

Strong Jordan separation and applications to rigidity.

Jean-François Lafont

October 3, 2006

Abstract

In this paper, we extend the results of [14] to higher dimension. We prove that simple, thick hyperbolic P-manifolds of dimension ≥ 3 exhibit Mostow rigidity. We also prove a quasi-isometry rigidity result for the fundamental groups of simple, thick hyperbolic P-manifolds of dimension ≥ 3 . The key tool in the proof of these rigidity results is a strong form of the Jordan separation theorem, for maps from $S^n \rightarrow S^{n+1}$ which are not necessarily injective. ¹

1 Introduction.

1.1 Basic definitions.

In recent years, there has been much interest in proving rigidity type theorems for non-positively curved spaces. All of these results originated from Mostow's seminal work, in which he showed that homotopy equivalent compact locally-symmetric spaces of rank one are always isometric [16].

In this paper, we are interested in showing rigidity results for certain singular CAT(-1) spaces that are "mostly" hyperbolic manifolds. Let us start by defining the various objects we will be discussing in the paper, and state the main theorems we will prove.

Definition 1.1. We define a closed n -dimensional *piecewise manifold* (henceforth abbreviated to P-manifold) to be a topological space which has a natural stratification into pieces which are manifolds. More precisely, we define a 1-dimensional P-manifold to be a finite graph. An n -dimensional P-manifold ($n \geq 2$) is defined inductively as a closed pair $X_{n-1} \subset X_n$ satisfying the following conditions:

¹MSC Codes: 20F67, 53C24, 57N35.

- Each connected component of X_{n-1} is either an $(n-1)$ -dimensional P-manifold, or an $(n-1)$ -dimensional manifold.
- The closure of each connected component of $X_n - X_{n-1}$ is homeomorphic to a compact orientable n -manifold with boundary, and the homeomorphism takes the component of $X_n - X_{n-1}$ to the interior of the n -manifold with boundary; the closure of such a component will be called a *chamber*.

Denoting the closures of the connected components of $X_n - X_{n-1}$ by W_i , we observe that we have a natural map $\rho : \coprod \partial(W_i) \longrightarrow X_{n-1}$ from the disjoint union of the boundary components of the chambers to the subspace X_{n-1} . We also require this map to be surjective, and a homeomorphism when restricted to each component. The P-manifold is said to be *thick* provided that each point in X_{n-1} has at least three pre-images under ρ . We will henceforth use a superscript X^n to refer to an n -dimensional P-manifold, and will reserve the use of subscripts X_{n-1}, \dots, X_1 to refer to the lower dimensional strata. For a thick n -dimensional P-manifold, we will call the X_{n-1} strata the *branching locus* of the P-manifold.

Intuitively, we can think of P-manifolds as being “built” by gluing manifolds with boundary together along lower dimensional pieces. Examples of P-manifolds include finite graphs and soap bubble clusters. Observe that compact manifolds can also be viewed as (non-thick) P-manifolds. Less trivial examples can be constructed more or less arbitrarily by finding families of manifolds with homeomorphic boundary and glueing them together along the boundary using arbitrary homeomorphisms. We now define the family of metrics we are interested in.

Definition 1.2. A Riemannian metric on a 1-dimensional P-manifold (finite graph) is merely a length function on the edge set. A Riemannian metric on an n -dimensional P-manifold X^n is obtained by first building a Riemannian metric on the X_{n-1} subspace, then picking, for each chamber W_i a Riemannian metric with totally geodesic boundary satisfying that the gluing map ρ is an isometry. We say that a Riemannian metric on a P-manifold is hyperbolic if at each step, the metric on each W_i is hyperbolic.

A hyperbolic P-manifold X^n is automatically a locally $CAT(-1)$ space (see Chapter II.11 in Bridson-Haefliger [5]). Furthermore, the lower dimensional strata X_i are totally geodesic subspaces of X^n . In particular, the universal cover \tilde{X}^n of a hyperbolic P-manifold X^n is a $CAT(-1)$ space (so is automatically δ -hyperbolic), and has a well-defined boundary at infinity $\partial^\infty \tilde{X}^n$. Finally we note that the fundamental group $\pi_1(X_n)$ is a δ -hyperbolic group. We refer the reader to [5] for background on $CAT(-1)$ and δ -hyperbolic spaces.

Definition 1.3. We say that an n -dimensional P-manifold X^n is *simple* provided its codimension two strata is empty. In other words, the $(n - 1)$ -dimensional strata X_{n-1} consists of a disjoint union of $(n - 1)$ -dimensional manifolds.

Next we introduce a locally defined topological invariant. We use \mathbb{D}^n to denote a closed n -dimensional disk, and \mathbb{D}_\circ^n to denote its interior.

Definition 1.4. Define the 1-*tripod* T to be the topological space obtained by taking the join of a one point set with a three point set. Denote by $*$ the point in T corresponding to the one point set. We define the n -*tripod* ($n \geq 2$) to be the space $T \times \mathbb{D}^{n-1}$, and call the subset $* \times \mathbb{D}^{n-1}$ the *spine* of the tripod $T \times \mathbb{D}^{n-1}$. The subset $* \times \mathbb{D}^{n-1}$ separates $T \times \mathbb{D}^{n-1}$ into three open sets, which we call the *open leaves* of the tripod. The union of an open leaf with the spine will be called a *closed leaf* of the tripod. We say that a point p in a topological space X is n -*branching* provided there is a topological embedding $f : T \times \mathbb{D}^{n-1} \rightarrow X$ such that $p \in f(* \times \mathbb{D}_\circ^{n-1})$.

It is clear that the property of being n -branching is invariant under homeomorphisms. Note that, in a simple, thick P-manifold of dimension n , points in the codimension one strata are automatically n -branching. One can ask whether this property can be detected at the level of the boundary at infinity. This motivated the author [14] to make the following:

Conjecture: Let X^n be a simple, thick hyperbolic P-manifold of dimension n , and let p be a point in the boundary at infinity of \tilde{X}^n . Then p is $(n - 1)$ -branching if and only if $p = \gamma(\infty)$ for some geodesic ray γ contained entirely in a connected lift of X_{n-1} .

Note that in general, the local structure of the boundary at infinity of a CAT(-1) space (or of a δ -hyperbolic group) is very hard to analyze. The conjecture above says that, with respect to branching, the boundary of a simple, thick hyperbolic P-manifold of dimension n is particularly easy to understand. In [14], the author showed that the above conjecture holds for $n = 3$.

1.2 Statement of results.

The main goal of this paper is to prove:

Theorem 1.1 (Mostow rigidity). *Let X_1, X_2 be a pair of simple, thick, n -dimensional ($n \geq 3$), hyperbolic P-manifolds, and assume that $\phi : \pi_1(X_1) \rightarrow \pi_1(X_2)$ is an isomorphism. Then there is an isometry $f : X_1 \rightarrow X_2$ which induces the map ϕ on the level of fundamental groups.*

Note that an immediate consequence of this result is that the canonical map $Isom(X) \rightarrow Out(\pi_1(X))$ from the isometry group to the outer automorphisms of the fundamental group is surjective when X is a simple, thick, n -dimensional ($n \geq 3$) hyperbolic P-manifold. Our second result determines the coarse geometry of the fundamental groups of these spaces:

Theorem 1.2 (Quasi-isometric rigidity). *Let X be a simple, thick, n -dimensional ($n \geq 3$), hyperbolic P-manifold, and assume that a finitely generated group Γ is quasi-isometric to $\pi_1(X)$. Then there exists a short exact sequence:*

$$0 \rightarrow F \rightarrow \Gamma \rightarrow \Lambda \rightarrow 0$$

with F a finite group, and Λ a discrete subgroup of $Isom(\tilde{X})$.

We remark that the 3-dimensional version of these two theorems were proved in the author's thesis [14]. In fact, the proofs given in that paper immediately extend to higher dimension, provided one knows that the conjecture in the previous section holds. The author's proof of the 3-dimensional case of the conjecture in [14], relied on the fact that the Schoenflies theorem holds in dimension 2. Since the corresponding theorem fails in higher dimensions, a different argument is needed.

Now recall that the classical Jordan separation theorem asserts that the image of an injective map $f : S^n \hookrightarrow S^{n+1}$ always separates S^{n+1} into precisely two connected components. The key result that allows us to extend the 3-dimensional arguments of [14] to higher dimensions is the following generalization of the Jordan separation theorem, which may be of independent interest.

Theorem 1.3 (Strong Jordan separation). *Let $f : S^n \rightarrow S^{n+1}$ be a continuous map, and let $I \subset S^n$ be the set of injective points (i.e. points $p \in S^n$ with the property that $f^{-1}(f(p)) = \{p\}$). If I contains an open set U , and $q \in U$, then:*

- $f(S^n)$ separates S^{n+1} into open subsets (we write $S^{n+1} - f(S^n)$ as a disjoint union $\amalg U_i$, with each U_i open),
- there are precisely two open subsets U_1, U_2 in the complement of $f(S^n)$ which contain $f(q)$ in their closure.
- if $F : \mathbb{D}^{n+1} \rightarrow S^{n+1}$ is an extension of the map f to the closed ball, then $F(\mathbb{D}^{n+1})$ surjects onto either U_1 or U_2 .

