

Diagram rigidity for geometric amalgamations of free groups.

Jean-François Lafont

September 16, 2006

Abstract

In this paper, we show that a class of 2-dimensional locally CAT(-1) spaces is topologically rigid: isomorphism of the fundamental groups is equivalent to the spaces being homeomorphic. An immediate application of this result is a diagram rigidity theorem for certain amalgamations of free groups. The direct limit of two such amalgamations are isomorphic if and only if there is an isomorphism between the respective diagrams.

1 Introduction.

In this paper, we introduce the notion of a *geometric amalgamation* of free groups. This is a class of diagrams of groups, with the property that they are *rigid* in the following sense:

Theorem 1.1 (Diagram rigidity). *Let $\mathcal{D}_1, \mathcal{D}_2$ be a pair of geometric amalgamations of free groups. Then $\varinjlim \mathcal{D}_1$ is isomorphic to $\varinjlim \mathcal{D}_2$ if and only if \mathcal{D}_1 is isomorphic to \mathcal{D}_2 (as diagrams of groups).*

In order to prove this theorem, we will first translate the question to a more topological setting. Associated to any geometric amalgamation of free groups, there is a canonically defined topological space, which we call a *simple, thick, 2-dimensional hyperbolic P-manifold*. The associated space will have fundamental group isomorphic to the direct limit of the geometric amalgamation, and has the property that the diagram can be “read off” from the topology of the space. The first theorem will then be a consequence of the following purely topological result:

Theorem 1.2 (Topological rigidity). *Let X_1, X_2 be a pair of simple, thick, 2-dimensional hyperbolic P-manifolds, and assume that $\phi : \pi_1(X_1) \rightarrow \pi_1(X_2)$ is an*

isomorphism. Then there exists a homeomorphism $\Phi : X_1 \rightarrow X_2$ that induces ϕ on the level of the fundamental groups.

Two consequences of this second theorem are the following:

Corollary 1.1 (Nielsen realization). *Let X be a simple, thick, 2-dimensional hyperbolic P -manifold. Then the canonical map from $\text{Homeo}(X)$ to $\text{Out}(\pi_1(X))$ is surjective.*

Corollary 1.2 (Weak co-Hopf property). *Let \mathcal{D} be a geometric amalgamation of free groups. Then $\varinjlim \mathcal{D}$ is weakly co-Hopfian, i.e. every injection $\varinjlim \mathcal{D} \hookrightarrow \varinjlim \mathcal{D}$ with image of finite index is in fact an isomorphism.*

We note that the second corollary also follows from a considerably more sophisticated result of Sela [11], who proved that a non-elementary torsion-free δ -hyperbolic groups is co-Hopfian if and only if it is freely indecomposable. For our groups, a simple geometric argument (see Section 3) gives the weaker conclusion.

The proof of Theorem 1.2 relies on a topological characterization of certain points in the boundary at infinity of the universal cover of a simple, thick, 2-dimensional hyperbolic P -manifold (these spaces are $\text{CAT}(-1)$).

Acknowledgements: The author would like to thank the anonymous referee for pointing out a gap in the preliminary version of this paper, as well as for catching various mathematical inaccuracies.

2 Some preliminaries.

In this section, we briefly define the various notions that are relevant to this paper, and recall some basic facts that will be used in the proofs.

Definition 2.1. For our purposes, a *diagram of groups* \mathcal{D} will consist of:

- a finite connected directed graph, which we will also denote by \mathcal{D} , with vertex set $V(\mathcal{D})$ and directed edges $E(\mathcal{D})$,
- an assignment of a group G_v to each vertex $v \in V(\mathcal{D})$
- an assignment of a homomorphism $\phi_e : G_{e_-} \rightarrow G_{e_+}$ to each directed edge $e \in E(\mathcal{D})$, where e_- and e_+ denote the initial and terminal endpoints of the directed edge e .


We will denote by $\varinjlim \mathcal{D}$ the direct limit of the diagram \mathcal{D} .

Observe that the above definition differs superficially from the notion of a graph of groups used in Bass-Serre theory (see for instance Serre [12]). The interested reader can easily translate the above definition into the language of Bass-Serre theory. We now refine the above definition to the diagrams we are really interested in:

Definition 2.2. We will say that a diagram of groups \mathcal{D} is a *geometric amalgamation of free groups* provided it satisfies the following properties:

- the vertex set $V(\mathcal{D})$ can be partitioned into $V_0(\mathcal{D})$ and $V_1(\mathcal{D})$, and every directed edge $e \in E(\mathcal{D})$ has $e_- \in V_0(\mathcal{D})$, $e_+ \in V_1(\mathcal{D})$. Furthermore, each vertex in $V_0(\mathcal{D})$ has degree at least three.
- the group associated to every $v \in V_0(\mathcal{D})$ is isomorphic to \mathbb{Z} , and the group associated to every $w \in V_1(\mathcal{D})$ is a free group with rank ≥ 2 .
- associated to every $w \in V_1(\mathcal{D})$, there is a compact surface M_w whose fundamental group is G_w .
- each edge morphism ϕ_e maps $G_{e_-} \cong \mathbb{Z}$ isomorphically onto a group conjugate to the fundamental group of a boundary component in $G_{e_+} \cong \pi_1(M_{e_+})$.
- for every vertex $w \in V_1(\mathcal{D})$, and every conjugacy class of \mathbb{Z} -subgroups of $G_w \cong \pi_1(M_w)$ corresponding to a boundary component of M_w , there is precisely one edge $e \in E(\mathcal{D})$ with $e_+ = w$, and $\phi_e(G_{e_-})$ lying within the conjugacy class.

More concisely, we can think of a geometric amalgamation of free groups as being a diagram of groups consisting of two rows:



where each \mathbb{Z} group in the bottom row injects into at least three of the free groups in the top row, and has image lying in a “boundary subgroup” of the free groups in the top row. In addition, each “boundary subgroup” in the top row lies in the image of precisely one \mathbb{Z} from the bottom row (up to conjugacy).

Now the main motivation behind our terminology lies in that the direct limit of a geometric amalgamation of free groups naturally corresponds to the fundamental group of an associated topological space. We now proceed to define these spaces.

