UCLA UCLA Electronic Theses and Dissertations

Title

A coarse entropy-rigidity theorem and discrete length-volume inequalities

Permalink https://escholarship.org/uc/item/3c04z060

Author Kinneberg, Kyle Edward

Publication Date 2014

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

A coarse entropy-rigidity theorem and discrete length-volume inequalities

A dissertation submitted in partial satisfaction of the requirements for the degree Doctor of Philosophy in Mathematics

by

Kyle Edward Kinneberg

© Copyright by Kyle Edward Kinneberg 2014

Abstract of the Dissertation

A coarse entropy-rigidity theorem and discrete length-volume inequalities

by

Kyle Edward Kinneberg

Doctor of Philosophy in Mathematics University of California, Los Angeles, 2014 Professor Mario Bonk, Chair

In [5], M. Bonk and B. Kleiner proved a rigidity theorem for expanding quasi-Möbius group actions on Ahlfors *n*-regular metric spaces with topological dimension *n*. This led naturally to a rigidity result for quasi-convex geometric actions on CAT(-1)-spaces that can be seen as a metric analog to the "entropy-rigidity" theorems of U. Hamenstädt [31] and M. Bourdon [10]. Building on the ideas developed in [5], we establish a rigidity theorem for certain expanding quasi-Möbius group actions on spaces with different metric and topological dimensions. This is motivated by a corresponding entropy-rigidity result in the coarse geometric setting.

Our analysis of these "fractal" metric spaces depends heavily on a combinatorial inequality that relates volume to lengths of curves within the space. We extend such inequalities to a broader metric setting and obtain discrete analogs of some results due to W. Derrick [23, 24]. In the process, we shed light on a related question of Y. Burago and V. Zalgaller about pseudometrics on the n-dimensional unit cube. The dissertation of Kyle Edward Kinneberg is approved.

Sheldon Smith John B. Garnett Mario Bonk, Committee Chair

University of California, Los Angeles

2014

To Celine ...

... and all those who will join us.

TABLE OF CONTENTS

1	Intr	oducti	ion	1
2	Rig	idity fo	or quasi-Möbius actions on fractal metric spaces	6
	2.1	Introd	luction	6
		2.1.1	Extending Mostow rigidity	7
		2.1.2	Toward a metric setting	8
		2.1.3	Rigidity on fractal spaces	11
	2.2	Defini	tions and Notation	15
	2.3	Strong	gly Quasi-Möbius Group Actions	18
	2.4	De-sno	owflaking	23
		2.4.1	The definition of d_{new}	24
		2.4.2	The upper bound	27
		2.4.3	A discrete length-volume inequality	29
		2.4.4	Proof of Lemma 2.17	35
	2.5	Proof	of Theorem 2.5	39
	2.6	Entroj	py Rigidity in Coarse Geometry	42
		2.6.1	The hyperbolic boundary	44
		2.6.2	Geometric actions on hyperbolic metric spaces	47
		2.6.3	Proof of Theorem 2.8	53
	2.7	A clos	er look at the case $n = 2 \dots \dots$	58
		2.7.1	Proof of Proposition 2.40	63
		2.7.2	Further remarks	67

3	Len	gth-volume inequalities revisited	68			
	3.1	Introduction	68			
	3.2	Notation and preliminaries	72			
	3.3	A topological length-volume inequality for cubes	76			
	3.4	Lower volume bounds in metric spaces	85			
	3.5	Some analogous considerations for simplices	89			
Bibliography						

Acknowledgments

As with any large undertaking, there are many who contributed to the completion of this project—certainly more than can be mentioned here. Among those who have contributed most directly, I ought first to thank Mario Bonk for his deep mathematical insights and his devotion to communicating those insights clearly. It has been a pleasure to study under him. I am also grateful to many others in the UCLA Mathematics department who have increased my understanding of the subject, especially John Garnett, Guy David, and Zhiqiang Li.

Beyond the mathematical content of this thesis, I am deeply indebted to my family and friends for their constant love, which continues to manifest things that *the eye has not seen* nor my mind has fully understood. In particular, I thank my parents, my brother, and Celine.

The second chapter of this thesis, excluding the final section, will appear in the Journal of Differential Geometry in 2014, published by International Press.

The author was partially supported by NSF grants DMS-1162471 and DMS-1344959 during the completion of this thesis.

Vita

2009–2014 Teaching assistant, Mathematics department, UCLA.

2011 M.Sc., Mathematics, UCLA.

2009 B.A., Mathematics, Claremont McKenna College, Summa Cum Laude.

PUBLICATIONS AND PRESENTATIONS

Rigidity for quasi-Möbius actions on fractal metric spaces, preprint, to appear in J. Differential Geom. (2014), available at arXiv:1308.0639v1.

Compact widths in metric trees, with A. G. Aksoy, Function spaces IX, Banach Center Publ., vol. 92, Polish Acad. Sci. Inst. Math., Warsaw, 2011, pp. 15–25.

"Recognizing snowflakes and an application to hyperbolic groups", Geometry Seminar, University of Jyväskylä, March 2014.

"The Loewner equation and Hölder geometry", Analysis and Dynamical Systems Seminar, Université d'Orléans, March 2014.

"Rigidity for quasiconformal actions on fractal spaces", Dynamical Systems Seminar, Stony Brook University, November 2013. "Entropy rigidity in coarse geometry", Session on Analysis, Dynamics, and Geometry in and around Teichmüller Spaces, AMS Sectional Meeting, Iowa State University, April 2013.

"Mostow rigidity: Further perspectives", Introductory workshop for *Interactions between* Analysis and Geometry, IPAM, March 2013.

CHAPTER 1

Introduction

In the past several decades, quasiconformal techniques have proven to be indispensable tools in various areas of analysis and geometry. Originally studied in the Euclidean and Riemannian settings, quasiconformal mappings played a prominent role, particularly, in the development of hyperbolic geometry. Teichmüller theory for hyperbolic surfaces is saturated with quasiconformal relationships (indeed, the Teichmüller metric is built explicitly using quasiconformal homeomorphisms); many rigidity results for hyperbolic manifolds of dimension ≥ 3 (including G. Mostow's classic theorem [47]) crucially use analytic properties of quasiconformal mappings; and quasiconformal group actions are closely related to ergodic properties of the geodesic flow on hyperbolic manifolds (as in the work of D. Sullivan [57]). These connections between quasiconformal analysis and hyperbolic structures often have roots in the following heuristic: hyperbolic geometry "at infinity" becomes conformal geometry on a sphere of one dimension less; and perturbations of hyperbolic structures often lead to quasiconformal relationships between the associated conformal structures.

More recently, the theory of quasiconformal mappings has been extended to general metric spaces, motivated by the need for various types of "quasi-conformal" analysis on spaces with no a priori smooth structure [41]. A primary example of this is the study of large-scale geometry of hyperbolic groups and the metric properties of their boundaries at infinity. In turn, such developments have led to interesting investigations into the intrinsic geometry of metric spaces, especially those with fractal properties. Due to the lack of local regularity, tools from discrete geometry have become essential to the metric analysis of these types of spaces. Let us back up slightly, to a more classical observation of W. Thurston regarding the relationship between circle packings in the plane and the Riemann mapping theorem. Let Ω be a simply connected region in \mathbb{C} , and approximate Ω by a hexagonal circle packing of mesh size 1/n, where each circle lies within Ω . Let γ_n be a Jordan curve in Ω that surrounds this packing and is tangent to the "boundary" circles. The Koebe-Andreev theorem then gives a circle packing of the unit disk \mathbb{D} , unique up to Möbius transformations, that is isomorphic to the original hexagonal circle packing of Ω (and also preserves boundary tangencies with respect to γ_n and $\partial \mathbb{D}$). This isomorphism induces a map $f_n \colon \mathbb{D} \to \Omega$, and Thurston conjectured that these f_n (properly normalized) converge uniformly on compact subsets of \mathbb{D} to a conformal map.

B. Rodin and D. Sullivan [51] confirmed Thurston's conjecture, using a mix of combinatorial arguments, properties of quasiconformal mappings, and hyperbolic geometry. In particular, they show that the maps f_n are K-quasiconformal, for some uniform K, and hence converge to a K-quasiconformal mapping, f. Moreover, the combinatorics ensures that f sends small equilateral triangles to triangles that become more and more equilateral as $n \to \infty$. Consequently, f must actually be 1-quasiconformal, hence conformal.

The work of Rodin and Sullivan initiated many investigations into more precise relationships between circle packings, conformal geometry, and quasiconformal geometry. Notably, it was proved by K. Stephenson [56] and by Z.-X. He and O. Schramm [34] that the convergence of Thurston's algorithm does not, in fact, require the circle packings of Ω to be hexagonal. These topics also play a crucial role in the work of M. Bonk and B. Kleiner [4] on quasisymmetric uniformization of 2-dimensional metric spheres (see Section 2.7 for further discussion about this result). Here, there is no a priori conformal structure at all; essentially one obtains it through the Koebe-Andreev theorem.

An important setting where similar difficulties arise is the asymptotic geometry of hyperbolic groups. A finitely generated group Γ is *hyperbolic* if there is a finite, symmetric generating set for which the corresponding Cayley graph, endowed with the word metric, has the following property: there is $\delta > 0$ such that each side of every geodesic triangle is contained in the δ -neighborhood of the union of the other two sides. This is a "coarse" notion of negative curvature that is well-suited to the study of large-scale geometric properties; in particular, it cannot see anything at scale $\ll \delta$. For example, free groups and fundamental groups of closed negatively-curved manifolds are hyperbolic.

We defer to the next chapter further discussion about the the foundational theory for hyperbolic groups and instead focus on one aspect of this rich subject. To every hyperbolic group, there is a natural boundary at infinity, denoted $\partial_{\infty}\Gamma$, which is equipped with a family of *visual metrics* that are closely related to the asymptotic geometry of Γ . Understanding the relationships between topological or metric properties of $\partial_{\infty}\Gamma$ and algebraic properties of Γ has been a central goal in geometric group theory for decades. The following is a significant open problem in the field.

Conjecture (J. Cannon [15]). A hyperbolic group Γ is a finite extension of a cocompact lattice in Isom(\mathbb{H}^3) if and only if $\partial_{\infty}\Gamma$ is homeomorphic to \mathbb{S}^2 .

Using a classical theorem of D. Sullivan [57] and P. Tukia [59], Cannon's conjecture is equivalent to the following statement: if Γ is a hyperbolic group with $\partial_{\infty}\Gamma$ homeomorphic to \mathbb{S}^2 , then $\partial_{\infty}\Gamma$ is quasisymmetric to \mathbb{S}^2 , when equipped with a visual metric. Thus, the conjecture is reduced to a quasiconformal uniformization problem, where one must somehow build the conformal structure of \mathbb{S}^2 , starting with the metric space $\partial_{\infty}\Gamma$.

In general, $\partial_{\infty}\Gamma$ will exhibit self-similar, fractal properties. On the one hand, it has many symmetries, coming from a natural action by Γ itself that mimics the way Möbius transformations act on \mathbb{S}^2 . On the other hand, its metric dimension is usually strictly larger than its topological dimension; for example, it may fail to contain *any* curves of finite length. In such a setting, discrete methods are indispensable. Not surprisingly, Cannon's original approach to his conjecture was to construct a "combinatorial conformal structure" on $\partial_{\infty}\Gamma$ using finite subdivision rules (see, for example, [17–19]). In this context, he established a combinatorial Riemann mapping theorem [14], which is similar in spirit to the work on circle packings that we discussed above. More recently, M. Bourdon and B. Kleiner introduced the *combinatorial Loewner property* (or CLP) as a tool to study self-similar metric spaces, with particular attention to the boundaries of hyperbolic groups [11]. The CLP is a discrete version of the classical *Loewner property* that has been fundamental to the development of analysis on metric measure spaces [35, Chapters 8 and 9]. This latter property ensures that any two compact sets in the metric space can be joined by a sufficiently rich family of (rectifiable) paths. Its combinatorial version is similar but uses discrete chains of points in place of paths. Using the CLP, Bourdon and Kleiner were able to verify Cannon's conjecture for Coxeter groups.

We should also mention, in the context of Cannon's work, that combinatorial methods, especially those connected to finite subdivision rules, have been very fruitful in the study of conformal dynamics [8, 16, 29].

The topics we study in this thesis revolve around the theme of discrete analysis on metric spaces that is motivated by coarse hyperbolic geometry, especially the metric geometry of boundaries of hyperbolic groups. This connection will be made more clear and more precise in Chapter 2, as will the connection between coarse hyperbolicity and the classical notion of negative curvature in Riemannian manifolds. Our main result in that chapter is Theorem 2.8—an entropy-rigidity theorem for hyperbolic groups that is analogous to a result of U. Hamenstädt [31] for negatively-curved manifolds. Most of the work needed to establish Theorem 2.8, however, will be a metric analysis of the canonical action of Γ on $\partial_{\infty}\Gamma$. In the particular setting that is relevant for Theorem 2.8, the metric and topological dimensions of $\partial_{\infty}\Gamma$ differ; moreover, $\partial_{\infty}\Gamma$ contains no rectifiable curves. As a result, our analysis must be discrete, and it relies heavily on a combinatorial inequality (found in Proposition 2.19) that relates volumes to lengths in *n*-dimensional cubes.

The statement in Proposition 2.19 is a combinatorial analog of similar length-volume inequalities established by W. Derrick [23, 24]. Viewed in this context, it cries out loudly for generalization, and we pacify this cry in Chapter 3. There, we prove discrete length-volume inequalities that parallel the results of Derrick and extend Proposition 2.19. In the process, we shed some light on a related question of Y. Burago and V. Zalgaller concerning

pseudometrics on the n-dimensional cube (see Corollary 3.14). These considerations then compel us to explore lower volume bounds in more general metric settings.

CHAPTER 2

Rigidity for quasi-Möbius actions on fractal metric spaces

2.1 Introduction

Let (M^n, g) be a compact Riemannian manifold of dimension n. If we assume that M is locally symmetric and negatively curved, then its universal cover is isometric to \mathbb{H}_F^k —one of the hyperbolic spaces defined over $F = \mathbb{R}, \mathbb{C}$, the quaternions, or the octonians (in the last case, only for k = 2, which corresponds to real dimension n = 16). We can therefore identify (M^n, g) with the quotient \mathbb{H}_F^k/Γ , where $\Gamma = \pi_1(M)$ acts on \mathbb{H}_F^k by deck transformations. A natural question then arises: is this hyperbolic structure uniquely determined by the topology of M?

G. Mostow's classic rigidity theorem [47] gives an affirmative answer to this question in dimensions $n \geq 3$. More specifically, he proves that if two locally symmetric compact manifolds, both with maximal sectional curvature -1, have isomorphic fundamental groups, then they are isometric. The curvature assumption here is simply a scaling normalization. In other words, the topology of a locally symmetric compact manifold determines its metric structure, up to scaling.

It is important to note, of course, that there is no analogous theorem for surfaces. Indeed, a compact surface of genus $g \ge 2$ has a (6g - 6)-dimensional moduli space of hyperbolic metrics (i.e., locally symmetric metrics of constant curvature -1). Such surfaces therefore have many metric deformations which would be ruled out by a rigidity theorem.

2.1.1 Extending Mostow rigidity

Let us now turn our attention to a different question: how can one determine when a negatively curved manifold is locally symmetric? A significant amount of work in this direction, much of it from the early and mid 1990s, sought to find a symmetric structure in manifolds that were extremal for certain metric quantities—volume, curvature bounds, geodesic lengths, entropy, etc. Most relevant for us here is the entropy-rigidity theorem of U. Hamenstädt [31].

Before stating this result, we must first define entropy. Let (M^n, g) be a compact Riemannian manifold and let (\tilde{M}, \tilde{g}) be its universal Riemannian cover with metric \tilde{g} . Let B(p, R)denote the ball of radius R centered at $p \in \tilde{M}$, and let $\operatorname{Vol}_{\tilde{g}} B(p, R)$ be the volume of this ball. We call

$$h_{\rm vol}(g) = \lim_{R \to \infty} \frac{\log\left(\operatorname{Vol}_{\tilde{g}} B(p, R)\right)}{R}$$

the volume entropy of g; this limit exists and is independent of the choice $p \in \tilde{M}$. For example, if (M^n, g) is hyperbolic, then (\tilde{M}, \tilde{g}) can be identified with real hyperbolic space \mathbb{H}^n of constant sectional curvature -1. Consequently,

$$\operatorname{Vol}_{\tilde{q}} B(p, R) = \operatorname{Vol}_{\mathbb{H}^n} B(p, R) \approx e^{(n-1)R},$$

so that h(g) = n - 1.

The following relationship indicates why this volume-growth quantity is considered to be a type of entropy. Let $h_{top}(g)$ denote the *topological entropy* of the geodesic flow on the unit tangent bundle of (M^n, g) (see [45, Section 3] for definitions). For general compact manifolds, Manning [45] showed that

$$h_{\text{top}}(g) \ge h_{\text{vol}}(g),$$

and if (M^n, g) has non-positive sectional curvature, then equality holds. As we will concern ourselves only with compact manifolds of negative sectional curvature, we can set

$$h(g) = h_{top}(g) = h_{vol}(g)$$

from now on and refer to it simply as the *entropy* of g.

Theorem 2.1 (Hamenstädt [31]). Let (M^n, g_0) be a locally symmetric compact manifold with maximal sectional curvature -1 and $n \ge 3$. Let g be another Riemannian metric on M, also with maximal sectional curvature -1. Then $h(g) \ge h(g_0)$, and equality holds if and only if g is locally symmetric. In particular, equality holds if and only if (M^n, g) is isometric to (M^n, g_0) .

In other words, the locally symmetric structures on M are precisely the minima of the entropy functional, at least among metrics suitably normalized by curvature. Note also that the "in particular" statement in this theorem follows from Mostow rigidity.

From Hamenstädt's theorem, there are various directions in which one may proceed (see, for example, the survey [55] on rigidity theory). Remaining in the Riemannian setting, we can ask if there are other pairs of normalizations and metric quantities for which rigidity theorems can be obtained. For example, suppose that (M, g_0) is a locally symmetric compact manifold of dimension ≥ 3 with unit volume and let g be another metric on M with unit volume. Then $h(g) \geq h(g_0)$ and equality holds precisely when (M, g) and (M, g_0) are isometric. A more general version of this was established by G. Besson, G. Courtois, and S. Gallot [1], along with several consequential rigidity statements, but the general theme is that the locally symmetric metrics on manifolds related to (M, g_0) can be identified by two quantities: volume and entropy. Incidentally, the methods used in their paper give a constructive proof of Mostow's original result by exhibiting the desired isometry.

2.1.2 Toward a metric setting

A different direction one may take (and the direction we wish to push further in this thesis) is to extend Hamenstädt's theorem to metric geometry. To motivate the comparison between the Riemannian and metric settings, let (M^n, g) be as in Theorem 2.1, and let $\Gamma = \pi_1(M)$ be its fundamental group. Also, let (X, d) be its Riemannian universal cover with metric d. Of course d is a Riemannian metric itself, but as we move away from the Riemannian setting, we want to think of d simply as a distance function. The negative curvature in (M^n, g) , which guarantees negative curvature in (X, d) as well, allows one to define an ideal boundary: the collection of asymptotic classes of geodesic rays emanating from a fixed base-point. Moreover, this boundary has a canonical metric structure that is closely related to the asymptotic geometry of X. For example, if g is hyperbolic, then its universal cover is, once again, the real hyperbolic space \mathbb{H}^n , whose ideal boundary is the Euclidean sphere \mathbb{S}^{n-1} .

The isometric action of Γ on (X, d) passes naturally to an action on the ideal boundary. Here the entropy of g plays an important role, as h(g) is the Hausdorff dimension of the canonical metric on the boundary. Equality of h(g) and $h(g_0)$ therefore guarantees that the boundary associated to g has metric properties similar to those of the boundary associated to g_0 , which is much better understood.

Let us make the comparison between Riemannian and metric geometry more explicit. The universal cover (X, d) has sectional curvature at most -1, so it satisfies the CAT(-1) condition [58, Théorème 9]. Recall that a geodesic metric space is called CAT(-1) if its geodesic triangles are thinner than their comparison triangles in the real hyperbolic plane. Moreover, the action of the fundamental group Γ on X is isometric, properly discontinuous, and cocompact. We will refer to such actions as *geometric* actions from now on. Actually, to deal with more general situations, it will be convenient to weaken the cocompactness property to *quasi-convex cocompactness*: there is a quasi-convex subset $Y \subset X$ on which Γ acts cocompactly. We call such actions *quasi-convex geometric*; see Section 2.6 for formal definitions.

For a CAT(-1)-space X, one can define a boundary at infinity, which we denote by $\partial_{\infty} X$. As with the ideal boundary, it will be a topological space with a canonical metric structure (again, see Section 2.6 for details). Let $\Lambda(\Gamma)$ be the limit set of Γ in $\partial_{\infty} X$. If the action is cocompact then $\Lambda(\Gamma) = \partial_{\infty} X$, but in general the limit set can be much smaller than the whole boundary. However, its Hausdorff dimension has a familiar form [9, Théorème 2.7.4]:

$$\dim_H \Lambda(\Gamma) = \limsup_{R \to \infty} \frac{\log(N(R))}{R}$$

where $p \in X$ is any point and $N(R) = \#\{\Gamma p \cap B_X(p, R)\}\$ is the number of points in the orbit Γp that lie at distance at most R from p. This Hausdorff dimension is therefore a metric analog of the entropy we considered earlier.

In this context, M. Bourdon [10] proved the following generalization of Theorem 2.1.

Theorem 2.2 (Bourdon [10]). Let $\Gamma = \pi_1(M^n, g_0)$ be the fundamental group of a locally symmetric compact manifold of maximal sectional curvature -1 and dimension $n \ge 3$. Suppose that Γ acts quasi-convex geometrically on a CAT(-1)-space X. Let S be the universal Riemannian cover of (M^n, g_0) . Then

$$\dim_H \Lambda(\Gamma) \ge \dim_H \partial_\infty S,$$

and equality holds if and only if there is an isometric embedding $F: S \to X$, equivariant with respect to the natural action of Γ on S, whose extension to the boundary has $F(\partial_{\infty}S) = \Lambda(\Gamma)$.

Although this theorem certainly points in the direction of metric geometry, it does not strictly fall in this category. Indeed, the restriction of Γ to fundamental groups of locally symmetric spaces and the use of dim_H $\partial_{\infty}S$ in the rigidity inequality seem to place this result, in some sense, between Riemannian geometry and metric geometry.

In [5], M. Bonk and B. Kleiner extended the real-hyperbolic version of Bourdon's theorem to the metric setting. By real-hyperbolic, we mean the case that (M^n, g_0) has constant sectional curvature -1, so that $S = \mathbb{H}^n$. Recall that $\partial_{\infty} \mathbb{H}^n = \mathbb{S}^{n-1}$, which has Hausdorff dimension n - 1.

Theorem 2.3 (Bonk–Kleiner [5,6]). Suppose that a group Γ acts quasi-convex geometrically on a CAT(-1) metric space X. Let $n \ge 1$ be the topological dimension of $\Lambda(\Gamma)$. Then

$$\dim_H \Lambda(\Gamma) \ge n,$$

and equality holds if and only if Γ acts geometrically on an isometric copy of \mathbb{H}^{n+1} in X.

The assertion $\dim_H \Lambda(\Gamma) \ge n$ here is nothing special, as the Hausdorff dimension of any metric space is bounded from below by its topological dimension [38, Chapter 7]. Let us focus on the case of equality, then, and briefly describe the method of proof. As in the rigidity theorems discussed above, the argument relies on a quasiconformal analysis of the limit set $\Lambda(\Gamma)$. The isometric action of Γ on X naturally passes to an action on $\Lambda(\Gamma)$ by uniformly quasi-Möbius maps. As Γ acts cocompactly on a quasi-convex subset of X, the induced action on $\Lambda(\Gamma)$ will be *cocompact on triples*: any three distinct points in the limit set can be uniformly separated by applying an element of the group. This property should be viewed as a type of expanding dynamics on $\Lambda(\Gamma)$. It also allows us to conclude that $\Lambda(\Gamma)$ is *Ahlfors regular* of dimension n: the n-dimensional Hausdorff measure of any metric ball B(x, r) in the limit set is $\approx r^n$ (for $0 \le r \le \text{diam } \Lambda(\Gamma)$).

The following theorem is the main result in [5].

Theorem 2.4 (Bonk–Kleiner [5]). Let Z be a compact, Ahlfors n-regular metric space with topological dimension $n \ge 1$. Suppose that $\Gamma \curvearrowright Z$ is a uniformly quasi-Möbius group action that is cocompact on triples. Then $\Gamma \curvearrowright Z$ is quasisymmetrically conjugate to an action of Γ on \mathbb{S}^n by Möbius transformations.

As Möbius transformations can be extended naturally to isometries of \mathbb{H}^{n+1} , we obtain a geometric action of Γ on \mathbb{H}^{n+1} . If $n \geq 2$, this puts us in the setting of Bourdon's theorem, which we apply to conclude that Γ acts cocompactly on an isometric copy of \mathbb{H}^{n+1} in X.

Actually, it turns out that appealing to Bourdon's theorem is not necessary. An alternative argument is given in [6], and it works just as well in the case that n = 1.

2.1.3 Rigidity on fractal spaces

Following Bonk and Kleiner, this thesis is primarily concerned with rigidity of expanding quasi-Möbius group actions. Indeed, results in this setting often lead to rigidity theorems that are more geometric. Reconsidering, then, Theorem 2.4, it is natural to wonder what one can say if the Hausdorff and topological dimensions differ.

A large collection of such examples are boundaries of Gromov hyperbolic groups equipped with a visual metric. In many important cases, the boundary is topologically a sphere; and always, it will be Ahlfors regular. Generally, though, the metric dimension is strictly larger than its topological dimension. In the case where the boundary is homeomorphic to S^2 , it is conjectured that there *exists* an Ahlfors regular metric of dimension 2, but this is a difficult problem (see [2, Section 5] for this formulation of Cannon's conjecture).

It is therefore of interest to obtain rigidity results for quasi-Möbius group actions on fractal metric spaces—spaces in which the metric dimension differs from the topological dimension. This is the general objective in the present thesis. In moving from such a broad goal to concrete theorems, we have kept an eye on applications to coarse hyperbolic geometry, which is a relevant setting for the study of Gromov hyperbolic groups. As a consequence, our main theorem will lead, via the work in [5], to an entropy-rigidity result for geometric group actions on Gromov hyperbolic metric spaces with an asymptotic upper curvature bound. Naturally, this can be seen as a "coarse" analog of the CAT(-1) rigidity theorem in [5] and therefore also as an analog of Hamenstädt's theorem and of Bourdon's theorem (in the real-hyperbolic cases).

The precise statement of our main result is the following. We will discuss terminology and notation in subsequent sections, but let us make one important remark now. Rather than considering general quasi-Möbius group actions, we restrict our attention to those that are *strongly quasi-Möbius*. In particular, each group element will act as a bi-Lipschitz homeomorphism. See Definition 2.10 for a formal definition.

Theorem 2.5. Let $n \in \mathbb{N}$, $0 < \epsilon \leq 1$, and let Z = (Z, d) be a compact metric space, homeomorphic to \mathbb{S}^n , and Ahlfors regular of dimension n/ϵ . Suppose that $\Gamma \curvearrowright Z$ is a strongly quasi-Möbius action that is cocompact on triples. Assume, moreover, that Z satisfies the following discrete length property:

each
$$\delta$$
-path between two points x, y has length $\geq c \left(\frac{d(x,y)}{\delta}\right)^{1/\epsilon}$. (2.1)

Then there is a metric d_{new} on Z satisfying

$$C^{-1}d(x,y)^{1/\epsilon} \le d_{new}(x,y) \le Cd(x,y)^{1/\epsilon}$$

for some $C \geq 1$ and a bi-Lipschitz homeomorphism between (Z, d_{new}) and \mathbb{S}^n . Moreover, if

 $n \geq 2$, then this map can be taken to conjugate the action of Γ on Z to an action on \mathbb{S}^n by Möbius transformations.

Remark 2.6. The assumption that Z is homeomorphic to \mathbb{S}^n can be replaced by the assumption that Z is an *n*-dimensional manifold. Indeed, in this case, the expanding behavior of the group action forces Z to be a topological *n*-sphere. See, for example, the proof of Theorem 4.4 in [40].

Remark 2.7. Recall that if ρ is a metric on Z, then ρ^{ϵ} is also a metric whenever $0 < \epsilon \leq 1$. The metric spaces (Z, ρ^{ϵ}) are typically called "snowflakes" of (Z, ρ) , in reference to the standard construction of the von Koch snowflake. In Theorem 2.5, we go in the opposite direction, "de-snowflaking" the original metric d on Z to a metric d_{new} with better regularity.

When the metric dimension and the topological dimension of Z coincide (i.e., if $\epsilon = 1$), the results in [5] give a bi-Lipschitz homeomorphism between Z and Sⁿ. Once these dimensions differ, relationships between the metric structures of Z and Sⁿ are more delicate. While Ahlfors regularity gives good control on volume, and the strongly quasi-Möbius action provides robust self-similarity structure in Z, additional assumptions are needed to obtain rigidity statements. We impose the condition (2.1) because, in the case where Z is the boundary of a hyperbolic metric space X, it arises naturally from upper curvature bounds on X.

In concise terms, the discrete length condition (2.1) is strong enough that it forces (Z, d) to be a "snowflake" of \mathbb{S}^n . Once we de-snowflake, we are able to pass almost directly through the theorem of Bonk and Kleiner. It is natural to ask, then, if there are weaker conditions one can place on Z that still guarantee it is, say, quasisymmetrically equivalent to \mathbb{S}^n . This would be of significant interest, in particular for n = 2.

As we suggested above, Theorem 2.5 leads to a rigidity theorem in a coarse geometric setting. The objects considered here are (Gromov) hyperbolic metric spaces with an appropriate asymptotic upper curvature bound. These curvature bounds, denoted by $AC_u(\kappa)$, were introduced by M. Bonk and T. Foertsch [3] as a coarse analog to the $CAT(\kappa)$ conditions. We will discuss this further in Section 2.6, but for now we only mention that $AC_u(-1)$ is an appropriate replacement for CAT(-1).

