5-2011

Discrete Geometric Homotopy Theory and Critical Values of Metric Spaces

Leonard Duane Wilkins
lwilkin2@utk.edu

To the Graduate Council:

I am submitting herewith a dissertation written by Leonard Duane Wilkins entitled "Discrete Geometric Homotopy Theory and Critical Values of Metric Spaces." I have examined the final electronic copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Mathematics.

Conrad P. Plaut, Major Professor

We have read this dissertation and recommend its acceptance:

James Conant, Fernando Schwartz, Michael Guidry

Accepted for the Council:

Carolyn R. Hodges

Vice Provost and Dean of the Graduate School

(Original signatures are on file with official student records.)

To the Graduate Council:

I am submitting herewith a dissertation written by Leonard Duane Wilkins entitled "Discrete Geometric Homotopy Theory and Critical Values of Metric Spaces." I have examined the final electronic copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Mathematics.

Conrad Plaut, Major Professor

We have read this dissertation
and recommend its acceptance:

James Conant

Fernando Schwartz

Michael Guidry

Accepted for the Council:

Carolyn R. Hodges
Vice Provost and Dean of the Graduate School

(Original signatures are on file with offical student records.)

Discrete Geometric Homotopy Theory and Critical Values of Metric Spaces

A Dissertation Presented for the
Doctor of Philosophy
Mathematics
The University of Tennessee, Knoxville

Leonard Duane Wilkins, Jr.
May 2011

Acknowledgments

I am very grateful to the members of my doctoral committee - Dr. Jim Conant, Dr. Fernando Schwarz, and Dr. Michael Guidry - for their advice and assistance. I am also deeply indebted to my family and my close friends - Anne and Duane Wilkins, Jill Wilkins, Christine Kraft, Shannon Mathis, Nicholas Underwood, and many others - for their patience and support during this process. Finally, none of this would have been possible without the guidance and support of my doctoral advisor, Dr. Conrad Plaut. He has been an invaluable source of information, advice, and knowledge. I will forever be thankful for having had the opportunity to work with him.

Abstract

Building on the work of Conrad Plaut and Valera Berestovskii regarding uniform spaces and the covering spectrum of Christina Sormani and Guofang Wei developed for geodesic spaces, the author defines and develops discrete homotopy theory for metric spaces, which can be thought of as a discrete analog of classical path-homotopy and covering space theory. Given a metric space, X, this leads to the construction of a collection of covering spaces of X - and corresponding covering groups - parameterized by the positive real numbers, which we call the ϵ-covers [epsilon-covers] and the ϵ-groups [epsilon-groups]. These covers and groups evolve dynamically as the parameter decreases, changing topological type at specific parameter values which depend on the topology and local geometry of X. This leads to the definition of a critical spectrum for metric spaces, which is the set of all values at which the topological type of the covers change. Several results are proved regarding the critical spectrum and its connections to topology and local geometry, particularly in the context of geodesic spaces, refinable spaces, and Gromov-Hausdorff limits of compact metric spaces. We investigate the relationship between the critical spectrum and covering spectrum in the case when X is geodesic, connections between the geometry of the ϵ-groups [epsilon-groups] and the metric and topological structure of the ϵ-covers [epsilon-covers], as well as the behavior of the ϵ-covers [epsilon-covers] and critical values under Gromov-Hausdorff convergence.

Contents

List of Figures

Chapter 1

Introduction and Background

1.1 Introduction

Spectral analysis has long been a very useful tool for studying objects in mathematics, from the eigenvalue spectrum students encounter in a basic linear algebra course to the operator spectra studied in functional analysis. Several of these spectra have proved useful in the study of geometry, including the eigenvalue spectrum of the Laplacian operator on a Riemannian manifold and the length spectrum of a geodesic space. In recent years, Christina Sormani and Guofang Wei developed a *covering spectrum* for geodesic spaces ([13],[14]), which is somewhat unique among other spectra in that it can be considered both a geometric and topological spectrum. More specifically, the covering spectrum detects fundamental group generators of geodesic spaces, and the values in this spectrum are directly related to the diameters of the loops that generate the fundamental group.

The work of Sormani and Wei grew out of their efforts to show that Gromov-Hausdorff limits of Riemannian manifolds with Ricci curvature uniformly bounded below have categorical universal covers ([12],[13],[14],[15]). This effort is part of another thread of research that has gained momentum in recent years, namely the geometric and topological study of metric spaces that are not necessarily smooth and may have singular structure. Classical differential geometry focuses on the study of smooth manifolds with various metric structures (e.g. Riemannian, Finsler, simplectic, etc.), and, in this context, one has a very rich structure of analytic and algebraic tools at his or her disposal. Recent uses of non-smooth spaces in various engineering and technological applications, however, have prompted the need for a more systematic study of singular geometry and topology. For instance, certain fractals have been used recently as design models for antennae, since such a structure allows a greater length to be enclosed in a more compact space. As another example, network and optimal transport theorists who study the flow of various materials and information between points in a space often use "simple looking" spaces as their underlying structures, but the metrics they employ are not always standard or intuitive and can lead to strange phenomena, such as the space having no rectifiable curves. The Gromov-Hausdorff limits studied by Sormani and Wei make up just a small part of the class of metric spaces that one might consider non-standard, since passing to a limit space may induce singular topology and/or geometry. Their goal of finding universal covers for these spaces is really a special case of the general effort to understand the topology of singular spaces, since the existence of such a cover implies certain topological properties of the space in question.

Following a similar goal, and around the same time, Conrad Plaut and Valera Berestovskii began an effort to generalize the classical universal cover of a topological space in the context of uniform spaces, a general class of spaces that includes, for example, metric spaces and topological groups ([1],[8]). Their efforts actually began several years earlier when they undertook a study of topological groups ([2]). In their work, they made a very interesting and productive use of the structure of uniform spaces, utilizing entourages, inverse limits, and a discretized analog of classical path-homotopy theory to construct what they call the *uniform universal cover* of a *coverable uniform space*. The uniform universal cover need not be a covering space in the traditional sense, but it does possess the lifting and universal properties of the classical universal cover. Moreover, coverable spaces include all geodesic spaces and all compact, connected, locally path-connected topological spaces (also known as compacta), so their work does, indeed, generalize the classical notion of a universal covering space.

The work contained herein grew out of a combination of influences deriving from the efforts of Berestovskii-Plaut and Sormani-Wei. As a tool for illuminating the topology of a geodesic space, the covering spectrum has a few drawbacks, most notably its seeming reliance on the requirement that every fundamental group element of the geodesic space in question have a rectifiable representative, which is not always the case. It also places some strong requirements on the local topology of the space in question, requirements that are not always satisfied by general metric spaces with singular geometry and/or topology. As will be shown, these problems can be eliminated or circumvented by employing the techniques of Berestovskii-Plaut in the context of metric spaces.

We will begin by briefly outlining the constructions of Sormani-Wei and Berestovskii-Plaut, since they will be referenced throughout for the purposes of comparing and contrasting our work to theirs. We will also introduce other basic definitions and concepts that will be needed. We will then introduce the fundamentals of what we call *discrete geometric homotopy theory*, which is essentially a translation of the uniform space methods of Berestovskii-Plaut into the specific case of metric spaces. This will lead us along a line similar to that of Sormani and Wei; given a metric space, X, we will construct a collection of covering spaces, $\{X_\epsilon\}_{\epsilon>0}$, parameterized by the positive real numbers. As the parameter ϵ decreases, these covers may change topological type, and the specific values at which these changes occur will make up the *critical spectrum of* X. We will show that, when X is a geodesic space, our critical spectrum differs from the covering spectrum by only a multiplicative constant. Thus, the critical spectrum does, indeed, generalize the covering spectrum, and, in the geodesic case, both spectra provide the same information. We will then proceed to study the ϵ-covers, themselves. In particular, we will prove some very interesting relationships between the metric and topological properties of the ϵ-covers and their corresponding covering groups. Finally, we will conclude with an investigation of the behavior of the ϵ-covers and the critical spectrum under Gromov-Hausdorff convergence, including structure and existence theorems for Gromov-Hausdorff limits of ϵ-covers and characterizations of the critical values of a limit space in terms of the critical values of the spaces in the sequence.

1.2 Metric Geometry Basics

Here we will introduce the basic definitions and terminology we need from the theory of metric and covering spaces. Given a metric space, (X, d), we will denote the open metric balls of radius ϵ and centered at $x \in X$ by $B(x, \epsilon) = \{y \in X \,:\, d(x, y) < \epsilon\}$. We will denote the *closed balls of radius* ϵ, $\{y \in X \,:\, d(x, y) \leq \epsilon\}$, by $C(x, \epsilon)$. When multiple metrics are given on one set, X, we will distinguish the different metric spaces with the notations (X, d_1), (X, d_2), etc.; otherwise, we will often just refer to a metric space, X, without explicitly listing the metric.

A map, $f : Y \to X$, between metric spaces is an *isometry* if f is surjective and preserves distances (i.e. $d_X(f(y_1), f(y_2)) = d_Y(y_1, y_2)$ for all $y_1, y_2 \in Y$), a *local isometry* if for every $y \in Y$ there is an open ball, $B(y, \epsilon)$, such that f restricted to $B(y, \epsilon)$ is an isometry onto its image in X, and a *uniform local isometry* if there exists some $\epsilon > 0$ such that for every $y \in Y$, f restricted to $B(y, \epsilon)$ is an isometry onto $B(f(y), \epsilon)$.

If X is a topological space, a *covering space of* X is a topological space, Y, and a continuous, surjective map, $f : Y \to X$, with the following property: for every $x \in X$, there is a neighborhood, U, of x, such that $f^{-1}(U)$ is a disjoint collection of open sets in X, each of which is mapped homeomorphically by f onto U. Such a neighborhood, U, is said to be *evenly covered.* When Y and X are metric spaces, we can always find metric balls that are evenly covered. Consequently, we sometimes refer to f, in that case, as a metric covering map.

A *length space* is a metric space, (X, d), with the property that for any two points, $x, y \in X$, the distance $d(x, y)$ is the infimum of the lengths of all rectifiable curves connecting x and y:

$$d(x, y) = \inf\{l(\gamma) \,:\, \gamma \text{ is a curve from } x \text{ to } y\}. \tag{1.1}$$

A *geodesic* in a length space, X, is a rectifiable curve, $\gamma : I \to X$, that is locally *distance minimizing* (here, I is any real interval, not necessarily open or closed). This means that for each $t \in I$, there is some interval $[t - \epsilon, t + \delta] \subset I$ such that $d(\gamma(t - \epsilon), \gamma(t + \delta)) = l(\gamma|_{[t-\epsilon,t+\delta]})$. It follows that this same distance-length relation holds for any subsegment of $\gamma|_{[t-\epsilon,t+\delta]}$, also. If $\gamma : [a, b] \to X$ satisfies $d(\gamma(a), \gamma(b)) = l(\gamma)$, then we will say that γ is a *minimal geodesic*; that is, a minimal geodesic is minimzing on its whole parameter domain. If X is also locally compact and complete, it turns out that the infimum in equation 1.1 is attained, meaning that there is at least one minimal geodesic between any two points.

Definition 1.2.1 *A* ***geodesic space*** *is a length space, (X, d), with the property that for any two points, $x, y \in X$, there is a rectifiable curve, γ, from x to y such that $d(x, y) = l(\gamma)$.*

Thus, any complete, locally compact length space - in particular, any compact length space - is a geodesic space, but there are geodesic spaces that are not complete (e.g. a connected, open interval of $\mathbb{R}$) or locally compact (e.g. an infinite collection of geodesic circles of the same circumference joined at a common point). Geodesic spaces are also metrically characterized by the following property: if a metric space, (X, d), is a geodesic space, then, for every pair of points, $x, y \in X$, there is a point, $z \in X$ - called a *midpoint of x and y* - such that $d(x, z) = d(z, y) = d(x, y)/2$. The converse if this is true if X is also complete. Geodesic spaces include a large class of metric spaces. All Riemannian manifolds, for instance, are length spaces, and all complete Riemannian manifolds are geodesic spaces. Moreover, all complete Gromov-Hausdorff limits (see next definition) of geodesic spaces are also geodesic spaces.

The final piece of background material we need concerns a commonly studied notion of convergence for metric spaces. In the setting of Riemannian manifolds, where there is a richer structure with which to work, there are many notions of convergence. A good survey of the various notions of convergence of manifolds is given in [11]. For general metric spaces, we must work with more general notions of convergence. This generality means that results are sometimes weaker, but they are also more widely applicable. The following formulation of Gromov-Hausdorff convergence borrows terminology and definitions from [3] and [7]. Both are excellent sources on the subject.

Let X be a metric space, and let A and B be subsets of X. The *Hausdorff distance in X between A and B*, denoted by $d_H(A, B)$, is defined to be

$$d_H(A, B) = \inf\{r > 0 \ : \ B \subset T_r(A) \text{ and } A \subset T_r(B)\},$$

where $T_r(A)$ denotes the r-neighborhood of A, or all points $y \in X$ such that there is some point $x \in A$ with $d(x, y) < r$. On the collection of all compact subsets of X, d_H is a metric. If X and Y are two compact metric spaces, then we define the *Gromov-Hausdorff distance between X and Y* by

$$d_{GH}(X, Y) := \inf\{d_H^Z(f(X), g(Y))\},$$

where the infimum is taken over all isometric embeddings, $f : X \to Z$, $g : Y \to Z$, of X and Y into a common metric space, Z, and d_H^Z denotes the Hausdorff distance in Z.

Definition 1.2.2 *A sequence of compact metric spaces, $\{X_n\}_{n=1}^{\infty}$, converges to a compact metric space, X, in the Gromov-Hausdorff sense if $d_{GH}(X_n, X) \to 0$ as $n \to \infty$, and we denote this by $X_n \xrightarrow{GH} X$.*

The function d_{GH} defines an actual metric on the set of compact metric spaces modulo isometry. We will not address any other mode of convergence of metric spaces here, so we will usually suppress the GH notation. A convenient consequence of this definition is the following: if $\{X_n\}$ is a Gromov-Hausdorff convergent sequence and $X_n \to X$, we may assume that X and all X_n are embedded in a common metric space, Z, and we may take the convergence in the Hausdorff sense in Z. See Proposition 43 and the subsequent remarks in [7] for a proof of this.

Definition 1.2.3 *Let X be a metric space, and let $\epsilon > 0$. A set $S \subset X$ is an ϵ-**net** in X if $dist(x, S) < \epsilon$ for every $x \in X$, where $dist(x, S) = \inf\{d(x, y) \ : \ y \in S\}$. We say that X is **totally bounded** if there is a finite ϵ-net in X for every $\epsilon > 0$.*

Lemma 1.2.4 *For compact metric spaces, X and $\{X_n\}$, $X_n \to X$ in the Gromov-Hausdorff sense if and only if the following holds: for every $\epsilon > 0$, there are finite ϵ-nets, $S \subset X$ and $S_n \subset X_n$, such that $S_n \to S$ in the Gromov-Hausdorff sense. Moreover, these ϵ-nets, S_n, can be chosen so that $card(S_n) = card(S)$ for all sufficiently large n.*

The following definition and theorem make up a version of the well-known *precompactness theorem* of M. Gromov. These versions appear in [3].

Definition 1.2.5 *A collection of compact metric spaces, $\mathscr{X}$, is **uniformly totally bounded** if the diameters of the spaces, $X \in \mathscr{X}$, are uniformly bounded above and for every $\epsilon > 0$ there is a natural number, $N(\epsilon)$, such that every $X \in \mathscr{X}$ contains an ϵ-net of $N(\epsilon)$ or fewer points.*

Theorem 1.2.6 *Any uniformly totally bounded class of compact metric spaces, $\mathscr{X}$, is precompact in the Gromov-Hausdorff metric topology, meaning that every sequence, $\{X_n\} \subset \mathscr{X}$, has a convergent subsequence.*

To illustrate the wide utility of this theorem, we point out two precompact classes of Riemannian manifolds that have been widely studied in Riemannian geometry. For any $r, V > 0$, the set of all compact Riemannian manifolds with volume uniformly bounded above by V and injectivity radius uniformly bounded below by r is precompact. Moreover, for any $D > 0$ and any $r \in \mathbb{R}$, the set of all Riemannian manifolds with diameter uniformly bounded above by D and Ricci curvature uniformly bounded below by r is also precompact.

These two examples of precompact classes of metric spaces bring up an interesting and important point. In the previous theorem, it should be noted that the limit, X, of a convergent sequence, $\{X_n\} \subset X$, need not be in $\mathscr{X}$. This is the essence of the *pre*-compact condition. If the class $\mathscr{X}$ is *compact* in the Gromov-Hausdorff topology, then this limit, X, does belong to $\mathscr{X}$. This distinction can be important. For example, a sequence of Riemannian manifolds need not converge to a Riemannian manifold. Let T_n be the torus $S^1 \times S^{1/n}$. Then it is not difficult to - intuitively, at least - see that $T_n \to S^1$, since we collapse one circle while leaving the other fixed. Now, extend this idea to the Riemannian double torus, which would converge to a figure-eight. This limit is not a Riemannian manifold. However, its metric shares the same geometric characterization as the geodesic (Riemannian) metrics on the double torus. This turns out to be true in general.

Theorem 1.2.7 *Let $\{X_n\}$ be a sequence of compact length spaces, and let X be a compact metric space. If $X_n \to X$, then X is a length space.*

Thus, if X is the compact Gromov-Hausdorff limit of a sequence of compact Riemannian manifolds - or, more generally, a sequence of geodesic spaces - then, X, itself, is a geodesic space.

We often need to work with *pointed* spaces where a fixed "point of reference" is specified. A *pointed metric space*, $(X, *)$, is simply a metric space, X, with a fixed base point $* \in X$.

Definition 1.2.8 *Let $\{(X_n, x_n)\}$ be a sequence of compact, pointed metric spaces. We say that $\{(X_n, x_n)\}$ converges to (X, x) in the pointed Gromov-Hausdorff sense if $d_{GH}(X_n, X) + d(x_n, x) \to 0$, where the second term, $d(x_n, x)$, is taken in the Hausdorff context arising from the previously mentioned consequence of the definition (i.e. one considers all of the spaces, X_n and X, as subspaces of a common metric space, Z, and the convergence in the Hausdorff sense).*

In other words, we not only require that the spaces, X_n, converge to X, but we also require that the points, x_n, converge to the point, x.

In the noncompact case, the only mode of convergence that makes real sense is a sort of truncated, pointed convergence.

Definition 1.2.9 *We say that a sequence of noncompact, pointed metric spaces, $\{(X_n, x_n)\}$, converges in the pointed Gromov-Hausdorff sense to the pointed metric space, (X, x), if for every $R > 0$ the closed metric balls $C(x_n, R) \subset X_n$ converge in the pointed sense to $C(x, R) \subset X$.*

Finally, we will conclude this section with the non-compact or pointed analogs of Theorems 1.2.6 and 1.2.7.

Theorem 1.2.10 *Let $(X_n, x_n) \to (X, x)$, where each X_n is a length space and X is complete. Then X is a length space. If, in addition, X is locally compact, then X is a geodesic space.*

Theorem 1.2.11 *Let $\mathscr{X}$ be a class of pointed metric spaces. Then the following are equivalent.*

1) *$\mathscr{X}$ is precompact.*

2) *For every $R, \epsilon > 0$, there exists a natural number, $N(R, \epsilon)$, with the property that for every $(X, x) \in \mathscr{X}$ the ball $B(x, R) \subset X$ contains an ϵ-net of $N(R, \epsilon)$ or fewer points.*

3) For every $R, \epsilon > 0$, there exists a natural number, $N(R, \epsilon)$, with the property that for every $(X, x) \in \mathscr{X}$ the minimal number of ϵ-balls it takes to cover $B(x, R) \subset X$ is less than or equal to $N(R, \epsilon)$.

4) For every $R, \epsilon > 0$, there exists a natural number, $N(R, \epsilon)$, with the property that for every $(X, x) \in \mathscr{X}$, the maximum number of disjoint ϵ-balls that can be fit inside $B(x, R)$ is less than or equal to $N(R, \epsilon)$.

In other words, if - for any $R, \epsilon > 0$ - we can cover all of the balls $B(x, R) \subset X$, $X \in \mathscr{X}$, by a uniform number of ϵ-balls, or if the number of disjoint ϵ-balls we can pack into each ball $B(x, R) \subset X$, $X \in \mathscr{X}$, is uniformly bounded, then $\mathscr{X}$ is precompact.

1.3 Uniform Spaces and Entourage Covers

If topological spaces are the most general setting in which one can discuss continuity of functions, then uniform spaces are the most general setting in which one can talk about *uniform continuity* of functions. They include, as special cases, metric spaces and topological groups. Since we will be concerned exclusively with metric spaces, we will give only a brief introduction to uniform spaces and the constructions of Berestovskii-Plaut. This will motivate the metric space constructions to follow in the next chapter. This particular introduction follows those given in [1] and [8].

Let X be a set, and let E, F be subsets of $X \times X$. We define the product of E and F to be

$$EF := \{(x, z) \in X \times X \ : \ \text{for some } y \in X,\ (x, y) \in E \text{ and } (y, z) \in F\}.$$

The diagonal of X is the set $\Delta X = \{(x, x) \ : \ x \in X\}$.

Definition 1.3.1 *A **uniform space** is a set, X, together with a non-empty collection, $\mathscr{E}$, of subsets of $X \times X$ containing the diagonal, ΔX, and satisfying the following properties.*

1) If E, $F \in \mathscr{E}$, then $E \cap F \in \mathscr{E}$.

2) If $E \in \mathscr{E}$ and $E \subset F$, then $F \in \mathscr{E}$.

3) If $E \in \mathscr{E}$, then $E^t := \{(x, y) \ : \ (y, x) \in E\} \in \mathscr{E}$.

4) For any $E \in \mathscr{E}$, there is some $F \in \mathscr{E}$ such that $F^2 = FF \subset E$.

*The collection $\mathscr{E}$ is called a **uniformity** on X, and, when it is necessary to emphasize the specific uniformity, we will denote a uniform space by $(X, \mathscr{E})$.*

If $E \in \mathscr{E}$ and $(x, y) \in E$, then we say that x and y are E-close. If $E \in \mathscr{E}$ is symmetric (i.e. $E = E^t$), and if $x \in X$, then we define the ***E-ball*** at x to be $B(x, E) := \{y \in X \ : \ (x, y) \in E\} = \{y \in X \ : \ (y, x) \in E\}$. Symmetric sets in $\mathscr{E}$ are called ***entourages***, and we will work primarily with entourages instead of arbitrary elements of $\mathscr{E}$. The reason it suffices to work with entourages is that they form a *base* for a naturally induced uniformity. A base for a uniformity, like a basis for a topology, is a simpler collection of subsets of $X \times X$ from which one can construct a full uniformity. Specifically, a uniformity base is a collection, $\mathscr{B}$, of subsets of $X \times X$ containing the diagonal and satisfying properties 3, 4, and the following additional property: for each E, $F \in \mathscr{B}$, there is some $G \in \mathscr{B}$ such that $G \subset E \cap F$. The uniformity, $\mathscr{E}$, is then formed by taking finite intersections and supersets of elements in $\mathscr{B}$. There is a natural way to topologize uniform spaces: a set $U \subset X$ is said to be open if and only if for every $x \in U$ there is some entourage, $E \in \mathscr{E}$, such that $B(x,E) \subset U$. Finally, if $(X,\mathscr{E})$, $(Y,\mathscr{F})$ are uniform spaces, then $f: X \to Y$ is *uniformly continuous* if for every entourage $F \in \mathscr{F}$ there is an entourage, $E \in \mathscr{E}$, such that $f(E) \subset F$, where we use $f(E)$ as a shorthand notation for $(f \times f)(E)$.

Now, let (X,d) be a metric space, and for each $\epsilon > 0$, let $E_\epsilon = \{(x,y) \in X \ : \ d(x,y) < \epsilon\}$. The symmetry of d implies that each E_ϵ is symmetric. These sets form a base for a uniformity, which we will refer to as the standard metric uniformity. The E-balls corresponding to sets in this base are precisely the metric balls induced by d, and uniform continuity of functions - according to the above definition - reduces to precisely the standard ϵ, δ-definition of uniform continuity. Hence, metric spaces are, very naturally, uniform spaces. Indeed, it is no accident that the terminology of uniform spaces (e.g. *E-ball*) follows closely that of metric spaces.

Let $(X,\mathscr{E})$ be a uniform space, and let E be an entourage in $\mathscr{E}$. An ***E-chain***, α, in X, is a finite, ordered set of points in X, $\alpha = \{x_0, x_1, \ldots, x_n\}$, such that for each $i = 1, \ldots, n$, we have $(x_{i-1}, x_i) \in E$ (intuitively, consecutive points are E-close). The inverse or reverse of an E-chain - because entourages are symmetric - is also an E-chain, and we denote this inverse chain by α^{-1}. If β is another E-chain with initial point equal to the terminal point of α, then the concatenation $\alpha * \beta$ is also an E-chain. We will simply denote this concatenation by the product notation $\alpha\beta$. An *E-loop* is an E-chain with equal initial and terminal points.

We say that X is E-connected if any pair of points, $x, y \in X$, can be joined by an E-chain $\{x = x_0, x_1, \ldots, x_{n-1}, x_n = y\}$. We further say that X is ***chain-connected*** if X is E-connected for every entourage, E. Chain-connectivity is a very mild condition, as we will see when we begin specifically addressing the metric case. On the other hand, it is also an essentially indispensable property if one wishes to carry out the covering space constructions that will be employed here. Thus, it will almost always be assumed that uniform spaces are chain connected, and, in the metric case, we will make this a standing assumption.

Next, we define a homotopy theory - comparable to classical path homotopy theory - for entourage chains. Let $\alpha = \{x_0, x_1, \ldots, x_n\}$ be an E-chain. A ***basic move*** on an E-chain consists of the removal or addition of a single point, with the added conditions that the endpoints remain fixed and that the new E-chain resulting from this basic move is still an E-chain. In other words, one can remove a point, x_i, from α as long as the "surrounding" points, x_{i-1} and x_{i+1}, satisfy $(x_{i-1}, x_{i+1}) \in E$, meaning that the resulting chain, $\alpha' = \{x_0, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n\}$ is still an E-chain. Likewise, one can add a point, x, between x_i and x_{i+1} as long as $(x_i, x) \in E$ and $(x, x_{i+1}) \in E$, so the resulting chain, $\alpha' = \{x_0, \ldots, x_i, x, x_{i+1}, \ldots, x_n\}$ is an E-chain.

Now, let $\beta = \{x_0 = y_0, \ldots, y_m = x_n\}$ be another E-chain beginning and ending at the same point as α, though not necessarily having the same number of points. We say that α and β are ***E-homotopic*** if there is a finite sequence of E-chains,

$$H = \{\alpha = \gamma_0, \gamma_1, \ldots, \gamma_{k-1}, \gamma_k = \beta\}$$

such that, for each $i = 1, \ldots, k$, γ_i differs from γ_{i-1} by a basic move. We call H an E-homotopy between α and β. Basically, this definition states that α and β are E-homotopic if one can transform α into β by removing and adding points in a finite, step-by-step process that maintains the E-chain property throughout and leaves the endpoints fixed. We will have no reason to consider free chain homotopies here, so all chain homotopies will be assumed to be fixed-endpoint homotopies. (However, in [9], the author and C. Plaut use free chain homotopies in the context of studying geodesic spaces.) Now, this does lead to a subtle point that should be addressed. If $\alpha = \{* = x_0, \ldots, x_n = *\}$ is an E-loop that is E-homotopic to the trivial chain $\{* = x_0, x_n = *\}$, then we say that α is *trivial*, *E-nullhomotopic*, or simply *E-null*. Since we leave the endpoints fixed, we cannot technically remove one of the duplicated endpoints to obtain the chain $\{*\}$. Nevertheless, even though this chain has two points, it is, literally for all intents and purposes, the same as the trivial chain $\{*\}$. Thus, we will adopt a slight abuse of notation and often write the two-point constant chain, $\{*, *\}$, as simply $\{*\}$. No confusion should arrive from this.

We will use the notation $\alpha \sim_E \beta$ to denote the relation "α is E-homotopic to β." Just like traditional fixed-endpoint path homotopies, $\sim_E$ is easily seen to be an equivalence relation, and, given an E-chain, α, we will denote the E-equivalence class of α by $[\alpha]_E$. The subscript here is essential, since we will eventually have to consider chains formed from several entourages. Concatenation of E-chains induces a well-defined concatenation operation on E-equivalence classes: if the initial point of β equals the terminal point of α, then, by our definition, the same holds for any chains E-homotopic to α and β. In that case, we define $[\alpha]_E[\beta]_E := [\alpha\beta]_E$. This concatenation of equivalence classes also satisfies an associative property: if α, β, and λ are E-chains, with the terminal point of α agreeing with the initial point of β and the terminal point of β agreeing with the initial point of λ, then we have

$$[\alpha]_E([\beta]_E[\lambda]_E) = [\alpha\beta\lambda]_E = ([\alpha]_E[\beta]_E)[\lambda]_E.$$

The results in the following lemma are not only technically useful, but they show that chain homotopies share many of the same basic properties as path homotopies.

Lemma 1.3.2 *Let X be a uniform space and E an entourage. Let α and β be E-chains in X.*

1) *If the initial point of β is the terminal point of α, and if $\alpha' \sim_E \alpha$ and $\beta' \sim_E \beta$, then $\alpha\beta \sim_E \alpha'\beta'$.*

2) *If $\alpha \sim_E \beta$, then $\alpha^{-1} \sim_E \beta^{-1}$.*

3) *$\alpha\alpha^{-1}$ is E-null.*

4) *If α and β have the same initial and terminal points, and if γ is the E-loop $\alpha\beta^{-1}$, then γ is E-null if and only if $\alpha \sim_E \beta$.*

Proof Part 1 is simply a restatement of the fact that concatenation of E-equivalence classes is well-defined. For part 2, first, suppose α and β differ by only a basic move, say $\alpha = \{x_0, \ldots, x_n\}$ and $\beta = \{x_0, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n\}$ results from removing x_i from α. We know that the reverse of an E-chain is an E-chain, so $\alpha^{-1} = \{x_n, \ldots, x_0\}$ and $\beta^{-1} = \{x_n, \ldots, x_{i+1}, x_{i-1}, \ldots, x_0\}$ are E-chains beginning and ending at the same points. Moreover, β^{-1} clearly differs from α^{-1} by the removal of a single point, so $\alpha^{-1} \sim_E \beta^{-1}$. The case where β is obtained by adding a point to α follows the same reasoning. So, in the general case, let $H = \{\alpha = \gamma_0, \gamma_1, \ldots, \gamma_k = \beta\}$ be an E-homotopy between α and β. Then γ_i differs from γ_{i-1} by a basic move for each i. Thus, applying the previous discussion and the transitivity of $\sim_E$, we have

$$\gamma_0^{-1} \sim_E \gamma_1^{-1} \sim_E \gamma_2^{-1} \sim_E \cdots \sim_E \gamma_{k-1}^{-1} \sim_E \gamma_k^{-1} \Rightarrow \alpha^{-1} \sim_E \beta^{-1}.$$

This proves part 2.

If $\alpha = \{x_0, \ldots, x_n\}$, then $\alpha\alpha^{-1} = \{x_0, \ldots, x_n, x_n, \ldots, x_0\}$. We can remove one of the x_n points, leaving us the chain $\{x_0, \ldots, x_{n-1}, x_n, x_{n-1}, \ldots, x_0\}$. We can then remove the other x_n point, since the surrounding points are equal. Repeating this successively, for $i = n-1, n-$

$2, \ldots, 1$, we can remove every point until we reach the chain $\{x_0, x_0\}$, which is the trivial chain. This proves part 3.

Now, if $\alpha \sim_E \beta$, then part 2 implies that $\alpha^{-1} \sim_E \beta^{-1}$. Likewise, part 1 implies that $\alpha\beta^{-1} \sim_E \alpha\alpha^{-1}$, showing that $\alpha\beta^{-1}$ is E-null. Conversely, suppose $\gamma = \alpha\beta^{-1}$ is E-null. Let $*$ denote the initial point of α, and let $\{*\}$ denote the resulting trivial chain. So, our assumption implies that $\alpha\beta^{-1} \sim_E \{*\}$. Then part 1 implies that $\alpha\beta^{-1}\beta \sim_E \{*\}\beta \sim_E \beta$. But parts 1 and 3 imply that $\alpha\beta^{-1}\beta \sim_E \alpha\{*\} \sim_E \alpha$. Thus, $\alpha \sim_E \beta$, proving the last part. ∎

We will not delve much further into chain homotopy theory in the context of general uniform spaces. Since our primary goal is to focus on metric spaces, we will reserve further discussion of these homotopies to that specific context. Consequently, we have only one more aspect of the Berestovskii-Plaut uniform covering theory to present. This the motivation for the covering space constructions we will present later.

Let X be a uniform space and E an entourage such that X is E-connected. Fix a base point, $* \in X$. Let X_E be the set of all E-equivalence classes of E-chains beginning at $*$, and let $\varphi_E : X_E \to X$ be the endpoint map, taking $[\alpha]_E = [\{* = x_0, x_1, \ldots, x_n\}]_E$ to $x_n \in X$. For reasons which will become clear soon, we call X_E the *E-entourage cover of* X. If we need to emphasize the base point, we will denote this set by $(X_E, \tilde{*})$, where $\tilde{*} = [\{*\}]_E$. There is a special subset of X_E that needs to be singled out. Some E-chains beginning at $*$ will also end at $*$, giving us equivalence classes of E-loops at $*$. Let $\pi_E(X)$ denote this subset. That is,

$$\pi_E(X) = \{[\alpha]_E \in X_E \ : \ \alpha \text{ is an } E\text{-loop at } *\}.$$

Concatenation of equivalence classes gives us a well-defined operation on this set. Since the concatenation of two loops at $*$ gives us another loop at $*$, we have $[\alpha]_E, [\beta]_E \in \pi_E(X) \Rightarrow [\alpha]_E[\beta]_E = [\alpha\beta]_E \in \pi_E(X)$. The equivalence class of the trivial chain, $[\{*\}]_E$, which contains all E-null loops based at $*$, is an identity for this operation. We also know that this operation is associative. Finally, noting that each element, $[\alpha]_E \in \pi_E(X)$, has an inverse element with respect to this operation and the identity, namely $[\alpha^{-1}]_E$ (the inverse of a loop at $*$ is also a loop at $*$), we see that $\pi_E(X)$ is actually a group, which we will call the *E-group.* We will see that, in the metric case, $\pi_E(X)$ can be thought of as a discrete analog to the fundamental group at a specific metric scale.

Lemma 1.3.3 *If X is E-connected, then $\varphi_E : X_E \to X$ is surjective.*

Proof Let $x \in X$ be given. Since X is E-connected, there is an E-chain connecting $*$ and x, say $\alpha = \{* = x_0, \ldots, x_n = x\}$. Then, $[\alpha]_E \in X_E$, and, by definition, we have $\varphi_E([\alpha]_E) = x$. ∎

Now, we can define a uniform, and therefore topological, structure on X_E. For any entourage $D \subseteq E$, define $D^* \subset X_E \times X_E$ as follows: $([\alpha]_E, [\beta]_E) \in D^*$ if and only if

$$([\alpha]_E, [\beta]_E) = ([* = x_0, \ldots, x_n, y]_E, [* = x_0, \ldots, x_n, z]_E) \text{ with } (y, z) \in D.$$

The sets D^*, for all entourages satisfying $D \subseteq E$, form a base for a uniformity on X_E, and each set D^* is, then, an entourage on X_E. The following result is proved as Proposition 16 and Theorem 39 in [1]. This result justifies calling X_E the E-entourage cover of X.

Theorem 1.3.4 *Let X be a uniform space. Suppose that for some entourage, E, on X, the map $\varphi_E : X_E \to X$ is surjective, or, equivalently, that X is E-connected. Then **1)** the map*

$\varphi_E : X_E \to X$ is uniformly continuous, ***2)*** *the restriction of φ_E to any D^*-ball with $D \subseteq E$ is a homeomorphism, and* ***3)*** *φ_E is a regular covering map. In particular, the covering group of φ_E is $\pi_E(X)$, and this group acts discretely and isomorphically on X_E.*

Now, the fact that X_E is an actual covering space of X is a desirable property, but the situation is, in general, not as bright as it might seem. Entourages can be unusual creatures, so, even if X is a "nice" space, X_E may have some strange structure. For instance, it may not even be connected, even if X is connected. We will give examples of this phenomenon later on in the case of metric spaces.

These entourage covers, and the chain homotopy theory that induces them, are what we will use to define our parameterized collection of covering spaces of a metric space. In fact, we've already seen that the entourages $E_\epsilon = \{(x, y) \in X \times X : d(x, y) < \epsilon\}$ induce a uniformity on a metric space, X, for which the E_ϵ-balls are precisely the usual metric balls of radius ϵ. We will be using these specific entourages, E_ϵ, for all $\epsilon > 0$, to define our covers.

1.4 Spanier Covers and the Covering Spectrum

Sormani and Wei also utilized parameterized covering space constructions to define the covering spectrum of a geodesic space. Their methods relied on a classical approach to constructing covering spaces found in the well-known book by Edwin Spanier ([10]). This method requires that the given space be connected and locally path-connected, and the resulting covering spaces are, themselves, connected and locally path-connected. In this section, we will outline the development of the covering spectrum as followed in [12], [13], and [10]. We will also prove some additional properties that are not in those sources; these properties will be used later on to compare our spectrum to the covering spectrum.

Let X be a connected, locally path-connected topological space. Keep in mind that we are working with traditional, continuous paths and path homotopies in this section. Let $\mathscr{U}$ be an open covering of X, and fix a basepoint, $*$, in X. Define $\pi_1(X, *, \mathscr{U})$ to be the subgroup of $\pi_1(X, *)$ consisting of all (fixed endpoint) path homotopy classes of loops at $*$ that are homotopic to a finite product of loops of the form $\alpha\beta\alpha^{-1}$, where β is a loop lying in some $U \in \mathscr{U}$ and α is a path from $*$ to the initial point of β. See Figure 1.1 below. Given a loop, γ, at $*$, we let $[\gamma]$ denote its fixed endpoint path homotopy equivalence class. Note, also, that $\pi_1(X, *, \mathscr{U})$ is a normal subgroup of the fundamental group. In fact, if λ is any other path loop at $*$, then

$$\lambda(\alpha_1\beta_1\alpha_1^{-1}) \cdots (\alpha_k\beta_k\alpha_k^{-1})\lambda^{-1}$$

is path-homotopic to

$$(\lambda\alpha_1\beta_1\alpha_1^{-1}\lambda^{-1})(\lambda\alpha_2\beta_2\alpha_2^{-1}\lambda^{-1}) \cdots (\lambda\alpha_k\beta_k\alpha_k^{-1}\lambda^{-1}).$$

This is a product of loops that are still of the desired form; all that has changed is that the paths from $*$ to the initial point of the loops, β_i, are now of the form $\lambda\alpha_i$. It is also evident that if $\mathscr{V}$ is a refinement of $\mathscr{U}$ - meaning that, for each $V \in \mathscr{V}$, there is some $U \in \mathscr{U}$ such that $V \subset U$ - then $\pi_1(X, *, \mathscr{V}) \subset \pi_1(X, *, \mathscr{U})$.

Drawing inspiration from the work of Berestovskii-Plaut, we will be interested in the case in which $\mathscr{U}$ is a *uniform* open covering, roughly meaning that there is some entourage, E, such that the sets, $U \in \mathscr{U}$, are of the form $U = B(x, E)$. Sormani and Wei used precisely a construction of this type, but without explicitly referencing any aspect of the theory of uniform spaces.

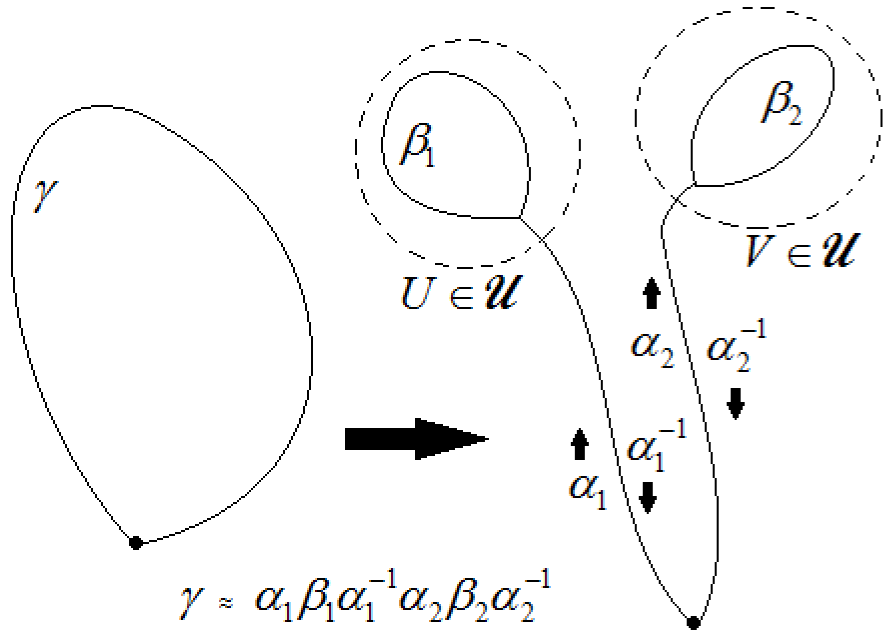

Figure 1.1: Loops generating $\pi_1(X, *, \mathscr{U})$

Now, let $\mathscr{P}$ be the set of all paths in X beginning at $*$. We may assume that all paths are parameterized on the interval $[0, 1]$, so that $\alpha(1)$ always denotes the terminal point of the path α. We define an equivalence relation on $\mathscr{P}$ by setting $\alpha \sim \beta$ if and only if $\alpha(1) = \beta(1)$ and $[\alpha\beta^{-1}] \in \pi_1(X, *, \mathscr{U})$. We define $X^{\mathscr{U}} = \mathscr{P}/\sim$, so that $X^{\mathscr{U}}$ is the set of all equivalence classes of paths beginning at $*$ under this relation. To distinguish between the path homotopy equivalence class of a path, α, at $*$ and the equivalence class of α under $\sim$, we will use $[\alpha]$ for the former and $\overline{\alpha}$ for the latter.

Define $\psi : X^{\mathscr{U}} \to X$ to be the endpoint map. In other words, given $\overline{\alpha} \in X^{\mathscr{U}}$, $\psi(\overline{\alpha}) = \alpha(1)$, the endpoint of α. Since all paths in $\overline{\alpha}$ end at the same point, this map is well-defined. Moreover, since X is path-connected, ψ is surjective.

Theorem 1.4.1 *$X^{\mathscr{U}}$ can be topologized so that it is path-connected and the endpoint map, ψ, is a regular covering map with corresponding covering group equal to $\pi_1(X, *, \mathscr{U})$.*

Proof The proof follows the standard construction of covering spaces in classical topology, so we will just outline it here. First, given a path $\alpha \in \mathscr{P}$ and a path-connected neighborhood, V, of $\alpha(1)$, we define the subsets of $X^{\mathscr{U}}$

$$B(V, \alpha) = \{\overline{\alpha\gamma} \,:\, \gamma \text{ is a path in } V \text{ beginning at } \alpha(1)\}.$$

One then shows that the collection of all such sets, ranging over all paths in $\mathscr{P}$ and path-connected neighborhoods of their endpoints, forms a basis for a topology. This topologizes $X^{\mathscr{U}}$. It is then straightforward to show that ψ is continuous and open, and that any path-connected neighborhood of $x \in X$ that lies in a set $U \in \mathscr{U}$ is evenly covered. Thus, ψ is a covering map.

The fact that $X^{\mathscr{U}}$ is locally path-connected now follows because it is locally homeomorphic to X. To show that $X^{\mathscr{U}}$ is path-connected, one first shows that the lift of a path, α, beginning at $*$ is given by $\hat{\alpha} : [0,1] \to X^{\mathscr{U}}$, where $\hat{\alpha}(t)$ is defined to be the equivalence class of $\alpha|_{[0,t]}$ under the relation $\sim$. With this characterization at hand, it follows that, given $\overline{\alpha} \in X^{\mathscr{U}}$, the lift of α to $\bar{*} \in X^{\mathscr{U}}$ - where $\bar{*}$ is the point of $X^{\mathscr{U}}$ containing the constant path at $*$ - is a path from $\bar{*}$ to $\overline{\alpha}$. Thus, $X^{\mathscr{U}}$ is path-connected. Finally, one shows that $\psi_*(\pi_1(Y, \bar{*})) = \pi_1(X, *, \mathscr{U})$ by showing that a loop, γ, at $*$ lifts closed to $X^{\mathscr{U}}$ if and only if $[\gamma] \in \pi_1(X, *, \mathscr{U})$, and the fact that ψ is regular follows because $\pi_1(X, *, \mathscr{U})$ is a normal subgroup of $\pi_1(X, *)$. ∎

We call the covering space defined in the previous theorem the *Spanier cover of X determined by the open covering $\mathscr{U}$*.

For the rest of this section, let X be a geodesic space. Given a fixed base point, $* \in X$, and $\delta > 0$, we define X^δ to be the Spanier cover of X determined by the open covering of X consisting of all balls of radius δ in X. Likewise, we let $\pi_1(X, *, \delta)$ be the corresponding covering group. Thus, the following are equivalent: **1)** $[\gamma] \in \pi_1(X, *, \delta)$, **2)** γ lifts as a closed loop to X^δ, and **3)** γ is homotopic to a product of loops of the form $\alpha\beta\alpha^{-1}$, where β is a loop lying in a ball of radius δ and α is a path from $*$ to the initial point of β. We let $\psi_\delta : X^\delta \to X$ denote the covering map, and we will refer to X^δ as the *δ-Spanier cover of* X, or simply just the δ-cover when no confusion will arise. This gives us a collection of covering spaces parameterized by $\mathbb{R}_+$.

It is well-known that if Y is a connected covering space of a length space, Z, then the length metric on Z can be lifted to Y so that Y is a length space and the covering map is a local isometry. Moreover, if Z is complete and locally compact, then Y is, also, making Y a geodesic space. Thus, each δ-cover can be naturally endowed with the lifted length metric from X. In fact, recalling the construction of X^δ in Theorem 1.4.1, each δ-ball in X is evenly covered. It can be shown, then, that this makes the covering map, ψ_δ, a bijection and radial isometry from δ-balls in Y onto δ-balls in X and an isometry from open $\delta/2$-balls in X^δ onto open $\delta/2$-balls in X. For the remainder of this section, we will assume that the geodesic space, X, is also locally compact and complete. Thus, each X^δ is a geodesic space.

Definition 1.4.2 *The **Covering Spectrum** of X is the set of all $\delta > 0$ such that $X^\delta \neq X^{\delta'}$ for all $\delta' > \delta$. We denote the set of all such positive values by $Cov(X)$.*

Intuitively, the δ-covers detect topology within X at specific metric scales. A more vivid description would be to say that the δ-cover unravels holes of diameter at least $\delta/2$ in X. The covering spectrum contains those δ-values where new topology is detected. If one pictures this process in a dynamic sense, with the δ-covers evolving as δ-decreases from ∞ to 0, then once an element of $Cov(X)$ is reached, a new loop is unraveled. In other words, if $\delta < \epsilon$, and if $X^\delta \neq X^\epsilon$, then there is a loop in X that lifts closed to X^ϵ but open to X^δ. Consider the following example.

Example 1.4.3 *Let S^1 be the circle of circumference 1 with its natural geodesic metric, and let $*$ be any fixed point in S^1. The fundamental group of S^1 is generated by the loop that traverses the circle once. This loop has diameter $1/2$; the furthest it extends away from $*$ is the antipodal point of $*$, which is a distance of $1/2$ away from $*$. Thus, this fundamental loop, anchored at $*$, lies in any ball of radius $\delta > 1/2$ centered at $*$. The same holds for any multiple of this loop or its inverse. So, if $\delta > 1/2$, every element of $\pi_1(S^1)$ can be expressed as a product of loops of the form $\alpha\beta\alpha^{-1}$, where β is a loop lying in a ball of radius δ and α is a path from $*$ to the initial point of β. (In fact, α is trivial in this case.) In other words, $\pi_1(S^1, *, \delta)$ is just the fundamental* *group for such δ. From the basic theory of covering spaces, it follows that, in this case, $(S^1)^\delta$ is the trivial cover: $(S^1)^\delta = S^1$.*

On the other hand, it is evident that this fundamental loop - and, thus, any multiple of it or its inverse - does not lie in any open ball of radius $\delta \leq 1/2$. That is, for $0 < \delta \leq 1/2$, no non-trivial element of $\pi_1(S^1)$ is in $\pi_1(S^1, *, \delta)$. It follows that the covering group of $(S^1)^\delta$ in this case is the trivial group, meaning that $(S^1)^\delta$ is a simply connected cover of S^1. Thus, for $0 < \delta \leq 1/2$, $(S^1)^\delta = \mathbb{R}$. Since the covers change from the trivial cover to the universal cover at $\delta = 1/2$, $Cov(S^1) = \{1/2\}$. ∎

Lemma 1.4.4 *The δ-covers of a geodesic space, X, are monotone in the following sense: if $\epsilon < \delta$, then X^ϵ covers X^δ.*

This lemma is proved in [13], but it follows immediately from the fact that $\pi_1(X, *, \epsilon) \subset \pi_1(X, *, \delta)$ when $\epsilon < \delta$. Hence, we can lift the map $\psi_\epsilon : X^\epsilon \to X$ to X^δ, and the lift is necessarily a covering map.

Many of the following results on the covering spectrum are actually not explicitly in [13], but these particular statements will prove useful for our efforts later on. Thus, we collect them here, now. The first lemma and corollary are really just restatements of the definition.

Lemma 1.4.5 *The following are equivalent.*

1. *$\delta \in Cov(X)$.*
2. *For every $\epsilon > 0$, there is some $\delta' \in (\delta, \delta + \epsilon)$ such that $X^\delta \neq X^{\delta'}$.*
3. *There is a sequence $\{\delta_n\}$, with $\delta_n \searrow \delta$, such that $X^\delta \neq X^{\delta_n}$ for all n.*

Corollary 1.4.6 *If $\delta > 0$, then $\delta \notin Cov(X)$ if and only if there is some interval $[\delta, \delta + \epsilon)$ such that $X^\delta = X^{\delta'}$ for all $\delta' \in [\delta, \delta + \epsilon)$.*

Lemma 1.4.7 (Lower Semi-continuity of the Covering Spectrum) *If $\{\delta_i\}_{i=1}^\infty \subset Cov(X)$ and $\delta_i \searrow \delta > 0$, then $\delta \in Cov(X)$.*

Proof Let $\delta' > \delta$ be given. Choose δ_i such that $\delta \leq \delta_i < \delta'$. Since $\delta_i \in Cov(X)$, $X^{\delta_i} \neq X^{\delta'}$. Thus, there is a loop, α, that lifts closed to $X^{\delta'}$ but open to X^{δ_i}. α must lift open to X^δ, since X^δ covers X^{δ_i}. Thus, $X^\delta \neq X^{\delta'}$. Since δ' was arbitrary, this shows that $\delta \in Cov(X)$. ∎

Lemma 1.4.8 *Suppose $\delta > 0$ is not in $Cov(X)$, and let $\epsilon = \inf\{\lambda > \delta \,:\, \lambda \in Cov(X)\}$. If this set is empty, we set $\epsilon = \infty$. Then $\epsilon > \delta$ and $X^\delta = X^{\delta'}$ for all $\delta' \in [\delta, \epsilon]$.*

Proof First, note that - if the set above is nonempty - $\epsilon \in Cov(X)$, because we can find a sequence in $Cov(X)$ converging down to ϵ. So, ϵ is the smallest element of $Cov(X)$ larger than δ. If the set is empty, then there are no elements of $Cov(X)$ larger than δ. The fact that $\delta < \epsilon$ is clear, for if not, then $\delta = \epsilon$, implying that δ is an element of the covering spectrum, contradicting our hypothesis. So, the important conclusion of this result that we must sill prove is that the covers do not change between δ and ϵ.

Since $\delta \notin Cov(X)$, there is some nonempty interval, $[\delta, t) \subseteq [\delta, \epsilon)$, such that $X^\delta = X^{\delta'}$ for all $\delta' \in [\delta, t)$. Let

$$S = \{t \in (\delta, \epsilon) \,:\, X^\delta = X^{\delta'} \text{for all } \delta' \in [\delta, t)\}$$

and let $\delta^* = \sup S$. Then S is nonempty, $\delta < \delta^* \leq \epsilon$, and $X^\delta = X^{\delta'}$ for all $\delta' \in [\delta, \delta^*)$. We need to show that $\delta^* = \epsilon$.

Suppose, toward a contradiction, that $\delta^* < \epsilon$. Then X^δ cannot equal X^{δ^*}; if it did, then, since $\delta^* \notin Cov(X)$, there would be some interval $[\delta^*, t)$, with $\delta^* < t < \epsilon$, such that $X^{\delta^*} = X^r$ for all $r \in [\delta^*, t)$, implying that $X^\delta = X^r$ for all $r \in [\delta, t)$. This would contradict that $\delta^* = \sup S$.

So, still assuming that $\delta^* < \epsilon$, we have $X^\delta = X^{\delta'}$ for all $\delta' \in [\delta, \delta^*)$, but $X^\delta \neq X^{\delta^*}$. For a fixed basepoint, $* \in X$, this means that the groups $\pi_1(X, *, \delta')$, for $\delta \leq \delta' < \delta^*$, are all equal, but that this group is a proper subgroup of $\pi_1(X, *, \delta^*)$. Let $[\gamma]$ be an element of $\pi_1(X, *, \delta^*)$ that is not in $\pi_1(X, *, \delta')$ for any $\delta' \in [\delta, \delta^*)$. Then there is a representative of this group element, say γ, that is a product of loops, $\alpha_1\beta_1\alpha_1^{-1} \cdots \alpha_m\beta_m\alpha_m^{-1}$, where each β_i is a loop lying in a ball of radius δ^*. But, since $[\gamma] \notin \pi_1(X, *, \delta')$ for any $\delta' \in [\delta, \delta^*)$, γ is not homotopic to any such product with β-loops lying in balls of radius δ' for any $\delta \leq \delta' < \delta^*$. However, each β_i is a compact set lying in an open ball of radius δ^*. This means that for each i, there is some r_i, $0 < r_i < \delta^*$, such that β_i actually lies in an open ball of radius $\delta^* - r_i$. Let $r = \min\{r_i\}$. Then $0 < r < \delta^*$, and each β_i lies in an open ball of radius $\delta^* - r$. Thus, γ is homotopic to a product of loops, $\alpha_1\beta_1\alpha_1^{-1} \cdots \alpha_m\beta_m\alpha_m^{-1}$, with each β_i lying in a ball of radius strictly less than δ^*. This contradicts the fact that $\gamma \notin \pi_1(X, *, \delta')$ for all $\delta' \in [\delta, \delta^*)$. Hence, our assumption that $\delta^* < \epsilon$ cannot hold. This shows that $\delta^* = \epsilon$, and $X^\delta = X^{\delta'}$ for $\delta' \in [\delta, \epsilon)$.

Finally, the same argument we just used shows that $X^\delta = X^\epsilon$. ∎

Lemma 1.4.9 *If $0 < \delta < \epsilon$ and $Cov(X) \cap [\delta, \epsilon] = \emptyset$, then $X^\delta = X^\epsilon$.*

Proof Let $\delta^* = \inf\{\lambda > \delta \,:\, \lambda \in Cov(X)\}$. Then $\delta^* \in Cov(X)$, and δ^* is the smallest element of $Cov(X)$ larger than δ. Since there are no elements of $Cov(X)$ in $[\delta, \epsilon]$, we must have $\delta^* > \epsilon$. By the previous lemma, $X^\delta = X^{\delta'}$ for all $\delta' \in [\delta, \delta^*]$, and this interval includes $[\delta, \epsilon]$. ∎

Lemma 1.4.10 *If $X^\delta \neq X^\epsilon$ for $0 < \delta < \epsilon$, then $Cov(X) \cap [\delta, \epsilon] \neq \emptyset$.*

Proof This is just the contrapositive of the previous Lemma. ∎

Lemma 1.4.11 *If $\delta > 0$ and $\delta \notin \overline{Cov(X)}$, there is some interval $(\delta - \epsilon, \delta + \epsilon) \subset \mathbb{R}_+$ such that $X^{\delta'} = X^\delta$ for all $\delta' \in (\delta - \epsilon, \delta + \epsilon)$.*

Proof Since $\delta \notin Cov(X)$, there is an interval $[\delta, \delta + \epsilon_1)$ such that $X^\delta = X^{\delta'}$ for all $\delta' \in [\delta, \delta + \epsilon_1)$. We claim that there is also an interval $(\delta - \epsilon_2, \delta]$ such that $X^\delta = X^{\delta'}$ for all $\delta' \in (\delta - \epsilon_2, \delta]$. Suppose, toward a contradiction, that no such interval existed. Then there is some $\lambda_1 \in (\delta - 1, \delta)$ such that $X^{\lambda_1} \neq X^\delta$. By Lemma 1.4.10, there is an element, $\delta_1 \in [\lambda_1, \delta]$ such that $\delta_1 \in Cov(X)$. Since $\delta \notin Cov(X)$, clearly we have $\delta_1 < \delta$. Likewise, there is some $\lambda_2 \in (\max\{\delta - 1/2, \delta_1\}, \delta)$ such that $X^{\lambda_2} \neq X^\delta$. Again, by Lemma 1.4.10, there is some $\delta_2 \in [\lambda_2, \delta]$ such that $\delta_2 \in Cov(X)$. Since $\delta \notin Cov(X)$, we have $\delta_1 < \delta_2 < \delta$ and $\delta - 1/2 < \delta_2 < \delta$. Continuing inductively, we get a sequence, $\{\delta_i\}$, such that $\delta_i \nearrow \delta$ and $\delta_i \in Cov(X)$ for all i. This means that $\delta \in \overline{Cov(X)}$, contradicting our hypothesis. So, there is some interval $(\delta - \epsilon_2, \delta]$ such that $X^\delta = X^{\delta'}$ for all $\delta' \in (\delta - \epsilon_2, \delta]$. ∎

Chapter 2

Discrete Homotopy Theory and the Critical Spectrum

2.1 Basic Definitions

In this section, we will develop what we call *discrete geometric homotopy theory.* It is here that we will translate the results from section 1.3 into the context of metric spaces. Specifically, we will work with a metric space, (X, d), and follow through the constructions described in that section. Since the results of section 1.3 hold for any entourage, E, in a uniform space, they obviously hold for the entourage $E_\epsilon = \{(x, y) \in X \times X \, : \, d(x, y) < \epsilon\}$.

Let X be a metric space, and let $\epsilon > 0$ be given. An ***ϵ-chain***, α, is a finite sequence of points, $\alpha = \{x_0, x_1, \ldots, x_{n-1}, x_n\}$, such that $d(x_{i-1}, x_i) < \epsilon$ for $i = 1, \ldots, n$. An *ϵ-loop* is an ϵ-chain with equal initial and terminal points. We say that X is ***ϵ-connected*** if any pair of points, $x, y \in X$, can be joined by some ϵ-chain. If X is ϵ-connected for every $\epsilon > 0$, then we say that X is ***chain-connected***. Chain-connectivity is a rather mild condition.

Lemma 2.1.1 *Let X be a metric space.*

1) If X is connected, then X is chain-connected.

2) If X is chain-connected and compact, then X is connected.

Proof Suppose X is connected, and fix a point, $* \in X$. Let $\epsilon > 0$ be given, and let S be the set of all $x \in X$ such that there is an ϵ-chain connecting $*$ to x. S contains $*$, so it is nonempty. If $x \in S$, then any $y \in B(x, \epsilon)$ is also in S, since we can take an ϵ-chain connecting $*$ to x and then add y to the end of this chain to get an ϵ-chain connecting $*$ and y. Thus, S is open. Now, suppose $\{x_n\} \subset S$ and $x_n \to x$. Choose n large enough so that $d(x_n, x) < \epsilon$. Then we can take an ϵ-chain connecting $*$ to x_n and then add x onto the end of this chain to get an ϵ-chain from $*$ to x. Thus, $x \in S$, showing that S is closed. Since X is connected, it follows that $S = X$. Since $\epsilon > 0$ was arbitrary, we see that X is chain-connected.

Next, assume the hypotheses of part 2. Suppose, toward a contradiction, that X is not connected. Then there are open sets, $A, B \subset X$, such that $A \cap B = \emptyset$ and $A \cup B = X$. Since X is compact, the closed sets $A^c = B$ and $B^c = A$ are compact. Hence, we must have $dist(A, B) > 0$. If we choose $\epsilon < dist(A, B)$, then no point in A can be joined by an ϵ-chain to a point in B, contradicting that X is chain-connected. Thus, X is connected. ∎

Note that the converse of part 1 of the above lemma is not true; the rationals, with their subspace metric inherited from $\mathbb{R}$, are chain-connected but not connected.