The main application of Strong Jordan separation will be to prove the conjecture mentioned in the previous section:

Theorem 1.4 (Characterization of branching points). *Let X^n be a simple, thick, n -dimensional ($n \geq 3$), hyperbolic P -manifold, and let p be a point in the boundary at infinity of \tilde{X}^n . Then p is $(n - 1)$ -branching if and only if $p = \gamma(\infty)$ for some geodesic ray γ contained entirely in a connected lift of X_{n-1} .*

We point out that the quasi-isometry result we have (Theorem 1.2) fits into the framework of understanding quasi-isometries of groups that have a splitting. In recent years, there has been an increasing interest in this question (see for instance [19], [13], [18], [7], [8], [15]). Our result differs from these in two ways. Firstly, the groups we allow as vertex and edge groups differ from those considered in these papers. Secondly, our method of proof involves a precise description of the local topology of the boundary at infinity (Theorem 1.4).

The outline of the paper is as follows: in Section 2, we will give a proof of Theorem 1.3 (Strong Jordan separation). Using Theorem 1.3, we will establish Theorem 1.4 (Characterization of branching points) in Section 3. In Section 4, we will give applications of Theorem 1.4 by outlining the proof (taken from [14]) of both Theorem 1.1 (Mostow rigidity) and Theorem 1.2 (Quasi-isometric rigidity). Some concluding remarks will be provided in Section 5.

Acknowledgements.

The author would like to thank his advisor, R. Spatzier, for many helpful discussions concerning this project. He would also like to thank F.T. Farrell for helpful suggestions on proving Theorem 1.3. F.D. Ancel informed the author of the locally-flat approximation theorem found in [1],[2], [3] (and used in our proof of Lemma 2.7). C. Westerland provided helpful comments on an early draft of this paper. Finally the author would like to thank the anonymous referee for pointing out simplifications for several of the author's original proofs.

2 Strong Jordan separation theorem.

In this section, we provide a proof of Theorem 1.3, which will subsequently be used in the proofs of the remaining results in this paper.

Before starting with the proof, we note that this theorem clearly generalizes the classical Jordan separation theorem (corresponding to the case $I = S^n$). The author does not know whether the hypotheses on I can be weakened to just assuming that I is measurable. Throughout our argument, \check{H}^* will refer to Čech cohomology with coefficients in \mathbb{Z} , while H_* will refer to reduced singular homology with coefficients in \mathbb{Z} . We use \mathbb{D}^n to refer to a closed n -dimensional disk, and \mathbb{D}_\circ^n to refer to its interior. We also occasionally use \mathbb{I} to denote a closed interval, and \mathbb{I}_\circ for its interior.

2.1 Global separation.

In this subsection, we provide a proof of the first claim in Theorem 1.3, and some other results which are of a global nature.

Lemma 2.1. *Under the hypotheses above, $f(S^n)$ separates S^{n+1} .*

Proof. We start by showing that $f(S^n)$ cannot surject onto S^{n+1} . Since I is assumed to contain an open set, we can find a $\mathbb{D}^n \subset I \subset S^n$, i.e. a (small) disk on which f is injective. Let $p \in \mathbb{D}^\circ$ be a point in the interior of the disk, and note that (since $p \in I$), we have that $d(f(p), f(q)) \neq 0$ for all $q \in S^n - \{p\}$. Since both p and $f(S^n - \mathbb{D}^\circ)$ are compact sets, the distance between them must be positive. Pick ϵ smaller than $d(f(p), f(S^n - \mathbb{D}^\circ))$, and note that an ϵ -ball centered at $f(p)$ can only intersect $f(S^n)$ in a subset of $f(\mathbb{D}^n)$. But on this set, f is injective, and it is well known that there is no injective and surjective map from a subset of \mathbb{D}^n to \mathbb{D}^{n+1} (the ϵ -ball centered at p). This shows that $f(S^n)$ is not surjective.

Since f is not surjective, Alexander duality tells us that $\tilde{H}_0(S^{n+1} - f(S^n)) \cong \check{H}^n(f(S^n))$, hence it is sufficient to show that $\check{H}^n(f(S^n)) \neq 0$. Now the space $f(S^n)$ can be expressed as a union $f(S^n) = f(\mathbb{D}^n) \cup_g f(S^n - \mathbb{D}^\circ)$ (where $g : \partial\mathbb{D}^n = S^{n-1} \rightarrow f(S^n - \mathbb{D}^\circ)$ is the gluing map); this follows from the fact that the defining property of the set I ensures that $f(\mathbb{D}^n)$ intersects $f(S^n - \mathbb{D}^\circ)$ in precisely $f(\partial\mathbb{D}^n)$. Since f is injective on \mathbb{D}^n , we have $f(\mathbb{D}^n)$ is homeomorphic to \mathbb{D}^n , and the gluing is along $\partial\mathbb{D}^n = S^{n-1}$. Furthermore, note that g is homotopic to a point in $f(S^n - \mathbb{D}^\circ)$, and since S^{n-1} is collarable in \mathbb{D}^n , we have that the space $f(S^n)$ is homotopy equivalent to the space $X = f(\mathbb{D}^n) \cup_h f(S^n - \mathbb{D}^\circ)$, where h maps $\partial\mathbb{D}^n$ to a point. In particular, we see that $f(S^n)$ is homotopy equivalent to the join $S^n \vee f(S^n - \mathbb{D}^\circ)$. This implies that $\check{H}^i(f(S^n)) \cong \check{H}^i(S^n \vee f(S^n - \mathbb{D}^\circ))$. But we can compute the latter by using the Mayer-Vietoris sequence in Čech cohomology, yielding $\check{H}^n(f(S^n)) \cong \mathbb{Z} \oplus \check{H}^n(f(S^n - \mathbb{D}^\circ)) \neq 0$. This gives us our claim.

We now have that the image does indeed separate S^{n+1} . We denote by U_i the connected components of $S^{n+1} - f(S^n)$, giving us the decomposition $S^{n+1} - f(S^n) = \coprod U_i$. Note that $S^{n+1} - f(S^n)$ is open, so that connected components coincide with path-connected components. Let us now focus on the local behavior near $f(\mathbb{D}^\circ)$.

Lemma 2.2. *The inclusion map $S^{n+1} - f(S^n) \hookrightarrow S^{n+1} - f(S^n - \mathbb{D}^\circ)$ induces a canonical splitting $\tilde{H}_0(S^{n+1} - f(S^n)) \cong \mathbb{Z} \oplus \tilde{H}_0(S^{n+1} - f(S^n - \mathbb{D}^\circ))$.*

Proof. Since Alexander duality is natural, we have a commutative diagram:

$$\begin{array}{ccc} \tilde{H}_0(S^{n+1} - f(S^n)) & \longrightarrow & \tilde{H}_0(S^{n+1} - f(S^n - \mathbb{D}^\circ)) \\ \downarrow \cong & & \downarrow \cong \\ \check{H}^n(f(S^n)) & \longrightarrow & \check{H}^n(f(S^n - \mathbb{D}^\circ)) \end{array}$$

where the two horizontal arrows are induced by the obvious inclusions, and the vertical arrows are given by duality. So it is sufficient to show that the morphism $\check{H}^n(f(S^n)) \rightarrow \check{H}^n(f(S^n - \mathbb{D}_\circ^n))$ gives the desired splitting. But this is clear, since we have $\check{H}^n(f(S^n)) \cong \mathbb{Z} \oplus \check{H}^n(f(S^n - \mathbb{D}_\circ^n))$, where the direct sum comes from the Mayer-Vietoris sequence, and the \mathbb{Z} comes from $\check{H}^n(S^n)$ (see previous Lemma). Combining this with the commutative diagram above gives us the desired splitting $\tilde{H}_0(S^{n+1} - f(S^n)) \cong \mathbb{Z} \oplus \tilde{H}_0(S^{n+1} - f(S^n - \mathbb{D}_\circ^n))$, completing the proof of the lemma.

This gives us a homological version of the statement “ $f(\mathbb{D}_\circ^n)$ locally separates S^{n+1} into pieces lying in precisely two components of $S^{n+1} - f(S^n)$ ”. We now want to replace this by a non-homological, equivalent statement.

Lemma 2.3. *Let $\{U_i\}$ be a collection of disjoint connected open sets in \mathbb{R}^{n+1} . Assume that there is a connected set Z which intersects the closure of each U_i non-trivially. Then the space $Y := Z \cup (\bigcup U_i)$ is a connected set.*

The proof of Lemma 2.3 is an easy exercise in point-set topology, which we leave to the reader.

Corollary 2.1. *There are precisely two connected components U_1, U_2 in $S^{n+1} - f(S^n)$ whose closure intersects $f(\mathbb{D}_\circ^n)$.*

Proof. Let r denote the number of connected components in $S^{n+1} - f(S^n)$ whose closure intersects $f(\mathbb{D}_\circ^n)$, and let us enumerate them by $\{U_i\}_{i=1}^r$. Picking points $x_i \in U_i$, note that for $1 \leq j < k \leq r$, all the elements $[x_j - x_k] \in \tilde{H}_0(S^{n+1} - f(S^n))$ represent distinct non-zero classes. On the other hand, Lemma 2.3 tells us that each U_j, U_k must lie in the same connected component of $S^{n+1} - f(S^n - \mathbb{D}_\circ^n)$, and hence that the classes $[x_j - x_k]$ map to zero under the homomorphism:

$$\tilde{H}_0(S^{n+1} - f(S^n)) \rightarrow \tilde{H}_0(S^{n+1} - f(S^n - \mathbb{D}_\circ^n))$$

induced by the inclusion. This implies that the kernel of the homomorphism is in fact generated by the classes $[x_i - x_{i+1}]$ ($1 \leq i < r$).

