RIGIDITY FOR QUASI-MÖBIUS GROUP ACTIONS

MARIO BONK AND BRUCE KLEINER

Abstract. If a group acts by uniformly quasi-Möbius homeomorphisms on a compact Ahlfors n-regular space of topological dimension n such that the induced action on the space of distinct triples is cocompact, then the action is quasi-symmetrically conjugate to an action on the standard n-sphere by Möbius transformations.

1. Introduction

It has been known since the time of Poincaré that the limit set of a subgroup of $\mathrm{PSL}(2,\mathbb{C})$ obtained by a small deformation of a discrete cocompact subgroup of $\mathrm{PSL}(2,\mathbb{R}) \subseteq \mathrm{PSL}(2,\mathbb{C})$ will be a nowhere differentiable curve unless it is round. Much later R. Bowen [3] made this more precise by proving that such a limit curve is either a round circle or has Hausdorff dimension strictly greater than 1. The group $\mathrm{PSL}(2,\mathbb{C})$ is isomorphic to the group of orientation preserving isometries of hyperbolic 3-space. Therefore, it is a natural question whether similar results hold for subgroups of the isometry group $\mathrm{Isom}(\mathbb{H}^{n+1})$ of hyperbolic $(n+1)$-space when $n \geq 2$, or, what is the same, for groups of Möbius transformations acting on the standard n-sphere $\mathbb{S}^n$. Rigidity results in this vein were obtained by Sullivan [11, p. 69] and Yue [14, Theorem 1.5].

In the present paper we generalize these results further by considering uniformly quasi-Möbius group actions on compact metric spaces Z that induce cocompact actions on the space $\mathrm{Tri}(Z)$ of distinct triples of Z. The following theorem is our main result.

Theorem 1.1. *Let $n \in \mathbb{N}$, and let Z be a compact, Ahlfors n-regular metric space of topological dimension n. Suppose $G \curvearrowright Z$ is a uniformly quasi-Möbius action of a group G on Z, where the induced action $G \curvearrowright \mathrm{Tri}(Z)$ is cocompact. Then $G \curvearrowright Z$ is quasi-symmetrically conjugate to an action of G on the standard sphere $\mathbb{S}^n$ by Möbius transformations.*

Date: December 14, 2001.

M.B. was supported by a Heisenberg fellowship of the Deutsche Forschungsgemeinschaft. B.K. was supported by NSF grant DMS-9972047.

The terminology will be explained in the body of the paper. Note that part of the conclusion is that Z is homeomorphic to $\mathbb{S}^n$.

When G is a hyperbolic group, the boundary $\partial_\infty G$ carries a metric d unique up to quasi-symmetry, with respect to which the canonical action $G \curvearrowright \partial_\infty G$ is uniformly quasi-Möbius. In this case the induced action on $\mathrm{Tri}(\partial_\infty G)$ is discrete and cocompact, so Theorem 1.1 may be applied if $(\partial_\infty G, d)$ is quasi-symmetric to an Ahlfors n-regular space whose topological dimension is equal to n. Note that $(\partial_\infty G, d)$ will always be Ahlfors Q-regular for some $Q > 0$, but in general Q will exceed the topological dimension of $\partial_\infty G$.

In order to state our next result, we recall (see the discussion in Section 7) that if X is a $\mathrm{CAT}(-1)$-space, then any point $p \in X$ determines a canonical metric on $\partial_\infty X$, and any two such metrics are bi-Lipschitz equivalent by the identity map. In particular, we may speak of the Hausdorff dimension of any subset of $\partial_\infty X$, since this number is independent of the choice of the canonical metric. We then have the following corollary of Theorem 1.1 which generalizes a result by Bourdon [2, 0.3 Théorème ($\mathbb{H}^n$ case)].

Theorem 1.2. *Suppose $n \in \mathbb{N}$, $n \geq 2$. Let $G \curvearrowright X$ be a properly discontinuous, quasi-convex cocompact, and isometric action on a $\mathrm{CAT}(-1)$-space X. If the Hausdorff dimension and topological dimension of the limit set $\Lambda(G) \subseteq \partial_\infty X$ are both equal to n, then X contains a convex, G-invariant subset Y isometric to $\mathbb{H}^{n+1}$ on which G acts cocompactly.*

The terminology and the notation will be explained in Section 7. Note that the ineffective kernel N of the induced action $G \curvearrowright Y$ is finite, and G/N is isomorphic to a uniform lattice in $\mathrm{Isom}(\mathbb{H}^{n+1})$.

In contrast to Theorem 1.1 where the case $n = 1$ is allowed, we assume $n > 1$ in the previous theorem, in order to be able to apply Bourdon's result. It is an interesting question whether the statement is also true for $n = 1$. See the discussion in Section 7.

The proof of Theorem 1.1 can be outlined as follows. First, we use the dimension assumption to get a Lipschitz map $f\colon Z \to \mathbb{S}^n$ such that the image of f has positive Lebesgue measure. According to a result by David and Semmes one can rescale f and extract a limit mapping $\phi\colon X \to \mathbb{R}^n$ defined on a weak tangent space of Z which has bounded multiplicity, i.e. point inverses $\phi^{-1}(y)$ have uniformly bounded cardinality. We then show that ϕ is locally bi-Lipschitz somewhere, and as a consequence some weak tangent of Z is bi-Lipschitz to $\mathbb{R}^n$. The assumptions on the group action can then be used to prove that Z is quasi-symmetric to $\mathbb{S}^n$. Once this is established, the theorem follows from a result by Tukia.

Our method of proving Theorem 1.1 can also be applied in other contexts. In [8, Question 5] Heinonen and Semmes ask whether every linearly locally contractible Ahlfors n-regular metric n-sphere Z that is quasi-symmetrically three point homogeneous is quasi-symmetrically equivalent to the standard n-sphere $\mathbb{S}^n$. One can show that the answer to this question is positive, if we make the stronger assumption that Z is three point homogeneous by uniform quasi-Möbius homeomorphisms. (see the discussion in Section 6).

Acknowledgement. A previous version of this paper was based on some rather deep results on the uniform rectifiability of metric spaces satisfying some topological nondegeneracy assumptions. The statements we needed are implicitly contained in the works of David and Semmes, but not stated explicitly. The approach taken in this version uses a much more elementary result by David and Semmes. The authors are indebted to Stephen Semmes for conversations about these issues and thank him especially for directing their attention to the results in Chapter 12 of [6].

Notation. The following notation will be used throughout the paper. Let Z be a metric space. The metric on Z will be denoted by d_Z, and the open and the closed ball of radius $r > 0$ centered at $a \in Z$ by $B_Z(a, r)$ and $\bar{B}_Z(a, r)$, respectively. We will drop the subscript Z if the space Z is understood. If $A \subseteq Z$ and $d = d_Z$, then $d|_A$ is the restriction of the metric d to A. We use $\operatorname{diam}(A)$ for the diameter, $\bar{A}$ for the closure, and $\#A$ for the cardinality of a set $A \subseteq Z$. If $z \in Z$ and $A, B \subseteq Z$, then $\operatorname{dist}(z, A)$ and $\operatorname{dist}(A, B)$ are the distances of z and A and of A and B, respectively. If $A \subseteq Z$ and $r > 0$, then we let $N_r(A) := \{z \in Z : \operatorname{dist}(z, A) < r\}$. The *Hausdorff distance* of two sets $A, B \subseteq Z$ is defined by

$$\operatorname{dist}_H(A, B) := \max\Big\{\sup_{a\in A} \operatorname{dist}(a, B),\ \sup_{b\in A} \operatorname{dist}(b, A)\Big\}.$$

Suppose X and Y are metric spaces. If $f\colon X \to Y$ is a map, then we let $\operatorname{Im}(f) := \{f(x) : x \in X\}$. If $A \subseteq X$, then $f|_A$ denotes the restriction of the map f to A. If $g\colon X \to Y$ is another map, we let

$$\operatorname{dist}(f, g) := \sup_{x\in X} \operatorname{dist}(f(x), g(x)).$$

The identity map on a set X will be denoted by id_X.

2. Quasi-Möbius maps and group actions

Let (Z,d) be a metric space. The *cross-ratio* of a four-tuple of distinct points (z_1,z_2,z_3,z_4) in Z is the quantity

$$[z_1,z_2,z_3,z_4] := \frac{d(z_1,z_3)d(z_2,z_4)}{d(z_1,z_4)d(z_2,z_3)}.$$

Suppose X and Y are metric spaces. Suppose $\eta\colon [0,\infty)\to[0,\infty)$ is a homeomorphism, and let $f\colon X\to Y$ be an injective map. The map f is an *η-quasi-Möbius map* if for every four-tuple (x_1,x_2,x_3,x_4) of distinct points in X, we have

$$[f(x_1),f(x_2),f(x_3),f(x_4)]\leq \eta([x_1,x_2,x_3,x_4]).$$

Note that by exchanging the roles of x_1 and x_2, one gets the lower bound

$$\eta([x_1,x_2,x_3,x_4]^{-1})^{-1}\leq [f(x_1),f(x_2),f(x_3),f(x_4)].$$

Hence f is a homeomorphism onto its image $f(X)$, and the inverse map $f^{-1}\colon f(X)\to X$ is also quasi-Möbius.

