

Marked length rigidity for one-dimensional spaces

David Constantine

*Department of Mathematics and Computer Science
Wesleyan University, 265 Church St.
Middletown, CT 06459, USA
dconstantine@wesleyan.edu*

Jean-François Lafont*

*Department of Mathematics, Ohio State University
231 West 18th Avenue, Columbus, OH 43210, USA
jlafont@math.ohio-state.edu*

Received 29 May 2013

Accepted 18 September 2017

Published 21 February 2018

In a compact geodesic metric space of topological dimension one, the minimal length of a loop in a free homotopy class is well-defined, and provides a function $l : \pi_1(X) \rightarrow \mathbb{R}^+ \cup \{\infty\}$ (the value ∞ being assigned to loops which are not freely homotopic to any rectifiable loops). This function is the *marked length spectrum*. We introduce a subset $\text{Conv}(X)$, which is the union of all non-constant minimal loops of finite length. We show that if X is a compact, non-contractible, geodesic space of topological dimension one, then X deformation retracts to $\text{Conv}(X)$. Moreover, $\text{Conv}(X)$ can be characterized as the minimal subset of X to which X deformation retracts. Let X_1, X_2 be a pair of compact, non-contractible, geodesic metric spaces of topological dimension one, and set $Y_i = \text{Conv}(X_i)$. We prove that any isomorphism $\phi : \pi_1(X_1) \rightarrow \pi_1(X_2)$ satisfying $l_2 \circ \phi = l_1$, forces the existence of an isometry $\Phi : Y_1 \rightarrow Y_2$ which induces the map ϕ on the level of fundamental groups. Thus, for compact, non-contractible, geodesic spaces of topological dimension one, the marked length spectrum completely determines the subset $\text{Conv}(X)$ up to isometry.

Keywords: Rigidity; marked length spectrum; one-dimensional space.

AMS Subject Classification: Primary 51F99; Secondary 52C25, 53C24

1. Introduction

This paper is motivated by a long-standing conjecture concerning negatively curved manifolds: that the length of closed geodesics on a closed negatively curved Riemannian manifold determines the space up to isometry. More precisely, in a closed negatively curved manifold (M^n, g) , there are unique geodesics in free homotopy

*Corresponding author

classes of loops. Assigning to each element in $\pi_1(M^n)$ the length of the corresponding minimal geodesic yields the function $l : \pi_1(M^n) \rightarrow \mathbb{R}^+$, which is called the marked length spectrum. The marked length spectrum conjecture states that, if we have a pair of negatively curved Riemannian metrics on M^n which yield the same length function l , then they are in fact isometric. In full generality, the conjecture is only known to hold for closed surfaces, which was independently established by Croke [10] and Otal [22] (see also Hersch and Paulin [19] for some extensions to singular metrics on surfaces). In the special case where one of the Riemannian metrics is locally symmetric, this result is due to Hamenstädt [18] (see also Dal’bo and Kim [14] for analogous results in the higher rank case).

In this paper, we consider compact geodesic spaces of topological dimension one. The starting observation is that these spaces share a lot of the properties of closed negatively curved manifolds. In particular, they are aspherical (see Curtis and Fort [12]), and they have unique minimal length representatives in each free homotopy class of loops (by Curtis and Fort [13], also shown by Cannon and Conner [5]). As such, it is reasonable to ask whether the marked length spectrum conjecture holds in the setting of compact geodesic spaces of topological dimension one. We define a subset $Conv(X)$ of any compact geodesic space X of topological dimension one. When X is non-contractible, we show that X deformation retracts to $Conv(X)$ (and the latter is the minimal subset with this property). We establish:

Main Theorem. *Let X_1, X_2 be a pair of compact, non-contractible, geodesic spaces of topological dimension one, and set $Y_i = Conv(X_i)$. Assume the two spaces have the same marked length spectrum, that is to say, there exists an isomorphism $\phi : \pi_1(X_1) \rightarrow \pi_1(X_2)$ such that the following diagram commutes:*

$$\begin{array}{ccc}
 \pi_1(Y_1) \cong \pi_1(X_1) & \xrightarrow{\phi} & \pi_1(X_2) \cong \pi_1(Y_2) \\
 & \searrow l_1 & \swarrow l_2 \\
 & \mathbb{R} &
 \end{array}$$

Then Y_1 is isometric to Y_2 , and the isometry induces (up to change of basepoints) the isomorphism ϕ .

Let us provide an outline of the proof, with reference to the next section for appropriate definitions. The idea behind the argument is to look at a certain subset of $Conv(X_1)$ consisting of *branch points*. For a pair of branch points, we consider a minimal geodesic joining them, and construct a pair of geodesic loops with the property that they intersect precisely in the given minimal geodesic. Now using the isomorphism of fundamental groups, we obtain a corresponding pair of geodesic loops in $Conv(X_2)$. Using the fact that the lengths are preserved, we show that the corresponding pair in $Conv(X_2)$ likewise intersects in a geodesic segment, and that furthermore, the length of the intersection is exactly equal to the length of the original geodesic segment. We then proceed to show that this correspondence is well-defined (i.e. does not depend on the pair of geodesic loops one constructs),

and preserves concatenations of geodesic segments. This is used to construct an isometry between the sets of branch points. Using completeness, we extend this to an isometry between the closures of the sets of branch points. Finally, we consider points in $Conv(X_1)$ which are *not* in the closure of the set of branch points. It is easy to see that each of these points lies on a unique maximal geodesic segment, with the property that the only branch points occur at the endpoints of the segment. The correspondence between geodesic segments can be used to see that there is a unique, well-defined, corresponding segment in $Conv(X_2)$ of precisely the same length, allowing us to extend our isometry to all of $Conv(X_1)$.

We conclude this introduction with a few remarks. If the spaces X_i are tame (i.e. are semi-locally simply connected), then our *Main Theorem* can also be deduced from some work of Culler and Morgan [11] (see also Alperin and Bass [1]). But of course, there are numerous examples of geodesic length spaces of topological dimension one which are not semi-locally simply connected (Hawaiian Earrings, Menger curves, Sierpinski curves, etc.), for which our result does not *a priori* follow from theirs. Another nice class of examples are Laakso spaces with Hausdorff dimension between one and two [21]. These spaces have nice analytic properties, and work regarding the spectra of the Laplacian has been carried out on them [23]. In view of the connections between such spectra and the length spectrum in other contexts, this seems like a most interesting family of examples.

In the course of our proof, we develop some structure theory for compact geodesic spaces of topological dimension one (see Sec. 3). Similar results have recently been obtained in the broader setting of Peano continua by Conner and Meilstrup [8, 9] (see Remarks 3.13 and 3.23). There has also been some work on rigidity of one-dimensional Peano continua, though with an emphasis on topological phenomena, see Eda [17] (building on the previous [15, 16]). Some of these results were also extended to planar continua by Conner and Kent (see for instance [6, 7]).

2. Preliminaries

In this section, we will define the various terms used in this paper, as well as quote certain results we will use in our proofs. We start by reminding the reader of a few basic notions on length spaces (and refer to Burago, Burago and Ivanov [4] for more details on the theory).

Definition 2.1. A *path* in a metric space (X, d) , is a continuous map $f : [a, b] \rightarrow X$ from a closed interval into X . A *loop* in X is a path f satisfying $f(a) = f(b)$. Observe that we can always view a loop as a based continuous map from $(S^1, *)$ to $(X, f(a))$.

Definition 2.2. Let (X, d) be a metric space. The *induced length structure* is a function on the set of paths, denoted by l_d , and defined as follows:

$$l_d(\gamma) := \sup \sum_{i=1}^n d(\gamma(x_{i-1}), \gamma(x_i)),$$

where $\gamma : [a, b] \rightarrow X$ is a path, and the supremum ranges over all finite collections of points $a = x_0 < x_1 < \dots < x_n = b$. A path γ is *rectifiable* provided $l_d(\gamma) < \infty$. We will often suppress the subscript and write $l(\gamma)$ if the original metric d is clear from the context.

Observe that rectifiability is preserved under finite concatenation of paths, and under restriction to subpaths. As a convenience, we will always parametrize rectifiable curves by their arclength.

Definition 2.3. Let (X, d) be a metric space, and l_d the induced length structure. We define the *intrinsic pseudo-metric* \hat{d} induced by d as follows. Let $p, q \in X$ be an arbitrary pair of points, and define the \hat{d} distance between them to be $\hat{d}(p, q) := \inf l_d(\gamma)$, where the infimum ranges over all paths $\gamma : [a, b] \rightarrow X$ satisfying $\gamma(a) = p, \gamma(b) = q$.

Note that the function \hat{d} actually maps $X \times X$ to $[0, \infty]$, where two points $p, q \in X$ have $\hat{d}(p, q) = \infty$ if and only if there are no rectifiable paths joining p to q .

Definition 2.4. Let (X, d) be a metric space, \hat{d} the corresponding intrinsic pseudo-metric. We call (X, d) a *length space* if $d = \hat{d}$ (in particular, \hat{d} has image in $(0, \infty)$).

Definition 2.5. Let (X, d) be a length space. We say that (X, d) is a *geodesic space* if, for every pair of points $p, q \in X$, there is a path $\gamma_{p,q}$ joining p to q , and having length precisely $d(p, q)$. Such a path is called a *distance minimizer*. Note that such a curve is clearly locally distance minimizing.

The usual terminology for such a path is geodesic. However, we will use that term for paths which minimize length in their homotopy class (see Definition 2.9 below), even though such paths need not be locally length-minimizing. The example following Definition 2.9 illustrates this distinction. In the literature, geodesic spaces are also sometimes referred to as *complete length spaces*.

In topology, one of the most important concepts is that of dimension. While there are many different notions of dimension, the one which will be of interest to us is that of *Lebesgue covering dimension*. We remind the reader of the definition.

Definition 2.6. Let X be a topological space. We say that X has topological (Lebesgue covering) dimension $\leq n$ if, for any open covering $\{U_i\}$ of X , there is a refinement $\{V_i\}$ with the property that every $x \in X$ lies in at most $n + 1$ of the V_i . We say that X is *n-dimensional* if X has dimension $\leq n$, but does not have dimension $\leq n - 1$.

We will denote the topological dimension of a space X by $\dim(X)$. Observe that any path connected topological space with at least two points has $\dim(X) \geq 1$. The spaces we will be interested in are those satisfying $\dim(X) = 1$. Examples of such

spaces are plentiful. In particular, we have the following criterion (see Chap. VII in Hurewicz and Wallman [20]):

Theorem 2.7. *Let (X, d) be a metric space, with Hausdorff dimension $\dim_H(X)$. Consider the induced topology on X . One has the inequality $\dim(X) \leq \dim_H(X)$.*

Hence any path-connected metric space with Hausdorff dimension less than two (and which is not a single point) has covering dimension one. We will henceforth focus exclusively on geodesic spaces of topological dimension one. Note that, in general, a one-dimensional geodesic space might not be negatively curved, and indeed, might not even be locally contractible.^a However, such spaces often exhibit properties which are quite similar to those of negatively curved spaces.

Definition 2.8. Let $\gamma : (S^1, *) \rightarrow (X, p)$ be a loop. We say that the loop is *reducible* provided that there is an open interval $I = (x, y) \subset S^1 - \{*\}$ such that $\gamma(x) = \gamma(y)$ and the loop $\gamma|_{[x,y]}$ is null-homotopic relative to its endpoints. We say that γ is *cyclically reducible* if the interval I is allowed to include the basepoint $*$. A loop which is not reducible (resp. cyclically reducible) will be said to be *reduced* (resp. *cyclically reduced*). We define a constant loop to be cyclically reduced.

Similarly, let $\gamma : [a, b] \rightarrow X$ be a path. We say that the path is *reducible* provided that there is an open interval $I = (x, y) \subseteq (a, b)$ such that $\gamma(x) = \gamma(y)$ and the loop $\gamma|_{[x,y]}$ is null-homotopic relative to its endpoints. We define a constant path to be reduced.

Definition 2.9. Let (X, d) be a 1-dimensional geodesic space. We say that a path is *geodesic* provided that it is rectifiable and minimizes the length in its homotopy class (rel. endpoints). We say that a loop γ is a *geodesic loop* provided γ is rectifiable and has minimal length in its free homotopy class. As per our convention, all geodesics will be parametrized by arclength.

As mentioned earlier, the reader who is more familiar with the Riemannian setting should beware that, for our highly singular spaces, geodesics might not be locally length-minimizing (or even locally injective), as the following example demonstrates.

Example. Let $X \subset \mathbb{R}^2$ be the union of the line segment $[-1, 1]$ with circles of radius $\frac{1}{2^n}$ for $n \in \mathbb{N}$ tangent to this segment at 0. Consider a path which follows the segment from -1 to 0 , then traverses the circles in order of descending radius, then from times $1 + 2\pi$ to $2 + 2\pi$ follows the segment from 0 to 1 . This is a geodesic connecting -1 to 1 , but around time $1 + 2\pi$ it is not length-minimizing.

^aRecall that negatively curved metric spaces are defined by the property that their universal cover satisfies the $CAT(-\delta^2)$ condition (a metric version of negative curvature, see Bridson and Haefliger [3]). Naturally, the problem in our setting is that spaces which are very singular might not have a universal cover.

We now state Cannon and Conner’s theorem (Theorem 3.9 in [5]), which establishes the uniqueness of geodesic loops within homotopy classes:

Theorem 2.10. *Let (X, d) be a compact, path-connected, 1-dimensional metric space. Then*

- every loop is homotopic to a reduced loop, which is unique up to reparametrization,
- every loop is freely homotopic to a cyclically reduced loop, which is unique up to cyclic reparametrization.

The homotopies are taken within the image set of the loops, so that the reduced loop (or cyclically reduced loop) always lies inside the image of the original curve.

Using this result, it is easy to show that in every free homotopy class of loops that contains a rectifiable curve, there is a unique minimal length representative. Note that any length space is automatically a path-connected metric space.

Corollary 2.11. *Let (X, d) be a compact 1-dimensional length space, γ a loop. Then one of the following two possibilities holds:*

- there is no rectifiable loop freely homotopic to γ ,
- there is a unique (up to reparametrization) minimal length rectifiable loop freely homotopic to γ .

Proof. Let us assume that there is a rectifiable loop freely homotopic to γ . Let $\hat{\gamma}$ denote one such loop. Applying Theorem 2.10, $\hat{\gamma}$ is freely homotopic to a cyclically reduced loop $\bar{\gamma}$ whose image is contained in the image of $\hat{\gamma}$. From this, it is easy to see that $l(\bar{\gamma}) \leq l(\hat{\gamma})$.

We claim that $\bar{\gamma}$ is of minimal length in its free homotopy class. Indeed, suppose that β is freely homotopic to γ with $l(\beta) < l(\bar{\gamma})$. Applying Theorem 2.10 to β we get $\bar{\beta}$, the unique (up to cyclic reparametrization) cyclically reduced loop in its free homotopy class. The uniqueness implies that $\bar{\beta}$ and $\bar{\gamma}$ agree, up to cyclic reparametrization, and therefore have the same length. However, $l(\bar{\beta}) \leq l(\beta) < l(\bar{\gamma})$, a contradiction. □

In view of Corollary 2.11, when looking for geodesic loops, it is sufficient to restrict throughout to cyclically reduced loops. We can now define:

Definition 2.12. Let (X, d) be a geodesic space. Assume that, in each free homotopy class of curves on X , there is at most one minimal length representative. The *marked length spectrum* is defined to be the function $l_d : \pi_1(X) \rightarrow \mathbb{R}^+ \cup \{\infty\}$ which assigns to each element of $\pi_1(X)$ the length of the corresponding minimal length loop (and assigns ∞ to the free homotopy classes that contain no rectifiable representatives).

For geodesic spaces of topological dimension one, Corollary 2.11 implies that the marked length spectrum is defined.

It is straightforward to prove an analogue of the Cannon and Conner result for paths in a 1-dimensional metric space.