Remark As we mentioned in section 1.3, even though chain-connectivity is a relatively weak condition, it is essential for most (but, admittedly, not all) of the constructions that follow. Hence, we make the following standing assumption: *from here on, it will be assumed - even when it is not specifically mentioned - that all metric spaces are chain-connected.*

Following our previous description, a ***basic move*** on an ϵ-chain is the addition or removal of a single point with the conditions that the endpoints remain fixed and the resulting chain is still an ϵ-chain. Thus, if $\alpha = \{x_0, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_n\}$ is an ϵ-chain, we can remove x_i to obtain the chain $\alpha' = \{x_0, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n\}$ as long as the surrounding points, x_{i-1}, x_{i+1}, satisfy $d(x_{i-1}, x_{i+1}) < \epsilon$. Likewise, we can add a point, x, between x_i and x_{i+1} to obtain the chain $\alpha' = \{x_0, \ldots, x_i, x, x_{i+1}, \ldots, x_n\}$ as long as $d(x_i, x) < \epsilon$ and $d(x, x_{i+1}) < \epsilon$. Two ϵ-chains, α and β, with the same initial and terminal points, are ϵ-***homotopic*** if there is a finite sequence of ϵ-chains

$$H = \{\alpha = \gamma_0, \gamma_1, \ldots, \gamma_{k-1}, \gamma_k = \beta\}$$

such that each γ_i differs from γ_{i-1} by a basic move. We call H an ϵ-homotopy between α and β. If an ϵ-loop, $\alpha = \{* = x_0, \ldots, x_n = *\}$, is ϵ-homotopic to the trivial loop, $\{*\}$, then we say that α is *trivial*, ϵ-*nullhomotopic*, or ϵ-*null*.

The notation $\alpha \sim_\epsilon \beta$ will denote the relation "α is ϵ-homotopic to β." As in the general case, this is an equivalence relation, and we denote the ϵ-equivalence class of an ϵ-chain, α, by $[\alpha]_\epsilon$. We have a well-defined concatenation operation on equivalence classes of ϵ-chains: if the initial point of β equals the terminal point of α, then we can define $[\alpha]_\epsilon[\beta]_\epsilon = [\alpha\beta]_\epsilon$. As before, this operation satisfies an associative property. The following lemma is the metric analog of Lemma 1.3.2. The proof of that lemma obviously implies this result.

Lemma 2.1.2 *Let X be a metric space, and let α and β ϵ-chains in X for some $\epsilon > 0$.*

1) *If the initial point of β is the terminal point of α, and if $\alpha' \sim_\epsilon \alpha$ and $\beta' \sim_\epsilon \beta$, then $\alpha\beta \sim_\epsilon \alpha'\beta'$.*

2) *If $\alpha \sim_\epsilon \beta$, then $\alpha^{-1} \sim_\epsilon \beta^{-1}$.*

3) *$\alpha\alpha^{-1}$ is ϵ-null.*

4) *If α and β have the same initial and terminal points, and if γ is the ϵ-loop $\alpha\beta^{-1}$, then γ is ϵ-null if and only if $\alpha \sim_\epsilon \beta$.*

There is a particularly important type of ϵ-homotopy that will be utilized frequently. If α is an ϵ-chain, $\delta < \epsilon$, and if α is ϵ-homotopic to a δ-chain, then we say that α ***can be ϵ-refined to a δ-chain***. We call such a homotopy an ϵ-***refinement***, and we refer to β as a δ-*refinement* - or simply a refinement - of α. As we move on, we will see that the ability, or inability, to refine chains is an important issue that we must deal with. In some spaces, this is not a problem. If α is an ϵ-chain in a geodesic space, for instance, then we can refine α to any desired degree by successively adding midpoints between each pair of consecutive points in the chain. There are other spaces, however, where refining chains will be less straightforward, if even possible.

Lastly, we finish up the introductory material with a result on the projections of chains and homotopies from one metric space to another.

Lemma 2.1.3 *Let $f : Y \to X$ be a Lipschitz map with Lipschitz constant, λ. Given a chain, $\alpha = \{y_0, \dots, y_n\}$, in Y, we define $f(\alpha) := \{f(y_0), \dots, f(y_n)\}$ in X. If α is an ϵ-chain in Y, then $f(\alpha)$ is a $(\lambda\epsilon)$-chain in X. If α and β are ϵ-chains in Y that are ϵ-homotopic, then $f(\alpha)$ and $f(\beta)$ are $(\lambda\epsilon)$-homotopic. In particular, if f is 1-Lipschitz, or a contraction, then ϵ-chains and homotopies in Y project to ϵ-chains and homotopies in X.*

Proof If $\alpha = \{y_0, \dots, y_n\}$ is an ϵ-chain, then

$$d(f(y_{i-1}), f(y_i)) \leq \lambda d(y_{i-1}, y_i) < \lambda\epsilon,$$

proving the first conclusion. To prove the conclusion regarding homotopies, it suffices to consider the case in which β differs from α by a basic move. The reason for this is that, in the general case, if $H = \{\alpha = \gamma_0, \dots, \gamma_k = \beta\}$ is an ϵ-homotopy taking α to β, then each γ_i differs from γ_{i-1} by a basic move. So, if $\gamma_{i-1} \sim_\epsilon \gamma_i \Rightarrow f(\gamma_{i-1}) \sim_{\lambda\epsilon} f(\gamma_i)$, then $f(H) := \{f(\alpha) = f(\gamma_0), \dots, f(\gamma_k) = f(\beta)\}$ will be a $(\lambda\epsilon)$-homotopy between $f(\alpha)$ and $f(\beta)$. So, suppose β is obtained by removing a point from α, say $\alpha = \{y_0, \dots, y_{i-1}, y_i, y_{i+1}, \dots, y_n\}$ and $\beta = \{y_0, \dots, y_{i-1}, y_{i+1}, \dots, y_n\}$. Then $d(y_{i-1}, y_{i+1}) < \epsilon$ and

$$d(f(y_{i-1}), f(y_{i+1})) \leq \lambda d(y_{i-1}, y_{i+1}) < \lambda\epsilon.$$

Thus, $f(\beta) = \{f(y_0), \dots, f(y_{i-1}), f(y_{i+1}), \dots, f(y_n)\}$ is obtained by removing $f(y_i)$ from $f(\alpha)$, and we have $f(\beta) \sim_{\lambda\epsilon} f(\alpha)$. Similar reasoning holds in the case where β is obtained by adding a point to α. Finally, the last conclusion follows by simply letting $\lambda = 1$. ∎

2.2 The ϵ-covers

Next, we will translate the entourage covers, X_E, into the metric context. Fix a base point, $* \in X$, and $\epsilon > 0$. Let X_ϵ be the set of all ϵ-equivalence classes of ϵ-chains in X beginning at $*$, and let $\varphi_\epsilon : X_\epsilon \to X$ be the endpoint map. Thus, X_ϵ is simply the set X_{E_ϵ}, where E_ϵ is the standard metric entourage, $E_\epsilon = \{(x, y) \in X \times X \ : \ d(x, y) < \epsilon\}$. Since X is chain-connected, Lemma 1.3.3 implies that φ_ϵ is surjective. Just like in the general case, the set of equivalence classes of ϵ-loops based at $*$ forms the ***ϵ-group***, $\pi_\epsilon(X) \subset X_\epsilon$.

We first define a metric on X_ϵ.

Definition 2.2.1 *The length of an ϵ-chain, $\alpha = \{x_0, \dots, x_n\}$, is defined by*

$$L(\alpha) := \sum_{i=1}^{n} d(x_{i-1}, x_i).$$

The length of an equivalence class of ϵ-chains, $[\alpha]_\epsilon \in X_\epsilon$, is defined to be

$$L([\alpha]_\epsilon) := \inf\{L(\beta) \ : \ \beta \in [\alpha]_\epsilon\}.$$

Note that if the initial point of β is the terminal point of α, then $L(\alpha\beta) = L(\alpha) + L(\beta)$. We also have $L(\alpha) = L(\alpha^{-1})$ for any ϵ-chain, α. Moreover, this implies that $L([\alpha^{-1}]_\epsilon) = L([\alpha]_\epsilon)$.

Lemma 2.2.2 *If X is a metric space and $\epsilon > 0$, then the following hold.*

1) *If $L([\alpha]_\epsilon) < \epsilon$, where $\alpha = \{x_0, \dots, x_n\}$, then $d(x_0, x_n) < \epsilon$, α is ϵ-homotopic to the two-point ϵ-chain $\{x_0, x_n\}$, and $L([\alpha]_\epsilon) = d(x_0, x_n)$.*

2) If α is an ϵ-loop with $L([\alpha]_\epsilon) < \epsilon$, then α is ϵ-null.

3) If α and β are ϵ-chains with the initial point of β equal to the terminal point of α, then $L([\alpha\beta]_\epsilon) \leq L([\alpha]_\epsilon) + L([\beta]_\epsilon)$.

Proof Assume the hypotheses of part 1. Then, by definition, there is some representative, $\alpha' = \{x_0 = y_0, y_1, \ldots, y_{m-1}, y_m = x_n\} \in [\alpha]_\epsilon$ such that $L(\alpha') < \epsilon$. This means that $\sum_{i=1}^{m} d(y_{i-1}, y_i) < \epsilon$. First, this and the triangle inequality imply that $d(x_0, x_n) = d(y_0, y_n) \leq \sum_{i=1}^{m} d(y_{i-1}, y_i) < \epsilon$. Furthermore, we have

$$d(y_0, y_2) \leq d(y_0, y_1) + d(y_1, y_2) \leq L(\alpha') < \epsilon.$$

This means that we can remove y_1 from α' to obtain the chain $\{y_0, y_2, \ldots, y_m\}$. But we also have

$$d(y_0, y_3) \leq d(y_0, y_1) + d(y_1, y_2) + d(y_2, y_3) \leq L(\alpha') < \epsilon.$$

Thus, we can then remove y_2 to obtain the chain $\{y_0, y_3, \ldots, y_m\}$. We can continue in this way, removing each y_i until we remove y_{m-1}, leaving us with $\{x_0 = y_0, y_m = x_n\}$, showing that α is homotopic to this two-point chain. By definition, we have $L([\alpha]_\epsilon) \leq L(\{x_0, x_n\}) = d(x_0, x_n)$. On the other hand, by the triangle inequality, any other chain in $[\alpha]_\epsilon$ must have length at least $d(x_0, x_n)$, so $d(x_0, x_n) \leq L([\alpha]_\epsilon)$. This proves part 1. Part 2 is an immediate consequence of part 1, since the loop assumption in part 2 just implies that $x_0 = x_n$ in part 1.

For part 3, let β' be a fixed but arbitrary representative of $[\beta]_\epsilon$, and let α' be any representative of $[\alpha]_\epsilon$. Then $\alpha'\beta' \in [\alpha\beta]_\epsilon$, so $L([\alpha\beta]_\epsilon) \leq L(\alpha'\beta') = L(\alpha') + L(\beta')$, which implies that $L([\alpha\beta]_\epsilon) - L(\beta') \leq L(\alpha')$. Since $\alpha' \in [\alpha]_\epsilon$ was arbitrary, we can take the infimum of the right-hand side, yielding $L([\alpha\beta]_\epsilon) - L(\beta') \leq L([\alpha]_\epsilon)$. We now have $L([\alpha\beta]_\epsilon) - L([\alpha]_\epsilon) \leq L(\beta')$, but $\beta' \in [\beta]_\epsilon$ was also arbitrary. Thus, taking the infimum of the right-hand side, again, we obtain the desired inequality. ∎

Lemma 2.2.3 *Given a metric space, X, and $\epsilon > 0$, define $d_\epsilon : X_\epsilon \times X_\epsilon \to [0, \infty)$ by*

$$d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) = L([\alpha^{-1}\beta]_\epsilon).$$

Then d_ϵ defines a metric on X_ϵ.

Proof Clearly, $d_\epsilon \geq 0$ and $d_\epsilon([\alpha]_\epsilon, [\alpha]_\epsilon) = L([\alpha^{-1}\alpha]_\epsilon) = 0$. Symmetry follows because

$$d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) = L([\alpha^{-1}\beta]_\epsilon) = L([(\alpha^{-1}\beta)^{-1}]_\epsilon) = L([\beta^{-1}\alpha]_\epsilon) = d_\epsilon([\beta]_\epsilon, [\alpha]_\epsilon).$$

Moreover, positive definiteness follows from part 2 of the previous lemma.

For the triangle inequality, let $[\alpha]_\epsilon$, $[\beta]_\epsilon$, $[\gamma]_\epsilon \in X_\epsilon$ be given. Then part 3 of the previous lemma implies

$$d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) = L([\alpha^{-1}\beta]_\epsilon) = L([\alpha^{-1}\gamma\gamma^{-1}\beta]_\epsilon) \leq L([\alpha^{-1}\gamma]_\epsilon) + L([\gamma^{-1}\beta]_\epsilon)$$

$$\Rightarrow d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) \leq d_\epsilon([\alpha]_\epsilon, [\gamma]_\epsilon) + d_\epsilon([\gamma]_\epsilon, [\beta]_\epsilon). \quad ∎$$

The following result will give us a way to characterize the metric balls in X_ϵ in terms of the uniform structure defined on X_ϵ in section 1.3. Recall that if D is an entourage satisfying $D \subset E_\epsilon$, then D^* is the set of all $([\alpha]_\epsilon, [\beta]_\epsilon) \in X_\epsilon \times X_\epsilon$ such that $([\alpha]_\epsilon, [\beta]_\epsilon) = ([* = x_0, \ldots, x_n, y]_\epsilon, [* = x_0, \ldots, x_n, z]_\epsilon)$ and $(y, z) \in D$.

Lemma 2.2.4 *Let X be a metric space, and let $\epsilon > 0$ be given. For any $\delta \leq \epsilon$, the E_δ^*-balls in the entourage covering space, X_ϵ, are precisely the metric balls of radius δ under the metric d_ϵ.*

Proof Suppose $[\beta]_\epsilon \in B([\alpha]_\epsilon, E_\delta^*)$. Then, by the definition of E_δ^*, there is a representative $\alpha' = \{* = x_0, x_1, \ldots, x_n\} \in [\alpha]_\epsilon$ such that β is ϵ-homotopic to $\{*, x_1, \ldots, x_{n-1}, x\}$, where x is the endpoint of β and $d(x, x_n) < \delta$. It follows that the chain $\alpha^{-1}\beta$ is ϵ-homotopic to the chain $\{x_n, x_{n-1}, x\}$. Moreover, since $d(x_n, x) < \delta \leq \epsilon$, we can remove x_{n-1} to get the chain $\{x_n, x\}$. Then this two-point chain is in the equivalence class $[\alpha^{-1}\beta]_\epsilon$ and has length less than δ. Thus, $L([\alpha^{-1}\beta]_\epsilon) < \delta$, implying that $d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) < \delta$. This shows that $B([\alpha]_\epsilon, E_\delta^*) \subset B([\alpha]_\epsilon, \delta)$.

Conversely, suppose $[\beta]_\epsilon \in B([\alpha]_\epsilon, \delta)$, and let x be the endpoint of β. Then $L([\alpha^{-1}\beta]_\epsilon) < \delta \leq \epsilon$. By lemma 2.2.2, it follows that $d(x, x_n) < \delta$ and $\alpha^{-1}\beta$ is ϵ-homotopic to the chain $\sigma := \{x_n, x\}$. Now, the ϵ-loop, $\alpha^{-1}\beta\sigma^{-1}$, is ϵ-nullhomotopic to $\{x_n\}$. It follows that the ϵ-chains α^{-1} and $\sigma\beta^{-1}$ are ϵ-homotopic. Equivalently, the ϵ-chains α and $\beta\sigma^{-1}$ are ϵ-homotpic. Clearly, we have $\beta \sim_\epsilon \beta\sigma^{-1}\sigma$, so $\beta \sim_\epsilon \beta\sigma^{-1}\sigma \sim_\epsilon \alpha\sigma$. But $[\alpha\sigma]_\epsilon \in B([\alpha]_\epsilon, E_\delta^*)$ by definition of E_δ^*, so $[\beta]_\epsilon$ is, also. This shows that $B([\alpha]_\epsilon, \delta) \subset B([\alpha]_\epsilon, E_\delta^*)$. ∎

The next theorem is the specialization of Proposition 16 and Theorem 39 in [1] - or Theorem 1.3.4 in section 1.3 - to the context of metric spaces. This is partly just a restatement of Theorem 1.3.4, but we will need to explicitly point out more of the specific properties of the covering space and map for later use.

Theorem 2.2.5 *Let X be a chain-connected metric space, and let $*$ be a fixed base point in X. Then, for any $\epsilon > 0$, if X_ϵ is the set of equivalence classes of ϵ-chains beginning at $*$ with metric d_ϵ, then the following hold.*

1) *For any $\delta \leq \epsilon$, $\varphi_\epsilon : X_\epsilon \to X$ is a homeomorphism from δ-balls, $B([\alpha]_\epsilon, \delta) \subset X_\epsilon$, onto δ-balls $B(\varphi_\epsilon([\alpha]_\epsilon), \delta) \subset X$.*

2) *For any $\delta \leq \epsilon$, φ_ϵ is a radial isometry on any δ-ball, $B([\alpha]_\epsilon, \delta) \subset X_\epsilon$, meaning that $d(\varphi_\epsilon([\alpha]_\epsilon), \varphi_\epsilon([\beta]_\epsilon)) = d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon)$ for any $[\beta]_\epsilon \in B([\alpha]_\epsilon, \delta)$.*

3) *φ_ϵ is an isometry on any $(\epsilon/2)$-ball (i.e. φ_ϵ is a uniform local isometry).*

4) *For any $x \in X$ and any distinct $[\alpha]_\epsilon, [\beta]_\epsilon \in \varphi_\epsilon^{-1}(x)$, $B([\alpha]_\epsilon, \epsilon/2) \cap B([\beta]_\epsilon, \epsilon/2) = \emptyset$.*

5) *$\varphi_\epsilon : X_\epsilon \to X$ is a 1-Lipschitz, regular metric covering map.*

Proof We can first reduce the proof as follows. Note that the covering map conclusion of 5 will follow from 3 and 4. Moreover, 4 follows from 1 with $\delta = \epsilon$, and 3 follows from 2 and the triangle inequality.

First, given $[\alpha]_\epsilon, [\beta]_\epsilon \in X_\epsilon$, with x the endpoint of α and y the endpoint of β, the triangle inequality implies that $d(x, y) \leq L(\lambda)$ for any $\lambda \in [\alpha^{-1}\beta]_\epsilon$. Thus, $d(x, y) \leq L([\alpha^{-1}\beta]_\epsilon) = d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon)$, and φ_ϵ is 1-Lipschitz. This shows that $\varphi_\epsilon(B([\alpha]_\epsilon, \delta)) \subset B(\varphi_\epsilon([\alpha]_\epsilon), \delta)$ for any $[\alpha]_\epsilon \in X_\epsilon$ and $0 < \delta \leq \epsilon$. On the other hand, let $\delta \leq \epsilon$ and $[\alpha]_\epsilon \in X_\epsilon$ be given. Let $\alpha = \{* = x_0, \ldots, x_n\}$, so that $\varphi_\epsilon([\alpha]_\epsilon) = x_n$. Given $y \in B(x_n, \delta)$, the chain $\beta := \{* = x_0, \ldots, x_n, y\}$ is an ϵ-chain, and $\alpha^{-1}\beta \sim_\epsilon \{x_n, y\}$, showing that $L([\alpha^{-1}\beta]_\epsilon) = d(x_n, y) < \delta \Rightarrow [\beta]_\epsilon \in B([\alpha]_\epsilon, \delta)$. Clearly, $\varphi_\epsilon([\beta]_\epsilon) = y$, proving the surjectivity condition in part 1. Now, suppose that $[\beta]_\epsilon$, $[\gamma]_\epsilon \in B([\alpha]_\epsilon, \delta)$ and $\varphi_\epsilon([\beta]_\epsilon) = \varphi_\epsilon([\gamma]_\epsilon) = y$. Then β and γ are ϵ-chains ending at y. But the conditions $[\beta]_\epsilon$, $[\gamma]_\epsilon \in B([\alpha]_\epsilon, \delta)$ imply that β and γ are each ϵ-homotopic to the chain

$\{* = x_0, \ldots, x_n, y\}$. In other words, β and γ are ϵ-homotopic, showing that $[\beta]_\epsilon = [\gamma]_\epsilon$. This proves the injectivity condition in part 1. To see that $\varphi_\epsilon^{-1} : B(x_n, \delta) \to B([\alpha]_\epsilon, \delta)$ is continuous, let $z \in B(x_n, \delta)$ be given, and let $\{z_i\} \subset B(x_n, \delta)$ be any sequence converging to z. For each i, let $\alpha_i = \{* = x_0, \ldots, x_n, z_i\}$, and let $\alpha_\infty = \{* = x_0, \ldots, x_n, z\}$. Then $[\alpha_i]_\epsilon, [\alpha_\infty]_\epsilon \in B([\alpha]_\epsilon, \delta)$ for all i, and $\varphi_\epsilon^{-1}(z_i) = [\alpha_i]_\epsilon$, $\varphi_\epsilon^{-1}(z) = [\alpha_\infty]_\epsilon$. For i large enough that $d(z_i, z) < \epsilon$, we have

$$\alpha_\infty^{-1}\alpha_i \sim_\epsilon \{z, x_n, z_i\} \sim_\epsilon \{z, z_i\} \Rightarrow d_\epsilon([\alpha_\infty]_\epsilon, [\alpha_i]_\epsilon) = L([\alpha_\infty^{-1}\alpha_i]_\epsilon) \leq d(z, z_i) \to 0.$$

For part 2, if $[\beta]_\epsilon \in B([\alpha]_\epsilon, \delta)$, and if $\beta = \{* = y_0, \ldots, y_m\}$, then $d_\epsilon([\beta]_\epsilon, [\alpha]_\epsilon) = L([\beta^{-1}\alpha]_\epsilon) < \delta \leq \epsilon$, implying that $d_\epsilon([\beta]_\epsilon, [\alpha]_\epsilon) = d(y_m, x_n) = d(\varphi_\epsilon([\beta]_\epsilon), \varphi_\epsilon([\alpha]_\epsilon))$.

The regularity of $\varphi_\epsilon : X_\epsilon \to X$ follows from the general result regarding general entourage covers, but it will also follow, in the metric case, from results we will prove later on. ∎

Given this result, we can now justify calling X_ϵ ***the ϵ-cover of*** X. Carrying this out for each $\epsilon > 0$ gives us a collection of covering spaces of X, $\{X_\epsilon\}_{\epsilon>0}$, parameterized by the positive reals. Note, also, that if X is a bounded metric space, then X_ϵ is isometric to X for all sufficiently large ϵ. This simply follows by choosing ϵ large enough (say, larger than $diam(X)$) and applying part 3 of the previous theorem.

Remark At this point, we have introduced two collections of covering spaces indexed by the positive reals: the Spanier or δ-covers of geodesic spaces and our ϵ-covers defined for general metric spaces. To distinguish between the two, we will adopt the following notational convention: we will use superscripts to denote the Spanier covers, X^δ, when X is a geodesic space and subscripts to denote the ϵ-covers, X_ϵ, in any case.

Next, we give our first example of ϵ-covers. We will work through this first example in detail, and the others that are presented hereafter will follow from similar reasoning.

Example 2.2.6 *Let S^1 be the geodesic circle of circumference 1, and fix a base point, $* \in S^1$. Let $\epsilon > 1/3$ be given. We will show that $S^1_\epsilon = S^1$, or that S^1_ϵ is the trivial cover. Let $\gamma = \{* = x_0, \ldots, x_n = *\}$ be any ϵ-loop at $*$, and let a and b denote the points in S^1 that, along with $*$, subdivide the circle into three segments of length $1/3$. So, we have $d(*, a) = d(*, b) = d(a, b) = 1/3$. Note that for any $i = 1, \ldots, n$, at least one of the points $*$, a, or b will be within ϵ of x_{i-1} and x_i. In fact, given any $x \in S^1$, there will be at least two of the points $*$, a, b that are a distance less than or equal to $1/3$ from x (hence, strictly within ϵ of x). So, taking those two points for each of x_{i-1} and x_i, the two pairs must have at least one point in common. Thus, for each $i = 1, \ldots, n$, we can insert a point, y_i, into γ between x_{i-1} and x_i, where y_i is one of $*$, a, or b, giving us a loop*

$$\gamma' = \{* = x_0, y_1, x_1, \ldots, x_{n-1}, y_n, x_n = *\}.$$

But since $d(y_j, y_k) = 1/3 < \epsilon$ for all j, k, we can then successively remove x_1, x_2, and so on up to x_{n-1}, giving us the loop $\gamma'' = \{* = x_0, y_1, \ldots, y_n, x_n = *\}$. Finally, by the same reasoning, we can successively remove y_1, y_2, and so on up to y_n, giving us the constant loop $\{*, *\}$. So, every ϵ-loop is trivial for $\epsilon > 1/3$. Equivalently, there is only one class of ϵ-chains from $*$ to any point $x \in S^1$, showing that $\varphi_\epsilon : S^1_\epsilon \to S^1$ is injective and, therefore, a homeomorphism. It is, in fact, an isometry, but this will follow from results we will prove later.

Now, suppose $\epsilon \leq 1/3$. We will show that S^1_ϵ is the universal cover, $\mathbb{R}$. We first choose an orientation. In fact, if we let C be the simple loop at $*$ traversing the circle one time in a given direction, then C and C^{-1} induce the orientation. Adopting an intuitive (though slight abuse *of) notation, given two points $x, y \in S^1$ with $d(x,y) < \epsilon$, we will let $[x,y]$ denote the geodesic segment of S^1 from x to y, including the two endpoints, and we likewise let $[x,y)$ and $(x,y]$ denote the half-open, half-closed geodesic segments from x to y including just the initial and final endpoints, respectively. Since $\epsilon \leq 1/3 < 1/2$, there is no ambiguity in the choice of this geodesic segment. Note, also, that this implicitly carries with it an assignment of orientation to the segment, depending on the loop - C or C^{-1} - with which the direction of the segment is consistent. We let $\omega([x,y])$ denote the orientation of a segment. Given $x, y \in S^1$, with $d(x,y) < \epsilon$, define*

$$|[x,y]| := \begin{cases} 0, & * \notin [x,y] \text{ or if } x = y \\ 1, & * \in [x,y) \text{ and } \omega([x,y]) = + \\ 0, & * = y \text{ and } \omega([x,y]) = + \\ -1, & * \in (x,y] \text{ and } \omega([x,y]) = - \\ 0, & * = x \text{ and } \omega([x,y]) = -. \end{cases}$$

*Then, given an ϵ-loop, $\gamma = \{* = x_0, \ldots, x_n = *\}$, we define the **winding number of** γ to be*

$$|\gamma| := \sum_{i=1}^{n} |[x_{i-1}, x_i]|.$$

It is a straightforward exercise in cases to verify that $|\gamma|$ is an ϵ-homotopy invariant, so the winding number of the equivalence class of an ϵ-loop, $|[\gamma]_\epsilon|$, is well-defined. Moreover, the winding number of a concatenation of two ϵ-loops is given by $|\gamma_1\gamma_2| = |\gamma_1| + |\gamma_2|$. One can also show that a non-trivial ϵ-loop has non-zero winding number. In fact, by utilizing appropriate steps, one can transform a given ϵ-loop via ϵ-homotopy to a loop that does not backtrack, or, equivalently, a loop that is monotone with respect to one of the two directions determined by the orientation. Such an ϵ-loop will determine segments, $[x_{i-1}, x_i]$, of a uniform orientation, and a non-trivial loop will have at least one segment such that $[x_{i-1}, x_i] \neq 0$. It then follows that an equivalence class of ϵ-loops has a winding number of 0 if and only if it is the trivial class. Now, define a map $\rho : \pi_\epsilon(S^1) \to \mathbb{Z}$ by $\rho([\gamma]_\epsilon) = |[\gamma]_\epsilon|$. The properties just stated show that ρ is an injective homomorphism. Surjectivity also holds, for if $m \in \mathbb{Z}$, we can take C^m or C^{-m} depending on the sign of m and subdivide this loop into segments of length less than ϵ, thus giving us an ϵ-loop that winds around m times in the desired direction. Thus, $\pi_\epsilon(S^1) \cong \mathbb{Z}$, and the preimage of any point in S^1 under the covering map φ_ϵ is isomorphic to $\mathbb{Z}$. Additionally, any $[\alpha]_\epsilon \in S^1_\epsilon$ - with x the endpoint of α - can be expressed uniquely as $[\alpha]_\epsilon = [\gamma]_\epsilon^{\pm m}[\beta]_\epsilon$, where $[\gamma]_\epsilon$ is the generator of $\pi_\epsilon(S^1)$, $|m|$ is minimal, and $[\beta]_\epsilon$ is one of the precisely two distinct classes ϵ-chains from $*$ to x having length less than 1. If we now define $f : S^1_\epsilon \to \mathbb{R}$ by $f([\gamma]_\epsilon^m[\beta]_\epsilon) = m + sign(m)L([\beta]_\epsilon)$, then f is an isometry.

Notice that the ϵ-covers changed topological type at $1/3$, or $1/3$ the circumference of the circle. This turns out to be a very general phenomenon in the case of geodesic spaces. ∎

Example 2.2.7 *As in the previous example, let S^1 be the geodesic circle of circumference 1. Fix two points, $a, b \in S^1$, such that $d(a,b) = 1/4$, and remove the open geodesic segment from a to b. Let X be the resulting set with this open segment removed and endowed with the subspace metric inherited from S^1, **not** the induced geodesic metric (which would just make X a line segment). Choose as the base point, $*$, the point that subdivides the longer segment from a to b into two subsegments of equal length. Reasoning as in the previous example, it follows that, for $\epsilon > 1/3$, X_ϵ is just the trivial cover. Essentially, since ϵ is larger than the length of the gap we created* *by removing the segment, an ϵ-loop can cross over the gap, intuitively meaning that the ϵ-cover does not recognize that a "piece of the space is missing." Said another way, the gap represents a geometric structure on a scale too small for the ϵ-cover to notice. For $1/4 < \epsilon \leq 1/3$, one can show - again, using the same arguments as in the previous example - that $\pi_\epsilon(X) \cong \mathbb{Z}$. In this case, X_ϵ is $\mathbb{R}$ with the open segments $(n/2 - 1/8, n/2 + 1/8)$ removed for every integer, n. Intuitively, we still unravel the circle as in the standard geodesic case, but the missing segment of the circle gets unraveled along with it. Again, this holds because an ϵ-loop can still cross over the gap, so that particular structure is not detected by the ϵ-cover. Note that X_ϵ is not connected in this case.*

Now, suppose $0 < \epsilon \leq 1/4$. Since the removed segment has length $1/4$, it becomes the case at $\epsilon = 1/4$ that no ϵ-loop can cross the gap. Suddenly, the ϵ-covers recognize the gap, because the scale at which they are detecting structure within X is now small enough. Intuitively, it is now impossible to travel around the circle, in either direction, via an ϵ-loop. Thus, all ϵ-loops are necessarily trivial, and $\pi_\epsilon(X)$ is the trivial group. However, we do not get the trivial cover again. In this case, X_ϵ is just a geodesic segment of length $3/4$, meaning that X_ϵ simply unravels or straightens out the partial circle making up X. See Figure 2.1 below. ∎

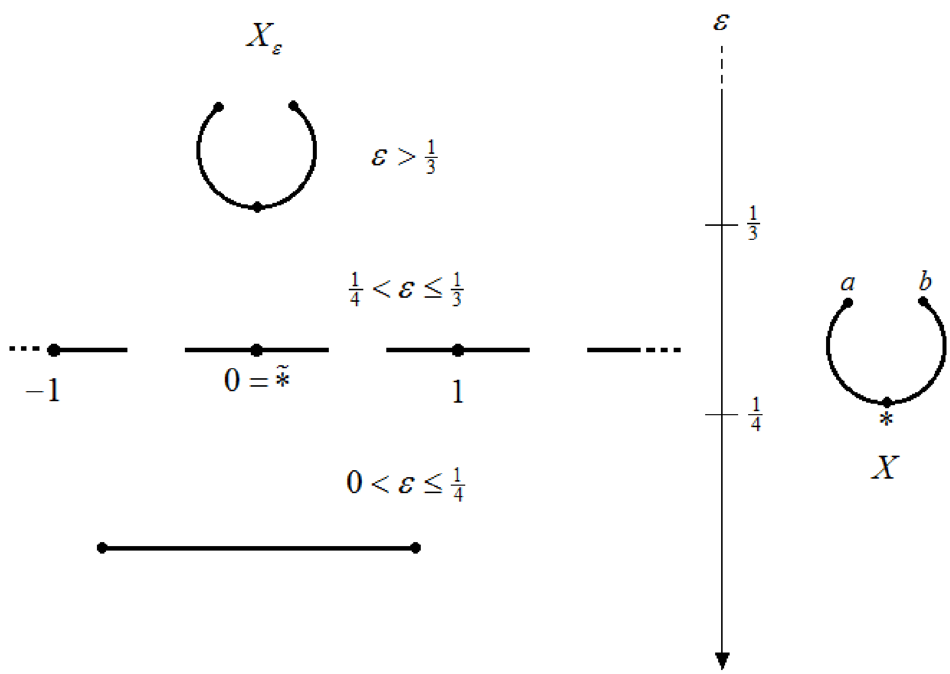

Figure 2.1: Geodesic Circle with a Gap (the ϵ-covers)

Example 2.2.8 *Let $T^2 = S^1_1 \times S^2_2$ be the geodesic torus formed by taking the metric product of geodesic circles of circumference 1 and 2, respectively. Applying the reasoning used in Example 2.2.6, we obtain the following.*

$$T^2_\epsilon = \begin{cases} T^2, & \epsilon > \frac{2}{3} \\ S^1_1 \times \mathbb{R}, & \frac{1}{3} < \epsilon \leq \frac{2}{3} \\ \mathbb{R}^2, & 0 < \epsilon \leq \frac{1}{3}. \end{cases}$$

In other words, when $\epsilon = 2/3$, the larger circle is unraveled while the smaller one is unaffected, giving us a cylinder. When $\epsilon = 1/3$, both circles are unraveled, giving us the universal cover.

If we let $T^2 = S^1_1 \times S^1_1$, then

$$T^2_\epsilon = \begin{cases} T^2, & \epsilon > \frac{1}{3} \\ \mathbb{R}^2, & 0 < \epsilon \leq \frac{1}{3}. \end{cases}$$

In this case, both circles are unraveled simultaneously. This example indicates that there is some notion of multiplicity lurking behind the values where the ϵ-covers change structure. ∎

Next, we will show that the choice of base point in the construction of the ϵ-cover is immaterial. This is one of the primary reasons for assuming chain-connectivity. In this lemma, the notation $(X_\epsilon, *)$ refers to the ϵ-cover determined by choosing $*$ as the base point.

Lemma 2.2.9 *Let X be a chain-connected metric space, $*_1$ and $*_2$ two base points in X, and $\epsilon > 0$. Let λ be an ϵ-chain from $*_1$ to $*_2$, and define maps $f : (X_\epsilon, *_1) \to (X_\epsilon, *_2)$, $\Phi : (\pi_\epsilon(X), *_1) \to (\pi_\epsilon(X), *_2)$ by*

$$f([\alpha]_\epsilon) = [\lambda^{-1}\alpha]_\epsilon \quad \textit{and} \quad \Phi([\gamma]_\epsilon) = [\lambda^{-1}\gamma\lambda].$$

Then f is an isometry, and Φ is an isomorphism.

Proof Let $[\alpha]_\epsilon, [\beta]_\epsilon \in (X_\epsilon, *_1)$ be given, and let d^1_ϵ, d^2_ϵ denote the metrics on $(X_\epsilon, *_1)$ and $(X_\epsilon, *_2)$, respectively. Then

$$d^2_\epsilon([\lambda^{-1}\alpha]_\epsilon, [\lambda^{-1}\beta]_\epsilon) = L([\alpha^{-1}\lambda\lambda^{-1}\beta]_\epsilon) = L([\alpha^{-1}\beta]_\epsilon)$$

$$\Rightarrow d^2_\epsilon(f([\alpha]_\epsilon), f([\beta]_\epsilon)) = d^1_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon).$$

This proves that f is an isometry onto its image. To see that f is surjective, note that if $[\sigma]_\epsilon \in (X_\epsilon, *_2)$, then $f([\lambda\sigma]_\epsilon) = [\sigma]_\epsilon$.

Now, assume further that α and β are ϵ-loops at $*_1$. Then

$$[\lambda\alpha\beta\lambda^{-1}]_\epsilon = [\lambda\alpha\lambda^{-1}\lambda\beta\lambda^{-1}]_\epsilon = [\lambda\alpha\lambda^{-1}]_\epsilon[\lambda\beta\lambda^{-1}]_\epsilon$$

$$\Rightarrow \Phi([\alpha]_\epsilon[\beta]_\epsilon) = \Phi([\alpha\beta]_\epsilon) = \Phi([\alpha]_\epsilon)\Phi([\beta]_\epsilon).$$

This shows that Φ is a homomorphism onto its image. If $[\sigma]_\epsilon \in (\pi_\epsilon(X), *_2)$, then $\Phi([\lambda\sigma\lambda^{-1}]_\epsilon) = [\sigma]_\epsilon$, showing that Φ is surjective. Finally, if $\Phi([\gamma]_\epsilon) = [\{*_2\}]_\epsilon$, then $\lambda^{-1}\gamma\lambda \sim_\epsilon \{*_2\}$, implying that $\gamma \sim_\epsilon \lambda\lambda^{-1}$. Thus, γ is ϵ-null, and Φ is injective. ∎

Remark As a consequence of this result, we will, henceforth, have few occasions to refer to the specific base point used to construct X_ϵ. Unless otherwise specified, we will always denote this base point by $*$, and when we speak of X_ϵ without referencing a base point, it will be assumed that we mean the ϵ-cover induced by the base point $*$. Moreover, we will also adopt a uniform notation for the point $[\{*\}]_\epsilon \in X_\epsilon$, or the point containing the trivial, constant chain at $*$. Unless specified otherwise, we will always denote the point, $[\{*\}]_\epsilon \in X_\epsilon$, simply by $\tilde{*}$, using the equivalence class notation only when it is necessary to emphasize the fact that this point is a class of ϵ-chains. If there are situations where reference to the base points is essential, we will use the notation $(X_\epsilon, *)$ or $(X_\epsilon, \tilde{*})$ to denote the ϵ-cover determined by the base point, $*$.

We conclude this section with some results regarding an analog of metric diameter. Let X be a compact, chain-connected metric space, and let $\epsilon > 0$ be given. Given $x, y \in X$, let $L^\epsilon_{x,y} = \inf\{L(\alpha)\}$, where the infimum is taken over all ϵ-chains from x to y. Clearly this infimum is well-defined. In fact, by the triangle inequality, we have $d(x,y) \leq L^\epsilon_{x,y}$. Now, given $x \in X$, let $r^\epsilon_x = \sup_{y\in X} L^\epsilon_{x,y}$. We call r^ϵ_x the *ϵ-radius of X at x*.

Lemma 2.2.10 *If X is compact and chain-connected and $x \in X$, then r^ϵ_x is finite for any $\epsilon > 0$.*

Proof If this were not true, then we could find points, $\{y_n\}_{n\geq 1} \subset X$, so that $L^\epsilon_{x,y_n} \nearrow \infty$. Since X is compact, there is a subsequence, $\{y_{n_k}\}$, that converges to some y. By reindexing, if necessary, we can just assume without loss of generality that $y_n \to y$. Now, let α be an ϵ-chain from x to y, and let $N \in \mathbb{N}$ be such that $n \geq N \Rightarrow d(y_n, y) < \epsilon$. Then for every $n \geq N$, we can form an ϵ-chain, α_n, from x to y_n by simply adding y_n onto the end of α. Moreover, we have

$$L(\alpha_n) = L(\alpha) + d(y, y_n) < L(\alpha) + \epsilon.$$

But this would imply that $L^\epsilon_{x,y_n} < L(\alpha) + \epsilon$ for all such n, a contradiction. ∎

Now, we define the *ϵ-diameter of X* to be

$$diam_\epsilon(X) = \sup_{x,y\in X} L^\epsilon_{x,y}.$$

A compactness argument just like that in the proof of the previous lemma - along with the triangle inequality - yields the following result.

Lemma 2.2.11 *If X is compact and chain-connected, then, for any $\epsilon > 0$, $diam_\epsilon(X) < \infty$ and $diam(X) \leq diam_\epsilon(X)$.*

Example 2.2.12 *Let X be the unit circle of radius 1 in $\mathbb{R}^2$, with its inherited Euclidean metric. The diameter of this space is 2. For $\epsilon > 2$, we have $diam_\epsilon(X) = diam(X)$, and any pair of antipodal points attains both diameters. For $\epsilon = 2$, we still have $diam_\epsilon(X) = diam(X) = 2$, but now there is no single ϵ-chain attaining the ϵ-diameter. If x and y are antipodal points, we can take the 2-chain, $\{x, z, y\}$, where z is a point arbitrarily close to - but not equal to - y. Taking the infimum of the lengths of these chains over all such z shows that $L^2_{x,y} = 2$, but any 2-chain from x to y must contain at least one other point. This should be contrasted with the regular diameter of a compact metric space, which is always attained by the distance between at least one pair of points.*

For $\epsilon < 2$, $diam_\epsilon(X)$ will be strictly greater than 2. For instance, for $\epsilon = \sqrt{2}$, we have $diam_\epsilon(X) = 2\sqrt{2}$. This can be seen by taking two antipodal points, x and y, and letting z be one *of the points that subdivides one of the arcs from x to y into two equal-length arcs. By taking points, u and v, that are arbitrarily close to z but closer to x than to y and closer to y than to x, respectively, we can form the $\sqrt{2}$-chain $\{x, u, v, y\}$. Taking the infimum of the lengths of these chains over all such u and v shows that $L^{\sqrt{2}}_{x,y} = 2\sqrt{2}$.* ∎

The final lemma deals with the number of points in ϵ-chains instead of lengths, and it proves to be useful on occasion.

Lemma 2.2.13 *If X is compact and chain-connected and $\epsilon > 0$, then there is a natural number, M, with the property that there is an ϵ-chain between any two points in X having cardinality bounded above by M.*

Proof First, we note that if x and y are any points in X, then there is an ϵ-chain from x to y of minimal cardinality. In fact, this is just a consequence of the well-ordering of the natural numbers. Such a chain may not be unique, but the minimal cardinality is unique, and there will be at least one chain with this minimal cardinality. We will actually show that this minimal cardinality is uniformly bounded above over all pairs of points in X.

Given $x, y \in X$, let $C^{\epsilon}_{x,y}$ be the minimal cardinality of an ϵ-chain from x to y. If the supremum of these values over all $x, y \in X$ is not bounded above, then we can find pairs of points in X, $\{(x_n, y_n)\}$, such that $C^{\epsilon}_{x_n,y_n} \nearrow \infty$. As before, since X is compact, we can assume that $x_n \to x$ and $y_n \to y$. Let α be an ϵ-chain of minimal cardinality from x to y, and let $N \in \mathbb{N}$ be such that $n \geq N \Rightarrow d(x_n, x) < \epsilon$ and $d(y_n, y) < \epsilon$. Then, for each $n \geq N$, we can form an ϵ-chain, α_n, from x_n to y_n by adding x_n at the beginning of α and y_n to the end of α. Thus, for each $n \geq N$, there is an ϵ-chain from x_n to y_n of cardinality $C^{\epsilon}_{x,y} + 2$, implying that $C^{\epsilon}_{x_n,y_n} \leq C^{\epsilon}_{x,y} + 2$. This is a contradiction. ∎

2.3 Critical Values of Metric Spaces

In this section, we will investigate the connection between ϵ-covers for different values of ϵ. This will finally lead us to the definition of the critical spectrum of a metric space.

Suppose $\delta < \epsilon$. Then every δ-chain is an ϵ-chain, and every δ-homotopy is an ϵ-homotopy. Consequently, given an equivalence class of δ-chains, $[\alpha]_\delta$, we can naturally consider the equivalence class, $[\alpha]_\epsilon$. This class contains every δ-chain, $\alpha' \in [\alpha]_\delta$, since any such α' will be δ-homotpic - and, thus, ϵ-homotopic - to α. This induces a natural map from X_δ to X_ϵ, and we will denote this map by $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$. Note the order of the subscripts, with the target listed first. We also have the following composition formula: if $\delta < \epsilon < \lambda$, then $\varphi_{\lambda\delta} = \varphi_{\lambda\epsilon} \circ \varphi_{\epsilon\delta}$.

Lemma 2.3.1 *If $\delta < \epsilon$, then $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is 1-Lipschitz.*

Proof It suffices to simply show that for any δ-equivalence class, $[\alpha]_\delta$, we have $L([\alpha]_\epsilon) \leq L([\alpha]_\delta)$. This is obvious, though, since the set of lengths of all ϵ-chains that are ϵ-homotopic to α includes the set of lengths of all δ-chains that are δ-homotopic to α. ∎

There are also corresponding maps between the groups, $\pi_\delta(X)$ and $\pi_\epsilon(X)$. In fact, this map is just the restriction of $\varphi_{\epsilon\delta}$ to $\pi_\delta(X)$, and we denote these maps by $\Phi_{\epsilon\delta} : \pi_\delta(X) \to \pi_\epsilon(X)$. These maps are also homomorphisms, since

$$\Phi_{\epsilon\delta}([\gamma_1]_\delta[\gamma_2]_\delta) = \Phi_{\epsilon\delta}([\gamma_1\gamma_2]_\epsilon) = [\gamma_1\gamma_2]_\epsilon = [\gamma_1]_\epsilon[\gamma_2]_\epsilon = \Phi_{\epsilon\delta}([\gamma_1]_\delta)\Phi_{\epsilon\delta}([\gamma_2]_\delta).$$

The surjectivity/injectivity of these homomorphisms is closely related to the corresponding properties of φ_ϵ.

Now, when X is a geodesic space and we can also discuss the Spanier covers, X^δ, we know that for any $\delta_1 < \delta_2$ there is always a surjective covering map from X^{δ_1} onto X^{δ_2}. Indeed, it is the failure of these maps to be *injective* that defines an element of the covering spectrum. In contrast, for a general chain-connected metric space, the maps $\varphi_{\epsilon\delta}$ need not be injective *or* surjective (though the preceding lemma implies that they are, at least, continuous). The failure of the maps, φ_ϵ, to be injective and/or surjective is what we will use to define critical values of a metric space. First, however, we will illustrate some geometric and algebraic characterizations of the injectivity and surjectivity of these maps. In particular, we will see that surjectivity of the maps, $\varphi_{\epsilon\delta}$, is closely tied to the ability (or inability) to refine chains.

Lemma 2.3.2 *Let X be a chain-connected metric space, and let $0 < \delta < \epsilon$ be given. The following are equivalent.*

1) $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is surjective.

2) $\Phi_{\epsilon\delta} : \pi_\delta(X) \to \pi_\epsilon(X)$ is surjective.

3) Every ϵ-chain in X can be ϵ-refined to a δ-chain.

4) Every two-point ϵ-chain in X can be ϵ-refined to a δ-chain.

In other words, $\varphi_{\epsilon\delta}$ is surjective if and only if we can refine ϵ-chains to δ-chains. Part 4 of this equivalence is particularly useful, since it shows that we only need to be able to refine two-point ϵ-chains to δ-chains.

Proof ($3 \Leftrightarrow 4$) That 3 implies 4 is trivial. The fact that 4 implies 3 is equally clear, since every ϵ-chain is just a concatenation of two-point ϵ-chains.

($1 \Rightarrow 3$) Let α be any ϵ-chain based at $*$. Then there is some δ-chain, β, such that $\varphi_{\epsilon\delta}([\beta]_\delta) = [\beta]_\epsilon = [\alpha]_\epsilon$. But this means that $\alpha \sim_\epsilon \beta$, so α is ϵ-homotopic to a δ-chain. This shows that every ϵ-chain anchored at the base point, $*$, can be refined to a δ-chain. Now, let α be an ϵ-chain anchored at another point, $\star$. Since X is chain-connected, we have isometries $f_\epsilon : (X_\epsilon, *) \to (X_\epsilon, \star)$ and $g_\delta : (X_\delta, *) \to (X_\delta, \star)$. In fact, if λ is a δ-chain from $*$ to $\star$, then λ is also an ϵ-chain. So, we can take $f_\epsilon([\alpha]_\epsilon) = [\lambda^{-1}\alpha]_\epsilon$ and $g_\delta([\beta]_\delta) = [\lambda^{-1}\beta]_\delta$. With these definitions, it is easy to see that $f_\epsilon \circ \varphi^*_{\epsilon\delta} = \varphi^\star_{\epsilon\delta} \circ g$, where $\varphi^*_{\epsilon\delta} : (X_\delta, *) \to (X_\epsilon, *)$ and $\varphi^\star_{\epsilon\delta} : (X_\delta, \star) \to (X_\epsilon, \star)$ are the corresponding maps. By our assumptions, we know that every map in the equation $f_\epsilon \circ \varphi^*_{\epsilon\delta} = \varphi^\star_{\epsilon\delta} \circ g_\delta$ is surjective except possibly $\varphi^\star_{\epsilon\delta}$. But since g_δ is an isometry, we have $\varphi^\star_{\epsilon\delta} = f_\epsilon \circ \varphi^*_{\epsilon\delta} \circ g_\delta^{-1}$, showing that $\varphi^\star_{\epsilon\delta}$ is surjective. Thus, the same reasoning we used before shows that every ϵ-chain at $\star$ can be ϵ-refined to a δ-chain. Since $\star$ was arbitrary, this shows that 1 implies 3.

($3 \Rightarrow 1$) Let $[\alpha]_\epsilon \in X_\epsilon$ be given. Then, by assumption, α is ϵ-homotopic to a δ-chain, β. Thus, $\varphi_{\epsilon\delta}([\beta]_\delta) = [\beta]_\epsilon = [\alpha]_\epsilon$.

($1 \Rightarrow 2$) Let $[\gamma]_\epsilon \in \pi_\epsilon(X)$ be given. Since $\varphi_{\epsilon\delta}$ is surjective, there is some $[\alpha]_\delta \in X_\delta$ such that $\varphi_{\epsilon\delta}([\alpha]_\delta) = [\gamma]_\epsilon$. This means that α is ϵ-homotopic to γ, which further means that the endpoint of α must be the same as that of γ. Thus, $[\alpha]_\delta \in \pi_\delta(X)$, showing that $\Phi_{\epsilon\delta}$ is surjective.

($2 \Rightarrow 1$) Let $\alpha = \{* = x_0, \ldots, x_n\}$ be an ϵ-chain. Since X is chain connected, there is some δ-chain, β, from $*$ to x_n. By assumption, there is some $[\gamma]_\delta \in \pi_\delta(X)$ such that $\Phi_{\epsilon\delta}([\gamma]_\delta) = [\alpha\beta^{-1}]_\epsilon$. So, we have $[\gamma]_\epsilon = [\alpha\beta^{-1}]_\epsilon \Rightarrow [\alpha]_\epsilon = [\gamma\beta]_\epsilon$. This implies that the δ-chain, $\gamma\beta$, is in $[\alpha]_\epsilon$. ∎

Lemma 2.3.3 *Let X be a chain-connected metric space, and let $0 < \delta < \epsilon$ be given. The following are equivalent.*

1) $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ *is injective.*

2) $\Phi_{\epsilon\delta} : \pi_\delta(X) \to \pi_\epsilon(X)$ *is injective.*

Proof ($1 \Rightarrow 2$) This follows because the restriction of an injective map is also injective.

($2 \Rightarrow 1$) Suppose $\varphi_{\epsilon\delta}([\alpha]_\delta) = \varphi_{\epsilon\delta}([\beta]_\epsilon)$. Then $[\alpha]_\epsilon = [\beta]_\epsilon$, meaning that the δ-chains, α and β, end at the same point. Let γ be the δ-loop, $\alpha\beta^{-1}$. Then $\Phi_{\epsilon\delta}([\gamma]_\delta) = [\gamma]_\epsilon = [\alpha\beta^{-1}]_\epsilon = \tilde{*}$, and the injectivity of $\Phi_{\epsilon\delta}$ implies, then, that $[\gamma]_\delta = \tilde{*} \Rightarrow [\alpha]_\delta = [\beta]_\delta$. This shows that $\varphi_{\epsilon\delta}$ is injective. ∎

This is the most common interpretation of non-injectivity we will use: $\varphi_{\epsilon\delta}$ is non-injective if and only if there is a non-trivial δ-loop that is ϵ-null.

Now, we can introduce the critical spectrum. Let $\mathbb{R}_+ = (0, \infty)$ and $\overline{\mathbb{R}}_+ = [0, \infty)$.

Definition 2.3.4 *Let X be a chain-connected metric space. A **non-critical interval of** X is a non-empty, open interval, $I \subset \mathbb{R}^+$, such that for each $\delta, \epsilon \in I$ with $\delta < \epsilon$, the map $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is bijective. A **maximal non-critical interval** is a non-critical interval that is not contained in any other non-critical interval. A positive real number, ϵ, is a **non-critical value of** X if and only if it lies in some non-critical interval. We call ϵ a **critical value of** X if it is not non-critical, or, equivalently, if it is not contained in any non-critical interval. We denote the subset of $\mathbb{R}_+$ consisting of all critical values of X by $Cr(X)$, and we call this the **Critical Spectrum** of X.*

An immediate consequence of the definition is the following.

Lemma 2.3.5 $\mathbb{R}_+ \setminus Cr(X)$ *is open in* $\mathbb{R}_+$.

This follows simply because every non-critical value, by definition, lies in a non-critical interval, which consists entirely of non-critical values. Note, however, that the complement of $Cr(X)$ in $\mathbb{R}$ may not be open in $\mathbb{R}$, or, equivalently, $Cr(X)$ may not be closed in $\mathbb{R}$. In particular, 0, by definition, cannot be a critical value, but it may be a limit point of $Cr(X)$. However, it does follow from the definition that $Cr(X)$ contains all of its positive limit points. In fact, if $\{\epsilon_n\} \subset Cr(X)$ and $\lim_n \epsilon_n = \epsilon > 0$ is a limit point of $Cr(X)$, then ϵ cannot lie in any non-critical interval. If it did, say $\epsilon \in (a, b)$, then ϵ_n would lie in (a, b) for large enough n, a contradiction. Thus, $Cr(X)$ is "almost" closed in $\mathbb{R}$. Indeed, if $\inf Cr(X) > 0$, then $Cr(X) = \overline{Cr(X)}$; if $\inf Cr(X) = 0$, then $\overline{Cr(X)} = Cr(X) \cup \{0\}$.

Lemma 2.3.6 *Let (a, b) be a maximal non-critical interval. Then for any $\epsilon \in (a, b)$, the map $\varphi_{b\epsilon} : X_\epsilon \to X_b$ is a bijection. Moreover, a and b are critical values of X.*

Proof Let $\epsilon \in (a, b)$ be given. Let $\{x, y\}$ be any two-point b-chain in X. Then $d(x, y) < b$, meaning that there is some δ such that $d(x, y) < \delta < b$. That is, $\{x, y\}$ is a δ-chain, also. We can assume that $\delta > \epsilon$. Since $\varphi_{\delta\epsilon} : X_\epsilon \to X_\delta$ is bijective, it follows that $\{x, y\}$ can be δ-refined to an ϵ-chain. This δ-homotopy will also be a b-homotopy, showing that $\{x, y\}$ can be b-refined to an ϵ-chain. In other words, $\varphi_{b\epsilon}$ is surjective.

Next, suppose there were some non-trivial ϵ-loop, γ, that was b-null. Then there is a b-homotopy taking γ to the chain $\{*\}$. But a b-homotopy is just a finite collection of b-chains, so it is also a δ-homotopy for some $\delta < b$. We may assume that $\epsilon < \delta$. That is, γ is δ-null. But this contradicts the fact that $\varphi_{\delta\epsilon} : X_\epsilon \to X_\delta$ is injective. Thus, $\varphi_{b\epsilon} : X_\epsilon \to X_b$ must be injective. This proves the first part of the lemma.

Now, if b were not a critical value, it would lie in some non-critical interval, say $(b-t_1, b+t_2)$. It would follow that $(a, b+t_2)$ is a non-critical interval containing (a,b), contradicting that (a,b) is maximal. Hence, $b \in Cr(X)$. Similar reasoning shows that $a \in Cr(X)$. ∎

An equivalent definition of critical value that is often more useful is given by the following lemma. First, however, we should make a few brief notes on terminology and notation. In the following, we may occasionally need to refer to the map between two covers, X_δ and X_ϵ, without knowing beforehand whether $\delta < \epsilon$ or vice versa. In this case, when we do not know the order of δ and ϵ and we speak of "the map between X_δ and X_ϵ," it should be understood that we mean either the map $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ if $\delta < \epsilon$ or the map $\varphi_{\delta\epsilon} : X_\epsilon \to X_\delta$ if $\epsilon < \delta$. Second, since we will be using these maps extensively, to avoid cumbersome and repetitive expressions, we will often suppress the map itself and refer to the map $\varphi_{\delta\epsilon} : X_\epsilon \to X_\delta$ simply with the symbol $X_\epsilon \to X_\delta$.

Lemma 2.3.7 *A positive number, ϵ, is in $Cr(X)$ if and only if there is a sequence, $\{\epsilon_n\}$, with $\epsilon_n \to \epsilon$, such that the map between X_ϵ and X_{ϵ_n} is not bijective for all n.*

Proof First, suppose such a sequence, $\{\epsilon_n\}$, exists. Then ϵ cannot lie in any non-critical interval. If it did, say $\epsilon \in (a,b)$, then ϵ_n would lie in this interval for all sufficiently large n. This would imply that for such n the map between X_ϵ and X_{ϵ_n} is a bijection, which is a contradiction. So, ϵ must be in $Cr(X)$.

Conversely, suppose $\epsilon \in Cr(X)$. Then there is no non-critical interval containing ϵ. So, there are two positive numbers, r_1 and r_2, in $(\epsilon-1, \epsilon+1)$ such that $r_1 < r_2$ and the map $X_{r_1} \to X_{r_2}$ is not bijective. We will show that at least one of the maps between X_ϵ and X_{r_1}, X_{r_2} must be non-bijective. If ϵ equals r_1 or r_2, we are done. If $r_1 < r_2 < \epsilon$, then $\varphi_{\epsilon r_1} = \varphi_{\epsilon r_2} \circ \varphi_{r_2 r_1}$. If both $\varphi_{\epsilon r_1}$ and $\varphi_{\epsilon r_2}$ were bijective, then $\varphi_{r_2 r_1}$ would be also, a contradiction. If $\epsilon < r_1 < r_2$, then $\varphi_{r_2\epsilon} = \varphi_{r_2 r_1} \circ \varphi_{r_1\epsilon}$. Again, if both $\varphi_{r_2\epsilon}$ and $\varphi_{r_1\epsilon}$ were bijections, $\varphi_{r_2 r_1}$ would be, also. Finally, if $r_1 < \epsilon < r_2$, then $\varphi_{r_2 r_1} = \varphi_{r_2\epsilon} \circ \varphi_{\epsilon r_1}$. If both $\varphi_{r_2\epsilon}$ and $\varphi_{\epsilon r_1}$ were bijections, $\varphi_{r_2 r_1}$ would be, also. Thus, in any of the possible cases, at least one of the maps between X_ϵ and X_{r_1}, X_{r_2} must not be bijective. So, we can choose $\epsilon_1 \in (\epsilon-1, \epsilon+1)$ - either r_1 or r_2 depending on the previous cases - such that the map between X_ϵ and X_{ϵ_1} is not bijective. Likewise, there are two positive numbers r_1 and r_2 in the interval $(\epsilon - 1/2, \epsilon + 1/2)$, with $r_1 < r_2$, such that map $X_{r_1} \to X_{r_2}$ is not bijective. Reasoning as before, we can conclude that at least one of the maps between X_ϵ and X_{r_1}, X_{r_2} is not a bijection. We choose ϵ_2 to be whichever of r_1, r_2 will make this so.

Continuing inductively, this yields a sequence, $\{\epsilon_n\}$, with $\epsilon_n \to \epsilon$, such that the map between X_ϵ and X_{ϵ_n} is non-bijective for each n. ∎

This last result gives us some more geometrically intuitive ways to characterize a critical value of X. Based on this result, there are four possible ways in which $\epsilon > 0$ might be a critical value of X.

1) There is a sequence, $\epsilon_n \searrow \epsilon$, such that the map $X_\epsilon \to X_{\epsilon_n}$ is not injective, in which case we say that X is *ϵ-upper non-injective.*

2) There is a sequence, $\epsilon_n \searrow \epsilon$, such that the map $X_\epsilon \to X_{\epsilon_n}$ is not surjective, in which case we say that X is *ϵ-upper non-surjective.*

3) There is a sequence, $\epsilon_n \nearrow \epsilon$, such that the map $X_{\epsilon_n} \to X_\epsilon$ is not injective, in which case we say that X is *ϵ-lower non-injective.*

4) There is a sequence, $\epsilon_n \nearrow \epsilon$, such that the map $X_{\epsilon_n} \to X_\epsilon$ is not surjective, in which case we say that X is *ϵ-lower non-surjective.*

Of course, two or more of these cases may simultaneously hold for a given ϵ.

