Triangle inequalities in path metric spaces

MICHAEL KAPOVICH

We study side-lengths of triangles in path metric spaces. We prove that unless such a space X is bounded, or quasi-isometric to $\mathbb{R}_+$ or to $\mathbb{R}$, every triple of real numbers satisfying the strict triangle inequalities, is realized by the side-lengths of a triangle in X. We construct an example of a complete path metric space quasi-isometric to $\mathbb{R}^2$ for which every degenerate triangle has one side which is shorter than a certain uniform constant.

51K05

1 Introduction

Given a metric space X define

$$K_3(X) := \{(a,b,c) \in \mathbb{R}^3_+ : \text{ there exist points } x, y, z \text{ with } d(x,y) = a,\ d(y,z) = b,\ d(z,x) = c\}.$$

Note that $K_3(\mathbb{R}^2)$ is the closed convex cone K in $\mathbb{R}^3_+$ given by the usual triangle inequalities. On the other hand, if $X = \mathbb{R}$ then $K_3(X)$ is the boundary of K since all triangles in X are degenerate. If X has finite diameter, $K_3(X)$ is a bounded set. We refer the reader to [3] and [6] for discussion of the sets $K_4(X)$.

Gromov [3, Page 18] (see also Roe [6]) raised the following question:

Question 1.1 Find *reasonable* conditions on path metric spaces X, under which $K_3(X) = K$.

It is not so difficult to see that for a path metric space X quasi-isometric to $\mathbb{R}_+$ or $\mathbb{R}$, the set $K_3(X)$ does not contain the interior of K, see Section 7. Moreover, every triangle in such X is D–degenerate for some $D < \infty$ and therefore $K_3(X)$ is contained in the D–neighborhood of ∂K.

Our main result is essentially the converse to the above observation:

Published: 2 August 2007 DOI: 10.2140/gt.2007.11.1653

Theorem 1.2 *Suppose that X is an unbounded path metric space not quasi-isometric to $\mathbb{R}_+$ or $\mathbb{R}$. Then:*

(1) *$K_3(X)$ contains the interior of the cone K.*

(2) *If, in addition, X contains arbitrary long geodesic segments, then $K_3(X) = K$.*

In particular, we obtain a complete answer to Gromov's question for geodesic metric spaces, since an unbounded geodesic metric space clearly contains arbitrarily long geodesic segments. In Section 6, we give an example of a (complete) path metric space X quasi-isometric to $\mathbb{R}^2$, for which

$$K_3(X) \neq K.$$

Therefore, Theorem 1.2 is the optimal result.

It appears that very little can be said about $K_3(X)$ for general metric spaces even under the assumption of uniform contractibility. For instance, if X is the paraboloid of revolution in $\mathbb{R}^3$ with the induced metric, then $K_3(X)$ does not contain the interior of K. The space X in this example is uniformly contractible and is not quasi-isometric to $\mathbb{R}$ and $\mathbb{R}_+$.

The proof of Theorem 1.2 is easier under the assumption that X is a proper metric space: In this case X is necessarily complete, geodesic metric space. Moreover, every unbounded sequence of geodesic segments $\overline{ox_i}$ in X yields a geodesic ray. The reader who does not care about the general path metric spaces can therefore assume that X is proper. The arguments using the ultralimits are then replaced by the Arcela–Ascoli theorem.

Below is a sketch of the proof of Theorem 1.2 under the extra assumption that X is proper. Since the second assertion of Theorem 1.2 is clear, we have to prove only the first statement. To motivate the use of *tripods* in the proof we note the following: Suppose that X is itself isometric to the tripod with infinitely long legs, i.e., three rays glued at their origins. Then it is easy to see that $K_3(X) = K$.

We define R*–tripods* $T \subset X$, as unions $\gamma \cup \mu$ of two geodesic segments $\gamma, \mu \subset X$, having the lengths $\geq R$ and $\geq 2R$ respectively, so that:

(1) $\gamma \cap \mu = o$ is the end-point of γ.

(2) o is distance $\geq R$ from the ends of μ.

(3) o is a nearest-point projection of γ to μ.

The space X is called R*–thin* if it contains no R–tripods. The space X is called *thick* if it is not R–thin for any $R < \infty$.

We break the proof of Theorem 1.2 in two parts: Theorem 1.3 and Theorem 1.4.

Theorem 1.3 *If X is thick then $K_3(X)$ contains the interior of $K_3(\mathbb{R}^2)$.*

The proof of this theorem is mostly the coarse topology. The side-lengths of triangles in X determine a continuous map

$$L\colon X^3 \to K$$

Then $K_3(X) = L(X^3)$. Given a point κ in the interior of K, we consider an R–tripod $T \subset X$ for sufficiently large R. We then restrict to triangles in X with vertices in T. We construct a 2–cycle $\Sigma \in Z_2(T^3, \mathbb{Z}_2)$ whose image under L_* determines a nontrivial element of $H_2(K \setminus \kappa, \mathbb{Z}_2)$. Since T^3 is contractible, there exists a 3–chain $\Gamma \in C_3(T^3, \mathbb{Z}_2)$ with the boundary Σ. Therefore the support of $L_*(\Gamma)$ contains the point κ, which implies that κ belongs to the image of L.

Remark Gromov observed in [3] that *uniformly contractible* metric spaces X have *large* $K_3(X)$. Although uniform contractibility is not relevant to our proof, the key argument here indeed has the coarse topology flavor.

Theorem 1.4 *If X is a thin unbounded path metric space, then X is quasi-isometric to $\mathbb{R}$ or $\mathbb{R}_+$.*

Assuming that X is thin, unbounded and is not quasi-isometric to $\mathbb{R}$ and to $\mathbb{R}_+$, we construct three diverging geodesic rays ρ_i in X, $i = 1, 2, 3$. Define $\mu_i \subset X$ to be the geodesic segment connecting $\rho_1(i)$ and $\rho_2(i)$. Take γ_i to be the shortest segment connecting $\rho_3(i)$ to μ_i. Then $\gamma_i \cup \mu_i$ is an R_i–tripod with $\lim_i R_i = \infty$, which contradicts the assumption that X is thin.

Acknowledgements During this work the author was partially supported by the NSF grants DMS-04-05180 and DMS-05-54349. Most of this work was done when the author was visiting the Max Plank Institute for Mathematics in Bonn. I am grateful to the referee for useful comments and corrections.

2 Preliminaries

Convention 2.1 All homology will be taken with the $\mathbb{Z}_2$–coefficients.

In the paper we will talk about *ends of a metric space* X. Instead of looking at the noncompact complementary components of *relatively compact open subsets* of X as it is usually done for topological spaces, we will define ends of X by considering unbounded complementary components of bounded subsets of X. With this modification, the usual definition goes through.

If x, y are points in a topological space X, we let $P(x,y)$ denote the set of continuous paths in X connecting x to y. For $\alpha \in P(x,y), \beta \in P(y,z)$ we let $\alpha * \beta \in P(x,z)$ denote the concatenation of α and β. Given a path $\alpha\colon [0,a] \to X$ we let $\overline{\alpha}$ denote the reverse path

$$\overline{\alpha}(t) = \alpha(a-t).$$

2.1 Triangles and their side-lengths

We set $K := K_3(\mathbb{R}^2)$; it is the cone in $\mathbb{R}^3$ given by

$$\{(a,b,c) : a \le b+c, b \le a+c, c \le a+b\}.$$

We metrize K by using the maximum-norm on $\mathbb{R}^3$.

By a *triangle* in a metric space X we will mean an ordered triple $\Delta = (x,y,z) \in X^3$. We will refer to the numbers $d(x,y), d(y,z), d(z,x)$ as the *side-lengths* of Δ, even though these points are not necessarily connected by geodesic segments. The sum of the side-lengths of Δ will be called the *perimeter* of Δ.

We have the continuous map

$$L\colon X^3 \to K$$

which sends the triple (x,y,z) of points in X to the triple of side-lengths

$$(d(x,y), d(y,z), d(z,x)).$$

Then $K_3(X)$ is the image of L.

Let $\epsilon \ge 0$. We say that a triple $(a,b,c) \in K$ is *ϵ–degenerate* if, after reordering if necessary the coordinates a,b,c, we obtain

$$a + \epsilon \ge b + c.$$

Therefore every ϵ–degenerate triple is within distance $\le \epsilon$ from the boundary of K. A triple which is not ϵ–degenerate is called ϵ–nondegenerate. A triangle in a metric space X whose side-lengths form an *ϵ–degenerate* triple, is called *ϵ–degenerate*. A 0–degenerate triangle is called *degenerate*.