Proposition 2.13. *Let (X, d) be a compact, path-connected, 1-dimensional metric space. Then every path is homotopic (relative to the endpoints of the path) to a reduced path, which is unique up to reparametrization. Moreover, the reduced path has image lying within the image of the original path.*

Proof. Let X be our space, \mathbf{p} our path, and p, q the two endpoints. Consider the space X' , defined by $X' = (X \cup [0, 1]) / \sim$ where $p \sim 0$ and $q \sim 1$ are the only nontrivial equivalences for \sim . $X \cup [0, 1]$ has topological dimension 1 and quotients can only decrease dimension. Since X' is still path-connected, $\dim(X') = 1$. It is compact and metric with respect to the induced length metric.

Consider the loop γ in X' obtained by concatenating \mathbf{p} with the added segment $[0, 1]$, and parametrized by arclength. By Theorem 2.10, γ is freely homotopic to a unique cyclically reduced loop $\bar{\gamma}$ with image in the image of γ . Let t_1 be the time at which $\gamma(t)$ passes from $[0, 1]$ to X through p and t_2 the time at which $\gamma(t)$ passes from X back to $[0, 1]$ through q . The path γ cannot reduce on any small interval containing t_1 or t_2 since $\gamma(t_i - \epsilon)$ and $\gamma(t_i + \epsilon)$ are distinct, the one lying in X and the other in $(0, 1)$. Therefore, p and q belong to $\bar{\gamma}$. Restricting $\bar{\gamma}$ to the reduction of $\gamma|_{[t_1, t_2]}$ gives the unique (up to parametrization) reduced path homotopic to \mathbf{p} relative to its endpoints. Its image lies in the image of \mathbf{p} because the image of $\bar{\gamma}$ lies in the image of γ . □

Corollary 2.14. *Let (X, d) be a compact 1-dimensional length space, \mathbf{p} a rectifiable path joining points p and q . Then the unique reduced path joining p to q homotopic (relative to the endpoints) to \mathbf{p} has minimal length amongst all paths with this property.*

Proof. This argument is identical to the one for loops: assume that \mathbf{p} is rectifiable, but not reduced, and is a mapping from $[a, b]$ into X . Then there is a subinterval $I = [a', b'] \subset [a, b]$ such that $\mathbf{p}(a') = \mathbf{p}(b')$ and the loop \mathbf{p}_1 obtained by restricting \mathbf{p} to I is contractible. Denoting by \mathbf{p}_2 the concatenation of \mathbf{p} restricted to $[a, a']$ and $[b', b]$, we see that \mathbf{p}_2 is homotopic to \mathbf{p} via a homotopy preserving the endpoints. Furthermore, $l(\mathbf{p}_2) < l(\mathbf{p})$, which yields our claim. □

Corollary 2.14 immediately yields:

Corollary 2.15. *Let (X, d) be a compact 1-dimensional geodesic space, γ a distance minimizer joining p to q . Then γ is a geodesic.*

For the remainder of this section we make the standing assumption that (X, d) is a compact 1-dimensional geodesic space.

The rest of this section will be devoted to analyzing the behavior of paths under concatenation. We fix the following notation: given a pair of paths \mathbf{p} and \mathbf{q} , with

the terminal endpoint of \mathbf{p} coinciding with the initial endpoint of \mathbf{q} , we will denote by $\mathbf{q} * \mathbf{p}$ the concatenation of the two paths (traversing \mathbf{p} first, followed by \mathbf{q}), and by \mathbf{p}^{-1} the path obtained by reversing \mathbf{p} . We will denote by γ^n the n -fold iterated concatenation of a loop γ . Such paths are always understood to be parametrized by arclength. We start by proving several lemmas concerning concatenations of various types of paths.

Lemma 2.16. *Let $\mathbf{p}_1, \mathbf{p}_2$ be a pair of reduced paths, parametrized by arclength, in X . Let $t_i = l(\mathbf{p}_i)$ be the length of the respective paths, and assume that $\mathbf{p}_1(t_1) = \mathbf{p}_2(0)$ (i.e. that they have a common endpoint). Then any reduced path \mathbf{q} homotopic to $\mathbf{p} := \mathbf{p}_2 * \mathbf{p}_1$ is of the form $\mathbf{q}_2 * \mathbf{q}_1$ where \mathbf{q}_i is a subpath of \mathbf{p}_i . Furthermore, we have a decomposition $\mathbf{p}_1 = \mathbf{r}_1 * \mathbf{q}_1$ and $\mathbf{p}_2 = \mathbf{q}_2 * \mathbf{r}_2$, satisfying $\mathbf{r}_1 = \mathbf{r}_2^{-1}$.*

Proof. We start by observing that the claim is trivial if the concatenation \mathbf{p} is reduced (just take $\mathbf{q}_i = \mathbf{p}_i$). So let us assume that the concatenation \mathbf{p} is *not* a reduced path, and view it as a map from $D := [0, t_1 + t_2]$ into X . Since this path is reducible, there exist closed intervals $U_j \subset D$ with the property that \mathbf{p} restricted to each U_j is a closed path which is null homotopic relative to the endpoints of $\mathbf{p}|_{U_j}$. Since each of the paths $\mathbf{p}_1, \mathbf{p}_2$ is reduced, this forces $t_1 \in U_j$.

We now claim that, under inclusion, the family of such closed intervals forms a totally ordered set. In order to see this, we show that any such set $U_j = [a_j, b_j]$ is a symmetric closed interval around t_1 (i.e. that $(a_j + b_j)/2 = t_1$). But this is clear: one can just consider the restriction of \mathbf{p} to the two sets $[a_j, t_1]$ and $[t_1, b_j]$. This yields a pair of paths, parametrized by arclength, joining the point $\mathbf{p}(a_j) = \mathbf{p}(b_j)$ to the point $\mathbf{p}(t_1)$. Furthermore, each of these paths is reduced (since they are subpaths of the reduced paths \mathbf{p}_1 and \mathbf{p}_2 respectively). But we know by Proposition 2.13 that there is a unique reduced path in each endpoint-preserving homotopy class of paths joining a pair of points. Hence the two paths have to coincide, and as they are parametrized by arclength, we immediately obtain our claim.

Next, we argue that this totally ordered chain has a maximal element. Indeed, consider the set U defined to be the union of our sets U_j . We claim that U is still within our family. To see this, we merely note that, by our previous observation on the $U_j = [a_j, b_j]$, the restriction of \mathbf{p} to each symmetric (about t_1) subinterval of U consists merely of traversing some reduced path on $[a_j, t_1]$, followed by backtracking along the same path on $[t_1, b_j]$. By continuity, the same must hold for the symmetric closed interval \bar{U} , so that the closure of U also lies within our family. Hence $U = \bar{U}$, and we have found our maximal element.

It is now easy to complete our proof: if $[a, b]$ is our maximal interval U , we can now define our \mathbf{q}_i and \mathbf{r}_i explicitly. We set $\mathbf{q}_1 := \mathbf{p}|_{[0, a]}$, $\mathbf{r}_1 := \mathbf{p}|_{[a, t_1]}$, $\mathbf{r}_2 := \mathbf{p}|_{[t_1, b]}$, and $\mathbf{q}_2 := \mathbf{p}|_{[b, t_1 + t_2]}$. We note that it is clear that $\mathbf{p}_1 = \mathbf{r}_1 * \mathbf{q}_1$ and $\mathbf{p}_2 = \mathbf{q}_2 * \mathbf{r}_2$. From our proof, it is also immediate that $\mathbf{r}_1 = \mathbf{r}_2^{-1}$. Finally, since U was picked to be maximal, the path $\mathbf{q}_2 * \mathbf{q}_1$ must be reduced, completing the proof. □

An easy inductive argument gives the following corollary, which will underpin our consideration of more complicated paths later in the paper.

Corollary 2.17. *Let \mathbf{p}_i be reduced paths and assume that $\mathbf{p}_{i+1} * \mathbf{p}_i$ is reduced for all i . Then the path $\mathbf{p}_n * \mathbf{p}_{n-1} * \cdots * \mathbf{p}_2 * \mathbf{p}_1$ is reduced.*

Corollary 2.18. *Let η be a reduced loop in X . Then η can be expressed as a concatenation $\mathbf{p}^{-1} * \gamma * \mathbf{p}$, where \mathbf{p} is a reduced path, and γ is a geodesic loop.*

Proof. Let us view η , parametrized by arclength, as a map from $[0, t]$ into X . Consider the point $p := \eta(t/2)$, and consider the pair of paths $\mathbf{q}_1 := \eta|_{[0, t/2]}$ and $\mathbf{q}_2 := \eta|_{[t/2, t]}$. Observe that each of these paths is reduced (being a subpath of η), that they have common endpoints, and that $\eta = \mathbf{q}_2 * \mathbf{q}_1$. Now consider the concatenation of paths $\mathbf{q}_1 * \mathbf{q}_2$ and apply Lemma 2.16. Our claim immediately follows. □

For the next lemma, we need the following definition.

Definition 2.19. A path \mathbf{p} with domain $[0, t]$ is *non-self-terminating* if $\mathbf{p}(t) \notin \mathbf{p}([0, t))$. It is *non-self-originating* if $\mathbf{p}(0) \notin \mathbf{p}((0, t])$.

We now analyze the reduced loop within the homotopy class of a path-loop-path concatenation.

Lemma 2.20. *Let \mathbf{p} be a reduced path parametrized by arclength $[0, t]$, γ a cyclically reduced loop based at $p := \mathbf{p}(t)$, parametrized by arclength $[0, s]$ with $s > 0$. Then the unique reduced loop η in the homotopy class of $\mathbf{p}^{-1} * \gamma * \mathbf{p}$ must pass through p .*

Assume further that \mathbf{p} is non-self-terminating. Then η either coincides with \mathbf{p} on the interval $[0, t]$ or coincides with \mathbf{p}^{-1} on the interval $[l(\eta) - t, l(\eta)]$.

Proof. By contradiction, suppose η avoids p . First, assume that \mathbf{p}^{-1} has the (reduced) form $\mathbf{p}^{-1} = \mathbf{c}^{-1} * \gamma^{-n}$ for some maximal $n > 0$, where $\mathbf{c}^{-1} * \mathbf{p}$ is reducible so that it avoids p . Then $\mathbf{p} = \gamma^n * \mathbf{c}$ (note that the non-self-terminating assumption we will use for the second part of the lemma rules such behavior out). Then we reduce:

$$\mathbf{p}^{-1} * \gamma * \mathbf{p} = \mathbf{c}^{-1} * \gamma * \mathbf{c}.$$

We know that \mathbf{c}^{-1} does not fully cancel γ by the choice of n maximal above. We may now replace \mathbf{p} with \mathbf{c} and proceed.

So we can now assume that there is no initial factor of γ^{-1} in \mathbf{p}^{-1} . Thus $\gamma * \mathbf{p}$ must be reducible as it hits p . By Lemma 2.16, $(\gamma|_{[0, \epsilon]})^{-1} = \mathbf{p}|_{[t-\epsilon, t]}$. Consider instead $\mathbf{p}^{-1} * \gamma$, again reducible as it hits p . By Lemma 2.16, $\gamma|_{[s-\epsilon', s]} = (\mathbf{p}^{-1}|_{[0, \epsilon']})^{-1} = \mathbf{p}|_{[t-\epsilon', t]}$. Picking $\delta < \epsilon, \epsilon'$, we conclude that

$$(\gamma|_{[0, \delta]})^{-1} = \mathbf{p}|_{[t-\delta, t]} = \gamma|_{[s-\delta, s]}.$$

This contradicts the assumption that γ is cyclically reduced. Thus $p \in \eta$, completing the first part of the lemma.

For the second part of the lemma, we now assume that \mathbf{p} is non-self-terminating. If $\mathbf{p}^{-1} * \gamma * \mathbf{p}$ is already reduced, we are done. If not, Corollary 2.17 tells us that there is a reduction of $\mathbf{p}^{-1} * \gamma * \mathbf{p}$ around time t or $t + s$ (as \mathbf{p} and γ are individually reduced). By the arguments for the first part of the lemma, if reductions at *both* of these times are possible, there is a contradiction to γ being cyclically reduced.

Let us assume first that there is no reduction around time t . Then $\gamma * \mathbf{p}$ is reduced, and the only way η can fail to coincide with \mathbf{p} over $[0, t]$ is if \mathbf{p}^{-1} completely cancels γ and part of \mathbf{p} . But this is the situation in the first part of our arguments above; as we noted, it is ruled out by the non-self-terminating assumption. In this case, η coincides with \mathbf{p} over $[0, t]$.

Finally, let us assume that $\mathbf{p}^{-1} * \gamma * \mathbf{p}$ does not reduce around time $t + s$, and hence $\mathbf{p}^{-1} * \gamma$ is reduced. In this case, η will coincide with \mathbf{p}^{-1} over $[l(\eta) - t, l(\eta)]$ unless the initial segment of \mathbf{p}^{-1} completely cancels γ . Again, as noted above, this is ruled out by the non-self-terminating assumption. Thus, the final segment of η coincides with \mathbf{p}^{-1} and the second part of the lemma is proven. \square

Note that, in our previous lemma, we can always (by reversing γ if need be) get the reduced loop η to coincide with \mathbf{p} on the interval $[0, t]$.

3. π_1 -Hull and Structure Theory

This section is devoted to understanding the structure of an arbitrary compact 1-dimensional geodesic space X . In the first subsection, we will introduce the π_1 -hull $\text{Conv}(X)$ of X , and see that $\text{Conv}(X)$ is the “homotopically essential” part of the space X . In the second subsection, we introduce the notion of *branch point* of X , and use the set of branch points in $\text{Conv}(X)$ to analyze the structure of the π_1 -hull.

3.1. Structure theory: general case

Definition 3.1. Given a compact geodesic space X of topological dimension one, we define the π_1 -hull of X , denoted $\text{Conv}(X)$, as the union of (the images of) all non-constant geodesic loops in X . A space satisfying $X = \text{Conv}(X)$ is said to be π_1 -convex.

Recall that geodesic loops are both rectifiable, and cyclically reduced (and hence have minimal finite length in their free homotopy class, see Corollary 2.11). Note that, in the special case where X is contractible (e.g. if X is an \mathbb{R} -tree), the π_1 -hull is empty. We will establish some structure theory for arbitrary compact 1-dimensional geodesic space X , and show that all homotopy information about X is actually carried by its π_1 -hull. To begin, let us show how to extend some geodesic paths to geodesic loops.

Lemma 3.2. *Let \mathbf{p} be a non-self-terminating and non-self-originating geodesic path whose endpoints lie in $\text{Conv}(X)$. Then \mathbf{p} can be extended to a geodesic loop.*

Proof. Subdivide \mathbf{p} as $\mathbf{p} = \mathbf{p}_2 * \mathbf{p}_1$ meeting at $\mathbf{p}(t/2)$. Note that \mathbf{p}_1^{-1} and \mathbf{p}_2 are non-self-terminating. As $\mathbf{p}(0), \mathbf{p}(t) \in \text{Conv}(X)$, there exist γ_1, γ_2 geodesic loops passing through (and parametrized with basepoints at) the points $\mathbf{p}(0)$ and $\mathbf{p}(t)$, respectively. Let η_1 and η_2 be the reduced loops in the homotopy classes (based at $\mathbf{p}(t/2)$) for $\mathbf{p}_1 * \gamma_1 * \mathbf{p}_1^{-1}$ and $\mathbf{p}_2^{-1} * \gamma_2 * \mathbf{p}_2$, respectively. After re-orienting the γ_i if necessary, by Lemma 2.20 we have

$$\eta_1|_{[l(\eta_1)-t/2, l(\eta_1)]} = \mathbf{p}_1|_{[0, t/2]},$$

$$\eta_2|_{[0, t/2]} = \mathbf{p}_2|_{[0, t/2]}.$$

We can write $\eta_1 = \mathbf{p}_1 * \mathbf{c}_1$, $\eta_2 = \mathbf{c}_2 * \mathbf{p}_2$; these are both cyclically reduced. We now have a few possible cases.