For hyperbolic metric spaces X, even with asymptotic upper curvature bounds, there is no canonical Hausdorff dimension of the boundary, as there was for CAT(-1)-spaces. Indeed, the visual metrics on $\partial_{\infty}X$ form a Hölder class, and there is not a natural choice of a bi-Lipschitz sub-class. Thus, to formulate an entropy-rigidity statement here, we must look back inside X and use the coarse version of volume entropy—the same quantity that bridged the results of Hamenstädt and Bourdon. Namely, if X is a hyperbolic metric space and Γ acts on X, the exponential growth rate of the action is

$$e(\Gamma) = \limsup_{R \to \infty} \frac{\log(N(R))}{R}$$

where $N(R) = \#\{\Gamma p \cap B_X(p, R)\}$ is the number of points in an orbit Γp of distance at most R from p. Once again, the limit is independent of $p \in X$. We then have the corresponding coarse rigidity theorem.

Theorem 2.8. Let X be a proper, geodesic, Gromov hyperbolic metric space, and let $\Gamma \curvearrowright X$ be a quasi-convex geometric group action. Suppose that $\Lambda(\Gamma)$ is homeomorphic to \mathbb{S}^n , with $n \ge 2$, and that there is an orbit Γp that is $\operatorname{AC}_u(-1)$. Then $e(\Gamma) \ge n$ and equality holds if and only if there is a rough isometry $\Phi \colon \mathbb{H}^{n+1} \to \Gamma p$ that is roughly equivariant with respect to a geometric action of Γ on \mathbb{H}^{n+1} .

Remark 2.9. Again, the assumption that $\Lambda(\Gamma)$ is a topological sphere can be weakened; it suffices to assume that $\Lambda(\Gamma)$ contains an open subset homeomorphic to \mathbb{R}^n . Indeed, this will imply that $\Lambda(\Gamma)$ is homeomorphic to \mathbb{S}^n (cf. Theorem 4.4 in [40]). We leave as an open question, though, whether it suffices to assume only that the topological dimension of $\Lambda(\Gamma)$ is n.

Moreover, one should ask about the case n = 1. We do not know if the conclusion in Theorem 2.8 holds in this case as well. See, however, Remark 2.36 at the end of Section 2.6, where we discuss what can be said in its place. This chapter is organized as follows. In Section 2.2, we will introduce the necessary definitions, terminology, and background for the consideration and proof of Theorem 2.5. Section 2.3 will be devoted to a slightly technical study of strongly quasi-Möbius group actions that will reveal some properties relevant for a "de-snowflaking" result. In Section 2.4, we will state and prove this general de-snowflaking theorem, which forms the heart of the proof of Theorem 2.5. In Section 2.5 we finish the proof of Theorem 2.5 by de-snowflaking and applying quantitative versions of theorems from [5] and [59]. We will also, of course, need to verify these quantitative versions. Lastly, in Section 2.6 we will prove Theorem 2.8 after discussing in more detail the terminology used in its statement.

2.2 Definitions and Notation

Let (Z, d) be a metric space. Occasionally, we will write d_Z for the metric on Z when this needs to be specified. If $x \in Z$ and r > 0, then we use

$$B(x,r) = \{y \in Z : d(x,y) < r\} \text{ and } \overline{B}(x,r) = \{y \in Z : d(x,y) \le r\}$$

to denote, respectively, the open and closed metric balls of radius r about x.

If $x_1, x_2, x_3, x_4 \in Z$ are distinct points, we define their *(metric) cross-ratio* as

$$[x_1, x_2, x_3, x_4] = \frac{d(x_1, x_3)d(x_2, x_4)}{d(x_1, x_4)d(x_2, x_3)}$$

We are interested in maps between metric spaces that distort cross-ratios in a controlled manner. To make this precise, let $\eta: [0, \infty) \to [0, \infty)$ be a homeomorphism. Then a homeomorphism $f: X \to Y$ is called η -quasi-Möbius if

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

for all distinct four-tuples $x_1, x_2, x_3, x_4 \in X$. Note that this definition makes sense for injective f as well, but we will be concerned only with homeomorphisms in what follows.

A second class of maps that arise naturally in quasiconformal geometry are the quasisymmetric maps, which distort relative distances by a controlled amount. A homeomorphism $f: X \to Y$ is η -quasisymmetric if

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

for all triples x_1, x_2, x_3 of distinct points in X.

The quasi-Möbius and quasisymmetric conditions are closely related, though there are subtle differences. For example, every η -quasisymmetric map is $\tilde{\eta}$ -quasi-Möbius, where $\tilde{\eta}$ depends only on η . Conversely, if X and Y are bounded, then each individual η -quasi-Möbius map will be $\tilde{\eta}$ -quasisymmetric for *some* $\tilde{\eta}$, but in general there is no quantitative relationship between η and $\tilde{\eta}$.

In this chapter, we are mostly interested in studying metric spaces on which there is a group action by maps belonging to a particular function class. In such a context, the quasi-Möbius and quasisymmetry conditions are very different. As quasi-Möbius maps are the weaker of these two types, it makes sense to focus on these actions. This choice is further motivated by the following fact about hyperbolic groups (which occupy center stage in studying the geometry of hyperbolic metric spaces). If G is a hyperbolic group and $\partial_{\infty}G$ is its boundary (i.e., the Gromov boundary of the Cayley graph of G with respect to a fixed finite generating set) equipped with a visual metric, then the isometric action of G on its Cayley graph by translations extends to an action on $\partial_{\infty}G$. Moreover, there is η for which each $g \in G$ acts as an η -quasi-Möbius map. Actually, something stronger is true: we can take η to be linear (see Section 2.6 for more details).

Quasi-Möbius maps with a linear distortion function will play an important role in our analysis. Thus, we give them a name.

Definition 2.10. A homeomorphism $f: X \to Y$ is called *strongly quasi-Möbius* if there is $C \ge 1$ for which

$$[f(x_1), f(x_2), f(x_3), f(x_4)] \le C[x_1, x_2, x_3, x_4]$$

whenever $x_1, x_2, x_3, x_4 \in X$ are distinct.

Each strongly quasi-Möbius map between bounded metric spaces is actually bi-Lipschitz:

there is a constant $C' \ge 1$ for which

$$\frac{1}{C'}d_X(x_1, x_2) \le d_Y(f(x_1), f(x_2)) \le C'd_X(x_1, x_2)$$

for all $x_1, x_2 \in X$ (see Remark 2.14). But again, the relationship between C and C' is not quantitative. We will study group actions by strongly quasi-Möbius maps in much greater detail in subsequent sections.

Most of the group actions we encounter here will be of an expanding type, in the following sense.

Definition 2.11. An action of a group Γ on a metric space (Z, d) is said to be *cocompact* on triples if there is $\delta > 0$ such that for every triple $x_1, x_2, x_3 \in Z$ of distinct points, there is a map $g \in \Gamma$ for which $d(gx_i, gx_j) \ge \delta$ if $i \ne j$.

It should be no surprise that this assumption is again motivated by the geometry of hyperbolic groups: the action of a hyperbolic group on its boundary (equipped with a visual metric) is indeed cocompact on triples. More generally, the expanding behavior of a group action, combined with an (assumed) regularity of maps in the group, often translates into self-similarity properties of the metric space. See Lemma 2.13 for a particular manifestation of this principle.

A final metric property that will commonly undergird our spaces is a standard type of volume regularity.

Definition 2.12. A compact metric space (Z, d) is Ahlfors α -regular (or Ahlfors regular of dimension $\alpha > 0$) if there is a Borel measure μ on Z and a constant $C \ge 1$ so that

$$\frac{1}{C}r^{\alpha} \le \mu(\overline{B}(x,r)) \le Cr^{\alpha} \tag{2.2}$$

for all $x \in Z$ and $0 < r \leq \text{diam } Z$.

Using standard covering arguments, it is not difficult to show that Z is Ahlfors α -regular if and only if (2.2) holds with μ replaced by Hausdorff measure of dimension α . In subsequent sections, we will frequently encounter the *n*-dimensional sphere \mathbb{S}^n . Unless otherwise specified, we give it the chordal metric—the restriction of the Euclidean metric when \mathbb{S}^n is viewed as the unit sphere in \mathbb{R}^{n+1} . However, every metric property of \mathbb{S}^n that we consider will be preserved under a bi-Lipschitz change of coordinates. Thus, any metric that is bi-Lipschitz equivalent to the chordal metric would work just as well.

Lastly, it will be convenient for us to suppress non-essential multiplicative constants in many inequalities. For quantities A and B that depend on some collection of input variables, we write $A \leq B$ to indicate that there is a constant C, independent of these variables, for which $A \leq CB$. When possible confusion could arise, we will indicate which data C may depend on. For example, the bi-Lipschitz condition can be expressed simply as

$$d_X(x_1, x_2) \lesssim d_Y(f(x_1), f(x_2)) \lesssim d_X(x_1, x_2),$$

where the constants are uniform over all $x_1, x_2 \in X$.

2.3 Strongly Quasi-Möbius Group Actions

We now focus our attention on strongly quasi-Möbius maps—those with a linear distortion function. Such maps tend to behave even more like traditional Möbius functions than general quasi-Möbius maps do. For example, each strongly quasi-Möbius homeomorphism between bounded metric spaces is bi-Lipschitz (cf. Remark 2.14).

Strongly quasi-Möbius maps are particularly important when they come in a group with uniform distortion constant. We will say that a group action $\Gamma \curvearrowright Z$ on a metric space Z is strongly quasi-Möbius if there is a constant $C \ge 1$ for which every $g \in \Gamma$ is an η -quasi-Möbius homeomorphism with $\eta(t) = Ct$.

The following lemma tells us that a strongly quasi-Möbius group action that is cocompact on triples gives Z locally self-similar structure: each ball can be blown up to a uniform scale by a homeomorphism that is essentially a scaling on that ball. See [13, Section 2.3] for a general discussion of local self-similarity in metric spaces. See also Lemma 5.1 in [5] for a statement similar to ours, albeit in a slightly different context.

Lemma 2.13. Let (Z, d) be a compact, connected metric space with at least two points, and let $\Gamma \curvearrowright Z$ be a strongly quasi-Möbius group action that is cocompact on triples. For fixed $p \in Z$, $0 < r \le \text{diam } Z$, and $L \ge 2$, let N = B(p, r) be a "near" set and $F = Z \setminus B(p, Lr)$ be a "far" set with respect to p. Then there is a map $g \in \Gamma$ satisfying the following:

(i)
$$r \cdot d(x,y) \lesssim d(gx,gy) \lesssim (1/r) \cdot d(x,y)$$
 for all $x, y \in Z$,

- (ii) $(1/r) \cdot d(x,y) \lesssim d(gx,gy) \lesssim (1/r) \cdot d(x,y)$ for all $x, y \in N$,
- (iii) there exists c > 0 such that $B(gx, c) \subset gN$ for each $x \in B(p, r/2)$,
- (iv) diam $gF \lesssim 1/L$.

Here, the implicit constants and c depend only on diam Z, the constant δ in Definition 2.11, and C from the strongly quasi-Möbius condition. In particular, they do not depend on p, r, or L.

Observe that property (ii) tells us that on N, the map g is basically a scaling by 1/r, in that g blows up B(p,r) to a uniform scale. Property (iii) guarantees that gN will contain large balls around images of points that are well-inside N; or, to put it negatively, points outside of N cannot get mapped nearby the images of points well within N. Property (iv) shows that if we take L to be large, we can map the "far" set F to something negligible.

Proof. Given p, r, and L, we first choose three points that we wish to δ -separate. Let $x_1 = p$, and choose x_2 to be a point for which $d(x_1, x_2) = r/2$. Then choose x_3 so that $d(x_1, x_3) = r/4$. Such points x_2 and x_3 exist by the assumption that Z is compact and connected. Take $g \in \Gamma$ so that gx_1, gx_2, gx_3 have pairwise distances at least δ . Now that we have chosen g, we use x' to refer to the image of points x under g.

(i) Let $x, y \in Z$ with $x \neq y$. Then there are $i, j \in \{1, 2, 3\}$ for which $d(x, x_i) \geq r/8$ and $d(x, x_j) \geq r/8$. Of these, either $d(y, x_i) \geq r/8$ or $d(y, x_j) \geq r/8$; without loss of generality,

say $d(y, x_i) \ge r/8$. We then have

$$\frac{d(x',y')d(x'_i,x'_j)}{d(x',x'_j)d(y',x'_i)} \lesssim \frac{d(x,y)d(x_i,x_j)}{d(x,x_j)d(y,x_i)} \lesssim \frac{d(x,y)\cdot r}{r/8\cdot r/8} \lesssim \frac{d(x,y)}{r}$$

and so

$$d(x',y') \lesssim \frac{d(x',x'_j)d(y',x'_i)}{d(x'_i,x'_j)} \cdot \frac{d(x,y)}{r} \lesssim \frac{1}{\delta} \cdot \frac{d(x,y)}{r} \lesssim \frac{d(x,y)}{r}$$

which is the second inequality in (i). Recall that the implicit constant is allowed to depend on diam Z and on δ .

For the first inequality, take $i, j \in \{1, 2, 3\}$ for which $d(x', x'_i) \ge \delta/2$ and $d(x', x'_j) \ge \delta/2$. Then either $d(y', x'_i) \ge \delta/2$ or $d(y', x'_j) \ge \delta/2$; without loss of generality, say $d(y', x'_i) \ge \delta/2$. Then

$$\frac{d(x',x'_j)d(y',x'_i)}{d(x',y')d(x'_i,x'_j)} \lesssim \frac{d(x,x_j)d(y,x_i)}{d(x,y)d(x_i,x_j)} \lesssim \frac{1}{d(x,y)\cdot r},$$

and so

$$d(x,y) \lesssim \frac{1}{r} \cdot \frac{d(x',y')d(x'_i,x'_j)}{d(x',x'_j)d(y',x'_i)} \lesssim \frac{d(x',y')}{r},$$

as desired.

Remark 2.14. The same reasoning can be used to show that any strongly quasi-Möbius map between bounded metric spaces is necessarily bi-Lipschitz. Indeed, notice that after choosing g, the arguments in (i) use only four facts: pairwise distances between x_1 , x_2 , and x_3 are $\geq r/4$; pairwise distances between x'_1 , x'_2 , and x'_3 are $\geq \delta$; the domain and image of g are both bounded; and g is strongly quasi-Möbius. In general, then, if f is a strongly quasi-Möbius homeomorphism between bounded metric spaces, choose distinct points x_1 , x_2 , and x_3 in the domain. These four facts will hold for some $r, \delta > 0$, so we can conclude that f is bi-Lipschitz.

(ii) Let $x, y \in N$. The second inequality here is directly from (i). For the first inequality, take i, j for which $d(x', x'_j), d(y', x'_i) \ge \delta/2$ as we did above. Then

$$\frac{d(x', x'_j)d(y', x'_i)}{d(x', y')d(x'_i, x'_j)} \lesssim \frac{d(x, x_j)d(y, x_i)}{d(x, y)d(x_i, x_j)} \lesssim \frac{2r \cdot 2r}{d(x, y) \cdot r/4} \lesssim \frac{r}{d(x, y)},$$

and so

$$d(x,y) \lesssim r \cdot \frac{d(x',y')d(x'_i,x'_j)}{d(x',x'_j)d(y',x'_i)} \lesssim r \cdot d(x',y'),$$

as claimed.

(iii) Fix $x \in B(p, r/2)$ and note that if $y \in Z \setminus N$, then $d(x, y) \ge r/2$. Taking $i \in \{2, 3\}$ for which $d(y', x'_i) \ge \delta/2$, we have

$$\frac{d(x',p')d(y',x'_i)}{d(x',y')d(p',x'_i)} \lesssim \frac{d(x,p)d(y,x_i)}{d(x,y)d(p,x_i)} \lesssim \frac{d(x,p)}{r} \cdot \frac{d(y,x_i)}{d(x,y)}.$$

As $d(x,y) \ge r/2$, we also have $d(y,x_i) \le d(x,y) + d(x,x_i) \le d(x,y) + 2r \le 5d(x,y)$, and so

$$\frac{d(x',p')d(y',x'_i)}{d(x',y')d(p',x'_i)} \lesssim \frac{d(x,p)}{r}$$

Thus,

$$d(x',y') \gtrsim \frac{d(x',p')d(y',x'_i)}{d(p',x'_i)} \cdot \frac{r}{d(x,p)} \gtrsim \frac{d(x',p')}{d(x,p)} \cdot r \gtrsim 1$$

by the bounds we established in (ii). Let c be the implicit constant in this last inequality. Then $B(x',c) \subset gN$ by the fact that g is surjective.

(iv) We may, of course, assume that B(p, Lr) is not all of Z. Then fix a point $x \in B(p, 2Lr) \setminus B(p, Lr)$. We claim that gF is contained in a small ball centered at x'. Indeed, let $y \in F$, and observe that

$$\begin{aligned} &d(x,y) \le d(x,x_1) + d(y,x_1) \le 2Lr + d(y,x_1) \le 3d(y,x_1), \\ &d(x_1,x_2) = r/2, \\ &d(x,x_1) \ge Lr, \\ &d(y,x_2) \ge d(y,x_1) - d(x_1,x_2) \ge d(y,x_1) - r. \end{aligned}$$

Thus, we have

$$\frac{d(x',y')d(x'_1,x'_2)}{d(x',x'_1)d(y',x'_2)} \lesssim \frac{d(x,y)d(x_1,x_2)}{d(x,x_1)d(y,x_2)} \lesssim \frac{d(y,x_1) \cdot r}{Lr \cdot (d(y,x_1)-r)}$$

As $d(y, x_1) \ge Lr$ and the function $t \mapsto t/(t-r)$ is decreasing for t > r, we obtain

$$\frac{d(y,x_1)\cdot r}{Lr\cdot (d(y,x_1)-r)} \le \frac{Lr}{L(Lr-r)} \le \frac{1}{L}$$

and so

$$d(x',y') \lesssim \frac{1}{L} \cdot \frac{d(x',x_1')d(y',x_2')}{d(x_1',x_2')} \lesssim \frac{1}{L}.$$

Consequently, gF is contained in a ball of radius $\lesssim L^{-1}$ centered at x', as needed.

The previous lemma gives us a good understanding of the type of expanding behavior found in a strongly quasi-Möbius group action, when it acts cocompactly on triples. This use of a group element to "blow up" a ball to a uniform scale is sometimes called a "conformal elevator" (see, for example, [29]). One can therefore view Lemma 2.13 as a particular type of conformal elevator that comes with a strongly quasi-Möbius group action. This elevator will be essential in the proof of our de-snowflaking result, coming in the next section.

Actually, in the proof of that result, it is the conformal elevator itself (rather than the strongly quasi-Möbius action generating it) that will be important. In order to work in greater generality, we make the following definition.

Definition 2.15. A metric space (Z, d) admits a *conformal elevator* if there exists a constant $C \ge 1$ and a function $\omega \colon (0, \infty) \to (0, \infty)$ with $\omega(t) \to 0$ as $t \to 0$ such that, for every choice of $p \in Z$, $0 < r \le \text{diam } Z$, and $\lambda \ge 2$, there is a homeomorphism $g \colon Z \to Z$ with the following properties:

- (i) $d(gx, gy) \leq Cd(x, y)/r$ for all $x, y \in B(p, \lambda r)$,
- (ii) $C^{-1}d(x,y)/r \le d(gx,gy)$ for all $x, y \in B(p,r)$,
- (iii) $B(gx, 1/C) \subset g(B(p, r))$ for all $x \in B(p, r/C)$,
- (iv) diam $(Z \setminus g(B(p, \lambda r))) \le \omega(1/\lambda)$.

The conclusions in Lemma 2.13 tell us that if Z admits a strongly quasi-Möbius group action that is cocompact on triples, then it admits a conformal elevator. We now turn our attention toward using the conformal elevator to de-snowflake a metric space.

2.4 De-snowflaking

This section is devoted to establishing the following proposition, which provides quantitative conditions under which a metric space can be de-snowflaked by a particular amount.

Proposition 2.16. Fix $n \in \mathbb{N}$ and $0 < \epsilon < 1$. Let (Z, d) be a metric space with the following properties:

- (i) Z is homeomorphic to \mathbb{S}^n ,
- (ii) Z admits a conformal elevator, in the sense of Definition 2.15,
- (iii) every δ -separated set in Z has size at most $C\delta^{-n/\epsilon}$,
- (iv) every discrete δ -path from x to y in Z has length at least $(1/C) \cdot (d(x,y)/\delta)^{1/\epsilon}.$

Then there is a metric d_{new} on Z satisfying

$$d(x,y)^{1/\epsilon} \lesssim d_{new}(x,y) \lesssim d(x,y)^{1/\epsilon}$$
(2.3)

where the implicit constant depends only on the data from assumptions (i)-(iv).

A " δ -separated set" is simply a set of points for which pairwise distances are at least δ . Also, by a "discrete δ -path from x to y" we mean a chain of points $x = z_0, z_1, \ldots, z_l = y$ in Z with $d(z_i, z_{i-1}) \leq \delta$. The length of such a chain is l; notice that this is one less than the number of points in the chain.

Some remarks on the assumptions in Proposition 2.16 are in order. The first condition gives Z non-trivial topological structure and guarantees that the metric structure of (Z, d)on large scales is similar to that of the standard sphere. The conformal elevator, as we have already said, gives us a way of moving from small scales to a uniformly large scale. The third assumption should be thought of as a volume condition; for example, it is easily implied by Ahlfors n/ϵ -regularity. In fact, our condition is similar to assuming that the Assound dimension of (Z, d) is n/ϵ . See [13, Chapter 9] for a discussion on various notions of metric dimension. Lastly, condition (iv) is exactly the discrete length assumption that appears in Theorem 2.5. It basically functions as a 1-dimensional metric condition. Thus, the essential ingredients to our de-snowflaking result are topological regularity, metric selfsimilarity, upper bounds on volume, and lower bounds on 1-dimensional metric structure.

Regarding the conclusion of the proposition, the "data from assumptions (i)–(iv)" include the following: the parameters ϵ and n; the diameter of Z; the constant C from conditions (ii), (iii), and (iv); the function ω from the conformal elevator; and the modulus of continuity of a fixed homeomorphism between Z and \mathbb{S}^n . We will also refer to these as the "data associated to Z."

The proof of Proposition 2.16 will proceed as follows. We begin by fixing, for each length scale $e^{-\epsilon k}$, a cover of Z by metric balls. For $x, y \in Z$, the smallest number of these balls needed to join x and y provides a "fuzzy" notion of distance at scale $e^{-\epsilon k}$. After a proper normalization of this fuzzy distance, we let k tend to infinity to obtain the metric d_{new} .

The lower bound in (2.3) will follow almost directly from the discrete length condition in (iv). The upper bound is more complicated, but in it we will see a nice interplay between the topological and metric structures of Z, which are linked together by the existence of the conformal elevator.

2.4.1 The definition of d_{new}

Fix notation as in the statement of the proposition. In particular, we let C be a constant large enough so that conditions (iii) and (iv) hold, as well as the conditions from the definition of a conformal elevator. We also let ω be the function associated with the conformal elevator on Z. This notation will remain fixed throughout the proof.

By scaling the metric d, we may assume for simplicity that diam Z = 1. Indeed, the implicit constant in the desired conclusion is allowed to depend on the diameter, so we lose no generality.

For each $k \in \mathbb{N}$, fix a maximal $e^{-\epsilon k}$ -separated set in Z and call it P_k . This also will remain fixed throughout the proof. It is not difficult to see that if P is any $e^{-\epsilon k}$ -separated set contained in a ball B(p, r), with $0 < r \le 1$, then

$$\#P \lesssim r^{n/\epsilon} \cdot e^{nk}.\tag{2.4}$$

Indeed, this follows from the volume bound in (iii) after applying the conformal elevator for $p, r, and \lambda = 2$. In particular, if $r = e^{-\epsilon m}$, then such a set has size $\leq e^{n(k-m)}$.

By maximality of P_k , we mean with respect to set inclusion. This is, of course, equivalent to

$$Z = \bigcup_{x \in P_k} B(x, e^{-\epsilon k})$$

for each k. It will be more convenient to work with balls of twice this radius, and so we refer to

$$\{B(x, 2e^{-\epsilon k}) : x \in P_k\}$$

as the set of k-balls. For notational simplicity, we may abbreviate $B(x, 2e^{-\epsilon k})$ by $B_k(x)$ when $x \in P_k$. Often, the center-point x of $B_k(x)$ is not important. As a result, we will usually denote k-balls simply by B or B_i , e.g., when dealing with a chain of such balls. In these cases, k will be understood from the context.

We first observe that for each k, the set of k-balls has controlled overlap, in that each k-ball intersects at most a uniformly bounded number of k-balls. Indeed, if $B_k(x)$ is a k-ball and $B_k(x_i)$, $1 \leq i \leq m$, are those that intersect $B_k(x)$ non-trivially, then the collection $\{x_1, \ldots, x_m\}$ is an $e^{-\epsilon k}$ -separated set in the ball $B(x, 4e^{-\epsilon k})$. By (2.4) above, we get

$$m \lesssim (4e^{-\epsilon k})^{n/\epsilon} \cdot e^{nk} \lesssim 1,$$
 (2.5)

where the implicit constant is allowed to depend on n and ϵ .

A sequence of k-balls B_1, B_2, \ldots, B_l with $B_i \cap B_{i+1} \neq \emptyset$ is called a k-ball chain. We say that such a chain connects two points x and y if $x \in B_1$ and $y \in B_l$. Observe that, as B_i may not be a connected set itself, chains may not be connected topologically. This will pose no problem for our later analysis, though.
The length of a k-ball chain is simply the number of balls appearing in it, counted with multiplicity. For each k, let

$$d_k(x, y) = (\text{length of shortest } k\text{-ball chain connecting } x \text{ and } y) \cdot e^{-k}.$$

The normalization by e^{-k} is appropriate; indeed, each k-ball has diameter approximately $e^{-\epsilon k}$ with respect to d, so its diameter with respect to the sought-after d_{new} should be approximately e^{-k} . Note that d_k is not actually a metric; for each x we have $d_k(x, x) = e^{-k}$. But d_k is symmetric and the triangle inequality clearly holds.

We now set

$$d_{new}(x,y) = \limsup_{k \to \infty} d_k(x,y).$$

It is not immediate that this is a metric either. It is certainly symmetric, has $d_{new}(x, x) = 0$ for all x, and satisfies the triangle inequality. The inequalities $0 < d_{new}(x, y) < \infty$ for $x \neq y$ will, however, be a consequence of proving that d_{new} is bi-Lipschitz equivalent to $d^{1/\epsilon}$. Thus, we wish to show that for each pair of distinct points $x, y \in Z$,

$$d(x,y)^{1/\epsilon} \lesssim d_k(x,y) \lesssim d(x,y)^{1/\epsilon}$$

for k large enough, where the implicit constants depend only on the data associated to Z.

As we mentioned before, the lower bound will be an easy consequence of assumption (iv) in the statement of the proposition—the lower bound on the length of discrete "paths" between points. We quickly verify this.

Let $x, y \in Z$ be distinct and let B_1, \ldots, B_l be a k-ball chain connecting x and y. Then $x \in B_1$ and $y \in B_l$. Choose $x_i \in B_i \cap B_{i+1}$ for each $1 \le i \le l-1$ and consider the discrete path

$$x = x_0, x_1, \dots, x_{l-1}, x_l = y$$

from x to y. Observe that $d(x_i, x_{i+1}) \leq 4e^{-\epsilon k}$, as diam $B_i \leq 4e^{-\epsilon k}$. Consequently,

$$l \gtrsim \left(\frac{d(x,y)}{4e^{-\epsilon k}}\right)^{1/\epsilon},$$

so that $l \gtrsim d(x, y)^{1/\epsilon} \cdot e^k$ where the implicit constant depends only on C and ϵ . This gives immediately that $d_k(x, y) \gtrsim d(x, y)^{1/\epsilon}$, as desired.

We now turn to the upper bound, which is much more subtle. To obtain it, we will use the lower bound, along with a discrete length-volume inequality for cubes.

2.4.2 The upper bound

We begin by stating the crucial lemma, which will almost immediately give the upper bound when applied iteratively. This method of proof was motivated by a similar argument in [61], where the author also sought to establish upper bounds on "tile" chains connecting two points.

Lemma 2.17. Let $x, y \in Z$ and $m \in \mathbb{N}$ with $d(x, y) \leq e^{-\epsilon(m-1)}$. Then for each $k \geq m$, there is a k-ball chain connecting $B(x, e^{-\epsilon m})$ and $B(y, e^{-\epsilon m})$ of length at most $C'e^{k-m}$. Here, C' depends only on the data associated to Z.

As should be clear, we say that a k-ball chain B_1, \ldots, B_ℓ connects two sets A and B if $B_1 \cap A$ and $B_\ell \cap B$ are non-empty. We will first see how this lemma implies the desired upper bound.

Proof of the Upper Bound. Suppose that Lemma 2.17 holds. Fix distinct points x and y in Z, and choose $h \in \mathbb{N}$ so that $e^{-\epsilon h} < d(x, y) \leq e^{-\epsilon(h-1)}$. Recall that we have normalized diam Z = 1. Also fix k > h; it is helpful to think of k being very large relative to h. We will temporarily use B_z^m to denote the ball $B(z, e^{-\epsilon m})$ for $z \in X$ and $m \in \mathbb{N}$. This should not be confused with the shorthand notation we used earlier for k-balls.

By the lemma, there is a k-ball chain connecting B_x^h and B_y^h of length at most $C'e^{k-h}$. This gives us points $x = x_{1,0}, x_{1,1}, x_{1,2}, x_{1,3} = y$ where $x_{1,0}, x_{1,1} \in B_x^h$ and $x_{1,2}, x_{1,3} \in B_y^h$, and $x_{1,1}$ is connected to $x_{1,2}$ by a k-ball chain of length at most $C'e^{k-h}$.

We now iterate this process. Observe that $d(x_{1,0}, x_{1,1}) \leq e^{-\epsilon h}$ so that we can apply the lemma again to obtain a k-ball chain of length at most $C'e^{k-(h+1)}$ connecting $B_{x_{1,0}}^{h+1}$ and $B_{x_{1,1}}^{h+1}$. This gives us points $x = x_{2,0}, x_{2,1}, x_{2,2}, x_{2,3} = x_{1,1}$ where $x_{2,0}, x_{2,1} \in B_{x_{1,0}}^{h+1}$ and $x_{2,2}, x_{2,3} \in B_{x_{1,1}}^{h+1}$, and $x_{2,1}$ is connected to $x_{2,2}$ by a k-ball chain of length at most $C'e^{k-(h+1)}$. Of course, we do a similar process to the pair of points $x_{1,2}$ and $x_{1,3}$.

