Generalizing the Croke–Kleiner construction

Christopher Mooney

Department of Mathematics, 2074 East Hall, 530 Church Street, Ann Arbor, MI 48109-1043, United States

A R T I C L E   I N F O

A B S T R A C T

It is well known that every word hyperbolic group has a well-defined visual boundary. An example of C. Croke and B. Kleiner shows that the same cannot be said for CAT(0) groups. All boundaries of a CAT(0) group are, however, *shape equivalent*, as observed by M. Bestvina and R. Geoghegan. Bestvina has asked if they also satisfy the stronger condition of being *cell-like equivalent*. This article describes a construction which will produce CAT(0) groups with multiple boundaries. These groups have very complicated boundaries in high dimensions. It is our hope that their study may provide insight into Bestvina's question.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

The CAT(0) condition is a geometric notion of nonpositive curvature, similar to the definition of Gromov δ-hyperbolicity. A complete geodesic space X is called CAT(0) if it has the property that geodesic triangles in X are "no fatter" than geodesic triangles in Euclidean space (see [4, Chapter II.1] for a precise definition). The *visual* or *ideal boundary* of X, denoted ∂X, is the collection of endpoints of geodesic rays emanating from a chosen basepoint. It is well known that ∂X is well defined and independent of choice of basepoint. Furthermore, when given the cone topology, $X \cup \partial X$ is a Z-set compactification for X. A group G is called CAT(0) if it acts geometrically (i.e. properly discontinuously and cocompactly by isometries) on some CAT(0) space X. In this setup, we call X a cocompact CAT(0) G-space and ∂X a CAT(0) boundary of G.

It is an important fact in geometric group theory that every negatively curved group (that is, every word hyperbolic group) has a well-defined visual boundary. Specifically, if a group G acts geometrically on two different CAT(-1) spaces, then the visual boundaries of these spaces will be homeomorphic. In the absence of strict negative curvature, the situation becomes more complicated.

We will call a CAT(0) group *rigid* if it has only one topologically distinct boundary. P.L. Bowers and K. Ruane showed that if G splits as the product of a negatively curved group with a free abelian group, then G is rigid [3]. Ruane proved later in [14] that if G splits as a product of two negatively curved groups, then G is rigid. T. Hosaka has extended this work to show that in fact it suffices to know that G splits as a product of rigid groups [8]. Another condition which guarantees rigidity is knowing that G acts on a CAT(0) space with isolated flats [9,10].

Not all CAT(0) groups are rigid, however: C. Croke and B. Kleiner constructed in [5] an example of a non-rigid CAT(0) group G. Specifically, they showed that G acts on two different CAT(0) spaces whose boundaries admit no homeomorphism. J. Wilson proved in [15] that this same group has *uncountably many* boundaries. More recently it has been shown in [12] that the knot group G of any connected sum of two nontrivial torus knots has uncountably many CAT(0) boundaries.

E-mail address: cpmooney@umich.edu.

On the other end of the spectrum, it has been observed by M. Bestvina [2], R. Geoghegan [6], and P. Ontaneda [13] that all boundaries of a given CAT(0) group are *shape equivalent*. Bestvina then posed the question of whether they satisfy the stronger condition of being *cell-like equivalent*. R. Ancel, C. Guilbault, and J. Wilson showed in [1] that all the currently known boundaries of Croke and Kleiner's original group satisfy this property; they are all cell-like equivalent to the Hawaiian earring.

Further progress on Bestvina's question has been hampered by a lack of examples of non-rigid CAT(0) groups. The results in [12] have boundaries which are similar to the boundaries of the Croke–Kleiner group and, as such, are unlikely to shed new light on Bestvina's question. One simple approach to producing other non-rigid CAT(0) groups is to take a direct product $G \times H$ of two CAT(0) groups where one of the factors is not rigid (either the Croke–Kleiner group or one of the groups from [12]). However, it was proven in [11] that the answer to Bestvina's question is "Yes" for CAT(0) groups of this form.

The goal of this article is to describe a construction which will yield a richer collection of non-rigid CAT(0) groups. Our work borrows freely from the main ideas of [5], but in the end, we have a flexible strategy for producing CAT(0) groups which have very complicated boundaries in high dimensions. We are not claiming new progress on Bestvina's question, but it is our hope that the study of this new collection may provide insight.

The main theorem of this paper is the following.

Theorem 1. *Let Γ_- and Γ_+ be infinite* CAT(0) *groups and m and n be positive integers. Then the free product with amalgamation*

$$G = \left(\Gamma_- \times \mathbb{Z}^m\right) *_{\mathbb{Z}^m} \left(\mathbb{Z}^m \times \mathbb{Z}^n\right) *_{\mathbb{Z}^n} \left(\mathbb{Z}^n \times \Gamma_+\right)$$

is a non-rigid CAT(0) *group.*

Observe that we get the Croke–Kleiner group if we take $\Gamma_- = \Gamma_+ = \mathbb{Z}$ and $m = n = 1$ in this theorem. In general, though, any boundary of G will contain spheres of dimension $m + n - 1$. Thus by choosing m and n to be large, we get a non-rigid CAT(0) group whose boundaries have high dimension.

2. Croke and Kleiner's original construction

Before diving into the proof of Theorem 1, we quickly sketch the proof of the main theorem of [5]. The CAT(0) spaces X constructed have the property that each is covered by a collection of closed convex subspaces, called *blocks*. The visual boundary ∂B of every block B is the suspension of a Cantor set. The suspension points are called *poles*. The intersection of two blocks is a Euclidean plane called a *wall*. We then have the following five statements for each X.

Theorem A. *([5, Section 1.4]) The nerve N of the collection of blocks is a tree.*

Theorem B. *([5, Lemma 3]) Let B_0 and B_1 be blocks, and D be the distance between the corresponding vertices in N. Then:*

(1) *If $D = 1$, then $\partial B_0 \cap \partial B_1 = \partial W$ where W is the wall $B_0 \cap B_1$.*
(2) *If $D = 2$, then $\partial B_0 \cap \partial B_1$ is the set of poles of $B_{\frac{1}{2}}$ where $B_{\frac{1}{2}}$ intersects B_0 and B_1.*
(3) *If $D > 2$, then $\partial B_0 \cap \partial B_1 = \emptyset$.*

A *local path component* of a point in a space is a path component of an open neighborhood of that point.

Theorem C. *([5, Lemma 4]) Let B be a block and $\zeta \in \partial B$ not be a pole of any neighboring block. Then ζ has a local path component which stays in ∂B.*

Theorem D. *([5, Corollary 8]) The union of block boundaries in ∂X is the unique dense safe path component of ∂X.*

The definition of *safe path* will be given in Section 6.3. For now, it suffices to understand that Theorem D gives a way to topologically distinguish the union of block boundaries in ∂X. With these thereoms in hand, it is not hard to prove that given two constructions X_1 and X_2, any homeomorphism $\partial X_1 \to \partial X_2$ takes poles to poles, block boundaries to block boundaries, and wall boundaries to wall boundaries. The last piece of the puzzle is Theorem E. Given $0 < \theta \leqslant \pi/2$, we can construct X_θ in such a way that the minimum distance between poles is θ. This distance is in the sense of the Tits path metric on the boundary of a block containing both poles. For a block B, we denote by ΠB the set of poles of blocks which intersect B at a wall.

Theorem E. *([5, Lemma 9], also [15, Proposition 2.2]) For a block B, the union of boundaries of walls of B is dense in ∂B and $\overline{\Pi B}$ is precisely the set of points of ∂B which are a Tits distance of θ from a pole of B.*

Using these five theorems, we get the following statement.

Theorem F. *Let B be a block and L be a suspension arc of ∂B. Then $|L \cap \overline{\Pi B}| = 1$ iff $\theta = \pi/2$.*

Therefore $\partial X_{\pi/2} \not\approx \partial X_\theta$ for any $\theta < \pi/2$, which is the main theorem of [5]. In this article, we produce for every group G in question a pair of cocompact CAT(0) G-spaces X and X' and show that Theorems A–D still hold. These four theorems, along with an analogue to Theorem E will be used to prove that there is no homeomorphism $\partial X \to \partial X'$.

