

Detecting sphere boundaries of hyperbolic groups

BENJAMIN BEEKER

NIR LAZAROVICH

We show that a one-ended simply connected at infinity hyperbolic group G with enough codimension-1 surface subgroups has $\partial G \cong \mathbb{S}^2$. By work of Markovic (2013), our result gives a new characterization of virtually fundamental groups of hyperbolic 3-manifolds.

20F65, 20F67, 20H10

1 Introduction

We recall the following well-known conjecture.

Conjecture 1.1 (Cannon’s conjecture [9]) *Let G be a hyperbolic group. If $\partial G \cong \mathbb{S}^2$ then G acts geometrically on the hyperbolic space \mathbb{H}^3 .*

Markovic [17] described the following criterion under which the conjecture is true.

Theorem 1.2 *Let G be a hyperbolic group that acts faithfully on its boundary $\partial G = \mathbb{S}^2$ and contains enough quasiconvex surface subgroups. Then G acts geometrically on the hyperbolic space \mathbb{H}^3 .*

Kahn and Markovic [14] showed that the fundamental group of a hyperbolic 3-manifold contains enough quasiconvex surface subgroups. This shows that Markovic’s criterion is also necessary. More about the history of Cannon’s conjecture and the work preceding Markovic’s criterion can be found in a survey about boundaries of hyperbolic groups by Kapovich and Benakli [15].

In this paper we prove that it is possible to replace the assumption that the boundary at infinity is homeomorphic to \mathbb{S}^2 by the assumption of vanishing of the first cohomology of G at infinity (see Section 2.3 for the definitions). In other words, we have the following result.

Theorem 1.3 (main result) *Let G be a one-ended hyperbolic group. Assume that G has vanishing first cohomology over $\mathbb{Z}/2$ at infinity and that G contains enough quasiconvex codimension-1 surface subgroups. Then $\partial G \cong \mathbb{S}^2$.*

Combining this result with Markovic [17] and Kahn and Markovic [14], we get the following.

Corollary 1.4 *Let G be a hyperbolic group. The following are equivalent:*

- (1) G acts geometrically on \mathbb{H}^3 .
- (2) $\partial G \cong \mathbb{S}^2$, and G contains enough quasiconvex surface subgroups.
- (3) G is one-ended, has vanishing first cohomology over $\mathbb{Z}/2$ at infinity and contains enough quasiconvex codimension-1 surface subgroups.

The main tool we use is the Kline sphere characterization, which was proven by Bing [5].

Theorem 1.5 (Kline sphere characterization theorem) *Let M be a topological space which is*

- (1) *nondegenerate, metrizable, compact, connected, locally connected,*
- (2) *not separated by any pair of points, and*
- (3) *separated by any Jordan curve.*

Then $M \cong \mathbb{S}^2$.

The outline of the proof of the main result, and of the paper, is as follows. In Section 2, we provide the necessary preliminaries for the proof. We end Section 2 with a summary of the additional assumptions one can make of the group G . In Lemma 3.2 in Section 3, we prove that no pair of points separates the boundary of such a group. In Sections 4–10, we prove that any Jordan curve separates, using ideas from Jordan’s original proof of the Jordan curve theorem which we outline in Section 4 (see Hales [13] for details).

Bestvina and Mess [4] studied the topological properties of boundaries of hyperbolic groups, and in particular showed that if G is hyperbolic and R is a ring, then $H^i(G, RG) \simeq \check{H}^{i-1}(\partial G, R)$ (Čech reduced cohomology). They then use their results to show that if moreover G is the fundamental group of a closed 3–manifold K , then $\partial G \cong \mathbb{S}^2$. In their proof they use a similar characterization of the sphere due to Zippin [22]. However, the proof that every Jordan curve on ∂G separates relies on extending it to a separating surface in the 3–manifold \tilde{K} . This step is impossible in the setting of our work, and a major part of this paper is devoted to replacing this step with a cube complex analogue; see Section 4 for details.




Figure 1: The torus link of the counterexample

We would like to mention the following question, which was raised by the referee. Does Theorem 1.3 remain true when replacing the assumption of enough quasiconvex codimension-1 surface subgroups by the assumption that any two points in ∂G are separated by some Jordan curve? We remark that our techniques rely heavily on the existence of a CAT(0) cube complex on which G acts, which is constructed using the quasiconvex codimension-1 surface subgroups, and thus do not apply to this more general setting.

We finish the introduction with an example of a one-ended hyperbolic group that has enough quasiconvex codimension-1 surface subgroups and whose boundary is not a sphere.

Let L be the flag complex triangulation of the 2–dimensional torus consisting of 96 triangles illustrated in Figure 1, where the edges of the hexagon labeled by a , b and c are glued accordingly to obtain a torus. Let G be the right-angled Coxeter group associated with L . That is, G is the group given by the presentation

$$G = \langle s \in L^{(0)} \mid s^2 = 1 \text{ for all } s \in L \text{ and } [s, s'] = 1 \text{ for all } \{s, s'\} \in L^{(1)} \rangle$$

The group G acts properly and cocompactly on the Davis complex X associated with L , which happens to be the unique CAT(0) cube complex whose link is isomorphic to L at each vertex (see Lazarovich [16]).

The group G is hyperbolic since L has no isometrically embedded geodesic of length 2π (when considered with the spherical metric), and it is one-ended since the link L does not have a separating simplex. The hyperplane stabilizers are isomorphic to the right-angled Coxeter groups associated with the link of a vertex in L , which are 6–cycle graphs in our case. Thus, the hyperplane stabilizers are Fuchsian groups, which implies that G has enough codimension-1 surface subgroups.

Fix some vertex $x \in X$, and let p be the projection map from ∂X to its link, $\text{Link}(x, X)$, which assigns to each boundary point ξ the direction in $\text{Link}(x, X)$ of the unique geodesic that connects x to ξ . Since X has extendable geodesics, one can lift any curve on $\text{Link}(x)$ to the boundary ∂X . Therefore, the induced map $\pi_1(\partial X) \rightarrow \pi_1(\text{Link}(x, X)) \simeq \mathbb{Z}^2$ is onto, and the boundary ∂G is not homeomorphic to S^2 .

Acknowledgements The authors would like to thank Michah Sageev for his support and helpful comments. We also thank Antoine Clais and Jason Manning for finding a mistake in a previous version of this paper. Beeker is supported by ISF grant 1941/14. Lazarovich acknowledges the support received by the ETH Zurich Postdoctoral Fellowship Program and the Marie Curie Actions for People COFUND Program.

2 Preliminaries

We begin by a survey of definition and results concerning CAT(0) cube complexes and quasiconvex subgroups of hyperbolic groups. For a more complete survey, see Sageev [19].

2.1 CAT(0) cube complexes and hyperplanes

Definition 2.1 (CAT(0) cube complexes) A *cube complex* is a complex made by gluing unit Euclidean cubes (of varying dimension) along their faces using isometries. A cube complex is CAT(0) if it is CAT(0) with respect to the quotient metric induced by endowing each cube with the Euclidean metric (see [8]).

As was first observed by Sageev [18], CAT(0) cube complexes naturally carry a combinatorial structure given by the associated hyperplanes and halfspaces. We now recall their definition and properties.

Definition 2.2 (hyperplanes) Let X be a CAT(0) cube complex. The equivalence relation on the edges of X generated by $e \sim e'$ when e and e' are parallel edges in a square of X is called the parallelism relation. The equivalence classes of edges under the parallelism relation are the *combinatorial hyperplanes* of X . The convex hull of the midpoints of the edges of a combinatorial hyperplane is called a *hyperplane*. We denote the set of hyperplanes in X by $\hat{\mathcal{H}} = \hat{\mathcal{H}}(X)$.

The main features of hyperplanes are summarized in the following.

Proposition 2.3 *Let X be a CAT(0) cube complex. Then every hyperplane $\hat{h} \in \hat{\mathcal{H}}$ is naturally a CAT(0) cube complex of codimension-1 in X , and $X \setminus \hat{h}$ has exactly two components.*

Definition 2.4 The components of $X \setminus \hat{h}$ are the *halfspaces* of X associated to (or bounded by) \hat{h} . The set of halfspaces is denoted by $\mathcal{H} = \mathcal{H}(X)$. There is a natural map $\hat{\cdot}: \mathcal{H} \rightarrow \hat{\mathcal{H}}$, which maps each halfspace to its bounding hyperplane. The set of halfspaces also carries a natural complementation involution $*$: $\mathcal{H} \rightarrow \mathcal{H}$, which maps each halfspace h to the other component h^* of $X \setminus \hat{h}$.

2.2 Cubulating hyperbolic groups

Recall the following definitions about quasiconvex subgroups of Gromov hyperbolic groups.

Definition 2.5 A quasiconvex subgroup H of a hyperbolic group G is *codimension-1* if G/H has more than one end. The group G has *enough codimension-1 subgroups* if every two distinct points in ∂G can be separated by the limit set of a codimension-1 quasiconvex subgroup.

By results of Sageev [18], Gitik, Mitra, Rips and Sageev [11] and Bergeron and Wise [3] we have the following.

Theorem 2.6 *Let G be a hyperbolic group with enough codimension-1 subgroups, then G acts properly cocompactly on a finite-dimensional CAT(0) cube complex whose hyperplane stabilizers belong to the family of codimension-1 subgroups.*

This theorem is the starting point for our proof: since the group G in Theorem 1.3 is assumed to have enough codimension-1 surfaces subgroups, it therefore acts properly and cocompactly on a finite-dimensional CAT(0) cube complex X with surface group hyperplane stabilizers. The proof of the main theorem is then based on understanding how the limit sets of hyperplanes interact and how they determine the topology of the boundary of G .

Recall that any quasiconvex subgroup H in G is itself a hyperbolic group and there is a well-defined homeomorphic embedding $\partial H \rightarrow \partial G$ whose image is the limit set of H in ∂G . In particular, in our situation, the stabilizer of each hyperplane \hat{h} in X acts properly cocompactly on the hyperplane and thus is a quasiconvex surface subgroup.

Its boundary is therefore a circle, which embeds as the limit set of \widehat{h} in $\partial G = \partial X$. We denote this limit set by $\partial\widehat{h}$.

However, for halfspaces we have to distinguish between two kinds of limit sets: the *closed limit set* ∂h and the *open limit set* $\partial^\circ h$. The former is simply the limit set of h as a subset of X , while the latter is the set of all geodesic rays in h which do not stay at a bounded distance from \widehat{h} . Equivalently, we have $\partial h = \partial\widehat{h} \sqcup \partial^\circ h$. As their names suggest the open (resp. closed) limit sets are indeed open (resp. closed) subsets of ∂G . Moreover, we have the following.