The *m*-th step in this process (for $1 \le m \le k-h$) proceeds as follows. From the (m-1)-th step, we have 2^m points

$$x = x_{m-1,0}, x_{m-1,1}, \dots, x_{m-1,2^m-1} = y$$

satisfying $d(x_{m-1,i}, x_{m-1,i+1}) \leq e^{-\epsilon(h+m-2)}$ for each even integer $i \in \{0, \ldots, 2^m - 2\}$. Moreover, there is a previously-constructed k-ball chain connecting $x_{m-1,i+1}$ to $x_{m-1,i+2}$.

It will be convenient to rename these 2^m points so that they appear in the *m*-th step. To do this, we let

$$x_{m,j} = \begin{cases} x_{m-1,j/2}, & \text{if } j \equiv 0 \mod 4\\ \\ x_{m-1,(j-1)/2}, & \text{if } j \equiv 3 \mod 4 \end{cases}$$

for $0 \le j \le 2^{m+1} - 1$. We do not yet define the points corresponding to $j \equiv 1, 2 \mod 4$ because we still have to find them.

To this end, observe that $d(x_{m,j}, x_{m,j+3}) \leq e^{-\epsilon(h+m-2)}$ for each $j \equiv 0 \mod 4$. Thus, applying Lemma 2.17 to these points, we find a k-ball chain of length at most $C'e^{k-(h+m-1)}$ connecting the balls $B_{x_{m,j}}^{h+m-1}$ and $B_{x_{m,j+3}}^{h+m-1}$. We can therefore choose points $x_{m,j+1} \in B_{x_{m,j+3}}^{h+m-1}$ and $x_{m,j+2} \in B_{x_{m,j+3}}^{h+m-1}$ that are connected by this k-ball chain.

We have now obtained 2^{m+1} points

$$x = x_{m,0}, x_{m,1}, \dots, x_{m,2^{m+1}-1} = y_{n,2^{m+1}-1}$$

such that

$$d(x_{m,i}, x_{m,i+1}) \le e^{-\epsilon(h+m-1)}$$

for each even integer $i \in \{0, \ldots, 2^{m+1} - 2\}$. Note that we have constructed 2^{m-1} different *k*-ball chains at this step, each of length at most $C'e^{k-(h+m-1)}$. Moreover, for each even $i \in \{0, \ldots, 2^{m+1} - 2\}$, there is a *k*-ball chain connecting $x_{m,i+1}$ to $x_{m,i+2}$, constructed either at this *m*-th step or at a previous one. Consider what happens at the end of the (k - h + 1)-th step. We obtain 2^{k-h+2} points

$$x = x_{k-h+1,0}, x_{k-h+1,1}, \dots, x_{k-h+1,2^{k-h+2}-1} = y$$

satisfying $d(x_{k-h+1,i}, x_{k-h+1,i+1}) \leq e^{-\epsilon k}$ for each even *i*. Consequently, there is a *k*-ball containing both $x_{k-h+1,i}$ and $x_{k-h+1,i+1}$. Indeed, if $z \in P_k$ with $d(z, x_{k-h+1,i}) < e^{-\epsilon k}$, then

$$x_{k-h+1,i}, x_{k-h+1,i+1} \in B(z, 2e^{-\epsilon k}) = B_k(z)$$

Moreover, there is a k-ball chain connecting $x_{k-h+1,i+1}$ to $x_{k-h+1,i+2}$ that was constructed at some step of the whole process.

Concatenating these k-ball chains, using the single k-balls containing both $x_{k-h+1,i}$ and $x_{k-h+1,i+1}$ to join them together, we end up with a k-ball chain from $x = x_{k-h+1,0}$ to $y = x_{k-h+1,2^{k-h+2}-1}$ of length at most

$$2^{k-h+1} + \sum_{m=1}^{k-h+1} 2^{m-1} \cdot C' e^{k-(h+m-1)} \lesssim e^{k-h}.$$

By the way we chose h, we obtain

$$d_k(x,y) \lesssim e^{k-h} \cdot e^{-k} \lesssim (e^{-\epsilon h})^{1/\epsilon} \lesssim d(x,y)^{1/\epsilon},$$

where the implicit constants depend only on C'. Taking limits then gives the desired bound $d_{new}(x,y) \lesssim d(x,y)^{1/\epsilon}$.

It therefore remains to prove Lemma 2.17. Broadly, our goal is to use the conformal elevator on Z to blow up the balls $B(x, e^{-\epsilon m})$ and $B(y, e^{-\epsilon m})$ to a uniform scale, so that we can essentially reduce to the case that $m \approx 1$. Establishing the analogous bound on this uniform scale will require some topological arguments, combined with the discrete volume and length bounds on Z. We will first develop the topological tools necessary to carry this out.

2.4.3 A discrete length-volume inequality

An important topic in metric geometry is the relationship between the volume of a space and the lengths of curves that, in some way, generate it. See [28, Chapter 4] for a survey of methods and results in this spirit. Among these is the following theorem, originally proved by W. Derrick [23]. We state it in the form cited in [28] in order to motivate more clearly what will follow.

Theorem 2.18 (Derrick [23, Theorem 3.4]). Let g be a Riemannian metric on the cube $[0,1]^n$, let F_k and G_k , $1 \le k \le n$, denote the pairs of opposite codimension-1 faces of the cube, and let d_k be the distance between F_k and G_k with respect to the metric g. Then $\operatorname{Vol}(g) \ge d_1 d_2 \cdots d_n$, where $\operatorname{Vol}(g)$ denotes the volume of $[0,1]^n$ with respect to g.

We will need a discrete/topological version of this theorem. Incidentally, the proof of the discrete version mimics the proof of the Riemannian version.

To set this up, let U_1, \ldots, U_N be an open cover of the cube $[0, 1]^n$. Again let F_k and G_k denote the pairs of opposite codimension-1 faces:

$$F_k = [0,1]^n \cap \pi_k^{-1}(\{0\})$$
 and $G_k = [0,1]^n \cap \pi_k^{-1}(\{1\}),$

where $\pi_k \colon \mathbb{R}^n \to \mathbb{R}$ is the projection onto the k-th coordinate axis. We say that U_{i_1}, \ldots, U_{i_l} is a chain if $U_{i_j} \cap U_{i_{j+1}} \neq \emptyset$ for each j. Moreover, such a chain is said to connect two sets A and B if $U_{i_1} \cap A \neq \emptyset$ and $U_{i_l} \cap B \neq \emptyset$.

Proposition 2.19. Let U_1, \ldots, U_N be as above, and let d_k denote the smallest number of sets U_i in a chain that connects F_k and G_k . Then $N \ge d_1 d_2 \cdots d_n$.

Proof. Without loss of generality, we may assume that no U_i is redundant, i.e., that for each i, there is a point $x_i \in U_i$ that belongs to no other U_j . Otherwise, we could delete U_i from the cover, thereby decreasing the total number of sets without reducing the numbers d_k .

We first define a map $f_0: \{x_1, \ldots, x_N\} \to \mathbb{Z}^n$ where the k-th component is

$$\pi_k(f_0(x_i)) = \begin{array}{l} \text{minimal number of sets } U_j \text{ in a chain} \\ \text{that connects } F_k \text{ and } \{x_i\}. \end{array}$$

The existence of such a chain follows from the fact that there is a path in $[0,1]^n$ from F_k to x_i and the collection $\{U_j\}_{j=1}^N$ is an open cover of this path.

Now we extend f_0 to a map on $[0,1]^n$ by using a partition of unity subordinate to $\{U_i\}_{i=1}^N$. More precisely, let $\{\varphi_i\}_{i=1}^N$ be a partition of unity such that $\{x : \varphi_i(x) \neq 0\} \subset U_i$. In particular, each φ_i is continuous and takes values in [0,1]. Let

$$f(x) = \sum_{i=1}^{N} \varphi_i(x) f_0(x_i) = \sum_{i=1}^{N} \varphi_i(x) y_i,$$

for $x \in [0,1]^n$, where we use y_i to denote $f_0(x_i)$. Observe that f does indeed extend f_0 because $\varphi_j(x_i) = 0$ for $j \neq i$ and $\varphi_i(x_i) = 1$. It is also, of course, continuous.

We claim that $\pi_k(f(F_k)) = \{1\}$ and $\pi_k(f(G_k)) \subset [d_k, \infty)$ for each k. For $x \in F_k$, let U_{i_1}, \ldots, U_{i_m} be the sets containing x so that

$$f(x) = \sum_{j=1}^{m} \varphi_{i_j}(x) y_{i_j}.$$

As $x \in U_{i_j} \cap F_k$, we see that the single set U_{i_j} connects F_k and $\{x_{i_j}\}$. Thus, the k-th coordinate of y_{i_j} is 1. Consequently,

$$\pi_k(f(x)) = \sum_{j=1}^m \varphi_{i_j}(x) = 1.$$

Similarly, if $x \in G_k$, again let U_{i_1}, \ldots, U_{i_m} be the sets containing x. Then $U_{i_j} \cap G_k$ is nonempty, so any chain connecting F_k to $\{x_{i_j}\}$ (which necessarily must end with the set U_{i_j}) actually connects F_k and G_k . By the definition of d_k , this chain has size at least d_k . Thus, the k-th coordinate of y_{i_j} is at least d_k . As a result,

$$\pi_k(f(x)) = \sum_{j=1}^m \varphi_{i_j}(x) \pi_k(y_{i_j}) \ge d_k \sum_{j=1}^m \varphi_{i_j}(x) = d_k.$$

We now claim that the image of f must contain the n-dimensional rectangle

$$S = \prod_{k=1}^{n} [1, d_k]$$

If not, there exists a point $y \in S \setminus f([0,1]^n)$. As f is continuous, $f([0,1]^n)$ is closed, so we may assume that y is in the interior of S. Let g be a homeomorphism from $[0,1]^n$ to S

that sends corresponding faces to corresponding faces (an affine map will do). Then by the previous claim,

$$f_t = (1-t)f|_{\partial [0,1]^n} + tg|_{\partial [0,1]^n}$$

gives a homotopy with values in $\mathbb{R}^n \setminus \operatorname{int}(S)$. In particular, $f|_{\partial[0,1]^n}$ is homotopic to $g|_{\partial[0,1]^n}$ in $\mathbb{R}^n \setminus \{y\}$.

Fix a simplicial decomposition of $[0,1]^n$; this gives a corresponding decomposition of $\partial [0,1]^n$. The latter decomposition allows us to express $f|_{\partial [0,1]^n}$ and $g|_{\partial [0,1]^n}$ as singular (n-1)chains with integer coefficients. Abusing notation we continue to denote the chains by $f|_{\partial [0,1]^n}$ and $g|_{\partial [0,1]^n}$. The homotopy given above implies that the corresponding classes $[f|_{\partial [0,1]^n}]$ and $[g|_{\partial [0,1]^n}]$ are equal in the singular homology group $H_{n-1}(\mathbb{R}^n \setminus \{y\})$. Notice, though, that $f|_{\partial [0,1]^n}$ extends to the map $f: [0,1]^n \to \mathbb{R}^n \setminus \{y\}$, which we can view as a singular *n*-chain
via the decomposition of $[0,1]^n$. Thus, the chain $f|_{\partial [0,1]^n}$ is the image of the chain f under
the boundary map, so $[f|_{\partial [0,1]^n}]$ is zero in $H_{n-1}(\mathbb{R}^n \setminus \{y\})$. In particular, $[g|_{\partial [0,1]^n}]$ is also
zero. This, however, contradicts the fact that $[g|_{\partial [0,1]^n}]$ generates $H_{n-1}(\mathbb{R}^n \setminus \{y\})$, which is
isomorphic to \mathbb{Z} . Hence, it must be that $S \subset f([0,1]^n)$.

Lastly, we claim that if $f(x) \in \mathbb{Z}^n$, then $f(x) = y_i$ for some *i*. Let U_{i_1}, \ldots, U_{i_m} be the sets for which $\varphi_{i_j}(x) > 0$, so that $x \in U_{i_1} \cap \cdots \cap U_{i_m}$. Then

$$f(x) = \sum_{j=1}^{m} \varphi_{i_j}(x) y_{i_j},$$

and for each $j, l \in \{1, \ldots, m\}$,

$$||y_{i_j} - y_{i_l}||_{\infty} \le 1,$$

where $||y||_{\infty} = \max\{|\pi_k(y)| : 1 \le k \le n\}$ is the ℓ^{∞} -norm. This inequality follows immediately from the fact that U_{i_j} and U_{i_ℓ} have non-trivial intersection. Consequently, for each k, there is an integer a_k such that

$$\pi_k(y_{i_j}) \in \{a_k, a_k+1\}$$

for all j; namely, $a_k = \min\{\pi_k(y_{i_j}) : 1 \le j \le m\}$. Now fix k and let

$$I = \{j : \pi_k(y_{i_j}) = a_k\} \text{ and } J = \{j : \pi_k(y_{i_j}) = a_k + 1\}.$$

By the definition of a_k , we have $I \neq \emptyset$. Then

$$\pi_k(f(x)) = \sum_{j=1}^m \varphi_{i_j}(x)\pi_k(y_{i_j}) = \sum_{j\in I}\varphi_{i_j}(x)a_k + \sum_{j\in J}\varphi_{i_j}(x)(a_k+1)$$
$$= a_k + \sum_{j\in J}\varphi_{i_j}(x).$$

By assumption, $\pi_k(f(x))$ is an integer, so $\sum_{j \in J} \varphi_{i_j}(x)$ is also an integer, necessarily equal to 0 or 1. This can happen only if $J = \emptyset$ or $J = \{1, \ldots, m\}$, but the latter implies that $I = \emptyset$, contrary to assumption. Thus, $J = \emptyset$, so each y_{i_j} has k-th coordinate a_k . In particular, $f(x) = (a_1, \ldots, a_n)$ as well, giving $f(x) = y_{i_1}$.

From the previous claims we obtain the desired conclusion immediately. Each integer lattice point in the cube $S = [1, d_1] \times \cdots \times [1, d_n]$ has some x_i in its pre-image under f. There are $d_1 d_2 \cdots d_n$ integer lattice points in S, so $N \ge d_1 d_2 \cdots d_n$.

It is easy to see that Proposition 2.19 still holds if all of the sets U_i are assumed to be closed. Indeed, we can enlarge U_i by a small amount to obtain open sets U'_i without changing the incidence structure. Then apply the proposition to these open sets.

More importantly for our later use, we point out that the proposition above remains true if we replace $[0,1]^n$ by a topological cube. Indeed, the assumptions and conclusions are entirely topological.

Before moving on to the proof of Lemma 2.17, we must establish a basic fact about finding topological cubes in the sphere \mathbb{S}^n .

Lemma 2.20. Let B_0 and B_1 be metric balls of radius $\delta > 0$ in \mathbb{S}^n for which $\operatorname{dist}(B_0, B_1) \ge \delta$. Suppose that $E \subset \mathbb{S}^n$ has $\operatorname{diam} E < \delta$ and $\operatorname{dist}(B_i, E) \ge \delta$ for i = 0, 1. Then there is a set $S \subset \mathbb{S}^n$ with the following properties:

- (i) S is homeomorphic to $[0, 1]^n$;
- (ii) the faces {0} × [0, 1]ⁿ⁻¹ and {1} × [0, 1]ⁿ⁻¹ correspond, under this homeomorphism, to sets C₀ ⊂ B₀ and C₁ ⊂ B₁, respectively;

- (iii) if x and y in S correspond to points that lie in opposite codimension-1 faces of $[0,1]^n$, then $d_{\mathbb{S}^n}(x,y) \ge c\delta^3$;
- (iv) S is disjoint from E.
- Here, c > 0 is an absolute constant.

We should remark that the bound in (iii) is certainly far from optimal. It is, however, sufficient for our purposes and makes the following proof much simpler.

Proof. By rotation, we may assume that E contains the north pole N = (0, ..., 0, 1). Let D be the metric ball of radius δ centered at N so that $E \subset D$ and $B_i \cap D = \emptyset$ for i = 0, 1.

Let $p: \mathbb{S}^n \to \mathbb{R}^n$ be the stereographic projection

$$(x_1, \dots, x_{n+1}) \mapsto \left(\frac{x_1}{1 - x_{n+1}}, \dots, \frac{x_n}{1 - x_{n+1}}\right)$$

Then $p(\partial D)$ is an (n-1)-dimensional sphere of radius $R \in [1/\delta, 2/\delta]$. Moreover, we can say that $p|_{\mathbb{S}^n \setminus D}$ is a bi-Lipschitz map onto the Euclidean ball $B_{\mathbb{R}^n}(0, R)$ with

$$d_{\mathbb{S}^n}(x,y) \lesssim ||p(x) - p(y)|| \lesssim \frac{1}{\delta^2} d_{\mathbb{S}^n}(x,y), \qquad (2.6)$$

where the implicit constants are absolute. This follows from the standard expression

$$d_{\mathbb{S}^n}(x,y) = \frac{2||p(x) - p(y)||}{\sqrt{1 + ||p(x)||^2} \cdot \sqrt{1 + ||p(y)||^2}}$$

of the chordal metric on \mathbb{S}^n in terms of the Euclidean norm $|| \cdot ||$ on \mathbb{R}^n . As a result, $p(B_0)$ and $p(B_1)$ are Euclidean balls in $B_{\mathbb{R}^n}(0, R)$ with

$$\operatorname{dist}(p(B_0), p(B_1)) \gtrsim \delta$$
 and $\operatorname{diam} p(B_i) \gtrsim \delta$

for i = 0, 1. It is then easy to find an *n*-dimensional topological cube $\hat{S} \subset B_{\mathbb{R}^n}(0, R)$ with a pair of opposite codimension-1 faces in $p(B_0)$ and $p(B_1)$, respectively, and for which any two opposite faces are at distance $\geq \delta$ from each other.

Now let $S = p^{-1}(\hat{S})$. Properties (i), (ii), and (iv) immediately follow from our choice of \hat{S} , and property (iii) is a consequence of the bounds in (2.6).

2.4.4 Proof of Lemma 2.17

We will prove a slight variant of the lemma, which easily implies the form stated above. Namely, we show that if x and y are distinct with $e^{-\epsilon m} < d(x, y) \le e^{-\epsilon(m-1)}$ then for each $k \ge m$, there is a k-ball chain connecting the balls $B(x, e^{-\epsilon m})$ and $B(y, e^{-\epsilon m})$ of length at most $C'e^{k-m}$.

It is straightforward to obtain Lemma 2.17 from this. Indeed, let $x, y \in Z$ be distinct, and let $m \in \mathbb{N}$ such that $d(x, y) \leq e^{-\epsilon(m-1)}$. Fix $k \geq m$ and let $m' \geq m$ be the integer for which $e^{-\epsilon m'} < d(x, y) \leq e^{-\epsilon(m'-1)}$. If $k \geq m'$, then the desired conclusion in Lemma 2.17 follows immediately from the conclusion of the variant. If k < m', then $d(x, y) \leq e^{-\epsilon k}$ so that x and y are contained in a common k-ball. Thus, it suffices to prove the variant.

To this end, let $x, y \in Z$ with $e^{-\epsilon m} < d(x, y) \le e^{-\epsilon(m-1)}$, and fix $k \ge m$. For ease of notation, let

$$B_x = B(x, e^{-\epsilon m})$$
 and $B_y = B(y, e^{-\epsilon m})$,

which again should not be confused with the earlier notation for k-balls. To find a short k-ball chain connecting B_x and B_y , we will proceed in the following way. We first restrict our attention to larger balls

$$B(p,r) \subset B(p,\lambda r),$$

containing both B_x and B_y but still of radius roughly $e^{-\epsilon m}$. By estimates we have discussed earlier, such a ball intersects $\leq e^{n(k-m)}$ k-balls. Applying the conformal elevator at this location and scale enlarges B_x and B_y to a uniform size so that we can find a "wide" topological cube inside the image of $B(p, \lambda r)$. This cube will have a pair of opposite codimension-1 faces in the images of B_x and B_y , respectively, and each pair of opposite faces will be uniformly far apart. Pulling this cube back down to scale $\approx e^{-\epsilon m}$, it will be covered by those k-balls that intersect $B(p, \lambda r)$. This puts us in the setting of Proposition 2.19, where the lower discrete length bound will imply that $d_i \gtrsim e^{k-m}$ for each *i*. As the size of this cover is $\leq e^{n(k-m)}$, Proposition 2.19 guarantees that $d_i \leq e^{k-m}$ as well. In particular, there is a chain connecting B_x and B_y of length roughly e^{k-m} . We must, of course, make these arguments rigorous; to do so, it will be convenient to set p = x and $r = 2Ce^{-\epsilon(m-1)}$, where C is the large constant we chose at the beginning of Section 2.4.1. Observe then that

$$B_x \cup B_y \subset B(p,r),$$

and moreover, that $x, y \in B(p, r/C)$. Let us also choose λ (for later use in applying the conformal elevator) in the following way. Fix a homeomorphism $F: Z \to \mathbb{S}^n$ and let $0 < \delta < 1$ be small enough that

$$d(z,w) \ge \frac{1}{2C^2e^{\epsilon}}$$
 implies $d_{\mathbb{S}^n}(F(z),F(w)) \ge 3\delta$ (2.7)

for $z, w \in \mathbb{Z}$. Then take $\lambda \geq 2$ large enough so that

$$d(z,w) < \omega\left(\frac{1}{\lambda}\right)$$
 implies $d_{\mathbb{S}^n}(F(z),F(w)) < c\delta^3$, (2.8)

where 0 < c < 1 is the constant from Lemma 2.20. Note that λ will depend on the modulus of continuity of F and F^{-1} .

The conformal elevator on Z gives a map g for this choice of p, r, and λ . If it happens that r > 1 (i.e., if m is small), then we simply choose g to be the identity map. All of the following estimates work equally well in this case.

Let x' = g(x), y' = g(y), and $K = g(Z \setminus B(p, \lambda r))$. Property (ii) of the conformal elevator guarantees that

$$d(x',y') \ge \frac{d(x,y)}{Cr} \ge \frac{e^{-\epsilon m}}{2C^2 e^{-\epsilon(m-1)}} = \frac{1}{2C^2 e^{\epsilon}},$$

so x' and y' are far apart. Property (iii) tells us that

$$B(x', 1/C) \cup B(y', 1/C) \subset g(B(p, r)),$$
 (2.9)

and as g is a homeomorphism, this implies that

$$\operatorname{dist}(\{x', y'\}, K) \ge \frac{1}{C}.$$

Moreover, we claim that

$$B\left(x', \frac{1}{2C^2e^{\epsilon}}\right) \subset g(B_x) \quad \text{and} \quad B\left(y', \frac{1}{2C^2e^{\epsilon}}\right) \subset g(B_y).$$
 (2.10)

Indeed, if $w \in B(x', 1/(2C^2e^{\epsilon}))$, then $z = g^{-1}(w)$ must be in B(p, r) by (2.9). Consequently, property (ii) again gives

$$d(x,z) \leq Cr \cdot d(gx,gz) = Cr \cdot d(x',w) < \frac{2C^2 e^{-\epsilon(m-1)}}{2C^2 e^{\epsilon}} = e^{-\epsilon m},$$

so that $z \in B_x$. Hence, w = g(z) is in $g(B_x)$. The same reasoning works also for B_y . Lastly, property (iv) of the conformal elevator guarantees that

diam
$$K \leq \omega\left(\frac{1}{\lambda}\right)$$
,

which we view as being very small.

Let us now use the homeomorphism $F: \mathbb{Z} \to \mathbb{S}^n$ to "regularize" this large-scale configuration. By our choice of δ from (2.7), we have

$$d_{\mathbb{S}^n}\left(F(x'), F(y')\right) \ge 3\delta \quad \text{and} \quad \operatorname{dist}\left(\{F(x'), F(y')\}, F(K)\right) \ge 3\delta$$

so that $B_0 = B_{\mathbb{S}^n}(F(x'), \delta)$ and $B_1 = B_{\mathbb{S}^n}(F(y'), \delta)$ are metric balls in \mathbb{S}^n with

$$\operatorname{dist}(B_0, B_1) \ge \delta,$$

and of distance at least δ from F(K). Moreover, observe that

$$B_0 \subset F\left(B(x', 1/(2C^2e^{\epsilon}))\right) \subset F\left(g(B_x)\right)$$
(2.11)

and

$$B_1 \subset F\left(B(y', 1/(2C^2e^{\epsilon}))\right) \subset F\left(g(B_y)\right), \qquad (2.12)$$

both of which follow from (2.7) and (2.10). Also note that by our choice of λ ,

$$\operatorname{diam} F(K) \le c\delta^3 < \delta.$$

The metric balls B_0 and B_1 and the set F(K) therefore satisfy the hypotheses in Lemma 2.20. Let $\hat{S} \subset \mathbb{S}^n \setminus F(K)$ be the *n*-dimensional topological cube given in the conclusion of this lemma. Then \hat{S} has a pair of opposite codimension-1 faces \hat{C}_0 and \hat{C}_1 in B_0 and B_1 , respectively; moreover, any two opposite faces have spherical distance $\geq c\delta^3$ from each other. Now send this set \hat{S} back to Z via the homeomorphism $(F \circ g)^{-1}$; that is, let

$$S = g^{-1} \circ F^{-1}(\hat{S})$$

so that S is a topological cube in the ball $B(p, \lambda r)$. Observe that it has a pair of opposite codimension-1 faces

$$C_0 = g^{-1} \circ F^{-1}(\hat{C}_0)$$
 and $C_1 = g^{-1} \circ F^{-1}(\hat{C}_1)$

that lie within B_x and B_y , respectively. This follows from the inclusions in (2.11) and (2.12).

Consider the set of k-balls that meet $B(p, \lambda r)$. Intersect each k-ball with S, and call the resulting collection \mathcal{U} . The estimate in (2.4) implies that

$$\#\mathcal{U} \lesssim e^{n(k-m)}.$$

Hence, \mathcal{U} is an open cover of the topological cube S by $\leq e^{n(k-m)}$ sets. In view of Proposition 2.19, we wish to show that each chain from \mathcal{U} that joins opposite codimension-1 faces of S must have $\geq e^{k-m}$ sets.

To this end, let U_1, \ldots, U_l be such a chain, so that $U_i \cap U_{i+1} \neq \emptyset$ for each *i*, and there are $a \in U_1$ and $b \in U_l$ in opposite faces of *S*. As $F \circ g(a)$ and $F \circ g(b)$ lie in opposite faces of \hat{S} , we know that

$$d_{\mathbb{S}^n}(F \circ g(a), F \circ g(b)) \ge c\delta^3.$$

By our choice of λ in (2.8), this implies that

$$d(ga, gb) \ge \omega\left(\frac{1}{\lambda}\right).$$

Property (i) of the conformal elevator then guarantees that

$$d(a,b) \ge \frac{r \cdot d(ga,gb)}{C} = 2e^{-\epsilon(m-1)}d(ga,gb) \gtrsim e^{-\epsilon m}$$

because $\omega(1/\lambda)$ is a uniform constant. The points

$$a = x_0, x_1, \dots, x_{l-1}, x_l = b$$

where $x_i \in U_i \cap U_{i+1}$ for each $1 \leq i \leq l-1$, form a discrete $4e^{-\epsilon k}$ -path from a to b. Consequently,

$$l \gtrsim \left(\frac{d(a,b)}{4e^{-\epsilon k}}\right)^{1/\epsilon} \gtrsim e^{k-m}$$

as desired.

Using the notation from Proposition 2.19, let d_1 denote the smallest number of sets in \mathcal{U} that form a chain connecting C_0 to C_1 . Similarly, for $2 \leq i \leq n$, let d_i be the smallest number of sets in a chain connecting the other (n-1) pairs of opposite faces in S. We have shown that $d_i \gtrsim e^{k-m}$ for each i, so Proposition 2.19 gives

$$e^{n(k-m)} \gtrsim \#\mathcal{U} \gtrsim d_1 \cdot e^{(n-1)(k-m)}.$$

Thus, there is a k-ball chain of length $\leq e^{k-m}$ joining C_0 and C_1 ; in particular, such a chain joins B_x and B_y . This completes the proof of Lemma 2.17.

2.5 Proof of Theorem 2.5

Let (Z, d) be a compact metric space satisfying the assumptions in Theorem 2.5. The strongly quasi-Möbius action $\Gamma \curvearrowright Z$ equips Z with a conformal elevator by Lemma 2.13 (see the remarks following the definition of a conformal elevator). The Ahlfors n/ϵ -regularity of Z immediately implies that every δ -separated set in Z has size at most $C\delta^{-n/\epsilon}$ for some uniform constant C. Lastly, the discrete length property we impose on Z is precisely the lower bound on discrete paths between points which appears in condition (iv) of Proposition 2.16. Thus, Z satisfies all four de-snowflaking conditions, so there is a metric d_{new} on Z for which

$$d(x,y)^{1/\epsilon} \lesssim d_{new}(x,y) \lesssim d(x,y)^{1/\epsilon}.$$

It is an easy exercise to see that the Ahlfors n/ϵ -regularity of (Z, d) translates into Ahlfors *n*-regularity of (Z, d_{new}) . Of course, (Z, d_{new}) remains homeomorphic to \mathbb{S}^n . More importantly, the action $\Gamma \curvearrowright Z$ remains strongly quasi-Möbius and cocompact on triples with respect to d_{new} . The following theorem, which we discussed in Section 2.1, is therefore relevant.

Theorem 2.21 ([5], Theorem 1.1). Let $n \in \mathbb{N}$, and let Z be a compact, Ahlfors n-regular metric space of topological dimension n. Suppose that $\Gamma \curvearrowright Z$ is a uniformly η -quasi-Möbius action on Z that is cocompact on triples. Then Z is $\tilde{\eta}$ -quasi-Möbius equivalent to the sphere \mathbb{S}^n , where $\tilde{\eta}(t) = C\eta(Ct)$ for some constant C.

Proof. The conclusion we state in this theorem is slightly different from that stated in [5]. The authors conclude that the action $\Gamma \curvearrowright Z$ is quasisymmetrically conjugate to a Möbius action on \mathbb{S}^n , but the above statement is implicit on the way to this conclusion.

We must point out, though, that the authors do not explicitly state the quantitative relationship between $\tilde{\eta}$ and η . However, the control on $\tilde{\eta}$ that we give here comes from their proof: first establish, as they do, that Z and Sⁿ have bi-Lipschitz equivalent weak-tangents; a quantitative version of [5, Lemma 2.1] gives a quantitative version of [5, Lemma 5.3], which guarantees that the compactification of a weak tangent of Z is η_1 -quasi-Möbius equivalent to Z, where $\eta_1(t) = C_1 \eta(C_1 t)$; the compactification of a weak tangent of Sⁿ is again Sⁿ; and the bi-Lipschitz equivalence between weak tangents translates into a strongly quasi-Möbius equivalence between the compactifications of weak tangents. Putting these facts together gives the desired function $\tilde{\eta}$.