The map f is *η-quasi-symmetric* if

$$\frac{d_Y(f(x_1),f(x_2))}{d_Y(f(x_1),f(x_3))}\leq \eta\left(\frac{d_X(x_1,x_2)}{d_X(x_1,x_3)}\right)$$

for every triple (x_1,x_2,x_3) of distinct points in X.

Finally, f is called *bi-Lipschitz* if there exists a constant $L\geq 1$ (the bi-Lipschitz constant of f) such that

$$(1/L)d_X(x_1,x_2)\leq d_Y(f(x_1),f(x_2))\leq Ld_X(x_1,x_2),$$

whenever $x_1,x_2\in X$.

We mention some basic properties of these maps.

(1) The post-composition of an η_1-quasi-Möbius map by an η_2-quasi-Möbius map is an $\eta_2\circ\eta_1$-quasi-Möbius map. Similar statements are true for quasi-symmetric maps and bi-Lipschitz maps.

(2) A bi-Lipschitz map is quasi-symmetric and quasi-Möbius. A quasi-symmetric map is quasi-Möbius. A quasi-Möbius map defined on a bounded space is quasi-symmetric.

(3) Let X and Y be compact metric spaces, and suppose $f_k\colon X\to Y$ is an η-quasi-Möbius map for $k\in\mathbb{N}$. Then we have that

(a) the sequence (f_k) subconverges uniformly to an η-quasi-Möbius map, or

(b) there is a point $x_0\in X$ so that the sequence $(f_k|_{X\setminus\{x_0\}})$ subconverges uniformly on compact subsets of $X\setminus\{x_0\}$ to a constant map.

The alternative (b) can be excluded by a normalization condition; namely, that each map f_k maps a uniformly separated triple of points in X to a uniformly separated triple in Y.

We will need the following extension of property (3).

Lemma 2.1. *Suppose (X, d_X) and (Y, d_Y) are compact metric spaces, and let $f_k\colon D_k \to Y$ for $k \in \mathbb{N}$ be an η-quasi-Möbius map defined on a subset D_k of X. Suppose*

$$\lim_{k\to\infty} \mathrm{dist}_H(D_k, X) = 0,$$

and that for $k \in \mathbb{N}$ there exist triples (x_k^1, x_k^2, x_k^3) and (y_k^1, y_k^2, y_k^3) of points in D_k and Y, respectively, such that

$$f_k(x_k^i) = y_k^i \quad \text{for} \quad k \in \mathbb{N},\ i \in \{1,2,3\},$$

$$d_X(x_k^i, x_k^j) \geq \delta \text{ and } d_Y(y_k^i, y_k^j) \geq \delta \text{ for } k \in \mathbb{N},\ i,j \in \{1,2,3\},\ i \neq j,$$

where $\delta > 0$ is independent of k.

Then the sequence (f_k) subconverges uniformly to a quasi-Möbius map $f\colon X \to Y$, i.e. there exists a monotonic sequence (k_ν) in $\mathbb{N}$ such that

$$\lim_{\nu\to\infty} \mathrm{dist}(f_{k_\nu}, f|_{D_{k_\nu}}) = 0.$$

Suppose in addition that

$$\lim_{k\to\infty} \mathrm{dist}_H(f_k(D_k), Y) = 0.$$

Then the sequence (f_k) subconverges uniformly to quasi-Möbius homeomorphism $f\colon X \to Y$.

The lemma says that a sequence (f_k) of uniformly quasi-Möbius maps defined on denser and denser subsets of a space X and mapping into the same space Y subconverges to a quasi-Möbius map defined on the whole space X, if each map f_k maps a uniformly separated triple in X to a uniformly separated triple in Y. Moreover, a surjective limiting map can be obtained if the images of the maps f_k Hausdorff converge to the space Y.

Proof. The assumptions imply that the functions f_k are equicontinuous (cf. [13, Thm. 2.1]). The proof of the first part of the lemma then follows from standard arguments based on the Arzelà-Ascoli theorem, and we leave the details to the reader.

To prove the second part, note that according to the first part, by passing to a subsequence if necessary, we may assume that

$$\mathrm{dist}(f_k, f|_{D_k}) \to 0 \quad \text{for} \quad k \to \infty.$$

Let $D'_k := f_k(D_k)$ and $g_k := f_k^{-1} \colon D'_k \to X$. The maps g_k are uniformly quasi-Möbius. Hence, by our additional assumption we can apply the first part of the lemma to the sequence (g_k). Again by selecting a subsequence of (g_k) if necessary, we may assume that

$$\operatorname{dist}(g_k, g|_{D'_k}) \to 0 \quad \text{for} \quad k \to \infty,$$

where $g \colon Y \to X$ is a quasi-Möbius map. Since $g_k \circ f_k = \mathrm{id}_{D_k}$ and $f_k \circ g_k = \mathrm{id}_{D'_k}$, we obtain from the uniform convergence of the sequences (f_k) and (g_k) that $g \circ f = \mathrm{id}_X$ and $g \circ f = \mathrm{id}_Y$. Hence f is a bijection and therefore a quasi-Möbius homeomorphism. $\square$

Let Z be an unbounded locally compact metric space with metric $d = d_Z$, let $p \in Z$ be a base point, and let $\hat{Z} = Z \cup \{\infty\}$ be the one-point compactification of Z. In order to define a metric on $\hat{Z}$ associated with the pointed space (Z, p) let $h_p \colon \hat{Z} \to [0, \infty)$ be given by

$$h_p(z) := \begin{cases} \dfrac{1}{1 + d(z,p)} & \text{for} \quad z \in Z, \\ 0 & \text{for} \quad z = \infty. \end{cases}$$

Moreover, let

$$\rho_p(x,y) = h_p(x) h_p(y) d(x,y) \quad \text{for} \quad x, y \in Z,$$

$\rho_p(x, \infty) = \rho_p(\infty, x) = h_p(x)$ for $x \in Z$, $\rho_p(\infty,\infty) = 0$. Note that if an argument of the functions h_p and ρ_p is the point at infinity, the corresponding value can be obtained as a limiting case of values at arguments in Z. Essentially, the function ρ_p is the metric on $\hat{Z}$ that we are looking for. This distance function is an analog of the chordal metric on the Riemann sphere. Unfortunately, ρ_p will not satisfy the triangle inequality in general. We remedy this problem by a standard procedure.

If $x, y \in \hat{Z}$ we define

$$\hat{d}_p(x,y) := \inf \sum_{i=0}^{k-1} \rho_p(x_i, x_{i+1}),$$

where the infimum is taken over all finite sequence of points $x_0, \dots, x_k \in \hat{Z}$ with $x_0 = x$ and $x_k = y$.

Lemma 2.2. *The function $\hat{d}_p$ is a metric on $\hat{Z}$ whose induced topology agrees with the topology of $\hat{Z}$. The identity map $\mathrm{id}_Z \colon (Z, d) \to (Z, \hat{d}_p|_Z)$ is an η-quasi-Möbius homeomorphism where $\eta(t) = 16t$.*

Proof. The first part of the lemma immediately follows if we can show that

$$\frac{1}{4}\rho_p(x,y) \le \hat{d}_p(x,y) \le \rho_p(x,y) \quad \text{for} \quad x,y \in \hat{Z}. \tag{2.3}$$

The second part also follows from this inequality by observing that if (z_1, z_2, z_3, z_4) is a four-tuple of distinct points in Z, then

$$\frac{\hat{d}_p(z_1,z_3)\hat{d}_p(z_2,z_4)}{\hat{d}_p(z_1,z_4)\hat{d}_p(z_2,z_3)} \le 16\frac{\rho_p(z_1,z_3)\rho_p(z_2,z_4)}{\rho_p(z_1,z_4)\rho_p(z_2,z_3)} = 16\frac{d(z_1,z_3)d(z_2,z_4)}{d(z_1,z_4)d(z_2,z_3)}.$$

The right hand inequality in (2.3) follows from the definition of $\hat{d}_p$. In order to prove the left hand inequality, we may assume $h_p(x) \ge h_p(y)$ without loss of generality. Moreover, we may assume $x \in Z$ and so $h_p(x) > 0$, because otherwise $x = y = \infty$ and the inequality is true.

If $x_0, \ldots, x_k$ is an arbitrary sequence with $x_0 = x$ and $x_k = y$, we consider two cases:

If $h_p(x_i) \ge \frac{1}{2}h_p(x) > 0$ for all $i \in \{0, \ldots, k\}$, then $x_i \in Z$, and the triangle inequality applied to d gives

$$\begin{aligned} \sum_{i=0}^{k-1} \rho_p(x_i, x_{i+1}) &\ge \frac{1}{4}h_p(x)^2 \sum_{i=0}^{k-1} d(x_i, x_{i+1}) \\ &\ge \frac{1}{4}d(x,y)h_p(x)h_p(y) = \frac{1}{4}\rho_p(x,y). \end{aligned} \tag{2.4}$$

Suppose there exists $j \in \{0, \ldots, k\}$ such that $h_p(x_j) < \frac{1}{2}h_p(x)$. Note that that it follows from the definitions that $|h_p(u) - h_p(v)| \le \rho_p(u,v)$ for $u, v \in \hat{Z}$. Moreover, since $h_p(y) \le h_p(x)$ we have $d(x,p) \le d(y,p)$ in case $y \in Z$. This implies

$$\frac{d(x,y)}{1+d(y,p)} \le 2\frac{d(y,p)}{1+d(y,p)} \le 2,$$

which leads to $\rho_p(x,y) \le 2h_p(x)$. This is also true if $y = \infty$. We arrive at

$$\sum_{i=0}^{k-1} \rho_p(x_i, x_{i+1}) \ge \sum_{i=0}^{k-1} |h_p(x_i) - h_p(x_{i+1})| \ge \frac{1}{2}h_p(x) \ge \frac{1}{4}\rho_p(x,y). \tag{2.5}$$

The desired inequality follows from (2.4) and (2.5). □

Let (Z, d) be a metric space. We write $G \curvearrowright Z$, if G is a group that acts on Z by homeomorphisms. The image of a point $z \in Z$ under the group element g is denoted by $g(z)$. The action $G \curvearrowright Z$ is called *faithful* if the only element in G that acts as the identity on Z is the unit element.