Lemma 2.7 *The open limit sets of halfspaces form a basis for the topology of $\partial G = \partial X$.*

Proof Let x_0 be a fixed vertex in X . Identify ∂X with the visual Gromov boundary from x_0 . Let $\xi \in \partial X$, and let γ be a geodesic ray such that $\gamma(\infty) = \xi$ and $\gamma(0) = x_0$. Let \mathcal{H}_γ be the set of halfspaces h such that $\gamma \setminus h$ is nonempty and bounded. The set \mathcal{H}_γ is infinite (since otherwise γ traverses only finitely many hyperplanes, which would contradict that γ is unbounded) and does not contain pairs of disjoint halfspaces. Since X is finite-dimensional, \mathcal{H}_γ must contain an infinite descending chain of halfspaces $h_1 \supset h_2 \supset \dots$. Moreover, $d(x_0, h_n) \rightarrow \infty$, and since halfspaces are convex and X is δ -hyperbolic, it follows that $\text{diam}(\partial h_n) \rightarrow 0$ (with respect to a visual metric on the boundary). Moreover, since the h_n form a descending chain of halfspaces that cross γ , it follows that γ does not remain within a bounded distance from the bounding hyperplanes \widehat{h}_n for all n . Thus, $\xi \in \partial^\circ h_n$ and $\partial^\circ h_n$ forms a local basis at ξ . \square

Recall from Caprace and Sageev [10] that a cube complex is *essential* if there is no halfspace which is at bounded distance from its bounding hyperplane. The rank rigidity results in [10] imply that if G is a hyperbolic group that acts properly and cocompactly on a CAT(0) cube complex X , then there is a convex G -invariant subcomplex $Y \subseteq X$ which is essential, and the open limit set of every halfspace is nonempty. For that reason, we assume from now on that X is essential.

2.3 Properties of groups at infinity

In this section we define topological properties of spaces and groups at infinity. We begin with the following general definition.

Definition 2.8 Let F be a contravariant (resp. covariant) functor from the category of topological spaces **Top** to the category of groups **Grp**. Then, let F_∞ (resp. F^∞) be the functor from the category of noncompact topological spaces (with proper continuous maps as morphisms) that assigns to a space X the direct limit (resp. inverse limit) of the directed set $F(X \setminus K)$, where K ranges over the compact subsets in X .

Applying the above definition to the contravariant functor $H^1(\cdot; R)$, we say that X has *vanishing first cohomology over a ring R at infinity* if for every compact set $K \subset X$ and every 1-cocycle α on $X \setminus K$ there exists a bigger compact set K' such that α restricted to $X \setminus K'$ is a coboundary.

Similar definitions can be defined by applying the above definition to π_1 , $H_n(\cdot; R)$ and $H^n(\cdot; R)$. We recall in particular the following well-studied notion.

A topological space X is said to be *simply connected at infinity* if for every compact set K there exists a compact set $K' \supseteq K$ such that any loop in $X \setminus K'$ is null homotopic in $X \setminus K$, ie if the map $\pi_1(X \setminus K') \rightarrow \pi_1(X \setminus K)$ is trivial.

We remark that any space which is simply connected at infinity has, in particular, vanishing first cohomology over $\mathbb{Z}/2$.

Definition 2.9 Let G be a group of type F (ie has a compact $K(G, 1)$), and let X be the universal cover of a compact $K(G, 1)$ of G . We say that G has *vanishing first cohomology over a ring R at infinity* (resp. is *simply connected at infinity*) if X has vanishing first cohomology over a ring R at infinity (resp. is simply connected at infinity).

More generally, if G has a finite-index subgroup H of type F, then we say that G has *vanishing first cohomology over a ring R at infinity* (resp. is *simply connected at infinity*) if H has vanishing first cohomology over a ring R at infinity (resp. is simply connected at infinity).

Brick [7] showed that being simply connected at infinity is a property of finitely presented groups, which is a quasi-isometric invariant.

Lemma 2.10 *The notions defined above do not depend on the compact $K(G, 1)$ and the choice of a finite-index subgroup.*

Proof Let X and Y be two compact $K(G, 1)$. Then there exists a G -equivariant homotopy equivalence $f: \tilde{X} \rightarrow \tilde{Y}$ which is proper. We prove it for the n^{th} cohomology

functor. Since this map is proper, it induces a map

$$H^n(Y \setminus K) \rightarrow H^n(X \setminus f^{-1}(K)) \rightarrow H_\infty^n(X),$$

and thus, a map $H_\infty^n(f): H_\infty^n(Y) \rightarrow H_\infty^n(X)$. Similarly, the G -equivariant homotopy inverse map $g: \tilde{Y} \rightarrow \tilde{X}$ induces a map $H_\infty^n(g): H_\infty^n(X) \rightarrow H_\infty^n(Y)$. The composition $f \circ g$ induces a map $H_\infty^n(f) \circ H_\infty^n(g): H_\infty^n(X) \rightarrow H_\infty^n(X)$. Since $f \circ g$ is homotopic to the identity by an homotopy of G -equivariant proper maps, it follows that the induced map in H_∞^n is the identity map. Similarly, $H_\infty^n(g) \circ H_\infty^n(f)$ is the identity map.

Let H_1 and H_2 be two finite-index type-F subgroups of G . Then $H_1 \cap H_2$ is also a finite-index type-F subgroup of G . Therefore, one can assume $H_1 \leq H_2 \leq G$. But in this case any compact $K(G, 1)$ for H_2 has a finite cover which is a $K(G, 1)$ of H_1 . This completes the proof, since the defined notions only depend on the universal cover of the finite $K(G, 1)$. \square

Since every torsion-free hyperbolic group G is of type F (for example, its Rips complex for sufficiently large parameter r is the universal cover of a compact $K(G, 1)$), the notions defined above can be defined for every torsion-free hyperbolic group. It is a very well-known question whether hyperbolic groups are virtually torsion-free. However, in the context of Theorem 1.3, G is a hyperbolic group that acts properly cocompactly on a CAT(0) cube complex. It follows from works of Haglund and Wise [12] and Agol [1] that G is virtually torsion-free. Therefore, the assumption that G has vanishing first cohomology over $\mathbb{Z}/2$ at infinity is well-defined for G . By replacing G by its finite-index torsion-free subgroup, it is enough to prove the theorem for torsion-free groups. Hence, in the remainder of the paper we assume that G is torsion-free.

We remark that, in the context of Theorem 1.3, since G is torsion-free and acts properly and cocompactly on the CAT(0) cube complex X it follows that X is the universal cover of a compact $K(G, 1)$ (namely X/G), and thus, by assumption, X has vanishing first cohomology over $\mathbb{Z}/2$ at infinity.

2.4 Summary of preliminaries

In the previous subsections we have seen that under the assumptions of Theorem 1.3, one can make further assumptions on G . In this subsection we collect the assumptions we made on G that will be used in the remainder of the paper:

- The group G is hyperbolic, one-ended and torsion-free.

- There exists a finite-dimensional CAT(0) cube complex X such that:
 - The group G acts freely and cocompactly on X .
 - The hyperplane stabilizers are surface subgroups.
 - The cube complex X is essential, and in particular, open limit sets of halfspaces are nonempty.
 - The cube complex X has vanishing first cohomology at infinity.

3 Connectivity and nonseparation by a pair of points

In this section we prove that no pair of points can separate the boundary of a group that satisfies the assumption of the main theorem (Theorem 1.3).

Recall the following result of Bowditch [6].

Theorem 3.1 *Let G be a one-ended word-hyperbolic group such that $\partial G \neq \mathbb{S}^1$. Then:*

- (1) *The boundary is path connected and locally path connected and has no global cutpoints.*
- (2) *The following are equivalent:*
 - (a) *The group G does not split essentially over a two-ended subgroup*
 - (b) *The boundary ∂G has no local cutpoint.*
 - (c) *The boundary ∂G is not separated by a pair of points.*
 - (d) *The boundary ∂G is not separated by a finite set of points.*

The following lemma and corollary apply to the setting of Theorem 1.3, since we assume that G has enough codimension-1 quasiconvex surface subgroups. However, we chose to phrase them in a more general setting.

Lemma 3.2 *Let G be a one-ended hyperbolic group that contains enough codimension-1 quasiconvex one-ended subgroups. Then ∂G is not separated by any pair of points.*

Proof Let $\xi, \zeta \in \partial G$ be distinct points. Let \widehat{h} be a hyperplane such that $\xi \in \partial^\circ h$ and $\zeta \in \partial^\circ h^*$. Let $A = \partial h$ and $B = \partial h^*$. The closed sets A and B satisfy $(A \setminus \{\xi\}) \cup B = \partial G \setminus \{\xi\}$ and $(A \setminus \{\xi\}) \cap B = \partial \widehat{h}$. By Theorem 3.1(1) we know that $\partial G \setminus \{\xi\}$ is connected, and thus each of the two halves $A \setminus \{\xi\}$ and B is connected (since we assumed $\partial \widehat{h}$ is connected). Similarly, A and $B \setminus \{\zeta\}$ are connected. Finally, $\partial G \setminus \{\xi, \zeta\} = (A \setminus \{\xi\}) \cup (B \setminus \{\zeta\})$ is connected as a union of connected intersecting sets. \square

By Theorem 3.1(2) we have the following.

Corollary 3.3 *Let G be a one-ended hyperbolic group that contains enough codimension-1 quasiconvex one-ended subgroups. Then ∂G does not have any local cutpoints and is not separated by any finite number of points.* \square

The following lemma is a direct corollary of the above, and will be used later on in Lemma 9.3.

Lemma 3.4 (no blob lemma) *Let M be a compact, path connected, locally path connected metric space with no cutpoints. Then for all $\epsilon > 0$ there exists $\delta > 0$ such that for all $x \in M$ and for all $A \subset B(x, \delta)$ there is at most one component of $M \setminus A$ which is not contained in $B(x, \epsilon)$.*

Proof Assume for contradiction that there is $\epsilon > 0$ such that for all $n \in \mathbb{N}$ there is a point $x_n \in M$ and a subset $A_n \subset B(x_n, \frac{1}{n})$ and two distinct components B_n and C_n of $M \setminus A_n$ which have points b_n and c_n , respectively, outside $B(x, \epsilon)$. By passing to a subsequence if necessary, we may assume that the sequences x_n , b_n and c_n converge to x , b and c , respectively. Since $d(b_n, x_n) > \epsilon$ (resp. $d(c_n, x_n) > \epsilon$) we deduce that $d(b, x) \geq \epsilon$ (resp. $d(c, x) \geq \epsilon$).

The space M has no cutpoint, therefore we can connect b and c with a path γ that does not pass through x . Since $A_n \rightarrow x$, for n big enough A_n is disjoint from γ . By local path connectivity, for n big enough b_n and b (resp. c and c_n) can be connected by a short path γ_b (resp. γ_c) that avoids A_n . The concatenated path $\gamma_b * \gamma * \gamma_c$ from b_n to c_n avoids A_n and thus contradicts the assumption that B_n and C_n are distinct components of $M \setminus A_n$. \square

4 Outline of the proof of Jordan's theorem

The aim of this section is to provide a short outline of the proof that any Jordan curve on ∂G separates. From this section on, let G and X be a group and a CAT(0) cube complex satisfying the assumptions made in Section 2.4.