Case I. $\mathbf{c}_1 * \mathbf{c}_2$ is a reduced path. Consider the closed path $\eta_2 * \eta_1 = \mathbf{c}_2 * \mathbf{p}_2 * \mathbf{p}_1 * \mathbf{c}_1$, which extends \mathbf{p} . It is a reduced path as η_i and \mathbf{p} are reduced. Under the assumption for this case, it is also cyclically reduced. This is the geodesic loop we were seeking.

Case II. $\mathbf{c}_1 * \mathbf{c}_2$ is not reduced. We will separately consider the two cyclic permutations of the path $\eta_2 * \eta_1$ which might fail to be reduced paths.

Case IIa. (see Fig. 1) Suppose the cyclic permutation of $\eta_2 * \eta_1$ given by $\mathbf{c}_1 * \mathbf{c}_2 * \mathbf{p}_2 * \mathbf{p}_1$ reduces to a geodesic loop not containing \mathbf{p} . This happens only if \mathbf{c}_1 totally cancels \mathbf{c}_2 and then part of \mathbf{p} . In such a case, we can write in reduced form $\mathbf{c}_1 = \mathbf{a} * \mathbf{c}_2^{-1}$ where $\mathbf{a} * \mathbf{p}_2$ is reducible. Write $\mathbf{a} = \mathbf{a}_2 * \mathbf{a}_1$ where \mathbf{a}_1 is the maximal sub-path with image along \mathbf{p} . We must have both \mathbf{a}_i nontrivial: \mathbf{a}_1 since \mathbf{a} partially cancels \mathbf{p} , and \mathbf{a}_2 by the fact that $\mathbf{p}_1 * \mathbf{c}_1 = \mathbf{p}_1 * \mathbf{a}_2 * \mathbf{a}_1 * \mathbf{c}_2^{-1}$ is reduced (if \mathbf{a}_2 is




Fig. 1. Path configuration for Case IIa.

trivial, $\mathbf{p}_1 * \mathbf{a}_1$ is reducible). Finally, let \mathbf{d} be the geodesic sub-path of \mathbf{p} connecting $\mathbf{p}(t/2)$ to the initial point of \mathbf{a}_2 .

We record the following facts: $\mathbf{c}_2 * \mathbf{p}_2 * \mathbf{p}_1$ is reduced. If \mathbf{d} lies along \mathbf{p}_2 , $\mathbf{d} * \mathbf{c}_2$ is reduced. If \mathbf{d} lies along \mathbf{p}_1 , $\mathbf{d} * \mathbf{c}_2$ is still reduced, as (in this subcase) \mathbf{c}_2 coincides with the initial segment of \mathbf{c}_1 , traversed backwards, and $\mathbf{c}_1 * \mathbf{p}_1$ is cyclically reduced. In addition, $\mathbf{a}_2 * \mathbf{d}$ is reduced, by definition of \mathbf{a}_2 . Finally, $\mathbf{p}_1 * \mathbf{a}_2$ is reduced as it is a sub-path of the reduced path $\mathbf{p}_1 * \mathbf{c}_1$.

Consider, then, the closed loop $\mathbf{a}_2 * \mathbf{d} * \mathbf{c}_2 * \mathbf{p}_2 * \mathbf{p}_1$. It is cyclically reduced, by the remarks in the previous paragraph; hence it is a geodesic loop and it proves the lemma.

Case IIb. Suppose the cyclic permutation of $\eta_2 * \eta_1$ given by $\mathbf{p}_2 * \mathbf{p}_1 * \mathbf{c}_1 * \mathbf{c}_2$ reduces to a geodesic loop not containing \mathbf{p} . Similarly to the case above, this happens only if $\mathbf{c}_2 = \mathbf{c}_1^{-1} * \mathbf{a}$ in reduced form where $\mathbf{p}_1 * \mathbf{a}$ is reducible. Write $\mathbf{a} = \mathbf{a}_2 * \mathbf{a}_1$ where \mathbf{a}_2 is the maximal sub-path of \mathbf{a} along \mathbf{p} . By the same arguments as in the previous case, both \mathbf{a}_1 and \mathbf{a}_2 are non-trivial. Again, let \mathbf{d} be the geodesic sub-path of \mathbf{p} connecting $\mathbf{p}(t/2)$ to the endpoint of \mathbf{a}_1 .

We record: $\mathbf{a}_1 * \mathbf{p}_2 * \mathbf{p}_1$ is reduced. $\mathbf{d}^{-1} * \mathbf{a}_1$ is reduced by definition of \mathbf{a}_1 . $\mathbf{c}_1 * \mathbf{d}^{-1}$ is reduced because η_1 and η_2 are cyclically reduced. Finally, $\mathbf{p}_1 * \mathbf{c}_1$ is reduced.

Consider $\mathbf{c}_1 * \mathbf{d}^{-1} * \mathbf{a}_1 * \mathbf{p}_2 * \mathbf{p}_1$. It is a geodesic loop extending \mathbf{p} by the facts presented in the previous paragraph. □

To illustrate the usefulness of the previous lemmas, we note the following immediate corollary:

Corollary 3.3. *Suppose X is a compact, 1-dimensional, geodesic metric space with $\text{Conv}(X) \neq \emptyset$. Then $\text{Conv}(X)$ is path connected. Furthermore, $\text{Conv}(X)$ is a strongly convex subset of X .^b*

Proof. Let $p, q \in \text{Conv}(X)$ be an arbitrary pair of distinct points, and let \mathbf{p} be a distance minimizer joining the two points. It is clearly non-self-terminating and non-self-originating. Since \mathbf{p} is a distance minimizer, it is geodesic (Corollary 2.15). So by Lemma 3.2, there is a geodesic loop extending it. This immediately shows that \mathbf{p} itself lies in $\text{Conv}(X)$. Both our claims follow. □

We can now establish some basic properties of the π_1 -hull.

Proposition 3.4. *Suppose X is a compact, 1-dimensional, geodesic metric space with $\text{Conv}(X) \neq \emptyset$. Then $\text{Conv}(X)$ is also a compact, 1-dimensional, geodesic metric space.^c*

^bWe say that a subset of a geodesic space is strongly convex provided that for every pair of points in the subset, every distance minimizer joining them also lies within the subset.

^cThe authors are indebted to J. W. Cannon for suggesting this result and the main idea of its proof.

Proof. Corollary 3.3 covers convexity; we need only show $\text{Conv}(X)$ is closed to prove the result. Let $p \in \overline{\text{Conv}(X)}$. The proof breaks down into two cases, according to whether there is an ϵ -neighborhood of p containing no geodesic loop.

Case I. For some $\epsilon > 0$, the ϵ -neighborhood N of p contains no closed geodesic. In this case, note that points in N are *uniquely* arcwise connected (by convention, arcs will be reduced, *rectifiable* paths, and uniqueness is of course up to reparametrization). For if not, we can cyclically reduce the concatenation of two such arcs to obtain a geodesic loop in N , a contradiction.

We must then have that for every i there exists a closed geodesic γ_i in X which intersects the ϵ/i -neighborhood of p . Let α_i be a component of $N \cap \gamma_i$; let $x_i \in \alpha_i$ be a point at minimum distance from p . This point is unique, as otherwise we could form a simple closed curve in N by connecting two such points to p with minimizing paths and to each other along α_i . By choosing the component α_i appropriately, we may assume $d(p, x_i) < \epsilon/i$.

The points x_i divide α_i into two arcs, A_i and B_i . By passing to a subsequence, we may assume that the arcs A_i converge to an arc A joining the exterior of N to p , and that the B_i converge to an arc B doing the same. As N is uniquely arc-connected, $A \cap B$ is either p alone or some geodesic segment ending at p .

If $A \cap B = \{p\}$, take i very large, so that A_i and A very nearly agree over a comparatively long segment of A . As X is a geodesic space, we must then be able to connect A_i to A with short geodesics near the end points. Unless A_i and A in fact agree over a long segment, this contradicts unique arc-connectedness of N . The same argument holds for B and B_i . We claim that for large i , A_i and B_i coincide with A and B *all the way to* p . If not, then one can look at the short arc along the corresponding α_i where the A_i, B_i differ from A, B . This gives a short path joining the segments A, B together. Concatenating this path with the portion of A and B going to p provides a closed curve in N which can be shortened to a geodesic, a contradiction. We conclude that A_i and B_i hit p , and p belongs to the corresponding closed geodesic γ_i , as desired.

Similarly, if $A \cap B = I$, an interval with p as an endpoint, take i very large so that A_i and B_i coincide with I over a comparatively large interval. If they coincide all the way to x_i , then this contradicts the fact that α_i was geodesic. Otherwise, we can use the portion of α_i near x_i where they differ to obtain a geodesic loop in N , again a contradiction. This completes the proof of Case I.

Case II. (see Fig. 2) For every $\epsilon > 0$, the ϵ -neighborhood N of p contains a closed geodesic. If p belongs to one such geodesic, we are done. If not, connect each closed geodesic to p by a distance-minimizing path. Any two such paths must coincide on some interval with p as an endpoint, otherwise p belongs to a closed geodesic formed by these paths and the closed geodesics to which they connect. Thus we may assume that p lies at the endpoint of a geodesic path \mathbf{p} to which a sequence of geodesic loops γ_i with quickly decreasing length are connected by geodesic segments \mathbf{t}_i , also with quickly decreasing length. We build a closed geodesic on which p lies




Fig. 2. The situation for Case II.

as follows. Start at p . Follow \mathbf{p} to its intersection with \mathbf{t}_1 . Follow $\mathbf{t}_1^{-1} * \gamma_1 * \mathbf{t}_1$. Follow \mathbf{p} to its intersection with \mathbf{t}_2 and repeat. Continue this process; picking the sequence of paths to decrease in length sufficiently quickly gives a rectifiable curve \mathbf{l} approaching p , defined over times $[0, t)$. Set $\mathbf{l}(t) = p$. This loop is geodesic assuming each \mathbf{t}_i is chosen to meet γ_i and \mathbf{p} in a single point each. \square

Having analyzed the π_1 -hull $Conv(X)$, we now turn our attention to the various connected components of $X \setminus Conv(X)$.

Definition 3.5. For Z a metric space, define an equivalence relation \sim on points of Z by setting $x \sim y$ if there exists a rectifiable path in Z joining x to y . A single \sim equivalence class is called a *rectifiable component* of Z , and if all points in Z are \sim equivalent to each other, we say that Z is *rectifiably connected*. If every point in Z has a neighborhood base consisting of open, rectifiably connected sets, we say that Z is *locally rectifiably connected*.

For example, any length space is rectifiably connected. Since open metric balls in a length space are obviously rectifiably connected, length spaces are also locally rectifiably connected.

Clearly, any rectifiable component of Z is entirely contained within a single path component of Z . But one could *a priori* have a path component of Z which breaks up into several distinct rectifiable components.

Lemma 3.6. *Let Z be a metric space, and consider the partition of Z into (i) connected components, (ii) path components, and (iii) rectifiable components. If Z is locally rectifiably connected, then these three partitions of Z coincide.*

Proof. Since Z is locally rectifiably connected, it is also locally path connected. A basic result in point set topology asserts that for locally path connected spaces, path components coincide with connected components, giving the equivalence of partitions (i) and (ii).

For the equivalence of (ii) and (iii), observe that since Z is locally rectifiably connected, each rectifiable component is open. If we had a path component P of

Z breaking up into several rectifiable components, this would give a partition of P into pairwise disjoint open sets. But since the partitions (i) and (ii) coincide, P is also a connected component of Z . We conclude that P must consist of a single rectifiable component, as desired. \square

Corollary 3.7. *Suppose X is a compact, 1-dimensional, geodesic metric space, and let $X \setminus Conv(X) = \coprod_{i \in I} Z_i$ be the decomposition of $X \setminus Conv(X)$ into connected components. Then each Z_i is rectifiably connected, and $\overline{Z_i} \cap Z_j = \emptyset$ if $i \neq j$.*

Proof. X is a geodesic space, so it is locally rectifiably connected. Proposition 3.4 tells us $Conv(X) \subseteq X$ is a closed subset. As $X \setminus Conv(X)$ is open, it inherits the property of being locally rectifiably connected. Lemma 3.6 implies that the connected components Z_i are all rectifiably connected. For the second statement, assume that $i \neq j$, and $x \in Z_j$. As Z_j is open, it is itself a neighborhood of x which is disjoint from Z_i , and hence $x \notin \overline{Z_i}$. \square

Proposition 3.8. *Suppose X is a compact, 1-dimensional, geodesic metric space. Then each connected component of $X \setminus Conv(X)$ is a strongly convex subset of X , isometric to an \mathbb{R} -tree.^d*

Proof. If Z is any connected component of $X \setminus Conv(X)$, Corollary 3.7 tells us Z is rectifiably connected. Let $x, y \in Z$, and let $\eta \subset X$ be any distance minimizer from x to y . To show that Z is strongly convex, we need to argue that η lies in Z . By way of contradiction, let us assume that η passes through $Conv(X)$. As Z is rectifiably connected, we can find a rectifiable path $\eta_Z^\circ \subset Z$ joining x to y . Viewing the path η_Z° as a path in X , we can apply Proposition 2.13 to obtain a reduced path η_Z homotopic (rel. endpoints) to η_Z° . The path η_Z has image contained within the image of η_Z° , forcing $\eta_Z \subset Z$. Now concatenating the two paths η and η_Z yields a closed, rectifiable, loop $\gamma^\circ = \eta * \eta_Z$. Consider the geodesic loop γ obtained by cyclically reducing the loop γ° . From the definition of $Conv(X)$, we have that $\gamma \subset Conv(X)$. But γ was obtained by cyclically reducing the concatenation $\eta * \eta_Z$, where η_Z was a reduced path contained entirely in $Z \subset X \setminus Conv(X)$. Since η_Z must be fully cancelled in the cyclic reduction, but is itself a reduced path, it follows that the image of η_Z must be contained in the image of η . The path η is a distance minimizer, hence an embedded path. As η_Z has image lying within the same set, and joins together the two endpoints, it must be a reparametrization of η . This yields a contradiction, as $\eta_Z \subset X \setminus Conv(X)$, while $\eta \cap Conv(X) \neq \emptyset$. We conclude that the distance minimizer η must satisfy $\eta \subset Z$, and hence that Z is indeed strongly convex.

Next we note that, given any two points $p, q \in Z$, there is a *unique* reduced rectifiable path \overline{pq} in Z (up to reparametrization) joining p to q , i.e. Z is uniquely

^dSee [2] for a reference on \mathbb{R} -trees.

arcwise connected. For if η, η' were two such paths, whose images in Z did not coincide, we could cyclically reduce $\eta * \eta'$ to obtain a geodesic loop in Z , contradicting $Z \cap \text{Conv}(X) = \emptyset$. It is now easy to see that Z is a 0-hyperbolic geodesic space: given any three points $x, y, z \in Z$, consider the distance minimizers $\overline{xy}, \overline{yz}$, and \overline{xz} . Reducing the concatenation $\overline{xy} * \overline{yz}$ gives us a reduced rectifiable path whose image lies within the set $\overline{xy} \cup \overline{yz}$, and joins x to z . Since \overline{xz} is another reduced rectifiable path joining these two points, the uniqueness kicks in and forces $\overline{xz} \subseteq \overline{xy} \cup \overline{yz}$. Finally, it is a well-known result that 0-hyperbolic geodesic spaces are precisely \mathbb{R} -trees, concluding our proof. \square

An immediate application of Proposition 3.8 is:

Corollary 3.9. *Suppose X is a compact, 1-dimensional, geodesic metric space. Then the following three statements are equivalent: (i) X is contractible, (ii) $\text{Conv}(X) = \emptyset$ and (iii) X is an \mathbb{R} -tree.*

Now that we understand the connected components of $X \setminus \text{Conv}(X)$, let us see how these attach together. In view of Corollary 3.7, distinct connected components of $X \setminus \text{Conv}(X)$ do not interact. We now study how they attach to $\text{Conv}(X)$.