In our situation, Γ acts on (Z, d_{new}) by strongly quasi-Möbius maps, so the distortion function $\tilde{\eta}$ that we obtain from Theorem 2.21 is also linear. Hence, (Z, d_{new}) is strongly quasi-Möbius equivalent to \mathbb{S}^n . As any strongly quasi-Möbius homeomorphism between compact sets is necessarily bi-Lipschitz (cf. Remark 2.14), we find that (Z, d_{new}) and \mathbb{S}^n are bi-Lipschitz equivalent. Let $\tilde{f}: Z \to \mathbb{S}^n$ be a map giving this equivalence, so that

$$d_{new}(x,y) \lesssim d_{\mathbb{S}^n}(f(x), f(y)) \lesssim d_{new}(x,y)$$

for all $x, y \in Z$. This completes the proof in the case that n = 1.

Suppose now that $n \ge 2$. Of course, the map \tilde{f} that we have chosen need not conjugate the action $\Gamma \curvearrowright Z$ to a Möbius action on \mathbb{S}^n . To correct this, we use a classical theorem of Tukia.

Theorem 2.22 (Tukia [59, Theorem G]). Let Γ be a group that acts on \mathbb{S}^n , $n \geq 2$, by η quasi-Möbius homeomorphisms and is cocompact on triples. Then there is an $\tilde{\eta}$ -quasi-Möbius map $\psi \colon \mathbb{S}^n \to \mathbb{S}^n$ for which $\psi \Gamma \psi^{-1}$ is a Möbius action; here, $\tilde{\eta}(t) = C\eta(t)$ for some constant C.

Proof. Again, Tukia's stated result does not include the quantitative relationship between $\tilde{\eta}$ and η that we give here. His proof, however, constructs ψ as a limit of maps whose cross-ratio distortion we can keep track of. More specifically, he finds a sequence $g_i \in \Gamma$, corresponding scaling factors $\lambda_i > 0$, and a linear map $\alpha \in \operatorname{GL}_n(\mathbb{R})$ for which

$$f_i(x) = \hat{\alpha}(\lambda_i \cdot g_i(x))$$

converges to the desired map, ψ . Here, $\hat{\alpha}$ is the bi-Lipschitz homeomorphism of \mathbb{S}^n obtained from α by conjugation by stereographic projection. Consequently,

$$[f_i(x_1), f_i(x_2), f_i(x_3), f_i(x_4)]$$

$$\leq \|\hat{\alpha}\|^4 [\lambda_i g_i(x_1), \lambda_i g_i(x_2), \lambda_i g_i(x_3), \lambda_i g_i(x_4)] \leq \|\hat{\alpha}\|^4 \eta([x_1, x_2, x_3, x_4]),$$

as scaling by λ_i does not change the cross-ratio. We use $\|\hat{\alpha}\|$ to denote the bi-Lipschitz constant of $\hat{\alpha}$.

Thus, each f_i is $\tilde{\eta}$ -quasi-Möbius with $\tilde{\eta}(t) = \|\hat{\alpha}\|^4 \eta(t)$, and so the limit function ψ is also $\tilde{\eta}$ -quasi-Möbius.

Applying this theorem to the strongly quasi-Möbius action $\tilde{f}\Gamma\tilde{f}^{-1}$ on \mathbb{S}^n , we obtain a strongly quasi-Möbius map ψ , which is therefore also bi-Lipschitz, such that

$$(\psi \circ \tilde{f})\Gamma(\tilde{f}^{-1} \circ \psi^{-1})$$

is a group of Möbius transformations on \mathbb{S}^n . Setting $f = \psi \circ \tilde{f}$ yields the desired f.

Remark 2.23. It is not clear whether the stronger conclusion (bi-Lipschitz conjugacy to a Möbius group) should hold in the case n = 1. Tukia's theorem has analogs in this setting; see, for example [37] and [46], which give us quasisymmetric conjugacy to a Möbius group. The problem is in choosing the "correct" conjugacy. Note that there are pairs of cocompact Möbius groups acting on S^1 that are quasisymmetrically conjugate but whose conjugating homeomorphism has non-zero derivative nowhere. See [39] for more information about the delicacy of such questions.

2.6 Entropy Rigidity in Coarse Geometry

We now turn our attention to Theorem 2.8, which is a rigidity result in the setting of Gromov hyperbolic geometry. We refer primarily to [13] and [25] for background on hyperbolic metric spaces.

Let (X, d) be a metric space. We say that X is *proper* if all closed balls $\overline{B}(x, r)$ are compact and that X is *geodesic* if any two points can be connected by an isometric image of an interval in \mathbb{R} .

Given two metric spaces (X, d_X) and (Y, d_Y) , a map $f: X \to Y$ is called a *quasi-isometric* embedding if there are constants $\lambda \ge 1$ and $k \ge 0$ such that

$$\frac{1}{\lambda}d_X(x,x') - k \le d_Y(f(x), f(x')) \le \lambda d_X(x,x') + k$$

for all $x, x' \in X$. If, in addition, each point $y \in Y$ lies in the k-neighborhood of the image f(X), then we say that f is a quasi-isometry. A rough isometric embedding or a rough isometry is defined in the same way by requiring that $\lambda = 1$. For the most part, we will be concerned with rough isometries. When it is necessary to specify the additive constant k, we will use the term k-rough isometry.

For any three points $x, y, p \in X$, let

$$(x,y)_p = \frac{1}{2} (d_X(x,p) + d_X(y,p) - d_X(x,y)).$$

This is the *Gromov product* of x and y based at p.

Definition 2.24. A metric space X is δ -hyperbolic if there is a base-point $p \in X$ so that

$$(x, y)_p \ge \min\{(x, z)_p, (y, z)_p\} - \delta$$
(2.13)

for every $x, y, z \in X$. We say that X is a *(Gromov) hyperbolic metric space* if it is δ -hyperbolic for some $\delta \geq 0$.

We will refer to the inequality in (2.13) as the δ -inequality. Although this definition may seem slightly esoteric, it has a concrete geometric meaning as a "thinness" condition on triangles. More precisely, if X is a δ -hyperbolic geodesic metric space, then for every geodesic triangle in X, each side is contained in the δ' -neighborhood of the union of the other two sides, where δ' is a constant multiple of δ (cf. [13, Proposition 2.1.3]).

Iterating the δ -inequality, one can obtain a corresponding condition on finite chains of points in X. Namely, if $x_0, x_1, \ldots, x_n \in X$, then

$$(x_0, x_n)_p \ge \min_{1\le i\le n} (x_i, x_{i-1})_p - \frac{\delta}{\log 2} \log n - c$$

where c is a uniform constant depending only on δ [25, Chapter 2, Lemma 14(i)]. Notice that the smaller we can take δ , the more negatively-curved X is. This leads to the following definition, given by M. Bonk and T. Foertsch in [3].

Definition 2.25. For $\kappa \in [-\infty, 0)$, we say that X has an asymptotic upper curvature bound κ if there is $p \in X$ and a constant $c \ge 0$ so that

$$(x_0, x_n)_p \ge \min_{1 \le i \le n} (x_i, x_{i-1})_p - \frac{1}{\sqrt{-\kappa}} \log n - c$$

for all chains x_0, \ldots, x_n in X.

Here, we use the convention that $1/\sqrt{\infty} = 0$. If X has an asymptotic upper curvature bound $\kappa < 0$, then we say that X is an $AC_u(\kappa)$ -space. By our discussion in the previous paragraph, every hyperbolic metric space is an $AC_u(\kappa)$ -space for some $\kappa < 0$. And conversely, the definitions immediately imply that every $AC_u(\kappa)$ -space is Gromov hyperbolic. Allowing the additive constant c in the definition of asymptotic upper curvature is what makes this notion *asymptotic*. A collection of uniformly bounded configurations in X will not affect the asymptotic curvature bounds, as one could simply make c larger. It makes sense, then, that the best way to study these curvature bounds is to pass to the boundary at infinity, which we now recall.

2.6.1 The hyperbolic boundary

To begin, we say that a sequence $\{x_n\}$ in X converges at infinity if

$$(x_n, x_m)_p \to \infty$$
 as $n, m \to \infty$.

It is immediate to see that this property is independent of p. We consider two such sequences $\{x_n\}$ and $\{y_n\}$ to be equivalent if

$$\lim_{n \to \infty} (x_n, y_n)_p = \infty,$$

and in this case, we write $\{x_n\} \sim \{y_n\}$. This is an equivalence relation on the set of sequences converging at infinity, and we let $\partial_{\infty} X$ denote the set of equivalence classes. Observe that if a sequence converges at infinity, then any subsequence also converges at infinity and, moreover, is equivalent to the original sequence.

The Gromov product on X extends to $\partial_{\infty} X$ by

$$(\xi,\eta)_p = \inf \liminf_{n \to \infty} (x_n, y_n)_p$$

where the infimum is taken over all $\{x_n\}$ and $\{y_n\}$ in the equivalence classes ξ and η , respectively. Although taking this infimum is necessary in general, the following lemma shows that it is not too restrictive.

Lemma 2.26 ([13], Lemma 2.2.2). Let X be δ -hyperbolic with base-point p and let $\xi, \eta, \zeta \in \partial_{\infty} X$.

(i) If $\{x_n\}$ represents ξ and $\{y_n\}$ represents η , then

$$(\xi,\eta)_p \le \liminf_{n\to\infty} (x_n,y_n)_p \le \limsup_{n\to\infty} (x_n,y_n)_p \le (\xi,\eta)_p + 2\delta.$$

(ii) The δ -inequality $(\xi, \eta)_p \ge \min\{(\xi, \zeta)_p, (\eta, \zeta)_p\} - \delta$ is satisfied.

When X is a CAT(-1)-space, Bourdon [9] showed that

$$\rho(\xi,\eta) = e^{-(\xi,\eta)_{\mathrm{F}}}$$

is a metric on $\partial_{\infty} X$ and thus gives the boundary a canonical metric. In the more general Gromov hyperbolic setting, however, this function may fail the triangle inequality. In its place, we have

$$\rho(\xi,\eta) \le K \max\{\rho(\xi,\zeta), \rho(\zeta,\eta)\}$$
(2.14)

for any $\xi, \eta, \zeta \in \partial_{\infty} X$, which follows immediately from part (ii) in the preceding lemma. Note that $K = e^{\delta}$ if X is δ -hyperbolic. A general procedure then produces, for ϵ small enough (depending only on δ), a metric d_{ϵ} on $\partial_{\infty} X$ satisfying

$$\frac{1}{4}e^{-\epsilon(\xi,\eta)_p} \le d_{\epsilon}(\xi,\eta) \le e^{-\epsilon(\xi,\eta)_p}$$

See [13, Section 2.2], especially Lemma 2.2.5, for details. This motivates the following definition.

Definition 2.27. A metric d on $\partial_{\infty} X$ is called a visual metric of parameter ϵ if there is a base-point $p \in X$ so that

$$e^{-\epsilon(\xi,\eta)_p} \lesssim d(\xi,\eta) \lesssim e^{-\epsilon(\xi,\eta)_p}$$

for all $\xi, \eta \in \partial_{\infty} X$. We say that d is visual if it is visual with respect to some $\epsilon > 0$.

The dependence on p is not important here; if d is visual with respect to p, then it will be visual with respect to any other base-point, with the same parameter ϵ . Observe that if $\partial_{\infty} X$ admits a visual metric of parameter ϵ , then it admits metrics of all parameters smaller than ϵ . Thus, if we set

 $\epsilon_0 = \epsilon_0(X) = \sup\{\epsilon : \text{there is a visual metric on } \partial_\infty X \text{ of parameter } \epsilon\},\$

then each $\epsilon \in (0, \epsilon_0)$ has an associated visual metric. We call this interval the visual interval. One should keep in mind the heuristic that the more negatively-curved X is, the larger ϵ_0 will be. The relationship between curvature in X and the length of this visual interval is more explicit in terms of asymptotic upper curvature bounds. Actually, we first need an additional assumption on X to guarantee that its boundary accurately reflects its geometry at large scales.

Definition 2.28. We say that X is visual if there is a constant k and a base-point $p \in X$ such that for every $x \in X$, there is a k-rough isometric embedding $\gamma: [0, \infty) \to X$ with $\gamma(0) = p$ and x in the image of γ .

We will refer to the image of such γ as a k-rough geodesic ray, starting at p. For visual metric spaces, the AC_u(κ) condition can be transferred to the boundary.

Proposition 2.29 ([3], Lemma 4.1). Let X be a visual, hyperbolic metric space and assume that there are constants a and c with

$$(\xi_0, \xi_n)_p \ge \min_{1 \le i \le n} (\xi_i, \xi_{i-1})_p - a \log n - c$$

for all chains ξ_0, \ldots, ξ_n in $\partial_{\infty} X$. Then there is a constant c' for which

$$(x_0, x_n)_p \ge \min_{1 \le i \le n} (x_i, x_{i-1})_p - a \log n - c'$$

for all chains x_0, \ldots, x_n in X. Conversely, if the inequality with chains in X holds for some c', then there is a constant c for which the inequality with boundary chains holds.

This condition on boundary chains gives more precise control on the type of inequality for ρ in (2.14). Indeed, we now have

$$\rho(\xi_0,\xi_n) \le Cn^a \max_{1 \le i \le n} \rho(\xi_i,\xi_{i-1})$$

for any chain ξ_0, \ldots, ξ_n . Arguments similar to those in [13, Lemma 2.2.5] allow one to build visual metrics on $\partial_{\infty} X$, but this time with more control on the optimal value of ϵ_0 . In the end, the authors obtain the following.

Proposition 2.30 ([3], Theorem 1.5). Let X be a visual, hyperbolic metric space. If X is $AC_u(\kappa)$, then for each $0 < \epsilon < \sqrt{-\kappa}$, there is a visual metric on $\partial_{\infty}X$ with parameter ϵ . Conversely, if there is a visual metric on $\partial_{\infty} X$ with parameter $\epsilon > 0$, then X is an $AC_u(-\epsilon^2)$ -space.

Together with other results in [3], this fact suggests that the correct analog of CAT(-1)in the coarse setting is $AC_u(-1)$. In the case where X is CAT(-1), the canonical metric on $\partial_{\infty}X$ is associated to the parameter $\epsilon_0 = 1$; in particular, there are visual metrics of parameter 1. Unfortunately, this may not happen for more general $AC_u(-1)$ -spaces, even though we know that visual metrics exist for all parameters $0 < \epsilon < 1$.

2.6.2 Geometric actions on hyperbolic metric spaces

Let X be a proper, geodesic, hyperbolic metric space. These basic assumptions guarantee two important "accessibility" properties for points in $\partial_{\infty} X$. First, for any base-point $p \in X$ and each $z \in \partial_{\infty} X$, there is an isometric embedding $\gamma \colon [0, \infty) \to X$ for which

$$\gamma(0) = p$$
 and $\{\gamma(t_n)\}$ represents z

whenever $t_n \to \infty$. We refer to images of such embeddings as geodesic rays and denote them by [p, z).

Similarly, for any two distinct points $z, z' \in \partial_{\infty} X$ there is an isometry $\gamma \colon \mathbb{R} \to X$ for which

$$\{\gamma(-t_n)\}$$
 represents z and $\{\gamma(t_n)\}$ represents z'

whenever $t_n \to \infty$. Naturally, we will denote such geodesic lines by (z, z'). The hyperbolicity of X then guarantees that there is a uniform constant C for which

$$|(z, z')_p - \operatorname{dist}(p, (z, z'))| \le C \tag{2.15}$$

whenever $z, z' \in X \cup \partial_{\infty} X$ are distinct and $p \in X$.

A subset $Y \subset X$ is called *quasi-convex* if there is a constant C for which every geodesic segment in X with endpoints in Y lies in the C-neighborhood of Y. We then say that an action $\Gamma \curvearrowright X$ is *quasi-convex geometric* if the action is

- (i) isometric: each $g \in \Gamma$ acts as an isometry;
- (ii) properly discontinuous: the set $\{g \in \Gamma : g(K) \cap K \neq \emptyset\}$ is finite for every compact set $K \subset X$;
- (iii) quasi-convex cocompact: there is a non-empty, Γ -invariant, quasi-convex set $Y \subset X$ and a compact set $K \subset Y$ for which $Y = \bigcup_{q \in \Gamma} g(K)$.

Let us fix such a group action $\Gamma \curvearrowright X$ and a corresponding quasi-convex set Y. As Y is Γ -invariant, the action $\Gamma \curvearrowright Y$ is isometric, properly discontinuous, and cocompact. Recall that such actions are said to be *geometric*.

For $p \in X$ fixed, the limit set $\Lambda(\Gamma)$ is the collection of points $z \in \partial_{\infty} X$ that can be represented by a sequence $\{x_n\} \subset \Gamma p$. Of course, this is independent of our choice of p. It is not difficult to see that the orbit Γp and the set Y are within finite Hausdorff distance from each other, so $\Lambda(\Gamma)$ coincides with $\partial_{\infty} Y$, viewed as a subset of $\partial_{\infty} X$. In particular, $\Lambda(\Gamma)$ is compact.

In fact, it will be convenient simply to replace Y with Γp . We lose no generality in doing this, as quasi-convexity of Y implies quasi-convexity of Γp . Thus, we take $Y = \Gamma p$ from now on.

Recall from earlier that the entropy of this action $\Gamma \curvearrowright X$ is

$$e(\Gamma) = \limsup_{R \to \infty} \frac{\log(N(R))}{R}$$
(2.16)

where $N(R) = \#\{\Gamma p \cap B_X(p, R)\}$. Under our assumptions, $e(\Gamma) < \infty$ and we can replace the "lim sup" with "lim"; in fact,

$$\exp(e(\Gamma)R) \lesssim N(R) \lesssim \exp(e(\Gamma)R)$$

(see [20, Théorème 7.2]). This quantity $e(\Gamma)$ is the coarse analog of volume entropy for Riemannian manifolds, and it is closely related to the metric regularity on $\Lambda(\Gamma)$.

Theorem 2.31 (Coornaert [20, Section 7]). When equipped with a visual metric of parameter $\epsilon > 0$, the limit set $\Lambda(\Gamma)$ is Ahlfors regular of dimension $e(\Gamma)/\epsilon$.

We now wish to transfer the action $\Gamma \curvearrowright X$ to a quasi-Möbius action on $\Lambda(\Gamma)$. Until we mention otherwise, we equip $\Lambda(\Gamma)$ with a visual metric d of parameter ϵ . The following lemma indicates that the induced action on $\Lambda(\Gamma)$ is strongly quasi-Möbius.

Lemma 2.32. Let $g \in \Gamma$. Then g extends naturally to an η -quasi-Möbius homeomorphism of $(\Lambda(\Gamma), d)$, where $\eta(t) = Ct$, and C depends only on the hyperbolicity constant of X and the constants in the visual metric d.

Proof. Let $z \in \Lambda(\Gamma)$ and let $\{x_n\} \subset Y = \Gamma p$ be a sequence representing the boundary point z. Then $\{gx_n\} \subset Y$ also converges at infinity; indeed $(x_n, x_m)_p = (gx_n, gx_m)_{gp}$ and convergence to infinity does not depend on base-points. Let z' be the point in $\Lambda(\Gamma)$ represented by $\{gx_n\}$ and define g(z) = z'. Note that this does not depend on the choice of $\{x_n\}$, as g is an isometry. Moreover, this extension is a bijection; its inverse is simply the extension of the isometry $g^{-1} \in \Gamma$.

It now suffices to prove that this extension is quasi-Möbius with a linear distortion function η (continuity follows easily from the quasi-Möbius condition). To this end, let $z_1, z_2, z_3, z_4 \in \mathbb{Z}$ be distinct and let $\{a_n\}, \{b_n\}, \{c_n\}, \{d_n\}$ be sequences in Y representing these points. For simplicity, let x' = gx denote the image of a point $x \in Y$ or $x \in \Lambda(\Gamma)$ under the map g. Then $\{a'_n\}, \{b'_n\}, \{c'_n\}, \{d'_n\}$ represent z'_1, z'_2, z'_3, z'_4 , respectively, and so by Lemma 2.26(i), we can estimate

$$\begin{aligned} &\frac{d(z'_1, z'_3)d(z'_2, z'_4)}{d(z'_1, z'_4)d(z'_2, z'_3)} \\ &\lesssim \limsup_{n \to \infty} \exp\left(-\epsilon\left((a'_n, c'_n)_p + (b'_n, d'_n)_p - (a'_n, d'_n)_p - (b'_n, c'_n)_p\right)\right) \\ &= \limsup_{n \to \infty} \exp\left(-\epsilon\left((a_n, c_n)_{g^{-1}p} + (b_n, d_n)_{g^{-1}p} - (a_n, d_n)_{g^{-1}p} - (b_n, c_n)_{g^{-1}p}\right)\right) \\ &= \limsup_{n \to \infty} \exp\left(-\epsilon\left((a_n, c_n)_p + (b_n, d_n)_p - (a_n, d_n)_p - (b_n, c_n)_p\right)\right) \\ &\lesssim \frac{d(z_1, z_3)d(z_2, z_4)}{d(z_1, z_4)d(z_2, z_3)}, \end{aligned}$$

where the implicit constants depend only on the hyperbolicity of X and on the multiplicative constant in the visual metric. Here, we have used the observation that for $x_1, x_2, x_3, x_4 \in X$, the cross-difference

$$(x_1, x_3)_p + (x_2, x_4)_p - (x_1, x_4)_p - (x_2, x_3)_p$$

= $\frac{1}{2} (d_X(x_1, x_4) + d_X(x_2, x_3) - d_X(x_1, x_3) - d_X(x_2, x_4)),$

is independent of the chosen base-point $p \in X$.

We should remark that this lemma and its proof are well-known, though most references deal with the more general case when g is assumed to be only a quasi-isometry (see, for example, [13, Chapter 5]). In that setting, g still extends to a quasi-Möbius homeomorphism of the boundary, but the distortion function η might not be linear. One does, however, recover a linear distortion function when g is a rough isometry; the proof is the same as the one above.

Abusing terminology, we will continue to let g denote the extension of $g \in \Gamma$ to $\Lambda(\Gamma)$. It is clear that composition is preserved in the extension, so we indeed obtain a strongly quasi-Möbius group action $\Gamma \curvearrowright \Lambda(\Gamma)$. The next lemma, well-known in this subject, shows that the cocompactness and proper discontinuity of $\Gamma \curvearrowright Y$ extends to cocompactness and proper discontinuity on triples for $\Gamma \curvearrowright \Lambda(\Gamma)$. For completeness, we include a proof, but see, for example, [27, Sections 8.2.K–8.2.Q] for further discussion.

Lemma 2.33. If $\Lambda(\Gamma)$ has at least three points, then the induced action $\Gamma \curvearrowright \Lambda(\Gamma)$ is

- (i) cocompact on triples,
- (ii) properly discontinuous on triples: for each triple $z_1, z_2, z_3 \in \Lambda(\Gamma)$ of distinct points, for every $\tau > 0$ there are only finitely many $g \in \Gamma$ for which gz_1, gz_2, gz_3 are τ -separated.

Proof. Let us first establish an important fact: if $z, z' \in \Lambda(\Gamma)$ are distinct and (z, z') is a geodesic line between them, then (z, z') lies in the C_1 -neighborhood of Y, where C_1 is a uniform constant. Essentially this follows from the quasi-convexity of Y. Indeed, there are sequences $\{x_n\}$ and $\{x'_n\}$ in Y that represent z and z', respectively. Quasi-convexity then implies that the geodesic segments $[x_n, x'_n]$ lie in the C_1 -neighborhood of Y. However, the parameterized geodesic line (z, z') is the limit of the segments $[x_n, x'_n]$, parameterized appropriately, in the topology of uniform convergence on compact sets. Thus, it must also lie in the C_1 -neighborhood of Y.

Now, fix three distinct points z_1, z_2, z_3 in $\Lambda(\Gamma)$, and let $(z_1, z_2), (z_2, z_3)$, and (z_1, z_3) be geodesic lines in X. Each (z_i, z_j) is then in the C_1 -neighborhood of Y. Taken together, the three lines form a geodesic triangle Δ with endpoints at infinity. The hyperbolicity of X guarantees that

$$\{x \in X : \operatorname{dist}(x, (z_i, z_j)) \le C_2 \text{ for each } i \ne j\} \ne \emptyset$$

for a large enough uniform constant C_2 (this is an easy consequence of the thinness condition for geodesic triangles). We will refer to this set as the C_2 -rough center of Δ . By taking C_2 slightly larger, we can find a point $x \in Y$ that lies in this rough center. Then, as $Y = \Gamma p$ is an orbit of the action, there is an isometry $g \in \Gamma$ with gx = p. Consequently, $\operatorname{dist}(p, (gz_i, gz_j)) \leq C_2$ for each $i \neq j$.

Recall also that the Gromov product of two points based at p roughly measures the distance between p and a geodesic line joining those points, in the sense of (2.15). Thus, we obtain

$$(gz_i, gz_j)_p \le \operatorname{dist}(p, (gz_i, gz_j)) + C_3 \le C_4,$$

where C_4 depends only on uniform quantities associated to the action $\Gamma \curvearrowright X$. By the definition of a visual metric on $\Lambda(\Gamma)$, this implies that

$$d(gz_i, gz_j) \ge \tau_0$$

for $i \neq j$, where $\tau_0 > 0$ is a uniform constant. This establishes cocompactness on triples.

Let us now prove the statement in part (ii). Again let $z_1, z_2, z_3 \in \Lambda(\Gamma)$ be distinct, and let $(z_1, z_2), (z_2, z_3)$, and (z_1, z_3) be geodesic lines in X. Fix $\tau > 0$ and suppose, for $g \in \Gamma$, that the points gz_1, gz_2, gz_3 are τ -separated. This implies that

$$(gz_i, gz_j)_p \le C_1(\tau)$$

for $i \neq j$, where $C_1(\tau)$ depends only on τ and on uniform constants. In particular, as $g(z_i, z_j)$ is a geodesic line between gz_i and gz_j , there is $C_2(\tau)$ such that

$$dist(g^{-1}p, (z_i, z_j)) = dist(p, g(z_i, z_j)) \le C_2(\tau).$$

Thus, $g^{-1}p$ is in the $C_2(\tau)$ -rough center of the geodesic triangle with sides (z_1, z_2) , (z_2, z_3) , and (z_1, z_3) .

As rough centers of geodesic triangles are necessarily bounded, the set

$$\{g^{-1}p: gz_1, gz_2, gz_3 \text{ are } \tau\text{-separated}\}$$

has finite diameter. Let $K \subset X$ be the closure of

$$\{g^{-1}p: gz_1, gz_2, gz_3 \text{ are } \tau\text{-separated}\} \cup \{p\}$$

in X, so K is compact and $g(K) \cap K \neq \emptyset$ whenever gz_1, gz_2 , and gz_3 are τ -separated. By proper discontinuity of the action $\Gamma \curvearrowright X$, there can be only finitely many such $g \in \Gamma$. This establishes proper discontinuity on triples.

Before setting out to prove Theorem 2.8, it is necessary to explain what it means for a rough isometry $\Phi: S \to Y$ to be "roughly equivariant" with respect to a geometric action of Γ on S. Of course, we will be interested in the case when $S = \mathbb{H}^{n+1}$.

Definition 2.34. A map $\Phi: S \to Y$ is *roughly equivariant* with respect to the actions $\Gamma \curvearrowright S$ and $\Gamma \curvearrowright Y$ if there is a constant *C* for which

$$d_X\left(\Phi(gx), g\Phi(x)\right) \le C$$

for each $x \in S$ and $g \in \Gamma$.

We will shortly need the fact that $\partial_{\infty} \mathbb{H}^{n+1}$ can be identified with \mathbb{S}^n . Under this identification, the chordal metric on \mathbb{S}^n is a visual metric of parameter 1, cf. [13, Section 2.4.3].

2.6.3 Proof of Theorem 2.8

Let us return now to the set-up in Theorem 2.8. Fix $\Gamma \curvearrowright X$ as in the statement of the theorem, and recall that $\Lambda(\Gamma)$ is assumed to be a topological sphere. Let $Y = \Gamma p$, so that Y is quasi-convex and is an $AC_u(-1)$ -space. Using the geometric action $\Gamma \curvearrowright Y$, we can verify the following lemma.

Lemma 2.35. There is a uniform constant C such that each $y \in Y$ lies in a C-rough geodesic ray in Y, starting at p. In other words, Y is visual, in the sense of Definition 2.28.

Proof. Fix $x \in Y$. We first want to find a geodesic line (z, z'), with $z, z' \in \Lambda(\Gamma)$, that passes close to x. To do this, choose two distinct points $w, w' \in \Lambda(\Gamma)$. The quasi-convexity of Y ensures that the geodesic line (w, w') in X lies in the C_1 -neighborhood of Y, for some uniform constant C_1 . In particular, there is a point $x' \in Y$ for which $dist(x', (w, w')) \leq C_1$, and there is $g \in \Gamma$ with gx' = x. Thus, $dist(x, (gw, gw')) \leq C_1$. Let z = gw and z' = gw' so that $dist(x, (z, z')) \leq C_1$.

Consider now the geodesic triangle with sides [p, z), [p, z'), and (z, z'). The δ -inequality in Lemma 2.26(ii) is valid for points in $X \cup \partial_{\infty} X$, and this translates into a thinness condition for geodesic triangles, even those with some vertices in $\partial_{\infty} X$. Consequently,

$$dist(x, [p, z) \cup [p, z')) \le dist(x, (z, z')) + C_2 \le C_1 + C_2,$$

where C_2 is uniform. Thus, we may assume that $dist(x, [p, z)) \leq C_3$ for a uniform constant C_3 .

It now suffices to show that [p, z) is in the C_4 -neighborhood of Y, where again C_4 is a uniform constant. Indeed, this easily implies that we can find a C-rough geodesic ray in Y, starting at p, and passing through x. As $z \in \Lambda(\Gamma)$, there is a sequence $\{x_n\} \subset Y$ that represents z, and by quasi-convexity of Y, the geodesic segments $[p, x_n]$ lie in the C_4 -neighborhood of Y. The parameterized geodesic ray [p, z) is simply the limit of the parameterized segments $[p, x_n]$ (in the topology of uniform convergence on compact sets), so we immediately see that [p, z) is also in the C_4 -neighborhood of Y. As Y is visual and $AC_u(-1)$, we can apply Proposition 2.30 to obtain, for each $0 < \epsilon < 1$, a visual metric on $\partial_{\infty} Y$ of parameter ϵ . Recall, though, that $\partial_{\infty} Y$ coincides with $\Lambda(\Gamma)$. Thus, there are visual metrics on $\Lambda(\Gamma)$ for all parameters $0 < \epsilon < 1$. By Theorem 2.31, these metrics are Ahlfors regular of dimension $e(\Gamma)/\epsilon$. In particular, the Hausdorff dimension of $\Lambda(\Gamma)$ with this metric is $e(\Gamma)/\epsilon$.