Definition 2.3. We say that a compact geodesic metric space X is a simple, thick, 2-dimensional hyperbolic P-manifold, provided that there exists a closed subset $Y \subset X$ with the property that:

- each connected component of Y is homeomorphic to S^1 , and forms a totally geodesic subspace of X (simplicity hypothesis).
- the closure of each connected component of $X - Y$ is homeomorphic to a compact surface with boundary, and the homeomorphism takes the component of $X - Y$ to the interior of the surface with boundary; the closure of such a component will be called a *chamber*.
- there exists a hyperbolic Riemannian metric on each chamber which coincides with the original metric.
- each connected component of Y lies in at least three distinct chambers (thickness hypothesis).

We will call the subset Y the *branching locus*, and will call the connected components of Y *branching geodesics*.

In the previous definition, the thickness hypothesis can be viewed as a non-triviality hypothesis. The simplicity hypothesis ensures that the codimension one strata Y is not too complicated. In the general definition of a 2-dimensional hyperbolic P-manifold (in [5]), the codimension one strata can be an arbitrary metric graph (as opposed to a disjoint union of circles).

Observe that simple, thick, 2-dimensional hyperbolic P-manifolds are locally CAT(-1) (see [2]), and hence their universal covers are CAT(-1) spaces. In particular, this implies that their fundamental groups are δ -hyperbolic groups, and that an abstract isomorphism between the fundamental groups of two such spaces naturally induces a quasi-isometry between their universal covers.

To make precise the correspondence between the previous two definitions, we show the following:

Lemma 2.1. *A group G is the fundamental group of a simple, thick, 2-dimensional hyperbolic P-manifold if and only if it is the direct limit of a geometric amalgamation of free groups.*

Proof. If G is the fundamental group of a simple, thick, 2-dimensional hyperbolic P-manifold X , we consider an open cover $\{\mathcal{U}_i\}$ of X by open sets which consist of ϵ -neighborhoods of the chambers of X , where ϵ is chosen to be small enough. An immediate application of the general form of the Siefert-Van Kampen Theorem (see Chapter 2, Section 7 in May [7]) that the fundamental group of X is isomorphic to the direct limit of a diagram of groups obtained from the covering $\{\mathcal{U}_i\}$. Furthermore, the diagram has vertex groups isomorphic to the fundamental groups of the various intersections of open sets in the covering $\{\mathcal{U}_i\}$, with edge morphisms induced by

inclusions. Note that the only non-trivial intersections have fundamental group \mathbb{Z} , one arising from each of the components of the branching locus. Furthermore, the second bullet in Definition 2.3 guarantees that, for each chamber, all the various boundary circles are part of the branching locus. This ensures that the last two conditions in Definition 2.2 are satisfied. In particular, we see that for the covering we've defined, the resulting diagram is a geometric amalgamation of free groups. Furthermore, the diagram is uniquely defined by the space X .

Conversely, assume that G is the direct limit of a geometric amalgamation of free groups, denoted by \mathcal{D} . Corresponding to the diagram \mathcal{D} , we can associate a diagram of topological spaces by associating to each vertex in $v \in V_0$ an S^1 , and to each vertex $w \in V_1$ the corresponding compact surface with boundary M_w . Note that to each edge, there corresponds a homeomorphism from one of the S^1 (corresponding to the initial vertex of the edge) to a boundary component of one of the M_w (corresponding to the terminal vertex of the edge). This homeomorphism maps the S^1 to the unique boundary component $S^1 \subset \partial M_w$ having the property that the image of the group \mathbb{Z} is conjugate to the subgroup $\pi_1(S^1) \subset \pi_1(M_w)$ (uniqueness of such a component follows from the fact that M_w supports a hyperbolic metric). Now consider the direct limit of this diagram of spaces in the category of topological spaces. It is immediate that this direct limit is a simple, thick, 2-dimensional P-manifold X , and, by the discussion in the previous paragraph, that $\pi_1(X) \cong G$.

To conclude, we need to show that X supports a hyperbolic metric. To see this, we make each S^1 isometric to the unit circle in \mathbb{R}^2 , and make each M_w isometric to a compact hyperbolic surface with all boundary components totally geodesic of length 2π . We further require the homeomorphisms from the S^1 to the boundary components of the M_w to be isometries. This immediately yields a hyperbolic metric on the P-manifold X .

Finally, we point out that the space X constructed above is unique up to homeomorphism. This follows from the fact that if M is a compact orientable surface with boundary, and $\phi : \partial M \rightarrow \partial M$ is an orientation preserving self-homeomorphism, then there is a self-homeomorphism $\hat{\phi} : M \rightarrow M$ that induces ϕ when restricted to the boundary. In particular, the choice of homeomorphisms used to identify the various S^1 (corresponding to the \mathbb{Z} groups) with the boundary components of the M_w (corresponding to the free groups) does not influence the topology of the resulting space.

Finally, to conclude this section, we reduce the proof of Theorem 1.1 to that of Theorem 1.2:

Proof (Diagram Rigidity). Let us start with a pair $\mathcal{D}_1, \mathcal{D}_2$ of geometric amalgamations of free groups, and assume that $\varinjlim \mathcal{D}_1$ is isomorphic to $\varinjlim \mathcal{D}_2$. From the previous Lemma, we can associate to each \mathcal{D}_i a simple, thick, 2-dimensional hyperbolic P-

manifold X_i , with the property that $\pi_1(X_i) \cong \varinjlim \mathcal{D}_i$. In particular, the isomorphism between direct limits yields an isomorphism $\phi : \pi_1(X_1) \rightarrow \pi_1(X_2)$. From the topological rigidity result, there is a homeomorphism $\Phi : X_1 \rightarrow X_2$ which induces ϕ on the level of fundamental groups.

Now note that Φ , being a homeomorphism, must map the branching locus in X_1 homeomorphically to the branching locus in X_2 , and hence maps the chambers of X_1 homeomorphically to the chambers of X_2 . Furthermore, the map Φ induces a bijection between the chambers in X_1 and those in X_2 . But by the uniqueness portion of the Lemma above, this implies that the diagrams G_1 and G_2 are isomorphic, concluding our proof.