Using Lemmas 2.3.2 and 2.3.3, there are even more descriptive ways of characterizing these cases. If case 1 holds, this means that for each δ larger than but sufficiently close to ϵ, there is a non-trivial ϵ-loop, α, that is δ-null. In other words, there are non-trivial ϵ-loops that suddenly "collapse" when considered as δ-loops for δ greater than ϵ. Likewise, if case 3 holds, this means that for δ less than but sufficiently close to ϵ, there are non-trivial δ-loops that are ϵ-null. On the other hand, if case 2 holds, this means that for δ sufficiently close to but greater than ϵ there is a δ-chain that cannot be δ-refined to an ϵ-chain. Finally, if case 4 holds, this means that for δ less than but sufficiently close to ϵ, there are ϵ-chains that cannot be ϵ-refined to δ-chains. For these reasons, we call critical values of type 1 and 3 ***homotopy critical values***, since they indicate sudden changes in the structure of the ϵ-groups, or changes in what one might call the *ϵ-topology of X*. Furthermore, we call critical values of type 2 and 4 ***refinement critical values***, since they indicate refinability obstructions within X.

The refinement issues arising from cases 2 and 4 are interesting in their own right, and we will address these later on in more detail. In fact, case 4 turns out to be important enough that we need a special name for the instance in which it does *not* occur.

Definition 2.3.8 *Let X be a chain-connected metric space, and let $\epsilon > 0$ be given. We say that X is **ϵ-surjective from below** if there is some $\delta < \epsilon$ such that the map $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is surjective, or, equivalently, if every ϵ-chain in X can be ϵ-refined to a δ-chain.*

*If, for any $0 < \delta < \epsilon$, the map $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is surjective, meaning that any ϵ-chain can be arbitrarily refined for any ϵ, then X is an **infinitely refinable** - or, simply, **refinable** - space.*

Refinable spaces seem to be the best context in which to discuss the connection between the critical spectrum and topological properties of X. As we will show throughout, this is a large class of spaces. Any geodesic space is refinable, since we can take any ϵ-chain and simply insert midpoints between each consecutive pair of points until we obtain the desired refinement. Thus, all Riemannian manifolds are refinable spaces. Many non-geodesic spaces are also refinable, like the sphere $S^n \subset \mathbb{R}^{n+1}$ with its inherited Euclidean metric. In addition, many continua and compacta, including several commonly-studied fractals, are refinable spaces.

Example 2.3.9 *Let S^1 be the geodesic circle of circumference 1 that was studied in Example 2.2.6. Since geodesic spaces are refinable, there are no refinement critical values in $Cr(S^1)$. However, we have seen that the ϵ-covers of S^1 change at $\epsilon = 1/3$ from the trivial cover to $\mathbb{R}$. In other words, the map $\varphi_{\epsilon,1/3} : X_{1/3} \to X_\epsilon$ is not injective for all $\epsilon > 1/3$. Thus, 1/3 is a homotopy critical value of S^1, and - since there are no more topological changes in the ϵ-covers - it is the only critical value. That is, $Cr(S^1) = \{1/3\}$.*

For the geodesic torus, $T^2 = S^1_1 \times S^1_2$, in Example 2.2.8, the critical spectrum is $Cr(T^2) = \{1/3, 2/3\}$. There are no refinement critical values, but the maps $\varphi_{\epsilon,2/3} : S^1_1 \times \mathbb{R} \to T^2$ and $\varphi_{\delta,1/3} : \mathbb{R}^2 \to S^1_1 \times \mathbb{R}$ are not injective for all $\epsilon > 2/3$ and $\delta > 1/3$, respectively.

Now, consider the space, X, from Example 2.2.7. This is a geodesic circle of circumference 1 with a gap of length 1/4 and endowed with the subspace metric inherited from S^1. There is a *homotopy critical value at* $1/3$*, since* $X_{1/3}$ *is the real line with periodic gaps of length* $1/4$ *while* X_ϵ*, for* $\epsilon > 1/3$*, is the trivial cover. However,* $1/4$ *is a refinement critical value of* X*, since the map* $\varphi_{\epsilon,1/4}$ *is not surjective for* $1/4 < \epsilon \leq 1/3$*. Another way to see this is to consider the chain,* $\{a, b\}$*, where - as in Example 2.2.7 -* a *and* b *are the endpoints of the gap. For* $\epsilon > 1/4$ *but sufficiently close to* $1/4$*,* $\{a, b\}$ *is an* ϵ*-chain that cannot be* ϵ*-refined to a* $1/4$*-chain. In fact, as we noted before, no* $1/4$*-chain can even cross the gap, so it should at least be intuitively clear that this chain cannot be refined to a* $1/4$*-chain. We will prove this rigorously in Chapter 4.* ∎

In the case of a geodesic space, Sormani and Wei showed that the covering spectrum of a compact geodesic space was discrete in $\mathbb{R}_+$ and bounded above. The discreteness property need not hold for the critical spectrum of a general compact metric space, even if the space is simply connected. Some of the details in constructing such a space can be technical, so we will postpone those examples until Chapter 4, where several examples will be presented.

2.4 Metric, Topological, and Lifting Properties of X_ϵ

Lemma 2.4.1 *Let* X *be a chain-connected metric space. If* X *is locally compact and/or complete, then* X_ϵ *inherits the same property (or properties).*

Proof This follows directly from the fact that $\varphi_\epsilon : X_\epsilon \to X$ is a uniform local isometry, or a local isometry on every $\epsilon/2$-ball. Thus, if X is locally compact and $x \in X$, then some sufficiently small ball at x of radius less than $\epsilon/2$ will have compact closure. The lift of this compact set to any preimage point in X_ϵ will also be compact.

Likewise, completeness follows by the same reasoning. Any Cauchy sequence in X_ϵ will eventually lie in some ball of radius less than $\epsilon/2$. Projecting this sequence to X shows that it will converge. ∎

A metric space is said to be ***proper*** if all closed metric balls are compact. It should be noted that X_ϵ need not be proper, even if X is. This turns out to be a very important property, and we will investigate this more thoroughly in Chapter 5.

The following result was proved in [1] for general uniform spaces and general entourage covers. Thus, it applies equally well here.

Lemma 2.4.2 *Let* X *be a chain connected metric space. Suppose the* ϵ*-balls in* X *possess any one of the following properties: connected, chain-connected, path-connected. Then the whole space,* X_ϵ *is, respectively, connected, chain-connected, path-connected.*

Before moving to the next lemma, we will introduce some more terminology. In the following lemma, the *lift* of an ϵ-chain is defined exactly as for paths in traditional covering theory: if $f : Y \to X$, and if $\alpha = \{x_0, \ldots, x_n\}$ is an ϵ-chain in X, we say that $\tilde{\alpha} = \{y_0, \ldots, y_n\}$ is a lift of α if $f(\tilde{\alpha}) := \{f(y_0), \ldots, f(y_n)\} = \alpha$. The lift of a chain must, necessarily, have the same number of points as the chain to which it projects. Secondly, an ϵ-loop of the form $\{x_0, x_1, x_2, x_3 = x_0\}$ is called an ϵ***-triangle***, since the three distinct points making up this loop can be thought of as the vertices of a triangle with side lengths less than ϵ. Note that ϵ-triangles are ϵ-null.

Lemma 2.4.3 (Chain and Homotopy Lifting) *Let* $f : Y \to X$ *be a surjective map between metric spaces that is a bijection from* ϵ*-balls in* Y *onto* ϵ*-balls in* X*. Let* $\alpha = \{x_0, x_1, \ldots, x_n\}$ *be an* ϵ*-chain in* X*, and let* $\tilde{x}_0$ *be any point in the preimage of* x_0 *under* f*. Then* α *lifts uniquely* *to an ϵ-chain, $\tilde{\alpha}$, beginning at $\tilde{x}_0$. If, in addition, f has the property that ϵ-triangles in X lift to ϵ-triangles in Y, and if β is an ϵ-chain that begins at x_0 and is ϵ-homotopic to α, then the lifts of α and β to $\tilde{x}_0$ end at the same point and are ϵ-homotopic.*

Proof The proof of the first part is by induction on the number of points in α. If α contains one point, the result is trivial. Suppose $\alpha = \{x_0, x_1\}$. Then $x_1 \in B(x_0, \epsilon)$, and f, restricted to $B(\tilde{x}_0, \epsilon)$, is a bijection onto $B(x_0, \epsilon)$. Let $\tilde{x}_1$ be the unique point in $B(\tilde{x}_0, \epsilon)$ mapping to x_1. Then $\tilde{\alpha} = \{\tilde{x}_0, \tilde{x}_1\}$ is an ϵ-chain, and $f(\tilde{\alpha}) = \alpha$. This lift is unique because f is a bijection on $B(\tilde{x}_0, \epsilon)$: the second point of $\tilde{\alpha}$ must lie in this ball, and $\tilde{x}_1$ is the only choice for such a point.

Now, assume the result holds for all chains with n or fewer points for some $n \geq 2$. Let $\alpha = \{x_0, \ldots, x_n\}$ be an ϵ-chain with $n+1$ points, and let $\lambda = \{x_0, \ldots, x_{n-1}\}$. Using the inductive hypothesis, let $\tilde{\lambda} = \{\tilde{x}_0, \ldots, \tilde{x}_{n-1}\}$ be the unique lift of λ to $\tilde{x}_0$. Then $f(\tilde{x}_{n-1}) = x_{n-1}$, f is a bijection from $B(\tilde{x}_{n-1}, \epsilon)$ onto $B(x_{n-1}, \epsilon)$, and $x_n \in B(x_{n-1}, \epsilon)$. Let $\tilde{x}_n$ be the unique point in $B(\tilde{x}_{n-1}, \epsilon)$ mapping to x_n under f, and let $\tilde{\alpha} = \tilde{\lambda}\{\tilde{x}_n\} = \{\tilde{x}_0, \ldots, \tilde{x}_{n-1}, \tilde{x}_n\}$. Then $\tilde{\alpha}$ is an ϵ-chain and $f(\tilde{\alpha}) = \alpha$. If there were some other ϵ-chain, $\bar{\alpha} = \{\tilde{x}_0 = \bar{x}_0, \ldots, \bar{x}_n\}$, beginning at $\tilde{x}_0$ and projecting to α, then by the uniqueness part of the inductive hypothesis, the first n points of $\bar{\alpha}$ must agree with $\tilde{\lambda}$: $\bar{x}_i = \tilde{x}_i$ for $0 \leq i \leq n-1$. Finally, by the same reasoning as in the two-point case, the points $\tilde{x}_n$ and $\bar{x}_n$ must be the same. This proves existence and uniqueness of chain liftings.

To prove that the lifts of homotopic chains are homotopic, it suffices to consider the case in which α and β differ by only a basic move, since a general ϵ-homotopy simply consists of a sequence of ϵ-chains with each one differing from its predecessor by a basic move. So, suppose β is obtained by removing a point from α, say $\alpha = \{x_0, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_n\}$ and $\beta = \{x_0, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n\}$. Let $\tilde{\alpha}$ and $\tilde{\beta}$ denote the unique lifts of α and β to $\tilde{x}_0$. By uniqueness, $\tilde{\alpha}$ and $\tilde{\beta}$ clearly must agree for their first i points. Denote $\tilde{\alpha}$ by $\{\tilde{x}_0, \ldots, \tilde{x}_{i-1}, \tilde{x}_i, \tilde{x}_{i+1}, \ldots, \tilde{x}_n\}$ and $\tilde{\beta}$ by $\{\tilde{x}_0, \ldots, \tilde{x}_{i-1}, \tilde{y}_{i+1}, \ldots, \tilde{y}_n\}$. Since we can remove x_i from α, $\{x_{i-1}, x_i, x_{i+1}, x_{i-1}\}$ is an ϵ-triangle. The lift of this ϵ-triangle to $\tilde{x}_{i-1}$ is an ϵ-triangle by hypothesis, and, by uniqueness of lifts, the first three points of that triangle must be $\tilde{x}_{i-1}$, $\tilde{x}_i$, and $\tilde{x}_{i+1}$. Since a triangle is a loop, the fourth point of this lift must be $\tilde{x}_{i-1}$. In other words, we have $d_Y(\tilde{x}_{i-1}, \tilde{x}_{i+1}) < \epsilon$. So, $\tilde{x}_{i+1}$ and $\tilde{y}_{i+1}$ both lie in $B(\tilde{x}_{i-1}, \epsilon)$ and project under f to x_{i+1}, implying that $\tilde{x}_{i+1} = \tilde{y}_{i+1}$. Finally, by the uniqueness of lifts, the rest of $\tilde{\beta}$ must agree with $\tilde{\alpha}$. Thus, $\tilde{\beta}$ is obtained by removing a point from $\tilde{\alpha}$, so $\tilde{\alpha} \sim_\epsilon \tilde{\beta}$. For the case where β is obtained by adding a point to α, we can simply use the symmetry of ϵ-homotopy and reverse the roles of α and β in the previous argument. ∎

Corollary 2.4.4 *Let X be a chain-connected metric space, and, for some $\epsilon > 0$, let α and β be ϵ-chains beginning at a common point, $x \in X$, that are ϵ-homotopic. Then their lifts, $\tilde{\alpha}$ and $\tilde{\beta}$, to any $\tilde{x} \in \varphi_\epsilon^{-1}(x)$ are ϵ-homotopic in X_ϵ.*

Proof Let $\{z_0, z_1, z_2, z_0\}$ be an ϵ-triangle in X, and let $[\alpha]_\epsilon$ be any point in $\varphi_\epsilon^{-1}(z_0)$, where $\alpha = \{* = x_0, \ldots, x_n = z_0\}$. Let $\beta = \{* = x_0, \ldots, x_n = z_0, z_1\}$ and $\lambda = \{* = x_0, \ldots, x_n = z_0, z_2\}$. Then $d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon)$ and $d_\epsilon([\alpha]_\epsilon, [\lambda]_\epsilon)$ are less than ϵ. Moreover, $\beta^{-1}\lambda \sim_\epsilon \{z_1, z_0, z_2\}$, and, since $d(z_1, z_2) < \epsilon$, we can remove z_0 from this chain to conclude that $\beta^{-1}\lambda \sim_\epsilon \{z_1, z_2\}$. Thus, $d_\epsilon([\beta]_\epsilon, [\lambda]_\epsilon) < \epsilon$. So, $\{[\alpha]_\epsilon, [\beta]_\epsilon, [\lambda]_\epsilon, [\alpha]_\epsilon\}$ is an ϵ-triangle, and it projects to $\{z_0, z_1, z_2, z_0\}$. The uniqueness of lifts now implies that ϵ-triangles lift to ϵ-triangles. ∎

Lemma 2.4.5 *Let X be a chain-connected metric space, and let $\epsilon > 0$ be given. If $\alpha = \{* = x_0, \ldots, x_n\}$ is an ϵ-chain beginning at the base point, $* \in X$, then the unique lift of α to $\tilde{*} =$*

$[\{*\}]_\epsilon \in X_\epsilon$ *is given by*

$$\tilde{\alpha} = \Big\{[\{*\}]_\epsilon, [\{x_0, x_1\}]_\epsilon, \ldots, [\{x_0, \ldots, x_{n-1}\}]_\epsilon, [\{x_0, \ldots, x_n\}]_\epsilon = [\alpha]_\epsilon\Big\}.$$

In particular, the endpoint of the lift of α is the ϵ-class, $[\alpha]_\epsilon$, and the distances between consecutive points, as well as the chain length, are preserved in the lift.

Proof First, note that, for $i = 1, \ldots, n$,

$$\begin{aligned} d_\epsilon([\{x_0, \ldots, x_{i-1}\}]_\epsilon, [\{x_0, \ldots, x_{i-1}, x_i\}]_\epsilon) &= [\{x_{i-1}, \ldots, x_1, x_0, x_1, \ldots, x_{i-1}, x_i]_\epsilon \\ &= [\{x_{i-1}, x_i\}]_\epsilon \\ &= d(x_{i-1}, x_i) < \epsilon, \end{aligned}$$

since we can successively remove x_0, then each x_1, and so on via an ϵ-homotopy. Thus, $\tilde{\alpha}$ is an ϵ-chain. Clearly, we have $\varphi_\epsilon(\tilde{\alpha}) = \alpha$. Uniqueness of lifts now yields the result. ∎

This immediately yields the following.

Corollary 2.4.6 *If X is a chain-connected metric space and $\epsilon > 0$, then an ϵ-loop, γ, based at $*$ lifts to an ϵ-loop at $\tilde{*} \in X_\epsilon$ if and only if γ is ϵ-null. Thus, any representative of a nontrivial element of $\pi_\epsilon(X)$ lifts open to $\tilde{*}$.*

Lemma 2.4.7 *For a chain-connected metric space, X, and $\epsilon > 0$, X_ϵ is ϵ-connected and ϵ-* ***simply connected*** *(i.e. every ϵ-loop based at $\tilde{*} \in X_\epsilon$ is ϵ-null, or $\pi_\epsilon(X_\epsilon)$ is trivial).*

Proof The ϵ-connectivity follows immediately from Lemma 2.4.5. Now, given an ϵ-loop, $\tilde{\gamma}$, at $\tilde{*} \in X_\epsilon$, it will project to a loop, $\gamma := \varphi_\epsilon(\tilde{\gamma})$, at $*$. Moreover, γ is an ϵ-loop, since φ_ϵ preserves distances for all points within ϵ of each other. Since γ lifts to a closed loop, it is ϵ-null. By corollary 2.4.4, this ϵ-nullhomotopy will lift to X_ϵ. ∎

Lemmas 2.4.5 and 2.4.7 yield another equivalent condition to add to those in Lemma 2.3.2.

Corollary 2.4.8 *If X is chain-connected and $0 < \delta < \epsilon$, then $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is surjective if and only if X_ϵ is δ-connected.*

Proof Suppose $\varphi_{\epsilon\delta}$ is surjective, and let $[\alpha]_\epsilon \in X_\epsilon$ be given. Then α can be ϵ-refined to a δ-chain, or, equivalently, there is a δ-chain, α', in $[\alpha]_\epsilon$. Since $\delta < \epsilon$, the lift of α' to $\tilde{*}$ will be a δ-chain from $\tilde{*}$ to $[\alpha]_\epsilon$, showing that X_ϵ is δ-connected. Conversely, suppose X_ϵ is δ-connected, and let α be an ϵ-chain beginning at $*$. There is a δ-chain, $\tilde{\alpha}'$, from $\tilde{*}$ to $[\alpha]_\epsilon$ in X_ϵ. Since $\delta < \epsilon$, $\tilde{\alpha}'$ is also an ϵ-chain from $\tilde{*}$ to $[\alpha]_\epsilon$. Since X_ϵ is ϵ-simply connected, we have $\tilde{\alpha}' \sim_\epsilon \tilde{\alpha}$ in X_ϵ, where $\tilde{\alpha}$ is the lift of α given in Lemma 2.4.5. Thus, the projections of these chains, α' and α, will be ϵ-homotopic in X. But $\alpha' := \varphi_\epsilon(\tilde{\alpha}')$ is a δ-chain, so α is ϵ-homotopic to the δ-chain, α'. ∎

Remark Even though we have proved that ϵ-loops at $\tilde{*}$ are ϵ-null, this result holds for loops based at *any* point in X_ϵ. Let $\tilde{\gamma}$ be an ϵ-loop based at any $[\alpha]_\epsilon \in X_\epsilon$. Then the lift of $\alpha \subset X$ is an ϵ-chain, $\tilde{\alpha}$, from $\tilde{*}$ to $[\alpha]_\epsilon$. By Lemma 2.4.7, $\tilde{\alpha}\tilde{\gamma}\tilde{\alpha}^{-1}$ is ϵ-null in X_ϵ. Thus,

$$\tilde{\alpha}\tilde{\gamma}\tilde{\alpha}^{-1} \sim_\epsilon \{\tilde{*}\} \Rightarrow \tilde{\gamma} \sim_\epsilon \tilde{\alpha}^{-1}\{\tilde{*}\}\tilde{\alpha} \sim_\epsilon \{[\alpha]_\epsilon\},$$

showing that $\tilde{\gamma}$ is ϵ-null at $[\alpha]_\epsilon$.

We conclude with a general result on the lifting of maps, which is analogous to the well-known lifting theorem for path-connected covering spaces in classical topology.

Lemma 2.4.9 *Let X be a chain-connected metric space with base point, $*$, and suppose $f : (Y, \bar{*}) \to (X, *)$ is a pointed, surjective map with the following properties:* ***1)*** *f is a bijection and radial isometry from open δ-balls in Y onto open δ-balls in X for every $0 < \delta \leq \epsilon$, and* ***2)*** *ϵ-triangles in X lift to ϵ-triangles in Y. Then there exists a unique map $\rho : (X_\epsilon, \tilde{*}) \to (Y, \bar{*})$ such that* ***1)*** *$f \circ \rho = \varphi_\epsilon$,* ***2)*** *for every $0 < \delta \leq \epsilon$, ρ is a bijection from δ-balls in X_ϵ onto δ-balls in Y, and* ***3)*** *ρ is a radial isometry on ϵ-balls and an isometry on $(\epsilon/2)$-balls. Moreover, Y is ϵ-connected if and only if ρ is surjective, in which case ρ is a covering map.*

Remark In light of this result, one may consider $\varphi_\epsilon : X_\epsilon \to X$ to be universal in the category of ϵ-connected covering spaces of X having the same local properties as φ_ϵ.

Proof Define $\rho : X_\epsilon \to Y$ as follows. Given $[\alpha]_\epsilon \in X_\epsilon$, since X_ϵ is ϵ-connected, take any ϵ-chain, $\tilde{\lambda}$, in X_ϵ from $\tilde{*}$ to $[\alpha]_\epsilon$. This chain projects down to an ϵ-chain, λ, at $*$ in X, and then that chain can be lifted to an ϵ-chain, $\bar{\lambda}$, at $\bar{*} \in Y$. We set $\rho([\alpha]_\epsilon)$ equal to the endpoint of the lift, $\bar{\lambda}$. To see that ρ is well-defined, we note that if we took any other ϵ-chain, $\tilde{\sigma}$, from $\tilde{*}$ to $[\alpha]_\epsilon$, then - by Lemma 2.4.7 - $\tilde{\lambda}$ and $\tilde{\sigma}$ are ϵ-homotopic. Thus, their projections to X and the resulting lifts to Y are also homotopic and end at the same point, and ρ is well defined. Given $[\alpha]_\epsilon$, if x is the endpoint of α and $\bar{\alpha}$ denotes the lift of α to $\bar{*} \in Y$, with $\bar{x}$ the endpoint of $\bar{\alpha}$, then

$$(f \circ \rho)([\alpha]_\epsilon) = f(\bar{x}) = x = \varphi_\epsilon([\alpha]_\epsilon).$$

Thus, $f \circ \rho = \varphi_\epsilon$.

Given $0 < \delta \leq \epsilon$ and a ball $B([\alpha]_\epsilon, \delta) \subset X_\epsilon$, with $y_0 = \rho([\alpha]_\epsilon)$ and $x_0 = \varphi_\epsilon([\alpha]_\epsilon) = f(\rho([\alpha]_\epsilon))$, let $[\beta]_\epsilon$ be in $B([\alpha]_\epsilon, \delta)$ with y the endpoint of β. Then, by Lemma 2.2.4, there are representatives, $\alpha' = \{* = z_0, \ldots, z_m = x_0\} \in [\alpha]_\epsilon$ and $\beta' = \{* = z_0, \ldots, z_{m-1}, y\} \in [\beta]_\epsilon$, such that $d(x, y) < \delta$. Hence, we can transform β' via ϵ-homotopy to $\{* = z_0, \ldots, z_{m-1}, z_m, y\}$, and we will relabel and call *this* representative β'. If $\bar{\alpha}$ and $\bar{\beta}$ denote the lifts of α' and β', respectively, to $\bar{*} \in Y$, then, by uniqueness, the first $m+1$ points of $\bar{\beta}$ must agree with $\bar{\alpha}$, meaning that the endpoint of $\bar{\beta}$ must be within δ of the endpoint of $\bar{\alpha}$. This shows that $\rho(B([\alpha]_\epsilon, \delta)) \subset B(y_0, \delta)$. On the other hand, if $y \in B(y_0, \delta)$, then $x := f(y) \in B(x_0, \delta)$. If $\alpha = \{* = z_0, \ldots, z_m = x_0\}$, then let $\beta = \{* = z_0, \ldots, z_m = x_0, x\}$. We clearly have $[\beta]_\epsilon \in B([\alpha]_\epsilon, \delta)$. Moreover, the respective lifts, $\bar{\alpha}$ and $\bar{\beta}$, of α and β to $\bar{*}$ must agree for the first $m+1$ points of $\bar{\beta}$, meaning that the endpoint of $\bar{\beta}$ must lie within δ of the endpoint of $\bar{\alpha}$, y_0, and map to x under f. Since f is a bijective radial isometry on δ-balls, the only point satisfying these conditions is y. Hence, y must be the endpoint of $\bar{\beta}$, or $\rho([\beta]_\epsilon) = y$, showing that $\rho(B([\alpha]_\epsilon, \delta)) = B(y_0, \delta)$. If we had $\rho([\beta_1]_\epsilon) = \rho([\beta_2]_\epsilon)$ for $[\beta_1]_\epsilon, [\beta_2]_\epsilon \in B([\alpha]_\epsilon, \delta)$, then we would also have

$$\varphi_\epsilon([\beta_1]_\epsilon) = f(\rho([\beta_1]_\epsilon)) = f(\rho([\beta_2]_\epsilon)) = \varphi_\epsilon([\beta_2]_\epsilon),$$

contradicting that φ_ϵ is bijective on $B([\alpha]_\epsilon, \delta)$. Thus, ρ is a bijection from $B([\alpha]_\epsilon, \delta)$ onto $B(y_0, \delta)$, proving part 2.

The isometry conditions of part 3 now follow from part 2, since, on any ball, $B([\alpha]_\epsilon, \epsilon) \subset X_\epsilon$, we have $\rho = f^{-1} \circ \varphi_\epsilon$, and both f and φ_ϵ satisfy the desired properties. It also follows from conclusion 2 that if $[\alpha]_\epsilon, [\beta]_\epsilon \in \rho^{-1}(y)$ for some $y \in Y$, then $B([\alpha]_\epsilon, \epsilon/2) \cap B([\beta]_\epsilon, \epsilon/2) = \emptyset$. Thus, ρ is a covering map onto its image.

To see that the lifted map, ρ, is unique, suppose we had another map, ρ', satisfying the same properties. Given any $[\alpha]_\epsilon \in X_\epsilon$, let $\tilde{\alpha}$ be the lift of α to $\tilde{*}$ (which ends at $[\alpha]_\epsilon$). Let $\bar{\alpha}$ be the lift of this chain to $\bar{*}$. We then have $\alpha = \varphi_\epsilon(\tilde{\alpha}) = f(\rho(\tilde{\alpha})) = f(\rho'(\tilde{\alpha}))$, so, by uniqueness of chain lifts, $\rho(\tilde{\alpha})$ and $\rho'(\tilde{\alpha})$ must equal $\bar{\alpha}$. In particular, the endpoints of $\rho(\tilde{\alpha})$ and $\rho'(\tilde{\alpha})$, which are $\rho([\alpha]_\epsilon)$ and $\rho'([\alpha]_\epsilon)$, respectively, must both be the endpoint of $\bar{\alpha}$. That is, $\rho'([\alpha]_\epsilon) = \rho([\alpha]_\epsilon)$. Since $[\alpha]_\epsilon \in X_\epsilon$ was arbitrary, this shows that $\rho' = \rho$.

Now, if ρ is surjective, then, since X_ϵ is ϵ-connected and ρ preserves distances for all points within ϵ of each other, Y will also be ϵ-connected. Conversely, suppose Y is ϵ-connected, and let $y \in Y$ be given. Let $\bar{\alpha}$ be an ϵ-chain from $\bar{*}$ to y, and let $\alpha = f(\bar{\alpha})$ be the projected ϵ-chain in X from $*$ to $x := f(y)$. The lift, $\tilde{\alpha}$, of α to $\tilde{*} \in X_\epsilon$ ends at $[\alpha]_\epsilon$. By definition of ρ, then, we have $\rho([\alpha]_\epsilon) = y$, and ρ is surjective. ∎

2.5 The ϵ-intrinsic Property

Let X be a chain-connected metric space, with metric, d. Define $D_\epsilon : X \times X \to [0, \infty)$ by

$$D_\epsilon(x, y) = \inf\{L(\alpha) \,:\, \alpha \text{ is an } \epsilon\text{-chain from } x \text{ to } y\}.$$

We use the notation D_ϵ to distinguish this (what we will show to be a) metric from d_ϵ. This will be particularly important when we apply this construction to the space, X_ϵ.

A subtle but important point here is that the chain length, $L(\alpha)$, in this definition is taken in terms of the given metric, d. Thus, D_ϵ depends strongly on the given metric, d, since the length of a chain, $L(\alpha) = \sum_{i=1}^{n} d(x_{i-1}, x_i)$, depends explicitly on this metric. This is a situation very similar to the *induced* length metric on a path-connected metric space, (X, d), where the induced length metric, d_l, is defined by setting $d_l(x, y)$ equal to the infimum of the lengths of all rectifiable curves connecting x and y. Of course, the *length* of a curve, $\gamma : [a, b] \to X$, in this definition is defined to be

$$l(\gamma) = \sup \sum_{i=1}^{n} d(\gamma(t_{i-1}), \gamma(t_i)),$$

where the supremum is taken over all partitions, $\{a = t_0, t_1, \ldots, t_n = b\}$, of $[a, b]$. This shows the dependence of d_l on d; in fact, a path-connected metric space, (X, d), is a length space if and only if $d = d_l$. Hence, with this comparison in hand, one can think of D_ϵ as a discrete analog of the induced length metric, but if there are two or more metrics being discussed on a space, we will usually need to specify the metric with respect to which D_ϵ is defined.

Lemma 2.5.1 *If (X, d) is chain-connected, then, for any $\epsilon > 0$, D_ϵ is a metric on X satisfying $d(x, y) \le D_\epsilon(x, y)$ for all x, y. We call D_ϵ the **induced ϵ-metric determined by** d.*

Proof Since X is chain-connected, D_ϵ is defined for every $\epsilon > 0$. Clearly, we have $D_\epsilon(x, y) \ge 0$ and $D_\epsilon(x, x) = 0$ for all x and y, and symmetry follows from the fact that $L(\alpha) = L(\alpha^{-1})$ for any ϵ-chain.

Suppose $D_\epsilon(x, y) = 0$. Then, for every $\tau > 0$, there is an ϵ-chain, $\alpha = \{x = x_0, \ldots, x_n = y\}$, with $L(\alpha) < \tau$. The triangle inequality implies that $d(x, y) \le L(\alpha) < \tau$, so $x = y$.

Now, let $x, y, z \in X$ be given. Let β be a fixed but arbitrary ϵ-chain from y to z. Then, let α be any ϵ-chain from x to y. The concatenated ϵ-chain, $\alpha\beta$, is an ϵ-chain from x to z and has length $L(\alpha) + L(\beta)$. So, we have

$$D_\epsilon(x, z) \le L(\alpha) + L(\beta) \;\Rightarrow\; D_\epsilon(x, z) - L(\beta) \le L(\alpha).$$

Since α from x to y was arbitrary, we can take the infimum of the right-hand side over all such chains, giving us

$$D_\epsilon(x,z) - L(\beta) \le D_\epsilon(x,y) \Rightarrow D_\epsilon(x,z) - D_\epsilon(x,y) \le L(\beta).$$

Just as before, since β from y to z was arbitrary, we can take the infimum of the right-hand side over all such chains, giving us

$$D_\epsilon(x,z) - D_\epsilon(x,y) \le D_\epsilon(y,z) \Rightarrow D_\epsilon(x,z) \le D_\epsilon(x,y) + D_\epsilon(y,z).$$

Finally, if $\alpha = \{x = x_0, \ldots, x_n = y\}$ is any ϵ-chain from x to y, then the triangle inequality implies that, for all $x, y \in X$, $d(x,y) \le L(\alpha) \Rightarrow d(x,y) \le D_\epsilon(x,y)$. ∎

Definition 2.5.2 *A metric space, (X,d), that is ϵ-connected and is such that $d = D_\epsilon$ is said to be **ϵ-intrinsic.***

Lemma 2.5.3 *If (X,d) is a length space, then (X,d) is ϵ-intrinsic for every $\epsilon > 0$.*

Proof Let $x, y \in X$ be given. By the previous lemma, we have $d \le D_\epsilon$, so we need only prove the other inequality. Let $\gamma : [a,b] \to X$ be any rectifiable curve from x to y. Choose a partition of $[a,b]$, say $\{a = t_0, t_1, \ldots, t_n = b\}$, fine enough so that the length of each subsegment, $\gamma|_{[t_{i-1},t_i]}$, is strictly less than ϵ. For $i = 0, 1, \ldots, n$, let $x_i = \gamma(t_i)$, and note that $x_0 = x$ and $x_n = y$. In any length space, simply by definition of arclength, the length of any curve is at least as great as the metric distance between the endpoints of the curve. Thus, for each $i = 1, \ldots, n$, we have

$$d(x_{i-1}, x_i) = d(\gamma(t_{i-1}), \gamma(t_i)) \le l(\gamma|_{[t_{i-1},t_i]}) < \epsilon.$$

That is, $\alpha := \{x_0, \ldots, x_n\}$ is an ϵ-chain from x to y. Moreover, the previous inequality and the additivity of arclength imply that

$$D_\epsilon(x,y) \le L(\alpha) = \sum_{i=1}^{n} d(x_{i-1}, x_i) \le \sum_{i=1}^{n} l(\gamma|_{[t_{i-1},t_i]}) = l(\gamma).$$

Since γ was an arbitrary rectifiable curve from x to y, this shows that $D_\epsilon(x,y) \le d(x,y)$. ∎

Perhaps a more surprising result is the following. Note that we do *not* require that X be geodesic or even a length space in this result.

Proposition 2.5.4 *Let X be a chain-connected metric space, and let $\epsilon > 0$ be given. Then (X_ϵ, d_ϵ) is ϵ-intrinsic.*

Proof We have shown that X_ϵ is ϵ-connected, so D_ϵ is well-defined. Moreover, we already know that $d_\epsilon \le D_\epsilon$, so, again, we only need to prove the opposite inequality.

Let $[\alpha]_\epsilon$, $[\beta]_\epsilon \in X_\epsilon$ be given, and let x, y be the endpoints of α and β, respectively. By definition, $d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) = L([\alpha^{-1}\beta]_\epsilon)$ is the infimum of the lengths of all ϵ-chains in X that are ϵ-homotopic to $\alpha^{-1}\beta$. Let λ be any ϵ-chain from x to y that is ϵ-homotopic to $\alpha^{-1}\beta$. Since $\varphi_\epsilon([\alpha]_\epsilon) = x$, we can uniquely lift λ to an ϵ-chain, $\tilde{\lambda}$, beginning at $[\alpha]_\epsilon$. We need to show that $\tilde{\lambda}$ ends at $[\beta]_\epsilon$. But the concatenated chain, $\alpha\lambda$, is ϵ-homotopic to β, and if $\tilde{\alpha}$ denotes the unique lift of α to $\tilde{*}$ - which must end at $[\alpha]_\epsilon$ - then $\tilde{\alpha}\tilde{\lambda}$ must be the unique lift of $\alpha\lambda$ to $\tilde{*}$. Since $\alpha\lambda \sim_\epsilon \beta$, it follows from our previous results that $\tilde{\alpha}\tilde{\lambda}$ must end at $[\beta]_\epsilon$. Hence, $\tilde{\lambda}$ ends at $[\beta]_\epsilon$.

Moreover, the length of $\tilde{\lambda}$ is the same as the length of λ, since φ_ϵ is a radial isometry on ϵ-balls. Thus, we have

$$D_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) \le L(\tilde{\lambda}) = L(\lambda).$$

Since $\lambda \in [\alpha^{-1}\beta]_\epsilon$ was arbitrary, taking the infimum of the right-hand side shows that

$$D_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) \le L([\alpha^{-1}\beta]_\epsilon) = d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon). \quad \blacksquare$$

The following example shows how the ϵ-intrinsic metric induced by a metric, d, can differ from d when (X, d) is not geodesic.

Example 2.5.5 *Let X be the metric subspace of $\mathbb{R}^2$ formed by the line segments from $(0,0)$ to $(0,1)$, from $(0,0)$ to $(1,0)$, and from $(1,0)$ to $(1,1)$. We call this space a tuning fork space. Let $x = (0,1)$ and $y = (1,1)$. See Figure 2.2. It is easy to see that, for any $\epsilon > 0$, all ϵ-loops in X are trivial. Thus, $\varphi_\epsilon : X_\epsilon \to X$ is a homeomorphism. Knowing that X_ϵ is ϵ-intrinsic, it follows that, in this case, X_ϵ is simply X with the ϵ-intrinsic metric, d_ϵ, induced by the Euclidean metric, d.*

Now, for large enough ϵ - say, greater than the diameter of X - d and d_ϵ agree. As ϵ decreases to 0, however, the metric d_ϵ becomes more and more distorted relative to d. For instance, when $\epsilon > 1$, the distance between x and y in X_ϵ is equal to $d(x, y)$, because $\{x, y\}$, itself, is an ϵ-chain from x to y. For $\epsilon \le 1$, however, an ϵ-chain from x to y must traverse around the tuning fork, meaning that it must have at least one point in each of the three segments making up X. In fact, it can be shown that, for $\epsilon = 1$, the distance between x and y in X_ϵ is the length of the chain $\{x, u, (1/2, 0), v, y\}$ shown in Figure 2.2. This chain is not a 1-chain; it is a chain where the distance between each consecutive pair of points is less than or equal to 1. However, the length of this chain is the infimum of the lengths of all 1-chains from x to y. Thus, for $\epsilon = 1$, we have $d(x, y) = 1$ while $d_\epsilon(x, y) = 4 - \sqrt{3}$. $\blacksquare$

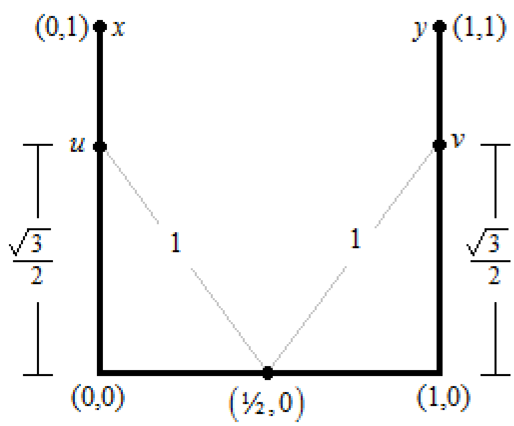

Figure 2.2: The Tuning Fork

Now, consider the case when X is a geodesic space, and let $\epsilon > 0$ be given. Since the ϵ-balls in a geodesic space are path-connected, we know that X_ϵ is path-connected and locally path-connected. As we mentioned in the section concerning the Spanier covers, we can lift the geodesic metric of X to a length metric on X_ϵ, which we will denote, for the time being, by d_l. The precise definition of this metric is as follows: we define the length of a curve, $\tilde{\gamma}$, in X_ϵ to be the length of its projection, $\gamma := \varphi_\epsilon \circ \tilde{\gamma}$, and we then define d_l in the usual way, letting $d_l([\alpha]_\epsilon, [\beta]_\epsilon)$ be the infimum of the lengths of all curves connecting $[\alpha]_\epsilon$ and $[\beta]_\epsilon$. (One can show - in the process of establishing that this is, indeed, a metric - that every pair of points can be joined by a rectifiable curve, so the metric is well-defined.) In particular, this means that lengths of curves are preserved under φ_ϵ. Moreover, it turns out that $\varphi_\epsilon : (X_\epsilon, d_l) \to X$ is a radial isometry on ϵ-balls and an isometry on $\epsilon/2$-balls, just like $\varphi_\epsilon : (X_\epsilon, d_\epsilon) \to X$. In fact, we have the following result.

Lemma 2.5.6 *If X is a geodesic space, and if d_l denotes the lifted length metric on X_ϵ, then $d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) < \epsilon$ if and only if $d_l([\alpha]_\epsilon, [\beta]_\epsilon) < \epsilon$, so the ϵ-balls in both metrics coincide. Moreover, if either of these inequalities holds, then the two distances are equal: $d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) = d_l([\alpha]_\epsilon, [\beta]_\epsilon)$.*

Proof Let $[\alpha]_\epsilon$, $[\beta]_\epsilon \in X_\epsilon$ be given, and let $x = \varphi_\epsilon([\alpha]_\epsilon)$, $y = \varphi_\epsilon([\beta]_\epsilon)$. Let $B_\epsilon([\alpha]_\epsilon, \epsilon)$ and $B_l([\alpha]_\epsilon, \epsilon)$ denote the metric balls with respect to the metrics d_ϵ and d_l, respectively. Suppose $d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) < \epsilon$. The map $\varphi_\epsilon : B_\epsilon([\alpha]_\epsilon, \epsilon) \to B(x, \epsilon)$ is a homeomorphism and radial isometry. Since X is geodesic, there is a minimal geodesic, γ, from x to y, and γ must necessarily lie in $B(x, \epsilon)$. The map $\varphi_\epsilon^{-1} : B(x, \epsilon) \to B_\epsilon([\alpha]_\epsilon, \epsilon)$ will take γ to a rectifiable curve, $\tilde{\gamma}$, from $[\alpha]_\epsilon$ to $[\beta]_\epsilon$ in $B([\alpha]_\epsilon, \epsilon)$, and $l(\tilde{\gamma}) = l(\gamma)$. By definition, we have $d_l([\alpha]_\epsilon, [\beta]_\epsilon) \leq l(\tilde{\gamma}) = l(\gamma) = d(x, y) = d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon)$. If this were a strict inequality, then there would be a rectifiable curve from $[\alpha]_\epsilon$ to $[\beta]_\epsilon$ of length strictly less than $d(x, y)$, and this curve would project to a rectifiable curve in X from x to y of length strictly less than $d(x, y)$, which cannot be the case. Thus, we actually have $d_l([\alpha]_\epsilon, [\beta]_\epsilon) = l(\tilde{\gamma}) = l(\gamma) = d(x, y) = d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon)$. Note that this also shows that when $d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) < \epsilon$, the two metrics agree, which proves the second statement of the lemma. This also shows that the length space (X_ϵ, d_l) is *locally* geodesic, even if it is not globally so.

Conversely, suppose $d_l([\alpha]_\epsilon, [\beta]_\epsilon) < \epsilon$. Then there is a rectifiable curve, $\tilde{\gamma}$, from $[\alpha]_\epsilon$ to $[\beta]_\epsilon$ of length strictly less than ϵ. The projected curve, $\gamma := \varphi_\epsilon(\tilde{\gamma})$, is, then, a rectifiable curve from x to y of length strictly less than ϵ. Thus, $y \in B(x, \epsilon)$ and γ must necessarily lie in this ball. Under the map $\varphi_\epsilon^{-1} : B(x, \epsilon) \to B_\epsilon([\alpha]_\epsilon, \epsilon)$, the curve $\varphi_\epsilon^{-1}(\gamma)$ will be a rectifiable lift of γ from $[\alpha]_\epsilon$ to $[\beta]_\epsilon$ lying in $B_\epsilon([\alpha]_\epsilon, \epsilon)$. By uniqueness of path lifts, this curve must agree with $\tilde{\gamma}$, showing that $\tilde{\gamma}$ lies in $B_\epsilon([\alpha]_\epsilon, \epsilon)$. Thus, $[\beta]_\epsilon \in B([\alpha]_\epsilon, \epsilon)$. ∎

Consequently, the metrics d_l and d_ϵ agree on all pairs of points within ϵ of each other with respect to either metric. With this result and the previous proposition in hand, we can now show that this lifted geodesic metric, d_l, agrees *globally* with d_ϵ when X is a geodesic space.

Theorem 2.5.7 *Let X be a geodesic space, and let $\epsilon > 0$ be given. Let d_l denote the lifted length metric on X_ϵ. Then $d_l = d_\epsilon$.*

Proof Let D_ϵ denote the ϵ-metric induced by d_ϵ. Then we know from Proposition 2.5.4 that $d_\epsilon = D_\epsilon$. So, we will actually show that $d_l = D_\epsilon$.

Let $[\alpha]_\epsilon, [\beta]_\epsilon \in X_\epsilon$ and $\tau > 0$ be given, and let $\tilde{\gamma} : [0, 1] \to X_\epsilon$ be a rectifiable curve between $[\alpha]_\epsilon$ and $[\beta]_\epsilon$ such that $d_l([\alpha]_\epsilon, [\beta]_\epsilon) \leq l(\tilde{\gamma}) < d_l([\alpha]_\epsilon, [\beta]_\epsilon) + \tau$. Choose a partition of $[0, 1]$, say $\{0 = t_0, \ldots, t_n = 1\}$, such that each subsegment, $\tilde{\gamma}|_{[t_{i-1}, t_i]}$, has length strictly shorter than ϵ.

For each $i = 0, 1, \ldots, n$, let $\tilde{x}_i = \tilde{\gamma}(t_i)$, and let $\tilde{\alpha} = \{[\alpha]_\epsilon = \tilde{x}_0, \ldots, \tilde{x}_n = [\beta]_\epsilon\}$. We have, for each $i = 1, \ldots, n$,

$$d_l(\tilde{x}_{i-1}, \tilde{x}_i) \leq l(\tilde{\gamma}|_{[t_{i-1},t_i]}) < \epsilon,$$

from which it follows that

$$d_\epsilon(\tilde{x}_{i-1}, \tilde{x}_i) = d_l(\tilde{x}_{i-1}, \tilde{x}_i) < \epsilon, \quad i = 1, \ldots, n.$$

Thus, $\tilde{\alpha}$ is an ϵ-chain in (X_ϵ, d_ϵ) from $[\alpha]_\epsilon$ to $[\beta]_\epsilon$. In addition, we have

$$L(\tilde{\alpha}) = \sum_{i=1}^{n} d_\epsilon(\tilde{x}_{i-1}, \tilde{x}_i) \leq \sum_{i=1}^{n} l(\tilde{\gamma}|_{[t_{i-1},t_i]}) = l(\tilde{\gamma}) = d_l([\alpha]_\epsilon, [\beta]_\epsilon) + \tau.$$

It now follows that $D_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) \leq d_l([\alpha]_\epsilon, [\beta]_\epsilon) + \tau$. Since τ was arbitrary, this shows that $D_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) \leq d_l([\alpha]_\epsilon, [\beta]_\epsilon)$.

On the other hand, let $\tilde{\lambda} = \{[\alpha]_\epsilon = \tilde{x}_0, \ldots, \tilde{x}_n = [\beta]_\epsilon\}$ be an ϵ-chain (with respect to d_ϵ) from $[\alpha]_\epsilon$ to $[\beta]_\epsilon$. Then, for each $i = 1, \ldots, n$, we have $d_l(\tilde{x}_{i-1}, \tilde{x}_i) = d_\epsilon(\tilde{x}_{i-1}, \tilde{x}_i) < \epsilon$. Let $\tau > 0$ be given, and, for each $i = 1, \ldots, n$, choose a rectifiable curve, $\tilde{\gamma}_i$, from $\tilde{x}_{i-1}$ to $\tilde{x}_i$ such that $d_l(\tilde{x}_{i-1}, \tilde{x}_i) \leq l(\tilde{\gamma}_i) < d_l(\tilde{x}_{i-1}, \tilde{x}_i) + \tau/n$. Let $\tilde{\gamma}$ be the piecewise curve from $[\alpha]_\epsilon$ to $[\beta]_\epsilon$ formed by joining the curves, $\tilde{\gamma}_i$. We then have

$$d_l([\alpha]_\epsilon, [\beta]_\epsilon) \leq l(\tilde{\gamma}) = \sum_{i=1}^{n} l(\tilde{\gamma}_i) < \tau + \sum_{i=1}^{n} d_\epsilon(\tilde{x}_{i-1}, \tilde{x}_i) = L(\tilde{\lambda}) + \tau.$$

Since τ was arbitrary, we have $d_l([\alpha]_\epsilon, [\beta]_\epsilon) \leq L(\tilde{\lambda})$. Since $\tilde{\lambda}$ was arbitrary, this shows that $d_l([\alpha]_\epsilon, [\beta]_\epsilon) \leq D_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon)$. ∎

As a consequence of this result, from here on, whenever X is a godesic space, we will drop the notation d_l for the length metric on X_ϵ and simply remember that d_ϵ is, in fact, the lifted length metric in this case.

2.6 Group Action of $\pi_\epsilon(X)$ on X_ϵ

In [1], Plaut and Berestovskii also proved a very useful connection between the ϵ-covers and ϵ-groups which we will use later on. In short, $\pi_\epsilon(X)$ acts on X_ϵ by pre-concatenation, and the action is *discrete* and, therefore, free and properly discontinuous. When we say that the action is discrete, we mean the following: if a group element maps a point of X_ϵ, say $[\alpha]_\epsilon$, to a point that is within ϵ of $[\alpha]_\epsilon$, then that group element must be the identity.

Theorem 2.6.1 *Given $[\gamma]_\epsilon \in \pi_\epsilon(X)$, define a map, $h_\gamma : X_\epsilon \to X_\epsilon$, by $h_\gamma([\alpha]_\epsilon) = [\gamma\alpha]_\epsilon$. Then h_γ is an isometry, and this yields a group action $\Theta : \pi_\epsilon(X) \times X_\epsilon \to X_\epsilon$ defined by $\Theta([\gamma]_\epsilon, [\alpha]_\epsilon) = h_\gamma([\alpha]_\epsilon) = [\gamma\alpha]_\epsilon$. This action satisifies the following properties.*

1) The action is discrete: if $d_\epsilon(h_{\gamma_1}([\alpha]_\epsilon), h_{\gamma_2}([\alpha]_\epsilon)) < \epsilon$, then $[\gamma_1]_\epsilon = [\gamma_2]_\epsilon$. This implies that the action is free and properly discontinuous.

2) If $[\alpha]_\epsilon$ and $[\beta]_\epsilon$ are such that α and β end at the same point, then $[\alpha]_\epsilon$ and $[\beta]_\epsilon$ are in the same orbit.

3) The metric, d_ϵ, is left invariant with respect to this action.

Proof To see that h_γ is surjective, note that if $[\beta]_\epsilon \in X_\epsilon$, then $h_\gamma([\gamma^{-1}\beta]_\epsilon) = [\beta]_\epsilon$. The fact that h_γ is an isometry follows because

$$d_\epsilon(h_\gamma([\alpha]_\epsilon), h_\gamma([\beta]_\epsilon)) = d_\epsilon([\gamma\alpha]_\epsilon, [\gamma\beta]_\epsilon) = L([\alpha^{-1}\gamma^{-1}\gamma\beta]_\epsilon) = L([\alpha^{-1}\beta]_\epsilon)$$
$$\Rightarrow d_\epsilon(h_\gamma([\alpha]_\epsilon), h_\gamma([\beta]_\epsilon)) = d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon).$$

Then Θ is a well-defined action, since the composition $h_{\gamma_1} \circ h_{\gamma_2}$ is just the isometry $h_{\gamma_1\gamma_2}$.

Now, suppose that $d_\epsilon(h_{\gamma_1}([\alpha]_\epsilon), h_{\gamma_2}([\alpha]_\epsilon)) < \epsilon$. Then $L([\alpha^{-1}\gamma_1^{-1}\gamma_2\alpha]_\epsilon) < \epsilon$, implying that $\alpha^{-1}\gamma_1^{-1}\gamma_2\alpha$ is ϵ-null. So, if x is the endpoint of α,

$$\alpha^{-1}\gamma_1^{-1}\gamma_2\alpha \sim_\epsilon \{x\} \Rightarrow \gamma_1^{-1}\gamma_2 \sim_\epsilon \alpha\alpha^{-1} \sim_\epsilon \{*\}.$$

Thus, $\gamma_1 \sim_\epsilon \gamma_2$, proving the first part of 1.

For part 2, if α and β end at the same point, then $\gamma = \alpha\beta^{-1}$ is an ϵ-loop at $*$. It follows that $h_\gamma([\beta]_\epsilon) = [\gamma\beta]_\epsilon = [\alpha\beta^{-1}\beta]_\epsilon = [\alpha]_\epsilon$. Left invariance follows easily, since $d_\epsilon(h_\gamma([\alpha]_\epsilon), h_\gamma([\beta]_\epsilon)) = L([\alpha^{-1}\gamma^{-1}\gamma\beta]_\epsilon) = L([\alpha^{-1}\beta]_\epsilon) = d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon)$.

Finally, a discrete action is clearly free. To see why it also implies proper discontinuity of the action, first let $[\alpha]_\epsilon \in X_\epsilon$ be given. The set of all $[\gamma]_\epsilon \in \pi_\epsilon(X)$ such that

$$h_\gamma(B([\alpha]_\epsilon, \epsilon/2)) \cap B([\alpha]_\epsilon, \epsilon/2) \neq \emptyset$$

contains only the identity, for if $[\gamma]_\epsilon$ satisfies this relation, then we have $d_\epsilon(h_\gamma([\alpha]_\epsilon), [\alpha]_\epsilon) < \epsilon$. The discreteness, then, implies that $[\gamma]_\epsilon = [\{*\}]_\epsilon$. Next, we claim that if $[\alpha]_\epsilon$ and $[\beta]_\epsilon$ are not in the same orbit, then there is some $r > 0$ such that the balls, $U := B([\alpha]_\epsilon, r)$ and $V := B([\beta]_\epsilon, r)$, satisfy $U \cap h_\gamma(V) = \emptyset$ for all $[\gamma]_\epsilon \in \pi_\epsilon(X)$. Suppose, toward a contradiction, that this were not the case. Then we could find a sequence, $\{[\gamma_n]_\epsilon\}$, such that $d_\epsilon([\alpha]_\epsilon, [\gamma_n\beta]_\epsilon) \to 0$. Let x and y be the endpoints of α and β, respectively. Then, for n large enough that $d_\epsilon([\alpha]_\epsilon, [\gamma_n\beta]_\epsilon) < \epsilon$, we have $L([\alpha^{-1}\gamma_n\beta]_\epsilon) < \epsilon$, implying that $L([\alpha^{-1}\gamma_n\beta]_\epsilon) = d(x, y)$ for such n. It follows that $d(x, y) = 0 \Rightarrow x = y$. Hence, α and β end at the same point. Part 2 and the condition $d_\epsilon([\alpha]_\epsilon, [\gamma_n\beta]_\epsilon) \to 0$ now imply that $[\alpha]_\epsilon$ and $[\beta]_\epsilon$ are in the same orbit, a contradiction. Hence, we have $U \cap h_\gamma(V) = \emptyset$ for all $[\gamma]_\epsilon \in \pi_\epsilon(X)$. ∎

Of course, one obvious consequence of this result is that the orbits resulting from this action are closed; this is a general consequence of free and properly discontinuous group actions. Moreover, two distinct orbits cannot have a zero distance between them. To see why, let $[[\alpha]]_\epsilon$ denote the orbit of $[\alpha]_\epsilon \in X_\epsilon$, and suppose we had

$$0 = dist([[\alpha]]_\epsilon, [[\beta]]_\epsilon) = \inf\{d_\epsilon([\lambda]_\epsilon, [\sigma]_\epsilon) \ : \ [\lambda]_\epsilon \in [[\alpha]]_\epsilon, \ [\sigma]_\epsilon \in [[\beta]]_\epsilon\}.$$

Let x, y denote the endpoints of α and β, respectively. Then, using the left-invariance of the action, for each $n \in \mathbb{N}$, we could find $[\gamma_n]_\epsilon \in \pi_\epsilon(X)$ such that $d_\epsilon([\alpha]_\epsilon, [\gamma_n]_\epsilon[\beta]_\epsilon) < 1/n$. Thus, for all n large enough so that $1/n < \epsilon$, we would have $d(x, y) < 1/n$, implying that $x = y$. That is, α and β end at the same point, meaning that $[\alpha]_\epsilon$ and $[\beta]_\epsilon$ are in the same orbit.

Hence, we can define a metric, d_q, on the quotient space, $X_\epsilon/\pi_\epsilon(X)$, by

$$d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) = \inf\{d_\epsilon([\lambda]_\epsilon, [\sigma]_\epsilon) \ : \ [\lambda]_\epsilon \in [[\alpha]]_\epsilon, \ [\sigma]_\epsilon \in [[\beta]]_\epsilon\}.$$

Symmetry and positivity follow easily, and positive definiteness follows from the preceding observation. To prove the triangle inequality, we first note the following simplification of the definition of d_q. Given $[[\alpha]]_\epsilon$, $[[\beta]]_\epsilon \in X_\epsilon/\pi_\epsilon(X)$,

$$d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) = \inf\{d_\epsilon([\alpha]_\epsilon, [\sigma]_\epsilon) \ : \ [\sigma]_\epsilon \in [[\beta]]_\epsilon\}. \tag{2.1}$$

Clearly, the left-hand side is less than or equal to the right-hand side by definition of d_q. To see the other inequality, let $[\lambda]_\epsilon \in [[\alpha]]_\epsilon$ and $[\sigma]_\epsilon \in [[\beta]]_\epsilon$ be given. Choose $[\gamma]_\epsilon \in \pi_\epsilon(X)$ so that $[\lambda]_\epsilon = [\gamma]_\epsilon[\alpha]_\epsilon$, and let L_α denote the infimum on the right side of equation 2.1. Then

$$d_\epsilon([\lambda]_\epsilon, [\sigma]_\epsilon) = d_\epsilon([\gamma]_\epsilon[\alpha]_\epsilon, [\sigma]_\epsilon) = d_\epsilon([\alpha]_\epsilon, [\gamma^{-1}]_\epsilon[\sigma]_\epsilon) \geq L_\alpha,$$

showing that $d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) \geq L_\alpha$. Now, given $[[\alpha]]_\epsilon$, $[[\beta]]_\epsilon$, $[[\lambda]]_\epsilon$, and any $\tau > 0$, choose $[\beta']_\epsilon \in [[\beta]]_\epsilon$ so that $d_q([[\beta]]_\epsilon, [[\lambda]]_\epsilon) \leq d_\epsilon([\beta']_\epsilon, [\lambda]_\epsilon) < d_q([[\beta]]_\epsilon, [[\lambda]]_\epsilon) + \tau$. Then choose $[\alpha']_\epsilon \in [[\alpha]]_\epsilon$ so that $d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) \leq d_\epsilon([\alpha']_\epsilon, [\beta']_\epsilon) < d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) + \tau$. We then have

$$d_\epsilon([\alpha']_\epsilon, [\beta']_\epsilon) + d_\epsilon([\beta']_\epsilon, [\lambda]_\epsilon) < d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) + d_q([[\beta]]_\epsilon, [[\lambda]]_\epsilon) + 2\tau$$

$$\Rightarrow d_q([[\alpha]]_\epsilon, [[\lambda]]_\epsilon) \leq d_\epsilon([\alpha']_\epsilon, [\lambda]_\epsilon) < d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) + d_q([[\beta]]_\epsilon, [[\lambda]]_\epsilon) + 2\tau.$$

This holds for all $\tau > 0$, so, letting $\tau \to 0$, we get the triangle inequality.

It is natural, then, to ask how the space $(X_\epsilon/\pi_\epsilon(X), d_q)$ compares with the given metric space, (X, d). First, we point out that there is a natural map, $f : X_\epsilon/\pi_\epsilon(X) \to X$, which is a bijection. Given $[[\alpha]]_\epsilon \in X_\epsilon/\pi_\epsilon(X)$, we let $f([[\alpha]]_\epsilon) = x$ where x is the endpoint of α. Any $[\alpha']_\epsilon \in [[\alpha]]_\epsilon$ can be expressed as $[\gamma]_\epsilon[\alpha]_\epsilon$ for some $[\gamma]_\epsilon \in \pi_\epsilon(X)$, meaning that any chain in any equivalence class contained in the orbit, $[[\alpha]]_\epsilon$, will end at the same point as α. Thus, f is well-defined. It is surjective, because - given $x \in X$ - we can take any ϵ-chain, α, from the base point, $*$, to x, and we will have $f([[\alpha]]_\epsilon) = x$. It is injective by part 2 of Theorem 2.6.1. Moreover, if $\pi : X_\epsilon \to X_\epsilon/\pi_\epsilon(X)$ denotes the quotient map, then we have $f \circ \pi = \varphi_\epsilon$.

Theorem 2.6.2 *Let X be a chain-connected metric space, and let $\epsilon > 0$ be given. Then $f : (X_\epsilon/\pi_\epsilon, d_q) \to (X, d)$ is a bijective, 1-Lipschitz, uniform local isometry (i.e. an isometry on $\epsilon/2$-balls) and, thus, a homeomorphism. Moreover, the space $X_\epsilon/\pi_\epsilon(X)$ satisfies the following properties.*

1) *$(X_\epsilon/\pi_\epsilon(X), d_q)$ is ϵ-intrinsic and chain-connected.*

2) *If X is locally compact, complete, proper, and/or compact, then $X_\epsilon/\pi_\epsilon(X)$ possesses the same property (or properties).*

3) *$f : X_\epsilon/\pi_\epsilon(X) \to X$ is an isometry if and only if (X, d) is ϵ-intrinsic.*

By part 3 of this result, we see that there are many metric spaces, X, and ϵ-values such that X is only homeomorphic and locally isometric - but not fully isometric - to $X_\epsilon/\pi_\epsilon(X)$. Of course, if X is geodesic, then it is also ϵ-intrinsic and f is an isometry. In general, what this result tells us is that $X_\epsilon/\pi_\epsilon(X)$ is topologically the same as X and "looks like X" locally, but their global geometry may differ.