But now Lemma 2.2 tells us that the homomorphism above induces a splitting $\mathbb{Z} \oplus \tilde{H}_0(S^{n+1} - f(S^n - \mathbb{D}_\circ^n)) \cong \tilde{H}_0(S^{n+1} - f(S^n))$. Since the kernel has rank 1, this implies that there are precisely two connected components in $S^{n+1} - f(S^n)$ whose closure intersects $f(\mathbb{D}_\circ^n)$, as desired.

2.2 Local separation.

While the separation results we’ve obtained so far are useful, they do not tell us much about the behavior of the embedding near a point in the injectivity set I . The purpose of this section is to fine-tune the information we have in order to capture local information concerning the map f . We start with a local version of Lemma 2.2:

Lemma 2.4. *Let $p \in f(\mathbb{D}_\circ^n)$ be an arbitrary point, and U_1, U_2 be the two connected components of $S^{n+1} - f(S^n)$ whose closure intersects $f(\mathbb{D}_\circ^n)$. Then p lies in the closure of both U_1 and U_2 .*

Proof. Given $p \in \mathbb{D}_\circ^n$, choose a sequence of open metric balls $D_i \subset \mathbb{D}_\circ^n$ centered at p , having the property that their radii tend to zero. Now since each D_i lies in the injectivity set I , Corollary 2.1 tells us that for each D_i , there are precisely two connected components $U_1^{(i)}, U_2^{(i)}$ in $S^{n+1} - f(S^n)$ whose closure intersects $f(D_i)$. But since $f(D_i) \subset f(\mathbb{D}_\circ^n)$, this forces (upto relabelling) $U_1^{(i)} = U_1, U_2^{(i)} = U_2$. This implies that the closure of U_1, U_2 intersects points in each $f(D_i)$. Since the radii of the D_i tends to zero, this forces $f(p)$ to likewise lie in the closure of both U_1, U_2 .

Next we introduce a group that measures, for an arbitrary point p in S^{n+1} , the number of components of $S^{n+1} - f(S^n)$ that contain p in their closure.

Definition 2.1. For a point $p \in S^{n+1}$, and a continuous map $f : S^n \rightarrow S^{n+1}$, we define the *group of local components of p relative to f* to be:

$$H(p, f) := \varprojlim \{ \tilde{H}_0(B(p, 1/m) - f(S^n)) \},$$

where $B(p, 1/m)$ is the open metric ball of radius $1/m$, centered at p , and the inverse limit is over the directed system $m \in \mathbb{N}$, with morphisms induced by inclusions.

Observe that the groups $H(p, f)$ are *local*, in the sense that if one has a pair of maps $f, g : S^n \rightarrow S^{n+1}$, and the point p has an open neighborhood N with the property that $f(S^n) \cap N = g(S^n) \cap N$, then the groups $H(p, f)$ and $H(p, g)$ are canonically isomorphic. Furthermore, note that if $p \notin f(S^n)$, then $H(p, f) = 0$.

Next we observe that for an arbitrary point $p \in S^{n+1}$, and each $m \in \mathbb{N}$, we have inclusions $B(p, 1/m) - f(S^n) \hookrightarrow S^{n+1} - f(S^n)$, which induce homomorphisms $\tilde{H}_0(B(p, 1/m) - f(S^n)) \rightarrow \tilde{H}_0(S^{n+1} - f(S^n))$. It is easy to see that these homomorphisms induce a homomorphism $\phi : H(p, f) \rightarrow \tilde{H}_0(S^{n+1} - f(S^n))$. Our next two lemmas serve to analyze the image of this map in the case where $p \in f(\mathbb{D}_\circ^n)$.

Lemma 2.5. *If $p \in f(\mathbb{D}_\circ^n)$, the $H(p, f) \cong \mathbb{Z}$.*

Proof. Fix $m \in \mathbb{N}$, and consider the map $H : S^{n+1} \rightarrow S^{n+1} = B(p, 1/m)^*$ obtained by collapsing the complement of $B(p, 1/m)$ to a point (so that the target can be identified with the one point compactification $B(p, 1/m)^*$ of the open ball $B(p, 1/m)$). Applying Corollary 2.1 and Lemma 2.4 to the map $H \circ f : S^n \rightarrow B(p, 1/m)^*$, we see that there are exactly two connected components of $B(p, 1/m)^* - (H \circ f)(S^n)$ containing p in their closure. But this immediately implies that there are precisely two connected components of $B(p, 1/m) - f(S^n)$ that contain p in their closure

Now consider what this tells us about the inverse system of groups $\{\tilde{H}_0(B(p, 1/m) - f(S^n))\}_{m \in \mathbb{N}}$. By an argument similar to the one given at the beginning of the proof of Lemma 2.4, we see that for each $m \in \mathbb{N}$, there exists a M with the property that $\forall k \geq M$, the set $B(p, 1/k) - f(S^n)$ lies in the union of the two connected components of $B(p, 1/m) - f(S^n)$ which contain p in their closure. In particular, on the level of \tilde{H}_0 , we see that for all $k \geq M$, the map $\tilde{H}_0(B(p, 1/k) - f(S^n)) \rightarrow \tilde{H}_0(B(p, 1/m) - f(S^n))$ has image lying in a fixed \mathbb{Z} -subgroup of $\tilde{H}_0(B(p, 1/m) - f(S^n))$. But the inverse limit of an inverse system of groups with this property is automatically isomorphic to \mathbb{Z} , concluding the proof of Lemma 2.5.

Lemma 2.6. *If $p \in f(\mathbb{D}_\circ^n)$, then $\phi(H(p, f)) \cong \mathbb{Z}$. Furthermore, the image of ϕ is in fact generated by the reduced homology class $[x_1 - x_2] \in \tilde{H}_0(S^{n+1} - f(S^n))$, where $x_i \in U_i$ (and U_1, U_2 are the two connected components of $S^{n+1} - f(S^n)$ whose closures intersect $f(\mathbb{D}_\circ^n)$).*

Proof. If $m \in \mathbb{N}$ is arbitrary, then $B(p, 1/m) - f(S^n)$ must contain points in both U_1 and U_2 . Hence the inclusion induces a map $\tilde{H}_0(B(p, 1/m) - f(S^n)) \rightarrow \tilde{H}_0(S^{n+1} - f(S^n))$ whose image contains the subgroup generated by the reduced homology class $[x_1 - x_2] \in \tilde{H}_0(S^{n+1} - f(S^n))$. Since this holds for all $m \in \mathbb{N}$, the induced map on the inverse limit $H(p, f) \rightarrow \tilde{H}_0(S^{n+1} - f(S^n))$ likewise has image containing the subgroup generated by $[x_1 - x_2]$.

Furthermore, if m is sufficiently large, then the image of the map $\tilde{H}_0(B(p, 1/m) - f(S^n)) \rightarrow \tilde{H}_0(S^{n+1} - f(S^n))$ actually lies in the subgroup generated by $[x_1 - x_2]$. Passing to the induced map from the inverse limit $H(f, p)$, we see that ϕ maps $H(f, p)$ into the subgroup generated by $[x_1 - x_2]$, completing our proof.

Morally speaking, the last two lemmas make precise the fact that *the components which are incident to a point in the set of injectivity can be detected purely locally*. Indeed, the combination of Lemma 2.5 and 2.6 tells us that the map ϕ is in fact an isomorphism from $H(p, f)$ to the subgroup generated by $[x_1 - x_2]$. This will be crucial to us in the next section, where we prove the third claim from our theorem.

2.3 Surjectivity of extensions.

Finally, to complete the proof of the theorem, we need to show that any extension $F : \mathbb{D}^{n+1} \rightarrow S^{n+1}$ of the map f must surject onto either U_1 or U_2 . Let us assume, without loss of generality, that F does not surject onto U_1 . Then using a stereographic projection, we can view the image of f as lying in \mathbb{R}^{n+1} , with U_1 unbounded and all other components in the complement of $f(S^n)$ being bounded. Furthermore, we also have that the image of F lies in \mathbb{R}^{n+1} (after stereographic projection).

We will also fix an orientation on \mathbb{R}^{n+1} , yielding for every point $z \in \mathbb{R}^{n+1}$ a canonical generator for $H_n(\mathbb{R}^{n+1} - z) \cong \mathbb{Z}$. In particular, using this identification, it makes sense to say that an n -cycle in $\mathbb{R}^{n+1} - \{z_1, z_2\}$ represents distinct elements in $H_n(\mathbb{R}^{n+1} - z_1)$ and $H_n(\mathbb{R}^{n+1} - z_2)$.

Now observe that for any pair of points z_1, z_2 which lie in a fixed connected component U of $\mathbb{R}^{n+1} - f(S^n)$, we have that the map f represents the same element in $H_n(\mathbb{R}^{n+1} - z_1)$ and $H_n(\mathbb{R}^{n+1} - z_2)$. Furthermore, if there exists an extension of the map f to a map $F : \mathbb{D}^{n+1} \rightarrow \mathbb{R}^{n+1}$ which fails to surject onto the component U , then denoting by $z \in U$ a point that does not lie in the image of F , we see that f represents 0 in the group $H_n(\mathbb{R}^{n+1} - z)$ (since f bounds in the complement of z), and hence that f represents 0 in every $H_n(\mathbb{R}^{n+1} - z')$ where $z' \in U$ is arbitrary. In particular, since the image of F is compact, and U_1 is unbounded, we have that f must represent 0 in $H_n(\mathbb{R}^{n+1} - z')$ whenever $z' \in U_1$.

