Journal of Topology and Analysis
Vol. 1, No. 4 (2009) 431–459

World Scientific
www.worldscientific.com

A BOUNDARY VERSION OF CARTAN–HADAMARD AND APPLICATIONS TO RIGIDITY

JEAN-FRANÇOIS LAFONT
Department of Mathematics, The Ohio State University,
Columbus, OH 43210, USA
jlafont@math.ohio-state.edu

Received 21 October 2009

The classical Cartan–Hadamard theorem asserts that a closed Riemannian manifold M^n with non-positive sectional curvature has universal cover $\tilde{M}^n$ diffeomorphic to $\mathbb{R}^n$, and a by-product of the proof is that $\partial^\infty \tilde{M}^n$ is homeomorphic to S^{n-1}. We prove analogues of these two results in the case where M^n has a non-empty totally geodesic boundary. More precisely, if M_1^n, M_2^n are two negatively curved Riemannian manifolds with non-empty totally geodesic boundary, of dimension $n \neq 5$, we show that $\partial^\infty \tilde{M}_1^n$ is homeomorphic to $\partial^\infty \tilde{M}_2^n$. We show that if M_1^n and M_2^n are a pair of non-positively curved Riemannian manifolds with totally geodesic boundary (possibly empty), then the universal covers $\tilde{M}_1^n$ and $\tilde{M}_2^n$ are diffeomorphic if and only if the universal covers have the same number of boundary components. We also show that the number of boundary components of the universal cover is either 0, 2 or ∞. As a sample application, we show that simple, thick, negatively curved P-manifolds of dimension ≥ 6 are topologically rigid. We include some straightforward consequences of topological rigidity (diagram rigidity, weak co-Hopf property, and the Nielson problem).

1. Introduction

The classical Cartan–Hadamard theorem states that, if M is a simply-connected, complete Riemannian manifold with non-positive sectional curvature, then the exponential mapping $T_pM \to M$ at the point p is a diffeomorphism from the tangent space at p to the entire manifold. This result has two important consequences which we propose to generalize in the present paper.

First of all, if M^n is a closed Riemannian manifold of non-positive sectional curvature, then the universal cover $\tilde{M}^n$ is diffeomorphic to $\mathbb{R}^n$. In particular, in each dimension n, there is a unique smooth manifold arising as the universal cover of such manifolds.

Secondly, if M^n is a closed Riemannian manifold of non-positive curvature, and $p \in \tilde{M}^n$ is a point in the universal cover, then distinct geodesic rays emanating from p only intersect at their common basepoint. In particular, geodesic projection yields a homeomorphism between the boundary at infinity $\partial^\infty \tilde{M}^n$ and the unit tangent space at p, and hence one obtains that $\partial^\infty \tilde{M}^n \cong S^{n-1}$.

Our goal is to obtain analogues of these two results, in the situation where M^n has non-empty, totally geodesic boundary. The second statement above is generalized, under a stronger curvature constraint, in the following:

Theorem 1.1. (Topological characterization of $\partial^\infty \tilde{M}$) *Assume M_1, M_2 are a pair of compact*, negatively curved *Riemannian manifolds of dimension* $n \neq 5$, *with non-empty, totally geodesic boundary. Then we have $\partial^\infty \tilde{M}_1$ is homeomorphic to $\partial^\infty \tilde{M}_2$, where $\tilde{M}_i$ is the universal cover of M_i.*

Note that if $n = 2$, then the boundaries at infinity of the $\tilde{M}_i$ are Cantor sets, and the theorem follows from the classical fact that any two Cantor sets are homeomorphic (Brouwer's characterization theorem). It is relatively easy to extend the homeomorphism between the $\partial^\infty \tilde{M}_i$ obtained in Theorem 1.1 to a homeomorphism between the universal covers $\tilde{M}_i$ (an outline of this argument is given at the end of Sec. 2). We do not provide the details for this argument, as in Sec. 3 the following stronger result will be obtained, via different methods:

Theorem 1.2. (Smooth characterization of $\tilde{M}$) *Assume M_1, M_2 are a pair of compact*, non-positively curved *Riemannian manifolds with totally geodesic boundary (possibly empty). Then the following two statements are equivalent:*

- *$\tilde{M}_1$ is diffeomorphic to $\tilde{M}_2$.*
- *$\tilde{M}_1$ has the same number of boundary components as $\tilde{M}_2$.*

Furthermore, the possible number of boundary components of such an $\tilde{M}$ is either

- *0 boundary components, which is clearly equivalent to M being closed, or*
- *2 boundary components, in which case the universal cover $\tilde{M}$ splits isometrically as the product of a totally geodesic, codimension one submanifold with a closed interval, or*
- *infinitely many boundary components, which is the generic case.*

In particular, in each dimension $n \geq 2$, there are up to diffeomorphism precisely three spaces that occur as the universal cover of a compact, non-positively curved Riemannian manifolds with totally geodesic boundary (possibly empty).

Finally, we conclude by providing the following application of Theorem 1.1 (see Sec. 4 for definitions):

Theorem 1.3. (Topological rigidity of negatively curved P-manifolds) *Let X_1, X_2 be a pair of simple, thick, negatively curved P-manifolds, of dimension ≥ 6. If $\pi_1(X_1)$ is isomorphic to $\pi_1(X_2)$, then X_1 is homeomorphic to X_2.*

This last result has a number of interesting consequences:

Corollary 1.1. (Diagram rigidity) *Let $\mathcal{D}_1, \mathcal{D}_2$ be a pair of diagrams of groups, corresponding to a pair of negatively curved, simple, thick P-manifolds of dimension*

$n \geq 6$. *Then* $\varinjlim \mathcal{D}_1$ *is isomorphic to* $\varinjlim \mathcal{D}_2$ *if and only if the two diagrams are isomorphic.*

Corollary 1.2. (Weak Co-Hopf property) *Let* X *be a simple, thick, negatively curved* P*-manifold of dimension* $n \geq 6$, *and assume that at least one of the chambers has a nonzero characteristic number. Then* $\Gamma = \pi_1(X)$ *is weakly co-Hopfian, i.e. every injection* $\Gamma \hookrightarrow \Gamma$ *with image of finite index is in fact an isomorphism.*

Corollary 1.3. (Nielson realization problem) *Let* X *be a simple, thick, negatively curved* P*-manifold of dimension* $n \geq 6$, *and* $\Gamma = \pi_1(X)$. *Then the canonical map* $\mathrm{Homeo}(X) \to \mathrm{Out}(\Gamma)$ *is surjective.*

We now outline the layout of this paper. In Sec. 2, we will give a proof of Theorem 1.1. The argument relies heavily on a characterization of n-dimensional Sierpinski curves ($n \neq 4$) due to Cannon [4]. The dimension restriction in Theorem 1.1 arises from the corresponding dimension restriction in Cannon's work. We note that Ruane [17] used Cannon's theorem in a similar manner to characterize CAT(0)-boundaries for non-uniform lattices $\Gamma \leq \mathrm{SO}(n, 1)$ acting co-compactly on a Γ-equivariantly truncated $\mathbb{H}^n$.

In Sec. 3, we will give a proof of Theorem 1.2. The argument relating the diffeomorphism type of the universal cover with the number of boundary components is Morse theoretic in nature. The analysis of the possible number of boundary components relies on some elementary geometric properties of non-positively curved spaces.

Finally in Sec. 4, we will discuss the proof of Theorem 1.3, as well as the proofs of the three corollaries. The arguments for these follow almost verbatim from previous results of the author [12, 13]. As such, we content ourselves with outlining the arguments from our previous paper, detailing how our Theorem 1.1 allows us to extend our previous results to the present setting.

2. Characterizations of Boundaries at Infinity

We now proceed to prove Theorem 1.1 from the Introduction. So let M_1, M_2 be a pair of compact, negatively curved manifolds of dimension $n \neq 5$, with non-empty totally geodesic boundary. We want to establish that $\partial^\infty \tilde{M}_1$ is homeomorphic to $\partial^\infty \tilde{M}_2$. In order to do this, we will make use of the characterization of Sierpinski curves due to Cannon [4] (generalizing a classic result of Whyburn [21] in dimension $n = 2$). We first start with a definition:

Definition 2.1. Let $\{U_i\}$ be a countable collection of pairwise disjoint subsets of S^n satisfying the following four conditions:

(1) the collection $\{U_i\}$ forms a null sequence, i.e. $\lim\{\mathrm{diam}(U_i)\} = 0$,
(2) $S^n - U_i$ is an n-cell for each i,
(3) $\mathrm{Cl}(U_i) \cap \mathrm{Cl}(U_j) = \emptyset$ for each $i \neq j$ (Cl denotes closure),
(4) $\mathrm{Cl}(\bigcup U_i) = S^n$.

Then we call the complement $S^n - \bigcup U_i$ an $(n-1)$-dimensional *Sierpinski curve* (abbreviated to $\mathcal{S}$-curve).

Theorem 2.1. (Cannon, [4]) *Let X, Y be an arbitrary pair of $(n-1)$-dimensional $\mathcal{S}$-curves $(n \neq 4)$. Then we have*:

- *X is homeomorphic to Y,*
- *if $i : X \to S^n$ is an arbitrary embedding, then $i(X) \subset S^n$ is an $(n-1)$-dimensional $\mathcal{S}$-curve,*
- *if $h : X \to Y$ is an arbitrary homeomorphism, then h extends to a homeomorphism of the ambient n-dimensional spheres.*

The scheme of the proof of Theorem 1.1 is now clear: considering the double DM_i of the manifold M_i across its boundary, we can view $\tilde{M}_i$ as a totally geodesic subset of $\widetilde{DM}_i$, and hence $\partial^\infty \tilde{M}_i$ as an embedded subset of $\partial^\infty \widetilde{DM}_i \cong S^{n-1}$. If we can establish that $\partial^\infty \tilde{M}_i$ is an $(n-2)$-dimensional $\mathcal{S}$-curve, Cannon's theorem will immediately imply that $\partial^\infty \tilde{M}_1$ is homeomorphic to $\partial^\infty \tilde{M}_2$. We now proceed to verify the four conditions of an $(n-2)$-dimensional $\mathcal{S}$-curve for $\partial^\infty \tilde{M} \subset \partial^\infty \widetilde{DM} \cong S^{n-1}$.

Let us first fix some notation: the collection $\{U_i\}$ will be the connected components of $\partial^\infty \widetilde{DM} - \partial^\infty \tilde{M}$ inside $\partial^\infty \widetilde{DM} \cong S^{n-1}$. We will denote by $\{Y_i\}$ the connected components of $\widetilde{DM} - \tilde{M}$. Note that each $Cl(Y_i)$ intersects $\tilde{M}$ along a boundary component, which is a totally geodesic codimension one submanifold of $\widetilde{DM}$. We will denote by $Z_i \subset \partial \tilde{M}$ the boundary component corresponding to $Y_i \subset \widetilde{DM} - \tilde{M}$. Finally, we observe that each U_i can be identified with a corresponding $\partial^\infty Y_i - \partial^\infty Z_i$, for some suitable component Y_i.

Condition 1. The collection $\{U_i\}$ forms a null sequence.

Proof. At the cost of rescaling the metric on DM, we may assume that the sectional curvature is bounded above by -1, and hence that $\widetilde{DM}$ is a $CAT(-1)$ space. In this situation, Bourdon [2] defined a metric on $\partial^\infty \widetilde{DM}$ inducing the standard topology on $\partial^\infty \widetilde{DM} \cong S^{n-1}$. The metric is given by:

$$d_\infty(p, q) = e^{-d(*, \gamma_{pq})},$$

where γ_{pq} is the unique geodesic joining the points $p, q \in \partial^\infty \widetilde{DM}$, $* \in DM$ a chosen basepoint (and d denotes the distance inside $\widetilde{DM}$). Note that different choices of basepoints result in metrics which are Lipschitz equivalent. For convenience, we will pick the basepoint $*$ to lie in the interior of the lift $\tilde{M}$.