Proof Given $[[\alpha]]_\epsilon$, $[[\beta]]_\epsilon \in X_\epsilon/\pi_\epsilon(X)$, let $x = f([[\alpha]]_\epsilon)$ and $y = f([[\beta]]_\epsilon)$. Then, given any $[\beta']_\epsilon \in [[\beta]]_\epsilon$, we have

$$d(x, y) \leq L([\alpha^{-1}\beta']_\epsilon) = d_\epsilon([\alpha]_\epsilon, [\beta']_\epsilon),$$

which, by equation 2.1, shows that

$$d(f([[\alpha]]_\epsilon), f([[\beta]]_\epsilon)) = d(x, y) \leq d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon).$$

This shows that f is 1-Lipschitz.

Next, we will show that $\pi : X_\epsilon \to X_\epsilon/\pi_\epsilon(X)$ is a covering map with the same local properties as φ_ϵ. Suppose $d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) < \epsilon$. There cannot be any other element, $[\beta']_\epsilon \in [[\beta]]_\epsilon$, such that $d_\epsilon([\alpha]_\epsilon, [\beta']_\epsilon) < d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon)$, for if such an element existed, it would lie in $B([\alpha]_\epsilon, \epsilon)$ and map, under φ_ϵ, to the same point as $[\beta]_\epsilon$ in $B(x, \epsilon)$, contradicting the fact that φ_ϵ is a bijection on ϵ-balls. Thus, for every $[\beta']_\epsilon \in [[\beta]]_\epsilon$, we have $d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) \leq d_\epsilon([\alpha]_\epsilon, [\beta']_\epsilon)$, from which it follows that the infimum of $\{d_\epsilon([\alpha]_\epsilon, [\beta']_\epsilon) \,:\, [\beta']_\epsilon \in [[\beta]]_\epsilon\}$ is actually attained by $[\beta]_\epsilon$. In other words, $d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) = d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon)$, and we see that π is an isometry on all pairs of points that are strictly within ϵ of each other. It now follows - just as for φ_ϵ - that $\pi : X_\epsilon \to X_\epsilon/\pi_\epsilon(X)$ is a radial isometry on ϵ-balls and an isometry on $\epsilon/2$-balls. Moreover, suppose $[\alpha]_\epsilon$ and $[\alpha']_\epsilon$ are distinct elements of the same orbit, and suppose further that $d_\epsilon([\alpha]_\epsilon, [\alpha']_\epsilon) < \epsilon$. Since these elements are in the same orbit, the chains α and α' end at the same point. Thus, $\alpha^{-1}\alpha'$ is an ϵ-loop. Since

$$\epsilon > d_\epsilon([\alpha]_\epsilon, [\alpha']_\epsilon) = L([\alpha^{-1}\alpha']_\epsilon),$$

it follows that $\alpha \sim_\epsilon \alpha'$, or $[\alpha]_\epsilon = [\alpha']_\epsilon$. Hence, two distinct elements of the same orbit must be at least a distance ϵ apart, which implies that the open balls of radius $\epsilon/2$ centered at the elements of a given orbit are disjoint. So, π is a covering map with the same properties as φ_ϵ. Since $f \circ \pi = \varphi_\epsilon$, f is also a radial isometry on ϵ-balls and an isometry on $\epsilon/2$-balls.

Since X is chain-connected, it now follows that $X_\epsilon/\pi_\epsilon(X)$ is, also, since any δ-chain, α, in X - for $\delta \leq \epsilon$ - will map to a δ-chain in $X_\epsilon/\pi_\epsilon(X)$ under f^{-1}. To see that $(X_\epsilon/\pi_\epsilon(X), d_q)$ is ϵ-intrinsic, we first recall that X_ϵ is ϵ-intrinsic, that $\pi : X_\epsilon \to X_\epsilon/\pi_\epsilon(X)$ is an isometry on all pairs of points within ϵ of each other, and that, consequently, π preserves the lengths of ϵ-chains. So, let $[[\alpha]]_\epsilon$, $[[\beta]]_\epsilon \in X_\epsilon/\pi_\epsilon(X)$ be given. For each $n \in \mathbb{N}$, choose $[\beta_n]_\epsilon \in [[\beta]]_\epsilon$ so that

$$d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) \leq d_\epsilon([\alpha]_\epsilon, [\beta_n]_\epsilon) < d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) + \frac{1}{n}.$$

Then choose an ϵ-chain, $\tilde{\lambda}_n$, in X_ϵ from $[\alpha]_\epsilon$ to $[\beta_n]_\epsilon$ such that

$$d_\epsilon([\alpha]_\epsilon, [\beta_n]_\epsilon) \leq L(\tilde{\lambda}_n) < d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) + \frac{1}{n}.$$

The projected ϵ-chain, $\lambda_n = \pi(\tilde{\lambda}_n)$, is an ϵ-chain (with respect to d_q) from $[[\alpha]]_\epsilon$ to $[[\beta]]_\epsilon$, and it has the same length as $\tilde{\lambda}_n$. Thus, we have

$$d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) \leq L(\lambda_n) = L(\tilde{\lambda}_n) < d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) + \frac{1}{n}.$$

Since $d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon)$ is a lower bound for the lengths of all ϵ-chains between these two points, the previous inequality shows that it must also be the infimum of those lengths. That is,

$$d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) = \inf\{L(\lambda) \,:\, \lambda \text{ is an } \epsilon\text{-chain from } [[\alpha]]_\epsilon \text{ to } [[\beta]]_\epsilon\},$$

showing that $(X_\epsilon/\pi_\epsilon(X), d_q)$ is ϵ-intrinsic. This proves part 1.

If X is locally compact, complete, proper, and/or compact, then $X_\epsilon/\pi_\epsilon(X)$ has the same property/properties because $f : X_\epsilon/\pi_\epsilon(X) \to X$ is a local isometry and a homeomorphism. Finally, let $x = f([[\alpha]]_\epsilon)$ and $y = f([[\beta]]_\epsilon)$ be given. Since f is 1-Lipschitz, we already know that $d(x, y) \leq d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon)$. Suppose that X is ϵ-intrinsic, and, for each $n \in \mathbb{N}$, choose an ϵ-chain from x to y, λ_n, so that $d(x, y) \leq L(\lambda) < d(x, y) + 1/n$. Then $f^{-1}(\lambda)$ is an ϵ-chain from $[[\alpha]]_\epsilon$ to $[[\beta]]_\epsilon$, and it has the same length as λ. Thus,

$$d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) \leq L(f^{-1}(\lambda)) = L(\lambda) < d(x, y) + \frac{1}{n}.$$

Letting $n \to \infty$, we get $d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) \le d(x,y)$, showing that f is an isometry when X is ϵ-intrinsic. Conversely, if f is an isometry, then $X = f(X_\epsilon/\pi_\epsilon(X))$ will be ϵ-intrinsic because $X_\epsilon/\pi_\epsilon(X)$ is. ∎

Next, we will need a generalization of the previous two theorems for subgroups, G, of $\pi_\epsilon(X)$. Since $\pi_\epsilon(X)$ acts discretely on X_ϵ, any subgroup $G \subset \pi_\epsilon(X)$ will, also. Thus, the two previous proofs follow through with essentially no changes in this case. The only significant difference is that the map $f : X_\epsilon/G \to X$ - while still a surjective, uniform, local isometry - need not be injective. Since the same local properties hold for, f, though, it is a covering map.

Theorem 2.6.3 *Let X be chain-connected, and let $\epsilon > 0$ be given. Let G be a subgroup of $\pi_\epsilon(X)$, and define $\Theta_G : G \times X_\epsilon \to X_\epsilon$ by $\Theta_G([\gamma]_\epsilon, [\alpha]_\epsilon) := \Theta([\gamma]_\epsilon, [\alpha]_\epsilon) = [\gamma\alpha]_\epsilon$. Then Θ_G defines a discrete (therefore, free and properly discontinuous) action by isometries on X_ϵ, and d_ϵ is left-invariant with respect to this action. Moreover, the resulting quotient space, X_ϵ/G - metrized as before - and quotient map, $\pi_G : X_\epsilon \to X/G$, satisfy the following.*

1) *π_G is a covering map, a radial isometry on ϵ-balls, and an isometry on $\epsilon/2$-balls.*

2) *There is a 1-Lipschitz covering map, $f : X_\epsilon/G \to X$, that is a radial isometry on ϵ-balls, an isometry on $\epsilon/2$-balls, and satisfies $f \circ \pi_G = \varphi_\epsilon$.*

3) *X_ϵ/G is ϵ-intrinsic.*

4) *If X is locally compact and/or complete, then X_ϵ/G possesses the same property/properties.*

Proof The proof of Theorem 2.6.1 goes through almost without change. The only part of that result that doesn't apply directly to this case is part 2; it is not necessarily true that if α and β end at the same point then they are in the same orbit, since $[\gamma]_\epsilon = [\alpha\beta^{-1}]_\epsilon$ may not be in G. However, the analog that is needed in the subgroup case is the following: *if α and β end at the same point and $d_\epsilon([\alpha]_\epsilon, [\gamma]_\epsilon[\beta]_\epsilon) < \epsilon$ for some $[\gamma]_\epsilon \in G$, then $[\alpha]_\epsilon$ and $[\beta]_\epsilon$ are in the same orbit.* To see why this is true, suppose the hypotheses hold, and let x be the common endpoint of α and β. Then $L([\alpha^{-1}\gamma\beta]_\epsilon) < \epsilon$, implying that

$$\alpha^{-1}\gamma\beta \sim_\epsilon \{x\} \Rightarrow \gamma\beta \sim_\epsilon \alpha \Rightarrow [\alpha]_\epsilon = [\gamma\beta]_\epsilon.$$

Since $[\gamma]_\epsilon \in G$, $[\alpha]_\epsilon$ and $[\beta]_\epsilon$ are in the same orbit. With this result replacing part 2 in Theorem 2.6.1, that proof can be directly translated to this case.

The quotient space, X_ϵ/G, is metrized exactly as before, and, the proof of Theorem 2.6.2 goes through without change, with the exception that $f : X_\epsilon/G \to X$ need not be injective. ∎

Finally, we end this section with some comments regarding the regularity of the covering map, φ_ϵ. We have seen in this section that each group element, $[\gamma]_\epsilon$, determines an isometry on X_ϵ. Moreover, each of these isometries is a covering equivalence for φ_ϵ, since $\varphi_\epsilon([\gamma]_\epsilon[\alpha]_\epsilon) = \varphi_\epsilon([\alpha]_\epsilon)$ (i.e. the endpoint of $[\gamma\alpha]_\epsilon$ is the same as the endpoint of $[\alpha]_\epsilon$). We claim that the group of covering transformations of φ_ϵ is, in fact, $\pi_\epsilon(X)$. To see this, suppose we have a covering equivalence, $f : X_\epsilon \to X_\epsilon$, such that $\varphi_\epsilon \circ f = \varphi_\epsilon$. Let $[\gamma]_\epsilon = f(\tilde{*})$. By Lemma 2.4.9, there is a *unique* lift of φ_ϵ mapping $\tilde{*}$ to $[\gamma]_\epsilon$, so f must be this lift. On the other hand, the map, h_γ, also satisfies the lift conditions, since $h_\gamma(\tilde{*}) = [\gamma]_\epsilon[\{*\}]_\epsilon = [\gamma]_\epsilon$. Thus, we must have $f = h_\gamma$. Hence, the set of covering transformations of φ_ϵ is exactly the group of isometries of X_ϵ determined by $\pi_\epsilon(X)$. Therefore, the covering is regular.

2.7 Convergence Results in X_ϵ

Here, we will prove some convergence results regarding chains and equivalence classes of chains. Intuitively, this first lemma states that if an ϵ-chain, β, is sufficiently *uniformly* (which, in this case, also means pointwise!) close to a fixed ϵ-chain, α, with the same initial and terminal points, then β is, in fact, ϵ-homotopic to α.

Lemma 2.7.1 *Let $\alpha = \{x_0, \ldots, x_n\}$ be an ϵ-chain in X, and let*

$$g_\alpha := \min_{1 \le i \le n} \{\epsilon - d(x_{i-1}, x_i)\}.$$

If $\beta = \{x_0 = y_0, y_1, \ldots, y_n\}$ is an ϵ-chain such that $d(x_i, y_i) < \tau$ for each $i = 0, 1, \ldots, n$, where $0 < \tau \le g_\alpha$, then β is ϵ-homotopic to the ϵ-chain $\{x_0, x_1, \ldots, x_{n-1}, y_n\}$. If, in addition, $y_n = x_n$, then $\beta \sim_\epsilon \alpha$.

Proof First, since $x_0 = y_0$ and $d(x_1, y_1) < \tau < \epsilon$, we can insert x_1 into β between y_0 and y_1. But we also have

$$d(x_1, y_2) \le d(x_1, x_2) + d(x_2, y_2) < d(x_1, x_2) + \tau \le d(x_1, x_2) + \epsilon - d(x_1, x_2) = \epsilon,$$

so we can then remove y_1 to obtain the chain $\beta^{(1)} := \{x_0 = y_0, x_1, y_2, \ldots, y_n\}$.

Continuing this inductively, we can successively insert the points, x_i, and then remove the corresponding point, y_i, until we obtain the chain $\beta^{(n-1)} = \{x_0 = y_0, x_1, \ldots, x_{n-2}, x_{n-1}, y_n\}$. This proves the first part of the lemma. But if $y_n = x_n$, then $\beta^{(n-1)} = \{x_0 = y_0, \ldots, x_n\}$, which proves the last part. ∎

We call the value g_α in the previous lemma the *gap differential of* α.

The following lemma can be thought of as an analog of Arzela-Ascoli for ϵ-chains.

Lemma 2.7.2 *Let X be a proper metric space, and let $\{\alpha_n\}$ be a sequence of ϵ-chains in X having the same number of points - say, $m + 1$ - and such that the sequence of initial points, $\{x_0^n\}$, is bounded. Then there is a subsequence, $\{\alpha_{n_k}\}$, and a chain, $\beta = \{x_0, \ldots, x_m\}$, satisfying $d(x_{i-1}, x_i) \le \epsilon$ for each $i = 1, \ldots, m$, such that α_{n_k} converges pointwise to β. That is, if $\alpha_{n_k} = \{x_0^{n_k}, \ldots, x_m^{n_k}\}$, then $x_i^{n_k} \to x_i$ as $k \to \infty$ for each $i = 0, 1, \ldots, m$.*

Proof Let $\alpha_n = \{x_0^n, x_1^n, \ldots, x_m^n\}$. We first claim that the chains, α_n, all lie in some closed ball in X, which - since X is proper - is compact. Since the sequence of initial points, $\{x_0^n\}$, is bounded, these points all lie in some ball, $B(x, R)$. Now, let $n \ge 1$ be given. For any $x_i^n \in \alpha_n$,

$$\begin{aligned} d(x, x_i^n) &\le d(x, x_0^n) + d(x_0^n, x_1^n) + \cdots + d(x_{i-1}^n, x_i^n) \\ &\le d(x, x_0^n) + L(\alpha_n) \\ &\le R + m\epsilon. \end{aligned}$$

Thus, every point of α_n, for any n, lies in the closed ball, $C(x, R + m\epsilon)$.

So, we have $m+1$ sequences, $\{x_i^n\}_{n=1}^\infty$, $0 \le i \le m$, all lying in a compact set in X. This means that each of them will have some convergent subsequence. By taking *successive* subsequences, we can obtain the desired result. That is, we first take a convergent subsequence of $\{x_0^n\}_{n \ge 1}$, say $\{x_0^{n(k_0)}\}_{k_0 \ge 1}$ with $x_0^{n(k_0)} \to x_0$ as $k_0 \to \infty$. Then we consider the corresponding subsequence of

$\{x_1^n\}$, $\{x_1^{n(k_0)}\}_{k_0 \geq 1}$, and we take a convergent subsequence of this sequence, say $\{x_1^{n(k_1)}\}_{k_1 \geq 1}$ with $x_1^{n(k_1)} \to x_1$ as $k_1 \to \infty$. Then the corresponding subsequence, $\{x_0^{n(k_1)}\}_{k_1 \geq 1}$ will still converge to x_0. Likewise, we consider the corresponding subsequence of $\{x_2^n\}$, $\{x_2^{n(k_1)}\}_{k_1 \geq 1}$, and we take a convergent subsequence of this sequence, say $\{x_2^{n(k_2)}\}_{k_2 \geq 1}$ with $x_2^{n(k_2)} \to x_2$ as $k_2 \to \infty$. Then the corresponding subsequences, $\{x_0^{n(k_2)}\}_{k_2 \geq 1}$ and $\{x_1^{n(k_2)}\}_{k_2 \geq 1}$, will still converge to x_0 and x_1, respectively.

We continue this process, and - since there are only $m+1$ sequences - it must stop at some point, giving us a collection of $m+1$ convergent sequences. By relabeling the index set after obtaining the last convergent subsequence, we can simply denote these by $\{x_i^{n_k}\}_{k=1}^{\infty}$, with $x_i^{n_k} \to x_i$ for each $i = 0, 1, \ldots, m$. Moreover, since we took successive subsequences as described above, the points $x_i^{n_k}$, for fixed k and $0 \leq i \leq m$, all belong to the same chain, α_{n_k}. Hence, we have a subsequence of ϵ-chains, $\alpha_{n_k} = \{x_0^{n_k}, x_1^{n_k}, \ldots, x_m^{n_k}\}$, and points x_i, $0 \leq i \leq m$, such that $x_i^{n_k} \to x_i$ as $k \to \infty$. Finally, given $1 \leq i \leq m$, since $d(x_{i-1}^{n_k}, x_i^{n_k}) < \epsilon$ for all k, we have

$$d(x_{i-1}^{n_k}, x_i^{n_k}) \to d(x_{i-1}, x_i) \Rightarrow d(x_{i-1}, x_i) \leq \epsilon. \quad \blacksquare$$

Typically, this result is slightly more general than we need. As one might imagine, in most cases, the chains will all be anchored at our fixed base point, $*$, so the condition that the sequence of initial points be bounded is unnecessary in that case. However, the requirement that X be proper cannot be eliminated. Otherwise, there is no way to guarantee that any of the sequences $\{x_i^n\}$ has a convergent subsequence. In addition, we see in this result another problem that may arise when talking about convergence of chains: the ϵ-chain property may not be preserved upon passing to the limit. We will have to address this issue later on.

The following counting lemma will be very useful. Its proof is due to Conrad Plaut and originally appeared in [9], though we reproduce it here. For this lemma, we adopt the following notation: if $\alpha = \{x_0, \ldots, x_n\}$ is an ϵ-chain in X, then we define $\nu(\alpha) := n+1$, the cardinality of α. We have been using Greek letters to indicate chains, but we will reserve the letter, ν, for this specific use.

Lemma 2.7.3 *Let L, $\epsilon > 0$ be given, and let α be an ϵ-chain in X with $L(\alpha) \leq L$. Then there is some $\alpha' \in [\alpha]_\epsilon$ such that $L(\alpha') \leq L(\alpha)$ and $\nu(\alpha') \leq \left\lfloor \frac{2L}{\epsilon} + 1 \right\rfloor$.*

In other words, given any equivalence class of ϵ-chains, we can find always a representative from the class with cardinality bounded solely in terms of ϵ and the length of the class.

Proof Let $\alpha = \{x_0, \ldots, x_n\}$. Suppose that for some i, we have $d(x_{i-1}, x_i) + d(x_i, x_{i+1}) < \epsilon$. Then $d(x_{i-1}, x_{i+1}) < \epsilon$, and we can remove x_i to form the ϵ-chain $\alpha^{(1)} = \{x_0, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n\}$. Moreover, we have $L(\alpha') \leq L(\alpha)$ by the triangle inequality. We can continue this process for every three-point subchain, $\{x_{k-1}, x_k, x_{k+1}\}$, satisfying $d(x_{k-1}, x_k) + d(x_k, x_{k+1}) < \epsilon$. After performing all of these steps (there can obviously be only finitely many such steps), we have obtained an ϵ-chain, α', with the property that $d(x_{i-1}, x_i) + d(x_i, x_{i+1}) \geq \epsilon$ for all i and with length bounded above by $L(\alpha)$. By pairing off terms, this last inequality implies that

$$L(\alpha') \geq \left\lfloor \frac{\nu(\alpha')}{2} \right\rfloor \epsilon,$$

from which it follows that

$$\nu(\alpha') \leq \left\lfloor \frac{2L(\alpha')}{\epsilon} + 1 \right\rfloor \leq \left\lfloor \frac{2L}{\epsilon} + 1 \right\rfloor. \quad \blacksquare$$

Corollary 2.7.4 *Given any ϵ-homotopy equivalence class of ϵ-chains, $[\alpha]_\epsilon$, let N be the unique positive integer such that*

$$\frac{(N-1)\epsilon}{2} \leq L([\alpha]_\epsilon) < \frac{N\epsilon}{2}.$$

Then there is a chain, $\alpha' \in [\alpha]_\epsilon$, such that $\nu(\alpha') \leq N+1$.

Proof By definition of $L([\alpha]_\epsilon)$, there is some ϵ-chain, $\bar{\alpha} \in [\alpha]_\epsilon$, such that $L(\bar{\alpha}) < N\epsilon/2$. By the previous lemma, there is, then, an ϵ-chain, $\alpha' \in [\alpha]_\epsilon$, such that

$$\nu(\alpha') \leq \frac{2}{\epsilon}\frac{N\epsilon}{2} + 1 = N+1. \quad \blacksquare$$

Corollary 2.7.5 *If $\{[\alpha_n]_\epsilon\}$ is a bounded sequence of elements in X_ϵ, then there is a representative, α_n', from each $[\alpha_n]_\epsilon$ such that $\nu(\alpha_n')$ is the same for all n.*

Proof Choose L such that $L([\alpha_n]_\epsilon) = d_\epsilon(\tilde{*}, [\alpha_n]_\epsilon) < L$, and then, for each n, choose a representative, $\alpha_n \in [\alpha_n]_\epsilon$, such that $L(\alpha_n) < L$. By the previous lemma, we can find, for each n, an ϵ-chain, $\alpha_n' \in [\alpha_n]_\epsilon$ such that

$$\nu(\alpha_n') \leq \left\lfloor \frac{2L}{\epsilon} + 1 \right\rfloor.$$

For any α_n' with fewer than $\lfloor 2L/\epsilon + 1 \rfloor$ points, we can simply repeat the initial point of α_n' enough times to give each α_n' a cardinality of exactly $\lfloor 2L/\epsilon + 1 \rfloor$, and doing so does not change the length or homotopy class of the chain. $\blacksquare$

Chapter 3

The Critical and Covering Spectra

In this section, we will compare the entourage covers, X_ϵ, to the Spanier covers, X^δ, used by Sormani and Wei for geodesic spaces. As a consequence, we will show that - in the case when X is a geodesic space - the two spectra differ only by a multiplicative constant. In other words, in the geodesic case, the covering and critical spectra yield the same information.

3.1 Path Lifting

To relate the critical spectrum of a geodesic space to the covering spectrum, we need some means of transferring from homotopies between chains to homotopies between paths. These first few technical results will help facilitate this. Since we will be working with both chains and continuous paths, to avoid confusion, we will use the word "path" to denote a continuous curve, $\gamma : [a, b] \to X$. When γ is a closed curve, we will use the phrase "path loop" to distinguish this from a chain-loop.

Definition 3.1.1 *Let $\gamma : [a, b] \to X$ be a path in a metric space, X, and let $\epsilon > 0$ be given. An **ϵ-chain along** γ is an ϵ-chain, $\{x_0, x_1, \ldots, x_n\}$ with the property that there exists a partition $\{t_0 = a, t_1, \ldots, t_n = b\}$ of $[a, b]$ such that $\gamma(t_i) = x_i$ for each $i = 0, 1, \ldots, n$. In particular, $x_0 = \gamma(a)$ and $x_n = \gamma(b)$. A **strong ϵ-chain along** γ is an ϵ-chain along γ with the additional property that $\gamma([t_{i-1}, t_i]) \subset B(\gamma(t_{i-1}), \epsilon)$ for each $i = 1, \ldots, n$. An **ultra ϵ-chain along** γ is an ϵ-chain along γ with the additional property that $\gamma([t_{i-1}, t_i]) \subset B(\gamma(t_{i-1}), \epsilon) \cap B(\gamma(t_i), \epsilon)$.*

Thus, for a strong ϵ-chain along γ, each path subsegment, $\gamma([t_{i-1}, t_i])$, lies in the ϵ-ball centered at the initial point of that subsegment. In the ultra case, this subsegment also lies in the ϵ-ball centered at the terminal point of the subsegment. Note that an ϵ-chain along γ need not be a strong ϵ-chain along γ, and a strong ϵ-chain along γ need not be an ultra ϵ-chain along γ.

Typically, ϵ-chains along paths are weaker than we need. Consequently, we will work mostly with strong and ultra ϵ-chains along paths.

Lemma 3.1.2 *Given any path $\gamma : [a, b] \to X$ in the metric space, X, and any $\epsilon > 0$, there exists a strong ϵ-chain along γ.*

Proof Consider the covering of the compact set, $\gamma([a, b])$, consisting of the $\epsilon/2$-balls, $B(\gamma(t), \epsilon/2)$, for all $t \in [a, b]$. The inverse images of these balls under γ form an open covering of $[a, b]$. Let δ be the Lebesgue number for this covering. Let $\{t_0 = a, \ldots, t_n = b\}$ be a partition of $[a, b]$ such that each subinterval $[t_i, t_{i+1}]$ has length less than δ. Given any i, $0 \le i \le n-1$, consider the point $\gamma(t)$ for any $t \in [t_i, t_{i+1}]$. Since this interval has diameter less than δ, there is some t^* such that $\gamma([t_i, t_{i+1}]) \subset B(\gamma(t^*), \epsilon/2)$. Thus, $d(\gamma(t^*), \gamma(t_i)) < \epsilon/2$ and $d(\gamma(t^*), \gamma(t)) < \epsilon/2$, implying that $d(\gamma(t_i), \gamma(t)) < \epsilon$. It follows that $\gamma(t) \in B(\gamma(t_i), \epsilon)$. ∎

Clearly, the previous result can be generalized to show that there is always an ultra ϵ-chain along a path. In fact, a strong $(\epsilon/2)$-chain along a path is an ultra ϵ-chain along that same path.

There are two obvious, but useful, results concerning strong ϵ-chains along paths that we will use often without mentioning them. Let $\gamma : [0,1] \to X$ and $\lambda : [0,1] \to X$ be two paths, with $\gamma(1) = \lambda(0)$. Let α and β be strong (resp. ultra) ϵ-chains along γ and λ, respectively. Then the concatenation, $\alpha\beta$, is a strong (resp. ultra) ϵ-chain along $\gamma\lambda$. Moreover, if $\alpha = \{x_0, \ldots, x_n\}$, then, for any $i = 1, \ldots, n$, the truncation of α to $\{x_0, \ldots, x_i\}$ is a strong (resp. ultra) ϵ-chain along the path $\gamma|_{[0,t_i]}$.

Remark Many of the proofs in this section are necessarily technical, since they require the manipulation of chains along paths, which, in turn, means inserting and removing points in some sort of ordered process. So, to simplify the presentation of these proofs, we will adopt some intuitive notation and terminology. First, by reparameterizing, we may obviously assume that all paths are defined on $[0,1]$. Second, note that an ϵ-chain along γ depends on the mapping $\gamma : [0,1] \to X$ and the particular partition, and not just on the image of this map. ϵ-chains are ordered structures by definition, but the parameterization of a path $\gamma : [0,1] \to X$ induces an order on $\gamma([0,1])$, namely $\gamma(t_1) \le \gamma(t_2)$ if and only if $t_1 \le t_2$. We will use this order liberally in the following, so some clarifying statements are in order. Even though the map $\gamma : [0,1] \to X$ may not be injective, we will identify points on $\gamma([0,1])$ with their corresponding parameter values. So, even if $x = \gamma(t_1) = \gamma(t_2) = y$ for two distinct values t_1 and t_2, we will still say that $x \le y$ - or x precedes y, or y follows x - if $t_1 \le t_2$ and $x < y$ if $t_1 < t_2$. Considered only as points of X, one may "reach" $\gamma(t_2)$ before $\gamma(t_1)$ as one traverses the path γ (if, for example, γ intersects itself), but since $t_1 \le t_2$, we will still say that $\gamma(t_1)$ precedes $\gamma(t_2)$. Consequently, we will accept some minor abuse of notation and terminology with regard to a chain, $\{x_0, \ldots, x_n\}$, along a path, γ. In particular, when we refer to a point, x_i, in the chain, we will really be referencing both the point, x_i, on the path and the value t_i such that $x_i = \gamma(t_i)$. More importantly, we will use $[x_i, x_{i+1}]$ to denote the path segment $\gamma([t_i, t_{i+1}])$ or the map $\gamma|_{[t_i,t_{i+1}]}$, and we will even refer to $[x_i, x_{i+1}]$ as *the interval* or *segment along* γ *from* x_i *to* x_{i+1}. One result, in particular, will be useful. Given an ϵ-chain, $\{x_0, \ldots, x_n\}$, along γ and any point $y = \gamma(t)$, there is a *unique* interval $[x_i, x_{i+1})$ such that $y \in [x_i, x_{i+1})$. Without reference to the parameterization, γ, this is not true, since the sets $\gamma([t_i, t_{i+1}))$, $0 \le i \le n-1$ need not disjointly partition the set $\gamma([0,1))$. But, when the parameterization is taken into account, there is some unique i such that $t_i \le t < t_{i+1}$.

The following lemma, when paired with the previous one, will give us a well-defined way of associating to any path an equivalence class of ϵ-chains.

Lemma 3.1.3 *Let X be a metric space, and let $\epsilon > 0$ be given. Let $\gamma : [0,1] \to X$ be any path in X. Then any two strong ϵ-chains along γ are ϵ-homotopic. Consequently, all strong and ultra ϵ-chains along γ are in the same ϵ-equivalence class.*

It is crucial to emphasize that we do not require the ϵ-homotopy in this lemma to preserve the strong ϵ-chain property at each step. This result simply says that if we have two strong ϵ-chains along a path γ, then they are ϵ-homotopic in the usual sense. Also, it should be noted that this lemma is not true for regular ϵ-chains along γ. The "strong" condition is essential.

Proof Let $* = \gamma(0)$ be the initial point of γ. Fix a strong ϵ-chain along γ, say $\lambda = \{* = x_0, x_1 = \gamma(t_1), \ldots, x_{n-1} = \gamma(t_{n-1}), x_n = \gamma(1)\}$. Let $\rho = \{* = y_0, y_1 = \gamma(s_1), \ldots, y_{m-1} = \gamma(s_{m-1}), y_m = \gamma(1)\}$ be any other strong ϵ-chain along γ. We will construct an ϵ-homotopy transforming λ to ρ. The proof is tedious, but not difficult.

The goal is to show that we can replace an initial segment of λ with an initial segment of ρ. Since λ and ρ agree at their endpoints, we only need to insert the points $y_1, \ldots, y_{m-1}$ into λ and remove the points $x_1, \ldots, x_{n-1}$. There is some interval, $[x_i, x_{i+1}]$, such that $y_1 \in [x_i, x_{i+1})$. Now, y_2 may also lie inside $[x_i, x_{i+1})$, along with other successive points of ρ, so let y_k be the first point in ρ such that $y_k \geq x_{i+1}$. We claim that we can insert the segment $\{y_1, \ldots, y_{k-1}\}$ in between x_i and x_{i+1}. We do this in a backwards fashion. Since y_{k-1} lies in $[x_i, x_{i+1}] \subset B(x_i, \epsilon)$, we have

$$d(y_{k-1}, x_i) < \epsilon. \tag{3.1}$$

Since y_k follows x_{i+1} and y_{k-1} precedes x_{i+1}, we have $x_{i+1} \in [y_{k-1}, y_k] \subset B(y_{k-1}, \epsilon)$, from which it follows that

$$d(y_{k-1}, x_{i+1}) < \epsilon. \tag{3.2}$$

Combining 3.1 and 3.2, we conclude that we can insert y_{k-1} into λ between x_i and x_{i+1}. The rest of the steps now follow similarly. We know that $d(y_{k-2}, y_{k-1}) < \epsilon$. Moreover, since y_{k-2} follows x_i and precedes x_{i+1}, it lies in $[x_i, x_{i+1}]$, implying that $d(x_i, y_{k-2}) < \epsilon$. Hence, we can insert y_{k-2} into the new λ (with y_{k-1} already added) between x_i and y_{k-1}. Continuing in this way, we can insert each of the points $y_1, \ldots, y_{k-1}$ successively so that the we have now transformed λ via an ϵ-homotopy to the ϵ-chain

$$\{* = x_0, x_1, \ldots, x_i, y_1, \ldots, y_{k-1}, x_{i+1}, \ldots, x_n = \gamma(1)\}.$$

Now, if $i = 0$, then we are done with this step. We have inserted an initial segment of ρ into λ and obtained the chain $\{x_0 = \gamma(0), y_1, \ldots, y_{k-1}, x_1, \ldots, x_n = \gamma(1)\}$. If $i \geq 1$, we can successively remove the points $x_1, \ldots, x_i$ as follows. The interval $[x_0, y_1]$ is the initial segment of γ determined by ρ, and since $y_1 \geq x_i$, the points $x_1, \ldots, x_i$ all lie in this segment, which is contained in $B(\gamma(0), \epsilon) = B(x_0, \epsilon)$. Thus, we can remove $x_1, \ldots, x_i$ successively, starting with x_1 and going up to x_i. This leaves us with the chain $\lambda^{(1)}$ given by

$$\{x_0 = \gamma(0), y_1, \ldots, y_{k-1}, x_{i+1}, \ldots, x_n = \gamma(1)\}.$$

Hence, we have inserted an initial segment of ρ into λ, removed the initial segment of λ that preceded the inserted one, and, thus, we have transformed λ to a chain that agrees with ρ up through the first k points. Now we apply exactly the same procedure, starting with the point y_k. This point lies in some interval $[x_j, x_{j+1})$, with $j \geq i + 1$. So, just as before, we can insert some segment $y_k, y_{k+1}, \ldots, y_l$ into $\lambda^{(1)}$ between x_j and x_{j+1}, and we can remove the points $x_{i+1}, \ldots, x_j$, also as before. This provides an algorithm, each step of which follows via an ϵ-homotopy, for successively transforming segments of λ into segments of ρ.

Finally, at each step of this process, we insert at least one point of ρ into λ, or rather into the chain we have obtained from λ via previous steps of this homotopy. Since ρ consists of only finitely many points, this process must stop at some point, namely when we insert y_{m-1} and remove any preceding points of λ. At this point, there may still be points of λ in the interval $[y_{m-1}, y_m] = [y_{m-1}, x_n]$, say the points $x_r, \ldots, x_{n-1}, x_n$. Those can be removed, starting with x_r and going up to x_{n-1}, because they all lie in the ball $B(y_{m-1}, \epsilon)$. When this step is completed, the result is an ϵ-homotopy transforming λ into ρ. ∎

Because of the previous two lemmas, given any path, $\gamma : [0,1] \to X$, and any $\epsilon > 0$, we can speak of *the* ϵ-equivalence class of strong ϵ-chains along γ. Such a chain always exists by lemma 3.1.2, and - by lemma 3.1.3 - any two such chains are ϵ-homotopic and, thus, are in the same ϵ-equivalence class.

Now, let $\gamma : [0,1] \to X$ be any path in X beginning at $*$. Then, there is a unique continuous lift of γ to $\tilde{*} \in X_\epsilon$. We can precisely characterize this lifted path. Define $\hat{\gamma} : [0,1] \to X_\epsilon$ by letting $\hat{\gamma}(0) = \tilde{*}$ and letting $\hat{\gamma}(t)$ - for each $t \in (0,1]$ - be the equivalence class of strong ϵ-chains along $\gamma|_{[0,t]}$. This is where the necessity of lemmas 3.1.2 and 3.1.3 becomes evident; without either of them, $\hat{\gamma}$ would not be well-defined. Since $\hat{\gamma}(t)$ is an equivalence class of ϵ-chains starting at $*$ and ending at $\gamma(t)$, it is clear that we have $\varphi_\epsilon \circ \hat{\gamma} = \gamma$ as functions. We would like to know that $\hat{\gamma}$ is the unique path lifting of γ to X_ϵ beginning at $\tilde{*}$, but we still need continuity of $\hat{\gamma}$.

Lemma 3.1.4 *Let X be a chain-connected metric space, and let $\epsilon > 0$ be given. Let $\gamma : [0,1] \to X$ be a path beginning at $*$. The function $\hat{\gamma} : [0,1] \to X_\epsilon$ defined above is continuous, and is, therefore, the unique lift of γ to X_ϵ beginning at $\tilde{*}$.*

Proof Let $\tilde{\gamma}$ be the unique continuous lift of γ to X_ϵ starting at $\tilde{*}$. We will actually show that $\hat{\gamma}$ is continuous by showing directly that it equals $\tilde{\gamma}$. Let $\{* = x_0, x_1, \ldots, x_n\}$ be a strong ϵ-chain along γ, with $x_i = \gamma(t_i)$, $i = 0, 1, \ldots, n$. By definition, we have $\tilde{\gamma}(0) = \hat{\gamma}(0) = \tilde{*}$. Now let t be any point in $(0, t_1]$. We know that φ_ϵ is a homeomorphism from the ball $B(\tilde{*}, \epsilon)$ onto the ball $B(x_0, \epsilon) = B(*, \epsilon)$, and since $\gamma([0, t_1])$ lies in $B(x_0, \epsilon)$, $\tilde{\gamma}([0, t_1])$ lies in $B(\tilde{*}, \epsilon)$. Since we also know that $\varphi_\epsilon \circ \tilde{\gamma} = \gamma = \varphi_\epsilon \circ \hat{\gamma}$, to show that $\tilde{\gamma}(t) = \hat{\gamma}(t)$, we need only show that $\hat{\gamma}(t) \in B(\tilde{*}, \epsilon)$. Now, $\hat{\gamma}(t)$ is the equivalence class of strong ϵ-chains along $\gamma|_{[0,t]}$, and since $\gamma([0, t_1])$ lies in $B(*, \epsilon)$, one such strong ϵ-chain is $\{\gamma(0), \gamma(t)\}$. Thus, $\hat{\gamma}(t) = [\{\gamma(0), \gamma(t)\}]_\epsilon$ by the preceding lemma. Clearly we have

$$d_\epsilon([\{\gamma(0), \gamma(t)\}]_\epsilon, [\{\gamma(0), \gamma(0)\}]_\epsilon) < \epsilon,$$

It follows that $[\{\gamma(0), \gamma(t)\}]_\epsilon \in= B(\tilde{*}, \epsilon)$. So, $\hat{\gamma}(t) \in B(\tilde{*}, \epsilon)$, implying that $\hat{\gamma}(t) = \tilde{\gamma}(t)$. This shows that $\hat{\gamma}$ and $\tilde{\gamma}$ agree on $[0, t_1]$.

Proceeding inductively, suppose that $\hat{\gamma}$ and $\tilde{\gamma}$ agree on $[0, t_k]$ for some k, $1 \leq k \leq n-1$. Let t be any point in $(t_k, t_{k+1}]$. Just as before, we have $\gamma([t_k, t_{k+1}]) \subset B(x_k, \epsilon)$, and, therefore, $\tilde{\gamma}([t_k, t_{k+1}])$ lies in $B(\tilde{\gamma}(t_k), \epsilon)$, since φ_ϵ is a homeomorphism of $B(\tilde{\gamma}(t_k), \epsilon)$ onto $B(x_k, \epsilon)$. So, again, we need only show that $\hat{\gamma}(t) \in B(\tilde{\gamma}(t_k), \epsilon)$. The chain $\{x_0, x_1, \ldots, x_k, \gamma(t)\}$ is a strong ϵ-chain along $\gamma|_{[0,t]}$. Therefore, $\hat{\gamma}(t) = [\{x_0, x_1, \ldots, x_k, \gamma(t)\}]_\epsilon$. Moreover,

$$d_\epsilon([\{x_0, \ldots, x_k, x_k\}]_\epsilon, [\{x_0, \ldots, x_k, \gamma(t)\}]_\epsilon) = d(x_k, \gamma(t)) < \epsilon,$$

so $[\{x_0, \ldots, x_k, \gamma(t)\}]_\epsilon$ is in the ϵ-ball centered at $[\{x_0, \ldots, x_k, x_k\}]_\epsilon$. But $\{x_0, \ldots, x_k, x_k\}$ is ϵ-homotopic to $\{x_0, \ldots, x_k\}$, and $[\{x_0, \ldots, x_k\}]_\epsilon = \hat{\gamma}(t_k) = \tilde{\gamma}(t_k)$ by the definition of $\hat{\gamma}$ and the inductive hypothesis. So,

$$\hat{\gamma}(t) = [\{x_0, \ldots, x_k, \gamma(t)\}]_\epsilon \in B(\tilde{\gamma}(t_k), \epsilon).$$

This shows that $\hat{\gamma}$ and $\tilde{\gamma}$ agree on $[0, t_{k+1}]$. By induction, they must agree on all of $[0,1]$, thus showing that $\hat{\gamma}$ is the unique lift of γ beginning at $\tilde{*}$. ∎

One of the most important consequences of the preceding results is the following theorem, which gives us an explicit connection between chain loops and continuous loops.

Theorem 3.1.5 *Let X be a chain-connected metric space, and let $\epsilon > 0$ be given. Let $\gamma : [0,1] \to X$ be a path loop at $* \in X$. Then γ lifts to a path loop at $\tilde{*}$ if and only if there is a strong ϵ-chain along γ that is ϵ-nullhomotopic. Consequently, γ lifts to a closed loop at $\tilde{*}$ if and only if* every *strong ϵ-chain along γ is ϵ-nullhomotopic.*

Proof Let γ be a loop at $*$, and let $\tilde{\gamma}$ be its lift to X_ϵ at $\tilde{*}$. Suppose $\tilde{\gamma}$ is closed. Let $\{* = x_0, x_1, \dots, x_n = *\}$ be a strong ϵ-chain (or strong ϵ-loop, in this case) along γ. Then by the previous lemma, $\tilde{\gamma}(1) = [\{* = x_0, \dots, x_n = *\}]_\epsilon$. Since γ lifts to a closed loop at $\tilde{*}$, we have $\tilde{\gamma}(1) = \tilde{\gamma}(0) = \tilde{*}$, from which it follows that $[\{* = x_0, \dots, x_n = *\}]_\epsilon = [\{*\}]_\epsilon$, and the strong ϵ-chain $\{* = x_0, \dots, x_n = *\}$ is ϵ-null-homotopic.

Conversely, suppose there is a strong ϵ-chain along γ that is ϵ-nullhomotopic, say $\{* = x_0, x_1, \dots, x_n = *\}$. Then $[\{x_0, \dots, x_n\}]_\epsilon = [\{*\}]_\epsilon$. But $\tilde{\gamma}(0) = [\{*\}]_\epsilon = [\{x_0, \dots, x_n\}]_\epsilon = \tilde{\gamma}(1)$. Hence, $\tilde{\gamma}$ is closed. ∎

Along with some technical details, Theorem 3.1.5 will be the primary tool used in deriving the relationship between the Spanier covers and entourage-covers of geodesic spaces, and this will lead us to the comparison of the critical spectrum with the covering spectrum. Before moving on to that comparison, however, we can derive some interesting applications of this connection between chain and path loops.

Lemma 3.1.6 *Let $\gamma : [0,1] \to X$ and $\lambda : [0,1] \to X$ be paths in a metric space, X, that are fixed endpoint path homotopic, and let $\epsilon > 0$. Then any strong ϵ-chain along γ is ϵ-homotopic to any strong ϵ-chain along λ.*

Proof Let $*$ be the initial point of γ and λ, and let X_ϵ be the ϵ-cover of X determined by $*$. By our hypothesis, the path loop $\gamma\lambda^{-1}$ is null-homotopic at $*$. Thus, $\gamma\lambda^{-1}$ lifts closed to $\tilde{*} \in X_\epsilon$. Let α be a strong ϵ-chain along γ, and let β be an *ultra* ϵ-chain along λ. Then the inverse of β is a strong ϵ-chain along λ^{-1}. This implies that $\alpha\beta^{-1}$ is a strong ϵ-chain along $\gamma\lambda^{-1}$. Since this path loop lifts closed to X_ϵ, Theorem 3.1.5 implies that the ϵ-loop, $\alpha\beta^{-1}$ is ϵ-null, or, equivalently, $\alpha \sim_\epsilon \beta$. Since an ultra ϵ-chain is also a strong ϵ-chain, any strong ϵ-chain along λ will be ϵ-homotopic to β - and, therefore, α - by Lemma 3.1.3. ∎

Lemma 3.1.7 *Let X be a chain-connected metric space, and let $\epsilon > 0$ be given. Then there exists a homomorphism $h : \pi_1(X) \to \pi_\epsilon(X)$. If the ϵ-balls of X are path-connected (e.g. X is geodesic), then h is surjective.*

Proof Obviously, we take $\pi_1(X)$ to be the fundamental group based at $*$, the same base point that induces $\pi_\epsilon(X)$. Given a path loop, γ, at $*$, we let $\tilde{\gamma} \in \pi_1(X)$ denote its fixed endpoint path homotopy equivalence class. We define $h : \pi_1(X) \to \pi_\epsilon(X)$ as follows. Let $h(\tilde{\gamma})$ be the equivalence class of strong ϵ-loops along γ. The previous lemma shows that h is well-defined. To see that it is a homomorphism, let $\tilde{\gamma}, \tilde{\lambda} \in \pi_1(X)$ be given. Let α and β be strong ϵ-loops along γ and λ, respectively. Then $\alpha\beta$ is a strong ϵ-loop along $\gamma\lambda$, and

$$h(\tilde{\gamma}\tilde{\lambda}) = h(\widetilde{\gamma\lambda}) = [\alpha\beta]_\epsilon = [\alpha]_\epsilon[\beta]_\epsilon = h(\tilde{\gamma})h(\tilde{\lambda}).$$

Thus, h is a homomorphism.

Lastly, suppose the ϵ-balls in X are path-connected. Let $[\alpha]_\epsilon \in \pi_\epsilon(X)$ be given, and let $\alpha = \{* = x_0, x_1, \dots, x_n = *\}$. For each $i = 0, 1, \dots, n-1$, the ball $B(x_i, \epsilon)$ is path-connected and contains x_{i+1}. So, for each i, let γ_i be a path in $B(x_i, \epsilon)$ from x_i to x_{i+1}. Then the path, γ, formed by concatenating the paths, γ_i, is a loop at $*$, and - by construction - α is a strong ϵ-chain along γ. Thus, $h(\tilde{\gamma}) = [\alpha]_\epsilon$, and h is surjective. ∎

We can also use these lifting properties to show that the critical spectrum of a compact geodesic space is discrete and consists of only a single type of critical value.

Lemma 3.1.8 *Let X be a geodesic space. If γ is a path loop at $*$ along which there is an ultra ϵ-chain that is ϵ-null (i.e. γ lifts closed to X_ϵ), then γ is path-homotopic to a finite product of path loops of the form $\alpha\beta\alpha^{-1}$, where α is a path from $*$ to some point, x, and β is a loop at x lying in an open ball of radius $3\epsilon/2$.*

Proof The proof is by induction on the number of steps in the nullhomotopy of the ultra ϵ-chain along γ. For the base step, suppose γ is a path loop at $*$ along which there is an ultra ϵ-chain that is ϵ-nullhomotopic via a 1-step homotopy. Then this one step must be removal of a point to obtain the constant chain $\{*\}$. This means that we have an ϵ-chain $\{* = \gamma(0), \gamma(t_1), \gamma(1) = *\}$ such that $[\gamma(0), \gamma(t_1)], [\gamma(t_1), \gamma(1)] \subset B(*, \epsilon)$. In other words, the entire path, γ, lies in $B(*, \epsilon)$, and γ is of the desired form.

Now, assume, for some $n \geq 1$, that if γ is any loop at $*$ along which there is an ultra ϵ-chain that is ϵ-nullhomotopic via a homotopy of n or fewer steps, then γ is homotopic to a product of loops of the form $\alpha\beta\alpha^{-1}$ as described above. Let γ be a path loop at $*$ such that there is an ultra ϵ-chain along γ that is ϵ-nullhomotopic via a homotopy of $n+1$ steps. We will assume that γ is parameterized on $[0,1]$, and we denote this chain by

$$\lambda = \{* = x_0 = \gamma(t_0), x_1 = \gamma(t_1), \ldots, x_{m-1} = \gamma(t_{m-1}), x_m = \gamma(t_m) = *\},$$

where $0 = t_0 < t_1 < \cdots < t_m = 1$. Then, for each $i = 1, \ldots, m$, we have $[x_{i-1}, x_i] \subset B(x_{i-1}, \epsilon) \cap B(x_i, \epsilon)$. There are two cases to consider, depending on whether the first step of the ϵ-nullhomotopy of λ adds or removes a point.

Suppose the first step of the homotopy removes x_i from λ. Choose a minimizing geodesic, η, from x_{i+1} to x_{i-1}. Since we can remove x_i, we know that $d(x_{i-1}, x_{i+1}) < \epsilon$, implying that $l(\eta) < \epsilon$. Let $\alpha := \gamma|_{[0,t_{i-1}]}$, $\beta := \gamma|_{[t_{i-1},t_{i+1}]}\eta$, and $\gamma' := \alpha\eta^{-1}\gamma|_{[t_{i+1},1]}$. Then γ is path-homotopic to $\alpha\beta\alpha^{-1}\gamma'$. We will show that the loop β lies in $B(x_i, 3\epsilon/2)$. If x lies on the segment $\gamma|_{[t_{i-1},t_{i+1}]}$, then the ultra ϵ-chain property implies that $x \in B(x_i, 3\epsilon/2)$. So, suppose $x = \eta(t)$ lies on η. Since η is a minimal geodesic connecting x_{i+1} to x_{i-1}, it has a midpoint, and each point on η is closer to one endpoint than the other, except the midpoint, which is equidistant from both. If x is closer to x_{i-1} than to x_{i+1}, or equidistant from both, then, since $l(\eta) < \epsilon$, we must have $d(x_{i-1}, x) < \epsilon/2$. It follows that

$$d(x, x_i) \leq d(x_i, x_{i-1}) + d(x_{i-1}, x) < \epsilon + \frac{\epsilon}{2} = \frac{3\epsilon}{2}.$$

If x is closer to x_{i+1} than to x_{i-1}, then we must have $d(x_{i+1}, x) < \epsilon/2$, so

$$d(x, x_i) \leq d(x_i, x_{i+1}) + d(x_{i+1}, x) < \epsilon + \frac{\epsilon}{2} = \frac{3\epsilon}{2}.$$

Thus, β lies in $B(x_i, 3\epsilon/2)$, proving that $\alpha\beta\alpha^{-1}$ is of the desired form. Moreover, since λ was ϵ-nullhomotopic in $n+1$ steps, and since we just completed the first one, the new chain $\lambda' = \{x_0, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n\}$ is an ϵ-chain along γ' that is ϵ-nullhomotopic via a homotopy of n steps. In addition, each segment of γ' determined by λ' is either a segment of the original chain/path pair, in which case it satisfies the ultra ϵ-chain property, or it is the segment $[x_{i-1}, x_{i+1}]$ along γ', which consists of the minimizing geodesic, η. Since $l(\eta) < \epsilon$, it follows that this segment is contained in $B(x_{i-1}, \epsilon) \cap B(x_{i+1}, \epsilon)$. Thus, λ' is an ultra ϵ-chain along γ' that is

ϵ-null via n steps, so γ' satisfies the inductive hypothesis. This implies that it can be expressed as a product of loops of the required form. Hence, γ can, also.

Now, if the first step of the nullhomotopy of λ adds a point, x, in between x_i and x_{i+1}, then a similar "triangle" argument involving x_i, x, and x_{i+1} holds. One simply chooses minimizing geodesics, η_1 and η_2, from x_i to x and from x to x_{i+1}, respectively. The same reasoning as in the previous case will show that the loop $\beta = \gamma|_{[t_i,t_{i+1}]}\eta_2^{-1}\eta_1^{-1}$ lies in $B(x_i, 3\epsilon/2)$, and this completes the proof. ∎

Lemma 3.1.9 *Let X be a compact geodesic space, and let N be the smallest natural number such that $N > \frac{2}{\delta}diam(X)$ for some $\delta > 0$. Suppose, for some $\epsilon > \delta$, the map $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is not injective, and let K be the smallest natural number such that $K > 3\epsilon/\delta$. Then there is a δ-loop at $*$ consisting of at most $2(N + K)$ points that is δ-nontrivial but ϵ-null.*

Proof By the hypothesis, there is a non-trivial δ-loop at $*$ that is ϵ-null. We will denote this loop by λ. By joining the consecutive points of λ by minimal geodesics and concatenating them, we obtain a path loop, γ, and λ is an ultra δ-chain along γ. Since $\delta < \epsilon$, λ is also an ultra ϵ-chain along γ. Since λ is δ-nontrivial, γ lifts as an open path loop to X_δ, but it lifts closed to X_ϵ since λ is ϵ-null. By the previous lemma, γ is path-homotopic to a finite product of loops, $\alpha_1\beta_1\alpha_1^{-1}\cdots\alpha_k\beta_k\alpha_k^{-1}$, where, for each i, α_i is a path from $*$ to a point, x_i, and β_i is a loop at x_i lying in an open ball of radius $3\epsilon/2$.

Fix i, $1 \le i \le k$, and consider the path loop $\alpha_i\beta_i\alpha_i^{-1}$. Let $B(\bar{x}, 3\epsilon/2)$ be the ball containing β_i. Since β_i is a compact subset of this ball, there is some τ, $0 < \tau < 3\epsilon/2$, such that β_i lies in $B(\bar{x}, 3\epsilon/2 - \tau)$. We may assume that $\tau < \min\{\delta, \epsilon/2\}$. Choose an ultra τ-chain along β_i, denoted by $\{x_i = y_0, \ldots, y_n = x_i\}$, and, for each $j = 1, \ldots, n$, let β_i^j denote the subsegment of β_i from y_{j-1} to y_j. For each $j = 0, 1, \ldots, n$, choose a minimal geodesic, ω_j, from $\bar{x}$ to y_j, being sure to choose $\omega_0 = \omega_n$. Note that $\alpha_i\beta_i\alpha_i^{-1}$ is path-homotopic to the path loop

$$(\alpha_i\omega_0^{-1})(\omega_0\beta_i^1\omega_1^{-1})(\omega_1\beta_i^2\omega_2^{-1})\cdots(\omega_{n-2}\beta_i^{n-1}\omega_{n-1}^{-1})(\omega_{n-1}\beta_i^n\omega_n^{-1})(\omega_n\alpha^{-1}).$$

By inserting the trivial loop $\omega_0\alpha^{-1}\alpha\omega_0^{-1}$ in between each pair

$$(\omega_{j-1}\beta_i^j\omega_j^{-1})(\omega_j\beta_i^{j+1}\omega_{j+1}^{-1}),$$

we see that $\alpha_i\beta_i\alpha_i^{-1}$ is further path-homotopic to the product

$$[(\alpha_i\omega_0^{-1})(\omega_0\beta_i^1\omega_1^{-1})(\omega_0\alpha_i^{-1})]\cdots[(\alpha_i\omega_0^{-1})(\omega_{n-1}\beta_i^n\omega_n^{-1})(\omega_0\alpha_i^{-1})].$$

Now, we claim that around each triangular path loop, $\omega_{j-1}\beta_i^j\omega_j^{-1}$, there is a strong ϵ-triangle. So, fix j, $1 \le j \le n$. Since ω_{j-1} is a minimal geodesic from $\bar{x}$ to $y_{j-1} \in B(\bar{x}, 3\epsilon/2 - \tau)$, it has length strictly less than $3\epsilon/2 - \tau$. The sets $\{z \in \omega_{j-1} : d(\bar{x}, z) < \epsilon\}$ and $\{z \in \omega_{j-1} : d(y_{j-1}, z) < \epsilon/2 - \tau\}$ must intersect, for, if not, then ω_{j-1} would have length at least $\epsilon + \epsilon/2 - \tau = 3\epsilon/2 - \tau$, a contradiction. Thus, there is some point on ω_{j-1} that is strictly within ϵ of $\bar{x}$ and strictly within $\epsilon/2 - \tau$ of y_{j-1}. Denote this point by u_1. Likewise, there is some point on ω_j, which we will denote by u_2, that is strictly within ϵ of $\bar{x}$ and strictly within $\epsilon/2 - \tau$ of y_j. Consider the chain $\{\bar{x} = u_0, u_1, u_2, u_3 = \bar{x}\}$. This is a chain along the path loop $\omega_{j-1}\beta_i^j\omega_j$. We've just shown that $d(u_0, u_1), d(u_3, u_2) < \epsilon$, and, since ω_{j-1} and ω_j are minimal geodesics, the path subsegments of $\omega_{j-1}\beta_i^j\omega_j$, $[u_0, u_1]$ and $[u_2, u_3]$, lie in $B(u_0, \epsilon)$ and $B(u_2, \epsilon)$, respectively. So, we need only show that the path subsegment, $[u_1, u_2]$, lies in $B(u_1, \epsilon)$ to show that this is a strong ϵ-triangle along

$\omega_{j-1}\beta_i^j\omega_j$. Note that the path subsegment, $[u_1, u_2]$, consists of the portion of ω_{j-1} between u_1 and y_{j-1}, followed by β_i^j, and then followed by the portion of ω_j between y_j and u_2. So, if z lies on $[u_1, u_2]$, then there are three possibilities. If z lies on ω_{j-1}, then z lies between u_1 and y_{j-1} on this minimal geodesic, meaning that $d(u_1, z) \leq d(u_1, y_{j-1}) < \epsilon/2 - \tau < \epsilon$. If z lies on β_i^j, then, since $\{y_0, \ldots, y_n\}$ is an ultra τ-chain along β, β_i^j lies in $B(y_{j-1}, \tau)$. Thus,

$$d(u_1, z) \leq d(u_1, y_{j-1}) + d(y_{j-1}, z) < \epsilon/2 - \tau + \tau < \epsilon.$$

Finally, if z lies on ω_j, then z lies between u_2 and y_j. Thus,

$$d(u_1, z) \leq d(u_1, y_{j-1}) + d(y_{j-1}, y_j) + d(y_j, z) < \epsilon/2 - \tau + \tau + \epsilon/2 - \tau < \epsilon.$$

So, $\{\bar{x} = u_0, u_1, u_2, u_3 = \bar{x}\}$ is a strong ϵ-triangle along $\omega_{j-1}\beta_i^j\omega_j$, and we can construct such a triangle for each j.

Since ϵ-triangles are ϵ-null, this implies that each of the path loops, $\omega_{j-1}\beta_i^j\omega_j$, lifts closed to X_ϵ, and this further implies that each path loop $(\alpha_i\omega_0^{-1})(\omega_{j-1}\beta_i^k\omega_j)(\omega_0\alpha_i^{-1})$ lifts closed to X_ϵ. Moreover, this can be carried out for each loop $\alpha_i\beta_i\alpha_i^{-1}$, $i = 1, \ldots, k$. So, to summarize, we have shown that γ is path-homotopic to a finite product of path loops of the form

$$(\alpha_i\omega_0^{-1})(\omega_{j-1}\beta_i^j\omega_j^{-1})(\omega_0\alpha_i^{-1}),$$

and each of these loops lifts closed to X_ϵ. Now, at least one of these path loops must lift open to X_δ, for if they all lifted closed to X_δ, then γ would, also, a contradiction. This further implies that at least one of the triangular path loops, $\omega_{j-1}\beta_i^j\omega_j^{-1}$, must lift open to X_δ. Indeed, the "lollipop" structure of the path loop $(\alpha_i\omega_0^{-1})(\omega_{j-1}\beta_i^j\omega_j^{-1})(\omega_0\alpha_i^{-1})$ is such that this path loop will lift closed or open if and only if the head - in this case, $\omega_{j-1}\beta_i^j\omega_j^{-1}$ - lifts the same way.