We first recall Jordan's original proof of the Jordan curve theorem. Let P be a polygonal curve. One can associate a parity function to P which, for a point $x \in \mathbb{R}^2 \setminus P$ and a generic ray l starting at x , counts the number of intersections of l with P mod 2. It is easy to see that this parity function is constant on the connected components of $\mathbb{R}^2 \setminus P$. Now, for a Jordan curve J , one approximates J with polygonal curves P_n , and proves that the sequence of the parity functions of the polygonal curves P_n has a

limit which is constant on connected components of $\mathbb{R}^2 \setminus J$ (and does not depend on the choice of approximating polygonal curves). The final crucial step is to show that this parity function indeed obtains two values, which shows that $\mathbb{R}^2 \setminus J$ has at least two components. A more subtle point in the proof of Jordan's theorem is the proof that each of the two parity regions is connected, but luckily we will not need to prove this, since we are only interested in showing that the Jordan curve separates. The details of Jordan's proof of his theorem can be found in a paper by Hales [13].

The main idea in our proof is to replace polygonal path approximations with "piecewise hyperplane paths" (or "PH paths"). Those are paths in ∂X which are piecewise subsegments of limit sets of hyperplanes (which are homeomorphic to circles). In Section 8 we give a precise definition and show that every Jordan curve can be approximated by PH curves.

In order to obtain a parity function we first approximate the PH paths by "grids". Roughly speaking a grid is a collection of hyperplanes which are connected along "connectors". Precise definitions and examples are given in Sections 5.1 and 5.2. In Section 6 we describe how, given an (oriented) grid, one can assign a 1-cocycle (over $\mathbb{Z}/2$) which is defined in X outside a large enough ball. Using the vanishing of the first cohomology over $\mathbb{Z}/2$ at infinity, one can find a 0-chain defined outside a bigger ball in X whose coboundary is this 1-cocycle. The 0-chain extends to a map which is defined on the boundary except at the grid. We call it "the parity function of the grid". Its construction and properties are discussed in Section 7.

In Section 9, we show how one obtains a parity function of a PH curve by taking the limit of the parity functions of approximating grids. We then define the parity function for a Jordan curve, by taking the limit of the parities of a sequence of approximating PH curves.

The crucial part, as in Jordan's proof, is to show that the limiting parity obtains two values. This is done in the final section, Section 10, by analyzing how (certain) PH curves intersect a given hyperplane. This is strongly based on Proposition 7.7, which analyzes the ways in which two hyperplanes can intersect.

5 Intersections of hyperplanes

Since our proof involves studying intersections of hyperplanes which have surface group stabilizers, we begin by recalling the following well-known elementary fact about surface groups.

Fact *Finitely generated subgroups of surface groups are one of the following: trivial, cyclic, free or of finite index.*

Recall that the intersection of two quasiconvex subgroup of a hyperbolic group is quasiconvex [20], and the limit set of the intersection is the intersection of limit sets of the two subgroups [11]. Since our hyperplane stabilizers are quasiconvex surface subgroups, any two intersecting hyperplane stabilizers do so on either a trivial, cyclic, free or finite-index quasiconvex subgroup. This means that the intersection of their limit sets is either empty, a pair of points, a Cantor set or the whole circle.

5.1 Connectors

Let \hat{h} and \hat{k} be two intersecting hyperplanes. A *connector* is a pair $(C, \{\hat{h}, \hat{k}\})$, which by abuse of notation we will simply denote by C , where C is a nonempty clopen proper subset $\emptyset \neq C \subsetneq \partial\hat{h} \cap \partial\hat{k}$. In particular, $\partial\hat{h} \cap \partial\hat{k}$ is not the whole circle. We say that the connector C is *supported on the hyperplanes* \hat{h} and \hat{k} .

Connectors $(C_1, \{\hat{h}_1, \hat{k}_1\})$ and $(C_2, \{\hat{h}_2, \hat{k}_2\})$ are *disjoint* if either $\{\hat{h}_1, \hat{k}_1\} \cap \{\hat{h}_2, \hat{k}_2\} = \emptyset$ or $C_1 \cap C_2 = \emptyset$.

Let I_1, I_2, \dots, I_n be disjoint intervals in $\partial\hat{h}$ whose endpoints are disjoint from $\partial\hat{h} \cap \partial\hat{k}$ and such that $C = \partial\hat{h} \cap \partial\hat{k} \cap \bigcup_j I_j$. To each interval I_j we assign the number $\text{type}_{\hat{h}}(I_j) = 1 \in \mathbb{Z}/2$ if the two endpoints of I_j are in different sides of \hat{k} , and $\text{type}_{\hat{h}}(I_j) = 0$ otherwise.

Next, we extend this definition to the connector C by assigning its type in \hat{h} to be

$$\text{type}_{\hat{h}}(C) = \sum_{j=1}^n \text{type}_{\hat{h}}(I_j) \in \mathbb{Z}/2.$$

It is easy to verify that this definition depends only on C , \hat{h} and \hat{k} and does not depend on the choice of the open cover by intervals I_1, \dots, I_n (by passing to a common subdivision of the two open covers). Notice that any nonempty clopen subset of a connector is a connector and that if \mathcal{P} is a clopen partition of a connector C , then $\text{type}_{\hat{h}}(C) = \sum_{D \in \mathcal{P}} \text{type}_{\hat{h}}(D)$.

Finally, we define the type of C by $\text{type}(C, \hat{h}, \hat{k}) = (\text{type}_{\hat{h}}(C), \text{type}_{\hat{k}}(C))$. We will later show in Proposition 7.7 that the only possible types, under our assumptions, are $(0, 0)$ and $(1, 1)$.

5.2 Grids and oriented grids

Let \widehat{h} be a hyperplane and let $\widehat{k}_1, \dots, \widehat{k}_n$ be a collection of (not necessarily distinct) hyperplanes that intersect \widehat{h} . Let $C_{\widehat{h}} = \{C_1, \dots, C_n\}$ be a set of disjoint connectors such that each C_i is supported on \widehat{h} and \widehat{k}_i . The choice of $C_{\widehat{h}}$ is *admissible* if $\sum_i \text{type}_{\widehat{h}}(C_i) = 0$.

Remark 5.1 The following examples show how this definition behaves for $n = 0, 1, 2$:

- If $n = 0$, then $C_{\widehat{h}} = \emptyset$, and it is admissible.
- If $n = 1$, then $C_{\widehat{h}} = \{C_1\}$ is admissible if and only if $\text{type}_{\widehat{h}}(C_1) = 0$.
- If $n = 2$, then $C_{\widehat{h}} = \{C_1, C_2\}$ is admissible if and only if $\text{type}_{\widehat{h}}(C_1) = \text{type}_{\widehat{h}}(C_2)$.

We comment at this point that admissible sets of connectors correspond to trivial classes in the first cohomology at infinity of \widehat{h} over $\mathbb{Z}/2$ (which is isomorphic to $\mathbb{Z}/2$), and in fact they are the coboundary (at infinity) of exactly two 0-chains. We call an extension of such a 0-chain to the boundary an “orientation” for \widehat{h} . We make this comment explicit, avoiding the use of cohomology at infinity, in the following definition and claim.

Definition 5.2 Let \widehat{h} be a hyperplane and let $C_{\widehat{h}}$ be an admissible collection of disjoint connectors on \widehat{h} . A function $\alpha_{\partial\widehat{h}}: \partial\widehat{h} \setminus \bigcup C_{\widehat{h}} \rightarrow \mathbb{Z}/2$ is called an *orientation* if for each $\xi \in \partial\widehat{h}$ there exists an open neighborhood V of ξ in $\partial\widehat{h}$ such that either

- V intersects exactly one connector $C \in C_{\widehat{h}}$ which is supported on \widehat{h} and another hyperplane \widehat{k} , and the function $\alpha_{\partial\widehat{h}}$ on $V \setminus C$ is the characteristic function of $\partial\widehat{k}$ of an orientation ξ of the hyperplane \widehat{k} (which depends on V), or
- V does not intersect any connector and $\alpha_{\partial\widehat{h}}$ is constant on V .

Claim 5.3 Given a hyperplane \widehat{h} and an admissible collection of disjoint connectors $C_{\widehat{h}}$, there are exactly two orientations $\alpha_{\partial\widehat{h}}: \partial\widehat{h} \setminus \bigcup C_{\widehat{h}} \rightarrow \mathbb{Z}/2$, and they differ by the constant function 1.

Proof Since $C_{\widehat{h}}$ is a finite set of disjoint closed subsets of $\partial\widehat{h}$, one can find a partition I_1, \dots, I_n of the circle $\partial\widehat{h}$ into subintervals whose endpoints $\xi_0, \xi_1, \dots, \xi_n = \xi_0$ are disjoint from the connectors $C_{\widehat{h}}$, and each interval intersects at most one connector. Take ξ_0 , the joint endpoint of I_1 and I_n , and define $\alpha_{\partial\widehat{h}}(\xi_0) \in \mathbb{Z}/2$ arbitrarily. We show that this choice defines $\alpha_{\partial\widehat{h}}$ uniquely.




Figure 2: Examples of grids, from left to right, Examples 5.5, 5.6 and 5.7. The hyperplanes are shown in dotted lines. The connectors and an arbitrary choice of orientation are the thick solid lines.

For each interval I_m , if I_m intersects a connector C then we define $\text{type}_{\hat{h}}(I_m) := \text{type}_{\hat{h}}(I_m \cap C)$, and if it does not intersect any connector we set $\text{type}_{\hat{h}}(I_m) = 0$. The assumption that $C_{\hat{h}}$ is admissible amounts to saying that $\sum_{j=1}^n \text{type}_{\hat{h}}(I_j) = 0$. Now one can define the function $\alpha_{\partial\hat{h}}$ on the endpoints of the partition by $\alpha_{\partial\hat{h}}(\xi_m) = \alpha_{\partial\hat{h}}(\xi_0) + \sum_{j=1}^m \text{type}_{\hat{h}}(I_j)$.

Let I_m be one of the intervals. If I_m does not intersect any connector then let $\alpha_{\partial\hat{h}}$ be the constant function $\alpha_{\partial\hat{h}}(\xi_{m-1}) = \alpha_{\partial\hat{h}}(\xi_m)$ on I_m . Otherwise, I_m intersects a connector C that is supported on \hat{h} and $\hat{\xi}$. Since $\alpha_{\partial\hat{h}}(\xi_m) - \alpha_{\partial\hat{h}}(\xi_{m-1}) = \text{type}_{\hat{h}}(I_m)$, there is an orientation ξ of $\hat{\xi}$ such that $\alpha_{\partial\hat{h}}$ is equal to $\mathbb{1}_{\partial\xi}$ on ξ_{m-1} and ξ_m . Thus we can define $\alpha_{\partial\hat{h}} = \mathbb{1}_{\partial\xi}$ on I_m .

Clearly this construction satisfies the required condition. Moreover, since the construction of this function is invariant under subdivisions of the partition I_1, \dots, I_m , we conclude that it is the unique function that satisfies the required conditions. \square

Definition 5.4 (grids and oriented grids) A *grid* Γ is a pair (H, C) of a finite collection of hyperplanes H and a finite collection of disjoint connectors C which are supported on hyperplanes in H , satisfying that for each $\hat{h} \in H$ the set $C_{\hat{h}}$ of the connectors in C which are supported on \hat{h} is admissible.