On the other hand, all visual metrics induce the same topology on $\Lambda(\Gamma)$; in our case, this is the topology of the standard *n*-dimensional sphere. Recalling that the topological dimension of a compact metric space always bounds the Hausdorff dimension from below (cf. [38, Theorem 7.2]), we obtain

$$\frac{e(\Gamma)}{\epsilon} = \dim_H(\Lambda(\Gamma), d_{\epsilon}) \ge \dim_{\mathrm{top}}(\Lambda(\Gamma), d_{\epsilon}) = n$$

for all $0 < \epsilon < 1$. This gives $e(\Gamma) \ge n$, which is the first part of the theorem.

It remains to prove the rigidity statement in Theorem 2.8, and this task will occupy us for the remainder of the section. The "if" part of the statement follows easily from standard facts about hyperbolic metric spaces. Namely, if $\Phi \colon \mathbb{H}^{n+1} \to Y$ is a rough isometry, then the fact that $\partial_{\infty} \mathbb{H}^{n+1} = \mathbb{S}^n$ admits a visual metric of parameter 1 implies that $\partial_{\infty} Y = \Lambda(\Gamma)$ does as well. Equipped with these metrics, we can extend Φ to a bi-Lipschitz map of the boundaries:

$$\Phi\colon \mathbb{S}^n\to \Lambda(\Gamma).$$

In particular,

$$e(\Gamma) = \dim_H \Lambda(\Gamma) = n.$$

Let us now address the converse statement. Thus, we assume that $e(\Gamma) = n$ and wish to construct the desired action $\Gamma \curvearrowright \mathbb{H}^{n+1}$ and map $\Phi \colon \mathbb{H}^{n+1} \to Y$.

Fix a visual parameter $0 < \epsilon < 1$ for $Z = \Lambda(\Gamma)$, which we also view as $\partial_{\infty}Y$, and let *d* denote a corresponding visual metric. Then (Z, d) is Ahlfors n/ϵ -regular and Lemma 2.32 implies that there is a strongly quasi-Möbius action $\Gamma \curvearrowright Z$. Moreover, Lemma 2.33(i) guarantees that this action is cocompact on triples. We claim that the discrete length condition appearing in Theorem 2.5 follows from the $AC_u(-1)$ assumption on Y. Indeed, if

$$u = z_0, z_1, \ldots, z_l = v$$

is a discrete δ_0 -path between u and v in Z, then Proposition 2.29 gives

$$(u, v)_p \ge \min_{1 \le i \le l} (z_i, z_{i-1})_p - \log l - c$$

for some uniform constant c. Translating this to the metric, we obtain

$$d(u,v) \lesssim l^{\epsilon} \cdot \max_{1 \leq i \leq l} d(z_i, z_{i-1}) \lesssim \delta_0 l^{\epsilon},$$

and rearranging gives $l \gtrsim (d(u, v)/\delta_0)^{1/\epsilon}$. The conditions in Theorem 2.5 are therefore satisfied, and so we obtain a metric d_{new} for which d and d_{new}^{ϵ} are bi-Lipschitz equivalent. In particular, d_{new} is a visual metric on Z of parameter 1. We also obtain a bi-Lipschitz map $f \colon \mathbb{S}^n \to (Z, d_{new})$ for which $f^{-1}\Gamma f$ is a Möbius action on the sphere. Observe that this action is cocompact on triples.

Furthermore, as $\Gamma \curvearrowright X$ is properly discontinuous, the induced boundary action $\Gamma \curvearrowright Z$ will be properly discontinuous on triples by Lemma 2.33(ii). This property is preserved under conjugation by homeomorphisms, so the Möbius action $f^{-1}\Gamma f$ will also be properly discontinuous on triples.

By the correspondence between Möbius transformations on \mathbb{S}^n and isometries of \mathbb{H}^{n+1} , for each $g \in \Gamma$, there is a unique isometry of \mathbb{H}^{n+1} that induces $f^{-1}gf$ on the boundary. This gives us a geometric action $\Gamma \curvearrowright \mathbb{H}^{n+1}$. Indeed, the cocompactness and proper discontinuity on triples for $f^{-1}\Gamma f$ translate into cocompactness and proper discontinuity for $\Gamma \curvearrowright \mathbb{H}^{n+1}$. This fact can be proved using arguments very similar to those in the proof of Lemma 2.33, so we do not repeat them. The main ingredients, though, are the following properties: every $x \in \mathbb{H}^{n+1}$ lies in the *C*-rough center of a geodesic triangle with vertices at infinity, where *C* is a uniform constant; and the three vertices of any such triangle are τ -separated if and only if its rough center is at distance $\leq C(\tau)$ from the fixed base-point. Here, we are strongly using the fact that the chordal metric on $\mathbb{S}^n = \partial_{\infty} \mathbb{H}^{n+1}$ is a visual metric. It remains to construct Φ . For this, we use standard arguments about extending bi-Lipschitz maps between boundaries of hyperbolic metric spaces to rough isometries of the hyperbolic spaces themselves. Actually, our argument will mimic the proof of Theorem 7.1.2 in [13]. Important to this construction is again the fact that the chordal metric on \mathbb{S}^n is a visual metric of parameter 1 under the identification of \mathbb{S}^n with $\partial_{\infty} \mathbb{H}^{n+1}$. Thus, $f: \mathbb{S}^n \to (Z, d_{new})$ is a bi-Lipschitz homeomorphism between two spaces whose metrics are of the form $e^{-(u,v)}$, up to multiplicative constants.

Fix $x \in \mathbb{H}^{n+1}$; for concreteness we will use the unit ball model of \mathbb{H}^{n+1} . Then there is a geodesic ray [0, z) in \mathbb{H}^{n+1} , ending at some $z \in \mathbb{S}^n$, with $x \in [0, z)$. Let $\gamma : [0, \infty) \to \mathbb{H}^{n+1}$ be the unit speed parameterization of this ray, and let $t = d_{\mathbb{H}^{n+1}}(0, x)$ so that $\gamma(t) = x$. Now, as f(z) is in $\Lambda(\Gamma)$, we also know that there is a geodesic ray [p, f(z)) in X that lies in the C_1 -neighborhood of Y (cf. the proof of Lemma 2.35). Let $\tilde{\gamma} : [0, \infty) \to X$ be the geodesic parameterization of this ray. We then define $\Phi(x)$ to be a point in Y that is of distance at most C_1 from $\tilde{\gamma}(t)$. Of course, this definition depends on the choice of a ray [p, f(z)) and on the choice of a point in Y. Making different choices, however, yields points that are within distance C_2 of each other, where C_2 is uniform.

The map $\Phi \colon \mathbb{H}^{n+1} \to Y$ thus defined induces, almost by definition, the homeomorphism f between \mathbb{S}^n and Z. Moreover, we claim that Φ is a rough isometry. To prove the desired bounds on $d_X(\Phi(x), \Phi(y))$, it suffices to show that

$$|(x,y)_0 - (\Phi(x),\Phi(y))_p| \le C_3 \tag{2.17}$$

for a constant C_3 independent of $x, y \in \mathbb{H}^{n+1}$. Indeed, the definition of Φ guarantees that

$$|d_{\mathbb{H}^{n+1}}(0,x) - d_X(p,\Phi(x))|, |d_{\mathbb{H}^{n+1}}(0,y) - d_X(p,\Phi(y))| \le C_4$$
(2.18)

for a uniform constant C_4 .

Fix $x, y \in \mathbb{H}^{n+1}$ and let $u, v \in \mathbb{S}^n$ be boundary points for which $x \in [0, u)$ and $y \in [0, v)$. The metric hyperbolicity of \mathbb{H}^{n+1} implies that

$$|(x,y)_0 - \min\{d_{\mathbb{H}^{n+1}}(0,x), d_{\mathbb{H}^{n+1}}(0,y), (u,v)_0\}|$$

is uniformly bounded (cf. [13, Lemma 7.1.3]). Similarly, the hyperbolicity of X ensures that

$$|(\Phi(x), \Phi(y))_p - \min\{d_X(p, \Phi(x)), d_X(p, \Phi(y)), (f(u), f(v))_p\}|$$

is uniformly bounded. Using (2.18) again, it is clear that (2.17) would follow from the uniform boundedness of

$$|(u, v)_0 - (f(u), f(v))_p|.$$

This, however, is an immediate consequence of the fact that f is bi-Lipschitz with respect to visual metrics of parameter 1; observe that the bound will depend on the bi-Lipschitz constant of f. Thus, (2.17) holds, and so

$$d_{\mathbb{H}^{n+1}}(x,y) - C_5 \le d_X(\Phi(x),\Phi(y)) \le d_{\mathbb{H}^{n+1}}(x,y) + C_5,$$

where C_5 is uniform. Note also that the definition of Φ , along with the facts that f is surjective and Y is visual, imply that each point in Y is of distance at most C_6 from $\Phi(\mathbb{H}^{n+1})$. Thus, Φ is a rough isometry.

Finally, we must show that Φ is roughly equivariant. To this end, let $g \in \Gamma$ and consider the rough isometry $g^{-1}\Phi g \colon \mathbb{H}^{n+1} \to Y$. Observe that it extends to the map

$$g^{-1}fg\colon \mathbb{S}^n \to \Lambda(\Gamma).$$

As f is equivariant with respect to the boundary actions, we know that $g^{-1}fg = f$. Thus, $g^{-1}\Phi g$ and Φ are rough isometries whose boundary extensions coincide. This implies that there is a uniform constant C for which

$$d_X(g^{-1}\Phi(gx),\Phi(x)) \le C$$

whenever $x \in \mathbb{H}^{n+1}$ (cf. [7, Proposition 9.1]). Hence

$$d_X(\Phi(gx), g\Phi(x)) \le C$$

for each $x \in \mathbb{H}^{n+1}$ and $g \in \Gamma$. This completes the proof of Theorem 2.8.

Remark 2.36. Many of the arguments we have used are valid in the case n = 1 as well. In particular, we can conclude that $e(\Gamma) \geq 1$ and that if Γp is roughly isometric to \mathbb{H}^2 , then $e(\Gamma) = 1$. The notable exception is the argument that allows us to conjugate the action $\Gamma \curvearrowright \Lambda(\Gamma)$ to a Möbius action on \mathbb{S}^1 by a bi-Lipschitz map. We discussed this issue in Remark 2.23, where we also indicated that the conjugation is possible with a quasisymmetric map. By standard extension arguments similar to those we used above to construct Φ , one can extend the quasisymmetric conjugation map to a quasi-isometry between \mathbb{H}^2 and Γp . This quasi-isometry will still be roughly equivariant with respect to the actions of Γ on \mathbb{H}^2 and on Γp by the same arguments we used above.

Thus, when n = 1, we can say only that Φ will be a quasi-isometry, rather than a rough isometry. We leave as an open question whether the stronger conclusion holds.

2.7 A closer look at the case n = 2

Our work in the previous sections broadly falls into the category of quasiconformal uniformization: given a metric space with certain properties, one attempts to find global parameterizations with some "analytic" structure. The theory that has developed around such problems is quite deep and is especially rich for two-dimensional spaces. The following theorem is characteristic of the subject. It gives conditions under which a topological sphere must actually be a *quasi-sphere*, i.e., must be a quasisymmetric image of the Euclidean sphere.

Theorem 2.37 (Bonk–Kleiner [4]). Suppose that Z is homeomorphic to \mathbb{S}^2 , is Ahlfors 2regular, and is linearly locally contractible. Then Z is quasisymmetrically equivalent to \mathbb{S}^2 .

Here, linear local contractibility is a connectivity property that arises naturally in the study of hyperbolic groups. Namely, for $\lambda \geq 1$, we say that Z is λ -linearly locally contractible if for each $x \in Z$ and $0 < r \leq \text{diam } Z/\lambda$, the metric ball B(x, r) can be contracted within $B(x, \lambda r)$ to a point. We say that Z is linearly locally contractible if it is λ -linearly locally contractible for some λ . As quasisymmetric maps distort relative distances by a controlled amount, they preserve linear local contractibility. Of course, the Euclidean sphere \mathbb{S}^2 is

linearly locally contractible, so every quasi-sphere is as well. Thus, the more restrictive assumption in Theorem 2.37 is that of Ahlfors 2-regularity.

There are other, related, connectivity conditions that also appear in the hyperbolic group setting. Let us list them and then discuss how they relate to each other. For $\lambda \geq 1$, we say that

- (i) Z is λ -LLC₁ if, for each $p \in Z$ and $0 < r \le \text{diam } Z$, any two points $x, y \in B(p, r)$ can be joined by a continuum in $B(p, \lambda r)$;
- (ii) Z is λ -LLC₂ if, for each $p \in Z$ and $0 < r \le \text{diam } Z$, any two points $x, y \in Z \setminus B(p, r)$ can be joined by a continuum in $Z \setminus B(p, r/\lambda)$;
- (iii) Z is λ -annularly linearly connected if it is connected, and for each $p \in Z$ and $0 < r \leq$ diam Z, any two points $x, y \in A(p, r, 2r)$ can be joined by a continuum in $A(p, r/\lambda, 2\lambda r)$.

Here, we use $A(p, r, R) = \overline{B}(z, R) \setminus B(z, r)$ to denote the (closed) metric annulus centered at p with inner radius r > 0 and outer radius R > r. Recall that a continuum is simply a compact connected set. The "LLC" acronym in LLC₁ and LLC₂ stands for "linearly locally connected" and should not be confused with linear local contractibility. For convenience, we will also use the acronym "ALC" in place of "annularly locally connected." This third condition is the strongest of the three; it was introduced by J. Mackay in [42].

Lemma 2.38. If (Z, d) is λ -ALC, then it is λ' -LLC₁ and λ' -LLC₂ for some λ' , depending only on λ .

Proof. We first verify the LLC₂ property. Let $p \in Z$, let $0 < r \leq \text{diam } Z$, and fix $x, y \in Z \setminus B(p, r)$. Without loss of generality, suppose that $R = d(p, x) \leq d(p, y)$, and let n be the largest integer for which $2^n R < d(p, y)$. Let $x_0 = x$, and for each $k \leq n$, choose $x_k \in Z$ with $d(p, x_k) = 2^k R$. This is possible because Z is connected. Finally, let $x_{n+1} = y$. The ALC condition guarantees that for each $0 \leq k \leq n$, there is a continuum

$$E_k \subset A\left(p, \frac{2^k R}{\lambda'}, 2^{k+1} R \lambda'\right)$$

connecting x_k and x_{k+1} , as long as $\lambda' > \lambda$. In particular, the continuum $E = E_0 \cup \cdots \cup E_n$ connects x and y and is contained in $Z \setminus B(p, r/\lambda')$. Thus, Z is λ' -LLC₂ for every $\lambda' > \lambda$.

To verify the LLC₁ condition, fix $x, y \in B(p, r)$. By the connectivity of Z, we may choose $q \in Z$ for which d(x,q) = d(x,y)/2. Then

$$x, y \in A\left(q, \frac{d(x,y)}{4}, \frac{3d(x,y)}{2}\right)$$

and using the same technique as in the previous paragraph, it is not difficult to show that xand y can be connected by a continuum

$$E \subset A\left(q, \frac{d(x,y)}{4\lambda'}, \frac{3\lambda'd(x,y)}{2}\right)$$

where λ' depends only on λ . In particular, $E \subset B(p, 5\lambda' r)$. Thus, Z is $5\lambda'$ -LLC₁.

When Z has some topological regularity, there are close relationships between the LLC_1 and LLC_2 conditions and the contractibility condition. For surfaces, there is also a correspondence between linear local contractibility and the ALC condition.

Lemma 2.39. If the metric space (Z, d) is a closed, connected manifold of dimension $n \ge 2$, then the following are true.

- (i) If Z is λ -linearly locally contractible, then it is λ' -LLC₁ and λ' -LLC₂ for all $\lambda' > \lambda$.
- (ii) If n = 2 and (Z, d) is both LLC₁ and LLC₂, then it is linearly locally contractible.
- (iii) If n = 2 and (Z, d) is λ-linearly locally contractible, then it is λ'-ALC for some λ', depending only on λ.

Proof. Parts (i) and (ii) are directly from [4, Lemma 2.5]. Part (iii) is effectively a consequence of the following non-trivial fact about linearly locally contractible surfaces [60, Proposition 4.28]: if (Z, d) is a topological surface that is λ -linearly locally contractible, then there is $\lambda' \geq 1$, depending only on λ , such that, for each $p \in Z$ and $0 < r \leq \text{diam } Z$, there is a neighborhood U of p, homeomorphic to a disk, with

$$B\left(p,\frac{r}{\lambda'}\right) \subset U \subset B(p,r).$$

Of course, we might as well take $\lambda' > \lambda$. We will use this fact to verify the ALC condition.

Fix $x, y \in A(p, r, 2r)$, where $0 < r \leq \text{diam } Z$. As a first case, suppose that $2\lambda' r > \text{diam } Z$. By (i), we know that Z is λ' -LLC₂, so as $x, y \in Z \setminus B(p, r)$, there is a continuum $E \subset Z \setminus B(p, r/\lambda')$ that connects x and y. The fact that $2\lambda' r > \text{diam } Z$ implies that $E \subset A(p, r/\lambda', 2\lambda' r)$, as needed. Thus, we may assume that $2\lambda' r \leq \text{diam } Z$. Consequently, there is a neighborhood U of p, homeomorphic to a disk, with

$$B(p,2r) \subset U \subset B(p,2\lambda'r);$$

in particular, $x, y \in U$. Applying this fact again, we can find a neighborhood V of p, also homeomorphic to a disk, with

$$B\left(p,\frac{r}{\lambda'}\right) \subset V \subset B(p,r)$$

Notice that $\overline{V} \subset \overline{B}(p,r) \subset U$. Moreover, as Z is a manifold, U is disjoint from ∂U . This implies that $\overline{V} \cap \partial U = \emptyset$, so in particular, ∂V is a continuum contained in U. Note that the connectivity of ∂V follows from the following fact: if $\phi \colon \mathbb{D} \to \Omega$ is a conformal map with $\Omega \subset \mathbb{C}$ bounded, then the Jordan curves $\phi(\partial B(0,r))$ converge, in the Hausdorff sense, to ∂V as $r \nearrow 1$.

By construction, we have $x, y \in U \setminus V$. Let γ be an arc in U from x to y. If $\gamma \subset U \setminus V$, then it is a continuum that joins x and y in $A(p, r/\lambda', 2\lambda'r)$. Otherwise, there are sub-arcs $\gamma_1, \gamma_2 \subset \gamma$ joining x to ∂V and joining y to ∂V , respectively, both of which lie in $U \setminus V$. Consequently, $E = \gamma_1 \cup \partial V \cup \gamma_2$ is a continuum in $U \setminus V \subset A(p, r/\lambda', 2\lambda'r)$ that joins x and y. Therefore, (Z, d) is λ' -ALC.

A major motivation for introducing these types of connectivity conditions is that they appear in the analysis of hyperbolic groups. Namely, if G is a hyperbolic group and $\partial_{\infty}G$ denotes its boundary at infinity equipped with a visual metric, then under suitable topological hypotheses, $\partial_{\infty}G$ will satisfy all of these conditions. For example, if $\partial_{\infty}G$ is non-empty, connected, and has no local cut points (equivalently, G does not split over a finite group or over a virtually cyclic group), then it is ALC [44, Proof of Corollary 1.2]. Similarly, if $\partial_{\infty}G$
is a connected manifold, then it is linearly locally contractible ([40, Theorem 4.4], along with [41, Theorem 3.3]).

Returning to Theorem 2.37, let us compare it to the n = 2 case of Theorem 2.3, which played such a central role in the previous sections. In both results, it is assumed that the metric space (Z, d) is Ahlfors 2-regular and has topological dimension 2. Theorem 2.37, then, essentially allows one to drop the additional hypothesis of self-similarity, replacing it by the much weaker linear local contractibility condition, and still obtain the same conclusion regarding quasisymmetric uniformization.

One might be tempted to ask whether this phenomenon occurs in all dimensions. Namely, if (Z, d) is homeomorphic to \mathbb{S}^n , is linearly locally contractible, and is Ahlfors *n*-regular, is it necessarily quasisymmetric to the Euclidean sphere? The answer is no for $n \geq 3$, by counterexamples due to S. Semmes [54], even though such spaces have good analytic properties [53]. See the Introduction of [4] for further discussion.

These considerations motivate us to reconsider our work from Section 2.4, at least in the case n = 2. The main result we needed to prove Theorems 2.5 and 2.8 was precisely the "de-snowflaking" statement that appeared in Proposition 2.16. There, it was required that (Z, d) admit a conformal elevator—quite a strong hypothesis for general metric spaces. Our goal here is to show that when n = 2, this assumption is not necessary, if one replaces it by connectivity conditions.

Proposition 2.40. Let $0 < \epsilon < 1$, and let (Z, d) be a complete metric space such that

- (i) Z is λ -ALC and λ -linearly locally contractible,
- (ii) every ball of radius $0 < R \le \text{diam } Z$ can be covered by at most $C(R/r)^{2/\epsilon}$ balls of radius 0 < r < R,
- (iii) every discrete δ -path from x to y in Z has length at least $C^{-1}(d(x,y)/\delta)^{1/\epsilon}$.

Then there is a metric ρ on Z satisfying $\rho \approx d^{1/\epsilon}$, where the implicit constant depends only on ϵ , λ , and C.

The condition in (ii) implies that (Z, d) has Assound dimension at most $2/\epsilon$. It is not difficult to show that if (Z, d) is compact and is Ahlfors $2/\epsilon$ -regular, it automatically satisfies this hypothesis. Thus, using Lemma 2.39 and Theorem 2.37, we obtain the following corollary.

Corollary 2.41. Let (Z, d) be a metric space, homeomorphic to \mathbb{S}^2 , that is linearly locally contractible. Suppose that Z is Ahlfors $2/\epsilon$ -regular and satisfies the discrete length condition in (iii) above. Then there is a metric ρ on Z for which $\rho \approx d^{1/\epsilon}$. In particular, (Z, ρ) is Ahlfors 2-regular and is quasisymmetrically equivalent to \mathbb{S}^2 .

The proof of Proposition 2.40 follows the proof of Proposition 2.16 very closely, but with the following input, used in place of the conformal elevator.

Theorem 2.42 (Mackay [43, Theorem 1.4]). Suppose that (Z, d) is an N-doubling, λ -ALC, complete metric space. For any $n \in \mathbb{N}$, there is $\alpha = \alpha(N, \lambda, n) \geq 1$ such that any two distinct $x, y \in Z$ can be joined by n different α -quasiarcs, such that the concatenation of any two of them is an α -quasicircle.

Recall that Z is said to be N-doubling if every ball of radius r can be covered by at most N balls of radius r/2. An arc γ in Z is called an α -quasiarc if for each pair $x, y \in \gamma$, the sub-arc between them has diameter at most $\alpha d(x, y)$. Similarly, a topological circle in Z is called an α -quasicircle if for each pair of points x, y on it, there is a sub-arc between them with diameter at most $\alpha d(x, y)$.

2.7.1 Proof of Proposition 2.40

Fix $0 < \epsilon < 1$ and (Z, d) as in the statement of the proposition. For each $k \in \mathbb{N}$, let P_k be a maximal $e^{-\epsilon k}$ -separated set in Z, and call $\{B(x, 2e^{-\epsilon k}) : x \in P_k\}$ the collection of k-balls, just as we did in the proof of Proposition 2.16. The one difference here is that we do not assume that Z is compact, so in particular, we have not normalized the metric to have diam Z = 1. However, assumption (ii) guarantees that Z is doubling, so P_k is countable for

each k. Just as in the proof of Proposition 2.16, we let

 $d_k(x, y) = ($ length of shortest k-ball chain connecting x and $y) \cdot e^{-k}$,

and then define

$$\rho(x,y) = \limsup_{k \to \infty} d_k(x,y)$$

We must show that $d_k(x, y) \approx d(x, y)^{1/\epsilon}$ for k large enough, with implicit constants depending only on ϵ , λ , and C.

As before, the lower bound follows immediately from the discrete length assumption (iii). In nearly the same way as for Proposition 2.16, the upper bound can be reduced to the following lemma, which is a variant of the statement in Lemma 2.17. To state this, we first must observe that assumption (ii) implies that Z is N-doubling, with N depending only on ϵ and C. Let $\alpha = \alpha(N, \lambda, 2) \ge 1$ be the "quasiarc" constant from Theorem 2.42. Once again, this depends only on ϵ , λ , and C. From now on, we refer to constants depending only on these three parameters as *uniform constants*.

Lemma 2.43. Let $x, y \in Z$ and $m \in \mathbb{Z}$ with $d(x, y) \leq e^{-\epsilon(m-1)} \leq \operatorname{diam} Z/(\alpha \lambda)$. Then for each $k \geq m$, there is a k-ball chain connecting $B(x, e^{-\epsilon m})$ and $B(y, e^{-\epsilon m})$ of length at most $C'e^{k-m}$, were C' is uniform.

Proof. We may assume that $B(x, e^{-\epsilon m})$ and $B(y, e^{-\epsilon m})$ are disjoint; otherwise, we connect them by a single k-ball. By Theorem 2.42, we can find two α -quasiarcs in Z connecting x and y such that the concatenation of the two is an α -quasicircle, C. Notice that C lies in the ball $B(x, \alpha d(x, y))$.

Let γ_1 be a subarc of \mathcal{C} with $\gamma_1 \subset B(x, e^{-\epsilon m}/4)$ and diam $\gamma_1 \geq e^{-\epsilon m}/8$. Similarly, let γ'_1 be a subarc of \mathcal{C} with $\gamma'_1 \subset B(y, e^{-\epsilon m}/4)$ and diam $\gamma'_1 \geq e^{-\epsilon m}/8$. Then, let γ_2 and γ'_2 be the sub-arcs of \mathcal{C} that connect γ_1 to γ'_1 . Observe that by construction, dist $(\gamma_1, \gamma'_1) \geq e^{-\epsilon m}/2$, and by the quasicircle condition,

$$\operatorname{dist}(\gamma_2, \gamma_2') \ge \frac{1}{\alpha} \min\{\operatorname{diam} \gamma_1, \operatorname{diam} \gamma_1'\} \ge \frac{e^{-\epsilon m}}{8\alpha}.$$

In particular, every k-ball chain connecting γ_i to γ'_i has length at least ce^{k-m} , where c is a uniform constant.

As (Z, d) is λ -linearly locally contractible and $\alpha d(x, y) \leq \operatorname{diam} Z/\lambda$, we know that the ball $B(x, \alpha d(x, y))$ is contractible inside $B(x, \alpha \lambda d(x, y))$. Thus, there is a continuous map

$$H: [0,1] \times B(x, \alpha d(x,y)) \to B(x, \alpha \lambda d(x,y))$$

for which $H_0 = \text{id}$ and $H_1 \equiv \text{const}$; here, we use the standard notation $H_t(z) = H(t, z)$ for $0 \leq t \leq 1$ and $z \in B(x, \alpha d(x, y))$. We may restrict this homotopy to the quasicircle, and pre-composing with a parameterization of \mathcal{C} , we obtain

$$\tilde{H}: [0,1] \times \mathbb{S}^1 \to B(x, \alpha \lambda d(x,y)),$$

where $\tilde{H}_0(\mathbb{S}^1) = \mathcal{C}$ and $\tilde{H}_1 \equiv \text{const.}$ This defines a continuous map $\tilde{g} \colon \overline{\mathbb{D}} \to B(x, \alpha \lambda d(x, y))$, where \mathbb{D} denotes the unit disk in \mathbb{R}^2 , and $\tilde{g}|_{\mathbb{S}^1} = \tilde{H}_0|_{\mathbb{S}^1}$ gives a parameterization of \mathcal{C} . More precisely,

$$\tilde{g}(re^{i\theta}) = \tilde{H}(1-r,e^{i\theta})$$

for each $0 \le r \le 1$ and $0 \le \theta \le 2\pi$. Notice that $\gamma_1, \gamma'_1, \gamma_2$, and γ'_2 correspond to sub-arcs of \mathbb{S}^1 under this parameterization. It is straightforward to see that we may pre-compose \tilde{g} with an appropriate homeomorphism from $[0, 1]^2$ to $\overline{\mathbb{D}}$ to obtain a continuous map

$$g\colon [0,1]^2 \to B(x,\alpha\lambda d(x,y)),$$

where $g(F_1) = \gamma_1$, $g(G_1) = \gamma'_1$, $g(F_2) = \gamma_1$, and $g(G_2) = \gamma'_2$. Recall that F_1, G_1, F_2, G_2 denote the sides of $[0, 1]^2$.

Let \mathcal{V} denote the collection of k-balls that intersect $B(x, \alpha \lambda d(x, y))$ non-trivially. The assumption (ii) easily implies that the number of such balls is at most

$$C\left(\frac{\alpha\lambda d(x,y)}{e^{-\epsilon k}/2}\right)^{2/\epsilon} \lesssim e^{2(k-m)}.$$

Now, let \mathcal{U} denote the open cover of $[0, 1]^2$ formed by pre-images of k-balls under the continuous map g. Using notation from earlier, for i = 1, 2, let d_i denote the minimal number of sets in \mathcal{U} that form a chain connecting F_i to G_i . As every k-ball chain connecting γ_i to γ'_i has length at least ce^{k-m} , we have $d_i \geq ce^{k-m}$. Applying Proposition 2.19 gives

$$ce^{k-m}d_1 \le d_1d_2 \le #\mathcal{U} = #\mathcal{V} \lesssim e^{2(k-m)},$$

so that $d_1 \leq e^{k-m}$. In particular, there is a k-ball chain connecting $\gamma_1 \subset B(x, e^{-\epsilon m})$ to $\gamma'_1 \subset B(x, e^{-\epsilon m})$ of length $\leq e^{k-m}$. This proves the lemma.

To obtain the bound $d_k(x, y) \leq d(x, y)^{1/\epsilon}$, we apply the same iteration argument as we did in Section 2.4. The only subtlety here is that the bound we have from Lemma 2.43 works only on length scales $e^{-\epsilon(m-1)} \leq \operatorname{diam} Z/(\alpha \lambda)$. Of course, if $\operatorname{diam} Z = \infty$, there is no need to make any change.