If $\eta\colon [0,\infty) \to [0,\infty)$ is a homeomorphism, then an action $G \curvearrowright Z$ is an *η-quasi-Möbius action* if each $g \in G$ induces an η-quasi-Möbius homeomorphism of Z. An action $G \curvearrowright Z$ is *uniformly quasi-Möbius* if it is η-quasi-Möbius for some homeomorphism $\eta\colon [0,\infty) \to [0,\infty)$. If Z is locally compact, then the action $G \curvearrowright Z$ is called *cocompact* if there exists a compact set $K \subseteq Z$ such that

$$Z = \bigcup_{g\in G} g(K).$$

We denote by

$$\mathrm{Tri}(Z) := \{(z_1, z_2, z_3) \in Z^3 : z_1 \neq z_2 \neq z_3 \neq z_1\}$$

the space of distinct triples in Z. If $G \curvearrowright Z$ is a group action, then there is a natural induced action $G \curvearrowright \mathrm{Tri}(Z)$ defined by

$$g(z_1, z_2, z_3) := (g(z_1), g(z_2), g(z_3))$$

for $g \in G$ and $(z_1, z_2, z_3) \in \mathrm{Tri}(Z)$.

Suppose $G \curvearrowright Z$ is an action on a compact space Z. Then the induced action $G \curvearrowright \mathrm{Tri}(Z)$ is cocompact if and only if there exists $\delta > 0$ such that for every triple $(z_1, z_2, z_3) \in \mathrm{Tri}(Z)$ there exists a group element $g \in G$ such that

$$d(g(z_i), g(z_j)) \geq \delta \quad \text{for} \quad i, j \in \{1, 2, 3\},\ i \neq j.$$

This condition means that every triple in $\mathrm{Tri}(Z)$ can be mapped to a uniformly separated triple by some map $g \in G$.

3. Maps of bounded multiplicity

The goal of this section is to study continuous maps of bounded multiplicity between a space of topological dimension n and $\mathbb{R}^n$. The main result is Theorem 3.4 which may be of independent interest.

Definition 3.1. If $f\colon X \to Y$ is a continuous map between metric spaces X and Y, then $y \in Y$ is a *stable value* of f if there is $\epsilon > 0$ such that $y \in \mathrm{Im}(g)$ for every continuous map $g\colon X \to Y$ with $\mathrm{dist}(f, g) < \epsilon$.

Note that the set of stable values of a map $f\colon X \to \mathbb{R}^n$ is an open subset of $\mathbb{R}^n$.

Recall that a map is *light* if all point inverses are totally disconnected. We will prove the following proposition.

Proposition 3.2. *Let X be a compact metric space of topological dimension at least n, and let $f\colon X \to \mathbb{R}^n$ be a light continuous map. Then f has stable values.*

As we will see, the proof is a slight amplification of the well-known argument that such a map f cannot decrease topological dimension.

Definition 3.3. A map $f\colon X \to Y$ between two spaces has *bounded multiplicity* if there is a constant $N \in \mathbb{N}$ such that $\#f^{-1}(y) \leq N$ for all $y \in Y$.

Using Proposition 3.2 we will prove:

Theorem 3.4. *Suppose X is a compact metric space, every nonempty open subset of X has topological dimension at least n, and $f\colon X \to \mathbb{R}^n$ is a continuous map of bounded multiplicity. Then there is an open subset $V \subseteq \mathrm{Im}(f)$ with $\bar{V} = \mathrm{Im}(f)$, such that $U := f^{-1}(V)$ is dense in X and $f|_U\colon U \to V$ is a covering map.*

In particular, there exist nonempty open sets $U_1 \subseteq X$ and $V_1 \subseteq \mathbb{R}^n$ such that $f|_{U_1}$ is a homeomorphism of U_1 onto V_1. It is in this form that we will use Theorem 3.4 in the proof of Theorem 1.1.

Let X be a topological space, and let $\mathcal{U} = \{U_i : i \in I\}$ be a cover of X by open subsets U_i indexed by some set I. The *nerve* of $\mathcal{U}$, denoted by $\mathrm{Ner}(\mathcal{U})$, is a simplicial complex whose simplices corresponds to the subsets $I' \subseteq I$ for which

$$U_{I'} := \bigcap_{i \in I'} U_i \neq \emptyset.$$

The *order* of $\mathcal{U}$ is the supremum of all numbers $\#I'$ such that $U_{I'} \neq \emptyset$. We denote the topological dimension of X by $\dim_{\mathrm{top}}(X)$ (cf. [9, Def. I.4]). A compact metric space X has topological dimension at most n, if and only if open covers of order at most $n+1$ are cofinal in the family of all open covers of X, i.e., every open cover has an open refinement which has order at most $n+1$. The order of an open cover $\mathcal{U}$ is equal to $\dim_{\mathrm{top}}(\mathrm{Ner}(\mathcal{U})) + 1$.

In order to prove Proposition 3.2 we discuss a general construction that associates a fine cover with a light continuous map $f\colon X \to Y$ from a compact metric space X to a separable metric space Y. Pick $\epsilon > 0$.

If $y \in Y$, then $f^{-1}(y)$ is compact and totally disconnected, so the diameter of connected components of $N_\delta(f^{-1}(y))$ tends to zero as $\delta \to 0$. Hence there is a number $r_y > 0$ such that $N_{r_y}(f^{-1}(y))$ can be decomposed as a finite disjoint union of open sets with diameter less than ϵ; moreover, there is a number $s_y > 0$ such that $f^{-1}(B(y, s_y)) \subseteq N_{r_y}(f^{-1}(y))$. Let $\mathcal{B}$ be a finite cover of $\mathrm{Im}(f)$ by balls of the form $B(y, s_y)$.

Suppose $\mathcal{U} = \{U_i : i \in I\}$ is a cover of $\mathrm{Im}(f)$ by open subsets of Y. Let $I' := \{i \in I : U_i \cap \mathrm{Im}(f) \neq \emptyset\}$, and assume that $\mathcal{U}' := \{U_i : i \in I'\}$ refines $\mathcal{B}$. Then $f^{-1}(\mathcal{U}') := \{f^{-1}(U_i) : i \in I'\}$ is an open cover of X such that for all $i \in I'$, we have $f^{-1}(U_i) \subseteq N_{r_y}(f^{-1}(y))$ for some $y \in Y$, which implies that $f^{-1}(U_i)$ may be written as a finite disjoint union of open subsets with diameter less than ϵ. Choosing such a decomposition of $f^{-1}(U_i)$ for each $i \in I'$ yields a collection of open sets $\mathcal{V} = \{V_j : j \in J\}$ which covers X, and a map $\alpha \colon J \to I' \subseteq I$ such that V_j is an open set appearing in the decomposition of $f^{-1}(U_{\alpha(j)})$. Note that α induces a simplicial map $\phi : \mathrm{Ner}(\mathcal{V}) \to \mathrm{Ner}(\mathcal{U})$ since

$$\begin{aligned} V_{j_1} \cap \ldots \cap V_{j_k} \neq \emptyset &\Rightarrow f^{-1}(U_{\alpha(j_1)}) \cap \ldots \cap f^{-1}(U_{\alpha(j_k)}) \neq \emptyset \\ &\Rightarrow U_{\alpha(j_1)} \cap \ldots \cap U_{\alpha(j_k)} \neq \emptyset. \end{aligned}$$

In fact, ϕ is injective on simplices, since if $j,\, j' \in J$ are distinct and $\alpha(j) = \alpha(j')$, then V_j and $V_{j'}$ are disjoint fragments of the same open set $f^{-1}(U_{\alpha(j)}) = f^{-1}(U_{\alpha(j')})$, and so $V_j \cap V_{j'} = \emptyset$. In particular, we have $\dim_{\mathrm{top}}(\mathrm{Ner}(\mathcal{V})) \leq \dim_{\mathrm{top}}(\mathrm{Ner}(\mathcal{U}))$.

Suppose $\{\rho_i : i \in I\}$ is a partition of unity in Y subordinate to $\mathcal{U}$. Here and in the following we interpret subordination in the sense that $\{\rho_i \neq 0\} \subseteq U_i$ for all $i \in I$. We can produce a partition of unity $\{\nu_j : j \in J\}$ in X subordinate to $\mathcal{V}$ as follows: let $\nu_j := \chi_{V_j} \cdot \big(\rho_{\alpha(j)} \circ f\big)$, where χ_{V_j} is the characteristic function of V_j. Using the functions $\{\rho_i : i \in I\}$ as barycentric coordinates in $\mathrm{Ner}(\mathcal{U})$, and the functions $\{\nu_j : j \in J\}$ as barycentric coordinates in $\mathrm{Ner}(\mathcal{V})$, we obtain induced continuous maps $\rho \colon Y \to \mathrm{Ner}(\mathcal{U})$ and $\nu \colon X \to \mathrm{Ner}(\mathcal{V})$ such that $\phi \circ \nu = \rho \circ f$.