An *oriented grid* $\Gamma = (H, C, \{\alpha_{\partial\hat{h}} \mid \hat{h} \in H\})$ is a grid (H, C) together with a choice of a orientations $\alpha_{\partial\hat{h}}$ for each $\hat{h} \in H$.

Before we proceed to the construction of the cocycle we give the three main examples of grids that we will use. See Figure 2 and compare with Remark 5.1.

Example 5.5 (a hyperplane grid) The grid consisting of one hyperplane $H = \{\hat{h}\}$ and no connectors.

Example 5.6 (an arc grid) The grid consisting of three hyperplanes $H = \{\hat{h}, \hat{k}_1, \hat{k}_2\}$ and two disjoint connectors C_1 and C_2 of \hat{h} with \hat{k}_1 and \hat{k}_2 , respectively, such that $\text{type}(C_i, \hat{h}, \hat{k}_i) = (1, 0)$. We remark that we will see in Proposition 7.7 that $(1, 0)$ connectors cannot exist under our assumptions, and hence arc grids cannot occur.

Example 5.7 (a cycle grid) The grid H consisting of cyclically intersecting hyperplanes $\{\hat{h}_1, \dots, \hat{h}_n\}$ and disjoint connectors $C = \{C_1, \dots, C_n\}$ such that each C_i is supported on \hat{h}_i and $\hat{h}_{i+1} \pmod n$ and satisfies $\text{type}(C_i, \hat{h}_i, \hat{h}_{i+1}) = (1, 1)$.

We end this section with the following remark.

Remark 5.8 If we replace a connector of a grid by its partition into disjoint clopen sets we obtain a new grid. An orientation for the original grid will remain an orientation for the new. In what follows, this operation will not make any substantial difference, thus we may identify two grids if they have the same hyperplanes and a common partitioning of their connectors.

6 Constructing cocycles

To an oriented grid $\Gamma = (H, C, \{\alpha_{\partial\hat{h}} \mid \hat{h} \in H\})$ we assign a 1-cocycle defined outside a ball in X in several steps.

Step 1 Fix $x_0 \in X$. Let $C \in C$ be a connector supported on \hat{h} and \hat{k} . Since $\hat{h} \cap \hat{k}$ is quasi-isometric to a tree, there exists a big enough ball $B(x_0, R)$ in X such that C can be written as a finite union of limit sets of the components of $\hat{h} \cap \hat{k} \setminus B(x_0, R)$. Let $R_0 > 0$ be big enough such that the above is true for all $C \in C$.

By Remark 5.8, we may assume that the connectors are exactly the limit sets of these components. We denote the component of $\hat{h} \cap \hat{k} \setminus B(x_0, R_0)$ that corresponds to C by \check{C} .

Step 2 For each $\hat{h} \in H$ let $\mathcal{T}_{\hat{h}}$ be the set of all halfspaces $t \in \hat{\mathcal{H}}(\hat{h})$ in the cube complex \hat{h} (ie $t = \hat{h} \cap \hat{h}'$ for some halfspace \hat{h}' in X) that satisfy:

- $t \cap B(x_0, R_0) = \emptyset$.
- ∂t is contained in a neighborhood V as in Definition 5.2, and in particular it intersects at most one connector C ;

- The halfspace t does not intersect any of the sets \check{C} of the connectors in $C_{\hat{h}}$ except if ∂t intersects a connector C supported on \hat{h} and \hat{k} , in which case $t \cap \hat{k} = t \cap \check{C}$, and in particular $t \cap \check{C}$ is a hyperplane in the cube complex t .

By Lemma 2.7, for every $\xi \in \partial \hat{h}$ there exists $t \in \mathcal{T}_{\hat{h}}$ such that $\xi \in \partial^\circ t$. Let $R_1 > R_0$ be such that $B(x_0, R_1)$ contains $\hat{h} \setminus \bigcup_{t \in \mathcal{T}_{\hat{h}}} t$ for all $\hat{h} \in H$.

Step 3 We define the function $\alpha_{\hat{h}}: \hat{h}^{(0)} \setminus B(x_0, R_1) \rightarrow \mathbb{Z}/2$ in the following way. On each halfspace $t \in \mathcal{T}_{\hat{h}}$ either ∂t does not intersect any connector, in which case $\alpha_{\partial \hat{h}}|_{\partial t}$ is constant, and we define $\alpha_{\hat{h}}|_t$ to be the constant function with the same value. Otherwise, ∂t intersects exactly one connector C which is supported on \hat{h} and \hat{k} , and there exists a halfspace k bounded by \hat{k} such that $\alpha_{\partial \hat{h}}|_{\partial t} = \mathbb{1}_{\partial k}|_{\partial t}$. We define $\alpha_{\hat{h}}|_t = \mathbb{1}_k|_t$.

Since vertices in a hyperplane correspond to edges in the cube complex, we identify the function $\alpha_{\hat{h}}$ with the corresponding function on the edges of $X \setminus B(x_0, R_1)$ which are transverse to \hat{h} .

Claim 6.1 *The sum $\alpha = \sum_{\hat{h} \in H} \alpha_{\hat{h}}$ is a 1-cocycle on $X \setminus B(x_0, R_1)$.*

Proof Let S be a square in $X \setminus B(x_0, R_1)$, and let \hat{h} and \hat{k} be the two hyperplanes which are transverse to the edges of S . Let $e_{\hat{h}}$ and $e'_{\hat{h}}$ (resp. $e_{\hat{k}}$ and $e'_{\hat{k}}$) be the pair of opposite edges of S which are transverse to \hat{h} (resp. \hat{k}) (see Figure 3). We have to show that

$$\alpha(e_{\hat{h}}) + \alpha(e'_{\hat{h}}) + \alpha(e_{\hat{k}}) + \alpha(e'_{\hat{k}}) = 0.$$

In other words, we have to show that if one pair of opposite edges in S has different values of α then so does the other pair.

Assume that the two opposite edges of S which are transverse to \hat{h} have different values of α , ie

$$\alpha(e_{\hat{h}}) + \alpha(e'_{\hat{h}}) = 1.$$

Then, if we view $e_{\hat{h}}$ and $e'_{\hat{h}}$ as vertices of \hat{h} , there exists a halfspace $t \in \mathcal{T}_{\hat{h}}$ of \hat{h} such that the two vertices are in t , and there exists a connector C such that these two vertices are separated by the hyperplane $\check{C} \cap t$ of t . Thus, the other pair of opposite edges $e_{\hat{k}}$ and $e'_{\hat{k}}$ of S , viewed as an edge in \hat{h} are transverse to the hyperplane $\check{C} \cap t$ of t . This implies that $\hat{k} \in H$ and the connector C which correspond to \check{C} is supported on \hat{h} and \hat{k} . Since \check{C} also separates $e_{\hat{k}}$ and $e'_{\hat{k}}$, now viewed as vertices in \hat{k} , they must also have different values of α . □




Figure 3: A square S , and the hyperplanes \hat{h} and \hat{k} that cross it. The 1-cocycle α obtains the value 1 on the top and right edges. If, on one pair of opposite edges, α obtains different values, then so it does on the other pair.

We remark that a priori the definition of α depends on R_0 and R_1 . However, any two such 1-cocycles coincide outside a large enough ball in X .

7 The parity function of a grid

Let $\Gamma = (\mathbf{H}, \mathbf{C}, \{\alpha_{\partial\hat{h}} \mid \hat{h} \in \mathbf{H}\})$ be an oriented grid, and let α be the 1-cocycle on $X \setminus B(\mathbf{x}_0, R_1)$ defined in the previous section. Since X has trivial first cohomology over $\mathbb{Z}/2$ at infinity, we deduce that there is $R > R_1$ such that α is a coboundary in $X \setminus B(\mathbf{x}_0, R)$. That is, there exists a 0-chain $\pi: X^{(0)} \setminus B(\mathbf{x}_0, R) \rightarrow \mathbb{Z}/2$ such that $\delta\pi = \alpha$. Since X is one-ended (ie connected at infinity), any two such 0-chains are equal up to a constant outside a possibly bigger ball. We call the function π the *parity function of the grid* Γ .

Let us denote by $\partial\Gamma$ the set $\{\xi \mid \alpha_{\partial\hat{h}}(\xi) = 1 \text{ for some } \hat{h} \in \mathbf{H}\} \cup \bigcup_{C \in \mathbf{C}} C \subset \partial X$. For example, in Figure 2, the set $\partial\Gamma$ are the thick solid lines.

Lemma 7.1 *The limit of the 0-chain π along geodesic rays defines a function on $\partial X \setminus \partial\Gamma$, which we denote as well by π . This function is constant on connected components of $\partial X \setminus \partial\Gamma$, and in particular continuous.*

Proof It is enough to prove that any point in $\partial X \setminus \partial\Gamma$ has a neighborhood V in $X \cup \partial X$ on which π is constant.

For a point $\xi \in \partial X \setminus \partial\Gamma$ there is a neighborhood halfspace \mathfrak{t} disjoint from $\partial\Gamma$ and from $B(\mathbf{x}_0, R)$. This halfspace, by definition, does not meet the cocycle α and thus the parity function on \mathfrak{t} is constant. □

Since the parity function is defined uniquely up to constants, it makes more sense to consider the function $\Delta\pi: (\partial X \setminus \partial\Gamma)^2 \rightarrow \mathbb{Z}/2$ (similarly, $\Delta\pi: (X^{(0)} \setminus B(x_0, R))^2 \rightarrow \mathbb{Z}/2$) given by $\Delta\pi(x, y) = \pi(x) - \pi(y)$. Moreover, for a path P whose endpoints are in the domain of π we define $\Delta\pi(P)$ to be the value of $\Delta\pi$ on the pair of endpoints of P .

Example 7.2 Let \hat{h} be a hyperplane and let $\Gamma = (H = \{\hat{h}\}, C = \emptyset, \alpha_{\partial\hat{h}} = 1)$ be the hyperplane grid described in Example 5.5. Then $\alpha = \mathbb{1}_{\hat{h}}$, $\partial\Gamma = \partial\hat{h}$, and the parity function is $\pi = \mathbb{1}_{\mathfrak{h}}$ for some choice of halfspace \mathfrak{h} of \hat{h} , which extends to the boundary to the function $\pi = \mathbb{1}_{\partial\mathfrak{h}|_{\partial X \setminus \partial\hat{h}}}$. We denote this parity function by $\pi_{\hat{h}}$. Similarly, the function $\Delta\pi_{\hat{h}}$ is the function that returns 0 if the two points are on the same side of \hat{h} , and returns 1 if they are separated by \hat{h} .

Remark 7.3 Using this notation, we can rewrite the notation introduced in Section 5.1 as follows. If C is a connector supported on \hat{h} , $\hat{\mathfrak{k}}$ and I_j is an interval on \hat{h} as in the definition of $\text{type}(C)$. Then

$$\text{type}_{\hat{h}}(I_j) = \Delta\pi_{\hat{\mathfrak{k}}}(I_j)$$

The next lemma follows from the definitions.