3. Block structures on CAT(0) spaces

We begin by observing that the work in Sections 1.4–5 of [5] does not depend on the specific construction used in [5]. The same observations apply if we replace their definition of a *block* with the following one.

Definition 3.1. Let X be a CAT(0) space and $\mathcal{B}$ be a collection of closed convex subspaces covering X. We call $\mathcal{B}$ a *block structure on X* and its elements *blocks* if $\mathcal{B}$ satisfies the following three properties:

(1) Every block intersects at least two other blocks.
(2) Every block has a $(+)$ or $(-)$ parity such that two blocks intersect only if they have opposite parity.
(3) There is an $\epsilon > 0$ such that two blocks intersect iff their ϵ-neighborhoods intersect.

If we refer to blocks as *left* or *right*, we mean that the former have parity $(-)$ and the latter have parity $(+)$. The *nerve* of a collection $\mathcal{C}$ of sets is the (abstract) simplicial complex with vertex set $\{v_B \mid B \in \mathcal{C}\}$ such that a simplex $\{v_{B_1}, \ldots, v_{B_n}\}$ is included whenever $\bigcap_{i=1}^n B_i \neq \emptyset$. In exactly the same way as in [5], we get that the nerve $\mathcal{N}$ of the collection of blocks is a tree, and we can define the itinerary of a geodesic. A geodesic α is said to *enter* a block if it passes through a point which is not in any other block. The itinerary of α is defined to be the list $[B_1, B_2, \ldots]$ where B_i is the ith block that α enters. This list is denoted by $\operatorname{Itin} \alpha$. The following lemma follows in exactly the same way as [5, Lemma 2], which simply uses the fact that a block B is convex and that its topological frontier is covered by the collection of blocks corresponding to the link in $\mathcal{N}$ of the vertex v_B.

Lemma 3.2. *If $\operatorname{Itin} \alpha = [B_1, B_2, \ldots]$, then $[v_{B_1}, v_{B_2}, \ldots]$ is a geodesic in $\mathcal{N}$.*

We may also talk about the *itinerary between two blocks*. If $[v_{B_1}, \ldots, v_{B_n}]$ is the geodesic edge path in $\mathcal{N}$ connecting two vertices $v_{B'_0}$ and $v_{B'_1}$, then we call $[B_1, \ldots, B_n]$ the itinerary between B'_0 and B'_1 and write

$$\operatorname{Itin}\left[B'_0, B'_1\right] = [B_1, \ldots, B_n].$$

The two notions of itineraries are related in the following way: The itinerary of a geodesic segment α is the shortest itinerary $\operatorname{Itin}[B'_0, B'_1]$ for which α begins in B'_0 and ends in B'_1. Note that the same observations which gave us Lemma 3.2 also provide the following:

Lemma 3.3. *Let B'_0 and B'_1 be blocks, write $\operatorname{Itin}[B'_0, B'_1] = [B_1, \ldots, B_n]$, and let α be a geodesic beginning in B'_0 and ending in B'_1. Then:*

(1) α *enters B_k for every $1 < k < n$.*
(2) α *passes through $B_k \cap B_{k+1}$ for every $1 \leqslant k < n$.*
(3) $\bigcup_{k=1}^n B_k$ *is convex.*

We call a geodesic ray *rational* if its itinerary is finite and *irrational* if its itinerary is infinite. A point of ∂X is called *irrational* if it is the endpoint of an irrational geodesic ray; otherwise we call it *rational*. We denote the set of rational points of ∂X by RX and the set of irrational points by IX.

Lemma 3.4. *Let α be an irrational geodesic ray. Then for any block B_0,*

$$\lim_{t \to \infty} d\big(\alpha(t), B_0\big) = \infty.$$

Proof. Write $\operatorname{Itin} \alpha = [B_1, B_2, \ldots]$. Since $\mathcal{N}$ is a tree, we can find $M > 1$ so that for every $m \geqslant M$, $\operatorname{Itin}[B_0, B_m] \ni B_M$. For $m \geqslant M$, choose a time t_m such that $\alpha(t_m) \in B_m$. Then

$$\lim_{t \to \infty} d\big(\alpha(t), B_0\big) \geqslant \lim_{t \to \infty} d\big(\alpha(t), B_M\big)$$
$$= \lim_{m \to \infty} d\big(\alpha(t_m), B_M\big)$$
$$\geqslant \lim_{m \to \infty} d(B_m, B_M).$$

Hence, it suffices to prove the following.

Claim. *Let ϵ be given as in condition* (3) *of Definition* 3.1. *Then whenever $d(v_B, v_{B'}) \geqslant 2k$, we have $d(B, B') \geqslant 2k\epsilon$.*

Note that whenever $d(v_B, v_{B'}) = 2$, then we have $d(B, B') \geqslant 2\epsilon$ because the ϵ-neighborhoods of B and B' do not overlap. Assume Itin$[B, B'] = [B_0, B_1, \ldots, B_n]$, where $n \geqslant 2k$. Then for any $x \in B$ and $x' \in B'$, the geodesic $[x, x']$ passes through B_{2i} for $0 \leqslant i \leqslant k$ at some point z_i. So

$$d\big(x, x'\big) = \sum_{i=0}^{k-1} d(z_i, z_{i+1})$$
$$\geqslant 2k\epsilon. \quad \square$$

Corollary 3.5.

(1) *RX is the union of block boundaries in ∂X, and IX is its complement.*
(2) *If $\zeta \in IX$, then every geodesic ray going out to ζ is irrational.*
(3) *If $\zeta \in IX$ and α and β are geodesic rays going out to ζ, then the itineraries of α and β eventually coincide.*

A geodesic space is said to have the *geodesic extension property* if every geodesic segment can be extended to a geodesic line. As is true with the original Croke–Kleiner construction, the blocks we construct will satisfy the geodesic extension property.

Lemma 3.6. *If blocks have the geodesic extension property, then RX is dense.*

Proof. Let α be an irrational geodesic ray and write Itin$\alpha = [B_1, B_2, \ldots]$. For each $n \geqslant 1$, let t_n be a time at which $\alpha(t_n) \in B_n$. Then every ray $\alpha|_{[0, t_n]}$ can be extended to a geodesic ray α_n which does not leave the block B_n. Then $\alpha_n \to \alpha$. $\square$

Given a space Y, we call a map $\phi : IX \to Y$ an *irrational map* if it satisfies the property that $\phi(a) = \phi(b)$ iff whenever α and β are geodesic rays going out to a and b respectively, then Itinα and Itinβ eventually coincide. The obvious candidate for such a map is the function $\phi : IX \to \partial\mathcal{N}$ which takes a to the boundary point in $\partial\mathcal{N}$ determined by the itinerary of a ray going out to a. This function is well defined by Corollary 3.5(3). All we need to know is that ϕ is continuous, which amounts to proving the following lemma:

Lemma 3.7. *Let (α_n) be a sequence of irrational rays with common basepoint converging to another irrational ray α. Then for every $B \in$ Itinα, we have $B \in$ Itinα_n for large enough n.*

Proof. Write Itin$\alpha = [B_1, B_2, \ldots]$, and choose $k \geqslant 1$. Then B_{k+1} is a neighborhood of $\alpha(t)$ for some time t, which means that for large enough n, $\alpha_n(t) \in B_{k+1}$. Since $\alpha_n|_{[0, t]}$ begins in B_1 and ends in B_{k+1}, Lemma 3.3(1) tells us that it must enter B_k. $\square$

Corollary 3.8. *The natural map $\phi : IX \to \partial\mathcal{N}$ determined by itineraries is an irrational map.*

4. Some local homology calculations

We will use singular homology, with [7, pp. 108–130] as our reference. Here are a couple of key technical lemmas. All homology will be computed using $\mathbb{Z}$ coefficients. Recall that saying that π is a local path component of a point x in a space X means that there is an open neighborhood U of x such that π is the path component of U containing x.