Now consider one of the components U_i, and let us try to estimate $\mathrm{diam}(U_i)$. Note that given any two points $p, q \in \mathrm{Cl}(U_i)$, we have that the geodesic $\gamma_{pq} \subset \mathrm{Cl}(Y_i)$, where Y_i is the component corresponding to U_i. In particular, we see that $d(*, \gamma_{pq}) \geq d(*, Z_i)$, and hence that for any $p, q \in \mathrm{Cl}(U_i)$ we have the upper bound:

$$d_\infty(p, q) = e^{-d(*, \gamma_{pq})} \leq e^{-d(*, Z_i)}.$$

Since $\text{diam}(U_i)$ is the supremum of $d_\infty(p,q)$, where $p,q \in \text{Cl}(U_i)$, the above bound yields $\text{diam}(U_i) \leq e^{-d(*,Z_i)}$. On the other hand, since $\tilde{M}$ is the universal cover of a *compact* negatively curved manifold with non-empty boundary, we have that $\lim\{d(*,Z_i)\} = \infty$, where Z_i ranges over the boundary components of $\tilde{M}$. This implies that the collection $\{U_i\}$ forms a null sequence in $\partial^\infty \widetilde{DM} \cong S^{n-1}$, as desired. □

Condition 2. $S^{n-1} - U_i$ is an $(n-1)$-cell for each i.

Proof. Recall that there exists a homeomorphism $\pi_x : S^{n-1} \cong \partial^\infty \widetilde{DM} \to T^1_x \widetilde{DM} \cong S^{n-1}$, obtained by mapping a point $p \in \partial^\infty \widetilde{DM}$ to the unit vector $\dot{\gamma}_{xp}(0)$, where γ_{xp} is the unit speed geodesic ray originating from x, in the direction $p \in \partial^\infty \widetilde{DM}$. Now let U_i be given, and pick x to lie on the corresponding Z_i. Note that under the homeomorphism π_x, we have that $\partial^\infty Z_i$ maps homeomorphically to a totally geodesic $S^{n-2} \subset S^{n-1} \cong T^1_x \widetilde{DM}$, while the subset U_i maps homeomorphically to one of the open hemispheres determined by $\pi_x(\partial^\infty Z_i)$. This forces $\partial^\infty \widetilde{DM} - U_i$ to map homeomorphically to one of the *closed* hemispheres determined by $\pi_x(\partial^\infty Z_i)$, and hence must be an $(n-1)$-cell, as desired. □

Condition 3. $\text{Cl}(U_i) \cap \text{Cl}(U_j) = \emptyset$ for all $i \neq j$.

Proof. Note that by definition we have that $U_i \cap U_j = \emptyset$, and that $\text{Cl}(U_i) = U_i \cup \partial^\infty Z_i$, $\text{Cl}(U_j) = U_j \cup \partial^\infty Z_j$. Hence it is sufficient to show that $\partial^\infty Z_i \cap \partial^\infty Z_j = \emptyset$ for $i \neq j$ (since these are codimension one spheres in $S^{n-1} \cong \partial^\infty \widetilde{DM}$, with the U_i, U_j connected components of the respective complements). But a pair of distinct boundary components of $\tilde{M}$, the universal cover of a *compact* negatively curved manifold with non-empty totally geodesic boundary, must diverge exponentially (with growth rate bounded below in terms of the upper bound on sectional curvature). In particular, no geodesic ray in Z_i is within bounded Hausdorff distance of a geodesic ray in Z_j, and hence the boundaries at infinity are pairwise disjoint, as desired. □

Condition 4. $\text{Cl}(\bigcup U_i) = S^{n-1}$.

Proof. Fix a point $x \in \tilde{M}$, and consider the homeomorphism $\pi_x : S^{n-1} \cong \partial^\infty \widetilde{DM} \to T^1_x \widetilde{DM} \cong S^{n-1}$. We will show that every point in $T^1_x \widetilde{DM} \cong S^{n-1}$ can be approximated by a sequence of points in $\pi_x(U_i)$. This will imply that $T^1_x \widetilde{DM} = \text{Cl}(\bigcup \pi_x(U_i))$, and since π_x is a homeomorphism, Condition 4 will follow.

Now if $p \in T^1_x \widetilde{DM}$ lies in one of the $\pi_x(U_i)$, we are done, so let us assume that $p \in T^1_x \widetilde{DM} - \bigcup \pi_x(U_i)$. Let γ be a unit speed geodesic ray originating from x with tangent vector p at the point x. Note that we have that $\gamma \subset \tilde{M} \subset \widetilde{DM}$, since we are assuming $p \in T^1_x \widetilde{DM} - \bigcup \pi_x(U_i)$. Now observe that $\tilde{M}$ is the universal cover of a *compact* negatively curved manifold with non-empty totally geodesic boundary, and hence there exists a constant K with the property that every point in $\tilde{M}$ is within distance K of $\partial \tilde{M} = \bigcup Z_i$ (for instance take $K = \text{diam}(M)$).

So for each integer $k \in \mathbb{N}$, we can find a point $y_k \in \partial\tilde{M}$ satisfying $d(\gamma(k), y_k) \leq K$. If η_k is the geodesic ray originating from x and passing through y_k, we have that $\eta_k(\infty) \in U_{i_k}$, where Z_{i_k} is the component of $\partial\tilde{M}$ containing the point y_k. This implies that $\dot{\eta}_k(0) \in T^1_x\widetilde{DM}$ lies in the corresponding $\pi_x(U_{i_k})$, i.e. that the sequence of vectors $\{\dot{\eta}_k(0)\} \subset T^1_x\widetilde{DM}$ lies in the set $\bigcup \pi_x(U_i)$. We are left with establishing that $\lim\{\dot{\eta}_k(0)\} = p$. To see this, we need to estimate the angle between the geodesics η_k and the geodesic γ. But this is easy to do: consider the geodesic triangle with vertices $(x, \gamma(k), y_k)$, and note that $d(x, \gamma(k)) = k$, while $d(\gamma(k), y_k) \leq K$. Applying the Alexandrov–Toponogov triangle comparison theorem, we see that the angle $\angle(\dot{\eta}_k(0), \dot{\gamma}(0))$ is bounded above by the angle of a comparison triangle in $\mathbb{H}^2$ (recall that we assumed the metrics have been scaled to have upper bound -1 on the sectional curvature). But an easy calculation in hyperbolic geometry shows that if one has a sequence of triangles in $\mathbb{H}^2$ of the form (A_k, B_k, C_k) with the property that $d(A_k, B_k) = k$ and $d(B_k, C_k) \leq K$, then the angle at the vertex A_k tends to zero as k tends to infinity. This implies that $\lim\{\angle(\dot{\eta}_k(0), \dot{\gamma}(0))\} = 0$, and hence completes the proof of Condition 4. □

Appealing to Cannon's theorem now immediately yields Theorem 1.1: if M_1, M_2 are a pair of compact, n-dimensional ($n \neq 5$), negatively curved manifolds with non-empty, totally geodesic boundary, then $\partial^\infty\tilde{M}_1, \partial^\infty\tilde{M}_2$ are a pair of $(n-2)$-dimensional $\mathcal{S}$-curves, and hence are homeomorphic to each other.

Remark. We point out that Theorem 1.1 can be used to give a proof of a weak form of Theorem 1.2 under some stricter dimension and curvature hypotheses. The rough outline of such an argument is as follows: taking two such manifolds M_1, M_2, Theorem 1.1 tells us that $\partial^\infty\tilde{M}_1$ is homeomorphic to $\partial^\infty\tilde{M}_2$. Fixing a pair of points $p_i \in \mathrm{Int}(\tilde{M}_i)$, one can use the homeomorphism between the pair of $\partial^\infty\tilde{M}_i$ to "radially extend" to a homeomorphism between a pair of subsets $C_i \subset \tilde{M}_i$, each of which is homeomorphic to the cone over the corresponding $\partial^\infty\tilde{M}_i$ (and where each p_i is the cone point of the corresponding C_i). Now when $n \geq 3$, the complements of C_i in $\tilde{M}_i$ can be easily seen to decompose into countably many connected components, one for each component of the boundary $\partial\tilde{M}_i$. Furthermore, the closure of each of these components can be shown to be homeomorphic to $\mathbb{R}^{n-1} \times [0,1]$, with the subset $\mathbb{R}^{n-1} \times \{1\}$ contained in C_i, and the subset $\mathbb{R}^{n-1} \times \{0\}$ corresponding to a component of $\partial\tilde{M}_i$. With some work, one can see that the complements of C_1 in $\tilde{M}_1$ attach to C_1 in precisely the same manner as the complements of C_2 attach to $\tilde{M}_2$, allowing the homeomorphism between the C_i to extend to a homeomorphism between the $\tilde{M}_i$. Note that the argument sketched out here can only *a priori* give homeomorphism information (though see the remark at the end of Sec. 3.1), since it is obtained by "extending inwards" the homeomorphism between the boundaries at infinity (which are fairly pathological spaces). We omit the details of this argument, since the considerably stronger Theorem 1.2 will be established (via completely different methods) in the next section.

3. Generalized Cartan–Hadamard Theorem

In this section, we provide a proof of Theorem 1.2. Let us first recall that there are two components to Theorem 1.2:

- a characterization of the diffeomorphism type of $\tilde{M}$ in terms of the number of boundary components of $\tilde{M}$, and
- a count of the possible number of boundary components of $\tilde{M}$

where $\tilde{M}$ is the universal cover of a compact Riemannian manifold M of non-positive curvature, with totally geodesic boundary. Note that the case where the manifold M is closed is classical, hence we will assume throughout this section that $\partial M \neq \emptyset$. We argue each of the two portions of Theorem 1.2 separately, as they require drastically different techniques.

3.1. *Characterization of universal covers*

In order to establish the characterization of universal covers in terms of the number of boundary components, we make use of Morse theory. This approach is philosophically very different from the argument sketched out in the remark at the end of the previous section, since instead of "extending inwards" from the boundary at infinity, we will be "growing outwards" our diffeomorphism.

We first observe that, since $\tilde{M}$ is a manifold with boundary, it has a canonical stratification with two strata: the interior $\mathrm{Int}(\tilde{M})$ of $\tilde{M}$, and the boundary $\partial\tilde{M}$. In addition, since the boundary ∂M is totally geodesic inside M, one can embed M as a totally geodesic codimension zero submanifold of the double DM. Lifting, we have a natural totally geodesic embedding of the universal cover $\tilde{M}$ inside the Riemannian manifold $\widetilde{DM}$ (which we know is diffeomorphic to $\mathbb{R}^n$). Our plan is now to use a suitable version of Morse theory to analyze the topology of $\tilde{M}$. The function we will use will be the square of the distance to a suitable point $p \in \mathrm{Int}(\tilde{M})$. The next two Claims establish the existence of a suitable point p.

Claim 1. *There exists a point $p \in \mathrm{Int}(\tilde{M})$ such that for every pair of distinct boundary components $N, N' \subset \partial\tilde{M}$, we have that $d(p, N) \neq d(p, N')$.*

Proof. (Claim 1) To see this, we first note that given any pair N, N' of distinct boundary components, the set of points $q \in \widetilde{DM}$ satisfying $d(q, N) = d(q, N')$ is a codimension one submanifold of $\widetilde{DM}$. Indeed, we can consider the smooth function $\phi : \widetilde{DM} \to \mathbb{R}$ given by $\phi(x) := d(x, N)^2 - d(x, N')^2$, and observe that the set of points we are interested in is just the pre-image set $\phi^{-1}(0)$. Hence to show that this is a submanifold, we just need to establish that 0 is a regular value of the smooth map ϕ. So let $x \in \widetilde{DM}$ satisfy $\phi(x) = 0$, and observe that, since N, N' are totally geodesic submanifolds and $\tilde{M}$ is simply connected of non-positive curvature, there exists a *unique* pair of minimal length geodesic segments γ, γ' emanating from x, and terminating on N, N' respectively.

Now consider the unit tangent vectors $v, v' \in T_x\widetilde{DM}$ tangent to γ, γ'. From the explicit form of ϕ, we observe that the corresponding differential $d\phi : T_x\widetilde{DM} \to T_0\mathbb{R} \cong \mathbb{R}$ is given by the concrete expression:

$$d\phi(w) = 2D \cdot \langle w, v - v' \rangle_x,$$

where $w \in T_x\widetilde{DM}$ is arbitrary, D is the distance from x to N, and the inner product is taken with respect to the Riemannian metric on $\tilde{M}$. Finally, we observe that if x was *not* a regular point for the map ϕ, then $d\phi$ would have to be identically zero on $T_x\widetilde{DM}$. This would imply that $v - v' = 0$, and hence that $v = v'$, which in turn would force $\gamma = \gamma'$. But this contradicts the fact that N, N' were *distinct* boundary components.