Finally, we let $\omega_{j-1}\beta_i^j\omega_j^{-1}$ be such a path loop that lifts closed to X_ϵ but open to X_δ. Choose a minimal geodesic, η, from $*$ to the initial point of ω_{j-1}. Then η has length at most $diam(X)$, so it can be subdivided into N or fewer minimal subsegments of length $\delta/2$. The geodesic, ω_{j-1}, has length less than $3\epsilon/2$, so we can divide it up into K or fewer minimal segments of length less than $\delta/2$, and likewise for ω_j. Since β_i^j lies in $B(y_{j-1}, \delta)$, we can take the partition points making up these subsegments and obtain a strong δ-chain along $\eta(\omega_{j-1}\beta_i^j\omega_j^{-1})\eta^{-1}$ consisting of, at most, $2(N+K)$ points. This strong δ-chain will also be a strong ϵ-chain. Since the path loop $\eta(\omega_{j-1}\beta_i^j\omega_j^{-1})\eta^{-1}$ lifts closed to X_ϵ, this strong δ-chain will be ϵ-null. On the other hand, since this path loop lifts open to X_δ, this chain will be δ-nontrivial, completing the proof. ∎

We have already seen that geodesic spaces do not have refinement critical values. The next lemma shows that, in the compact case, they can only have one type of critical value.

Lemma 3.1.10 *If X is a compact geodesic space, then the only types of critical values in $Cr(X)$ are upper non-injective critical values.*

Proof Suppose there were some $\epsilon > 0$ in $Cr(X)$ that was a lower non-injective critical value. Using the limit characterization of a critical value, along with the fact that ϵ-homotopies (in this case, nullhomotopies) are also $(\epsilon - \tau)$-homotopies for sufficiently small τ, this means that we could find a sequence of positive real numbers, $\{\epsilon_n\}$, such that $\epsilon/2 < \epsilon_n < \epsilon$ for all n, $\epsilon_n \nearrow \epsilon$, and the map $X_{\epsilon_{n-1}} \to X_{\epsilon_n}$ is not injective for every $n \geq 2$. Now, since $\epsilon/2 < \epsilon_n < \epsilon$ for all n, if we choose $N > \frac{4}{\epsilon} diam(X)$ and $K > 6$, then we have $N > \frac{2}{\epsilon_n} diam(X)$ and $K > 3\epsilon_{n+1}/\epsilon_n$ for all

n. So, by the previous lemma, for each n, there is an ϵ_n-loop, γ_n, at $*$ consisting of $2(N+K)$ or fewer points that is ϵ_n-nontrivial but ϵ_{n+1}-null. In particular, γ_n and γ_m, for $m > n$, cannot be ϵ_m-homotopic, since one is ϵ_m-null and the other is not. By inserting midpoints, we can ϵ_n-refine each γ_n to an $\epsilon/2$-chain, and this refined chain will still be ϵ_n-nontrivial and ϵ_{n+1}-null. Moreover, this refinement no more than doubles the number of points in γ_n, since $\epsilon/2 < \epsilon_n < \epsilon$. We will just assume that γ_n is already so refined.

This gives us a sequence of $\epsilon/2$-chains, $\{\gamma_n\}$, with a uniform bound on their cardinalities. By repeating the initial points of each chain a finite number of times, if necessary, we can assume that they all have the same number of points. Now, fix some δ such that $\epsilon/2 < \delta < \epsilon_1$. Then each γ_n is also a δ-loop. By Lemma 2.7.2, there is a subsequence, $\{\gamma_{n_k}\}$, converging pointwise to a loop in which the distances between consecutive points is less than or equal to $\epsilon/2$. That is, this limiting loop is also a δ-loop. By Lemma 2.7.1, it follows that for all sufficiently large k, γ_{n_k} will be δ-homotopic to this limiting chain. In particular, for all sufficiently large k, the loops γ_{n_k} will be δ-homotopic to each other, so there is some M such that $k \geq M$ implies $\gamma_{n_{k+1}} \sim_\delta \gamma_{n_k}$. But a δ-homotopy is also an ϵ_{n_k}-homotopy for any k, since $\delta < \epsilon_1$. In particular, this would imply that γ_{n_k} is $\epsilon_{n_{k+1}}$-homotopic to $\gamma_{n_{k+1}}$, which is a contradiction, because γ_{n_k} is $\epsilon_{n_{k+1}}$-null while $\gamma_{n_{k+1}}$ is not. ∎

Theorem 3.1.11 *Let X be a compact geodesic space. Then $Cr(X)$ is discrete and bounded above in $\mathbb{R}^+$.*

Proof We only need to show that there are not any positive limit points of $Cr(X)$. First, suppose there is a sequence of critical values, $\{\epsilon_n\}$, converging up to some $\epsilon > 0$. This means that, for each n, there is some δ_n satisfying $\epsilon_n < \delta_n < \epsilon_{n+1}$ and such that the map $X_{\epsilon_n} \to X_{\delta_n}$ is not injective, and this further implies that the map $X_{\epsilon_n} \to X_{\epsilon_{n+1}}$ is not injective. But this is precisely the situation we just showed could not occur in the proof of the previous lemma. Hence, there can be no such sequence of critical values converging up to a positive number.

Suppose, next, that there is a sequence of critical values of X, $\{\epsilon_n\}$, so that $\epsilon_n \searrow \epsilon$, where $\epsilon > 0$. We may assume that $\epsilon < \epsilon_{n+1} < \epsilon_n < 2\epsilon$ for all n. This means that the map $X_{\epsilon_{n+1}} \to X_{\epsilon_n}$ is not injective for all n. Since $\epsilon < \epsilon_{n+1} < \epsilon_n < 2\epsilon$ for all n, if we choose $N > \frac{2}{\epsilon} diam(X)$ and $K > 3$, then we have $N > \frac{2}{\epsilon_n} diam(X)$ and $K > 3\epsilon_{n+1}/\epsilon_n$ for all n. So, by Lemma 3.1.9, there is, for each $n \geq 2$, an ϵ_n-loop, γ_n, that is ϵ_n-nontrivial but ϵ_{n-1}-null and consists of, at most $2(N+K)$ points. In particular, for $m > n$, γ_m cannot be ϵ_n-homotopic to γ_n, since γ_m is ϵ_n-trivial while γ_n is not.

Now, we can ϵ_n-refine each γ_n to an $\epsilon/2$-loop by adding midpoints, and this does not change any of the above homotopy properties. Moreover, since $\epsilon < \epsilon_n < 2\epsilon$, the cardinalities of the refined loops still remain uniformly bounded. As before, we may assume that they each have exactly the same number of points by repeating the initial point finitely many times, if necessary. We will assume that each γ_n is so refined. Then, by Lemma 2.7.2, there will be some subsequence, $\{\gamma_{n_k}\}$, converging pointwise to a loop in which the distances between consecutive points is less than or equal to $\epsilon/2$. This limiting loop is an ϵ-loop, as is each γ_{n_k}. Thus, reasoning as before, there is some M such that $k \geq M$ implies that $\gamma_{n_{k+1}} \sim_\epsilon \gamma_{n_k}$. This would imply, however, that $\gamma_{n_{k+1}} \sim_{\epsilon_{n_k}} \gamma_{n_k}$ for $k \geq M$, a contradiction.

That $Cr(X)$ is bounded above follows because - as we have noted previously - for all sufficiently large ϵ, X_ϵ is isometric to X. ∎

3.2 Covering Comparison for Geodesic Spaces

The next step in deriving the spectral comparison is to formulate some results comparing the ϵ-covers and the Spanier covers of a geodesic space. For the remainder of this chapter, we will assume that our given metric space, X, is a geodesic space. We also remind the reader of our convention of denoting entourage covers by subscripts and Spanier covers by superscripts. In particular, by earlier remarks and results in the previous chapter, X^δ and X_ϵ are both path-connected and locally path-connected, X^δ is endowed with the lifted length metric, and d_ϵ is the lifted length metric on X_ϵ. We will continue to use $\varphi_\epsilon : X_\epsilon \to X$ and $\psi_\delta : X^\delta \to X$ to denote the entourage and Spanier covering maps, respectively.

Lemma 3.2.1 *Let X be a geodesic space, and let $\delta > 0$ be given. A path loop at $*$ lifts to a closed loop in X^δ if and only if it lifts closed in $X_{2\delta/3}$.*

Proof First, suppose γ is a path loop at $*$ that lifts closed to $X_{2\delta/3}$. Then there is an ultra $(2\delta/3)$-chain along γ that is $(2\delta/3)$-null. By Lemma 3.1.8, it follows that γ is path-homotopic to a product of loops of the form $\alpha\beta\alpha^{-1}$, where β is a loop lying in a ball of radius $(3/2)(2\delta/3) = \delta$. These are precisely the loops that lift closed to X^δ by construction.

For the converse, let $\hat{*}$ be the point in X^δ containing the equivalence class of the constant path at $*$ (recall the construction of X^δ). As we just noted, any path loop lifting to a closed path loop at $\hat{*} \in X^\delta$ is homotopic to a finite product of path loops of the form $\alpha\beta\alpha^{-1}$, where β is a loop lying in a ball of radius δ and α is a path from $*$ to the initial point of β. If we can show that any such loop, $\alpha\beta\alpha^{-1}$, lifts to a closed loop at $\tilde{*} \in X_\epsilon$, then any product of such loops will, also. Thus, we've reduced the proof to verifying this claim.

Let $\gamma = \alpha\beta\alpha^{-1}$ be a path loop at $*$ as described above, so that β lies in a ball of radius δ centered at some point $\bar{x} \in X$. Since the image of β is compact and lies in the open ball $B(\bar{x}, \delta)$, there is some λ, $0 < \lambda < \delta$, such that β actually lies in $B(\bar{x}, \delta - \lambda)$. We may assume without loss of generality that $\lambda < \delta/3$. Choose an ultra λ-chain along β, and partition β into the resulting subsegments, $[\beta(t_0), \beta(t_1)]$, $[\beta(t_1), \beta(t_2)]$, $\ldots$, $[\beta(t_{n-1}), \beta(t_n)]$, with $0 = t_0 < t_1 < \cdots < t_{n-1} < t_n = 1$. Let β_i denote the subsegment $[\beta(t_{i-1}), \beta(t_i)]$ for each $i = 1, \ldots, n$. Then $\beta_i \subset B(\beta(t_{i-1}), \lambda) \cap B(\beta(t_i), \lambda)$ for each $i = 1, \ldots, n$. For each $i = 0, 1, \ldots, n$, let ω_i be a minimizing geodesic from $\bar{x}$ to $\beta(t_i)$, being sure to choose $\omega_n = \omega_0$. Then each ω_i lies in $B(\bar{x}, \delta - \lambda)$. Note that $\gamma = \alpha\beta\alpha^{-1}$ is path-homotopic to the loop

$$\gamma' = \alpha\omega_0^{-1}(\omega_0\beta_1\omega_1^{-1})(\omega_1\beta_2\omega_2^{-1})\cdots(\omega_{n-1}\beta_n\omega_n^{-1})\omega_0\alpha^{-1}.$$

The path $\alpha\omega_0^{-1}$ is a path from $*$ to $\bar{x}$, and each $\omega_{i-1}\beta_i\omega_i^{-1}$ is a loop at $\bar{x}$.

Now, since γ is path-homotopic to γ', γ will lift to a closed loop in X_ϵ if and only if γ' does, so we need only show that. The rest of the proof now follows exactly as in the proof of Lemma 3.1.9. Following the same steps as in that proof - just with $3\epsilon/2$ replaced by δ and ϵ replaced by $2\delta/3$ - one constructs a strong $(2\delta/3)$-triangle along each of the triangular path loops $\omega_{i-1}\beta_i\omega_i$, say σ_i. Then we take an *ultra* $2\delta/3$-chain along $\alpha\omega_0^{-1}$ - say, μ - so that its inverse is a strong $2\delta/3$-chain along $\omega_0\alpha^{-1}$. Concatenating these chains yields $\mu\sigma_1\cdots\sigma_n\mu^{-1}$, a strong $2\delta/3$-loop along γ' that is $2\delta/3$-nullhomotopic. By Theorem 3.1.5, it will follow that γ' - and, thus, γ - lifts to a closed loop in $X_{2\delta/3}$. ∎

Lemma 3.2.2 *Let X be a geodesic space, and let $\delta > 0$ be given. If $\epsilon = 2\delta/3$, then ϵ-triangles in X lift to ϵ-triangles in X^δ.*

Proof Let $\{x_0, x_1, x_2, x_3 = x_0\}$ be an ϵ-triangle in X, with $\epsilon = 2\delta/3$. Let γ be a minimal geodesic from x_1 to x_2, and assume γ is parameterized proportional to arclength on $[0,1]$. We claim that γ lies in $B(x_0, \delta)$.

Let $\gamma(1/2)$ be the midpoint of γ, so, for any $0 \leq t \leq 1/2$, we have

$$d(x_1, \gamma(t)) = d(\gamma(0), \gamma(t)) \leq d(\gamma(0), \gamma(1/2)) = \frac{1}{2} l(\gamma) < \frac{1}{2}\frac{2\delta}{3} = \frac{\delta}{3}.$$

It follows that, for any such t,

$$d(x_0, \gamma(t)) \leq d(x_0, x_1) + d(x_1, \gamma(t)) < \frac{2\delta}{3} + \frac{\delta}{3} = \delta.$$

Likewise, we have $d(x_2, \gamma(t)) < \delta/3$ for any $1/2 \leq t \leq 1$, and, for any such t,

$$d(x_0, \gamma(t)) =\leq d(x_0, x_2) + d(x_2, \gamma(t)) < \frac{2\delta}{3} + \frac{\delta}{3} = \delta.$$

Now, $\psi_\delta : X^\delta \to X$ is a bijection and radial isometry on δ-balls. Let $\tilde{x}_0$ be any point in $\psi_\delta^{-1}(x_0)$. There are unique points, $\tilde{x}_1$ and $\tilde{x}_2$, lying in $B(\tilde{x}_0, \delta)$ and projecting under ψ_δ to x_1 and x_2, respectively. We clearly have $d(\tilde{x}_0, \tilde{x}_1) < \epsilon$ and $d(\tilde{x}_0, \tilde{x}_2) < \epsilon$ in X^δ by the radial isometry property. The geodesic γ will lift to a geodesic in $B(\tilde{x}_0, \delta)$ between $\tilde{x}_1$ and $\tilde{x}_2$, showing that $d(\tilde{x}_1, \tilde{x}_2) < \epsilon$ in X^δ. Thus, $\{\tilde{x}_0, \tilde{x}_1, \tilde{x}_2, \tilde{x}_0\}$ is an ϵ-triangle in X^δ projecting to $\{x_0, x_1, x_2, x_0\}$. ∎

Theorem 3.2.3 *Let X be a geodesic space, and let $\delta > 0$ be given. Then $X_{2\delta/3}$ is isometric to X^δ, and the isometry is a covering equivalence between $\varphi_{2\delta/3}$ and ψ_δ.*

Proof Let $\epsilon = 2\delta/3$. Since ψ_δ is a radial isometry and bijection from δ-balls in X^δ onto δ-balls in X, it is also a radial isometry and bijection from ϵ-balls in X^δ onto ϵ-balls in X. The previous lemma also shows that ϵ-triangles in X lift to ϵ-triangles in X^δ. Let $\tilde{*}$ be the usual base point in X_ϵ, and let $\hat{*}$ be the point (i.e. equivalence class) in X^δ containing the constant path at $*$. So, we have pointed covering maps $\varphi_\epsilon : (X_\epsilon, \tilde{*}) \to (X, *)$ and $\psi_\delta : (X^\delta, \hat{*}) \to (X, *)$. By Lemma 2.4.9, there is a surjective map, $\rho : (X_\epsilon, \tilde{*}) \to (X^\delta, \hat{*})$, that is a uniform local isometry and satisfies $\psi_\delta \circ \rho = \varphi_\epsilon$. We need only further show that ρ is injective to prove the desired result. In fact, since ρ is a uniform local isometry, and since X_ϵ and X^δ are geodesic spaces, ρ preserves the lengths of curves. So, if it is injective, it will be an isometry.

Suppose $\rho([\alpha]_\epsilon) = \rho([\beta]_\epsilon)$. Choose rectifiable curves, $\tilde{\gamma}$ and $\tilde{\sigma}$, in X_ϵ from $\tilde{*}$ to $[\alpha]_\epsilon$ and $[\beta]_\epsilon$, respectively. Choose strong ϵ-chains along $\tilde{\gamma}$ and $\tilde{\sigma}$, denoted, respectively, by $\tilde{\gamma}_c$ and $\tilde{\sigma}_c$. Let γ, σ, γ_c, and σ_c denote the projections of these geodesics and ϵ-chains to $*$ in X, and note that γ_c and σ_c are *strong* ϵ-chains along γ and σ, respectively. Then let $\hat{\gamma}$, $\hat{\sigma}$, $\hat{\gamma}_c$, and $\hat{\sigma}_c$ denote the lifts of these geodesics and ϵ-chains to $\hat{*} \in X^\delta$. We first claim that the endpoints of $\hat{\gamma}$ and $\hat{\gamma}_c$ are the same, and similarly for the endpoints of $\hat{\sigma}$ and $\hat{\sigma}_c$. To prove this, we choose a strong ϵ-chain along $\hat{\gamma}$, say $\hat{\lambda}$. The projection of $\hat{\lambda}$ to X will give us a strong ϵ-chain along γ. By Lemma 3.1.3, it will follow that this projected chain is ϵ-homotopic to γ_c in X. Since the lifts of homotopic chains are homotopic and end at the same point, it follows that the endpoint of $\hat{\gamma}_c$ is the same as the endpoint of $\hat{\lambda}$, which is the endpoint of $\hat{\gamma}$. The same argument holds for $\hat{\sigma}$ and $\hat{\sigma}_c$.

Now, by definition of ρ (see the proof of Lemma 2.4.9), the assumption that $\rho([\alpha]_\epsilon) = \rho([\beta]_\epsilon)$ implies that $\hat{\gamma}_c$ and $\hat{\sigma}_c$ end at the same point in X^δ. Thus, $\hat{\gamma}$ and $\hat{\sigma}$ end at the same point in X^δ. It follows from uniqueness of lifts that γ and σ must end at the same point in X. So, $\gamma\sigma^{-1}$ is a loop in X that lifts to a loop in X^δ. By Lemma 3.2.1, it follows that $\gamma\sigma^{-1}$ lifts to a loop at $* \in X_\epsilon$. It follows, again, from uniqueness of lifts that $\tilde{\gamma}$ and $\tilde{\sigma}$ end at the same point in X_ϵ. Hence, $[\alpha]_\epsilon = [\beta]_\epsilon$. ∎

3.3 Spectral Comparison

This covering comparison will finally allow us to compare the covering spectrum and critical spectrum of a geodesic space. The fact that X^δ is isometric to $X_{2\delta/3}$ (or, equivalently, that X_ϵ is isometric to $X^{3\epsilon/2}$) immediately yields the following corollary, which is really the key to comparing the spectra.

Corollary 3.3.1 *For any $\delta > 0$, a loop at $* \in X$ lifts closed to $\hat{*} \in X^\delta$ if and only if it lifts closed to $\tilde{*} \in X_{2\delta/3}$. Equivalently, a loop at $* \in X$ lifts closed to $\tilde{*} \in X_\epsilon$ if and only if it lifts closed to $\hat{*} \in X^{3\epsilon/2}$.*

Note, also, that because the Spanier and ϵ-covers are *regular* covers, lifts of loops are either always closed or always open, so lifting closed to $\tilde{*}$ or $\hat{*}$ will mean that the loop lifts closed to any point in the preimage of $*$.

Before we proceed with the next theorem, we need to discuss some simple but subtle details. Given a subset $S \subset \mathbb{R}_+$, we let $\bar{S}$ denote its closure in $\mathbb{R}$, or equivalently, in $\overline{\mathbb{R}}_+$. It is easy to see that $\lambda\bar{S} = \overline{\lambda S}$ for any $S \subset \mathbb{R}$ and $\lambda > 0$. Now, the critical spectrum and covering spectrum, both, by definition, consist of only positive real numbers, though either may have 0 as a limit point. We have already seen that $Cr(X)$ is either closed in $\mathbb{R}$ or is "almost" closed in the sense that the only limit point of $Cr(X)$ that may not be in $Cr(X)$ is 0. The same reasoning does not hold for $Cov(X)$, due to one difference: $Cov(X)$ may not contain all of its positive limit points (though this can only happen if X is not compact).

Example 3.3.2 *Consider the sequence $\{1 - 1/n\}_{n=2}^\infty$. Obviously, we have $1 - 1/n < 1$ for all n and $1 - 1/n \nearrow 1$. Define a geodesic space, X, as follows. Take $\mathbb{R}$, and at each integer $n \geq 2$, attach a circle, C_n, of circumference $3(1 - 1/n)$. Give this space the obvious geodesic metric. This space is complete and locally compact, but not compact. For each individual circle, C_n, $1 - 1/n$ is a critical value. So, $1 - 1/n \in Cr(X)$ for all $n \geq 2$. Thus, we also have $1 \in Cr(X)$. Moreover, for each individual circle, C_n, $3/2(1 - 1/n)$ is in the covering spectrum of that circle. So, $3/2(1 - 1/n) \in Cov(X)$ for all $n \geq 2$. However, $3/2 = \lim_{n\to\infty} 3/2(1 - 1/n)$ is not in $Cov(X)$. In fact, since each circle, C_n, is contained in a ball of radius $3/2$, any loop in X based, say, at 0, will be homotopic to a product of loops of the form $\alpha\beta\alpha^{-1}$ where β is a loop lying in a ball of radius $3/2$. Thus, $\pi_1(X, 0, 3/2)$ is the whole fundamental group of X, and $X^\delta = X$ for all $\delta \geq 3/2$. This implies that $3/2$ is not in $Cov(X)$, even though it is a positive limit point of $Cov(X)$. This phenomenon cannot occur, however, for a compact geodesic space.* ∎

Lemma 3.3.3 *Let X be a geodesic space. Then*

$$\frac{2}{3}Cov(X) \subseteq Cr(X).$$

Proof Suppose $\delta \in Cov(X)$. Then there is a sequence $\{\delta_i\}$, with $\delta_i \searrow \delta$, such that $X^\delta \neq X^{\delta_i}$ for all i. This means that, for each i, there is a path loop, γ_i, that lifts closed to X^{δ_i} but open to X^δ. We also have, by Corollary 3.3.1, that γ_i lifts closed to $X_{2\delta_i/3}$ and open to $X_{2\delta/3}$. So, for each i, there is a strong $2\delta/3$-chain, α_i, along γ_i, and α_i is $2\delta/3$-nontrivial. However, since $\delta < \delta_i$, α_i is also a strong $2\delta_i/3$-chain along γ_i, and since γ_i lifts closed to $X_{2\delta_i/3}$, α_i is $2\delta_i/3$-nullhomotopic. Hence, the map $X_{2\delta/3} \to X_{2\delta_i/3}$ is not injective for each i. Since $2\delta_i/3 \searrow 2\delta/3$, it follows that $2\delta/3 \in Cr(X)$. ∎

Lemma 3.3.4 *If X is a geodesic space, then*

$$\frac{3}{2}Cr(X) \subseteq \overline{Cov(X)}.$$

Proof Assume $\delta \in Cr(X)$. Then there is a sequence, $\{\delta_i\}$, with $\delta_i \to \delta$, such that the map between X_δ and X_{δ_i} is not bijective for each i. There are two possibilities: either $\{\delta_i\}$ contains a subsequence converging down to δ, or there are at most finitely many $\delta_i > \delta$, in which case there is a subsequence of $\{\delta_i\}$ convering up to δ.

Consider the former case first, and, for simplicity, denote the subsequence simply by $\{\delta_i\}$. So, each map $X_\delta \to X_{\delta_i}$ is not bijective. For a geodesic space, each such map is necessarily surjective, so each map $X_\delta \to X_{\delta_i}$ is non-injective. This means that, for each i, there is a δ-loop at $*$, α_i, that is δ-nontrivial but δ_i-nullhomotopic. Connect each consecutive pair of points of α_i by a minimal geodesic, and let γ_i denote the resulting path loop at $*$. Then α_i is a strong δ-chain along γ_i. Since α_i is δ-nontrivial, γ_i lifts open to X_δ. But, since $\delta < \delta_i$, α_i is also a strong δ_i-chain along γ_i. Since α_i is δ_i-nullhomotopic, γ_i lifts closed to X_{δ_i}. By Corollary 3.3.1, we have that γ_i lifts open to $X^{3\delta/2}$ but closed to $X^{3\delta_i/2}$. Hence, $X^{3\delta/2} \neq X^{3\delta_i/2}$ for all i. Since $3\delta_i/2 \searrow 3\delta/2$, it follows that $3\delta/2 \in Cov(X)$.

Next, consider the latter case, so we get a sequence $\delta_i \nearrow \delta$ such that each map $X_{\delta_i} \to X_\delta$ is not injective. This means that there is, for each i, a δ_i-loop at $*$, α_i, that is δ_i-nontrivial but δ-nullhomotopic. Connect each consecutive pair of points of α_i by a minimal geodesic, and let γ_i denote the resulting path loop at $*$. Then α_i is a strong δ_i-chain along γ_i. Since α_i is δ_i-nontrivial, γ_i lifts open to X_{δ_i}. But α_i is also a strong δ-chain along γ_i, and, since α_i is δ-nullhomotopic, γ_i lifts closed to X_δ. By Corollary 3.3.1, γ_i lifts closed to $X^{3\delta/2}$ and open to $X^{3\delta_i/2}$. This means that $X^{3\delta_i/2} \neq X^{3\delta/2}$ for all i, and $3\delta_i/2 \nearrow 3\delta/2$. Suppose, now, that $3\delta/2$ was not in $\overline{Cov(X)}$. By Lemma 1.4.11, there would be an interval $(3\delta/2 - \epsilon, 3\delta/2 + \epsilon)$ such that $X^{\delta'} = X^{3\delta/2}$ for all δ' in this interval. But $3\delta_i/2 \in (3\delta/2 - \epsilon, 3\delta/2]$ for all sufficiently large i, and $X^{3\delta_i/2} \neq X^{3\delta/2}$ for all i. This is a contradiction. Hence, we have $3\delta/2 \in \overline{Cov(X)}$. ∎

Corollary 3.3.5 *If X is a compact geodesic space, then*

$$\frac{2}{3}Cov(X) = Cr(X).$$

Proof Lemma 3.3.3 holds whether X is compact or not, and that containment - along with the discreteness and boundedness of $Cr(X)$ - shows that $Cov(X)$ is also discrete and bounded above in $\mathbb{R}^+$. Thus, $Cov(X) = \overline{Cov(X)}$ if $\inf Co(X) > 0$, and $\overline{Cov(X)} = Cov(X) \cup \{0\}$ if $\inf Cov(X) = 0$. The previous lemma now completes the proof. ∎

This shows that the critical spectrum is, indeed, a genuine generalization of the covering spectrum. In the case of geodesic spaces, our critical spectrum yields the same information as the covering spectrum; the only difference is in the exact geometric measurements at which each spectrum detects a new non-trivial loop in the space. Very roughly speaking, the covering spectrum detects a loop or fundamental group element when δ reaches $1/2$ the circumference, or the diameter, of that loop. The critical spectrum detects loops when ϵ reaches $1/3$ the circumference, or $2/3$ the diameter, of that loop. Of course, these are intuitive measurements in general, since non-trivial loops in geodesic spaces need not be the circular looking loops this rough description is using. For nice enough geodesic spaces, however, including all Riemannian manifolds, it turns out that this description is, indeed, valid (see [9]).

Chapter 4

Some Examples

In this chapter we will present several examples illustrating some of the phenomena that can occur within the critical spectrum of a general compact metric space. The examples in this chapter will focus on refinement critical values.

4.1 Essential Gaps

We will first prove some technical results that facilitate finding refinement critical values more efficiently. Most of these technical results apply only to a particular class of spaces possessing a special type of structure. Nevertheless, they are a fruitful source of examples illustrating the various behaviors one can encounter in the critical spectrum.

Let X be a chain-connected metric space. Assume there are two points, $x, y \in X$, with $d(x, y) = l > 0$, and a number $\epsilon^* > l$ such that the following holds: for each $l < \epsilon \leq \epsilon^*$, if we let $B_x = B(x, \epsilon - l)$ and $B_y = B(y, \epsilon - l)$, then we can express X as a disjoint union, $X = Z \cup Y$, such that

1) $B_x \subset Z$ and $B_y \subset Y$ (hence $B_x \cap B_y = \emptyset$),

2) the only points in Z that are strictly within ϵ of a point in B_y lie in B_x, and the only points of Y that are strictly within ϵ of a point in B_x lie in B_y.

Given any such ϵ and any ϵ-chain, $\gamma = \{x_0, \ldots, x_n\}$, a pair of consecutive points, (x_{i-1}, x_i), $1 \leq i \leq n$, will be said to ***contain*** or ***cross the*** x, y***-gap*** if and only if x_{i-1} lies in either B_x or B_y and x_i lies in the other ball. Assign each pair of consecutive points a value, denoted by $|x_{i-1}, x_i|$, as follows:

$$|x_{i-1}, x_i| = \begin{cases} 0, & (x_{i-1}, x_i) \text{ does not contain the } x, y\text{-gap} \\ 1, & x_{i-1} \in B_x, \;\; x_i \in B_y \\ -1, & x_{i-1} \in B_y, \;\; x_i \in B_x. \end{cases}$$

Note that the order of the points in the notation $|x_{i-1}, x_i|$ does matter; the second case in this definition, for instance, occurs when the first point of the pair lies in B_x and the second point lies in B_y, while the third case occurs when the opposite holds. If all of these conditions hold, we call $\{x, y\}$ an ***essential gap***. Now, define $\mathscr{G}(\gamma; x, y, \epsilon) := \sum_{i=1}^{n} |x_{i-1}, x_i|$. We call this the (x, y, ϵ)***-gap number of*** γ; it measures the net number of times γ crosses the x, y-gap.

Lemma 4.1.1 *Assume that the above conditions hold for some $\epsilon^* > l$, so that $\{x, y\}$ is an essential gap. Given ϵ such that $l < \epsilon \leq \epsilon^*$, the integer $\mathscr{G}(\gamma; x, y, \epsilon)$ is an ϵ-homotopy invariant. That is, for fixed $\epsilon \in (l, \epsilon^*]$, if α and γ are ϵ-chains such that $\alpha \sim_\epsilon \gamma$, then $\mathscr{G}(\gamma; x, y, \epsilon) = \mathscr{G}(\alpha; x, y, \epsilon)$.*

Proof Since any ϵ-homotopy taking γ to α will consist of a finite sequence of basic moves, it suffices to prove the result in the case where α is obtained by adding or removing a single point to/from γ. The proof is not difficult, but it is a tedious process in working through all the possible cases. We will prove one case to illustrate the reasoning used. The rest of the cases follow in exactly the same manner.

Let $\gamma = \{x_0, x_1, \ldots, x_n\}$, and assume that α is obtained by adding z between x_{i-1} and x_i. Since this basic move only affects three different pairs of points in the sums defining the (x, y, ϵ)-gap numbers of γ and α, we only need to show that $|x_{i-1}, x_i| = |x_{i-1}, z| + |z, x_i|$.

(**Case:** $|x_{i-1}, x_i| = 0$.) If $|x_{i-1}, z| = |z, x_i| = 0$, then the result is clear. If $|x_{i-1}, z| = 1$ and $|z, x_i| = -1$ (or $|x_{i-1}, z| = -1$ and $|z, x_i| = 1$), the result is also clear. The subcase $|x_{i-1}, z| = 1 = |z, x_i|$ cannot occur, for the first equality would imply that $x_{i-1} \in B_x$ and $z \in B_y$, while the second would imply that $z \in B_x$ and $x_i \in B_y$, which would further imply that $z \in B_x \cap B_y$, a contradiction. The case $|x_{i-1}, z| = -1 = |z, x_i|$ also cannot occur, for the first equality would imply $x_{i-1} \in B_y$ and $z \in B_x$, while the second would imply that $z \in B_y$ and $x_i \in B_x$, another contradiction. Suppose $|x_{i-1}, z| = 1$ and $|z, x_i| = 0$. Then $x_{i-1} \in B_x$, $z \in B_y$, and x_i cannot be in B_x or B_y (or else we would have $|x_{i-1}, x_i| = 1$ in the latter case and $|z, x_i| = -1$ in the former). But x_i must lie in Z or Y, and x_i is strictly within ϵ of x_{i-1}, a point in B_x, and strictly within ϵ of z, a point in B_y. If $x_i \in Z$, then, since $d(x_i, z) < \epsilon$, condition 2 above implies that $x_i \in B_x$, a contradiction. If $x_i \in Y$, then $d(x_i, x_{i-1}) < \epsilon$ implies that $x_i \in B_y$, another contradiction. Similar reasoning applies to the cases $|x_{i-1}, z| = -1$ and $|z, x_i| = 0$, $|x_{i-1}, z| = 0$ and $|z, x_i| = 1$, and $|x_{i-1}, z| = 0$ and $|z, x_i| = -1$. Thus, given that $|x_{i-1}, x_i| = 0$, the only possible cases that can occur result in the equality $|x_{i-1}, x_i| = |x_{i-1}, z| + |z, x_i|$.

Proceeding, one would argue simliarly for the cases $|x_{i-1}, x_i| = \pm 1$, and then work through the same procedure in the case where a point is removed from γ to obtain α. All cases that can occur lead to the desired equality, thus proving the result. ∎

This is a rather technical set-up and result, and it is probably not yet clear why we call $\{x, y\}$ an essential gap. It turns out that essential gaps yield refinement critical values. The following example can be taken as a sort of canonical example illustrating this concept.

Example 4.1.2 *Let L, l_1, l_2, and h be positive real numbers such that L is significantly larger than l_1 (say, $L > 3l_1$), $l_2 \leq l_1$, and $h^2 + (l_1 + l_2)^2/4 > l_1^2$. Let X be the metric subspace of $\mathbb{R}^2$ shown in Figure 4.1, and let x, y, u, and v be the points $((L - l_1)/2, h)$, $((L + l_1)/2, h)$, $((L - l_2)/2, 0)$, and $((L + l_2)/2, 0)$, respectively. Let d be the length of the diagonal from x to v (or y to u, by symmetry). The condition $h^2 + (l_1 + l_2)^2/4 > l_1^2$ implies that $d > l_1$. Let ϵ^* be such that $l_1 < \epsilon^* < \min\{d, 2l_1, (L - l_1)/2\}$. Now, fix any ϵ such that $l_1 < \epsilon \leq \epsilon^*$. Let Z be the left half of X, and let Y be the right half. The balls $B_x := B(x, \epsilon - l_1)$ and $B_y := B(y, \epsilon - l_1)$ are just the segments shown below. The conditions $\epsilon < 2l_1$ and $\epsilon < (L - l_2)/2$ ensure that these balls do not intersect or extend to the vertical sides of X.*

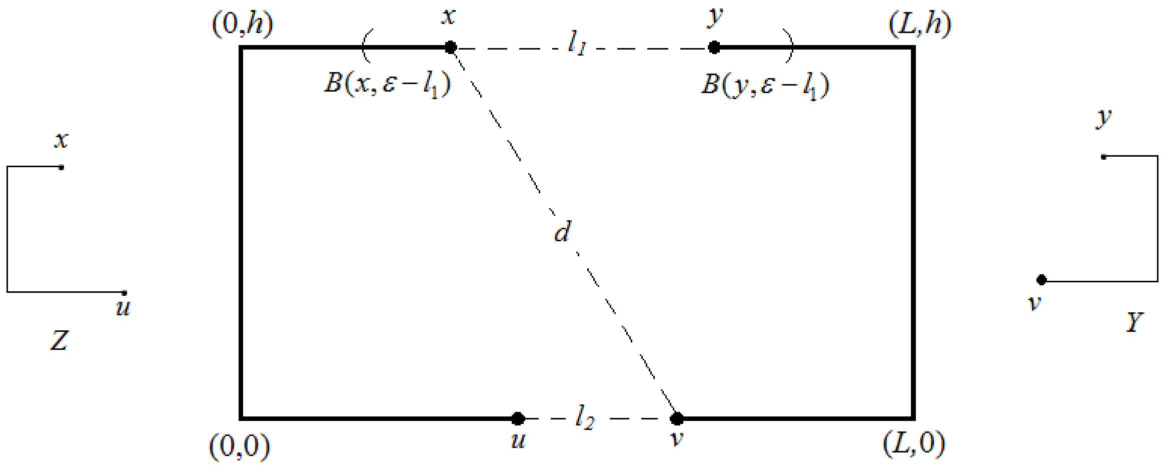

Figure 4.1: An essential gap.

Suppose $z \in Z$ and z is strictly within ϵ of a point in B_y. Clearly, z cannot lie on the vertical segment of Z. If z lies on the lower boundary of Z, then the closest it could be to any point in B_y is d, which occurs when $z = u$. But $d > \epsilon^* \geq \epsilon$, so z cannot, in fact, lie on this lower boundary. If z lies on the upper boundary of Z but outside of B_x, then it is at least $\epsilon - l_1 + l_1 = \epsilon$ away from any point of B_y. Thus, it must hold that z lies in B_x. Likewise, by symmetry, if $z \in Y$ and is strictly within ϵ of a point in B_x, then $z \in B_y$. Therefore, $\{x, y\}$ is an essential gap. What does this mean in terms of critical values? For $l_1 < \epsilon \leq \epsilon^*$, $\gamma := \{x, y\}$ is an ϵ-chain, and $\mathscr{G}(\gamma; x, y, \epsilon) = 1$. It follows from the ϵ-homotopy invariance of the (x, y, ϵ)-gap number that γ is not ϵ-homotopic to any l_1-chain, since an l_1-chain cannot cross the x, y-gap and, therefore, must have (x, y, ϵ)-gap number 0. Since this holds for all $\epsilon \in (l_1, \epsilon^*)$, we see that l_1 is a refinement critical value, or, more specifically, an upper non-surjective critical value.

To see intuitively what makes this essential gap phenomenon occur, consider the trapezoid $\{x, u, v, y\}$. The diagonals of this trapezoid are longer than the longest base of the trapezoid. This, essentially, is why $\{x, y\}$ cannot be ϵ-refined to an l_1-chain for $\epsilon \in (l_1, \epsilon^*)$. One cannot "jump down" from y to v, because $\epsilon^* < d$; this would violate the ϵ-chain requirement. Likewise, one cannot jump from x to u for the same reason. In other words, because the diagonal is too long, one cannot overcome the $\{x, y\}$-gap by going around it, at least via "hops" that are sufficiently close to l_1 in length. This is why we call $\{x, y\}$ an essential gap. Note that if d were less than or equal to l_1, then we could, in fact, go around the $\{x, y\}$-gap. Indeed, for any $\epsilon > l_1$, we could then transform the ϵ-chain, $\{x, y\}$, via ϵ-homotopy by adding v and then u. So, it is the diagonal length that makes this gap essential.

Finally, X is obviously not connected, but we could clearly attach a long joining curve to X to make it path-connected. Moreover, we could do so without affecting the critical value. Thus, one should certainly not think that it is the disconnectivity of X that results in this critical value. ∎

The previous example illustrates the following result, to which we have already alluded.

Lemma 4.1.3 (Essential Gap Lemma) *Let X be a chain-connected metric space, and suppose $\{x, y\}$ is an essential gap. If $dist(B(x,r), B(y,r)) = d(x,y)$ for all sufficiently small r, then $l := d(x,y)$ is a refinement critical value.*

Note that the condition $dist(B(x,r), B(y,r)) = d(x,y)$ is satisfied in the previous example for all sufficiently small r.

Proof Let $\epsilon^* > l$, Z, and Y be as in the definition of an essential gap, and we may assume that $\epsilon^* - l$ is small enough that $dist(B(x,r), B(y,r)) = d(x,y)$ for all $r \leq \epsilon^* - l$. Fix ϵ so that $l < \epsilon \leq \epsilon^*$. Then $\gamma := \{x, y\}$ is an ϵ-chain, and $\mathscr{G}(\gamma; x, y, \epsilon) = 1$. No l-chain can cross the x, y-gap. In fact, if $\{z_0, \ldots, z_n\}$ is an l-chain, and if we had $z_{i-1} \in B_x = B(x, \epsilon - l)$ and $z_{i-1} \in B_y = B(y, \epsilon - l)$, then we would have $dist(B_x, B_y) \leq d(z_{i-1}, z_i) < l$, contradicting the fact that $dist(B_x, B_y) = d(x,y) = l$. Thus, the (x, y, ϵ)-gap number of any l-chain must be 0. The ϵ-homotopy invariance of this value then implies that $\{x, y\}$ is not ϵ-homotopic to an l-chain. Since $\epsilon \in (l, \epsilon^*)$ was arbitrary, it follows that l is a refinement critical value. ∎

4.2 Variations on a Theme

Here we will use the Essential Gap Lemma to produce several examples of metric spaces having critical spectra with positive limit points. Moreover, these examples will show that critical values of one type (i.e. homotopy or refinement critical values) can converge to critical values of the other type.

Example 4.2.1 *We define the following sets.*

For $n \geq 0$, $A_n = \{(x,y) \in \mathbb{R}^2 : 0 \leq x \leq 1, y = 1/2^n\} \cup \{(x,y) \in \mathbb{R}^2 : 2 \leq x \leq 3, y = 1/2^n\}$.

$A_\infty = \{(x,y) \in \mathbb{R}^2 \, : \, 0 \leq x \leq 1, \; y = 0\} \cup \{(x,y) \in \mathbb{R}^2 \, : \, 2 \leq x \leq 3, \; y = 0\}$.

$B_1 = \{(x,y) \in \mathbb{R}^2 \, : \, x = 0, \; 0 \leq y \leq 2\}$.

$B_2 = \{(x,y) \in \mathbb{R}^2 \, : \, x = 3, \; 0 \leq y \leq 2\}$.

$C = \{(x,y) \in \mathbb{R}^2 \, : \, 0 \leq x \leq 3, \; y = 2\}$.

Define a metric subspace of $\mathbb{R}^2$ by

$$X = \Big(\bigcup_{n=0}^{\infty} A_n \Big) \cup A_\infty \cup B_1 \cup B_2 \cup C.$$

For $n \geq 0$, let $x_n = (1, 1/2^n)$ and $y_n = (2, 1/2^n)$, and let $x_\infty = (1, 0)$, $y_\infty = (2, 0)$, $z_0 = (3/2, 2)$. Let $d_0 = d(x_0, z_0)$, and, for $n \geq 1$, let $d_n = d(x_{n-1}, y_n)$. Note that $d_0 = d_1$. For $m > n \geq 0$, let $d^n_m = d(x_n, y_m)$, and note that $d^{n-1}_n = d_n$ for $n \geq 1$. See Figure 4.2 below. We call X "Rapunzel's Comb," and this space is a variation of a construction originally carried out by Maria Walpole as part of The University of Tennessee's 2009 Research Experience for Undergraduates (REU) program in mathematics (see [4]).

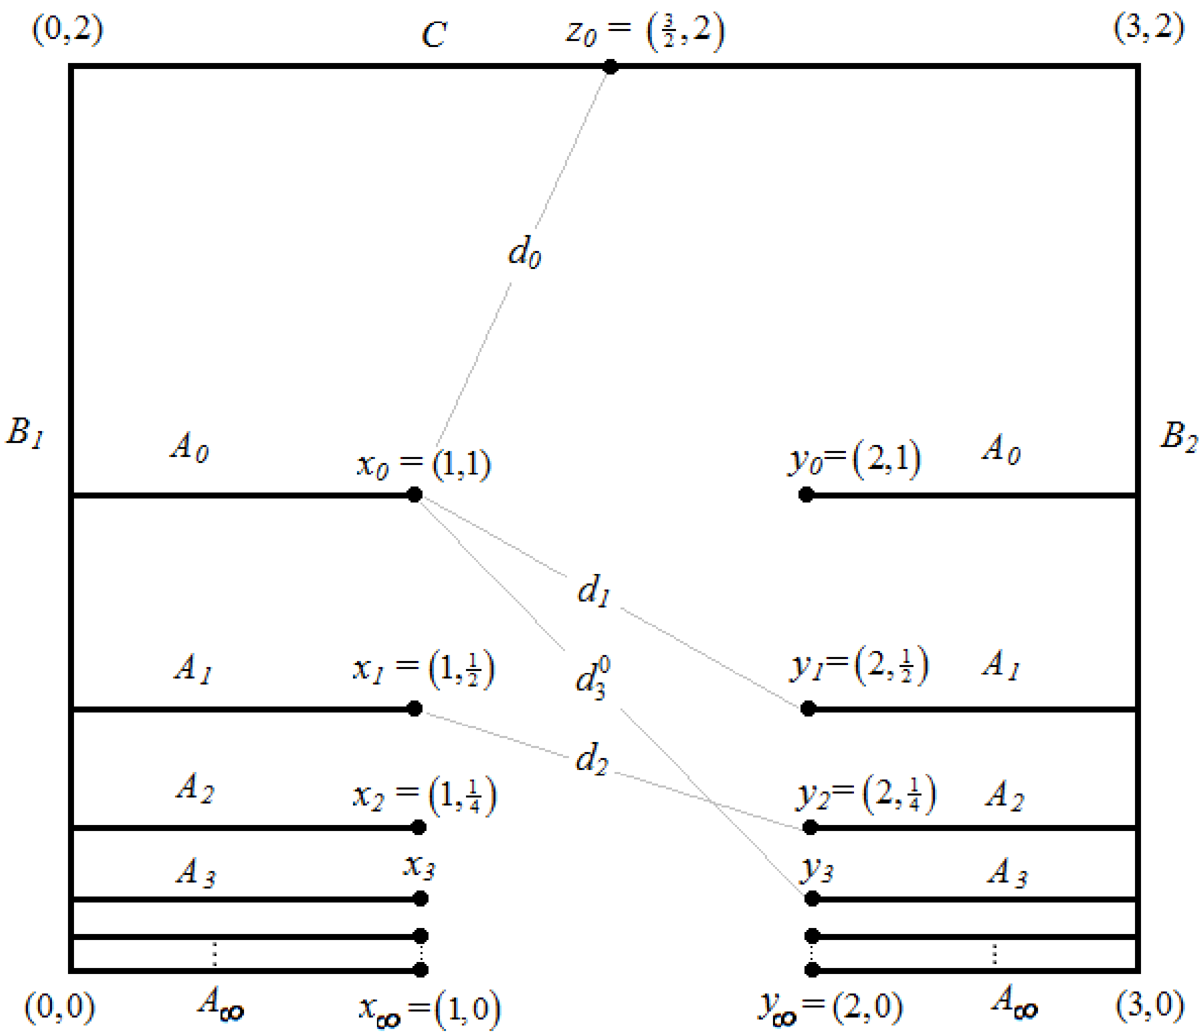

Figure 4.2: Rapunzel's Comb

The following results can easily be verified by direct computation.

1) $1 < d_n < d_{n-1}$ $\forall\, n \geq 1$, *and* $d_n \searrow 1$ *as* $n \to \infty$.

2) $d_m^n > 1$ $\forall m > n \geq 0$, *and, for fixed* n, d_m^n *is minimized when* $m = n + 1$.

Now, fix $n \geq 1$. *It is evident that* $dist(B(x_n, r), B(y_n, r)) = d(x_n, y_n)$ *for sufficiently small* r. *We also claim that* $\{x_n, y_n\}$ *is an essential gap. Fix any* ϵ *such that* $d(x_n, y_n) = 1 < \epsilon \leq d_{n+1}$. *Using the obvious symmetry of* X, *let* Z *be the left half of* X, *including the left halves of* C *and each* A_n, *as well as* B_1, *and assume that* Z *includes* z_0 *(so* Z *is closed). Let* Y *be the rest of the space (so* Y *is open). Let* $B_{x_n} = B(x_n, \epsilon - 1)$ *and* $B_{y_n} = B(y_n, \epsilon - 1)$, *so that* $B_{x_n} \subset Z$ *and* $B_{y_n} \subset Y$. *Suppose that* $z \in Z$ *and* z *lies strictly within* ϵ *of a point in* B_{y_n}. *Clearly* z *cannot lie in* B_1. *Moreover,* z *cannot lie in* C. *In fact, the closest any point of* $C \cap Z$ *can be to any point of* B_{y_n} *is the distance from* z_0 *to* y_n, *which is greater than* d_0. *But* $d_0 = d_1 > d_{n+1} \geq \epsilon$, *so* z *cannot lie in* C. *Thus,* z *must lie on the left half of one of the sets,* A_k. *However,* z *cannot be in* A_k *for* $0 \leq k < n$, *since - in that case - the distance between* z *and any point of* B_{y_n} *would be at least* d_n^k, *which, in turn, is at least as great as* $d_n^{n-1} = d_n$. *Since* $d_n > d_{n+1} \geq \epsilon$, *this shows that this case cannot occur. Hence,* z *must lie on* A_m *for some* $m \geq n$. *But if* z *were in* A_m *for* $m > n$, *the distance between* z *and any point of* B_{y_n} *would be at least* d_m^n. *For fixed* n, d_m^n *is minimized when* $m = n + 1$, *so the distance between* z *and any point of* B_{y_n} *is at least* $d_{n+1}^n = d_{n+1} \geq \epsilon$. *This contradicts that* z *is strictly within* ϵ *of a point of* B_{y_n}. *Therefore,* z *must lie in* A_n, *and since it is within* ϵ *of a point of* B_{y_n}, *it must lie in* B_{x_n}. *By symmetry of* X, *the same argument holds for* Y: *if* $z \in Y$ *and is strictly within* ϵ *of a point of* B_{x_n}, *then it must lie in* B_{y_n}. *Hence,* $\{x_n, y_n\}$ *is an essential gap, and the Essential Gap Lemma now yields the desired conclusion. More specifically, for all* ϵ *such that* $1 < \epsilon \leq d_{n+1}$, $\{x_n, y_n\}$ *is an* ϵ*-chain that cannot be* ϵ*-refined to a* 1*-chain.*

Finally, fix $n \geq 1$, *and consider the loop* $\gamma_n = \{x_n, x_{n+1}, y_{n+1}, y_n, x_n\}$. *For* $1 < \epsilon \leq d_{n+1}$, *this is an* ϵ*-loop. Moreover, its* (x_n, y_n, ϵ)*-gap number is* -1. *An* ϵ*-null loop would be* ϵ*-homotopic to a loop that does not cross the* x_n, y_n*-gap, and, so, the* ϵ*-homotopy invariance of the* (x_n, y_n, ϵ)*-gap number implies that* γ_n *is* ϵ*-nontrivial. This holds for all* $1 < \epsilon \leq d_{n+1}$. *However, for* $\epsilon > d_{n+1}$, γ_n *is trivial. Indeed, in the rectangle formed by* γ_n, *the diagonals have length* d_{n+1}, *so, when* $\epsilon > d_{n+1}$, *we can successively remove* y_{n+1}, x_{n+1}, *and* y_n, *in that order, giving us the trivial loop at* x_n. *This implies that* d_{n+1} *is an upper non-injective critical value of* X, *since there is a non-trivial* d_{n+1}*-loop that is* ϵ*-null for all* $\epsilon > d_{n+1}$. *Since* $d_n \searrow 1$, *and since the critical spectrum contains all of its positive limit points, we have a sequence of critical values converging down to a critical value. The critical value,* 1, *however, is of a different nature than the critical values* d_n. *In fact, there are no non-trivial* 1*-loops in* X. *So,* 1 *is an upper non-surjective critical value. Of course, we already knew this from showing that each* $\{x_n, y_n\}$ *is an essential gap. Note, also, that* X *is compact, path-connected, and simply-connected in the traditional sense.* ∎

There are many different variations on Rapunzel's Comb that one can use to illustrate critical value limiting behavior. All of them use the Essential Gap Lemma in some form, and the details follow in much the same manner as in the previous example. So, we will just briefly exhibit two more examples along these lines.

Example 4.2.2 *For* $n \geq 1$, *let* $h_n = 2^{-n/2} = 1/(\sqrt{2})^n$, *and let*

$$H = \sum_{n=1}^{\infty} h_n = \sum_{n=1}^{\infty} \Big(\frac{1}{\sqrt{2}}\Big)^n = \frac{1}{\sqrt{2} - 1} = 1 + \sqrt{2}.$$

Define the following sets.

For $n \geq 1$, $A_n = \{(x,y) \in \mathbb{R}^2 : 0 \leq x \leq 1 + \frac{1}{2^{n+1}},\, y = \sum_{i=1}^{n-1} h_i\} \cup \{(x,y) \in \mathbb{R}^2 : 2 - \frac{1}{2^{n+1}} \leq x \leq 3,\, y = \sum_{i=1}^{n-1} h_i\}$*; if* $n = 1$*, we define* $\sum_{i=1}^{n-1} h_i$ *to be* 0.

$A_\infty = \{(x,y) \in \mathbb{R}^2 : 0 \leq x \leq 1,\, y = H\} \cup \{(x,y) \in \mathbb{R}^2 : 2 \leq x \leq 3,\, y = H\}$.

$B_1 = \{(x,y) \in \mathbb{R}^2 : x = 0,\, 0 \leq y \leq H + 2\}$.

$B_2 = \{(x,y) \in \mathbb{R}^2 : x = 3,\, 0 \leq y \leq H + 2\}$.

$C = \{(x,y) \in \mathbb{R}^2 : 0 \leq x \leq 3,\, y = H + 2\}$.

Let X *be the metric subspace of* $\mathbb{R}^2$ *defined by*

$$X = \Big(\bigcup_{n=1}^{\infty} A_n\Big) \cup A_\infty \cup B_1 \cup B_2 \cup C.$$

We also define the following:

$$x_n = \left(1 + \frac{1}{2^{n+1}}, \sum_{i=1}^{n-1} h_i\right),\ \ y_n = \left(2 - \frac{1}{2^{n+1}}, \sum_{i=1}^{n-1} h_i\right),\ \forall\, n \geq 1,$$

$$z_0 = \Big(\frac{3}{2}, H + 2\Big),\quad x_\infty = (1, H),\quad y_\infty = (2, H),$$
$$d_0 = d(x_\infty, z_0) = d(y_\infty, z_0),\ \ d_n = d(x_n, y_{n+1})\ \forall\, n \geq 1.$$

See Figure 4.3 on the following page. Here, the gaps increase in length to a gap of length 1.

Reasoning as in the previous example, one can show that each $\{x_n, y_n\}$*, for* $n \geq 1$*, is an essential gap. In fact, it is straightforward to show that* $d_n > 1 - \frac{1}{2^{n+1}}$*, so the diagonal from* x_n *to* y_{n+1} *(or* y_n *to* x_{n+1}*) is greater than the gap above it. We also clearly have that* $dist(B(x_n, r), B(y_n, r)) = d(x_n, y_n)$ *for sufficiently small* r*. Thus, for each* n *and all* ϵ *greater than but sufficiently close to* $1 - \frac{1}{2^n} = d(x_n, y_n)$*,* $\{x_n, y_n\}$ *is an* ϵ*-chain that cannot be* ϵ*-refined to a* $(1 - \frac{1}{2^n})$*-chain. Hence, for each* $n \geq 1$*,* $1 - \frac{1}{2^n}$ *is an upper non-surjective critical value. These values converge up to* 1*, but* 1 *is, in fact, not an upper non-surjective critical value. Since, for each* n*, the diagonals between* x_n *and* y_{n+1} *(and between* x_{n+1} *and* y_n*) are strictly less than* 1 *in length, it turns out that, for every* $\epsilon > 1$*, every* ϵ*-chain can be* ϵ*-refined to a* 1*-chain. But positive limit points of critical values are critical values. So, what type of critical value is* 1*?*

Fix $n \geq 2$*, and let* γ_n *be the loop* $\{x_n, x_{n-1}, y_{n-1}, y_n, x_n\}$*. For all* ϵ *greater than* $1 - \frac{1}{2^n}$*,* γ_n *is an* ϵ*-loop. Since* $\{x_n, y_n\}$ *is an essential gap, we also know that, for each* ϵ *greater than but sufficiently close to* $1 - \frac{1}{2^n}$*, the* (x_n, y_n, ϵ)*-gap number is an* ϵ*-homotopy invariant. Fixing any such* ϵ*, we see that the* ϵ*-chain* $\alpha_n := \{x_n, y_n\}$ *has non-zero* (x_n, y_n, ϵ)*-gap number, while the* ϵ*-chain* $\beta_n := \{x_n, x_{n-1}, y_{n-1}, y_n\}$ *does not cross the* x_n, y_n*-gap at all. Hence,* γ_n *cannot be* ϵ*-null for such an* ϵ*-value, or else we would have* $\alpha_n \sim_\epsilon \beta_n$*, further implying that these chains have equal* (x_n, y_n, ϵ)*-gap numbers, a contradiction. So,* γ_n *is* ϵ*-nontrivial for all* ϵ *greater than but sufficiently close to* $1 - \frac{1}{2^n}$*. On the other hand,* γ_n *is* 1*-null. In fact, since the diagonals between* x_n *and* y_{n-1} *are less than* 1 *in length, we can successively remove* y_{n-1}*,* x_{n-1}*, and* y_n*, giving us the trivial chain. Therefore, the map* $\varphi_{1,1-1/2^n} : X_{1-1/2^n} \to X_1$ *is non-injective for all* n*, showing that* 1 *is a lower non-injective critical value. Note that, as before,* X *is compact, path-connected, and simply-connected.* ∎

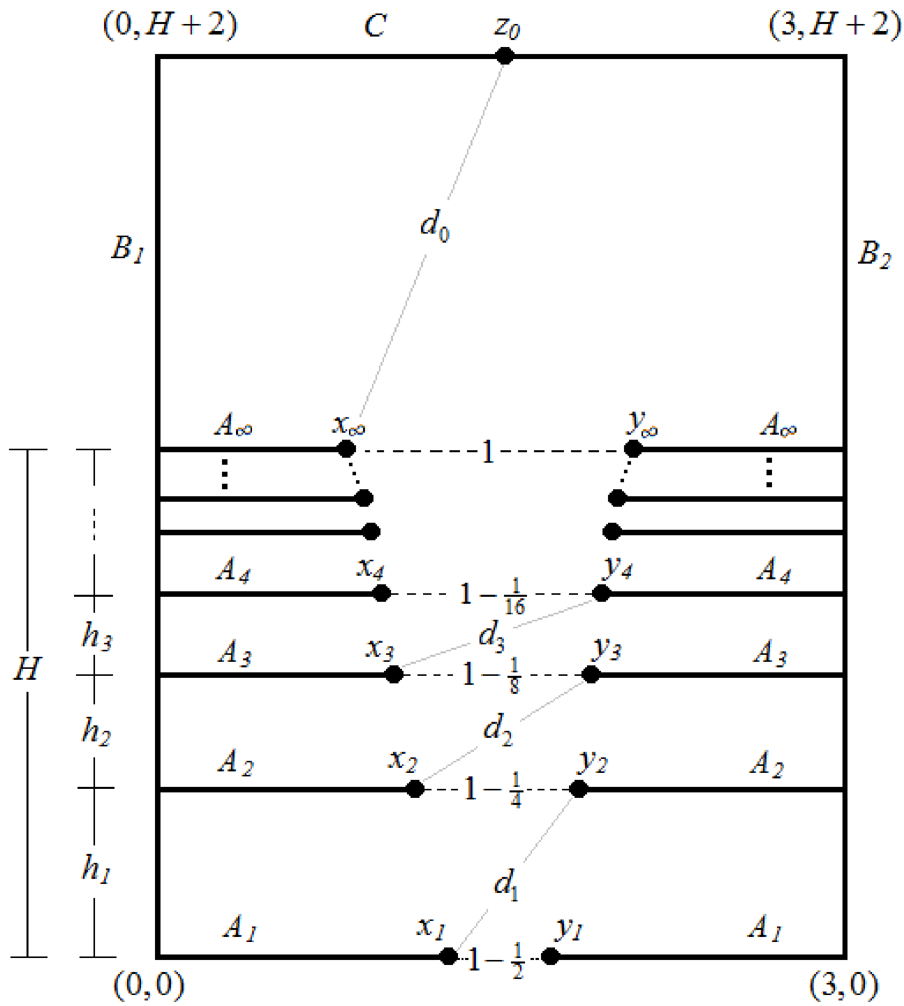

Figure 4.3: Rapunzel's Comb - Variation 1

In the next example, we will construct yet another variation of Rapunzel's Comb in which we have a sequence of upper non-surjective critical values, $\{\epsilon_n\}$, converging up to a *lower* non-surjective critical value, ϵ. In other words, there is no $\delta < \epsilon$ such that we can refine all ϵ-chains to δ-chains. We have already mentioned that the ability to refine chains is important when studying the critical spectrum of a metric space, and results in subsequent chapters will further emphasize this fact. It is natural, then, to wonder how common the *inability* to refine chains actually is. As the next example will show, it takes some intricate calculation to construct a space in which chains cannot be refined to any degree. This should give the reader some sense of the notion that refinability is rather common, even outside the context of geodesic spaces. In fact, the following example can be taken as a sort of canonical example of a space with a lower non-surjective critical value, because - as it turns out - the type of behavior exhibited by the space in this example is *necessary* for such a critical value. That is, ϵ is a lower non-surjective critical value of a metric space X only if there is a sequence of upper non-surjective critical values converging up to ϵ.