Proposition 3.10. *Suppose X is a compact, non-contractible, 1-dimensional, geodesic metric space (so $\text{Conv}(X) \neq \emptyset$). Let $X \setminus \text{Conv}(X) = \coprod_{i \in I} Z_i$ be the decomposition of $X \setminus \text{Conv}(X)$ into connected components. Then we have that each $\overline{Z_i} \cap \text{Conv}(X)$ consists of a single point x_i , and the (metric) completion of each Z_i is precisely $Z_i \cup \{x_i\}$. Moreover, the index set I is countable, and $\lim_{i \rightarrow \infty} \text{diam}(Z_i) = 0$.*

Proof. Let Z'_i denote the metric completion of the space Z_i . Observe that the metric completion of an \mathbb{R} -tree is again an \mathbb{R} -tree (this follows easily from the 0-hyperbolicity characterization of \mathbb{R} -trees), so Z'_i is a bounded, complete, \mathbb{R} -tree. Since Z_i is a connected dense subset of the \mathbb{R} -tree Z'_i , we see that for each pair of distinct points $p \neq q \in Z'_i \setminus Z_i$, the distance minimizer $\overline{pq} \subset Z'_i$ joining them satisfies $\overline{pq} \cap (Z'_i \setminus Z_i) = \{p, q\}$, and hence $\overline{pq} \cap Z_i \neq \emptyset$.

Since $Z_i \subset X$ is strongly convex, there is a natural surjective map $\rho : Z'_i \rightarrow \overline{Z_i}$ which extends the identity map on Z_i . Clearly ρ restricts to a surjection from $Z'_i \setminus Z_i$ to the set $\overline{Z_i} \cap \text{Conv}(X)$. If $p, q \in Z'_i \setminus Z_i$ with $p \neq q$ satisfy $\rho(p) = \rho(q)$, then the ρ -image of the distance minimizer $\overline{pq} \subset Z'_i$ joining p to q gives us a geodesic loop in X which passes through points in $Z_i \subset X \setminus \text{Conv}(X)$, a contradiction. We conclude that $\rho : Z'_i \setminus Z_i \rightarrow \overline{Z_i} \cap \text{Conv}(X)$ is also an *injective* map, and hence a bijection.

Assume $p \neq q$ are distinct points lying in the set $\rho(Z'_i \setminus Z_i)$, and let $p', q' \in Z'_i$ be their ρ -preimages. Since $p' \neq q'$, the distance minimizer $\eta \subset Z'_i$ joining them satisfies $\eta \setminus \{p', q'\} \subset Z_i$. On the other hand, the points $p \neq q$ lie in $\text{Conv}(X)$, so by Lemma 3.3, we can find a distance minimizer η° joining them within the set $\text{Conv}(X)$. Look at the concatenation $\rho(\eta) * \eta^\circ$ of the reduced paths $\rho(\eta)$ and η° .

These give a closed rectifiable loop, which in view of the discussion above (and of Lemma 2.16) is cyclically reduced. So this defines a geodesic loop, which passes through points in Z_i , a contradiction. Thus $\rho(Z_i' \setminus Z_i)$ consists of at most one point. But $\overline{Z_i} \setminus Z_i$ must have at least one point, for otherwise Z_i would be both open and closed in the connected space X , a contradiction (as we are assuming that $\text{Conv}(X) \neq \emptyset$). We conclude that $\overline{Z_i} \setminus Z_i = \overline{Z_i} \cap \text{Conv}(X)$ consists of a single point x_i , as desired.

Finally, for each natural number $n \in \mathbb{N}$, consider the set $I_n \subset I$ of indices such that the corresponding connected components Z_i ($i \in I_n$) have diameter $\geq 1/n$. We claim this set is *finite*. For if not, one has an injection $i : \mathbb{N} \hookrightarrow I_n$. Choose a point $x_k \in Z_{i(k)}$ with the property that $d(x_k, \text{Conv}(X)) \geq 1/n$. Then the sequence $\{x_k\}$ in X has no convergent subsequence, contradicting compactness. Since $I = \bigcup_{n \in \mathbb{N}} I_n$ is a countable union of finite sets, it is itself countable. The statement concerning the diameters of Z_i also follows. \square

Summarizing what we have so far, we see that an arbitrary compact 1-dimensional geodesic space X consists of:

- its π_1 -hull $\text{Conv}(X)$, which is itself a (π_1 -convex) compact 1-dimensional geodesic space, sitting as a strongly convex subset of X (see Corollary 3.3 and Proposition 3.4), and
- a countable collection of compact \mathbb{R} -trees $\overline{Z_i}$ (whose diameters are shrinking to zero), each of which is attached to the π_1 -hull $\text{Conv}(X)$ along a single terminal vertex x_i (see Propositions 3.8 and 3.10).

This structural result has a few nice consequences.

Corollary 3.11. *Let X be a compact 1-dimensional geodesic space, and assume X is not contractible (so $\text{Conv}(X) \neq \emptyset$). Then X deformation retracts onto its π_1 -hull $\text{Conv}(X)$. In particular, the inclusion $\text{Conv}(X) \hookrightarrow X$ induces an isomorphism $\pi_1(\text{Conv}(X)) \cong \pi_1(X)$.*

Proof. Each compact \mathbb{R} -tree $\overline{Z_i}$ deformation retracts to the corresponding terminal vertex x_i . It is an easy exercise to check that these homotopies glue together to define a deformation retraction of X to $\text{Conv}(X)$; that the *diam*(Z_i) shrink to zero is key to the proof. \square

As another application, we can now provide an alternate characterization of the π_1 -hull of X : it is the unique minimal deformation retract of X .

Corollary 3.12. *Let X be a compact 1-dimensional geodesic space, and assume X is not contractible (so $\text{Conv}(X) \neq \emptyset$). Assume we have a subset $X_\circ \subseteq X$ satisfying the following two properties: (i) X deformation retracts to X_\circ , and (ii) if X deformation retracts to some subset $Y \subseteq X$, then $X_\circ \subseteq Y$. Then X_\circ coincides with the π_1 -hull $\text{Conv}(X)$.*

Proof. If a subset X_\circ satisfying properties (i) and (ii) exists, it must be unique. The fact that $\text{Conv}(X)$ satisfies (i) is just Corollary 3.11 above. Now assume that X deformation retracts to Y , and let us argue that $\text{Conv}(X) \subseteq Y$. Let $\gamma \subset X$ be an arbitrary geodesic loop. Under the deformation retraction $\rho_t : X \rightarrow Y$, the geodesic loop γ maps to a loop $\rho_1(\gamma)$ which is freely homotopic to γ (via the homotopy ρ_t). By Theorem 2.10, we have the containments of sets $\gamma \subseteq \rho_1(\gamma) \subseteq Y$, which gives us $\gamma \subset Y$. We conclude $\text{Conv}(X) \subseteq Y$, showing $\text{Conv}(X)$ satisfies (ii). \square

Remark 3.13. Conner and Meilstrup showed in [9, Corollary 4.4] that *any* Peano continua has a *core*, a unique minimal strong deformation retract. In view of our Corollary 3.12, we see that, in the case where S is a compact 1-dimensional geodesic space, our π_1 -hull $\text{Conv}(X)$ coincides with the Conner–Meilstrup core of X . They also establish, in the case of 1-dimensional Peano continua, a structure theory (see [9, Theorem 4.3]) that is similar to our Proposition 3.10. Our arguments rely on the geodesic space structure, and as a result, are also suited to establishing metric properties of the π_1 -hull (for instance, our Proposition 3.4), results which do not seem accessible from Conner–Meilstrup’s techniques. Conversely, we do not know whether our results can recover the 1-dimensional version of the Conner–Meilstrup theorem (as we do not know whether every 1-dimensional Peano continua supports a geodesic space structure).

3.2. Structure theory: the π_1 -hull

In the previous subsection, we reduced the study of general compact 1-dimensional geodesic spaces to the study of their π_1 -hull. In this subsection, we focus on understanding the structure of the π_1 -hull $\text{Conv}(X)$. Our analysis starts with the notion of branch point.

Definition 3.14. Let X be a compact 1-dimensional geodesic space, p a point in X . We say that X has *branching* at p provided there exists a triple of geodesic paths $\gamma_i : [0, \epsilon] \rightarrow X$ with the following properties:

- $\gamma_i(0) = p$ for all three paths,
- each concatenated path $\gamma_i * \gamma_j^{-1}$ is a reduced (and hence geodesic) path.

In other words, there are at least three distinct germs of geodesic paths originating at the point p . If X has branching at p , we call p a *branch point* of X .

Away from the set of branch points, the local topology of X is fairly simple, as indicated in the following proposition.

Proposition 3.15. *Let X be a compact 1-dimensional geodesic space, and $\mathcal{B}(X) \subset X$ the subset of all branch points of X . Assume that the point p does not lie in the closure $\overline{\mathcal{B}(X)}$ of the set of branch points (i.e. $p \in X \setminus \overline{\mathcal{B}(X)}$). Then for ϵ small*

enough, the metric ϵ -neighborhood of p is isometric to either:

- (1) the half-open interval $[0, \epsilon)$, with the point p corresponding to 0, or
- (2) an open interval $(-\epsilon, \epsilon)$, with the point p corresponding to 0.

Proof. First, we claim that for δ small enough, the open δ -ball $B_p(\delta)$ centered at p is isometric to an \mathbb{R} -tree. This will follow from the fact that the δ -ball contains no geodesic loops (see the end of the proof of Proposition 3.8). Indeed, since p does not lie in the closure $\overline{\mathcal{B}(X)}$, by choosing δ small, we can ensure that $B_p(\delta)$ contains no branch points, and that the complement $X \setminus B_p(\delta)$ is non-empty. If γ is a geodesic loop contained in $B_p(\delta)$, take a point x outside the ball, and let η be a distance minimizer from x to the curve γ . Let L be the length of γ , and choose the parametrization $\gamma : [-L/2, L/2] \rightarrow X$ so that $\gamma(0)$ denotes the endpoint of η on γ . Then the point $\gamma(0) \in B_p(\delta)$ is a branch point: the three geodesics $\gamma|_{[0, L/2]}$, $(\gamma|_{[-L/2, 0]})^{-1}$, and η^{-1} satisfy the conditions of Definition 3.14. This is a contradiction, hence $B_p(\delta)$ contains no geodesic loops, and so must be isometric to an \mathbb{R} -tree.

Now consider the connected components of $B_p(\delta) \setminus \{p\}$. If there were ≥ 3 such connected components, take points x_1, x_2, x_3 in three distinct connected components, and let η_i be a distance minimizer from p to x_i . It is immediate that the η_i satisfy the conditions of Definition 3.14, showing that p is a branch point, a contradiction. So we have that there are either one or two connected components in $B_p(\delta) \setminus \{p\}$. We consider each of these cases separately.

Case 1. If there is only one connected component, then p must be a vertex of the \mathbb{R} -tree $B_p(\delta)$. Choose $\gamma : [0, \epsilon] \rightarrow X$ a distance minimizer from p to some point in $B_p(\delta) \setminus \{p\}$. Given any point $q \in B_p(\delta) \setminus \{p\}$, we can consider the distance minimizer $\eta_q : [0, \delta'] \rightarrow B_p(\delta)$ from p to q . Since p is a vertex of the \mathbb{R} -tree $B_p(\delta)$, we have that η_q and γ must coincide on some neighborhood of p , i.e. there exists a corresponding real number $0 < \delta_q \leq \min(\epsilon, \delta')$, with the property that $\gamma \equiv \eta_q$ on the interval $[0, \delta_q]$. If δ_q were *strictly* smaller than $\min(\epsilon, \delta')$, then the point $\gamma(\delta_q)$ would be a branch point: the three geodesics $(\gamma|_{[0, \delta_q]})^{-1}$, $\gamma|_{[\delta_q, \epsilon]}$, and $\eta|_{[\delta_q, \delta']}$ satisfy the conditions of Definition 3.14. But the neighborhood $B_p(\delta)$ was chosen to contain no branch points, forcing $\delta_q = \min(\epsilon, \delta')$. So if q is any point at distance $< \epsilon$ from p , then $\delta_q = \min(\epsilon, \delta') = \delta'$, and we have that $\eta_q \equiv \gamma|_{[0, \delta']}$, i.e. the point q lies on γ . This immediately implies that the metric ϵ -neighborhood of p consists precisely of the points along the distance minimizer γ , giving the first statement in the proposition.

Case 2. If there are two connected components, choose $\gamma_1, \gamma_2 : [0, \epsilon] \rightarrow X$ to be a pair of distance minimizers from p to points in the two distinct components of $B_p(\delta) \setminus \{p\}$. The concatenation $\gamma := \gamma_1 * \gamma_2^{-1}$ is a distance minimizer of length 2ϵ , passing through the point p , and entirely contained in $B_p(\delta)$. If q is any point at distance $< \epsilon$ from p , the distance minimizer $\eta_q : [0, \delta'] \rightarrow X$ must coincide

with either $\gamma_1|_{[0,\delta']}$, or with $\gamma_1|_{[0,\delta]}$ (otherwise, as in Case 1, the first point from which they start differing would give a branch point in $B_p(\delta)$, a contradiction). We conclude that the metric ϵ -neighborhood of p consists precisely of the points along the distance minimizer γ , giving the second statement in the proposition. \square

Corollary 3.16. *Let X be a compact 1-dimensional geodesic space, and $\mathcal{B}(X) \subset X$ the subset of all branch points of X . If X is π_1 -convex, and $\mathcal{B}(X) = \emptyset$, then X is isometric to the circle S^1 of some radius $r > 0$.*

Proof. From Proposition 3.15, we know that each point $p \in X$ has a (metric) neighborhood $B_\epsilon(p)$ isometric to either (i) a half-open interval, or (ii) an open interval. Since X is assumed to be π_1 -convex, we can rule out (i), for otherwise we could deformation retract X onto a proper subset of itself, contradicting Corollary 3.12. Since X is a compact geodesic space, it is second countable and Hausdorff. Hence it is a compact connected 1-dimensional manifold, so must be homeomorphic to S^1 . Finally, it is easy to see that geodesic metric space structures on S^1 are completely determined (up to isometry) by their diameter. \square

With this result in hand, we can now study the complement of the set of branch points in $Conv(X)$.

Lemma 3.17. *Let X be a compact 1-dimensional geodesic space, and $\mathcal{B}(X) \subset X$ the subset of all branch points of X . Assume $\gamma : [0, L] \rightarrow X \setminus \mathcal{B}(X)$ is a geodesic path. Then γ is locally a distance minimizer, and $\gamma((0, L)) \subset X \setminus \overline{\mathcal{B}(X)}$.*

Proof. Take any $t \in (0, L)$, and consider the numbers $\sup_{s \in [0,t]} d(\gamma(t), \gamma(s))$ and $\sup_{s \in [t,L]} d(\gamma(t), \gamma(s))$. As γ is parametrized by arclength, both these numbers are > 0 , and we can choose ϵ so that 2ϵ is smaller than both of these. Now consider the metric ϵ -ball $B_{\gamma(t)}(\epsilon)$ centered at $\gamma(t)$. We first claim that, as a set, this metric ball is entirely contained in the image of γ . If not, we can find a point $p \in X$ which satisfies $d(p, \gamma(t)) < \epsilon$, and which does *not* lie on the image of γ . Let η be a distance minimizer from p to the image of γ (a compact set). By the choice of ϵ , η terminates at a point on the image of γ which is distinct from $\gamma(0), \gamma(L)$, so yields a branch point on the image of γ . This contradicts the fact that γ lies in $X \setminus \mathcal{B}(X)$.

Now that we know that the set $B_{\gamma(t)}(\epsilon)$ is contained in the image of γ , we proceed to show that (possibly after shrinking ϵ) it in fact coincides with $\gamma((t - \epsilon, t + \epsilon))$. Indeed, take a point x_1 at distance ϵ from $\gamma(0)$, and let $\eta_1 : [0, \epsilon] \rightarrow X$ be a distance minimizer from $\gamma(t)$ to x_1 . Since $\gamma(t)$ is not a branch point, η_1 must initially coincide with one of the two geodesics $\gamma|_{[t,s]}$ and $(\gamma|_{[0,t]})^{-1}$. In fact, we must have either $\eta_1 \equiv \gamma|_{[t,t+\epsilon]}$ or $\eta_1 \equiv (\gamma|_{[t-\epsilon,t]})^{-1}$, for otherwise, the first point from which they start disagreeing would be a branch point on the curve γ , a contradiction. Without loss of generality, we may now assume that $\gamma|_{[t,t+\epsilon]}$ is a distance minimizer.