Otherwise, let m_0 be the largest integer for which $\alpha \lambda e^{-\epsilon(m_0-1)} > \text{diam } Z$. Fix $x, y \in Z$, and let m be the integer for which

$$e^{-\epsilon(m+1)} < d(x,y) \le e^{-\epsilon m}.$$

If $m \ge m_0$, then Lemma 2.43 can be applied at all relevant length scales, and it follows from the iteration argument that

$$d_k(x,y) \lesssim e^{-m} \lesssim d(x,y)^{1/\epsilon}$$

If, instead, $m < m_0$, then we construct a chain in the following way. Notice that $e^{-\epsilon m_0}$ is roughly equal to diam Z, so we may cover Z by a uniform number of balls of radius at most $e^{-\epsilon m_0}/2$. There is a chain of such balls, necessarily of uniformly bounded length, connecting x and y. Observe that if B_1, B_2, B_3 are consecutive balls in this chain, we can find points $x_1 \in B_1 \cap B_2$ and $x_2 \in B_2 \cap B_3$ with

$$d(x_1, x_2) \le 2 \operatorname{diam} B_2 \le e^{-\epsilon m_0}$$

Consequently, we can connect x_1 and x_2 by a k-ball chain of length $\leq e^{k-m_0} \leq e^{k-m}$. Concatenating a uniformly bounded number of such k-ball chains together, we obtain a chain from x to y with length $\leq e^{k-m}$. This gives $d_k(x, y) \leq e^{-m} \leq d(x, y)^{1/\epsilon}$, which finishes the proof of Proposition 2.40.

2.7.2 Further remarks

Notice that the essential ingredient in the proof of Proposition 2.40 was the construction of a continuous map $g: [0,1]^2 \to B(x, \alpha \lambda d(x, y))$ for which the images of opposite sides of the unit square are at distance $\gtrsim d(x, y)$ from each other. In this light, we make the following definition.

Definition 2.44. A metric space (Z, d) is said to admit fat connecting n-cubes if there is $\lambda \geq 1$ such that for any two distinct points $x, y \in Z$, one can find a continuous map $g: [0,1]^n \to B(x, \lambda d(x, y))$ with $g(F_1) \subset B(x, d(x, y)/4)$ and $g(G_1) \subset B(y, d(x, y)/4)$, and which has $\operatorname{dist}(g(F_k), g(G_k)) \geq d(x, y)/\lambda$ for each $1 \leq k \leq n$.

We have shown in this section that if (Z, d) is complete, ALC, and linearly locally contractible, then it admits fat connecting squares. Moreover, the arguments we made in Section 2.4 show that if (Z, d) is a topological *n*-sphere and admits a conformal elevator, then it admits fat connecting *n*-cubes. This definition may seem a bit contrived, but we will see a similar condition arise in the next chapter. For now, we only mention that our arguments in the proof of Proposition 2.40 can be easily adapted to give the following generalization.

Proposition 2.45. Let $0 < \epsilon < 1$, and let (Z, d) be a metric space with the following properties:

- (i) Z admits fat connecting n-cubes,
- (ii) every ball of radius 0 < R ≤ diam Z can be covered by at most C(R/r)^{n/ϵ} balls of radius 0 < r < R,
- (iii) every discrete δ -path from x to y in Z has length at least $C^{-1}(d(x,y)/\delta)^{1/\epsilon}$.

Then there is a metric ρ on Z satisfying $\rho \approx d^{1/\epsilon}$, where the implicit constant depends only on ϵ , λ , and C.

CHAPTER 3

Length-volume inequalities revisited

3.1 Introduction

There is a deep and well-studied relationship in metric geometry between the volume of a space and the lengths of curves that, in some way, generate it. An early example of such a relationship is due to K. Loewner (unpublished, but see [50] for a discussion) and deals with conformal structures on the torus \mathbb{T}^2 .

Theorem 3.1. Let (\mathbb{T}^2, g) be the 2-dimensional torus, equipped with a Riemannian metric g, let $\ell(g)$ denote the infimal length of a closed curve on \mathbb{T}^2 that is not homotopically trivial, and let $\operatorname{Vol}(g)$ denote the volume of \mathbb{T}^2 with respect to the metric g. Then $\operatorname{Vol}(g) \geq \frac{\sqrt{3}}{2}\ell(g)^2$, and equality holds if and only if (\mathbb{T}^2, g) is isometric to the flat torus \mathbb{R}^2/Λ , where Λ is the lattice generated by (1,0) and $(1/2, \sqrt{3}/2)$.

Loewner's inequality is only the beginning of a very rich body of work that has sought to understand similar phenomena for more general spaces and in dimensions greater than two. We refer to [28, Chapter 4] for a broad survey of methods and results in this area. Of particular interest to us is the following theorem, originally proved by W. Derrick [23, Theorem 3.4]. Here we state it in the form cited in [28].

Theorem 3.2. Let $([0,1]^n, g)$ be the n-dimensional unit cube, equipped with a Riemannian metric g. Let F_k and F'_k denote the pairs of opposite codimension-1 faces of $[0,1]^n$, for $1 \le k \le n$, and let d_k be the distance between F_k and F'_k with respect to the metric g. Then $\operatorname{Vol}(g) \ge d_1 \cdots d_n$. Inequalities such as those in Theorems 3.1 and 3.2 are interesting in themselves, and in this chapter we will be concerned primarily with inequalities of a similar flavor. We should mention, though, that understanding the case of equality is also very desirable. This theme appears strongly, for example, in the study of the marked length spectra of negativelycurved manifolds. Namely, if (M^n, g) is compact and negatively-curved, then each free homotopy class contains a unique geodesic of minimal length. The map associating this length to each class is called the marked length spectrum of g. It is conjectured that this spectrum determines the metric g, up to isometry. C. Croke [21] and J.-P. Otal [48] showed, independently, that this is true for surfaces. In higher dimensions, much less is known, though there are some positive results. For example, if M admits a locally symmetric metric g_0 , then any metric whose marked length spectrum coincides with that of g_0 is necessarily isometric to g_0 . This follows from two theorems of U. Hamenstädt [30, 32], along with the work of G. Besson, G. Courtois, and S. Gallot on "entropy vs. volume" rigidity [1]. See [55] for more discussion on these topics.

In the previous chapter, we proved a combinatorial version of Theorem 3.2 for open covers of the unit cube $[0,1]^n$. Incidentally, the purpose of the inequality was, ultimately, to prove a coarse geometric analog of U. Hamenstädt's "entropy vs. curvature" rigidity theorem [31]. More specifically, it was used in the construction of a metric with certain regularity properties on the boundary of a Gromov hyperbolic metric space. The set-up was as follows. Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of $[0,1]^n$, and again let F_k, F'_k denote the pairs of opposite codimension-1 faces. We say that U_{i_1}, \ldots, U_{i_m} is a *chain* if $U_{i_j} \cap U_{i_{j+1}} \neq \emptyset$ for each j. Moreover, such a chain is said to connect two sets A and B if $U_{i_1} \cap A \neq \emptyset$ and $U_{i_m} \cap B \neq \emptyset$.

Theorem 3.3 (Chapter 2, Proposition 2.19). Let d_k denote the smallest number of sets U_i in a chain that connects F_k and F'_k . Then $\#\mathcal{U} \ge d_1 \cdots d_n$.

Note that although this result is analogous to Derrick's theorem, it does not parallel the Riemannian inequality. Indeed, the sets U_i in \mathcal{U} are essentially treated as if they all had diameter 1. The primary purpose of this chapter is to extend the preceding result to a weighted version, which is closer to Theorem 3.2 and also generalizes Theorem 3.3.

To this end, let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of $[0, 1]^n$ as before, and let $w \colon I \to [0, \infty)$ be a corresponding weight function. One should think of w(i) as the weight associated to the set U_i . Together, \mathcal{U} and w give us a discrete notion of distance on $[0, 1]^n$. Namely, we define

$$\operatorname{dist}_{w}(A,B) = \inf \left\{ \sum_{j=1}^{m} w(i_{j}) : \begin{array}{c} U_{i_{1}}, \dots, U_{i_{m}} \text{ is a chain} \\ \text{that connects } A \text{ and } B \end{array} \right\}$$

The sum $w(i_1) + \ldots + w(i_m)$ is said to be the length of the corresponding chain. By path connectedness of $[0, 1]^n$ and compactness of paths, it is easy to see that any two points in the cube can be connected by a chain. In particular, $dist_w(A, B)$ is finite for $A, B \subset [0, 1]^n$. We should note that a chain might be disconnected topologically, as we have made no assumption on the connectedness of the sets in \mathcal{U} . Our main theorem will be the following.

Theorem 3.4. Let \mathcal{U} be an open cover of $[0,1]^n$, let w be a corresponding weight function, and let $d_k = \operatorname{dist}_w(F_k, F'_k)$ for each $1 \leq k \leq n$. Then

$$\sum_{i\in I} w(i)^n \ge d_1 \cdots d_n.$$

In fact, we will prove a more general version of this inequality that has more in common with the results of Derrick in [22] and [24]: we allow the discrete distance between F_k and F'_k to be taken with respect to (possibly) different weight functions for different values of k.

Theorem 3.5. Let \mathcal{U} be an open cover of $[0,1]^n$, and let w_k be associated weight functions for $1 \leq k \leq n$. If $d_k = \operatorname{dist}_{w_k}(F_k, F'_k)$ for each k, then

$$\sum_{i \in I} \left(\prod_{k=1}^n w_k(i) \right) \ge d_1 \cdots d_n.$$

It is clear that Theorem 3.4 follows immediately from Theorem 3.5 by setting $w_k = w$ for each k.

As a corollary to Theorem 3.4, we will easily obtain lower Hausdorff content bounds for continuous images of $[0, 1]^n$ in arbitrary metric spaces. Recall that if (X, d) is a metric space, the Q-dimensional Hausdorff content of a compact subset $E \subset X$ is defined to be

$$\mathcal{H}_Q^{\infty}(E) = \inf \left\{ \sum_{i \in I} (\operatorname{diam} U_i)^Q : \{U_i\}_{i \in I} \text{ is an open cover of } E \right\}.$$

We will show the following bound.

Corollary 3.6. Let $g: [0,1]^n \to X$ be continuous. Then

$$\mathcal{H}_n^{\infty}(g([0,1]^n)) \ge \prod_{k=1}^n \operatorname{dist}(g(F_k), g(F'_k)),$$

where dist denotes the metric distance between sets in X.

For example, if $([0,1]^n, d)$ is the unit cube equipped with an arbitrary metric whose topology coincides with the Euclidean topology, then one can apply this inequality to the identity function. We remark here, though, that our definition of Hausdorff content does not include the normalizing multiplicative factor $\operatorname{Vol}(B^n)2^{-n}$, where B^n is the unit ball in \mathbb{R}^n and $\operatorname{Vol}(\cdot)$ is *n*-dimensional Lebesgue measure, as is standard in geometric measure theory. In particular, Corollary 3.6 does not immediately recover Derrick's estimates.

As will become clear, the methods we use in this chapter are well-suited to study relationships between lengths and volumes for other polyhedral objects. In fact, we will prove versions of Theorems 3.3 and 3.4 for simplices, and so also obtain an analog of Corollary 3.6. There are, in addition, interesting extremal questions related to the inequalities stated above. We have not sufficiently explored this topic, so we leave it as an open direction.

The outline of the chapter is as follows. In Section 3.2, we will introduce some techniques that appear frequently: partitions of unity, the nerve of open covers, and a topological nondegeneracy lemma that will play a central role in our proofs of the length-volume inequalities. Section 3.3 will be devoted entirely to the proof of Theorem 3.5. In Section 3.4, we take up the topic of lower Hausdorff content bounds in metric spaces, and the final section deals with the extension of our methods to diameter-volume inequalities for simplices.

3.2 Notation and preliminaries

Let (X, d) be a metric space. As is standard, we will use B(x, r) to denote the open ball centered at $x \in X$ and of radius r > 0. For subsets $A, B \subset X$, we let

$$dist(A, B) = \inf\{d(x, y) : x \in A \text{ and } y \in B\}$$

be the distance between A and B. In the case that $A = \{x\}$ is a singleton, we will allow a slight abuse of notation and simply write dist(x, B). We also let

$$\operatorname{diam} A = \sup\{d(x, y) : x, y \in A\}$$

denote the diameter of the subset A. Following common notation, we use int(A) and A to denote, respectively, the interior and closure of A (the ambient space for the closure operation will be understood from context). Also, we let $\partial A = \overline{A} \setminus int(A)$ denote the boundary of A. Lastly, for $\epsilon > 0$, let

$$N_{\epsilon}(A) = \{x \in X : \operatorname{dist}(x, A) < \epsilon\}$$

be the open ϵ -neighborhood of A.

Suppose that (X, d) is compact, and let $\mathcal{U} = \{U_i\}_{i \in I}$ be a finite open cover. By the Lebesgue lemma, there is a positive constant $\delta > 0$ such that each ball $B(x, 2\delta)$ lies entirely in some set U_i . Let

$$f_i(x) = \min\left\{1, \frac{1}{\delta}\operatorname{dist}(x, N_{\delta}(X \setminus U_i)\right\},\$$

which is a $1/\delta$ -Lipschitz function with values in [0, 1] and whose support is contained in U_i . Moreover, for each $x \in X$, we have

$$f(x) := \sum_{i \in I} f_i(x) \ge 1,$$

as $B(x, 2\delta) \subset U_i$ for some *i*. If $N = \max\{\sum_i \chi_{U_i}(x) : x \in X\}$, then also $f(x) \leq N$ for all *x*. Define $\phi_i(x) = f_i(x)/f(x)$ so that the following properties hold:

(i) ϕ_i is $(2N+1)/\delta$ -Lipschitz with support contained in U_i ;

- (ii) $0 \le \phi_i(x) \le 1$ for all $x \in X$;
- (iii) $\sum_{i} \phi_i(x) = 1$ for all $x \in X$.

The family $\{\phi_i\}$ is therefore a $2N/\delta$ -Lipschitz partition of unity subordinate to \mathcal{U} [36, Chapter 2]. Partitions of unity are useful in general metric settings to produce a type of proxy for linear structure.

There is a canonical way to associate a simplicial complex to the cover \mathcal{U} whose combinatorics mimics the combinatorics of \mathcal{U} . Often, this simplicial complex is defined as an abstract complex that encodes the intersections among the sets in \mathcal{U} . We prefer to work with a geometric realization of this complex in Euclidean space. For ease, then, let us index the collection $\{U_i\}$ by the integers $1, \ldots, M$, and let e_i be the *i*-th standard basis vector in \mathbb{R}^M .

Definition 3.7. The nerve of \mathcal{U} , denoted by Ner(\mathcal{U}), is

$$\operatorname{Ner}(\mathcal{U}) = \bigcup \{ \operatorname{conv}(e_{i_0}, \dots, e_{i_m}) : U_{i_0} \cap \dots \cap U_{i_m} \neq \emptyset \}$$

where the union runs over collections of sets in \mathcal{U} that have non-empty intersection.

Here, and in general, we use conv(A) to denote the convex hull of a set $A \subset \mathbb{R}^M$. When $A = \{a_0, \ldots, a_m\}$ is a finite set, we can express

$$\operatorname{conv}(A) = \left\{ \sum_{i=0}^{m} \lambda_i a_i : \lambda_i \ge 0 \text{ and } \lambda_0 + \ldots + \lambda_m = 1 \right\},$$
(3.1)

and if $m \leq M$, then this is a (possibly degenerate) *m*-dimensional simplex in \mathbb{R}^M . Thus, the simplex spanned by e_{i_0}, \ldots, e_{i_m} in \mathbb{R}^M is in the nerve of \mathcal{U} if, and only if, the corresponding sets U_{i_0}, \ldots, U_{i_m} have a common intersection.

The partition of unity $\{\phi_i\}_{i \in I}$ allows us to map X naturally to $\operatorname{Ner}(\mathcal{U})$. Namely, define $\phi \colon X \to \operatorname{Ner}(\mathcal{U})$ by

$$\phi(x) = \sum_{i \in I} \phi_i(x) \cdot e_i, \quad x \in X,$$
(3.2)

and note that ϕ is continuous. In fact, as each ϕ_i is Lipschitz, the map ϕ will be Lipschitz as well. The fact that $\phi(x) \in \operatorname{Ner}(\mathcal{U})$ follows immediately from the definition of the nerve, the characterization in (3.1), and the properties (ii) and (iii) above.

It will be useful for us later to subdivide the simplices in the nerve without changing $Ner(\mathcal{U})$ as a set in \mathbb{R}^M . The *barycentric subdivision* allows us to do this in a canonical way. Once again, we work with a geometric realization of the relevant complexes.

Let $S \subset \mathbb{R}^M$ be a simplicial complex whose simplices are convex hulls of the standard basis vectors e_i . For each collection $\{e_{i_0}, \ldots, e_{i_m}\}$ of vertices which generate a simplex in S, we define its *barycenter* to be the point

$$bc(e_{i_0},\ldots,e_{i_m}) = \frac{1}{m+1} (e_{i_0} + \ldots + e_{i_m}).$$

The subdivision proceeds inductively, by dimension, on the simplices in S. Intuitively, we may think about it in the following way. First, subdivide each edge by adding a vertex at $bc(e_i, e_j)$ whenever $conv(e_i, e_j)$ is in S. Second, subdivide each 2-dimensional simplex by adding a vertex at $bc(e_i, e_j, e_k)$ whenever $conv(e_i, e_j, e_k)$ is in S, and then add edges from $bc(e_i, e_j, e_k)$ to each vertex on the boundary of $conv(e_i, e_j, e_k)$ (these vertices may come from S itself or from the first step in the subdivision). Continue in the same way, until each simplex in S has been subdivided. For further reference, see [33, p. 119-20].

The resulting simplicial complex is called the first barycentric subdivision of S. Observe that the geometric realizations of the complexes we obtain throughout this process, including in the final step, coincide as sets in \mathbb{R}^M with S. We will, however, use S_b to denote the geometric realization of this new complex to emphasize the fact that we have a refined simplicial structure. The following fact will be important in later arguments. If $\operatorname{conv}(e_{i_0},\ldots,e_{i_m})$ is an *m*-dimensional simplex in S, then after the barycentric subdivision, it is a union of *m*-dimensional simplices in S_b with geometric form

$$\operatorname{conv}(p_0,\ldots,p_m),\tag{3.3}$$

where $p_j = bc(e_{\sigma(i_0)}, \ldots, e_{\sigma(i_j)})$ for each $0 \le j \le m$, and each σ is a permutation of the indices i_0, \ldots, i_m . In particular, $p_0 = e_{\sigma(i_0)}$.

Before concluding this section, let us record a topological lemma that will be useful more than once in subsequent sections. Let P be a compact, convex set in \mathbb{R}^n . We say that a closed half-space H of \mathbb{R}^n supports the set P if $P \cap H$ is non-empty and is contained in ∂H .

Lemma 3.8. Let $P \subset \mathbb{R}^n$ be compact and convex, with non-empty interior. Suppose that $f: P \to \mathbb{R}^n$ is continuous, and for each $x \in \partial P$ there is a closed half-space H which supports P, with $x, f(x) \in H$. Then $P \subset f(P)$.

Proof. As f(P) is compact and P is the closure of int(P), it suffices to show that $int(P) \subset f(P)$. Aiming for a contradiction, let us suppose that there is $y \in int(P) \setminus f(P)$. By translation, we may assume, without loss of generality, that y = 0.

The compactness and convexity of P imply that for each $x \in \mathbb{R}^n \setminus \{0\}$, the ray from 0 through x intersects ∂P in exactly one point, which we denote by p(x). We claim that the map $p: \mathbb{R}^n \setminus \{0\} \to \partial P$ is continuous. To verify this, fix $\epsilon > 0$ small enough that $\overline{B}(0,\epsilon) \subset \operatorname{int}(P)$, and let $\pi: \mathbb{R}^n \setminus \{0\} \to \partial B(0,\epsilon)$ be the canonical projection onto the sphere of radius ϵ . It is clear that π is continuous. Observe also that $p(x) = p(\pi(x))$ for all $x \in \mathbb{R}^n \setminus \{0\}$, so it suffices to show that p, restricted to $\partial B(0,\epsilon)$, is continuous. For this, we note that $p|_{\partial B(0,\epsilon)}$ is the inverse map of $\pi|_{\partial P}$, the latter of which is a continuous bijection from the compact set ∂P to $\partial B(0,\epsilon)$. Consequently, it is a homeomorphism, so its inverse $p|_{\partial B(0,\epsilon)}$ is also continuous.

Consider the map $g: P \to \partial P$ defined by g(x) = p(-f(x)). This is continuous by the assumption that $0 \notin f(P)$. We claim that g has no fixed point. Indeed, if g(x) = x, then necessarily $x \in \partial P$. By hypothesis, there is a half-space H of \mathbb{R}^n which supports P and has $x, f(x) \in H$. As 0 lies in the complement of H, we know that -f(x) is in the complement of H as well. The point p(-f(x)) lies on the segment joining 0 and -f(x), and so it also fails to lie in H. This, however, contradicts the fact that $g(x) = x \in H$, so g can have no fixed points. The existence of such a map g contradicts the Brouwer fixed point theorem: any continuous map from a compact, convex set in \mathbb{R}^n to itself has a fixed point. Thus, we obtain $int(P) \subset f(P)$, as desired.

3.3 A topological length-volume inequality for cubes

Let $[0,1]^n$ be the standard Euclidean unit cube of dimension $n \ge 1$. We will use F_k and F'_k , for $1 \le k \le n$, to denote the pairs of opposite codimension-1 faces of $[0,1]^n$:

$$F_k = [0,1]^n \cap \pi_k^{-1}(\{0\})$$
 and $F'_k = [0,1]^n \cap \pi_k^{-1}(\{1\}),$

where $\pi_k \colon \mathbb{R}^n \to \mathbb{R}$ is the projection to the k-th coordinate axis.

Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of $[0, 1]^n$, and let $w_k \colon I \to [0, \infty)$ be corresponding weight functions for $1 \leq k \leq n$. Here, and in what follows, an open cover of $[0, 1]^n$ will always mean that the sets are open in the relative topology on $[0, 1]^n$, unless otherwise explicitly stated. Using the notation from Section 2.1, we let $d_k = \operatorname{dist}_{w_k}(F_k, F'_k)$. To prove Theorem 3.5, we must show that

$$\sum_{i \in I} \left(\prod_{k=1}^{n} w_k(i) \right) \ge d_1 \cdots d_n.$$
(3.4)

To obtain this inequality, we first work under an additional technical assumption on the open cover \mathcal{U} . Namely, if there exists $i \in I$ for which $U_i \cap F_k \neq \emptyset$ and $U_i \cap F'_k \neq \emptyset$ for some k, then we say that \mathcal{U} is *spanning*. We will verify the desired inequality first in the case that \mathcal{U} is *non-spanning*. After doing this, we will treat the general case by modifying slightly the open cover under consideration. Let us state the intermediate result as a separate proposition.

Proposition 3.9. The inequality in (3.4) holds under the hypothesis that the cover \mathcal{U} is non-spanning.

Proof. We may, of course, assume that $d_k > 0$ for each k; otherwise, the desired inequality is trivial. We may also assume that the cover \mathcal{U} is finite. Indeed, compactness guarantees that any cover \mathcal{U} of the cube contains a finite sub-cover. Removing the "redundant" sets from this collection does not increase the left-hand side of the desired inequality and also does not decrease the distances d_k . Our proof will now proceed in several steps, which we explicitly indicate. Step 1: Associate a rectangle to each U_i . For $i \in I$, define

$$d_k(i) = \begin{cases} 0 & \text{if } U_i \cap F_k \neq \emptyset, \\ \text{dist}_{w_k}(F_k, U_i) & \text{otherwise} \end{cases}$$

for $1 \le k \le n$. Of course, we have $d_k(i) \ge 0$ for each *i* and *k*. The more important property is the following: if $U_i \cap F'_k \ne \emptyset$, then $d_k \le d_k(i) + w_k(i)$. This follows immediately from the relevant definitions.

Now let

$$R_{i} = \prod_{k=1}^{n} [d_{k}(i), d_{k}(i) + w_{k}(i)],$$

which is an *n*-dimensional rectangle with side lengths $w_k(i)$. To simplify notation, we let

$$I_k(i) = \pi_k(R_i) = [d_k(i), d_k(i) + w_k(i)]_{i=1}^{k}$$

so that $R_i = \prod_k I_k(i)$.

We will use R_i as a sort of proxy for the set U_i . It will therefore be important that the combinatorics of the rectangles $\{R_i\}_{i \in I}$ mimic those of the sets $\{U_i\}_{i \in I}$, in the following sense.

Lemma 3.10. If $U_i \cap U_j \neq \emptyset$, then $R_i \cap R_j \neq \emptyset$.

Proof. We simply need to show that $I_k(i) \cap I_k(j) \neq \emptyset$ for each k. To this end, fix k and without loss of generality, assume that $d_k(i) \leq d_k(j)$. We claim that $d_k(j) \leq d_k(i) + w_k(i)$. Indeed, there is a chain U_{i_1}, \ldots, U_{i_l} that connects F_k and U_i of length $d_k(i) = \text{dist}_{w_k}(F_k, U_i)$; in case $U_i \cap F_k \neq \emptyset$, this is simply the empty chain. As $U_i \cap U_j \neq \emptyset$, the augmented chain $U_{i_1}, \ldots, U_{i_l}, U_i$ connects F_k and U_j . Thus, $d_k(j) \leq d_k(i) + w_k(i)$, which, along with the assumption that $d_k(i) \leq d_k(j)$, immediately gives $I_k(i) \cap I_k(j) \neq \emptyset$.

We should remark here that the converse need not hold; there are many configurations, in fact, for which two rectangles intersect even though the corresponding open sets are disjoint.

Corollary 3.11. If $U_{i_0} \cap \cdots \cap U_{i_m} \neq \emptyset$, then $R_{i_0} \cap \cdots \cap R_{i_m} \neq \emptyset$.

Proof. As before, it suffices to show that $I_k(i_0) \cap \cdots \cap I_k(i_m) \neq \emptyset$ for each k. From the previous lemma, we know that $R_{i_j} \cap R_{i_{j'}} \neq \emptyset$ for any pair j, j'; in particular, $I_k(i_j) \cap I_k(i_{j'}) \neq \emptyset$. Thus, the m+1 intervals $I_k(i_0), \ldots, I_k(i_m)$ have pairwise non-empty intersections. This immediately implies that $I_k(i_0), \ldots, I_k(i_m)$ have a point of common intersection: indeed, the maximum among their left endpoints is at most the minimum among their right endpoints.

Now that we have established the correspondence between the combinatorics of $\{U_i\}_{i\in I}$ and $\{R_i\}_{i\in I}$, we wish to map the unit cube $[0,1]^n$ continuously into $\bigcup_i R_i$. To ensure that the image is within the union of the R_i 's, it is technically convenient to pass through the nerve of the cover \mathcal{U} . Recall from Section 3.2 that by enumerating $\mathcal{U} = \{U_1, \ldots, U_M\}$, we can express

$$\operatorname{Ner}(\mathcal{U}) = \bigcup \{ \operatorname{conv}(e_{i_0}, \dots, e_{i_m}) : U_{i_0} \cap \dots \cap U_{i_m} \neq \emptyset \}.$$

The associated partition of unity $\{\phi_i\}$ subordinate to \mathcal{U} , which we constructed in the previous section, gives the continuous map

$$\phi \colon [0,1]^n \to \operatorname{Ner}(\mathcal{U})$$

that was introduced in (3.2).

Step 2: Map $\operatorname{Ner}(\mathcal{U})$ into $\bigcup_i R_i$. In order to map $\operatorname{Ner}(\mathcal{U})$ into the union of the rectangles R_i , we will pass to the first barycentric subdivision of the nerve and then define our map simplicially. For ease, we use S to denote the complex $\operatorname{Ner}(\mathcal{U})$ and, consistent with earlier notation, the complex obtained after the subdivision will be denoted by S_b . As sets in \mathbb{R}^M , the complexes S and S_b coincide; moreover, each vertex in S_b arises as the barycenter of a simplex in S.

To define $\psi: S_b \to \bigcup_i R_i$, let us first determine where it sends the vertices. Fix such a vertex p, so that $p = bc(e_{i_0}, \ldots, e_{i_m})$ for some simplex, $conv(e_{i_0}, \ldots, e_{i_m})$, in the nerve S. Note that the choice of e_{i_0}, \ldots, e_{i_m} is uniquely determined by p, up to order. Then, as $U_{i_0} \cap \cdots \cap U_{i_m} \neq \emptyset$, Corollary 3.11 guarantees that

$$R_{i_0} \cap \dots \cap R_{i_m} \neq \emptyset.$$

We want to send p to a point z_p in this intersection, but we must be careful how to choose it. Recall that

$$R_{i_0} \cap \dots \cap R_{i_m} = \left(\prod_{k=1}^n I_k(i_0)\right) \cap \dots \cap \left(\prod_{k=1}^n I_k(i_m)\right)$$

$$= \prod_{k=1}^n I_k(i_0) \cap \dots \cap I_k(i_m),$$

(3.5)

so choosing z_p in $R_{i_0} \cap \cdots \cap R_{i_m}$ amounts to choosing each coordinate $\pi_k(z_p)$ in the interval

$$[a_k, b_k] := I_k(i_0) \cap \dots \cap I_k(i_m)$$

We do this according to the following rule. If $U_{i_j} \cap F_k \neq \emptyset$ for each j, then we choose $\pi_k(z_p) = a_k$; observe that in this case, $a_k = 0$. Otherwise, we choose $\pi_k(z_p) = b_k$.

Let $\psi(p) = z_p$ be as above for the vertices p of S_b . Extend ψ to be affine on each simplex in S_b so that $\psi: S_b \to \mathbb{R}^n$ is continuous. We claim that the image is contained in $\bigcup_i R_i$. To see this, first observe that we may express S_b as a union of simplices Δ that are obtained by subdividing a simplex in S of the same dimension. By (3.3) in the previous section, such simplices have geometric form

$$\Delta = \operatorname{conv}(p_0, \ldots, p_m),$$

where $p_j = bc(e_{i_0}, \ldots, e_{i_j})$ and $conv(e_{i_0}, \ldots, e_{i_m})$ is a simplex in S. Consequently,

$$\psi(\Delta) = \operatorname{conv}(\psi(p_0), \dots, \psi(p_m)) = \operatorname{conv}(z_{p_0}, \dots, z_{p_m})$$

The choice of z_{p_j} guarantees that

$$z_{p_i} \in R_{i_0} \cap \dots \cap R_{i_i} \subset R_{i_0}$$

for each $j = 0, \ldots, m$. As R_{i_0} is convex, we find $\psi(\Delta) \subset R_{i_0}$.