Example 4.2.3 (Rapunzel's Comb - Variation 2) *The construction of this example is very similar to the previous case. In fact, the lengths of the gaps will be the same. The key difference will be changing the heights between the gaps in the comb. We want to increase them slightly enough so that the diagonal lengths between x_n and y_{n+1} are greater than 1 for each n, but still small enough so that the sum of the heights is finite.*

So, for $n \geq 1$, let $h_n = \sqrt{3}/(\sqrt{2})^n$, and let

$$H = \sum_{n=1}^{\infty} h_n = \sqrt{3} \sum_{n=1}^{\infty} \Big(\frac{1}{\sqrt{2}} \Big)^n = \sqrt{3} + \sqrt{6}.$$

We define X exactly as in the previous example except for the different values h_n, and we similarly define the points x_n, y_n, x_∞, y_∞, and z_0. Also as before, we let $d_n = d(x_n, y_{n+1})$, so that d_n is the length of the diagonal between x_n and y_{n+1}. Since the construction is the same, Figure 4.3 holds equally well for this example. We just need to keep in mind that the heights, h_n, and, therefore, the diagonals, d_n, are larger in this case. One can verify by direct computation the following:

$$d_n^2 = 1 + \frac{3}{2^{n+1}} + \frac{9}{2^{2n+4}} \text{ and } d_{n+1} < d_n \, \forall \, n \Rightarrow d_n \searrow 1,$$

$$d(x_n, y_m) > 1 \text{ for all } 1 \leq n < m.$$

In addition, for fixed n and $m > n$, the diagonal lengths, $d(x_n, y_m)$, increase as m increases.

Now, fix $n \geq 2$, and, recalling that $d(x_n, y_n) = 1 - \frac{1}{2^n}$, let ϵ be such that

$$1 - \frac{1}{2^n} < \epsilon \leq \min\Big\{ d(x_1, y_n), \ldots, d(x_{n-1}, y_n), d(x_{n+1}, y_n), 1 + h_n - \frac{1}{2^n} \Big\}.$$

The condition that ϵ be less than or equal to $1 + h_n - \frac{1}{2^n}$ is to ensure that the ball of radius $\epsilon - (1 - \frac{1}{2^n})$ centered at x_n (or y_n) does not intersect any nearby teeth of the comb or either of the vertical sides of X. That is, these balls are just segments of the teeth of the comb formed by A_n. As before, we let Z be the left half of X, and we let Y be the right half. Suppose $z \in Z$ and z lies within ϵ of a point of $B_{y_n} := B(y_n, \epsilon - (1 - \frac{1}{2^n}))$. Clearly, z cannot lie in C or B_1. If z were in A_m for some $m \leq n-1$, the distance between z and any point of B_{y_n} would be at least $d(x_m, y_n)$. But $\epsilon \leq d(x_m, y_n)$ for such m, so this cannot occur. If z were in A_{n+1}, then the closest z could be to any point of B_{y_n} is $d_n = d(x_{n+1}, y_n)$, but, again, we have $\epsilon \leq d(x_{n+1}, y_n)$.

So, this cannot occur either. Niether can z be in A_m for $m > n+1$, since the diagonal lengths, $d(x_m, y_n)$, are greater than d_n for $m > n+1$. Hence, z must lie in A_n, and, in fact, it must lie in $B(x_n, \epsilon - (1 - \frac{1}{2^n}))$. By symmetry, the same result holds if $z \in Y$ and lies within ϵ of a point in $B(x_n, \epsilon - (1 - \frac{1}{2^n}))$. Thus, each $\{x_n, y_n\}$ is an essential gap, and we also have that $dist(B(x_n, r), B(y_n, r)) = d(x_n, y_n)$ for all sufficiently small r. It follows that $1 - \frac{1}{2^n} = d(x_n, y_n)$ is an upper non-surjective critical value; for all ϵ greater than but sufficiently close to $1 - \frac{1}{2^n}$, $\{x_n, y_n\}$ is an ϵ-chain that cannot be ϵ-refined to a $(1 - \frac{1}{2^n})$-chain.

Finally, since $1 - \frac{1}{2^n} \nearrow 1$, we know that 1 is a critical value. Note that

$$1 < \min\Big\{d(x_1, y_n), \ldots, d(x_{n-1}, y_n), d(x_{n+1}, y_n), 1 + h_n - \frac{1}{2^n}\Big\},$$

because all diagonals have length greater than 1 and

$$h_n > \frac{1}{(\sqrt{2})^n} > \frac{1}{2^n} \Rightarrow 1 + h_n - \frac{1}{2^n} > 1.$$

Thus, for $\epsilon = 1$, the (x_n, y_n, ϵ)-gap number of an ϵ-chain is an ϵ-homotopy invariant. Now, $\{x_n, y_n\}$ is a 1-chain, and its $(x_n, y_n, 1)$-gap number is 1. However, no $(1 - \frac{1}{2^n})$-chain can cross the x_n, y_n-gap. Thus, $\{x_n, y_n\}$ cannot be 1-homotopic to a $(1 - \frac{1}{2^n})$-chain. In other words, the map $\varphi_{1,1-1/2^n} : X_{1-1/2^n} \to X_1$ is not surjective, and this holds for all $n \geq 1$. Hence, 1 is a lower non-surjective critical value. ∎

4.3 A Space with a Dense Critical Spectrum

Now we have seen examples of compact, path-connected metric spaces with non-discrete critical spectra. This, of course, distinguishes the critical spectrum of a general compact metric space from the covering spectrum of a compact geodesic space, which is discrete in $\mathbb{R}^+$. Furthermore, this brings up the question of the extent to which the critical spectrum can differ in its structure from the covering spectrum. This entire section will be devoted to constructing an example showing that the difference can, indeed, be extreme. We will construct a metric space, X, with the property that every dyadic rational between 0 and 1 is a critical value of X. Since these numbers are dense in the interval $(0, 1)$, and since positive limit points of critical values are critical values, it will follow that the critical spectrum of X contains the entire interval $(0, 1]$.

The construction of this space is technical, and it requires some significant computation. We will actually construct X inductively by constructing a sequence of spaces, X_n, such that X_n converges to X in the Gromov-Hausdorff sense. Each X_n will be a subspace of separable Hilbert space, and, therefore, X will, also. Moreover, each X_n will isometrically imbed into X_{n+1}, and the critical values of each X_n will be maintained as we construct each subsequent space in the sequence. The final space, X, will actually be defined as the closure of the limit of the spaces, X_n. Hence, we need to know that the critical spectrum is not altered by taking the closure of a possibly non-compact (but still bounded) metric space. So, we will first prove that if S is a dense subset of a metric space, X, then the critical spectra of S and X are equal.

Let X be a chain-connected metric space, and let $S \subset X$ be dense. It is easy to see that this implies that S is chain-connected, also.

Lemma 4.3.1 *If α and β are ϵ-chains in S that are ϵ-homotopic in X, then they are also ϵ-homotopic in S.*

Proof Let $H = \{\alpha = \gamma_0, \ldots, \gamma_N = \beta\}$ be an ϵ-homotopy in X between α and β. Denote γ_k by $\gamma_k = \{x_0^k, \ldots, x_{m_k}^k\}$, $k = 0, \ldots, N$, and let

$$\tau = \max_{\substack{0 \leq k \leq N \\ 1 \leq i \leq m_k}} d(x_{i-1}^k, x_i^k).$$

Then $\tau < \epsilon$, and $\tau + (\epsilon - \tau)/2 < \epsilon$.

For each $k = 0, \ldots, N$, we define a new chain, $\gamma_k' = \{\bar{x}_0^k, \ldots, \bar{x}_{m_k}^k\}$, as follows. For $k = 0$ and $0 \leq i \leq m_0$, if $x_i^0 \in S$, we let $\bar{x}_i^0 = x_i^0$. If $x_i^0 \notin S$, choose $\bar{x}_i^0 \in S$ so that $d(x_i^0, \bar{x}_i^0) < (\epsilon - \tau)/4$. Then γ_0' is an ϵ-chain, since

$$d(\bar{x}_{i-1}^0, \bar{x}_i^0) \leq d(\bar{x}_{i-1}^0, x_{i-1}^0) + d(x_{i-1}^0, x_i^0) + d(x_i^0, \bar{x}_i^0) < \tau + 2\frac{\epsilon - \tau}{4} < \epsilon.$$

Now, γ_1 differs from γ_0 by the addition or removal of a single point. If γ_1 is obtained by removing a point from γ_0, say x_i^0, then we set $\gamma_1' = \{\bar{x}_0^0, \ldots, \bar{x}_{i-1}^0, \bar{x}_{i+1}^0, \ldots, \bar{x}_{m_0}^0\}$. If γ_1 is obtained by adding a point, x, into γ_0 between x_i and x_{i+1}, then we set $\gamma_1' = \{\bar{x}_0^0, \ldots, \bar{x}_i^0, \bar{x}, \bar{x}_{i+1}^0, \ldots, \bar{x}_{m_0}^0\}$, where $\bar{x} = x$ if $x \in S$ and - if $x \notin S$ - $\bar{x}$ is a point in S such that $d(x, \bar{x}) < (\epsilon - \tau)/4$. In either case, γ_1' is an ϵ-chain, since, in the former case, we have

$$d(\bar{x}_{i-1}^0, \bar{x}_{i+1}^0) \leq d(\bar{x}_{i-1}^0, x_{i-1}^0) + d(x_{i-1}^0, x_{i+1}^0) + d(x_{i+1}^0, \bar{x}_{i+1}^0) < \tau + \frac{\epsilon - \tau}{2} < \epsilon,$$

and, in the latter case, we have

$$d(\bar{x}_i^0, \bar{x}) \leq d(\bar{x}_i^0, x_i^0) + d(x_i^0, x) + d(x, \bar{x}) \leq \tau + \frac{\epsilon - \tau}{2} < \epsilon$$

$$d(\bar{x}, \bar{x}_{i+1}^0) \leq d(\bar{x}, x) + d(x, x_{i+1}) + d(x_{i+1}, \bar{x}_{i+1}^0) < \tau + \frac{\epsilon - \tau}{2} < \epsilon.$$

Continuing this process inductively, we finish up when we construct γ_N'. By construction, each chain in the sequence, $H' = \{\gamma_0', \ldots, \gamma_N'\}$, is an ϵ-chain and differs from its predecessor and/or successor by a basic move. Also by construction, each γ_k', $k = 0, \ldots, N$, lies in S. Thus, H' is an ϵ-homotopy in S between γ_0' and γ_N'. But $\alpha = \gamma_0$ and $\beta = \gamma_N$ are in S by hypothesis, and, again by construction, when a point in a chain in H lies in S, we choose that point as its correspondent in H'. Thus, γ_0' must equal γ_0, and γ_N must equal γ_N', showing that H' is an ϵ-homotopy between α and β. ∎

Lemma 4.3.2 *If $\alpha = \{x_0, \ldots, x_n\}$ is an ϵ-chain with endpoints in S, then α is ϵ-homotopic to an ϵ-chain in S.*

Proof The proof is by induction on n, the number of points in α. If α is a two-point ϵ-chain, then, by hypothesis, all of its points lie in S, and the result is trivial. So, assume, for some $n \geq 2$, that every ϵ-chain consisting of n or fewer points and having endpoints in S is ϵ-homotopic to an ϵ-chain in S. Let $\alpha = \{x_0, \ldots, x_n,\}$ be an $(n+1)$-point ϵ-chain with endpoints in S, and let $\tau = \max_{1 \leq i \leq n} d(x_{i-1}, x_i)$. Then $\tau < \epsilon$. Choose $x_{n-1}' \in S$ such that $d(x_{n-1}', x_{n-1}) < (\epsilon - \tau)/2$. Then

$$d(x_n, x_{n-1}') \leq d(x_n, x_{n-1}) + d(x_{n-1}, x_{n-1}') < \tau + \frac{\epsilon - \tau}{2} < \epsilon.$$

So, we can insert x_{n-1}' into α between x_{n-1} and x_n. But

$$d(x_{n-2}, x_{n-1}') \leq d(x_{n-2}, x_{n-1}) + d(x_{n-1}, x_{n-1}') < \tau + \frac{\epsilon - \tau}{2} < \epsilon,$$

so we can then remove x_{n-1}. This gives us an ϵ-homotopy transforming α to the chain $\alpha' := \{x_0, \dots, x_{n-2}, x'_{n-1}, x_n\}$. Now, the chain $\{x_0, \dots, x_{n-2}, x'_{n-1}\}$ is an n-point ϵ-chain with endpoints in S, so, by the inductive hypothesis, it can be transformed via ϵ-homotopy to an ϵ-chain in S. This homotopy will leave x_0 and x'_{n-1} fixed, so it extends to an ϵ-homotopy of α' and, therefore, α. ∎

Now, let $\iota : S \to X$ be the inclusion map of S into X. This map induces a map $\iota_\epsilon : S_\epsilon \to X_\epsilon$ for any $\epsilon > 0$ as follows. We may assume without loss of generality that our base point, $*$, is in S. Let $[\alpha]^S_\epsilon$ denote the equivalence class of an ϵ-chain in S beginning at $*$, and we will continue to let $[\alpha]_\epsilon$ denote the equivalence class of α in X. Given $[\alpha]^S_\epsilon \in S_\epsilon$, define $\iota_\epsilon([\alpha]^S_\epsilon) = [\alpha]_\epsilon$. That is, we just take an equivalence class of chains in S and consider the equivalence class of those chains in X. This is a well-defined map, for if α and β are ϵ-chains in S such that $\alpha \sim_\epsilon \beta$ in S, then clearly they are also ϵ-homotopic in X. This map is also injective. In fact, suppose $\iota_\epsilon([\alpha]^S_\epsilon) = \iota_\epsilon([\beta]^S_\epsilon)$, which implies that α and β are chains in S such that $\alpha \sim_\epsilon \beta$ in X. Then, by Lemma 4.3.1, α and β are ϵ-homotopic in S, also, showing that $[\alpha]^S_\epsilon = [\beta]^S_\epsilon$. Moreover, for any $\delta < \epsilon$, we have the commutativity relation $\varphi_{\epsilon\delta} \circ \iota_\delta = \iota_\epsilon \circ \varphi^S_{\epsilon\delta}$. This follows because if $[\alpha]^S_\delta \in S_\delta$, then

$$\iota_\epsilon\big(\varphi^S_{\epsilon\delta}\big([\alpha]^S_\delta\big)\big) = \iota_\epsilon([\alpha]^S_\epsilon) = [\alpha]_\epsilon = \varphi_{\epsilon\delta}([\alpha]_\delta) = \varphi_{\epsilon\delta}\big(\iota_\delta\big([\alpha]^S_\delta\big)\big).$$

Lemma 4.3.3 *Given $\delta < \epsilon$, $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is injective if and only if $\varphi^S_{\epsilon\delta} : S_\delta \to S_\epsilon$ is injective.*

Proof Suppose $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is injective. Then the map $\varphi_{\epsilon\delta} \circ \iota_\delta$ is injective, and this map equals $\iota_\epsilon \circ \varphi^S_{\epsilon\delta}$. Hence, $\iota_\epsilon \circ \varphi^S_{\epsilon\delta}$ is injective, which implies that $\varphi^S_{\epsilon\delta}$ must be injective, also.

For the other direction, we work with the ϵ and δ-groups. Suppose $\varphi^S_{\epsilon\delta} : S_\delta \to S_\epsilon$ is injective. Then $\Phi^S_{\epsilon\delta} : \pi_\delta(S) \to \pi_\epsilon(S)$ is also injective. Let γ be any δ-loop at $*$ in X that is ϵ-nullhomotopic in X (i.e. $[\gamma]_\delta \in \ker(\Phi_{\epsilon\delta})$). By Lemma 4.3.2, we can transform γ via δ-homotopy to a δ-loop in S. Let γ' denote this δ-loop. Note that γ' is ϵ-null in X, since it is δ-homotopic - and, thus, ϵ-homotopic - to γ. That is, γ' and $\{*\}$ are ϵ-chains in S that are ϵ-homotopic in X. By Lemma 4.3.1, γ' is ϵ-homotopic to $\{*\}$ in S. In other words, $[\gamma']^S_\epsilon$ is the trivial element in $\pi_\epsilon(S)$. But $[\gamma']^S_\epsilon = \Phi^S_{\epsilon\delta}([\gamma']^S_\delta)$, and, since $\Phi^S_{\epsilon\delta}$ is injective, this implies that $[\gamma']^S_\delta$ is also trivial. So, γ' is δ-null in S. It follows that γ' is δ-null in X, also. Finally, since γ is δ-homotopic in X to γ', γ is also δ-null in X, showing that $\Phi_{\epsilon\delta} : \pi_\delta(X) \to \pi_\epsilon(X)$ - and, hence, $\varphi_{\epsilon\delta}$ - is injective. ∎

Lemma 4.3.4 *Given $\delta < \epsilon$, $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is surjective if and only if $\varphi^S_{\epsilon\delta} : S_\delta \to S_\epsilon$ is surjective.*

Proof Suppose $\varphi^S_{\epsilon\delta}$ is surjective. Then every ϵ-chain in S can be ϵ-refined, in S, to a δ-chain. Let $\alpha = \{* = x_0, \dots, x_n\}$ be an ϵ-chain in X beginning at $*$. Let $\tau = \max_{1 \le i \le n} d(x_{i-1}, x_i)$, so that $\tau < \epsilon$. Choose $x'_n \in S$ so that $d(x_n, x'_n) < \min\{\delta, (\epsilon - \tau)/2\}$. Then

$$d(x_{n-1}, x'_n) \le d(x_{n-1}, x_n) + d(x_n, x'_n) < \tau + \frac{\epsilon - \tau}{2} < \epsilon.$$

Thus, we can insert x'_n into α between x_{n-1} and x_n, giving us $\alpha' := \{* = x_0, \dots, x_{n-1}, x'_n, x_n\}$. Now, the chain $\{* = x_0, \dots, x_{n-1}, x'_n\}$ is an ϵ-chain with endpoints in S. By Lemma 4.3.2, this chain is ϵ-homotopic to an ϵ-chain in S, and, by assumption, that resulting chain is, then, ϵ-homotopic to a δ-chain in S. By leaving x'_n and x_n fixed, this yields an ϵ-homotopy taking α to a chain $\{* = y_0, \dots, y_k, x'_n, x_n\}$, where $d(y_{i-1}, y_i) < \delta$ for $i = 1, \dots, k$, $d(y_k, x'_n) < \delta$, and $d(x'_n, x_n) < \delta$. Thus, α is ϵ-homotopic to a δ-chain, showing that $\varphi_{\epsilon\delta}$ is surjective.

Conversely, suppose $\varphi_{\epsilon\delta}$ is surjective. Let $\alpha = \{* = x_0, \ldots, x_n\}$ be an ϵ-chain in S. Then α is an ϵ-chain in X, also, and, by assumption, it is ϵ-homotopic in X to a δ-chain, α'. Since α' is a δ-chain in X with endpoints in S, by Lemma 4.3.2, α' is δ-homotopic - and, therefore, ϵ-homotopic - to a δ-chain with all points in S. Thus, α is ϵ-homotopic in X to a δ chain, β, with $\beta \subset S$. By Lemma 4.3.1, α and β are also ϵ-homotopic in S, showing that α can be ϵ-refined in S to a δ-chain. So, $\varphi^S_{\epsilon,\delta}$ is surjective. ∎

It now follows that the map $\varphi^S_{\epsilon\delta} : S_\delta \to S_\epsilon$ is bijective if and only if $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is bijective. Since non-critical values are defined precisely by intervals over which these maps are bijective, we immediately see that the non-critical values - and, thus, the critical values - of S and X agree. Hence, we obtain the following result.

Theorem 4.3.5 *If S is a dense subset of a metric space, X, then $Cr(X) = Cr(S)$.*

We are now ready to proceed with the construction of a compact metric space, X, with a dense critical spectrum. To facilitate the construction, we will arrange the dyadic rationals from the fractal or binary tree point of view. That is, we take the midpoint of $[0, 1]$, then take the midpoint of each of the resulting two intervals, then the midpoint of each of the resulting four intervals, and so on. So, if we set

$$S_n = \left\{ \frac{2k-1}{2^n} \ : \ 1 \le k \le 2^{n-1} \right\}$$

for $n \ge 1$, then we can express dyadic rationals in $(0, 1)$ by $\cup_{n=1}^{\infty} S_n$. We can express this decomposition with a binary tree as in Figure 4.4. The n^{th} generation of this tree is S_n.

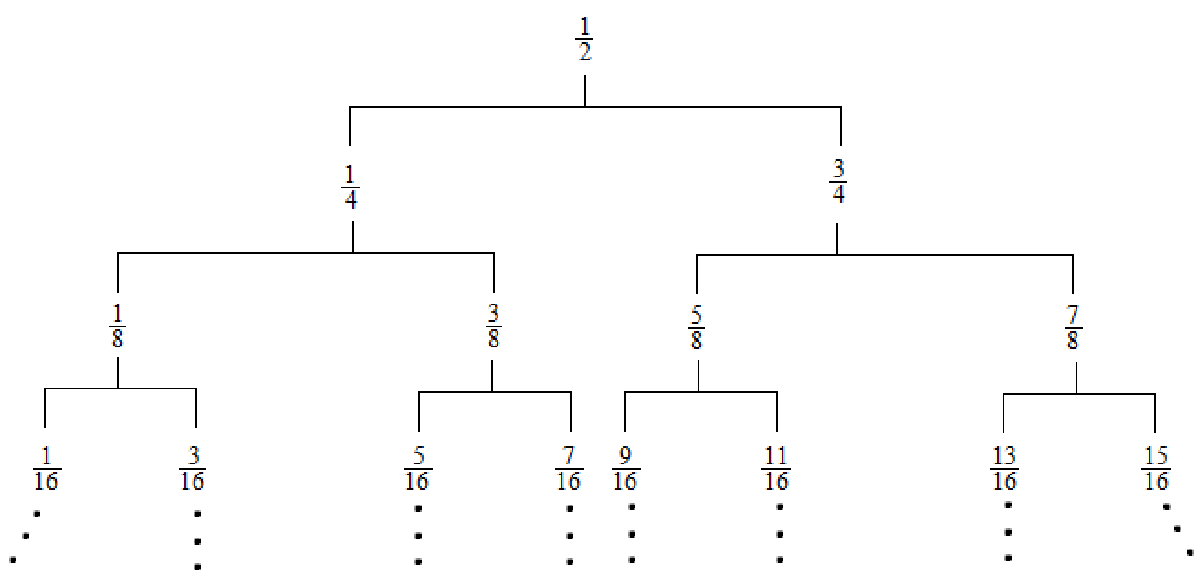

Figure 4.4: Unit Interval Dyadic Rationals

Note that each dyadic rational has two successors in this tree arrangement. Specifically, the successors of the rational $(2k-1)/2^n$ are $(4k-3)/2^{n+1}$ and $(4k-1)/2^{n+1}$. This particular arrangement literally dictates how we will go about constructing the sequence of spaces, $\{X_n\}$, converging to X. To outline the process, we will first construct X_1 with critical values $1/2$, $1/4$, and $3/4$. We will then form X_2 by attaching two new "pieces" to X_1, with one of these spaces adding the critical values $1/8$ and $3/8$, and the other piece adding the critical values $5/8$ and $7/8$. Proceeding, we will then form X_3 by attaching four new pieces to X_2, each piece adding the critical values $1/16$ and $3/16$, $5/16$ and $7/16$, $9/16$ and $11/16$, and $13/16$ and $15/16$, respectively. Continuing, this yields the inductive process by which X is formed.

We start with a real, separable Hilbert space, $\mathscr{H}$, with orthonormal basis, $\{e_n\}_{n=0}^{\infty}$. Each two-dimensional subspace, $span\{e_k, e_j\}$, $k \neq j$, is isometric to $\mathbb{R}^2$, and we will implicitly make this identification without mentioning it from here on. We index the basis beginning with $n = 0$, even though we will construct X in $span\{e_n\}_{n\geq 1}$. The reason for this is that, at the end of the construction, we will come back and add a piece to X in the $\{e_0, e_1\}$ plane to make it path-connected. So, we leave one basis vector alone for that purpose. Now, for $n \geq 1$ and $1 \leq k \leq 2^{n-1}$, consider what we will call *the basic piece* (shown below in Figure 4.5) with the subspace metric inherited from $\mathbb{R}^2$. We denote this space by $A_{n,k}$. Note that the upper and lower gaps of $A_{n,k}$ correspond to the two successors of the central gap as indicated in the tree in Figure 4.4.

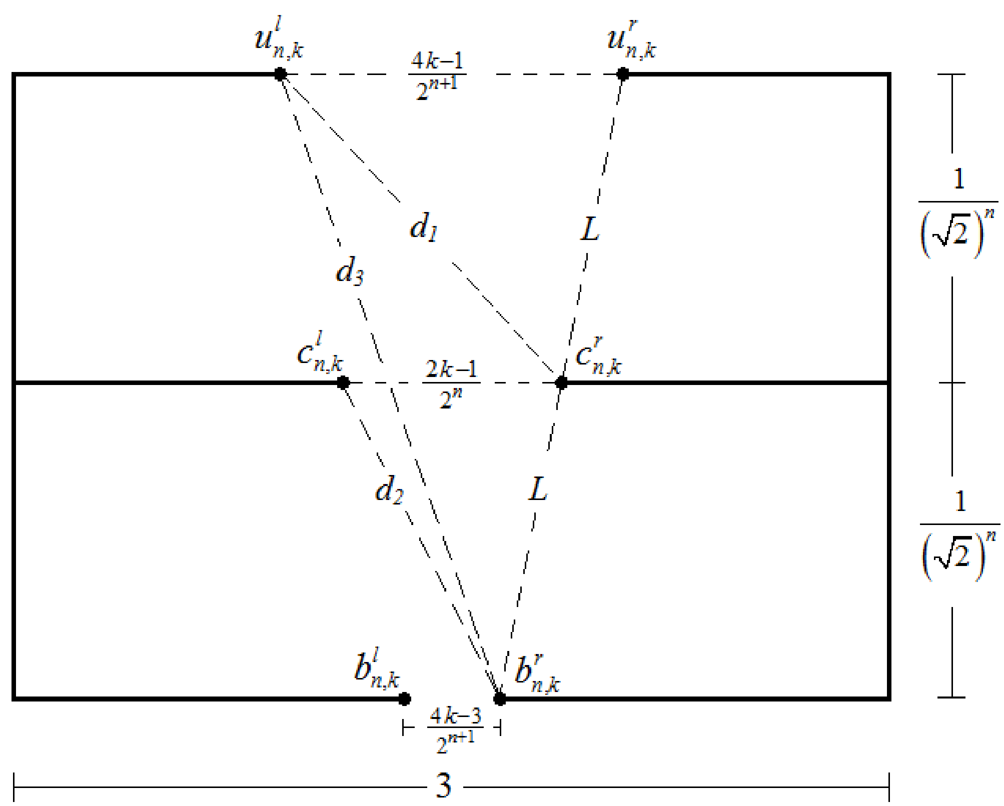

Figure 4.5: The Basic Piece

One can prove the following via direct computation.

$$d_1 > \frac{4k-1}{2^{n+1}}, \;\; d_2 > \frac{2k-1}{2^n}, \;\; d_3 > \frac{4k-1}{2^{n+1}},$$

$$d\big(u^r_{n,k}, c^r_{n,k}\big) = d\big(c^r_{n,k}, b^r_{n,k}\big) = L < 2^{-\frac{n-2}{2}}.$$

Thus, since the diagonals are all longer than than the gaps above them, it follows from our previous examples and discussions that each gap length in $A_{n,k}$ is a critical value of $A_{n,k}$.

Now, consider an isometric copy of $A_{n,k}$ imbedded in $span\{e_1, e_{2^{n-1}+k}\}$ as in Figure 4.6. We will abuse notation slightly and refer to this imbedded copy by $A_{n,k}$, also. We will view this space as it is in this figure, and we will use the corresponding terminology when referring to it. For instance, the *right central gap point of* $A_{n,k}$ is $c^r_{n,k}$; the *left upper gap point of* $A_{n,k}$ is $u^l_{n,k}$, and so on. Note that the upper and lower gap lengths of $A_{n,k}$ correspond, respectively, to the central gap lengths of $A_{n+1,2k}$ and $A_{n+1,2k-1}$. This correspondence illustrates how we will attach each new piece. In fact, the basic pieces form a tree structure directly corresponding to the dyadic rational tree above. If we replace each dyadic rational in the tree above with the basic piece, $A_{n,k}$, having that rational as its central gap length, then we obtain Figure 4.7. This actually illustrates the construction process; X_n will consist of all of the pieces from generations 1 through n, and we construct X_{n+1} by attaching the pieces in generation $n+1$.

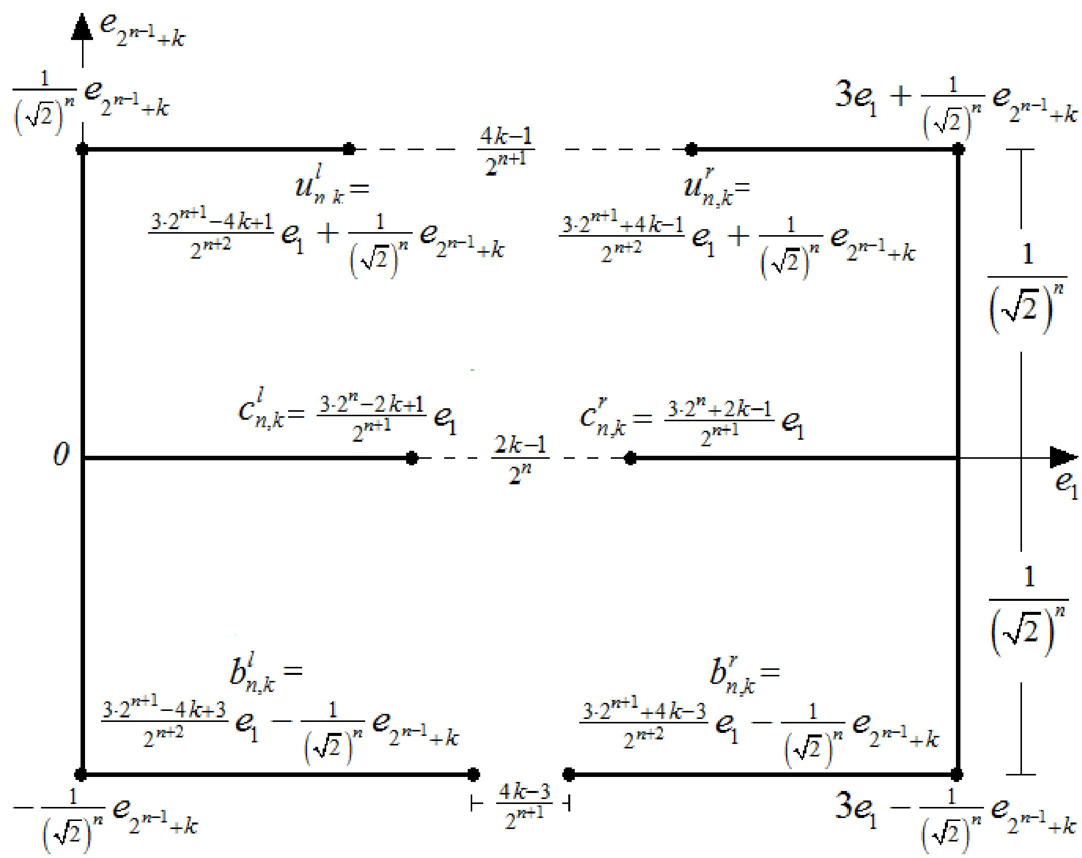

Figure 4.6: The Imbedded Basic Piece

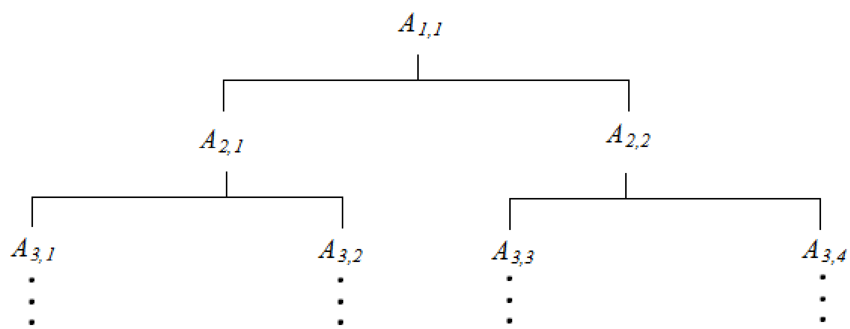

Figure 4.7: Organization and Attachment Order of the Basic Pieces

We will describe the first couple of steps of the construction process in detail. After that, the general inductive process should be clear. We begin by letting $X_1 = A_{1,1}$. Since the central gap length of $A_{2,1}$ is the same as the lower gap length of $A_{1,1}$, we want to attach $A_{2,1}$ so that these two gaps coincide. But $A_{2,1}$ is already in $span\{e_1, e_3\}$, so the attachment can be effected simply by translating or shifting $A_{2,1}$ by the vector $(-1/\sqrt{2})e_2$. Likewise, the central gap length of $A_{2,2}$ is the same as the upper gap length of $A_{1,1}$, and since $A_{2,2}$ lies in $span\{e_1, e_4\}$, we can attach it to $A_{1,1}$ simply by shifting it by the vector $(1/\sqrt{2})e_2$. After attaching $A_{2,1}$ and $A_{2,2}$, to X_1 in this way, the resulting space is X_2. Note that X_1 isometrically imbeds into X_2. Moreover, each point on $A_{2,1}$ (and $A_{2,2}$) is less than a distance of $2^{-(2-2)/2} = 1$ away from a point of X_1. (See L in Figure 4.5.) Thus, we have

$$d_H(X_1, X_2) < 2^{-(2-2)/2} \ \text{ and } \ diam(X_2) \leq diam(X_1) + 2 \cdot 2^{-\frac{2-2}{2}}.$$

Since the construction is carried out entirely in $\mathscr{H}$, it suffices to just use the Hausdorff metric, d_H, in $\mathscr{H}$ instead of the general Gromov-Hausdorff metric. We also note that X_2 now has four *ends* corresponding to the gap lengths 1/8, 3/8, 5/8, and 7/8. The upper and lower left corners of $A_{2,1}$ now (i.e. after attaching it to X_1) lie, respectively, at

$$-\frac{1}{\sqrt{2}}e_2 + \frac{1}{(\sqrt{2})^2}e_3 \ \text{ and } \ -\frac{1}{\sqrt{2}}e_2 - \frac{1}{(\sqrt{2})^2}e_3.$$

Likewise, the upper and lower left corners of $A_{2,2}$ now lie at

$$\frac{1}{\sqrt{2}}e_2 + \frac{1}{(\sqrt{2})^2}e_4 \ \text{ and } \ \frac{1}{\sqrt{2}}e_2 - \frac{1}{(\sqrt{2})^2}e_4.$$

Now, we construct X_3 by attaching each $A_{3,k}$, $1 \leq k \leq 2^{3-1}$, to the end of X_2 corresponding to its central gap length. To attach $A_{3,1}$ so that its central gap lines up with the lower gap of $A_{2,1}$, we need only shift $A_{3,1}$ - since it already lies in $span\{e_1, e_5\}$ - by $-(1/\sqrt{2})e_2 - (1/2)e_3$. This ensures that the central gap of $A_{3,1}$, after shifting, then coincides with the lower gap of $A_{2,1}$. Note, also, that the vector by which we shifted $A_{3,1}$ is precisely the vector we obtained above as the lower left corner of $A_{2,1}$ after attaching it to X_1. Similarly, we attach $A_{3,2}$ so that its central gap coincides with the upper gap of of $A_{2,1}$, and this is effected by shifting $A_{3,2}$ by the vector $-(1/\sqrt{2})e_2 + (1/2)e_3$, which, again, is just the new upper left corner of $A_{2,1}$ after attaching it to X_1. Finally, we attach $A_{3,3}$ and $A_{3,4}$ to the lower and upper gaps, respectively, of $A_{2,2}$, and this is effected by shifting $A_{3,3}$ and $A_{3,4}$, respectively, by the vectors

$$\frac{1}{\sqrt{2}}e_2 - \frac{1}{(\sqrt{2})^2}e_4 \text{ and } \frac{1}{\sqrt{2}}e_2 + \frac{1}{(\sqrt{2})^2}e_4.$$

The resulting space is X_3. As in the first step, X_2 isometrically imbeds into X_3, and every point of X_3 is less than a distance of $2^{-(3-2)/2}$ away from a point of X_2. It follows that $d_H(X_2, X_3) < 2^{-(3-2)/2}$ and

$$diam(X_3) \leq diam(X_2) + 2 \cdot 2^{-\frac{3-2}{2}} \leq diam(X_1) + 2\left(\sum_{i=2}^{3} 2^{-\frac{i-2}{2}}\right).$$

Now, for the general process, we need a way to determine - in terms of n and k - what vectors to shift each piece by in order to effect its attachment. Toward this end, we set $s_{1,1} = 0$, and we inductively define $s_{n,k}$, $n \geq 2$, $1 \leq k \leq 2^{n-1}$ by the following. Given $s_{n,k}$, we define

$$s_{n+1,2k-1} = s_{n,k} + \frac{(-1)^{2k-1}}{(\sqrt{2})^n}e_{2^{n-1}+k} = s_{n,k} - \frac{1}{(\sqrt{2})^n}e_{2^{n-1}+k}$$

$$s_{n+1,2k} = s_{n,k} + \frac{(-1)^{2k}}{(\sqrt{2})^n}e_{2^{n-1}+k} = s_{n,k} + \frac{1}{(\sqrt{2})^n}e_{2^{n-1}+k}.$$

Then $s_{n,k}$ is the vector by which we shift or translate $A_{n,k}$ to attach it to X_{n-1} for $n \geq 2$, $1 \leq k \leq 2^{n-1}$. So, given X_n for some $n \geq 3$, we construct X_{n+1} by shifting each $A_{n+1,k}$, $1 \leq k \leq 2^{n+1-1} = 2^n$, by $s_{n+1,k}$, and this has the effect of lining up the central gap of $A_{n+1,k}$ with the corresponding end-gap of X_n having the same length. This yields a sequence of compact metric spaces, $\{X_n\}$, with the following properties.

1) X_n isometrically imbeds into X_m for all $1 \leq n < m$.

2) $d_H(X_{n-1}, X_n) < 2^{-\frac{n-2}{2}}$ for all $n \geq 1$.

3) $diam(X_n) \leq diam(X_{n-1}) + 2 \cdot 2^{-\frac{n-2}{2}}$ for all $n \geq 1$.

From induction on n in the third inequality, we also see that

$$diam(X_n) \leq diam(X_1) + 2\sum_{i=2}^{n} 2^{-\frac{i-2}{2}} \leq diam(X_1) + 2\sum_{i=2}^{\infty} 2^{-\frac{i-2}{2}}.$$

Thus, the diameters of the spaces in this sequence are uniformly bounded. Moreover, property 2 above, taken with the fact that $\sum_{i=1}^{\infty} 2^{-(i-2)/2} < \infty$, implies that the sequence $\{X_n\}$ is Cauchy. Since $\mathscr{H}$ is complete, the corresponding metric space of compact subspaces of $\mathscr{H}$, endowed with the Hausdorff metric, is also complete. Thus, there is a metric space, $X \subset \mathscr{H}$, such that $X_n \to X$. In fact, since $X_n \subset X_{n+1}$ for all n, it follows that $X = \overline{\cup_{n=1}^{\infty} X_n}$. Now, X is complete, because it is a closed subspace of $\mathscr{H}$. To show that it is compact, we need only show that X is totally bounded. For this, it suffices to show that $X' := \cup_{n=1}^{\infty} X_n$ is totally bounded, but this, in turn, is a simple consequence of the fact that each X_n is compact and $d_H(X_{n-1}, X_n) < 2^{-(n-2)/2}$. In fact, given $\epsilon > 0$, if we choose N large enough so that $\sum_{i=N}^{\infty} 2^{-(i-2)/2} < \epsilon$, and if we choose a finite number of points in X_N - say $x_1, \ldots, x_k$ - such that the $\epsilon/2$-balls centered at these points cover X_N (and, thus, $\cup_{n=1}^{N} X_n$), then the corresponding collection of ϵ-balls centered at these points covers $\cup_{n=1}^{\infty} X_n$. Hence, X is compact.

Note, also, that X' is chain-connected, for if $\epsilon > 0$ is given, we can choose n large enough and $1 \leq k \leq 2^{n-1}$ so that $(2k-1)/2^n < \epsilon$. But $A_{n,k}$ has central gap length equal to $(2k-1)/2^n$, so we can cross from one "side" of X' to the other by crossing over this gap. Consequently, X is chain-connected and, therefore, connected. Additionally, we can attach a long connecting curve to X' in the $\{e_0, e_1\}$-plane (recall that this is why we saved the initial basis vector) to make X' path-connected, and we can make sure this curve is long enough so as not to interfere with any of the critical values. Thus, X is a compact, connected metric space with a dense, path-connected subspace.

Now, there is one last detail; we have not yet shown that the critical values of each basic piece are maintained during the construction process. This may be intuitively evident, since we attached each piece in a different dimension than the rest of the pieces. Indeed, this was the reason for doing so; the fact that each piece extends in a different dimension than all other pieces means that the pieces do not interfere with each other with regard to the gaps and the diagonals. This must still be proved, however, and it is the most cumbersome part of the proof. The computations are not difficult, but they are technical and tedious. For the sake of brevity, we will supress most of the actual computations and just state the results that indicate that each gap length does, in fact, remain a critical value of X.

The goal is to show that each pair of central gap points, $\{c^l_{n,k}, c^r_{n,k}\}$, is an essential gap. It is easy to see that $dist(B(c^l_{n,k}, \tau), B(c^r_{n,k}, \tau)) = d(c^l_{n,k}, c^r_{n,k})$ for all sufficiently small τ. Thus, the Essential Gap Lemma will imply that $\{c^l_{n,k}, c^r_{n,k}\}$ is an essential gap, implying that $(2k-1)/2^n$ is a critical value. Furthermore, by Theorem 4.3.5, it suffices to work only in X'.

From the discussions and examples given in this chapter, it should be evident that we only need to show that there is some $\tau > d(c^l_{n,k}, c^r_{n,k}) = (2k-1)/2^n$ such that the distance between $c^r_{n,k}$ and the *left* endpoints of all other gaps in X' is greater than or equal to τ. This would mean that lengths of all of the diagonals from $c^r_{n,k}$ to all left endpoints of all other gaps are greater than and bounded away from $(2k-1)/2^n$, intuitively meaning that one cannot "jump across" from $c^r_{n,k}$ to the left side of X' with a jump that is greater than, but sufficiently close to, $(2k-1)/2^n$ in length. Finally, since every upper and lower left gap point of any $A_{n,k}$ is also the left central gap point of some other $A_{n',k'}$, it suffices to work only with central gap points. In other words, the problem of showing that $\{c^l_{n,k}, c^r_{n,k}\}$ is an essential gap reduces to the proving the following: for fixed $n \geq 1$, $1 \leq k \leq 2^{n-1}$, there is some $\tau > (2k-1)/2^n$ such that $d(c^r_{n,k}, c^l_{m,j}) \geq \tau$ for all m, j with $m \neq n$ and all $j \neq k$ when $m = n$.

First, one proves the following results, most of which actually follow easily from the definitions and induction. The last two still follow by induction, but they take a bit more work. We use the symbols $\langle , \rangle$ and $\perp$ to denote the Hilbert space inner product and the notion of orthogonality.

1) $s_{n,k} \perp e_1$ for all $n \geq 1$, $1 \leq k \leq 2^{n-1}$.

2) $s_{n,k} \perp e_{2^{m-1}+j}$ for all $m \geq n$, $1 \leq j \leq 2^{m-1}$.

3) $\langle s_{n,k}, s_{n,k} \rangle = \sum_{i=1}^{n-1} 2^{-i}$ for all $n \geq 2$, $1 \leq k \leq 2^{n-1}$, and $\langle s_{1,1}, s_{1,1} \rangle = 0$.

4) For $n \geq 2$ and $1 \leq j < k \leq 2^{n-1}$,

$$\|s_{n,k} - s_{n,j}\|^2 > \Big(\frac{3k+j-2}{2^n}\Big)\Big(\frac{k-j}{2^n}\Big).$$

5) For $n \geq 2$, $1 \leq k \leq 2^{n-1}$, $m > n$, and $1 \leq j \leq 2^{m-1}$,

$$\|s_{n,k} - s_{m,j}\|^2 > \left(\frac{2^{m-n} - 2^{m-n+1}k + 2j - 1}{2^{m+1}}\right)\left(\frac{2^{m-n+1}k - 2^{m-n} + 6j - 3}{2^{m+1}}\right).$$

Now, for the gap $\{c^l_{1,1}, c^r_{1,1}\} = \{(5/4)e_1, (7/4)e_1\}$ (see Figure 4.6), the desired result can be shown directly. We will show this result to give some illustration of the nature of the computations involved in this process. The rest of the computations are carried out in a similar manner, but they are much more lengthy. We want to show that the distance from $(7/4)e_1$ to any $c^l_{n,k}$ - with $n \geq 2$ and $1 \leq k \leq 2^{n-1}$ - is at least $\frac{1}{2} + \tau$ for some $\tau > 0$. We reason as follows.

$$n \geq 2 \Rightarrow 1 - \frac{1}{2^{n-1}} \geq \frac{1}{2} > \frac{3}{8} \Rightarrow \left(\frac{2^n + 4k - 2}{2^{n+2}}\right)^2 + 1 - \frac{1}{2^{n-1}} > \frac{3}{8}$$

$$\begin{aligned} \Rightarrow \frac{3}{8} &< \left(\frac{3 \cdot 2^{n+1} - 4k + 2 - 7 \cdot 2^n}{2^{n+2}}\right)^2 + 1 - \frac{1}{2^{n-1}} \\ &< \left(\frac{3 \cdot 2^{n+1} - 2k + 1}{2^{n+1}} - \frac{7}{4}\right)^2 + 1 - \frac{1}{2^{n-1}} \end{aligned}$$

From this, we obtain

$$\frac{49}{16} - 2 \cdot \frac{7}{4} \cdot \frac{3 \cdot 2^n - 2k + 1}{2^{n+1}} + \left(\frac{3 \cdot 2^n - 2k + 1}{2^{n+1}}\right)^2 + \sum_{i=1}^{n-1} \frac{1}{2^i} > \frac{3}{8}.$$

Using property 3 above, substituting $\langle s_{n,k}, s_{n,k} \rangle$ for the sum on the left-hand side, and condensing the inner product, this implies that

$$\left\langle \frac{7}{4}e_1 - \frac{3 \cdot 2^n - 2k + 1}{2^{n+1}}e_1 - s_{n,k}, \frac{7}{4}e_1 - \frac{3 \cdot 2^n - 2k + 1}{2^{n+1}}e_1 - s_{n,k} \right\rangle > \frac{3}{8}$$

$$\Rightarrow d\left(\frac{7}{4}e_1, \frac{3 \cdot 2^n - 2k + 1}{2^{n+1}}e_1 + s_{n,k}\right)^2 > \frac{3}{8}.$$

Thus, $d((7/4)e_1, c^l_{n,k}) > \sqrt{3/8} > \frac{1}{2}$, and this holds for any $n \geq 2$, $1 \leq k \leq 2^{n-1}$. Hence, $\frac{1}{2}$ is a critical value of X' and, therefore, of X.

Now, for the general case, we fix $n \geq 2$ and $1 \leq k \leq 2^{n-1}$. First, using inequality 4 above, one can show that, for $j \neq k$,

$$d\left(\frac{3 \cdot 2^n + 2k - 1}{2^{n+1}}e_1 + s_{n,k}, \frac{3 \cdot 2^n - 2j + 1}{2^{n+1}}e_1 + s_{n,j}\right) > \frac{2k - 1}{2^n}.$$

Note that the point on the left is $c^r_{n,k}$ after $A_{n,k}$ has been shifted and attached to X_{n-1}, and the point on the right is the left center gap point of $A_{n,j}$ after it has been shifted. Thus, all of the diagonals from $c^r_{n,k}$ to the left center gap point of any other $A_{n,j}$ are strictly greater than $(2k - 1)/2^n$.

Using 4 again, one can prove by induction on m - starting by directly proving the result for $m = n + 1$ - that

$$d(c^r_{n,k}, c^l_{m,j}) = d\left(c^r_{n,k}, \frac{3 \cdot 2^m - 2j + 1}{2^{m+1}} e_1 + s_{m,j}\right) \geq \frac{1}{2^n} + \left(\frac{8k-5}{2^{n+2}}\right)^2,$$

for any $m > n$ and $1 \leq j \leq 2^{m-1}$. Moreover, $\frac{1}{2^n} + ((8k-5)/(2^{n+2}))^2$ is the length of the lower diagonal of $A_{n,k}$ (i.e. the length d_2 in Figure 8), which is strictly greater than $(2k-1)/2^n$.

Finally, using inequality 5 above, one then proves that $d(c^r_{n,k}, c^l_{m,j})$ - or the distance from the left center gap point of $A_{m,j}$ to $c^r_{n,k}$ - is strictly greater than $(2k-1)/2^n$ for $m < n$ and $1 \leq j \leq 2^{m-1}$. This is also proved by induction. Assuming $n \geq 2$, it can easily be shown for $m = j = 1$ and $m = 2$, $1 \leq j \leq 2$. Then, inequality 5 is the key to the induction step.

Putting all of this together, we conclude the following. First, we know that $d(c^r_{n,k}, c^l_{m,j}) > (2k-1)/2^n$ for all $m \leq n$ and $1 \leq j \leq 2^{m-1}$ ($j \neq k$ when $n = m$). But there are only finitely many such $c^l_{m,j}$ for $m \leq n$. Furthermore, for $m > n$, we know that $d(c^r_{n,k}, c^l_{m,j})$ is at least $\frac{1}{2^n} + ((8k-5)/(2^{n+2}))^2$, which is greater than $(2k-1)/2^n$. Thus, by choosing the minimum of $\frac{1}{2^n} + ((8k-5)/(2^{n+2}))^2$ and the distances $d(c^r_{n,k}, c^l_{m,j})$ - where m ranges over $1, \ldots, n$, $1 \leq j \leq 2^{m-1}$ (again, $j \neq k$ when $n = m$) - we obtain a $\tau > (2k-1)/2^n$ such that $d(c^r_{n,k}, c^l_{m,j}) \geq \tau$ for all m and j. Hence, $(2k-1)/2^n$ is a critical value of X' and, thus, X.

Summing up, we have constructed a compact, connected metric space, X, with a path-connected, dense subspace, X', such that $(0, 1] \subset Cr(X)$. Another interesting point to make about X' is that it is topologically self-similar. Except for the fact that each piece extends in a different dimension in $\mathscr{H}$, X' essentially has a binary tree structure.

Chapter 5

Refinability and ϵ-group Geometry

In this section, we will investigate some interesting connections between the ability/inability to refine chains, the metric structure of the ϵ-covers, and the geometry of the ϵ-groups.

5.1 Refinability and Generators of $\pi_\epsilon(X)$

We begin with a theorem relating refinability to the property of $\pi_\epsilon(X)$ being finitely generated.

Theorem 5.1.1 *Let X be a chain-connected metric space, and let $\epsilon > 0$ be given. If $\pi_\epsilon(X)$ is finitely generated, then X is ϵ-surjective from below. If X is compact, the converse also holds.*

Proof Suppose $\pi_\epsilon(X)$ is finitely generated, and let $\{[\gamma_1]_\epsilon, \dots, [\gamma_n]_\epsilon\}$ be a generating set. Each loop, γ_i, $1 \le i \le n$, is an ϵ-loop, which means that each loop is also a δ-loop for some δ less than but sufficiently close to ϵ. Since there are only finitely many generators, we can choose $\delta < \epsilon$ such that each γ_i is a δ-loop. Let γ be any ϵ-loop at $* \in X$. Then we can express $[\gamma]_\epsilon$ as a finite product of generators,

$$[\gamma]_\epsilon = [\gamma_{i_1}]_\epsilon^{\pm 1} \cdots [\gamma_{i_k}]_\epsilon^{\pm 1}.$$

But this means that γ is ϵ-homotopic to the product $\gamma_{i_1}^{\pm 1} \cdots \gamma_{i_k}^{\pm 1}$, which is a δ-loop. Thus, γ is ϵ-homotopic to a δ-loop, showing that the map $\Phi_{\epsilon\delta} : \pi_\delta(X) \to \pi_\epsilon(X)$ is surjective. By Lemma 2.3.2, it follows that $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is surjective.

Now, assume that X is compact and ϵ-surjective from below. Since X is ϵ-surjective from below, there is some λ, with $0 < \lambda < \epsilon$, such that the map $\varphi_{\epsilon,\epsilon-\lambda} : X_{\epsilon-\lambda} \to X_\epsilon$ is surjective. Since X is compact, we can find a finite $\lambda/3$-net, $S \subset X$, meaning that every point, $x \in X$, is a distance less than $\lambda/3$ from at least one point of S. Let $S = \{z_1, \dots, z_m\}$. We may assume without loss of generality that our base point, $*$, is in S.

We first prove the following claim: if $\alpha = \{x = x_0, \dots, x_n = y\}$ is any ϵ-chain in X with endpoints, x and y, in S, then α is ϵ-homotopic to an ϵ-chain

$$\alpha' = \{x = x_0, z_{i_1}, z_{i_2}, \dots, z_{i_{n-1}}, x_n = y\}$$

with each $z_{i_j} \in S$ and such that $x_j \in B(z_{i_j}, \lambda/3)$ for $j = 1, \dots, n-1$. First, by hypothesis, α can be ϵ-refined to an $(\epsilon - \lambda)$-chain, so we may simply assume that α is already an $(\epsilon - \lambda)$-chain. There is some $z_{i_1} \in S$ such that $x_1 \in B(z_{i_1}, \lambda/3)$, implying that $d(x_1, z_{i_1}) < \lambda/3 < \epsilon$. We also have

$$d(x_0, z_{i_1}) \le d(x_0, x_1) + d(x_1, z_{i_1}) < \epsilon - \lambda + \frac{\lambda}{3} = \epsilon - \frac{2\lambda}{3} < \epsilon.$$

Thus, we can insert z_{i_1} into α between x_0 and x_1. But we also have

$$d(x_2, z_{i_1}) \leq d(x_2, x_1) + d(x_1, z_{i_1}) < \epsilon - \lambda + \frac{\lambda}{3} = \epsilon - \frac{2\lambda}{3} < \epsilon.$$

So, we can then remove x_1 to obtain the new chain

$$\alpha^{(1)} = \{x = x_0, z_{i_1}, x_2, \ldots, x_n = y\}.$$

Now, if α has three points, then we are done. If not, there is some $z_{i_2} \in S$ such that $d(x_2, z_{i_2}) < \lambda/3 < \epsilon$,

$$d(z_{i_2}, x_3) \leq d(z_{i_2}, x_2) + d(x_2, x_3) < \frac{\lambda}{3} + \epsilon - \lambda = \epsilon - \frac{2\lambda}{3} < \epsilon,$$

$$d(z_{i_1}, z_{i_2}) \leq d(z_{i_1}, x_2) + d(x_2, z_{i_2}) < \epsilon - \frac{2\lambda}{3} + \frac{\lambda}{3} = \epsilon - \frac{\lambda}{3} < \epsilon.$$

The first two inequalities imply that we can insert z_{i_2} into $\alpha^{(1)}$ between x_2 and x_3, and the third implies that we can, then, remove x_2 to obtain the chain $\alpha^{(2)} = \{x = x_0, z_{i_1}, z_{i_2}, x_3, \ldots, x_n = y\}$.

If α has four points, then we are done. If not, then we can continue this process inductively, each step following exactly as before. This process must clearly stop once we add in $z_{i_{n-1}}$ and remove x_{n-1}, proving the claim.

Now, define a set, $\mathscr{L}$, of ϵ-loops at $* \in X$ as follows: $\mathscr{L}$ is the set of all ϵ-loops at $*$ of the form $\sigma\gamma\sigma^{-1}$, where $\sigma = \{* = z_{i_1}, z_{i_2}, \ldots, z_{i_k}\}$, $\gamma = \{z_{i_k}, z_{i_{k+1}}, \ldots, z_{i_{k+r}} = z_{i_k}\}$, each $z_{i_j} \in S$, and the points

$$\{z_{i_1}, \ldots, z_{i_k}, z_{i_{k+1}}, \ldots, z_{i_{k+r-1}}\}$$

are all distinct. That is, $\sigma\gamma\sigma^{-1}$ is formed by taking an ϵ-chain, σ, based at $*$ and consisting of distinct points of S, concatenating that with an ϵ-loop, γ, having the property that the points of γ come from S and are not only distinct from each other but also - with the exception of the initial and terminal points of γ agreeing with the terminal point of σ - distinct from those in σ, and then concatenating that with σ^{-1}. Hence, an element of $\mathscr{L}$ can be written as

$$\{\underbrace{* = z_{i_1}, \ldots, z_{i_k}}_{\sigma}, \underbrace{z_{i_k}, z_{i_{k+1}}, \ldots, z_{i_{k+r-1}}, z_{i_k} = z_{i_{k+r}}}_{\gamma}, \underbrace{z_{i_k}, z_{i_{k-1}}, \ldots, z_{i_1} = *}_{\sigma^{-1}}\},$$

where for any $1 \leq j < l \leq k + r - 1$, $z_{i_j} \neq z_{i_l}$.

Note that the finiteness of S and the conditions on σ and γ in the definition of $\mathscr{L}$ imply that $\mathscr{L}$ is finite. Moreover, since $* \in S$, the trivial loop, $\{*\}$ is in S. It follows that the set

$$\mathscr{L}_\epsilon := \{[\gamma]_\epsilon \in \pi_\epsilon(X) \, : \, \gamma \in \mathscr{L}\}$$

is also finite and contains $[\{*\}]_\epsilon$. We will show that $\mathscr{L}_\epsilon$ generates $\pi_\epsilon(X)$. It suffices, of course, to show that any ϵ-loop at $*$ is ϵ-homotopic to a loop in $\mathscr{L}$. The proof is by induction. Any ϵ-loop, α, can - by assumption - be refined to an $(\epsilon - \lambda)$-loop, and we will prove the result by strong induction on the number of points in an $(\epsilon - \lambda)$-loop in $[\alpha]_\epsilon$. For the base step, we note that any ϵ-loop consisting of four or fewer points is ϵ-null. So, if α is an ϵ-loop that can be ϵ-refined to an $(\epsilon - \lambda)$-loop consisting of four or fewer points, then $[\alpha]_\epsilon = [\{*\}]_\epsilon \in \mathscr{L}_\epsilon$ and we are done.

Now, assume the following inductive hypothesis for some $n \geq 4$: any ϵ-loop, α, at $*$ that can be ϵ-refined to an $(\epsilon - \lambda)$-loop consisting of n or fewer points is ϵ-homotopic to a product of loops from $\mathscr{L}$. Let α be an ϵ-loop that can be ϵ-refined to an $(\epsilon - \lambda)$-loop consisting of $n+1$ points. For simplicity of notation, we will just assume that α already is an $(\epsilon - \lambda)$-loop and denote it by $\alpha = \{* = x_0, \ldots, x_n = *\}$. By the claim above, α is ϵ-homotopic to an ϵ-loop,

$$\alpha' = \{* = z_{i_0}, z_{i_1}, \ldots, z_{i_{n-1}}, * = z_{i_n}\},$$

with each $z_{i_j} \in S$. If the points $\{*, z_{i_1}, \ldots, z_{i_{n-1}}\}$ are all distinct then we are done, since α' is, then, an element of $\mathscr{L}$. If not, let j be the smallest index such that $z_{i_k} = z_{i_j}$ for some $k < j$. Let β be the ϵ-loop

$$\beta := \{*, z_{i_1}, \ldots, z_{i_k}, z_{i_{k+1}}, \ldots, z_{i_j}, z_{i_{k-1}}, \ldots, z_{i_1}, *\}.$$

Then β is an element of $\mathscr{L}$, where

$$\sigma = \{*, z_{i_1}, \ldots, z_{i_k}\} \text{ and } \gamma = \{z_{i_k}, z_{i_{k+1}}, \ldots, z_{i_j} = z_{i_k}\}.$$

If we let $\mu = \{*, z_{i_1}, \ldots, z_{i_k}, z_{i_{j+1}}, \ldots, z_{i_{n-1}}, *\}$, then μ is an ϵ-loop because $d(z_{i_k}, z_{i_{j+1}}) = d(z_{i_j}, z_{i_{j+1}}) < \epsilon$, and we have $\alpha' = \beta\mu$. Moreover, μ has n or fewer points, so, by the inductive hypothesis, μ can be ϵ-refined to a product of elements from $\mathscr{L}$. It follows that α' - and, therefore, α - can, also. ∎

5.2 Refinability and Proper ϵ-covers

In this section, we will strengthen the previous result slightly. Though we have not mentioned it explicitly, yet, we have already seen an example where the ϵ-cover of a metric space, X, need not be proper even when X is compact. In fact, Example 4.2.3 has this property, and this will become clear in the proof of Theorem 5.2.6 below. We saw, in that example, that the space in question was not ϵ-surjective from below. The main theorem we will prove in this section - Theorem 5.2.6 - will clarify the connection between these two phenomena. We need some preliminary results beforehand, however. First, we need a stronger refinement property than just ϵ-surjectivity from below. Recall from Lemma 2.3.2 that the ability to refine ϵ-chains depends solely on the ability to refine two-point ϵ-chains. We use this in the following.