Next, note that the metric ball $B_{\gamma(t)}(\epsilon/2)$ *cannot* consist solely of the points on the curve $\gamma|_{[t,t+\epsilon/2]}$, for otherwise $(\gamma|_{[t-\epsilon/2,t]})^{-1}$ would have to coincide with

$\gamma|_{[t, t+\epsilon/2]}$, contradicting the fact that γ is reduced. Let x_2 be a point in $B_{\gamma(t)}(\epsilon/2)$ which does not lie on $\gamma|_{[t, t+\epsilon/2]}$, and let η_2 be a distance minimizer from x_2 to the (compact) set $\gamma([t, t+\epsilon])$. If η_2 terminates on a point in $\gamma((t, t+\epsilon))$, we would obtain a branch point on γ , a contradiction. The triangle inequality implies that η_2 cannot terminate at the point $\gamma(t+\epsilon)$. Hence η_2 must terminate at $\gamma(t)$. This gives us a distance minimizer η_2^{-1} which intersects the distance minimizer $\gamma|_{[t, t+\epsilon]}$ only at their common initial point $\gamma(t)$.

Since $\gamma(t)$ is not branching, the geodesic $(\gamma|_{[0, t]})^{-1}$ must initially coincide with either η_2^{-1} or with $\gamma|_{[t, t+\epsilon]}$. As γ is reduced, we see that $(\gamma|_{[0, t]})^{-1}$ must coincide with η_2^{-1} . So at the cost of further shrinking ϵ , we can assume that both $\gamma|_{[t-\epsilon, t]}$ and $\gamma|_{[t, t+\epsilon]}$ are distance minimizers, that only intersect at the point $\gamma(t)$. Finally, we can conclude that their union $\gamma((t-\epsilon, t+\epsilon))$ is exactly the metric ball $B_{\gamma(t)}(\epsilon)$. For if not, then taking a distance minimizer η_3 from a point $x_3 \in B_{\gamma(t)}(\epsilon) \setminus \gamma((t-\epsilon, t+\epsilon))$ to the closest point on $\gamma((t-\epsilon, t+\epsilon))$ would yield a branch point on γ .

So each point in $\gamma((0, L))$ has a metric neighborhood isometric to an open interval contained entirely in the set $\gamma((0, L)) \subset X \setminus \mathcal{B}(X)$. As this is a neighborhood which is disjoint from $\mathcal{B}(X)$, we conclude that each of these points lies in the complement of $\overline{\mathcal{B}(X)}$, completing the proof. □

Lemma 3.18. *Suppose X is a π_1 -convex, compact, 1-dimensional, geodesic metric space, and let $X \setminus \overline{\mathcal{B}(X)} = \coprod_{i \in I} W_i$ be the decomposition of $X \setminus \overline{\mathcal{B}(X)}$ into connected components. Then each W_i is rectifiably connected, and $\overline{W_i} \cap W_j = \emptyset$ if $i \neq j$.*

Proof. The proof is identical to that of Corollary 3.7. □

The fact that the W_i are rectifiably connected tells us that they have a well-defined intrinsic length space structure.

Lemma 3.19. *Suppose X is a π_1 -convex compact 1-dimensional geodesic metric space, and assume $\mathcal{B}(X) \neq \emptyset$ (so X is not homeomorphic to S^1). Then each connected component W of $X \setminus \overline{\mathcal{B}(X)}$, equipped with the induced intrinsic geodesic metric, is isometric to an open interval of finite length.*

Proof. The argument from Corollary 3.16 applies verbatim to give that W is a connected 1-dimensional manifold, so is either homeomorphic to S^1 or to an open interval. We can rule out S^1 , for otherwise W would be both closed (being a compact subset of the Hausdorff space X) and open (being a connected component of the open set $X \setminus \overline{\mathcal{B}(X)}$) proper subset of X (since by hypothesis $\mathcal{B}(X) \neq \emptyset$). But this would violate the fact that X is connected.

Next, note that Lemma 3.18 tells us that each W inherits a well-defined intrinsic length space structure. Since W is homeomorphic to an open interval, this is actually a geodesic metric space: given any two points, there is a unique (up to reparametrization) embedded path joining them, which must be rectifiable (by Lemma 3.18) and is of minimal length amongst all rectifiable paths joining the two

points. To conclude, we merely observe that a geodesic metric space structure on an open interval is completely determined (up to isometry) by the diameter of the interval. In the case of W , this diameter must be finite, for if W were isometric to \mathbb{R} , then there would be no rectifiable path joining the point corresponding to $0 \in \mathbb{R}$ with a point in $\overline{\mathcal{B}(X)}$, contradicting the fact that X is a geodesic space. \square

Proposition 3.20. *Suppose X is a π_1 -convex compact 1-dimensional geodesic metric space, and assume $\mathcal{B}(X) \neq \emptyset$ (so X is not homeomorphic to S^1). Let $X \setminus \overline{\mathcal{B}(X)} = \coprod_{i \in J} W_i$ be the decomposition of $X \setminus \overline{\mathcal{B}(X)}$ into connected components. Then we have that each $\overline{W_i} \cap \overline{\mathcal{B}(X)}$ consists of at most two points. Moreover, the index set J is countable, and $\lim_{i \rightarrow \infty} \text{diam}(W_i) = 0$.*

Proof. The metric completion W'_i of an open interval of finite length is a closed interval of the same length. There is a natural surjective map $W'_i \rightarrow \overline{W_i}$ extending the identity map on W_i . Since $W'_i \setminus W_i$ consists of two points, and maps to $\overline{W_i} \cap \overline{\mathcal{B}(X)}$, the first claim follows. The argument in the last paragraph of the proof of Proposition 3.10 applies almost verbatim to give the statement concerning the cardinality of the indexing set and the limit of the diameters. \square

Lemma 3.21. *Suppose X is a π_1 -convex compact 1-dimensional geodesic metric space, and assume $\overline{\mathcal{B}(X)} \neq \emptyset$ (so X is not homeomorphic to S^1). Let W be a component of $X \setminus \overline{\mathcal{B}(X)}$ which is attached along one of its endpoints to a point $p \in \overline{\mathcal{B}(X)} \setminus \mathcal{B}(X)$. Let $\gamma : [0, r] \rightarrow X$ coincide with \overline{W} , parametrized by arclength, and satisfying $\gamma(0) = p$. Assume $\hat{\gamma} : [-\epsilon, r] \rightarrow X$ is any geodesic path extending the geodesic γ . Then for small enough ϵ' , the curve $\hat{\gamma}|_{[-\epsilon', 0]}$ is a distance minimizer. Moreover, there exists a strictly increasing sequence $\{t_i\}_{i \in \mathbb{N}} \subset [-\epsilon', 0)$ with the property that $\lim t_i = 0$, and each point $\hat{\gamma}(t_i)$ is a branch point.*

Proof. Let d denote the maximal distance from p to a point on $\hat{\gamma}([-\epsilon, 0])$, and choose ϵ_1 so that $0 < \epsilon_1 < d$. Since $p \in \overline{\mathcal{B}(X)} \setminus \mathcal{B}(X)$, there exists a branching point x_1 satisfying $0 < d(p, x_1) < \epsilon_1$. Let η_1 be a distance minimizer from p to x_1 . Since p is not a branch point, η_1 starts out coinciding with either $(\hat{\gamma}|_{[-\epsilon, 0]})^{-1}$ or with $\hat{\gamma}|_{[0, r]} \equiv \gamma$. If η_1 lies entirely along one of these curves, then x_1 lies on $\hat{\gamma}$, and we can let t_1 satisfy $\hat{\gamma}(t_1) = x_1$. Otherwise, there is a first occurrence after which the two curves are distinct. We can then set $t_1 \in [-\epsilon, r]$ to be the parameter at which η_1 diverges from the $\hat{\gamma}$ curve. Then the point $\hat{\gamma}(t_1)$ is branching, by an argument identical to that in Proposition 3.15. Note that, in both cases, t_1 must in fact satisfy $t_1 < 0$, as there are no branch points on $\gamma \equiv \hat{\gamma}|_{[0, r]}$. So in either case, we have obtained a $t_1 < 0$ with the property that $\hat{\gamma}(t_1)$ is branching. Note that setting $\epsilon' := -t_1$, we have by construction that $\hat{\gamma}|_{[-\epsilon', 0]}$ coincides with the distance minimizer $(\eta_1|_{[0, \epsilon']})^{-1}$.

By induction, assume that we have already chosen t_1, \dots, t_{i-1} , and set $\epsilon_i = -\frac{1}{2}t_{i-1} > 0$. Let x_i be a branch point satisfying $0 < d(p, x_i) < \epsilon_i$, and let η_i be a

distance minimizer from p to x_i . Proceeding as in the last paragraph, we let $-t_i$ be the largest positive number so that $\eta_i|_{[0, -t_i]} \equiv (\hat{\gamma}|_{[t_i, 0]})^{-1}$. Then $\hat{\gamma}(t_i)$ is the desired branching point. This defines our sequence $\{t_i\}_{i \in \mathbb{N}}$. Moreover, our inductive step ensures that each t_i satisfies $|t_i| < \frac{1}{2}|t_{i-1}|$, hence the increasing sequence limits to zero. \square

Corollary 3.22. *Suppose X is a π_1 -convex compact 1-dimensional geodesic metric space, and assume $\mathcal{B}(X) \neq \emptyset$ (so X is not homeomorphic to S^1). Let $X \setminus \overline{\mathcal{B}(X)} = \coprod_{i \in I} W_i$ be the decomposition of $X \setminus \overline{\mathcal{B}(X)}$ into connected components, and denote by $\rho : \coprod_{i \in I} \overline{\partial W'_i} \rightarrow \overline{\mathcal{B}(X)}$ the attaching map from the metric completion of the W_i to the set $\mathcal{B}(X)$. Then for any point $x \in \overline{\mathcal{B}(X)} \setminus \mathcal{B}(X)$, we have that $|\rho^{-1}(x)| \leq 1$.*

Proof. If there is a point $x \in \overline{\mathcal{B}(X)} \setminus \mathcal{B}(X)$ with $|\rho^{-1}(x)| \geq 2$, then one can concatenate the geodesic path γ traveling along one of the incident W_i with a small path along the incident W_j (note that $i = j$ could *a priori* happen, if both endpoints of W_i are attached to the same point). The resulting extension $\hat{\gamma}$ contains no branch points, contradicting Lemma 3.21. \square

This gives us a fairly good picture of how a π_1 -convex compact 1-dimensional geodesic metric space is built. Specifically, such an X consists of:

- its set of branch points $\mathcal{B}(X)$, which are points where there are ≥ 3 germs of geodesic paths emanating from the point (see Definition 3.14),
- the set of points $\overline{\mathcal{B}(X)} \setminus \mathcal{B}(X)$, consisting of points which are not branching, but are limits of branch points (see Fig. 2 for an illustration of such a point),
- a countable collection of closed intervals W'_i , whose lengths shrink to zero, each of which is attached to $\overline{\mathcal{B}(X)}$ along its endpoints (see Lemma 3.19 and Proposition 3.20), and
- each point in $\overline{\mathcal{B}(X)} \setminus \mathcal{B}(X)$ has at most *one* W'_i attached to it (see Corollary 3.22).

Remark 3.23. A somewhat similar structure theorem in the context of Peano continua was obtained by Conner and Meilstrup. In [8, Theorem 3.1] they show that every Peano continuum X is *homotopy equivalent* to a Peano continuum Y which is *arc reduced*. In the 1-dimensional setting, this means that either Y is a finite bouquet of circles, or contains a specific subset $B(Y)$ with $Y \setminus B(Y)$ a disjoint union of a null sequence of open arcs (compare with our Lemma 3.19 and Proposition 3.20).

We work in the more restrictive setting of π_1 -convex compact 1-dimensional geodesic metric spaces, but obtain a structure theorem on the actual space X . Their work applies to a more general setting, but the cost is that they replace the space X by a homotopy equivalent Y , and show Y has a certain structure theory.

4. Marked Length Spectrum Rigidity

This section is devoted to proving our *Main Theorem*. Let us recall the general setup. We are given two compact 1-dimensional geodesic spaces X_1, X_2 , and an

isomorphism $\phi : \pi_1(X_1) \rightarrow \pi_1(X_2)$ which preserves the marked length spectrum. Since the spaces X_i are by hypothesis not contractible, we know that $\text{Conv}(X_i) \neq \emptyset$ (see Corollary 3.9). We want to conclude that $\text{Conv}(X_1)$ is isometric to $\text{Conv}(X_2)$. From the analysis in the last section, we know that the inclusions $j_i : \text{Conv}(X_i) \hookrightarrow X_i$ induce isomorphisms $(j_i)_*$ on π_1 (see Corollary 3.11). It follows from the structure theory of these spaces that composing the isomorphism $(j_i)_*$ with the length function l_i gives the length function on the space $\text{Conv}(X_i)$. We have established that, under the hypotheses of the *Main Theorem*, we actually have an induced isomorphism $\phi : \pi_1(\text{Conv}(X_1)) \rightarrow \pi_1(\text{Conv}(X_2))$, preserving the marked length spectrum of the π_1 -convex spaces $\text{Conv}(X_i)$. Thus the *Main Theorem* will follow immediately from the following special case:

Theorem 4.1. *Let X_1, X_2 be a pair of compact, geodesic spaces of topological dimension one, and assume each X_i is π_1 -convex (i.e. $X_i = \text{Conv}(X_i)$). Assume the two spaces have the same marked length spectrum, that is to say, there exists an isomorphism $\phi : \pi_1(X_1) \rightarrow \pi_1(X_2)$ satisfying $l_2 \circ \phi = l_1$. Then X_1 is isometric to X_2 , and the isometry induces (up to change of basepoints) the isomorphism ϕ on $\pi_1(X_i)$.*

The rest of this section will be devoted to establishing Theorem 4.1. We start by introducing a few definitions and proving an important lemma.

Definition 4.2. Given a geodesic path \mathbf{p} in a 1-dimensional geodesic space X joining points p to q , we say that a pair of geodesic loops γ_1, γ_2 based at p and parametrized by arclength are **\mathbf{p} -distinguishing** provided that $\gamma_1|_{[0, l(\mathbf{p})]} \equiv \mathbf{p} \equiv \gamma_2|_{[0, l(\mathbf{p})]}$, where $l(\mathbf{p})$ is the length of the path \mathbf{p} (in particular, the geodesic loops start out by respecting the orientation on \mathbf{p}). Furthermore, we require that $[0, l(\mathbf{p})]$ be a *maximal* subinterval (with respect to inclusion) on which the loops γ_1 and γ_2 coincide.

If a geodesic path \mathbf{p} has a pair of **\mathbf{p} -distinguishing** geodesic loops, then we say the geodesic path \mathbf{p} is *distinguished*. The collection of distinguished paths inside a 1-dimensional geodesic space X will be denoted by $\mathcal{D}(X)$.

The importance of **\mathbf{p} -distinguishing** loops lies in the fact that, if γ_1 and γ_2 are **\mathbf{p} distinguishing**, and if we use an overline to denote the geodesic loop freely homotopic to a given loop, then we automatically have (using Lemma 2.16):

$$l(\overline{\gamma_2 * \gamma_1^{-1}}) = l(\gamma_1) + l(\gamma_2) - 2l(\mathbf{p}).$$

In particular, since the concatenated loop represents the product of the elements corresponding to γ_i in $\pi_1(X, \mathbf{p}(0))$, we see that the length of the geodesic path \mathbf{p} can be recovered from the marked length spectrum. It is also easy to verify that the endpoints of the path \mathbf{p} are branching (the germs of the paths \mathbf{p} and the

appropriately oriented γ_1, γ_2 will all be distinct), showing that:

Lemma 4.3. *If $\mathbf{p} \in \mathcal{D}(X)$ joins the points $p, q \in X$, then both endpoints are branch points, i.e. $p, q \in \mathcal{B}(X)$.*

We do not know whether or not the collection $\mathcal{D}(X)$ of distinguished geodesics coincides with the set of all geodesics whose endpoints are branch points. Our next result aims at showing that geodesics which are injective near their endpoints are indeed distinguished.