3 Topological rigidity.

In this section, we provide a proof of Theorem 1.2. Let us start by fixing some notation. X will always denote a simple, thick, 2-dimensional hyperbolic P-manifolds, \tilde{X} the universal cover of X , and $\partial^\infty \tilde{X}$ the boundary at infinity of \tilde{X} . We will let Γ denote the fundamental group of X . \mathcal{G} will denote the collection of geodesics in \tilde{X} , and $\mathcal{BG} \subset \mathcal{G}$ will denote the collection of lifts of branching geodesics in \tilde{X} . We will let \mathcal{C} denote the collection of lifts of chambers in \tilde{X} . Finally, we will let $\partial^\infty \mathcal{BG} \subset \partial^\infty \tilde{X}$ be the collection of points of the form $\gamma(\pm\infty)$ where $\gamma \in \mathcal{BG}$. In the portions of this section where we deal with a pair of simple, thick, 2-dimensional hyperbolic P-manifolds, we will append subscripts to keep track of which of the two spaces we are referring to.

The first step in our argument consists of identifying the separation properties of individual points in $\partial^\infty \tilde{X}$. We start with an easy:

Lemma 3.1. *The boundary at infinity $\partial^\infty \tilde{X}$ is path-connected.*

Proof. To see this, let $p_1, p_2 \in \partial^\infty \tilde{X}$ be an arbitrary pair of points. Let γ be a geodesic in \tilde{X} satisfying $\gamma(-\infty) = p_1$, $\gamma(\infty) = p_2$. The simplicity and thickness hypotheses on X ensure that there exists an isometrically embedded $f : \mathbb{H}^2 \hookrightarrow \tilde{X}$ with the property that $\gamma \subset \mathbb{H}^2$. In particular, we see that the pair of points p_1, p_2 lie on an embedded $S^1 = \partial^\infty \mathbb{H}^2 \hookrightarrow \partial^\infty \tilde{X}$. Hence there exists a path in $\partial^\infty \tilde{X}$ joining p_1 to p_2 , concluding the proof of the Lemma.

Let us recall some basic definitions. A subset S in a topological space X is said to *locally separate* provided there exists a neighborhood N of S with the property that $N - S$ is disconnected. If there exists a neighborhood N such that $N - S$ consists of $\geq m$ connected components, we say that S locally separates into $\geq m$ components. We say that S locally separates into M components provided M is the supremum of

the integers m with the property that S locally separates into $\geq m$ components. We will be interested in the case where S consists either of a single point, or of a pair of points.

In the case where $S = \{x\}$, and S locally separates the space X , we will say that x is a *local cutpoint* of X . If in addition $X - \{x\}$ is disconnected, we will say that the point x is a *global cutpoint* of X .

We observe that the previous Lemma in particular implies that $\partial^\infty \tilde{X}$ is connected. But Bowditch [1] and Swarup [13] showed that the boundary of a one-ended hyperbolic group has no global cutpoints, which immediately gives:

Corollary 3.1. *The boundary at infinity $\partial^\infty \tilde{X}_i$ does not contain any global cutpoints.*

We note that the previous Corollary tells us that the *global* separation properties of individual points in $\partial^\infty \tilde{X}_i$ are uninteresting. On the other hand, the *local* separation properties of points in $\partial^\infty \tilde{X}_i$ are quite interesting. Our next step is to consider a notion which is slightly weaker than “local separation into ≥ 3 components”.

We now collect some basic facts concerning separation and connectedness properties in simple, thick, hyperbolic P-manifolds. The proofs of the following four lemmas can be found in [5] in the 3-dimensional setting, but the arguments given there extend verbatim to the 2-dimensional setting.

Lemma 3.2 (Lemma 2.1 in [5]). *Let $\gamma \in \mathcal{BG}$ be a branching geodesic in \tilde{X} , and let $C_1, C_2 \in \mathcal{C}$ be two lifts of chambers which are both incident to γ . Then $C_1 - \gamma$ and $C_2 - \gamma$ lie in different connected components of $\tilde{X} - \gamma$.*

Lemma 3.3 (Lemma 2.2 in [5]). *Let $C \in \mathcal{C}$ be a lift of a chamber, and let $\gamma_1, \gamma_2 \in \mathcal{BG}$ be two branching geodesics which are both incident to C . Then γ_1 and γ_2 lie in different connected components of $\tilde{X} - \text{Int}(C)$.*

Lemma 3.4 (Lemma 2.3 in [5]). *Let $\{\gamma(\pm\infty)\}$ be the pair of points in $\partial^\infty \tilde{X}$ corresponding to some $\gamma \in \mathcal{BG}$, and let $\partial^\infty C_1, \partial^\infty C_2$ be the boundaries at infinity of two lifts of chambers $C_1, C_2 \in \mathcal{C}$ which are both incident to γ . Then $\partial^\infty C_1 - \{\gamma(\pm\infty)\}$ and $\partial^\infty C_2 - \{\gamma(\pm\infty)\}$ lie in different connected components of $\partial^\infty \tilde{X} - \{\gamma(\pm\infty)\}$.*

Lemma 3.5 (Lemma 2.4 in [5]). *Let $\partial^\infty C$ be the boundary at infinity corresponding to a connected lift of a chamber $C \in \mathcal{C}$, and let $\{\gamma_1(\pm\infty)\}, \{\gamma_2(\pm\infty)\}$ be the boundary at infinity of two branching geodesics $\gamma_1, \gamma_2 \in \mathcal{BG}$ which are both incident to C . Then $\{\gamma_1(\pm\infty)\}$ and $\{\gamma_2(\pm\infty)\}$ lie in different connected components of $\partial^\infty \tilde{X} - (\partial^\infty C - \cup \partial^\infty \eta_i)$, where the union is over all $\eta_i \in \mathcal{BG}$ which are boundary components of C .*

We define the *tripod* to be the space T obtained by taking the cone of a 3-point set. The cone point will be denoted by $* \in T$. We say that a point x is a *branching point*

in a topological space Y provided there exists an injective map $f : T \rightarrow Y$ satisfying $f(*) = x$. Note that this notion is uninteresting in high-dimensional spaces, as every point will be branching. The relevance of the notion of branching in the setting we are looking at is due to the following observation: if a point x in a geodesic space Y locally separates into ≥ 3 components, then the point x is a branching point in Y . The notion of branching was introduced by the author in [5], [6] in order to study the local topology of the boundary at infinity of simple, thick hyperbolic P-manifolds. The proof of the following Proposition closely parallels the arguments given in the paper [5]:

Proposition 3.1. *For a point $p \in \partial^\infty \tilde{X}$, we have that p is branching if and only if $p = \gamma(\infty)$ for some branching geodesic $\gamma \in \mathcal{BG}$.*

Proof. We first observe that one implication in the Proposition is immediate: if $p = \gamma(\infty)$ for some branching geodesic $\gamma \in \mathcal{BG}$, then p is branching. So let us focus on the converse.