Now the inverse function theorem implies that the set of points $\phi^{-1}(0)$ we are interested in is in fact a smooth submanifold of codimension one. Finally, since there are only countably many pairs of boundary components, one sees that the set of points E where some $d(q, N) = d(q, N')$ lies on a countable union of codimension one submanifolds, and hence has measure zero in $\widetilde{DM}$. Since $\mathrm{Int}(\tilde{M})$ ia an open set in $\widetilde{DM}$, this implies that there exists a point in $p \in \mathrm{Int}(\tilde{M}) - E$, and it is immediate from the definition of E that the point p has the desired property. □

Claim 2. *For the point p chosen above, the set of distances from p to the connected components of* $\partial\tilde{M}$ *forms a discrete subset of* $\mathbb{R}^+$.

Proof. (Claim 2) Let us assume that the set of distances from p to the connected components of $\tilde{M}$ have an accumulation point, and argue by contradiction. Pick $r > 0$ such that the metric ball $B_p(r)$ intersects infinitely many boundary components $\{N_i\}$. Now for each N_i, define the subset U_i to be the set of directions, in $T_p^1(\tilde{M})$, corresponding to geodesic segments joining p to points in N_i. Note that each $U_i \subset T_p^1\tilde{M}$ is (topologically) an open ball inside $T_p^1\tilde{M} \cong S^{n-1}$, and that the collection of subsets $\{U_i\}$ are pairwise disjoint in $T_p^1\tilde{M}$. We now argue that each U_i *contains* a metric ball V_i of radius a fixed $\delta > 0$, which will obviously give us a contradiction, as the entire sphere $T_p^1\tilde{M}$ has finite volume, and hence can only contain finitely many such pairwise disjoint metric balls.

To establish this result, we first note that every U_i contains a distinguished point x_i, consisting of the direction corresponding to the *unique* minimal length geodesic joining p to the corresponding N_i. We will use the point x_i as the center for our metric balls V_i. Now note that each of the open sets U_i can be uniquely identified by its boundary $\partial U_i \subset T_p^1\tilde{M}$ (homeomorphic to S^{n-2}), hence it is sufficient for us to establish that the distance from x_i to ∂U_i is uniformly bounded from below. Observe that the distance in the unit tangent space $T_p^1\tilde{M}$ is given by the angle between the corresponding vectors. To bound this angle from below, we make use of the Alexander–Toponogov triangle comparison theorem: a point in ∂U_i is a limit of points inside U_i, corresponding to a sequence of points $\{y_k\}$ in N_i whose distance from the point x_i tends to infinity. Considering the sequence of triangles with vertices $\{p, x_i, y_k\}$ (corresponding to the sequence of points $\{y_k\}$) one can use the *lower bound* κ on sectional curvatures (recall that $\tilde{M}$ is the universal cover of a compact manifold M) to construct a sequence of comparison triangles $\{\bar{p}, \bar{x}_i, \bar{y}_k\}$ in $\mathbb{H}^2_\kappa$, the constant κ-curvature space. These comparison triangles are built to have $d_{\mathbb{H}^2_\kappa}(\bar{p}, \bar{x}_i) = d(p, x_i), d_{\mathbb{H}^2_\kappa}(\bar{x}_i, \bar{y}_k) = d(x_i, y_k)$, and $\angle_{\mathbb{H}^2_\kappa}(\bar{x}_i) = \angle(x_i)\pi/2$. The triangle comparison theorem tells us that the angle $\angle(p)$ of the triangle $\{p, x_i, y_k\}$ at the vertex p is at least as large as the angle $\angle_{\mathbb{H}^2_\kappa}(\bar{p})$ of the comparison triangle $\{\bar{p}, \bar{x}_i, \bar{y}_k\}$ at the vertex $\bar{p}$. But observe that we have $d_{\mathbb{H}^2_\kappa}(\bar{p}, \bar{x}_i) = d(p, x_i) \leq r$, while $d_{\mathbb{H}^2_\kappa}(\bar{x}_i, \bar{y}_k) = d(x_i, y_k) \to \infty$ as $k \to \infty$. A direct computation shows that for such triangles in $\mathbb{H}^2_\kappa$, the angles at the vertex $\bar{p}$ approach a limiting value $\delta > 0$. This implies that each of the sets U_i contains an open metric ball, centered at x_i, of radius $\delta > 0$, giving us the desired contradiction.

We conclude that each ball centered at p intersects only finitely many boundary components, and hence the collection of distances from p to the boundary components does indeed form a discrete subset in $\mathbb{R}^+$. □

Having established the existence of a point p as in Claim 1, we can now consider the function $\phi : \widetilde{DM} \to \mathbb{R}$ given by $\phi(-) = d^2(p, -)$. Note that ϕ is a proper function, and by the classical Cartan–Hadamard theorem is smooth on $\widetilde{DM}$, with a single critical point (a minimum) at $p \in \mathrm{Int}(\tilde{M}) \subset \widetilde{DM}$. In particular, ϕ defines a proper Morse function on $\widetilde{DM}$. Let us denote by f the restriction of ϕ to $\tilde{M}$. We now plan on using the function f to analyze the topology of $\tilde{M}$, a non-compact manifold with boundary. In order to do this, we will use Morse theory for manifolds with boundary.

Let us now briefly recall the definition of a Morse function in the setting of manifolds with boundary. Given a manifold with boundary M, embedded as a smooth submanifold of $\mathbb{R}^N$, and a smooth function $\phi : \mathbb{R}^N \to \mathbb{R}$, we denote by f the restriction $f := \phi|_M : M \to \mathbb{R}$. We will denote by ∂f the further restriction of f to the smooth submanifold $\partial M \subset \mathbb{R}^N$. Restricting f to the interior of M, we obtain a smooth function f° on a non-compact manifold $\mathrm{Int}(M)$ with empty boundary; let $\mathrm{Crit}(f^\circ) \subset \mathrm{Int}(M)$ denote the critical points of this function. Furthermore, we can also consider the smooth function ∂f on ∂M; let $\mathrm{Crit}(\partial f) \subset \partial M$ denote the critical points of this function. We now say that f is a Morse function provided that:

(1) the restriction f° is a Morse function on $\mathrm{Int}(M)$ (in the classical sense),
(2) the restriction ∂f is a Morse function on ∂M (in the classical sense), and
(3) $\mathrm{Crit}(f^\circ)$ has no accumulation points in M.

For such a Morse function f, the set of *critical points* of f is just the union $\mathrm{Crit}(f) := \mathrm{Crit}(f^\circ) \coprod \mathrm{Crit}(\partial f)$.

Our next step is to verify that the function f, obtained by restricting the function $\phi(-) := d^2(p, -)$ from $\widetilde{DM} \cong \mathbb{R}^n$ to $\tilde{M}$, is indeed a Morse function on the manifold with boundary $\tilde{M}$. We first observe that the function ϕ on $\widetilde{DM}$ is Morse, and has a unique critical point, which is a minimum occurring at $p \in \mathrm{Int}(\tilde{M})$. In particular, we have that $\mathrm{Crit}(f^\circ) = \{p\}$. We now need to identify the set $\mathrm{Crit}(\partial f)$. Note that since $\mathrm{Crit}(f^\circ)$ consists of a single point, condition (3) holds vacuously.

Claim 3. *The function $\partial f : \partial\tilde{M} \to \mathbb{R}$ has one critical point on each component N of $\partial\tilde{M}$. Furthermore, each of these critical points is a minimum.*

Proof. (Claim 3) To see this, we first observe that, since $N \subset \tilde{M}$ is a totally geodesic submanifold, the non-positive curvature hypothesis forces the existence of a unique point x realizing $d(p,x) = d(p,N)$. This point will clearly be the unique global minimum of the function f restricted to N, completing the second point of the Claim. So we are left with arguing that f has no other critical points. This is of course equivalent to showing that for all $y \neq x$ with $y \in N$, the restriction $f|_N : N \to \mathbb{R}^+$ has a nonzero gradient at the point y. But observe that the gradient $\nabla f|_N(y)$ of the restricted function $f|_N$ is simply the projection of the gradient $\nabla f(y)$ of the original function f to the tangent space T_yN. Hence it is sufficient to argue that $\nabla f(y)$ fails to be perpendicular to T_yN. But this is easy to do: take the geodesic triangle $\{p,x,y\}$, and consider the comparison triangle $\{\bar{p},\bar{x},\bar{y}\}$ in $\mathbb{R}^2$. By the Alexander–Toponogov triangle comparison theorem, we know that all the angles in the triangle $\{p,x,y\}$ must be *smaller* than the corresponding angles in the triangle $\{\bar{p},\bar{x},\bar{y}\}$. Note that the angle at the vertex x is $\pi/2$, since x minimizes the distance from p to N (and applying the first variation of energy formula), which tells us that the angle at vertex $\bar{x}$ is $\geq \pi/2$. But the sum of the angles in the Euclidean triangle $\{\bar{p},\bar{x},\bar{y}\}$ is π, hence both the remaining angles must be $< \pi/2$. Since the angle at y is smaller than the angle at $\bar{y}$, we immediately get that the angle at y is likewise $< \pi/2$. Finally, we observe the initial vector of the geodesic segment $\overline{yp}$ is a scalar multiple of the vector $\nabla f(y)$, while the initial vector of the geodesic segment $\overline{yx}$ lies in T_yN. This yields that $\nabla f(y)$ is not perpendicular to T_yN, and hence that y cannot be a critical point of $f|_N$, as desired. □

Having established that the function f is a Morse function, we now want to use this function to understand the topology of $\tilde{M}$. Note that, by the choice of the point p, the critical values of the Morse function f form a discrete subset of $\mathbb{R}^+$, and each critical value corresponds to a unique critical point. Let us denote by $\tilde{M}_r$ the sublevel set $f^{-1}(-\infty, r]$. An illustration of such a sublevel set is given in Fig. 1: $\tilde{M}$ is drawn as a submanifold in $\widetilde{DM}$, and the subset $\tilde{M}_R \subset \tilde{M}$ is shaded. Note that $\tilde{M}_R$ is naturally a manifold with corners, as well as a stratified space, with the codimension one strata (corresponding to $\partial\tilde{M}$) drawn in a darker shade.

Before stating our Morse theoretic result, let us briefly elaborate on the structure of the sublevel sets for a Morse function on a manifold with boundary. First recall that an n-dimensional manifold with corners is a space locally modeled (in the obvious sense) on the subspaces

$$\mathbb{R}^n_k := \{(x_1,\ldots,x_n) \in \mathbb{R}^n \mid x_1 \geq 0,\ldots,x_k \geq 0\} \subset \mathbb{R}^n,$$

where $0 \leq k \leq n$. Observe that $\mathbb{R}^n_0$ is just the usual $\mathbb{R}^n$, while $\mathbb{R}^n_1$ is a standard half-space. The subset of points which locally correspond to the origin in $\mathbb{R}^n_k$ form the

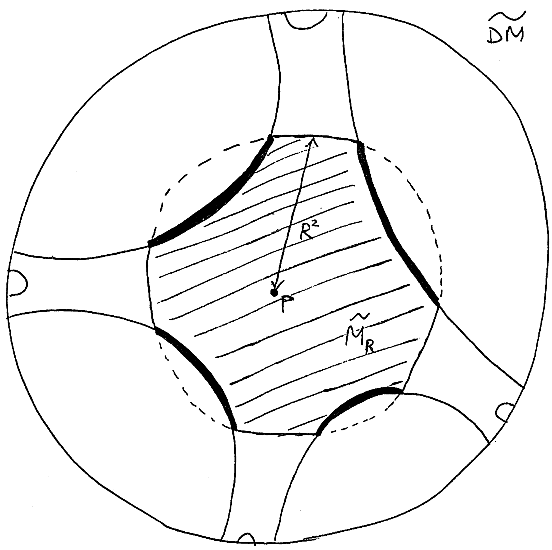

Fig. 1. Stratified manifold with corners $\tilde{M}_R := f^{-1}(-\infty, R]$.

codimension k corner. With this convention in place, a manifold without boundary can be viewed as a manifold with corners, where all corners are of codimension 0. A manifold with boundary can be viewed as a manifold with corners, where all corners are of codimension ≤ 1 (and the boundary of the manifold coresponds precisely to the one-dimensional corner). In particular, we see that in the classical Morse theory, the generic sublevel sets of a Morse function on a manifold with corners of codimension 0 (a manifold without boundary) naturally have the structure of a manifold with corners of codimension ≤ 1 (a manifold with boundary). Similarly, for a Morse function on a manifold with corners of codimension ≤ 1 (a manifold with boundary), generic sublevel sets will have a natural structure of a manifold with corners of codimension ≤ 2. This structure can readily be seen in Fig. 1: the illustration shows $\tilde{M}_R$ as a manifold with corners, with exactly eight points forming the corners of codimension 2.