Lemma 4.4. *Let X be a π_1 -convex compact 1-dimensional geodesic metric space, and let $\mathbf{p} : [0, L] \rightarrow X$ be a geodesic path joining a pair of branch points lying in $X = \text{Conv}(X)$. Assume that, for ϵ_0 small enough, the set $[0, \epsilon_0] \cup [L - \epsilon_0, L]$ lies in the set of injectivity of the map \mathbf{p} , i.e. for any t in this set, $\mathbf{p}^{-1}(\mathbf{p}(t)) = \{t\}$. Then there exists a pair of \mathbf{p} -distinguishing geodesic loops, i.e. $\mathbf{p} \in \mathcal{D}(X)$.*

Proof. Since our path \mathbf{p} joins a pair of branch points, it is easy to see that there is a pair of geodesic paths γ_1, γ_2 in X which intersect precisely in \mathbf{p} . Indeed, let us start by considering $p := \mathbf{p}(0)$, and note that since this point is branching, there exists a triple of geodesic paths $\gamma_i : [0, \epsilon] \rightarrow X$ emanating from the point p with the property that each concatenated path $\gamma_i * \gamma_j^{-1}$ is a geodesic path. Now consider the three possible concatenations $\mathbf{p} * \gamma_i^{-1}$. We claim that at least two of them have to be geodesic paths. Indeed, if not, then two of these concatenations, say $\mathbf{p} * \gamma_1^{-1}$ and $\mathbf{p} * \gamma_2^{-1}$ have to be reducible. Using Lemma 2.16, this forces \mathbf{p} , γ_1 , and γ_2 to coincide in a small enough interval $[0, \epsilon']$. But this contradicts the fact that $\gamma_2 * \gamma_1^{-1}$ is reduced. So we can extend \mathbf{p} past p in two distinct ways, and still have a reduced path. Similarly, we can extend \mathbf{p} past the point $q := \mathbf{p}(L)$ in two distinct ways, and still have a reduced path. This gives us a pair of geodesic segments which are distinct, then come together and agree precisely along \mathbf{p} , and then separate again.

Shrinking the two geodesic segments if need be, we can assume that they are defined on $[-\epsilon, L + \epsilon]$, and that the geodesic \mathbf{p} corresponds to the image of $[0, L]$ in both geodesics. Now we claim that, perhaps by further shrinking the geodesic segments γ_i , we can ensure that the reduced paths we find above are non-self-terminating and non-self-originating. Without loss of generality, assume that for our triple of geodesics γ_i , it is γ_1 and γ_2 that geodesically extend \mathbf{p} . As $\mathbf{p} * \gamma_j^{-1}$ ($j = 1, 2$) is geodesic, from Lemma 2.16 we see that there exist arbitrarily small t for which $\gamma_j(t) \neq \mathbf{p}(t)$. We require something slightly stronger, i.e. a small value of t for which $\gamma_j(t) \notin \mathbf{p}([0, L])$.

By way of contradiction, suppose this were not the case. We choose a $0 < \delta$, with the property that $\gamma_j(t) \in \mathbf{p}([0, T])$ for all $t \in [0, \delta]$. But recall that, by hypothesis, \mathbf{p} is injective on $[0, \epsilon_0]$, so there is a definite positive distance between p and the image set $\mathbf{p}([\epsilon_0, T])$. So at the cost of shrinking δ , we can in fact assume that the set $\gamma_j([0, \delta])$ has image in the set $\mathbf{p}([0, \epsilon_0])$. Then the path $\gamma_j|_{[0, \delta]}$ has image contained entirely in the embedded path \mathbf{p} , and satisfies $\gamma_j(0) = \mathbf{p}(0)$. Since both curves

γ_j and \mathbf{p} are geodesics parametrized by arclength, this forces $\gamma_j|_{[0,\delta]} \equiv \mathbf{p}|_{[0,\delta]}$, contradicting the fact that $\mathbf{p} * \gamma_j^{-1}$ is irreducible. Thus we find arbitrarily small values of t satisfying $\gamma_j(t) \notin \mathbf{p}([0, T])$; let δ_0 be such a value. The curve γ_j might not be an embedded path, so conceivably we could have points $t \in [0, \delta_0)$ with the property that $\gamma_j(t) = \gamma_j(\delta_0)$. But the set of such values forms a closed subset of $[0, \delta_0)$, which is bounded away from zero. Hence, there exists a smallest $\delta \in (0, \delta_0]$ with the property that $\gamma_j(\delta) = \gamma_j(\delta_0)$. Then by the choice of δ , we have $\gamma_j(\delta) \notin \gamma_j([0, \delta))$. We conclude that the concatenation $\mathbf{p} * (\gamma_j|_{[0,\delta]})^{-1}$ is non-self-originating.

We can run a symmetric argument at the other endpoint $q = \mathbf{p}(L)$ of the path \mathbf{p} . This yields a short geodesic $\gamma'_j|_{[0,\delta']}$ originating at q , with the property that the concatenation $\gamma'_j|_{[0,\delta']} * \mathbf{p}$ is irreducible, and non-self-terminating. Finally, consider the concatenation $\gamma'_j|_{[0,\delta']} * \mathbf{p} * (\gamma_j|_{[0,\delta]})^{-1}$. In view of our discussion above, the only way this concatenation could fail to be non-self-originating (respectively, non-self-terminating) is if $\gamma_j(\delta) \in \gamma'_j|_{[0,\delta']}$ (resp. $\gamma'_j(\delta') \in \gamma_j|_{[0,\delta]}$). But recall that δ, δ' could be chosen arbitrarily close to zero. By the injectivity assumption, the two endpoints p, q are at a positive distance $d(p, q) > 0$ apart. Now choose δ, δ' small enough to satisfy $\delta + \delta' < d(p, q)$. We verify the condition for non-self-originating: if $\gamma_j(\delta) \in \gamma'_j|_{[0,\delta']}$, then concatenating γ_j with a subpath of γ'_j gives us a path of length at most $\delta + \delta' < d(p, q)$ joining the points p and q , a contradiction. The condition for non-self-terminating is completely analogous. This confirms that the concatenations $\gamma'_j|_{[0,\delta']} * \mathbf{p} * (\gamma_j|_{[0,\delta]})^{-1}$ (for both $j = 1, 2$) are a pair of geodesic paths which are both non-self-originating and non-self-terminating.

To finish, we want to extend these geodesic paths to closed geodesic loops. But that is precisely what Lemma 3.2 guarantees (the hypotheses of the lemma are satisfied because X is π_1 -convex). Hence we obtain a pair of \mathbf{p} -distinguishing geodesic loops, and we are done. □

Clearly, distance minimizers are geodesic paths which are *globally* injective, hence satisfy the conditions of Lemma 4.4. This immediately implies:

Corollary 4.5. *Let $p, q \in \mathcal{B}(X)$ be a pair of distinct branch points, and let \mathbf{p} be a distance minimizer joining p to q . Then \mathbf{p} is distinguished, i.e. $\mathbf{p} \in \mathcal{D}(X)$.*

Note that in Lemma 4.4, the local injectivity condition in particular forces the endpoints of the curve \mathbf{p} to be distinct. Next we consider geodesic loops based at a branch point, and provide a condition for them to be distinguished.

Lemma 4.6. *Let X be a π_1 -convex compact 1-dimensional geodesic metric space, and let $\mathbf{p} : [0, L] \rightarrow X$ be a geodesic path satisfying $\mathbf{p}(0) = \mathbf{p}(L) = p \in \mathcal{B}(X)$. Viewing \mathbf{p} instead as a map $(S^1, *) \rightarrow (X, p)$, assume that there exists an ϵ_0 so that the ϵ_0 -neighborhood of the basepoint $*$ lies in the set of injectivity of the map \mathbf{p} (in the sense of Lemma 4.4). Then the path \mathbf{p} is distinguished, i.e. $\mathbf{p} \in \mathcal{D}(X)$.*

Proof. One loop in our distinguishing pair is \mathbf{p}^2 . For the second loop, invoke the local injectivity condition as in the proof of Lemma 4.4 to find a geodesic γ_1 leaving

p such that $\gamma_1 * \mathbf{p}$ is non-self-terminating. By Lemma 3.2 we may extend it to a loop $\mathbf{c} * \gamma_1 * \mathbf{p}$ which is reduced, although it may not be cyclically reduced (as $\gamma_1 * \mathbf{p}$ is self-originating). If \mathbf{c} does not intersect \mathbf{p} in any interval properly containing p , it is cyclically reduced and forms a distinguishing pair for \mathbf{p} with \mathbf{p}^2 . If it does intersect \mathbf{p} in such an interval, let \mathbf{c}' be the sub-path of c up to its first intersection with \mathbf{p} . Then $\gamma_1^{-1} * \mathbf{c}'^{-1} * \mathbf{p} * \mathbf{c}' * \gamma_1 * \mathbf{p}$ forms a distinguishing pair for \mathbf{p} with \mathbf{p}^2 . \square

A nice consequence of the previous few results is the following:

Corollary 4.7. *Let X be a π_1 -convex compact 1-dimensional geodesic metric space, and let $p \in \mathcal{B}(X)$ be an arbitrary branch point. Then there exists a distinguished geodesic $\mathbf{p} \in \mathcal{D}(X)$ originating at p .*

Proof. If $|\mathcal{B}(X)| \geq 2$, then given any branch point $p \in \mathcal{B}(X)$, we can find a branch point $q \in \mathcal{B}(X)$ with $q \neq p$. Let \mathbf{p} be a distance minimizer from p to q , and apply Corollary 4.5 to see that this \mathbf{p} is a distinguished geodesic.

If $\mathcal{B}(X)$ consists of the single point p , then $\mathcal{B}(X) = \mathcal{B}(X) = \{p\}$, and the structure theory tells us that $X \setminus \{p\}$ consists of a countable collection of open intervals, of diameter shrinking to zero, each of which is attached to p at both endpoints (see Lemma 3.19 and Proposition 3.20). In other words, X is either a bouquet of finitely many circles (with lengths attached to each loop), or a generalized Hawaiian earring space. In either case, we can take a geodesic \mathbf{p} in X which loops through a single connected component of $X \setminus \{p\}$. Lemma 4.6 tells us \mathbf{p} is a distinguished geodesic. \square

Before proving our main proposition, we give one last definition.

Definition 4.8. Let X be a compact geodesic space of topological dimension one, and $\mathbf{p}_1, \mathbf{p}_2$ a pair of geodesics in the space, parametrized by the intervals $[a_1, b_1]$ and $[a_2, b_2]$ respectively. We say that \mathbf{p}_1 and \mathbf{p}_2 are *incident*, provided that $\mathbf{p}_1(b_1) = \mathbf{p}_2(a_2)$. We say that they are *geodesically incident* provided that, in addition to being incident, the concatenated path $\mathbf{p}_2 * \mathbf{p}_1$ is geodesic.

Proposition 4.9. *Let X_1, X_2 be a pair of π_1 -convex spaces. If they have the same marked length spectrum, then there is an isometry from the set $\mathcal{B}(X_1)$ of branch points of X_1 to the set $\mathcal{B}(X_2)$ of branch points of X_2 .*

Proof. We start out by defining a length preserving map from $\mathcal{D}(X_1)$ to $\mathcal{D}(X_2)$, where we recall that $\mathcal{D}(X)$ denotes the set of distinguished geodesics in X . Let $\mathbf{p} \in \mathcal{D}(X_1)$ be given. Then by definition, there exists a pair of \mathbf{p} -distinguishing geodesic loops; call them γ_1 and γ_2 . Without loss of generality, we can assume the base point p_1 for $\pi_1(X_1, p_1)$ is the common vertex $\gamma_i(0)$. Corresponding to the homomorphism $\Phi : \pi_1(X_1, p_1) \rightarrow \pi_1(X_2, p_2)$, we can find a pair of closed geodesic paths $\Phi(\gamma_1)$ and $\Phi(\gamma_2)$ (i.e. reduced paths, but not necessarily cyclically reduced)

based at $p_2 \in X_2$ having precisely the same lengths of minimal representatives in their free homotopy class. We will use an overline to denote the geodesic loop (i.e. cyclically reduced loop) in the free homotopy class of a loop. Observe that by our choice of γ_1, γ_2 being \mathbf{p} -distinguishing, we have that:

$$l_1(\overline{\gamma_2 * \gamma_1^{-1}}) = l_1(\gamma_1) + l_1(\gamma_2) - 2l_1(\mathbf{p}).$$

Furthermore, since the isomorphism preserves the marked length spectrum, we have that $l_2(\overline{\Phi(\gamma_i)}) = l_1(\gamma_i)$ and $l_2(\overline{\Phi(\gamma_2) * \Phi(\gamma_1^{-1})}) = l_1(\overline{\gamma_2 * \gamma_1^{-1}})$. Let $\overline{\Phi(\gamma_i)} = \eta_i$, so in particular, by Corollary 2.18, we have that $\overline{\Phi(\gamma_1)} = \alpha^{-1} * \eta_1 * \alpha$ and $\overline{\Phi(\gamma_2)} = \beta^{-1} * \eta_2 * \beta$ where α, β are geodesic paths in X_2 .

Claim 1. *Using the notation from the previous paragraph, we must have: $\alpha = \mathbf{n}_2 * \beta$ or $\beta = \mathbf{n}_1 * \alpha$, where \mathbf{n}_i is a sub-path of η_i .*

Consider the path $\overline{\Phi(\gamma_2) * \Phi(\gamma_1^{-1})} = \beta^{-1} * \eta_2 * \beta * \alpha^{-1} * \eta_1^{-1} * \alpha$. Unless the concatenation $\beta * \alpha^{-1}$ completely reduces and eliminates some portion of η_i , we have the inequalities:

$$l_2(\eta_1) + l_2(\eta_2) < l_2(\overline{\Phi(\gamma_2) * \Phi(\gamma_1^{-1})}) = l_1(\overline{\gamma_2 * \gamma_1^{-1}}) < l_1(\gamma_1) + l_1(\gamma_2).$$

But we have by the marked length spectrum being preserved, and the definition of η_i , that $l_2(\eta_i) = l_2(\overline{\Phi(\gamma_i)}) = l_1(\gamma_i)$ which gives us a contradiction (see Fig. 3 for an illustration of this phenomenon). Denote by $q \in X_2$ the endpoint of whichever of α or β contains the other. We can, without loss of generality, assume that $p_2 = q$ (by taking a change of basepoint for $\pi_1(X_2)$ if necessary). So we have reduced to the image being a pair of geodesic loops η_1 and η_2 based at q .

Claim 2. *The geodesic loops η_1 and η_2 intersect in a path of length precisely $l_1(\mathbf{p})$ passing through the point q (i.e. $\eta_1([0, l_1(\mathbf{p})]) = \eta_2([0, l_1(\mathbf{p})])$), but no such relationship holds for any larger interval).*




Fig. 3. Initial segments forced to agree: minimal representative of composite curve in first picture has shorter length than corresponding one in the second picture.

In order to see this, let us assume that we can write $\eta_i = \sigma_i * \nu$, where ν is a path corresponding to the largest interval $[0, r]$ satisfying $\eta_1([0, r]) = \eta_2([0, r])$, and σ_i is the path $\eta_i([r, l_1(\gamma_i)])$ (in other words, ν is the longest path along which the two image curves agree, and σ_i is the rest of the respective curves). We claim that $l_2(\nu) = l_1(\mathbf{p})$.

