UNIVERSITY OF CALIFORNIA
RIVERSIDE

The Curtain Model of CAT(0) Spaces and its Relationship to the Sublinearly Morse Boundary

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Mathematics

by

Elliott Scott Vest

September 2024

Dissertation Committee:

Dr. Matthew Gentry Durham, Chairperson
Dr. Stefano Vidussi
Dr. Peter Samuelson

The Dissertation of Elliott Scott Vest is approved:

Committee Chairperson

University of California, Riverside

Acknowledgments

First and foremost, I'd like the thank Dr. Matthew Gentry Durham for their unwavering honesty, dedication, and kindness towards advising me over the last four years. Without a doubt, this thesis wouldn't have been possible without them, and their guidance helped me survive graduate school.

Among the other faculty at UCR, an especial thanks goes to Margarita Roman, as without her the entire mathematics program would go into disarray. Another thanks goes to Stefano Vidussi and Peter Samuelson. Both helped pave my enjoyment of topology and algebra, and I am happy to have them as part of my defense committee.

I'd also like the thank the many peers and professors that I crossed paths with during my decent into geometric group theory. Their names include Abdul Zalloum, Harry Petyt, Davide Spriano, Alexandra Edletzberger, Vivian He, Arianna Zikos, Jagerynn Verano, and Jake Murphy. I tried to hide it, but I felt incredibly nervous during my visits to conferences or attending talks abroad. These individuals in particular were extremely welcoming and helped combat my imposter syndrome. An especial thanks goes to Abdul and Harry, who in some ways, I view as heroes in my research. Both helped immensely with my understanding towards the curtain model and sublinearly Morse boundaries.

Incoming into graduate school, I was lucky to have the best cohort of friends I could ever ask for. Raymond Matson, Kevin Su, and William Hoffer, my roommates during my stay at UCR, were my greatest source of happiness during the highs and lows of graduate school. We survived Covid together, and I would've went insane without their support during these extremely tough years. Jialin Wang and Zachary Virgilio were also great sources of joy during my more senior years at UCR. Someday, I hope to have as much happiness as Jialin, or as much awesomeness as Zachary, but I'll settle for being glad they were a part of my life. Lastly, I nor anyone else will ever be able to forget Neima Ghandian. Any story worth telling in my five years here has Neima at its center. He's a perfect jar of wild and laughs, and no social event is as good unless Neima is there.

Jacob Garcia and William Sablan were not only my friends, but also extremely helpful academic brothers. Jacob was essentially a secondary advisor for me, as his understanding of Morse boundaries paved my own understanding. Also, I depended on William more than he realized as we travelled to conferences together.

If you know me, you know my family is more important to me than anything else. I am so proud to be equal parts my mom and equal parts my dad, Lisa and Scott Vest. Everything great in my life is because of them. They gave me endless love, time, effort, kindness, empathy, support and all the other terms to describe the best parents ever. Dylan Vest and his wife, Kayla, are also the greatest brother and sister-in-law I could ever know. I'm so lucky to have them in my life, to know that adventure for me awaits one call away. Plus, they also have the cutest baby to ever crawl the face of the earth (who this thesis is dedicated to). With my graduate journey coming to a close, all I want more than anything is to be with these people on a warm day at Bass Lake. I'm coming home now.

Lastly, I could not end these acknowledgments without mentioning Sarah Toomey. Sarah, I love you, and this thesis was so much easier to write with you in my life.

To Gianna Ellie Vest.

The world is yours to discover.

ABSTRACT OF THE DISSERTATION

The Curtain Model of CAT(0) Spaces and its Relationship to the Sublinearly Morse Boundary

by

Elliott Scott Vest

Doctor of Philosophy, Graduate Program in Mathematics
University of California, Riverside, September 2024
Dr. Matthew Gentry Durham, Chairperson

We show that the sublinear Morse boundary of every CAT(0) space continuously injects into the Gromov boundary of a hyperbolic space, which was not previously known even for all CAT(0) cube complexes. Our work utilizes the curtain machinery introduced by Petyt-Spriano-Zalloum. Curtains are more general combinatorial analogues of hyperplanes in cube complexes, and we develop multiple curtain characterizations of the sublinear Morse property along the way. The hyperbolic space mentioned is the *curtain model*, and its role for a CAT(0) space has shown a striking comparison to the curve graph for a mapping class group of a finite type surface. We show that the curtain model is not a quasi-isometry invariant for all CAT(0) spaces and that quasi-flats are of bounded diameter in the curtain model. Our results answer multiple questions of Petyt-Spriano-Zalloum in [PSZ22].

Contents

List of Figures

Chapter 1

Introduction

1.1 Summary of Results

The *sublinearly Morse boundary*, denoted $\partial_\kappa X$, was introduced by Qing-Rafi-Tiozzo in [QR22, QRT20] as an extension of the Morse boundary [Cor17]. It remains a quasi-isometry invariant while also capturing the asymptotic behavior of random walks on finitely generated groups. A motivating application they use is showing that, for each finite-type surface S and some $p \in \mathbb{N}$ depending on S, the $\log^p$-Morse boundary for the mapping class group of S serves as a topological model for the Poisson boundary for its random walks. This boundary has now subsequently been studied in the CAT(0) cube complex setting [MQZ22, IMZ23], and more generally the hierarchically hyperbolic setting [DZ22, NQ22].

When in regards to CAT(0) spaces, authors Petyt-Spriano-Zalloum in [PSZ22] introduce a combinatorial tool named a *curtain* that serves as an analogue to a hyperplane for a CAT(0) cube complex. Building off of "hyperplane-separation" metrics introduced by Genevois [Gen19], the authors utilize curtains in a CAT(0) space X to build the *curtain* *model* - a hyperbolic space whose isometry group contains Isom X. Denoting $\widehat{X}$ as the curtain model and $\partial\widehat{X}$ as its Gromov boundary, we summarize our main results.

Main Results: *Let X be a proper CAT(0) space and κ be a sublinear function. We show the following,*

1. *If κ^4 is sublinear and $\partial_\kappa X$ is endowed with the sublinearly Morse topology, then $\partial_\kappa X$ continuously injects into the Gromov boundary $\partial\widehat{X}$. Endowed with the cone topology, $\partial_\kappa X$ topologically embeds into $\partial\widehat{X}$. (Theorem 1.3.1.)*

2. *Any κ-Morse ray can be characterized by a dual chain of κ-separated curtains crossing the ray at a sublinear rate. (Theorem 1.4.1.)*

3. *Any κ-Morse ray can be characterized by the κ-persistence of its projection into $\widehat{X}$. (Theorem 1.4.2.)*

Theorem 1.3.1 upgrades a similar theorem by Petyt-Spriano-Zalloum in [PSZ22] from the Morse setting. Since sublinearly Morse rays have been shown to be unparametrized quasi-geodesics in the hyperbolic space $\widehat{X}$, a key ingredient we prove is that these rays are also unbounded in $\widehat{X}$ and, thus, define a point in $\partial\widehat{X}$. Moreover, as random walks sublinearly track κ-Morse rays in CAT(0) spaces [Cho23], Theorem 1.3.1 immediately concludes that the Gromov boundary of the curtain model for a CAT(0) group is a model for its Poisson boundary (Corollary 1.3.2). The central characterization used to give Theorem 1.3.1 is Theorem 1.4.1 - a CAT(0) analogue of a result in [MQZ22] in the cube complex setting that allows directions of κ-Morse rays to be described in a combinatorial fashion. Lastly, a version of Theorem 1.4.2 was proven for hierarchically hyperbolic spaces in [DZ22], including mapping class groups of finite-type surfaces. In particular, they proved that κ-Morse rays in the mapping class group make persistent and fast (compared to κ) progress in the curve graph of the whole surface. Hence, Theorem 1.4.2 further strengthens observations that the curtain model for a CAT(0) space gives a similar role of the curve graph for a mapping class group of a finite-type surface [Zal23]. These main results are from the author's paper [Ves23a] and are to appear in the refereed journal *Groups, Geometry, and Dynamics*.

In addition to our main results, Petyt-Spriano-Zalloum also asked questions concerning hyperbolicity criterion and quasi-isometry invariance in their seminal paper on the curtain model. We give answers to some of these questions through the following results.

Supplemental Results:

1. *There exists a CAT(0) space X and a self quasi-isometry $\phi : X \longrightarrow X$ such that ϕ does not descend to a quasi-isometry for $\widehat{X}$. Further, there exists two quasi-isometric CAT(0) spaces W, Z whose curtain models $\widehat{W}, \widehat{Z}$ are not quasi-isometric. (Theorem 1.5.1.)*

2. *Quasi-flats of CAT(0) cube complexes project to bounded diameter sets in the curtain model. (Theorem 1.5.3.)*

3. *Let X be a CAT(0) space. Then X is hyperbolic if and only if every curtain grid is E-thin for some uniform $E > 0$. (Theorem 1.5.4.)*

Though a counter example for quasi-isometry invariance is unfortunate, it is still an open problem whether CAT(0) spaces with cocompact actions have quasi-isometric curtain models. Theorem 1.5.4 is also from [Ves23a] and is to appear in *Groups, Geometry, and Dynamics*. Theorem 1.5.1 can be found in [Ves23b] and is being reviewed by the journal

Proceedings of the American Mathematical Society. Further motivation, details, relevant citations, and the full statement of the theorems will be given in the rest of the introduction.

1.2 Hyperbolic-like Boundaries and Cube Complexes

In recent years, there have been numerous advancements in investigating the boundaries of spaces, either to gain insights into the spaces themselves or to understand groups that act geometrically on these spaces. A foundational result is due to Gromov [Gro87], who demonstrated that quasi-isometries between hyperbolic spaces induce homeomorphisms on their visual boundaries, leading to a well-defined notion of a boundary of a hyperbolic group. However, Croke and Kleiner showed that this result does not hold for CAT(0) spaces, as they discovered two quasi-isometric CAT(0) spaces that lack homeomorphic visual boundaries [CK00]. To address this problem in the CAT(0) setting, Charney and Sultan introduced the contracting boundary, a quasi-isometry invariant restriction of the visual boundary that only looks at the "hyperbolic" directions of a space [CS15]. Subsequently, Cordes extended this result to any proper geodesic space by defining Morse geodesics and the Morse boundary [Cor17]. It is worth noting that the contracting and Morse conditions are equivalent in the CAT(0) setting, and if the underlying space is hyperbolic, then all mentioned boundaries are equivalent.

The *sublinear Morse boundary* (Definition 3.3.6), denoted $\partial_\kappa X$ where κ is a sublinear function, was introduced in [QR22, QRT20] with the motivation to preserve a notion of a hyperbolic-like boundary that is quasi-isometry invariant while also capturing generic directions of the space in question. More specifically, random walks can sublinearly track the sublinearly Morse directions of a group G to serve as a topological model for the Poisson boundary when G is a right angled Artin group [QR22] or a mapping class group of a finite-type surface [QRT20]. This result has also been shown to hold for rank-1 CAT(0) spaces and Teichmüller spaces of finite-type surfaces [GQR22], then later with CAT(0) admissible groups with mild assumptions [NQ22].

Apart from the aforementioned developments, CAT(0) cube complexes have also been of particular interest due to their combinatorial nature [Sag95]. When one *cubulates* a group, i.e. shows that the group acts geometrically on a CAT(0) cube complex, one can import the various combinatorial information of the cube complex to the group. Notable applications of this strategy are shown in the resolutions of the virtual Haken and fibering conjectures in 3-manifold theory [Ago13, Wis21]. There are many examples of interesting groups that can be cubulated such as right-angled Artin groups [CD95], Coxeter groups [NR03], small cancellation groups [Wis04], hyperbolic 3-manifold groups [KM12], and others. Furthermore, CAT(0) cube complexes have become a useful tool for studying the geometry of mapping class groups of finite-type surfaces, and hierarchically hyperbolic spaces more generally [BHS21, DMS20, HHP20, DZ22].

1.3 The Curtain Model and a Continuous Injection

Recently, Petyt-Spriano-Zalloum introduced an analogue for a hyperplane in any CAT(0) space [PSZ22]. Hyperplanes are the basic combinatorial objects in a cube complex (Definition 2.3.4), and their general analogues, *curtains*, have a simple definition: each is the preimage of a unit length interval of a geodesic under closest point projection (Definition 4.1.1).

Like hyperplanes, curtains separate the ambient space into two components, and Petyt-Spriano-Zalloum prove that an analogous notion of separation allows one to build a hyperbolic space, called *the curtain model* (Definition 6.1.1), which encodes much of the ambient hyperbolic geometry of the space. Such a space serves as an analogue of curve graphs for mapping class groups [Zal23], as the curve graph of a finite-type surface also encodes the hyperbolic geometry of its associated mapping class group [MM99]. In addition, ongoing work of Le Bars shows the utility curtains can have when investigating asymptotic behavior of random walks in CAT(0) spaces [Bar22a, Bar22b]. This segues to our main result, which deals in the extension of the projection of a CAT(0) space $X \longrightarrow \widehat{X}$ to a continuous injection of its sublinearly Morse boundary.

Theorem 1.3.1. *Let X be a proper CAT(0) space, $\widehat{X}$ its curtain model, and κ be a sublinear function such that κ^4 is sublinear. Endow $\partial_\kappa X$ with the sublinear Morse topology and denote $\partial\widehat{X}$ as the Gromov boundary of $\widehat{X}$. Then the projection map $X \to \widehat{X}$ extends to an Isom X-equivariant continuous injection $\varphi : \partial_\kappa X \hookrightarrow \partial\widehat{X}$. Moreover, endowing $\partial_\kappa X$ with the subspace topology of the cone topology makes φ a homeomorphism onto its image.*

In other words, the curtain model of a CAT(0) space will capture all the generic directions of the CAT(0) space. We note that Theorem 1.3.1 relies on the condition that κ also have powers that are sublinear. Given that sublinear Morse boundaries are commonly employed in applications where random walks converge to $\partial_\kappa X$ for $\kappa(t) = \log^p(t)$ [NQ22, QRT20], the cost of our restriction on κ is justifiable for cleaner arguments. In addition, the assumption that the CAT(0) space be proper is only used for the characterization that κ-Morse rays are κ-contracting [QR22]. If one worked with the sublinearly contracting boundary instead of $\partial_\kappa X$, our work could yield the same result without the assumption that the CAT(0) space be proper.

In the recent paper [CFFT22], the authors show if a group G acts on a hyperbolic space $\widehat{X}$ with a WPD element and μ is a probability measure with finite entropy, then $\partial\widehat{X}$ with the hitting measure is a model for the Poisson boundary of (G, μ). In the context of CAT(0) spaces, work of Choi in [Cho23, Cho22] has shown that random walks with finite pth moment will sublinearly track $o(n^{1/p})$-Morse geodesics. Moreover, [QRT20] shows that whenever random walks sublinearly track κ-Morse rays, $\partial_\kappa X$ with its hitting measure is a model for the Poisson boundary of (G, μ). Thus, the injection created in Theorem 1.3.1 immediately gives the following corollary. This corollary recovers the aforementioned result in [CFFT22], but with a stronger assumption on the probability measure μ.

Corollary 1.3.2. *Let G be a group that acts properly on a proper CAT(0) space X, denote $\widehat{X}$ as its curtain model, and let μ be a non-elementary probability measure on G with finite 5th moment. Let ν be the hitting measure on the Gromov boundary $\partial\widehat{X}$. Then $(\varphi(\partial\widehat{X}), \nu)$ is a model for the Poisson boundary of (G, μ)*

Endowing $\partial_\kappa X$ with the sublinearly Morse topology is nice because it makes $\partial_\kappa X$ metrizable and a quasi-isometry invariant [QR22]. However, our proof of continuity only uses open sets in the *cone topology* which is strictly coarser than the sublinearly Morse topology [IMZ23]. We use techniques of [AIM22] to show that φ can be a homeomorphism on its image when $\partial_\kappa X$ is endowed with the cone topology (this is Theorem 6.3.8 in our paper).

Theorem 1.3.1 has been shown in the CAT(0) cube complex setting [IMZ23], where the authors project to a hyperbolic space inspired by Genevois's work [Gen19]. However, their proof relied on the assumption that the cube complex admits a *factor system* - namely that the cube complex be a hierarchically hyperbolic space [BHS17]. Recently, Shepherd found an example of a cocompact CAT(0) cube complex that does not admit a factor system [She22]. As Theorem 1.3.1 applies to the CAT(0) setting, it in particular applies to any CAT(0) cube complex - including those not covered in prior literature. Similar theorems like Theorem 1.3.1 have been shown for sublinear Morse boundaries of hierarchically hyperbolic spaces [DZ22] and the Morse boundary of CAT(0) spaces [PSZ22], both of such results requiring a leverage of cubical techniques. Since He in [He23] has shown the Morse boundary with the Cashen-Mackay topology to be homeomorphic to $\partial_\kappa X$ for $\kappa \equiv 1$, Theorem 1.3.1 recovers the Morse boundary result in [PSZ22] when in regards to the Cashen-Mackay topology. Additional related work of Theorem 1.3.1 can be found in [AIM22], where Abbott and Incerti-Medici study classes of spaces that have *hyperbolic projections* and are κ*-injective*. Hence, Theorem 1.3.1 shows that rank-one CAT(0) spaces and groups fall into the class of objects that Abbott and Incerti-Medici build a framework for studying.

1.4 Characterizations of κ-Morse Rays and Hyperbolicity

The key arguments needed to prove our main result involve characterizing κ-Morse geodesics in the same combinatorial fashion as done in the CAT(0) cube complex setting [CS15, MQZ22].

Theorem 1.4.1. *In a CAT(0) space, a geodesic ray b is κ-contracting if and only if there exists $C > 0$ such that b meets a chain of curtains $\{h_i\}$ at points $b(t_i) \in h_i$ satisfying:*

- $t_{i+1} - t_i \leq C\kappa(t_{i+1})$
- *h_i and h_{i+1} are $C\kappa(t_{i+1})$-separated*

A geodesic ray with a dual chain of curtains satisfying the above conditions will be defined as a *κ-curtain excursion geodesic* (Definition 5.2.1) in reference to similarly defined rays in the CAT(0) cube complex setting [MQZ22]. Thus, since the proof of the forward direction finds a dual chain (Proposition 5.2.5), Theorem 1.4.1 states that κ-contracting rays are κ-curtain excursion rays and vice versa. One doesn't need the curtains to be dual to the geodesic to prove the reverse direction, as Proposition 5.3.2 shows, but a dual chain is usually preferred for more straightforward arguments. In fact, Theorem 1.4.1 indirectly shows that a geodesic meeting a chain as above will also give a dual chain that meets the geodesic with the same properties as well (see Corollary 5.3.3). Theorem 1.4.1 plays a critical role in the proof of Theorem 1.3.1 similarly to how the CAT(0) cube complex version of Theorem 1.4.1 in [MQZ22] is used to prove main results in [IMZ23, DZ22, AIM22]. Now such a characterization can be applied to all CAT(0) groups.

Other characterizations of κ-contracting rays were shown in [MQZ22] in the CAT(0) cube complex setting, one of which was later defined as a *κ-persistent shadow* [DZ22]. A geodesic ray b in a CAT(0) space X with infinite diameter in the curtain model has a κ-persistent shadow if there exists a $C > 0$ such that for all $s < t$,

$$\hat{d}(b(s), b(t)) \geq C \cdot \frac{t-s}{\kappa(t)} - C.$$

where $\hat{d}$ is the distance in the curtain model $\widehat{X}$. We also give a similar characterization in the CAT(0) setting.

Theorem 1.4.2. *Let b be geodesic ray emanating from $\mathfrak{o} \in X$ with infinite diameter in the curtain model $\widehat{X}$.*

- *If b is κ-contracting and κ^4 is sublinear, then b has a κ^4-persistent shadow in the $\hat{d}$ metric.*
- *If b has a κ persistent shadow in the $\hat{d}$ metric and κ^2 is sublinear, then b is κ^2-contracting.*

In other words, κ-Morse rays make persistent and fast progress in the curtain model (relative to κ). Note that when $\kappa \equiv 1$, Theorem 1.4.2 is equivalent to a geodesic being contracting if and only if the geodesic projects to a *parametrized* quasi-geodesic in $\widehat{X}$. This is a known result in the mapping class group setting when projecting to its curve graph [Beh06, DR09], and [DZ22] extend the characterization to the sublinear Morse setting in hierarchically hyperbolic spaces. Thus, Theorem 1.4.2 shows yet another comparison that the curve graph is to the mapping class group as the curtain model is to its CAT(0) space (see Section 8 of [Zal23]).

1.5 Quasi-Isometry Invariance and Hyperbolicity Criterion

Petyt-Spriano-Zalloum asked in [PSZ22] if a quasi-isometry between CAT(0) spaces always induces a quasi-isometry between their corresponding curtain models. We answer this question in the negative. For a CAT(0) space X, we denote $\widehat{X}$ to be its curtain model.

Theorem 1.5.1. *There exists a CAT(0) space X and a self quasi-isometry $\phi : X \longrightarrow X$ such that ϕ does not descend to a quasi-isometry for $\widehat{X}$. Further, there exists two quasi-isometric CAT(0) spaces W, Z whose curtain models $\widehat{W}, \widehat{Z}$ are not quasi-isometric.*

Our counterexample is based on a counterexample due to Cashen [Cas16], which he used to show that quasi-isometries of CAT(0) spaces need not induce homeomorphisms of their contracting boundaries when equipped with the Gromov product topology. Thus, it also follows that we get an analogous result for the curtain models of CAT(0) spaces.

Corollary 1.5.2. *There exist quasi-isometric CAT(0) spaces W, Z whose curtain models have non-homeomorphic Gromov boundaries.*

The counterexample for both Theorem 1.5.1 and Corollary 1.5.2 involve the gluing of quarter-flats along their axes. In this counterexample, the curtain model crunches the flat parts of the glued space resulting into a quasi-line. Berhstock-Hagen-Sisto show in [BHS21] that a similar flavor happens to quasi-flats in hierarchically hyperbolic spaces. In particular, they show that quasi-flats project to finite diameter sets in the top level curve graph of the hierarchically hyperbolic space. We show an analogous result for a CAT(0) cubes complex and its curtain model.

Theorem 1.5.3. *Let X be a CAT(0) cube complex and $Q \subset X$ be an n-dimensional quasi-flat. Then Q has bounded projection in $\widehat{X}$.*

Lastly, Genevois has shown a hyperbolicity criterion for CAT(0) cubes complexes through grids of hyperplanes. As an application of Theorem 1.4.1 in the Morse setting, we create a curtain version for this criterion. Two collections of chains $\mathcal{H}, \mathcal{K}$ are said to be a *curtain grid* if all curtains in $\mathcal{H}$ cross all curtains in $\mathcal{K}$. We prove the following.

Theorem 1.5.4. *Let X be a CAT(0) space. Then X is hyperbolic if and only if every curtain grid is E-thin for some uniform $E > 0$.*

A curtain grid $(\mathcal{H}, \mathcal{K})$ is *E-thin* if $\min\{|\mathcal{H}|, |\mathcal{K}|\} \leq E$. So, similar to the cube complex setting, the intuition behind Theorem 1.5.4 is that large regions of crossing curtain chains imply large regions of flatness in the space. Thus, when one bounds the thickness of all curtain grids in the space, one obtains hyperbolicity.

1.6 Outline of Thesis

Chapter 2 gives a small introduction to geometric group theory and serves as motivation for why the problems answered in this thesis are relevant. This chapter is meant for those with no familiarity with the field, and experienced readers are encouraged to skip. Chapter 3 gives a historical journey of relevant boundaries of metric spaces in the context of this thesis. It starts with defining the visual boundary of a geodesic metric space, then the

Morse and sublinearly Morse boundaries. Both Chapter 2 and Chapter 3 are introductory in nature, so they are written with the audience of an inexperienced reader in mind.

Chapter 4 gives the relevant definitions for the curtain machinery introduced in [PSZ22]. Then, we give the counterexample to prove Theorem 1.5.1 (Theorem 4.2.3). The counterexample given involves an augmentation of the universal cover for $\mathbb{R}^2$ minus a disk. This segues to the proof of Theorem 1.5.3. This is done in two parts. First, we show orthants are of bounded diameter in the curtain model. Next, we use Huang's main theorem in [Hua17] that tells us quasi-flats are of bounded Hausdorff distance from some union of orthants. We conclude the chapter by proving Theorem 1.5.4 (Theorem 4.4.2). The forward and converse assumptions both give that curtains dual to the same geodesic at a sufficiently far distance from each other must be L-separated for some uniform L. This connection leads to the equivalence in our statement.

Chapter 5 proves Theorem 1.4.1, where the forward and backward directions are Propositions 5.2.5 and 5.3.2, respectively. Our proofs required synthesizing the curtain machinery along with techniques of [PSZ22, MQZ22] to give a more generalized characterization in the CAT(0) setting. Prior techniques for similar characterizations of Theorem 1.4.1 would give a weaker converse direction in our context, so we use different techniques than prior literature to prove rays *meeting* a κ-chain of curtains will be κ-contracting.

Chapter 6 shows the continuous injection given in Theorem 1.3.1 with Theorem 6.3.5 stating continuity with respect to the sublinearly Morse topology and Theorem 6.3.8 giving a homeomorphic notion when $\partial_\kappa X$ is endowed with the cone topology. A major contribution of our work is showing κ-contracting rays are unbounded in the curtain model

(Lemma 6.1.9), and it's why we need the assumption that κ^4 is also a sublinear function (as Example 6.1.4 shows). This gives that the identity map $X \longrightarrow \widehat{X}$ can extend to a well defined map $\varphi : \partial_\kappa X \longrightarrow \partial \widehat{X}$. Our arguments for injectivity rely heavily on the characterization stated in Theorem 1.4.1 in order to give proofs of a combinatorial nature. Chapter 6 also proves Theorem 1.4.2 (Theorem 6.4.3). The work in showing κ-contracting rays are unbounded in the curtain model also shows the forward direction of Theorem 1.4.2, and its argument has a similar flavor to the analogous arguments made in [MQZ22] for the cube complex setting. The reverse direction of Theorem 1.4.2 is entirely original.

Chapter 2

Preliminaries

The following will be a summary in the background information introduced at the start of a geometric group theorist's career. The goal of this chapter is for an aspiring math student to be able to get a good intuition for what a geometric group theorist cares about as well as the classic tools one uses in this field. Anyone with knowledge of the field already should feel free to skip or reference when needed.