With this in mind, the last claim of the theorem will follow immediately from:

Lemma 2.7. *The map $f : S^n \rightarrow \mathbb{R}^{n+1}$, ($n \geq 2$) cannot represent 0 in $H_n(\mathbb{R}^{n+1} - z) \cong \mathbb{Z}$, where $z \in U_2$. In particular, F must surject onto U_2 .*

Proof. Note that we may assume that the set $\mathbb{D}^n \subset S^n$ is actually a round disk; let $p \in I$ be the center of \mathbb{D}^n , r the radius of \mathbb{D}^n . Denote by $\hat{\mathbb{D}}^n$ the closed round disk with center p , and radius $r/2$, and for $t \in [0, r/2]$, let $d(t)$ be the distance from $f(S_t^{n-1})$ to $f(S^n - \hat{\mathbb{D}}_o^n)$, where S_t^{n-1} is the sphere of radius t centered at p . Note that, for each $t \in [0, r/2]$, the sphere S_t^{n-1} lies in \mathbb{D}^n , hence is contained in the set of injective points. Let $q \in \mathbb{D}_o^n - \hat{\mathbb{D}}^n$ be an arbitrary point, and let $[x_m - y_m] \in \tilde{H}_0(B(f(q), 1/m) - f(S^n))$ be a sequence of group elements defining the generator in $H(f(q), f)$ which maps to $[z' - z] \in \tilde{H}_0(S^{n+1} - f(S^n))$ (recall that $z' \in U_1$ and $z \in U_2$).

We would now like to approximate the map f by a new map g which has the following properties:

- $f = g$ on the complement of $\hat{\mathbb{D}}_o^n$,
- g is locally flat embedding on $\hat{\mathbb{D}}_o^n$,
- g is an embedding on \mathbb{D}_o^n

To see that such an approximation exists, we take a continuous function $\delta : [0, r/2] \rightarrow \mathbb{R}^+$ with the property that $\delta(t) < d(t)$. A result of Ancel-Cannon [2] ($n \geq 4$), Ancel [1] ($n = 3$), and Bing [3] ($n = 2$), states that we can approximate the map f on $\hat{\mathbb{D}}_o^n$, by a map $\hat{g} : \hat{\mathbb{D}}_o^n \rightarrow \mathbb{R}^{n+1}$ having the property that \hat{g} is a locally flat embedding, and $d(\hat{g}, f) \leq \delta(t)$ in the sup norm on each S_t^{n-1} ($t \in [0, r/2]$). Finally, we note that the concatenation of \hat{g} on the S_t^{n-1} ($t \in [0, r/2]$) with the map f on the S_t^{n-1} ($t \in [r/2, r]$) has the desired properties; let us call this map g .

Applying Corollary 2.1 to the map g , let us denote by U'_1 and U'_2 the two components of $S^{n+1} - g(S^n)$ that contain $g(p)$ in their closure. From Lemma 2.4, we know that $f(q) = g(q)$ also lies in the closure of both U'_1 and U'_2 . Next we note that, for a small enough neighborhood N of $f(q)$, we have that $N \cap f(S^n) = N \cap g(S^n)$, hence $H(f(q), f)$ is canonically isomorphic to $H(f(q), g)$. In particular, the sequence $[x_m - y_m] \in \tilde{H}_0(B(f(q), 1/m) - f(S^n))$ still defines a generator for $H(f(q), g)$. This implies that, for m large enough $x_m \in U'_1$ and $y_m \in U'_2$ (upto reindexing).

Now observe that, since g is locally flat near $\hat{\mathbb{D}}^n_\circ$, the map g is *actually flat* on $\hat{\mathbb{D}}^n$. A proof of this old folklore result can be found in Rushing [20] Flattening Theorem 3.4.1. This implies that there is a homeomorphism $\phi : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^{n+1}$ with the property that $(\phi \circ g)(\hat{\mathbb{D}}^n_\circ) \subset \mathbb{R}^n \times \{0\}$. Furthermore, this homeomorphism takes the U'_i to the two open sets W^+ (respectively W^-) containing $(\phi \circ g)(\hat{\mathbb{D}}^n_\circ)$ in their closure, where W^\pm consists of points lying (locally near $(\phi \circ g)(\hat{\mathbb{D}}^n_\circ)$) above/below the hyperplane $\mathbb{R}^n \times \{0\}$. But now it is classical that for a pair of points $x \in U'_1$, $y \in U'_2$, that the map $\phi \circ g$ must represent distinct elements in $H_n(\mathbb{R}^{n+1} - \phi(x))$ and $H_n(\mathbb{R}^{n+1} - \phi(y))$. Indeed, there is a PL-curve joining $\phi(x)$ to $\phi(y)$ and intersecting $\phi \circ g$ transversally in a single point. Pulling back along ϕ , we see that g represents distinct elements in $H_n(\mathbb{R}^{n+1} - x)$ and $H_n(\mathbb{R}^{n+1} - y)$.

This tells us that for m large enough, the map g represents distinct elements in $H_n(\mathbb{R}^{n+1} - x_m)$ and $H_n(\mathbb{R}^{n+1} - y_m)$. Finally, note that by our choice of the control function δ , for m large enough we have that g is homotopic to f in the complement of $\{x_m, y_m\}$. Hence f must represent distinct elements in $H_n(\mathbb{R}^{n+1} - x_m)$ and $H_n(\mathbb{R}^{n+1} - y_m)$. Since f represents 0 in $H_n(\mathbb{R}^{n+1} - x_m)$ (by the comment before this proof), we get the desired result.

Combining Lemma 2.1, Corollary 2.1, and Lemma 2.7 proves Theorem 1.3.

3 Recognizing branching.

In this section, we will provide a proof of Theorem 1.4 from the introduction. For the convenience of the reader, we restate the Theorem:

Theorem 1.4. *Let X^n be a simple, thick hyperbolic P -manifold of dimension $n \geq 3$, and let p be a point in the boundary at infinity of \tilde{X}^n . Then p is $(n-1)$ -branching if and only if $p = \gamma(\infty)$ for some geodesic ray γ contained entirely in a connected lift of X_{n-1} (the $(n-1)$ -dimensional strata).*

As mentioned in the introduction, the 3-dimensional version of this result was shown by the author in [14]. The key point where the hypothesis $n = 3$ was used was in the 2-dimensional version of the Strong Jordan Separation Theorem (Theorem 2.1 in [14]). Indeed, in that proof, we made use of the Schoenflies Theorem, which holds

in dimension 2, but is known to fail in dimension ≥ 3 . Since the majority of the results in this section follow immediately from the proofs of the corresponding results in [14], we will merely emphasize the arguments that need to be modified. We start by mentioning:

Proposition 3.1. *Let X^n be a simple, thick n -dimensional P -manifold, and let p be a point in the boundary at infinity of \tilde{X}^n . If γ is a geodesic ray contained entirely in a connected lift \tilde{B} of X_{n-1} , then $\gamma(\infty)$ is $(n-1)$ -branching.*

This was Proposition 2.1 in [14]. The idea was to use the simplicity and thickness hypothesis to ensure that there are three disjoint totally geodesic subsets, each isometric to a “half” \mathbb{H}^n , whose boundaries all coincide with \tilde{B} . This forces all points in $\partial^\infty \tilde{B}$ to be $(n-1)$ -branching.

Note that points in $\partial^\infty \tilde{X}^n$ which *do not* lie in any of the $\partial^\infty \tilde{B}_i$ are of two types. They are either:

- endpoints of geodesic rays that eventually stay trapped in a connected lift \tilde{W} of a chamber W , and are *not* asymptotic to any boundary component of \tilde{W} , or
- endpoints of geodesic rays that pass through infinitely many connected lifts \tilde{W}_i of chambers W_i .

We will discuss these two cases separately.

3.1 The geodesic ray is trapped in a chamber.

We now explain how to deal with the first of the two cases.

Proposition 3.2. *Let X^n be a simple, thick n -dimensional hyperbolic P -manifold ($n \geq 3$). Let $\gamma \subset \tilde{X}^n$ be a geodesic ray lying entirely within a connected lift \tilde{W} of a chamber W , and not asymptotic to any boundary component of \tilde{W} . Then $\gamma(\infty)$ is **not** $(n-1)$ -branching.*

This was Proposition 2.2 in [14]. The argument relies on picking a point $x \in \gamma$, and considering the link of the point x in \tilde{X}^n . Note that since x is in a chamber, the link $lk(x)$ is isometric to S^{n-1} . We denote by $\pi_x : \partial^\infty \tilde{X}^n \rightarrow lk(x) \cong S^{n-1}$ the geodesic retraction from the boundary at infinity to the link at x . Finally, we denote by I_x the set $\{p \in lk(x) : |\pi_x^{-1}(p)| = 1\} \subset lk(x)$, in other words, the set of points in the link where the projection map is actually injective.

Now by way of contradiction, assume that $\gamma(\infty)$ is $(n-1)$ -branching. Then there exists an injective map $f : T \times \mathbb{D}^{n-2} \rightarrow \partial^\infty \tilde{X}^n$ with the property that $\gamma(\infty) \in f(* \times \mathbb{D}_o^{n-1})$. One can now consider the composite map $\pi_x \circ f : T \times \mathbb{D}^{n-1} \rightarrow lk(x)$.