Lemma 4.1. *Local homology can be computed using local path components. That is, for a point x in a topological space X with local path component π, we have*
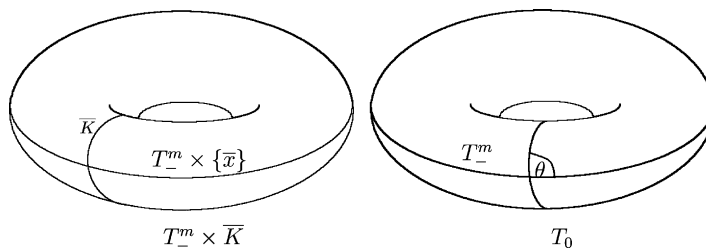
$$H_*(X, X - x) = H_*(\pi, \pi - x).$$

Fig. 1. $\overline{Y}$.

Proof. Let U be an open neighborhood of x which has π as the path component containing x. Using excision, we get

$$H_*(X, X - x) = H_*(U, U - x).$$

Now, since the image of a singular simplex σ is path connected, the chain complex $C_*(U)$ splits as $C_*(U - \pi) \oplus C_*(\pi)$. Passing to the relative chain complex $C_*(U, U - x)$ kills off the entire first factor. Thus $C_*(U, U - x) = C_*(\pi, \pi - x)$ and $H_*(U, U - x) = H_*(\pi, \pi - x)$. □

Lemma 4.2. *Let X be a path connected space, $x_0 \in X$, and B be the open n-ball with $z_0 \in B$. Then the local homology $H_*(B \times X, B \times X - (z_0, x_0))$ is zero whenever $* < n$. When $* = n$, it is nonzero only when X is a single point, in which case it is $\mathbb{Z}$.*

Proof. The statement is obvious in the case where $X = \{x_0\}$. If $X \neq \{x_0\}$, then $H_0(X, X - x_0) = 0$. $H_*(B, B - z_0)$ is $\mathbb{Z}$ when $* = n$ and 0 everywhere else. So the relative Künneth formula gives an isomorphism

$$H_*\big(B \times X, B \times X - (z_0, x_0)\big) = \bigoplus_{i=0}^{*} H_i(B, B - z_0) \otimes H_{*-i}(X, X - x_0) = 0. □$$

5. Hyperblocks

For positive integers m and n and an infinite CAT(0) group Γ, define

$$G_0 = \big(\Gamma \times \mathbb{Z}^m\big) *_{\mathbb{Z}^m} \big(\mathbb{Z}^m \times \mathbb{Z}^n\big) = \mathbb{Z}^m \times \big(\Gamma * \mathbb{Z}^n\big),$$

and choose a geometric action of $\mathbb{Z}^m \times \mathbb{Z}^n$ on $E_0 = \mathbb{R}^m \times \mathbb{R}^n$ (by translations). We will denote the convex hull of the $\mathbb{Z}^m$-orbit of the origin by E_-^m and the convex hull of the $\mathbb{Z}^n$-orbit of the origin by E_+^n (these are just isometric copies of $\mathbb{R}^m$ and $\mathbb{R}^n$). The minimal angle θ between geodesics in the two subspaces is called the *skew*, which can be any number $0 < \theta \leqslant \pi/2$. The quotient of E_0 by the group action is an $m + n$ torus T_0 with m- and n-tori $T_-^m = E_-^m/\mathbb{Z}^m$ and $T_+^n = E_+^n/\mathbb{Z}^n$; we call these the *left-* and *right-hand* subtori of T_0.

Choose a CAT(0) space K on which Γ acts geometrically. In order to simplify matters, we will temporarily assume that the action is free so that $\overline{K} = K/\Gamma$ is a $K(\Gamma, 1)$. Choose a point $\bar{x} \in \overline{K}$ and glue $T_-^m \times \overline{K}$ to T_0 via the isometry $T_-^m \times \{\bar{x}\} \to T_-^m$. The resulting space is nonpositively curved [4, Proposition II.11.6(2)] and is a $K(G_0, 1)$; we denote it by $\overline{Y} = \overline{Y}(\Gamma, m, n)$ (see Fig. 1). Its universal cover Y is a CAT(0) G_0-space, which we call a *hyperblock*. Path components of $p^{-1}(T_0)$ are called *walls*; these are isometric copies of E_0. The names *hyperblock* and *wall* are given because these spaces will play the role in this paper that "blocks" and "walls" play in [5].

5.1. The hedge

Lemma 5.1. *Y splits as a product $E_-^m \times H$ where H comes with a natural block structure in the sense of Section 3. In this block structure, left blocks are isometric copies of K and right blocks are isometric copies of $\mathbb{R}^n$. The intersection of two blocks is at most a single point.*

Proof. In the same way as in [5], Y comes with a natural block structure in which left blocks are path components of $p^{-1}(T_-^m \times \overline{K})$ and right blocks are path components of $p^{-1}(T_0)$. We begin with a left block B_0 and an isometry

$$\phi_0 : E_-^m \times K \to B_0$$

which comes naturally from the splitting of the downstairs space. Let $\mathcal{B}_1$ denote the collection of right blocks which intersect B_0. For each $B \in \mathcal{B}_1$, we have

$$B \cap B_0 = \phi_0\big(E_-^m \times x_B\big)$$

for some point $x_B \in K$. Now $(E_-^m)^\perp$ is the vector subspace of E orthogonal to E_-^m. There is also an isometry

$$\phi_1^B : E_-^m \times \left(E_-^m\right)^\perp \to B$$

such that

$$B \cap B_0 = \phi_1^B\left(E_-^m \times y_B\right)$$

for some $y_B \in (E_-^m)^\perp$. We choose this isometry so that for every $z \in E_-^m$,

$$\phi_1^B(z, y_B) = \phi_0(z, x_B).$$

Define $L_B = \phi_1^B((E_-^m)^\perp)$, $y_B' = \phi_1^B(y_B)$,

$$D_1 = B_0 \cup \bigcup \mathcal{B}_1,$$

and

$$H_1 = K \cup \bigcup_{B \in \mathcal{B}_1} L_B / \sim$$

where $\sim$ is generated by the rule $x_B \sim y_B'$. Then we can extend ϕ_0 to an isometry

$$\phi_1 : E_-^m \times H_1 \to D_1$$

by letting $\phi_1(z, x) = \phi_1^B(z, x)$ whenever $x \in L_B$.

Now, let $\mathcal{B}_2$ be the collection of left blocks intersecting D_1. As before, every $B \in \mathcal{B}_2$ intersects D_1 at a subspace of the form $\phi_1(E_-^m \times x_B)$ for some point $x_B \in H_1$ and any natural isometry $E_-^m \times K \to B$ has the property that the intersection $B \cap D_1$ is the image of $E_-^m \times y_B$ for some $y_B \in K$. Denote the image of $0 \times K$ under this map by L_B and the image of $(0, y_B)$ by y_B'. Let

$$D_2 = D_1 \cup \bigcup \mathcal{B}_2$$

and

$$H_2 = H_1 \cup \bigcup_{B \in \mathcal{B}_2} L_B / \sim$$

where $\sim$ is generated by the rule $x_B \sim y_B'$ for $B \in \mathcal{B}_2$. In the same was above, ϕ_1 can be extended to an isometry

$$\phi_2 : E_-^m \times H_2 \to D_2.$$

Continue in this manner to get for every n an isometry $\phi_n : E_-^m \times H_n \to D_n$ where H_n has the desired block structure. $\quad\square$

It is now easy to see that we can remove the assumption that the action of Γ is free. We have an action of $\mathbb{Z}^m$ on E_-^m which comes from the action on E_0. Take the action of $\Gamma \times \mathbb{Z}^m$ on $K \times E_-^m$ to be the corresponding product action and use the proof of the previous lemma as prescription for construction Y by gluing copies of $K \times E_-^m$ to copies of E_0 along equivariant collections of m-flats.