Lemma 7.4 Let Γ_1 and Γ_2 be two oriented grids with disjoint or identical connectors. Then we denote by $\Gamma_1 + \Gamma_2$ the pair $\Gamma = (H_1 \cup H_2, C_1 \Delta C_2)$. The pair Γ is a grid, and there exists an orientation on Γ which satisfies $\alpha_{\partial\hat{h}} = \alpha_{1,\partial\hat{h}} + \alpha_{2,\partial\hat{h}}$ when all of the functions are defined. Moreover, the cocycle (resp. the parity function) associated to Γ is the sum of the cocycles (resp. the parity functions) associated with Γ_1 and Γ_2 . \square

Lemma 7.5 Let $\xi \in \partial\Gamma \setminus \bigcup_{C \in \mathcal{C}} C$ and let $\hat{h}_1, \dots, \hat{h}_n \in H$ be all the hyperplanes such that $\xi \in \partial\hat{h}_i$ and $\alpha_{\partial\hat{h}_i}(\xi) = 1$. Then there is a neighborhood V of ξ in $X \cup \partial X \setminus \partial\Gamma$ on which

$$\Delta\pi = \sum_{i=1}^n \Delta\pi_{\hat{h}_i}.$$

Proof Let \mathfrak{k} be a halfspace neighborhood of ξ , disjoint from the connectors C and from $B(x_0, R)$ such that for $i = 1, \dots, n$, $\mathfrak{k} \cap \hat{h}_i \in \mathcal{T}_{\hat{h}_i}$ (see Step 2 in Section 6). The definition of α and $\alpha_{\partial\hat{h}_i}(\xi) = 1$ imply that $\alpha(\hat{h} \cap \mathfrak{k}) = 1$. Thus, on \mathfrak{k} the cocycle α is the sum of the corresponding hyperplane cocycles (see Example 7.2), ie $\alpha = \sum_{i=1}^n \mathbb{1}_{\hat{h}_i \cap \mathfrak{k}}$, and the lemma follows. \square

By subdividing a path which avoids the connectors of a grid into small enough segments and applying the previous lemma, one can deduce the following.

Corollary 7.6 *Let P be a path in ∂X whose endpoints are in $\partial X \setminus \partial \Gamma$ and which is disjoint from $\bigcup_{C \in \mathcal{C}} C$. For each $\hat{h} \in \mathbf{H}$ let $J_{\hat{h},1}, \dots, J_{\hat{h},n}$ be $n = n(\hat{h})$ disjoint subintervals of P such that $P \cap \alpha_{\partial \hat{h}}^{-1}(1) \subset \bigcup_{j=1}^n J_{\hat{h},j}$, where the endpoints of each subinterval $J_{\hat{h},i}$ are disjoint from $\partial \hat{h}$ and its interior intersects $\partial \hat{h}$ in $\alpha_{\partial \hat{h}}^{-1}(1)$. Then*

$$\Delta\pi(P) = \sum_{\hat{h} \in \mathbf{H}} \sum_{j=1}^n \Delta\pi_{\hat{h}}(J_{\hat{h},j}). \quad \square$$

Proposition 7.7 *Under the assumptions of Section 2.4, there is no connector of type $(1, 0)$.*

Proof Assume that \hat{h} and \hat{k} have a connector C such that $\text{type}(C, \hat{h}, \hat{k}) = (1, 0)$. By the dynamics of $\text{Stab}(\hat{h})$ on $\partial \hat{h}$ we can find such \hat{k} and C in any small open set in $\partial \hat{h}$.




Figure 4: The hyperplanes \hat{h} and \hat{k}_i and \hat{k}'_i are drawn as dotted lines. The paths P_1 and P_2 between ζ and ξ are shown. The boundary of the grid which contradicts the existence of type- $(1, 0)$ connectors is the pair of solid horizontal lines.

Since X is assumed to be essential, the open limit sets of halfspaces $\partial^\circ \hat{h}$ and $\partial^\circ \hat{h}^*$ are nonempty. Let ξ and ζ be points in $\partial^\circ \hat{h}$ and $\partial^\circ \hat{h}^*$, respectively. By the cyclic connectivity theorem (by Ayres [2] and Whyburn [21]), since our space has no cutpoints we can find two disjoint paths P_1 and P_2 connecting ξ and ζ (see Figure 4). Let $A_i = \partial \hat{h} \cap P_i$ for $i = 1, 2$.

Let I_1, I_2, \dots, I_n be disjoint intervals of \widehat{h} disjoint from A_2 and such that the union of their interiors contain A_1 . Let $\delta > 0$ be such that δ -neighborhoods of the endpoints of I_1, \dots, I_n are pairwise disjoint and disjoint from $P_1 \cup P_2$. For each I_i find two hyperplanes \widehat{k}_i and \widehat{k}'_i contained in the δ -neighborhoods of the two endpoints of I_i , and two connectors C_i and C'_i supported on \widehat{h} and on \widehat{k}_i and \widehat{k}'_i , respectively, with $\text{type}(C_i, \widehat{h}, \widehat{k}_i) = \text{type}(C'_i, \widehat{h}, \widehat{k}'_i) = (1, 0)$, as described above, in each such open neighborhood.

The grid $(H = \{\widehat{h}_1, \widehat{k}_1, \widehat{k}'_1, \dots, \widehat{k}_n, \widehat{k}'_n\}, C = \{C_1, C'_1, \dots, C_n, C'_n\})$ is the sum of arc grids described in Example 5.6. Orient the grid such that $\alpha_{\partial\widehat{h}}$ is 1 on A_1 (and 0 on A_2).

By Corollary 7.6 applied to P_1 (with $n(\widehat{h}) = 1, J_{\widehat{h},1} = P_1$ and $n(\widehat{k}_i) = n(\widehat{k}'_i) = 0$ for $i = 1, \dots, n$), we get that $\pi(\xi) - \pi(\zeta) = \Delta\pi(P_1) = \Delta\pi_{\widehat{h}}(P_1) = 1$ because P_1 intersects the grid exactly as it would intersect the hyperplane grid \widehat{h} , and we recall that ξ and ζ are on different sides of \widehat{h} , thus have different parity.

On the other hand, P_2 does not intersect the grid, and thus $\pi(\xi) - \pi(\zeta) = \Delta\pi(P_2) = 0$. This contradicts the existence of connectors of type $(1, 0)$. □

The following useful corollary is a direct consequence of Proposition 7.7.

Corollary 7.8 *For two hyperplanes \widehat{h} and \widehat{k} , assume that there exists two points of $\partial\widehat{h}$ on different sides of $\partial\widehat{k}$. Then the hyperplanes \widehat{h} and \widehat{k} share a $(1, 1)$ connector. □*

8 Approximations of curves and arcs

Let J be a Jordan curve or arc. In this section, we assume that the Jordan arc J is parametrized by $[0, 1]$. Given a parameter t , the associated point is $J(t)$; conversely, if x is a point on J its parameter will be denoted by $J^{-1}(x)$. A segment between two points ζ and ξ on the arc J is denoted $[\zeta, \xi]_J$.

Definition 8.1 *A piecewise-hyperplane curve or PH curve (resp. PH arc) is a (not necessarily simple) parametrized curve (resp. arc) P on ∂X which has a partition into finitely many oriented segments I_1, \dots, I_n such that:*

- (1) Each segment I_i is a segment of the boundary $\partial\widehat{h}$ of some hyperplane \widehat{h} . In fact, the choice of the hyperplane \widehat{h} is implicitly part of the data, but to avoid cumbersome notation we will simply write $I_i \subset \partial\widehat{h}$ and say that I_i is supported on \widehat{h} .

- (2) Any two segments which are supported on the same hyperplane \widehat{h} are disjoint.
- (3) For all $1 \leq i \leq n$ (resp. $1 \leq i < n$) the last point of I_i is the first point of I_{i+1} . It is called a *vertex* of P , and we say that it is *supported on* the pair $(\widehat{h}, \widehat{k})$, where I_i and I_{i+1} are supported on \widehat{h} and \widehat{k} , respectively.
- (4) For every vertex ξ supported on $(\widehat{h}, \widehat{k})$, the pair $(\widehat{h}, \widehat{k})$ supports a $(1, 1)$ connector.

An ϵ -approximating PH curve (resp. arc) for J is a PH curve (resp. arc) which is at distance less than ϵ from J with respect to the sup metric.

The following lemmas show that given a Jordan curve one can construct PH approximations of the curve with certain technical restrictions that will be useful later on. We first prove this in the case of an arc.

A set \mathcal{H} of limit sets of halfspaces is a δ -cover of a set S in the boundary if it is a cover of S and if every element in \mathcal{H} has diameter less than δ .

Lemma 8.2 *Let J be a Jordan arc such that the two endpoints ζ_0 and ζ_1 of J are on the boundaries of the hyperplanes $\partial\widehat{k}_0$ and $\partial\widehat{k}_1$, respectively. Let \widehat{h} be a hyperplane such that $\partial\widehat{h}$ does not intersect J except maybe for its endpoints. Let \mathcal{K} be a finite set of hyperplanes. For every $\epsilon > 0$, there exists δ such that for any two distinct halfspaces $\partial^\circ h_0$ and $\partial^\circ h_1$ not in \mathcal{K} of diameter less than δ covering ζ_0 and ζ_1 , respectively, there exists a δ -cover $\mathcal{H} = \{\partial^\circ h_i\}$ of J and an ϵ -approximating PH arc P supported on $\bigcup \partial\widehat{h}_i$ such that*

- the set \mathcal{H} is disjoint from \mathcal{K} ,
- P has at most one segment on each $\partial\widehat{h}_i$,
- the endpoints of P lie on $\partial\widehat{k}_0$ and $\partial\widehat{k}_1$,
- the first and last segments of P are supported on $\partial\widehat{h}_0$ and $\partial\widehat{h}_1$, respectively,
- the hyperplanes \widehat{h}_0 and \widehat{h}_1 share $(1, 1)$ connectors with \widehat{k}_0 and \widehat{k}_1 , respectively, and
- P intersects $\partial\widehat{h}$ only on $\partial\widehat{h}_0$ and $\partial\widehat{h}_1$.

Proof Let $0 < \delta < \frac{\epsilon}{3}$ be such that for any two points ζ and ξ on J at distance less than δ , the segment $[\zeta, \xi]_J$ has diameter less than $\frac{\epsilon}{2}$, and such that for $i = 0, 1$ there exists a point on $\partial\widehat{k}_i$ at distance greater than δ from ζ_i .

Let $\partial^\circ h_0$ and $\partial^\circ h_1$ be open limit sets of halfspaces of diameter less than δ covering ζ_0 and ζ_1 , respectively. Note that by the choice of δ , there are points of $\partial \hat{\mathfrak{E}}_i$ in both sides of $\partial \hat{h}_i$, which by Corollary 7.8 ensures that $\partial \hat{h}_i$ and $\partial \hat{\mathfrak{E}}_i$ share a $(1, 1)$ connector.