Assume that $p \in \partial^\infty \tilde{X}$ is a branching point, and that p is not a limit point at infinity of any branching geodesic. Let γ be a geodesic ray with $\gamma(\infty) = p$, and observe that for the geodesic ray γ we have either:

1. γ passes through finitely many lifts of chambers, or
2. γ passes through infinitely many lifts of chambers.

For each of the two cases above, we need to show that p cannot be branching. We argue by contradiction. Assume that $f : T \rightarrow \partial^\infty \tilde{X}$ be an injective mapping of the tripod into $\partial^\infty \tilde{X}$, satisfying $f(*) = p$.

Case 1: Since γ passes through finitely many lifts of chambers, and since the lifts of chambers are totally geodesic subsets of \tilde{X} , we have that there exists a fixed lift C of a chamber with the property that $\gamma(t) \in C$ for all t sufficiently large. Now pick a basepoint x in the *interior* of C , and pick ϵ small enough so that the ϵ metric ball $B_\epsilon(x)$ centered at x is contained entirely in the interior of C . Denote by $lk(x)$ the boundary of this metric ball, and observe that $lk(x)$ is homeomorphic to S^1 . Let $\rho : \partial^\infty \tilde{X} \rightarrow lk(x)$ be the geodesic projection, and consider the composite map $\rho \circ f : T \rightarrow lk(x) \cong S^1$.

We first note that there is no injective map from T to S^1 , hence the composite $\rho \circ f$ must fail to be injective at some point. Let $I \subset S^1$ be the subset of points in $lk(x)$ where the map ρ is injective. Our goal is now to show that the composite $\rho \circ f$ fails to be injective at some point $z \in I$. This immediately implies that f fails to be injective at the point $\rho^{-1}(z)$, which would yield the desired contradiction.

Let us start by observing that I consists of a Cantor set in S^1 . Indeed, if a point w lies in the complement of I , then there exist a pair of geodesic rays γ_1, γ_2 emanating from x , both of which pass through $w \in lk(x)$, but satisfying $\gamma_1(\infty) \neq \gamma_2(\infty)$. In particular, the geodesic rays γ_1, γ_2 must coincide for a period of time, and subsequently diverge. This implies that $\gamma_i \cap \partial C \neq \emptyset$. Hence the point w lies in the image $\bar{\rho}(\partial C)$ of ∂C under the geodesic retraction map $\bar{\rho} : \tilde{X} - B_\epsilon(x) \rightarrow lk(x)$. Conversely, given a point w in $\bar{\rho}(\partial C)$, one can easily construct a pair of geodesic rays γ_1, γ_2 originating from x , passing through w , but having $\gamma_1(\infty) \neq \gamma_2(\infty)$. This forces the complement of I to coincide with the set $\bar{\rho}(\partial C)$. But the complement of the set $\bar{\rho}(\partial C)$ can naturally be identified with $\partial^\infty C$. Since C is the universal cover of a compact negatively curved surface with non-empty boundary, it is quasi-isometric to a free group F_k . This implies that I is homeomorphic to $\partial^\infty F_k$, which is known to be a Cantor set.

Now observe that the complement of the set $I \subset lk(x)$ consists of a countable dense union of open intervals. Let $I_\partial \subset I$ denote the subset of I consisting of the boundary points of these intervals. Note that, by the discussion above, the set I_∂ coincides with the set $\rho(\partial^\infty(\partial C))$, and since we are assuming that the point $p \in \partial^\infty X$ is not the limit point at infinity of a branching geodesic, we have that $(\rho \circ f)(*) = \rho(p) \in I - I_\partial$.

Let $L_i \cong [0, 1)$ ($1 \leq i \leq 3$) denote the three components of $T - *$, which we will call the *open leaves* of the tripod T . Since $(\rho \circ f)(*) \in I$, we have that $(\rho \circ f)(*) \notin (\rho \circ f)(L_i)$ for each i . Let U denote a small open connected neighborhood of $(\rho \circ f)(*)$ in $lk(x) \cong S^1$, and observe that $(\rho \circ f)(*)$ locally separates U into a pair of open intervals U_1, U_2 . If U is chosen small enough, we must have that a pair of leaves surjects onto one of the U_j . We assume, without loss of generality that $U_1 \subset (\rho \circ f)(L_1) \cap (\rho \circ f)(L_2)$. But now we note that in the Cantor set I , every point in $I - I_\partial$ can be approximated on both sides by points in I_∂ . This implies that $[(\rho \circ f)(L_1) \cap (\rho \circ f)(L_2)] \cap I_\partial \neq \emptyset$, and as ρ is injective on I_∂ , that there exist a pair of points $q_1 \in L_1, q_2 \in L_2$ with $f(q_1) = f(q_2) \in \partial^\infty \tilde{X}$. But this contradicts the fact that f is injective, concluding the proof for the first case.

Case 2: For the second case, we assume that γ is a geodesic ray with $\gamma(\infty) = p$, and which passes through infinitely many lifts of chambers. Note that this forces the geodesic ray γ to intersect infinitely many branching geodesics. Let $\{\eta_i\}$ be the collection of branching geodesics intersected by γ , indexed in the order in which their intersections occur along γ . We now recall two facts:

1. each $\partial^\infty \eta_i$ separates $\partial^\infty \tilde{X}$ (Lemma 3.4 above), and
2. if U_i denotes the path-connected component of $\partial^\infty \tilde{X} - \partial^\infty \eta_i$ containing p , then the collection $\{U_i\}$ forms an open, path-connected, neighborhood base of p in

$\partial^\infty \tilde{X}$ (see the proof of Proposition 2.3 in [5]).