2.1 Motivation in Geometric Group Theory

The flavor of many fields of math, in general, involves thinking about a complicated task by recontextualizing the problem in a different setting. For example, an algebraic topologist has a primary interest of studying topological spaces. However, if one associated the topological space to an algebraic object (such as its fundamental group, homology groups, cohomology groups, etc.), one can learn about the topological space through the lens of algebra. An algebraic topologist might first and foremost consider themselves a topologist, but they specifically use algebraic objects to learn about the topology (hence, the name).

For a geometric group theorist, the objects of interest are groups (in particular, finitely generated groups), which can be tough to study in their own right. If one associates the group to a metric space, we could possibly learn about the group by how we understand the geometry of the metric space. The term "associates" is ill defined right now, but most classical groups already have a natural metric space that even undergraduates readily associate with.

Take the group of integers $\mathbb{Z}$. If one thought of all integers as symbols in a bowl of soup, it would be pretty tough to sift how certain numbers are related to each other. However, we already have a canonical ordering of the integers that can be seen by drawing the *real line*. Just by looking at the picture of the real line, one can already, say, compare distances between integers. This natural association can actually always be found for any group we can think of. In particular, one can associate any group to its Cayley Graph in order to see the group in the light of a metric space.

Definition 2.1.1. Let G be a finitely generated group, with generating set S. (We assume $1 \notin S$ and $S^{-1} = S$.) Then the *Cayley Graph* of G with respect to S, denoted $Cay(G, S)$, is a graph where the vertices are elements of G, and two vertices $g, h \in G$ are connected by an edge if $g^{-1}h \in S$. The metric on $Cay(G, S)$ is the path length metric, where every edge has length 1.

Now, for the example $G = \mathbb{Z}$, choosing $S = \{\pm 1\}$ will give the a graph that is exactly a line. We give other classic examples below.

Example 2.1.2. Consider the free group on two generators $\mathbb{F}_2 = \langle a, b \rangle$ Choosing $S = \{a, b, a^{-1}, b^{-1}\}$ gives the fractal tree displayed in Figure 2.1. Indeed, if one starts with the identity element, e, in the middle of the figure, the surrounding vertices must be the generating set by definition of a Cayley Graph. Thus, we can view the generator a as representing movement in the horizontal direction while the generator b represents movement in the vertical direction. So the element a^2b would be the vertex two notches to the right and one notch up from the center, e, in Figure 2.1. Compare this Cayley Graph with the Cayley Graph of $\mathbb{Z}^2 = \langle a, b : [a, b] = 1 \rangle$. Again, choosing the generators $S = \{a, b, a^{-1}, b^{-1}\}$, the Cayley Graph of $\mathbb{Z}^2$ becomes the grid lattice of $\mathbb{R}^2$. See Figure 2.1.

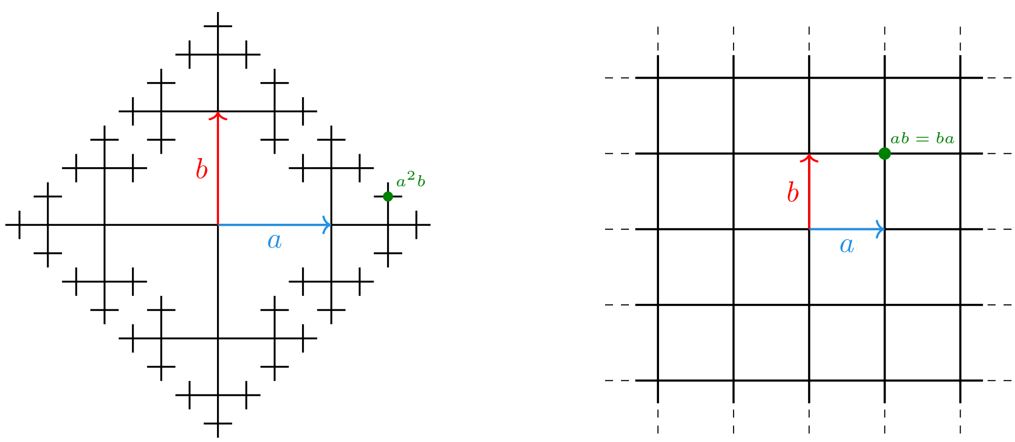

Figure 2.1: The Cayley graph of $\mathbb{F}_2$ and the Cayley graph of $\mathbb{Z}^2$ with respect to the natural generating set $S = \{a, b, a^{-1}, b^{-1}\}$.

Notice that, since $ab = ba$ in $\mathbb{Z}^2$ (unlike in $\mathbb{F}_2$), the ending vertex of the path ab is the same as the ending vertex as the path ba in $Cay(\mathbb{Z}^2, S)$. This is why $Cay(\mathbb{Z}^2, S)$ creates a grid. On the other hand, $\mathbb{F}_2$ has no relations, so $Cay(\mathbb{F}^2, S)$ must not have any cycles.

This means that $Cay(\mathbb{F}^2, S)$ must remain a tree. It is important to note that the notion of commuting generators creates a grid lattice of $\mathbb{R}^2$ in a Cayley Graph.

Example 2.1.3. As the number of generators increase, the complexity of the corresponding metric space will also get more complicated. Take the group $G = \mathbb{Z}^2 * \mathbb{Z} = \langle a, b, c : [a, b] = 1 \rangle$. Since the group has a copy of of $\mathbb{Z}^2$ in it, one can expect, when using the generating set $S = \{a, b, c, a^{-1}, b^{-1}, c^{-1}\}$, $Cay(G, S)$ will contain a grid lattice of $\mathbb{R}^2$. However, the generator c tells us that on each vertex in our grid lattice, we must attach a c edge and c^{-1} edge. Furthermore the endpoints of these edges must also connect to a different grid lattice. The resulting space $Cay(G, S)$ is depicted in Figure 2.2. The group, in a way, looks like a combination of the two prior examples given in Example 2.1.2. Thus, it might make sense that the corresponding Cayley Graph might have both the fractal-like nature of $\mathbb{F}_2$ combined with the Euclidean-like nature of $\mathbb{Z}^2$.

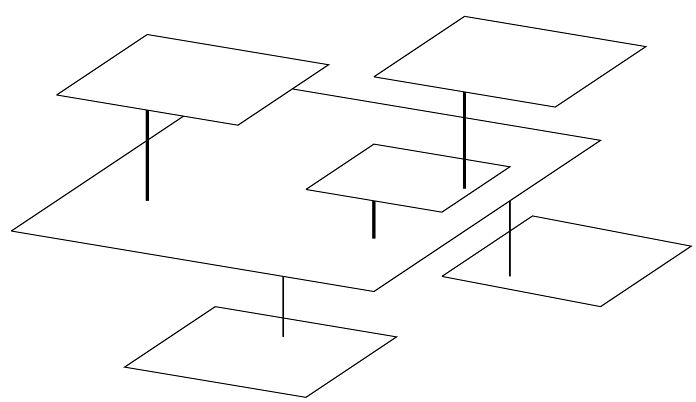

Figure 2.2: The Cayley graph of $\mathbb{Z}^2 * \mathbb{Z}$ with respect to the generating set $S = \{a, b, c, a^{-1}, b^{-1}, c^{-1}\}$. Each c edge will have a grid lattice attached to both of their vertices. Similarly, each vertex on a grid lattice will have a c edge and a c^{-1} edge emanating from it. This space is commonly known as the *tree of flats*.

Notice that the Cayley Graph of a group is dependant on the generating set we choose. One can see in Figure 2.3 that the choice of generating set can make the resulting Cayley Graph of the group look different. However, the difference is not as substantial as one would think. Though the two Cayley graphs $Cay(\mathbb{Z}, \pm 1)$(with black edges) and $Cay(\mathbb{Z}, \pm 2, \pm 3)$(with red and blue edges) in Figure 2.3 look different, if one were to zoom out farther and farther away, the weaving of the red and blue edges become minute compared to the "large scale" geometry of the space. The following definition rigorously interprets this idea.

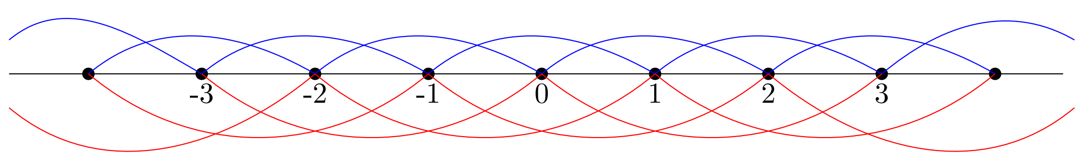

Figure 2.3: Two different Cayley graphs of $\mathbb{Z}$: $Cay(\mathbb{Z}, \pm 1)$ (in black edges) and $Cay(\mathbb{Z}, \pm 2, \pm 3)$ (in red and blue edges). Figure by Jacob Garcia.

Definition 2.1.4 (Quasi-isometric embedding, Quasi-isometry)**.** Let (X, d_X) and (Y, d_Y) be metric spaces. A function $f : X \to Y$ is called a (K, C)-*quasi-isometric embedding* if, for every pair of points $x, x' \in X$ we have

$$\frac{1}{K} d_X(x, x') - C \leq d_Y(f(x), f(x')) \leq K d_X(x, x') + C.$$

If for all $y \in Y$, $d_Y(f(X), y) \leq C$, we call f a *quasi-isometry*.

We see that when $(K, C) = (1, 0)$, Definition 2.1.4 defines an exact isometry, so the role of (K, C) determines how close to a isometry the function f is. A quasi-isometry is often the map that geometric group theorists care about. This, among many other reasons, is due to the following fact (which is proven in Lemma 2.1.7).

Fact 2.1.5. Given any group G, and any two generating sets S and S'. We have $Cay(G,S)$ is quasi-isometric to $Cay(G,S')$.

Thus, up to quasi-isometry, all Cayley Graphs of a group are the same. This fact might not be obvious to someone outside the field, but we can even do better, as the following definitions show.

Definition 2.1.6 (Properly Discontinuous, Cocompact, Geometric Action)**.** Let G be a finitely generated group, and X be a geodesic metric space such that G acts on X by isometries. We say the action is *properly discontinuous* if for any compact set $K \subset X$, the set $\{g \in G : gK \cap K \neq \emptyset\}$ is finite. We say the action is *cocompact* if for any basepoint $x \in X$, there is an $R > 0$ such that the closed R-neighborhood of the G-orbit of x is all of X. We say that the action of G on X is *geometric* if the action is both properly discontinuous and cocompact.

One can think of a properly discontinuous action as making sure for any $x \in X$, the orbit $G \cdot x$ doesn't "clump up" around x (in other words, x has finite stabilizers). On the other hand, an action being cocompact ensures that any orbit $G \cdot x$ sufficiently "spreads out" across the entire space. Together, these nice properties lead to the following foundational lemma known as the Milnor-Švarc Lemma.

Lemma 2.1.7 (Milnor-Švarc Lemma)**.** *Let G act by isometries on a geodesic metric space X such that the action is geometric. Then the group is finitely generated and for every* *generating set S and any $p \in X$, the orbit map*

$$f_p : (G, d_S) \longrightarrow X$$

$$g \longmapsto g \cdot p$$

where d_S is the word metric to G corresponding to S, is a quasi-isometry.

In particular, this lemma proves Fact 2.1.5 as the action of G on any of its Cayley Graphs via left translation is geometric. Lemma 2.1.7 shows us that quasi-isometry invariant properties of metric spaces are a valuable commodity of groups. For example, as Theorem 2.2.4 shows, hyperbolicity is a quasi-isometry invariant property. Thus, if a group G acts geometrically on *some* hyperbolic space, then *all* geometric actions of G are on a hyperbolic space. There are also various kinds of boundaries of metric spaces that are quasi-isometry invariant depending on the context (as Chapter 3 shows), and this can serve as a way to characterize differences between groups.

The theme we try to leverage in Lemma 2.1.7 is as follows: If X is among a class of spaces that has nice properties and G acts geometrically on X, there will likely be nice properties of G that correspond to the nice properties of X. From Example 2.1.2, we see $\mathbb{Z}^2$ acts geometrically on $\mathbb{R}^2$, which gives us some intuition commuting generators of a group equate to quasi-flat in the metric space. We can similarly see this observation in Figure 2.2,commonly known as the the *tree of flats*. The following instructive example shows an application of this observation, which eventually leads us to an extremely useful tool in geometric group theory.

Example 2.1.8. Let Σ be a genus two surface. Consider the mapping class group of Σ, $\mathrm{Mod}(\Sigma)$. That is, $\mathrm{Mod}(\Sigma)$ is the group of orientation preserving homeomorphisms up to isotopy. For those unfamiliar with mapping class groups, [FM11] is a great reference. However, for this example we can take as a fact that $\mathrm{Mod}(\Sigma)$ is generated by *Dehn twists,* that is, homeomorphisms by twisting a full revolution along a simple closed curve. See Figure 2.4

If one takes two simple closed curves in Σ that are disjoint, such as the red and blue curves in Figure 2.4, we see that neither Dehn twist affects the other. That is, doing a Dehn twist homeomorphism along the red curve and then a Dehn twist homeomorphism along the blue curve is equivalent to first Dehn twisting along the blue curve and then Dehn twisting along the red curve. In other words, if $a, b \in \mathrm{Mod}(\Sigma)$ are Dehn twists around the red and blue curves, then a and b commute. As seen in Example 2.1.2 and 2.1.3, commuting generators equate to a quasi-isometric copy of a Euclidean plane in the Cayley Graph. So, even though $\mathrm{Mod}(\Sigma)$ is a rather complicated group, this observation already tells us that certain group elements mimic some sort of zero curvature, or "flatness" behavior. If one wanted to, say, cone off all zero curvature areas in a mapping class group, the resulting space would be the *curve graph* — whose discovery of hyperbolicity is one of the greatest recent tools developed in geometric group theory [MM99, BHS17].

We have now shown through numerous examples how one associates a group to a metric space. Likely, if the metric space satisfies nice properties, then the group will also satisfy nice properties. In general, the metric spaces of interest for this thesis are spaces

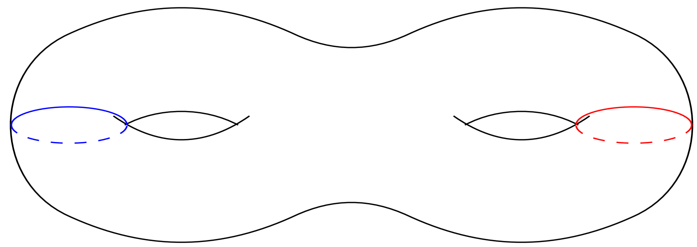

Figure 2.4: A genus two surface with two disjoint simple closed curves. One can create a homeomorphism, called a Dehn twist, by cutting along the blue curve, performing a full revolution along the incision, and then gluing back together the incision. Similarly for the red curve. Notice that these two homeomorphisms do not interact with each other, so the homeomorphisms commute.

of a certain kind of *curvature*, and the following section gives the reasoning for why these spaces are interesting.

2.2 Curvature

In the early 1900s, Max Dehn proved many results regarding groups that have actions on the hyperbolic plane. For example, for any genus g surface with $g \geq 2$, its fundamental group has a natural action on the hyperbolic plane through the fact that the hyperbolic plane is the universal cover of a genus surface with $g \geq 2$. A notable result Dehn showed is that, by utilizing this action on the hyperbolic plane, these groups can be proven to have solvable word and conjugacy problems in linear time [Deh12]. It was much later when Gromov made the observation in [Gro87] that most of the arguments Dehn made were due to the hyperbolic plane having *thin triangles*. This inspired the following definition of a δ-hyperbolic space.

Definition 2.2.1. A proper geodesic metric space X is said to be *δ-hyperbolic* for some $\delta \geq 0$ if, for every three points $x, y, z \in X$, we have that $[x, z] \subseteq Nbhd([x, y] \cup [y, z], \delta)$. We call such a triangle *δ-thin*.

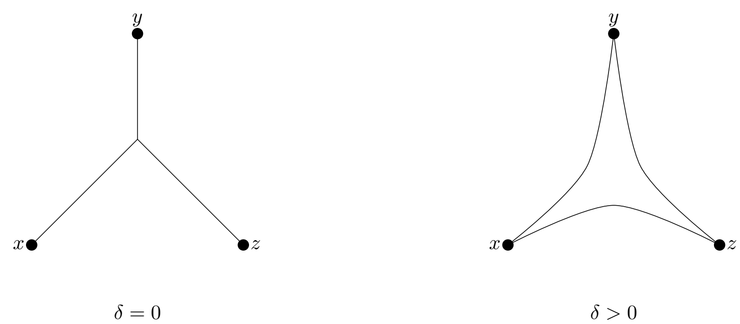

Figure 2.5: Two δ-thin triangles, one with $\delta = 0$ and the other with $\delta > 0$.

Note, this is often the most common definition, for equivalent definitions see [CM17, Chapter 9]. There are many reasons to like hyperbolic spaces and groups (for a survey showing some reasons not discussed here, see [BK02]), but the main reason it is particularly useful for this thesis is due to the following definition and lemma.

Definition 2.2.2 (Geodesics, Quasi-geodesics). A *geodesic ray* in a space X is an isometric embedding $b : [0, \infty) \to X$. Given two points $a, b \in X$, we often denote a geodesic between a and b as $[a, b]$. Similarly, a *quasi-geodesic ray* is a quasi-isometric embedding $\beta : [0, \infty) \to X$.

Lemma 2.2.3 (Morse lemma). *Let X be a δ-hyperbolic space. For every $K, C \geq 0$ there exists $N = N(K, C) \geq 0$ such that, for every geodesic γ in X and every (K, C)-quasi-geodesic ϕ with endpoints on $[a, b]$, we have $\phi \subset Nbhd(\gamma, N)$. Similarly, $\gamma \subset Nbhd(\phi, N)$.*

The Morse Lemma essentially states that quasi-geodesics always fellow travel geodesics and vice versa. When formulating proofs in metric spaces, it is often natural for a

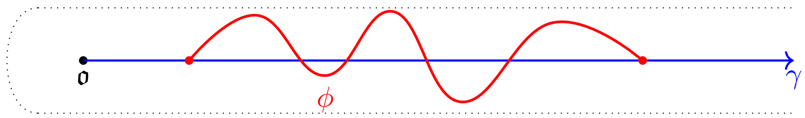

Figure 2.6: The Morse Lemma: If X is a hyperbolic space, then quasi-geodesics with endpoints on geodesics must stay within the N-neighborhood of the geodesic.

geodesic to map to a quasi-geodesic if the mapping we are working with is a quasi-isometry. Because, of this, it is often that geometric group theorists work with quasi-geodesics instead of geodesics at some point in proofs. However, the Morse lemma states that quasi-geodesics travel "close enough" to true geodesics in hyperbolic spaces, which allows us to find sufficient information regarding distance. Theorem 2.2.4 is an instructive example.

Theorem 2.2.4. *Let X be a geodesic metric space, Y be a δ-hyperbolic metric space, and $\phi : X \longrightarrow Y$ be a (K, C)-quasi-isometry. Then X is a δ'-hyperbolic space, where δ' depends on $\delta, K,$ and C.*

Proof. Since ϕ is a quasi-isometry, there exists an (K, C)-quasi-isometry $\phi^{-1} : Y \longrightarrow X$ such that $d(x, \phi^{-1} \circ \phi(x)) \leq C$ for all $x \in X$. Let $\triangle abc$ be a geodesic triangle in X. Put $\overline{a} = \phi(a), \overline{b} = \phi(b)$, and $\overline{c} = \phi(c)$ Then $\phi(\triangle abc)$ comprises of three (K, C)-quasi-geodesic paths: one from $\overline{a}$ to $\overline{b}$, one from $\overline{b}$ to $\overline{c}$, and one from $\overline{c}$ to $\overline{a}$. By the Morse lemma, each of these paths is within $N = N(K, C)$ distance of a corresponding geodesic in the triangle $\triangle\overline{abc}$ in Y. Since geodesic triangles are δ-slim in Y, any point on one geodesic in $\triangle\overline{abc}$ is within δ distance of the other two geodesics in $\triangle\overline{abc}$.

Now consider some point $p \in [a, b]$. Note, by the construction in the previous paragraph, $d_Y(\phi(p), [\overline{a}, \overline{b}] \leq N$ and $[\overline{a}, \overline{b}] \subset Nbhd([\overline{b}, \overline{c}] \cup [\overline{a}, \overline{c}], \delta)$. Since $[\overline{b}, \overline{c}] \cup [\overline{a}, \overline{c}]$ is in the N-neighborhood of $\phi([b, c] \cup [a, c])$, we get that there exists a $q \in [b, c] \cup [a, c]$ such that

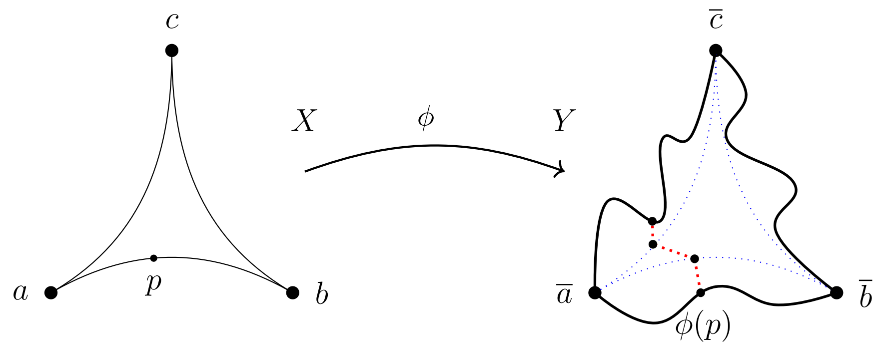

Figure 2.7: A picture for the proof of Theorem 2.2.4. The red path shows that the distance from $\phi(p)$ to $[\overline{a},\overline{c}] \cup [\overline{b},\overline{c}]$ is bounded by $2N+\delta$.

$d_Y(\phi(p),\phi(q)) \le 2N+\delta$. Hence, by the inequalities of a quasi isometry,

$$d_X(p,q) \le K d_Y(\phi(p),\phi(q)) + C \le K(2N+\delta)+C.$$

Setting $\delta' = K(2N+\delta)+C$ completes the proof. See Figure 2.7. □

This classical result captures the use of the Morse Lemma quite nicely. In order to leverage the hyperbolicity of Y, the proof of Theorem 2.2.4 required mapping geodesics to Y via a quasi-isometry. Since these geodesics get mapped to quasi-geodesics in Y, the Morse Lemma ensures that we have control over these quasi-geodesics.

Of course, the Morse Lemma is only a lemma that applies in δ-hyperbolic spaces — spaces of *negative curvature* — and Chapter 3 works in extending this content to investigate boundaries of proper geodesic metric spaces. Either way, this shows that hyperbolic spaces are a quasi-isometry invariant, so hyperbolic spaces are spaces of interest in geometric group theory due to Lemma 2.1.7. If one were to extend spaces of interest outside of negatively curved spaces, a natural step forward would be to investigate spaces of *non-positive curvature*. This is precisely what the following definition formulates.

Definition 2.2.5 (CAT(0) Space)**.** A geodesic metric space (X, d_X) is said to be *CAT(0)* if geodesic triangles in X are at least as thin as corresponding representative triangles in Euclidean space. More precisely, given any triangle $\triangle xyz$, one can create a representative triangle in the Euclidean plane $\triangle \overline{xyz}$ where matching sides of both triangles have the same lengths. See Figure 2.8. If one picks any point p on $\triangle xyz$, say p is on edge $[y, z]$, there exists a corresponding point $\overline{p}$ on $\triangle \overline{xyz}$ such that $d_X(y, p) = d_{\mathbb{E}^2}(\overline{y}, \overline{p})$ and $d_X(p, z) = d_{\mathbb{E}^2}(\overline{p}, \overline{z})$. Now, picking any two points p, q on $\triangle xyz$, a CAT(0) space must have the relationship:

$$d_X(p, q) \leq d_{\mathbb{E}^2}(\overline{p}, \overline{q}).$$

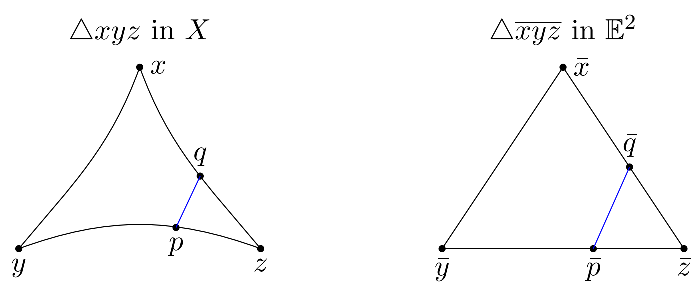

Figure 2.8: For a CAT(0) space X, the geodesic $[x, z]$ in X is the same length as the geodesic $[\overline{x}, \overline{z}]$ in $\mathbb{E}^2$. The same relationship holds for the other two sides of both triangles.

Lemma 2.2.6 (See [BH99])**.** *A CAT(0) space X has the following properties:*

1. *It is uniquely geodesic, that is, for any two points x, y in X, there exists exactly one geodesic connecting them.*

2. *The nearest-point projection from a point x to a geodesic b is a unique point. In fact, the closest-point projection map $\pi_b : X \to b$ is 1-Lipschitz.*

3. *Convexity: For any convex set $Z \in X$, the distance function $f : X \to \mathbb{R}^+$ given by $f(x) = d(x, Z)$ is convex in the sense of [BH99, II.2.1]*

2.3 CAT(0) Cube Complexes

As many of our main results involve importing cubical techniques to non-cube complex spaces, it is useful to have proper a background in cube complexes along with their useful tools. We give the introduction and motivation necessary for CAT(0) cube complexes in this section.

Given some n-dimensional cube $[0, 1]^n$, we define the *faces* of the cube by isometrically embedded copies of $[0, 1]^{n-1}$ along the boundary of the n-cube. For instance, a 3-cube has six faces (the six sides of the three dimensional cube). Likewise, a 2-cube, *i.e.* a square, has four faces (the four sides of the square). A *cube complex* is defined by gluing Euclidean cubes of varying dimensions via isometries along their faces. We call the 0-cubes of the cube complex the *vertices.* For this thesis, it suffices to only consider cube complexes that are *simply connected* — *i.e.* the cube complex has trivial fundamental group (see Definition 2.3.3). All these terms are the normal vocabulary for what's to come. We now give definitions that build up to CAT(0) cube complexes.

Definition 2.3.1 (simplicial complex, clique, flag complex)**.** A simplicial complex, like a cube complex, is formed by gluing varying dimensional simplicies together via isometries of their faces. We also call the 0-simplicies vertices. A *clique* is a set of of k vertices for some

k such that each pair of vertices has an edge connecting them. A simplicial complex is a *flag complex* if each clique of k vertices spans a $(k-1)$-simplex.