Let $\delta' < \delta$ be such that $J \setminus (\partial^\circ h_0 \cup \partial^\circ h_1)$ is at distance more than δ' from $\partial \hat{h}$, ζ_0 and ζ_1 , and such that no halfspace in \mathcal{K} has diameter less than δ' . Let $\mathcal{H}' = \{\partial^\circ h_2, \dots, \partial^\circ h_n\}$ be a minimal finite δ' -cover of $J \setminus (\partial^\circ h_0 \cup \partial^\circ h_1)$. Then $\mathcal{H} = \{\partial^\circ h_0, \dots, \partial^\circ h_n\}$ is a minimal finite δ -cover of J such that

- the set \mathcal{H} disjoint from \mathcal{K} ,
- the only limit set containing ζ_0 is $\partial^\circ h_0$,
- the only limit set containing ζ_1 is $\partial^\circ h_1$, and
- the only limit sets that may intersect $\partial \hat{h}$ are $\partial^\circ h_0$ and $\partial^\circ h_1$.

Let δ'' be smaller than the diameter of any $\partial \hat{h}_i$.

For $i = 0, \dots, n$ define $y_i = \sup J^{-1}(\partial^\circ h_i) \in [0, 1]$. We define a successor function S as follows: $S(i) = j$, where $J(y_i) \in \partial^\circ h_j$ and y_j is maximal under this condition. Note that the function satisfies $y_{S(i)} > y_i$, unless $i = 1$ (in which case $y_i = 1$). These conditions ensure that there exists p such that $\zeta_1 \in \partial^\circ h_{S^p(0)} = \partial^\circ h_1$.

For all $1 < i \leq p$, let v_i be an intersection point of $\partial \hat{h}_{S^{i-1}(0)}$ and $\partial \hat{h}_{S^i(0)}$. Let $\partial^\circ h'_1, \dots, \partial^\circ h'_p$ be distinct hyperplanes of diameter less than δ'' that cover v_1, \dots, v_p , respectively. By the choice of δ'' they do not belong to \mathcal{K} , do not intersect $\partial \hat{h}$ and do not contain ζ_0 or ζ_1 . Moreover, each of $\partial^\circ h_{S^{i-1}(0)}$ and $\partial^\circ h_{S^i(0)}$ contains points on both sides of $\partial \hat{h}'_i$. By Corollary 7.8, $\partial \hat{h}'_i$ shares $(1, 1)$ connectors with both $\partial \hat{h}_{S^{i-1}(0)}$ and $\partial \hat{h}_{S^i(0)}$. Let v'_i and v''_i be in the intersection of $\partial \hat{h}'_i$ with $\partial \hat{h}_{S^{i-1}(0)}$ and $\partial \hat{h}_{S^i(0)}$ respectively. Let v''_0 (resp. v'_{p+1}) be the point in the intersection of $\partial \hat{h}_0$ and $\partial \hat{\mathfrak{E}}_0$ (resp. $\partial \hat{h}_1$ and $\partial \hat{\mathfrak{E}}_1$) which is closest to ζ_0 (resp. ζ_1). Let I_i be one of the two intervals of $\partial \hat{h}_{S^i(0)}$ with endpoints v''_i and v'_{i+1} . Let I'_i be one of the two intervals of $\partial \hat{h}'_i$ with endpoints v'_i and v''_i .

Let η be such that subpaths of J parametrized by $[x - \eta, x + \eta]$ have diameter less than δ and $y_{S(i)} - y_i > 2\eta$ for all $i \neq 1$. We obtained a PH arc $P = (I_0, I'_1, I_1, \dots, I'_p, I_p)$, which we parametrize continuously, such that v'_i and v''_i have parameters $y_{S^{i-1}(0)} - \eta$ and $y_{S^i(0)} + \eta$, respectively (with the convention that the parameter of v''_0 is 0 and of v'_{p+1} is 1).

It remains to show that the path P is at distance ϵ from J with respect to the sup metric. Points in I_i and $J(y_{S^i(0)})$ are in $\partial \hat{h}_{S^{i-1}(0)}$ and are thus at distance less than δ . The

points $J(y_{S^{i-1}(0)})$ and $J(y_{S^i(0)})$ belong to $\partial^\circ h_{S^i(0)}$ and thus are at a distance less than δ , which by the definition of δ implies that the path $J([y_{S^{i-1}(0)}, y_{S^i(0)}])$ connecting them on J has diameter less than $\frac{\epsilon}{2}$, and thus also $J([y_{S^{i-1}(0)} + \eta, y_{S^i(0)} - \eta])$. Hence I_i and $J([y_{S^{i-1}(0)} + \eta, y_{S^i(0)} - \eta])$ are at distance less than $\frac{\epsilon}{2} + 3\delta < \epsilon$.

Similarly, I'_i is at distance less than 2δ from $J(y_{S^{i-1}(0)})$, and so

$$J([y_{S^{i-1}(0)} - \eta, y_{S^{i-1}(0)} + \eta])$$

are at distance less than $3\delta < \epsilon$ from I'_i . This completes the proof that P is an ϵ -approximation of J that satisfies the requirements of the lemma. \square

In what follows we will denote by $[x, y]_{\hat{h}}$ one of the two subsegments of the limit set of a hyperplane $\partial\hat{h}$ that connect x and y (which are not necessarily distinct).

Definition 8.3 Given a path J and a hyperplane $\partial\hat{h}$, a δ -bypass of $x \in J \cap \partial\hat{h}$ on $\partial\hat{h}$ is a segment $[x, y]_{\hat{h}}$ of $\partial\hat{h}$ with $y \in J \cap \partial\hat{h}$ such that $[x, y]_{\hat{h}}$ can be partitioned into a finite union $\bigcup_{i=0}^{n-1} [x^i, x^{i+1}]_{\hat{h}}$ (with $x^0 = x$ and $x^n = y$) of disjoint segments (except at their extremities) of diameter less than δ , with extremities in $J \cap \partial\hat{h}$ and such that for every $\zeta \in [x^i, x^{i+1}]_{\partial\hat{h}} \cap J$ we have $J^{-1}(\zeta) \leq J^{-1}(x^{i+1})$. A δ -bypass $[x, y]_{\hat{h}}$ is maximal if $J^{-1}(y)$ is maximal amongst all δ -bypasses of x (ie if $[x, y']_{\hat{h}}$ is another δ -bypass then $J^{-1}(y') \leq J^{-1}(y)$). A maximal bypass is *degenerate* if $x = y$.

Remarks 8.4 Let J be a path.

- (1) If $[x, y]_{\hat{h}}$ is a δ -bypass, then $J^{-1}(\zeta) \leq J^{-1}(y)$ for all $\zeta \in [x, y]_{\hat{h}} \cap J$.
- (2) If $[x, y]_{\hat{h}}$ is a maximal δ -bypass, then for any $z \in [x, y]_{\hat{h}} \cap J$, the segment $[z, y]_{\hat{h}}$ is a maximal δ -bypass.
- (3) There exists ν such that for all maximal δ -bypass $[x, y]_{\hat{h}}$, the segment of J parametrized by $(J^{-1}(y), J^{-1}(y) + \nu)$ does not intersect $\partial\hat{h}$.
- (4) If two maximal δ -bypasses $[x, y]_{\hat{h}}$ and $[x', y']_{\hat{h}}$ intersect, then $y = y'$. Indeed, otherwise we can suppose that $J^{-1}(y') < J^{-1}(y)$. From (1), the element y cannot belong to $[x', y']_{\hat{h}}$, so $[x, y]_{\hat{h}} \not\subset [x', y']_{\hat{h}}$. Therefore either x' or y' belongs to $[x, y]_{\hat{h}}$; using (2), we get a contradiction.
- (5) For every $x \in J \cap \partial\hat{h}$ there exists a maximal δ -bypass. Indeed, the set of endpoints y of δ -bypasses $[x, y]_{\hat{h}}$ is closed.

Definition 8.5 A *detour* of δ -bypasses is a set of maximal bypasses $\{[x_i, y_i]_{\widehat{h}}\}$ such that for any two bypasses $[x_i, y_i]_{\widehat{h}}$ and $[x_j, y_j]_{\widehat{h}}$, either $J^{-1}(y_i) < J^{-1}(x_j)$ or $J^{-1}(y_j) < J^{-1}(x_i)$.

A detour is *covering* if, for any element $z \in J \cap \partial\widehat{h}$, there is one bypass $[x, y]_{\widehat{h}}$ of the detour such that $J^{-1}(x) \leq J^{-1}(z) \leq J^{-1}(y)$, or, equivalently, $J \cap \partial\widehat{h} \subseteq \bigcup_i [x_i, y_i]_J$.

From Remark 8.4(4), bypasses of a detour are disjoint.

Lemma 8.6 *There exists a finite covering detour.*

Proof We construct the detour by induction: take the element x of smallest parameter that is not covered by the detour, and by Remark 8.4(5) add a maximal bypass of x . Remark 8.4(3) implies that the process finishes in a finite number of steps. \square

Lemma 8.7 *Let J be an arc such that the two endpoints ζ_0 and ζ_1 of J are on the boundaries of the hyperplanes $\partial\widehat{\xi}_0$ and $\partial\widehat{\xi}_1$, respectively. Let $\epsilon > 0$ and let \widehat{h} be a hyperplane. There exists an ϵ -approximating PH arc P which satisfies the following two conditions:*

- *The hyperplanes supporting the first and last segments of P share a $(1, 1)$ connector with $\widehat{\xi}_0$ and $\widehat{\xi}_1$, respectively.*
- *Any intersection of P with $\partial\widehat{h}$ is along a segment of $P \cap \partial\widehat{h}$.*

Proof Let $\delta_1 < \frac{\epsilon}{9}$ such that for any two points ζ and ξ on J at distance less than δ_1 the segment $[\zeta, \xi]_J$ has diameter less than $\frac{\epsilon}{3}$. This ensures that a δ_1 -bypass $[x, y]_{\widehat{h}}$ is an $\frac{\epsilon}{3} + \delta_1$ -approximation of the segment $[x, y]_J$ (with the natural parametrization of $[x, y]_{\widehat{h}}$ by $[J^{-1}(x), J^{-1}(y)]$ which agrees with J on the endpoints $x^i \in J \cap \partial\widehat{h}$ of its subdivision as it appears in the definition of a δ -bypass). Indeed, $[x, y]_{\widehat{h}}$ is a union of segments $[x^i, x^{i+1}]_{\widehat{h}}$ of diameter less than δ_1 and with extremities on J with increasing parameter. By the choice of δ_1 , each $[x^i, x^{i+1}]_J$ has diameter less than $\frac{\epsilon}{3}$. Since the parameters are increasing, the segments $[x^i, x^{i+1}]_J$ on J are disjoint (except at their extremities), and each $[x^i, x^{i+1}]_{\widehat{h}}$ is an $\frac{\epsilon}{3} + \delta_1$ -approximation of $[x^i, x^{i+1}]_J$.