Note that this splitting corresponds directly to the group splitting $G_0 = \mathbb{Z}^m \times [\Gamma * \mathbb{Z}^n]$ where $\mathbb{Z}^m$ acts only in the E_-^m-coordinate. Then the projected action of $\Gamma * \mathbb{Z}^n$ on H is a geometric action and the stabilizer of a block of H is a subgroup conjugate to either Γ or $\mathbb{Z}^n$. A good example to keep in mind at this point is $G_0 = \mathbb{Z} \times [\mathbb{Z} * \mathbb{Z}]$ (as in the Croke–Kleiner group). In this case $Y = \mathbb{R} \times H$ where H is the infinite 4-valent tree.

From the splitting $Y = E_-^m \times H$, we get that ∂Y is the join $\partial E_-^m * \partial H$. We call the points of ∂E_-^m *poles* of Y. The set of poles, which we denote PY, is an $(m-1)$-sphere. We call H a *hedge* and the blocks described in Lemma 5.1 of the hedge *leaves*. Those leaves L for which the space $E_-^m \times L$ is a wall (a path component of $p^{-1}(T_0)$) are called *right leaves* (these are isometric copies of $\mathbb{R}^n$). The other leaves are called *left leaves* (these are isometric copies of K). The points of intersection of leaves are called *gluing points*. Since the set of gluing points in a leaf is the orbit of a single one by a geometric action (of either $\mathbb{Z}^n$ or Γ), this set is discrete and quasi-dense in the leaf. With this block structure in hand, we can talk about itineraries of geodesics and geodesic rays in the hedge in terms of leaves. We will denote the itinerary of a geodesic α in H by $\text{Itin}_H \alpha$. The set of irrational points of ∂H will be denoted by IH and the set of rational points by RH. Geodesics in the hedge are easy to compute: If $\text{Itin}_H[x, y] = [L_1, \ldots, L_n]$ and $y_i = L_i \cap L_{i+1}$ for $1 \leqslant i < n$, then

$$[x, y] = [x, y_1] \cup [y_1, y_2] \cup \cdots \cup [y_{n-1}, y]$$

where each of these segments is taken in a leaf.

Given a CAT(0) space X with points $p, x, y \in X$, the *Alexandrov angle between the geodesics* $[p, x]$ *and* $[p, y]$ is defined to be the angle between the initial velocities of these geodesics; this number is denoted by $\angle_p(x, y)$. If α and β are geodesic

segments [rays] based at p, then we may denote the Alexandrov angle between them by $\angle_p(\alpha, \beta)$. If α and β are rays with endpoints ζ and $\eta \in \partial X$, then we may also write $\angle_p(\zeta, \eta) = \angle_p(\alpha, \beta)$. Given $\zeta, \eta \in \partial X$, the *Tits* angle between ζ and η is defined to be

$$\angle_{Tits}(\zeta, \eta) = \sup_{p \in X} \angle_p(\zeta, \eta).$$

In [4, Corollary 9.9], it is shown that for geodesic rays α and β,

$$\angle_{Tits}\big(\alpha(\infty), \beta(\infty)\big) = \angle_p(\alpha, \beta)$$

iff α and β bound a flat sector (or their union is a geodesic line, in which case this angle is $\pi/2$). For closed subspaces C and D of ∂X, we may also write

$$\angle_{Tits}(C, D) = \min\big\{ \angle_{Tits}(\zeta, \eta) \mid \zeta \in C, \ \eta \in D \big\}.$$

In the context of this section, we will let

$$q : PY \times \partial H \times [0, \pi/2] \to \partial Y$$

denote the natural quotient map where $q(\eta, \zeta, 0) = \eta$ and $q(\eta, \zeta, \pi/2) = \zeta$. That is, ∂H disappears at level 0 and PY disappears at level $\pi/2$, and t is the Tits angle between $q(\eta, \zeta, 0)$ and $q(\eta, \zeta, t)$.

Lemma 5.2. *Consider the induced actions of G_0 on ∂Y and $\Gamma * \mathbb{Z}^n$ on ∂H. For $g \in G_0$, and $q(\eta, \zeta, t) \in \partial Y$, we have*

$$gq(\eta, \zeta, t) = q(\eta, g_2\zeta, t)$$

where g_2 is the image of g under the isomorphism $G_0/\mathbb{Z}^m \to \Gamma * \mathbb{Z}^n$.

Proof. Given a point $p = (p_1, p_2) \in Y = E_-^m \times H$ and points $\eta \in PY$ and $\zeta \in \partial H$, let α and β be geodesic rays based at p going out to η and ζ respectively. Then we can write $\alpha(t) = (\alpha_1(t), p_2)$ and $\beta(t) = (p_1, \beta_2(t))$ where α_1 is a ray in E_-^m and β_2 is a ray in H. Because $Y = \text{Min}\,\mathbb{Z}^m$, every $g \in G_0$ can be written coordinate-wise as (g^E, g^H) where g^E is a translation of E_-^m and g^H is an isometry of H. Of course, since g_1 acts only in the E_-^m-coordinate of Y, $g^H = g_2^H$. Using this, we compute

$$\begin{aligned}
d\big(g\alpha(t), \alpha(t)\big) &= \sqrt{d_{E_-^m}\big(g^E\alpha_1(t), \alpha_1(t)\big)^2 + d_H\big(g^H p_2, p_2\big)^2} \\
&= \sqrt{d_{E_-^m}\big(g^E p_1, p_1\big)^2 + d_H\big(g^H p_2, p_2\big)^2} \\
&= d(gp, p),
\end{aligned}$$

which means that $g\eta = \eta$, and

$$\begin{aligned}
d\big(g\beta(t), g_2\beta(t)\big) &= \sqrt{d_{E_-^m}\big(g^E p_1, p_1\big)^2 + d_H\big(g^H \beta_2(t), g_2^H \beta_2(t)\big)^2} \\
&= d_{E_-^m}\big(g^E p_1, p_1\big),
\end{aligned}$$

which means that $g\zeta = g_2\zeta$. Since g takes the flat quadrant $\alpha[0, \infty) \times \beta[0, \infty)$ to the flat quadrant $g\alpha[0, \infty) \times g\beta[0, \infty)$, the lemma follows. $\square$

When $g \in \Gamma * \mathbb{Z}^n$, we will simply write the equation in this lemma as

$$gq(\eta, \zeta, t) = q(\eta, g\zeta, t),$$

since there is no confusion.

Lemma 5.3. *IH is dense in ∂H.*

Proof. We begin by proving the following claim.

Claim. *For distinct points $x, y_1 \in H$ such that y_1 is a gluing point, the geodesic $[x, y_1]$ can be extended to an irrational ray.*

Let L_1 be the leaf which intersects the geodesic $[x, y_1]$ only at the point y_1. Since the set of gluing points in L_1 is quasi-dense (and Γ is infinite), we can choose another gluing point $y_2 \in L_1 - y_1$, say $y_2 = L_1 \cap L_2$. Then $[x, y_2] = [x, y_1] \cup [y_1, y_2]$. Get another gluing point $y_3 \in L_2 - y_2$, say $y_3 = L_2 \cap L_3$, so that

$$[x, y_3] = [x, y_1] \cup [y_1, y_2] \cup [y_2, y_3].$$

Continue in this manner, always extending the geodesic into a new leaf to get an irrational ray. This proves the claim.

Now, let α be any rational ray in H based at a point x, where L is the last leaf α enters. Get a sequence z_n of gluing points in L converging to $\alpha(\infty)$ in $L \cup \partial L$. By the claim, we can extend every geodesic $[x, z_n]$ to an irrational ray α_n. Then $\alpha_n \to \alpha$. $\square$

It turns out that RH is also dense in ∂H. In fact, something much stronger is true:

Lemma 5.4. *Let $\mathcal{L}$ be either the collection of left leaves or right leaves and $\mathcal{S} = \{\zeta_L\}_{L \in \mathcal{L}}$ be a subset of ∂H obtained by choosing for each $L \in \mathcal{L}$ a point $\zeta_L \in \partial L$. Then $\mathcal{S}$ is dense in ∂H.*

Proof. We show that every irrational point is a limit point of $\mathcal{S}$. Choose any irrational ray α based at a non-gluing point $x \in L_1$ with infinite itinerary $[L_1, L_2, \ldots]$. For each i, denote the gluing point $L_i \cap L_{i+1}$ by y_i. Then either all the even leaves or all the odd leaves are from $\mathcal{L}$; we'll assume it's the evens. Let α_i be the geodesic ray based at x which goes out to $\zeta_{L_{2i}}$. Then since each α_i eventually stays in L_{2i}, it agrees with α along the segment $[x, y_{2i-1}]$, and $\alpha_i \to \alpha$. $\square$

Lemma 5.5. *If two geodesic rays in a hedge have different itineraries, then they lie in different path components of ∂H.*

Proof. It suffices to prove the following claim.