Definition 5.2.1 *Given $\epsilon > 0$, a chain-connected metric space, X, has the **ϵ-Bounded Minimal Refinement Property** - or the **ϵ-BMR property** - if there exist $0 < \delta < \epsilon$ and a natural number, N, such that every two-point ϵ-chain in X can be ϵ-refined to a δ-chain consisting of N or fewer points.*

This property is clearly stronger than simply requiring that X be ϵ-surjective from below. The point of this property is that it not only allows us to refine ϵ-chains to δ-chains but to also control the *lengths* of the refinements in terms of the original chains. If X has the ϵ-BMR property, and if $\alpha = \{x_0, \ldots, x_n\}$ is an ϵ-chain in X, then, for each $i = 1, \ldots, n$, we can refine the subchain $\{x_{i-1}, x_i\}$ to a δ-chain of N or fewer points. Doing this for each of the n subchains, while leaving the points of α fixed, gives us an ϵ-refinement of α to a δ-chain with at most $nN - (n-1) = n(N-1) + 1$ points. The word "minimal" in this definition comes from the following observation: if we can refine an ϵ-chain, α, to a δ-chain for $\delta < \epsilon$, then there may be many such δ-chains that are refinements of α. However, there will be a *minimal refinement*, or one of minimal cardinality. This follows from the well-ordering of $\mathbb{N}$. The ϵ-BMR property gives us control over the length of this minimal refinement. In particular, if X has the ϵ-BMR property and, in addition, the map $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is injective (hence, bijective), then the inverse of this map is Lipschitz, and X_ϵ and X_δ are not just homeomorphic but also bi-Lipschitz equivalent.

All geodesic spaces have the ϵ-BMR property, since, by simply adding the midpoint to any two-point ϵ-chain, we can let $\delta = \epsilon/2$ and $N = 3$ in the definition above. There are many non-geodesic spaces, however, that also satisfiy this property.

The next concept we need will prove useful in later chapters, also. We will define what will essentially turn out to be a fundamental domain for the action of $\pi_\epsilon(X)$ on X_ϵ in the case when X_ϵ is proper. Of course, we could appeal to well-known results to conclude that such a fundamental domain exists, but, as will be seen, it will be helpful to have this concrete example we can refer to when needed. Also, for brevity, we will adopt a standard notation for the endpoint of an ϵ-chain. Given an ϵ-chain, α, with endpoint, x, we set $\alpha_t := x$ (t for *terminal*).

Now, assume, for a chain-connected metric space, X, and some $\epsilon > 0$, that X_ϵ is proper, so that closed metric balls in X_ϵ are compact. As usual, let $*$ be our base point in X and $\tilde{*} = [\{*\}]_\epsilon$ the corresponding lifted base point in X_ϵ. Given $x \in X$, define

$$L_{[\epsilon]}(x) := \inf\{L([\alpha]_\epsilon) \,:\, \alpha \text{ is an } \epsilon\text{-chain at } * \text{ with } \alpha_t = x\}.$$

Note that we're taking the infimum over *ϵ-homotopy equivalence classes* of ϵ-chains from $*$ to x. We want to know that there exists at least one class of chains, $[\alpha]_\epsilon$, that actually attains this infimum.

To see why this is true, we first note the following: there are only finitely many classes of chains, $[\alpha]_\epsilon \in X_\epsilon$, with $\alpha_t = x$ and

$$L_{[\epsilon]}(x) \leq L([\alpha]_\epsilon) < L_{[\epsilon]}(x) + 1.$$

This holds because X_ϵ is proper. In fact, if we had infinitely many *distinct* classes satisfying this relation, then we could find a sequence, say $\{[\alpha_n]_\epsilon\} \subset X_\epsilon$, with $(\alpha_n)_t = x$, $L([\alpha_n]_\epsilon) < L_{[\epsilon]}(x)+1$, and $[\alpha_n]_\epsilon \neq [\alpha_m]_\epsilon$ for all $n \neq m$. Then, since X_ϵ is proper and $d_\epsilon(\tilde{*}, [\alpha_n]_\epsilon) = L([\alpha_n]_\epsilon) < L_{[\epsilon]}(x)+1$ for all n, there will be a convergent subsequence, $[\alpha_{n_k}]_\epsilon \to [\alpha]_\epsilon$. In particular, we must have $\alpha_t = x$ and $d_\epsilon([\alpha_{n_k}]_\epsilon, [\alpha]_\epsilon) < \epsilon$ for all sufficiently large k. This further implies that the loop, $\alpha_{n_k}\alpha^{-1}$, is ϵ-null, or $\alpha_{n_k} \sim_\epsilon \alpha$, for sufficiently large k. In other words, we have $[\alpha_{n_k}]_\epsilon = [\alpha]_\epsilon$ for all sufficiently large k, contradicting that the elements of the sequence $\{[\alpha_n]_\epsilon\}$ are distinct. It follows, then, that $L_{[\epsilon]}(x)$ must be attained by at least one element, $[\alpha]_\epsilon$, with $\alpha_t = x$. In fact, this also shows that the number of elements attaining this infimum is finite.

We will call any $[\alpha]_\epsilon \in X_\epsilon$ satisfying $\alpha_t = x$ and $L([\alpha]_\epsilon) = L_{[\epsilon]}(x)$ a **minimal ϵ-class from $*$ to x**. The preceding argument shows that, for every $x \in X$, there is at least one minimal ϵ-class from $*$ to x and that there are only finitely many minimal ϵ-classes from $*$ to x.

Definition 5.2.2 *Given a chain-connected metric space, X, and $\epsilon > 0$ such that X_ϵ is proper, we define $\tilde{X}_\epsilon$ to be the set of all minimal ϵ-classes in X_ϵ, or, equivalently, if $m(x, \epsilon)$ is the set of all minimal ϵ-classes from $*$ to x, then*

$$\tilde{X}_\epsilon := \bigcup_{x\in X} m(x, \epsilon).$$

Note that requiring that X_ϵ be proper was essential in defining these sets. Without this property, the infimum, $L_{[\epsilon]}(x)$, is still well-defined, but it need not be attained by any class of ϵ-chains.

Lemma 5.2.3 *Given a chain-connected metric space, X, and $\epsilon > 0$ such that X_ϵ is proper, the following properties hold for $\tilde{X}_\epsilon$.*

1) $\tilde{X}_\epsilon$ is closed.

2) The collection of translates, $\{[\gamma]_\epsilon \tilde{X}_\epsilon\}_{[\gamma]_\epsilon \in \pi_\epsilon(X)}$, covers X_ϵ.

3) If, in addition, X is compact, then $\tilde{X}_\epsilon$ is compact.

Proof Suppose $\{[\alpha_n]_\epsilon\} \subset \tilde{X}_\epsilon$ and $[\alpha_n]_\epsilon \to [\alpha]_\epsilon$ in X_ϵ. Let $x = \alpha_t$ and $x_n = (\alpha_n)_t$ for all n. Then $x_n \to x$ in X. Suppose, toward a contradiction, that $[\alpha]_\epsilon$ is not a minimal ϵ-class from $*$ to x. Then we can find $[\beta]_\epsilon$ with $\beta_t = x$ and $L([\beta]_\epsilon) < L([\alpha]_\epsilon)$. Let $\tau > 0$ be such that $0 < \tau < \min\{\epsilon, L([\alpha]_\epsilon) - L([\beta]_\epsilon)\}$. Then choose n large enough so that $d_\epsilon([\alpha_n]_\epsilon, [\alpha]_\epsilon) < \tau/3$ and $|L([\alpha_n]_\epsilon) - L([\alpha]_\epsilon)| < \tau/3$. (It is straightforward to see that the length functional, $[\alpha]_\epsilon \mapsto L([\alpha]_\epsilon) = d_\epsilon(\tilde{*}, [\alpha]_\epsilon)$, is continuous on (X_ϵ, d_ϵ), making the last inequality possible.) Next, choose a representative, $\beta' \in [\beta]_\epsilon$ such that $L([\beta]_\epsilon) \leq L(\beta') < L([\beta]_\epsilon) + \tau/3$. Denote β' by $\{* = y_0, \ldots, y_m = x\}$, and let $\alpha'_n = \{* = y_0, \ldots, y_m = x, x_n\}$. Note that $d(x_n, x) \leq d([\alpha_n]_\epsilon, [\alpha]_\epsilon) < \tau/3$, so α'_n is an ϵ-chain from $*$ to x_n. Moreover, we have

$$L(\alpha'_n) = L(\beta') + d(x, x_n) < L([\beta]_\epsilon) + \frac{\tau}{3} + \frac{\tau}{3} < L([\alpha]_\epsilon) - \frac{\tau}{3}$$

$$\Rightarrow L([\alpha'_n]_\epsilon) < L([\alpha]_\epsilon) - \frac{\tau}{3}.$$

But $L([\alpha_n]_\epsilon) \leq L([\alpha'_n]_\epsilon) < L([\alpha]_\epsilon) - \tau/3$, contradicting that

$$-\frac{\tau}{3} < L([\alpha_n]_\epsilon) - L([\alpha]_\epsilon) < \frac{\tau}{3} \Rightarrow L([\alpha]_\epsilon) - \frac{\tau}{3} < L([\alpha_n]_\epsilon).$$

This proves part 1.

Next, let $[\alpha]_\epsilon \in X_\epsilon$ be given, and let $\alpha_t = x$. Let $[\beta]_\epsilon$ be a minimal ϵ-class from $*$ to x, and let $[\gamma]_\epsilon = [\alpha\beta^{-1}]_\epsilon \in \pi_\epsilon(X)$. Then $[\gamma]_\epsilon[\beta]_\epsilon = [\alpha]_\epsilon$, showing that $[\alpha]_\epsilon \in [\gamma]_\epsilon \tilde{X}_\epsilon$. Thus, the translates, $\{[\gamma]_\epsilon \tilde{X}_\epsilon\}_{[\gamma]_\epsilon \in \pi_\epsilon(X)}$, cover X_ϵ.

Finally, if X is compact, then we can set $D = diam_\epsilon(X)$, and it follows that, for any $x \in X$, there is $[\alpha]_\epsilon \in X_\epsilon$ with $\alpha_t = x$ and $L([\alpha]_\epsilon) < D + 1$. Thus, any minimal ϵ-class from $*$ to x must have length less than $D + 1$. It follows that if $[\alpha]_\epsilon \in \tilde{X}_\epsilon$, then $d_\epsilon(\tilde{*}, [\alpha]_\epsilon) = L([\alpha]_\epsilon) < D + 1$, showing that $\tilde{X}_\epsilon$ is bounded. So, by part 1, $\tilde{X}_\epsilon$ is closed and bounded in the proper space, X_ϵ. Hence, it must be compact, also, proving part 3. ∎

In light of the previous lemma, we will abuse terminology slightly and refer to $\tilde{X}_\epsilon$ as the *standard fundamental domain* for the action of $\pi_\epsilon(X)$ on X_ϵ. The accepted definition of a fundamental domain for a group action $G \times Y \to Y$ is an *open* set in Y with disjoint translates and such that the translates of the closure of the domain form a covering of Y. One can show that the set of all minimal ϵ-classes that are *unique* from $*$ to their endpoints - that is, the set of all $[\alpha]_\epsilon \in \tilde{X}_\epsilon$ with the property that if $[\beta]_\epsilon$ satisfies $\beta_t = \alpha_t$, then either $[\alpha]_\epsilon = [\beta]_\epsilon$ or $L([\alpha]_\epsilon) < L([\beta]_\epsilon)$ - is open, has disjoint translates under $\pi_\epsilon(X)$, and even satisfies the requirements of a Dirichlet domain at $\tilde{*}$. Thus, this may seem like the correct choice for a true fundamental domain. However, because of the possibility of unusual connectivity properties of X_ϵ, it is not necessarily true that the closure of this set equals $\tilde{X}_\epsilon$. Said another way, $\tilde{X}_\epsilon$ may contain isolated points, and the translates of the closure of the set of unique minimal ϵ-classes may not cover X_ϵ. This will be a particularly important property for our purposes. It should be noted, though, that if X is a compact or proper geodesic space, then $\tilde{X}_\epsilon$ is, in fact, the closure of the set of unique minimal ϵ-classes and a true fundamental domain.

This issue will not present any significant problems, however, since $\tilde{X}_\epsilon$ still possesses all of the useful properties of a fundamental domain for the action of $\pi_\epsilon(X)$ on X_ϵ. For instance, even though we will not need it, one such interesting property is the following: if we let Y be the Cartesian set product, $Y := \pi_\epsilon(X) \times \tilde{X}_\epsilon$, and if we metrize Y by

$$d_Y\big(([\gamma_1]_\epsilon, [\alpha_1]_\epsilon), ([\gamma_2]_\epsilon, [\alpha_2]_\epsilon)\big) := d_\epsilon([\gamma_1\alpha_1]_\epsilon, [\gamma_2\alpha_2]_\epsilon),$$

then Y is isometric to X_ϵ, showing that we can express each element of X_ϵ as an ordered pair consisting of a group element and a minimal ϵ-class. In other words, *every* element of X_ϵ is, intuitively, the unraveling of a chain formed by concatenating a loop with a minimal class.

Lemma 5.2.4 *Let X be a chain-connected metric space, and let $\epsilon > 0$ be such that X_ϵ is proper. Define $\mathscr{G}_\epsilon := \{[\gamma]_\epsilon \in \pi_\epsilon(X) \ : \ dist\big([\gamma]_\epsilon\tilde{X}_\epsilon, \tilde{X}_\epsilon\big) < \epsilon\}$. Then $\mathscr{G}_\epsilon^{-1} = \mathscr{G}_\epsilon$, and $\mathscr{G}_\epsilon$ generates $\pi_\epsilon(X)$. If, in addition, X is compact, then $\mathscr{G}_\epsilon$ is finite, and $\pi_\epsilon(X)$, taken with the word metric induced by any finite generating set, is bi-Lipschitz equivalent to $(\pi_\epsilon(X), d_\epsilon)$.*

We call $\mathscr{G}_\epsilon$ the ***ϵ-generating set of*** X.

Proof The fact that $\mathscr{G}_\epsilon^{-1} = \mathscr{G}_\epsilon$ follows because $\pi_\epsilon(X)$ acts by isometries and d_ϵ is left-invariant with respect to the action, which means that

$$dist([\gamma]_\epsilon\tilde{X}_\epsilon, \tilde{X}_\epsilon) = dist(\tilde{X}_\epsilon, [\gamma^{-1}]_\epsilon\tilde{X}_\epsilon).$$

Next, let $[\gamma]_\epsilon \in \pi_\epsilon(X)$ be given, and let N be the unique positive integer such that $(N-1)\epsilon/2 \le L([\gamma]_\epsilon) < N\epsilon/2$. Then there is a representative, $\gamma' \in [\gamma]_\epsilon$ with $\nu(\gamma') \le N+1$ by Corollary 2.7.4. The unique lift of γ' to $\tilde{*}$ gives us an ϵ-chain in X_ϵ from $\tilde{*}$ to $[\gamma]_\epsilon$. Denote this lifted chain by

$$\tilde{\gamma}' = \{\tilde{*} = z_1, \dots, z_m = [\gamma]_\epsilon\},$$

where $m \le N+1$. For each $i = 2, \dots, m-1$, choose $[\gamma_i]_\epsilon \in \pi_\epsilon(X)$ such that $z_i \in [\gamma_i]_\epsilon\tilde{X}_\epsilon$. For $i=1$ and $i=m$, we choose $[\gamma_1]_\epsilon = \tilde{*}$ and $[\gamma_m]_\epsilon = [\gamma]_\epsilon$. Now, let $g_i = [\gamma_{i-1}^{-1}]_\epsilon[\gamma_i]_\epsilon$ for each $i = 2, \dots, m$, and let $g_1 = \tilde{*}$. Then

$$g_1g_2\cdots g_m = \tilde{*}[\gamma_1^{-1}]_\epsilon[\gamma_2]_\epsilon[\gamma_2^{-1}]_\epsilon[\gamma_3]_\epsilon\cdots[\gamma_{m-2}^{-1}]_\epsilon[\gamma_{m-1}]_\epsilon[\gamma_{m-1}^{-1}]_\epsilon[\gamma_m]_\epsilon = [\gamma]_\epsilon.$$

Moreover, we claim that each g_i is in $\mathscr{G}_\epsilon$. Clearly $g_1 = \tilde{*}$ is. If $2 \le i \le m$, then we have

$$z_i \in [\gamma_i]\tilde{X}_\epsilon \Rightarrow [\gamma_i^{-1}]_\epsilon z_i \in \tilde{X}_\epsilon$$

and

$$z_{i-1} \in [\gamma_{i-1}]_\epsilon\tilde{X}_\epsilon \Rightarrow [\gamma_{i-1}^{-1}]_\epsilon z_{i-1} \in \tilde{X}_\epsilon.$$

It follows that

$$\begin{aligned} dist(\tilde{X}_\epsilon, g_i\tilde{X}_\epsilon) &= dist\big(\tilde{X}_\epsilon, [\gamma_{i-1}^{-1}]_\epsilon[\gamma_i]_\epsilon\tilde{X}_\epsilon\big) \\ &\le d_\epsilon\big([\gamma_{i-1}^{-1}]_\epsilon z_{i-1}, [\gamma_{i-1}^{-1}]_\epsilon[\gamma_i]_\epsilon[\gamma_i^{-1}]_\epsilon z_i\big) \\ &\le d_\epsilon(z_{i-1}, z_i) \\ &< \epsilon. \end{aligned}$$

To see that $\mathscr{G}_\epsilon$ is finite when X is compact, we first observe that, since the action is properly discontinuous (discrete, in fact) and X_ϵ is proper, we have the following result: for any compact set $K \subset X_\epsilon$, the set of all $[\gamma]_\epsilon \in \pi_\epsilon(X)$ such that $[\gamma]_\epsilon K \cap K \neq \emptyset$ is finite. Since X is compact, the previous lemma shows that $\tilde{X}_\epsilon$ is compact, also. Choose $R > 0$ so that $\tilde{X}_\epsilon$ is contained in the closed - hence, compact - ball of radius R centered at $\tilde{*}$, and consider the closed ball, $K := C(\tilde{*}, R + \epsilon)$, which is also compact. If $[\gamma]_\epsilon \in \mathscr{G}_\epsilon$, then $[\gamma]_\epsilon K \cap K \neq \emptyset$. Hence, $\mathscr{G}_\epsilon$ is finite.

Taken together, the previous two results and the compactness of X show not only that the finite set $\mathscr{G}_\epsilon$ generates $\pi_\epsilon(X)$ but that any $[\gamma]_\epsilon$ with $(N-1)\epsilon/2 \leq L([\gamma]_\epsilon) < N\epsilon/2$ can be represented as a k-fold product of elements from $\mathscr{G}_\epsilon$, where $k \leq N+1$. Let d_w be the word metric on $\pi_\epsilon(X)$ determined by $\mathscr{G}_\epsilon$. By a well-known result first due to J. Milnor (see [3] or [6]), any two word metrics on a group determined by finite generating sets are bi-Lipschitz equivalent, and bi-Lipschitz equivalence is an equivalence relation on the set of metrics on a given set. So it suffices to show that $(\pi_\epsilon(X), d_\epsilon)$ is bi-Lipschitz equivalent to $(\pi_\epsilon(X), d_w)$.

Let $[\gamma]_\epsilon \in \pi_\epsilon(X)$ be given, and let N be the unique positive integer such that $(N-1)\epsilon/2 \leq L([\gamma]_\epsilon) < N\epsilon/2$. By the first statement of the preceding paragraph, we have $d_w(\tilde{*}, [\gamma]_\epsilon) \leq N+1$, which implies that

$$\frac{\epsilon(d_w(\tilde{*}, [\gamma]_\epsilon) - 2)}{2} \leq \frac{(N-1)\epsilon}{2} \leq L([\gamma]_\epsilon) = d_\epsilon(\tilde{*}, [\gamma]_\epsilon)$$
$$\Rightarrow d_w(\tilde{*}, [\gamma]_\epsilon) \leq \frac{2}{\epsilon} d_\epsilon(\tilde{*}, [\gamma]_\epsilon) + 2.$$

So, given any $[\gamma_1]_\epsilon$, $[\gamma_2]_\epsilon \in \pi_\epsilon(X)$, and using the left-invariance of the word metric, we have

$$d_w([\gamma_1]_\epsilon, [\gamma_2]_\epsilon) \leq \frac{2}{\epsilon} d_\epsilon([\gamma_1]_\epsilon, [\gamma_2]_\epsilon) + 2.$$

Now, there cannot be non-trivial elements of $\pi_\epsilon(X)$ arbitrarily close (with respect to d_ϵ) to $\tilde{*}$, since any $[\gamma]_\epsilon \in \pi_\epsilon(X)$ satisfying $d_\epsilon(\tilde{*}, [\gamma]_\epsilon) < \epsilon$ is necessarily trivial. Thus, $d_\epsilon(\tilde{*}, [\gamma]_\epsilon) \geq \epsilon$ for all non-trivial $[\gamma]_\epsilon \in \pi_\epsilon(X)$. The left-invariance of d_ϵ, of course, implies that $d_\epsilon([\gamma_1]_\epsilon, [\gamma_2]_\epsilon) \geq \epsilon$ for all $[\gamma_1]_\epsilon \neq [\gamma_2]_\epsilon \in \pi_\epsilon(X)$. Choose $C > 0$ large enough that $C > \frac{4}{\epsilon}$. Then, if $[\gamma_1]_\epsilon \neq [\gamma_2]_\epsilon$,

$$C - \frac{2}{\epsilon} > \frac{2}{\epsilon} \geq \frac{2}{d_\epsilon([\gamma_1]_\epsilon, [\gamma_2]_\epsilon)}$$
$$\Rightarrow \left(C - \frac{2}{\epsilon}\right) d_\epsilon([\gamma_1]_\epsilon, [\gamma_2]_\epsilon) \geq 2 \;\Rightarrow\; C d_\epsilon([\gamma_1]_\epsilon, [\gamma_2]_\epsilon) \geq \frac{2}{\epsilon} d_\epsilon([\gamma_1]_\epsilon, [\gamma_2]_\epsilon) + 2,$$

and, from this, it follows that $d_w \leq C d_\epsilon$.

To prove the other direction, let $\tau = \max\{d_\epsilon(\tilde{*}, [\gamma]_\epsilon) \,:\, [\gamma]_\epsilon \in \mathscr{G}_\epsilon\}$. For each $[\gamma]_\epsilon \in \mathscr{G}_\epsilon$,

$$\frac{1}{\tau} d_\epsilon(\tilde{*}, [\gamma]_\epsilon) \leq 1 = d_w(\tilde{*}, [\gamma]_\epsilon). \tag{5.1}$$

So, the inequality $\frac{1}{\tau} d_\epsilon(\tilde{*}, [\gamma]_\epsilon) \leq d_w(\tilde{*}, [\gamma]_\epsilon)$ holds for all elements of $\pi_\epsilon(X)$ that have minimal word length equal to 1. Proceeding inductively, suppose it holds for all elements of $\pi_\epsilon(X)$ with minimal word length n for some $n \geq 1$. Let $[\gamma]_\epsilon \in \pi_\epsilon(X)$ be an element that has a minimal word representation of length $n+1$, say $[\gamma]_\epsilon = [\gamma_1]_\epsilon \cdots [\gamma_n]_\epsilon [\gamma_{n+1}]_\epsilon$, where $[\gamma_i]_\epsilon \in \mathscr{G}_\epsilon$ for each $i = 1, \ldots, n+1$. Then $\frac{1}{\tau} d_\epsilon(\tilde{*}, [\gamma]_\epsilon) = \frac{1}{\tau} d_\epsilon([\gamma_1^{-1}]_\epsilon, [\gamma_2]_\epsilon \cdots [\gamma_{n+1}]_\epsilon)$. From this, it follows that

$$\begin{aligned}
\frac{1}{\tau} d_\epsilon(\tilde{*}, [\gamma]_\epsilon) &\leq \frac{1}{\tau}\Big(d_\epsilon([\gamma_1^{-1}]_\epsilon, \tilde{*}) + d_\epsilon(\tilde{*}, [\gamma_2]_\epsilon \cdots [\gamma_{n+1}]_\epsilon) \Big) \\
&\leq 1 + d_w(\tilde{*}, [\gamma_2]_\epsilon \cdots [\gamma_{n+1}]_\epsilon) \\
&\leq 1 + n = d_w(\tilde{*}, [\gamma]_\epsilon).
\end{aligned}$$

So, by induction, inequality 5.1 holds for all $[\gamma]_\epsilon \in \pi_\epsilon(X)$. Finally, the same left-invariance argument used before shows that this inequality holds in general. Thus, we have $\frac{1}{\tau}d_\epsilon \le d_w \le C d_\epsilon$, and the two metrics are bi-Lipschitz equivalent. ∎

Corollary 5.2.5 *Under the conditions of the previous lemma - including compactness - if we let D denote the ϵ-diameter of X, then $\pi_\epsilon(X)$ is generated by $C(\tilde{*}, 2D+\epsilon) \cap \pi_\epsilon(X)$, and this set is finite. If we let r_*^ϵ denote the ϵ-radius of X at $*$, then $\pi_\epsilon(X)$ is generated by $C(\tilde{*}, 2r_*^\epsilon+\epsilon) \cap \pi_\epsilon(X)$, and this set is finite.*

The point of including both statements is that it is often easier to compute, or at least bound, the ϵ-diameter of a space than the ϵ-radius of a space at a given base point. Hence, while the second statement is stronger, the first one seems more practical.

Proof First, we note that, when X_ϵ is proper, the fact that $\pi_\epsilon(X)$ is discrete in X_ϵ implies that *any* ball in X_ϵ can only contain, at most, finitely many distinct elements of $\pi_\epsilon(X)$. We will prove this formally in the next theorem, without referring to any result of this corollary, so there is no circular reasoning being used here.

We will prove the second statement first. By definition of r_*^ϵ, we have, for any $x \in X$,

$$\inf\{L(\alpha) \,:\, \alpha \text{ is an } \epsilon\text{-chain from } * \text{ to } x\} \le r_*^\epsilon.$$

If this is a strict inequality, then we can find an ϵ-chain from $*$ to x with length strictly less than r_*^ϵ. However, if this is an equality, then - since there need not be, in general, a minimal length chain attaining this infimum - we can only conclude that, for any $\tau > 0$, there is some ϵ-chain, α, from $*$ to x such that $L(\alpha) < r_*^\epsilon + \tau$. This holds for any $x \in X$, implying that a minimal ϵ-class from $*$ to any x will satisfy $L([\alpha]_\epsilon) < r_*^\epsilon + \tau$ for all $\tau > 0$. Thus, every minimal ϵ-class in $\tilde{X}_\epsilon$ will satisfy $L([\alpha]_\epsilon) \le r_*^\epsilon$. Said another way, $\tilde{X}_\epsilon$ is a subset of the closed ball $C(\tilde{*}, r_*^\epsilon)$. If $[\gamma]_\epsilon \in \mathscr{G}_\epsilon$, then $dist(\tilde{X}_\epsilon, [\gamma]_\epsilon\tilde{X}_\epsilon) < \epsilon$, and the compactness of $\tilde{X}_\epsilon$ implies that there are points $[\alpha]_\epsilon, [\beta]_\epsilon \in \tilde{X}_\epsilon$ such that

$$d_\epsilon([\alpha]_\epsilon, [\gamma]_\epsilon[\beta]_\epsilon) = dist(\tilde{X}_\epsilon, [\gamma]_\epsilon\tilde{X}_\epsilon) < \epsilon.$$

It follows from the left-invariance of d_ϵ that

$$\begin{aligned} d_\epsilon(\tilde{*}, [\gamma]_\epsilon) &\le d_\epsilon(\tilde{*}, [\alpha]_\epsilon) + d_\epsilon([\alpha]_\epsilon, [\gamma]_\epsilon[\beta]_\epsilon) + d_\epsilon([\gamma]_\epsilon[\beta]_\epsilon, [\gamma]_\epsilon) \\ &\le r_*^\epsilon + \epsilon + d_\epsilon([\beta]_\epsilon, \tilde{*}) \\ &\le r_*^\epsilon + \epsilon + r_*^\epsilon. \end{aligned}$$

That is, $\mathscr{G}_\epsilon \subset C(\tilde{*}, 2r_*^\epsilon + \epsilon)$, proving the second statement.

Since we clearly have $r_*^\epsilon \le D$, the first statement follows, also. ∎

We now come to the main result of this chapter.

Theorem 5.2.6 *Let X be a compact, chain-connected metric space, and let $\epsilon > 0$ be given. The following are equivalent.*

1) X_ϵ is proper.

2) X has the ϵ-BMR property.

3) For any two points, $x, y \in X$, and any $L > 0$, there are only finitely many ϵ-homotopy equivalence classes of ϵ-chains connecting x and y and having length bounded above by L.

4) There are, at most, finitely many group elements, $[\gamma]_\epsilon \in \pi_\epsilon(X)$, contained in any open (or closed) ball, $B([\alpha]_\epsilon, r) \subset X_\epsilon$.

5) $\pi_\epsilon(X)$ is finitely generated, and, if d_w denotes the word metric on $\pi_\epsilon(X)$ with respect to any finite generating set, then $(\pi_\epsilon(X), d_w)$ is bi-Lipschitz equivalent to $(\pi_\epsilon(X), d_\epsilon)$.

Moreover, if any of these conditions hold, then (X_ϵ, d_ϵ) is quasi-isometric to $(\pi_\epsilon(X), d_\epsilon)$ and to $(\pi_\epsilon(X), d_w)$, where d_w denotes the word metric determined by any finite generating set.

Remark Before we prove this theorem, we need to make a few remarks. First, part 2 implies that a necessary condition for any of these properties to hold is that X be ϵ-surjective from below. This should further emphasize how important refinability is. Second, the result is stated for compact spaces, but several of the implications only require that X be proper. In the proof below, we have only used compactness where it is necessary, so the other implications can be seen to hold in the more general case. In fact, the implications $1 \Rightarrow 4$ and $4 \Rightarrow 3$ do not require X to be compact or proper. The implications $2 \Rightarrow 1$ and $3 \Rightarrow 1$ only require that X be proper. The implications $1 \Rightarrow 2$ and $1 \Rightarrow 5$ do require compactness.

Proof ($1 \Rightarrow 4$) Assume X_ϵ is proper. Suppose, towards a contradiction, that there is some ball containing infinitely many distinct elements of $\pi_\epsilon(X)$. We may assume without loss of generality that this ball is centered at $\tilde{*}$. Then, in particular, there would be a sequence of distinct elements of $\pi_\epsilon(X)$ contained in this ball, say $\{[\gamma_n]_\epsilon\} \subset \pi_\epsilon(X) \cap B(\tilde{*}, r)$. Since X_ϵ is proper, this sequence will contain a convergent subsequence, say $\{[\gamma_{n_k}]_\epsilon\}$ with $[\gamma_{n_k}]_\epsilon \to [\gamma]_\epsilon$. The ϵ-chain, γ, must necessarily be an ϵ-loop at $*$ since each γ_{n_k} is. Moreover, we know that if an ϵ-group element has length less than ϵ, then it is trivial. Thus, for all sufficiently large k, we have

$$d_\epsilon([\gamma_{n_k}]_\epsilon, [\gamma]_\epsilon) < \epsilon \Rightarrow L([\gamma_{n_k}^{-1}\gamma]_\epsilon) < \epsilon$$

$$\Rightarrow [\gamma_{n_k}^{-1}\gamma]_\epsilon = \tilde{*} \;\Rightarrow\; [\gamma_{n_k}^{-1}]_\epsilon = [\gamma]_\epsilon.$$

But this contradicts that the elements of the sequence $\{[\gamma_n]_\epsilon\}$ are all distinct.

($4 \Rightarrow 3$) Assume every ball in X_ϵ contains, at most, finitely many elements of $\pi_\epsilon(X)$. Suppose, toward a contradiction, that there are points $x, y \in X$ and infinitely many distinct homotopy classes of ϵ-chains from x to y with lengths uniformly bounded above. Then, in particular, there would be a sequence of ϵ-chains from x to y, say $\{\alpha_n\}$, and some $M > 0$, such that $[\alpha_n]_\epsilon \neq [\alpha_m]_\epsilon$ when $n \neq m$ and $L([\alpha_n]_\epsilon) \leq M$ for all n. Fix an ϵ-chain, α, from the base point, $*$, to x, and let $[\beta_n]_\epsilon = [\alpha\alpha_n]_\epsilon$. Then $[\beta_n]_\epsilon \in X_\epsilon$ for each n, and we also have

$$L([\beta_n]_\epsilon) \leq L([\alpha]_\epsilon) + L([\alpha_n]_\epsilon) \leq L([\alpha]_\epsilon) + M.$$

Thus, the lengths of the classes, $[\beta_n]_\epsilon$ are uniformly bounded above. Moreover, they are all distinct. In fact, if we had $\beta_n \sim_\epsilon \beta_m$ for some $n \neq m$, then we would have $\alpha\alpha_n \sim_\epsilon \alpha\alpha_m \Rightarrow \alpha_n \sim_\epsilon \alpha_m$, a contradiction.

Now, for each $n \geq 2$, define $\gamma_n = \beta_1\beta_n^{-1}$. This gives us a sequence of ϵ-loops at $*$. They are all homotopically distinct, because if we had $\gamma_n \sim_\epsilon \gamma_m$ for some $n \neq m$, this would imply that

$$\beta_1\beta_n^{-1} \sim_\epsilon \beta_1\beta_m^{-1} \Rightarrow \beta_n^{-1} \sim_\epsilon \beta_m^{-1} \Rightarrow \beta_n \sim_\epsilon \beta_m,$$

a contradiction. Moreover, the elements $[\gamma_n]_\epsilon$ are uniformly bounded in length, since $L([\gamma_n]_\epsilon) \leq L([\beta_1]_\epsilon) + L([\beta_n]_\epsilon) \leq 2(L([\alpha]_\epsilon) + M)$, but this is another contradiction.

$(3 \Rightarrow 1)$ Assume that 3 holds, and let $*$ be our usual base point. Then, in particular, the conclusion of 3 holds for $*$ and any $x \in X$. Let $\{[\alpha_n]_\epsilon\} \subset X_\epsilon$ be a sequence contained in some ball, $B(\tilde{*}, r)$. Then the lengths, $L([\alpha_n]_\epsilon)$, are uniformly bounded above by r. We will show that this sequence contains a convergent subsequence, thus showing that any closed ball centered at $\tilde{*}$ is compact. This, in turn, will show that any closed ball in X_ϵ is compact.

For each n, let $x_n = (\alpha_n)_t$. The sequence $\{x_n\} \subset X$ is bounded. In fact, since $\{L([\alpha_n]_\epsilon)\}$ is bounded above by r, we can find representatives, $\alpha_n' \in [\alpha_n]_\epsilon$ such that $L(\alpha_n') < r+1$ for each n. It follows from the triangle inequality that $d(*, x_n) \leq L(\alpha_n') < r+1$. So, since X is proper, there is a convergent subsequence of $\{x_n\}$, say $\{x_{n_k}\}$ with $x_{n_k} \to x$. Let $\{[\alpha_{n_k}]_\epsilon\}$ denote the corresponding subsequence of $\{[\alpha_n]_\epsilon\}$. By cutting off, if necessary, a finite initial subsegment of the subsequence $\{[\alpha_{n_k}]_\epsilon\}$, we may assume that $d(x_{n_k}, x) < \epsilon$ for all k. Then, for each $k \geq 1$, let σ_{n_k} be the ϵ-chain $\{x_{n_k}, x\}$, and let $\beta_{n_k} = \alpha_{n_k}\sigma_{n_k}$. So, $\{[\beta_{n_k}]_\epsilon\}$ is a sequence of homotopy classes of ϵ-chains from $*$ to x, and they are uniformly bounded above in length since

$$L([\beta_{n_k}]_\epsilon) \leq L([\alpha_{n_k}]_\epsilon) + L([\sigma_{n_k}]_\epsilon) \leq r + \epsilon.$$

By assumption, there are only finitely many such classes, so there must be at least one subsequence of $\{[\beta_{n_k}]_\epsilon\}$, say $\{[\beta_{n_{k_j}}]_\epsilon\}$, that is constant. That is, $[\beta_{n_{k_j}}]_\epsilon = [\beta_{n_{k_l}}]_\epsilon$ for all $j, l \geq 1$. We will show that the corresponding subsequence, $\{[\alpha_{n_{k_j}}]_\epsilon\}$, is Cauchy. Since X is proper - hence, complete, also - X_ϵ is complete, and this will imply that the Cauchy subsequence converges.

Let $\tau > 0$ be given, and choose $N \in \mathbb{N}$ so that $j \geq N \Rightarrow d(x_{n_j}, x) < \tau/2$. Fix $j, l \geq N$. Then $\beta_{n_{k_j}} \sim_\epsilon \beta_{n_{k_l}}$, which implies that

$$\alpha_{n_{k_j}}\sigma_{n_{k_j}} \sim_\epsilon \alpha_{n_{k_l}}\sigma_{n_{k_l}} \Rightarrow \sigma_{n_{k_j}}\sigma_{n_{k_l}}^{-1} \sim_\epsilon \alpha_{n_{k_j}}^{-1}\alpha_{n_{k_l}}.$$

Since $d_\epsilon([\alpha_{n_{k_j}}]_\epsilon, [\alpha_{n_{k_l}}]_\epsilon) = L([\alpha_{n_{k_j}}^{-1}\alpha_{n_{k_l}}]_\epsilon) = L([\sigma_{n_{k_j}}\sigma_{n_{k_l}}^{-1}]_\epsilon)$, this further implies

$$d_\epsilon([\alpha_{n_{k_j}}]_\epsilon, [\alpha_{n_{k_l}}]_\epsilon) \leq L(\{x_{n_{k_j}}, x, x, x_{n_{k_l}}\}) = d(x_{n_{k_j}}, x) + d(x, x_{n_{k_l}}) < \tau.$$

This shows that the sequence is Cauchy.

Next, note that the implication $1 \Rightarrow 5$ is just the previous lemma. To prove $5 \Rightarrow 1$, assume that $\pi_\epsilon(X)$ is finitely generated and that $(\pi_\epsilon(X), d_\epsilon)$ is bi-Lipschitz equivalent to $(\pi_\epsilon(X), d_w)$, where d_w represents any word metric on $\pi_\epsilon(X)$ determined by a finite generating set. Then there are constants $C, c > 0$ such that, for all $[\gamma_1]_\epsilon, [\gamma_2]_\epsilon \in \pi_\epsilon(X)$,

$$cd_\epsilon([\gamma_1]_\epsilon, [\gamma_2]_\epsilon) \leq d_w([\gamma_1]_\epsilon, [\gamma_2]_\epsilon) \leq Cd_\epsilon([\gamma_1]_\epsilon, [\gamma_2]_\epsilon).$$

Let $B(\tilde{*}, r) \subset X_\epsilon$ be any open ball centered at $\tilde{*}$. If $[\gamma]_\epsilon \in \pi_\epsilon(X)$ is in this ball, then $d_w(\tilde{*}, [\gamma]_\epsilon) \leq Cd_\epsilon(\tilde{*}, [\gamma]_\epsilon) < Cr$. In other words, the set of all $[\gamma]_\epsilon \in \pi_\epsilon(X)$ that are contained in the ball $B(\tilde{*}, r)$ is a subset of the ball $B_w(\tilde{*}, Cr)$, where B_w denotes the fact that this is a ball in the metric space $(\pi_\epsilon(X), d_w)$. Because d_w is determined by a finite generating set, there are only finitely many elements of $\pi_\epsilon(X)$ in *any* ball centered at the identity in $(\pi_\epsilon(X), d_w)$. Thus, there are only finitely many elements of $\pi_\epsilon(X)$ in the ball $B(\tilde{*}, r)$. Since r was arbitrary, it follows that there are only finitely many elements of $\pi_\epsilon(X)$ in any ball in X_ϵ. Thus, the implication 4 $\Rightarrow$ 1, which holds independently of this result, yields the desired conclusion.

$(2 \Rightarrow 1)$ It suffices to consider closed balls centered at $\tilde{*}$ and show that they are compact. Let $\{[\alpha_n]_\epsilon\}$ be a sequence contained in the closed ball $C(\tilde{*}, r)$. Then the lengths of these equivalence classes are all uniformly bounded above. Thus, by Corollary 2.7.5, we may choose a representative from each $[\alpha_n]_\epsilon$ so that these representatives all have the same number of points. By relabeling these representatives if necessary, we may assume that each α_n has the same number of points, say $m+1$. Denote α_n by $\alpha_n = \{* = x_0^n, x_1^n, \ldots, x_m^n\}$. By the bounded minimal refinement assumption, there is some $\delta < \epsilon$ and $N \in \mathbb{N}$ such that each two-point subchain, $\{x_{i-1}^n, x_i^n\}$, for any $1 \leq i \leq m$ and $n \geq 1$, can be ϵ-refined to a δ-chain consisting of N or fewer points. This means that, for any $n \geq 1$, we can ϵ-refine α_n to a δ-chain consisting of fewer than Nm points. Now, we want each refinement to have the same number of points, so, if we let $k+1 \leq Nm$ be the maximum number of points among our refined chains, and if there is a refinement with fewer than $k+1$ points, we can simply repeat the initial point of this chain a finite number of times until it has $k+1$ points. This does not change the length or the ϵ-homotopy class of the chain. For each n, let α_n' denote this refinement of α_n to a δ-chain with - if necessary - any repetitions of the initial point so that it has $k+1$ points. This gives us a sequence of δ-chains, $\{\alpha_n'\}$, such that $[\alpha_n']_\epsilon = [\alpha_n]_\epsilon$ and $\nu(\alpha_n') = k+1$ for all n. We will denote these chains by

$$\alpha_n' = \{* = z_0^n, \ldots, z_k^n = x_m^n\}.$$

Next, by Lemma 2.7.2, there is a subsequence, $\{\alpha_{n_l}'\}$, that converges pointwise to a chain, $\alpha = \{* = z_0, \ldots, z_k\}$, such that $d(z_{i-1}, z_i) \leq \delta < \epsilon$ for each $i = 1, \ldots, k$. That is, we have $z_i^{n_l} \to z_i$ as $l \to \infty$ for each $i = 0, 1, \ldots, k$. Thus, α is an ϵ-chain with the same number of points as each α_{n_l}'. We claim that $[\alpha_{n_l}]_\epsilon = [\alpha_{n_l}']_\epsilon \to [\alpha]_\epsilon$. But this just follows from Lemma 2.7.1. If we choose $\tau \leq g_\alpha$, and then choose $M \in \mathbb{N}$ so that $l \geq M$ implies $d(z_i^{n_l}, z_i) < \tau$ for each $i = 0, 1, \ldots, k$, then that lemma shows that α_{n_l}' is ϵ-homotopic to $\{* = z_0, z_1, \ldots, z_{k-1}, z_k^{n_l}\}$. This, then, implies that

$$(\alpha_{n_l}')^{-1}\alpha \sim_\epsilon \{z_k^{n_l}, z_{k-1}, \ldots, z_1, *, z_1, \ldots, z_{k-1}, z_k\} \sim_\epsilon \{z_k^{n_1}, z_k\}$$

$$\Rightarrow d_\epsilon([\alpha_{n_l}]_\epsilon, [\alpha]_\epsilon) = L([\alpha_{n_l}^{-1}\alpha]_\epsilon) \leq d(z_k^{n_l}, z_k) \to 0.$$

This yields a convergent subsequence of $\{[\alpha_n]_\epsilon\}$, showing that X_ϵ is proper.

($1 \Rightarrow 2$) This is the most technical part of the proof, and we will prove this result by contradiction. Assume X_ϵ is proper, and suppose, toward a contradiction, that X does *not* have the ϵ-BMR property. We will explicitly construct a bounded sequence in X_ϵ with no convergent subsequence, contradicting that X_ϵ is proper. Since we've already shown that X_ϵ proper $\Rightarrow$ $\pi_\epsilon(X)$ is finitely generated $\Rightarrow$ X is ϵ-surjective below, we know that there is some $\delta < \epsilon$ such that every ϵ-chain can be ϵ-refined to a δ-chain. By Lemma 2.3.2, this is equivalent to the same result for *two-point* ϵ-chains. So, we can refine all two-point ϵ-chains to δ-chains, but X does not have the ϵ-BMR property. This means that there is no pair, (λ, N), with $\lambda < \epsilon$ and $N \in \mathbb{N}$, such that all two-point ϵ-chains can be ϵ-refined to λ-chains with N or fewer points. In particular, this holds for any pair (δ, N). In other words, we have the refinement property, but no uniformly bounded refinement.

We need a quick reminder on terminology before proceeding. Recall that a *minimal refinement* of an ϵ-chain to a δ-chain is a refinement of minimal cardinality. Minimal refinements need not be unique, but the cardinality of a minimal refinement is.

Now, there must be some ϵ-chain, $\alpha_1 = \{x_1, y_1\}$, such that the cardinality of a minimal refinement of α_1 to a δ-chain consists of more than two points. If not, this would imply that X has the ϵ-BMR property. We will let $\sharp_\delta\alpha$ denote the cardinality of a minimal ϵ-refinement of α to a δ-chain. Thus, $\sharp_\delta\alpha_1 > 2$. Likewise, there must be some two point ϵ-chain, $\alpha_2 = \{x_2, y_2\}$, such that $\sharp_\delta\alpha_2 > \sharp_\delta\alpha_1 + 2 > 4$. Proceeding inductively, suppose we have two point ϵ-chains, α_i, $1 \leq i \leq n$, such that $\sharp_\delta\alpha_i > \sharp_\delta\alpha_{i-1} + 2 > 2i$ for each $i = 2, \ldots, n$. Again, since X does not have the ϵ-BMR property, there must be some ϵ-chain, $\alpha_{n+1} = \{x_{n+1}, y_{n+1}\}$, such that

$\sharp_\delta\alpha_{n+1} > \sharp_\delta\alpha_n + 2 > 2(n+1)$. This gives us a sequence of two-point ϵ-chains, $\alpha_n = \{x_n, y_n\}$, such that $\sharp_\delta\alpha_n > \sharp_\delta\alpha_{n-1} + 2 > 2n$ for $n \geq 2$.

Since X is compact, there are convergent subsequences of $\{x_n\}$ and $\{y_n\}$. By reindexing if necessary, we can simply assume that $x_n \to x$ and $y_n \to y$, and the inductive inequalities from the previous paragraph still hold. We first note that we cannot have $d(x,y) < \epsilon$. To see why, suppose this were true. Then choose $N \in \mathbb{N}$ so that

$$n \geq N \;\Rightarrow\; d(x_n, x),\; d(y_n, y) < \min\{\delta, \epsilon - d(x,y)\}.$$

We then have, for any $n \geq N$,

$$d(y_n, x) \leq d(y_n, y) + d(y, x) < \epsilon - d(x,y) + d(x,y) = \epsilon$$

$$d(x_n, y) \leq d(x_n, x) + d(x, y) < \epsilon - d(x,y) + d(x,y) = \epsilon.$$

The first inequality implies that we can insert x into $\alpha_n = \{x_n, y_n\}$ between x_n and y_n. The assumption that $d(x,y) < \epsilon$ and $d(y_n, y) < \delta < \epsilon$ then implies that we can insert y between x and y_n. In other words, α_n is ϵ-homotopic to $\{x_n, x, y, y_n\}$. Since $d(x_n, x), d(y_n, y) < \delta$, and since we could, by hypothesis, refine $\{x, y\}$ (*if* $d(x,y) < \epsilon$) to a δ-chain with, say, p points, this means that we could ϵ-refine α_n to a δ-chain consisting of $p+2$ points, and this holds for any $n \geq N$. But this contradicts our conclusion that $\sharp_\delta\alpha_n \nearrow \infty$. So, $d(x,y) \geq \epsilon$, and, in particular, $x \neq y$. On the other hand, the fact that $x_n \to x$, $y_n \to y$ implies that $d(x,y) \leq \epsilon$. Thus, we have $d(x,y) = \epsilon$.

Next, for reasons that will become clear soon, we want each x_n to be within δ of x and every other x_m, and we want the analogous conclusion for $\{y_n\}$ and y. So, we choose $N \in \mathbb{N}$ so that $n \geq N \Rightarrow d(x_n, x), d(y_n, y) < \delta/2 < \epsilon$. Then, for every $n, m \geq N$, we also have $d(x_n, x_m) < \delta < \epsilon$ and $d(y_n, y_m) < \delta < \epsilon$. For simplicity of notation, we then truncate our sequence at N and reindex again (so $N \sim 1$, $N+1 \sim 2$, etc.). This, finally, gives us a sequence of two-point ϵ-chains, $\alpha_n = \{x_n, y_n\}$, with the following properties: **1)** $\sharp_\delta\alpha_n > \sharp_\delta\alpha_{n-1} + 2 > 2n$ for all $n \geq 2$, **2)** $x_n \to x$ and $y_n \to y$, **3)** $d(x,y) = \epsilon$, **4)** $d(x_n, x), d(y_n, y) < \delta/2$ for all n, and **5)** $d(x_n, x_m), d(y_n, y_m) < \delta$ for all $n, m \geq 1$. This yields a construction we call an **ϵ-ladder**; see Figure 5.1 below.

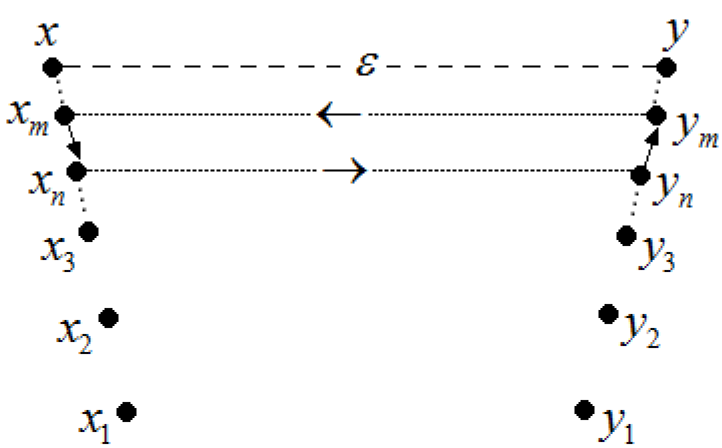

Figure 5.1: An ϵ-ladder

We now claim that, for any $1 \leq n < m$, the ϵ-chains $\alpha_m = \{x_m, y_m\}$, $\alpha_n^m := \{x_m, x_n, y_n, y_m\}$ are not ϵ-homotopic, or, equivalently, that the ϵ-loop $\{x_m, x_n, y_n, y_m, x_m\}$ is ϵ-nontrivial. (See Figure 5.1.) To see why, fix $1 \leq n < m$, and suppose we had an ϵ-homotopy transforming α_m to α_n^m. Then, by concatenating this homotopy with an ϵ-homotopy that ϵ-refines $\alpha_n = \{x_n, y_n\}$ to a minimal δ-chain, this would give us an ϵ-refinement of α_m to a δ-chain with $\sharp_\delta \alpha_n + 2$ points. This is why we required $d(x_m, x_n), d(y_n, y_m) < \delta$ in the previous paragraph. But this would imply that $\sharp_\delta \alpha_m \leq \sharp_\delta \alpha_n + 2$ with $m > n$, a contradiction. Hence, α_m is not ϵ-homotopic to α_n^m.

Finally, we can construct a bounded sequence in X_ϵ that has no convergent subsequence. Since X is chain-connected, the ϵ-covers for different base points are all isometric. In particular, the condition that X is ϵ-surjective from below, or that the map $X_\delta \to X_\epsilon$ is surjective, is independent of what base point we choose. Thus, we can choose whatever base point is most convenient. In this case, we will choose y to be our base point, and we will let X_ϵ be the ϵ-cover determined by homotopy equivalence classes of ϵ-chains beginning at y. If the cover, X_ϵ, determined by y is not proper, then the ϵ-cover determined by any other base point will also not be proper.

For each $n \geq 1$, let $\beta_n = \{y, y_n, x_n, x\}$. Then $[\beta_n]_\epsilon \in X_\epsilon$ for each n, and we further have

$$d_\epsilon(\tilde{y}, [\beta_n]_\epsilon) = L([\beta_n]_\epsilon) \leq L(\beta_n) < \frac{\delta}{2} + \epsilon + \frac{\delta}{2} < \epsilon + \delta.$$

Thus, $\{[\beta_n]_\epsilon\}$ is a bounded sequence, and, in fact, it is contained in the closed ball $C(\tilde{y}, 2\epsilon)$. Consider $d_\epsilon([\beta_n]_\epsilon, [\beta_m]_\epsilon) = L([\beta_n^{-1}\beta_m]_\epsilon)$ for any $1 \leq n < m$. The chain $\beta_n^{-1}\beta_m$ is given by $\{x, x_n, y_n, y, y, y_m, x_m, x\}$, which can be transformed via an ϵ-homotopy as follows.

$$\beta_n^{-1}\beta_m \to \{x, x_n, y_n, y, y_m, x_m, x\} \to \{x, x_n, y_n, y_m, x_m, x\}$$

$$\to \{x, x_m, x_n, y_n, y_m, x_m, x\}.$$

Note that this last loop is just the loop $\{x_m, x_n, y_n, y_m, x_m\}$ pulled back - as in the proof of base point independence of the isomorphism class of the ϵ-groups, Lemma 2.2.9 - to be anchored at x. But, we just showed that the loop $\{x_m, x_n, y_n, y_m, x_m\}$ is non-trivial. Thus, the loop $\{x, x_m, x_n, y_n, y_m, x_m, x\}$ is non-trivial. It follows that $\beta_n^{-1}\beta_m$ is non-trivial, and $L([\beta_n^{-1}\beta_m]_\epsilon) \geq \epsilon$. This shows that for any $1 \leq n < m$ we have $d_\epsilon([\beta_n]_\epsilon, [\beta_m]_\epsilon) \geq \epsilon$. This, in turn, shows that $\{[\beta_n]_\epsilon\}$ can have no convergent subsequence. If it did, say $\{[\beta_{n_k}]_\epsilon\}$, then this subsequence would be Cauchy, implying that we could find two distinct indices so that $d_\epsilon([\beta_{n_k}]_\epsilon, [\beta_{n_{k'}}]_\epsilon) < \epsilon$, which - as we just showed - cannot occur. This contradicts our assumption that X_ϵ is proper.

Lastly, we will prove the final statement: when properties 1 through 5 hold, (X_ϵ, d_ϵ) is quasi-isometric to $(\pi_\epsilon(X), d_\epsilon)$ and to $(\pi_\epsilon(X), d_w)$, where d_w is the word metric determined by any finite generating set. By assumption, $\pi_\epsilon(X)$ is finitely generated, and $(\pi_\epsilon(X), d_\epsilon)$ is bi-Lipschitz equivalent to - hence, quasi-isometric to - $(\pi_\epsilon(X), d_w)$, where d_w is the word metric determined by any finite generating set. So, since "quasi-isometric" is an equivalence relation, we need only show that $(\pi_\epsilon(X), d_\epsilon)$ is quasi-isometric to (X_ϵ, d_ϵ). This means we need to show two things: **1)** there is a quasi-isometry, $f : (\pi_\epsilon(X), d_\epsilon) \to (X_\epsilon, d_\epsilon)$ - meaning that there are constants $\lambda > 0$, $C \geq 0$ such that

$$\frac{1}{\lambda} d_\epsilon([\gamma_1]_\epsilon, [\gamma_2]_\epsilon) - C \leq d_\epsilon(f([\gamma_1]_\epsilon), f([\gamma_2]_\epsilon)) \leq \lambda d_\epsilon([\gamma_1]_\epsilon, [\gamma_2]_\epsilon) + C$$

for all $[\gamma_1]_\epsilon, [\gamma_2]_\epsilon \in \pi_\epsilon(X)$ - and **2)** that there is some constant, $D \geq 0$, such that every point of X_ϵ is within D of some point of $f(\pi_\epsilon(X))$ (i.e. $f(\pi_\epsilon(X))$ is a *D-net* in (X_ϵ, d_ϵ)). But this is essentially trivial. Define $f : \pi_\epsilon(X) \to X_\epsilon$ to be the inclusion map. Then, when we take the metric d_ϵ on $\pi_\epsilon(X)$, f is an isometric embedding. So, it is trivially a quasi-isometry. For the last part, let D be the radius at $\tilde{*}$ of $\tilde{X}_\epsilon$ with respect to d_ϵ. That is,

$$D = \sup_{[\alpha]_\epsilon \in \tilde{X}_\epsilon} d_\epsilon(\tilde{*}, [\alpha]_\epsilon).$$

Since X is compact, $\tilde{X}_\epsilon$ is compact, also, and D is, therefore, finite. Given any $[\alpha]_\epsilon \in X_\epsilon$ with $x = \alpha_t$, let $[\beta]_\epsilon$ be a minimal ϵ-class from $*$ to x, and let $[\gamma]_\epsilon = [\alpha\beta^{-1}]_\epsilon$. Then $[\gamma]_\epsilon \in \pi_\epsilon(X) = f(\pi_\epsilon(X))$ and

$$\begin{aligned} d_\epsilon([\gamma]_\epsilon, [\alpha]_\epsilon) = L([\gamma^{-1}\alpha]_\epsilon) = L([\beta\alpha^{-1}\alpha]_\epsilon) = L([\beta]_\epsilon) \\ \Rightarrow d_\epsilon([\gamma]_\epsilon, [\alpha]_\epsilon) = d_\epsilon(\tilde{*}, [\beta]_\epsilon) \le D. \end{aligned}$$

Thus, $f(\pi_\epsilon(X))$ is a D-net in X_ϵ, and this concludes the proof. ∎

In proving the implication $1 \Rightarrow 2$ in the previous proof, note that we built up a structure (i.e. the ϵ-ladder) exactly like that in Example 4.2.3. Moreover, this proof showed that ϵ-surjectivity from below is a necessary condition for X_ϵ to be proper. In other words, the space, X, in Example 4.2.3 is such that the cover X_1 is not proper, since that space is not 1-surjective from below. The construction of the ϵ-ladder in the previous proof justifies - to some extent - the remarks made in the paragraph just prior to Example 4.2.3, namely that a necessary condition for a space to not be ϵ-surjective from below is that it contain an ϵ-ladder. We technically did not prove that claim, but its proof is carried out in a manner very similar to the construction above. Hence, this phenomenon of being unable to refine chains can only occur in a very specific way. This should be further motivation of the idea that refinability is a common property among metric spaces.

The last statement of the preceding theorem, that $\pi_\epsilon(X)$ is quasi-isometric to X_ϵ, is an important one, particularly when X_ϵ and $\pi_\epsilon(X)$ are non-compact. The class of non-compact metric spaces is essentially the only context in which quasi-isometries play a significant role. This can be observed by noting that any bounded metric space is quasi-isometric to a one-point space. Given that "quasi-isometric" is an equivalence relation, it follows that all bounded (e.g. compact) metric spaces are quasi-isometric. In the non-compact case, however, it can be a useful tool. (There are two very good discussions of quasi-isometries in [3] and [5].) Roughly speaking, non-compact quasi-isometric spaces have similar large-scale geometry, and there are many useful results in this direction. For instance, any two left-invariant Riemannian metrics on a connected real Lie group, G, are quasi-isometric, and the fundamental group of a compact Riemannian manifold - which must be finitely generated - is quasi-isometric, with its word metric, to the universal Riemannian covering space. In our context, the previous results mean that we can study large-scale geometric properties of X_ϵ by studying corresponding properties of $\pi_\epsilon(X)$, and vice versa. The following result provides a simple illustration of this idea.

Lemma 5.2.7 *Let X be a compact, chain-connected metric space, and assume that X_ϵ is proper. Then X_ϵ is compact if and only if $\pi_\epsilon(X)$ is finite.*

Proof First, if X_ϵ is compact, then implication $1 \Rightarrow 4$ in the previous theorem implies that $\pi_\epsilon(X)$ is finite. Conversely, assume that $\pi_\epsilon(X)$ is finite, and suppose, toward a contradiction, that X_ϵ was not bounded. Then $\pi_\epsilon(X)$ could not be a D-net in X_ϵ for any D, since we could find points of X_ϵ arbitrarily far away from points in $\pi_\epsilon(X)$. Thus, X_ϵ must be bounded, and, since it is proper, it must also be compact. ∎

Chapter 6

Convergence of ϵ-covers and Critical Values

In this chapter, we will study the behavior of ϵ-covers and critical values under Gromov-Hausdorff convergence. Convergence of covering spaces is an important tool in the study of topology of singular spaces which can be expressed as Gromov-Hausdorff limits of simpler spaces. On the other hand, convergence of covering spaces, in general, can pose some tricky problems, as well. Take, for instance, a sequence of geodesic tori, $T_n = S^1 \times S^1_{1/n}$, where the circumference of one circle remains fixed while the circumference of the other shrinks to 0. The spaces, T_n, converge in the Gromov-Hausdorff sense to S^1. The universal cover of each T_n is $\mathbb{R}^2$, so these covers trivially converge to $\mathbb{R}^2$. However, not only is the limit of these covers *not* the universal cover of the limit space, it is not a cover of the limit space at all. We will see, though, that the ϵ-covers provide a context in which convergence of covers is, intuitively, stable and well-behaved.