Let $\mathcal{D} = \{[x_1, y_1]_{\widehat{h}}, \dots, [x_n, y_n]_{\widehat{h}}\}$ be a finite covering detour of δ_1 -bypasses, ordered by their parameter along the curve J . Let $\delta_2 < \delta_1$ be such that each bypass is at distance at least $2\delta_2$ from any other. Let $\delta_3 < \delta_2$ be such that each pair of points ζ and ξ on $\partial\widehat{h}$ at distance less than δ_3 has a path from ζ to ξ on $\partial\widehat{h}$ with diameter less than δ_2 .

Let J_i be the segment $[y_i, x_{i+1}]_J$ on J for all $i = 0, \dots, n$ (where $y_0 = \zeta_0$ and $x_{n+1} = \zeta_1$). Each of the endpoints of J_i are on $\partial\widehat{h}$ except maybe for J_0 and J_n , which are on $\partial\widehat{k}_0$ and $\partial\widehat{k}_1$.

Let $\delta < \delta_3$ be obtained from Lemma 8.2 for all the J_i . Take covers $\partial^\circ h_{i,x}$ and $\partial^\circ h_{i,y}$ of x_i and y_i , respectively, of diameter less than δ , such that the limit sets of their bounding hyperplanes are all disjoint. It may happen that $x_i = y_i$ if a bypass is degenerate, but by taking $\partial^\circ h_{i,y}$ small enough, we can ensure that its boundary does not intersect that of $\partial^\circ h_{i,x}$. Since \mathcal{D} is covering, the interior of each J_i does not intersect $\partial\widehat{h}$.

Using Lemma 8.2, let P_i be an ϵ -approximation of J_i whose first and last segments are supported on $\partial\widehat{h}_{i,y}$ and $\partial\widehat{h}_{i+1,x}$, respectively. Moreover, we may assume that the ϵ -approximations P_i are supported on distinct hyperplanes. Indeed we can construct them one after the other and at each step adding all the used hyperplanes to the set \mathcal{K} of Lemma 8.2.

Let ζ_i and ξ_i be the first and last points of P_i , respectively. By construction, the point ζ_i (resp. ξ_i) is at distance less than δ from y_{i-1} (resp. x_i). Thus there is a segment $[\xi_i, \zeta_{i+1}]_{\widehat{h}}$ on $\partial\widehat{h}$ at distance less than δ_2 from $[x_i, y_i]_{\widehat{h}}$, and by the choice of δ_2 these paths are disjoint and not reduced to a point (since $\xi_i \neq \zeta_{i+1}$). Thus $[\xi_i, \zeta_{i+1}]_{\widehat{h}}$ is an $\frac{\epsilon}{3} + \delta_1 + \delta_2 < \epsilon$ -approximation of $[x_i, y_i]_{\widehat{h}}$. By concatenating P_i and $[\xi_i, \zeta_{i+1}]_{\widehat{h}}$ alternately, we obtain an ϵ -approximating PH arc of J satisfying the required properties. □

Lemma 8.8 *Let J be a Jordan curve, let $\epsilon > 0$ and let \widehat{h} be a hyperplane. Then there exists an ϵ -approximating PH curve P such that any intersection of P with $\partial\widehat{h}$ is along a segment of P .*

Proof Up to taking ϵ small enough we can assume that the diameter of J is larger than ϵ .

To approximate a curve, we do the following: if $J = \partial\widehat{h}$, then the approximation is $\partial\widehat{h}$; otherwise there is a point $\xi \in J \setminus \partial\widehat{h}$. Let $\eta < \frac{\epsilon}{2}$ be such that for any two points ζ and ξ on J at distance less than η , one of the segments of J with extremities ζ and ξ has diameter less than $\frac{\epsilon}{2}$.

Let k be a halfspace such that $\partial^\circ k$ contains ξ , it has diameter less than η and such that $\partial\widehat{k}$ does not intersect $\partial\widehat{h}$. This is possible since ξ is not on $\partial\widehat{h}$. Let ζ_0 and ζ_1 be two

points of $J \cap \partial \hat{\mathfrak{k}}$. Let J_0 be the path between ζ_0 and ζ_1 of diameter less than $\frac{\epsilon}{2}$, and let J_1 be the complementary path. Notice that any segment of $\partial \hat{\mathfrak{k}}$ is an ϵ -approximation of J_0 . Applying Lemma 8.7 to J_1 (and the hyperplanes $\mathfrak{k}_1 = \mathfrak{k}_2 = \mathfrak{k}$), let P_1 be a δ -approximation of J_1 with $0 < \delta < \epsilon$ which is less than the diameter of \mathfrak{k} . The approximation P_1 of J_1 begins and ends on \mathfrak{k} and thus can be closed by a segment to form an approximation of J . Notice that (2) of Definition 8.1 is ensured by the fact that \mathfrak{k} cannot be used in P_1 , and (4) follows from Lemma 8.7. \square

9 The parity function of a Jordan curve

9.1 The parity function of a PH curve

Let P be a PH curve. We assign a parity function to P in the following way. Let $\delta > 0$ be smaller than half the minimal distance between the vertices of P . For each vertex ξ of P which is supported on the two hyperplanes $\hat{\mathfrak{h}}$ and $\hat{\mathfrak{k}}$, let C be a type-(1, 1) connector which is small enough that in each of $\partial \hat{\mathfrak{h}}$ and $\partial \hat{\mathfrak{k}}$ the connector is contained in an interval which stays in the δ -neighborhood of ξ (one can find such a connector using the dynamics of $\text{Stab}_G(\hat{\mathfrak{h}} \cap \hat{\mathfrak{k}})$ on $\partial \hat{\mathfrak{h}} \cap \partial \hat{\mathfrak{k}}$). Let Γ be the grid consisting of the hyperplanes on which P is supported and the connectors which we assigned to each vertex of P . We choose the orientation so that outside of the δ -neighborhoods of the vertices of P , we have $P = \partial \Gamma$ (one can do so, since each connector is of type (1, 1) and contained in an interval which stays in the δ -neighborhood of the corresponding vertex).

Lemma 9.1 *For $\epsilon > 0$ small enough, there exists $\delta > 0$ such that outside an ϵ -neighborhood of the vertices of P and outside P , the parity function of Γ (associated to δ) does not depend on the choice of connectors (up to constants).*

Proof By Corollary 3.3, for any $\epsilon > 0$ there exists $\delta > 0$ small enough such that any two points in ∂X outside the ϵ -neighborhood of the vertices of P can be connected with a path that stays outside the δ -neighborhood of the vertices of P . We note that this path satisfies the conditions of Corollary 7.6 and, since it does not enter the δ -neighborhood of the vertices of P , it does not depend on the choice of connectors. \square

We deduce the following:

Corollary 9.2 *For a PH curve P , the parity $\Delta\pi$ of the grid constructed above for δ has a limit as $\delta \rightarrow 0$ which is defined outside P .* \square




Figure 5: The Jordan curve J (middle), two approximating PH curves P (top) and P' (bottom) and the transverse auxiliary PH curves Q_i

9.2 The parity function of a Jordan curve

We would like to follow the same idea, approximating the Jordan curve J with PH curves, and taking the limit of their parity functions as the definition of the parity function of J .

Lemma 9.3 *Let J be a Jordan curve. Let $\eta > 0$, there exists $\epsilon > 0$ such that for any two ϵ -approximating PH curves P and \tilde{P} for J the parity functions of P and \tilde{P} are equal (up to constants) outside an η -neighborhood of J .*

Proof By the no blob lemma (Lemma 3.4) let η' be such that outside any set that is contained in a ball $B(\xi, \eta')$ there is at most one component which is not contained in $B(\xi, \eta)$. By local connectedness, let $0 < \epsilon < \frac{1}{4}\eta'$ be such that any two points in $B(\xi, \epsilon)$ can be connected by a path in $B(\xi, \frac{1}{2}\eta')$.

Let P and \tilde{P} be two ϵ -approximating PH curves for J . Let $\epsilon' > 0$ be smaller than the diameter of the hyperplanes of P and \tilde{P} , and let P' be an ϵ' approximating PH curve for J such that the diameter of the hyperplanes of P' are smaller than ϵ' .

It suffices to show that the parity functions of P' and P (similarly, \tilde{P}) coincide (up to constants) outside the η -neighborhood of J .

Let $\xi_0, \xi_1, \dots, \xi_n = \xi_0$ be a partition of J into intervals of diameter less than ϵ . Let ζ_i and ζ'_i be the corresponding points on P and P' at distance less than ϵ from ξ_i . By perturbing these points if necessary, we may assume that $\{\zeta_i, \zeta'_i\}_{i=0}^{n-1}$ are $2n$ distinct points, and that they are not vertices of P and P' (see Figure 5). Since $\epsilon < \frac{1}{4}\eta'$, the interval between ζ_i and ζ_{i+1} on P (resp. ζ'_i and ζ'_{i+1} on P') is contained in $B(\xi_i, \frac{1}{2}\eta')$.

By Lemma 8.8 and the definition of ϵ , connect ζ_i and ζ'_i by a PH arc Q_i in $B(\xi_i, \frac{1}{2}\eta')$. Form the short closed path P_i by connecting $\zeta_i, \zeta_{i+1}, \zeta'_i, \zeta'_{i+1}$ along P, Q_{i+1}, P', Q_i . We note that by the construction, each P_i is contained in $B(\xi_i, \frac{1}{2}\eta')$.

Let $\delta < \frac{1}{2}\eta'$ be such that the δ -neighborhoods of all the vertices of both P and P' and the arcs Q_i are disjoint. Choose a connector in the δ -neighborhood of each vertex (including connectors for the endpoints of Q_i with P and P'). Let Γ, Γ' and Γ_i be the grids described in Section 9.1 for P, P' and P_i . Note that, by construction, $\Gamma + \Gamma' = \sum_i \Gamma_i$ (in the notation of Lemma 7.4).

Since each P_i is contained in $B(\xi_i, \frac{1}{2}\eta')$ and $\delta < \frac{1}{2}\eta'$, the boundary $\partial\Gamma_i$ of the corresponding grid is contained in $B(\xi_i, \eta')$. This implies that the complement of each $\partial\Gamma_i$ has at most one component which is not contained in the η -neighborhood of J . Thus, the parity function π_i of Γ_i is constant outside the η -neighborhood of J . If we denote by π and π' the parity functions of P and P' , respectively, then, by Lemma 7.4, $\pi - \pi' = \sum_i \pi_i$. We deduce that $\pi - \pi'$ is constant outside an η -neighborhood of J . □

This implies the following.

Corollary 9.4 *For a Jordan curve J , the parity $\Delta\pi$ for an ϵ approximating PH curve has a limit as $\epsilon \rightarrow 0$, which is defined outside J .* □

10 Jordan’s theorem

Proposition 10.1 *Let P be a PH curve, and let ξ_i be a vertex of P . Let J_i and J_{i+1} be the segments of P on $\partial\hat{h}_i$ and $\partial\hat{h}_{i+1}$, respectively, which are incident to ξ_i . Assume J_i does not intersect $\partial\hat{h}_{i+1}$ except at ξ_i . Let I_i be an open interval in $\partial\hat{h}_{i+1}$ which is disjoint from P and such that one of its endpoints is ξ_i . Then points on I_i have the same parity with respect to P as any point in h_{i+1}^* which is close enough to J_{i+1} , where h_{i+1} is the halfspace of \hat{h}_{i+1} which contains J_i .*

Proof We write $\xi = \xi_i, J' = J_i, J = J_{i+1}, I = I_i, \hat{h}' = \hat{h}_i$ and $\hat{h} = \hat{h}_{i+1}$ (see Figure 6).