Note that, if we were to forget the "corner" structure, we can view the sublevel set as a manifold with boundary. The boundary of the sublevel set $\tilde{M}_R = f^{-1}(-\infty, R]$ naturally decomposes into two sets: the set $(\partial f)^{-1}(-\infty, R] \subset \partial\tilde{M}$, along with the set $f^{-1}(R)$. Each of these two sets are $(n-1)$-dimensional manifolds with boundary, and they intersect in the subset $(\partial f)^{-1}(R) = \partial(f^{-1}(R))$. Since we will be considering the sublevel sets $\tilde{M}_R$ for larger and larger values of R, we will need to keep track of the portion of $\partial\tilde{M}_R$ that lies inside the set $\partial\tilde{M}$. This is achieved by imposing a stratification on $\tilde{M}_R$, where the codimension one strata is the subset $(\partial f)^{-1}(-\infty, R] \subset \tilde{M}$.

Next, let us recall the basic results concerning the topology of sublevel sets in the classical setting of Morse functions on closed manifolds. If $f : M \to \mathbb{R}$ is a Morse function, and M_r denotes the sublevel set $M_r = f^{-1}(-\infty, r]$, then we have:

- if the interval $[a, b]$ contains no critical values of f, then there is a diffeomorphism $M_a \cong M_b$,
- if v is the only critical value in the interval $[v-\epsilon, v+\epsilon]$, with a unique corresponding critical point of index k, then $M_{v+\epsilon}$ is diffeomorphic to the space obtained from $M_{v-\epsilon}$ by attaching a k-handle (i.e. $\mathbb{D}^k \times \mathbb{D}^{n-k}$ attached along $\partial\mathbb{D}^k \times \mathbb{D}^{n-k}$), with the attaching corner "smoothed out".

For our purposes, we will need a version of Morse theory for manifolds with boundary. Such a theory has been studied and developed by a variety of authors, including Baiada-Morse [1], Hamm [8], Hamm-Le [9], Siersma [20], and of course, Goresky–MacPherson [7]. Most of these authors have focused on applications of Morse theoretic techniques to problems in algebraic geometry (topology of Stein spaces, Lefschetz theorems), and as such they focus primarily on "coarse" topological data (recognizing Betti numbers, homology, or homotopy type). In our situation, we are seeking more refined data, as we would like to recognize the sublevel sets up to *diffeomorphism*.

We were unable to locate the precise statements we needed in the literature. However, these results seem to be well known to experts, and follow relatively easily from the methods used in Milnor's book [14]. For the convenience of the reader, we provide a brief sketch of the proofs, leaving the details to the interested reader.

Claim 4. *If the interval $[a, b]$ contains no critical values of f, then there is a diffeomorphism of manifolds with corners $\tilde{M}_a \cong \tilde{M}_b$, which furthermore preserves the stratification of these two spaces.*

Proof. (Claim 4) This is shown in a manner similar to the corresponding statement in the classical setting, namely, the diffeomorphism is constructed as the time one flow associated to a suitable vector field. In our situation, we first assert that there exists a smooth vector field X defined on $\tilde{M}_b$ having the following four properties:

(1) X vanishes outside a compact neighborhood K of $f^{-1}[a, b]$, chosen so f has no critical points on K,
(2) at all points p where $X(p) \neq 0$, we have $\langle X, \nabla f \rangle < 0$,
(3) at all points in the codimension one strata X is tangential to the strata,
(4) for the associated flow $\varphi_t : \tilde{M}_b \to \tilde{M}_b$, one has that the time one map takes $\varphi_1(f^{-1}(b)) \subset f^{-1}(a)$.

To see this, we first recall that in the classical setting, an analogous vector field is constructed by taking the negative gradient vector field of the function f, multiplying it by a positive function which vanishes outside of K, and then suitably renormalizing (see e.g. [14, pp. 12–13]). Now the same argument almost works in the setting of manifolds with boundary: one just starts with the vector field $-\nabla\phi$.

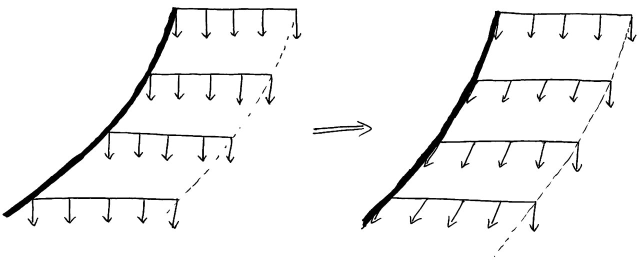

Fig. 2. Modifying $-\nabla f$ to obtain the vector field X.

The only subtlety lies in the fact that $-\nabla\phi$, restricted to $\partial\tilde{M}$, is *not* tangential to $\partial\tilde{M}$, i.e. fails property (3) in our Claim. But this is easy to remedy, since one can use a partition of unity to smoothly transition from $-\nabla\phi$ away from $\partial\tilde{M}$ to $-\nabla(\partial f)$ along the submanifold $\partial\tilde{M}$ (an illustration of the modification is given in Fig. 2). One can again multiply by a function vanishing outside of K, resulting in a vector field satisfying conditions (1)–(3) of our Claim. Finally, at the cost of renormalizing this new vector field we can ensure that the associated flow takes the level set $f^{-1}(b)$ into the level $f^{-1}(a)$, giving us property (4).

Now that we have the vector field X, we proceed to show that the map φ_1 defines a diffeomorphism from the stratified space $\tilde{M}_b$ to the stratified space $\tilde{M}_a$. From the existence and uniqueness of solutions to ODEs, we know that the map φ_1 is injective. Since solutions depend smoothly on the initial conditions, the map φ_1 is also smooth, and by reversing the flow, has smooth inverse. From compactness of $\tilde{M}_b$ (recall that f is proper), we have that ϕ_1 is a diffeomorphism onto its image. So we are left with arguing that $\varphi_1(\tilde{M}_b) = \tilde{M}_a$.

First we argue that $\varphi_1(\tilde{M}_b) \subset \tilde{M}_a$. Property (2) of the vector field X ensures that f is strictly decreasing along flow lines, so that we clearly have $\varphi_1(\tilde{M}_a) \subset \tilde{M}_a$. For points $x \in \tilde{M}_b - \tilde{M}_a$, we note that property (3) ensures that x lies on the flow line of a well-defined $p \in f^{-1}(b)$, i.e. there exists a $0 \le t < 1$ with $\varphi_t(p) = x$. Since $\varphi_1(p) \in \tilde{M}_a$, and f is strictly decreasing along flow lines, we get that

$$f(\varphi_1(x)) = f(\varphi_{1+t}(p)) \le f(\varphi_1(p)) = a \Rightarrow \varphi_1(x) \in \tilde{M}_a.$$

This gives the desired containment $\varphi_1(\tilde{M}_b) \subset \tilde{M}_a$. For later use, we also point out that the argument above establishes that $f^{-1}[a,b]$ is diffeomorphic, as a stratified manifold with corners, to the manifold $f^{-1}(b) \times [0,1]$ (where the codimension one strata is given by $\partial(f^{-1}(b)) \times [0,1]$).

Next, to see that $\varphi_1(\tilde{M}_b) = \tilde{M}_a$, we need to argue surjectivity of the map $\varphi_1 : \tilde{M}_b \to \tilde{M}_a$. This is achieved as follows: forgetting the stratification and the corner structure, we can view $\tilde{M}_b, \tilde{M}_a$ as a pair of oriented manifolds with boundary. We first argue that φ_1 restricts to a homeomorphism between the boundaries. As we discussed earlier, there are natural decompositions: $\tilde{M}_b = (\partial f)^{-1}(-\infty, b] \cup f^{-1}(b)$, and $\tilde{M}_a = (\partial f)^{-1}(-\infty, a] \cup f^{-1}(a)$. By construction, we see that φ_1, restricted to $\partial\tilde{M}$, coincides with the diffeomorphism from the classical Morse setting (see

[14, pp. 12–13]) from $(\partial f)^{-1}(-\infty, b]$ to $(\partial f)^{-1}(-\infty, a]$. Property (4) of the vector field X tells us that φ_1 maps the manifold with boundary $f^{-1}(b)$ *into* $f^{-1}(a)$. These are manifolds with boundary, and φ_1 restricts to a diffeomorphism between their boundaries (as these coincide with $(\partial f)^{-1}(b)$, $(\partial f)^{-1}(a)$ respectively). A degree argument now tells us that φ_1 maps $f^{-1}(b)$ *onto* $f^{-1}(a)$. This now tells us that φ_1 restricts to a homeomorphism from $\partial\tilde{M}_b$ to $\partial\tilde{M}_a$, and again, a degree argument allows us to conclude that φ_1 is surjective. The fact that φ_1 is strata preserving follows immediately from property (3) of the vector field X. We furthermore observe that the collection of maps φ_t, $0 \leq t \leq 1$, define a smooth, strata preserving, deformation retraction from $\tilde{M}_b$ to $\tilde{M}_a$. This concludes the sketch of our proof of Claim 4. □

Our next goal is to relate the diffeomorphism type of $\tilde{M}_{v+\epsilon}$ with that of $\tilde{M}_{v-\epsilon}$, when the interval $[v-\epsilon, v+\epsilon]$ contains the single critical value v. This is the content of our:

Claim 5. *If v is the only critical value in the interval $[v-\epsilon, v+\epsilon]$, with a unique corresponding critical point x lying on $\partial\tilde{M}$, then $\tilde{M}_{v+\epsilon}$ is diffeomorphic to the stratified manifold with corners obtained from $\tilde{M}_{v-\epsilon}$ by attaching the stratified manifold with corners $[0,1] \times \mathbb{D}^{n-1}$ along an embedding of the subspace $\{1\} \times \mathbb{D}^{n-1} \hookrightarrow \mathrm{Int}(f^{-1}(v-\epsilon))$, with the attaching corner "smoothed out". The codimension one strata of $[0,1] \times \mathbb{D}^{n-1}$ consists of the set $\{0\} \times \mathbb{D}^{n-1}$.*

Proof. (Claim 5) We now sketch out how this result can be deduced from the analogous statement in the classical form of Morse theory. Let N denote the boundary component of $\tilde{M}$ containing the critical point x. Take a second copy of $\tilde{M}$, which we denote $\tilde{M}'$. Corresponding to the boundary component N, we have a boundary component $N' \subset \tilde{M}'$. We define $\bar{M}$ to be the smooth manifold obtained by gluing together $\tilde{M}$ and $\tilde{M}'$, where the gluing is obtained by identifying N with N'. Observe that there is a natural $\mathbb{Z}_2$-action on $\bar{M}$, which interchanges the two copies of $\tilde{M}$; if $w \in \bar{M}$, we will denote by $w' \in \bar{M}$ the image of w under the canonical involution. We can now define a natural $\mathbb{Z}_2$-invariant function $f \cup f'$ on $\bar{M}$, defined by:

$$(f \cup f')(w) = \begin{cases} f(w) & w \in \tilde{M}, \\ f(w') & w \in \tilde{M}'. \end{cases}$$

Note that the function $f \cup f'$ is smooth on the complement of $N \subset \bar{M}$. We can now equivariantly smooth $f \cup f'$ in an arbitrarily small neighborhood of N, resulting in a $\mathbb{Z}_2$-equivariant, smooth function $\bar{f} : \bar{M} \to \mathbb{R}$. This smoothing can also be chosen so as to *not* introduce any new critical points in the subset $\tilde{M} \subset \bar{M}$ (and hence, by equivariance, inside $\tilde{M}' \subset \bar{M}$). We now have the following three observations:

- the sublevel set $\bar{M}_{v-\epsilon} := \bar{f}^{-1}(-\infty, v-\epsilon]$ is diffeomorphic to the disjoint union of two copies of the sublevel set $\tilde{M}_{v-\epsilon}$,

- the sublevel set $\bar{M}_{v+\epsilon} := \bar{f}^{-1}(-\infty, v+\epsilon]$ is diffeomorphic to two copies of the sublevel set $\tilde{M}_{v+\epsilon}$, with the two copies glued together along the two copies of $N \cap \tilde{M}_{v+\epsilon}$,
- the function $\bar{f}$ contains a single critical value in the interval $[v-\epsilon, v+\epsilon]$, with the unique corresponding critical point $x \in N \subset \bar{M}$ having index $= 1$.