We note that since $\epsilon > 0$ was chosen arbitrarily, if we have a cofinal family of covers $\mathcal{U}'$ of $\mathrm{Im}(f)$ of order at most N, then the corresponding family of covers $\mathcal{V}$ of X will be cofinal and its members will have order at most N; this implies that $\dim_{\mathrm{top}}(Y) \geq \dim_{\mathrm{top}}(\mathrm{Im}(f)) \geq \dim_{\mathrm{top}}(X)$.

Proof of Proposition 3.2. For $\epsilon > 0$ we now apply the construction above in the special case that $Y = \mathbb{R}^n$, $\dim_{\mathrm{top}}(X) \geq n$, and the open cover $\mathcal{U} = \{U_i : i \in I\}$ of $\mathbb{R}^n$ is the open star cover associated with a triangulation of $\mathbb{R}^n$. Since $f(X)$ is compact, the associated cover $\mathcal{U}'$ will refine a given cover of $f(X)$ if the triangulation of $\mathbb{R}^n$ is chosen fine enough. We have a homeomorphism $\rho \colon \mathbb{R}^n \to \mathrm{Ner}(\mathcal{U})$ (we conflate simplicial complexes with their geometric realizations), and an induced partition of unity $\{\rho_i : i \in I\}$ coming from the barycentric coordinate functions of the map ρ.

Since the family of open covers of X induced by our construction is cofinal in the family of all open covers of X, we can choose $\epsilon > 0$

small enough so that the induced cover $\mathcal{V}$ of X does not admit an open refinement $\mathcal{W}$ of order at most n.

Lemma 3.5. *Some n-simplex σ of* $\mathrm{Ner}(\mathcal{V})$ *has an interior point ξ which is a stable value of $\nu\colon X \to \mathrm{Ner}(\mathcal{V})$.*

Proof. Suppose not. Then we may form a set S by choosing one interior point from each n-simplex of $\mathrm{Ner}(\mathcal{V})$, and perturb ν slightly on a small neighborhood of $\nu^{-1}(S)$ to get a map $\nu'\colon X \to \mathrm{Ner}(\mathcal{V})$ such that its barycentric coordinate functions are subordinate to $\mathcal{V}$, and $\mathrm{Im}(\nu') \cap S = \emptyset$. (See the first part of the proof of Lemma 3.7 for the idea of how to construct this perturbation.) Then we may compose ν' with the "radial projection" in each n-simplex to get a map ν'' that maps $\mathrm{Ner}(\mathcal{V}) \setminus S$ to the $(n-1)$-skeleton $[\mathrm{Ner}(\mathcal{V})]_{n-1}$ of $\mathrm{Ner}(\mathcal{V})$ and whose barycentric coordinates are subordinate to $\mathcal{V}$; pulling back the open star cover of $\mathrm{Ner}(\mathcal{V})$ by ν'', we get a refinement of $\mathcal{V}$ of order at most n, which is a contradiction. □

If ξ is as in the lemma, then $\phi(\xi) \in \mathrm{Ner}(\mathcal{U})$ is clearly a stable value of $\phi \circ \nu\colon X \to \mathrm{Ner}(\mathcal{U})$; but $f = \rho^{-1} \circ \phi \circ \nu$ where ρ^{-1} is a homeomorphism, so $\rho^{-1}(\phi(\xi))$ is a stable value of f. This completes the proof of Proposition 3.2. □

Definition 3.6. Let X be a topological space, and $f\colon X \to \mathbb{R}^n$ be a map. Then $x \in X$ is a *stable point* of f if $f(x)$ is a stable value of $f|_U$ for every neighborhood U of x.

Lemma 3.7. *Suppose X is metric space, and $f\colon X \to \mathbb{R}^n$ is a continuous map. Then $y \in \mathbb{R}^n$ is a stable value of f if and only if y is a stable value of $f|_{f^{-1}(W)}$ for every neighborhood W of y. When X is a compact metric space and $f^{-1}(y)$ is totally disconnected, then y is a stable value of f if and only if the fiber $f^{-1}(y)$ contains a stable point.*

Proof. We will only prove the "only if" implications; the other implications are immediate.

Suppose $W \subseteq \mathbb{R}^n$ is an open neighborhood of y, and y is an unstable value of $f|_U$, where $U := f^{-1}(W)$. Choose $\delta > 0$ such that $\bar{B}(y,\delta) \subseteq W$, and let $V := f^{-1}(\mathbb{R}^n \setminus \bar{B}(y,\delta))$. Pick $\epsilon > 0$. As y is an unstable value of $f|_U$, we can find a map $g_U\colon U \to \mathbb{R}^n$ such that $\mathrm{dist}(g_U, f|_U) < \min(\epsilon,\delta)$ and $y \notin \mathrm{Im}(g_U)$. Define $g_V\colon V \to \mathbb{R}^n$ to be the restriction of f to V. Combining g_U and g_V using a partition of unity subordinate to the cover $\{U, V\}$, we get a continuous map $g\colon X \to \mathbb{R}^n$ such that $\mathrm{dist}(g, f) < \epsilon$ and $g^{-1}(y) = \emptyset$. Since $\epsilon > 0$ was arbitrary, we have shown that y is not a stable value of f.

Now suppose X is compact, $f^{-1}(y)$ is totally disconnected, and every point $x \in f^{-1}(y)$ is unstable. By the compactness of $f^{-1}(y)$ we can find a finite cover $\mathcal{B} = \{B(x_1, r_1), \dots, B(x_k, r_k)\}$ of $f^{-1}(y)$ by balls where $x_i \in f^{-1}(y)$ and y is an unstable value of $f|_{B(x_i,r_i)}$ for each $1 \le i \le k$. When $\delta > 0$ is sufficiently small, then $f^{-1}(B(y,\delta))$ can be decomposed into a disjoint union of open sets $U_1, \dots, U_j$ so that the cover $\{U_i\}$ of $f^{-1}(y)$ refines $\mathcal{B}$. This means that y is an unstable value of $f|_{U_i}$ for each i, which implies that y is an unstable value of $f|_{f^{-1}(B(y,\delta))}$. This is a contradiction to what we proved in the first part of the proof. $\square$

Now let X be a compact metric space such that $\dim_{\mathrm{top}}(U) \ge n$ for all nonempty open subsets $U \subseteq X$, and $f\colon X \to \mathbb{R}^n$ be a continuous map of bounded multiplicity.

Lemma 3.8. *For all $y \in \mathbb{R}^n$ and all $\epsilon > 0$, there is $\delta > 0$ such that for all $y' \in B(y,\delta)$ and all stable points $x \in f^{-1}(y)$, there is a stable point in $f^{-1}(y') \cap B(x,\epsilon)$.*

Proof. Let $\{x_1, \dots, x_k\}$ be the stable points in $f^{-1}(y)$ and pick $i \in \{1, \dots, k\}$. Since x_i is stable point, y is a stable value of $f|_{B(x_i,\epsilon)}$. So any y' sufficiently close to y is also a stable value of $f|_{B(x_i,\epsilon)}$ and by Lemma 3.7 for such y' we will have a stable point in $f^{-1}(y') \cap B(x_i,\epsilon)$. This holds for all i, so the lemma follows. $\square$

We define the *stable multiplicity function* $\mu\colon \mathbb{R}^n \to \mathbb{N}$ by letting $\mu(y)$ be the number of stable points in $f^{-1}(y)$.

Lemma 3.9. *If μ is locally maximal at $y \in \mathbb{R}^n$, then every $x \in f^{-1}(y)$ is stable.*

Proof. Let $U \subseteq \mathbb{R}^n$ be a neighborhood of y such that $\mu(y') \le \mu(y)$ for all $y' \in U$. Let $x_1, \dots, x_k$ be the stable points in $f^{-1}(y)$, and suppose $x \in f^{-1}(y) \setminus \{x_1, \dots, x_k\}$. Pick $\epsilon > 0$ such that the balls $B(x,\epsilon)$, $B(x_1,\epsilon), \dots, B(x_k,\epsilon)$ are disjoint.

Choose $\delta > 0$ as in the previous lemma. Let y' be a stable value of $f|_{B(x,\epsilon)}$ lying in $U \cap B(y,\delta)$; such a y' exists since by Proposition 3.2 stable values of $f|_{B(x,\epsilon)}$ are dense in $\mathrm{Im}(f|_{B(x,\epsilon)})$. Then $f^{-1}(y')$ has a stable point in each of the balls $B(x,\epsilon)$, $B(x_1,\epsilon), \dots, B(x_k,\epsilon)$, so $\mu(y') \ge k+1$; this is a contradiction. $\square$

Proof of Theorem 3.4. Let $V \subseteq \mathrm{Im}(f) \subseteq \mathbb{R}^n$ be the set where the stable multiplicity function μ is locally maximal; clearly V is dense in $\mathrm{Im}(f)$. By Lemma 3.8, V is an open subset of $\mathbb{R}^n$, and μ is locally constant on V. By Lemma 3.9, the map $y \mapsto \# f^{-1}(y)$ is a locally constant function on V. It is therefore clear by Lemma 3.8 that f is locally injective near any $x \in U := f^{-1}(V)$, and hence $f|_U$ is a covering map. If W is a nonempty open set in X, then $f(W)$ has nonempty interior by Proposition 3.2. Hence $f(W)$ meets V, since V is dense in $\mathrm{Im}(f)$.