2.2 Basic notions of metric geometry

For a subset E in a metric space X and $R < \infty$ we let $N_R(E)$ denote the metric R–neighborhood of E in X:

$$N_R(E) = \{x \in X : d(x, E) \leq R\}.$$

Definition 2.2 Given a subset E in a metric space X and $\epsilon > 0$, we define the *ϵ–nearest–point projection* $p = p_{E,\epsilon}$ as the map which sends X to the set 2^E of subsets in E:

$$y \in p(x) \Longleftrightarrow d(x, y) \leq d(x, z) + \epsilon, \quad \forall z \in E.$$

If $\epsilon = 0$, we will abbreviate $p_{E,0}$ to p_E.

2.2.1 Quasi-isometries Let X, Y be metric spaces. A map $f\colon X \to Y$ is called an (L, A)–quasi-isometric embedding (for $L \geq 1$ and $A \in \mathbb{R}$) if for every pair of points $x_1, x_2 \in X$ we have

$$L^{-1} d(x_1, x_2) - A \leq d(f(x_1), f(x_2)) \leq L d(x_1, x_2) + A.$$

A map f is called an (L, A)–*quasi-isometry* if it is an (L, A)–quasi-isometric embedding so that $N_A(f(X)) = Y$. Given an (L, A)–quasi-isometry, we have the *quasi-inverse* map

$$\bar{f}\colon Y \to X$$

which is defined by choosing for each $y \in Y$ a point $x \in X$ so that $d(f(x), y) \leq A$. The quasi-inverse map $\bar{f}$ is an $(L, 3A)$–quasi-isometry. An (L, A)–*quasi-isometric embedding* f of an interval $I \subset \mathbb{R}$ into a metric space X is called an (L, A)–*quasi-geodesic* in X. If $I = \mathbb{R}$, then f is called a *complete* quasi-geodesic.

A map $f\colon X \to Y$ is called a *quasi-isometric embedding* (resp. a *quasi-isometry*) if it is an (L, A)–quasi-isometric embedding (resp. (L, A)–quasi-isometry) for some $L \geq 1, A \in \mathbb{R}$.

Every quasi-isometric embedding $\mathbb{R}^n \to \mathbb{R}^n$ is a quasi-isometry, see for instance Kapovich–Leeb [5].

2.2.2 Geodesics and path metric spaces A *geodesic* in a metric space is an isometric embedding of an interval into X. By abusing the notation, we will identify geodesics and their images. A metric space is called *geodesic* if any two points in X can be connected by a geodesic. By abusing the notation we let $\overline{xy}$ denote a geodesic connecting x to y, even though this geodesic is not necessarily unique.

The length of a continuous curve $\gamma\colon [a,b] \to X$ in a metric space, is defined as

$$\text{length}(\gamma) = \sup \left\{ \sum_{i=1}^{n} d(\gamma(t_{i-1}), \gamma(t_i)) : a = t_0 < t_1 < \cdots < t_n = b \right\}.$$

A path γ is called *rectifiable* if $\text{length}(\gamma) < \infty$.

A metric space X is called a *path metric space* (or a *length space*) if for every pair of points $x, y \in X$ we have

$$d(x,y) = \inf\{\text{length}(\gamma) : \gamma \in P(x,y)\}.$$

We say that a curve $\gamma\colon [a,b] \to X$ is ϵ–geodesic if

$$\text{length}(\gamma) \le d(\gamma(a), \gamma(b)) + \epsilon.$$

It follows that every ϵ–geodesic is $(1,\epsilon)$–quasi-geodesic. We refer the reader to Burago–Ivanov [2] and Gromov [3] for the further details on path metric spaces.

2.3 Ultralimits

Our discussion of ultralimits of sequences of metric space will be somewhat brief, we refer the reader to Burago–Ivanov [2], Gromov [3], Kapovich [4], Kapovich–Leeb [5] and Roe [6] for the detailed definitions and discussion.

Choose a nonprincipal ultrafilter ω on $\mathbb{N}$. Suppose that we are given a sequence of pointed metric spaces (X_i, o_i), where $o_i \in X_i$. The *ultralimit*

$$(X_\omega, o_\omega) = \omega\text{–}\lim(X_i, o_i)$$

is a pointed metric space whose elements are equivalence classes x_ω of sequences $x_i \in X_i$. The distance in X_ω is the ω–limit:

$$d(x_\omega, y_\omega) = \omega\text{–}\lim d(x_i, y_i).$$

One of the key properties of ultralimits which we will use repeatedly is the following. Suppose that (Y_i, p_i) is a sequence of pointed metric spaces. Assume that we are given a sequence of (L_i, A_i)–quasi-isometric embeddings

$$f_i\colon X_i \to Y_i$$

so that $\omega\text{–}\lim d(f(o_i), p_i)) < \infty$ and

$$\omega\text{–}\lim L_i = L < \infty, \quad \omega\text{–}\lim A_i = 0.$$

Then there exists the ultralimit f_ω of the maps f_i, which is an $(L,0)$–quasi-isometric embedding

$$f_\omega\colon X_\omega \to Y_\omega.$$

In particular, if $L=1$, then f_ω is an isometric embedding.

2.3.1 Ultralimits of constant sequences of metric spaces Suppose that X is a path metric space. Consider the constant sequence $X_i = X$ for all i. If X is a proper metric space and o_i is a bounded sequence, the ultralimit X_ω is nothing but X itself. In general, however, it could be much larger. The point of taking the ultralimit is that some properties of X improve after passing to X_ω.

Lemma 2.3 *X_ω is a geodesic metric space.*

Proof Points x_ω, y_ω in X_ω are represented by sequences $(x_i), (y_i)$ in X. For each i choose a $\frac{1}{i}$–geodesic curve γ_i in X connecting x_i to y_i. Then

$$\gamma_\omega := \omega\text{–}\lim \gamma_i$$

is a geodesic connecting x_ω to y_ω. □

Similarly, every sequence of $\frac{1}{i}$–geodesic segments $\overline{yx_i}$ in X satisfying

$$\omega\text{–}\lim d(y, x_i) = \infty,$$

yields a geodesic ray γ_ω in X_ω emanating from $y_\omega = (y)$.

If $o_i \in X$ is a bounded sequence, then we have a natural (diagonal) isometric embedding $X \to X_\omega$, given by the map which sends $x \in X$ to the constant sequence (x).

Lemma 2.4 *For every geodesic segment $\gamma_\omega = \overline{x_\omega y_\omega}$ in X_ω there exists a sequence of $1/i$–geodesics $\gamma_i \subset X_i$, so that*

$$\omega\text{-}\lim \gamma_i = \gamma_\omega.$$

Proof Subdivide the segment γ_ω into n equal subsegments

$$\overline{z_{\omega,j} z_{\omega,j+1}}, \quad j = 1, \dots, n,$$

where $x_\omega = z_{\omega,1}$, $y_\omega = z_{\omega,n+1}$. Then the points $z_{\omega,j}$ are represented by sequences $(z_{k,j}) \in X$. It follows that for ω–all k, we have

$$\left| \sum_{j=1}^{n} d(z_{k,j}, z_{k,j+1}) - d(x_k, y_k) \right| < \frac{1}{2i}.$$

Connect the points $z_{k,j}, z_{k,j+1}$ by $\frac{1}{2i}$–geodesic segments $\alpha_{k,j}$. Then the concatenation

$$\alpha_n = \alpha_{k,1} * \cdots * \alpha_{k,n}$$

is an $\frac{1}{i}$–geodesic connecting x_k and y_k, where

$$x_\omega = (x_k), \quad y_\omega = (y_k).$$

It is clear from the construction, that, if given i we choose sufficiently large $n = n(i)$, then

$$\omega\text{–}\lim \alpha_{n(i)} = \gamma.$$

Therefore we take $\gamma_i := \alpha_{n(i)}$. □

2.4 Tripods

Our next goal is to define *tripods* in X, which will be our main technical tool. Suppose that x, y, z, o are points in X and μ is an ϵ–geodesic segment connecting x to y, so that $o \in \mu$ and $o \in p_{\mu,\epsilon}(z)$. Then the path μ is the concatenation $\alpha \cup \beta$, where α, β are ϵ–geodesics connecting x, y to o. Let γ be an ϵ–geodesic connecting z to o.

Definition 2.5 (1) We refer to $\alpha \cup \beta \cup \gamma$ as a *tripod* T with the vertices x, y, z, legs α, β, γ, and the center o.

(2) Suppose that the length of α, β, γ is at least R. Then we refer to the tripod T as (R,ϵ)–tripod. An $(R,0)$–tripod will be called simply an *R–tripod*.

The reader who prefers to work with proper geodesic metric spaces can safely assume that $\epsilon = 0$ in the above definition and thus T is a geodesic tripod.

Definition 2.6 Let $R \in [0,\infty), \epsilon \in [0,\infty)$. A metric space is called *(R,ϵ)–thin* if it contains no (R,ϵ)–tripods. We will refer to $(R,0)$–thin spaces as R–thin. A metric space which is not (R,ϵ)–thin for any $R < \infty$, $\epsilon > 0$ is called *thick*.