Claim. *Choose a basepoint $p \in H$ (not a gluing point), and let L be a leaf not containing p. Let Ω be the collection of geodesic rays that have L in their itineraries, and $\Omega(\infty) \subset \partial H$ be the set of endpoints of rays from Ω. Then $\Omega(\infty)$ is both open and closed.*

Let y be the gluing point at which every ray in Ω enters L, and get $\epsilon > 0$ so that the open ϵ-ball $B_\epsilon(y)$ in L based at y contains no other gluing points. Then $\Omega(\infty)$ is open because Ω is the collection of rays passing through the open space $B_\epsilon(y) \cap L - y$. It is closed, since whenever (α_n) is a sequence of rays in Ω converging to a ray α, the sequence of points where α_n intersects the closed space $\partial B_\epsilon(y) \cap L$ converges to a point on $\partial B_\epsilon(y) \cap L$, showing that α also enters L. $\square$

Corollary 5.6.

(1) *Every irrational point is a path component of ∂H.*
(2) *∂H is nowhere locally path connected.*
(3) *∂Y is locally path connected precisely at its poles.*

For $\zeta \in \partial H$, we define the *longitude* of Y at ζ to be the subspace

$$l(\zeta) := PY * \zeta$$

of ∂Y. Since PY is an $(m-1)$-sphere, $l(\zeta)$ is an m-ball. $\Gamma * \mathbb{Z}^n$ permutes right leaves transitively and for every right leaf L of H, the wall $W_L = E_-^m \times L$ is stabilized by the conjugate subgroup

$$G_L = \left[\mathbb{Z}^m \times \mathbb{Z}^n\right]_{g_L} = g_L \left[\mathbb{Z}^m \times \mathbb{Z}^n\right] g_L^{-1}$$

for some element $g_L \in \Gamma * \mathbb{Z}^n$. The limit set of $[\mathbb{Z}^n]_{g_L}$ (that is, the set of limit points in ∂W_L of the orbit of a single point in W_L) will be denoted by S_L; this is an $(n-1)$-sphere. Note that since the action of G_L on W_L comes from the original action of $\mathbb{Z}^n \times \mathbb{Z}^m$ on E, we have

$$\angle_{Tits}(PY, S_L) = \theta,$$

where θ is the skew. Consider the subset

$$\Pi Y = \bigcup_{\text{right leaves } L} S_L.$$

Later we will see that ΠY is the collection of poles of hyperblocks which intersect Y at a wall.

Lemma 5.7. $\Gamma * \mathbb{Z}^n$ *permutes the spheres of ΠY transitively.*

Proof. This follows from the fact that $\Gamma * \mathbb{Z}^n$ permutes right leaves transitively and the equation

$$\lim\left[\mathbb{Z}^n\right]_{g_L} = g_L \lim \mathbb{Z}^n. \qquad \square$$

The topological object which we will use to distinguish between boundaries is called the *watermark* $\mathcal{W}(Y)$. It is defined as follows: let L_0 be the right leaf which is stabilized by $\mathbb{Z}^n$. Then $\mathcal{W}(Y)$ is the image of S_{L_0} under the map $\partial W_{L_0} \to l(\zeta_0)$ which is given by the rule

$$q(\eta, \zeta, t) \mapsto q(\eta, \zeta_0, t)$$

for some (any) fixed $\zeta_0 \in \partial H$. The following lemma guarantees that this is well defined.

Lemma 5.8. *For every $\zeta_0 \in \partial H$, $l(\zeta_0) \cap \overline{\Pi Y} \approx \mathcal{W}(Y)$.*

Proof. We begin with the forward inclusion; assume we have been given $\nu \in l(\zeta_0) \cap \overline{\Pi Y}$, say

$$\nu = q(\eta, \zeta_0, \lambda).$$

Then there is a sequence $(\nu_n) \subset \Pi Y$ converging to ν, say $\nu_n = q(\eta_n, \zeta_n, \lambda_n)$. For each n, let S_{L_n} be the sphere of ΠY containing ν_n. By Lemma 5.7, we can get a sequence $(g_n) \subset \Gamma * \mathbb{Z}^n$ for which $g_n \nu_n \in S_{L_0}$ for all n. By passing to a subsequence, we may assume that $g_n \nu_n \to \nu' \in S_{L_0}$. Write $\nu' = q(\eta', \zeta', \lambda')$. Using Lemma 5.2, we have $g_n \nu_n = q(\eta_n, g_n \zeta_n, \lambda_n)$. If $\lambda < \pi/2$, then we must have $\eta_n \to \eta$ and $\eta' = \eta$. If $\lambda = \pi/2$, then we have $q(\eta', \zeta', \lambda) = q(\eta, \zeta', \lambda)$. Either way,

$$\begin{aligned}\nu' &= \lim_{n \to \infty} g_n \nu_n \\ &= q(\eta, \zeta', \lambda),\end{aligned}$$

showing that $q(\eta, \zeta', \lambda) \in S_{L_0}$ and $\nu \in \psi(S_{L_0})$.

For the reverse inclusion, assume we have $\nu = q(\eta, \zeta', \lambda) \in S_{L_0}$. By Lemma 5.4, the $\Gamma * \mathbb{Z}^n$-orbit of ζ' is dense in ∂H, which means that we can get a sequence $(g_n) \subset \Gamma * \mathbb{Z}^n$ such that $g_n \zeta' \to \zeta_0$. This done, we have

$$\begin{aligned}\lim_{n \to \infty} g_n \nu &= \lim_{n \to \infty} q(\eta, g_n \zeta', \lambda) \\ &= q(\eta, \zeta_0, \lambda),\end{aligned}$$

showing that $q(\eta, \zeta_0, \lambda) \in \overline{\Pi Y}$, as desired. $\square$

In the original Croke–Kleiner construction, the watermark of every block is exactly two points if the skew is less than $\pi/2$ and one point if the skew is precisely $\pi/2$. The appropriate analogue in this setting is the following.

Proposition 5.9. *The watermark of Y contains exactly one point iff the skew is $\pi/2$.*

5.2. Local homology of ∂Y

Given a topological space A, we let CA denote the cone on A. We will also denote by $C^o A$ the "open cone" on A; that is,

$$C^o A = CA - A.$$

Lemma 5.10. *Let ζ be a pole of Y. Then $H_k(\partial Y, \partial Y - \zeta)$ is zero when $k < m$ and uncountably generated when $k = m$.*

Proof. Get an open neighborhood $V \subset PY$ of ζ which is homeomorphic to the open $(m-1)$-ball and consider

$$U = q\big(V \times \partial H \times [0, \pi/2)\big).$$

This is an open neighborhood of ζ homeomorphic to

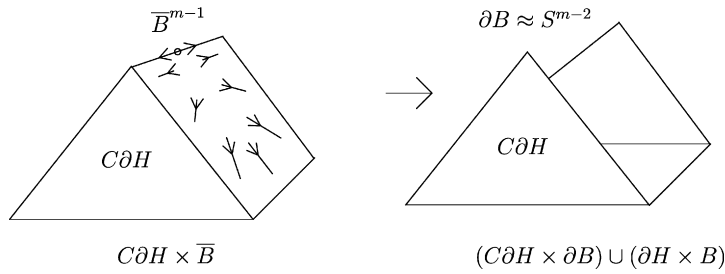$$(B * \partial H) - \partial H \approx C^o \partial H \times B$$