Definition 2.3.2 (vertex link)**.** Let X be a cube complex, for each vertex $v \in X$, we create the simplicial complex $lk(v)$ as follows.

1. For each 1-cube e_i with v as a face, create a vertex for $lk(v)$. We will also label this vertex e_i. Thus, the set of vertices for $lk(v)$ is $\mathfrak{E} = \{e_i : e_i$ is an edge that has v as a face$\}$

2. If $e_i, e_j \in \mathfrak{E}$ are both among a shared n-cube in X, connect an edge between e_i and e_j in $lk(v)$

3. Similarly, if $\{e_1, \cdots, e_n\} \subset \mathfrak{E}$ are all among a shared n-cube in X, then $\{e_1, \cdots, e_n\}$ span a simplex in $lk(v)$. See Figure 2.9 for examples.

Definition 2.3.3 (CAT(0) cube complex)**.** Let X be a cube complex. X is a *CAT(0) cube complex* if and only if

1. X is simply connected.

2. For each vertex v in X, its link $lk(v)$ is a flag complex.

This equivalent definition to a CAT(0) cube complex is often nice because it can be tough to establish the standard CAT(0) condition, in general. Under this definition, one can restrict to just looking at vertices and their links, which is significantly easier. There are many reasons to care about CAT(0) cube complexes. Among which are a few.

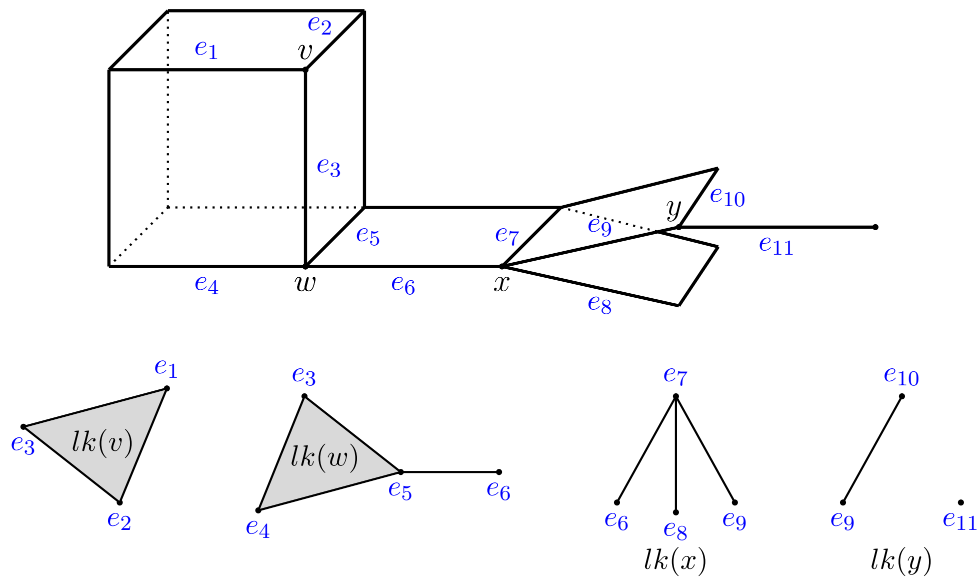

Figure 2.9: A CAT(0) complex comprising of one 3-cube, three 2-cubes, and one 1-cube glued together via isometries of their faces. For the vertices v, w, x, y, we see that all the links are flag complexes. If, say, the 3-cube was instead not filled in, then $lk(v)$ and $lk(w)$ would not have a filled in 2-simplex either. This would result in $lk(v)$ and $lk(w)$ not being flag complexes and the cube complex not being CAT(0).

1. There are many groups that act geometrically on CAT(0) Cube Complexes. A few popular examples are right angled Artin groups [CD95], some Coxter groups [NR03], small cancellation groups [Wis04], and hyperbolic 3-manifold groups [KM12]. Thus, as the beginning of this chapter showed, when one studies the metric behavior of a CAT(0) cube complex, one also learns more about the above groups.

2. The use of CAT(0) cube complexes helped resolve the resolutions of the virtual Haken and fibering conjectures in 3-manifold theory [Ago13, Wis21].

3. CAT(0) Cube Complexes have a strong relationship with Mapping Class Groups. For one, they both fall into a class of spaces called *hierarchically hyperbolic spaces*, so they already have comparable features. Also, mapping class groups can “locally”

look like a CAT(0) cube complex. That is, given a finite set of points and unbounded geodesic rays, the *geodesic hull* of this set is quasi-isometric to a CAT(0) cube complex [BHS17, Dur23]. The usefulness of this property has turned into an investigation of *locally quasi-cubical* spaces, *i.e.* the collection of spaces whose finite hulls are quasi-isometric to a CAT(0) cube complex [DZ22].

4. CAT(0) cube complexes have a rank rigidity theorem [CS11]. Rank rigidity has been conjectured to be true for all CAT(0) spaces, but it has only been proven for CAT(0) cube complex due to the extra cubical structure.

All of these are very relevant reasons for the significance of CAT(0) cube complexes, but perhaps the most important reason one likes CAT(0) cube complexes is because their geometry is *easy*. This is due to the construction of *hyperplanes*, which we now discuss. Given a CAT(0) cube complex X, edges e_1, e_2 in X will be among the same equivalence class if e_1 and e_2 are on opposite ends of a square in X.

Definition 2.3.4 (Midcube, Hyperplane)**.** for a cube $[0,1]^n$, a *midcube* is the subset formed by restricting one of the intervals $[0,1]$ to $\frac{1}{2}$. So, each n-dimensional cube will have n midcubes. Given an equivalence class of edges $[e]$, the *hyperplane* dual to $[e]$ is the collection of midcubes which intersect edges in $[e]$.

There are many nice features of hyperplanes, including but not limited to:

- Hyperplanes are CAT(0) cube (sub)complexes.

- Hyperplanes are closed as a set.

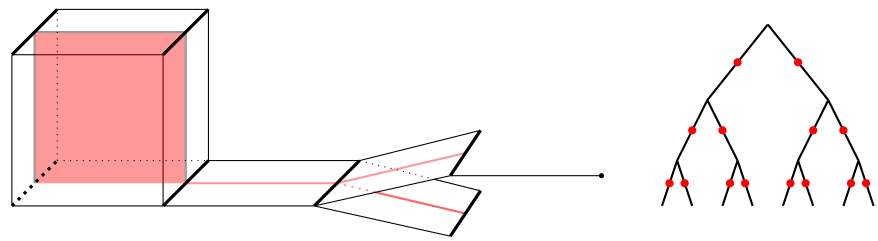

Figure 2.10: Hyperplanes in CAT(0) cube complexes. Notice how hyperplanes of an edge are just a point, and midcubes are codimension 1 in the cubes they live in.

- Each hyperplane creates two complement halfspace components. For example, in the first figure in Figure 2.10, one halfspace is the part of the cube complex in front of the hyperplane. The other halfspace is the part of the cube complex behind the hyperplane.

- Hyperplanes are *convex*: for any two points in a hyperplane, the geodesic connecting these two points is also contained in the hyperplane.

- Hyperplanes can determine the distance between two points. That is, for a CAT(0) cube complex X, and any $x, y \in X$, we can define the metric

$$d_{\mathcal{H}}(x,y) = \big|\{\text{number of hyperplanes separating}\, x \,\text{and}\, y\}\big|.$$

  Denoting (X, d) as the CAT(0) cube complex with the CAT(0) metric, it turns out that $id : (X, d_{\mathcal{H}}) \longrightarrow (X, d)$ is a quasi-isometry. Since the field of geometric group theory cares about quasi-isometry invariants, it suffices to uses the $d_{\mathcal{H}}$ metric instead of the CAT(0) metric.

This last point above might be the most relevant to us. Hyperplanes are a discrete set of objects and are often easy to count. This allows us to view certain situations in a combinatorial fashion. The following theorem is an example that shows this theme.

Theorem 2.3.5 ([CS15])**.** *Let X be a CAT(0) cube complex. A geodesic ray α is D-contracting if and only if there exists a $C > 0$ and an infinite sequence of hyperplanes $\{h_1, h_2, \cdots\}$ crossing α such that*

- *h_i, h_{i+1} are C-separated for all i.*
- *$d(h_i, h_{i+1}) \leq C$ for all i.*

We will cover the definitions of *contracting* in Chapter 3 and *C-separated* in a Chapter 4. For now, take as a fact that a contracting geodesic is a geodesic traveling through negative curvature. Thus, Theorem 2.3.5 tells us one can find negative curvature of a CAT(0) cube complex by only looking at hyperplanes. In fact, the list of reasons for the relevance of CAT(0) cube complexes are all directly or indirectly because of the complex's hyperplane structure. This segues into a natural question:

Question 2.3.6. Hyperplanes in CAT(0) cube complexes are very useful. Is there a analogue of a hyperplane in non-cube complex spaces that would also be similarly useful?

The answer of this is yes, and authors Petyt-Spriano-Zalloum in [PSZ22] create one in the context of CAT(0) spaces (with foregoing work to generalize this further). A recap of this work is given in Chapter 4. It is the utility of this new combinatorial tool that allows us to generalize Theorem 2.3.5 in Chapter 5.

Chapter 3

Quasi-Isometry Invariant Boundaries

3.1 Visual Boundary

Depending on the context, there are many useful boundaries that are relevant for proper geodesic metric spaces. The first boundary one likely gets exposed to is the *visual boundary.* For a proper geodesic metric space X, fix a basepoint $\mathfrak{o} \in X$. Consider the collection of all geodesic rays that start at $\mathfrak{o}$ and extend out towards a direction of infinity. For a picture, one can imagine these rays being eyesight lines that stretch out towards a horizon. Each direction one looks at becomes a new eyesight line and, hence, a new ray. For a geodesic ray γ, denote $N_m(\gamma)$ as the set of all points within m distance of γ. Define an equivalence class $\sim$ among geodesic rays based at $\mathfrak{o}$ where $\gamma_1 \sim \gamma_2$ if $\gamma_1 \subset N_m(\gamma_2)$ for some m and vice versa. When $\gamma_1 \sim \gamma_2$, we say that γ_1 and γ_2 *fellow travel* each other, or are of *asymptotic equivalence*.

Definition 3.1.1 (Visual Boundary)**.** Given a proper geodesic metric space X and basepoint $\mathfrak{o} \in X$, the *visual boundary*, denoted $\partial_\infty X$, is the set of equivalence classes of rays based at $\mathfrak{o}$.

Note, it turns out that the visual boundary is not dependant on basepoint, so we often omit $\mathfrak{o}$ in the notation. The natural topology to endow on $\partial_\infty X$ is as follows:

Definition 3.1.2 (Topology of the Visual Boundary)**.** The *cone topology* of $\partial_\infty X$ is generated by the basic open sets

$$U_{T,\epsilon}(\gamma) = \{\gamma' \,:\, d(\gamma(t), \gamma'(t) \leq \epsilon \text{ for all } t \leq T\}$$

In other words, γ and γ' are close in $\partial_\infty X$ if they fellow travel each other for their initial segments of length T. The larger T becomes, the longer the rays must fellow travel.

Example 3.1.3. For $X = \mathbb{R}^2$, choosing the basepoint as the origin for convenience, we see that geodesic rays based at the origin linearly diverge from each other. Thus, any geodesic ray γ based at the origin will be among its own equivalence class. This means that each $[\gamma] \in \partial_\infty \mathbb{R}^2$ is synonymous with the angle γ has relative to the x-axis. With angles relative to the x-axis ranging from 0 to 2π, the cone topology endowed on $\partial_\infty \mathbb{R}^2$ gives us that $\partial_\infty X$ is a circle.

Example 3.1.4. Similarly for $X = \mathbb{H}^2$, the hyperbolic plane, given any basepoint, no two geodesic rays emanating from the same base point will fellow travel each other. Similarly, we get that $\partial_\infty \mathbb{H}^2$ is also a circle. This is intuitive since the Poincare disk model of $\mathbb{H}^2$ has a circle that bounds $\mathbb{H}^2$ in the model.

Example 3.1.5. In the first example of Example 2.1.2, $X = Cay(\mathbb{F}_2, S)$ will have a visual boundary that is a Cantor set.

Now, in the context of hyperbolic groups, Gromov showed the following in [Gro87].

Theorem 3.1.6. *If a finitely generated group G acts geometrically on two δ-hyperbolic metric spaces X_1, X_2, then there is a G-equivariant homeomorphism $\partial_\infty X_1 \longrightarrow \partial_\infty X_2$.*

In other words, Gromov showed that visual boundaries are a quasi-isometry invariant in the context of δ-hyperbolic spaces. This allows us to give a well-defined notion of a visual boundary for a group. Gromov also asked whether the same result holds for non-δ-hyperbolic spaces in [Gro92]. Croke and Kleiner show in [CK00] that this is not the case even for CAT(0) spaces. We give non-exhaustive description of their counterexample now. Let $G = \langle a, b, c, d : [a, b], [b, c], [c, d]\rangle$. There are two spaces that G acts geometrically on. The first space is the universal cover of its Salvetti Complex (See [Cha07]), and the second space is a slight perturbation of the first space:

Example 3.1.7. Let T_2 be a flat torus. Let b', c' be cannonical π_1-generating simple closed curves of length 1 in T_2 that meet at a single point at angle $\alpha = \frac{\pi}{2}$. Let T_1, T_3 also be flat tori with a, b being π_1-generating simple closed curves of length one in T_1 and c, d being π_1-generating simple closed curves of length one in T_3. Now, let $X_{\frac{\pi}{2}}$ be the union of T_1, T_2, T_3 with b', c' identified isometrically with b, c, respectively. See Figure 3.1. Let $\widetilde{X}_{\frac{\pi}{2}}$ be the universal cover of $X_{\frac{\pi}{2}}$. Note $\widetilde{X}_{\frac{\pi}{2}}$ is a CAT(0) cube complex.

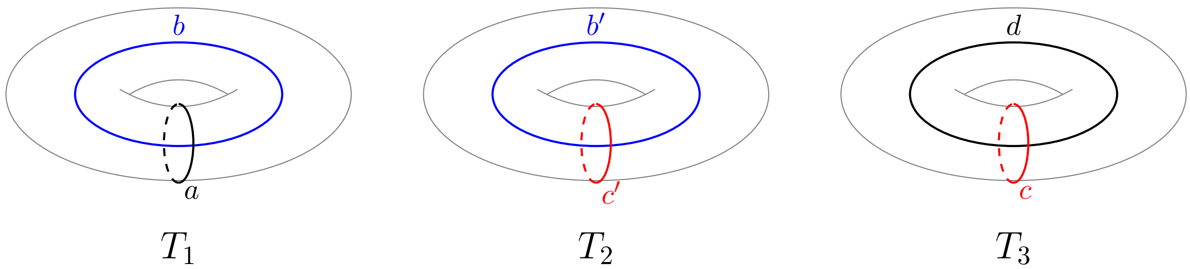

Figure 3.1: The space $X_{\frac{\pi}{2}}$. We identify the blue curves b and b' as well as the red curves c andc'. $\widetilde{X}_{\frac{\pi}{2}}$ is the universal cover of this space.

It might not be hard to see that G acts geometrically on $\widetilde{X}_{\frac{\pi}{2}}$. Indeed, if one coincides the natural generators of G with lifts of the simple closed curves of $X_{\frac{\pi}{2}}$ (abusing notation, we denote both the generators and the curves by a, b, c, d), one sees that the Cayley graph of G becomes the 1-skeleton of the CAT(0) cube complex structure of $\widetilde{X}_{\frac{\pi}{2}}$.

Now, let $\widetilde{X}_\alpha$ be a similarly constructed space, but with the simple closed curves b', c' meeting at a single point at angle α for any $0 < \alpha < \frac{\pi}{2}$. See Figure 3.2. It's not tough to also show that G acts geometrically on $\widetilde{X}_\alpha$, and Croke and Kleiner showed, surprisingly, that $\widetilde{X}_{\frac{\pi}{2}}$ and $\widetilde{X}_\alpha$ have visual boundaries of different homeomorphism type. In fact, Wilson showed in [Wil05] that even though G acts geometrically on $\widetilde{X}_\alpha$ for every $0 < \alpha \leq \frac{\pi}{2}$, each $\widetilde{X}_\alpha$ produces a different visual boundary.

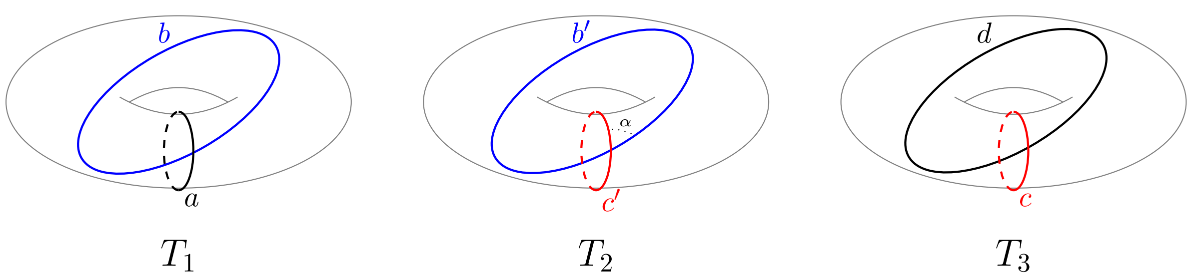

Figure 3.2: The space X_α. When one chooses a different b' that meets c' at a different angle, its universal cover will have a visual boundary of different homeomorphism type from $\partial_\infty \widetilde{X}_{\frac{\pi}{2}}$

So, even for CAT(0) groups, there is not a notion of a well-defined visual boundary and Theorem 3.1.6 suggests that hyperbolicity is a key factor for a boundary that is quasi-isometry invariant. So, in order to create a quasi-invariant boundary for a non-hyperbolic space, one might instead restrict to only the "hyperbolic directions" of the non-hyperbolic space. We explore this in the next section.

3.2 Morse Boundaries

More specifically, since the visual boundary was defined through geodesic rays emanating from a chosen basepoint, restricting to geodesic rays that behave similarly to the geodesic rays in δ-hyperbolic space might result in a boundary that is quasi-isometry invariant. The Morse lemma (Lemma 2.2.3) gives a characterization of all geodesics in a δ-hyperbolic space. In any proper geodesic metric space, we characterize geodesics that satisfy the same property through the following definition.

Definition 3.2.1. Let X be a proper metric space. A set Z in X is said to be N-*Morse* if there exists a function $N = N(K, C)$ such that for any (K, C)-quasi-geodesic ϕ with endpoints on Z, we have $\phi \subset Nbhd(Z, N)$. We call Z *Morse* if it is N-Morse for some N. We call N the *Morse gauge.* Note, in this paper, all sets Z will be either a geodesic or a quasi-geodesic.

Example 3.2.2. When $X = \mathbb{R}^2$, there are no Morse geodesic rays. As Figure 3.4 shows, any geodesic ray γ in $\mathbb{R}^2$ has a class of $(3, 0)$-quasi-geodesic rays whose Hausdorff distance grows farther and farther away from γ. However, every finite geodesic segment γ' is Morse

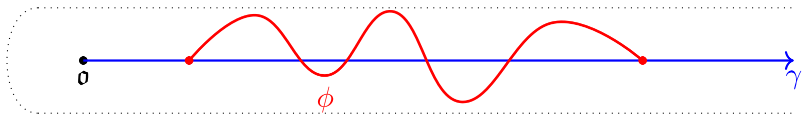

Figure 3.3: A N-Morse geodesic γ. Given any (K, C)-quasi-geodesic ϕ, the $N(K, C)$-neighborhood of γ will contain ϕ. Not surprisingly, this figure is a replica of Figure 2.6, which represented the Morse Lemma.

in $\mathbb{R}^2$. For intuition behind this, if we tried our little trick of $(3, 0)$-quasi-geodesics again for γ', eventually the $(3, 0)$-quasi geodesics will not have endpoints on γ'. It follows that for any (K, C), one can find a bounded neighborhood (namely $N(K, C)$) that contains all (K, C)-quasi-geodesics. Thus, in the tree of flats example, a geodesic ray is Morse if and only if there exists a uniform bounded length on the segments of the geodesic traveling through a single flat. See Figure 3.4.

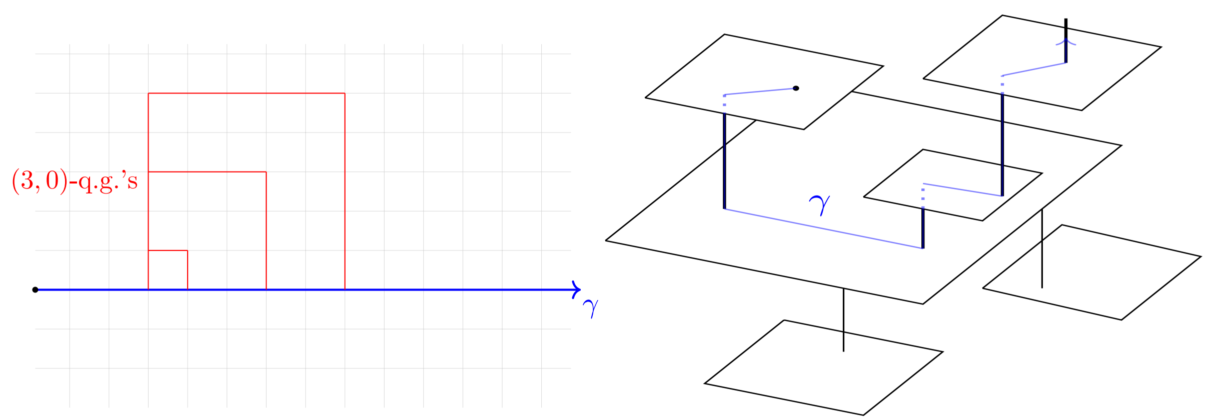

Figure 3.4: In the case of $X = \mathbb{R}^2$, there does not exist any N-Morse geodesic rays since any geodesic ray has a class of $(3, 0)$-quasi-geodesics at a arbitrarily far Hausdorff distance away from the geodesic. In the case X is the tree of flats, a geodesic is N-Morse provided it has uniform bounded diameter in each flat it travels through.

Definition 3.2.3 (Morse Boundary)**.** Let $\mathfrak{o} \in X$ be a base point. The *Morse boundary*, denoted $\partial_M X_{\mathfrak{o}}$ is the set of all Morse quasi-geodesic rays emanating from $\mathfrak{o}$ up to asymptotic equivalence. Note, there will always be a geodesic in each equivalence class.

One might decide to endow $\partial_M X_{\mathfrak{o}}$ with the subspace topology of the cone topology on the visual boundary. However, in [Cas16], Cashen provides two quasi-isometric CAT(0) spaces that, under the cone topology, have Morse boundaries of different homeomorphism type. This goes against the motivation to create the Morse boundary in the first place. When the Morse boundary was first created by Cordes in [Cor17], he made the *direct limit topology* which did indeed become a quasi-isometry invariant. However, instead of focusing on the direct limit topology, we now make steps towards a different topology on the Morse boundary we eventually generalize in future sections. We say a function $\rho : [0, \infty) \to [1, \infty)$ is sublinear if $\lim_{x\to\infty} \frac{\rho(x)}{x} = 0$. We also require ρ to be increasing and concave.

Definition 3.2.4 (contracting)**.** Let Z be a closed subset of X and $\pi_Z : X \to 2^Z$ be the closest point projection to Z. We say that Z is *contracting* if there is a sublinear function ρ such that for all x and y in X,

$$d(x, y) \leq d(x, Z) \quad \Longrightarrow \quad \operatorname{diam}(\pi_Z(x) \cup \pi_Z(y)) \leq \rho(d(x, Z))$$

The authors of [ACGH17] prove that for every Morse gauge N, there exists a ρ depending only on N so that every N-Morse set is contracting with respect to the sublinear function ρ. Conversely, for every ρ, there exists an N so that every contracting ray with respect to ρ is N-Morse. In other words, Morse and contracting are equivalent definitions.

The following descriptions were taken from [IM21, He23]. Given a sublinear function ρ and constants $K \geq 1, C \geq 0$, define the constant

$$\theta(\rho, K, C) = \max\left\{3K, 3C^2, 1 + \inf\left\{R > 0 \mid \forall r \geq R, 3K^2\rho(r) < r\right\}\right\}.$$

Now fix $\mathfrak{o} \in X$. Let $\xi \in \partial_M X_{\mathfrak{o}}$. Then there exists ρ such that all geodesic representatives of ξ based at $\mathfrak{o}$ are contracting with sublinear function ρ. Choose one geodesic representative $\gamma \in [\xi]$ that is based at $\mathfrak{o}$. Let $R \geq 0$. We say that a (K, C)-quasi-geodesic β fellow travels along γ for distance R, if $\beta \cap N_{\theta(\rho,K,C)}(\gamma \backslash B_R(\mathfrak{o})) \neq \emptyset$, where $B_R(\mathfrak{o})$ denotes the open ball of radius R centered at $\mathfrak{o}$. This is saying that β is within θ distance of γ at some point past $\gamma(R)$. We define the set $U_{\mathfrak{o},R}(\xi) := \{\eta \in \partial_M X_{\mathfrak{o}} \mid$ All quasi-geodesic representatives of η that are based at $\mathfrak{o}$ fellow-travel along γ for distance $R\}$.

Definition 3.2.5 (Cashen-Mackay Topology)**.** For a proper geodesic metric space X and fixed basepoint $\mathfrak{o} \in X$, the *Cashen-Mackay topology* of $\partial_M X_{\mathfrak{o}}$ is the topology whose neighborhood basis is formed by the sets $\{U_{\mathfrak{o},R}(\xi)\}_{R,\xi}$. We denote this topology by $\mathcal{FQ}$ for "fellow-travelling quasi-geodesics".

This gives the following useful properties for the Morse boundary, which have been proven through [Cor17, CM19].

Theorem 3.2.6. *Given a proper geodesic metric space X and a fixed basepoint $\mathfrak{o} \in X$, the Morse boundary $\partial_M X_{\mathfrak{o}}$ equipped with the $\mathcal{FQ}$ topology is*

1. *a visibility space, i.e., any two points in the Morse boundary can be joined by a bi-infinite Morse Geodesic.*

2. independent of choice of basepoint.

3. a quasi-isometry invariant.

4. homeomorphic to the visual boundary if X is hyperbolic.

5. metrizable.