Since $lk(x)$ is topologically S^{n-1} , we know that the composite map cannot be an embedding. The strategy lies in showing that there exists a point in I_x which has two distinct pre-images under the composite $\pi_x \circ f$. Since π_x is injective on points in I_x , this would imply that f fails to be injective, yielding a contradiction.

To see that the composite $\pi_x \circ f$ fails to be injective at a point in I_x , we proceed via a series of Claims.

Claim 1. *The complement of I_x has the following properties:*

- *it consists of a countable union of open disks U_i in S^{n-1} ,*
- *the U_i are the interiors of a family of pairwise disjoint closed disks,*
- *the U_i are dense in S^{n-1} .*

Claim 2. *For any point $p \in I_x - \cup(\partial U_i)$, and any neighborhood N_p of p , there exist arbitrarily small U_i with $U_i \subset N_p$.*

Claim 3. *The image $(\pi_x \circ f)(\partial(T \times \mathbb{D}^{n-2}))$ is a bounded distance away from $\pi_x(\gamma(\infty))$.*

The proof of these three claims given in [14] for the case $n = 3$ extend verbatim to the case $n \geq 3$. For the convenience of the reader, we briefly summarize the arguments.

Points in the complement of I_x correspond to directions in which π_x fails to be injective. This means that there exists a pair of geodesic rays which coincide on a small neighborhood of x , and subsequently diverge. So the complement of I_x is precisely the image of $\partial\tilde{W}$ under the geodesic retraction. Since \tilde{W} is the universal cover of a compact hyperbolic manifold with boundary, standard methods yield the first claim.

For the second claim, we use the fact that the U_i are actually *round* disks in $lk(x) = S^{n-1}$. This implies that, for the standard measure μ on S^{n-1} , the radius of the U_i is proportional to $\mu(U_i)^{\frac{1}{n-1}}$. In particular, there are only finitely many U_i of radius $\geq r$ for any $r > 0$. Using this fact, one can find a suitably smaller neighborhood $N'_p \subset N_p$ with the property that the U_i that intersect N'_p have radius less than half the distance between $\partial N'_p$ and $S^{n-1} - N_p$, giving the second claim.

The third claim follows immediately from the compactness of the set $(\pi_x \circ f)(\partial(T \times \mathbb{D}^{n-2}))$, along with the fact that $\pi_x(\gamma(\infty)) \notin (\pi_x \circ f)(\partial(T \times \mathbb{D}^{n-2}))$.

Continuing with our argument we now have:

Claim 4. *There is a neighborhood N of the point $f^{-1}(\gamma(\infty))$ in the spine $* \times \mathbb{D}_\circ^{n-2}$ whose image is entirely contained in I_x (i.e. $(\pi_x \circ f)(N) \subset I_x$).*

This claim requires a substantial modification from that given in [14]. We first note that, denoting by r_i the radius of the respective U_i , the argument for Claim 2 shows that for any $R > 0$, there are only finitely many U_i with the property that $r_i > R$. Note that, by hypothesis, we have that $\pi_x(\gamma(\infty)) \notin \cup \bar{U}_i$. Hence it is sufficient to show that, for a small enough ϵ , there are only finitely many U_i intersecting $(\pi_x \circ f)(\ast \times \mathbb{D}_\circ^{n-2}) \cap B(\pi_x(\gamma(\infty)))$, where $B(\pi_x(\gamma(\infty)))$ denotes the open ϵ ball centered at $\pi_x(\gamma(\infty))$.

Now if this were not the case, then there exists, for each $\epsilon > 0$, infinitely many U_i intersecting $(\pi_x \circ f)(\ast \times \mathbb{D}_\circ^{n-2})$ at a distance $< \epsilon$ from $\pi_x(\gamma(\infty))$. Now pick ϵ to be strictly smaller than the distance D from $\pi_x(\gamma(\infty))$ to the set $(\pi_x \circ f)(\partial(T \times \mathbb{D}^{n-2}))$. We can now find a U_i such that $2r_i + \epsilon < D$, and U_i intersects $(\pi_x \circ f)(\ast \times \mathbb{D}_\circ^{n-2})$ at a distance $< \epsilon$ from $\pi_x(\gamma(\infty))$.

Let us focus now on this U_i . The condition that $2r_i + \epsilon < D$ ensures that $\bar{U}_i \cap (\pi_x \circ f)(\partial(T \times \mathbb{D}^{n-2})) = \emptyset$ (i.e. that the image of the boundary of the $(n-1)$ -tripod fails to intersect the closure of the U_i). Furthermore, note that $\partial U_i \subset I_x$, and is homeomorphic to S^{n-2} . This implies that $(\pi_x \circ f)^{-1}(\partial U_i)$ is a closed subset of S^{n-2} (topologically) embedded in $T \times \mathbb{D}_\circ^{n-2}$, and which *fails* to intersect $\partial(T \times \mathbb{D}^{n-2})$. Since ∂U_i separates $lk(x)$, we have that $(\pi_x \circ f)^{-1}(\partial U_i)$ likewise separates $T \times \mathbb{D}_\circ^{n-2}$.

Let us denote by S the subset $(\pi_x \circ f)^{-1}(\partial U_i)$ in $T \times \mathbb{D}_\circ^{n-2}$, and denote by L_1, L_2, L_3 the three closed leaves of $T \times \mathbb{D}_\circ^{n-2}$. We now observe that $L_j \cup L_k$ is homeomorphic to \mathbb{D}^{n-1} ($1 \leq j < k \leq 3$), and that the intersection $S_{j,k} := S \cap (L_j \cup L_k)$ is a closed subset of $L_j \cup L_k$. Furthermore, we note that at least two of the three sets $S_{1,2}, S_{1,3}, S_{2,3}$ must separate the corresponding $L_j \cup L_k$. Without loss of generality, assume that $S_{1,2}, S_{2,3}$ separate the corresponding $L_j \cup L_k$. Finally, observe that since $S \cap \partial(T \times \mathbb{D}^{n-2}) = \emptyset$, we can view each of the sets $S_{1,2}, S_{2,3}$ as lying in the interior of the corresponding $L_i \cup L_j$, which is naturally homeomorphic to \mathbb{R}^{n-1} .

Now note that both $S_{1,2}$ and $S_{2,3}$ are closed subsets which separate $Int(L_1 \cup L_2) \cong \mathbb{R}^{n-1}$ and $Int(L_2 \cup L_3) \cong \mathbb{R}^{n-1}$ respectively. By Alexander duality, this implies that $\check{H}^{n-2}(S_{1,2}) \neq 0$ and $\check{H}^{n-2}(S_{2,3}) \neq 0$. However, both these sets are homeomorphic to closed subsets of S^{n-2} , and the only closed subset A of S^{n-2} satisfying $\check{H}^{n-2}(A) \neq 0$ is S^{n-2} itself. This implies that both $S_{1,2}$ and $S_{2,3}$ must be homeomorphic to S^{n-2} , yielding that $S_{1,2} = S_{2,3} = S$. From the definition of the sets $S_{j,k}$, we immediately get that $S \subset L_2$.

So we have now reduced the problem to showing that the pre-image $S = (\pi_x \circ f)^{-1}(\partial U_i)$ of ∂U_i cannot lie entirely in L_2 . Note that S separates L_2 into two components, one of which must map into U_i , while the other one must map into $lk(x) - U_i$. Now by hypothesis, there is a point in the spine $\ast \times \mathbb{D}^{n-2}$ of $T \times \mathbb{D}^{n-2}$ that maps into U_i , which implies that there is a point on the boundary of $L_2 = \mathbb{I} \times \mathbb{D}^{n-2}$ which maps into U_i . This implies that the entire spine must in fact map either to U_i , or to

∂U_i . Since $\pi_x(\gamma(\infty))$ lies in the image of the spine, we contradict the hypothesis that $\pi_x(\gamma(\infty)) \in I_x$ does not lie in the closure of any U_i . This concludes the argument for Claim 4.

Using Claim 4, and focusing on a sufficiently small neighborhood of $f^{-1}(\gamma(\infty))$ in $T \times \mathbb{D}^{n-2}$, we may assume that the map $\pi_x \circ f$ maps the spine into the set I_x . We will henceforth work under this assumption.

Let us denote by F_i the restriction of the composition $\pi_x \circ f$ to each of the closed leaves in $T \times \mathbb{D}^{n-2}$, and by G_i the restriction of $\pi_x \circ f$ to the boundary of the closed leaves. We now have the following:

Claim 5. *There is a connected open set $V \subset S^{n-1}$ with the property that:*

- *at least two of the maps F_i surject onto V*
- *the closure of V contains the point $\pi_x(\gamma(\infty))$*

As in [14], Claim 5 is a consequence of the Strong Jordan separation theorem. Observe that each leaf is homeomorphic to $\mathbb{I} \times \mathbb{D}^{n-2} \cong \mathbb{D}^{n-1}$, and from Claim 4, we know that the maps G_i is injective on the open set $* \times \mathbb{D}_o^{n-2}$ (the spine). Hence the Strong Jordan separation tells us that each $G_i(S^{n-2})$ separates S^{n-1} , and that there are precisely two connected open sets $U_i, V_i \subset S^{n-1} - G_i(S^{n-2})$ which contain $G_i(\mathbb{D}^{n-2})$ in their closure. Furthermore, each of the maps F_i surjects onto either V_i or U_i .