Step 3: Map $[0,1]^n$ into $\bigcup_i R_i$. We now, of course, want to compose $\phi \colon [0,1]^n \to S$ and $\psi \colon S_b \to \bigcup_i R_i$ to obtain a map from the unit cube into the collection of rectangles. Recall that the complexes S and S_b coincide as sets in \mathbb{R}^M , so we can define

$$f = \psi \circ \phi \colon [0, 1]^n \to \bigcup_{i \in I} R_i,$$

which is continuous. Our goal now is to show that the image of f contains the *n*-dimensional rectangle $R = \prod_k [0, d_k]$.

The main claim that we must establish toward this end is that

$$f(F_k) \subset \pi_k^{-1}(\{0\})$$
 and $f(F'_k) \subset \pi_k^{-1}([d_k, \infty)).$

From here, Lemma 3.8 almost immediately implies that $R \subset f([0,1]^n)$. To begin, let $x \in F_k \cup F'_k$, and let U_{i_0}, \ldots, U_{i_m} be the sets in \mathcal{U} that contain x. Then

$$\phi(x) = \sum_{j=0}^{m} \phi_{i_j}(x) e_{i_j},$$

and $x \in \bigcap_j U_{i_j}$ implies that $\operatorname{conv}(e_{i_0}, \ldots, e_{i_m})$ is a simplex in $\operatorname{Ner}(\mathcal{U})$. Also observe that if $x \in F_k$, then $U_{i_j} \cap F_k \neq \emptyset$ for each j; similarly, if $x \in F'_k$, then $U_{i_j} \cap F'_k \neq \emptyset$ for each j.

As $\phi(x) \in \operatorname{conv}(e_{i_0}, \ldots, e_{i_m})$, after the barycentric subdivision, we know that

$$\phi(x) \in \Delta = \operatorname{conv}(p_0, \dots, p_m),$$

where $p_j = bc(e_{i_0}, \ldots, e_{i_j})$ for each j (without loss of generality, we may re-order the indices so that σ is the identity permutation). Consequently, $f(x) = \psi(\phi(x))$ is a convex combination of the points $\psi(p_j) = z_{p_j}$. It therefore suffices to show that $\pi_k(z_{p_j}) = 0$ for each j if $x \in F_k$, and that $\pi_k(z_{p_j}) \ge d_k$ for each j if $x \in F'_k$.

In the former case, we have $U_{i_j} \cap F_k \neq \emptyset$, so that $d_k(i_j) = 0$ for each j. Consequently, we know that $I_k(i_j) = [0, w_k(i_j)]$, so

$$I_k(i_0) \cap \cdots \cap I_k(i_j) = [0, b_k(j)]$$

for some $b_k(j) \ge 0$. The choice of z_{p_j} then guarantees that $\pi_k(z_{p_j}) = 0$ for each j.

In the latter case, we have $U_{i_j} \cap F'_k \neq \emptyset$, so that

$$d_k \le d_k(i_j) + w_k(i_j)$$

for each j. In particular,

$$I_k(i_0) \cap \dots \cap I_k(i_j) = [a_k(j), b_k(j)]$$

for some $a_k(j) \ge 0$ and $b_k(j) \ge d_k$. As \mathcal{U} is non-spanning, we know that $U_{i_j} \cap F_k = \emptyset$ for each j. By the choice of z_{p_j} , we therefore have

$$\pi_k(z_{p_j}) = b_k(j) \ge d_k$$

for each j, as desired.

It is now straightforward to conclude the proof using Lemma 3.8. Namely, let H_k be the half-space $\pi_k^{-1}([0,\infty))$, and let H'_k be the half-space $\pi_k^{-1}((-\infty, d_k])$, so that

$$R = \left(\bigcap_{k=1}^{n} H_k\right) \cap \left(\bigcap_{k=1}^{n} H'_k\right).$$

Let $G_k = R \cap \partial H_k$ and $G'_k = R \cap \partial H'_k$ be the faces of R corresponding to F_k and F'_k , respectively, and let $g \colon R \to [0,1]^n$ be the linear map with $g(G_k) = F_k$ and $g(G'_k) = F'_k$. We showed above that $f(F_k) \subset H^c_k$ and $f(F'_k) \subset H'^c_k$, so the composition $f \circ g \colon R \to \mathbb{R}^n$ has $f \circ g(G_k) \subset H^c_k$ and $f \circ g(G'_k) \subset H'^c_k$. Lemma 3.8 then guarantees that

$$R \subset f \circ g(R) = f([0,1]^n).$$

As $f([0,1]^n) \subset \bigcup_i R_i$, volume considerations immediately give

$$d_1 \cdots d_n = \operatorname{Vol}(R) \le \sum_{i \in I} \operatorname{Vol}(R_i) = \sum_{i \in I} \left(\prod_{k=1}^n w_k(i) \right),$$

as desired.

It is not difficult now to prove Theorem 3.5; we only need to argue that the non-spanning assumption in Proposition 3.9 is not necessary.

Proof of Theorem 3.5. Let \mathcal{U} be an open cover of $[0,1]^n$, let w_k be associated weight functions, and let $d_k = \operatorname{dist}_{w_k}(F_k, F'_k)$ be the corresponding distances, as in the statement of the theorem. Just as in the beginning of the proof of Proposition 3.9, it suffices to assume that \mathcal{U} is finite. Our goal is to modify the cover and the weights slightly in order to obtain a new cover to which we can apply Proposition 3.9. We will do this in such a way that the "volume" and the "lengths" associated to the new cover are very close to the original quantities. We will perform this modification in multiple steps.

First, we wish to modify \mathcal{U} to obtain an open cover of the cube so that any two sets either intersect or have strictly positive distance from each other. To this end, let $\delta_1 > 0$ be small enough that for each $x \in [0, 1]^n$, the ball $B(x, \delta_1)$ is entirely contained in some U_i . Also, let $\delta_2 > 0$ be small enough so that whenever $U_i \cap U_j \neq \emptyset$, there is some point z with $B(z, \delta_2) \subset U_i \cap U_j$. Similarly, let $\delta_3 > 0$ be small enough so that whenever $U_i \cap F_k \neq \emptyset$, there is $z \in F_k$ with $B(z, \delta_3) \subset U_i$, and whenever $U_i \cap F'_k \neq \emptyset$, there is $z \in F'_k$ with $B(z, \delta_3) \subset U_i$. Now define $\delta = \min\{\delta_1, \delta_2, \delta_3\}$.

For each i, let $A_i = \{j : U_i \cap U_j = \emptyset\}$, $B_i = \{k : U_i \cap F_k = \emptyset\}$, and $B'_i = \{k : U_i \cap F'_k = \emptyset\}$. We then form the sets

$$\tilde{U}_i = U_i \setminus \left(\bigcup_{j \in A_i} \overline{N}_{\delta/2}(U_j) \cup \bigcup_{k \in B_i} \overline{N}_{\delta/2}(F_k) \cup \bigcup_{k \in B'_i} \overline{N}_{\delta/2}(F'_k) \right)$$

where $\overline{N}_{\epsilon}(V)$ denotes the closed ϵ -neighborhood of V. Each \tilde{U}_i is open, and as $\delta \leq \delta_1$, it is clear that $\bigcup_i \tilde{U}_i$ contains $[0, 1]^n$.

We claim that for each i and j, either $\tilde{U}_i \cap \tilde{U}_j \neq \emptyset$ or $\operatorname{dist}(\tilde{U}_i, \tilde{U}_j) \geq \delta/2$. Indeed, if $\operatorname{dist}(\tilde{U}_i, \tilde{U}_j) < \delta/2$, then there are $x \in \tilde{U}_i$ and $y \in \tilde{U}_j$ with $|x - y| < \delta/2$. This implies that $U_i \cap U_j \neq \emptyset$, for if not, then $j \in A_i$ so that x could not be in \tilde{U}_i . Choose $z \in U_i \cap U_j$ with $B(z, \delta) \subset U_i \cap U_j$, which is possible because $\delta \leq \delta_2$. Then it must be that $z \in \tilde{U}_i \cap \tilde{U}_j$. Indeed, suppose that $z \notin \tilde{U}_i$. Then either there is $l \in A_i$ with $z \in \overline{N}_{\delta/2}(U_l)$ or there is $k \in B_i \cup B'_i$ with $z \in \overline{N}_{\delta/2}(F_k) \cup \overline{N}_{\delta/2}(F'_k)$. However, as $B(z, \delta) \subset U_i$, the distance from z to any of these U_l, F_k , or F'_k is strictly larger than $\delta/2$. This immediately rules out $z \in \overline{N}_{\delta/2}(U_l)$ or $z \in \overline{N}_{\delta/2}(F_k) \cup \overline{N}_{\delta/2}(F'_k)$. The argument for $z \in \tilde{U}_j$ is the same.

Similarly, we can also show that for each i and k, either $\tilde{U}_i \cap F_k \neq \emptyset$ or $\operatorname{dist}(\tilde{U}_i, F_k) \geq \delta/2$. Indeed, if $\operatorname{dist}(\tilde{U}_i, F_k) < \delta/2$, then $U_i \cap F_k \neq \emptyset$. This implies that there is $z \in F_k$ for which $B(z, \delta) \subset U_i$. As $z \notin \overline{N}_{\delta/2}(U_j)$ for each $j \in A_i$, we necessarily have $z \in \tilde{U}_i$. Hence, $\tilde{U}_i \cap F_k \neq \emptyset$. The same arguments also show that for each i and k, either $\tilde{U}_i \cap F'_k \neq \emptyset$ or $\operatorname{dist}(\tilde{U}_i, F'_k) \geq \delta/2$. Thus, the collection $\{\tilde{U}_i\}_{i \in I}$ has the convenient property that for every incidence relevant to the calculation of a combinatorial distance, the associated sets either intersect or are of distance $\geq \delta/2$ from each other.

Let us again modify the collection $\{\tilde{U}_i\}_{i\in I}$ slightly, in a way dependent on a parameter $\epsilon > 0$ that we will eventually send to 0. Namely, let $0 < \epsilon < \delta/(8\sqrt{n})$ be very small. For each i, let $V_i = N_{\epsilon/2}(\tilde{U}_i)$, where the neighborhood is now taken in \mathbb{R}^n . Thus, each V_i is open in \mathbb{R}^n , and the union $\bigcup_i V_i$ contains $[0, 1]^n$ but does not intersect any of the half-spaces $\pi_k^{-1}([1 + \epsilon/2, \infty))$.

To the collection $\{V_i\}_{i \in I}$ we add small Euclidean balls to produce a cover of the cube $[0, 1 + \epsilon]^n$. Namely, we can find a collection of points $\{x_j\}_{j \in J}$ with the following properties: $\#J \leq C_n(1/\epsilon)^{n-1}$, where C_n is a dimensional constant; each point x_j lies in one of the codimension-1 spaces $\{x \in \mathbb{R}^n : \pi_k(x) = 1 + \epsilon\}$; and the balls $B_j = B(x_j, \sqrt{n\epsilon})$ have

$$(1, 1+\epsilon]^n \subset \bigcup_{j\in J} B_j.$$

Let \mathcal{V} denote the collection of open sets $\{V_i\}_{i \in I} \cup \{B_j\}_{j \in J}$ so that \mathcal{V} is an open cover of the cube $[0, 1 + \epsilon]^n$. Let

$$G_k = \pi_k^{-1}(1+\epsilon) \cap [0, 1+\epsilon]^n$$

be the codimension-1 face of $[0, 1+\epsilon]^n$ opposite to F_k , and observe that no set in \mathcal{V} intersects both F_k and G_k . In other words, \mathcal{V} is non-spanning (the fact that we are covering a slightly larger cube is not a problem; indeed, Proposition 3.9 applies equally well to topological cubes).

To obtain weight functions for \mathcal{V} , we of course want to use the original weights w_k associated to the cover \mathcal{U} . Namely, let v_k be weight functions for \mathcal{V} defined by

$$v_k(V_i) = w_k(U_i)$$
 and $v_k(B_j) = \epsilon$,

and let $\tilde{d}_k = \operatorname{dist}_{v_k}(F_k, G_k)$ be the associated distance between opposite faces of the cube. We claim that $d_k \leq \tilde{d}_k$ for each k. To see this, let us first establish that any chain in \mathcal{V} of minimal v_k -length that connects F_k to G_k has the form

$$V_{i_1},\ldots,V_{i_m},B_j$$

for some collection i_1, \ldots, i_m and some j. It is clear that the chain must end with some B_j , as none of the V_i intersect G_k . Also, each ball B_j intersects G_k , so the penultimate set in the chain cannot be some other ball B_l . Lastly, note that if $V_i, B_j, V_{i'}$ appears in the chain, then $\operatorname{dist}(V_i, V_{i'}) \leq \operatorname{diam} B_j \leq 2\sqrt{n\epsilon}$, so that

$$\operatorname{dist}(\tilde{U}_i, \tilde{U}_{i'}) \le 4\sqrt{n}\epsilon < \delta/2$$

Consequently, $\tilde{U}_i \cap \tilde{U}_{i'} \neq \emptyset$, which also means that $V_i \cap V_{i'} \neq \emptyset$. Thus, in a minimal chain, a ball B_j never appears between two of the V_i 's.

Let $V_{i_1}, \ldots, V_{i_m}, B_j$ be a chain of minimal v_k -length from F_k to G_k . As $V_{i_l} \cap V_{i_{l+1}} \neq \emptyset$ for each l, we know that $\operatorname{dist}(\tilde{U}_{i_l}, \tilde{U}_{i_{l+1}}) \leq \epsilon < \delta/2$. Thus, $\tilde{U}_{i_l} \cap \tilde{U}_{i_{l+1}} \neq \emptyset$, so also $U_{i_l} \cap U_{i_{l+1}} \neq \emptyset$. Hence, U_{i_1}, \ldots, U_{i_m} is a chain in the collection \mathcal{U} . Moreover, $\operatorname{dist}(\tilde{U}_{i_1}, F_k) \leq \epsilon/2 < \delta/2$, so $\tilde{U}_{i_1} \cap F_k \neq \emptyset$, and also $U_{i_1} \cap F_k \neq \emptyset$. Similarly,

$$\operatorname{dist}(\tilde{U}_{i_m}, F'_k) \le \epsilon/2 + \operatorname{diam} B_j \le \epsilon/2 + 2\sqrt{n\epsilon} < \delta/2,$$

so that $\tilde{U}_{i_m} \cap F'_k \neq \emptyset$, and also $U_{i_m} \cap F'_k \neq \emptyset$. Therefore, the chain U_{i_1}, \ldots, U_{i_m} connects F_k and F'_k , which implies that

$$d_k \le \sum_{l=1}^m w_k(U_{i_l}) = \sum_{l=1}^m v_k(V_{i_l}) = \tilde{d}_k - \epsilon \le \tilde{d}_k.$$

Applying Proposition 3.9 to the collection \mathcal{V} with weight functions v_k gives

$$\prod_{k=1}^{n} d_k \le \prod_{k=1}^{n} \tilde{d}_k \le \sum_{i \in I} \left(\prod_{k=1}^{n} v_k(V_i) \right) + \sum_{j \in J} \epsilon^n \le \sum_{i \in I} \left(\prod_{k=1}^{n} w_k(U_i) \right) + C_n \epsilon,$$

where the last inequality follows from the bound $\#J \leq C_n(1/\epsilon)^{n-1}$. As this holds for any $0 < \epsilon < \delta/(8\sqrt{n})$, we send ϵ to zero to obtain

$$\prod_{k=1}^{n} d_k \le \sum_{i \in I} \left(\prod_{k=1}^{n} w_k(i) \right).$$

3.4 Lower volume bounds in metric spaces

Using Theorem 3.5, we can prove a similar length-volume inequality for *images* of the Euclidean cube in a metric space. To set this up, let (X,d) be a metric space and let $g: [0,1]^n \to X$ be a continuous map. Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of $g([0,1]^n)$ with corresponding non-negative weight functions w_k for $1 \leq k \leq n$. If $A, B \subset g([0,1]^n)$ are subsets, then we define

$$\operatorname{dist}_{w_k}(A, B) = \inf \left\{ \sum_{j=1}^m w_k(i_j) : \begin{array}{c} U_{i_1}, \dots, U_{i_m} \text{ is a chain} \\ \text{that connects } A \text{ and } B \end{array} \right\}$$

as before, where *chains* are finite sequences of the $\{U_i\}_{i \in I}$ whose consecutive sets have nonempty intersection. Let

$$d_k(g) = \operatorname{dist}_{w_k}(g(F_k), g(F'_k))$$

be the discrete distance between the images of opposite faces. Of course, this distance depends strongly on \mathcal{U} and w_k as well.

Proposition 3.12. For such $g \colon [0,1]^n \to X$, we have

$$\sum_{i \in I} \left(\prod_{k=1}^{n} w_k(i) \right) \ge \prod_{k=1}^{n} d_k(g).$$

Proof. For each $i \in I$, let $V_i = g^{-1}(U_i)$, so that the collection $\mathcal{V} = \{V_i\}_{i \in I}$ forms an open cover of $[0,1]^n$. Let $d_k = \operatorname{dist}_{w_k}(F_k, F'_k)$ be the discrete distance associated to the cover \mathcal{V} . Observe that if V_{i_1}, \ldots, V_{i_m} is a chain in \mathcal{V} that connects F_k and F'_k , then U_{i_1}, \ldots, U_{i_m} is a chain in \mathcal{U} that connects $g(F_k)$ and $g(F'_k)$. Consequently, we have $d_k(g) \leq d_k$ for each k. Applying Theorem 3.5 to the cover \mathcal{V} with weights w_k gives

$$\prod_{k=1}^{n} d_k(g) \le \prod_{k=1}^{n} d_k \le \sum_{i \in I} \left(\prod_{k=1}^{n} w_k(i) \right),$$

as desired.

Using this proposition, we can establish a similar inequality relating more standard metric quantities such as Hausdorff measure and metric distance between sets. Recall from Section 3.1 that if (X, d) is a metric space and $E \subset X$ is compact, the Q-dimensional Hausdorff content of E is

$$\mathcal{H}_Q^{\infty}(E) = \inf \left\{ \sum_{i \in I} (\operatorname{diam} U_i)^Q : \{U_i\}_{i \in I} \text{ is an open cover of } E \right\}.$$

The associated Hausdorff Q-dimensional measure is defined to be

$$\mathcal{H}_Q(E) = \liminf_{\epsilon \searrow 0} \left\{ \sum_{i \in I} (\operatorname{diam} U_i)^Q : \{U_i\}_{i \in I} \text{ covers } E \text{ and } \operatorname{diam} U_i < \epsilon \right\},\$$

and it is clear that $\mathcal{H}_Q(E) \geq \mathcal{H}_Q^{\infty}(E)$. Thus, lower bounds on Hausdorff content are also lower bounds on Hausdorff measure.

Corollary 3.13. If $g: [0,1]^n \to X$ is continuous, then

$$\mathcal{H}_n^{\infty}(g([0,1]^n)) \ge \prod_{k=1}^n \operatorname{dist}(g(F_k), g(F'_k)).$$

Proof. Fix an open cover $\{U_i\}_{i\in I}$ of $g([0,1]^n)$, and let $w_k(i) = \operatorname{diam} U_i$ for each i and each $1 \leq k \leq n$. Observe that if U_{i_1}, \ldots, U_{i_m} is a chain connecting $g(F_k)$ and $g(F'_k)$, then

$$dist(g(F_k), g(F'_k)) \le \sum_{j=1}^m diam U_{i_j} = \sum_{j=1}^m w_k(i_j).$$

Thus, $\operatorname{dist}(g(F_k), g(F'_k)) \leq \operatorname{dist}_{w_k}(g(F_k), g(F'_k))$ for each k. By Proposition 3.12, we have

$$\sum_{i \in I} (\operatorname{diam} U_i)^n \ge \prod_{k=1}^n \operatorname{dist}_{w_k}(g(F_k), g(F'_k)) \ge \prod_{k=1}^n \operatorname{dist}(g(F_k), g(F'_k))$$

As this holds for any open cover $\{U_i\}_{i\in I}$ of $g([0,1]^n)$, we obtain the desired inequality. \Box

Observe that in both Proposition 3.12 and Corollary 3.13, it is not necessary that X be a metric space. The same arguments hold if X is a *pseudometric space*: the distance between distinct points is allowed to be zero. We illustrate this with the following result, which is closely related to a question of Y. Burago and V. Zalgaller [12, p. 296].

Corollary 3.14. Let ρ be a pseudometric on $[0,1]^n$, and assume that every open set in the topology determined by ρ is also open in the Euclidean topology. Then

$$\mathcal{H}_{n,\rho}^{\infty}([0,1]^n) \ge \prod_{k=1}^n \operatorname{dist}_{\rho}(F_k, F'_k).$$

Here, $\mathcal{H}_{n,\rho}^{\infty}([0,1]^n)$ and $\operatorname{dist}_{\rho}(F_k, F'_k)$ are defined in the same manner as the usual Hausdorff content and distance, but using the pseudometric ρ instead of an actual metric. Once again, the definition of Hausdorff content and measure that we use differs from that in [12] by a multiplicative constant. As a result, Corollary 3.14 does not resolve their question.

Proof. The pseudometric ρ on $X = [0,1]^n$ canonically induces a metric $\tilde{\rho}$ on the quotient space $\tilde{X} = X/\sim$, where \sim is the equivalence relation $x \sim x'$ if $\rho(x, x') = 0$. Let $\pi \colon X \to \tilde{X}$ be the associated projection. It is straightforward to see that

$$\operatorname{dist}_{\rho}(F_k, F'_k) = \operatorname{dist}_{\tilde{\rho}}(\pi(F_k), \pi(F'_k)) \quad \text{and} \quad \mathcal{H}^{\infty}_{n,\rho}(X) = \mathcal{H}^{\infty}_{n,\tilde{\rho}}(\tilde{X}).$$
(3.6)

Moreover, $U \subset \tilde{X}$ is open in the metric topology if and only if $\pi^{-1}(U)$ is open in the topology on $[0,1]^n$ determined by ρ . The hypothesis of topological compatibility then ensures that π , when viewed as a map from the Euclidean cube $[0,1]^n$ to the metric space \tilde{X} , is continuous. Corollary 3.13, along with (3.6), gives the desired conclusion.

Corollary 3.13 points us in the following direction: in what generality can one obtain Euclidean-type lower volume bounds in metric spaces? More precisely, for which metric spaces (X, d), does one have

$$\mathcal{H}_n^{\infty}(B(x,r)) \gtrsim r^n \tag{3.7}$$

for all metric balls B(x, r), with $0 < r \le \text{diam } X$? An immediate consequence of Corollary 3.13 is that (3.7) holds whenever X satisfies the following property.

Definition 3.15. A metric space (X, d) is said to *admit fat n-cubes* if there is $\lambda \ge 1$ such that, for each $x \in X$ and $0 < r \le \text{diam } Z$, there is a continuous map $g: [0, 1]^n \to B(x, r)$ with $\text{dist}(g(F_k), g(F'_k)) \ge r/\lambda$ for each k.

Notice that this is a weaker variant of *admitting fat connecting n-cubes*, which was introduced in the previous chapter. In particular, if X admits fat connecting *n*-cubes, then (3.7) holds, with constants depending only on the fat connecting *n*-cube condition. Recall that, for n = 2, this condition is satisfied when (X, d) is complete, annularly linearly connected, and linearly locally contractible. Euclidean-type lower volume bounds cannot, of course, hold in complete generality: the existence of "thin necks" or "thin fingers" in X would hinder large volume in certain regions. Thus, it makes sense to impose connectivity conditions such as annular linear connectivity or linear local contractibility when addressing such questions. In this context, it turns out that there is a fairly general result on lower volume bounds, which relies on the following deep fact proved by S. Semmes.

Theorem 3.16 (Semmes [53, Theorem 1.29(a)]). Let (X, d) be a closed manifold of dimension $n \ge 2$ that is N-doubling and λ -linearly locally contractible. Then for each $x \in X$ and each $0 < r \le \text{diam } X$, there is a surjective map $f: X \to \mathbb{S}^n$ that is C/r-Lipschitz and is constant outside of B(x, r/2). Here, C depends only on n, N, and λ .

From this theorem, we immediately obtain a bound of the form in (3.7). Indeed, as f has non-zero degree, $f(B(x,r)) = \mathbb{S}^n$, so the Lipschitz bounds imply that

$$\mathcal{H}_n^{\infty}(B(x,r)) \ge \left(\frac{r}{C}\right)^n \mathcal{H}_n^{\infty}(\mathbb{S}^n) \gtrsim r^n,$$

where the implicit constant depends only on n, N, and λ .

This "Semmes approach" to lower volume bounds is, in a sense, dual to our original approach, which sought to map a nice space *into* X, rather than map X into some other controlled space. The relative ease of building such maps from a general metric space into Euclidean spaces (for example, as we did in Section 3.2) makes the Semmes method viable. Still, it would be desirable to make our dual argument work. To this end, let us observe that it is generally easier to construct controlled maps from simplices into metric spaces than it is to construct such maps on cubes. In the final section, we therefore take up the topic of diameter-volume bounds on simplices, and we use these to verify lower volume bounds for linearly locally contractible metric spaces, in the case that n = 2.

3.5 Some analogous considerations for simplices

Not surprisingly, the methods we used to prove Theorem 3.5 can be adapted to prove similar inequalities on other convex polyhedra. In the Riemannian setting, Derrick's methods were extended to a much more general framework by M. Gromov, and this includes a diameter-volume inequality for simplices [26, Section 7]. We will not attempt to build an analogous framework here, but motivated by the discussion in the previous section, let us touch upon this problem for simplices.

Let Δ^n denote the standard *n*-dimensional simplex, which we generally view in the form

$$\Delta^n = \operatorname{conv}(e_1, \dots, e_{n+1}) \subset \mathbb{R}^{n+1}.$$

We will denote the coordinate faces by T_k ; these are simply the intersections of Δ^n with the codimension-1 planes where the k-th coordinate is zero. Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of Δ^n , and let $w: I \to [0, \infty)$ be a corresponding weight function. We define the *diameter* of \mathcal{U} , with respect to w, to be

$$d_w(\mathcal{U}) = \inf \left\{ \sum_{j=1}^m w(i_j) : \begin{array}{l} \text{the collection } U_{i_1}, \dots, U_{i_m} \text{ contains a chain} \\ \text{that connects } T_k \text{ and } T_l \text{ for each pair } k, l \end{array} \right\}$$

In this section, we will prove the following theorem.

Theorem 3.17. For an open cover $\mathcal{U} = \{U_i\}_{i \in I}$ of Δ^n and an associated weight function w, we have

$$\sum_{i \in I} w(i)^n \ge \frac{d_w(\mathcal{U})^n}{n!}.$$

Moreover if $w \equiv 1$ is constant, then we have the improved estimate

$$\#\mathcal{U} \ge \binom{n+d_w(\mathcal{U})-1}{n}.$$

The proof of the "weighted" statement will mimic the proof of Theorem 3.5 very closely; in fact, we will only indicate what changes must be made for the simplex. The proof of the "unweighted" statement, however, will more closely mimic the proof of Theorem 3.3. Let us treat this case first. Proof of the unweighted statement. Without loss of generality, we may assume that for each $i \in I$, there is a point $x_i \in U_i$ that lies in no other U_j . Indeed, deleting redundant sets in the cover decreases the left-hand side of the desired inequality without reducing the right hand side. In particular, we may assume that $\#\mathcal{U}$ is finite.

Let us define a map $f_0: \{x_i : i \in I\} \to \mathbb{Z}^{n+1}$, where the k-th component is

$$\pi_k(f_0(x_i)) = \begin{cases} \operatorname{dist}_w(T_k, U_i) & \text{if } U_i \cap T_k = \emptyset, \\ 0 & \text{if } U_i \cap T_k \neq \emptyset. \end{cases}$$

Note that, as $w \equiv 1$, the combinatorial distance $\operatorname{dist}_w(T_k, U_i)$ is simply the smallest number of sets U_j in a chain that connects T_k and U_i . For notational ease, let $y_i = f_0(x_i)$. Observe that the definition of $d = d_w(\mathcal{U})$ implies that

$$\pi_1(y_i) + \ldots + \pi_{n+1}(y_i) \ge d - 1.$$

Indeed, choosing minimal-length chains from U_i to T_k for each k, and adding in the set U_i , gives a collection of at most $\pi_1(y_i) + \ldots + \pi_{n+1}(y_i) + 1$ sets in \mathcal{U} which, for each k and l, contains a chain that connects T_k and T_l . Let

$$C = \{ z \in \mathbb{R}^{n+1} : \pi_k(z) \ge 0 \text{ for all } k, \text{ and } \pi_1(z) + \ldots + \pi_{n+1}(z) \ge d-1 \}.$$

Note that C is convex and that $y_i \in C$ for each $i \in I$.

Now extend f_0 to a continuous map on Δ^n via a partition of unity $\{\phi_i\}_{i\in I}$ subordinate to $\mathcal{U} = \{U_i\}_{i\in I}$:

$$f(x) = \sum_{i \in I} \phi_i(x) f_0(x_i) = \sum_{i \in I} \phi_i(x) y_i, \quad x \in \Delta^n.$$

Observe that f indeed extends f_0 because $\phi_j(x_i) = 0$ for $j \neq i$. Moreover, f(x) is a convex combination of the points y_i , so that $f(x) \in C$. Also, if $x \in T_k$, then $\pi_k(f(x)) = 0$. Indeed, if U_{i_1}, \ldots, U_{i_m} are the sets in \mathcal{U} that contain x, then $f(x) = \sum_{j=1}^m \phi_{i_j}(x)y_{i_j}$. In particular,

$$\pi_k(f(x)) = \sum_{j=1}^m \phi_{i_j}(x) \pi_k(y_{i_j}).$$

The fact that $x \in U_{i_j} \cap T_k$ implies that $\pi_k(y_{i_j}) = 0$ for each j, so we obtain $\pi_k(f(x)) = 0$.