Therefore, a path metric space is thick if and only if it contains a sequence of (R_i,ϵ_i)–tripods with

$$\lim_i R_i = \infty, \quad \lim_i \epsilon_i = 0.$$

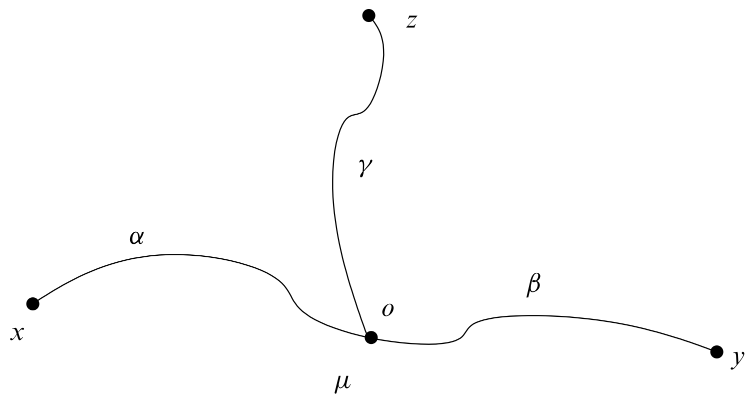

Figure 1: A tripod

2.5 Tripods and ultralimits

Suppose that a path metric space X is thick. Thus, X contains a sequence of (R_i, ϵ_i)–tripods T_i with

$$\lim_i R_i = \infty, \quad \lim_i \epsilon_i = 0,$$

so that the center of T_i is o_i and the legs are $\alpha_i, \beta_i, \gamma_i$. Then the tripods T_i clearly yield a geodesic $(\infty, 0)$–tripod T_ω in $(X_\omega, o_\omega) = \omega\text{–}\lim(X, o_i)$. The tripod T_ω is the union of three geodesic rays $\alpha_\omega, \beta_\omega, \gamma_\omega$ emanating from o_ω, so that

$$o_\omega = p_{\mu_\omega}(\gamma_\omega).$$

Here $\mu_\omega = \alpha_\omega \cup \beta_\omega$. In particular, X_ω is thick.

Conversely, in view of Lemma 2.4, we have:

Lemma 2.7 *If X is (R, ϵ)–thin for $\epsilon > 0$ and $R < \infty$, then X_ω is R'–thin for every $R' > R$.*

Proof Suppose that X_ω contains an R'–tripod T_ω. Then T_ω appears as the ultralimit of $\left(R' - \frac{1}{i}, \frac{1}{i}\right)$–tripods in X. This contradicts the assumption that X is (R, ϵ)–thin. □

Let $\sigma\colon [a,b] \to X$ be a rectifiable curve in X parameterized by its arc-length. We let d_σ denote the path metric on $[a,b]$ which is the pull-back of the path metric on X. By abusing the notation, we denote by d the restriction to σ of the metric d. Note that, in general, d is only a pseudo-metric on $[a,b]$ since σ need not be injective. However, if σ is injective then d is a metric.

We repeat this construction with respect to the tripods: Given a tripod $T \subset X$, define an abstract tripod T_{mod} whose legs have the same length as the legs of T. We have a natural map

$$\tau\colon T_{\mathrm{mod}} \to X$$

which sends the legs of T_{mod} to the respective legs of T, parameterizing them by the arc-length. Then T_{mod} has the path metric d_{mod} obtained by pull-back of the path metric from X via τ. We also have the restriction pseudo-metric d on T_{mod}:

$$d(A,B) = d(\tau(A), \tau(B)).$$

Observe that if $\epsilon = 0$ and X is a tree then the metrics d_{mod} and d on T agree.

Lemma 2.8 $d \le d_{\mathrm{mod}} \le 3d + 6\epsilon$.

Proof The inequality $d \le d_{\mathrm{mod}}$ is clear. We will prove the second inequality. If $A, B \in \alpha \cup \beta$ or $A, B \in \gamma$ then, clearly,

$$d_{\mathrm{mod}}(A,B) \le d(A,B) + \epsilon,$$

since these curves are ϵ–geodesics. Therefore, consider the case when $A \in \gamma$ and $B \in \beta$. Then

$$D := d_{\mathrm{mod}}(A,B) = t + s,$$

where $t = d_\gamma(A,o), s = d_\beta(o,B)$.

Case 1 $t \ge \frac{1}{3}D$. Then, since $o \in \alpha \cup \beta$ is ϵ–nearest to A, it follows that

$$\tfrac{1}{3}D \le t \le d(A,o) + \epsilon \le d(A,B) + 2\epsilon.$$

Hence

$$d_{\mathrm{mod}}(A,B) = \tfrac{3D}{3} \le 3(d(A,B) + 2\epsilon) = 3d(A,B) + 6\epsilon,$$

and the assertion follows in this case.

Case 2 $t < \frac{1}{3}D$. By the triangle inequality,

$$D - t = s \le d(o,B) + \epsilon \le d(o,A) + d(A,B) + \epsilon \le t + 2\epsilon + d(A,B).$$

Hence

$$\tfrac{1}{3}D = D - \tfrac{2}{3}D \le D - 2t \le 2\epsilon + d(A,B),$$

and

$$d_{\mathrm{mod}}(A,B) = \tfrac{3D}{3} \le 3d(A,B) + 6\epsilon. \qquad \square$$

3 Topology of configuration spaces of tripods

We begin with the model tripod T with the legs α_i, $i = 1, 2, 3$, and the center o. Consider the configuration space $Z := T^3 \setminus diag$, where $diag$ is the small diagonal

$$\{(x_1, x_2, x_3) \in T^3 : x_1 = x_2 = x_3\}.$$

We recall that the homology is taken with the $\mathbb{Z}_2$–coefficients.

Proposition 3.1 $H_1(Z) = 0$.

Proof T^3 is the union of cubes

$$Q_{ijk} = \alpha_i \times \alpha_j \times \alpha_k,$$

where $i, j, k \in \{1, 2, 3\}$. Identify each cube Q_{ijk} with the unit cube in the positive octant in $\mathbb{R}^3$. Then in the cube Q_{ijk} $(i, j, k \in \{1, 2, 3\})$ we choose the equilateral triangle σ_{ijk} given by the intersection of Q_{ijk} with the hyperplane

$$x + y + z = 1$$

in $\mathbb{R}^3$. We adopt the convention that if exactly one of the indices i, j, k is zero (say, i), then σ_{ijk} stands for the 1–simplex

$$\{(0, y, z) : y + z = 1\} \cap \{o\} \times \alpha_j \times \alpha_k.$$

Therefore,

$$\partial\sigma_{ijk} = \sigma_{0jk} + \sigma_{i0k} + \sigma_{ij0}.$$

Define the 2–dimensional simplicial complex

$$S := \bigcup_{ijk} \sigma_{ijk}.$$

This complex is homeomorphic to the link of (o, o, o) in T^3. Therefore Z is homotopy-equivalent to

$$W := S \setminus (\sigma_{111} \cup \sigma_{222} \cup \sigma_{333}).$$

Consider the loops $\gamma_i := \partial\sigma_{iii}$, $i = 1, 2, 3$.

Lemma 3.2 (1) *The homology classes* $[\gamma_i], i = 1, 2, 3$ *generate* $H_1(W)$.

(2) $[\gamma_1] = [\gamma_2] = [\gamma_3]$ *in* $H_1(W)$.

Proof of Lemma 3.2 (1) We first observe that S is the 3–fold join of a 3–element set with itself and, therefore, is simply-connected. Alternatively, note that S a 2–dimensional spherical building. Hence, S is homotopy-equivalent to a bouquet of 2–spheres (see Brown [1, Theorem 2, page 93]), which implies that $H_1(S) = 0$. Now the first assertion follows from the long exact sequence of the pair (S, W).

(2) Let us verify that $[\gamma_1] = [\gamma_2]$. The subcomplex

$$S_{12} = S \cap (\alpha_1 \cup \alpha_2)^3$$

is homeomorphic to the 2–sphere. Therefore $S_{12} \cap W$ is the annulus bounded by the circles γ_1 and γ_2. Hence $[\gamma_1] = [\gamma_2]$. □

Lemma 3.3

$$[\gamma_1] + [\gamma_2] + [\gamma_3] = 0$$

in $H_1(W)$.

Proof of Lemma 3.3 Let B' denote the 2–chain

$$\sum_{\{ijk\} \in A} \sigma_{ijk},$$

where A is the set of triples of distinct indices $i, j, k \in \{1, 2, 3\}$. Let

$$B'' := \sum_{i=1}^{3} (\sigma_{ii(i+1)} + \sigma_{i(i+1)i} + \sigma_{(i+1)ii})$$

where we set $3 + 1 := 1$. We note that

$$\gamma_1 + \gamma_2 + \gamma_3 = \partial\Delta,$$

where

$$\Delta = \sum_{i=1}^{3} \sigma_{iii}.$$

Hence, the assertion of lemma is equivalent to

$$\partial(B' + B'' + \Delta) = 0.$$

To prove this, it suffices to show that every 1–simplex in S, appears in $\partial(B' + B'' + \Delta)$ exactly twice. Since the chain $B' + B'' + \Delta$ is preserved by the permutation of the indices i, j, k, it suffices to consider the 1–simplex σ_{ij0} where $j = i + 1$ or $i = j$.