The first two observations are obtained by suitably choosing the smoothing $\bar{f}$ (close enough to $f \cup f'$ and having the same critical points). The third observation can be seen as follows: since no new critical points are introduced, we know that the only potential critical point in the set $\bar{f}^{-1}[v-\epsilon, v+\epsilon]$ occurs at the point $x \in N \subset \bar{M}$ (which we recall was the unique critical point of f located on the boundary component N). On the other hand, we know that their is a change in the topology of the sublevel sets, and hence there *must* exist a critical point in the set $\bar{f}^{-1}[v-\epsilon, v+\epsilon]$ (by Claim 4), telling us that x is indeed a critical point. Since x lies in the *interior* of the manifold with boundary $\bar{M}$, we find ourselves back in the classical setting. Now the sublevel sets for $\bar{f}$ go from being disconnected (at height $v-\epsilon$) to being connected (at height $v+\epsilon$), so we conclude that the critical point x must have index $= 1$.

Classical Morse theory tells us that there is a diffeomorphism between $\bar{M}_{v+\epsilon}$ and the space obtained from $\bar{M}_{v-\epsilon}$ by attaching a 1-handle (see [14, pp. 14–17]. More precisely, the classical proof constructs a submanifold of $\bar{M}_{v+\epsilon}$ which is (1) a smooth deformation retract, and (2) diffeomorphic to $\bar{M}_{v-\epsilon}$ along with a 1-handle attached. Now from the fact that $\bar{f}$ is $\mathbb{Z}_2$-equivariant, each sublevel set is automatically $\mathbb{Z}_2$-invariant. But now we observe that the proof given in Milnor [14, pp. 14–17], when applied to our *equivariant* function, actually guarantees $\mathbb{Z}_2$-equivariance of the smooth deformation retraction, as well as $\mathbb{Z}_2$-invariance of the submanifold. To achieve this, we merely need to ensure that the local coordinate chart chosen in [14, p. 15] satisfies the obvious $\mathbb{Z}_2$-invariance, i.e. in terms of the local coordinates $\{u^1, \dots, u^n\}$, the involution takes the form $u^1 \mapsto -u^1$. With such a choice of local coordinate chart, it is easy to verify that equivariance is preserved throughout the rest of the argument.

Finally, to conclude our sketch, we note that we can recover $\tilde{M}_{v+\epsilon}$ from $\bar{M}_{v+\epsilon}$, since the $\mathbb{Z}_2$-action merely interchanges the two copies of $\tilde{M}_{v+\epsilon}$ inside $\bar{M}_{v+\epsilon}$ by reflecting across the fixed set $N \cap \tilde{M}_{v+\epsilon}$. But the sublevel set $\bar{M}_{v+\epsilon}$ can be $\mathbb{Z}_2$-equivariantly smoothly retracted onto a subset diffeomorphic to two copies of $\tilde{M}_{v-\epsilon}$, joined by a 1-handle $\mathbb{D}^1 \times \mathbb{D}^{n-1}$. Recall that in terms of the local coordinate system, the $\mathbb{D}^1$ factor corresponds to the u^1-coordinate. In particular, we see that the $\mathbb{Z}_2$-action on this subset interchanges the two copies of $\tilde{M}_{v-\epsilon}$, and on the 1-handle, acts via a flip ($u^1 \mapsto -u^1$) on the $\mathbb{D}^1$-factor. The fixed set of the involution is thus the subset $\{0\} \times \mathbb{D}^{n-1} \subset \mathbb{D}^1 \times \mathbb{D}^{n-1}$, and the two half spaces determined by the reflection across this fixed set are each diffeomorphic to $\tilde{M}_{v-\epsilon}$ with a "half" 1-handle attached, as asserted in our Claim. □

An illustration of this retraction is given in Fig. 3. The shaded region represents the "half" 1-handle $[0, 1] \times \mathbb{D}^{n-1}$, attached to the sublevel set $\tilde{M}_{v-\epsilon}$, all lying within

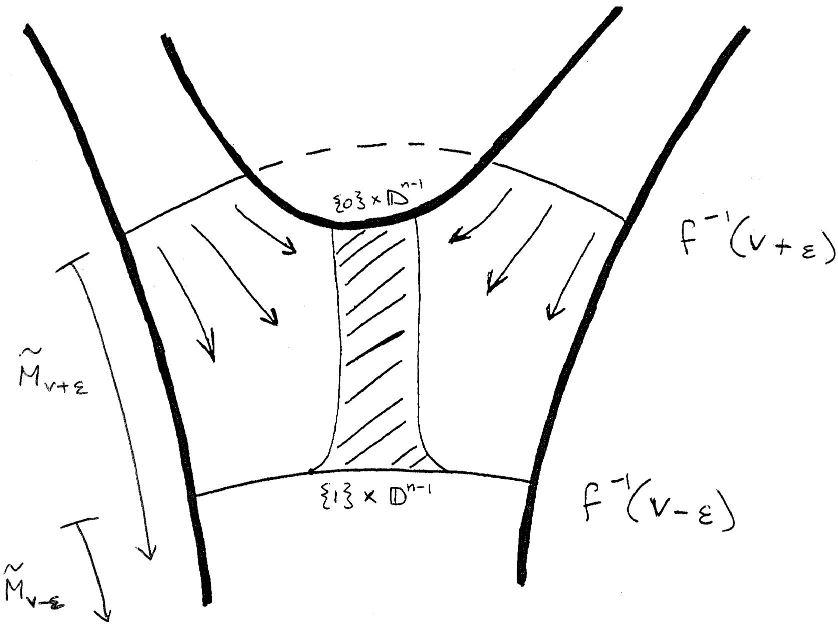

Fig. 3. Change of topology across a boundary critical point.

the ambient sublevel set $\tilde{M}_{v+\epsilon}$. The codimension one strata (corresponding to $\partial\tilde{M}$) is indicated in a heavier shade. Finally, the arrows indicate the deformation retraction from the sublevel set $\tilde{M}_{v+\epsilon}$ to the set $\tilde{M}_{v-\epsilon}$ with the "half" 1-handle attached.

At this point, we have an efficient way to describe the diffeomorphism type of $\tilde{M}$ via the Morse function f. We now return to our original purpose: given two manifolds M_1, M_2 satisfying the hypotheses of our theorem, with $\tilde{M}_1$ having the same number of boundary components as $\tilde{M}_2$, we want to establish a diffeomorphism between the universal covers.

To start out, we note that we can choose points $p_i \in \tilde{M}_i$ so that the corresponding Morse functions f_i have precisely the same number of critical points (by hypothesis, combined with Claim 3). In particular, since the set of critical values for each of the two functions f_i is a discrete subset of $[0, \infty)$ (Claim 2), one can choose a diffeomorphism $r : [0, \infty) \to [0, \infty)$ with the property that $x \in [0, \infty)$ is a critical value of f_1 if and only if $r(x) \in [0, \infty)$ is a critical value of f_2. We let $0 = \lambda_0 < \lambda_1 < \lambda_2 < \cdots$ be the sequence of critical values of the Morse function f_1, and let $\mu_i = (\lambda_i + \lambda_{i+1})/2$. We denote by $(\tilde{M}_1)_i$ the stratified manifold with corners $f_1^{-1}[0, \mu_i]$, and by $(\tilde{M}_2)_i$ the stratified manifold with corners $f_2^{-1}[0, r(\mu_i)]$. Note that we have that the collection of codimension zero submanifolds $\{(\tilde{M}_1)_i\}$ form an exhaustion of $\tilde{M}_1$, and likewise for $\tilde{M}_2$. Our main result will now follow from:

Claim 6. *For each value of $i \geq 0$, there is a diffeomorphism ψ_i between the stratified manifolds with corners $(\tilde{M}_1)_i$ and $(\tilde{M}_2)_i$. Furthermore, for $i \geq 1$, the diffeomorphism ψ_i can be chosen to coincide with ψ_{i-1} on the submanifold $(\tilde{M}_1)_{i-1} \subset (\tilde{M}_1)_i$.*

Proof. (Claim 6) In order to do this, we first observe that this is clearly true for $i = 0$, since in this case, both $(\tilde{M}_1)_0 = f_1^{-1}[0, \mu_0]$ and $(\tilde{M}_2)_0 = f_2^{-1}[0, r(\mu_0)]$ are diffeomorphic to a standard closed disk $\mathbb{D}^n$ (with empty codimension one strata). Inductively, let us now assume that we have a diffeomorphism $\psi_i : (\tilde{M}_1)_i \to (\tilde{M}_2)_i$, and we would now like to extend it to a diffeomorphism ψ_{i+1} in order to obtain a commutative diagram:

$$\begin{array}{ccc} (\tilde{M}_1)_{i+1} & \xrightarrow{\psi_{i+1}} & (\tilde{M}_2)_{i+1} \\ \uparrow & & \uparrow \\ (\tilde{M}_1)_i & \xrightarrow{\psi_i} & (\tilde{M}_2)_i \end{array}$$

where the vertical maps are the obvious inclusions. But recall that the change of topology of sublevel sets of a Morse function as one transits through a critical value are well understood. In our setting, since there is a unique critical value in the interval (μ_i, μ_{i+1}), with a unique corresponding critical point, the manifold $(\tilde{M}_1)_{i+1} = f_1^{-1}[0, \mu_{i+1}]$ is diffeomorphic to $(\tilde{M}_1)_i = f_1^{-1}[0, \mu_i]$ with a "half 1-handle" attached (see Claim 5). Using the Morse function f_2, we have that a similar statement holds for the corresponding exhaustion of the manifold $\tilde{M}_2$.

Concretely, we have that $(\tilde{M})_{i+1} = f_1^{-1}[0, \mu_{i+1}]$ is diffeomorphic (see Claim 5) to the stratified manifold with corners obtained from $(\tilde{M})_i = f_1^{-1}[0, \mu_i]$ by attaching a copy of the stratified space $[0, 1] \times \mathbb{D}^{n-1}$ via a diffeomorphism $\rho : \{1\} \times \mathbb{D}^{n-1} \to \mathrm{Int}(f_1^{-1}(\mu_i))$, and "smoothing out" the attaching map. Now note that the image of the attaching map $\rho(\mathbb{D}^{n-1}) \subset \mathrm{Int}(f_1^{-1}(\mu_i))$ is a smoothly embedded codimension zero submanifold in the *interior* of the compact manifold with boundary $f^{-1}(\mu_i)$. Similarly, $(\tilde{M}_2)_{i+1}$ is diffeomorphic to the stratified manifold with corners obtained from $(\tilde{M}_2)_i$ by smoothly attaching $[0, 1] \times \mathbb{D}^{n-1}$ via a diffeomorphism $\hat{\rho} : \{1\} \times \mathbb{D}^{n-1} \to \mathrm{Int}(f_2^{-1}(r(\mu_i)))$.

If the attaching map $\hat{\rho}$ coincided with the composite $\psi_i \circ \rho$, then one could immediately extend the diffeomorphism $\psi_i : (\tilde{M}_1)_i \to (\tilde{M}_2)_i$ to a diffeomorphism:

$$\psi_{i+1} : (\tilde{M}_1)_i \bigcup_{\rho} ([0, 1] \times \mathbb{D}^{n-1}) \to (\tilde{M}_2)_i \bigcup_{\psi_i \circ \rho} ([0, 1] \times \mathbb{D}^{n-1})$$

by setting $\psi_{i+1} \cong \psi_i$ on $(\tilde{M}_1)_i$, setting ψ_{i+1} to be the identity on the $[0, 1] \times \mathbb{D}^{n-1}$ term, and using the same smoothing map on both gluings. Of course, in general the maps $\hat{\rho}$ and $\psi_i \circ \rho$ define distinct smooth embeddings of $\mathbb{D}^{n-1}$ into $\mathrm{Int}(f_2^{-1}(r(\mu_i)))$. We now proceed to reduce the general case to the special case where $\hat{\rho} = \psi_i \circ \rho$.