Armed with these two facts, the argument for Case 2 is easy: let $L_i \cong [0, 1)$ again denote the three open leaves of the tripod. Since $f : T \rightarrow \partial^\infty \tilde{X}$ is injective, we have that $p \notin f(\partial T)$, and hence there exists a small enough neighborhood N of p with the property that $f(\partial T) \subset \partial^\infty \tilde{X} - \bar{N}$. From Fact (2) above, we have that $U_i \subset N$ for i sufficiently large, and hence that $f(\partial T) \subset \partial^\infty \tilde{X} - \bar{N}$. From Fact (1) above, we have that the corresponding $\partial^\infty \eta_i = \{\eta_i(\pm\infty)\}$ separates $f(\partial T)$ from p . But this implies that, for each $1 \leq i \leq 3$, we have that $f(L_i) \cap \{\eta_i(\pm\infty)\} \neq \emptyset$. The pigeon-hole principle forces the image of a pair of leaves to pass through one of the two points $\{\eta_i(\pm\infty)\}$. But this contradicts the fact that f is injective, concluding the argument for Case 2, and completing the proof of the Proposition.

We observe an immediate corollary of the above proposition:

Corollary 3.2. *If $f : \tilde{X}_1 \rightarrow \tilde{X}_2$ is a quasi-isometry, and $f_\infty : \partial^\infty \tilde{X}_1 \rightarrow \partial^\infty \tilde{X}_2$ the induced map on the boundary at infinity. Then f_∞ restricts to a bijection from $\partial^\infty \mathcal{BG}_1$ to $\partial^\infty \mathcal{BG}_2$.*

Proof. Note that the induced map $f_\infty : \partial^\infty \tilde{X}_1 \rightarrow \partial^\infty \tilde{X}_2$ is a homeomorphism. But the previous Proposition characterizes the subsets $\partial^\infty \mathcal{BG}_i \subset \partial^\infty \tilde{X}_i$ purely topologically, yielding the corollary.

We would now like to focus on the specific situation at hand, namely we will assume that we are given a pair X_1, X_2 of simple, thick, 2-dimensional hyperbolic P-manifolds, and an abstract isomorphism ϕ from $\Gamma_1 := \pi_1(X_1)$ to $\Gamma_2 := \pi_1(X_2)$. We observe the following facts that hold in this setting:

- the groups Γ_i act by homeomorphisms on $\partial^\infty \tilde{X}_i$,
- the isomorphism ϕ induces a quasi-isometry $\bar{\phi}$ from \tilde{X}_1 to \tilde{X}_2 ,
- the quasi-isometry $\bar{\phi}$ induces a homeomorphism $\partial^\infty \bar{\phi} : \partial^\infty \tilde{X}_1 \rightarrow \partial^\infty \tilde{X}_2$ which is equivariant with respect to the Γ_i actions on the $\partial^\infty \tilde{X}_i$.

The previous corollary tells us that $\partial^\infty \bar{\phi}$ restricts to a bijection between from the set $\partial^\infty \mathcal{BG}_1$ to the set $\partial^\infty \mathcal{BG}_2$. Now notice that any branching geodesic $\gamma \in \mathcal{BG}_1$ naturally corresponds to a pair of points $\{\gamma(\pm\infty)\} \subset \partial^\infty \mathcal{BG}_1$. We would like to ensure that, under our homeomorphism ϕ_∞ , the pair $\{\gamma(\pm\infty)\}$ maps to a pair $\{\gamma'(\pm\infty)\}$ for some branching geodesic $\gamma' \in \mathcal{BG}_2$. In order to achieve this, our next step is to characterize the endpoints of branching geodesics in a purely topological manner. This is the content of our:

Proposition 3.2. *Let $\{p, q\} \subset \partial^\infty \mathcal{BG} \subset \partial^\infty \tilde{X}$ be an arbitrary pair of distinct points. Then we have that:*

1. *if there exists a $\gamma \in \mathcal{BG}$ with the property that $\{\gamma(\pm\infty)\} = \{p, q\}$, then $\{p, q\}$ separates $\partial^\infty \tilde{X}$ into ≥ 3 components.*
2. *if there exists a geodesic γ contained in the interior of a single lift of a chamber, with the property that $\{\gamma(\pm\infty)\} = \{p, q\}$, then $\{p, q\}$ separates $\partial^\infty \tilde{X}$ into exactly 2 components.*
3. *in all other cases, $\{p, q\}$ does not separate $\partial^\infty \tilde{X}$.*

Proof. Statement (1) is an immediate consequence of Lemma 3.4 and the thickness hypothesis.

To see statement (2), one starts with a $\gamma \notin \mathcal{BG}$, and satisfying $\gamma \subset \text{Int}(C)$ with $C \in \mathcal{C}$. Note that this implies that γ separates C into precisely two open components, denoted Z_1 and Z_2 . Furthermore, the closure of each component is a closed, totally geodesic subset of \tilde{X} . Now for $i = 1, 2$, define the sets $(Z_i)_j$ ($j \geq 1$) inductively by:

- the initial condition $(Z_i)_1 = Z_i$, and
- $(Z_i)_{j+1}$ is the union of $(Z_i)_j$ along with all lifts of chambers which are incident to $(Z_i)_j$.

We observe that we have proper inclusions $(Z_i)_j \subset (Z_i)_{j+1}$, and that each of the subsets $(Z_i)_j$ is totally geodesic and path-connected.

Now form the sets $Y_i := \bigcup_{j \in \mathbb{N}} (Z_i)_j$, and observe that each Y_i is a path-connected, totally geodesic subspace of \tilde{X} (as the latter properties are preserved under increasing unions). Furthermore each of the sets Y_i has the property that $\partial^\infty Y_i - \{\gamma(\pm\infty)\}$ is path-connected. Indeed, given a pair of points in $\partial^\infty Y_i$, one can consider the geodesic η corresponding to the pair of points. It is easy to see that there is an isometrically embedded, totally geodesic “half- \mathbb{H}^2 ” $H \subset Y_i$ with boundary the given geodesic η . This implies that there exists $\partial^\infty H \cong \mathbb{I} \subset \partial^\infty Y_i - \{\gamma(\pm\infty)\}$ whose endpoints correspond precisely to $\{\eta(\pm\infty)\}$.