Suppose that $j = i + 1$. Then the 1–simplex σ_{ij0} appears in $\partial(B' + B'' + \Delta)$ exactly twice: in $\partial\sigma_{ijk}$ (where $k \neq i \neq j$) and in $\partial\sigma_{i(i+1)i}$.

Similarly, if $i = j$, then the 1–simplex σ_{ii0} also appears in $\partial(B' + B'' + \Delta)$ exactly twice: in $\partial\sigma_{iii}$ and in $\partial\sigma_{ii(i+1)}$. □

By combining these lemmata we obtain the assertion of the theorem. □

3.0.1 Application to tripods in metric spaces Consider an (R, ϵ)–tripod T in a metric space X and its standard parametrization $\tau\colon T_{\mathrm{mod}} \to T$.

There is an obvious scaling operation

$$u \mapsto r \cdot u$$

on the space $(T_{\mathrm{mod}}, d_{\mathrm{mod}})$ which sends each leg to itself and scales all distances by $r \in [0, \infty)$. It induces the map $T^3_{\mathrm{mod}} \to T^3_{\mathrm{mod}}$, denoted $t \mapsto r \cdot t$, $t \in T^3_{\mathrm{mod}}$.

We have the functions

$$L_{\mathrm{mod}}\colon T^3_{\mathrm{mod}} \to K \qquad L_{\mathrm{mod}}(x, y, z) = (d_{\mathrm{mod}}(x, y), d_{\mathrm{mod}}(y, z), d_{\mathrm{mod}}(z, x)),$$
$$L\colon T^3_{\mathrm{mod}} \to K \qquad L(x, y, z) = (d(x, y), d(y, z), d(z, x))$$

computing side-lengths of triangles with respect to the metrics d_{mod} and d.

For $\rho \geq 0$ set

$$K_\rho := \{(a, b, c) \in K : a + b + c > \rho\}.$$

Define

$$T^3(\rho) := L^{-1}(K_\rho), \qquad T^3_{\mathrm{mod}}(\rho) := L^{-1}_{\mathrm{mod}}(K_\rho).$$

Thus

$$T^3_{\mathrm{mod}}(0) = T^3(0) = T^3 \setminus \mathrm{diag}.$$

Lemma 3.4 *For every $\rho \geq 0$, the space $T^3_{\mathrm{mod}}(\rho)$ is homeomorphic to $T^3_{\mathrm{mod}}(0)$.*

Proof Recall that S is the link of (o, o, o) in T^3. Then scaling defines homeomorphisms

$$T^3_{\mathrm{mod}}(\rho) \to S \times \mathbb{R} \to T^3_{\mathrm{mod}}(0). \qquad \square$$

Corollary 3.5 *For every $\rho \geq 0$, $H_1(T^3_{\mathrm{mod}}(\rho), \mathbb{Z}_2) = 0$.*

Corollary 3.6 *The map induced by inclusion*

$$H_1(T^3(3\rho + 18\epsilon)) \to H_1(T^3(\rho))$$

is zero.

Proof Recall that

$$d \leq d_{\mathrm{mod}} \leq 3d + 6\epsilon.$$

Therefore

$$T^3(3\rho + 18\epsilon) \subset T^3_{\mathrm{mod}}(\rho) \subset T^3(\rho).$$

Now the assertion follows from the previous corollary. □

4 Proof of Theorem 1.3

Suppose that X is thick. Then for every $R < \infty, \epsilon > 0$ there exists an (R, ϵ)–tripod T with the legs α, β, γ. Without loss of generality we may assume that the legs of T have length R. Let $\tau\colon T_{\mathrm{mod}} \to T$ denote the standard map from the model tripod onto T. We will continue with the notation of the previous section.

Given a space E and map $f\colon E \to T^3_{\mathrm{mod}}$ (or a chain $\sigma \in C_*(T^3_{\mathrm{mod}})$), let $\widehat{f}$ (resp. $\widehat{\sigma}$) denote the map $L \circ f$ from E to K (resp. the chain $L_*(\sigma) \in C_*(K)$). Similarly, we define $\widehat{f}_{\mathrm{mod}}$ and $\widehat{\sigma}_{\mathrm{mod}}$ using the map L_{mod} instead of L.

Every loop $\lambda\colon S^1 \to T^3_{\mathrm{mod}}$, determines the map of the 2–disk

$$\Lambda\colon D^2 \to T^3_{\mathrm{mod}},$$

given by

$$\Lambda(r, \theta) = r \cdot \lambda(\theta)$$

where we are using the polar coordinates (r, θ) on the unit disk D^2. Triangulating both S^1 and D^2 and assigning the coefficient $1 \in \mathbb{Z}_2$ to each simplex, we regard both λ and Λ as singular chains in $C_*(T^3_{\mathrm{mod}})$.

We let a, b, c denote the coordinates on the space $\mathbb{R}^3$ containing the cone K. Let $\kappa = (a_0, b_0, c_0)$ be a δ–nondegenerate point in the interior of K for some $\delta > 0$; set $r := a_0 + b_0 + c_0$.

Suppose that there exists a loop λ in T^3_{mod} such that:

(1) $\widehat{\lambda}(\theta)$ is ϵ–degenerate for each θ. Moreover, each triangle $\lambda(\theta)$ is either contained in $\alpha_{\mathrm{mod}} \cup \beta_{\mathrm{mod}}$ or has only two distinct vertices.

In particular, the image of $\widehat{\lambda}$ is contained in

$$K \setminus \mathbb{R}_+ \cdot \kappa.$$

(2) The image of $\widehat{\lambda}$ is contained in K_ρ, where $\rho = 3r + 18\epsilon$.

(3) The homology class $[\widehat{\lambda}]$ is nontrivial in $H_1(K \setminus \mathbb{R}_+ \cdot \kappa)$.

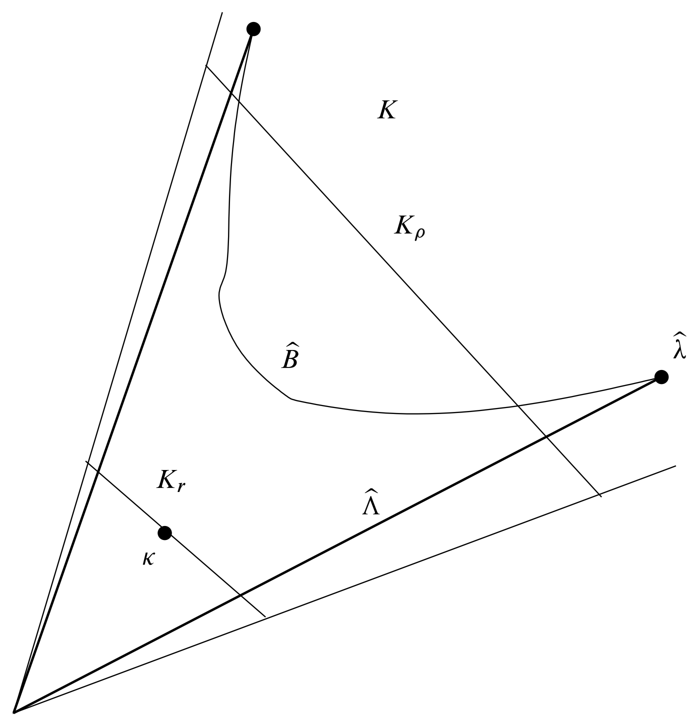

Figure 2: Chains $\widehat{\Lambda}$ and $\widehat{B}$

Lemma 4.1 *If there exists a loop λ satisfying the assumptions (1)–(3), and $\epsilon < \delta/2$, then κ belongs to $K_3(X)$.*

Proof We have the 2–chains

$$\widehat{\Lambda}, \widehat{\Lambda}_{\mathrm{mod}} \in C_2(K \setminus \kappa),$$

with

$$\widehat{\lambda} = \partial \widehat{\Lambda}, \widehat{\lambda}_{\mathrm{mod}} = \partial \widehat{\Lambda}_{\mathrm{mod}} \in C_1(K_\rho).$$

Note that the support of $\widehat{\lambda}_{\mathrm{mod}}$ is contained in ∂K and the 2–chain $\widehat{\Lambda}_{\mathrm{mod}}$ is obtained by coning off $\widehat{\lambda}_{\mathrm{mod}}$ from the origin. Then, by Assumption (1), for every $w \in D^2$:

(i) Either $d(\widehat{\Lambda}(w), \widehat{\Lambda}_{\mathrm{mod}}(w)) \le \epsilon$.

(ii) Or $\widehat{\Lambda}(w), \widehat{\Lambda}_{\mathrm{mod}}(w)$ belong to the common ray in ∂K.