It follows that W meets $U = f^{-1}(V)$. This implies that U is dense in X. $\square$

4. Weak Tangents

In this section we briefly review some results on weak tangents. For more details see [6] and [4].

A pointed metric space is a pair (Z, p), where Z is a metric space (with metric d_Z) and $p \in Z$. A sequence (Z_k, p_k) of pointed metric spaces is said to converge to a pointed metric space (Z, p), if for every $R > 0$ and for every $\epsilon > 0$ there exist $N \in \mathbb{N}$, a subset $M \subseteq B_Z(p, R)$, subsets $M_k \subseteq B_{Z_k}(R)$ and bijections $f_k \colon M_k \to M$ such that for $k \geq N$

(i) $p \in M$, $p_k \in M_k$, and $f_k(p_k) = p$,
(ii) the set M is ϵ-dense in $B_Z(p, R)$, and the sets M_k are ϵ-dense in $B_{Z_k}(p_k, R)$,
(iii) $|d_{Z_k}(x, y) - d_Z(f_k(x), f_k(y))| < \epsilon$ whenever $x, y \in M_k$.

The definitions for pointed space convergence given in [6] and [4] are different, but equivalent.

A complete metric space S is called a *weak tangent* of the metric space Z, if there exist a sequence of numbers $\lambda_k > 0$ with $\lambda_k \to \infty$ for $k \to \infty$ and points $q \in S$, $p_k \in Z$ such that the sequence of pointed spaces $(\lambda_k Z, p_k)$ converges to the pointed space (S, q). Here we denote by λZ for $\lambda > 0$ the metric space $(Z, \lambda d_Z)$. In other words, λZ agrees with Z as a set, but is equipped with the metric obtained by rescaling the original metric by the factor $\lambda > 0$. The set of all weak tangents of a metric space Z is denoted by $\mathrm{WT}(Z)$. If X, Y, Z are metric spaces, and X is a weak tangent of Y and Y is a weak tangent of Z, then X is a weak tangent of Z, i.e., $X \in \mathrm{WT}(Y)$ and $Y \in \mathrm{WT}(Z)$ imply $X \in \mathrm{WT}(Z)$.

A metric space Z is called *uniformly perfect* if there exists a constant $\lambda \geq 1$ such that for every $z \in Z$ and $0 < R \leq \operatorname{diam}(Z)$ we have $\bar{B}(z, R) \setminus B(z, R/\lambda) \neq \emptyset$.

For $Q > 0$ we denote by $\mathcal{H}^Q$ the Q-dimensional Hausdorff measure on a metric space Z. A complete metric space Z of positive diameter is called *Ahlfors Q-regular*, where $Q > 0$, if there exists a constant $C \geq 1$ such that

$$\frac{1}{C} R^Q \leq \mathcal{H}^Q(B(z, R)) \leq C R^Q,$$

whenever $z \in Z$ and $0 < R \leq \operatorname{diam}(Z)$.

A metric space Z is called *doubling*, if there exists a number $N \in \mathbb{N}$ such that every open ball of radius R in Z can be covered by at most N

open balls of radius $R/2$. The space Z is called *proper*, if closed balls in Z are compact.

Every Ahlfors regular space is uniformly perfect and doubling. A complete doubling space is proper. If Z is a compact metric space that is uniformly perfect and doubling, and $X \in \mathrm{WT}(Z)$, then X is an unbounded doubling metric space. Since X is also complete by definition, this space will be proper.

Suppose $f\colon X \to Y$ is a map between a metric space X and a doubling metric space Y. The map is called *regular* if it is Lipschitz and there exists a constant $N \in \mathbb{N}$ such that the inverse image of every open ball B in Y can be covered by at most N open balls in X with the same radius as B.

Note that this last condition implies that f is of bounded multiplicity. Indeed, we have $\#f^{-1}(y) \leq N$ for $y \in Y$. For suppose that there are $N+1$ distinct points $x_1, \dots, x_{N+1} \in f^{-1}(y)$. Let $\epsilon > 0$ be the minimum of the distances $d_X(x_i, x_j)$ for $i \neq j$. Consider the ball $B = B(y, \epsilon/2)$. By our assumption on f the preimage $f^{-1}(B) \supseteq f^{-1}(y)$ can be covered by N open balls $B_1, \dots, B_N \subseteq X$ of radius $\epsilon/2$. But this is impossible, because each ball B_i can contain at most one of the points $x_1, \dots, x_{N+1}$.

The proof of the following proposition can be found in [6, Prop. 12.8].

Proposition 4.1. *Let X and Y be metric spaces, and $f\colon X \to Y$ be a Lipschitz map. Suppose that X is compact and Ahlfors Q-regular, where $Q > 0$, Y is complete and doubling, and $\mathcal{H}^Q(f(X)) > 0$.*

Then there exist weak tangents $S \in \mathrm{WT}(X)$, $T \in \mathrm{WT}(Y)$, and a regular map $g\colon S \to T$.

We will need the following lemmas.

Lemma 4.2. *Suppose X is a metric space, and $f\colon X \to \mathbb{R}^n$ is regular. Assume that there is an open ball $B \subseteq \mathbb{R}^n$ and a set $U \subseteq f^{-1}(B)$ such that the map $g := f|_U\colon U \to B$ is a homeomorphism. Then g is a bi-Lipschitz map.*

It is understood that U is equipped with the restriction of the metric d_X to U, and B with the Euclidean metric.

Proof. Since f is Lipschitz, the map g is also Lipschitz. It remains to obtain an upper bound for $d_X(x, y)$ in terms of $|f(x) - f(y)|$, whenever $x, y \in U$, $x \neq y$. Let $R := 2|f(x) - f(y)| > 0$, $B' := B(x, R)$ and $S \subseteq B' \cap B$ be the Euclidean line segment connecting $f(x)$ and $f(y)$. Then $E := g^{-1}(S)$ is a compact connected set in U containing x and y. On the other hand, $E \subseteq f^{-1}(B')$. If $N \in \mathbb{N}$ is associated with f as in the definition of a regular map, then it follows that E can be covered by N open balls of radius R. Now we invoke the following elementary fact whose proof is left to the reader: If E is a compact connected set in a metric space covered by open balls, then the diameter of E is at most twice the sum of the radii of the balls.

In our situation we get the estimate

$$d_X(x,y) \le \operatorname{diam}(E) \le 2NR = 4N|f(x)-f(y)|,$$

which proves that g is a bi-Lipschitz homeomorphism. □

Lemma 4.3. *Suppose X and Y are complete doubling metric spaces. Suppose there exists a point $x \in X$, a neighborhood U of x and a bi-Lipschitz map $f\colon U \to V := f(U)$ such that V is a neighborhood of $y := f(x)$.*

Then there exist $S \in \mathrm{WT}(X)$, $T \in \mathrm{WT}(Y)$, and a bi-Lipschitz homeomorphism $g\colon S \to T$.

The lemma says that under the given hypotheses the spaces X and Y have bi-Lipschitz equivalent weak tangents.

Proof. For $\lambda > 0$ consider the pointed metric spaces $(\lambda U, x)$ and $(\lambda V, y)$, where λU and λV denote the metric spaces whose underlying sets are U and V equipped with the restrictions of the metric d_X and d_Y, respectively, rescaled by the factor $\lambda > 0$. The map f considered as a map between $(\lambda U, x)$ and $(\lambda V, y)$ preserves base points and is bi-Lipschitz with a constant independent of λ. Since X and Y are complete and doubling, it follows that in the terminology of David and Semmes [6, Sect. 8.5] the mapping packages $f\colon (\lambda U, x) \to (\lambda V, y)$ subconverge for $\lambda \to \infty$ to a mapping $g\colon S \to T$. Here S and T are limits of the pointed spaces $(\lambda_k U, x)$ and $(\lambda_k V, y)$, respectively, where λ_k is a sequence of positive numbers with $\lambda_k \to \infty$ as $k \to \infty$. Since U and V are neighborhoods of x and y, respectively, it follows that $S \in \mathrm{WT}(X)$ and $Y \in \mathrm{WT}(Y)$ (cf. [6, Lem. 9.12]). Moreover, since the bi-Lipschitz constant of $f\colon (\lambda U, x) \to (\lambda V, y)$ is independent of λ, the map g will be bi-Lipschitz. There is a slight problem here, because it is not clear whether g will be surjective. This problem can be addressed similarly as in the second part of the proof of Theorem 2.1. We may assume that the sequence λ_k is such that not only the mapping packages $f\colon (\lambda_k U, x) \to (\lambda_k V, y)$ converge, but also the mapping packages $f^{-1}\colon (\lambda_k V, y) \to (\lambda_k U, x)$, to $h\colon T \to S$, say. Then $g \circ h = \mathrm{id}_T$ which implies that g is onto, and hence a bi-Lipschitz homeomorphism. □

5. Weak tangents and quasi-Möbius actions

In this section we study weak tangents of compact metric spaces which admit a uniformly quasi-Möbius action for which the induced action $G \curvearrowright \mathrm{Tri}(Z)$ is cocompact. As the reader will notice, all the results in this section remain true under the weaker assumption that every triple of distinct points in Z can be blown up to a uniformly separated triple by a uniform quasi-Möbius homeomorphism of Z, i.e., an η-quasi-Möbius homeomorphism with η independent of the triple.