Fig. 2. A deformation retraction.

where B is the open $(m-1)$-ball. Assume $k > 1$ and look at the long exact sequence for the pair $(U, U - \zeta)$. Since U is contractible, $H_k(U) = H_{k-1}(U) = 0$ and we get

$$H_k(U, U - \zeta) = H_{k-1}(U - \zeta).$$

Now,

$$U - \zeta \stackrel{HE}{\simeq} C\partial H \times \overline{B} - (p, x)$$

where p is the cone point of $C\partial H$ and $x \in B$. This last deformation retracts onto the subspace

$$(C\partial H \times \partial B) \cup (\partial H \times B) \approx \partial H * S^{m-2}$$
$$= \Sigma^{m-1}\partial H$$

(see Fig. 2). So

$$
\begin{aligned}
H_k(U, U - \zeta) &= H_{k-1}\big(\Sigma^{m-1}\partial H\big) \\
&= \widetilde{H}_{k-1}\big(\Sigma^{m-1}\partial H\big) \\
&= \widetilde{H}_0\big(\Sigma^{m-k}\partial H\big) \\
&= 0
\end{aligned}
$$

where $\widetilde{H}_*$ denotes reduced homology. Now consider what happens to the long exact sequence for the pair $(U, U - \zeta)$ when $k = 0, 1$:

$$\to H_1(U) \to H_1(U, U - \zeta) \to H_0(U - \zeta) \to H_0(U) \to H_0(U, U - \zeta) \to 0.$$

Because U is path connected, we have $H_1(U) = H_0(U, U - \zeta) = 0$ and $H_0(U) = \mathbb{Z}$. So we are left with the short exact sequence

$$0 \to H_1(U, U - \zeta) \to H_0\big(\Sigma^{m-1}\partial H\big) \to \mathbb{Z} \to 0.$$

If $m > 1$, then $H_0(\Sigma^{m-1}\partial H) = \mathbb{Z}$ and $H_1(U, U - \zeta) = 0$. If $m = 1$, then $H_0(\Sigma^{m-1}\partial H) = H_0(\partial H)$ is uncountably generated and $H_1(U, U - \zeta)$ is also uncountably generated. □

Lemma 5.11. *Let $\zeta \in \partial Y - PY$. Then $H_k(\partial Y, \partial Y - \zeta)$ is finitely generated for $k \leqslant m$.*

Proof. We begin by finding a local path component of ζ which is homeomorphic to $B \times V_\zeta$ where B is the open m-ball and V_ζ is a local path component in ∂H. If $\zeta \notin \partial H$, then this is easy, since q is a homeomorphism on $PY \times \partial H \times (0, \pi/2)$. If $\zeta \in \partial H$, then we can get an open neighborhood of ζ of the form

$$U = q\big(PY \times V \times (0, \pi/2]\big)$$

where V is an open neighborhood of ζ in ∂H. Then

$$U \approx C^o PY \times V \approx B \times V,$$

and the path component of ζ in U has the form $B \times V_\zeta$ where V_ζ is the path component of ζ in V. Therefore

$$H_k(\partial Y, \partial Y - \zeta) = H_k(B \times V_\zeta, B \times V_\zeta - \zeta).$$

The conclusion now follows from Lemma 4.2. □

Corollary 5.12. *The local homology $H_k(\partial Y, \partial Y - \zeta)$ at a point $\zeta \in \partial Y$ is finitely generated for $k < m$. For $k = m$, it's infinitely generated iff ζ is a pole.*
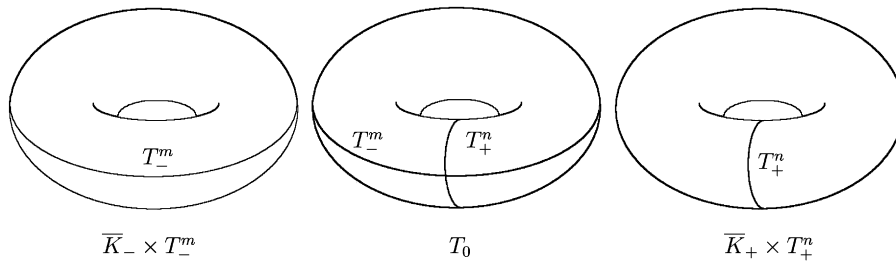
Fig. 3. $\overline{X}$.

6. The generalized Croke–Kleiner construction

Let Γ_- and Γ_+ be infinite CAT(0) groups and $m \leqslant n$ be positive integers, and define

$$G = \left(\Gamma_- \times \mathbb{Z}^m\right) *_{\mathbb{Z}^m} \left(\mathbb{Z}^m \times \mathbb{Z}^n\right) *_{\mathbb{Z}^n} \left(\mathbb{Z}^n \times \Gamma_+\right).$$

Choose $T_0(m, n) = E/\mathbb{Z}^m \times \mathbb{Z}^n$ with skew θ and form the spaces $\overline{K}_- = \overline{K}(\Gamma_-)$, $\overline{Y}_- = \overline{Y}(\Gamma_-, m, n)$. We also form $\overline{Y}_+ = \overline{Y}(\Gamma_+, n, m)$ where we use the same torus T_0 but with "left" and "right" swapped (this corresponds to a change of coordinates in K) so that $\overline{Y}_- \cap \overline{Y}_+ = T_0$. We define $\overline{X} = \overline{Y}_- \cup \overline{Y}_+$ (see Fig. 3).

The universal cover X of $\overline{X}$ is a CAT(0) G-space. We denote the universal covering projection by p and call X a *generalized Croke–Kleiner construction for G from the spaces T_0, $\overline{K}_-$, and $\overline{K}_+$.* The path components of $p^{-1}(T_0)$ are isometric copies of E; we call these *walls*. The path components of $p^{-1}(Y_-)$ and $p^{-1}(Y_+)$ are hyperblocks; we call them *left* and *right blocks* respectively, and denote them by B_- and B_+ if we wish to designate parity. The following lemmas are easy.

Lemma 6.1. *If a wall intersects a block, then that wall is contained in that block.*

Lemma 6.2. *If two blocks intersect, then they have opposite parity and their intersection is precisely a wall.*

Lemma 6.3. *Every block B splits as a product $E \times H$ where E is a Euclidean space and H is a hedge.*

In this last lemma, E and H are of two types depending on the parity of B. If B is a left block, then E is an isometric copy of E_-^m and H is a hedge whose leaves are copies of K_- and E_-^m. If B is a right block, then E is a copy of E_+^n and the leaves of H are copies of K_+ and E_+^n. We call H the *hedge factor* of B.

6.1. The block structure

To see that the collection of hyperblocks satisfies Definition 3.1, we need only the following.

Lemma 6.4. *There is an $\epsilon > 0$ such that two blocks intersect iff their ϵ-neighborhoods intersect.*

Proof. Take $\epsilon > 0$ to be smaller than half the lengths of the shortest nontrivial loops in $\overline{K}_-$ and $\overline{K}_+$. Suppose B_1 and B_2 are disjoint blocks, and let γ be a geodesic starting in B_1 and ending in B_2. Without loss of generality, assume B_1 is a left block. Then $\gamma' = p\gamma$ is a local geodesic in $\overline{X}$ which leaves $\overline{Y}_-$ at a point $\gamma'(t_0) \in T_+^n$. If $\gamma'|_{(t_0,1]}$ stays in $\overline{K}_+ \times T_+^n - T_+^n$, then B_2 intersects B_1 at $\gamma(t_0)$. So $\gamma'|_{(t_0,1]}$ reenters $\overline{Y}_-$, at another point $\gamma'(t_1) \in T_+^n$. Then $\gamma'|_{(t_0,t_1)} \subset \overline{K}_+ \times T_+^n$ and its projection onto the $\overline{K}_+$-coordinate is a nontrivial loop in $\overline{K}_+$. Since projections do not increase distance, it follows that the length of γ is also at least 2ϵ, guaranteeing that $d(B_1, B_2) \geqslant \epsilon$. $\square$

As a corollary, we get

Theorem A′. *The nerve $\mathcal{N}$ of blocks is a tree.*

The following is also clear.