Since $d(\kappa, \partial K) > \delta \ge 2\epsilon$, it follows that the straight-line homotopy H_t between the maps

$$\widehat{\Lambda}, \widehat{\Lambda}_{\mathrm{mod}} \colon D^2 \to K$$

misses κ. Since K_ρ is convex, $H_t(S^1) \subset K_\rho$ for each $t \in [0,1]$, and we obtain

$$[\widehat{\Lambda}_{\text{mod}}] = [\widehat{\Lambda}] \in H_2(K \setminus \kappa, K_\rho).$$

Assumptions (2) and (3) imply that the relative homology class

$$[\widehat{\Lambda}_{\text{mod}}] \in H_2(K \setminus \kappa, K_\rho)$$

is nontrivial. Hence

$$[\widehat{\Lambda}] \in H_2(K \setminus \kappa, K_\rho)$$

is nontrivial as well. Since $\rho = 3r + 18\epsilon$, according to Corollary 3.6, λ bounds a 2–chain

$$\mathrm{B} \in C_2(T^3(r)).$$

Set $\Sigma := \mathrm{B} + \Lambda$. Then the absolute class

$$[\widehat{\Sigma}] = [\widehat{\Lambda} + \widehat{\mathrm{B}}] \in H_2(K \setminus \kappa)$$

is also nontrivial. Since T^3_{mod} is contractible, there exists a 3–chain $\Gamma \in C_3(T^3_{\text{mod}})$ such that

$$\partial\Gamma = \Sigma.$$

Therefore the support of $\widehat{\Gamma}$ contains the point κ. Since the map

$$L\colon T^3 \to K$$

is the composition of the continuous map $\tau^3\colon T^3 \to X^3$ with the continuous map $L\colon X^3 \to K$, it follows that κ belongs to the image of the map $L\colon X^3 \to K$ and hence $\kappa \in K_3(X)$. □

Our goal therefore is to construct a loop λ, satisfying Assumptions (1)–(3).

Let $T \subset X$ be an (R,ϵ)–tripod with the legs α, β, γ of the length R, where $\epsilon \le \delta/2$. We let $\tau\colon T_{\text{mod}} \to T$ denote the standard parametrization of T. Let x, y, z, o denote the vertices and the center of T_{mod}. We let $\alpha_{\text{mod}}(s), \beta_{\text{mod}}(s), \gamma_{\text{mod}}(s)\colon [0,R] \to T_{\text{mod}}$ denote the arc-length parameterizations of the legs of T_{mod}, so that $\alpha(R) = \beta(R) = \gamma(R) = o$.

We will describe the loop λ as the concatenation of seven paths

$$p_i(s) = (x_1(s), x_2(s), x_3(s)), i = 1, \ldots, 7.$$

We let $a = d(x_2, x_3), b = d(x_3, x_1), c = d(x_1, x_2)$.

(1) $p_1(s)$ is the path starting at (x, x, o) and ending at (o, x, o), given by

$$p_1(s) = (\alpha_{\text{mod}}(s), x, o).$$

Note that for $p_1(0)$ and $p_1(R)$ we have $c = 0$ and $b = 0$ respectively.

(2) $p_2(s)$ is the path starting at (o, x, o) and ending at (y, x, o), given by

$$p_2(s) = (\overline{\beta}_{\mathrm{mod}}(s), x, o).$$

(3) $p_3(s)$ is the path starting at (y, x, o) and ending at (y, o, o), given by

$$p_3(s) = (y, \alpha_{\mathrm{mod}}(s), o).$$

Note that for $p_3(R)$ we have $a = 0$.

(4) $p_4(s)$ is the path starting at (y, o, o) and ending at (y, y, o), given by

$$p_4(s) = (y, \overline{\beta}_{\mathrm{mod}}(s), o).$$

Note that for $p_4(R)$ we have $c = 0$. Moreover, if $\alpha * \overline{\beta}$ is a geodesic, then

$$d(\tau(x), \tau(o)) = d(\tau(y), \tau(o)) \Rightarrow \widehat{p}_4(R) = \widehat{p}_1(0)$$

and therefore $\widehat{p}_1 * \cdots * \widehat{p}_4$ is a loop.

(5) $p_5(s)$ is the path starting at (y, y, o) and ending at (y, y, z) given by

$$(y, y, \overline{\gamma}_{\mathrm{mod}}(s)).$$

(6) $p_6(s)$ is the path starting at (y, y, z) and ending at (x, x, z) given by

$$(\beta_{\mathrm{mod}} * \overline{\alpha}_{\mathrm{mod}}, \beta_{\mathrm{mod}} * \overline{\alpha}_{\mathrm{mod}}, z).$$

(7) $p_7(s)$ is the path starting at (x, x, z) and ending at (x, x, o) given by

$$(x, x, \gamma_{\mathrm{mod}}(s)).$$

Thus

$$\lambda := p_1 * \cdots * p_7$$

is a loop.

Since $\alpha * \beta$ and γ are ϵ–geodesics in X, each path $p_i(s)$ determines a family of ϵ–degenerate triangles in (T_{mod}, d). It is clear that Assumption (1) is satisfied.

The class $[\widehat{\lambda}_{\mathrm{mod}}]$ is clearly nontrivial in $H_1(\partial K \setminus 0)$. See Figure 3. Therefore, since $\epsilon \leq \delta/2$,

$$[\widehat{\lambda}] = [\widehat{\lambda}_{\mathrm{mod}}] \in H_1(K \setminus \mathbb{R}_+ \cdot \kappa) \setminus \{0\},$$

see the proof of Lemma 4.1. Thus Assumption (2) holds.

Lemma 4.2 *The image of $\widehat{\lambda}$ is contained in the closure of $K_{\rho'}$, where*

$$\rho' = \frac{2}{3} R - 4\epsilon.$$

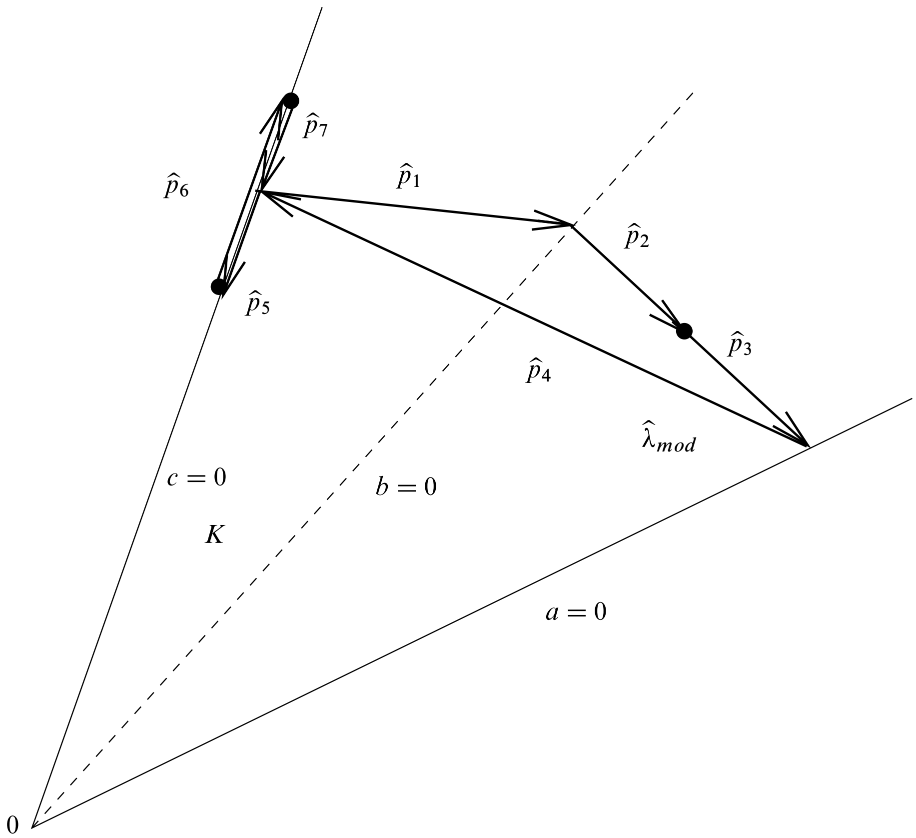

Figure 3: The loop $\widehat{\lambda}_{\mathrm{mod}}$

Proof We have to verify that for each $i = 1, \dots, 7$ and every $s \in [0, R]$, the perimeter (with respect to the metric d) of each triangle $p_i(s) \in T^3_{\mathrm{mod}}$ is at least ρ'. These inequalities follow directly from Lemma 2.8 and the description of the paths p_i. □

Therefore, if we take

$$R > \frac{9}{2}r - 33\epsilon$$

then the image of $\widehat{\lambda}$ is contained in

$$K_{3r+18\epsilon}$$

and Assumption (3) is satisfied. Theorem 1.3 follows. □

5 Quasi-isometric characterization of thin spaces

The goal of this section is to prove Theorem 1.4. Suppose that X is thin. The proof is easier if X is a proper geodesic metric space, in which case there is no need considering the ultralimits. Therefore, we recommend the reader uncomfortable with this technique to assume that X is a proper geodesic metric space.