Now let $p := \pi_x(\gamma(\infty))$, and note that we have canonical isomorphisms:

$$H(p, G_1) \cong H(p, G_2) \cong H(p, G_3) \cong \mathbb{Z}.$$

Indeed, this follows from the fact that all three maps have to coincide in a small enough neighborhood of p . Now fix a generator $\{[x_m - y_m]\}$ for $H(p, G_1)$ as in Lemma 2.7. Now Lemmas 2.5 and 2.6 imply that (upto re-indexing), we may assume that for m sufficiently large we have $x_m \in U_i, y_m \in V_i$ for all $1 \leq i \leq 3$. We now assume, without loss of generality, that F_1 and F_2 surject onto V_1 and V_2 respectively.

Next we note that if ϵ is small enough, we will have that the ball D of radius ϵ centered at p only intersects $G_i(\mathbb{D}^{n-2})$ (this follows from the argument in Claim 3). In particular, each connected component of $D - G_i(\mathbb{D}^{n-2})$ is contained in either U_i or in V_i . Now observe that the argument given in Lemma 2.5 ensures that there are precisely two connected components of $D - G_i(\mathbb{D}^{n-2})$ which contain p in their closure. Denote these two components by U, V , labelled so that for m sufficiently large, $x_m \in U, y_m \in V$. By the choice of our labelling, we have that $U \subset \cap U_i$, and $V \subset \cap V_i$.

Finally, we note that V is connected and contains p in its closure (by construction), and since $V \subset V_1 \cap V_2$, we have that both F_1 and F_2 surject onto V (since they surject onto V_1, V_2 respectively). This concludes the proof of Claim 5.

Claim 6. *Let $V \subset S^{n-1}$ be a connected open set, containing the point $p := \pi_x(\gamma(\infty))$ in its closure. Then V contains a connected open set U_j lying in the complement of the set I_x .*

The proof of Claim 6 is identical to the one given in [14]. We reproduce it here for completeness. We first claim that the connected open set V contains a point from I_x . Indeed, if not, then V would lie entirely in the complement of I_x , hence would lie in some U_i . Since $\pi_x(\gamma(\infty))$ lies in the closure of V , it would also lie in the closure of U_i , contradicting the fact that γ is *not* asymptotic to any of the boundary components of the chamber containing γ .

So not only does V contain the point $\pi_x(\gamma(\infty))$ in its closure, it also contains some point q in I_x . We claim it in fact contains a point in $I_x - \cup(\partial U_i)$. If q itself lies in $I_x - \cup(\partial U_i)$ then we are done. The other possibility is that q lies in the boundary of one of the U_i . Now since V is connected, there exists a path η joining q to $\pi_x(\gamma(\infty))$. Now assume that $\eta \cap (I_x - \cup(\partial U_i)) = \{\pi_x(\gamma(\infty))\}$. Let \bar{U}_i denote the closed disks (closure of the U_i), and note that the complement of the set $I_x - \cup(\partial U_i)$ is the set $\cup(\bar{U}_i)$.

A result of Sierpinski [21] states the following: let X be an arbitrary topological space, $\{C_i\}$ a countable collection of disjoint path connected closed subsets in X . Then the path connected components of $\cup C_i$ are precisely the individual C_i .

Applying Sierpinski's result to the set $\cup(\bar{U}_i)$ shows that the path connected component of this union are precisely the individual \bar{U}_i . So if $\eta \cap (I_x - \cup(\partial U_i)) = \{\pi_x(\gamma(\infty))\}$, we see that the path η must lie entirely within the \bar{U}_i containing q . This again contradicts the fact that $\pi_x(\gamma(\infty)) \notin \cup(\partial U_i)$. Finally, the fact that V contains a point in $I_x - \cup(\partial U_i)$ allows us to invoke Claim 2, which tells us that there is some U_j which is contained entirely within the set V , completing our argument for Claim 6.

We now proceed to complete the proof of Proposition 3.2. From Claims 5 and 6, there exists an open set $V \subset S^{n-1} = lk(x)$ with the property that:

- the maps F_1, F_2 surject onto V .
- V contains a connected open set U_j lying in the complement of the set I_x .

This of course implies that both F_1 and F_2 must surject onto the boundary ∂U_i of the U_i . But now we observe that ∂U_i lies in the set of injectivity I_x . Since ∂U_i lies in the image of F_1 and F_2 , the pre-image of ∂U_i must lie in $L_1 \cap L_2 = * \times \mathbb{D}^{n-2}$. Furthermore, the pre-image of ∂U_i is homeomorphic to an S^{n-2} . So we obtain an

embedded S^{n-2} inside $* \times \mathbb{D}^{n-2}$, which immediately gives us a contradiction. This completes the proof of Proposition 3.2.

3.2 Connectedness and separation properties.

We briefly state four results that will be used to deal with the second case of Theorem 1.4, and will also be used in the arguments for Theorems 1.1 and 1.2. The proofs of these results given in [14] under the hypothesis $n = 3$ extend immediately to the more general setting $n \geq 2$.

Lemma 3.1. *Let \tilde{B}_i be a connected lift of the branching locus, and let \tilde{W}_j, \tilde{W}_k be two lifts of chambers which are both incident to \tilde{B}_i . Then $\tilde{W}_j - \tilde{B}_i$ and $\tilde{W}_k - \tilde{B}_i$ lie in different connected components of $\tilde{X}^n - \tilde{B}_i$.*

Lemma 3.2. *Let \tilde{W}_i be a connected lift of a chamber, and let \tilde{B}_j, \tilde{B}_k be two connected lifts of the branching locus which are both incident to \tilde{W}_i . Then \tilde{B}_j and \tilde{B}_k lie in different connected components of $\tilde{X}^n - \text{Int}(\tilde{W}_i)$.*

Lemma 3.3. *Let $\partial^\infty \tilde{B}_i$ be the boundary at infinity of a connected lift of the branching locus, and let $\partial^\infty \tilde{W}_j, \partial^\infty \tilde{W}_k$ be the boundaries at infinity of two lifts of chambers which are both incident to \tilde{B}_i . Then $\partial^\infty \tilde{W}_j - \partial^\infty \tilde{B}_i$ and $\partial^\infty \tilde{W}_k - \partial^\infty \tilde{B}_i$ lie in different connected components of $\partial^\infty \tilde{X}^n - \partial^\infty \tilde{B}_i$.*

Lemma 3.4. *Let $\partial^\infty \tilde{W}_i$ be the boundary at infinity corresponding to a connected lift of a chamber, and let $\partial^\infty \tilde{B}_j, \partial^\infty \tilde{B}_k$ be the boundary at infinity of two connected lifts of the branching locus which are both incident to \tilde{W}_i . Then $\partial^\infty \tilde{B}_j$ and $\partial^\infty \tilde{B}_k$ lie in different connected components of $\partial^\infty \tilde{X}^n - (\partial^\infty \tilde{W}_i - \cup \partial^\infty \tilde{B}_l)$, where the union is over all \tilde{B}_l which are boundary components of \tilde{W}_i .*

3.3 The geodesic ray passes through infinitely many chambers.

We now explain how to deal with the second case of the Theorem 1.4, namely we show the following:

Proposition 3.3. *Let X^n be a simple, thick n -dimensional hyperbolic P -manifold, with $n \geq 3$. Let $\gamma \subset \tilde{X}^n$ be a geodesic that passes through infinitely many connected lifts \tilde{W}_i . Then $\gamma(\infty)$ is **not** $(n - 1)$ -branching.*

In the proof of Proposition 3.3, we will make use of a family of nice metrics on the boundary at infinity of an arbitrary n -dimensional hyperbolic P -manifold (in fact, on the boundary at infinity of any CAT(-1) space).

Definition 3.1. Given an n -dimensional hyperbolic P-manifold, and a basepoint $*$ in \tilde{X}^n , we can define a metric on the boundary at infinity by setting $d_\infty(p, q) = e^{-d(*, \gamma_{pq})}$, where γ_{pq} is the unique geodesic joining the points p, q (and d denotes the distance inside \tilde{X}_n).

The fact that d_∞ is a metric on the boundary at infinity of a proper $CAT(-1)$ space follows from Bourdon (Section 2.5 in [4]). Note that changing the basepoint from $*$ to $*'$ changes the metric, but that for any $p, q \in \partial^\infty(X^n)$, we have the inequalities:

$$A^{-1} \cdot d_{\infty,*}(p, q) \leq d_{\infty,*'}(p, q) \leq A \cdot d_{\infty,*}(p, q)$$

where $A = e^{d(*, *')}$, and the subscripts on the d_∞ refers to the choice of basepoint used in defining the metric. In particular, different choices for the basepoint induce the same topology on $\partial^\infty \tilde{X}^n$, and this topology coincides with the standard topology on $\partial^\infty \tilde{X}^n$ (the quotient topology inherited from the compact-open topology in the definition of $\partial^\infty \tilde{X}^n$ as equivalence classes of geodesic rays in \tilde{X}^n). This gives us the freedom to select basepoints at our convenience when looking for *topological* properties of the boundary at infinity. We now proceed to prove Proposition 3.3:

Proof. The approach here consists of reducing to the situation covered in proposition 3.2. We start by re-indexing the various consecutive connected lifts \tilde{W}_i that γ passes through by the integers. Fix a basepoint $x \in \tilde{W}_0$ interior to the connected lift \tilde{W}_0 , and lying on γ . Now assume that there is an injective map $f : T \times \mathbb{D}^{n-2} \rightarrow \partial^\infty \tilde{X}^n$ with $\gamma(\infty) \in f(* \times \mathbb{D}_\circ^{n-2})$.