In order to do this, we recall that fundamental work of Palais [16] (see also Cerf [5, Chap. II]) implies that the two embeddings given above are *smoothly isotopic* (rel. boundary), i.e. there exists a diffeomorphism $H : f_2^{-1}(r(\mu_i)) \times [0, 1] \to f_2^{-1}(r(\mu_i)) \times [0, 1]$, with the property that

(1) each $H_t : f_2^{-1}(r(\mu_i)) \times \{t\} \to f_2^{-1}(r(\mu_i)) \times \{t\}$ is a diffeomorphism,
(2) H_0 is the identity, and
(3) $H_1 \circ \psi_i \circ \rho = \hat{\rho}$.

Choosing a real number μ_i' lying in the interval $\mu_i < \mu_i' < \lambda_{i+1}$, we have that there are no critical values of f_1 in a neighborhood of the interval $[\mu_i, \mu_i']$, and similarly that f_2 has no critical values in a neighborhood of the corresponding interval $[r(\mu_i), r(\mu_i')]$. In particular, from Claim 4 we see that there are diffeomorphisms $f_1^{-1}[\mu_i, \mu_i'] \cong f_1^{-1}(\mu_i) \times [0,1]$ and $f_2^{-1}[r(\mu_i), r(\mu_i')] \cong f_2^{-1}(r(\mu_i)) \times [0,1]$. Using this product structure, we can now extend the diffeomorphism $\psi_i : f_1^{-1}(\mu_i) \to f_2^{-1}(r(\mu_i))$ to a diffeomorphism $\psi_1 \times \mathrm{Id} : f_1^{-1}[\mu_i, \mu_i'] \to f_2^{-1}[r(\mu_i), r(\mu_i')]$. Finally, we can compose this map with the smooth isotopy H, resulting in a new diffeomorphism $H \circ (\psi \times \mathrm{Id})$ from $f_1^{-1}[\mu_i, \mu_i']$ to $f_2^{-1}[r(\mu_i), r(\mu_i')]$. Now observe that, since H_0 is the identity, we have that this new map restricted to $f_1^{-1}(\mu_i) \times \{0\} = f_1^{-1}(\mu_i)$ coincides with ψ_i, hence we can glue this map to the previously defined ψ_i. This gives us a diffeomorphism $\psi' : (\tilde{M}_1)_i' \to (\tilde{M}_2)_i'$, where the two spaces are defined by $(\tilde{M}_1)_i' := f_1^{-1}[0, \mu_i']$, and $(\tilde{M}_2)_i' := f_2^{-1}[0, r(\mu_i')]$.

Now since $(\tilde{M}_1)_i' \cong (\tilde{M}_1)_i$, we can think of the space $(\tilde{M}_1)_{i+1}$ as being obtained by attaching $[0,1] \times \mathbb{D}^{n-1}$ to $(\tilde{M}_1)_i'$ rather than to $(\tilde{M}_1)_i$, and likewise with $(\tilde{M}_2)_{i+1}$. Furthermore, by construction we have that the diffeomorphism $\psi' : (\tilde{M}_1)_i' \to (\tilde{M}_2)_i'$ satisfies $\hat{\rho} = \psi' \circ \rho$. But this now reduces the general case to the special case we had previously discussed. We conclude that there exists a map $\psi_{i+1} : (\tilde{M}_1)_{i+1} \to (\tilde{M}_2)_{i+1}$ having the property that ψ_{i+1}, when restricted to $(\tilde{M}_1)_i'$, coincides with the map ψ'. In particular, the further restriction of ψ_{i+1} to $(\tilde{M}_1)_i \subset (\tilde{M}_1)_i'$ coincides with the restriction of ψ' to $(\tilde{M}_1)_i$, and hence is just the map ψ_i. This concludes the proof of Claim 6. □

Finally, we obtain a globally defined map $\Psi : \tilde{M}_1 \to \tilde{M}_2$ in the obvious manner: given $x \in \tilde{M}_1$, the fact that $\{(\tilde{M}_1)_i\}$ form an exhaustion of $\tilde{M}_1$ guarantees that there exists an i such that $x \in (\tilde{M}_1)_i$. We now define the image of x to be the point $\Psi(x) := \psi_i(x) \in (\tilde{M}_2)_i \subset \tilde{M}_2$. The compatibility condition on the collection of maps $\{\psi_i\}$ ensures that this is well defined. Furthermore, since each ψ_i is a diffeomorphism onto its image, and since $\{(\tilde{M}_2)_i\}$ form an exhaustion of $\tilde{M}_2$, we conclude that the map Ψ must likewise be a diffeomorphism. Finally, by construction, it is clear that each of the ψ_i preserves the induced stratification of the sublevel sets, hence the globally defined map Ψ will also preserve the stratification. This concludes the proof of the first part of Theorem 1.2, giving us a characterization of the diffeomorphism type of the universal cover $\tilde{M}$ in terms of the number of components of $\partial\tilde{M}$.

Remark. At the end of the previous section, the author sketched out how one could obtain a somewhat weaker form of this theorem. The argument, relying on our Theorem 1.1, required the stronger hypothesis of strictly negative curvature, as well as requiring that the dimension $n \geq 3$ and $n \neq 5$. The conclusion was the *a priori* weaker statement that, if the boundary was non-empty, then the universal covers had to be *homeomorphic*. We remark that, in principle, we could in fact conclude directly from that argument that the universal covers were diffeomorphic.

Indeed, the work of Kirby–Siebenmann (see [10]) translates the smoothing problem in high dimensions (≥ 5) into a homotopy lifting problem. But the universal covers $\tilde{M}_i$ are contractible, which immediately implies that the obstructions to lifting (and hence to smoothing) vanish.

The main subtlety in this approach is that the work in [10] seems to focus exclusively on the case of manifolds without boundary. While we certainly believe that (analogues of) these results hold for manifolds with boundary (perhaps at the cost of requiring dimension ≥ 6 rather than ≥ 5), we were unable to locate a reference discussing this case. Rather than trying to extend [10] to cover the boundary case, we chose to give the argument in the present section for three reasons: (1) it is probably accessible to a broader audience (having some familiarity with Morse theory), (2) it works even in dimensions ≤ 5, and (3) it gives information in the non-positively curved setting as well.

3.2. *Number of boundary components*

We now have a Riemannian manifold M of non-positive curvature, and would like to identify the number of boundary components of the universal cover $\tilde{M}$. It is clear that if M is closed, the universal cover will have no boundary component, so let us assume that $\partial M \neq \emptyset$. Let $N \subset M$ be a connected component of ∂M; our first step will be to analyze the number of connected components in the full lift of N to $\tilde{M}$. Let $\Gamma = \pi_1(M)$, $\Lambda = \pi_1(N)$, and recall that the map induced by inclusion $\Lambda \to \Gamma$ is an embedding (since N is totally geodesic in M, and M has non-positive curvature). We will identify Λ with its image in Γ. Now note that the number of connected components in the full pre-image of N in $\tilde{M}$ coincides with the index $[\Gamma : \Lambda]$ of the group Λ in the group Γ. In particular, if $[\Gamma : \Lambda] = \infty$, then we immediately obtain that the number of connected components of $\partial^\infty \tilde{M}$ is infinite. To establish our result, we first make:

Assertion 1. If $[\Gamma : \Lambda] < \infty$, then $[\Gamma : \Lambda] \leq 2$.

Proof. (Assertion 1) To see the assertion, let us assume that the full lift of $N \subset M$ in the universal cover $\tilde{M}$ has finitely many connected components $\tilde{N}_1, \ldots, \tilde{N}_k$, with $k > 1$. Without loss of generality, we may assume that the subgroup of Γ that stabilizes $\tilde{N}_1$ is precisely Λ. Letting $g \in \Gamma$ be an element satisfying $g\tilde{N}_1 = \tilde{N}_2$, we have that the stabilizer of $\tilde{N}_2$ is precisely $g\Lambda g^{-1}$. But we have that both Λ and $g\Lambda g^{-1}$ are finite index subgroups of the group Γ, hence the intersection $\Lambda \cap g\Lambda g^{-1}$ has finite index in both Λ and $g\Lambda g^{-1}$. Furthermore, the intersection $\Lambda \cap g\Lambda g^{-1}$ stabilizes both $\tilde{N}_1$ and $\tilde{N}_2$.

Now consider the two boundary components $\tilde{N}_1, \tilde{N}_2$, and observe that there exists at least one geodesic segment $\gamma : [0, D] \to \tilde{M}$ satisfying $\gamma(0) \in \tilde{N}_1$, $\gamma(1) \in \tilde{N}_2$, and realizing the distance between $\tilde{N}_1$ and $\tilde{N}_2$. Indeed, take any curve joining $\tilde{N}_1$ to $\tilde{N}_2$, and consider the projection α to the *compact* manifold M. Now take a sequence of curves, within the homotopy class of α (rel ∂M) whose length tends to the infimum within the homotopy class. Since M is compact, Arzela–Ascoli implies that there is a curve realizing this minimum, and it is immediate that such a curve is a geodesic in M. The lift will give the desired γ.

Next, observe that for all $h \in \Lambda \cap g\Lambda g^{-1}$, we have that $h \cdot \gamma$ is also a geodesic joining $\tilde{N}_1$ to $\tilde{N}_2$ (since Λ stabilizes both these subspaces) having the same length as γ (since we have an isometric action). But $\Lambda \cap g\Lambda g^{-1}$ acts co-compactly on $\tilde{N}_1$, and hence we see that $d(-, \tilde{N}_2) : \tilde{N}_1 \to \mathbb{R}^+$ is a bounded function on $\tilde{N}_1$. Since $\tilde{N}_1$ and $\tilde{N}_2$ are both totally geodesic in $\tilde{M}$, the function $d(-, \tilde{N}_2)$ is convex on $\tilde{N}_1$, and hence must be constant. The flat strip theorem (see [3]) now implies that $\tilde{M}$ is isometric to $\tilde{N}_1 \times [0, D]$, where $D = d(\tilde{N}_1, \tilde{N}_2)$. In particular, we see that $\partial\tilde{M}$ consists of precisely the disjoint union of $\tilde{N}_1$ and $\tilde{N}_2$, forcing $k = 2$, as desired. □

So we are now left with considering the case where $[\Gamma : \Lambda] \leq 2$. We analyze each of the two possibilities separately. Note that the argument given in the proof of Assertion 1 immediately implies:

Assertion 1′. If $[\Gamma : \Lambda] = 2$, then $\tilde{M}$ is isometric to $\tilde{N} \times [0, D]$ for a suitable $D > 0$. In particular, $\tilde{M}$ has two boundary components, each of which is a connected lift of the single boundary component of M.

Hence we are merely left with establishing:

Assertion 2. If $[\Gamma : \Lambda] = 1$, then $\tilde{M}$ is isometric to $\tilde{N} \times [0, D]$ for a suitable $D > 0$, and M itself is isometric to $N \times [0, D]$. In particular, $\tilde{M}$ has two boundary components.

Proof. (Assertion 2) In order to see this, we first note that from the compactness of M, we have the existence of a constant K such that *every* point in $\tilde{M}$ lies at distance $\leq K$ from a point on a lift of N. Furthermore, since $\Gamma = \Lambda$, we have that the lift of N has a single connected component $\tilde{N}$. Combining the two observations above, we see that $\tilde{M}$ lies in the K-neighborhood of $\tilde{N} \subset \tilde{M}$.

Next we recall that since $\tilde{M}$ is non-positively curved, and $\tilde{N} \subset \tilde{M}$ is totally geodesic, there is a projection map $\pi : \tilde{M} \to \tilde{N}$ sending each point $p \in \tilde{M}$ to the unique point $\pi(p) \in \tilde{N}$ which satisfies $d(p, \pi(p)) = d(p, \tilde{N})$. Note that the pre-image of a point $q \in \widetilde{DM}$ under the map π is precisely the geodesic η_q satisfying $\eta_q(0) = q$, $\dot{\eta}_q(0) \perp T_q\tilde{N}$. From the observation in the previous paragraph, we have that for each $q \in \tilde{N}$ the geodesic η_q is actually a geodesic segment of length $\leq K$, joining $\tilde{N}$ to a unique second component $\tilde{N}'$ of $\partial\tilde{M}$. Now focusing on the convexity of the distance function from $\tilde{N}$ to $\tilde{N}'$ as in the previous claim, we see that $\tilde{M}$ splits isometrically as a product $\tilde{M} = \tilde{N} \times [0, D]$, with $D = d(\tilde{N}, \tilde{N}')$. Furthermore, since $\Gamma = \Lambda$ acts isometrically and stabilizes $\tilde{N}$, we immediately obtain that M itself is isometric to $N \times [0, D]$, concluding the proof of Assertion 2. □

Putting all this together, we see that the number of boundary components of $\tilde{M}$ is either:

- 0: corresponding to the case where M is a closed manifold,
- 2: corresponding to the non-generic case where $\tilde{M}$ splits isometrically as a product with an interval, or
- ∞: the generic case corresponding to all other $\tilde{M}$.