Lemma 6.5. *Given a block B, ΠB, as defined in the previous section, is precisely the set of poles of blocks neighboring B.*

Theorem B′. *Let B_0 and B_1 be blocks, and D be the distance between the corresponding vertices in N. Then:*

(1) *If $D = 1$, then $\partial B_0 \cap \partial B_1 = \partial W$ where W is the wall $B_0 \cap B_1$.*

(2) *If $D = 2$, then $\partial B_0 \cap \partial B_1 = PB_{1/2}$ where $B_{1/2}$ intersects B_0 and B_1.*
(3) *If $D > 2$, then $\partial B_0 \cap \partial B_1 = \emptyset$.*

Proof. (1) If $D = 1$, then $B_0 \cap B_1$ is a wall W. That $\partial W \subset \partial B_0 \cap \partial B_1$ is obvious. To see the reverse inclusion, suppose we have asymptotic geodesic rays $\alpha_0 \subset B_0$ and $\alpha_1 \subset B_1$. Every geodesic from α_0 to α_1 intersects the wall W (Lemma 3.3). Thus we can get a sequence of points in W which remain asymptotic to α_0 and α_1.

(2) If $D = 2$, then there is one vertex between v_{B_0} and v_{B_1}; call it $v_{B_{1/2}}$. We will show that

$$PB_{1/2} \subset \partial B_0 \cap \partial B_1 \subset \partial W_0 \cap \partial W_1 \subset PB_{1/2}$$

where $W_i = B_{1/2} \cap B_i$ for $i = 0, 1$. The first inclusion follows naturally from the fact that geodesic rays in $B_{1/2}$ which go to poles are precisely those whose projections onto the hedge coordinate of $B_{1/2}$ are constant, and can be therefore be constructed easily in W_0 and W_1. The second inclusion follows by the same argument as in (1). For the third inclusion, suppose $\alpha_0 \subset W_0$ and $\alpha_1 \subset W_1$ are asymptotic geodesic rays, and let $\bar{\alpha}_0$ and $\bar{\alpha}_1$ be their projections onto the hedge coordinate of $B_{1/2}$. Let L_0 and L_1 be the leaves containing the images of these two maps. Since L_0 and L_1 are disjoint, every geodesic from $\bar{\alpha}_0$ to $\bar{\alpha}_1$ leaves L_0 at the same gluing point and enters L_1 at the same gluing point. Hence, the only way for α_0 and α_1 to be asymptotic is if $\bar{\alpha}_0$ and $\bar{\alpha}_1$ are both constant. Therefore α_0 and α_1 go to a pole of $B_{1/2}$.

Finally, we show (3) by contradiction: Suppose $\zeta \in \partial B_0 \cap \partial B_1$, and write $\text{Itin}[B_0, B_1] = [\bar{B}_1, \ldots, \bar{B}_n]$ where $n = D + 1$ by hypothesis. By the same argument as in (1), we actually have that $\zeta \in \partial B_i$ for every $1 \leqslant i \leqslant n$. By (2), then, it follows that $\zeta \in PB_i$ for every $1 < i < n$. But $PB_2 \cap PB_3 = \emptyset$, because $\angle_{Tits}(PB_2, PB_3) = \theta$, giving us a contradiction! □

6.2. Poles and n-vertices

In [5], a boundary point ζ is called a *vertex* if it has a local path component π and a homeomorphism from π to the open cone on the cantor set taking ζ to the cone point. An appropriate analogue in this context is the following: A point $\zeta \in \partial X$ is called a *vertex* if the local homology of ∂X at ζ is uncountably generated in some dimension. If n is the smallest dimension in which this local homology is uncountably generated, then we say that ζ is an *n-vertex*. The goal of this section is to distinguish topologically which vertices in RX are poles. A key tool is the following:

Theorem C′. *Let B be a block and $\zeta \in \partial B$ not be a pole of any neighboring block. Then ζ has a local path component which stays in ∂B.*

The proof of this theorem is the same as the proof of Theorem C [5, Lemma 4] with the following observation. The topological frontier of a left block is a subcollection of path components of $p^{-1}(T_+^n)$, and the topological frontier of a right block is a subcollection of path components of $p^{-1}(T_-^m)$; these path components are isometric copies of E_+^n and E_-^m and are the appropriate replacements for the "singular geodesics" given the original proof.

Recall that $1 \leqslant m \leqslant n$.

Lemma 6.6.

(1) *If $m = n$, then m-vertices in RX are poles.*
(2) *If $m < n$, then m-vertices in RX are poles of left blocks.*
(3) *If $m < n$, then n-vertices in right block boundaries are poles of right blocks.*

Proof. Choose any $\zeta \in RX$. Recall that in Section 3, we showed that the collection of rational points RX of ∂X is the same as the union of block boundaries. So there is a block B such that $\zeta \in \partial B - \Pi B$. If B is a left block, then let $k = m$, and if B is a right block, let $k = n$. Applying Theorem C′, Lemma 4.1, and Corollary 5.12, we know that ζ is a k-vertex iff ζ is a pole (of B). □

Now, it is conceivable in the case where $m < n$ that some n-vertices in left block boundaries are not poles. Here is the last resort for dealing with this situation.

Lemma 6.7. *Assume $m < n$ and let π be a path component in the set of n-vertices in RX. Then:*

(1) *If π is compact, then every point of π is a pole.*
(2) *If π is not compact, then no point of π is a pole.*

Proof. We begin by proving that if one point of π is a pole, then every other point of π is a pole as well: Suppose $\alpha : [0, 1] \to \pi$ is a path such that $\alpha(0)$ is a pole and $\alpha(1)$ is not. Since the collection of poles is closed, we may write

$$\bar{t} = \max\{0 \leqslant t \leqslant 1 \mid \alpha(t) \text{ is a pole}\};$$

say $\alpha(\bar{t}) \in PB_+$. By Theorem C′, there is an $s \in (\bar{t}, 1)$ so that $\alpha(\bar{t}, s) \subset \partial B_+$ and contains no n-vertices, giving us a contradiction.

We now prove (1) and (2) by contrapositive; if π contains all poles, then π is a sphere, giving us (2). For (1), assume π contains no poles; then

$$\pi \subset \partial B - PB$$
$$= PB * \partial H(B) - PB$$
$$\approx \partial H \times B^m$$

where B is a left block, B^m is the open m-ball, and H is the hedge factor of B. If $(\zeta, z) \in \pi$, then since every other point of $\{\zeta\} \times B^m$ has a local path component homeomorphic to that of (ζ, z), we know that in fact $\{\zeta\} \times B^m \subset \pi$. But the frontier of $\{\zeta\} \times B^m$ in ∂B is PB, which is disjoint from π. This proves (1). □

6.3. Safe path components

In their original work, Croke and Kleiner did not prove that RX was a path component. However, in order for a path starting in RX to leave it, it has to pass through infinitely many block boundaries on its way out. The same is true here. In this section, we define a *safe path* to be a path $\alpha(t)$ in ∂X which passes through m-vertices for only finitely many times t. Recall that under our assumption that $m \leqslant n$, it follows from Corollary 5.12 that m-vertices are poles. If $m < n$, then these are poles of left blocks.

Theorem D′. *RX is the unique dense safe path component of ∂X.*

Proof. We begin by showing that RX is safe path connected; choose $\zeta, \eta \in RX$, say $\zeta \in \partial B_0$ and $\eta \in \partial B_0'$. We will construct a safe path between ζ and η in RX by induction on the length of $\mathrm{Itin}[B_0, B_0'] = [B_1, \ldots, B_n]$. Assume $n = 1$ (that is, that $B_0 = B_0'$), and let $[\zeta_P, \zeta_H]$ and $[\eta_P, \eta_H]$ be join arcs of $\partial B_0 = PB_0 * \partial H_0$ containing ζ and η. Note that each of these join arcs contains at most two poles, one of which is a pole of B_0. For, if one contained two poles of a neighboring block, then we would not have $E_-^m \cap E_+^n = \{0\}$. Take α to be the path from ζ to η

$$[\zeta, \zeta_P] \cup [\zeta_P, \eta_H] \cup [\eta_H, \eta].$$

Then α passes through only finitely many left poles, and hence is safe. In general, if $n > 1$, then choose a point $\zeta' \in \partial B_{n-1} \cap \partial B_n$ and concatenate a safe path from ζ to ζ' in $\partial B_1 \cup \cdots \cup \partial B_{n-1}$ to a safe path from ζ' to η in ∂B_n.