Finally, observe that $\tilde{X} - \gamma = Y_1 \coprod Y_2$, and that the closure of Y_i is precisely $Y_i \cup \gamma$. Hence we have that $\partial^\infty \tilde{X} = \partial^\infty Y_1 \cup_{\{\gamma(\pm\infty)\}} \partial^\infty Y_2$, expressing $\partial^\infty \tilde{X}$ as a union of a pair of closed sets (as each Y_i is totally geodesic) whose intersection is precisely $\{\gamma(\pm\infty)\}$, and with the property that each $\partial^\infty Y_i - \{\gamma(\pm\infty)\}$ is path-connected. This immediately implies statement (2) of the proposition.

So we are now left with showing statement (3). In order to do this, we first make two observations concerning branching geodesics. Note that if $\rho \in \mathcal{BG}$, then we have

that $\partial^\infty \tilde{X} - \{\rho(\pm\infty)\}$ splits into $k \geq 3$ path-connected components U_1, \dots, U_k . We now observe:

Fact 1: The closure of each U_i is $\bar{U}_i = U_i \cup \{\rho(\pm\infty)\}$, and is path-connected.

Fact 2: For every pair of distinct points $x, y \in Y_i$ there is a path $\eta_x \subset \bar{U}_i - y$ joining x to one of the points $\{\rho(\pm\infty)\}$.

Now assuming these two facts, we proceed with the proof of statement (3). For the points $\{p, q\}$ satisfying the hypotheses of statement (3), we have that the geodesic γ satisfying $\gamma(\pm\infty) = \{p, q\}$ must intersect a branching geodesic $\rho \in \mathcal{BG}$. Up to re-indexing, we may assume that $p \in U_1$, and $q \in U_2$.

Now pick an arbitrary pair of points $\{x, y\} \subset \partial^\infty \tilde{X} - \{p, q\}$, and consider the subsets U_i, U_j satisfying $x \in U_i, y \in U_j$. If $i \geq 3$, then from **Fact 1**, there exists a path in \bar{U}_i joining x to $\rho(\infty)$. On the other hand, if $i = 1, 2$, then from **Fact 2**, there is a path joining x to one of the points $\rho(\pm\infty)$, avoiding the point p (if $i = 1$) or q (if $i = 2$). In either case, denote this path by η_x . Now applying the same reasoning to y , we find a path η_y joining y to one of the points $\rho(\pm\infty)$, and avoiding the pair $\{p, q\}$.

If the endpoints of the paths η_x, η_y coincide, concatenation gives us a path connecting x to y . Otherwise, from **Fact 1**, we note that there is a path η_ρ in Y_3 joining $\rho(\infty)$ to $\rho(-\infty)$. Concatenating the three paths η_x, η_ρ , and η_y yields a path joining x to y in $\partial^\infty \tilde{X} - \{p, q\}$. Since this holds for arbitrary $x, y \in \partial^\infty \tilde{X} - \{p, q\}$, we conclude that $\partial^\infty \tilde{X} - \{p, q\}$ is path-connected. So to complete the proof of the Proposition, we are left with verifying **Fact 1** and **Fact 2**.

To see **Fact 1**, we first note that the complement of $U_i \cup \{\rho(\pm\infty)\}$ consists of the union $\coprod_{j \neq i} U_j$. Since all the U_j are open, this implies that the closure of U_i is contained in the set $U_i \cup \{\rho(\pm\infty)\}$. To see the converse, we observe that we can construct, as in the argument for statement (2), totally geodesic subspaces Y_i with the property that $\partial^\infty Y_i = U_i \cup \{\rho(\pm\infty)\}$ and with $\partial Y_i = \rho$. But within the Y_i , it is easy to see that there exist totally geodesic embedded ‘half’ \mathbb{H}^2 ’s whose boundary is precisely ρ . At the level of the boundary at infinity, this yields an embedded interval in $U_i \cup \{\rho(\pm\infty)\}$ with endpoints precisely $\{\rho(\pm\infty)\}$. This immediately implies that $\{\rho(\pm\infty)\}$ lies in the closure of the U_i , completing the proof of **Fact 1**.

To see **Fact 2**, we note that given $p \in U_i$, the embedded interval $f : \mathbb{I} \hookrightarrow U_i \cup \{\rho(\pm\infty)\}$ mentioned in the previous paragraph can be chosen to satisfy $f(0) = \rho(+\infty)$, $f(1) = \rho(-\infty)$, and $f(1/2) = p$. Now note that if $q \notin f(\mathbb{I})$, we are done. On the other hand, if $q \in f(\mathbb{I})$, then the hypothesis that $p \neq q$ ensures that $q = f(r)$ where either $0 < r < 1/2$ or $1/2 < r < 1$. In both cases we can use f restricted to a suitable subinterval of \mathbb{I} to get the desired path. This completes the proof of **Fact 2**, and hence, of Proposition 3.2.

Since separation properties are purely topological, we obtain the immediate corollary:

Corollary 3.3. *Every quasi-isometry $f : \tilde{X}_1 \rightarrow \tilde{X}_2$ naturally induces a bijective correspondence between \mathcal{BG}_1 and \mathcal{BG}_2 .*

Now observe that for the universal cover \tilde{X} of a simple, thick, 2-dimensional hyperbolic P-manifold, we can naturally define an adjacency relation on the set \mathcal{BG} . We say that a pair of elements γ_1, γ_2 of \mathcal{BG} are *adjacent*, denoted by $\gamma_1 \sim \gamma_2$ provided there exists a geodesic joining a point in γ_1 to a point in γ_2 , and *lying entirely within a single chamber*. Note that the above relation is symmetric, but *not* transitive. The next step is to establish that a quasi-isometry preserves the adjacency relation.

Proposition 3.3. *If $f : \tilde{X}_1 \rightarrow \tilde{X}_2$ is a quasi-isometry, then for any pair $\gamma_1, \gamma_2 \in \mathcal{BG}_1$, we have:*

$$\gamma_1 \sim \gamma_2 \iff f_*(\gamma_1) \sim f_*(\gamma_2),$$

where $f_*(\gamma)$ in \mathcal{BG}_2 is the branching geodesic bijectively associated with $\gamma \in \mathcal{BG}_1$.