Pick a base-point $o \in X$, a nonprincipal ultrafilter ω and consider the ultralimit

$$X_\omega = \omega\text{–}\lim(X, o)$$

of the constant sequence of pointed metric spaces. If X is a proper geodesic metric space then, of course, $X_\omega = X$. In view of Lemma 2.7, the space X_ω is R–thin for some R.

Assume that X is unbounded. Then X contains a sequence of $1/i$–geodesic paths $\gamma_i = \overline{ox_i}$ with

$$\omega\text{–}\lim d(o, x_i) = \infty,$$

which yields a geodesic ray ρ_1 in X_ω emanating from the point o_ω.

Lemma 5.1 *Let ρ be a geodesic ray in X_ω emanating from a point O. Then the neighborhood $E = N_R(\rho)$ is an end $E(\rho)$ of X_ω.*

Proof Suppose that α is a path in $X_\omega \setminus B_{2R}(O)$ connecting a point $y \in X_\omega \setminus E$ to a point $x \in E$. Then there exists a point $z \in \alpha$ such that $d(z, \rho) = R$. Since X_ω contains no R–tripods,

$$d(p_\rho(z), O) < R.$$

Therefore $d(z, O) < 2R$. Contradiction. □

Set $E_1 := E(\rho_1)$. If the image of the natural embedding $\iota\colon X \to X_\omega$ is contained in a finite metric neighborhood of ρ_1, then we are done, as X is quasi-isometric to $\mathbb{R}_+$. Otherwise, there exists a sequence $y_n \in X$ such that:

$$\omega\text{–}\lim d(\iota(y_n), \rho_1) = \infty.$$

Consider the $\frac{1}{n}$–geodesic paths $\alpha_n \in P(o, y_n)$. The sequence (α_n) determines a geodesic ray $\rho_2 \subset X_\omega$ emanating from o_ω. Then there exists $s \geq 4R$ such that

$$d(\alpha_n(s), \gamma_i) \geq 2R$$

for ω–all n and ω–all i. Therefore, for $t \geq s$, $\rho_2(t) \notin E(\rho_1)$. By applying Lemma 5.1 to ρ_2 we conclude that X_ω has an end $E_2 = E(\rho_2) = N_R(\rho_2)$. Since E_1, E_2 are distinct ends of X_ω, $E_1 \cap E_2$ is a bounded subset. Let D denote the diameter of this intersection.

Lemma 5.2 (1) *For every pair of points* $x_i = \rho_i(t_i)$, $i = 1, 2$, *we have*

$$\overline{x_1 x_2} \subset N_{D/2+2R}(\rho_1 \cup \rho_2).$$

(2) $\rho_1 \cup \rho_2$ *is a quasi-geodesic.*

Proof Consider the points x_i as in Part 1. Our goal is to get a lower bound on $d(x_1, x_2)$. A geodesic segment $\overline{x_1 x_2}$ has to pass through the ball $B(o_\omega, 2R), i = 1, 2,$ since this ball separates the ends E_1, E_2. Let $y_i \in \overline{x_1 x_2} \cap B(o_\omega, 2R)$ be such that

$$\overline{x_i y_i} \subset E_i, \qquad i = 1, 2.$$

Then

$$d(y_1, y_2) \le D + 4R,$$
$$d(x_i, y_i) \ge t_i - 2R,$$

and

$$\overline{x_i y_i} \subset N_R(\rho_i), \qquad i = 1, 2.$$

This implies the first assertion of Lemma. Moreover,

$$d(x_1, x_2) \ge d(x_1, y_1) + d(x_2, y_2) \ge t_1 + t_2 - 4R = d(x_1, x_2) - 4R.$$

Therefore $\rho_1 \cup \rho_2$ is a $(1, 4R)$–quasi-geodesic. □

If $\iota(X)$ is contained in a finite metric neighborhood of $\rho_1 \cup \rho_2$, then, by Lemma 5.2, X is quasi-isometric to $\mathbb{R}$. Otherwise, there exists a sequence $z_k \in X$ such that

$$\omega\text{–}\lim d(\iota(z_k), \rho_1 \cup \rho_2) = \infty.$$

By repeating the construction of the ray ρ_2, we obtain a geodesic ray $\rho_3 \subset X_\omega$ emanating from the point o_ω, so that ρ_3 is not contained in a finite metric neighborhood of $\rho_1 \cup \rho_2$. For every t_3, the nearest-point projection of $\rho_3(t_3)$ to

$$N_{D/2+2R}(\rho_1 \cup \rho_2)$$

is contained in

$$B_{2R}(o_\omega).$$

Therefore, in view of Lemma 5.2, for every pair of points $\rho_i(t_i)$ as in that lemma, the nearest-point projection of $\rho_3(t_3)$ to $\overline{\rho_1(t_1)\rho_2(t_2)}$ is contained in

$$B_{4R+D}(o_\omega).$$

Hence, for sufficiently large t_1, t_2, t_3, the points $\rho_i(t_i)$, $i = 1, 2, 3$ are vertices of an R–tripod in X. This contradicts the assumption that X_ω is R–thin.

Therefore X is either bounded, or is quasi-isometric to a $\mathbb{R}_+$ or to $\mathbb{R}$. □

6 Examples

Theorem 6.1 *There exist an (incomplete) 2–dimensional Riemannian manifold M quasi-isometric to $\mathbb{R}$, so that:*

(1) *$K_3(M)$ does not contain $\partial K_3(\mathbb{R}^2)$.*

(2) *For the Riemannian product $M^2 = M \times M$, $K_3(M^2)$ does not contain $\partial K_3(\mathbb{R}^2)$ either.*

Moreover, there exists $D < \infty$ such that for every degenerate triangle in M and M^2, at least one side is $\leq D$.

Proof (1) We start with the open concentric annulus $A \subset \mathbb{R}^2$, which has the inner radius $R_1 > 0$ and the outer radius $R_2 < \infty$. We give A the flat Riemannian metric induced from $\mathbb{R}^2$. Let M be the universal cover of A, with the pull-back Riemannian metric. Since M admits a properly discontinuous isometric action of $\mathbb{Z}$ with the quotient of finite diameter, it follows that M is quasi-isometric to $\mathbb{R}$. The metric completion $\overline{M}$ of M is diffeomorphic to the closed bi-infinite flat strip. Let $\partial_1 M$ denote the component of the boundary of $\overline{M}$ which covers the inner boundary of A under the map of metric completions

$$\overline{M} \to \overline{A}.$$

As a metric space, $\overline{M}$ is $CAT(0)$, therefore it contains a unique geodesic between any pair of points. However, for any pair of points $x, y \in M$, the geodesic $\gamma = \overline{xy} \subset \overline{M}$ is the union of subsegments

$$\gamma_1 \cup \gamma_2 \cup \gamma_3$$

where $\gamma_1, \gamma_3 \subset M$, $\gamma_2 \subset \partial_1 M$, and the lengths of γ_1, γ_3 are at most $D_0 = \sqrt{R_2^2 - R_1^2}$. Hence, for every degenerate triangle (x, y, z) in M, at least one side is $\leq D_0$.

(2) We observe that the metric completion of M^2 is $\overline{M} \times \overline{M}$; in particular, it is again a $CAT(0)$ space. Therefore it has a unique geodesic between any pair of points. Moreover, geodesics in $\overline{M} \times \overline{M}$ are of the form

$$(\gamma_1(t), \gamma_2(t))$$

where γ_i, $i = 1, 2$ are geodesics in $\overline{M}$. Hence for every geodesic segment $\gamma \subset \overline{M} \times \overline{M}$, the complement $\gamma \setminus \partial \overline{M}^2$ is the union of two subsegments of length $\leq \sqrt{2} D_0$ each. Therefore for every degenerate triangle in M^2, at least one side is $\leq \sqrt{2} D_0$. □

Remark The manifold M^2 is, of course, quasi-isometric to $\mathbb{R}^2$.