By our choice of γ_1, γ_2 being \mathbf{p} -distinguishing, we have the relation:

$$2l_1(\mathbf{p}) = l_1(\gamma_1) + l_1(\gamma_2) - l_1(\overline{\gamma_2 * \gamma_1^{-1}}). \tag{1}$$

Since we have that η_i are the geodesic loops in the free homotopy class of $\Phi(\gamma_i)$, and as our isomorphism preserves lengths, we have that:

$$l_2(\eta_i) = l_1(\gamma_i). \tag{2}$$

Furthermore, the composite $\overline{\gamma_2 * \gamma_1^{-1}}$ corresponds to the composite $\eta_2 * \eta_1^{-1}$, which forces the equality $l_1(\overline{\gamma_2 * \gamma_1^{-1}}) = l_2(\eta_2 * \eta_1^{-1})$, and the latter is freely homotopic to the geodesic loop $\sigma_2 * \sigma_1^{-1}$. This gives us that:

$$l_1(\overline{\gamma_1 * \gamma_2^{-1}}) = l_2(\overline{\eta_1 * \eta_2^{-1}}) = l_2(\eta_1) + l_2(\eta_2) - 2l_2(\nu). \tag{3}$$

Substituting Eqs. (2) and (3) into Eq. (1), we obtain $2l_1(\mathbf{p}) = 2l_2(\nu)$ which immediately gives us the desired equality.

We denote the path in X_2 identified in this way by $\Phi^{(\gamma_1, \gamma_2)}\mathbf{p}$, in order to emphasize the dependence on the pair of \mathbf{p} -distinguishing loops (γ_1, γ_2) . Note that we clearly have that $\Phi^{(\gamma_1, \gamma_2)}\mathbf{p}$ lies in $\mathcal{D}(X_2)$, as the pair of loops (η_1, η_2) are $(\Phi^{(\gamma_1, \gamma_2)}\mathbf{p})$ -distinguishing.

Claim 3. *The path $\Phi^{(\gamma_1, \gamma_2)}\mathbf{p}$ is independent of the choice of \mathbf{p} -distinguishing loops.*

We have two possibilities, one of which is immediate: let $\gamma_1, \gamma_2, \gamma_3$ be geodesics based at p_1 an endpoint of \mathbf{p} which are pairwise \mathbf{p} -distinguishing. Then the three image geodesic loops $\overline{\Phi(\gamma_1)}, \overline{\Phi(\gamma_2)}, \overline{\Phi(\gamma_3)}$ are all based at q , and pairwise have the property that $\overline{\Phi(\gamma_i)}|_{[0, l_1(\mathbf{p})]} = \overline{\Phi(\gamma_j)}|_{[0, l_1(\mathbf{p})]}$. It is now immediate that all three of $\Phi^{(\gamma_i, \gamma_j)}\mathbf{p}$ must coincide.

The other possibility to account for occurs if we have two distinct pairs (γ_1, γ_2) and (η_1, η_2) of \mathbf{p} -distinguishing geodesics, but none of the pairs (γ_i, η_j) are \mathbf{p} -distinguishing. Since all four geodesics pass through \mathbf{p} , this means that the intersections $\mathbf{p}_{i,j} := \gamma_i \cap \eta_j$ are all geodesic segments which extend the original \mathbf{p} (see Fig. 4 for an illustration of two such pairs near the geodesic). In fact, this immediately forces the geodesic loops to have a local picture near \mathbf{p} as in Fig. 4.

Consider first $\overline{\Phi(\eta_1) * \Phi(\gamma_1^{-1})}$, the loop used as in Claim 2 to find $\Phi^{(\gamma_1, \eta_1)}\mathbf{p}_{1,1}$. By Claim 1 (after fixing the basepoint for $\pi_1(X_2)$ to be the initial point of $\overline{\Phi(\gamma_1)}$) we must have $\overline{\Phi(\eta_1)} = \alpha^{-1} * \hat{\eta}_1 * \alpha$ for a geodesic loop $\hat{\eta}_1$ and a geodesic α which lies along $\overline{\Phi(\gamma_1)}$. By Claim 2, $\hat{\eta}_1$ intersects $\overline{\Phi(\gamma_1)}$ in a segment of length $l(\mathbf{p}_{1,1})$; note that if α is non-trivial, $\hat{\eta}_1$ must begin by following $\overline{\Phi(\gamma_1)}$ for a distance precisely $l(\mathbf{p}_{1,1})$.




Fig. 4. Two pairs of \mathbf{p} -distinguishing curves with no cross-pair \mathbf{p} -distinguishing.

Now undertake the same considerations for $\overline{\Phi(\eta_1) * \Phi(\gamma_2^{-1})}$. We can maintain the same basepoint, and we know by Claim 1 that α lies along $\overline{\Phi(\gamma_2)}$ as well, hence is a sub-path of $\overline{\Phi(\gamma_1, \gamma_2)} \mathbf{p}$. If α is non-trivial, $\hat{\eta}_1$ must begin by following $\overline{\Phi(\gamma_2)}$ for a distance precisely $l(\mathbf{p}_{2,1})$. This, however, contradicts the corresponding fact noted at the end of the previous paragraph, as $\mathbf{p}_{1,1}$ and $\mathbf{p}_{2,1}$ are both strictly longer than $\overline{\Phi(\gamma_1, \gamma_2)} \mathbf{p}$ and $\overline{\Phi(\gamma_1)}$ branches from $\overline{\Phi(\gamma_2)}$ after traversing $\overline{\Phi(\gamma_1, \gamma_2)} \mathbf{p}$. (See Fig. 5.)

We conclude then that α is trivial, i.e. that $\overline{\Phi(\eta_1)}$ is a geodesic loop based at p_2 . Because of the branching of $\overline{\Phi(\gamma_1)}$ and $\overline{\Phi(\gamma_2)}$ at the end points of $\overline{\Phi(\gamma_1, \gamma_2)} \mathbf{p}$, $\overline{\Phi(\gamma_1, \eta_1)} \mathbf{p}_{1,1}$ and $\overline{\Phi(\gamma_2, \eta_1)} \mathbf{p}_{2,1}$ must be geodesic segments that extend $\overline{\Phi(\gamma_1, \gamma_2)} \mathbf{p}$ on opposite sides. Applying the same arguments with η_2 replacing η_1 we see that $\overline{\Phi(\eta_1)}$ and $\overline{\Phi(\eta_2)}$ agree at least along $\overline{\Phi(\gamma_1, \gamma_2)} \mathbf{p}$. But by Claim 2, this is the largest segment they can agree along, hence $\overline{\Phi(\gamma_1, \gamma_2)} \mathbf{p} = \overline{\Phi(\eta_1, \eta_2)} \mathbf{p}$, establishing Claim 3.




Fig. 5. An arrangement ruled out by the proof of Claim 3.

We have established that the map $\Phi^{(\gamma_1, \gamma_2)} \mathbf{p}$ depends solely on the distinguished path \mathbf{p} , and not on the choice of \mathbf{p} -distinguishing loops (γ_1, γ_2) . We can now suppress the superscript, and simply write $\Phi \mathbf{p}$ for the image path. This gives us a well-defined map $\Phi : \mathcal{D}(X_1) \rightarrow \mathcal{D}(X_2)$.

Claim 4. *The map $\Phi : \mathcal{D}(X_1) \rightarrow \mathcal{D}(X_2)$ is bijective.*

Applying the same procedure to the inverse group homomorphism $\pi_1(X_2) \rightarrow \pi_1(X_1)$ yields a corresponding map $\Psi : \mathcal{D}(X_2) \rightarrow \mathcal{D}(X_1)$. We verify that $\Psi \circ \Phi$ is the identity map on $\mathcal{D}(X_1)$. Let $\mathbf{p} \in \mathcal{D}(X_1)$ be a distinguished path, and (γ_1, γ_2) a pair of \mathbf{p} -distinguishing loops. From our construction, the image path $\Phi \mathbf{p} \in \mathcal{D}(X_2)$ naturally comes equipped with a pair of $\Phi \mathbf{p}$ -distinguishing loops (η_1, η_2) . Recall from the argument in Claim 2 how each η_i is obtained: we start with the corresponding γ_i , viewed as an element in $\pi_1(X_1, p)$, and use the isomorphism between fundamental groups to obtain a corresponding element in $\pi_1(X_2, q)$ (where for simplicity p, q are taken to be the initial points of \mathbf{p} and $\Phi \mathbf{p}$ respectively). Then η_i is the cyclically reduced loop in the free homotopy class represented by the image element in $\pi_1(X_2, q)$.

Reversing the argument, we see that if we start with $\mathbf{q} := \Phi \mathbf{p} \in \mathcal{D}(X_2)$ and pick (η_1, η_2) as the pair of \mathbf{q} -distinguishing loops, then the image $\Psi \mathbf{q}$ is identified via the pair (γ_1, γ_2) of geodesic loops, and hence coincides with \mathbf{p} . This verifies that $\Psi \circ \Phi$ is the identity on $\mathcal{D}(X_1)$, and hence that Φ is injective. But an identical argument shows that $\Phi \circ \Psi$ is the identity on $\mathcal{D}(X_2)$, establishing that Φ is also surjective.

Claim 5. *Let $\mathbf{p}_1, \mathbf{p}_2 \in \mathcal{D}(X_1)$ be a pair of geodesically incident (see Definition 4.8) distinguished paths in X_1 , meeting at a common vertex q which we will take as the basepoint for $\pi_1(X_1)$. Then the corresponding pair of geodesic paths $\Phi \mathbf{p}_1$ and $\Phi \mathbf{p}_2$ is also geodesically incident.*

Without loss of generality (by well-definedness of our map), we can assume that one of the closed loops used to find $\Phi \mathbf{p}_1, \Phi \mathbf{p}_2$ passes through both \mathbf{p}_1 and \mathbf{p}_2 , hence can be used as one of both pairs of \mathbf{p}_i -distinguishing loops. We refer to Fig. 6 (top left) to illustrate our situation. In the top left figure, we have a pair of geodesically incident paths, with the big geodesic representing the common loop, and the two smaller ones intersecting the large one in \mathbf{p}_1 and \mathbf{p}_2 respectively. The paths \mathbf{p}_i are oriented counterclockwise along the common loop γ , so that \mathbf{p}_2 precedes \mathbf{p}_1 along γ .

Now consider the image loops (see Fig. 6, remaining three pictures). If the resulting curves are not incident, we have that the two geodesic loops $\overline{\Phi(\gamma_1)}$ and $\overline{\Phi(\gamma_2)}$, which must intersect $\overline{\Phi(\gamma)}$ in geodesic segments of length $l_1(\mathbf{p}_1), l_1(\mathbf{p}_2)$ respectively, have an intersection which does not represent incident subpaths of the geodesic loop $\overline{\Phi(\gamma)}$. We have three possible cases, which we label (a), (b), (c).

Case (a). First, the intersections with γ might be entirely disjoint (as in Fig. 6, top right). Following the construction of Claim 1, we may assume the basepoint




Fig. 6. Incidence relations are preserved: the three possible cases.

p_2 for $\pi_1(X_2)$ is on the intersection of $\overline{\Phi(\gamma)}$ and $\overline{\Phi(\gamma_1)}$. Then $\Phi(\gamma_2)$ has the form $\delta^{-1} * \overline{\Phi(\gamma_2)} * \delta$ for some geodesic starting at p_2 . Consider $\overline{\Phi(\gamma_2 * \gamma^{-1})}$. Using the fact that the marked length spectrum is preserved, we have the inequalities:

$$l_2(\overline{\Phi(\gamma_2 * \gamma^{-1})}) = l_1(\gamma_2 * \gamma^{-1}) < l_1(\gamma) + l_1(\gamma_2) = l_2(\overline{\Phi(\gamma)}) + l_2(\overline{\Phi(\gamma_2)}).$$

This forces δ to lie entirely along $\overline{\Phi(\gamma)}$. Let d denote the distance along δ between $\Phi_{\mathbf{p}_1}$ and $\Phi_{\mathbf{p}_2}$ and note that we have:

$$l_2(\overline{\Phi(\gamma_1) * \Phi(\gamma_2)}) = l_2(\overline{\Phi(\gamma_1)}) + l_2(\overline{\Phi(\gamma_2)}) + 2d.$$

But observe that the geodesic loop $\overline{\gamma_1 * \gamma_2}$ has length which is bounded above by $l_1(\gamma_1) + l_1(\gamma_2)$. Combined with the fact that the isomorphism preserves the marked length spectrum, this gives us a contradiction if $d > 0$. Thus $\Phi_{\mathbf{p}_1}$ and $\Phi_{\mathbf{p}_2}$ are adjacent on $\overline{\Phi(\gamma)}$.

Finally, suppose their order is reversed, i.e. $\Phi_{\mathbf{p}_1}$ precedes $\Phi_{\mathbf{p}_2}$ when following $\overline{\Phi(\gamma)}$. In X_1 we have that

$$l_1(\gamma_2 * \gamma_1 * \gamma^{-1}) = l_1(\gamma) + l_1(\overline{\gamma_1 * \gamma_2}) - 2l_1(\mathbf{p}_1) - 2l_1(\mathbf{p}_2).$$

Whereas in X_2 , the corresponding loop satisfies

$$l_2(\overline{\Phi(\gamma_2) * \Phi(\gamma_1) * \Phi(\gamma^{-1})}) = l_2(\overline{\Phi(\gamma)}) + l_2(\overline{\Phi(\gamma_1)}) + l_2(\overline{\Phi(\gamma_2)}),$$

a strictly greater length, providing a contradiction. We conclude that the paths $\Phi_{\mathbf{p}_1}$ and $\Phi_{\mathbf{p}_2}$ are geodesically incident in X_2 in the same order that they are in X_1 .

Case (b). The second possibility is that $\overline{\Phi(\gamma_1)}$ and $\overline{\Phi(\gamma_2)}$ intersect in a subinterval of the geodesic loop $\overline{\Phi(\gamma)}$ of length d (as in Fig. 6, bottom left). As above, we may place the basepoint for $\pi_1(X_2)$ on $\overline{\Phi(\gamma_1)}$ and write $\Phi(\gamma_2) = \alpha^{-1} * \overline{\Phi(\gamma_2)} * \alpha$ with α a

geodesic along $\overline{\Phi(\gamma)}$. We consider the geodesic loop corresponding to $\gamma_2^{-1} * \gamma * \gamma_1^{-1}$, and observe that it has length:

$$l_1(\overline{\gamma_2^{-1} * \gamma * \gamma_1^{-1}}) = l_1(\gamma) + l_1(\overline{\gamma_2 * \gamma_1}) - 2l_1(\mathbf{p}_1) - 2l_1(\mathbf{p}_2). \tag{4}$$

Looking at the corresponding geodesic loop $\overline{\Phi(\gamma_2^{-1}) * \Phi(\gamma) * \Phi(\gamma_1^{-1})}$ in the image, we find that it has length:

$$l_2(\Phi(\gamma_1)) + l_2(\Phi(\gamma_2)) + l_2(\Phi(\gamma)) - 2l_2(\Phi\mathbf{p}_1) - 2l_2(\Phi\mathbf{p}_2) + 2d$$

if $\overline{\Phi(\gamma_2)}$ lies to the right of $\overline{\Phi(\gamma_1)}$ or

$$l_2(\Phi(\gamma_1)) + l_2(\Phi(\gamma_2)) + l_2(\Phi(\gamma)) - 2d$$

if $\overline{\Phi(\gamma_2)}$ lies to the left of $\overline{\Phi(\gamma_1)}$. In either case, using that the isomorphism preserves the marked length spectrum, and comparing with equation (4) we get a contradiction (in the second case, because $d < l_1(\mathbf{p}_1) + l_1(\mathbf{p}_2)$).