Proof. This follows immediately from Proposition 3.2. Indeed, from the definition of the relation \sim , we see that $\gamma_1 \sim \gamma_2$ if and only if the geodesics form a pair of distinct boundary geodesics of a single chamber C . Given a 4-tuple of distinct points $\{x_-, x_+, y_-, y_+\} \subset \partial^\infty \tilde{X}$, Proposition 3.2 (parts (1) and (2)) tells us there is a pair γ_1, γ_2 satisfying $\gamma_1(\pm\infty) = x_\pm, \gamma_2(\pm\infty) = y_\pm$ if and only if we have:

- the two pairs of points $\{x_\pm\}, \{y_\pm\}$ each separate $\partial^\infty \tilde{X}$ into ≥ 3 components,
- each of the four pairs of points $\{x_+, y_+\}, \{x_+, y_-\}, \{x_-, y_+\}, \{x_-, y_-\}$, separate $\partial^\infty \tilde{X}$ into precisely two components.

Since the quasi-isometry f induces a homeomorphism f_∞ between the two boundaries at infinity $\partial^\infty \tilde{X}_1$ and $\partial^\infty \tilde{X}_2$, the above topological characterization of endpoints of adjacent branching geodesics immediately yield the Proposition.

Next observe that the adjacency relation can be used to keep track of the chambers. This is the content of the following:

Lemma 3.6. *Elements of \mathcal{C} correspond bijectively to maximal subsets of \mathcal{BG} on which the adjacency relation is transitive.*

Proof. Let $C \in \mathcal{C}$ be a chamber, and associate to it the collection of $\gamma \in \mathcal{BG}$ which arise as the boundary components of C ; denote this set by B_C . It is immediate from the definition of the relation \sim that the adjacency relation is transitive on B_C .

Conversely, let $B \subset \mathcal{BG}$ be a subset on which the adjacency relation is transitive. We claim that there is a $C \in \mathcal{C}$ satisfying $B \subset B_C$. To see this, pick $\gamma_1, \gamma_2, \gamma_3 \in B$, and observe that the condition $\gamma_1 \sim \gamma_2$ implies that the two branching geodesics are boundary components of a fixed chamber C_{12} . Similarly, $\gamma_2 \sim \gamma_3$ implies that they are both boundary components of a chamber C_{23} . Now note that if $C_{12} \neq C_{23}$, then they form two distinct chambers both incident to γ_2 . From Lemma 3.2, this forces γ_1 and γ_3 to lie in distinct connected components of $\tilde{X} - \gamma_2$. Hence, if η is an arbitrary geodesic segment joining a point on γ_1 to a point on γ_3 , we have that $\eta \cap \gamma_2 \neq \emptyset$. Since η is assumed to be geodesic, by restricting we can view η as a concatenation of a geodesic joining a point in γ_1 to a point in γ_2 , together with a geodesic joining a point in γ_2 to a point in γ_3 . Now γ_1, γ_2 are distinct boundary components of C_{12} , and the space C_{12} is the universal cover of a hyperbolic surface with non-empty, totally geodesic boundary. This implies that the first geodesic segment must intersect the interior of C_{12} non-trivially. Likewise, the second geodesic segment must intersect the interior of C_{23} non-trivially. But this contradicts the assumption that $\gamma_1 \sim \gamma_3$.

Since the adjacency relation is transitive on the subsets B_C , and since every subset on which the adjacency relation is transitive is contained in one of the B_C , we conclude that the latter are precisely the maximal subsets on which \sim is transitive, concluding the proof of the Lemma.

By combining the previous Lemma with the previous proposition, we immediately obtain the:

Corollary 3.4. *If $f : \tilde{X}_1 \rightarrow \tilde{X}_2$ is a quasi-isometry, then f induces a bijection f_* from \mathcal{C}_1 to \mathcal{C}_2 . Furthermore, if $\gamma \in \mathcal{BG}_1, C \in \mathcal{C}_1$ satisfy $\gamma \subset C$, then $f_*(\gamma) \subset f_*(C)$.*

We now know that a quasi-isometry between the universal covers of a pair of simple, thick, 2-dimensional hyperbolic P-manifolds induces a bijection between the lifts of chambers. Now recall that in the situation we are interested in, the quasi-isometry $\bar{\phi}_\infty$ from \tilde{X}_1 to \tilde{X}_2 has the additional property that it is (Γ_1, Γ_2) -equivariant. In particular, each lift of a chamber $C \in \mathcal{C}_i$ has a stabilizer under the action of Γ_i on \tilde{X}_i .

Our next step is to identify the stabilizers of the various $C \in \mathcal{C}_i$ from the boundary at infinity. This is made precise in the following:

Proposition 3.4. *Consider the natural action of Γ on \tilde{X} (where $\Gamma = \pi_1(X)$), and the corresponding induced action on $\partial^\infty \tilde{X}$. Then for every $C \in \mathcal{C}$ we have that the stabilizer of C under the Γ -action on \tilde{X} coincides with the stabilizer of $\partial^\infty C$ under the induced Γ -action on $\partial^\infty \tilde{X}$.*

The argument for this Proposition can be found in [5] (see the Assertion on pg. 212). Since an abstract isomorphism between Γ_1 and Γ_2 yields a (Γ_1, Γ_2) -equivariant homeomorphism between $\partial^\infty \tilde{X}_1$ and $\partial^\infty \tilde{X}_2$, this immediately yields the:

Corollary 3.5. *The bijection $\bar{\phi}_* : \mathcal{C}_1 \rightarrow \mathcal{C}_2$ induced by the isomorphism ϕ has the property that, for every $C \in \mathcal{C}_1$, one has $Stab_{\Gamma_1}(C) \cong Stab_{\Gamma_2}(\bar{\phi}_*(C))$.*

Armed with this information, it is now easy to complete the proof of the topological rigidity Theorem 1.2. We first note that, by the (Γ_1, Γ_2) -equivariance of the homeomorphism from $\partial^\infty \tilde{X}_1$ to $\partial^\infty \tilde{X}_2$, and in view of Proposition 3.3, we conclude that there are an equal number of orbits of branching geodesics in \tilde{X}_1 and \tilde{X}_2 . Since each such orbit corresponds to precisely one circle in the branching locus of the respective X_i , we conclude that there is a bijective correspondence between the components of the branching locus of X_1 and X_2 .