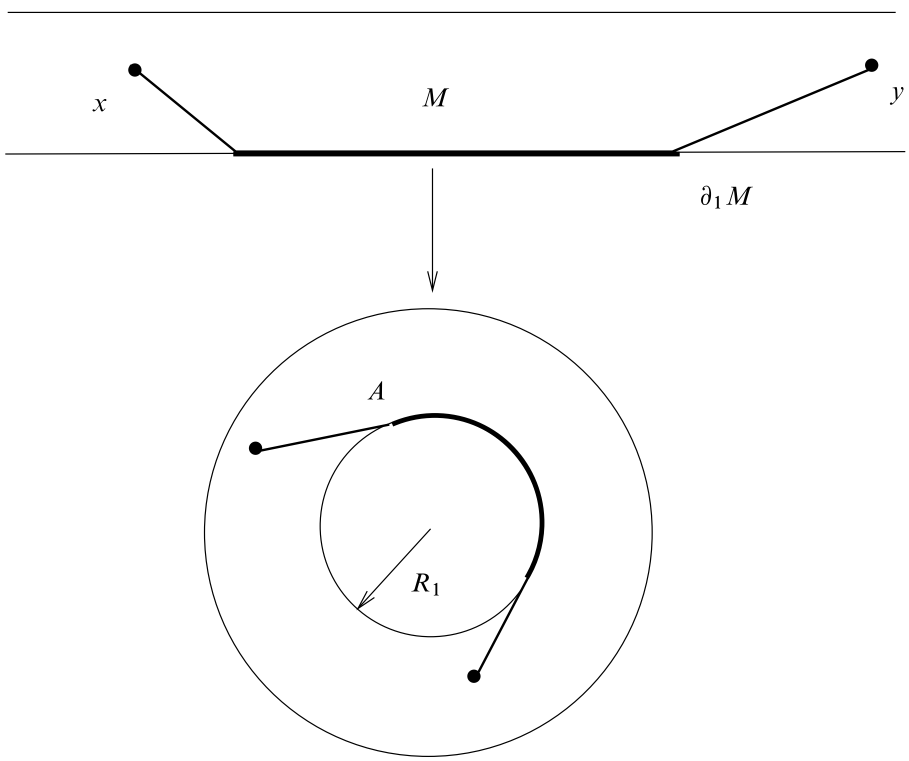

Figure 4: Geodesics in $\overline{M}$

Our second example is a graph-theoretic analogue of the Riemannian manifold M.

Theorem 6.2 *There exists a complete path metric space X (a metric graph) quasi-isometric to $\mathbb{R}$ so that:*

(1) *$K_3(X)$ does not contain $\partial K_3(\mathbb{R}^2)$.*

(2) *$K_3(X^2)$ does not contain $\partial K_3(\mathbb{R}^2)$.*

Moreover, there exists $D < \infty$ such that for every degenerate triangle in X and X^2, at least one side is $\le D$.

Proof (1) We start with the disjoint union of oriented circles α_i of the length $1 + \frac{1}{i}$, $i \in I = \mathbb{N} \setminus \{2\}$. We regard each α_i as a path metric space. For each i pick a point $o_i \in \alpha_i$ and its antipodal point $b_i \in \alpha_i$. We let α_i^+ be the positively oriented arc of α_i connecting o_i to b_i. Let α_i^- be the complementary arc.

Consider the bouquet Z of α_i's by gluing them all at the points o_i. Let $o \in Z$ be the image of the points o_i. Next, for every pair $i, j \in I$ attach to Z the oriented arc β_{ij} of the length

$$\tfrac{1}{2} + \tfrac{1}{4}\left(\tfrac{1}{i} + \tfrac{1}{j}\right)$$

connecting b_i and b_j and oriented from b_i to b_j if $i < j$. Let Y denote the resulting graph. We give Y the path metric. Then Y is a complete metric space, since it is a metric graph where the length of every edge is at least $1/2 > 0$. Note also that the length of every edge in Y is at most 1.

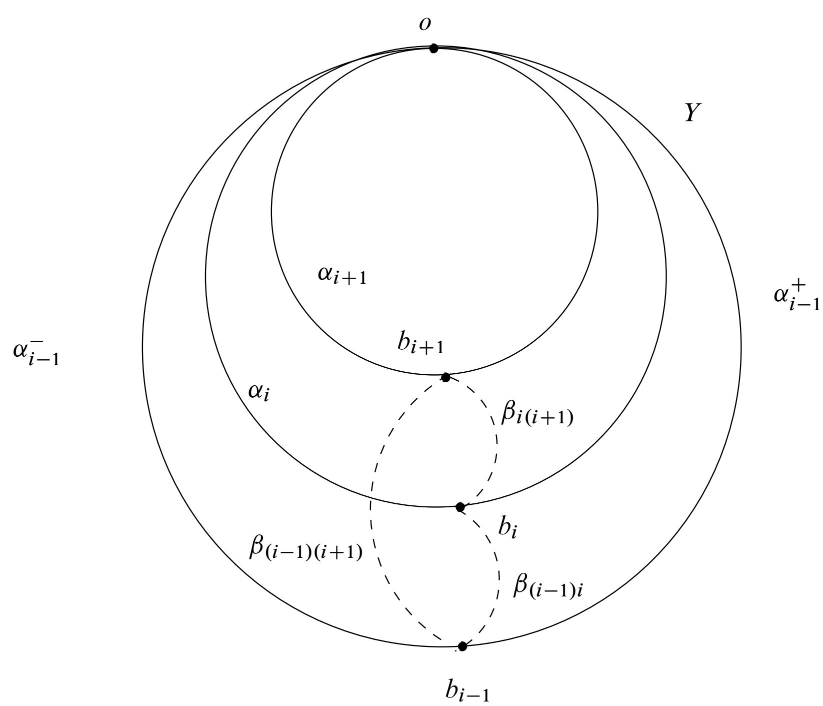

Figure 5: The metric space Y

The space X is the infinite cyclic regular cover over Y defined as follows. Take the maximal subtree

$$T = \bigcup_{i \in I} \alpha_i^+ \subset Y.$$

Every oriented edge of $Y \setminus T$ determines a free generator of $G = \pi_1(Y, o)$. Define the homomorphism $\rho\colon G \to \mathbb{Z}$ by sending every free generator to 1. Then the covering

$X \to Y$ is associated with the kernel of ρ. (This covering exists since Y is locally contractible.)

We lift the path metric from Y to X, thereby making X a complete metric graph. We label vertices and edges of X as follows.

(i) Vertices a_n which project to o. The cyclic group $\mathbb{Z}$ acts simply transitively on the set of these vertices thereby giving them the indices $n \in \mathbb{Z}$.

(ii) The edges $\alpha_i^{\pm}$ lift to the edges $\alpha_{in}^{+}, \alpha_{in}^{-}$ incident to the vertices a_n and a_{n+1} respectively.

(iii) The intersection $\alpha_{in}^{+} \cap \alpha_{i(n+1)}^{-}$ is the vertex b_{in} which projects to the vertex $b_i \in \alpha_i$.

(iv) The edge β_{ijn} connecting b_{in} to $b_{j(n+1)}$ which projects to the edge $\beta_{ij} \subset Y$.

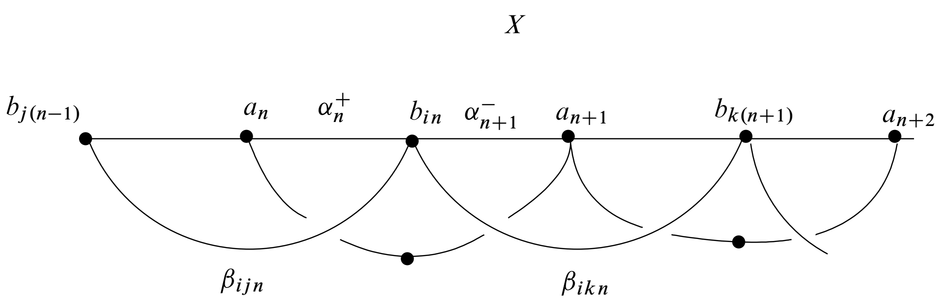

Figure 6: The metric space X

Lemma 6.3 *X contains no degenerate triangles (x, y, v), so that v is a vertex,*

$$d(x, v) + d(v, y) = d(x, y)$$

and $\min(d(x, v), d(v, y)) > 2$.

Proof of Lemma 6.3 Suppose that such degenerate triangles exist.

Case 1 $(v = b_{in})$ Since the triangle (x, y, v) is degenerate, for all sufficiently small $\epsilon > 0$ there exist ϵ–geodesics σ connecting x to y and passing through v.

Since $d(x, v), d(v, y) > 2$, it follows that for sufficiently small $\epsilon > 0$, $\sigma = \sigma(\epsilon)$ also passes through $b_{j(n-1)}$ and $b_{k(n+1)}$ for some j, k depending on σ. We will assume that as $\epsilon \to 0$, both j and k diverge to infinity, leaving the other cases to the reader.

Therefore

$$d(x,v) = \lim_{j\to\infty}(d(x,b_{j(n-1)}) + d(b_{j(n-1)},v)),$$
$$d(v,y) = \lim_{k\to\infty}(d(y,b_{k(n+1)}) + d(b_{k(n+1)},v)).$$

Then

$$\lim_{j\to\infty} d(b_{j(n-1)},v) + \lim_{k\to\infty} d(b_{k(n+1)},v) = 1 + \frac{1}{2i}.$$

On the other hand, clearly,

$$\lim_{j,k\to\infty} d(b_{j(n-1)},b_{k(n+1)}) = 1.$$

Hence

$$d(x,y) = \lim_{j\to\infty} d(x,b_{j(n-1)}) + \lim_{k\to\infty} d(y,b_{k(n+1)}) + 1 < d(x,v) + d(v,y).$$

Contradiction.