So RX is safe path connected. The next step is to show that RX is a safe path component. We do this by contradiction: Suppose we have a safe path α which starts in RX but ends outside. Set $\bar{t} = \inf\{t \mid \alpha(t) \notin RX\}$. If $\alpha(\bar{t}) \in RX$, then $\alpha(\bar{t}) \in \partial B$ for some block B and $\alpha(\bar{t}, \bar{t} + \epsilon) \subset \partial B$ for some small $\epsilon > 0$. So $\alpha(\bar{t}) \notin RX$, but $\alpha[0, \bar{t}) \subset RX$. Let s be the last time at which α passes through an m-vertex, say $\alpha(s) \in PB$. Then for some small $\epsilon > 0$, $\alpha(s, s + \epsilon)$ is contained in a path component of $\partial B - PB$ of the form $(PB * V) - PB$ where V is a path component of the boundary of the hedge factor of B. If $\alpha(s, s + \epsilon)$ not contained in the boundary of a wall, then $\alpha(s, \bar{t}) \subset \partial B$. But $\alpha|_{(s,\bar{t})}$ cannot stay inside ∂B for *any* B, for if it did, then $\alpha(\bar{t}) \in \partial B$ as well, since block boundaries are closed! So $\alpha(s, s + \epsilon) \subset \partial W$ for some wall $W = B \cap B_+$, and $\alpha|_{(s,t)}$ enters ∂B_+. Since this path must also leave ∂B_+, it must pass through another pole, giving us a contradiction.

Finally, RX is dense by Lemma 3.6 and no other safe path component of ∂X can be dense, since the irrational map provided in Section 3 takes components of IX to points of the Cantor set. □

7. The proof of Theorem 1

The *watermark* of X is defined to be the unordered pair $\{\mathcal{W}(Y_-), \mathcal{W}(Y_+)\}$. Let X_1 and X_2 be two generalized Croke–Kleiner constructions for G from the same spaces and suppose $\phi : \partial X_1 \to \partial X_2$ is a homeomorphism. By Proposition 5.9, there are at least two possible watermarks. Therefore Theorem 1 will follow from the following proposition.

Proposition 7.1. *X_1 and X_2 have the homeomorphic watermarks.*

The following are immediate from Theorem D′, Corollary 5.12, and Lemma 6.7.

Lemma 7.2. $\phi(RX_1) = RX_2$.

Lemma 7.3. *ϕ takes poles to poles.*

Lemma 7.4. *Let B_1 be a block of X_1. Then there is a block B_2 of X_2 such that*:

(1) $\phi(PB_1) = PB_2$.
(2) $\phi(\Pi B_1) = \Pi B_2$.
(3) $\phi(\partial B_1) = \partial B_2$.

Proof. Choose $\zeta \in PB_1$ and let B_2 be such that $\phi(\zeta) \in PB_2$. If the dimension of PB_1 (or PB_2) is at least 1, then PB_1 is connected and since ϕ takes poles to poles, we get that $\phi(PB_1) = PB_2$. Assume PB_1 (and PB_2) both have dimension zero. Then they both consist of exactly two points and write $PB_2 = \{\zeta, \eta\}$. Choose $\zeta \in \partial H(B_1)$ such that $\{\zeta\}$ is a path component of $\partial H(B_1)$ and let $\alpha : [0, 1] \to l(\zeta)$ parameterize the longitude in such a way that $\alpha(0) = \zeta$ to $\alpha(1) = \eta$. If $\phi(\eta)$ is not a pole of B_2, then it must be a pole of a neighboring block and $RX_2 - \phi(PB_1)$ has two path components: $\phi(\alpha(0, 1))$ and $RX_2 - \phi(\alpha[0, 1])$. But $RX_1 - PB_1$ has an infinite number of path components, which gives us a contradiction. This shows that $\phi(PB_1) = PB_2$.

We get that $\phi(\Pi B_1) = \Pi B_2$ and $\phi(\partial B_1) = \partial B_2$ by the following argument: If $\alpha : [0, 1] \to RX_1$ is a path such that $\alpha(0) \in PB_1$, $\alpha(1)$ is a pole, and $\alpha(0, 1)$ contains no poles, then $\alpha \subset \partial B_1$ and $\alpha(1)$ is either a pole of B_1 or a pole of a neighboring block. $\quad \square$

Lemma 7.5. *Let B_1 be a block of X_1 and B_2 be a block of X_2 such that $\phi(PB_1) = PB_2$. If we have $\zeta_1 \in \partial H(B_1)$ such that $\{\zeta_1\}$ is a path component of $\partial H(B_1)$, then there is a $\zeta_2 \in \partial H(B_2)$ such that $\phi(l(\zeta_1)) = l(\zeta_2)$.*

Proof. Whenever $\{\zeta_i\}$ is a path component of the boundary of the hedge factor of B_i, then the longitude $l(\zeta_i)$ is a ball in ∂B_i whose frontier is precisely PB_i and whose interior is a path component of $\partial B_i - PB_i$. $\quad \square$

Combining these lemmas, we get Proposition 7.1.

Acknowledgements

References

[1] F.D. Ancel, C. Guilbault, J. Wilson, The Croke–Kleiner boundaries are cell-like equivalent, preprint.
[2] M. Bestvina, Local homology properties of boundaries of groups, Michigan Math. J. 43 (1996) 123–139.
[3] P.L. Bowers, K. Ruane, Boundaries of nonpositively curved groups of the form $G \times \mathbb{Z}^n$, Glasg. Math. J. 38 (1996) 177–189.
[4] M.R. Bridson, A. Haefliger, Metric Spaces of Non-Positive Curvature, Grundlehren Math. Wiss. (Fundamental Principles of Mathematical Sciences), vol. 319, Springer-Verlag, Berlin, 1999.
[5] C.B. Croke, B. Kleiner, Spaces with nonpositive curvature and their ideal boundaries, Topology 39 (2000) 549–556.
[6] R. Geoghegan, The shape of a group—connections between shape theory and the homology of groups, in: Geometric and Algebraic Topology, in: Banach Center Publ., vol. 18, PWN, Warsaw, 1986, pp. 271–280.
[7] A. Hatcher, Algebraic Topology, Cambridge University Press, Cambridge, 2002.
[8] T. Hosaka, A splitting theorem for CAT(0) spaces with the geodesic extension property, Tsukuba J. Math. 27 (2003) 289–293.
[9] G.C. Hruska, Geometric invariants of spaces with isolated flats, Topology 44 (2005) 441–458.
[10] G.C. Hruska, B. Kleiner, Hadamard spaces with isolated flats, Geom. Topol. 9 (2005) 1501–1538 (electronic), with an appendix by the authors and Mohamad Hindawi.
[11] C. Mooney, All CAT(0) boundaries of a group of the form $H \times K$ are CE equivalent, Fund. Math. 203 (2009) 97–106.
[12] C.P. Mooney, Examples of non-rigid CAT(0) groups from the category of knot groups, Algebr. Geom. Topol. 8 (2008) 1666–1689.
[13] P. Ontaneda, Cocompact CAT(0) spaces are almost geodesically complete, Topology 44 (2005) 47–62.
[14] K.E. Ruane, Boundaries of CAT(0) groups of the form $\Gamma = G \times H$, Topology Appl. 92 (1999) 131–151.
[15] J.M. Wilson, A CAT(0) group with uncountably many distinct boundaries, J. Group Theory 8 (2005) 229–238.