Now given any chamber in X_1 , one can consider the family of connected lifts of the chamber. These form a single orbit under the Γ_1 -action on \mathcal{C}_1 . Each lift in this orbit, by Corollary 3.4 maps to the lift of a corresponding chamber in \mathcal{C}_2 . But (Γ_1, Γ_2) -equivariance ensures that the lifts of chambers one gets in \mathcal{C}_2 lie in a single Γ_2 -orbit under the corresponding Γ_2 -action on \mathcal{C}_2 . Finally, Corollary 3.5 implies that the original chamber in X_1 , and the corresponding chamber in X_2 have the same fundamental group.

Next we note that both the number of boundary components the chamber has can be entirely determined by the number of distinct orbits of points $\{\gamma(\pm\infty)\}$ lying in $\partial^\infty C$, where $C \in \mathcal{C}_i$ is a lift of the chamber. Again, by (Γ_1, Γ_2) -equivariance of the homeomorphism, we obtain that pairs of corresponding chambers in X_1 and X_2 have exactly the same number of boundary components. Now observe that a surface with boundary is uniquely determined by its fundamental group and the number of boundary components it has. This tells us that the correspondence between chambers in X_1 and X_2 preserves the homeomorphism type of the chambers. To conclude, we note that the dynamics on the boundaries at infinity also allow us to keep track of how each chamber is attached to the branching strata. Putting all this together, we obtain a homeomorphism from X_1 to X_2 . It is immediate by construction that this homeomorphism induces the original isomorphism on the level of the fundamental groups, completing the proof of Theorem 1.2.

We now proceed to show how Corollary 1.2 (weak co-Hopf property) follows from Theorem 1.2.

Proof (Corollary 1.2). Let $\Gamma = \varinjlim \mathcal{D}$, and assume first that $i : \Gamma \hookrightarrow \Gamma$ is an injection with image of finite index. Consider the simple, thick, 2-dimensional hyperbolic P-manifold X whose fundamental group is Γ , and note that corresponding to the subgroup $i(\Gamma) \leq \Gamma$, we have a finite index covering $h : \tilde{X} \rightarrow X$ of degree equal to the index k of the subgroup $i(\Gamma)$. We now proceed to argue that $k = 1$; in order to see this, we note that by the topological rigidity theorem, the P-manifolds \tilde{X} and X are homeomorphic, and hence have the same number of chambers. Since chambers map

to chambers under covering maps, this implies that each chamber $W \subset X$ has pre-image consisting of a single chamber $h^{-1}(W) \subset \bar{X}$. But each chamber in X supports a hyperbolic metric, and hence has negative Euler characteristic. So picking a chamber $W \subset X$ which has *minimal* Euler characteristic amongst the finitely many chambers of X , we see that $\chi(W) \geq \chi(h^{-1}(W)) = k \cdot \chi(W)$, which immediately implies that $k = 1$, as desired.

Let us finish this section by explaining why we cannot obtain the full co-Hopf property for the groups $\Gamma = \varinjlim \mathcal{D}$. If the injection $i : \Gamma \hookrightarrow \Gamma$ has image with infinite index, one can again look at the *non-compact* cover \bar{X} corresponding to the subgroup $i(\Gamma)$. The non-compactness of \bar{X} is what prevents the argument above from going through: in order to apply the same argument one would need a non-compact analogue of the topological rigidity theorem.

The problem is that the proof of the topological rigidity theorem relies heavily on the fact that the spaces under consideration are compact. Indeed, this was used in the first step of the proof to obtain, from an abstract isomorphism of fundamental groups, a homeomorphism between the boundaries at infinity. As such, the proof of Theorem 1.1 fails in the non-compact setting, preventing us from obtaining the co-Hopf property in full generality.

4 Concluding remarks.

Our main theorem states that, within a certain class of diagrams of groups, each group that appears as a direct limit has a unique representative. An interesting question is the following:

Question: Which classes of diagrams of groups have the property that any group that occurs as a direct limit arises as the limit of a unique diagram?

Forester [3] has given criterions under which a Bass-Serre splitting of a group is unique (see also Guirardel [4]). We note that the Bass-Serre trees naturally associated to geometric amalgamations of free groups *do not* satisfy the hypotheses in Forester's work.

Another interesting aspect of the groups we are considering lies in the fact that these groups are essentially combinatorially determined. Indeed, in order to recognize the isomorphism type of these groups, it is sufficient (by the main theorem) to keep track of:

- the ranks of the free groups arising as fundamental groups of chambers,
- the number of boundary components of each chamber

- how the chambers get glued

Since this information consists of a finite amount of data, these groups form a class of δ -hyperbolic groups (or $CAT(-1)$ groups) for which one can look at various decision type problems. We can ask:

Question: When are the direct limits of a pair of geometric amalgamations of free groups quasi-isometric? When are they bi-Lipschitz equivalent?

References

- [1] B.H. Bowditch, ‘Connectedness properties of limit sets’, *Trans. Amer. Math. Soc.* 351 (1999), pgs. 3673-3686.
- [2] M. Bridson & A. Haefliger, *Metric spaces of non-positive curvature*, Springer-Verlag, Berlin, 1999.
- [3] M. Forester, ‘Deformation and rigidity of simplicial group actions on trees’, *Geom. Topol.* 6 (2002), pgs. 219-267.
- [4] V. Guirardel, ‘A very short proof of Forester’s rigidity result’, *Geom. Topol.* 7 (2003), pgs. 321-328.
- [5] J.-F. Lafont, ‘Rigidity results for certain 3-dimensional singular spaces and their fundamental groups’, *Geom. Dedicata* 109 (2004), pgs. 197-219.
- [6] J.-F. Lafont, ‘Strong Jordan separation and applications to rigidity’, preprint available at <http://www.math.arxiv.org>.
- [7] P. May, *A concise introduction to algebraic topology*, Univ. Chicago Press, Chicago, 1999.
- [8] Z. Sela, ‘Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups. II.’, *Geom. Funct. Anal.* 7 (1997), pgs. 561-593.
- [9] J.-P. Serre, *Trees*, Springer-Verlag, Berlin, 1980.
- [10] G.A. Swarup, ‘On the cut-point conjecture’, *Electron. Res. Announc. Amer. Math. Soc.* 2 (1996), pgs. 98-100.