Lemma 5.1. *Suppose Z is a uniformly perfect compact metric space, and $G \curvearrowright Z$ is an η-quasi-Möbius action.*

(i) *Suppose that for each $k \in \mathbb{N}$ we are given a set D_k in a ball $B_k = B(p_k, R_k) \subseteq Z$ that is $(\epsilon_k R_k)$-dense in B_k, where $\epsilon_k > 0$, distinct points $x_k^1, x_k^2, x_3^k \in B(p_k, \lambda_k R_k)$, where $\lambda_k > 0$, with*

$$d_Z(x_k^i, x_k^j) > \delta_k R_k \quad \text{for} \quad i, j \in \{1,2,3\},\, i \neq j,$$

where $\delta_k > 0$, and group elements $g_k \in G$ such that for $y_k^i := g_k(x_k^i)$ we have

$$d_Z(y_k^i, y_k^j) > \delta' \quad \text{for} \quad i, j \in \{1,2,3\},\, i \neq j,$$

where $\delta' > 0$ is independent of k.

Let $D_k' := g_k(D_k)$, and suppose $\lambda_k \to 0$ for $k \to \infty$ and that the sequence (ϵ_k/δ_k^2) is bounded. Then

$$\operatorname{dist}_H(D_k', Z) \to 0 \quad \text{for} \quad k \to \infty.$$

(ii) *Suppose in addition that $G \curvearrowright \operatorname{Tri}(Z)$ is cocompact. If $U \subseteq Z$ is a nonempty open set, then there exists a sequence (g_k) in G such that*

$$\operatorname{diam}(Z \setminus g_k(U)) \to 0 \quad \text{for} \quad k \to \infty.$$

In plain words (i) essentially says that if we blow up a triple (x^1, x^2, x^3) that lies in a ball B to a uniformly separated triple, then a set D in B will be blown up to a rather dense set in Z, if the triple (x^1, x^2, x^3) lies deep inside B and its separation is much larger than $\operatorname{dist}_H(D, B)$.

Proof of (i). Let $d = d_Z$. Consider fixed $k \in \mathbb{N}$ and drop the subscript k for simplicity. The image of a point $z \in Z$ under $g = g_k$ will be denoted by $z' := g(z)$. Pick an arbitrary point in Z, and write it in the form $x' = g(x)$ where $x \in Z$. We have to find a point in D' close to x'.

Case 1: $x \in B(p, R)$. There is a point $y \in D \cap B$ with $d(x, y) \leq \epsilon R$. Since the minimal distance between the points x_1, x_2, x_3 is at least δR, we can find two of them, call them a and b, so that $d(y, a) \geq \delta R/2$ and $d(x, b) \geq \delta R/2$. Hence

$$\frac{d(x', y') d(a', b')}{d(x', b') d(a', y')} \leq \eta \left(\frac{d(x, y) d(a, b)}{d(x, b) d(a, y)} \right) \leq \eta(8\epsilon\lambda/\delta^2).$$

Rearranging factors, this implies that

$$d(x', y') \leq \operatorname{diam}(Z)^2 \eta(8\epsilon\lambda/\delta^2)/\delta' \leq C_1 \eta(C_2 \lambda).$$

The last expression becomes uniformly small as $\lambda \to 0$.

Case 2: $x \notin B(p,R)$. Since $\epsilon \lesssim \delta^2 \lesssim \lambda^2$, we may assume that $\epsilon > 0$ is small. Then by the uniform perfectness of Z and the (ϵR)-density of D in B, we can find a point $y \in D \cap B$ so that $d(y,p)/R$ is uniformly bounded away from zero, $d(y,p)/R \geq c_0 > 0$ say. Note that c_0 does not depend on k. We may assume that $\lambda < c_0/2 \leq 1/2$. Then setting $a = x^1$ and $b = x^2$ we get

$$\begin{aligned} \frac{d(x',y')d(a',b')}{d(x',b')d(a',y')} &\leq \eta\left(\frac{d(x,y)d(a,b)}{d(x,b)d(a,y)}\right) \\ &\leq \eta\left(\frac{4\lambda d(x,p)}{(d(x,p)-\lambda R)(c_0-\lambda)}\right) \\ &\leq \eta(16\lambda/c_0). \end{aligned}$$

Rearranging factors, this implies that

$$d(x',y') \leq \operatorname{diam}(Z)^2\eta(16\lambda/c_0)/\delta' \leq C_3\eta(C_4\lambda).$$

Again the last expression becomes uniformly small as $\lambda \to 0$.

Since $y' \in D'$, the first part of the lemma follows.

Proof of (ii). Let $B = B(p,R)$ be a ball in U with small radius $R \in (0,1/2]$. By the uniform perfectness of Z we can find a triple (x^1,x^2,x^3) of distinct points in $B(p,R^2)$ whose separation is comparable to R^2. Now use the cocompactness of $G \curvearrowright \operatorname{Tri}(Z)$ to find $g \in G$ mapping (x^1,x^2,x^3) to a uniformly separated triple.

Arguing as in Case 2 above, we find that whenever x' and y' are points in $Z \setminus g(B(p,R))$, then $d_Z(x',y') \lesssim \eta(CR)$. Hence $\operatorname{diam}(Z \setminus g(U)) \lesssim \eta(CR)$, and the claim follows by making R arbitrarily small. $\square$

Before we state the next lemma we recall that in Section 2 we have defined a metric $\hat{d}_p$ on the one-point compactification $\hat{X}$ of an unbounded locally compact pointed metric space (X,p) associated with the metric $d = d_X$ and the base point p.

Lemma 5.2. *Suppose Z is a compact metric space that is uniformly perfect and doubling, and $G \curvearrowright Z$ is a uniformly quasi-Möbius action for which the induced action $G \curvearrowright \operatorname{Tri}(Z)$ is cocompact.*

If $(S,p) \in \operatorname{WT}(Z)$, then there exist a quasi-Möbius homeomorphism $h\colon (\hat{S},\hat{d}_p) \to Z$. Moreover, $h|_S\colon S \to Z\setminus\{h(\infty)\}$ is also a quasi-Möbius homeomorphism.

In other words, up to quasi-Möbius homeomorphism the space Z is equivalent to the one-point compactification $\hat{S}$ of a weak tangent (S,p) of Z if we equip $\hat{S}$ with the canonical metric $\hat{d}_p$. Conversely, up to quasi-Möbius homeomorphism any weak tangent of Z is equivalent to Z with one point removed.

Proof. Note that as a weak tangent of a uniformly perfect doubling metric space, S is unbounded and proper.

From the definition of pointed space convergence it follows that for $k \in \mathbb{N}$ there exist subsets $\tilde{D}_k \subseteq B_S(p,k) \subseteq S$ that are $(1/k)$-dense in $B_S(p,k)$, numbers $\lambda_k > 0$ with $\lambda_k \to \infty$, points $p_k \in Z$, sets $D_k \subseteq B_{\lambda_k Z}(p_k,k) \subseteq \lambda_k Z$ that are $(1/k)$-dense in $B_{\lambda_k Z}(p_k,k)$ with respect to the metric $d_{\lambda_k Z} = \lambda_k d_Z$ and bijections $f_k \colon \tilde{D}_k \to D_k$ such that

$$\frac{1}{2} d_S(x,y) \leq \lambda_k d_Z(f_k(x), f_k(y)) \leq 2 d_S(x,y) \quad \text{for} \quad x,y \in \tilde{D}_k. \tag{5.3}$$

Moreover, it can be arranged that each set $\tilde{D}_k$ contains the points of a fixed triple $(q_1,q_2,q_3) \in \mathrm{Tri}(S)$.

Let $x_k^i := f_k(q_i)$ for $i \in \{1,2,3\}$ and $k \in \mathbb{N}$. Since the action $G \curvearrowright \mathrm{Tri}(Z)$ is cocompact, for $k \in \mathbb{N}$ we can find $g_k \in G$ such that the triples

$$(y_k^1, y_k^2, y_k^3) := g_k(x_k^1, x_k^2, x_k^3) \in \mathrm{Tri}(Z)$$

are uniformly separated.

The density condition for the sets D_k rephrased in terms of the metric d_Z says that D_k is (λ_k/k)-dense in $B_Z(p_k, \lambda_k k)$ with respect to d_Z. Moreover, in terms of the metric d_Z, the triple (x_k^1, x_k^2, x_k^3) has separation comparable to λ_k and is contained in a ball centered at p_k whose radius is also comparable to λ_k. It follows from Lemma 5.1 that for $D_k' := g_k(D_k)$ we have

$$\lim_{k\to\infty} \mathrm{dist}_H(D_k', Z) = 0, \tag{5.4}$$

where dist_H refers to the Hausdorff distance in Z.

The density condition for the sets $\tilde{D}_k \subseteq S \subseteq \hat{S}$ and the inequality (2.3) for the metric $\hat{d}_p$ imply that

$$\lim_{k\to\infty} \mathrm{dist}_H(\tilde{D}_k, \hat{S}) = 0, \tag{5.5}$$

where dist_H refers to the Hausdorff distance in $(\hat{S}, \hat{d}_p)$.