Note, all these results still hold under the direct limit topology as well besides being metrizable, which was a motivation for the creation of the $\mathcal{FQ}$ topology [CM19]. Since the Morse boundary is independent of basepoint, we often omit the basepoint in the notation and write $\partial_M X$. Again, the most relevant result is the quasi-isometry invariance. Indeed, since every finitely generated group acts geometrically on a proper geodesic metric space, quasi-isometry invariance gives us the ability to define a Morse boundary *for the group*.

Now, one could be content here with the Morse boundary, as it accomplishes the goal of creating a well defined boundary of a finitely generated group. However, one can still push for improvements.

Question 3.2.7. Is there a larger notion of a boundary compared to the Morse boundary that still remains a quasi-isometry invariant?

The reasoning for such a question is that a boundary that is larger as a set than the Morse boundary can likely contain more information than the Morse boundary. In fact, one can ask a better question: what is the "largest" quasi-isometry invariant boundary for proper geodesic metric spaces. This next section aims to step towards this direction.

3.3 Sublinear Morse Boundaries

For further details and expansion on below definitions, see [QR22].

Definition 3.3.1 (Sublinear function)**.** We fix a function

$$\kappa : [0, \infty) \to [1, \infty) \quad \text{such that} \quad \lim_{t \to \infty} \frac{\kappa(t)}{t} = 0.$$

This second requirement makes κ *sublinear*. We also add the convention that κ be monotone increasing and concave. For the chosen fixed basepoint $\mathfrak{o} \in X$, we denote $\kappa(x) = \kappa(d(\mathfrak{o}, x))$.

Oftentimes, we create upper bounds of distances in terms of κ with a specific input. The following computational lemma gives us a way to change our inputs of κ for a more suitable upper bound, and it is used throughout the thesis.

Lemma 3.3.2 (Lemma 3.2 in [QR22])**.** *For any $D_0 > 0$, there exist $D_1, D_2 > 0$ depending only on D_0 and κ so that for $x, y \in X$, we have*

$$d(x, y) \leq D_0 \kappa(x) \Rightarrow D_1 \kappa(x) \leq \kappa(y) \leq D_2 \kappa(x).$$

Definition 3.3.3 (κ-neighborhood)**.** Fix a basepoint $\mathfrak{o} \in X$ once and for all. For a closed set Z and a constant $\mathrm{n} \geq 0$, define the (κ, n)*-neighbourhood* of Z to be

$$\mathcal{N}_\kappa(Z, \mathrm{n}) = \{x \in X \mid d_X(x, Z) \leq \mathrm{n} \cdot \kappa(x)\}.$$

Note: We often abbreviate to κ*-neighborhood*, and our closed set Z will always be a geodesic or quasi-geodesic in this thesis.

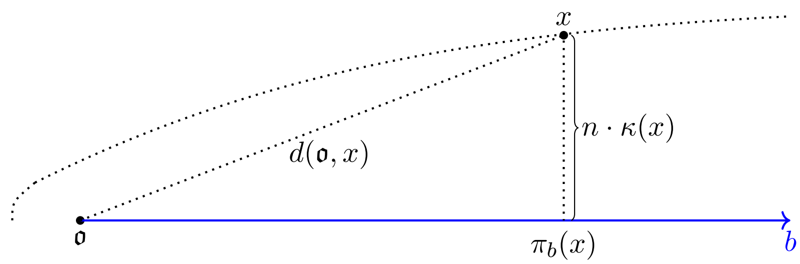

Figure 3.5: A (κ, n)-neighborhood of geodesic b.

Definition 3.3.4 (κ-fellow travelling)**.** Let α and β be two infinite quasi-geodesic rays in X. If α is contained in some κ-neighbourhood of β and β is contained in some κ-neighbourhood of α, we say that α and β *κ-fellow travel* each other. This defines an equivalence relation on the set of quasi-geodesic rays in X. It is known that, in CAT(0) spaces, each equivalence class of quasi-geodesics contains a unique geodesic ray emanating from $\mathfrak{o}$. (See [QR22]).

Definition 3.3.5 (κ-Morse geodesics)**.** A geodesic b is *κ-Morse* if there is a function $m_b : \mathbb{R}_+^2 \to \mathbb{R}_+$ so that if $\alpha : [s,t] \to X$ is a (K, C)-quasi-geodesic with end points on b then

$$\alpha[s,t] \subset \mathcal{N}_\kappa\left(b, m_b(K,C)\right).$$

We refer to m_b as the *Morse gauge* for b. We also always assume that $m_b(K,C)$ is the largest element in the set $\{K, C, m_b(K,C)\}$. Note that when $\kappa \equiv 1$ we recover the standard definition of a Morse geodesic, so this definition can truly be seen as a sublinear generalization of Morse.

Extending Morseness in a sublinear fashion is a rather intuitive try to create a larger quasi-invariant boundary. The construction of the Morse boundary was not the whole history of creating quasi-isometrically invariant boundaries. Relooking at Definition

3.2.4, a geodesic is *D-strongly contracting* if the geodesic is contracting for $\rho \equiv D$ where D is a positive constant. All geodesics in a δ-hyperbolic space are uniformly D-strongly contracting for some $D = D(\delta)$, so one could instead use the strongly contracting definition instead of the Morse definition as a characterization for "hyperbolic" geodesics in non-hyperbolic spaces. Charney and Sultan prove success of this in [CS15] as they created a quasi-isometry invariant *strongly contracting boundary* in the context of CAT(0) spaces. So, we have four kinds of characterizations for hyperbolic geodesics: strongly contracting, contracting, Morse, and κ-Morse. Now, a strongly contracting geodesic is also contracting, as Definition 3.2.4 is a sublinear extension of strongly contracting. Also, It is shown by [ACGH17] that contracting is equivalent to Morse. Thus, as all Morse geodesics are also κ-Morse, we have that the set of κ-Morse geodesic rays will be the *largest collection* of rays compared to the three prior definitions. See Figure 3.6. So, if broadening from strongly contracting to contracting (aka Morse) geodesics produces a larger quasi-isometry invariant boundary, then possibly broadening from Morse to κ-Morse geodesics will also produce a larger quasi-invariant boundary.

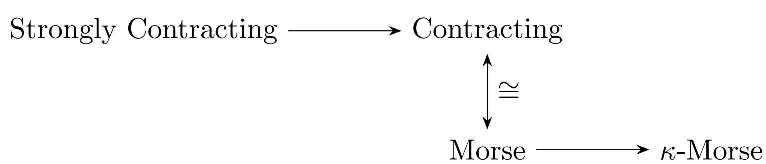

Figure 3.6: A diagram comparing different definitions of "hyperbolic" geodesics. Strongly contracting is the strongest definition. Contracting and Morse are equivalent via [ACGH17]. Lastly, κ-Morse is the weakest definition. That is, all other definitions imply κ-Morse.

Definition 3.3.6 (Sublinearly Morse boundary)**.** Let κ be a sublinear function and let X be a proper metric space. We define the *κ-Morse boundary*, as a set, by

$$\partial_\kappa X := \{ \text{ all } \kappa\text{-Morse quasi-geodesic rays } \}/\kappa\text{-fellow travelling} .$$

Definition 3.3.7 (Sublinearly Morse boundary, Sublinearly Morse topology)**.** The κ-Morse boundary, $\partial_\kappa X$, can be equipped with the *sublinearly Morse topology*: Fix a base point $\mathfrak{o}$, let $\xi \in \partial_\kappa X$, and let b be the unique geodesic representative of ξ that starts at $\mathfrak{o}$. For all $r > 0$, we define $U_\kappa(b, r)$ to be the set of all points $\eta \in \partial_\kappa X$ such that for every $(K, C)-$quasi-geodesic β representing η, starting at $\mathfrak{o}$, and satisfying $m_b(K, C) \leq \frac{r}{2\kappa(r)}$, we have $\beta|_{[0,r]} \subset \mathcal{N}_\kappa\left(a, m_a(K, C)\right)$. we denote the sublinear Morse topology as $\mathcal{SM}$.

It's clear that $\partial_M X \subset \partial_\kappa X$ and the $\mathcal{SM}$ topology is heavily based on the $\mathcal{FQ}$ topology for the Morse Boundary. Thus, we get the following theorem for $\partial_\kappa X$.

Theorem 3.3.8 ([QR22, QRT20])**.** *Given a proper geodesic metric space X and a fixed basepoint $\mathfrak{o} \in X$, the Morse boundary $\partial_\kappa X$ equipped with the $\mathcal{SM}$ topology is*

1. *a visibility space, i.e., any two points in the Morse boundary can be joined by a bi-infinite κ-Morse Geodesic. [DZ22]*
2. *independent of choice of basepoint.*
3. *a quasi-isometry invariant.*
4. *homeomorphic to the visual boundary if X is hyperbolic.*
5. *metrizable.*

6. homeomorphic to $\partial_M X$ with the $\mathcal{FQ}$ topology if $\kappa \equiv 1$ [He23]

Example 3.3.9. Just like in Example 3.2.2, $\partial_\kappa X = \emptyset$ when $X = \mathbb{R}^2$. In the example where X is the tree of flats, a geodesic γ is sublinearly Morse if it's diameter in individual flats is sublinearly bounded with respect to $\mathfrak{o}$. That is, the farther γ travels, the longer γ can stay in a flat.

A great tool between the cone topology and the $\mathcal{SM}$ topology on $\partial_\kappa X$ is the following lemma. As many of our arguments in Chapter 6 use open sets from the cone topology, the proofs also immediately transfer over to the $\mathcal{SM}$ topology due to this lemma.

Lemma 3.3.10 (Lemma 2.12 in [IMZ23])**.** *Let X be a proper CAT(0) space. The cone topology restricted to the set $\partial_\kappa X$ is coarser than the $\mathcal{SM}$ topology.*

The main new application in the κ-Morse boundary setting that prior quasi-isometry invariant boundaries have failed in is it's relationship to random walks and the Poisson Boundary. For background on the Poisson boundary and random walks, see [Woe94]. Being brief, Qing-Rafi-Tiozzo show in [QRT20] that if for almost every random walk there exists a κ-Morse ray that the random walk sublinearly tracks, then $\partial_\kappa X$ is a topological model for the Poisson boundary. Since this result, random walks have been shown to sublinearly track κ-Morse geodesics in right angle Artin groups [QR22], mapping class groups and relatively hyperbolic groups [QRT20], in hierarchically hyperbolic groups [NQ22] and in CAT(0) groups [Cho22, GQR22] under various assumptions — all such groups of great interest in geometric group theory. Thus, unlike the Morse boundary, $\partial_\kappa X$ serves as a bridge between two areas of research in geometric group theory: quasi-invariant boundaries and genericity through random walks.

Chapter 4

Curtains and the Curtain Model

We now import many cubical results created in [PSZ22]. The most important of which is the definition of a *curtain* as defined below. Curtains in CAT(0) spaces become the analogue of hyperplanes in CAT(0) cube complexes we desire. The main comparison we are interested in is how geodesics crossing curtains in CAT(0) spaces mimic behavior of geodesics crossing hyperplanes in CAT(0) cube complexes. In order to show this comparison in Section 5, we review the lemmas and definitions below.

4.1 Curtain Machinery

Definition 4.1.1 (Curtain, Pole)**.** Let X be a CAT(0) space and let $b : I \to X$ be a geodesic. For any number r such that $[r - \frac{1}{2}, r + \frac{1}{2}]$ in in the interior of I, the *curtain dual to* b at r is

$$h = h_b = h_{b,r} = h_{b,P} = \pi_b^{-1}(b[r - \tfrac{1}{2}, r + \tfrac{1}{2}])$$

where π_b is the closest point projection to b. We call the segment $b[r - \frac{1}{2}, r + \frac{1}{2}]$ the *pole* of the curtain which we denote as P when needed.

It is worth noting that curtains $h_{b,r}$ are defined from some geodesic b at time r, but we often use the simpler notation h when such information is not needed or already implied. There are certain properties that curtains and hyperplanes have in common. An example is that both curtains and hyperplanes separate their complements into two component *half spaces* which we denote as h^- and h^+. Also, both curtains and hyperplanes are closed as sets.

Remark 4.1.2. A notable difference between curtains and hyperplanes is that curtains are not convex, see [PSZ22, Remark 2.4] for more details.

Definition 4.1.3 (Chain, Separates)**.** A curtain h *separates* sets $A, B \subset X$ if $A \subset h^-$ and $B \subset h^+$. A set $\{h_i\}$ is a *chain* if each of the h_i are disjoint and h_i separates h_{i-1} and h_{i+1} for all i. We say a chain $\{h_i\}$ *separates* sets $A, B \subset X$ if each h_i separates A and B.

The notion of chains separating two sets A and B can give a description of the distance between sets A and B. More specifically, a maximal chain that separates two points $x, y \in X$ tells us the distance between x and y, as shown in the following lemma.

Lemma 4.1.4 (Lemma 2.10 in [PSZ22])**.** *For any $x, y \in X$, there is a chain c of curtains dual to $[x, y]$ such that $1 + |c| = \lceil d(x, y) \rceil$.*

Many of our arguments involve a geodesic crossing a curtain. A geodesic b will *cross* a curtain h if there exists $s < t$ such that $b(s)$ and $b(t)$ are separated by h. The next lemma describes how b interacts with the pole of h when b crosses h, and we use this lemma implicitly throughout the paper.

Lemma 4.1.5 (Lemma 2.5 in [PSZ22])**.** *Let $h = h_{b,r}$ be a curtain, and let $x \in h^-, y \in h^+$. For any continuous path $\gamma : [c,d] \to X$ from x to y and any $t \in [r - 1/2, r + 1/2]$, there is some $p \in [c,d]$ such that $\pi_b(\gamma(p)) = b(t)$.*

In particular, if a geodesic a crosses a curtain $h_{b,r}$, there exists some $p \in a$ such that $\pi_b(a(p)) = b(r)$.

Definition 4.1.6 (L-separated, L-chain)**.** Let $L \in \mathbb{N}$. Disjoint curtains h and h' are said to be *L-separated* if every chain meeting both h and h' has cardinality at most L. Two disjoint curtains are said to be *separated* if they are L-separated for some L. If c is a chain of curtains such that each pair is L-separated, then we refer to c as an *L-chain.* See Figure 4.1 for an example of L-separation.

Lemma 4.1.7 (Lemma 2.21 in [PSZ22])**.** *Let $L, n \in \mathbb{N}$, let $\{h_1, \ldots, h_{(4L+10)n}\}$ be an L-chain, and suppose that $A, B \subset X$ are separated by every h_i. For any $x \in A$ and $y \in B$, the sets A and B are separated by an L-chain of length at least $n+1$ all of whose elements are dual to $[x,y]$ and separate h_1 from $h_{(4L+10)n}$.*

Definition 4.1.8 (L-metric)**.** Denote X_L for the metric space (X, d_L), where d_L is the metric defined as

$$\mathrm{d}_L(x,y) = 1 + \max\{|c| : c \text{ is an } L\text{-chain separating } x \text{ from } y\}$$

with $d_L(x,x) = 0$. The function is indeed a metric by [PSZ22, Lemma 2.17], and one can replace the 1 in the above function by any $\epsilon > 0$. Note that, by Remark 2.16 in [PSZ22], we have that for any $x, y \in X$, it follows that $d_L(x,y) < 1 + d(x,y)$.

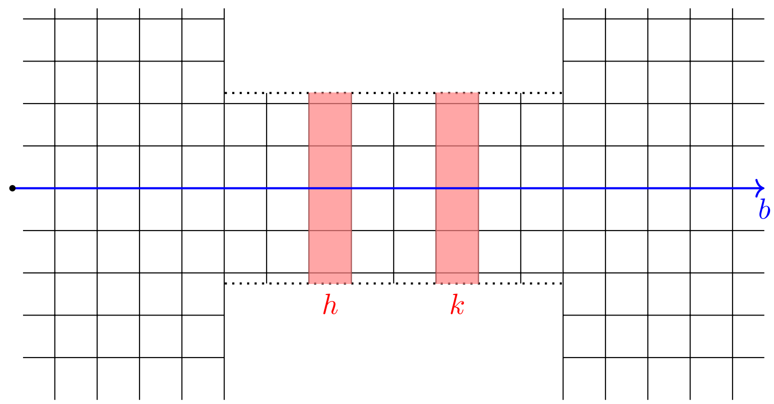

Figure 4.1: The above figure shows $\mathbb{R}^2$ with 2 strips cut out. We see the pair of curtains h, k dual to the geodesic b are 4-separated. A maximal chain that crosses both curtains would be 4 horizontal curtains. Thus, $\{h, k\}$ is a 4-chain.

These X_L spaces will be used as auxiliary spaces to define the curtain model (See Definition 6.1.1).

Theorem 4.1.9 (Theorem 3.5 in [PSZ22])**.** *For each $L < \infty$, the space X_L is a quasigeodesic hyperbolic space. Moreover, Isom $X \leqslant$ Isom X_L.*

For a CAT(0) cube complex X, these X_L spaces were first formed using hyperplanes instead of curtains and were also hyperbolic [Gen19]. With a sufficiently large L, the authors of [IMZ23] use this hyperbolic space to continuously inject the sublinearly Morse boundary of X into the Gromov boundary of X_L. However, they used the additional assumption that X admits a *factor system* (see [BHS17]). This assumption makes these X_L spaces equal for all $L \geq L_0$, where L_0 is some constant dependent on the hierarchical structure of X. Without these X_L spaces stabilizing, it is possible for for κ-contracting geodesics to be of bounded diameter in each X_L for all L. This the main motivation for projecting to the *curtain model* (see Definition 6.1.1) in Theorem 1.3.1 instead of X_L for some L.

Definition 4.1.10 (Curtain Model)**.** Fix a sequence of number $\lambda_L \in (0,1)$ such that

$$\sum_{L=1}^{\infty} \lambda_L < \sum_{L=1}^{\infty} L\lambda_L < \sum_{L=1}^{\infty} L^2\lambda_L < \infty$$

We consider the space $(X, \hat{d})$, where the distance between two points $x, y \in X$ is defined by

$$\hat{d}(x,y) = \sum_{L=1}^{\infty} \lambda_L d_L(x,y)$$

and d_L is the L-metric defined in Definition 4.1.8. We call $(X, \hat{d})$ the *curtain model* of X and denote it as $\widehat{X}$.

4.2 Quasi-Isometry Invariance Counterexample

Petyt-Spriano-Zalloum asked in [PSZ22] if a quasi-isometry between CAT(0) spaces always induces a quasi-isometry between their corresponding curtain models. We answer this question in the negative. Both of the following definitions will also help in the construction of the counterexample.

Definition 4.2.1 (Angles in CAT(0) spaces, Section II.3.1 in [BH99])**.** Let X be a CAT(0) space and let $\alpha : [0,a] \to X$ and $\alpha' : [0,a'] \to X$ be two geodesic paths issuing from the same point $\alpha(0) = \alpha'(0)$. Then the comparison angle $\angle_{\mathbb{E}}(\alpha(t), \alpha'(t'))$ is a non-decreasing function of both $t, t' \geq 0$, and the *Alexandrov angle* $\angle(\alpha, \alpha')$ is equal to

$$\lim_{t,t' \to 0} \angle_{\mathbb{E}}\left(\alpha(t), \alpha'\left(t'\right)\right) = \lim_{t \to 0} \angle_{\mathbb{E}}\left(\alpha(t), \alpha'(t)\right).$$

Hence, we define:

$$\angle\left(\alpha,\alpha'\right)=\lim_{t\to 0}2\arcsin\frac{1}{2t}d\left(\alpha(t),\alpha'(t)\right).$$

Definition 4.2.2 (Strongly Contracting)**.** A geodesic α is D-strongly contracting if for any ball B disjoint from α we have $diam(\pi_\alpha(B))\le D$, where π_α is the closest point projection to α.

The following counterexample was used in [Cas16] to show that two quasi-isometric CAT(0) spaces can have contracting boundaries of different homeomorphism type when equipped with the Gromov product topology. We first introduce this space and its curtain model.

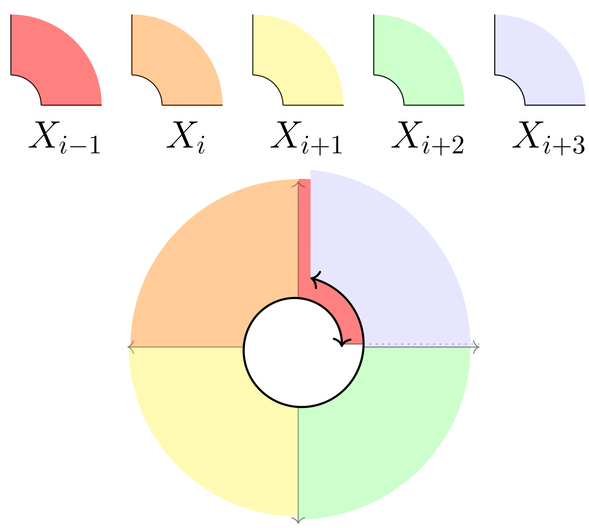

Figure 4.2: One level of the infinite parking lot $X=\cup_i X_i/\sim$. The space would continue to spiral upward and downward. Notice, since we are viewing the space from a birds eye view, X_{i-1} gets shadowed by X_{i+3}. This is due to X_{i+3} forming the "next level" of the infinite parking lot.

Let Y be $\mathbb{R}^2$ with a disc of radius one centered at the origin removed. Denote X as the universal cover of Y. We can view X in the following way: Take X_i to be a quarter flat with the quarter disc centered at the origin removed. Then $X = \cup_i X_i / \sim$ where $\sim$ denotes gluing the y-axis of X_i to the x-axis of X_{i+1} for all $i \in \mathbb{Z}$. One informally calls X the "infinite parking lot" as it can be viewed as a collection of quarter flats glued together that are spiraling up and down, giving the "infinite levels" of a parking lot. See Figure 4.2.

X is indeed a CAT(0) space since it is a gluing of CAT(0) spaces along single geodesic lines. The result of this space is that a half flat with a half disc of radius one removed at the origin can be isometrically embedded into each $X_i \cup X_{i+1}/\sim$. In fact, for any isometric spiral up or down by an angle θ, we get the same isometry of the half flat with a half disc removed at the origin. Parameterize X via its natural polar coordinates $\mathbb{R} \times [1, \infty)$, and define *the spiral* to be the line $\mathbb{R} \times \{1\}$. We now explain why X's curtain model $\widehat{X}$ is a quasi-line.

Take any geodesic ray γ such that $\gamma(0)$ is on the spiral and the Alexandrov angle between γ and the spiral is $\frac{\pi}{2}$. Up to an isometric rotation of X by some θ along the spiral, γ is the y-axis of some X_i. Since γ is the y-axis of some isometrically embedded half flat (with a half disc removed), all curtains dual to γ will stay in its half flat, $X_i \cup X_{i+1}/\sim$. As seen in Figure 4.3, if h_1, h_2 are two disjoint curtains dual to γ, then h_1, h_2 will be two parallel, infinitely long strips of width one in $X_i \cup X_{i+1}/\sim$. All curtains dual to the x-axis of X_i will meet h_1 and h_2, which means h_1 and h_2 are not L-separated for any L. The same is true for any two disjoint curtains dual to γ. Also, by Lemma 4.1.7, the max L-chain that can cross γ is bounded above by $4L + 10$. Thus, the diameter of γ is

$$\widehat{diam}(\gamma) = \sum_{L=1}^{\infty} \lambda_L diam_L(\gamma) \leq \sum_{L=1}^{\infty} \lambda_L (4L + 10) < \infty.$$

This is true for any geodesic ray that starts at the spiral and whose Alexandrov angle with the spiral is $\frac{\pi}{2}$. In particular, if we denote the spiral as α, then for any $x \in X$, $\hat{d}(x, \pi_\alpha(x)) \leq 4L + 10$.

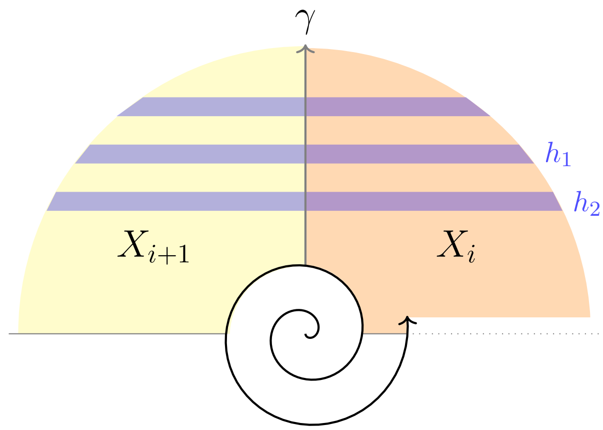

Figure 4.3: Since γ can always be seen as in the middle of a half flat, its dual curtains will behave like curtains in a half flat. This implies that any pair of curtains dual to γ will not be L-separated for any L.

Now, fix some origin $\mathfrak{o} \in X$ on α, and let α^+ denote the positive spiral direction and α^- the negative spiral direction emanating from $\mathfrak{o}$. Both directions are π-strongly contracting as balls disjoint from the axis can only project to half of the circumference of one of the circles in the spiral. By Theorem 1.4.1 for $\kappa \equiv 1$ (this is equivalent to Theorem 4.2 in [PSZ22]), there exists an infinite L-chain dual to α^+ for some L (similarly for α^-). Thus, in the curtain model $\widehat{X}$, the diameters of α^+ and α^- will both be unbounded. By [PSZ22, Proposition 9.5], both α^+ and α^- are unparameterized quasi-geodesics in $\widehat{X}$. This concludes α is a quasi-line in $\widehat{X}$. Since for any $x \in X$, $\hat{d}(x, \pi_\alpha(x)) \leq 4L + 10$, this yields that $\widehat{X}$ is a quasi-line.

Theorem 4.2.3. *Given X above, there exists a self quasi-isometry $\phi : X \longrightarrow X$ such that ϕ does not descend to a quasi-isometry for $\widehat{X}$. Further, there exists two quasi-isometric CAT(0) spaces W, Z whose curtain models $\widehat{W}, \widehat{Z}$ are not quasi-isometric.*

Proof. For some $\mathfrak{o} \in X$ on the spiral, denote the points of X by (θ, r), where θ is the angle traveled around the spiral starting at $\mathfrak{o}$, and r is the "radius" distance away from the spiral. Consider the points $(i, 2^i)$ and $(0, 2^i)$ for all $i \in \mathbb{N}$. Through a variation of the logarithmic spiral quasi-isometry of the Euclidean plane

$$\phi : X \longrightarrow X$$

$$(t, r) \longmapsto (t - \log_2(r), r),$$

we see that $\phi\left((i, 2^i)\right) = (0, 2^i)$. However, in the curtain model $\widehat{X}$, $\{(0, 2^i)\}_i$ represents a quasi-point, and $\{(i, 2^i)\}_i$ represents a quasi-line. This means that the self-quasi-isometry ϕ will not descend to a quasi-isometry for $\widehat{X}$.

