bar{h}\langle \bar{g} \rangle$ in $\bar{G}$. Extend the $\bar{G}$-action on $\hat{X}$ to $\bar{X}$ by letting each $\langle \bar{g}^{\bar{h}} \rangle$ act with a fixed point at the center of $\bar{h}R$. Regard R as a Euclidean n-gon whose ith side is the translate of σ_1 by $\bar{g}^i$. Now we claim that for sufficiently large n, the complex $\bar{X}$ has negative sectional curvature.

Observe that the links of vertices of $\bar{X}$ are independent of n. By making all angles of X slightly larger, we can assume that the angle distances between the initial and terminal vertices of corners along ℓ all exceed π. By Lemma 9.2, for sufficiently large n we can assign an angle of $(n-2)\pi/n$ to the corners of R and its translates. We assign an angle of π to the other corners along the interior of σ_1.

By subdividing R into n 2-cells using its barycenter and then making its corners slightly smaller, we can assume that the barycenter has negative curvature and that G acts without inversions on $\bar{X}$. □

References

[1] M. R. Bridson, *Geodesics and curvature in metric simplicial complexes,* Group theory from a geometrical viewpoint (Trieste, 1990), pp. 373–463, World Scientific, River Edge, NJ, 1991.

[2] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature,* Grundlehren Math. Wiss., 319, Springer-Verlag, Berlin, 1999.

[3] M. Feighn and M. Handel, *Mapping tori of free group automorphisms are coherent,* Ann. of Math. (2) 149 (1999), 1061–1077.

[4] S. M. Gersten, *Reducible diagrams and equations over groups,* Essays in group theory, Math. Sci. Res. Inst. Publ., 8, pp. 15–73, Springer-Verlag, New York, 1987.

[5] M. Gromov, *Hyperbolic groups,* Essays in group theory, Math. Sci. Res. Inst. Publ., 8, pp. 75–263, Springer-Verlag, New York, 1987.

[6] W. Lück, *L^2-invariants: Theory and applications to geometry and K-theory,* Ergeb. Math. Grenzgeb. (3), 44, Springer-Verlag, Berlin, 2002.

[7] A. M. Macbeath, *Groups of homeomorphisms of a simply connected space,* Ann. of Math. (2) 79 (1964), 473–488.

[8] E. Martínez-Pedroza and D. T. Wise, *Local quasiconvexity of groups acting on small cancellation complexes,* J. Pure Appl. Algebra 215 (2011), 2396–2405.

[9] K. Matsuzaki and M. Taniguchi, *Hyperbolic manifolds and Kleinian groups,* Oxford Math. Monogr., Oxford Univ. Press, New York, 1998.

[10] J. P. McCammond and D. T. Wise, *Fans and ladders in small cancellation theory,* Proc. London Math. Soc. (3) 84 (2002), 599–644.

[11] ———, *Coherence, local quasiconvexity, and the perimeter of 2-complexes,* Geom. Funct. Anal. 15 (2005), 859–927.

[12] J. W. Morgan, *On Thurston's uniformization theorem for three-dimensional manifolds,* The Smith conjecture (New York, 1979), Pure Appl. Math., 112, pp. 37–125, Academic Press, Orlando, FL, 1984.

[13] G. P. Scott, *Finitely generated 3-manifold groups are finitely presented,* J. London Math. Soc. (2) 6 (1973), 437–440.

[14] H. Short, *Groups and combings,* preprint, ENS Lyon, 1990.
[15] D. T. Wise, *Sectional curvature, compact cores, and local quasiconvexity,* Geom. Funct. Anal. 14 (2004), 433–468.
[16] ———, *Nonpositive sectional curvature for (p, q, r)-complexes,* Proc. Amer. Math. Soc. 136 (2008), 41–48.

E. Martínez-Pedroza
Memorial University
Saint John's
Newfoundland and Labrador, A1C 5S7
Canada

emartinezped@mun.ca

D. T. Wise
McGill University
Montreal
Quebec H3A 2K6
Canada

wise@math.mcgill.ca