Case 2 $(v = a_n)$ Since the triangle (x,y,v) is degenerate, for all sufficiently small $\epsilon > 0$ there exist ϵ–geodesics σ connecting x to y and passing through v. Then for sufficiently small $\epsilon > 0$, every σ also passes through $b_{j(n-1)}$ and b_{kn} for some j,k depending on σ. However, since $j,k \geq 2$,

$$d(b_{j(n-1)},b_{kn}) = \frac{1}{2} + \frac{1}{4j} + \frac{1}{4i} \leq \frac{3}{4} < 1 = \inf_{j,k}(d(b_{j(n-1)},v) + d(v,b_{kn})).$$

Therefore $d(x,y) < d(x,v) + d(v,y)$. Contradiction. □

Corollary 6.4 *X contains no degenerate triangles (x,y,z), such that*

$$d(x,z) + d(z,y) = d(x,y)$$

and $\min(d(x,z), d(z,y)) \geq 3$.

Proof of Corollary 6.4 Suppose that such a degenerate triangle exists. We can assume that z is not a vertex. The point z belongs to an edge $e \subset X$. Since $\mathrm{length}(e) \leq 1$, for one of the vertices v of e

$$d(z,v) \leq 1/2.$$

Since the triangle (x,y,z) is degenerate, for all ϵ–geodesics $\sigma \in P(x,z)$, $\eta \in P(z,y)$ we have:

$$e \subset \sigma \cup \eta,$$

provided that $\epsilon > 0$ is sufficiently small. Therefore the triangle (x, y, v) is also degenerate. Clearly,

$$\min(d(x, v), d(y, v)) \geq \min(d(x, z), d(y, z)) - 1/2 \geq 2.5.$$

This contradicts Lemma 6.3. □

Hence part (1) of Theorem 6.2 follows.

(2) We consider $X^2 = X \times X$ with the product metric

$$d^2((x_1, y_1), (x_2, y_2)) = d^2(x_1, x_2) + d^2(y_1, y_2).$$

Then X^2 is a complete path-metric space. Every degenerate triangle in X^2 projects to degenerate triangles in both factors. It therefore follows from part (1) that X contains no degenerate triangles with all sides ≥ 18. We leave the details to the reader. □

7 Exceptional cases

Theorem 7.1 *Suppose that X is a path metric space quasi-isometric to a metric space X', which is either $\mathbb{R}$ or $\mathbb{R}_+$. Then there exists a $(1, A)$–quasi-isometry $X' \to X$.*

Proof We first consider the case $X' = \mathbb{R}$. The proof is simpler if X is proper, therefore we sketch it first under this assumption. Since X is quasi-isometric to $\mathbb{R}$, it is 2–ended with the ends E_+, E_-. Pick two divergent sequences $x_i \in E_+, y_i \in E_-$. Then there exists a compact subset $C \subset X$ so that all geodesic segments $\gamma_i := \overline{x_i y_i}$ intersect C. It then follows from the Arcela-Ascoli theorem that the sequence of segments γ_i subconverges to a complete geodesic $\gamma \subset X$. Since X is quasi-isometric to $\mathbb{R}$, there exists $R < \infty$ such that $X = N_R(\gamma)$. We define the $(1, R)$–quasi-isometry $f\colon \gamma \to X$ to be the identity (isometric) embedding.

We now give a proof in the general case. Pick a non-principal ultrafilter ω on $\mathbb{N}$ and a base-point $o \in X$. Define X_ω as the ω–limit of (X, o). The quasi-isometry $f\colon \mathbb{R} \to X$ yields a quasi-isometry $f_\omega\colon \mathbb{R} = \mathbb{R}_\omega \to X_\omega$. Therefore X_ω is also quasi-isometric to $\mathbb{R}$.

We have the natural isometric embedding $\iota\colon X \to X_\omega$. As above, let E_+, E_- denote the ends of X and choose divergent sequences $x_i \in E_+, y_i \in E_-$. Let γ_i denote an $\frac{1}{i}$–geodesic segment in X connecting x_i to y_i. Then each γ_i intersects a bounded subset $B \subset X$. Therefore, by taking the ultralimit of γ_i's, we obtain a complete geodesic $\gamma \subset X_\omega$. Since X_ω is quasi-isometric to $\mathbb{R}$, the embedding $\eta\colon \gamma \to X_\omega$ is a quasi-isometry. Hence $X_\omega = N_R(\gamma)$ for some $R < \infty$.

For the same reason,

$$X_\omega = N_D(\iota(X))$$

for some $D < \infty$. Therefore the isometric embeddings

$$\eta\colon \gamma \to X_\omega, \quad \iota\colon X \to X_\omega$$

are $(1, R)$ and $(1, D)$–quasi-isometries respectively. By composing η with the quasi-inverse to ι, we obtain a $(1, R+3D)$–quasi-isometry $\mathbb{R} \to X$.

The case when X is quasi-isometric to $\mathbb{R}_+$ can be treated as follows. Pick a point $o \in X$ and glue two copies of X at o. Let Y be the resulting path metric space. It is easy to see that Y is quasi-isometric to $\mathbb{R}$ and the inclusion $X \to Y$ is an isometric embedding. Therefore, there exists a $(1, A)$–quasi-isometry $h\colon Y \to \mathbb{R}$ and the restriction of h to X yields the $(1, A)$–quasi-isometry from X to the half-line. □

Note that the conclusion of Theorem 7.1 is false for path metric spaces quasi-isometric to $\mathbb{R}^n$, $n \geq 2$.

Corollary 7.2 *Suppose that X is a path metric space quasi-isometric to $\mathbb{R}$ or $\mathbb{R}_+$. Then $K_3(X)$ is contained in the D–neighborhood of ∂K for some $D < \infty$. In particular, $K_3(X)$ does not contain the interior of $K = K_3(\mathbb{R}^2)$.*

Proof Suppose that $f\colon X \to X'$ is an (L, A)–quasi-isometry, where X' is either $\mathbb{R}$ or $\mathbb{R}_+$. According to Theorem 7.1, we can assume that $L = 1$. For every triple of points $x, y, z \in X$, after relabeling, we obtain

$$d(x,y) + d(y,z) \leq d(x,z) + D,$$

where $D = 3A$. Then every triangle in X is D–degenerate. Hence $K_3(X)$ is contained in the D–neighborhood of ∂K. □

Remark One can construct a metric space X quasi-isometric to $\mathbb{R}$ such that $K_3(X) = K$. Moreover, X is isometric to a curve in $\mathbb{R}^2$ (with the metric obtained by the restriction of the metric on $\mathbb{R}^2$). Of course, the metric on X is not a path metric.

Corollary 7.3 *Suppose that X is a path metric space. Then the following are equivalent:*

(1) *$K_3(X)$ contains the interior of $K = K_3(\mathbb{R}^2)$.*

(2) *X is not quasi-isometric to the point, $\mathbb{R}_+$ and $\mathbb{R}$.*

(3) *X is thick.*

Proof (1)⇒(2) by Corollary 7.2. (2)⇒(3) by Theorem 1.4. (3)⇒(1) by Theorem 1.3. □

Remark The above corollary remains valid under the following assumption on the metric on X, which is weaker than being a path metric:

For every pair of points $x, y \in X$ and every $\epsilon > 0$, there exists a $(1, \epsilon)$–quasi-geodesic path $\alpha \in P(x, y)$.

References

[1] **K S Brown**, *Buildings*, Springer, New York (1989) MR969123

[2] **D Burago**, **Y Burago**, **S Ivanov**, *A course in metric geometry*, Graduate Studies in Mathematics 33, American Mathematical Society, Providence, RI (2001) MR1835418

[3] **M Gromov**, *Metric structures for Riemannian and non-Riemannian spaces*, Progress in Mathematics 152, Birkhäuser, Boston (1999) MR1699320

[4] **M Kapovich**, *Hyperbolic manifolds and discrete groups*, Progress in Mathematics 183, Birkhäuser, Boston (2001) MR1792613

[5] **M Kapovich**, **B Leeb**, *On asymptotic cones and quasi-isometry classes of fundamental groups of 3–manifolds*, Geom. Funct. Anal. 5 (1995) 582–603 MR1339818

[6] **J Roe**, *Lectures on coarse geometry*, University Lecture Series 31, American Mathematical Society, Providence, RI (2003) MR2007488

Department of Mathematics, University of California, Davis
Davis CA 95616, USA

kapovich@math.ucdavis.edu

http://www.math.ucdavis.edu/~kapovich/

Proposed: Walter Neumann
Seconded: Yasha Eliashberg, Martin Bridson

Received: 6 December 2006
Accepted: 30 July 2007