Clearly all the points on I have the same parity, because they are connected with a path that does not meet P , namely a subinterval of I . By Lemma 7.5, around any point of J different of ξ , the parity function is the characteristic function of h or h^* . Thus, all the points in h^* which are close enough to J have the same parity.




Figure 6: The two possible cases of the arcs of $\partial\hat{t}$ drawn on the same figure. Beneath the horizontal, an arc A which does not cross J' and connects I and J . Above the horizontal, a collection of arcs A_i which cross J' .

Thus it suffices to find two points, one on I and the other in \mathfrak{h}^* close enough to J , that have the same parity.

Let \mathfrak{t} be a halfspace neighborhood of ξ which is small enough that $\mathfrak{t} \cap (P \cup \partial\hat{h}) \subset I \cup J \cup J'$. The connected components of $\partial\hat{t} \setminus \partial\hat{h}$ are arcs which stay on one side of \hat{h} . If one of these arcs A connects I with J and is contained in \mathfrak{h}^* , then since it does not cross P we deduce that its endpoint on I has the same parity as any other point on A , and in particular points which are arbitrarily close to J in \mathfrak{h}^* .

If not, let us analyze how the arcs of $\partial\hat{t} \setminus \partial\hat{h}'$ intersect J' . First we observe that only finitely many arcs of $\partial\hat{t} \setminus \partial\hat{h}$ intersect J' . We denote them by A_1, \dots, A_n .

Let π be the parity function of P , and let $\pi_{\hat{t}}$ be the parity function of the single hyperplane grid \hat{t} , ie $\pi_{\hat{t}}$ is (up to a constant) the characteristic function of $\partial\mathfrak{t}$.

If we follow the segment J' , we see that one of its endpoints is in $\partial\mathfrak{t}$ and the other is in $\partial\mathfrak{t}^*$. Therefore the difference in the value of $\pi_{\hat{t}}$ between the two endpoints is 1. On the other hand, by Corollary 7.6, it is the sum over the types of intersections of J' with the hyperplane \hat{t} . We break this sum apart by arcs A_i :

$$1 = \Delta\pi_{\hat{t}}(J') = \sum_i \text{type}_{\hat{t}}(A_i \cap J') = \sum_i \text{type}_{\hat{t}}(A_i \cap J') = \sum_i \Delta\pi_{\hat{t}}(A_i),$$

where the third equality follows from Proposition 7.7. Note that any arc A_i whose endpoints lie on the same side of ξ , ie both on I or both on J , contributes 0 to the sum because its endpoints have the same parity with respect to π .

Therefore, there exists i such that the arc A_i whose endpoints are on I and J and such that $\Delta\pi_{\hat{t}}(A_i \cap J') = 1$. Since A_i intersects P only in J' and is contained in $\partial\mathfrak{h}$, it

follows that points on I have different parity than points in $\partial\hat{h}$ which are close enough to J , from which the desired conclusion follows. \square

Corollary 10.2 *Let P be a PH curve and let ξ_i and ξ_{i+1} be two consecutive vertices of P . Let J_i, J_{i+1} and J_{i+2} be the segments of P on $\partial\hat{h}_i, \partial\hat{h}_{i+1}$ and $\partial\hat{h}_{i+1}$, respectively, which are incident with ξ_i and ξ_{i+1} in the obvious way. Assume J_i and J_{i+2} do not intersect $\partial\hat{h}_{i+1}$ except at ξ_i and ξ_{i+1} , respectively, and P intersects J_{i+1} only at the segment J_{i+1} . Let I_i (resp. I_{i+2}) be an open interval on $\partial\hat{h}_{i+1}$ which is disjoint from P and one of its endpoints is ξ_i (resp. ξ_{i+1}).*

Then J_i and J_{i+2} are on the same side of \hat{h}_{i+1} if and only if I_i and I_{i+2} have the same parity with respect to P . \square

In other words, if we denote by $\Delta\pi(J_{i+1})$ the difference of parities between I_i and I_{i+2} with respect to P and by $\Delta\pi_{\hat{h}}(J_{i+1})$ the difference of parities between J_i and J_{i+2} with respect to the parity function $\pi_{\hat{h}}$ of the grid defined by \hat{h} (see Example 7.2), then the previous corollary shows that $\Delta\pi(J_{i+1}) = \Delta\pi_{\hat{h}}(J_{i+1})$.

Theorem 10.3 *Let J be a Jordan curve. There are two points on which the associated parity function π_J takes different values. In particular, J separates ∂X into at least two components.*

Proof Let ζ_1 and ζ_2 be two distinct points on J , and let A and B be the two arcs on J that connect ζ_1 to ζ_2 . Let \hat{h} be a hyperplane that separates ζ_1 and ζ_2 .

Let r be the distance between $A \cap \partial\hat{h}$ and $B \cap \partial\hat{h}$, and let $\epsilon < \frac{r}{2}$. By Lemma 8.8, let P be an approximating PH curve for J such that the parity functions π_J and π of J and P , respectively, coincide outside an ϵ -neighborhood of J or P , and such that P intersects \hat{h} along segments of P . It suffices to find two points on $\partial\hat{h}$ which are at distance ϵ from P with different parity with respect to π . By abuse of notation we denote by A and B the corresponding approximating arcs on P .

We partition $\partial\hat{h}$ into intervals $A_1, B_1, A_2, B_2, \dots, A_m, B_m$ such that the endpoints of each A_i and B_i are at distance ϵ from P , and each A_i intersects J only in A and B_i intersects J only in B .

We denote by $\pi_{\hat{h}}$ the parity function defined by the hyperplane \hat{h} . The difference $\Delta\pi_{\hat{h}}(A) = \Delta\pi_{\hat{h}}(B) = \pi_{\hat{h}}(\zeta_1) - \pi_{\hat{h}}(\zeta_2) = 1$ since ζ_1 and ζ_2 are on different sides of \hat{h} .

On the other hand, we can compute it using Corollary 7.6 and the notation introduced above, and get

$$1 = \Delta\pi_{\hat{\sigma}}(A) = \sum_{j=1}^m \sum_{J_i \subseteq A_j} \Delta\pi_{\hat{\sigma}}(J_i),$$

where the second sum runs over all segments J_i of P which are contained in A_j .

By Corollary 10.2, we can replace the sum by

$$1 = \sum_{j=1}^m \sum_{J_i \subseteq A_j} \Delta\pi(J_i) = \sum_{j=1}^m \Delta\pi(A_j),$$

where $\Delta\pi(A_j)$ denotes the difference in the parity of the two endpoints of A_j with respect to π .

From this it follows that one of the intervals A_j has $\Delta\pi(A_j) = 1$, and thus its endpoints have different parity, as desired. \square

References

- [1] **I Agol**, *The virtual Haken conjecture*, Doc. Math. 18 (2013) 1045–1087 MR
- [2] **W L Ayres**, *Concerning continuous curves in metric space*, Amer. J. Math. 51 (1929) 577–594 MR
- [3] **N Bergeron, D T Wise**, *A boundary criterion for cubulation*, Amer. J. Math. 134 (2012) 843–859 MR
- [4] **M Bestvina, G Mess**, *The boundary of negatively curved groups*, J. Amer. Math. Soc. 4 (1991) 469–481 MR
- [5] **R H Bing**, *The Kline sphere characterization problem*, Bull. Amer. Math. Soc. 52 (1946) 644–653 MR
- [6] **B H Bowditch**, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. 180 (1998) 145–186 MR
- [7] **S G Brick**, *Quasi-isometries and ends of groups*, J. Pure Appl. Algebra 86 (1993) 23–33 MR
- [8] **MR Bridson, A Haefliger**, *Metric spaces of non-positive curvature*, Grundle. Math. Wissen. 319, Springer (1999) MR
- [9] **J W Cannon**, *The theory of negatively curved spaces and groups*, from “Ergodic theory, symbolic dynamics, and hyperbolic spaces” (T Bedford, M Keane, C Series, editors), Oxford Univ. Press (1991) 315–369 MR

- [10] **P-E Caprace, M Sageev**, *Rank rigidity for CAT(0) cube complexes*, *Geom. Funct. Anal.* 21 (2011) 851–891 MR
- [11] **R Gitik, M Mitra, E Rips, M Sageev**, *Widths of subgroups*, *Trans. Amer. Math. Soc.* 350 (1998) 321–329 MR
- [12] **F Haglund, D T Wise**, *Special cube complexes*, *Geom. Funct. Anal.* 17 (2008) 1551–1620 MR
- [13] **T C Hales**, *Jordan’s proof of the Jordan curve theorem*, *Studies in Logic, Grammar and Rhetoric* 10 (2007) 45–60
- [14] **J Kahn, V Markovic**, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, *Ann. of Math.* 175 (2012) 1127–1190 MR
- [15] **I Kapovich, N Benakli**, *Boundaries of hyperbolic groups*, from “Combinatorial and geometric group theory” (S Cleary, R Gilman, A G Myasnikov, V Shpilrain, editors), *Contemp. Math.* 296, Amer. Math. Soc., Providence, RI (2002) 39–93 MR
- [16] **N Lazarovich**, *On regular CAT(0) cube complexes*, preprint (2014) arXiv
- [17] **V Markovic**, *Criterion for Cannon’s conjecture*, *Geom. Funct. Anal.* 23 (2013) 1035–1061 MR
- [18] **M Sageev**, *Ends of group pairs and non-positively curved cube complexes*, *Proc. London Math. Soc.* 71 (1995) 585–617 MR
- [19] **M Sageev**, *CAT(0) cube complexes and groups*, from “Geometric group theory” (M Bestvina, M Sageev, K Vogtmann, editors), *IAS/Park City Math. Ser.* 21, Amer. Math. Soc., Providence, RI (2014) 7–54 MR
- [20] **H Short**, *Quasiconvexity and a theorem of Howson’s*, from “Group theory from a geometrical viewpoint” (E Ghys, A Haefliger, A Verjovsky, editors), *World Sci. Publ.*, River Edge, NJ (1991) 168–176 MR
- [21] **G T Whyburn**, *On the cyclic connectivity theorem*, *Bull. Amer. Math. Soc.* 37 (1931) 429–433 MR
- [22] **R L Wilder**, *Topology of manifolds*, *American Mathematical Society Colloquium Publications* 32, Amer. Math. Soc., New York (1949) MR

*Department of Mathematics, Haifa University
Jerusalem, Israel*

*Department of Mathematics, ETH Zürich
Zürich, Switzerland*

beeker@tx.technion.ac.il, nir.lazarovich@math.ethz.ch

Proposed: Martin Bridson
Seconded: Walter Neumann, Peter Teichner

Received: 14 February 2016
Revised: 16 February 2017