We start by noting that, between successive connected lifts \tilde{W}_i and \tilde{W}_{i+1} that γ passes through, lies a connected lift of the branching locus, which we denote \tilde{B}_i . Observe that distinct connected lifts of the branching locus stay a uniformly bounded distance apart. Indeed, any minimal geodesic joining two distinct lifts of the branching locus must descend to a minimal geodesic in a W_i with endpoints in the branching locus. But the length of any such geodesic is bounded below by the injectivity radius of DW_i , the double of W_i across its boundary. By setting δ to be the infimum, over all the finitely many chambers W_i , of the injectivity radius of the doubles DW_i , we have $\delta > 0$. Let K_i be the connected component of $\partial^\infty \tilde{X}^n - \partial^\infty \tilde{B}_i$ containing $\gamma(\infty)$. Then for every $p \in K_i$ ($i \geq 1$), we have:

$$d_x(p, \gamma(\infty)) < e^{-\delta(i-1)}.$$

Indeed, by Lemma 3.1, \tilde{B}_i separates \tilde{X}^n into (at least) two totally geodesic components. Furthermore, the component containing $\gamma(\infty)$ is *distinct* from that containing x . Hence, the distance from x to the geodesic joining p to $\gamma(\infty)$ is at least as large as the distance from x to \tilde{B}_i . But the later is bounded below by $\delta(i-1)$. Using

the definition of the metric at infinity, and picking x as our basepoint, our estimate follows.

Since our estimate shrinks to zero, and since the distance from $\gamma(\infty)$ to $f(\partial(T \times \mathbb{D}^{n-2}))$ is positive, we must have a point $q \in f(* \times \mathbb{D}_\circ^{n-2})$ satisfying $d_x(q, \gamma(\infty)) > e^{-\delta(i-1)}$ for i sufficiently large. Since $\partial^\infty \tilde{B}_i$ separates (by Lemma 3.3), we see that for i sufficiently large, $f(* \times \mathbb{D}_\circ^{n-2})$ contains points in two distinct components of $\partial^\infty \tilde{X}^n - \partial^\infty \tilde{B}_i$. This implies that there is a point $q' \in f(* \times \mathbb{D}_\circ^{n-2})$ that lies within some $\partial^\infty \tilde{W}_k - \cup \partial^\infty \tilde{B}_l$, where the union runs over all \tilde{B}_l which are boundary components of \tilde{W}_k . But such a point corresponds to a geodesic ray lying entirely within \tilde{W}_k , and *not* asymptotic to any of the lifts of the branching locus. Finally, we note that *any* point in the image $f(* \times \mathbb{D}_\circ^{n-2})$ can be considered $(n-1)$ -branching, so in particular the point q' is $(n-1)$ -branching. But in Proposition 3.2, we showed this is impossible. Our claim follows.

Finally, combining Proposition 3.1, 3.2, and 3.3 gives us Theorem 1.4.

4 Rigidity results.

In this section, we discuss two types of rigidity results (Theorems 1.1 and 1.2 from the introduction) for simple, thick hyperbolic P-manifolds of dimension ≥ 3 . In [14], the author proved both these theorems for the case $n = 3$, by making use of the 3-dimensional analogue of Theorem 1.4. But the argument given in that previous paper works equally well in higher dimensions (once one knows that Theorem 1.4 holds). As such, we merely outline the arguments in this section, and refer the interested reader to [14] for more details. We start by recalling a useful result from [14]:

Lemma 4.1. *Let X^n be a simple hyperbolic P-manifold of dimension at least three. Let $\partial^\infty \tilde{B} \subset \partial^\infty \tilde{X}^n$ consist of all limit points of geodesics in the branching locus. If $n \geq 3$, then the maximal path-connected components of $\partial^\infty \tilde{B}$ are precisely the sets of the form $\partial^\infty \tilde{B}_i$, where $\tilde{B}_i \subset \tilde{B}$ is a single connected component of the lifts of the branching locus.*

Indeed, each $\partial^\infty \tilde{B}_i$ is path-connected (since $n \geq 3$ and each \tilde{B}_i is isometric to \mathbb{H}^{n-1}) and closed (since each \tilde{B}_i is totally geodesic), hence $\partial^\infty \tilde{B}$ is the union of a disjoint family of closed, path-connected subspaces. The fact that the maximal path-connected components of $\partial^\infty \tilde{B}$ are the individual $\partial^\infty \tilde{B}_i$ now follows from an easy application of Sierpinski's result [21].

Note that in the previous Lemma, the hypothesis that the dimension be ≥ 3 really is necessary, since in the 2-dimensional case one would have the codimension one strata consisting of geodesics. Hence for each connected component of a lift of

the branching locus, we would have that $\partial^\infty \tilde{B}_i$ consists of a pair of points (and hence, would not be path connected).

4.1 Mostow rigidity.

We now explain how Theorem 1.1 follows from Theorem 1.4. Theorem 1.1 is clearly an analogue of Mostow rigidity for simple, thick hyperbolic P-manifolds. Let us start with a pair X_1, X_2 of simple, thick hyperbolic P-manifolds of dimension ≥ 3 , and an isomorphism $\phi : \pi_1(X_1) \rightarrow \pi_1(X_2)$ on the level of the fundamental groups. Our goal is to show that this implies that X_1 is isometric to X_2 .

In order to do this, we note that an abstract isomorphism between the fundamental groups induces a quasi-isometry between the universal covers $\tilde{\phi} : \tilde{X}_1 \rightarrow \tilde{X}_2$. One of the basic facts about quasi-isometries between δ -hyperbolic spaces is that they extend to give *homeomorphisms* between the boundaries at infinity of the spaces. In particular, the quasi-isometry $\tilde{\phi} : \tilde{X}_1 \rightarrow \tilde{X}_2$ extends to give a homeomorphism, denoted ϕ_∞ from $\partial^\infty \tilde{X}_1$ to $\partial^\infty \tilde{X}_2$. Furthermore, this homeomorphism intertwines the actions of the groups $\pi_1(X_1)$ and $\pi_1(X_2)$ on the respective boundaries at infinity.

Now recall that the property of being $(n-1)$ -branching is a topological invariant, ϕ_∞ must map branching points in $\partial^\infty \tilde{X}_1$ to branching points in $\partial^\infty \tilde{X}_2$. Since connectedness is preserved under homeomorphisms, ϕ_∞ must map connected components of the set of branching points in $\partial^\infty \tilde{X}_1$ to connected components of the set of branching points in $\partial^\infty \tilde{X}_2$.

Now Theorem 1.4 and Lemma 4.1 combine to tell us precisely what these connected components are: each connected component corresponds precisely to the points in the boundary at infinity of a single connected lift of the branching locus. This forces ϕ_∞ to induce a bijection between the lifts of the branching locus in \tilde{X}_1 and the lifts of the branching locus in \tilde{X}_2 .

The separation argument (Lemmas 3.3 and 3.4) in the previous section ensure that the map ϕ_∞ takes the boundaries of the lifts of chambers in $\partial^\infty \tilde{X}_1$ to the boundaries of the lifts of chambers in $\partial^\infty \tilde{X}_2$. This naturally induces a bijection between the lifts of chambers in \tilde{X}_1 and the lifts of chambers in \tilde{X}_2 . Furthermore, we note that this bijection preserves the adjacency relation between chambers (again, this follows from Lemmas 3.3 and 3.4).

Finally, we note that the fundamental group of the chambers in a simple, thick hyperbolic P-manifold X can be detected from the boundary at infinity: indeed, the fundamental group of a chamber W_i is isomorphic to the stabilizer of $\partial^\infty \tilde{W}_i$ under the action of $\pi_1(X)$ on $\partial^\infty \tilde{X}$. Since the map ϕ_∞ intertwines the actions, we conclude that the bijection takes lifts of chambers to lifts of chambers with isomorphic fundamental groups.

Mostow rigidity for hyperbolic manifolds with totally geodesic boundary (see Frigerio [10]) now allows us to conclude that the individual chambers must be isometric. Furthermore, the gluings between the chambers are detected, since lifts of adjacent chambers map to lifts of adjacent chambers, and the dynamics of $\pi_1(X)$ on $\partial^\infty \tilde{X}$ allow us to identify the subgroups of the individual chambers (i.e. the boundary components) that get identified together. Patching this information together, we obtain an equivariant isometry from \tilde{X}_1 to \tilde{X}_2 . This naturally descends to an isometry from X_1 to X_2 . Furthermore, by construction this isometry must induce ϕ on the level of π_1 , concluding the argument for Theorem 1.1.

4.2 Quasi-Isometric rigidity.

In this section, we outline the argument for showing the quasi-isometric rigidity result (Theorem 1.2). The starting point is the following well known result (for a proof, see for instance Prop. 3.1 in [6]):

Lemma 4.2. *Let X be a proper geodesic metric space, and assume that every quasi-isometry from X to itself is in fact a bounded distance (in the sup norm) from an isometry. Furthermore, assume that a finitely generated group G is quasi-isometric to X . Then there exists a cocompact lattice $\Gamma \subset \text{Isom}(X)$, and a finite group F which fit into a short exact sequence:*

$$0 \longrightarrow F \longrightarrow G \longrightarrow \Gamma \longrightarrow 0$$

In particular, if one can show that on the universal cover \tilde{X}^n of a simple, thick hyperbolic P-manifold of dimension $n \geq 3$, every quasi-isometry is a bounded distance from an isometry, then applying Lemma 4.2 would yield Theorem 1.2.