Now, following the same vein as [Cas16], we construct two quasi-isometric CAT(0) spaces whose curtain models are not quasi-isometric. Construct the space W by gluing a geodesic ray γ_i to X at each $(i, 2^i)$ point. Similarly, construct the space Z by gluing a geodesic ray γ_i' to X at each $(0, 2^i)$ point. These spaces are quasi-isometric via the quasi-isometry

$$\overline{\phi} : W \longrightarrow Z$$

$$(t, r) \longmapsto (t - \log_2(r), r)$$

$$\gamma_i \longmapsto \gamma_i'.$$

However, the curtain models will not be quasi-isometric. See Figure 4.4. Indeed, as $\{(0, 2^i)\}_i$ is a quasi-point in $\widehat{Z}$, each of the geodesic rays in $\{\gamma_i'\}_i$ emanate from a point which is within bounded distance of $\mathfrak{o}$ on the quasi-line $\widehat{X}$. Thus, $\widehat{Z}$ is quasi-isometric to an infinite wedge of rays. On the other hand, $\{(i, 2^i)\}_i$ represents some sub-quasi-line in $\widehat{X}$, so the geodesic rays $\{\gamma_i\}_i$ have starting points at increasing distance away from $\mathfrak{o}$ in $\widehat{X}$ as i increases. So, $\widehat{W}$ is quasi-isometric to $\mathbb{R}$ with a ray attached to each positive integer. These two spaces are not quasi-isometric. □

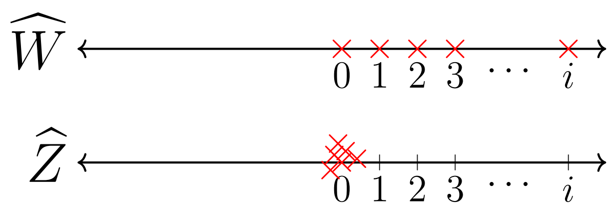

Figure 4.4: In this picture, the red marks represent the starting points of the geodesic rays that were glued to X. Notice that for $\widehat{W}$, there is a geodesic ray starting at each natural number. On the other hand, the geodesic rays in $\widehat{Z}$ are all clumped at the origin. The curtain models both crunch the flatness of the parking lot, but the geodesic rays still have recognizable distance at their starting points in $\widehat{W}$. Thus, the two curtain models are not quasi-isometric.

Corollary 4.2.4. *There exist quasi-isometric CAT(0) spaces W, Z whose curtain models have non-homeomorphic Gromov boundaries.*

Proof. The same logic can also apply to show $\widehat{W}$ and $\widehat{Z}$ have Gromov boundaries of different homeomorphism type. The sequence $\{\gamma_i\}_i$ in the Gromov boundary of $\widehat{W}$ converges to α^+.

No such converging sequence exists in $\widehat{Z}$. Thus, the Gromov boundaries for $\widehat{W}, \widehat{Z}$ are not homeomorphic. □

4.3 Quasi-flats are Bounded in the Curtain Model

Work Behrstock, Hagen, and Sisto show in [BHS21] that quasi-flats have finite diameter projections in top level curve graphs of mapping class groups. More generally, they also showed the same applies for quasi-flats in any hierarchically hyperbolic space projecting to its top level hyperbolic space. A *quasi-flat* is a quasi-isometric embedding of $\mathbb{R}^n$ for some n. In this section, we prove the same result as in [BHS21] for quasi-flats in CAT(0) spaces and their projections to the curtain model. Both arguments here and in [BHS21] leverage Huang's theory on orthants in [Hua17]. This is yet another result showing that the curtain model shows striking similarities to the top level curve graphs of hierarchically hyperbolic spaces. For this section, we set $\Lambda = \sum_{L=1}^{\infty} \frac{4L+10}{L^3}$.

Lemma 4.3.1. *$\widehat{\mathbb{R}^n_{\geq 0}}$ is bounded by $n\Lambda$.*

Proof. Let $n = 2$ and denote the origin as $\mathfrak{o}$. Let a be one of the two geodesic axes in $\mathbb{R}^2_{\geq 0}$. Let b be a geodesic ray perpendicular to a. Then, all of the curtains dual to b cross all of the curtains dual to a and vice versa. Thus, any pair of curtains h_1, h_2 dual to b will not be L-separated for any L. By Lemma 4.1.7, the max L-chain that can cross b is bounded above by $4L + 10$. Thus, we get a bound on the diameter of b:

$$\widehat{diam}(b) = \sum_{L=1}^{\infty} \frac{diam_L(b)}{L^3} \leq \sum_{L=1}^{\infty} \frac{4L+10}{L^3} = \Lambda.$$

Now, let $y \in \mathbb{R}^2_{\geq 0}$. Then y is on a geodesic that is perpendicular to some axis, a. Since a is perpendicular to the other axis of $\mathbb{R}^2_{\geq 0}$, and these axes connect at $\mathfrak{o}$, we have that $\hat{d}(\mathfrak{o}, y) \leq 2\Lambda$. This gives the proof for $n = 2$. For any $n > 2$, any $y \in \mathbb{R}^n$ is an element of a geodesic that is normal to some $n-1$-quarter flat that contains $n-1$ axes of $\mathbb{R}^n_{\geq 0}$. Thus, for similar reasons from the $n = 2$ case, y is of distance at most $4L + 10$ from the quarter flat. By a standard induction argument, $\widehat{\mathbb{R}^n_{\geq 0}}$ is bounded by $n\Lambda$. □

A *orthant* O of X is a convex subset which is isometric to the Cartesian product of finitely many half-lines $\mathbb{R}_{\geq 0}$. If O is both a subcomplex and an orthant, then O is called an *orthant subcomplex.*

Lemma 4.3.2. *Let X be a CAT(0) cube complex, O be an n-dimensional orthant contained in X, and $\phi : \mathbb{R}^n_{\geq 0} \longrightarrow O \subset X$ be the isometric embedding between $\mathbb{R}^n_{\geq 0}$ and O. Let a, b be geodesic rays in $\mathbb{R}^n_{\geq 0}$ and h_a, h_b be curtains dual to a, b with poles P_a, P_b, respectively. If h_1 and h_2 meet in $\mathbb{R}^n_{\geq 0}$, then the curtains $h_{\phi(a),\phi(P_a)}, h_{\phi(b),\phi(P_b)}$ meet in X*

Proof. Let $y \in h_a \cap h_b$. Then, $\pi_a(y) \in P_a$ and $\pi_b(y) \in P_b$. Since ϕ is an isometry, $\pi_{\phi(a)}(\phi(y)) \in \phi(P_a)$ and $\pi_{\phi(b)}(\phi(y)) \in \phi(P_b)$. Hence, $\phi(y) \in h_{\phi(a),\phi(P_a)} \cap h_{\phi(b),\phi(P_b)}$. □

Theorem 4.3.3 (Theorem 1.1 in [Hua17]). *If X is a CAT(0) cube complex of dimension n, then for every n-quasiflat Q in X, there is a finite collection $O_1, \ldots, O_k$ of n-dimensional orthant subcomplexes in X such that $d_H\left(Q, \cup_{i=1}^k O_k\right) < \infty$ where d_H denotes the Hausdorff distance.*

Theorem 4.3.4. *Let X be a CAT(0) cube complex and $Q \subset X$ be an n-dimensional quasi-flat. Then Q has bounded projection in $\widehat{X}$.*

Proof. Both Lemma 4.3.1 and Lemma 4.3.2 together tell us that each n-dimensional orthant in X has diameter bounded above by $n\Lambda$ in $\widehat{X}$. By Theorem 4.3.3, there is a collection of orthants $O_1, \dots, O_k$ in X so that $d_H\left(Q, \cup_{i=1}^k O_k\right) < \infty$. Each O_i is a quasi-point in $\widehat{X}$, so Theorem 4.3.3 tells us that Q is within finite Hausdorff distance of a finite collection of quasi-points in $\widehat{X}$. Let D be the largest $\hat{d}$ distance between two elements in $\cup_{i=1}^k O_k$. Then, for any $x, y \in Q$, $\hat{d}(x,y) \leq 2d_H\left(Q, \cup_{i=1}^k O_k\right) + D$. Thus, Q has bounded projection in $\widehat{X}$ □

4.4 Genevois's Hyperbolicity Criterion

Work of Genevois in [Gen16] has shown a hyperbolicity criterion for CAT(0) cube complexes. We work to generalize this criterion to the CAT(0) setting.

Definition 4.4.1 (Curtain Grid, E-thin)**.** Two chains of curtains $\mathcal{H} = \{h_1, \cdots, h_n\}$ and $\mathcal{K} = \{k_1, \cdots, k_m\}$ form a *curtain grid* if every curtain of $\mathcal{H}$ crosses every curtain of $\mathcal{K}$. We denote a curtain grid as $(\mathcal{H}, \mathcal{K})$. Given $E > 0$, a curtain grid is said to be *E-thin* if $\min\{|\mathcal{H}|, |\mathcal{K}|\} \leq E$.

Comparing to the cube complex setting, two chains of hyperplanes forming a grid will equate to a region of "flatness" with the intuition of larger grids equating to larger areas of "flatness". Thus, if one has an upper bound on how large these hyperplane grids can get, one could expect a notion of hyperbolicity. This is precisely what Genevois proves in [Gen16], and the following theorem uses curtains to get a similar criterion for the CAT(0) setting.

Theorem 4.4.2. *Let X be a CAT(0) space. Then X is hyperbolic if and only if every curtain grid is E-thin for some uniform $E > 0$.*

Proof. If X is a hyperbolic space, then all geodesics in X are uniformly D-contracting for some constant D. By Theorem C for $\kappa \equiv 1$ (this is equivalent to Theorem 4.2 in [PSZ22]), we have there exists an $L = L(D)$ such for any two curtains l_1, l_2 dual to the same geodesic that are also of distance at least L apart, we get that l_1 and l_2 are L-separated.

Consider any curtain grid $(\mathcal{H}, \mathcal{K})$. So $\mathcal{H} = \{h_1, \cdots, h_n\}$ for some n. Suppose $n > 5L + 8$. Let $x \in h_1^-$ and $y \in h_n^+$ and denote the unique geodesic between x and y as $[x, y]$. Each h_i crosses $[x, y]$ by nature of $\mathcal{H}$ being a chain. Since all of the curtains in $\mathcal{H}$ are disjoint, there exists $a_i \in [x, y] \cap h_i^+ \cap h_{i+1}^-$ for all i. Denote the curtains dual to $[x, y]$ and centered at a_i by l_i. We now consider the chain $\mathcal{L} = \{l_{L+3}, l_{2L+4}, l_{n-(2L+3)}, l_{n-(L+2)}\}$. Each curtain in $\mathcal{L}$ is distance at least L apart from the next curtain in $\mathcal{L}$. Thus, $\mathcal{L}$ is an L-chain. Notice the subchain of $\mathcal{H}$ that is $\{h_2, \cdots, h_{L+2}\}$ intersects nontrivially with l_{L+3}^-. Thus, h_1 cannot intersect l_{2L+4} or else we would contradict L-separation between l_{L+3} and l_{2L+4}. Similarly, h_n cannot intersect $l_{n-(2L+3)}$. See Figure 4.5. Thus, l_{2L+4} and $l_{n-(2L+3)}$ are L-separated curtains that both separate h_1 and h_n. Since all curtains in $\mathcal{K}$ cross both h_1 and h_n, they must also cross both l_{2L+4} and $l_{n-(2L+4)}$. The L-separability of l_{2L+4} and $l_{n-(2L+3)}$ implies that $|\mathcal{K}| \leq L \leq 5L+8$. Hence, all curtain grids are uniformly $(5L+8)$-thin.

For the reverse direction, let all curtain grids be E-thin for some uniform integer E. This means that any two curtains h_1, h_j dual to the same geodesic that are also greater than E distance away from each other must be E-separated. Indeed, such a situation would give a chain of $E + 1$ curtains $\{h_1, h_2, \cdots, h_j\}$ all dual to the same geodesic, so any grid made with this chain and some other chain $\mathcal{K}$ must give $|\mathcal{K}| \leq E$. Consider any $x, y \in X$.

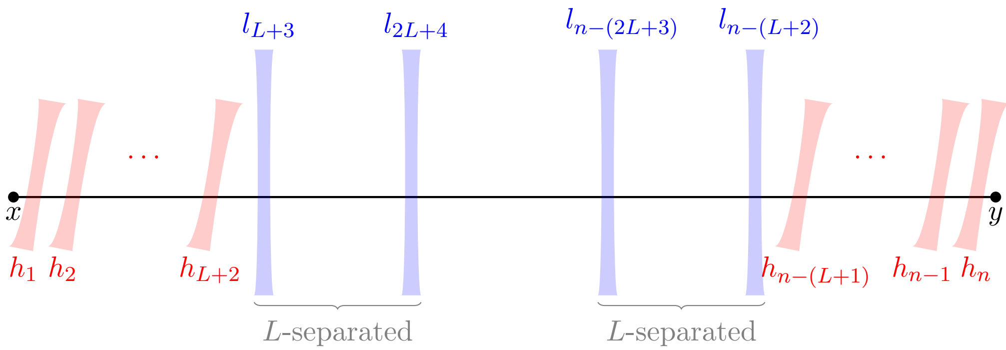

Figure 4.5: Since l_{L+3} and l_{2L+4} are L-separated and $\{h_2, \cdots, h_{L+2}\}$ is a chain of length L, we must have that h_1 cannot meet l_{2L+4}. Similarly, h_n cannot meet $l_{n-(2L+3)}$.

Then, there exists an $n \in \mathbb{Z}_{\geq 0}$ such that

$$n(E+2) \leq d(x,y) \leq (n+1)(E+2).$$

Thus, there exist a chain $c = \{h_1, \cdots, h_{n(E+2)-1}\}$ such that each h_i is dual to $[x,y]$. We can then conclude that $\{h_1, h_{(E+2)}, h_{2(E+2)}, \cdots, h_{(n-1)(E+2)}\}$ is an E-chain of length n. So $d_E(x,y) \geq n$. Hence, we get

$$d_E(x,y) \leq d(x,y) \leq d_E(x,y)(E+2) + (E+2).$$

This gives that X is quasi-isometric to X_E, a hyperbolic space. we conclude that X is hyperbolic. □

Chapter 5

Curtain Characterizations of Sublinearly Morse Geodesics

The point of this section is to give a characterization of κ-Morse rays in a more combinatorial light. Though sublinear Morseness is a useful definition, it can often be tough to prove things using this definition. In this chapter, we will be creating a equivalent notion of a κ-Morse geodesic ray that can be described solely on the curtains the ray crosses. This characterization was inspired from prior literature [CS15, MQZ22] that gave the same characterization in the CAT(0) cube complex setting via hyperplanes. For the remainder of the paper, we assume that X is a proper CAT(0) space.

5.1 Prior κ-Morse Characterizations and Notation

Since Chapter 3 has shown the contracting definition to be a useful characterization for Morse geodesics, we centrally use the following definition instead to characterize κ-Morse geodesics.

Definition 5.1.1 (κ-contracting geodesic, contracting constant)**.** A geodesic ray b is said to be κ-*contracting* if there exists a constant $C \geq 0$ such that for any $x \in X$ and any ball B centered at x with $B \cap b = \emptyset$, we have $\operatorname{diam}(\pi_b(B)) < C \cdot \kappa(x)$. We call the constant C the *contracting constant.*

Note, this is different than the definition of contracting given in Definition 3.2.4, as that definition requires $\operatorname{diam}(\pi_b(B)) < C \cdot \kappa(d(x, b))$ - the difference being that the upper bound in Definition 5.1.1 involves the distance from x to $\mathfrak{o}$ where the upper bound in Definition 3.2.4 involves the distance from x to b. However, with this new definition, we get the following characterization.

Theorem 5.1.2 (Theorem 3.8 in [QR22])**.** *Let X be a proper CAT(0) space. A geodesic ray b is κ-contracting if and only if it is κ-Morse.*

The following definition from [MQZ22] gives another characterization of κ-contracting rays that will be useful in the reverse direction of Theorem 5.2.2

Definition 5.1.3 (κ-slim geodesic)**.** We say an infinite geodesic ray b is κ-*slim* if there exists some $C \geq 0$ such that for any $x \in X$, $y \in b$, we have $d(\pi_b(x), [x, y]) \leq C\kappa(\pi_b(x))$.

Lemma 5.1.4 (Proposition 3.6/Corollary 3.7 in [MQZ22])**.** *In a CAT(0) space, a geodesic ray b is κ-contracting if and only if it is κ-slim.*

As seen in Figure 5.1, an updated version of Figure 3.6, There have been a lot of definitions attempting to categorize "hyperbolic" geodesics in non-hyperbolic spaces. In [QR22, IMZ23, MQZ22], the characterizations between κ-Morse, κ-contracting, and κ-slim were useful in applications to proofs in their respective papers. We introduce a new definition, *κ-curtain-excursion*(Definition 5.2.1), that is also inspired from [MQZ22]. Like similar characterizations before, we use our new definition and characterization in Chapter 6 to give continuous injection from $\partial_\kappa X$ to the curtain model of X.

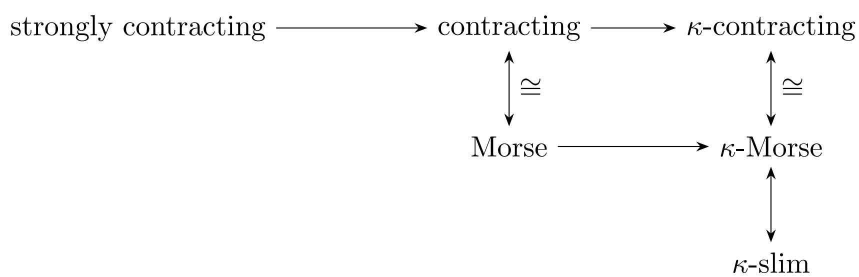

Figure 5.1: An update for the diagram given in Figure 3.6. In the CAT(0) setting, κ-Morse, κ-contracting, and κ-slim are equivalent definitions while still being our weakest definition for a "hyperbolic" geodesic in a non-hyperbolic space.

Notation 5.1.5. For the remainder of this paper, we will denote elements of $\partial_\kappa X$ as a^∞, b^∞, and so on, where a and b will be the unique geodesics based at $\mathfrak{o}$ that is in the equivalence class of a^∞ and b^∞, respectively. When relevant, quasi-geodesics in the same equivalence class as a^∞ and b^∞ will be denoted as α and β, respectively. We also often identify geodesics and quasi-geodesics with their images in X and treat them as subset of X when convenient. Lastly, we often use C and D to represent contracting and/or Morse constants.

5.2 A Characterization via Dual Curtains

Definition 5.2.1 (κ-chain, κ-curtain-excursion geodesic, excursion constant)**.** A *κ-chain* meeting some geodesic b is a chain of curtains $\{h_i\}$ meeting b at points $b(t_i)$ such that

- $t_{i+1} - t_i \leq C\kappa(t_{i+1})$
- h_i and h_{i+1} are $C\kappa(t_{i+1})$-separated

for some $C > 0$. If such $\{h_i\}$ are dual to b, then (up to a small increase in C) we choose $b(t_i)$ to be the centers of the poles of each h_i. A geodesic b is a *κ-curtain-excursion geodesic* when it is dual to a κ-chain. We refer to C as the *excursion constant.*

In the CAT(0) cube complex setting, a geodesic crossing a κ-chain of hyperplanes was defined as a κ-excursion geodesic in [MQZ22], so our above name of κ-curtain-excursion geodesics as a curtain analogue is fitting. This section works to prove Theorem 5.2.2, a dualized version of Theorem C in the introduction. Subsection 3.2 will recover all of Theorem C by working with non-dual chains.

Theorem 5.2.2. *A geodesic ray b is κ-contracting if and only if it is κ-curtain-excursion.*

The forward direction is Proposition 5.2.5 whereas the backward direction is Proposition 5.2.9. Before we prove the forward direction, we recall the following two lemmas.

Lemma 5.2.3 (Lemma 2.6 in [PSZ22])**.** *Let h be a curtain with pole P. For every $x \in h$, the geodesic $[x, \pi_P(x)]$ is contained in h. In particular, h is path-connected.*

Lemma 5.2.4 (Lemma 4.14 in [MQZ22])**.** *Assume X is a CAT(0) space. Let b be a κ-contracting geodesic ray with contracting constant D starting at $\mathfrak{o}$, and let $x, y \in X$ and* *not in b such that $d(\mathfrak{o}, \pi_b(x)) \leq d(\mathfrak{o}, \pi_b(y))$. If the projection of $[x, y]$ to b is larger than $4D\kappa(\pi_b(y))$, then $\pi_b([x, y]) \subseteq N_{5D\kappa(\pi_b(y))}([x, y])$.*

Lemma 5.2.3 tells us that for any curtain h and any $x \in h$, there will exist a geodesic in h connecting x to the pole of h, and Lemma 5.2.4 gives us that geodesics with large projections to a κ-contracting ray must get sublinearly close to the κ-contracting ray.

Proposition 5.2.5. *Let b be a κ-contracting ray with contracting constant D, then there exists $t_i \in \mathbb{R}$ such that b is dual to a κ-chain $\{h_i\}$ at points $b(t_i) \in h_i$ and*

- $t_{i+1} - t_i \leq C\kappa(t_{i+1})$
- *h_i and h_{i+1} are $C\kappa(t_{i+1})$-separated*

for some $C \geq 0$ depending only on D. In other words, b is a κ-curtain-excursion geodesic.

Proof. Since κ is a sublinear function, we can choose some t_0 such that $\kappa(t) \leq t$ for all $t \geq t_0$. For $i \in \mathbb{Z}_{\geq 0}$, choose t_{i+1} such that

$$t_{i+1} - t_i = 10D\kappa(t_{i+1}).$$

Note this is possible to do by the nature of κ being sublinear. Consider the chain $\{h_i = h_{b,i}\}$. What is left is to show $\{h_i\}$ have the second condition of a κ-chain. Let k be a curtain meeting both h_i and h_{i+1} and let P be its pole. Notice $\pi_b(P)$ has diameter less than 1 since π_b is 1−Lipschitz. Let $x \in h_i \cap k$, $y \in h_{i+1} \cap k$. There exists $x', y' \in P$ such that $[x, x'], [y, y'] \subset k$ by Lemma 5.2.3. Namely, $x' = \pi_P(x)$ and $y' = \pi_P(y)$. Now, the projection of the concatenation $[x, x'] * [x', y'] * [y', y]$ onto b will have diameter greater than $10D\kappa(t_{i+1}) - 1$. Since $\pi_b([x', y']) \subset \pi_b(P)$ has diameter less than 1, it must be that at least

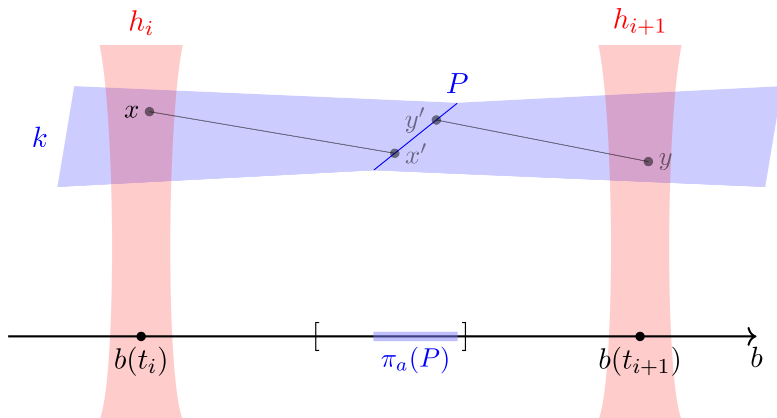

Figure 5.2: A picture of the argument of Proposition 5.2.5. We have h_i and h_{i+1} are curtains dual to b such that $d\left(b\left(t_i\right), b\left(t_{i+1}\right)\right) = 10D\kappa(t_{i+1})$, k is a curtain crossing both h_i and h_{i+1}, and P is the pole of k. Note that, no matter where P projects down to b, $\pi_b\left([x,x'] * [x',y'] * [y',y]\right)$ will contain $\left[b\left(t_i + \frac{1}{2}\right), b\left(t_{i+1} - \frac{1}{2}\right)\right]$ and will be of length at least $10\kappa(t_{i+1}) - 1$.

one of the projections of $[x, x']$ or $[y, y']$ has diameter greater than $4D\kappa(t_{i+1})$. See Figure 5.2.

Without loss of generality, we assume $d(\pi_b(x), \pi_b(x')) \geq 4D\kappa(t_{i+1})$. Thus, there exists a $p \in [t_i + 5D\kappa(t_{i+1}) - 1, t_i + 5D\kappa(t_{i+1}) + 1]$ such that $p \notin \pi_b(P)$, but also $p \in \pi_b([x, x'])$. By Lemma 5.2.4, we have that p is within $5D\kappa(t_{i+1})$ of $[x, x'] \subset k$. Hence, $b(t_i + 5D\kappa(t_{i+1}))$ is within $5D\kappa(t_{i+1}) + 1$ of k. Since this is true for any curtain k meeting both h_i and h_{i+1}, any chain that meets h_i and h_{i+1} must be bounded by $10D\kappa(t_{i+1}) + 3$. This gives the well-separation bound for κ-excursion. □

Remark 5.2.6. We have to double our bound since it is possible to have two chains of length $5D\kappa(t_{i+1}) + 1$ that are disjoint from each other. For example, in Figure 5.2 there can be one chain "above" b and another chain "below" b.

The above argument follows from the argument of [PSZ22, Theorem 4.2], which showed that a D-contracting ray has a $(10D+3)$-chain of curtains dual to the ray. The reverse direction follows arguments in [MQZ22]. The following lemma shows that certain curtains do not create bigons, and it will be used in Lemma 5.2.8.