Let $p: \mathbb{R}^{n+1} \to \mathbb{R}^n$ be defined by

$$p(z_1,\ldots,z_n,z_{n+1})=(z_1,\ldots,z_n),$$

and consider the composition $p \circ f$. We claim that $(p \circ f)(\Delta^n)$ contains the *n*-dimensional simplex

$$P = \{(z_1, \dots, z_n) \in \mathbb{R}^n : \text{ all } z_k \ge 0 \text{ and } z_1 + \dots + z_n \le d - 1\}.$$

To see this, let P_k be the face of P with k-th coordinate equal to 0, for $1 \le k \le n$, and let P_{n+1} be the other face, which is defined by

$$P_{n+1} = \{(z_1, \dots, z_n) \in \mathbb{R}^n : \text{ all } z_k \ge 0 \text{ and } z_1 + \dots + z_n = d - 1\}.$$

Observe that there is an affine map $g: P \to \Delta^n$ which sends the faces P_k of P to the faces T_k of Δ^n . Also, let H_k denote the closed half-space of \mathbb{R}^n determined by P_k that does not contain P, namely $H_k = \{(z_1, \ldots, z_n) \in \mathbb{R}^n : z_k \leq 0\}$ for $1 \leq k \leq n$ and

$$H_{n+1} = \{ (z_1, \dots, z_n) \in \mathbb{R}^n : z_1 + \dots + z_n \ge d - 1 \}.$$

As $\pi_k(f(x)) = 0$ for $x \in T_k$, we immediately have $(p \circ f)(T_k) \subset H_k$ for $1 \leq k \leq n$. If $x \in T_{n+1}$, then $\pi_{n+1}(f(x)) = 0$, and as $f(x) \in C$, this guarantees that

$$\pi_1(f(x)) + \ldots + \pi_n(f(x)) \ge d - 1.$$

Consequently, $p \circ f(x) \in H_{n+1}$. We can therefore apply Lemma 3.8 to the map $p \circ f \circ g$, and this gives $P \subset (p \circ f)(\Delta^n)$.

We now claim that if $z \in P \cap \mathbb{Z}^n$, then there is $i \in I$ for which $p \circ f(x_i) = z$. By the previous paragraph, we know that there is $x \in \Delta^n$ with $p \circ f(x) = z$. Let $U_{i_1}, \ldots, U_{i_m} \in \mathcal{U}$ be those sets which contain x, so that

$$f(x) = \sum_{j=1}^{m} \phi(x) y_{i_j}.$$

The fact that $U_{i_j} \cap U_{i_l} \neq \emptyset$ implies immediately that

$$||y_{i_j} - y_{i_l}||_{\infty} \le 1$$

for each $1 \leq j, l \leq m$, where $||y||_{\infty} = \max\{|\pi_k(y)| : 1 \leq k \leq n+1\}$ is the ℓ^{∞} -norm. Thus, for each k, there is an integer a_k such that

$$\pi_k(y_{i_j}) \in \{a_k, a_k+1\}$$

for all $1 \leq j \leq m$, namely, $a_k = \min\{\pi_k(y_{i_j}) : 1 \leq j \leq m\}$. Let

$$I_k = \{j : \pi_k(y_{i_j}) = a_k\}$$

so that $I_k \neq \emptyset$ for each k.

As $p \circ f(x) \in \mathbb{Z}^n$, we know that the first *n* coordinates of f(x) are integers. For $1 \le k \le n$, we can write

$$\pi_k(f(x)) = \sum_{j=1}^m \phi_{i_j}(x)\pi_k(y_{i_j}) = \sum_{j\in I_k} \phi_{i_j}(x)a_k + \sum_{j\notin I_k} \phi_{i_j}(x)(a_k+1)$$
$$= a_k + \sum_{j\notin I_k} \phi_{i_j}(x).$$

As $\pi_k(f(x))$ is an integer, we know that

$$\sum_{j \notin I_k} \phi_{i_j}(x)$$

is also an integer, necessarily equal to 0 or 1. This can happen only if $I_k = \emptyset$ or $I_k = \{1, \ldots, m\}$, but the former is ruled out by the definition of a_k . Thus, $I_k = \{1, \ldots, m\}$ for each $1 \le k \le n$, which implies that each y_{i_j} has k-th coordinate a_k . In particular,

$$z = p \circ f(x) = (a_1, \dots, a_n) = y_{i_1} = p \circ f(x_{i_1}),$$

as desired.

From the previous claim, we can conclude that $\#\mathcal{U} \ge \#(P \cap \mathbb{Z}^n)$. It is now essentially a combinatorial exercise to show that

$$\#(P \cap \mathbb{Z}^n) = \binom{n+d-1}{n}.$$

Indeed, if p_r denotes the number of integer lattice points $z = (z_1, \ldots, z_n)$ with

$$z_k \ge 0$$
 for all k and $z_1 + \ldots + z_n = r$,

then $\#(P \cap \mathbb{Z}^n) = p_0 + \ldots + p_{d-1}$. As p_r is the number of ways of placing r indistinguishable objects into n distinct bins, we have $p_r = \binom{n+r-1}{n-1}$. Thus,

$$\#(P \cap \mathbb{Z}^n) = \sum_{r=0}^{d-1} \binom{n+r-1}{n-1},$$

and an easy induction on d shows that this sum is $\binom{n+d-1}{n}$.

Before moving on to the proof of the weighted statement, let us make a few remarks about sharpness for the unweighted inequality. Fix an integer $m \ge 1$, and let

$$P = \{(x, y) \in \mathbb{R}^2 : x, y \ge 0 \text{ and } x + y \le m\},\$$

which is affinely equivalent to Δ^2 . From the preceding proof, we have

$$\#(P \cap \mathbb{Z}^2) = \binom{2+m}{2} =: N.$$

Let z_1, \ldots, z_N be the points in this lattice. It is easy to construct a corresponding open cover $\mathcal{U} = \{U_1, \ldots, U_N\}$ of P such that $U_i \cap \mathbb{Z}^2 = \{z_i\}$ for each i, the set U_i intersects ∂P if and only if $z_i \in \partial P$, and

$$U_i \cap U_j \neq \emptyset$$
 if and only if $z_i - z_j \in \{\pm(1,0), \pm(0,1), \pm(1,-1)\}$.

More colloquially, the configuration of these open sets resembles the configuration of bowling pins. Let d denote the diameter of \mathcal{U} , with respect to the constant weight function $w \equiv 1$. We claim that d = m + 1.

To see this, let P_1 , P_2 , and P_3 denote the edges of P on the x-axis, the y-axis, and the diagonal, respectively. The inequality $d \leq m + 1$ is immediate, as the collection of sets in \mathcal{U} corresponding to the m + 1 lattice points on P_1 contains chains connecting each pair of edges. For the opposite inequality, let $\mathcal{C} \subset \mathcal{U}$ be a collection that contains chains connecting each pair of edges. If \mathcal{C} contains any of the three open sets corresponding to the vertices of P, then it is clear that $\#\mathcal{C} \geq m + 1$. Indeed, each vertex has combinatorial distance m + 1 to the opposite edge. Assume, then, that \mathcal{C} contains none of these three sets. This

implies that there is $U_{i_0} \in \mathcal{C}$, along with three chains $V_1^k, \ldots, V_{l_k}^k \in \mathcal{C}$, for k = 1, 2, 3, that connect U_{i_0} to P_k , such that any two of these chains have no sets in common. Here, we allow $V_1^k, \ldots, V_{l_k}^k$ to be an empty chain, in which case $l_k = 0$; this happens only if z_{i_0} lies on P_k . If $z_{i_0} = (x_0, y_0)$, then $\operatorname{dist}_w(U_{i_0}, P_1) = y_0$ and $\operatorname{dist}_w(U_{i_0}, P_2) = x_0$, so $l_1 \ge y_0$ and $l_2 \ge x_0$. Moreover, $\operatorname{dist}_w(U_{i_0}, P_3) = m - x_0 - y_0$, so that $l_3 \ge m - x_0 - y_0$. As the three chains have no sets in common, $\#\mathcal{C} \ge 1 + l_1 + l_2 + l_3 \ge m + 1$. This gives $d \ge m + 1$, as claimed.

Hence, for the open cover \mathcal{U} of P, we obtain $\#\mathcal{U} = \binom{2+m}{2} = \binom{2+d-1}{2}$, so the unweighted inequality is sharp in dimension 2. In higher dimensions, one can form analogous open covers of Δ^n , but computing the combinatorial diameter is more subtle. Let us instead turn to the weighted inequality in Theorem 3.17.

Proof of the weighted statement. Fix a cover $\mathcal{U} = \{U_i\}_{i \in I}$ of the simplex Δ^n , and let $w(i) \geq 0$ be corresponding weights. Again we let $d = d_w(\mathcal{U})$ denote the diameter. Observe that if there is $U_i \in \mathcal{U}$ that intersects each face T_k of Δ^n , then the singleton U_i is a chain that connects T_k and T_l for each pair k, l. In particular, $d \leq w(i)$, so that the desired inequality trivially holds.

We may, therefore, assume that no U_i intersects all of the faces T_k . Moreover, it suffices to assume that $\#\mathcal{U}$ is finite, as deleting redundant sets from the cover can only decrease the left-hand side and increase the right-hand side of the desired inequality. We now run an argument very similar to the proof of Proposition 3.9. Namely, for each $1 \leq k \leq n+1$, we define

$$d_k(i) = \begin{cases} \operatorname{dist}_w(T_k, U_i) & \text{if } U_i \cap T_k = \emptyset \\ 0 & \text{if } U_i \cap T_k \neq \emptyset \end{cases}$$

and then let

$$R_{i} = \prod_{k=1}^{n} [d_{k}(i), d_{k}(i) + w(i)]$$

be an *n*-dimensional cube with side length w(i). Notice that we do not include the interval $[d_{n+1}(i), d_{n+1}(i) + w(i)]$, but we will see that the quantity $d_{n+1}(i)$ is still important in our analysis. By the same reasoning as in Lemmas 3.10 and 3.11, we have that $R_{i_0} \cap \cdots \cap R_{i_m} \neq \emptyset$

whenever $U_{i_0} \cap \cdots \cap U_{i_m} \neq \emptyset$.

Let $\phi: \Delta^n \to \operatorname{Ner}(\mathcal{U})$ be the map that we introduced in (3.2), built from a partition of unity subordinate to the cover \mathcal{U} . As in the proof of Proposition 3.9, we wish to define a map from $\operatorname{Ner}(\mathcal{U})$ to $\bigcup_i R_i$. Using the same notation as before, let $S = \operatorname{Ner}(\mathcal{U})$ and let S_b denote the first barycentric subdivision of S. Each vertex $p \in S_b$ can be expressed uniquely as $p = \operatorname{bc}(e_{i_0}, \ldots, e_{i_m})$ for some simplex $\operatorname{conv}(e_{i_1}, \ldots, e_{i_m})$ in the nerve S. We want this map to send p to a point z_p in the intersection

$$R_{i_0} \cap \dots \cap R_{i_m} = \left(\prod_{k=1}^n [d_k(i_0), d_k(i_0) + w(i_0)] \right) \cap \dots \cap \left(\prod_{k=1}^n [d_k(i_m), d_k(i_m) + w(i_m)] \right)$$
$$= \prod_{k=1}^n [d_k(i_0), d_k(i_0) + w(i_0)] \cap \dots \cap [d_k(i_m), d_k(i_m) + w(i_m)]$$
$$=: \prod_{k=1}^n [a_k, b_k].$$

Let us choose the k-th coordinate of z_p in the following way. If $U_{ij} \cap T_k \neq \emptyset$ for each $1 \leq j \leq m$, then we choose $\pi_k(z_p) = a_k$; note that in this case, $a_k = 0$. Otherwise, we choose $\pi_k(z_p) = b_k$. These choices define z_p , and we let $\psi(p) = z_p$ for vertices $p \in S_b$. Extend ψ to be affine on each simplex in S_b , so that $\psi: S_b \to \mathbb{R}^n$ is continuous. By the same reasoning as in the proof of Proposition 3.9, the image $\psi(S_b)$ is contained in the union $\bigcup_i R_i$.

As before, let us view ψ simply as a map on S, and define

$$f = \psi \circ \phi \colon \Delta^n \to \bigcup_{i \in I} R_i,$$

which is continuous. The same reasoning we used earlier will easily show that if $x \in T_k$ then $\pi_k(f(x)) = 0$, for each $1 \le k \le n$. Slightly more involved is the following claim.

Lemma 3.18. If $x \in T_{n+1}$, then $\pi_1(f(x)) + \ldots + \pi_n(f(x)) \ge d$.

Proof of the lemma. Let U_{i_1}, \ldots, U_{i_m} be the sets in \mathcal{U} that contain $x \in T_{n+1}$, so that

$$\phi(x) = \sum_{j=1}^{m} \phi_{i_j}(x) e_{i_j},$$

and $\operatorname{conv}(e_{i_1}, \ldots, e_{i_m})$ is a simplex in $\operatorname{Ner}(\mathcal{U})$. In particular, $U_{i_j} \cap T_{n+1} \neq \emptyset$ for each j, and this implies that

$$w(i_j) + \sum_{k=1}^n d_k(i_j) \ge d.$$
 (3.8)

Indeed, choosing minimal-length chains from U_{i_j} to T_k for each $1 \le k \le n$, and adding in the set U_{i_j} , gives a collection of sets for which the sum of corresponding weights is at most $d_1(i_j) + \ldots + d_n(i_j) + w(i_j)$, but also which, for each k and l, contains a chain that connects T_k and T_l .

As $\phi(x) \in \operatorname{conv}(e_{i_1}, \ldots, e_{i_m})$, we can write (after possibly re-ordering the indices)

$$\phi(x) \in \operatorname{conv}(p_1, \ldots, p_m),$$

where $p_j = bc(e_{i_1}, \ldots, e_{i_j})$. As ψ is affine on the simplex $conv(p_1, \ldots, p_m)$, it suffices to show that

$$\pi_1(\psi(p_j)) + \ldots + \pi_n(\psi(p_j)) \ge d$$

for each $1 \leq j \leq m$. To this end, fix $1 \leq j \leq m$, and for notational ease let $p = p_j$. The way we chose z_p ensures that

$$\pi_k(\psi(p)) = \pi_k(z_p) \in [d_k(i_1), d_k(i_1) + w(i_1)] \cap \dots \cap [d_k(i_j), d_k(i_j) + w(i_j)],$$

so certainly $\pi_k(z_p) \ge \max\{d_k(i_l) : 1 \le l \le j\}$. Moreover, as U_{i_1} intersects T_{n+1} non-trivially, there must be $k_0 \in \{1, \ldots, n\}$ for which $U_{i_1} \cap T_{k_0} = \emptyset$. For this value of k_0 , our choice of z_p guarantees that $\pi_{k_0}(z_p)$ is the right endpoint of the interval

$$[d_{k_0}(i_1), d_{k_0}(i_1) + w(i_1)] \cap \dots \cap [d_{k_0}(i_j), d_{k_0}(i_j) + w(i_j)],$$

so that $\pi_{k_0}(z_p) = d_{k_0}(i_l) + w(i_l)$ for some $1 \le l \le j$. Thus, we have

$$\sum_{k=1}^{n} \pi_k(z_p) \ge (d_{k_0}(i_l) + w(i_l)) + \sum_{k \ne k_0} \max\{d_k(i_l) : 1 \le t \le j\}$$
$$\ge w(i_l) + \sum_{k=1}^{n} d_k(i_l) \ge d,$$

where the final inequality comes from (3.8).

As in the proof of the unweighted statement, let

$$P = \{(z_1, \dots, z_n) \in \mathbb{R}^n : \text{ all } z_k \ge 0 \text{ and } z_1 + \dots + z_n \le d\}.$$

For $1 \le k \le n$, let P_k denote the codimension-1 face on which the k-th coordinate is 0, and let

$$P_{n+1} = \{(z_1, \dots, z_n) \in \mathbb{R}^n : \text{ all } z_k \ge 0 \text{ and } z_1 + \dots + z_n = d\}$$

be the other codimension-1 face. Let $g: P \to \Delta^n$ be the affine map which sends P_k to T_k for each $1 \le k \le n+1$. Lastly, let $H_k = \{(z_1, \ldots, z_n) \in \mathbb{R}^n : z_k \le 0\}$ for $1 \le k \le n$ and

$$H_{n+1} = \{(z_1, \dots, z_n) \in \mathbb{R}^n : z_1 + \dots + z_n \ge d\},\$$

so that H_1, \ldots, H_{n+1} are the closed half-spaces determined by P_1, \ldots, P_{n+1} that do not contain P.

We know that if $x \in T_k$, with $1 \le k \le n$, then $\pi_k(f(x)) = 0$, and so $f(x) \in H_k$. Lemma 3.18 implies that if $x \in T_{n+1}$, then $f(x) \in H_{n+1}$. Consequently, we can apply Lemma 3.8 to the map $f \circ g$ to conclude that

$$P \subset f(\Delta^n) \subset \bigcup_{i \in I} R_i.$$

In this way, we see that

$$\frac{d^n}{n!} = \operatorname{Vol}(P) \le \sum_{i \in I} \operatorname{Vol}(R_i) = \sum_{i \in I} w(i)^n,$$

as desired.

As was the case for covers of cubes, discrete length-volume inequalities for covers of Δ^n easily translate into similar inequalities for continuous images of Δ^n in metric spaces. Namely, let (X, d) be a metric space and $g: \Delta^n \to X$ be a continuous map. Fix an open cover $\mathcal{U} = \{U_i\}_{i \in I}$ of $g(\Delta^n)$ and a corresponding weight function $w: I \to [0, \infty)$. The diameter of g, with respect to \mathcal{U} and w, is

$$d_w(g) = \inf \left\{ \sum_{j=1}^m w(i_j) : \begin{array}{c} \text{the collection } U_{i_1}, \dots, U_{i_m} \text{ contains a chain that} \\ \text{connects } g(T_k) \text{ and } g(T_l) \text{ for each pair } k, l \end{array} \right\}$$
The following corollary can be deduced from Theorem 3.17 using the same arguments that appeared in the proof of Proposition 3.12.

Corollary 3.19. For such g, U, and w, we have

$$\sum_{i \in I} w(i)^n \ge \frac{d_w(g)^n}{n!}.$$

With this result, we can bound the Hausdorff content of $g(\Delta^n)$, just as we did in Corollary 3.13. More precisely, let

$$d(g) = \inf\{\operatorname{diam} K : K \subset g(\Delta^n) \text{ intersects } g(T_k) \text{ for all } 1 \le k \le n+1\},\$$

where diam K is the *metric* diameter of K. Note that this gives some notion of "thickness" for the map g. The following inequality can be proved using the same arguments found in the proof of Corollary 3.13; once again, we omit them.

Corollary 3.20. If $g: \Delta^n \to X$ is continuous, then $\mathcal{H}^{\infty}_n(g(\Delta^n)) \geq \frac{d(g)^n}{n!}$.

Let us briefly revisit the discussion from the end of the previous section regarding lower volume bounds in metric spaces. We remarked there that it is often easier to build nondegenerate maps $g: \Delta^n \to X$ than to build non-degenerate maps $g: [0, 1]^n \to X$. Motivated by Corollary 3.20, we make the following definition, similar to Definition 3.15 from earlier.

Definition 3.21. A metric space (X, d) is said to *admit fat n-simplices* if there is $\lambda \ge 1$ such that, for each $x \in X$ and each $0 < r \le \text{diam } X$, there is a continuous map $g: \Delta^n \to B(x, r)$ with $d(g) \ge r/\lambda$.

Corollary 3.20 immediately implies that if (X, d) admits fat *n*-simplices, then each metric ball B(x, r), with $0 < r \le \text{diam } X$, has

$$\mathcal{H}_n^\infty(B(x,r)) \gtrsim r^n,$$

where the implicit constant depends only on n and λ . Using the same techniques that were used to show that annularly linearly connected and linearly locally contractible metric spaces admit fat connecting squares, one can verify the following. **Lemma 3.22.** Let (X, d) be a complete metric space that is N-doubling, λ -ALC, and λ linearly locally contractible. Then (X, d) admits fat triangles, with constant depending only on N and λ .

In higher dimensions, the properties of admitting fat cubes or fat simplices is subtle, and we know very little in this regard. A reasonable place to begin is with Semmes spaces: closed n-manifolds with a metric that is linearly locally contractible and Ahlfors n-regular. While these spaces might lack global quasisymmetric parameterizations, they still have strong analytic properties, including Poincaré inequalities. It would be interesting to know whether they admit fat n-simplices, and if so, whether the hypothesis of Ahlfors regularity can be dropped. If the answer to both questions is "yes," our inequalities in this section would recover the lower volume bounds we deduced from Semmes's result, Theorem 3.16.

We finish this section by noting that the linear local contractibility condition guarantees bounds of this form for closed surfaces. Indeed, a λ -linearly locally contractible closed surface is λ' -ALC, by Lemma 2.39(iii) from the previous chapter, where λ' depends only on λ . Thus, Lemma 3.22 and its preceding discussion gives the following fact.

Proposition 3.23. Let (X, d) be a closed surface that is N-doubling and λ -linearly locally contractible. Then

$$\mathcal{H}_2^\infty(B(x,r)) \ge cr^2$$

for each $x \in X$ and each $0 < r \leq \text{diam } Z$, where c depends only on N and λ .

BIBLIOGRAPHY

- G. Besson, G. Courtois, and S. Gallot, *Minimal entropy and Mostow's rigidity theorems*, Ergodic Theory Dynam. Systems 16 (1996), no. 4, 623–649.
- M. Bonk, Quasiconformal geometry of fractals, Proceedings Internat. Congress Math. (Madrid, 2006), Eur. Math. Soc., Zürich, 2006, pp. 1349–1373.
- [3] M. Bonk and T. Foertsch, Asymptotic upper curvature bounds in coarse geometry, Math. Z. 253 (2006), no. 4, 753–785.
- [4] M. Bonk and B. Kleiner, Quasisymmetric parametrizations of two-dimensional metric spheres, Invent. Math. 150 (2002), no. 1, 127–183.
- [5] _____, Rigidity for quasi-Möbius group actions, J. Differential Geom. **61** (2002), no. 1, 81–106.
- [6] _____, Rigidity for quasi-Fuchsian actions on negatively curved spaces, Int. Math. Res. Not. 61 (2004), 3309–3316.
- [7] M. Bonk, B. Kleiner, and S. Merenkov, *Rigidity of Schottky sets*, Amer. J. Math. **131** (2009), no. 2, 409–443.
- [8] M. Bonk and D. Meyer, *Expanding Thurston maps*, preprint (2010), available at arXiv:1009.3647v1.
- [9] M. Bourdon, Structure conforme au bord et flot géodésique d'un CAT(-1)-espace, Enseign. Math. (2) 41 (1995), no. 1–2, 63–102.
- [10] _____, Sur le birapport au bord des CAT(-1)-espaces, Inst. Hautes Études Sci. Publ. Math. 83 (1996), 95–104.
- [11] M. Bourdon and B. Kleiner, Combinatorial modulus, the combinatorial Loewner property, and Coxeter groups, Groups Geom. Dyn. 7 (2013), no. 1, 39–107.
- [12] Y. D. Burago and V. A. Zalgaller, *Geometric inequalities*, Springer-Verlag, Berlin, 1988.
- [13] S. Buyalo and V. Schroeder, *Elements of asymptotic geometry*, EMS Monographs in Mathematics, Eur. Math. Soc., Zürich, 2007.
- [14] J. W. Cannon, The combinatorial Riemann mapping theorem, Acta Math. 173 (1994), no. 2, 155–234.
- [15] _____, The theory of negatively curved spaces and groups, Ergodic theory, symbolic dynamics, and hyperbolic spaces (Trieste, 1989), Oxford Sci. Publ., Oxford Univ. Press, New York, 1991, pp. 315–369.
- [16] J. W. Cannon, W. J. Floyd, and W. R. Parry, Constructing subdivision rules from rational maps, Conform. Geom. Dyn. 11 (2007), 128–136.

- [17] _____, Finite subdivision rules, Conform. Geom. Dyn. 5 (2001), 153–196.
- [18] _____, Sufficiently rich families of planar rings, Ann. Acad. Sci. Fenn. Math. 24 (1999), no. 2, 265–304.
- [19] J. W. Cannon and E. L. Swenson, Recognizing constant curvature discrete groups in dimension 3, Trans. Amer. Math. Soc. 350 (1998), no. 2, 809–849.
- [20] M. Coornaert, Mesures de Patterson-Sullivan sur le bord d'un espace hyperbolique au sens de Gromov, Pacific J. Math. 159 (1993), no. 2, 241–270.
- [21] C. Croke, Rigidity for surfaces of nonpositive curvature, Comment. Math. Helv. 65 (1990), no. 1, 150–169.
- [22] W. R. Derrick, Inequalities for p-modules of curve families on Lipschitz surfaces, Math. Z. 119 (1971), 1–10.
- [23] _____, A volume-diameter inequality for n-cubes, J. Analyse Math. 22 (1969), 1–36.
- [24] _____, A weighted volume-diameter inequality for N-cubes, J. Math. Mech. 18 (1968/1969), 453–472.
- [25] É. Ghys and P. de la Harpe (eds.), Sur les groupes hyperboliques d'après Mikhael Gromov, Progress in Mathematics, vol. 83, Birkhäuser Boston, Boston, MA, 1990.
- [26] M. Gromov, Filling Riemannian manifolds, J. Differential Geom. 18 (1983), no. 1, 1–147.
- [27] _____, Hyperbolic groups, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 75–263.
- [28] _____, Metric structures for Riemannian and non-Riemannian spaces, Progress in Mathematics, vol. 152, Birkhäuser Boston, Boston, MA, 1999.
- [29] P. Haïssinsky and K. M. Pilgrim, Coarse expanding conformal dynamics, Astérisque **325** (2009).
- [30] U. Hamenstädt, Cocycles, symplectic structures and intersection, Geom. Funct. Anal. 9 (1999), no. 1, 90–140.
- [31] _____, Entropy-rigidity of locally symmetric spaces of negative curvature, Ann. of Math. (2) 131 (1990), no. 1, 35–51.
- [32] _____, Time-preserving conjugacies of geodesic flows, Ergodic Theory Dynam. Systems 12 (1992), no. 1, 67–74.
- [33] A. Hatcher, Algebraic topology, Cambridge University Press, Cambridge, 2002.
- [34] Z.-X. He and O. Schramm, On the convergence of circle packings to the Riemann map, Invent. Math. 125 (1996), no. 2, 285–305.

- [35] J. Heinonen, Lectures on analysis on metric spaces, Universitext, Springer-Verlag, New York, 2001.
- [36] _____, Lectures on Lipschitz analysis, Report. University of Jyväskylä Department of Mathematics and Statistics, vol. 100, University of Jyväskylä, Jyväskylä, 2005.
- [37] A. Hinkkanen, The structure of certain quasisymmetric groups, Mem. Amer. Math. Soc. 83 (1990), no. 422, iv+87.
- [38] W. Hurewicz and H. Wallman, *Dimension Theory*, Princeton Mathematical Series, v. 4, Princeton Univ. Press, Princeton, NJ, 1941.
- [39] N. V. Ivanov, Action of Möbius transformations on homeomorphisms: stability and rigidity, Geom. Funct. Anal. 6 (1996), no. 1, 79–119.
- [40] I. Kapovich and N. Benakli, Boundaries of hyperbolic groups, Combinatorial and geometric group theory, Contemp. Math., vol. 296, Amer. Math. Soc., Providence, RI, 2002, pp. 39–93.
- [41] B. Kleiner, The asymptotic geometry of negatively curved spaces: uniformization, geometrization and rigidity, Proceedings Internat. Congress Math. (Madrid, 2006), Eur. Math. Soc., Zürich, 2006, pp. 743– 768.
- [42] J. M. Mackay, Existence of quasi-arcs, Proc. Amer. Math. Soc. 136 (2008), no. 11, 3975–3981.
- [43] _____, Quasi-circles through prescribed points, preprint, to appear in Indiana Univ. Math. J. (2013), available at arXiv:1210.5119v2.
- [44] _____, Spaces and groups with conformal dimension greater than one, Duke Math. J. **153** (2010), no. 2, 211–227.
- [45] A. Manning, Topological entropy for geodesic flows, Ann. of Math. (2) 110 (1979), no. 3, 567–573.
- [46] V. Markovic, Quasisymmetric groups, J. Amer. Math. Soc. 19 (2006), no. 3, 673–715.
- [47] G. D. Mostow, Strong rigidity of locally symmetric spaces, Princeton Univ. Press, Princeton, NJ, 1973.
- [48] J.-P. Otal, Le spectre marqué des longueurs des surfaces à courbure négative, Ann. of Math. (2) 131 (1990), no. 1, 151–162.
- [49] F. Paulin, Un groupe hyperbolique est déterminé par son bord, J. London Math. Soc. (2) 54 (1996), no. 1, 50–74.
- [50] P. M. Pu, Some inequalities in certain nonorientable Riemannian manifolds, Pacific J. Math. 2 (1952), 55–71.
- [51] B. Rodin and D. Sullivan, The convergence of circle packings to the Riemann mapping, J. Differential Geom. 26 (1987), no. 2, 349–360.

- [52] O. Schramm, Square tilings with prescribed combinatorics, Israel J. Math. 84 (1993), no. 1–2, 97–118.
- [53] S. Semmes, Finding curves on general spaces through quantitative topology, with applications to Sobolev and Poincaré inequalities, Selecta Math. (N.S.) 2 (1996), no. 2, 155–295.
- [54] _____, Good metric spaces without good parameterizations, Rev. Mat. Iberoamericana 12 (1996), no. 1, 187–275.
- [55] R. J. Spatzier, An invitation to rigidity theory, Modern dynamical systems and applications, Cambridge Univ. Press, 2004, pp. 211–231.
- [56] K. Stephenson, Circle packings in the approximation of conformal mappings, Bull. Amer. Math. Soc. (N.S.) 23 (1990), no. 2, 407–415.
- [57] D. Sullivan, On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions, Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference, Ann. of Math. Stud., vol. 97, Princeton Univ. Press, Princeton, NJ, 1981, pp. 465–496.
- [58] M. Troyanov, Espaces à courbure négative et groupes hyperboliques, Sur les groupes hyperboliques d'après Mikhael Gromov, Progr. Math., vol. 83, Birkhäuser Boston, 1990, pp. 47–66.
- [59] P. Tukia, On quasiconformal groups, J. Analyse Math. 46 (1986), 318–346.
- [60] K. Wildrick, Quasisymmetric structures on surfaces, Trans. Amer. Math. Soc. 362 (2010), no. 2, 623– 659.
- [61] Q. Yin, Lattès maps and combinatorial expansion, preprint (2011), available at arXiv:1109.2664v2.