Case (c). Finally, the third possibility is that one of $\overline{\Phi\mathbf{p}_i}$ lies entirely within the other (Fig. 6, bottom right). First, assume that $\overline{\Phi(\gamma_2)}$ is the small inner loop, while $\overline{\Phi(\gamma_1)}$ is the outer loop (so in particular $\Phi\mathbf{p}_2$ is a subpath of $\Phi\mathbf{p}_1$). If we let $0 \leq d < l_2(\Phi\mathbf{p}_1)$ be the distance between the left endpoints of $\Phi\mathbf{p}_i$, a simple calculation will show that:

$$l_2(\overline{\Phi(\gamma_2^{-1}) * \Phi(\gamma) * \Phi(\gamma_1^{-1})}) = l_2(\overline{\Phi(\gamma_1)}) + l_2(\overline{\Phi(\gamma)}) + l_2(\overline{\Phi(\gamma_2)}) - 2d - 2l_2(\Phi\mathbf{p}_2)$$

which we can compare with the expression in Eq. (4) to again obtain a contradiction. If $\overline{\Phi(\gamma_1)}$ is the inner loop, let $0 \leq d < l_2(\Phi\mathbf{p}_2)$ be the distance between the right endpoints of $\Phi(\mathbf{p}_i)$. We calculate:

$$l_2(\overline{\Phi(\gamma_2^{-1}) * \Phi(\gamma) * \Phi(\gamma_1^{-1})}) = l_2(\overline{\Phi(\gamma_1)}) + l_2(\overline{\Phi(\gamma)}) + l_2(\overline{\Phi(\gamma_2)}) - 2d - 2l_2(\Phi\mathbf{p}_1)$$

providing the same contradiction.

This gives us that the image paths $\Phi\mathbf{p}_1$ and $\Phi\mathbf{p}_2$ are subpaths of the geodesic loop γ which agree at one endpoint, but not in any larger neighborhood of the endpoint. Since γ is cyclically reduced, this immediately forces the concatenation $\Phi\mathbf{p}_1 * \Phi\mathbf{p}_2$ to be a geodesic path, hence geodesically incident paths map to geodesically incident paths. It is now clear that if $\mathbf{p} = \mathbf{q}_2 * \mathbf{q}_1$ is a geodesic path written as a concatenation of subpaths (all in $\mathcal{D}(X_1)$), then $\Phi\mathbf{p} = \Phi\mathbf{q}_2 * \Phi\mathbf{q}_1$ (since the \mathbf{q}_i are geodesically incident).

Claim 6. *Let \mathbf{p}_1 and \mathbf{p}_2 be a pair of incident geodesic paths in $\mathcal{D}(X_1)$, meeting at a common vertex q which we will take as the basepoint for $\pi_1(X)$. Then the corresponding geodesic paths $\Phi\mathbf{p}_1$ and $\Phi\mathbf{p}_2$ are also incident.*

To see this, note that if the incident paths are *geodesically* incident, we are done by the previous claim. So let us assume not. Then by Lemma 2.16, we have that

the reduced path corresponding to the concatenation $\mathbf{p}_2 * \mathbf{p}_1$ is of the form $\mathbf{r}_2 * \mathbf{r}_1$, where \mathbf{r}_i is a subpath of \mathbf{p}_i . Furthermore, $\mathbf{p}_2 = \mathbf{r}_2 * \mathbf{q}^{-1}$, while $\mathbf{p}_1 = \mathbf{q} * \mathbf{r}_1$. We now argue that the geodesics $\mathbf{q}, \mathbf{r}_1, \mathbf{r}_2$ all lie in $\mathcal{D}(X_1)$. Indeed, assume that (γ_1, γ_2) is a \mathbf{p}_1 -distinguishing pair, and (η_1, η_2) is a \mathbf{p}_2 -distinguishing pair. Then we can immediately write out distinguishing pairs as follows:

- a distinguishing pair for \mathbf{q} is given by $(\eta_1^{-1}, \overline{\gamma_1 * \eta_2^{-1}})$.
- a distinguishing pair for \mathbf{r}_1 is given by either (a) $(\gamma_1, \overline{\gamma_2 * \eta_1})$ if $\gamma_2 * \eta_1$ is not reducible as a concatenation of paths based at q (it is cyclically reducible), or (b) $(\gamma_1, \overline{\gamma_2 * \eta_2})$ if $\gamma_2 * \eta_2$ is not reducible as a concatenation of paths based at q . (Note that these are mutually exclusive options.)
- a distinguishing pair for \mathbf{r}_2 is given by either (a) $(\eta_1, \overline{\gamma_1 * \eta_2})$ if $\gamma_1 * \eta_2$ is not reducible as a concatenation of paths based at q , or (b) $(\eta_1, \overline{\gamma_2 * \eta_2})$ if $\gamma_2 * \eta_2$ is not reducible as a concatenation of paths based at q . (Again, these are mutually exclusive options.)

As usual, the overline means one should cyclically reduce. Above we have been somewhat cavalier with the starting points of the various loops. They should all be reparametrized so that they are based at the appropriate points (initial point of $\mathbf{q}, \mathbf{r}_1, \mathbf{r}_2$ respectively), by shifting the basepoint along the paths \mathbf{q}, \mathbf{r}_1 , and/or \mathbf{r}_2 as needed.

Since \mathbf{q} is geodesically incident to both $\mathbf{r}_1, \mathbf{r}_2$, the previous claim implies that $\Phi \mathbf{p}_2 = \Phi \mathbf{r}_2 * \Phi \mathbf{q}^{-1}$ and $\Phi \mathbf{p}_1 = \Phi \mathbf{q} * \Phi \mathbf{r}_1$. However, reversal of paths is preserved under the map Φ we have constructed (since we can take the same pair of distinguishing loops with reversed orientations). This immediately yields our last claim.

We can now complete the proof of the proposition. We define a map $f : \mathcal{B}(X_1) \rightarrow \mathcal{B}(X_2)$ as follows: given a point $x \in \mathcal{B}(X_1)$, we consider the subset $\mathcal{D}(x) \subset \mathcal{D}(X_1)$ consisting of all distinguished paths which originate at x . Corollary 4.7 guarantees that $\mathcal{D}(x) \neq \emptyset$. We can apply the map Φ to all the elements in $\mathcal{D}(x)$, obtaining a subset $\Phi(\mathcal{D}(x))$ of $\mathcal{D}(X_2)$. In view of Claim 6, all the distinguished paths given by the elements $\Phi(\mathcal{D}(x))$ originate at the same point in $\mathcal{B}(X_2)$. We define this point to be $f(x)$, i.e. $f(x) \in X_2$ is the unique point with the property $\Phi(\mathcal{D}(x)) \subset \mathcal{D}(f(x))$.

Next we argue that the map f is distance non-increasing. If $x \neq y$ are a pair of distinct points in $\mathcal{B}(X_1)$, we let \mathbf{p} denote a distance minimizer from x to y . Corollary 4.5 tells us that $\mathbf{p}, \mathbf{p}^{-1} \in \mathcal{D}(X_1)$. This gives us elements $\mathbf{p} \in \mathcal{D}(x)$ and $\mathbf{p}^{-1} \in \mathcal{D}(y)$, so by definition of the map f , the image path $\Phi \mathbf{p}$ is a geodesic path originating at $f(x)$ and ending at $f(y)$. But from Claim 2, we know that the map Φ preserves the length of paths. We conclude that

$$d(f(x), f(y)) \leq l_2(\Phi \mathbf{p}) = l_1(\mathbf{p}) = d(x, y).$$

Applying the same argument to the reverse isomorphism $\pi_1(X_2) \rightarrow \pi_1(X_1)$, we have the inverse map $\Psi : \mathcal{D}(X_2) \rightarrow \mathcal{D}(X_1)$. Applying the construction described above, we obtain an induced map $g : \mathcal{B}(X_2) \rightarrow \mathcal{B}(X_1)$. The argument in the previous paragraph tells us that g is also distance non-increasing, and by construction,

we have that f, g are inverse maps of each other (compare with the discussion in Claim 4). Composing the two maps, we obtain

$$d(x, y) = d(g(f(x)), g(f(y))) \leq d(f(x), f(y)) \leq d(x, y).$$

Hence all the inequalities are actually equalities, and the maps f, g are isometries. This completes the proof of Proposition 4.9. □

We now prove the main theorem of this section:

Proof of Theorem 4.1. Consider the set $\mathcal{B}(X_1)$. If this set is empty, then Proposition 4.9 forces the set $\mathcal{B}(X_2)$ to likewise be empty. Lemma 3.16 tells us that each X_i is isometric to a circle S^1 of some radius r_i . The function l_i , applied to one of the generators of $\pi_1(X_i) \cong \mathbb{Z}$ evaluates to $2\pi r_i$. Since the marked length spectrum is preserved, we conclude $2\pi r_1 = 2\pi r_2$, and the two circles are isometric.

So now we may assume that $\mathcal{B}(X_1) \neq \emptyset$. By Proposition 4.9, we already know that there is an isometry $f : \mathcal{B}(X_1) \rightarrow \mathcal{B}(X_2)$ between the sets of branch points. All we need to show is that we can extend this isometry to a global isometry. Note that since we are working with compact (hence complete) metric spaces, an isometry between the subsets $\mathcal{B}(X_i)$ extends to an isometry $f : \overline{\mathcal{B}(X_1)} \rightarrow \overline{\mathcal{B}(X_2)}$ between their closures. We are left with extending our map to points that do not lie in the closure of the branch points.

Let $p \in X_1 \setminus \overline{\mathcal{B}(X_1)}$ lie outside of the closure of the branching locus. We know p lies in a connected component W of $X_1 \setminus \overline{\mathcal{B}(X_1)}$. From Lemma 3.19, W equipped with its intrinsic geodesic structure is isometric to an open interval of some finite length r , with endpoints lying in the set $\overline{\mathcal{B}(X_1)}$. Let $x, y \in \overline{\mathcal{B}(X_1)}$ be the two points to which the interval W gets attached. There are three possibilities, according to whether (i) $x, y \in \mathcal{B}(X_1)$, (ii) exactly one of the two points x, y lies in $\overline{\mathcal{B}(X_1)} \setminus \mathcal{B}(X_1)$, or (iii) both points x, y lie in $\overline{\mathcal{B}(X_1)} \setminus \mathcal{B}(X_1)$. We present a proof that works in all three cases, though in cases (i) and (ii) it can be simplified.

First, extend W to a slightly larger, yet still embedded, geodesic W' whose endpoints x_1 and y_1 lie in $\mathcal{B}(X_1)$. (In case (i) or (ii) the extension to both or one side, respectively, can be trivial, with corresponding simplifications in the argument to follow.) By Lemma 3.21, there exist branch points x_i and y_i on W' with $x_i \rightarrow x$, $y_i \rightarrow y$. Let W_i denote the subpath of W' connecting x_i and y_i . As the W_i are embedded geodesic paths, we may use the construction of Proposition 4.9 to find corresponding paths ΦW_i in X_2 with lengths equal to $l(W_i)$ connecting $f(x_i)$ and $f(y_i)$. By continuity of f defined on $\overline{\mathcal{B}(X_1)}$, $f(x_i) \rightarrow f(x)$ and $f(y_i) \rightarrow f(y)$. In addition, it is clear from the construction of Proposition 4.9 that $\Phi W_{i+1} \subseteq \Phi W_i$ (as subpaths). Therefore, the paths ΦW_i converge to a path W^* in X_2 of length r connecting $f(x)$ and $f(y)$. By Lemma 3.17, W^* coincides with one connected component of $X_2 \setminus \overline{\mathcal{B}(X_2)}$ and we can extend the map f to W by isometrically sending W to W^* .

To ensure that this extension of f is well-defined, consider any other geodesic extension of W , together with any other choice of the points x_i, y_i . Regardless of these choices, we will still have that $f(x_i) \rightarrow f(x)$ and $f(y_i) \rightarrow f(y)$. Then Corollary 3.22 ensures that we obtain the same W^* (in the case where both endpoints are in $\mathcal{B}(X_1)$, use Claim 3 in the proof of Proposition 4.9 instead). Therefore, the extension of f is well-defined, and it is straightforward to check that the extension is still an isometry onto its image.

The map f we have constructed is an isometric embedding from Y_1 into Y_2 . But note that we can apply the same construction to ϕ^{-1} , yielding an isometric embedding from Y_2 into Y_1 . Furthermore, the composite of the two maps corresponds to the map from Y_1 to itself obtained by applying this construction to the identity isomorphism and hence must be the identity map on Y_1 . This implies the map is an isometry from Y_1 to Y_2 . By the naturality of the construction, we see that the map we constructed induces the isomorphism (up to change of basepoint) between the $\pi_1(X_i)$. This completes the proof of Theorem 4.1. \square

As discussed at the beginning of this section, our *Main Theorem* now follows immediately from the special case established in Theorem 4.1.

Acknowledgments

A preliminary version of this paper was posted by the second author on the arXiv in 2003. Many thanks are due to J. W. Cannon, who found a gap in the original argument, and provided many helpful suggestions (in particular, the idea for Proposition 3.4). The second author would also like to take this opportunity to thank the first author for resurrecting this project, which would otherwise have probably remained indefinitely in limbo.

The second author was partially supported by the NSF, under grants DMS-1207782, DMS-1510640, and by an Alfred P. Sloan Research Fellowship.

References

1. R. Alperin and H. Bass, Length functions of group actions on λ -trees, in *Combinatorial Group Theory and Topology* (Alta, Utah, 1984) Ann. of Math. Stud., Vol. 111 (Princeton Univ. Press, 1987), pp. 265–378.
2. M. Bestvina, \mathbb{R} -trees in topology, geometry, and group theory, in *Handbook of Geometric Topology* (North-Holland, 2002), pp. 55–91.
3. M. R. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Grundlehren der Mathematischen Wissenschaften, Vol. 319 (Springer-Verlag, 1999).
4. D. Burago, Y. Burago and S. Ivanov, *A Course in Metric Geometry*, Graduate Studies in Mathematics, Vol. 33 (Amer. Math. Soc., 2001).
5. J. W. Cannon and G. R. Conner, On the fundamental groups of one-dimensional spaces, *Topology Appl.* **153** (2006) 2648–2672.
6. G. R. Conner and C. Kent, Fundamental groups of locally connected subsets of the plane, arXiv:1105.0405.

7. G. R. Conner and C. Kent, Homotopy dimension of planar continua, arXiv: 1702.05194.
8. G. R. Conner and M. Meilstrup, Arc-reduced forms for Peano continua, *Topology Appl.* **159** (2012) 3538–3543.
9. G. R. Conner and M. Meilstrup, Deforestation of Peano continua and minimal deformation retracts, *Topology Appl.* **159** (2012) 3253–3262.
10. C. B. Croke, Rigidity for surfaces of nonpositive curvature, *Comment. Math. Helv.* **65** (1990) 150–169.
11. M. Culler and J. W. Morgan, Group actions on \mathbb{R} -trees, *Proc. London Math. Soc.* (3) **55** (1987) 571–604.
12. M. L. Curtis and M. K. Fort, Jr., Homotopy groups of one-dimensional spaces, *Proc. Amer. Math. Soc.* **8** (1957) 577–579.
13. M. L. Curtis and M. K. Fort, Jr., The fundamental group of one-dimensional spaces, *Proc. Amer. Math. Soc.* **10** (1959) 140–148.
14. F. Dal’Bo and I. Kim, Marked length rigidity for symmetric spaces, *Comment. Math. Helv.* **77** (2002) 399–407.
15. K. Eda, Free σ -products and fundamental groups of subspaces of the plane, *Topology Appl.* **84** (1998) 283–306.
16. K. Eda, The fundamental groups of one-dimensional spaces and spatial homomorphisms, *Topology Appl.* **123** (2002) 479–505.
17. K. Eda, Homotopy types of one-dimensional Peano continua, *Fund. Math.* **209** (2010) 27–42.
18. U. Hamenstädt, Entropy-rigidity of locally symmetric spaces of negative curvature, *Ann. of Math.* (2) **131** (1990) 35–51.
19. S. Hersonsky and F. Paulin, On the rigidity of discrete isometry groups of negatively curved spaces, *Comment. Math. Helv.* **72** (1997) 349–388.
20. W. Hurewicz and H. Wallman, *Dimension Theory*, Princeton Mathematical Series, Vol. 4 (Princeton Univ. Press, 1941).
21. T. J. Laakso, Ahlfors g -regular spaces with arbitrary $q > 1$ admitting weak Poincaré inequality, *Geom. Funct. Anal.* **10** (2000) 111–123.
22. J.-P. Otal, Le spectre marqué des longueurs des surfaces à courbure négative, *Ann. of Math.* (2) **131** (1990) 151–162.
23. K. Romeo and B. Steinhurst, Eigenmodes of the Laplacian on some Laakso spaces, *Complex Vari. Elliptic Equ.* **54** (2009) 623–637.