Lemma 5.2.7 (Lemma 2.7 in [PSZ22])**.** *Let $b = [x_1, x_3]$ be a geodesic and let $x \notin b$. For any $x_2 \in b$, if h is a curtain dual to $[x_2, x]$ that meets $[x_1, x_2]$, then h does not meet $[x_2, x_3]$.*

Lemma 5.2.8. *Let b be a κ-curtain-excursion geodesic ray with excursion constant $C > 0$ and κ-chain $\{h_i\}$. If a is another geodesic and crosses h_{i-1}, h_i, h_{i+1} with $\pi_b(a(s_i)) = b(t_i)$ for some $s_i > 0$, then there exists a $D > 0$ depending only on C such that $d(a(s_i), b(t_i)) \leq D\kappa(t_i)$.*

Proof. The following argument is illustrated in Figure 5.3. Let $s_{i-1} < s_i < s_{i+1}$ such that $\pi_b(a(s_j)) = b(t_j)$ for $j \in \{i-1, i, i+1\}$. Consider geodesics $[a(s_i), a(s_{i+1})]$, $[a(s_{i+1}), b(t_{i+1})]$, $[b(t_i), b(t_{i+1})]$, and $[a(s_i), b(t_i)]$. This is a quadrilateral in our space. Let c be a maximal chain dual to $[a(s_i), b(t_i)]$. Note that all curtains in c must cross at least one of $[a(s_i), a(s_{i+1})]$, $[a(s_{i+1}), b(t_{i+1})]$, or $[b(t_i), b(t_{i+1})]$. Thus, denote c_1, c_2, and c_3 as the collections of curtains in c as seen in Figure 5.3. That is, c_1 are the curtains in c that also meet $[b(t_i), b(t_{i+1})]$ and so on. We get $|c| \leq |c_1| + |c_2| + |c_3|$. Also, define c' as the curtains in c_3 that also meet $[a(s_{i-1}), b(t_{i-1})]$ and c'' as the curtains of c_3 that also meet $[b(t_{i-1}), b(t_{i+1})]$. No curtains in c_3 meet $[a(s_{i-1}), a(s_i)]$ by Lemma 5.2.7. We have that $|c_3| \leq |c'| + |c''|$.

Now, $|c_1| \leq C\kappa(t_{i+1})$ since, by assumption of the κ-curtain-excursion chain, the length of $[b(t_i), b(t_{i+1})]$ is bounded by $C\kappa(t_{i+1})$. Also, we have $|c_2| \leq C\kappa(t_{i+1})$ because h_i

and h_{i+1} are $C\kappa(t_{i+1})$-separated. Similarly, we have $|c'| \leq C\kappa(t_i)$, and $|c''| \leq C\kappa(t_i)$. Thus,

$$
\begin{aligned}
|c| &\leq |c_1| + |c_2| + |c_3| \\
&\leq |c_1| + |c_2| + |c'| + |c''| \\
&\leq C\kappa(t_{i+1}) + C\kappa(t_{i+1}) + C\kappa(t_i) + C\kappa(t_i) \\
&\leq 4C\kappa(t_{i+1}).
\end{aligned}
$$

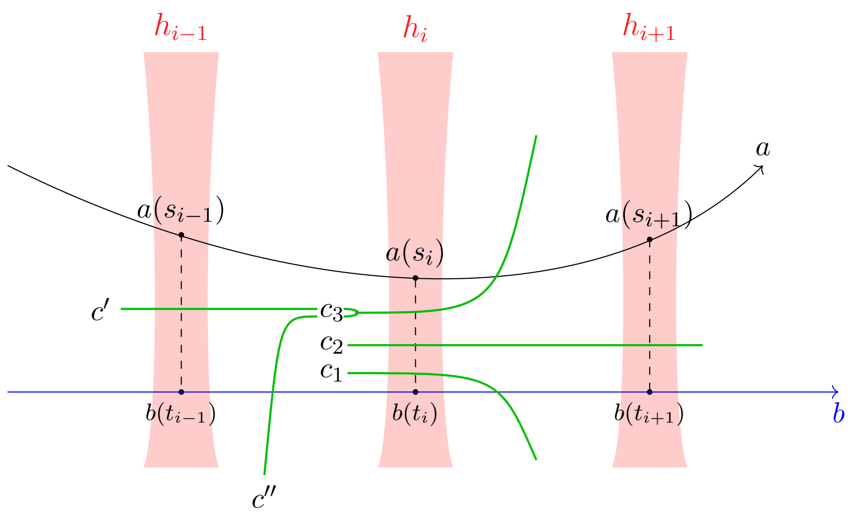

Figure 5.3: The set up of Lemma 5.2.8 with subchains c_1, c_2, and c_3 included. The bounds of c_1, c_2 and c_3 will show that $d(a(s_i), b(t_i)) \leq D\kappa(t_i)$ for some D depending on C.

Since $|c|$ is a maximal chain of $[a(s_i), b(t_i)]$, we have that $a(s_i)$ is within a distance of $(4C\kappa(t_{i+1})+1)$ from $b(t_i)$. By Lemma 3.3.2, we have that there exists a $D \geq 0$ such that $d(a(s_i), b(t_i)) \leq D\kappa(t_i)$. □

Proposition 5.2.9. *Let b be a κ-curtain-excursion geodesic with κ-chain $\{h_i\}$ dual to b. Denote $b(t_i)$ as the centers of the poles of each h_i, and put $C \geq 0$ the excursion constant.*

Then b is κ-slim for constant $D' \geq 0$, depending only on C (and κ-contracting by Lemma 5.1.4).

Proof. Let $x \in X$. Then, there exists a minimal i such that $x \in h_i^-$. Let $y \in b$. We split where y can be placed on b into three cases: $y \in h_{i+2}^+$, $y \in h_{i-4}^-$ and $y \in [b(t_{i-4} - \frac{1}{2}), b(t_{i+2} + \frac{1}{2})]$.

When $y \in h_{i+2}^+$, take $[x, y]$ which will cross h_i, h_{i+1} and h_{i+2}. By Lemma 5.2.8, we have $b(t_{i+1})$ is within $D\kappa(t_{i+1})$ of $[x, y]$ for some $D > 0$ depending only on C. We know $\pi_b(x) \in h_{i-2}^+$. Thus, $d(\pi_b(x), b(t_{i+1})) \leq 3C\kappa(t_{i+1})$. This gives

$$d(\pi_b(x), [x, y]) \leq (3C + D)\kappa(t_{i+1}).$$

Since $\pi_b(x)$ is on the geodesic $[b(t_{i-2}), b(t_i)]$, we get that Lemma 3.3.2 implies there exists a D' such that $d(\pi_b(x), [x, y]) \leq D'\kappa(\pi_b(x))$. See Figure 5.4.

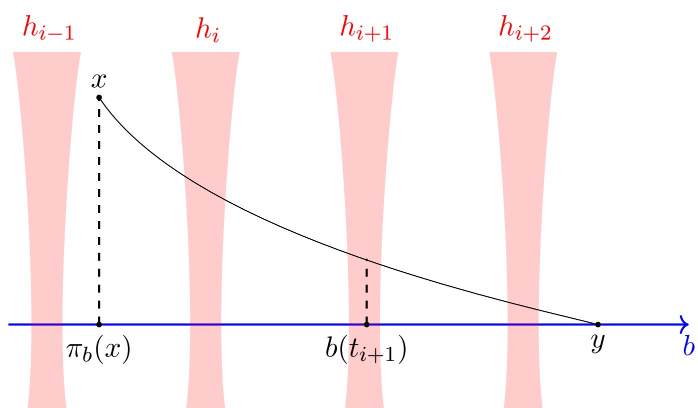

Figure 5.4: Picture of the first case of Proposition 5.2.9. Since $d(b(t_{i+1}), [x, y]) \leq D\kappa(t_{i+1})$ and $d(\pi_b(x), b(t_{i+1})) \leq 3C\kappa(t_{i+1})$, we see that the $\pi_b(x)$ will also be sublinearly close to $[x, y]$.

The proof in the case of $y \in h_{i-4}^-$ is the same and its corresponding picture will be a mirrored version of Figure 5.4.

Lastly, when $y \in [b(t_{i-4} - \frac{1}{2}), b(t_{i+2} + \frac{1}{2})]$, we similarly have that $\pi_b(x) \in h_{i-2}^+ \cap h_i^-$ will be within $5C\kappa(t_{i+2})$ of y. Thus, again by Lemma 3.3.2, there will exist a D' depending only on C such that $d(\pi_b(x), [x, y]) \leq D'\kappa(\pi_b(x))$. Hence the proof (and also the proof of Theorem 5.2.2). □

5.3 Dualizing a κ-Chain

In [PSZ22], the authors shows that if a chain of L-separated curtains ***meets*** a geodesic, one can follow the process of [PSZ22, Lemma 4.5] to find a chain of L-separated curtains dual to the geodesic. If one were to follow this process in the sublinear case, it is likely that the dual curtains will be at a κ^2 distance apart resulting in a dual κ^2-chain instead of a κ-chain. Instead, we rework Lemma 5.2.8 and Proposition 5.2.9 to allow for a κ-chain that is not necessarily dual to the geodesic. This is Proposition 5.3.1 and Proposition 5.3.2, respectively. Then, Proposition 5.2.5 finds a κ-chain dual to the geodesic.

Proposition 5.3.1. *Let b be a geodesic ray that meets a κ-chain $\{h_i\}$ (not necessarily dual) at points $b(t_i)$ with excursion constant $C > 0$. Put P_i as the poles of each h_i. If a is another geodesic and crosses h_{i-1}, h_i, h_{i+1} with $\pi_{P_i}(a(s_i)) = \pi_{P_i}(b(t_i))$ for some $s_i > 0$, then there exists a $D > 0$ depending only on C such that $d(a(s_i), b(t_i)) \leq D\kappa(t_i)$.*

Proof. Again, similar to Lemma 5.2.8, consider the geodesic $[a(s_i), b(t_i)]$ and let c be a maximal chain dual to $[a(s_i), b(t_i)]$. Since $\pi_{P_i}(a(s_i)) = \pi_{P_i}(b(t_i))$, then due to Lemma 5.2.3, the concatenation $[a(s_i), \pi_{P_i}(b(t_i))] * [\pi_{P_i}(b(t_i)), b(t_i)]$ will be inside of h_i. Since any curtain in the chain c must meet $[a(s_i), \pi_{P_i}(b(t_i))] * [\pi_{P_i}(b(t_i)), b(t_i)]$, we must have that any curtain in c must meet h_i (even though $[a(s_i), b(t_i)]$ might not stay inside of h_i). On the other hand, all curtains in c must also meet the concatenation $[a(s_i), a(s_{i+1})] * [a(s_{i+1}), b(t_{i+1})] * [b(t_{i+1}), b(t_i)]$. So, like Lemma 5.2.8, one gets that $d(a(s_i), b(t_i) \leq D\kappa(t_i)$ where D depends only on C. □

Proposition 5.3.2. *Let b be geodesic that meets a κ-chain $\{h_i\}$ (not necessarily dual) at points $b(t_i)$ with excursion constant $C > 0$. Then b is κ-slim for constant $D' \geq 0$, depending only on C (and κ-contracting by Lemma 5.1.4).*

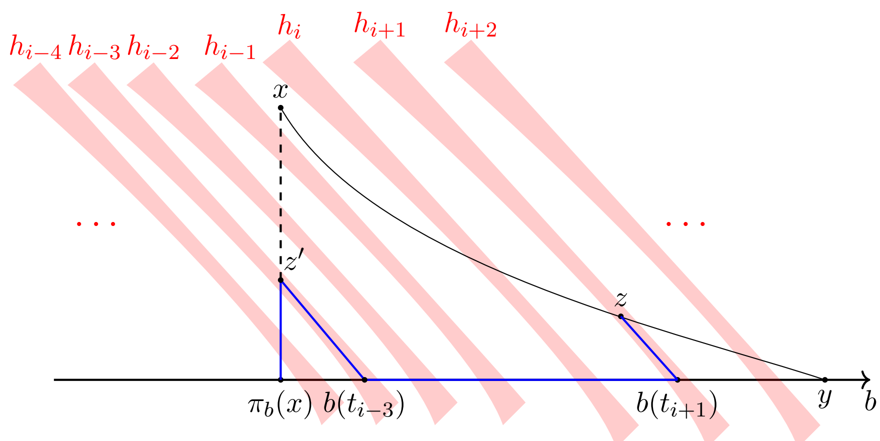

Figure 5.5: A picture for the proof of Proposition 5.3.2. No matter how many curtains in our κ-chain meet $[x, \pi_b(x)]$, we get that z' and $b(t_{i-3})$ will still be bounded by κ. The blue path represents a path from $\pi_b(x)$ to z, and all geodesic subpaths of the blue path are bounded by κ via the proof of Proposition 5.3.2.

Proof. Again we follow arguments in Proposition 5.2.9. Let $x \in X$ and $y \in b$. There exists a minimal i such that $x \in h_i^-$. To prove b is κ-slim, we again split where $y \in b$ can be placed on b into three cases: $y \in h_{i+2}^+$, $y \in h_{i-4}^-$, and $y \in [b(t_{i-4} - \frac{1}{2}), b(t_{i+2} + \frac{1}{2})]$.

When $y \in h^+_{i+2}$, we get by Proposition 5.3.1 that there exists $z \in [x, y] \cap h_{i+1}$ such that $d(z, b(t_{i+1})) \leq D\kappa(t_{i+1})$ where D depends on C. Now, since the $\{h_i\}$ are not necessarily dual to b, it is possible that $[x, \pi_b(x)]$ will cross elements in our κ-chain $\{h_i\}$. This gives two subcases. In the subcase that $\pi_b(x) \in h^+_{i-5}$ Then $d(\pi_b(x), b(t_{i+1}) \leq 5C\kappa(t_{i+1})$ and so $d(\pi_b(x), z) \leq (5C + D)\kappa(t_{i+1}) \leq D'\kappa(\pi_b(x))$ for some D' depending on C by Lemma 3.3.2. In the subcase that $\pi_b(x) \notin h^+_{i-5}$, then $[x, \pi_b(x)]$ crosses $h_{i-4}, h_{i-3}, h_{i-2}$ and, by Proposition 5.3.1, there exists a $z' \in [x, \pi_b(x)] \cap h_{i-3}$ such that $d(z', b(t_{i-3}) \leq D\kappa(t_{i-3})$ for some D depending on C. Since $\pi_b(z') = \pi_b(x)$, it must be true that also $d(z', \pi_b(x)) \leq D\kappa(t_{i-3})$. Thus, we get,

$$\begin{aligned}
d(\pi(x), z) &\leq d(\pi(x), z') + d(z', b(t_{i-3})) + d(b(t_{i-3}), b(t_{i+1})) + d(b(t_{i+1}), z) \\
&\leq D\kappa(t_{i-3}) + D\kappa(t_{i-3}) + 4C\kappa(t_{i+1}) + C\kappa(t_{i+1}) \\
&\leq (2D + 5C)\kappa(t_{i+1}).
\end{aligned}$$

This gives, by Lemma 3.3.2, there exists a D'' such that $d(\pi(x), z) \leq D''\kappa(\pi_b(x))$. The cases of $y \in h^-_{i-4}$ and $y \in [b(t_{i-4} - \frac{1}{2}), b(t_{i+2} + \frac{1}{2})]$ are done with a similar argument of the above case and Lemma 5.2.9, so we will omit them. See Figure 5.5. □

Corollary 5.3.3. *If a κ-chain $\{h_i\}$ meets a geodesic b at points $\{b(t_i)\}$ such that $t_{i+1} - t_i \leq C\kappa(t_{i+1})$ for some $C > 0$, then b is a κ-curtain-excursion geodesic.*

Proof. Proposition 5.3.2 shows b is κ-contracting, and Proposition 5.2.5 finds a dual κ-chain to b. □

Chapter 6

Genericity of Sublinear Morse Geodesics in the Curtain Model

6.1 The Behavior of Projected Geodesics in the Curtain Model

We now reintroduce the curtain model, a hyperbolic space with a Gromov boundary that $\partial_\kappa X$ can continuously inject into as formulated in [PSZ22]. As a set, we fix the CAT(0) space X but change the metric.

Definition 6.1.1 (Curtain Model)**.** We consider the space $(X, \hat{d})$ where the distance between two points $x, y \in X$ is defined by

$$\hat{d}(x, y) = \sum_{L=1}^{\infty} \frac{d_L(x, y)}{L^3}$$

where d_L is the L-metric defined in Definition 4.1.8. We call $(X, \hat{d})$ the *curtain model* of X and denote it as $\widehat{X}$ as seen in the introduction.

Remark 6.1.2. In [PSZ22] and Chapter 4, the curtain model metric used was $D(x,y) = \sum_{L=1}^{\infty} \lambda_L d_L(x,y)$ where $\lambda_L \in (0,1)$ such that $\sum_{L=1}^{\infty} \lambda_L < \sum_{L=1}^{\infty} L\lambda_L < \sum_{L=1}^{\infty} L^2\lambda_L < \infty$. By [PSZ22, Remark 9.1], the condition that $\sum_{L=1}^{\infty} L^2\lambda_L < \infty$ was only used in Section 9.3 of [PSZ22], and none of Chapter 6 uses this content. Thus, the sequence $\lambda_L = \frac{1}{L^3}$ is a sufficient sequence to use for this paper. It is also a desirable sequence since choosing $\lambda_L = \frac{1}{L^3}$ allows our theory to work when regarding all κ such that κ^4 is sublinear. This is the largest class of sublinear functions our proofs can work for — if one uses $\lambda'_L = \frac{1}{L^n}$ for any $n > 2$, we would require κ to be so that κ^{n+1} is sublinear. (See the proof of Lemma 6.1.9 to see this relationship.) This observation suggests that choosing $\lambda_L = \frac{1}{L^3}$ might be the preferred choice in the majority of contexts. Intuitively, it makes sense to choose the slowest converging sequence possible because the slower our sequence of λ_L's converges to 0, the more geometry each d_L sees for larger L. Given the conditions needed for λ_L, our choice of sequence is essentially the slowest converging sequence to 0 that fits the necessary conditions.

The function $\hat{d}$ is indeed a metric since each of the d_L's are metrics. Also, every isometry of X is also an isometry of X_L by Theorem 4.1.9. This implies also that any isometry of X is an isometry of $\widehat{X}$. Hence Isom X acts on $\widehat{X}$ by isometries [PSZ22, Lemma 9.2].

Theorem 6.1.3 (Theorem 9.10 in [PSZ22])**.** *There exists a δ such that $\widehat{X}$ is a quasigeodesic δ-hyperbolic space.*

Since $\widehat{X}$ is not a proper metric space, we cannot define its boundary via equivalence classes of geodesic directions as we do for $\partial_\kappa X$ or in Chapter 3. Rather, we will define the Gromov-boundary of $\widehat{X}$ in terms of equivalence classes of sequences. For more of the following information, see [BK02] or [BH99]. Recall that a sequence $\{x_n\}$ *Gromov-converges to infinity* if $\liminf_{i,j\to\infty} (x_i \cdot x_j)_{\mathfrak{o}} = \infty$ where $(x_i \cdot x_j)_{\mathfrak{o}}$ is the *Gromov product*, defined as

$$(x_i \cdot x_j)_{\mathfrak{o}} = \frac{1}{2}\Big(\hat{d}\,(\mathfrak{o}, x_i) + \hat{d}\,(\mathfrak{o}, x_j) - \hat{d}\,(x_i, x_j)\Big).$$

Two sequences $\{x_i\}$ and $\{y_j\}$ are said to be in the same equivalence class if

$$\liminf_{i,j\to\infty} (x_i \cdot y_j)_{\mathfrak{o}} = \infty.$$

Thus, we can define the *Gromov boundary* $\partial\widehat{X}$ as the set of equivalence classes of sequences Gromov-converging to infinity. Given $s \in \partial\widehat{X}$ and $r > 0$ the sets

$$U(s,r) = \left\{ [\{y_i\}] \;\middle|\; \liminf_{i,j\to\infty} (x_i \cdot y_j)_{\mathfrak{o}} \geq r \text{ for some } \{x_i\} \in s \right\}$$

form a basis for the standard topology on the Gromov boundary $\partial\widehat{X}$. Furthermore, isometries of $\widehat{X}$ induce homeomorphisms of $\partial\widehat{X}$ with respect to this topology.

We desire to create a map $\partial_\kappa X \longrightarrow \partial\widehat{X}$ that sends $b^\infty \in \partial_\kappa X$ to some $s \in \partial\widehat{X}$. Ideally, the geodesic b will project to some unbounded set in $\widehat{X}$, and increasing sequences of points on b will Gromov-converge to infinity in $\partial\widehat{X}$. However, it is not obvious that b will always be unbounded in $\widehat{X}$. A potential problem is the strength of our sublinear function κ, as the following example will show.

Example 6.1.4. Let $\kappa = \sqrt{\cdot}$, i.e. the square root function, and consider the following space:

Put $X_{\sqrt{\cdot}}$ to be the subset of $\mathbb{R}^2$ between the square root function and the x-axis. Put $X_i = \left[i^2 + \frac{1}{2}, (i+1)^2 - \frac{1}{2}\right] \times [0, \infty)$. Now denote $X = X_{\sqrt{\cdot}} \cup \bigcup_{i=1}^{\infty} X_i$. Notice that, when b is the x-axis, b is κ-contracting, and a κ-chain dual to b is $\{h_{b,t_i}\}$ where $b(t_i)$ is the point $\left(i^2, 0\right)$. See Figure 6.1, where X is the blue-shaded space.

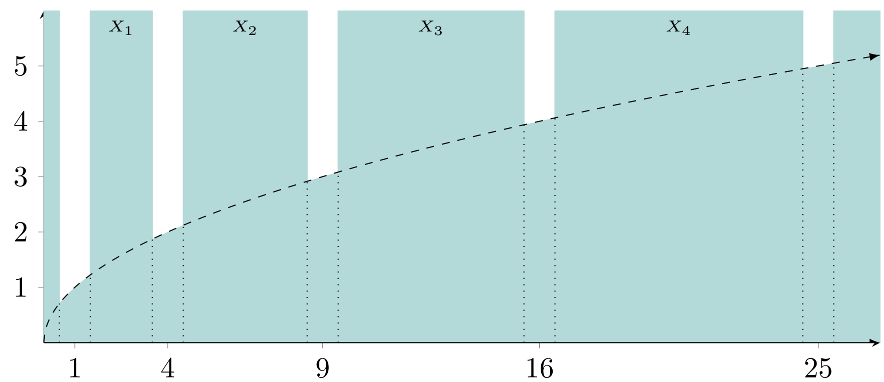

Figure 6.1: The space $X = X_{\sqrt{\cdot}} \cup \bigcup_{i=1}^{\infty} X_i$

In this constructed space, any L-chain dual to b must only have one curtain in each X_i. This is because any two curtains in X_i dual to b are not L-separated for any L. Though it is possible to find a larger L-chain such that its curtains are not dual to b, the remainder of Chapter 6 uses dual chains as a lower bound for distance. By observation we have that, for any t with $t_i \leq t \leq t_{i+1}$, the largest L-chain dual to $[\mathfrak{o}, b(t)]$ is $2L$ if $L \leq i$ and $2i$ if $L > i$.

Thus, as $i \to \infty$, we have

$$\lim_{i \to \infty} \hat{d}\left(\mathfrak{o}, b\left(t_i\right)\right) = \lim_{i \to \infty} \sum_{L=1}^{\infty} \frac{d_L\left(\mathfrak{o}, b\left(t_i\right)\right)}{L^3} \geq \sum_{L=1}^{\infty} \frac{2L}{L^3}.$$

Notice how the right hand side of the inequality is a finite value. This argument is not sufficient to show that $diam(b)$ is unbounded in $\widehat{X}$. Thus, it is possible b would not define a point in $\partial \widehat{X}$. It may still be possible to find collections of L-chains meeting b that show b has infinite diameter in $\widehat{X}$, but finding such L-chains will be tougher. Since most applications of sublinear functions involve a κ such that κ^4 is sublinear, such as logarithmic functions, we can afford to impose this stronger condition for cleaner arguments. Thus, for the remainder of the paper, we assume κ is any continuous, bijective, and monotonically increasing function such that κ^4 is sublinear.

The following lemmas show that geodesics in X project to unparameterized quasi-geodesics in $\widehat{X}$. Lemmas 6.1.5 and 6.1.6 assist in the proof of Lemma 6.1.7. Lemma 6.1.7 shows us that projections behave like rough geodesics. It should be noted that Lemma 6.1.7 shows the same result as [PSZ22, Propostion 9.5] but for when $\sum_{L=1}^{\infty} \lambda_L = \sum_{L=1}^{\infty} \frac{1}{L^3}$. Since the proof of Lemma 6.1.7 was written before the the updated version of [PSZ22], we still include Lemma 6.1.7 for completeness.

Lemma 6.1.5 (Lemma 2.13 in [PSZ22])**.** *Suppose that* $c = \{h_1, \dots, h_n\}$ *and* $c' = \{h'_1, \dots h'_m\}$ *are* L *chains with* $n > 1$ *and* $m > L+1$*, where* $h_1'^{-}$ *is the halfspace not containing* h'_2*. If* $h_1'^{-} \cap h_j \neq \varnothing$ *for all* j *and* $h_n^+ \cap h'_i \neq \emptyset$ *for all* i*, then* $c'' = \left\{h_1, \dots, h_{n-1}, h'_{L+2}, \dots, h'_m\right\}$ *is an* L*-chain of cardinality* $n + m - L - 2$*.*

Lemma 6.1.6 (Corollary 3.2 in [PSZ22])**.** *If* b *is a CAT(0) geodesic and* $t_1 < t_2 < t_3$*, then any* L*-chain* c *separating* $b(t_2)$ *from* $\{b(t_1), b(t_3)\}$ *has cardinality at most* $L' = 1 + \left\lfloor \frac{L}{2} \right\rfloor$*.*

Lemma 6.1.7. *Let* $b : I \to X$ *be a geodesic ray in* (X, d)*, then for any* $t_1, t_2, t_3 \in I$ *such that* $t_1 < t_2 < t_3$ *with* $b(t_1) = x, b(t_2) = y, b(t_3) = z$*, we have* $\hat{d}(x, z) \geq \hat{d}(x, y) + \hat{d}(x, z) - C$ *where* C *is a constant independent of* b*.*

Proof. Let $L \in \mathbb{N}$, and let c, c' be maximal L-chains realizing $d_L(x,y) = 1+|c|$ and $d_L(y,z) = 1+|c'|$. Assume $|c'| \geq \frac{3L}{2} + 3$. By Lemma 6.1.6, the maximum number of curtains in c' that cross $[x,y]$ is $1 + \lfloor \frac{L}{2} \rfloor$. Denote $c'' \subseteq c'$ the set with such curtains deleted so that $|c''| \geq |c'| - 1 - \lfloor \frac{L}{2} \rfloor$. Ordering $c'' = \{h_1, h_2 \ldots, h_n\}$, we denote h_1^+ as the halfspace containing h_2. Notice that h_1^- must contain $[x,y]$, which means every curtain in c meets h_1^-. Thus, by Lemma 6.1.5, $d_L(x,z) \geq |c| + (|c'| - 1 - \lfloor \frac{L}{2} \rfloor) - (L+2) \geq d_L(x,y) + d_L(y,z) - \frac{3L}{2} - 5$. Note that this inequality is also trivially true if $|c'| \leq \frac{3L}{2} + 3$. Thus, we get that

$$\begin{aligned} \hat{d}(x,z) &= \sum_{L=1}^{\infty} \frac{d_L(x,z)}{L^3} \\ &\geq \sum_{L=1}^{\infty} \frac{d_L(x,y) + d_L(y,z) - \frac{3L}{2} - 5}{L^3} \\ &= \hat{d}(x,y) + \hat{d}(y,z) - \sum_{L=1}^{\infty} (\frac{3}{2L^2} + \frac{5}{L^3}). \end{aligned}$$

Setting $C = \sum_{L=1}^{\infty} (\frac{3}{2L^2} + \frac{5}{L^3})$ gives the desired result. □

Remark 6.1.8. Lemma 6.1.7 shows a "coarse reverse triangle inequality" for geodesics projecting into $\widehat{X}$. Looking at the the definition of an unparameterized quasi geodesic (see [MMS12, Section 2.1]), one can use Lemma 6.1.7 to parameterize a projected geodesic into a quasi-geodesic in $\widehat{X}$.