In dimension two, these three possibilities are illustrated by taking, for instance: a flat torus, a flat cylinder, and a torus with an open disc removed. By taking products with S^1, we obtain corresponding examples in all dimensions ≥ 2.

4. Topological Rigidity and Applications

A key aspect in the study of non-positively curved Riemannian manifolds is the large number of *rigidity theorems* known to hold for these spaces. Two outstanding such theorems are (1) Mostow rigidity [15], stating that in dimension ≥ 3, homotopy equivalence of irreducible locally symmetric spaces of non-compact type implies isometry of the spaces, and (2) Farrell–Jones topological rigidity [6], stating that in dimension ≥ 5, homotopy equivalence of non-positively curved Riemannian manifolds implies homeomorphism of the spaces.

A natural question is how to extend these theorems to the context of singular spaces satisfying a metric analogue of "non-positive curvature". In some earlier papers ([11, 12]), the author introduced the class of hyperbolic P-manifolds, which one can view as some of the simplest non-manifold CAT(-1) spaces, and established Mostow rigidity within this class of spaces. In the present section, we establish Theorem 1.3, showing topological rigidity for negatively curved P-manifolds. The key point is that our Theorem 1.1 allows the arguments given in [12] to extend verbatim to the present setting. For the convenience of the reader, we first review the terminology we use, then provide a proof of the various corollaries, and finally outline the proof of Theorem 1.1 (referring the interested reader to [12] for more details).

4.1. *Basic definitions*

Let us recall the definition of a P-manifold:

Definition 4.1. A closed n-dimensional *piecewise manifold* (henceforth abbreviated to P-manifold) is a topological space which has a natural stratification into pieces which are manifolds. More precisely, we define a one-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:

- 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 $\rho : \coprod \partial W_i \to X_{n-1}$ 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 of $\coprod \partial W_i$. 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 $X_{n-1}, \ldots, X_1$ 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 4.2. A Riemannian metric on a one-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 non-empty totally geodesic boundary satisfying that the gluing map ρ is an isometry when restricted to each component of ∂W_i. We say that a Riemannian metric on a P-manifold is *negatively curved* if at each step, the metric on each W_i is negatively curved.

Observe that, at the cost of scaling the metric of the P-manifold X by a constant, one can assume that the metric on each W_i has sectional curvature bounded above by -1. Such a metric on the P-manifold will automatically be locally CAT(-1), and hence the fundamental group of a negatively curved P-manifold is a δ-hyperbolic group. In particular, the universal cover $\tilde{X}$ has a well-defined boundary at infinity, denoted $\partial^\infty \tilde{X}$.

Definition 4.3. 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.

We now recall the statement of our Theorem 1.3:

Theorem 4.1. (Topological rigidity of negatively curved P-manifolds) *Let X_1, X_2 be a pair of simple, thick, negatively curved P-manifolds, of dimension ≥ 6. If $\pi_1(X_1)$ is isomorphic to $\pi_1(X_2)$, then X_1 is homeomorphic to X_2.*

We note that, corresponding to the stratification of a negatively curved P-manifold, there is a natural diagram of groups having the property that the direct limit of the diagram is precisely the fundamental group of the P-manifold (by the generalized Seifert–Van Kampen theorem).

Remark. We note that topological rigidity fails (trivially) in dimension $n = 1$. In dimension $n = 2$, topological rigidity was proved in [13]. In dimension $n = 3$, the argument given in the present paper could be extended, provided one had an analogue of Farrell–Jones [6] for three-dimensional manifolds. This analogue is a well-known consequence of Thurston's hyperbolization conjecture. A proof of the hyperbolization conjecture is expected to follow from G. Perelman's work on the Ricci flow method. In dimension $n = 4$, topological rigidity for negatively curved P-manifolds reduces to topological rigidity for negatively curved 4-manifolds with totally geodesic boundary. In dimension $n = 5$, we are additionally lacking a characterization of the boundary at infinity, due to the dimension hypothesis in Cannon's characterization of Sierpinski curves [4].

4.2. *Consequences of topological rigidity*

Assuming for the time being our Theorem 1.3, let us first establish Corollaries 1.1 to 1.3. For the convenience of the reader, we restate each corollary before explaining it's proof.

Corollary 4.1. (Diagram rigidity) *Let $\mathcal{D}_1, \mathcal{D}_2$ be a pair of diagrams of groups, corresponding to a pair of negatively curved, simple, thick P-manifolds of dimension $n \geq 6$. Then $\varinjlim \mathcal{D}_1$ is isomorphic to $\varinjlim \mathcal{D}_2$ if and only if the two diagrams are isomorphic.*

Proof. To obtain Corollary 1.1, we merely note that the generalized Seifert–Van Kampen theorem implies that both $\pi_1(X_i)$ can be expressed as the direct limit of a diagram of groups, with vertex groups given by the fundamental groups of the chambers (and of the components of the branching locus), and edge morphisms induced by the inclusion of the components of the branching locus into the incident chambers. Now an abstract isomorphism between the direct limits corresponds to an isomorphism from $\pi_1(X_1)$ to $\pi_1(X_2)$. From Theorem 1.3, this isomorphism is induced by a homeomorphism from X_1 to X_2, and hence must take chambers to chambers and components of the branching locus to components of the branching locus. This implies the existence of isomorphism between the groups attached to the vertices in the diagram for $\pi_1(X_1)$ to the groups attached to the corresponding vertices in the diagram for $\pi_1(X_2)$. Furthermore, these isomorphisms commute (up to inner automorphisms, due to choice of base points) with the corresponding edge morphisms. But this is precisely the definition of diagram rigidity. This concludes the sketch of Corollary 1.1. □

Corollary 4.2. (Weak Co-Hopf property) *Let X be a simple, thick, negatively curved P-manifold of dimension $n \geq 6$, and assume that at least one of the chambers* *has a nonzero characteristic number. Then* $\Gamma = \pi_1(X)$ *is weakly co-Hopfian, i.e. every injection* $\Gamma \hookrightarrow \Gamma$ *with image of finite index is in fact an isomorphism.*

Proof. Since the space X is a $K(\Gamma, 1)$, any injection $i : \Gamma \hookrightarrow \Gamma$ with image of finite index yields a finite cover $\hat{i} : \bar{X} \to X$ with $\pi_1(\bar{X}) \cong \Gamma$, and $\hat{i}_*(\pi_1(\bar{X})) = i(\Gamma)$. Now Theorem 1.3 implies that $\bar{X}$ is homeomorphic to X, so this yields a covering map $\hat{i} : X \to X$, whose degree coincides with the index of the group $i(\Gamma)$ in Γ. Hence it is sufficient to show that this covering has degree one. But we know that X contains a chamber with a nonzero characteristic number. Since there are finitely many chambers, consider the finitely many chambers $W_1, \ldots, W_k$ for which this characteristic number has the largest possible magnitude $|r| \neq 0$. Then we know that under a covering of degree d, characteristic numbers scale by the degree, so we conclude that the full pre-image $\hat{i}^{-1}(W_i)$ of each W_i has characteristic number of magnitude $d \cdot |r|$. By maximality of $|r|$, we conclude that each connected component of $\hat{i}^{-1}(W_i)$ must also have characteristic number equal to $|r|$, and hence must be one of the chambers $W_1, \ldots, W_k$. In particular, the pre-image $\hat{i}^{-1}(W_i)$ of each W_i in the list $W_i, \ldots, W_k$ consists of d distinct chambers in the list $W_1, \ldots, W_k$. Since the list is finite, this forces $d = 1$, as desired. □

Corollary 4.3. (Nielson realization problem) *Let* X *be a simple, thick, negatively curved* P*-manifold of dimension* $n \geq 6$, *and* $\Gamma = \pi_1(X)$. *Then the canonical map* $\text{Homeo}(X) \to \text{Out}(\Gamma)$ *is surjective.*

Proof. Take any element $\alpha \in \text{Out}(\Gamma)$. Then there exists an element $\bar{\alpha} \in \text{Aut}(\Gamma)$ which projects to α under the canonical map $\text{Aut}(\Gamma) \twoheadrightarrow \text{Out}(\Gamma)$. From Theorem 1.3, we have a self-homeomorphism $\phi \in \text{Homeo}(X)$ with the property that $\phi_* = \alpha$, concluding the proof of Corollary 1.3. □

Remark. Concerning the hypothesis in Corollary 4.2 on the existence of a nonzero characteristic number for one of the chambers, we point out that the famous Hopf Conjecture on the sign of the Euler characteristic asserts that for a closed, negatively curved, even dimensional manifold M^{2n}, we have the inequality $(-1)^n \chi(M^{2n}) > 0$. It is easy to see (using a doubling argument) that the Hopf conjecture, if true, implies that for any compact negatively curved manifold M with *non-empty* totally geodesic boundary, we have $\chi(M) \neq 0$. In particular, the validity of the Hopf conjecture would yield the desired nonzero characteristic number. We also point out that a much stronger result is known, namely Sela [18] has shown that a non-elementary δ-hyperbolic group is co-Hopfian if and only if if is freely indecomposable.

4.3. *Proof of topological rigidity*

Let us now sketch out the proof of Theorem 1.3 from the introduction. We first start with a definition:

Definition 4.4. 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 $f : T \times \mathbb{D}^{n-1} \to X$ such that $p \in f(* \times \mathbb{D}^{n-1}_{\circ})$.

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 is the content of the following:

Proposition 4.1. (Characterization of branching points) *Let X be an n-dimensional, simple, thick, negatively curved P-manifold, and $p \in \partial^\infty \tilde{X}$. Then p is $(n-1)$-branching if and only if there exists a geodesic ray γ, entirely contained in the lift of the branching locus, and satisfying $\gamma(\infty) = p$.*

Proof. First observe that if $p \in \partial^\infty \tilde{X}$ coincides with $\gamma(\infty)$, for some γ entirely contained in a connected component B of the lift of the branching locus, then from the thickness hypothesis, there exist ≥ 3 lifts of chambers that contain γ in their common intersection B. Focusing on three such lifts of chambers, call them Y_1, Y_1', Y_1'', we can successively extend each of these in the following manner: form subspaces $Y_{i+1}, Y_{i+1}', Y_{i+1}''$ from the subspaces Y_i, Y_i', Y_i'' by choosing, for each boundary component of Y_i, Y_i', Y_i'' *distinct* from B, an incident lift of a chamber (note that each boundary component is a connected component of the lift of the branching locus). Finally, set $Y_\infty := \cup_i Y_i$, and similarly for Y_∞', Y_∞''. Now observe that, by construction, the three subsets $Y_\infty, Y_\infty', Y_\infty''$ have the following properties:

- they are totally geodesic subsets of $\tilde{X}$,
- their pairwise intersection is precisely B, their (common, totally geodesic) boundary component,
- doubling them across their boundary B results in a simply connected, negatively curved, complete Riemannian manifold.

The first property ensures that the boundary at infinity of the space $Y_\infty \cup Y_\infty' \cup Y_\infty''$ embeds in $\partial^\infty \tilde{X}$. The third property ensures that $\partial^\infty Y_\infty \cong \partial^\infty Y_\infty' \cong \partial^\infty Y_\infty'' \cong \mathbb{D}^{n-1}$. The second property ensures that $S^{n-2} \cong \partial^\infty B \subset \partial^\infty \tilde{X}$ coincides with the boundary of the three embedded $\mathbb{D}^{n-1}$. Since $p \in \partial^\infty B$, this immediately implies that p is $(n-1)$-branching, yielding one of the two desired implications.