So let us assume that $\phi : \tilde{X}^n \rightarrow \tilde{X}^n$ is a quasi-isometry, and observe that this quasi-isometry induces a homeomorphism $\partial^\infty \phi : \partial^\infty \tilde{X}^n \rightarrow \partial^\infty \tilde{X}^n$. Arguing as in the previous section, this quasi-isometry must map each $\partial^\infty \tilde{W}_i$ (W_i a chamber in X^n) homeomorphically to $\partial^\infty \tilde{W}_j$ (where W_j is again a chamber in X^n). In particular, the quasi-isometry ϕ restricts (upto a bounded amount) to a quasi-isometry from \tilde{W}_i to \tilde{W}_j . We now recall the well-known:

Folklore Theorem: *Let M_1^n, M_2^n be a pair of n -dimensional ($n \geq 3$) hyperbolic manifolds with non-empty, totally geodesic boundaries. If $\phi : \tilde{M}_1^n \rightarrow \tilde{M}_2^n$ is a quasi-isometry, then there exists an isometry $\psi : \tilde{M}_1^n \rightarrow \tilde{M}_2^n$ at a bounded distance from ϕ (in the sup norm).*

Proofs of this result have been discovered at various times by Farb, Kapovich, Kleiner, Leeb, Schwarz, Wilkinson, and others, though to the author's knowledge the only written proof of this theorem is to be found in Frigerio [11].

By appealing to the previous Theorem, we see that the quasi-isometry of \tilde{X}^n has the property that, for each lift of a chamber of \tilde{X}^n , there is an isometric mapping to the lift of some other chamber that is bounded distance away from the original quasi-isometry. There are two points to verify:

1. that the individual isometries on the lifts of the chambers glue together to give a global isometry,
2. that the distance between the quasi-isometry and the isometry on each lift of a chamber is uniformly bounded.

Concerning point (1), we first note that the quasi-isometry maps adjacent lifts of chambers to adjacent lifts of chambers (this follows from Lemmas 3.3 and 3.4). The fact that the isometries coincide on the common boundary follows from the fact that the isometries induced by the incident chambers are all at a bounded distance from each other (since they are all at a bounded distance from the original quasi-isometry). But the common boundary is isometric to \mathbb{H}^{n-1} (by the simplicity hypothesis), and any two isometries of \mathbb{H}^{n-1} which are bounded distance apart have to coincide. This allows us to glue together the isometries, taking care of point (1).

Concerning point (2), one sees that it is possible, on the universal covers of compact hyperbolic manifolds with totally geodesic boundary, to give a uniform bound (depending only on the quasi-isometry constants) on the distance between a quasi-isometry and the isometry that is at finite distance from it (see the end of Section 3.2 in [14]). Since the quasi-isometries on all the chambers were induced by a global quasi-isometry of \tilde{X}^n , they all have the same quasi-isometry constants, and hence one obtains point (2). This completes our discussion of Theorem 1.2.

5 Concluding remarks.

It would be interesting to find some further applications of the rigidity results presented in this paper. One potential application is to the Gromov-Thurston examples of negatively curved manifolds. Recall that these manifolds \bar{M}^n are topologically ramified coverings of hyperbolic manifolds M^n , with the ramification occurring over a totally geodesic codimension two submanifold N^{n-2} (see [12]). These manifolds naturally inherit a (singular) CAT(-1) metric from the hyperbolic metric on M^n , with the singular set for the metric consisting of the pre-image of the ramification locus N^{n-2} . Gromov-Thurston showed that this singular CAT(-1) metric can be smoothed out to a negatively curved Riemannian metric.

Now consider the case where N^{n-2} bounds a totally geodesic codimension one submanifold K^{n-1} , and observe that the pre-image of K^{n-1} form a totally geodesic

subset in \bar{M}^n (equipped with the singular CAT(-1) metric) isometric to a simple, thick (assuming the ramified cover has degree at least three) hyperbolic P-manifold. It is tempting to see whether the rigidity results presented in this paper can be used to analyze the Gromov-Thurston examples. One basic questions along these lines is the following:

Question: Let M^n be a Gromov-Thurston negatively curved manifold, $N^{n-2} \subset M^n$ the ramification locus, and $\Gamma := \pi_1(M^n)$, $\Lambda := \pi_1(N^{n-2})$ their respective fundamental groups. If $\phi : \Gamma \rightarrow \Gamma$ is an abstract isomorphism, is it true that $\phi(\Lambda)$ is always conjugate to Λ ?

Next, we note that the argument for Theorem 1.4 (Characterization of branching points) extends (with suitable modifications in the proofs) to the more general setting of P-manifolds with *negatively curved* metrics. We did not pursue this in the present paper because we could not think of suitable applications of this result. Of course, there is no hope of proving Mostow type rigidity in this setting (though topological rigidity à la Farrell-Jones [9] should still hold). As for quasi-isometric rigidity, what would be needed is the following open question:

Question: If M^n is a negatively curved Riemannian manifold with non-empty, totally geodesic boundary. Is every quasi-isometry of the universal cover \tilde{M}^n bounded distance away from an isometry?

The argument for the Folklore Theorem (see [11]) cannot possibly work in this more general setting, since it relies heavily on the relationship between quasi-isometries (and isometries) of \mathbb{H}^n on the one hand, and quasi-conformal (and conformal) maps of $S^{n-1} = \partial^\infty \mathbb{H}^n$ on the other. Since no such correspondance exists in the general variable curvature case (though see Pansu [17]), an entirely new approach would need to be given. Finally, concerning Theorem 1.1, we ask the following:

Question: If one has a map $f : S^n \rightarrow S^{n+1}$ with the property that the set of injectivity has *positive measure*, does that imply that $f(S^n)$ separates S^{n+1} ?

It is clear that the approach we use really requires the presence of an open disk in the set of injectivity, and hence is of no use in answering this last question.

References

- [1] F.D. Ancel, ‘Resolving wild embeddings of codimension-one manifolds in manifolds of dimensions greater than 3’, Special volume in honor of R. H. Bing (1914–1986), *Topology Appl.* 24 (1986), pp. 13–40.

- [2] F.D. Ancel and J.W. Cannon, ‘The locally flat approximation of cell-like embedding relations’, *Ann. of Math. (2)* 109 (1979), pp. 61–86.
- [3] R.H. Bing, ‘Approximating surfaces by polyhedral ones’, *Ann. of Math. (2)* 65 (1957), pp. 465–483.
- [4] M. Bourdon, ‘Structure conforme au bord et flot géodésiques d’un $CAT(-1)$ -espace’, *Enseign. Math. (2)* 41 (1995), pp. 63–102.
- [5] M.R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature* (Springer-Verlag, Berlin, 1999).
- [6] B. Farb, ‘The quasi-isometry classification of lattices in semisimple Lie groups’, *Math. Res. Lett.* 4 (1997), pp. 705–717.
- [7] B. Farb and L. Mosher, ‘Quasi-isometric rigidity for the solvable Baumslag-Solitar groups. II.’, *Invent. Math.* 137 (1999), pp. 613–649.
- [8] B. Farb and L. Mosher, ‘The geometry of surface-by-free groups’, *Geom. Funct. Anal.* 12 (2002), pp. 915–963.
- [9] F.T. Farrell and L.E. Jones, ‘A topological analogue of Mostow’s rigidity theorem’, *J. Amer. Math. Soc.* 2 (1989), pp. 257–370.
- [10] Frigerio R., ‘Hyperbolic manifolds with geodesic boundary which are determined by their fundamental group’, *Topology Appl.* 145 (2004), pp. 69–81.
- [11] Frigerio R., ‘Commensurability of hyperbolic manifolds with geodesic boundary’, preprint available on the ArXiv at <http://front.math.ucdavis.edu/math.GT/0502209>.
- [12] M. Gromov and W. Thurston, ‘Pinching constants for hyperbolic manifolds’, *Invent. Math.* 89 (1987), pp. 1–12.
- [13] M. Kapovich and B. Leeb, ‘Quasi-isometries preserve the geometric decomposition of Haken manifolds’, *Invent. Math.* 128 (1997), pp. 393–416.
- [14] J.-F. Lafont, ‘Rigidity result for certain 3-dimensional singular spaces and their fundamental groups’, *Geom. Dedicata* 109 (2004), pp. 197–219.
- [15] L. Mosher, M. Sageev, and K. Whyte, ‘Quasi-actions on trees. I. Bounded valence’, *Ann. of Math. (2)* 158 (2003), pp. 115–164.

- [16] G.D. Mostow, *Strong rigidity of locally symmetric spaces* (Princeton University Press, Princeton, N.J., 1973).
- [17] P. Pansu, 'Dimension conforme et sphère à l'infini des variétés à courbure négative', *Ann. Acad. Sci. Fenn. Ser. A I Math.* 14 (1989), pp. 177–212.
- [18] P. Papasoglu, 'Group splittings and asymptotic topology', preprint available on the ArXiv at <http://front.math.ucdavis.edu/math.GR/0201312>.
- [19] P. Papasoglu and K. Whyte, 'Quasi-isometries between groups with infinitely many ends', *Comment. Math. Helv.* 77 (2002), pp. 133–144.
- [20] T.B. Rushing, *Topological embeddings* (Academic Press, New York-London, 1973).
- [21] W. Sierpinski, 'Un théorème sur les continus', *Tohoku Math. Journ.* 13 (1918), pp. 300-303.