Consider the maps $h_k \colon (\tilde{D}_k, \hat{d}_p|_{\tilde{D}_k}) \to Z$ defined by $h_k(x) = g_k(f_k(x))$ for $x \in \tilde{D}_k$. Note that it follows from Lemma 2.2, inequality (5.3) and the fact that the action $G \curvearrowright Z$ is uniformly quasi-Möbius that the maps h_k are η-quasi-Möbius with η independent of k. Moreover, each map h_k maps the triple (q_1,q_2,q_3) to the uniformly separated triple (y_k^1, y_k^2, y_k^3). Finally, $D_k' = h_k(\tilde{D}_k)$ and so by (5.4) and (5.5) we can apply Lemma 2.1. It follows that the sequence (h_k) subconverges to a quasi-Möbius homeomorphism $h \colon (\hat{S}, \hat{d}_p) \to Z$.

The second part of the lemma follows by observing $h|_S\colon S \to Z \setminus \{f(\infty)\}$ is quasi-Möbius, since this map the composition of the maps $\mathrm{id}_S\colon S \to (S, \hat{d}_p|_S)$ which is quasi-Möbius by Lemma 2.2 and the map $h|_S\colon (S, \hat{d}_p|_S) \to Z \setminus \{h(\infty)\}$ which is quasi-Möbius by the first part of the proof. ☐

Lemma 5.6. *Suppose Z is a compact metric space that is uniformly perfect and doubling, and $G \curvearrowright Z$ is a uniformly quasi-Möbius action for which the induced action $G \curvearrowright \mathrm{Tri}(Z)$ is cocompact.*

If $\dim_{\mathrm{top}}(Z) = n \in \mathbb{N}$, then $\dim_{\mathrm{top}}(U) = n$ whenever U is a nonempty open subset of Z or of any weak tangent of Z.

Proof. If $U \subseteq Z$ is a nonempty open set, we can find a nonempty open set V with $\bar{V} \subseteq U$. By Lemma 5.1 there is a sequence (g_k) in G such that $\mathrm{diam}(Z \setminus g_k(\bar{V})) \to 0$ for $k \to \infty$. Hence the complement of $\bigcup_{k\in N} g_k(\bar{V})$ in Z can contain at most one point. Topological dimension is invariant under homeomorphisms, and and does not increase under a countable union of closed sets (cf. [9, Thm. II. 1]). So we get $\dim_{\mathrm{top}}(Z) \le \dim_{\mathrm{top}}(\bar{V}) \le \dim_{\mathrm{top}}(U) \le \dim_{\mathrm{top}}(Z)$.

If U is a nonempty open subset of any weak tangent S of Z, then U is also an open subset of the one-point compactification of S. Hence by Lemma 5.2, the set U is homeomorphic to a nonempty open subset of Z. Therefore $\dim_{\mathrm{top}}(U) = \dim_{\mathrm{top}}(Z)$ by the first part of the proof. ☐

Lemma 5.7. *Suppose X and Y are compact metric spaces that are uniformly perfect and doubling, and suppose $G \curvearrowright Z$ and $H \curvearrowright X$ are uniformly quasi-Möbius actions for which the induced actions $G \curvearrowright \mathrm{Tri}(X)$, $H \curvearrowright \mathrm{Tri}(Y)$ are cocompact.*

If there exist $S \in \mathrm{WT}(X)$ and $T \in \mathrm{WT}(Y)$ and a quasi-symmetric homeomorphism $f\colon S \to T$, then there exists a quasi-Möbius homeomorphism $g\colon X \to Y$.

So if X and Y have weak tangents that are quasi-symmetrically equivalent, then X and Y are equivalent up to a quasi-Möbius homeomorphism.

Proof. Let p and q be the base points in S and T, respectively, and consider the one-point compactifications $(\hat{S}, \hat{d}_p)$ and $(\hat{T}, \hat{d}_q)$. If we define $\hat{f}(x) = f(x)$ for $x \in X$, and $\hat{f}(\infty) = \infty$, then (2.3) implies that $\hat{f}\colon (\hat{S}, \hat{d}_p) \to (\hat{T}, \hat{d}_q)$ is a quasi-Möbius homeomorphism. Since $(\hat{S}, \hat{d}_p)$ is equivalent to X and $(\hat{T}, \hat{d}_q)$ is equivalent to Y up to quasi-Möbius homeomorphisms by Lemma 5.2, the claim follows. ☐

6. Proof of Theorem 1.1

Let Z and $G \curvearrowright Z$ be as in the statement of Theorem 1.1.

We are given that $\dim_{\mathrm{top}}(Z) = n$. This implies [9, Thm. III. 1] that there is a continuous map $f_0 \colon Z \to \mathbb{S}^n$ with a stable value $y \in \mathbb{S}^n$; in fact any continuous map $f_1 \colon Z \to \mathbb{S}^n$ for which $\mathrm{dist}(f_0, f_1)$ is sufficiently small will also have y as a stable value.

Every continuous function $g_0 \colon Z \to \mathbb{R}$ can be approximated by a Lipschitz function $g_1 \colon Z \to \mathbb{R}$ such that $\mathrm{dist}(g_0, g_1)$ is arbitrarily small. This standard fact can be established by using Lipschitz partitions of unity in Z subordinate to a cover of Z by small balls with controlled overlap. We apply this to the $n+1$ coordinate functions of the map $f_0 \colon Z \to \mathbb{S}^n \subseteq \mathbb{R}^{n+1}$ to obtain Lipschitz maps on Z which are arbitrarily close to f_0 and map Z into small neighborhoods of $\mathbb{S}^n$ in $\mathbb{R}^{n+1}$. Composing these maps with the radial projection from the origin in $\mathbb{R}^{n+1}$ to $\mathbb{S}^n$, we can find Lipschitz maps from Z into $\mathbb{S}^n$ arbitrarily close to f_0. In particular, there exists a Lipschitz map $f \colon Z \to \mathbb{S}^n$ such that y is a stable value of f. Then $\mathrm{Im}(f)$ is a neighborhood of y, and so $\mathcal{H}^n(\mathrm{Im}(f)) > 0$.

We now apply Proposition 4.1 to obtain a weak tangent S of Z, a weak tangent T of $\mathbb{S}^n$ and a regular map $\phi \colon S \to T$. Note that every weak tangent of $\mathbb{S}^n$ is isometric to $\mathbb{R}^n$, and so $T = \mathbb{R}^n$.

As we have seen, the fact that ϕ is regular implies that ϕ has bounded multiplicity. By Lemma 5.6, every nonempty open subset of S has topological dimension n. Therefore, by Theorem 3.4 (applied to the closure of the some bounded nonempty open set in S as the space X) there is a nonempty open subset $U \subseteq S$ such that $\psi := \phi|_U$ is a homeomorphism onto an open subset of $\mathbb{R}^n$. Shrinking the open set U if necessary, we may assume that ψ is a homeomorphism onto an open ball B in $\mathbb{R}^n$. Now Lemma 4.2 shows that ψ is bi-Lipschitz. Choosing $x \in U$ and setting $y := \psi(x)$ we are in the situation of Lemma 4.3. We conclude that S has a weak tangent bi-Lipschitz equivalent to a weak tangent of $\mathbb{R}^n$. Since the weak tangents of S are also weak tangents of Z, and all weak tangents of $\mathbb{R}^n$ are isometric to $\mathbb{R}^n$, we see that Z and $\mathbb{S}^n$ have bi-Lipschitz equivalent weak tangents. Since the group of Möbius transformations induces a uniformly quasi-Möbius action on $\mathbb{S}^n$ and a cocompact action on $\mathrm{Tri}(\mathbb{S}^n)$, Lemma 5.7 implies that there exists a quasi-Möbius homeomorphism $h \colon Z \to \mathbb{S}^n$. As a quasi-Möbius homeomorphism between bounded spaces, the map h will also be quasisymmetric. Conjugating the uniformly quasi-Möbius action $G \curvearrowright Z$ by h, we get a uniformly quasi-Möbius action $G \curvearrowright \mathbb{S}^n$ such that the induced action $G \curvearrowright \mathrm{Tri}(\mathbb{S}^n)$ is cocompact. By a result of Tukia [12, Cor. G(a)], this action is conjugate by a quasiconformal homeomorphism to an action by Möbius transformations. Since quasiconformal homeomorphisms of $\mathbb{S}^n$ onto itself are quasi-symmetric, Theorem 1.1 follows. $\square$

The method of proving Theorem 1.1 also leads to the following result.

Theorem 6.1. *Let $n \in \mathbb{N}$, and let Z be a compact, Ahlfors n-regular metric space of topological dimension n. Suppose every triple of distinct points in Z can be mapped to a uniformly separated triple by a uniform quasi-Möbius homeomorphism of Z. Then Z is quasi-symmetrically equivalent to the standard sphere $\mathbb{S}^n$.*

Proof. In the same way as in the proof of Theorem 1.1, we see that Z has a weak tangent bi-Lipschitz equivalent to $\mathbb{R}^n$. As we remarked in the beginning of Section 5, the results in this section remain true if the assumption on the group action is replaced by the assumption that every triple of distinct points in the space under consideration can be mapped to a uniformly separated triple by a uniform quasi-Möbius homeomorphism. So by the analog of Lemma 5.7, we again obtain a quasi-Möbius, and hence quasi-symmetric, homeomorphism $h\colon Z \to \mathbb{S}^n$. $\square$

This theorem justifies the remark in the introduction about the question of Heinonen and Semmes—recall that quasi-Möbius homeomorphisms of compact metric spaces are quasi-symmetric. We see that the three point homogeneity condition can be relaxed to a "cocompact on triples" condition, at the cost of requiring the homeomorphisms to be uniformly quasi-Möbius.