Conversely, assume that $p \in \tilde{X}$ is *not* of the form $\gamma(\infty)$, where γ is contained entirely in a connected component of the lift of the branching locus. Consider a geodesic ray γ satisfying $\gamma(\infty) = p$, and note that there are two possibilities:

- there exists a connected lift W of a chamber with the property that γ eventually lies in the *interior* of W, and is *not* asymptotic to any boundary component of W, or
- γ intersects infinitely many connected lifts of chambers.

In both these cases, we would like to argue that p *cannot* be $(n-1)$-branching.

Let us consider the first of these two cases, and assume that there exists an embedding $f : T \times \mathbb{D}^{n-2} \to \partial^\infty \tilde{X}$ satisfying $p \in f(\{*\} \times \mathbb{D}^{n-2}_\circ)$. Picking a point x in the interior of W, one can consider the composition $\pi_x \circ f : T \times \mathbb{D}^{n-2} \to lk_x \cong S^{n-1}$, where lk_x denotes a small enough ϵ-sphere centered at the point x, and the map π_x is induced by geodesic retraction. Note that the map π_x is *not* injective: the points in lk_x where π_x is injective coincides with $\pi_x(\partial^\infty W)$ (i.e. for every $q \in \partial^\infty W$, we have $\pi_x^{-1}(\pi_x q) = \{q\}$, and the latter are the only points in $\partial^\infty \tilde{X}$ with this property). Note that, from Theorem 1.1, along with part (2) of Cannon's theorem (see Theorem 2.1), this subset of injective points $I \subset lk_x$ is an $(n-2)$-dimensional Sierpinski curve. Furthermore, the hypothesis on the point p ensures that $\pi_x p$ does *not* lie on one of the boundary spheres of the $(n-2)$-dimensional Sierpinski curve I. But now in [12, Sec. 3.1] the following result was established:

Theorem. *Let $F : T \times \mathbb{D}^{n-2} \to S^{n-1}$ be a continuous map, and assume that the sphere S^{n-1} contains an $(n-2)$-dimensional Sierpinski curve I. Let $\{U_i\}$ be the collection of embedded open $(n-1)$-cells whose complement yield I, and let $Inj(F) \subset S^{n-1}$ denote the subset of points in the target where the map F is injective. Then $F(\{*\} \times \mathbb{D}^{n-2}_\circ) \cap [I - \cup_i(\partial U_i)] \neq \emptyset$, implies that $[\cup_i(\partial U_i)] - Inj(F) \neq \emptyset$. In other words, this forces the existence of a point in some ∂U_i which has* at least two *pre-images under F.*

Actually, in [12] this theorem was proved using *purely topological arguments* under some further hypotheses on the open cells U_i. But parts (1) and (3) of Cannon's Theorem allows the exact same proof to apply in the more general setting, just by composing with a homeomorphism taking the arbitrary Sierpinski curve to the one used in the proof in [12].

To conclude, we apply the theorem above to the composite map $F := \pi_x \circ f : T \times \mathbb{D}^{n-2} \to lk_x$. The point $f^{-1}(p) \in \{*\} \times \mathbb{D}^{n-2}_\circ$ has image lying in $I - \cup_i(\partial U_i)$, which tells us that $F(\{*\} \times \mathbb{D}^{n-2}_\circ) \cap [I - \cup_i(\partial U_i)] \neq \emptyset$. The theorem implies that there exists a point q in some $\partial U_i \subset I$ which has *at least two* pre-images under the composite map $F = \pi_x \circ f$. Since the map π_x is actually *injective* on the set I, this implies that the map f had to have two pre-images at the point $\pi_x^{-1}(q) \in \partial^\infty \tilde{X}$, contradicting the fact that f was an embedding. This resolves the first of the two possible cases.

For the second of the two cases (where the geodesic ray γ passes through infinitely many lifts of chambers), a simple separation argument (see Secs. 3.2, 3.3 in [12]) shows that if there exists a branching point of the second type, there must also exist a branching point of the first type. But we saw above that there cannot exist any branching points of the first type. This concludes the proof of Proposition 3.1. □

Now given the characterization of branching points, let us see how to show Theorem 1.3. So assume that we are given a pair X_1, X_2 of simple, thick, negatively curved P-manifolds of dimension $n \geq 6$, and that we are told that $\pi_1(X_1) \cong \pi_1(X_2)$. This immediately implies that $\tilde{X}_1$ is quasi-isometric to $\tilde{X}_2$, and hence that $\partial^\infty \tilde{X}_1$ is *homeomorphic* to $\partial^\infty \tilde{X}_2$. Let $\mathcal{B}_i$ denote the union, in each respective $\partial^\infty \tilde{X}_i$, of the boundaries at infinity of the individual connected components of the lift of the branching locus. Note that each $\mathcal{B}_i$ is a union of countably many, pairwise disjoint, embedded S^{n-2} inside $\partial^\infty \tilde{X}_i$ (each S^{n-2} arising as the boundary at infinity of a single connected component of the lift of the branching locus). Now the characterization of branching points in Proposition 4.1 implies that, under the homeomorphism between $\partial^\infty \tilde{X}_1$ and $\partial^\infty \tilde{X}_2$, we have that $\mathcal{B}_1$ must map homeomorphically to $\mathcal{B}_2$.

In particular, connected components of $\mathcal{B}_1$ must map homeomorphically to connected components of $\mathcal{B}_2$. A result of Sierpinski [19] implies that the connected components in each case are precisely the individual S^{n-2} in the countable union. This yields a bijection between connected components of the lift of the branching locus in the respective $\tilde{X}_i$. Furthermore, the homeomorphism must restrict to a homeomorphism between the *complements* of the $\mathcal{B}_i$ in the respective $\partial^\infty \tilde{X}_i$. The connected components of this complement are either:

- isolated points, corresponding to $\gamma(\infty)$, where γ is a geodesic ray passing through infinitely many connected lifts of chambers, and
- components with ≥ 2 points, which are in bijective correspondence with connected lifts of chambers in the respective $\tilde{X}_i$ (see [12, Sec. 3.2]).

This yields a bijective correspondence between lifts of chambers in $\tilde{X}_1$ and lifts of chambers in $\tilde{X}_2$. Furthermore, the *closure* of the components containing ≥ 2 points correspond canonically with $\partial^\infty W_i$, where W_i is the bijectively associated connected lift of a chamber.

Now recall that the homeomorphisms between $\partial^\infty \tilde{X}_1$ and $\partial^\infty \tilde{X}_2$ has the additional property that it is *equivariant* with respect to the respective $\pi_1(X_i)$ actions on the $\partial^\infty \tilde{X}_i$. We also have the following Lemma relating the action on $\partial^\infty \tilde{X}$ with the action on $\tilde{X}$ (the argument is identical to that given in [11, p. 212]):

Lemma 4.1. *Let B_i be a connected component of the lift of the branching locus in $\tilde{X}$, and let W_i be a connected lift of a chamber in $\tilde{X}$. Then we have*:

- $Stab_{\pi_1(X)}(B_i) = Stab_{\pi_1(X)}(\partial^\infty B_i)$, *and*
- $Stab_{\pi_1(X)}(W_i) = Stab_{\pi_1(X)}(\partial^\infty W_i)$,

where the action on the left-hand side is the obvious action of $\pi_1(X)$ on $\tilde{X}$ by deck transformations, and the action on the right-hand side is the induced action of $\pi_1(X)$ on $\partial^\infty \tilde{X}$.

Observe that equivariance of the homeomorphism implies that the bijective correspondence between connected lifts of chambers descends to a bijective correspondence between the chambers in X_1 and the chambers in X_2 (since two connected lifts of chambers cover the same chamber in X_i if and only if the two lifts have stabilizers which are conjugate in $\pi_1(X_i)$). Similarly, the bijective correspondence between connected components of the lifts of the branching loci descends to a bijective correspondence between the connected components of the branching loci in X_1 with those in X_2. Furthermore, by equivariance of the homeomorphism, we have that chambers (or connected components of the branching loci) that are bijectively identified have isomorphic fundamental groups. Separation arguments identical to the ones in [11, Lemmas 2.1–2.4] ensures that the bijective correspondence also preserves the incidence relation between chambers and components of the codimension one strata (and that the isomorphisms between the various fundamental groups respect the incidence structure).

To conclude, we apply the celebrated Farrell–Jones topological rigidity theorem for non-positively curved manifolds [6]. This implies that, corresponding to the bijections between chambers (and components of the branching loci), one has *homeomorphisms* between the corresponding chambers that induce the isomorphisms on the level of the fundamental groups. Note that, *a priori*, the various homeomorphisms between chambers might not be compatible with the gluing maps. But by construction, the attaching maps all induce the same maps on the fundamental group $\pi_1(B_i)$ of each individual component B_i of the branching locus. By Farrell–Jones, this implies that the restriction to B_i of the maps induced by the various homeomorphisms of incident chambers are *all pairwise pseudoisotopic*. Hence at the cost of deforming the homeomorphism in a collared neighborhood of the boundary of each chamber, we may assume that the homeomorphisms respect the gluing maps. But attaching together these individual homeomorphisms on chambers now induces a *globally defined* homeomorphism from X_1 to X_2. This concludes the sketch of Theorem 1.3.

Acknowledgments

The author would like to thank M. Hindawi and T. Januszkiewicz for raising the question of determining the topology of the universal cover of a compact non-positively curved Riemannian manifold with non-empty totally geodesic boundary.

The portion of Theorem 1.2 addressing the possible number of boundary components in $\partial^\infty \tilde{M}$ is in answer to a question of F. T. Farrell raised in a seminar given by the author at Binghamton University.

The work done here was partly supported by the National Science Foundation under grants DMS-0606002 and DMS-0906483, and by an Alfred P. Sloan research fellowship.

References

1. E. Baiada and M. Morse, Homotopy and homology related to the Schoenflies problem, *Ann. Math.* **58** (1953) 142–165.
2. M. Bourdon, Structure conforme au bord et flot géodésique d'un $CAT(-1)$-espace, *Enseign. Math.* **41** (1995) 63–102.
3. M. R. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature* (Springer-Verlag, 1999).
4. J. W. Cannon, A positional characterization of the $(n-1)$-dimensional Sierpiński curve in S^n $(n \neq 4)$, *Fund. Math.* **79** (1973) 107–112.
5. J. Cerf, Topologie de certains espaces de plongements, *Bull. Soc. Math. France* **89** (1961) 227–380.
6. F. T. Farrell and L. E. Jones, A topological analogue of Mostow's rigidity theorem, *J. Amer. Math. Soc.* **2** (1989) 257–370.
7. M. Goresky and R. MacPherson, *Stratified Morse Theory* (Springer-Verlag, 1988).
8. H. Hamm, Zur homotopietyp Steinscher räume, *J. Reine Angew. Math.* **338** (1983) 121–135.
9. H. Hamm and D. T. Lê, Un théorème de Zariski du type de Lefschetz, *Ann. Sci. Ecole Norm. Sup.* **6** (1973) 317–355.
10. R. C. Kirby and L. C. Siebenmann, *Foundational Essays on Topological Manifolds, Smoothings, and Triangulations* (Princeton Univ. Press, 1977).
11. J.-F. Lafont, Rigidity result for certain 3-dimensional singular spaces and their fundamental groups, *Geom. Dedicata* **109** (2004) 197–219.
12. J.-F. Lafont, Strong Jordan separation and applications to rigidity, *J. London Math. Soc.* **73** (2006) 681–700.
13. J.-F. Lafont, Rigidity of geometric amalgamations of free groups, *J. Pure App. Alg.* **209** (2007) 771–780.
14. J. Milnor, *Morse Theory* (Princeton Univ. Press, 1963).
15. G. D. Mostow, *Strong Rigidity of Locally Symmetric Spaces* (Princeton Univ. Press, Princeton, 1973).
16. R. Palais, Extending diffeomorphisms, *Proc. Amer. Math. Soc.* **11** (1960) 274–277.
17. K. Ruane, CAT(0) boundaries of truncated hyperbolic space, *Topol. Proc.* **29** (2005) 317–331.
18. Z. Sela, Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups. II, *Geom. Funct. Anal.* **7** (1997) 561–593.
19. W. Sierpinski, Un théorème sur les continus, *Tohoku Math. J.* **13** (1918) 300–303.
20. D. Siersma, Singularities of functions on boundaries, corners, etc., *Quart. J. Math. Oxford Ser.* **32** (1981) 119–127.
21. G. T. Whyburn, Topological characterization of the Sierpiński curve, *Fund. Math.* **45** (1958) 320–324.