In order the show that any κ-contracting geodesic defines a point in $\partial \widehat{X}$, we must also show that the geodesic will have infinite diameter with respect to $\hat{d}$.

Lemma 6.1.9. *Let X be a CAT(0) space and b be a κ-contracting geodesic ray based at $\mathfrak{o}$ where κ is a sublinear function such that κ^4 is sublinear. Then b has infinite diameter with respect to the $\hat{d}$ metric.*

Proof. Denote $\{h_i\}$ the κ-chain dual to b and denote $b(t_i)$ the center of the poles of each h_i. Put $C > 0$ as the excursion constant. We know that $\lfloor C\kappa(t_1) \rfloor = n$ for some $n \in \mathbb{N}$. Thus, for any $i > 1$, there exists some $m \geq n$ such that $\lceil C\kappa(t_i) \rceil = m$. Since, by definition of a κ-chain, $t_{j+1} - t_j \leq C\kappa(t_{j+1})$, we have

$$\begin{aligned} t_i - t_1 &= \sum_{j=1}^{i} (t_{j+1} - t_j) \\ &\leq \sum_{j=1}^{i} C\kappa(t_{j+1}) \\ &\leq \sum_{j=1}^{i} C\kappa(t_i) \\ &= iC\kappa(t_i) \\ &\leq im. \end{aligned}$$

Thus, we have $i \geq \frac{t_i - t_1}{m}$. Notice, since $\lceil C\kappa(t_i) \rceil = m$, this implies $C\kappa(t_i) + 1 \geq m$. Now, we claim $d_m(\mathfrak{o}, b(t_i)) \geq i - 1 \geq \dfrac{t_i - t_1}{m} - 1$. Indeed, $[\mathfrak{o}, b(t_i)]$ crosses curtains $\{h_1, h_2, \ldots, h_{i-1}\}$, and since each h_{j-1} and h_j are $C\kappa(t_j)$-separated, they are also m-separated since $C\kappa(t_j) \leq m$ for all $j \leq i - 1$. Thus, it follows that

$$\frac{d_m(\mathfrak{o}, b(t_i))}{m^3} \geq \frac{\frac{t_i - t_1}{m} - 1}{m^3} \geq \frac{t_i - t_1}{m^4} - 1 \geq \frac{t_i - t_1}{(C\kappa(t_i) + 1)^4} - 1.$$

By assumption, κ^4 is a sublinear function. Hence, as i grows to infinity, the right hand side of the above inequality grows to infinity. Since $\hat{d}(\mathfrak{o}, b(t_i)) \geq \frac{d_m(\mathfrak{o}, b(t_i))}{m^3}$, we conclude that

$$\lim_{i \to \infty} \hat{d}(\mathfrak{o}, b(t_i)) = \infty.$$

Hence, b is unbounded in $\widehat{X}$. □

Remark 6.1.10. Notice, in the above proof, we showed that for each t_i, we can find an $m \in \mathbb{N}$ such that

$$\frac{d_m(\mathfrak{o}, b(t_i))}{m^3} \geq \frac{t_i - t_1}{(c\kappa(t_i) + 1)^4} - 1.$$

This inequality also shows the possibility of κ-contracting rays having a κ-*persistent shadow* as defined in [DZ22]. We show a notion of equivalence between a ray being κ-contracting and having a κ-persistent shadow in Section 6.4, and we postpone further conversation to that section.

6.2 Creating an Injective Map

Assume X is a proper CAT(0) space with $\widehat{X}$ its curtain model and let κ be a sublinear function such that κ^4 is sublinear. We have that κ-contracting geodesics in X will be unbounded quasi-geodesics in $\widehat{X}$. Since each $b^\infty \in \partial_\kappa X$ has only one geodesic emanating from $\mathfrak{o}$ in its equivalence class, we define φ by

$$\begin{aligned} \varphi : \partial_\kappa X &\longrightarrow \partial\widehat{X} \\ b^\infty &\longmapsto \left[\{b(n)\}_{n\in\mathbb{N}}\right]. \end{aligned}$$

Notably, the map is well defined since b^∞ only has one geodesic emanating from $\mathfrak{o}$ (namely, b). Also, given any increasing sequence $\{t_i\}$ with $t_i \to \infty$, we get that $\{b(t_i)\}$ Gromov-converges to infinity and $[\{b(t_i)\}] = [\{b(n)\}]$. This fact is useful because we can choose any increasing sequence $\{b(t_i)\} \subset b$ that Gromov-converges to infinity to represent $\varphi(b^\infty)$ in our proofs. We now work to make φ injective.

Lemma 6.2.1. *Let b_1 be a κ-contracting geodesic ray. Denote $\{h_i\}$ the κ-chain dual to b_1 with center of poles $b_1(t_i)$. If a geodesic b_2 meets infinitely many of $\{h_i\}$, then b_2 is in a κ-neighborhood of b_1.*

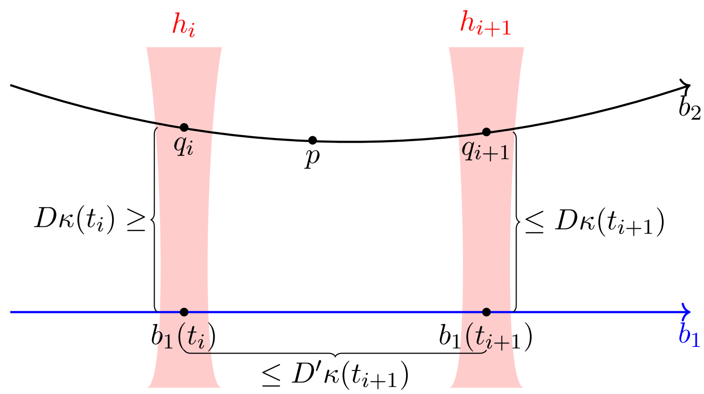

Figure 6.2: A picture for the proof of Proposition 6.2.1. We show that b_2 is in a κ-neighborhood of b_1.

Proof. Since the $\{h_i\}$ are ordered, then b_2 crossing h_m and h_n implies that b_2 also crosses all curtains of $\{h_i\}$ between h_m and h_n. Thus, b_2 must cross all but finitely many h_i. Removing such finitely many curtains will still result in a κ-chain dual to b_1, so we assume b_2 meets all of $\{h_i\}$ at points $q_i \in b_2$ such that $\pi_{b_1}(q_i) = b_1(t_i)$. Now, Lemma 5.2.8 tells us that $d(q_i, b_1(t_i)) \leq D\kappa(t_i)$ for some $D > 0$ depending only on the contracting constant C. Also, by definition of a κ-chain and Lemma 3.3.2, $d(b_1(t_i), b_1(t_{i+1})) \leq D'\kappa(t_i)$ for some $D' > 0$ depending only on C. Thus, for any $p \in b_2$ between q_i and q_{i+1}, we have

$$\begin{aligned} d(p, q_i) &\leq d(q_i, q_{i+1}) \\ &\leq d(q_i, b_1(t_i)) + d(b_1(t_i), b_1(t_{i+1})) + d(b_1(t_{i+1}), q_{i+1}) \\ &\leq D\kappa(t_i) + D'\kappa(t_i) + D\kappa(t_{i+1}) \\ &\leq D''\kappa(t_i) \end{aligned}$$

for some $D'' > 0$ by Lemma 3.3.2. This gives,

$$\begin{aligned} d(p, b_1) &\leq d(p, q_i) + d(q_i, b_1(t_i)) \\ &\leq \left(D'' + D\right)\kappa(t_i) \\ &\leq \left(D'' + D\right)\kappa(p). \end{aligned}$$

This is true for any $p \in b_2$ between some q_i and q_{i+1}, which gives us that b_2 is in a κ-neighborhood of b_1. See Figure 6.2. □

Remark 6.2.2. Note that, by the contrapositive, if b_1 and b_2 are κ-contracting rays that are not in the same κ-equivalence class, then each only crosses up to finitely many curtains in the other's κ-chain.

Proposition 6.2.3 (Injectivity of φ)**.** *Let $b_1^\infty, b_2^\infty \in \partial_\kappa X$ and let b_1, b_2 be the corresponding κ-contracting geodesic rays based at $\mathfrak{o}$. If $b_1^\infty \neq b_2^\infty$, then $\varphi(b_1^\infty) \neq \varphi(b_2^\infty)$ in $\partial\widehat{X}$.*

Proof. Since b_1 and b_2 are both κ-contracting rays, they both have κ-chains dual to them. Denote $\{h_i\}$ and $\left\{h'_j\right\}$ as the κ-chains dual to b_1 and b_2 with centers of poles $\{b_1(t_i)\}$ and

$\left\{b_2\left(t'_j\right)\right\}$, respectively. By Remark 6.2.2, b_1 and b_2 only cross finitely many of the other's κ-chain. Put h_m to be the last curtain of $\{h_i\}$ that b_2 crosses and h'_n to be the last curtain of $\{h_j\}$ that b_1 crosses. Consider the sequence $\{x_i\}$ where $x_i = b_1(t_i + 1)$ and $\{y_j\}$ where $y_j = b_2\left(t'_j + 1\right)$.

Now, h_{m+1} and h_{m+2} are L_1-separated for $L_1 = \lceil\kappa(t_{m+2})\rceil$ and, for $i \geq m+2$, separate x_i from b_2. Similarly, h'_{n+1} and h'_{n+2} are L_2-separated for $L_2 = \lceil\kappa(t'_{n+2})\rceil$ and for $j \geqslant n+2$, separate y_j from b_1. See the Figure 6.3.

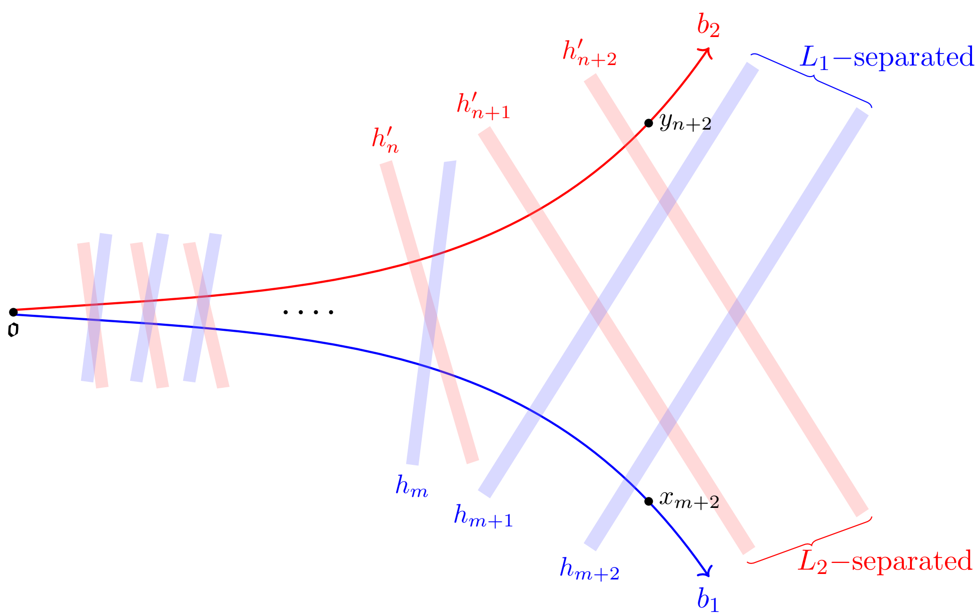

Figure 6.3: Two L_1-separated curtains separate x_{m+2} from b_2. Similarly, two L_2-separated curtains separate y_{n+2} from b_1.

We now investigate $d_L(x_i, y_j)$ for $i > m+2, j > n+2$. Let c be a maximal L-chain that realizes $d_L(x_i, x_{m+2}) = 1 + |c|$. Due to where h_{m+1} and h_{m+2} are placed, at least $|c| - L_1$ curtains in c will not intersect b_2. Let c' be a maximal chain realizing $d_L(y_j, y_{n+2}) = 1 + |c'|$. By a similar argument involving h'_{n+1} and h'_{n+2}, we get that at least

$|c'| - L_2$ curtains in c' will not intersect b_1. By gluing these two chains together via Lemma 6.1.5, we obtain an L-chain of length at least $|c| + |c'| - L_1 - L_2 - L - 2$ that separates x_i and y_j. Thus,

$$d_L(x_i, y_j) \geq d_L(x_i, x_{m+2}) + d_L(y_j, y_{n+2}) - L_1 - L_2 - L - 4.$$

Hence,

$$\begin{aligned}
(x_i \cdot y_j)_{\mathfrak{o}} &= \frac{1}{2}\left(\hat{d}(\mathfrak{o}, x_i) + \hat{d}(\mathfrak{o}, y_j) - \hat{d}(x_i, y_j)\right) \\
&= \frac{1}{2}\left(\hat{d}(\mathfrak{o}, x_i) + \hat{d}(\mathfrak{o}, y_j) - \sum_{L=1}^{\infty} \frac{d_L(x_i, y_j)}{L^3}\right) \\
&\leq \frac{1}{2}\left(\hat{d}(\mathfrak{o}, x_i) + \hat{d}(\mathfrak{o}, y_j) - \sum_{L=1}^{\infty} \frac{d_L(x_i, x_{m+2}) + d_2(y_i, y_{n+2}) - L_1 - L_2 - L - 4}{L^3}\right) \\
&= \frac{1}{2}\left(\hat{d}(\mathfrak{o}, x_i) + \hat{d}(\mathfrak{o}, y_j) - \hat{d}(x_i, x_{m+2}) - \hat{d}(y_i, y_{n+2}) + \sum_{L=1}^{\infty} \frac{L_1 + L_2 + L + 4}{L^3}\right) \\
&\leq \frac{1}{2}\left(\hat{d}(\mathfrak{o}, x_{m+2}) + \hat{d}(\mathfrak{o}, y_{n+2}) + \sum_{L=1}^{\infty} \frac{L_1 + L_2 + L + 4}{L^3}\right)
\end{aligned}$$

where the last inequality is due to the triangle inequality. Thus, for any $i \geq m+2$ and $j \geq n+2$, we have that $(x_i \cdot y_j)_{\mathfrak{o}}$ is bounded by the above constant. This gives

$$\liminf_{i,j \to \infty} (x_i \cdot y_j)_{\mathfrak{o}} < \infty,$$

so $\{x_i\}$ and $\{y_j\}$ represent different equivalence classes in $\partial \widehat{X}$. This implies the same for $\varphi(b_1^\infty)$ and $\varphi(b_2^\infty)$.

□

6.3 Continuity of $\varphi : \partial_\kappa X \longrightarrow \partial \widehat{X}$

We now review the topologies of both $\partial_\kappa X$ and $\partial \widehat{X}$ in preparation for showing continuity of φ. Recall that we defined the topology of $\partial \widehat{X}$ by the basic open sets

$$U(s,r) = \left\{ [\{y_i\}] \;\middle|\; \liminf_{i,j\to\infty} (x_i \cdot y_j)_{\mathfrak{o}} \geq r \text{ for some } \{x_i\} \in s \right\}$$

where $s \in \partial \widehat{X}$ and $r > 0$. Showing continuity will involve using intersections of cone topology bases (Definition 6.3.1) and Curtain topology sets (Definition 6.3.2).

Definition 6.3.1 (Visual boundary, cone topology as a subspace topology)**.** The *visual boundary* of a proper CAT(0) space X, denoted ∂X as a set, is the set of all geodesic rays emanating from $\mathfrak{o}$. The *cone topology* of ∂X is generated by the basic open sets

$$V_{R,\epsilon}(\xi) := \{\eta \in \partial X \mid d(\xi(R), \eta(R)) < \epsilon\}.$$

When ∂X is equipped with the cone topology, we denote it $\partial_\infty X$. For convenience, we denote such sets in the subspace $\partial_\kappa X$ as

$$U_{R,\epsilon}(\xi) := V_{R,\epsilon}(\xi) \cap \partial_\kappa X.$$

This defines a subspace topology on $\partial_\kappa X$.

The following topology on $\partial_\kappa X$ is a useful topology to be able leverage curtain machinery in the same way hyperplanes were leveraged in the analogous topology defined in [IMZ23].

Definition 6.3.2 (Curtain Topology Sets)**.** Let $b^\infty \in \partial_\kappa X$ be a geodesic ray emanating from $\mathfrak{o}$. We say a curtain h *separates* $\mathfrak{o}$ from a^∞ if there exists a $T > 0$ such that h separates $\mathfrak{o}$ from $a_{[T,\infty]}$. For each curtain h dual to b, we define

$$U_h(b^\infty) = \{a^\infty \in \partial X : h \text{ separates } \mathfrak{o} = a(0) \text{ from } a^\infty\} \cap \partial_\kappa X.$$

We define the *curtain topology on* $\partial_\kappa X$ as follows: a set $O \subset \partial_\kappa X$ is open if for each $b^\infty \in O$, there exists a curtain h dual to b such that $U_h(b^\infty) \subset O$. This yields a topology on $\partial_\kappa X$. Note, such a topology is actually the subspace topology of the curtain topology on all of ∂X. However, for convenience of notation, we define $U_h(b^\infty)$ to be intersecting with $\partial_\kappa X$ as above since we are only interested in $\partial_\kappa X$ for this paper.

Lemma 6.3.3. *The cone topology and the curtain topology agree on* $\partial_\kappa X$*.*

Proof. Showing the curtain topology is coarser than the cone topology is given in the proof of Theorem 8.8 in [PSZ22]. The reverse direction is a recreation of the proof of Theorem 4.2 in [IMZ23], but in the curtain setting.

Consider some $U_{R,\epsilon}(b^\infty)$ for some κ-contracting geodesic b and $R, \epsilon > 0$. Then there is an infinite κ-chain $\{h_i\}$ dual to b. Denote the center of poles of each h_i as $b(t_i)$. Put D as the constant in Lemma 5.2.8 (note that D only depends on the excursion constant of b). Fix m large enough so that $t_m > R$ and $2D\kappa(t_m) \leq \frac{t_m}{R}\epsilon$. We claim $U_{h_{m+1}}(b^\infty) \subset U_{R,\epsilon}(b^\infty)$.

Indeed, let $a^\infty \in U_{h_{m+1}}(b^\infty)$, so a is a geodesic emanating from $\mathfrak{o}$ that crosses h_{m+1}. By Lemma 5.2.8, there exists an $s > 0$ such that $d(a(s), b(t_m)) \leq D\kappa(t_m)$. Note that, since a and b are both geodesics emanating from $\mathfrak{o}$, we get that $|s - t_m| \leq D\kappa(t_m)$ Thus,

$$d(a(t_m), b(t_m)) \leq d(a(t_m), a(s)) + d(a(s), b(t_m)) \leq 2D\kappa(t_m) \leq \frac{t_m}{R}\epsilon$$

Due to convexity of the CAT(0) metric (as shown in [BH99, II.2.1] and its proof), $d(a(R), b(R)) \leq \epsilon$. Thus, $a^\infty \in U_{R,\epsilon}(b^\infty)$.

□

Theorem 6.3.4. *The map $\varphi : \partial_\kappa X \longrightarrow \partial\widehat{X}$ is a well-defined, injective, and continuous map when $\partial_\kappa X$ is endowed with subspace topology of the cone topology.*

Proof. Proposition 6.2.3 shows that φ is injective. What is left is to show continuity. Let $U(s, r)$ be an open set in $\partial\widehat{X}$. If its preimage is nonempty, we can assume that there exists a $b^\infty \in \partial_\kappa X$ such that $\varphi(b^\infty) = s$ and it suffices to show that there exists a $U_h(b^\infty) \subset \varphi^{-1}(U(s, r))$ for some curtain h dual to b.

Let $r' > 0$ and $\epsilon > 0$. Since b is a κ-contracting geodesic, denote $\{h_i\}$ the κ-chain dual to b with $b\,(t_i)$ the center of the pole of each h_i. Since b is unbounded in $\widehat{X}$, there exists an $i = i_{r'}$ such that $\hat{d}\left(\mathfrak{o}, b\left(t_{i_{r'}}\right)\right) > r'$. Consider the set $U_{h_{i_{r'}}}(b^\infty)$ and the open set $U_{R,\epsilon}(b^\infty)$ where $R > t_{i_{r'}}$. Since $U_{R,\epsilon}(b^\infty)$ is open, there exists a curtain k dual to b such that $U_k(b^\infty) \subseteq U_{R,\epsilon}(b^\infty)$ by Lemma 6.3.3 . Without loss of generality, the pole of the curtain k is farther distance away from $\mathfrak{o}$ in the CAT(0) metric than the pole of $h_{t_{i_{r'}}}$, so $U_k(b^\infty) \subset U_{h_{t_{i_{r'}}}}(b^\infty)$.

We claim $U_k(b^\infty) \subseteq \varphi^{-1}(U(s,r))$. Indeed, let $a^\infty \in U_k(b^\infty)$. So a is the geodesic in a^∞ emanating from $\mathfrak{o}$. Then, a crosses the first $i_{r'}$ curtains of $\{h_i\}$ and also $d(a(t), b) < \epsilon$ for all $t \leq R$. See Figure 6.4. Now, consider any unbounded and increasing sequence $\{x_i\} \subseteq b$ and $\{y_j\} \subseteq a$ with the added condition that, for some n, $x_n = b\left(t_{i_{r'}} + \frac{1}{2}\right)$ and y_n such that $d\left(x_n, y_n\right) < \epsilon$. What's left to show is that $\left(x_m \cdot y_{m'}\right)_{\mathfrak{o}} \geq r$ for all $m, m' > n$. By Lemma 6.1.7, we have there exists a constant C such that,

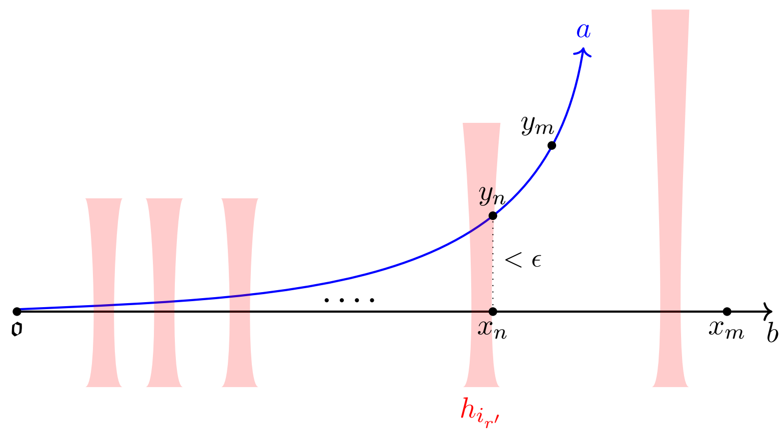

Figure 6.4: Picture of b crossing $h_{i_{r'}}$ as well as where x_n and y_n are placed.

$$\hat{d}\left(\mathfrak{o}, x_m\right) \geq \hat{d}\left(\mathfrak{o}, x_n\right) + \hat{d}\left(x_n, x_m\right) - C$$

$$\hat{d}\left(\mathfrak{o}, y_{m'}\right) \geq \hat{d}\left(\mathfrak{o}, y_n\right) + \hat{d}\left(y_n, y_{m'}\right) - C.$$

Also, by the triangle inequality,

$$\hat{d}\left(x_m, y_{m'}\right) \leq \hat{d}\left(x_m, x_n\right) + \hat{d}\left(x_n, y_n\right) + \hat{d}\left(y_n, y_{m'}\right).$$

Thus, by the above inequalities, we get

$$\begin{aligned}(x_m \cdot y_{m'})_{\mathfrak{o}} &= \frac{1}{2}\left(\hat{d}\,(\mathfrak{o}, x_m) + \hat{d}\,(\mathfrak{o}, y_{m'}) - \hat{d}\,(x_m, y_{m'})\right) \\ &\geq \frac{1}{2}\left(\hat{d}\,(\mathfrak{o}, x_n) + \hat{d}\,(\mathfrak{o}, y_n) - \hat{d}\,(x_n, y_n) - 2C\right)\end{aligned}$$

Since $[\mathfrak{o}, x_n]$ and $[\mathfrak{o}, y_n]$ both cross the curtains $\{h_1, h_2, \cdots, h_{i_{r'}}\}$ we see that $\hat{d}\,(\mathfrak{o}, x_n) > r'$ and $\hat{d}\,(\mathfrak{o}, y_n) > r'$. Thus,

$$\begin{aligned}(x_m \cdot y_{m'})_{\mathfrak{o}} &\geq \frac{1}{2}\left(r' + r' - \epsilon - 2C\right) \\ &= r' - \frac{1}{2}\epsilon - C.\end{aligned}$$

Since r' and ϵ were arbitrary, we can fix ϵ and choose r' such that $(x_m \cdot y_{m'})_{\mathfrak{o}} > r$. Thus, $\varphi(a^\infty) \in U(s, r)$. This is true for any $a^\infty \in U_k(b^\infty)$. Hence, φ is continuous.

□

Since the underlying sets of X and $\widehat{X}$ are the same, φ is an Isom X-equivariant map. Due to Lemma 3.3.10, we get the same result when $\partial_\kappa X$ is endowed with the $\mathcal{SM}$ topology (see Definition 3.3.7).