7. CAT(−1)-spaces and isometric group actions

We refer the reader to [7] for general background on Gromov hyperbolic spaces.

A metric space X is called *geodesic*, if any two points $x, y \in X$ can be joined by a geodesic segment in X, i.e., a curve whose length is equal to the distance of x and y. In the following we will always assume that X is proper and geodesic.

Let X be a Gromov hyperbolic space, and $\partial_\infty X$ be its boundary at infinity. There is a natural topology on $X \cup \partial_\infty X$ making this union compact. If $p \in X$, $a, b \in \partial_\infty X$, we let $[a, b]_p$ denote the Gromov product of $a, b \in \partial_\infty X$ with respect to the base point p. When $c > 0$ is sufficiently small, the function

$$d(a, b) := \exp(-c[a, b]_p) \tag{7.1}$$

is equivalent up to a multiplicative factor to a metric on $\partial_\infty X$; any two metrics of this type are quasi-symmetrically equivalent by the identity map. Fix one such metric on $\partial_\infty X$. If we denote the group of isometries of X by $\mathrm{Isom}(X)$, then we get an induced action $\mathrm{Isom}(X) \curvearrowright \partial_\infty X$ which is a uniformly quasi-Möbius action, [10, Prop. 4.5]. In fact, every quasi-isometry $f\colon X \to X$ induces an η-quasi-Möbius homeomorphism $\partial_\infty X \to \partial_\infty X$ where η depends only the parameters of the quasi-isometry and the hyperbolicity constant of X.

Now suppose that X is a $\mathrm{CAT}(-1)$-space (see [1] for more details on the topics discussed in the following). Then for every $p \in X$ we get a canonical metric on $\partial_\infty X$ as follows. For every point $a \in \partial_\infty X$, there is a unique geodesic ray $\overline{pa}$ starting at p whose asymptotic class represents a. Let $a, b \in \partial_\infty X$, and consider points $x \in \overline{pa}$, $y \in \overline{pb}$. Let $\Delta\tilde{p}\tilde{x}\tilde{y}$ be a comparison triangle (in the hyperbolic plane) for the triangle Δpxy, and let $\tilde{\angle}_p(x, y)$ denote the angle at $\tilde{p}$. When x and y tend to infinity along the rays $\overline{pa}$ and $\overline{pb}$, respectively, the comparison angle $\tilde{\angle}_p(x, y)$ has a limit, which we define to be the distance between a and b. This metric agrees up to a bounded factor with the expression in (7.1) when $c = 1$.

Suppose $G \curvearrowright X$ is an isometric action of a group on a $\mathrm{CAT}(-1)$-space X. If $x \in X$, then we denote its orbit under G by

$$Gx := \{g(x) : g \in G\}.$$

The *limit set* $\Lambda(G) \subseteq \partial_\infty X$ of G is by definition the set of all accumulation points of an orbit Gx on $\partial_\infty X$. This set is independent of $x \in X$. The group action $G \curvearrowright X$ is called *properly discontinuous* if

$$\{g \in G : g(K) \cap K \neq \emptyset\}$$

is finite for every compact subset K of X.

A subset $Y \subseteq X$ is *quasi-convex* if there is a constant C such that any geodesic segment with endpoints in Y lies in the C-neighborhood of Y. The action $G \curvearrowright X$ is *quasi-convex cocompact* if there is a G-invariant quasi-convex subset $Y \subseteq X$ on which G acts with compact quotient Y/G. The group G is quasi-convex cocompact if and only if all orbits Gx are quasi-convex.

We will need the following result due to Bourdon [2, 0.3 Théorème ($\mathbb{H}^n$ case)].

Theorem 7.2. *Let $n \geq 2$, G be a group, and X a $\mathrm{CAT}(-1)$-space. Suppose we have isometric group actions $G \curvearrowright X$ and $G \curvearrowright \mathbb{H}^{n+1}$ which are properly discontinuous. Suppose that $G \curvearrowright X$ is quasi-convex cocompact and $G \curvearrowright \mathbb{H}^{n+1}$ is cocompact. If the Hausdorff dimension of*

$\Lambda(G) \subseteq \partial_\infty X$ is equal to n, then there exists a G-equivariant isometry of $\mathbb{H}^{n+1}$ onto a convex, G-invariant set $Y \subseteq X$.

Actually, Bourdon proved this under the additional assumption that the group action $G \curvearrowright \mathbb{H}^{n+1}$ is faithful. In this case G is isomorphic to a uniform lattice in $\mathrm{Isom}(\mathbb{H}^{n+1})$. The proof of the above more general version is the same as the proof of his original result.

Proof of Theorem 1.2. Consider the induced actions $G \curvearrowright \Lambda(G)$ and $G \curvearrowright \mathrm{Tri}(\Lambda(G))$. Since $G \curvearrowright X$ is isometric, $G \curvearrowright \Lambda(G)$ is uniformly quasi-Möbius. Since the action $G \curvearrowright X$ is properly discontinuous, the same is true for $G \curvearrowright \mathrm{Tri}(\Lambda(G))$. Moreover, since $G \curvearrowright X$ is quasi-convex cocompact, $G \curvearrowright \mathrm{Tri}(\Lambda(G))$ is cocompact.

Since the Hausdorff dimension of $\Lambda(G)$ is n, this space will actually be Ahlfors n-regular (cf. [5, Section 7]). Now n is also the topological dimension of $\Lambda(G)$ by assumption. By Theorem 1.1, the action $G \curvearrowright \Lambda(G)$ is quasi-symmetrically conjugate to an action $G \curvearrowright \mathbb{S}^n$ by Möbius transformations. The action $G \curvearrowright \mathrm{Tri}(\mathbb{S}^n)$ is properly discontinuous and cocompact. This implies that there is a properly discontinuous, cocompact, and isometric action $G \curvearrowright \mathbb{H}^{n+1}$ which induces the action $G \curvearrowright \mathbb{S}^n = \partial_\infty \mathbb{H}^{n+1}$. Since $n \geq 2$ we can apply Bourdon's theorem, and conclude that there exists a G-equivariant isometric embedding of $\mathbb{H}^{n+1}$ onto a convex, G-invariant set $Y \subseteq X$ on which G acts cocompactly. The result follows. $\square$

As the proof shows, $n \geq 2$ is only used in the last step. In particular, even in the case $n = 1$ we can still conclude that $\Lambda(G)$ is quasi-symmetrically equivalent to $\mathbb{S}^1$, and that there is an action $G \curvearrowright \mathbb{H}^2$ which isometric, properly discontinuous and cocompact.

References

[1] M. Bourdon, *Structure conforme au bord et flot géodésique d'un CAT(-1)-espace*, Enseign. Math. (2), 41 (1995), pp. 63–102.

[2] ———, *Sur le birapport au bord des CAT(-1)-espaces*, Inst. Hautes Études Sci. Publ. Math., 83 (1996), pp. 95–104.

[3] R. Bowen, *Hausdorff dimension of quasicircles*, Inst. Hautes Études Sci. Publ. Math., 50 (1979), pp. 11–25.

[4] D. Burago, Y. Burago, and S. Ivanov, *A course in metric geometry*, American Mathematical Society, Providence, RI, 2001.

[5] M. Coornaert, *Mesures de Patterson-Sullivan sur le bord d'un espace hyperbolique au sens de Gromov*, Pacific J. Math., 159 (1993), pp. 241–270.

[6] G. David and S. Semmes, *Fractured fractals and broken dreams*, Oxford University Press, New York, 1997.

[7] É. Ghys and P. de la Harpe, eds., *Sur les groupes hyperboliques d'après Mikhael Gromov*, Birkhäuser, Boston, MA, 1990.

[8] J. Heinonen and S. Semmes, *Thirty-three yes or no questions about mappings, measures, and metrics*, Conform. Geom. Dyn., 1 (1997), pp. 1–12 (electronic).
[9] J. Nagata, *Modern Dimension Theory*, Heldermann Verlag, Berlin, 1983.
[10] F. Paulin, *Un groupe hyperbolique est déterminé par son bord*, J. London Math. Soc. (2), 54 (1996), pp. 50–74.
[11] D. Sullivan, *Discrete conformal groups and measurable dynamics*, Bull. Amer. Math. Soc. (N.S.), 6 (1982), pp. 57–73.
[12] P. Tukia, *On quasiconformal groups*, J. Analyse Math., 46 (1986), pp. 318–346.
[13] J. Väisälä, *Quasi-Möbius maps*, J. Analyse Math., 44 (1984/85), pp. 218–234.
[14] C. Yue, *Dimension and rigidity of quasi-Fuchsian representations*, Ann. of Math. (2), 143 (1996), pp. 331–355.

Department of Mathematics, University of Michigan, Ann Arbor, MI, 48109-1109

E-mail address: `mbonk@math.lsa.umich.edu`

Department of Mathematics, University of Michigan, Ann Arbor, MI, 48109-1109

E-mail address: `bkleiner@math.lsa.umich.edu`