Corollary 6.3.5. *The map $\varphi : \partial_\kappa X \longrightarrow \partial \widehat{X}$ is a well-defined, injective, and continuous map when $\partial_\kappa X$ is endowed with the sublinear Morse topology.*

Recent work of [AIM22] has shown new proving techniques when regarding cobounded projections to hyperbolic spaces. With Definition 6.3.6 and Lemma 6.3.7, we can follow [AIM22][Lemma 3.32] to upgrade φ in Theorem 6.3.4 to a homeomorphism.

Definition 6.3.6. Let X be a metric space, $C \geq 0$, and $I \subset \mathbb{R}$ a (possibly unbounded) closed interval. A path $\gamma : I \to X$ is an *unparameterized C-quasi-ruler* if it satisfies the following conditions.

- $\forall t < s < r$, we have $d(\gamma(t), \gamma(s)) + d(\gamma(s), \gamma(r)) \leq d(\gamma(t), \gamma(r)) + C$.
- $\forall t_0 \in I$, we have $\limsup_{|t-t_0| \to 0} d(\gamma(t), \gamma(t_0)) < C$.

We see that, for any κ-contracting ray $b : [0, \infty) \longrightarrow X$, and any $t_0 \in [0, \infty)$, we have $\limsup_{|t-t_0| \to 0} \hat{d}(b(t), b(t_0)) \leq \sum \frac{1}{L^3}$. This, along with Lemma 6.1.7 show that b is a C-quasi-ruler for the constant C in Lemma 6.1.7.

Lemma 6.3.7 (Lemma A.11 in [AIM22]). *Let $\widehat{X}$ be a δ-hyperbolic space with basepoint $\mathfrak{o} \in X$, and let $\gamma, \gamma' : [0, \infty) \to \widehat{X}$ be two C-quasi-rulers with unbounded image that start at $\mathfrak{o}$ and define points $[\{\gamma(n)\}], [\{\gamma'(m)\}]$ in $\partial \widehat{X}$. If $x' \in \gamma$ and $y' \in \gamma'$ are such that $d(\mathfrak{o}, x'), d(\mathfrak{o}, y') \leq \liminf_{m,n \to \infty} (\gamma(n) \cdot \gamma'(m))_{\mathfrak{o}} + C$, then*

$$(x' \cdot y')_{\mathfrak{o}} \geq \min \left\{ d(\mathfrak{o}, x'), d(\mathfrak{o}, y') \right\} - C - 2\delta$$

and

$$d(x', y') \leq C + 2\delta + \left| d(\mathfrak{o}, x') - d(\mathfrak{o}, y') \right|$$

Theorem 6.3.8. *When $\partial_\kappa X$ is endowed with the subspace topology of the cone topology of ∂X, the map $\varphi : \partial_\kappa X \longrightarrow \partial \widehat{X}$ is a homeomorphism onto its image.*

Proof. Continuity is already shown in Theorem 6.3.4. We only need to show that, for any open set $O \subset \partial_\kappa X$, $\varphi(O)$ is open in $\varphi(\partial_\kappa X)$ when endowed with the subspace topology. Let

$b^\infty \in O$. By the definition of being open in the curtain topology, there exists a curtain k dual to b such that $U_k(b^\infty) \subset O$. We have b is κ-contracting, so there exists a κ-chain $\{h_i\}$ dual to b with centers of poles $b(t_i)$ for each h_i. Since k is dual to b, we have that there exists a J such that k separates h_j from $\mathfrak{o}$ for all $j \geq J$.

Let $r > 0$. Choose t_r such that for all $y \in X$ on the same side of k as $\mathfrak{o}$, $\hat{d}(y, b(t_r)) > r$. Note, this is possible since, for large enough t, $b(t)$ will be on the side of h_J opposite of $\mathfrak{o}$. Since $b([J, \infty))$ is unbounded, such a $b(t_r)$ will exist. Put $R = \hat{d}(\mathfrak{o}, b(t_r))$. Set $V = U(\varphi(b^\infty), R) \cap \varphi(\partial_\kappa X)$. We claim $\varphi(b^\infty) \in V \subseteq \varphi(U_k(b^\infty)) \subseteq \varphi(O)$.

Indeed, let $\eta \in V$, and put a as the geodesic emanating from $\mathfrak{o}$ such that $\varphi(a^\infty) = \eta$. So $[\{a(n)\}] = \eta$. Let $m \in \mathbb{N}$ be the smallest m such that $\hat{d}(\mathfrak{o}, a(m)) \geq R$ for $a(m) \in \{a(n)\}$. Since a is a C-quasi-ruler, we have that there exists a C such that $\hat{d}(\mathfrak{o}, a(m)) \leq R + C$. By Lemma 6.3.7, we get that

$$\hat{d}(b(t_r), a(m)) \leq C + 2\delta + |\hat{d}(\mathfrak{o}, b(t_r)) - \hat{d}(\mathfrak{o}, a(m))| \leq 2C + 2\delta,$$

where δ is the hyperbolicity constant of $\widehat{X}$. Since r was arbitrary, we can choose r large enough to force $a(m)$ to be on the same side of k as $b(t_r)$. This means a crosses k. In other words, $a^\infty \in U_k(b)$. This completes the proof. □

6.4 A κ-Persistent Shadow Characterization

We now give our second characterization of κ-Morse rays (Theorem D in the introduction). Just as how κ-Morse rays in mapping class groups project to curve graphs in a sublinearly scaled way, we find an equivalent notion in the CAT(0) setting.

Definition 6.4.1 (κ-persistent shadow, persistent shadow constant)**.** A geodesic ray b with infinite diameter has a *κ-persistent shadow* in $\widehat{X}$ if there exists a $C > 0$ such that for all $s < t$,

$$\hat{d}(b(s), b(t)) \geq C \cdot \frac{t-s}{\kappa(t)} - C.$$

We refer to C above as the *persistent shadow constant.*

We show how this characterization is connected to κ-contracting in the next Theorem 6.4.3. We repeat the following lemma first given in Chapter 4, which works in dualizing L-chains that meet a geodesic. This Lemma is useful in the proof of Theorem 6.4.3.

Lemma 6.4.2 (Lemma 2.21 in [PSZ22])**.** *Let $L, n \in \mathbb{N}$, let $\{h_1, \dots, h_{(4L+10)n}\}$ be an L-chain, and suppose that $A, B \subset X$ are separated by every h_i. For any $x \in A$ and $y \in B$, the sets A and B are separated by an L-chain of length at least $n+1$ all of whose elements are dual to $[x, y]$ and separate h_1 from $h_{(4L+10)n}$.*

Theorem 6.4.3. *Let b be geodesic ray in a CAT(0) space X emanating from $\mathfrak{o}$ with infinite diameter projection onto $\widehat{X}$.*

- *If b is κ-contracting and κ^4 is sublinear, then b has a κ^4-persistent shadow in the $\hat{d}$ metric.*

- *If b has a κ-persistent shadow in the $\hat{d}$ metric and κ^2 is sublinear, then b is κ^2-contracting.*

Proof. For the forward direction, put $D > 0$ as the excursion constant, and denote $\{b(t_i)\}$ as the center of poles of the dual curtains in the κ-chain $\{h_i\}$. We follow the same process as in Lemma 6.1.9. That is, for any $s \leq t$, we have that there exists a maximal j and minimal i such that $t_j \leq s \leq t \leq t_i$ (if $s \leq t_1$, we choose $t_0 = 0$). Following the proof of Lemma 6.1.9 gives us

$$i - j \geq \frac{t_i - t_j}{D\kappa(t_i)}.$$

Thus, since $\lceil D\kappa(t_i) \rceil = m$ for some $m \in \mathbb{N}$, we get that $d_m(b(s), b(t)) \geq i - j - 2$ and

$$\begin{aligned}
\frac{d_m(b(s), b(t))}{m^3} &\geq \frac{\frac{t_i - t_j}{D\kappa(t_i)} - 2}{m^3} \\
&\geq \frac{t_i - t_j}{D\kappa(t_i)(D\kappa(t_i) + 1)^3} - 2 \\
&\geq \frac{t - s}{D'\kappa(t)(D'\kappa(t) + 1)^3} - 2.
\end{aligned}$$

for some $D' > 0$ dependent on D by Lemma 3.3.2. Thus, b has a κ^4-persistent shadow

For the reverse direction, put C as the persistent shadow constant. Since b has infinite diameter and κ^2 is sublinear, consider a sequence $\{t_i\}$ such that $t_{i+1} - t_i = \frac{D}{C^2}\kappa^2(t_{i+1})$ for some finite $D > 0$. Note that, for any i, $\lceil \kappa(t_{i+1}) \rceil \geq 2$, and $\sum_{L=2}^{\infty} \frac{1}{L^3} \leq \sum_{L=2}^{\infty} \frac{1}{L^2} \leq \frac{2}{3}$. Thus,

$$\begin{aligned}
\frac{D}{C}\kappa(t_{i+1}) - C &= C \cdot \frac{t_{i+1} - t_i}{\kappa(t_{i+1})} - C \\
&\leq \hat{d}(b(t_i), b(t_{i+1})) \\
&= \sum_{L=1}^{\lceil \kappa(t_{i+1}) \rceil - 1} \frac{d_L(b(t_i), b(t_{i+1}))}{L^3} + \sum_{L=\lceil \kappa(t_{i+1}) \rceil}^{\infty} \frac{d_L(b(t_i), b(t_{i+1}))}{L^3} \\
&\leq \sum_{L=1}^{\lceil \kappa(t_{i+1}) \rceil - 1} \frac{d_L(b(t_i), b(t_{i+1}))}{L^3} + \sum_{L=\lceil \kappa(t_{i+1}) \rceil}^{\infty} \frac{d(b(t_i), b(t_{i+1})) + 1}{L^3}
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{L=1}^{\lceil \kappa(t_{i+1}) \rceil - 1} \frac{d_L(b(t_i), b(t_{i+1}))}{L^3} + \sum_{L=\lceil \kappa(t_{i+1}) \rceil}^{\infty} \frac{\frac{D}{C}\kappa^2(t_{i+1}) + 1}{L^3} \\
&\leq \sum_{L=1}^{\lceil \kappa(t_{i+1}) \rceil - 1} \frac{d_L(b(t_i), b(t_{i+1}))}{L^3} + \frac{D}{C}\kappa(t_{i+1}) \cdot \left(\sum_{L=\lceil \kappa(t_{i+1}) \rceil}^{\infty} \frac{1}{L^2} \right) + 1 \\
&\leq \sum_{L=1}^{\lceil \kappa(t_{i+1}) \rceil - 1} \frac{d_L(b(t_i), b(t_{i+1}))}{L^3} + \frac{D}{C}\kappa(t_{i+1}) \cdot \frac{2}{3} + 1.
\end{aligned}
$$

That is,

$$
\frac{1}{3C} D\kappa(t_{i+1}) - C - 1 \leq \sum_{L=1}^{\lceil \kappa(t_{i+1}) \rceil - 1} \frac{d_L(b(t_i), b(t_{i+1}))}{L^3}. \tag{6.1}
$$

Since D was arbitrary, we now choose $D = 3003C$ so that the left hand side of (6.1) is greater than $1000\kappa(t_{i+1})$. We claim that for some L with $1 \leq L \leq \lceil \kappa(t_{i+1}) \rceil - 1$, we get $d_L(b(t_i), b(t_{i+1})) \geq 8\kappa(t_{i+1}) + 21$. Indeed, if not, then

$$
\sum_{L=1}^{\lceil \kappa(t_{i+1}) \rceil - 1} \frac{d_L(b(t_i), b(t_{i+1}))}{L^3} \leq \sum_{L=1}^{\lceil \kappa(t_{i+1}) \rceil - 1} \frac{8\kappa(t_{i+1}) + 21}{L^3} \leq 2(8\kappa(t_{i+1}) + 21)
$$

which contradicts (6.1). By Lemma 6.4.2, we get that there are three L-separated curtains dual to $[b(t_i), b(t_{i+1})]$. These curtains, by construction, are also $\kappa(t_{i+1})$ separated. This is true for all i. We choose h_i to be the dual curtain of $[b(t_i), b(t_{i+1})]$ that is closest to $b(t_i)$, and put s_i to be the center of the pole of h_i. Then, we see that each of the pairs $\{h_i, h_{i+1}\}$ are $\kappa(t_{i+1})$-separated since there are two $\kappa(t_{i+1})$-separated curtains that separate h_i and h_{i+1}. See Figure 6.5. Also,

$$
s_{i+1} - s_i \leq t_{i+2} - t_i \leq 2\frac{D}{C^2}\kappa^2(t_{i+2}) \leq D'\kappa^2(t_{i+1}) \leq D'\kappa^2(t_{s_{i+1}})
$$

for some $D' > 0$ dependent on C by Lemma 3.3.2. Thus, we have an infinite chain of dual curtains $\{h_i\}$ such that h_i and h_{i+1} are $\kappa(t_{i+1})$-separated (which implies $\kappa^2(t_{i+1})$-separated) and $t_{i+1} - t_i = \frac{D}{C^2}\kappa^2(t_{i+1})$ for all i. We conclude that b is κ^2-contracting. □

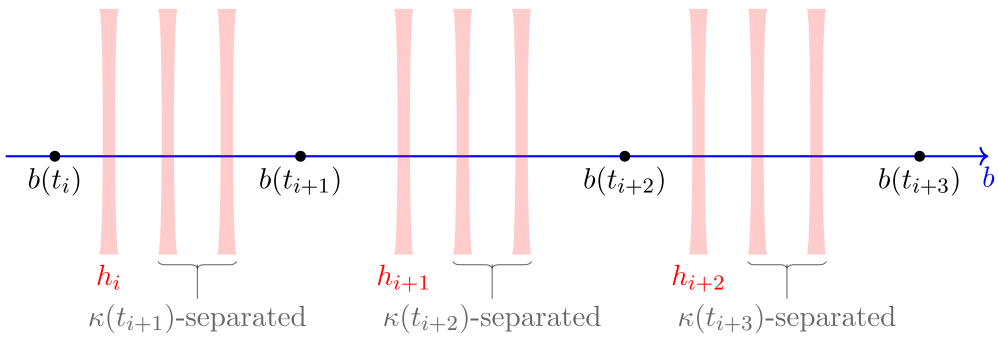

Figure 6.5: Creating a dual κ-chain to b. Since each pair $\{h_i, h_{i+1}\}$ will have two $\kappa(t_{i+1})$-separated curtains that separate the pair, we get that h_i and h_{i+1} will be $\kappa(t_{i+1})$-separated.

6.5 Further Questions

Question 1. Future work of Petyt, Spriano, and Zalloum will extend curtain machinery to a larger class of spaces than CAT(0) spaces. Zalloum gives a great survey discussing the recent developments of curtains and their relationships with hierarchically hyperbolic spaces in [Zal23]. If the curtain model can be extended to a larger class of spaces, will the results and techniques in this paper also extend as well? In what generality can sublinear Morseness be characterized by curtain excursion?

Question 2. As seen in Definition 6.1.1, the distance function in the curtain model is described through a family of d_L metrics inspired by Genevois's hyperplane-separation metrics in [Gen19]. It is likely that, in the cube complex setting, one could create a hyperbolic "hyperplane model" that uses hyperplane-separating metrics instead of curtain-separating metrics. If such a hyperplane model can be created, would this model be preferred to the curtain model in the CAT(0) cube complex setting? Would the curtain model be quasi-isometric to the "hyperplane model"?

Question 3. The study of acylindrical actions, as initiated by Osin in [Osi16], has been a great tool for studying groups with aspects of non-positive curvature. In the context of hierarchically hyperbolic spaces, Abbott-Behrstock-Durham show in [ABD21] that hierarchically hyperbolic groups have acylindrical actions which are largest and universal. Could the same conclusion be made about CAT(0) groups via an action on their curtain model?

Question 4. The *simplicial boundary* of a CAT(0) cube complex was first created by Hagen in [Hag13]. It is defined by classes of *unidirectional boundary sets*, i.e. nested sets of halfspaces diverging in a single direction of infinity. In a CAT(0) cube complex, the halfspaces mentioned are associated only to corresponding hyperplanes in the cube complex structure. Could one create a well-defined unidirectional boundary set with curtains in a CAT(0) space? Could this be extended to create a simplicial boundary for a CAT(0) space? Could one use this simplicial boundary to study quasi-isometry invariants like *divergence* [Hag13] and *thickness* [BH16]?

Question 5. Further work of Yulan Qing and Kasra Rafi is the creation of a *quasi-redirecting boundary* — a quasi-invariant boundary even larger than the κ-Morse boundary for a proper geodesic metric space. What can be said about the quasi-redirecting boundary? For example, do groups have minimal actions on the quasi-redirecting boundary like they do for the κ-Morse boundary?

Question 6. The counterexample for quasi-isometry invariance given in Theorem 4.2.3 is somewhat unsatisfying as the CAT(0) space does not have a cocompact action. Does there exist two quasi-isometric spaces with cocompact actions with non-quasi-isometric curtain models?

Bibliography

[ABD21] Carolyn Abbott, Jason Behrstock, and Matthew Durham. Largest acylindrical actions and stability in hierarchically hyperbolic groups. *Transactions of the American Mathematical Society, Series B*, 8(3):66–104, 2021.

[ACGH17] Goulnara N Arzhantseva, Christopher H Cashen, Dominik Gruber, and David Hume. Characterizations of morse quasi-geodesics via superlinear divergence and sublinear contraction. *Documenta Mathematica*, 22:1193–1224, 2017.

[Ago13] Ian Agol. The virtual haken conjecture (with an appendix by ian agol, daniel groves and jason manning). *Documenta Mathematica*, 18:1045–1087, 2013.

[AIM22] Carolyn Abbott and Merlin Incerti-Medici. Hyperbolic projections and topological invariance of sublinearly morse boundaries. *arXiv preprint arXiv:2212.09539*, 2022.

[Bar22a] Corentin Le Bars. Central limit theorem on cat (0) spaces with contracting isometries. *arXiv preprint arXiv:2209.11648*, 2022.

[Bar22b] Corentin Le Bars. Random walks and rank one isometries on cat (0) spaces. *arXiv preprint arXiv:2205.07594*, 2022.

[Beh06] Jason A Behrstock. Asymptotic geometry of the mapping class group and teichmüller space. *Geometry & Topology*, 10(3):1523–1578, 2006.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[BH16] Jason Behrstock and Mark F Hagen. Cubulated groups: thickness, relative hyperbolicity, and simplicial boundaries. *Groups, Geometry, and Dynamics*, 10(2):649–707, 2016.

[BHS17] Jason Behrstock, Mark Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, i: Curve complexes for cubical groups. *Geometry & Topology*, 21(3):1731–1804, 2017.

[BHS21] Jason Behrstock, Mark F Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *Duke Mathematical Journal*, 170(5):909–996, 2021.

[BK02] Nadia Benakli and Ilya Kapovich. Boundaries of hyperbolic groups. *arXiv preprint math/0202286*, 2002.

[Cas16] Christopher H. Cashen. Quasi-isometries need not induce homeomorphisms of contracting boundaries with the gromov product topology. *Analysis and Geometry in Metric Spaces*, 4(1):000010151520160011, 2016.

[CD95] Ruth Charney and Michael W Davis. Finite k (; 1) s for artin groups. *Ann. of Math. Stud., Princeton Univ. Press*, 138:110–124, 1995.

[CFFT22] Kunal Chawla, Behrang Forghani, Joshua Frisch, and Giulio Tiozzo. The poisson boundary of hyperbolic groups without moment conditions. *arXiv preprint arXiv:2209.02114*, 2022.

[Cha07] Ruth Charney. An introduction to right-angled artin groups. *Geometriae Dedicata*, 125:141–158, 2007.

[Cho22] Inhyeok Choi. Random walks and contracting elements i: Deviation inequality and limit laws. *arXiv preprint arXiv:2207.06597v1*, 2022.

[Cho23] Inhyeok Choi. Random walks and contracting elements iv: Sublinearly morse boundary. *to appear*, 2023.

[CK00] Christopher B. Croke and Bruce Kleiner. Spaces with nonpositive curvature and their ideal boundaries. *Topology*, 39(3):549–556, 2000.

[CM17] Matt Clay and Dan Margalit. *Office hours with a geometric group theorist.* Princeton University Press, 2017.

[CM19] Christopher H. Cashen and John M. Mackay. A metrizable topology on the contracting boundary of a group. *Trans. Amer. Math. Soc.*, 372(3):1555–1600, 2019.

[Cor17] Matthew Cordes. Morse boundaries of proper geodesic metric spaces. *Groups Geom. Dyn.*, 11(4):1281–1306, 2017.

[Cor19] Matthew Cordes. A survey on Morse boundaries and stability. In *Beyond hyperbolicity*, volume 454 of *London Math. Soc. Lecture Note Ser.*, pages 83–116. Cambridge Univ. Press, Cambridge, 2019.

[CS11] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for cat (0) cube complexes. *Geometric and functional analysis*, 21(4):851–891, 2011.

[CS15] Ruth Charney and Harold Sultan. Contracting boundaries of CAT(0) spaces. *J. Topol.*, 8(1):93–117, 2015.

[Deh12] Max Dehn. Transformation der kurven auf zweiseitigen flächen. *Mathematische Annalen*, 72(3):413–421, 1912.

[DMS20] Matthew G Durham, Yair N Minsky, and Alessandro Sisto. Stable cubulations, bicombings, and barycenters. *arXiv preprint arXiv:2009.13647*, 2020.

[DR09] Moon Duchin and Kasra Rafi. Divergence of geodesics in teichmüller space and the mapping class group. *Geometric and Functional Analysis*, 19:722–742, 2009.

[Dur23] Matthew Gentry Durham. Cubulating infinity in hierarchically hyperbolic spaces, 2023.

[DZ22] Matthew Gentry Durham and Abdul Zalloum. The geometry of genericity in mapping class groups and teichmuller spaces via cat (0) cube complexes. *arXiv preprint arXiv:2207.06516*, 2022.

[FM11] Benson Farb and Dan Margalit. *A primer on mapping class groups (pms-49)*, volume 41. Princeton university press, 2011.

[Gen16] Anthony Genevois. Coning-off cat (0) cube complexes. *arXiv preprint arXiv:1603.06513*, 2016.

[Gen19] Anthony Genevois. Hyperbolicities in CAT(0) cube complexes. *Enseign. Math.*, 65(1-2):33–100, 2019.

[GQR22] Ilya Gekhtman, Yulan Qing, and Kasra Rafi. Genericity of sublinearly morse directions in cat (0) spaces and the teichm\" uller space. *arXiv preprint arXiv:2208.04778*, 2022.

[Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[Gro92] Mikhael Gromov. Asymptotic invariants of infinite groups. Technical report, P00001028, 1992.

[Hag13] Mark F Hagen. The simplicial boundary of a cat (0) cube complex. *Algebraic & Geometric Topology*, 13(3):1299–1367, 2013.

[He23] Vivian He. Equivalent topologies on the contracting boundary. *Glasnik matematički*, 58(1):75–83, 2023.

[HHP20] Thomas Haettel, Nima Hoda, and Harry Petyt. Coarse injectivity, hierarchical hyperbolicity, and semihyperbolicity. *arXiv preprint arXiv:2009.14053*, 2020.

[Hua17] Jingyin Huang. Top-dimensional quasiflats in cat (0) cube complexes. *Geometry & Topology*, 21(4):2281–2352, 2017.

[IM21] Merlin Incerti-Medici. Comparing topologies on the morse boundary and quasi-isometry invariance. *Geometriae Dedicata*, 212(1):153–176, 2021.

[IMZ23] Merlin Incerti-Medici and Abdul Zalloum. Sublinearly morse boundaries from the viewpoint of combinatorics. *Forum Mathematicum*, 35(4):1077–1103, 2023.

[KM12] Jeremy Kahn and Vladimir Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Annals of Mathematics*, pages 1127–1190, 2012.

[MM99] Howard Masur and Yair Minsky. Geometry of the complex of curves, i: Hyperbolicity. *Invent. Math.*, 138:103–149, 1999.

[MMS12] Howard Masur, Lee Mosher, and Saul Schleimer. On train-track splitting sequences. *Duke Mathematical Journal*, 161(9):1613 – 1656, 2012.

[MQZ22] Devin Murray, Yulan Qing, and Abdul Zalloum. Sublinearly morse geodesics in CAT(0) spaces: lower divergence and hyperplane characterization. *Algebraic & Geometric Topology*, 22(3):1337–1374, aug 2022.

[NQ22] Hoang Thanh Nguyen and Yulan Qing. Sublinearly morse boundary of cat (0) admissible groups. *arXiv preprint arXiv:2203.00935*, 2022.

[NR03] Graham A Niblo and Lawrence D Reeves. Coxeter groups act on cat (0) cube complexes. *Journal of Group Theory*, 6(3):399–413, 2003.

[Osi16] Denis Osin. Acylindrically hyperbolic groups. *Transactions of the American Mathematical Society*, 368(2):851–888, 2016.

[PSZ22] Harry Petyt, Davide Spriano, and Abdul Zalloum. Hyperbolic models for cat (0) spaces. *arXiv preprint arXiv:2207.14127*, 2022.

[QR22] Yulan Qing and Kasra Rafi. Sublinearly morse boundary i: Cat (0) spaces. *Advances in Mathematics*, 404:108442, 2022.

[QRT20] Yulan Qing, Kasra Rafi, and Giulio Tiozzo. Sublinearly morse boundary ii: Proper geodesic spaces. *arXiv preprint arXiv:2011.03481*, 2020.

[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proceedings of the London Mathematical Society*, 3(3):585–617, 1995.

[She22] Sam Shepherd. A cubulation with no factor system. *arXiv preprint arXiv:2208.10421*, 2022.

[Ves23a] Elliott Vest. Curtain characterization of sublinearly morse geodesics in cat (0) spaces. *arXiv preprint arXiv:2309.02725*, 2023.

[Ves23b] Elliott Vest. The curtain model is not a quasi-isometry invariant of cat (0) spaces. *arXiv preprint arXiv:2312.16325*, 2023.

[Wil05] Julia M Wilson. A cat (0) group with uncountably many distinct boundaries. 2005.

[Wis04] Daniel T Wise. Cubulating small cancellation groups. *Geometric & Functional Analysis GAFA*, 14(1):150–214, 2004.

[Wis21] Daniel T Wise. *The structure of groups with a quasiconvex hierarchy:(ams-209)*. Princeton University Press, 2021.

[Woe94] Wolfgang Woess. Random walks on infinite graphs and groups—a survey on selected topics. *Bulletin of the London Mathematical Society*, 26(1):1–60, 1994.

[Zal23] Abdul Zalloum. Injectivity, cubical approximations and equivariant wall structures beyond cat (0) cube complexes. *arXiv preprint arXiv:2305.02951*, 2023.