We will split this material up into four sections. In the first section, we will collect some more results on Gromov-Hausdorff convergence to go along with the background material in Chapter 1. In particular, we will collect some basic results on convergence of points and functions that will make our results easier to describe.

In the second section, we will assume that we have a sequence of compact, chain-connected metric spaces, $\{X_n\}$, converging to a compact metric space, X, and we will further assume that the ϵ-covers, $(X^n_\epsilon, \tilde{*}_n)$, converge in the pointed sense to a space $(Y, *)$. We will, then, derive a structure theorem for the limit, Y, of the ϵ-covers. In the third section, we will derive conditions under which the covers $(X_{n,\epsilon}, \tilde{*}_n)$, do, in fact, converge in the pointed sense to a space, $(Y, *)$. Lastly, we will investigate the behavior of the critical spectrum, itself, under convergence.

6.1 More on Gromov-Hausdorff Convergence

Let $\{X_n\}$ be a sequence of compact metric spaces converging, in the Gromov-Hausdorff sense, to a compact space, X. We will only be discussing Gromov-Hausdorff convergence in these sections, so we will, henceforth, cease using the phrase “in the Gromov-Hausdorff sense” when talking about convergence of spaces. All convergence will be in this sense.

First, we make the following observation. Given any $x \in X$, we can find a sequence of points, $\{x_n\}$, with $x_n \in X_n$, such that $x_n \to x$ in the following sense. We can choose a strictly increasing sequence of natural numbers, $\{N_k\}_{k=1}^{\infty}$, such that $n \geq N_k \Rightarrow d_{GH}(X_n, X) < 1/k$. This means that, in some metric space, Z, into which we isometrically imbed X_n and X, X_n

lies in the $1/k$-neighborhood of X and X lies in the $1/k$-neighborhood of X_n. So, for each $N_k \le n < N_{k+1}$, we choose a point, $x_n \in X_n$ so that $d(x, x_n) < 1/k$. We clearly, then, have $x_n \to x$. Hence, given a convergent sequence of metric spaces, $X_n \to X$, and a point in X, we will often use this observation without comment. Of course, the "convergence" we have described seems to be a loosely defined concept, since we have not really defined the space in which the convergence is taking place. If we recall our discussion in Chapter 1, however, of Gromov-Hausdorff convergence, particularly the fact that one can assume that the spaces X_n and X are all imbedded in a common metric space, Z, and the convergence is in the Hausdorff sense, then we do have a somewhat firm notion of what this convergence means. It is not entirely firm since the space, Z, need not be unique. Nevertheless, we do have a clear idea of what it means for $\{x_n\}$ to converge to $x \in X$, and we will denote this phenomenon by $x_n \to x$ or $d(x_n, x) \to 0$, where the metric, d, is the metric on any space, Z, in which we imbed the spaces X_n and X so that appropriate neighborhoods of each space contain the other(s).

Remark We will frequently make this slight abuse of notation, using d to represent the metric on a space, Z, into which we imbed two metric spaces, X and Y, and we will almost always do so without explicitly mentioning Z. This should not cause any confusion, since the context should be clear in each situation. However, it should be kept in the back of the reader's mind that whenever we discuss two or more metric spaces in the context of Gromov-Hausdorff convergence and we use a single metric to describe the metrics on all of them, we are implicitly assuming that they are all isometrically imbedded in a common metric space, Z, and that d is the metric on Z. Since the imbeddings are isometric, the restriction of d to the image, in Z, of any of the spaces in question agrees with the given metric on that space. This justifies using this single metric to describe the metrics on all of the spaces being imbedded.

Extending the previous idea, we make another observation: given two compact metric spaces, X and Y, such that $d_{GH}(X, Y) < \tau$, and given an ϵ-chain, α, in X, we can find a corresponding $(\epsilon + 2\tau)$-chain, α', in Y that is "close" to α. Letting $\alpha = \{x_1, \dots, x_n\}$, we can, as described in the previous paragraph, find, for each $i = 1, \dots, n$, a point $y_i \in Y$ such that $d(x_i, y_i) < \tau$. Letting $\alpha' = \{y_1, \dots, y_n\}$, we see that α' is an $(\epsilon + 2\tau)$-chain, since, for each i,

$$d(y_{i-1}, y_i) \le d(y_{i-1}, x_{i-1}) + d(x_{i-1}, x_i) + d(x_i, y_i) < \tau + \epsilon + \tau = \epsilon + 2\tau.$$

Lemma 6.1.1 *Let $\{X_n\}$ be a sequence of compact, chain-connected metric spaces converging to a compact metric space, X. Then X is chain-connected, and, for any $\epsilon > 0$, $\{diam_\epsilon(X_n)\} \cup \{diam_\epsilon(X)\}$ is uniformly bounded above.*

Proof Let $x, y \in X$ and $\epsilon > 0$ be given. Choose n large enough so that $d_{GH}(X_n, X) < \frac{\epsilon}{4}$, then choose points $x_n, y_n \in X_n$ so that $d(x_n, x), d(y_n, y) < \frac{\epsilon}{4}$. Since each X_n is chain-connected, we can join x_n and y_n by an $\frac{\epsilon}{2}$-chain, say $\alpha = \{x_n = z_1, \dots, z_m = y_n\}$. Then, applying the observation preceding this lemma, we can choose an $(\frac{\epsilon}{2} + 2\frac{\epsilon}{4})$-chain in X between x and y. Since $\epsilon > 0$ was arbitrary, we see that X is chain-connected.

Next, since $\{X_n\}$ is Cauchy, there is some $K \in \mathbb{N}$ such that $k \ge K \Rightarrow d_{GH}(X_K, X_k) < \frac{\epsilon}{4}$. Since X_K is compact and chain-connected, Lemma 2.2.13 implies that there is some natural number, M, with the property that every two points in X_K can be joined by an $\frac{\epsilon}{2}$-chain of M or fewer points. Now, fix $k \ge K$, and choose any $x, y \in X_k$. Then choose $\bar{x}, \bar{y} \in X_K$ so that $d(x, \bar{x}), d(y, \bar{y}) < \frac{\epsilon}{4}$. There is an $\frac{\epsilon}{2}$-chain in X_K from $\bar{x}$ to $\bar{y}$ that contains M or fewer points. Reasoning as before, we can find an ϵ-chain in X_k from x to y consisting of M or fewer points.

Such a chain must have length less than or equal to $(M-1)\epsilon$, showing that there is an ϵ-chain of length less than or equal to $(M-1)\epsilon$ joining any pair of points in X_k. Thus, we must have $diam_\epsilon(X_k) \leq (M-1)\epsilon$. Since M did not depend on k, and since $k \geq K$ was arbitrary, this shows that $\{diam_\epsilon(X_k)\}_{k=K}^\infty$ is uniformly bounded above. The same method of proof shows that $diam_\epsilon(X)$ is bounded, also. ∎

The next result is primarily important in the non-compact case. We will need this result in the following section, but its proof - a standard diagonalization argument - is already known in the context of Gromov-Hausdorff convergence. Thus, we only state it here.

Lemma 6.1.2 *Let $\{(X_n, *_n)\}$ be proper metric spaces converging in the pointed sense to a complete metric space, $(X, *)$. Then X is proper.*

The last preliminary results we need before proceeding to the next section deal with convergence of maps in the Gromov-Hausdorff context. That is, suppose we have convergent sequences $Y_n \to Y$ and $X_n \to X$ and, for each n, a map $f_n : Y_n \to X_n$. In what sense, and under what conditions, can we say that $\{f_n\}$ converges to a map $f : Y \to X$? The following definitions and lemma can be found in the discussion of convergence given by P. Petersen in [7].

Definition 6.1.3 *Let $\{(X_n, *_n)\}$ and $\{(Y_n, \tilde{*}_n)\}$ be sequences of metric spaces converging in the pointed sense to $(X, *)$ and $(Y, \tilde{*})$, respectively. Suppose we have maps, $f_n : Y_n \to X_n$, for each n, and assume, further, that these maps are pointed, meaning that $f_n(\tilde{*}_n) = *_n$ for each n. We say that $\{f_n\}$ converges to the map $f : Y \to X$ if for every sequence $y_n \in Y_n$ such that $y_n \to y \in Y$, we have $f_n(y_n) \to f(y)$. In particular, we must have $f(\tilde{*}) = *$.*

Remark Note that this definition is making use of the conventions we discussed at the beginning of this section regarding convergence of points within a convergent sequence of spaces.

Definition 6.1.4 *A sequence of functions, $\{f_n\}$, as in the previous definition is said to be **equicontinuous** if for every $\epsilon > 0$ there is a $\delta > 0$ such that $f_n(B(x_n, \delta)) \subset B(f_n(x_n), \epsilon)$ for all n and $x_n \in X_n$.*

One consequence of this definition of which we will make use is the fact that if each function, f_n, in this sequence is Lipschitz - with the same Lipschitz constant - then that sequence is equicontinuous (Chapter 10, [7]).

The final result is an analog of the Arzela-Ascoli theorem in this context. The proof of this result is also found in Chapter 10 of [7].

Lemma 6.1.5 *Let $\{(X_n, *_n)\}$ and $\{(Y_n, \tilde{*}_n)\}$ be sequences of metric spaces converging in the pointed sense to $(X, *)$ and $(Y, \tilde{*})$, respectively. Suppose we have maps $f_n : Y_n \to X_n$ for each n, and assume, further, that these maps are pointed, meaning that $f_n(\tilde{*}_n) = *_n$ for each n. If the sequence $\{f_n\}$ is equicontinuous, then there is a subsequence, $\{f_{n_k}\}$, that converges to a pointed map $f : (Y, \tilde{*}) \to (X, *)$.*

6.2 Structure of the Limit Cover

Throughout this section, we assume the following situation: we have a sequence of compact, chain-connected, pointed metric spaces, $\{(X_n, *_n)\}$, converging in the pointed sense to a compact metric space, $(X,*)$, and, for $\epsilon > 0$, the corresponding sequence of ϵ-covers - which we denote by $\{(X^n_\epsilon, \tilde{*}_n)\}$ - converges in the pointed sense to a complete metric space, $(Y, \bar{*})$. For each n, we assume that $*_n$ is our chosen base point and that $\tilde{*}_n$, as usual, is the lifted base point in X^n_ϵ containing the trivial chain, $\{*_n\}$. The pointed convergence, then, implies that $*_n \to *$ and $\tilde{*}_n \to \bar{*}$. In addition, as one might imagine, it is difficult to obtain any meaningful results without assuming that the covers, X^n_ϵ, are proper. So, we assume this, as well; in particular, each X_n is ϵ-surjective from below. Furthermore, we denote, for each n, the ϵ-covering map by $\varphi^n_\epsilon : X^n_\epsilon \to X_n$, and these maps are pointed, taking $\tilde{*}_n$ to $*_n$. To avoid cumbersome repetition, we will simply assume these conditions for the rest of this section.

The first things we need to establish are convergence of the maps, φ^n_ϵ, and the passage of the local properties of these maps to the limit map.

Lemma 6.2.1 *There is a subsequence of the ϵ-covering maps, $\{\varphi^{n_k}_\epsilon\}$, converging to a pointed, 1-Lipschitz map $\varphi : (Y, \bar{*}) \to (X, *)$. Moreover, φ satisfies the following properties.*

1) *For any $\delta \le \epsilon$ and $y \in Y$, φ is a bijection and radial isometry from $B(y, \delta)$ onto $B(f(y), \delta)$.*

2) *φ is an isometry from $\epsilon/2$-balls in Y onto $\epsilon/2$-balls in X.*

3) *Given $x \in X$, two distinct points $y_1, y_2 \in \varphi^{-1}(x)$ satisfy $d_Y(y_1, y_2) \ge \epsilon$.*

4) *$\varphi : Y \to X$ is a surjective covering map.*

Proof The ϵ-covering maps, φ^n_ϵ, are 1-Lipschitz, so they are equicontinuous according to Definition 6.1.4. Hence, by Lemma 6.1.5, there is a convergent subsequence $\varphi^{n_k}_\epsilon \to \varphi$, $\varphi : (Y, \bar{*}) \to (X, *)$. The subsequences $\{(X_{n_k}, *_{n_k})\}$ and $\{(X^{n_k}_\epsilon, \tilde{*}_{n_k})\}$ will still converge to $(X,*)$ and $(Y, \bar{*})$, respectively, so, for simplicity of notation, we will simply reindex and assume that $\varphi^n_\epsilon \to \varphi$. Now, the fact that the local properties of the maps φ^n_ϵ carry over to φ follows from a series of standard arguments all using the same basic reasoning. To give the reader an idea of the process, we will show that φ is 1-Lipschitz and surjective. Properties 1 through 3 follow in exactly the same manner as these results will be proved, since each φ^n_ϵ possesses those same properties. Property 4 is, then, a consequence of 2 and 3. (In fact, 3 follows from 1.)

Let $x, y \in Y$ be given. Choose sequences, $\{x_n\}$ and $\{y_n\}$, such that $x_n, y_n \in X^n_\epsilon$, $x_n \to x$, and $y_n \to y$. For all $n \ge 1$, we have $d(\varphi^n_\epsilon(x_n), \varphi^n_\epsilon(y_n)) \le d(x_n, y_n)$, and we also have $\varphi^n_\epsilon(x_n) \to \varphi(x)$ and $\varphi^n_\epsilon(y_n) \to \varphi(y)$. Let $\tau > 0$ be given, and choose n large enough that $d(x_n, x)$, $d(y_n, y)$, $d(\varphi^n_\epsilon(x_n), \varphi(x))$, and $d(\varphi^n_\epsilon(y_n), \varphi(y))$ are all less than $\frac{\tau}{4}$. Then,

$$\begin{aligned} d(\varphi(x), \varphi(y)) &\le d(\varphi(x), \varphi^n_\epsilon(x_n)) + d(\varphi^n_\epsilon(x_n), \varphi^n_\epsilon(y_n)) + d(\varphi^n_\epsilon(y_n), \varphi(y)) \\ &< \frac{\tau}{2} + d(x_n, y_n) \\ &< \frac{\tau}{2} + d(x_n, x) + d(x, y) + d(y, y_n) \\ &< \tau + d(x, y). \end{aligned}$$

Thus, since τ was arbitrary, we have $d(\varphi(x), \varphi(y)) \le d(x, y)$, and φ is 1-Lipschitz.

We next show that φ is surjective. Let $x \in X$ be given, and let $D > 0$ be such that $diam_\epsilon(X)$ and $\sup_{n \ge 1}\{diam_\epsilon(X_n)\}$ are strictly less than D. Choose a subsequence, $\{X_{n_k}\}$, so that

$$d_{GH}((X_{n_k}, *_{n_k}), (X, *)) < \frac{1}{k} \quad \text{and} \quad d_{GH}(C(\tilde{*}_{n_k}, D), C(\bar{*}, D)) < \frac{1}{k}.$$

For each k, choose $x_k \in X_{n_k}$ so that $d(x, x_k) < \frac{1}{k}$. Then we have $x_k \to x$. Also for each k, there is some $[\alpha_k]_\epsilon \in X_\epsilon^{n_k}$ so that $(\alpha_k)_t = x_k$ and $L([\alpha_k]_\epsilon) < D$. Thus, $[\alpha_k]_\epsilon \in C(\tilde{*}_{n_k}, D)$, and we can, for each k, choose $y_k \in C(\bar{*}, D)$ so that $d([\alpha_k]_\epsilon, y_k) < \frac{1}{k}$. Now, since the covers X_ϵ^n are proper and Y is complete, it follows from Lemma 6.1.2 that Y is proper, also. Since $\{y_k\}$ is a bounded sequence in Y, it has a convergent subsequence, say $\{y_{k_j}\}$ with $y_{k_j} \to y$. We have

$$d(y, [\alpha_{k_j}]_\epsilon) \leq d(y, y_{k_j}) + d(y_{k_j}, [\alpha_{k_j}]_\epsilon) \leq d(y, y_{k_j}) + \frac{1}{k_j}.$$

Letting $j \to \infty$, we see that $[\alpha_{k_j}]_\epsilon \to y$, which means that $\varphi_\epsilon^{n_{k_j}}([\alpha_{k_j}]_\epsilon) \to \varphi(y)$. But $\varphi_\epsilon^{n_{k_j}}([\alpha_{k_j}]_\epsilon) = x_{k_j}$ and $x_{k_j} \to x$. Thus, $x = \varphi(y)$, and φ is surjective.

The rest of the properties follow similarly by taking sequences of points in the spaces X_ϵ^n and X_n and using the properties of the maps φ_ϵ^n to verify that they also hold for $\varphi : Y \to X$. ∎

Thus, $\varphi : Y \to X$ is a covering map with the same essential properties as an ϵ-covering map. We next need to know that we can lift ϵ-chains and homotopies from X to Y.

Lemma 6.2.2 *Any ϵ-chain in X beginning at x has a unique lift to any point in $\varphi^{-1}(x) \subset Y$. In addition, if α and β are homotopic ϵ-chains in X beginning at x, then their lifts to any point in $\varphi^{-1}(x)$ end at the same point and are ϵ-homotopic.*

Proof Since $\varphi : Y \to X$ is a bijection from ϵ-balls in Y onto ϵ-balls in X, it follows from Lemma 2.4.3 that we have unique lifts of ϵ-chains in X to Y. The fact that ϵ-homotopies lift will follow from the same lemma once we verify that ϵ-triangles in X lift to ϵ-triangles in Y.

Let $\gamma := \{a, b, c, a\}$ be an ϵ-triangle in X, and let $\bar{a}$ be any point in $\varphi^{-1}(a) \subset Y$. There is some $\delta < \epsilon$ such that γ is actually a δ-triangle. Choose $R > d_Y(\bar{*}, \bar{a}) + \epsilon$, and then choose subsequences, $\{(X_{n_k}, *_{n_k})\}$ and $\{(C(\tilde{*}_{n_k}, R), \tilde{*}_{n_k})\}$, so that

$$d_{GH}\Big((X_{n_k}, *_{n_k}), (X, *)\Big) < \frac{1}{k} \text{ and } d_{GH}\Big((C(\tilde{*}_{n_k}, R), \tilde{*}_{n_k}), (C(\bar{*}, R), \bar{*})\Big) < \frac{1}{k}.$$

We can then choose points $\tilde{a}_k \in C(\tilde{*}_{n_k}, R)$ so that $d(\tilde{a}_k, \bar{a}) < \frac{1}{k}$, from which it follows that $\varphi_\epsilon^{n_k}(\tilde{a}_k) \to \varphi(\bar{a}) = a$. Let $a_k = \varphi_\epsilon^{n_k}(\tilde{a}_k) \in X_{n_k}$, so that $a_k \to a$. By choosing another subsequence, if necessary, and reindexing, we may assume that $d(a, a_k) < \frac{1}{k}$.

Now, choose K large enough that $1/K < (\epsilon - \delta)/4$. Then, for each $k \geq K$, choose $b_k, c_k \in X_{n_k}$ so that $d(b_k, b), d(c_k, c) < 1/k$. Note that

$$d(a_k, b_k) \leq d(a_k, a) + d(a, b) + d(b, b_k) < \delta + 2\frac{\epsilon - \delta}{4} < \epsilon,$$

and likewise for $d(a_k, c_k)$ and $d(b_k, c_k)$. Thus, for each $k \geq K$, $\gamma_k := \{a_k, b_k, c_k, a_k\}$ is a $[\delta + (\epsilon - \delta)/2]$-triangle. This triangle will lift isometrically to a $[\delta + (\epsilon - \delta)/2]$-triangle at $\tilde{a}_k \in C(\tilde{*}_{n_k}, R)$, say $\tilde{\gamma}_k := \{\tilde{a}_k, \tilde{b}_k, \tilde{c}_k, \tilde{a}_k\}$. Moreover, for large enough k, this triangle will lie inside $C(\tilde{*}_{n_k}, R)$, since, for k large enough that $2/k < (\epsilon - \delta)/2$, we have

$$\begin{aligned} d(\tilde{*}_{n_k}, \tilde{b}_k) &\leq d(\tilde{*}_{n_k}, \bar{*}) + d(\bar{*}, \bar{a}) + d(\bar{a}, \tilde{a}_k) + d(\tilde{a}_k, \tilde{b}_k) \\ &< d(\bar{*}, \bar{a}) + \frac{2}{k} + \delta + \frac{\epsilon - \delta}{2} \\ &< d(\bar{*}, \bar{a}) + \epsilon. \end{aligned}$$

and likewise for $d(\tilde{*}_{n_k}, \tilde{c_k})$. We can simply assume that K is also large enough that each $\tilde{\gamma}_k$ lies in $C(\tilde{*}_{n_k}, R)$.

Now, for each $k \geq K$, choose $\bar{b}_k, \bar{c}_k \in C(\bar{*}, R)$ so that $d(\tilde{b}_k, \bar{b}_k), d(\tilde{c}_k, \bar{c}_k) < 1/k$. Since Y is proper, each sequence, $\{\bar{b}_k\}$ and $\{\bar{c}_k\}$, will have a convergent subsequence, and we can assume that the subsequences have corresponding indices, $\{\bar{b}_{k_j}\}$ and $\{\bar{c}_{k_j}\}$. Let $\bar{b} = \lim_j \bar{b}_{k_j}$, $\bar{c} = \lim_j \bar{c}_{k_j}$, and note that

$$d(\bar{a}, \bar{b}) \leq d(\bar{a}, \tilde{a}_{k_j}) + d(\tilde{a}_{k_j}, \tilde{b}_{k_j}) + d(\tilde{b}_{k_j}, \bar{b}_{k_j}) + d(\bar{b}_{k_j}, \bar{b}),$$

and likewise for $d(\bar{a}, \bar{c})$ and $d(\bar{b}, \bar{c})$. Letting $j \to \infty$ and recalling that $\{\tilde{a}_k, \tilde{b}_k, \tilde{c}_k, \tilde{a}_k\}$ is a $[\delta + (\epsilon - \delta)/2]$-triangle for all $k \geq K$, it follows that $d(\bar{a}, \bar{b})$ will be less than ϵ for large enough j, and likewise for $d(\bar{a}, \bar{c})$ and $d(\bar{b}, \bar{c})$. Hence, $\{\bar{a}, \bar{b}, \bar{c}, \bar{a}\}$ is an ϵ-triangle at $\bar{a} \in Y$.

Finally, we already have $\varphi(\bar{a}) = a$. We also have

$$d(\bar{b}, \tilde{b}_{k_j}) \leq d(\bar{b}, \bar{b}_{k_j}) + d(\bar{b}_{k_j}, \tilde{b}_{k_j}) \to 0,$$

and we likewise have that $\tilde{c}_{k_j} \to \bar{c}$. Thus, $\varphi_\epsilon^{n_{k_j}}(\tilde{b}_{k_j}) \to \varphi(\bar{b})$, which implies that $b_{k_j} \to \varphi(\bar{b})$. But $b_{k_j} \to b$, so we must have $\varphi(\bar{b}) = b$. The same reasoning shows that $\varphi(\bar{c}) = c$. Thus, $\{\bar{a}, \bar{b}, \bar{c}, \bar{a}\}$ is an ϵ-triangle that projects to γ. ∎

Now, let $(X_\epsilon, \tilde{*})$ be the ϵ-cover of the limit space, X, where - following our convention - we let $\tilde{*}$ denote the lifted base point in X_ϵ containing the trivial chain, $\{*\}$. By Lemma 2.4.9, there is a unique lift of $\varphi_\epsilon : X_\epsilon \to X$ to a map, $\psi : (X_\epsilon, \tilde{*}) \to (Y, \bar{*})$, such that $\varphi \circ \psi = \varphi_\epsilon$. Moreover, this map satisfies the same basic properties we have seen several times, now: it is a bijection from δ-balls in X_ϵ onto δ-balls in Y for any $\delta \leq \epsilon$, a radial isometry on ϵ-balls and, therefore, an isometry on $(\epsilon/2)$-balls, and distinct points in the preimage of any point $y \in Y$ are at least ϵ apart. Thus, ψ is a covering map onto its image in Y. In addition, ψ is surjective if and only if Y is ϵ-connected. If it is surjective, then it is a covering map. However, Y need not be ϵ-connected in general, as the following example shows.

Example 6.2.3 *This example extends the ideas illustrated in Example 2.2.7. For each $n \geq 1$, let X_n be the metric subspace formed by removing from the geodesic circle of circumference 1 an open segment of length $\frac{1}{4}(1 - \frac{1}{n})$. These spaces converge to the metric space, X, formed by removing an open segment of length $\frac{1}{4}$ from the geodesic circle of circumference 1. Let $\epsilon = \frac{1}{4}$. Then the ϵ-covers, X_ϵ^n, are formed by removing periodic open segments of length $\frac{1}{4}(1 - \frac{1}{n})$ from $\mathbb{R}$, and they converge to $\mathbb{R}$ with open segments of length $\frac{1}{4}$ removed. Thus, even though the covers X_ϵ^n are $\frac{1}{4}$-connected, their limit is not, since the gaps in the limiting space have length exactly $\frac{1}{4}$, meaning that no $\frac{1}{4}$-chain can cross such a gap. Note that the ϵ-covers in this case are all proper, also.* ∎

Thus, even for relatively simple metric spaces as we had in this example, the limit, Y, of the ϵ-covers may not be ϵ-connected. It does turn out that Y is what we call $\bar{\epsilon}$-connected, meaning that any two points can be joined by a chain, $\{y_0, \ldots, y_n\}$, such that $d(y_{i-1}, y_i) \leq \epsilon$. And Lemma 6.2.1 shows that Y does at least cover X. However, in order to obtain any meaningful results beyond this, particularly any connection between Y and X_ϵ, we need this limit to be ϵ-connected. Consequently, we will assume for the remainder of this section that Y is ϵ-connected. In the next section, we will derive some sufficient geometric conditions for this to hold. In short, Y will be ϵ-connected if the spaces, $\{X_n\}$, satisfy a *uniform* refinability property, meaning that there is some $\delta < \epsilon$ so that every ϵ-chain in *every* X_n can be ϵ-refined to a δ-chain. Any convergent sequence of compact geodesic spaces will have this property.

Now, assuming that Y is ϵ-connected, we will proceed to show, first, that the group of covering transformations of $\psi : X_\epsilon \to Y$ is isomorphic to a subgroup of $\pi_\epsilon(X)$ and, second, that Y is homeomorphic and locally isometric - though not necessarily globally isometric - to a quotient of X_ϵ by this same subgroup.

Let $\pi_\epsilon(Y)$ be the set of all $[\gamma]_\epsilon \in \pi_\epsilon(X)$ such that γ lifts as a closed loop to $\bar{*} \in Y$. Since we have chain and homotopy lifting properties from X to Y, it is easy to see that $\pi_\epsilon(Y)$ is a subgroup of $\pi_\epsilon(X)$. In addition, it is isomorphic to the ϵ-group of Y at $\bar{*}$, which is why we use the notation, $\pi_\epsilon(Y)$.

Lemma 6.2.4 *Assume Y is ϵ-connected, and let $G(\psi, X_\epsilon, Y)$ denote the group of covering transformations of $\psi : X_\epsilon \to Y$. That is, $G(\psi, X_\epsilon, Y)$ is the set of all homeomorphisms, $f : X_\epsilon \to X_\epsilon$, such that $\psi \circ h = \psi$. Then $G(\psi, X_\epsilon, Y)$ is naturally isomorphic to $\pi_\epsilon(Y)$.*

Proof Define a map $\Psi : \pi_\epsilon(Y) \to G(\psi, X_\epsilon, Y)$ as follows. Given $[\gamma]_\epsilon \in \pi_\epsilon(Y) \subset \pi_\epsilon(X)$, let $\Psi([\gamma]_\epsilon) = h_\gamma$, where, as in our discussion of the action of $\pi_\epsilon(X)$, h_γ is the isometry, $h_\gamma : X_\epsilon \to X_\epsilon$, defined by $h_\gamma([\alpha]_\epsilon) = [\gamma\alpha]_\epsilon$.

To show that Ψ is well-defined, we need to show that $[\gamma]_\epsilon \in \pi_\epsilon(Y) \Rightarrow \psi \circ h_\gamma = \psi$. Let $[\alpha]_\epsilon \in X_\epsilon$ be given. Let $x = \alpha_t$, and let $\bar{\alpha}$ be the unique lift of α to $\bar{*} \in Y$, with y the endpoint of $\bar{\alpha}$. Then $\psi([\alpha]_\epsilon) = y$. On the other hand, $h_\gamma([\alpha]_\epsilon) = [\gamma\alpha]_\epsilon$. Since γ lifts as a closed loop to $\bar{*} \in Y$, the lift of $\gamma\alpha$ to Y is just the lift of γ - which ends at $\bar{*}$ - followed by $\bar{\alpha}$. Thus, this lift ends at y, also, showing that $\psi(h_\gamma([\alpha]_\epsilon)) = y$. Hence, $\psi \circ h_\gamma = \psi$, and Ψ is well-defined. Injectivity follows because if $\Psi([\gamma]_\epsilon) = h_\gamma = id_{X_\epsilon}$, then

$$[\gamma]_\epsilon = [\gamma]_\epsilon[\{*\}]_\epsilon = h_\gamma([\{*\}]_\epsilon) = [\{*\}]_\epsilon.$$

The homomorphism property also follows directly, since

$$\Psi([\gamma_1]_\epsilon[\gamma_2]_\epsilon) = \Psi([\gamma_1\gamma_2]_\epsilon) = h_{\gamma_1\gamma_2} = h_{\gamma_1} \circ h_{\gamma_2} = \Psi([\gamma_1]_\epsilon) \circ \Psi([\gamma_2]_\epsilon).$$

To prove surjectivity, let $f : X_\epsilon \to X_\epsilon$ be a homeomorphism such that $\psi \circ f = \psi$, and let $[\gamma]_\epsilon = f(\tilde{*})$. Then $\psi([\gamma]_\epsilon) = \psi(f(\tilde{*})) = \psi(\tilde{*}) = \bar{*}$ and $\varphi_\epsilon([\gamma]_\epsilon) = \varphi(\psi([\gamma]_\epsilon)) = \varphi(\bar{*}) = *$. That is, γ is a loop at $* \in X$, and, so, $[\gamma]_\epsilon$ is an element of $\pi_\epsilon(X)$. Moreover, the equality $\psi([\gamma]_\epsilon) = \bar{*}$ shows that γ lifts closed to $\bar{*} \in Y$. Hence, $[\gamma]_\epsilon \in \pi_\epsilon(Y)$. We claim that $f = h_\gamma$. In fact, since ψ has the same local and covering properties as φ_ϵ, the proof of our map lifting lemma, Lemma 2.4.9, goes through equally well here - with only some minor additional comments, which we will mention - to show that there is a unique lift of $\psi : (X_\epsilon, \tilde{*}) \to (Y, \bar{*})$ to a map $h : (X_\epsilon, \tilde{*}) \to (X_\epsilon, [\gamma]_\epsilon)$. (Lemma 2.4.9 was proved specifically for lifts of φ_ϵ, so we cannot technically call upon it without some additional remarks.) It will follow that $f = h_\gamma = \Psi([\gamma]_\epsilon)$.

To see that Lemma 2.4.9 applies, we note that we have two pointed maps, $\psi : (X_\epsilon, \tilde{*}) \to (Y, \bar{*})$ and $\psi : (X_\epsilon, [\gamma]_\epsilon) \to (Y, \bar{*})$, and we want to lift the former to a map $h : (X_\epsilon, \tilde{*}) \to (X_\epsilon, [\gamma]_\epsilon)$ so that $\psi \circ h = \psi$. To define h, we take $[\alpha]_\epsilon \in X_\epsilon$ and let $\tilde{\lambda}$ be any ϵ-chain from $\tilde{*}$ to $[\alpha]_\epsilon$. We can project $\tilde{\lambda}$, via ψ, to an ϵ-chain at $\bar{*} \in Y$, and then lift that projection to an ϵ-chain at $[\gamma]_\epsilon \in X_\epsilon$. By the same reasoning used in the proof of Lemma 2.4.9, along with the fact that we have unique chain and homotopy lifting from Y to X_ϵ, this map is well-defined. To see that $\psi \circ h = \psi$, we let $[\alpha]_\epsilon \in X_\epsilon$ be given, with $x = \alpha_t \in X$. Let $\tilde{\alpha}$ be the lift of α to $\tilde{*} \in X_\epsilon$, $\bar{\alpha}$ the lift of α to $\bar{*} \in Y$ with $y = \bar{\alpha}_t$, and let $\hat{\alpha}$ be the lift of α to $[\gamma]_\epsilon \in X_\epsilon$. Also, let $\tilde{\gamma}$ be the lift of γ to $\tilde{*} \in X_\epsilon$, which ends at $[\gamma]_\epsilon$. Since $\varphi_\epsilon(\tilde{\gamma}\hat{\alpha}) = \gamma\alpha$, $\tilde{\gamma}\hat{\alpha}$ must be the unique lift of $\gamma\alpha$ to $\tilde{*} \in X_\epsilon$. But this lift ends at $[\gamma\alpha]_\epsilon$, so $\hat{\alpha}$ must end at that point, also. We also have

$$\varphi(\psi(\tilde{\gamma}))\varphi(\psi(\hat{\alpha})) = \varphi(\psi(\tilde{\gamma})\psi(\hat{\alpha})) = \varphi(\psi(\tilde{\gamma}\hat{\alpha})) = \varphi_\epsilon(\tilde{\gamma}\hat{\alpha}) = \gamma\alpha.$$

Since $\varphi(\psi(\tilde{\gamma})) = \gamma$, it follows that $\varphi(\psi(\hat{\alpha})) = \alpha$. By uniqueness of lifts, this means that $\psi(\hat{\alpha}) = \bar{\alpha}$. Applying uniqueness again, this, in turn, implies $\hat{\alpha}$ is the unique lift of $\bar{\alpha}$ to $[\gamma]_\epsilon \in X_\epsilon$. In other words, the lift of $\bar{\alpha}$ to $[\gamma]_\epsilon \in X_\epsilon$ ends at $[\gamma\alpha]_\epsilon$. By definition of h, this means that $h([\alpha]_\epsilon) = [\gamma\alpha]_\epsilon$. Furthermore, since γ lifts closed to Y, the lift of $\gamma\alpha$ to $\bar{*} \in Y$ ends at the same point as $\bar{\alpha}$. So, $\psi([\gamma\alpha]_\epsilon) = \psi([\alpha]_\epsilon)$. Putting all of this together yields

$$\psi(h([\alpha]_\epsilon)) = \psi([\gamma\alpha]_\epsilon) = \psi([\alpha]_\epsilon) \Rightarrow \psi \circ h = \psi.$$

In fact, since $[\alpha]_\epsilon$ was arbitrary, this actually shows that $h = h_\gamma$, which immediately verfies all of the local properties in Lemma 2.4.9. The uniqueness part of Lemma 2.4.9 now implies that f must be h_γ. ∎

Theorem 6.2.5 *Let $\{(X_n, *_n)\}$ be a sequence of compact, chain-connected metric spaces converging in the pointed sense to a compact metric space, $(X, *)$, and assume, for some $\epsilon > 0$, that the sequence of ϵ-covers, $\{(X^n_\epsilon, \tilde{*}_n)\}$ - which we assume to be proper - converges in the pointed sense to a complete metric space, $(Y, \bar{*})$. If Y is ϵ-connected, then Y is homeomorphic and locally isometric to $X_\epsilon/\pi_\epsilon(Y)$. Moreover, Y is isometric to $X_\epsilon/\pi_\epsilon(Y)$ if and only if Y is ϵ-intrinsic.*

Before we proceed with the proof, we should make some remarks on this result. First, given the results we derived in the section on ϵ-intrinsic spaces, the last statement of the theorem should not be too surprising. We saw in that section that the quotient space $X_\epsilon/\pi_\epsilon(X)$ was, in general, only homeomorphic and locally isometric to X. The two spaces differ only in their *global* metric structure. Thus, we have the same phenomenon occurring here; Y and $X_\epsilon/\pi_\epsilon(Y)$ differ only in their global geometry. Locally and topologically, they are the same. Also, we know from Theorem 2.6.3 that $X_\epsilon/\pi_\epsilon(Y)$, with the quotient metric, is a well-defined ϵ-intrinsic metric space. On the other hand, as the following example shows, limits of ϵ-intrinsic spaces need *not* be ϵ-intrinsic, so, even though $X^n_\epsilon \to Y$ and each X^n_ϵ is ϵ-intrinsic, Y does not have to be. Thus, one should expect that Y would not be isometric to $X_\epsilon/\pi_\epsilon(Y)$ unless Y is ϵ-intrinsic. Of course, geodesic spaces are ϵ-intrinsic, and a complete limit of a sequence of geodesic spaces is, itself, geodesic. Moreover, as we will see in the next section, Y is always ϵ-connected in the case of a geodesic sequence. Thus, when the spaces in question are geodesic spaces, this theorem does, in fact, tell us that Y is isometric to $X_\epsilon/\pi_\epsilon(Y)$.

Example 6.2.6 *This example extends the ideas in Example 2.5.5. For each $n \geq 1$, let X_n be the tuning fork space consisting of the line segments in $\mathbb{R}^2$ from $(0,0)$ to $(1 - \frac{1}{n}, 0)$, from $(0,0)$ to $(0,1)$, and from $(1 - \frac{1}{n}, 0)$ to $(1 - \frac{1}{n}, 1)$. As before, we endow these spaces with their inherited subspace metrics from $\mathbb{R}^2$. For each n, we choose as our base point, $*_n$, the center of the bottom segment of the tuning fork. In addition, we let $x_n = (0,1)$ and $y_n = (1 - \frac{1}{n}, 1)$ denote the upper left and right endpoints of the tuning fork, respectively. These spaces converge to the tuning fork, $(X, *)$, consisting of the segments from $(0,0)$ to $(1,0)$, from $(0,0)$ to $(0,1)$, and from $(1,0)$ to $(1,1)$, and where $* = (\frac{1}{2}, 0)$. We let $x = (0,1)$ and $y = (1,1)$ be the upper left and right endpoints of the limit space. As in Example 2.5.5, for any $\epsilon > 0$, X and each X_n are ϵ-simply connected. Moreover, all chains are arbitrarily refinable in all of the spaces, so there are no refinement critical values. Thus, the ϵ-covers of the spaces X_n are just the sets, X_n, with the ϵ-intrinsic metric induced by the Euclidean metric. For each n, let d_n denote the given Euclidean metric on X_n, and let d^n_ϵ denote the induced ϵ-intrinsic metric on X_n (i.e. the metric on X^n_ϵ).*

Fix $\epsilon = 1$. For each n, we have $d_n(x_n, y_n) < 1$, so there is an ϵ-chain between x_n and y_n with length precisely equal to the metric distance between the two points. Thus, in the ϵ-cover, X^n_ϵ, *we have $d_\epsilon^n(x_n, y_n) = d_n(x_n, y_n) < 1$. The ϵ-intrinsic metrics, d_ϵ^n, will converge to a metric, d, on X, and - since $d_\epsilon^n(x_n, y_n) = 1 - \frac{1}{n}$ - it is straightforward to see that we will have $d(x, y) = 1$ for this limit metric. However, this particular distance will not be the same as for the ϵ-intrinsic metric induced by d on X. For $\epsilon = 1$, any ϵ-chain from x to y in X must traverse around the tuning fork and will, consequently, have length not only greater than 1 but bounded away from 1. Thus, in the ϵ-intrinsic metric on X induced by the limit of the covering metrics, the distance between x and y must be strictly greater than 1, and this disagrees with the distance between x and y in the limiting metric. Hence, the ϵ-covers converge but not to a space that is ϵ-intrinsic.* ∎

Now, we proceed with the proof of Theorem 6.2.5.

Proof We first note that, given any orbit $[[\alpha]]_\epsilon \in X_\epsilon/\pi_\epsilon(Y)$, there is some $y \in Y$ such that $\psi^{-1}(y) = [[\alpha]]_\epsilon$. In fact, if the lift of α to $\bar{*} \in Y$ ends at y, then, by definition, $\psi([\alpha]_\epsilon) = y$. So, if $[\beta]_\epsilon \in [[\alpha]]_\epsilon$, then $[\beta]_\epsilon = [\gamma\alpha]_\epsilon$ for some $[\gamma]_\epsilon \in \pi_\epsilon(Y)$. Since γ, then, lifts closed to $\bar{*} \in Y$, the lift of $\gamma\alpha$ to Y will end at y, and since $\beta \sim_\epsilon \gamma\alpha$, the lift of β will end at y, also. Thus, $\psi([\beta]_\epsilon) = y$, and $[[\alpha]]_\epsilon \subset \psi^{-1}(y)$. On the other hand, if $[\beta]_\epsilon \in \psi^{-1}(y)$, then the lift of β to $\bar{*} \in Y$ ends at y. This means that α and β end at the same point in X, also. If we let $[\gamma]_\epsilon = [\beta\alpha^{-1}]_\epsilon$, then γ lifts closed to $\bar{*} \in Y$ by constuction, and $[\beta]_\epsilon = [\gamma\alpha]_\epsilon \Rightarrow [\beta]_\epsilon \in [[\alpha]]_\epsilon$. This shows that any orbit $[[\alpha]]_\epsilon$ is the preimage of some point in Y under ψ.

In the other direction, given any $y \in Y$, we can choose - since ψ is surjective - some $[\alpha]_\epsilon \in \psi^{-1}(y)$. Then the lift of α to $\bar{*} \in Y$ ends at y. If $[\beta]_\epsilon = [\gamma\alpha]_\epsilon$ for some $[\gamma]_\epsilon \in \pi_\epsilon(Y)$, then, as before, the lift of β to $\bar{*}$ will end at y, also. So, $\psi([\beta]_\epsilon) = y$, showing that $[[\alpha]]_\epsilon \subset \psi^{-1}(y)$. Conversely, if $[\beta]_\epsilon \in \psi^{-1}(y)$, then the lift of β to $\bar{*} \in Y$ ends at y. Just as before, it follows that $[\beta]_\epsilon \in [[\alpha]]_\epsilon$, showing that $[[\alpha]]_\epsilon = \psi^{-1}(y)$. In other words, the orbits of the action of $\pi_\epsilon(Y)$ on X_ϵ and the inverse images of points in Y under ψ exactly coincide: every orbit is the preimage of a unique point in Y, and every preimage corresponds to a unique orbit of some point in X_ϵ. This gives a bijective correspondence between the orbits in $X_\epsilon/\pi_\epsilon(Y)$ and the preimages, $\psi^{-1}(y)$.

So, we have a well-defined map, $f : X_\epsilon/\pi_\epsilon(Y) \to Y$, defined by $f([[\alpha]]_\epsilon) = y$ where $\psi([\alpha]_\epsilon) = y$, and this map is bijective. Moreover, if $\pi_Y : X_\epsilon \to X_\epsilon/\pi_\epsilon(Y)$ is the quotient map - which, by Proposition 2.6.3, is a radial isometry on ϵ-balls and an isometry on $\epsilon/2$-balls - then we have $f \circ \pi_Y = \psi$. Given $[[\alpha]]_\epsilon \in X_\epsilon/\pi_\epsilon(Y)$, let $y_0 = f([[\alpha]]_\epsilon) = \psi([\alpha]_\epsilon)$. Let d_q denote the quotient metric on $X_\epsilon/\pi_\epsilon(Y)$. If $d_q([[\beta]]_\epsilon, [[\alpha]]_\epsilon) < \delta \leq \epsilon$, then there is some $[\beta']_\epsilon \in [[\beta]]_\epsilon$ such that $d_\epsilon([\alpha]_\epsilon, [\beta']_\epsilon) < \delta$. Since ψ is bijection on δ-balls, $\psi([\beta]_\epsilon) = f([[\beta]]_\epsilon)$ lies in the δ-ball centered at $\psi([\alpha]_\epsilon) = y_0$. Thus, f maps δ-balls into δ-balls. On the other hand, let $y \in B_Y(y_0, \delta)$ be given, and let $[\beta]_\epsilon$ be the unique element in $B([\alpha]_\epsilon, \delta)$ mapping to y via ψ. Then $d_q([[\alpha]]_\epsilon, [[\beta]]_\epsilon) \leq d_\epsilon([\alpha]_\epsilon, [\beta]_\epsilon) < \delta$, and $f([[\beta]]_\epsilon) = \psi([\beta]_\epsilon) = y$. Hence, f maps δ-balls onto δ-balls for any $\delta \leq \epsilon$.

The other properties of f now follow directly. On any ϵ-ball, $B([[\alpha]]_\epsilon, \epsilon)$, we have $f = \psi \circ \pi_Y^{-1}$, and, since both maps on the right are radial isometries on ϵ-balls and isometries on $\epsilon/2$-balls, f is also. Thus, f is a local isometry and a homeomorphism. We also know from Proposition 2.6.3 that $X_\epsilon/\pi_\epsilon(Y)$ is ϵ-intrinsic. Thus, if Y is isometric to $X_\epsilon/\pi_\epsilon(Y)$, then Y will be ϵ-intrinsic, also. Conversely, if Y is ϵ-intrinsic, then - since f and f^{-1} will preserve the lengths of ϵ-chains - f will be an isometry. ∎

6.3 Precompactness Results

In the last section, we assumed that the ϵ-covers converged and derived the structure of their limits. In this section, we will actually derive conditions under which we do, in fact, have convergence. In fact, we will derive a precompactness result for the ϵ-covers of a precompact class of metric spaces.

We assume that we have a precompact class of pointed, compact, chain-connected metric spaces, $\mathscr{X}$. Given $\epsilon > 0$, we let $\mathscr{X}_\epsilon$ denote the corresponding class of ϵ-covers of spaces in X. In particular, we assume that this class is pointed, so that if $(X_\epsilon, \tilde{*}) \in \mathscr{X}_\epsilon$ is the ϵ-cover of $(X, *)$, then $\tilde{*} = [\{*\}]_\epsilon$. We also assume that the ϵ-covers, $(X_\epsilon, \tilde{*})$, are proper.

Recall that we denote the standard fundamental domain by $\tilde{X}_\epsilon$. We also have the corresponding ϵ-generating set

$$\mathscr{G}_\epsilon = \{[\gamma]_\epsilon \in \pi_\epsilon(X) \,:\, dist([\gamma]_\epsilon \tilde{X}_\epsilon, \tilde{X}_\epsilon) < \epsilon\}.$$

Since the ϵ-covers are proper, we know that $\mathscr{G}_\epsilon$ generates $\pi_\epsilon(X)$, whether it is finite or not. When X is also compact, $\mathscr{G}_\epsilon$ is finite.

To prove precompactness of $\mathscr{X}_\epsilon$, we will employ part 3 of Theorem 1.2.11, showing that the number of δ-balls needed to cover each closed ball $C(\tilde{*}, R) \subset X_\epsilon$, $X_\epsilon \in \mathscr{X}$, is uniformly bounded for any $R, \delta > 0$. The first result will be a very general one. After that, we will prove some corollaries of this result in special cases.

Theorem 6.3.1 *Let $\mathscr{X}$ be a precompact class of pointed, compact, chain-connected metric spaces, and let $\epsilon > 0$ be such that the ϵ-covers of each $X \in \mathscr{X}$ are proper. Let $\mathscr{X}_\epsilon$ be the corresponding pointed collection of ϵ-covers, $(X_\epsilon, \tilde{*})$, and let $\mathscr{G}_\epsilon(X)$ denote the ϵ-generating set of $X \in \mathscr{X}$. If the cardinalities of the generating sets, $|\mathscr{G}_\epsilon(X)|$, are uniformly bounded above, then $\mathscr{X}_\epsilon$ is precompact.*

Proof We need to show that any sequence in $\mathscr{X}$ has a convergent subsequence. Let $\{(X_\epsilon^n, \tilde{*}_n)\} \subset \mathscr{X}_\epsilon$ be such a sequence. Any sequence of ϵ-covers in $\mathscr{X}_\epsilon$ must come from a corresponding sequence in X, say $\{(X_n, *_n)\}$, and the precompactness of $\mathscr{X}$ implies that some subsequence of this sequence will converge. By just working with this subsequence, we get a sequence of ϵ-covers such that the corresponding sequence of base spaces converges. So, since we're looking for a subsequence anyway, we may as well just assume that the sequence $\{(X_n, *_n)\}$ converges.

For each n, let $\tilde{X}_\epsilon^n$ denote the standard fundamental domain in X_ϵ^n, $\mathscr{G}_\epsilon^n$ the ϵ-generating set, and d_ϵ^n the metric on X_ϵ^n. By our assumptions, each $\tilde{X}_\epsilon^n$ is compact and each $\mathscr{G}_\epsilon^n$ is finite. Moreover, there is some $K \in \mathbb{N}$ such that $|\mathscr{G}_\epsilon^n| \leq K$ for all n.

Given any $R, \delta > 0$, we want to show that the minimal number of δ-balls required to cover $C(\tilde{*}_n, R) \subset X_\epsilon^n$ is uniformly bounded in n. We will first show that it suffices to prove this result in the case $\delta < \epsilon < R$. So, assume we know that, for any $\delta < \epsilon < R$, there is some natural number $N(R, \delta)$ such that each ball, $C(\tilde{*}_n, R)$, can be covered by $N(R, \delta)$ or fewer δ-balls. If $R \leq \delta$, the result is trivial. If $\epsilon \leq \delta < R$, then we can cover $C(\tilde{*}_n, R)$ by $N(R, \epsilon/2)$ balls of radius $\epsilon/2$. Then the δ-balls centered at these same points will cover each $\epsilon/2$-ball and, therefore, $C(\tilde{*}_n, R)$. If $\delta < R \leq \epsilon$, then we can cover $C(\tilde{*}_n, 2\epsilon)$ by $N(2\epsilon, \delta)$ δ-balls, and this covering will also cover $C(\tilde{*}_n, R)$. So, we need only prove the covering result for the case $\delta < \epsilon < R$. The rest of the proof is just a series of counting and covering arguments, which we will prove as subclaims.

Claim: there exists $N \in \mathbb{N}$ - independent of n - so that each $C(\tilde{*}_n, R)$ can be covered by K^N or fewer translates, $[\gamma]_\epsilon \tilde{X}_\epsilon^n$. Let $D = \sup_n\{diam_\epsilon(X_n)\}$. Then, for any n, if $[\alpha]_\epsilon \in \tilde{X}_\epsilon^n$, we have $d_\epsilon(\tilde{*}_n, [\alpha]_\epsilon) = L([\alpha]_\epsilon) \leq D$. Let N be the smallest natural number such that

$$N > \frac{2(R+D)}{\epsilon} + 1.$$

We claim that the number of translates, $[\gamma]_\epsilon \tilde{X}_\epsilon^n$, that intersect $C(\tilde{*}_n, R)$ is bounded above by K^N and that this bound holds for every n. To see why this holds, fix $n \geq 1$. Suppose, for $[\gamma]_\epsilon \in \pi_\epsilon(X_n)$, we have $[\gamma]_\epsilon \tilde{X}_\epsilon^n \cap C(\tilde{*}_n, R) \neq \emptyset$. If $[\alpha]_\epsilon \in [\gamma]_\epsilon \tilde{X}_\epsilon^n \cap C(\tilde{*}_n, R)$ and $[\alpha]_\epsilon = [\gamma]_\epsilon[\beta]_\epsilon$ for $[\beta]_\epsilon \in \tilde{X}_\epsilon^n$, then

$$\begin{aligned} d_\epsilon^n(\tilde{*}_n, [\gamma]_\epsilon) &\leq d_\epsilon^n(\tilde{*}_n, [\alpha]_\epsilon) + d_\epsilon^n([\alpha]_\epsilon, [\gamma]_\epsilon) \\ &\leq R + d_\epsilon^n([\gamma]_\epsilon[\beta]_\epsilon, [\gamma]_\epsilon) \\ &\leq R + d_\epsilon^n([\beta]_\epsilon, \tilde{*}_n) \\ &\leq R + D. \end{aligned}$$

But $N > 2(R+D)/\epsilon + 1$, implying that

$$\frac{(N-1)\epsilon}{2} > R + D \geq d_\epsilon^n(\tilde{*}_n, [\gamma]_\epsilon) = L([\gamma]_\epsilon).$$

So, by the proof of Lemma 5.2.4, $[\gamma]_\epsilon$ can be expressed as a k-fold product of elements in $\mathscr{G}_\epsilon^n$, where $k \leq N$. There are, at most, $|\mathscr{G}_\epsilon^n|^N$ such products, and, since $|\mathscr{G}_\epsilon^n| \leq K$ for all n, we get the desired bound.

Now, the translates, $[\gamma]_\epsilon \tilde{X}_\epsilon^n$, cover X_ϵ^n, so some subcollection of these translates also covers $C(\tilde{*}_n, R)$. Specifically, we can take the set of all such translates that intersect $C(\tilde{*}_n, R)$. This collection has at most K^N elements by the previous argument. It also covers $C(\tilde{*}_n, R)$, since every $[\alpha]_\epsilon \in C(\tilde{*}_n, R)$ lies in some translate which must obviously intersect $C(\tilde{*}_n, R)$. This proves the claim.

Claim: there exists $M \in \mathbb{N}$ - independent of n - so that each $\tilde{X}_\epsilon^n$ can be covered by MK^2 δ-balls. Fix $n \geq 1$. By our assumptions, the restricted map $\varphi_\epsilon^n : \tilde{X}_\epsilon^n \to X_n$ is surjective, and the preimage of each point $x \in X_n$ is finite. We claim that the number of elements in the preimage of any point is actually bounded above by K. Given any $x \in X_n$, let $S_x = \{[\alpha_1]_\epsilon, \ldots, [\alpha_k]_\epsilon\} = (\varphi_\epsilon^n)^{-1}(x)$, and note that this set is just all of the minimal ϵ-classes from $*_n$ to x. Let $\mathscr{G}^n$ denote the set of all $[\gamma]_\epsilon \in \pi_\epsilon(X_n)$ such that $[\gamma]_\epsilon \tilde{X}_\epsilon^n \cap \tilde{X}_\epsilon^n \neq \emptyset$. Clearly we have $\mathscr{G}^n \subset \mathscr{G}_\epsilon^n$. Define a map, $f : S_x \to \mathscr{G}^n$ by $f([\alpha_i]_\epsilon) = [\alpha_i \alpha_1^{-1}]_\epsilon$ for $i = 1, \ldots, k$. To see that f is well-defined, notice that $[\gamma]_\epsilon := [\alpha_i \alpha_1^{-1}]_\epsilon$ takes $[\alpha_1]_\epsilon$ to $[\alpha_i]_\epsilon$. Thus, $[\gamma]_\epsilon \tilde{X}_\epsilon^n \cap \tilde{X}_\epsilon^n \neq \emptyset$, and f does map S_x into $\mathscr{G}^n$. Moreover, f is injective, for if $f([\alpha_i]_\epsilon) = f([\alpha_j]_\epsilon)$ for $1 \leq i, j \leq k$, then $[\alpha_i \alpha_1^{-1}]_\epsilon = [\alpha_j \alpha_1^{-1}]_\epsilon$ $\Rightarrow [\alpha_i]_\epsilon = [\alpha_j]_\epsilon$. Thus, the cardinality of S_x must be bounded above by the cardinality of $\mathscr{G}^n$, which, in turn, is bounded above by $|\mathscr{G}_\epsilon^n| \leq K$.

Since $\{(X_n, *_n)\}$ converges to some compact metric space, X, there is some $M \in \mathbb{N}$ such that every X_n can be covered by a collection of M or fewer δ-balls. We next claim that, for any n, $\tilde{X}_\epsilon^n$ can be covered by MK^2 or fewer δ-balls. To prove this, we first choose a covering of X_n consisting of M or fewer δ-balls, and let $x_1, \ldots, x_m \in X_n$ denote the centers of these balls, where $m \leq M$. Then, for each $i = 1, \ldots, m$, take all of the points in the preimage of x_i in $\tilde{X}_\epsilon^n$ (i.e. all of the minimal ϵ-classes from $*_n$ to x_i), and consider the collection of δ-balls centered at those points. This gives us a collection of, at most, MK δ-balls centered at points in $\tilde{X}_\epsilon^n$.

Proceeding further, for each of those δ-balls, take all of the translates of that ball by elements in $\mathscr{G}_\epsilon^n$. Since $\mathscr{G}_\epsilon^n$ has, at most, K elements, doing this for each ball gives us a collection of, at most, MK^2 δ-balls, which we will denote by $\mathscr{C}_n$. We claim that $\mathscr{C}_n$ covers $\tilde{X}_\epsilon^n$.

To be clear, a δ-ball, B, is in $\mathscr{C}_n$ if and only if there is some $x_i \in \{x_1, \ldots, x_m\}$ and some $[\alpha]_\epsilon \in \tilde{X}_\epsilon^n$ with $\alpha_t = x$ such that B is the image of $B([\alpha]_\epsilon, \delta)$ under some $[\gamma]_\epsilon \in \mathscr{G}_\epsilon^n$. So, let $[\alpha]_\epsilon \in \tilde{X}_\epsilon^n$ be given, and let $x = \alpha_t$. There is some x_i such that $x \in B(x_i, \delta) \subset X_n$. Choose some $[\beta]_\epsilon \in \tilde{X}_\epsilon^n$ so that $\beta_t = x_i$. By construction, the ball $B([\beta]_\epsilon, \delta)$ and all of its translates via elements in $\mathscr{G}_\epsilon^n$ are in $\mathscr{C}_n$. Let $\sigma = \{x, x_i\}$, an ϵ-chain because $\delta < \epsilon$, and let $[\gamma]_\epsilon = [\alpha\sigma\beta^{-1}]_\epsilon$. Since $[\alpha\sigma]_\epsilon = [\gamma\beta]_\epsilon \in [\gamma]_\epsilon \tilde{X}_\epsilon^n$, $[\alpha]_\epsilon \in \tilde{X}_\epsilon^n$, and

$$d_\epsilon^n([\alpha]_\epsilon, [\alpha\sigma]_\epsilon) = L([\alpha^{-1}\alpha\sigma]_\epsilon) = L([\sigma]_\epsilon) < \delta,$$

we have $dist([\gamma]_\epsilon \tilde{X}_\epsilon^n, \tilde{X}_\epsilon^n) < \delta < \epsilon$, showing that $[\gamma]_\epsilon \in \mathscr{G}_\epsilon^n$. Moreover, by the same reasoning, we have

$$[\alpha]_\epsilon \in B([\alpha\sigma]_\epsilon, \delta) = B([\gamma]_\epsilon[\beta]_\epsilon, \delta).$$

So, $[\alpha]_\epsilon$ lies in an element of $\mathscr{C}_n$, showing that this collection covers $\tilde{X}_\epsilon^n$. Hence, for any n, $\tilde{X}_\epsilon^n$ can be covered by MK^2 or fewer δ-balls. This proves the second claim.

Finally, we have shown that we can cover each $C(\tilde{*}_n, R)$ by K^N or fewer translates of $\tilde{X}_\epsilon^n$, where N does not depend on n. Likewise, we have shown that we can cover $\tilde{X}_\epsilon^n$ - and, thus, any of the translates $[\gamma]_\epsilon \tilde{X}_\epsilon^n$ - by MK^2 or fewer δ-balls, where M does not depend on n. It follows that we can cover $C(\tilde{*}_n, R)$ by MK^{2+N} or fewer δ-balls. Since this bound does not depend on n, it follows that the sequence $\{(X_\epsilon^n, \tilde{*}_n)\}$ is precompact and, therefore, contains a convergent subsequence. This, in turn, implies that $\mathscr{X}_\epsilon$ is precompact. ∎

We immediately have the following corollary, which, in fact, is what we actually proved in the previous result.

Corollary 6.3.2 *If $\{(X_n, *_n)\}$ is sequence of compact, chain-connected metric spaces converging to a compact metric space, X, and if $\epsilon > 0$ is such that the ϵ-covers, X_ϵ^n, are proper and $\{|\mathscr{G}_\epsilon^n|\}$ is uniformly bounded, then there is some subsequence of the corresponding pointed collection of ϵ-covers, $\{(X_\epsilon^n, \tilde{*}_n)\}$, that converges in the pointed sense to a complete metric space, $(Y, \bar{*})$. Moreover, Y is a covering space of X, and, if Y is ϵ-connected, then it is homeomorphic and locally isometric to a quotient of X_ϵ by a subgroup of $\pi_\epsilon(X)$.*

We will now prove some corollaries of Theorem 6.3.1 in certain special cases. The conditions in these contexts can sometimes be easier to check than a uniform bound on the ϵ-generating set cardinalities. To facilitate this, we will prove the following useful lemma. This lemma appears technical, but it turns out to be very helpful. Essentially, it states that if X and Y are sufficiently close, and if X is sufficiently refinable at and around some ϵ value, then we can define a surjective homomorphism from $\pi_\delta(Y)$ onto $\pi_\epsilon(X)$ for all δ less than but sufficiently close to ϵ. This lemma seems to be useful enough that we will refer to it by a specific name: the *the close homomorphism lemma*.

Lemma 6.3.3 (The Close Homomorphism Lemma) *Let X be a compact, chain-connected metric space, and, for some $\epsilon > 0$, assume that X is ϵ-surjective from below. Specifically, let $\lambda < \epsilon$ be such that every ϵ-chain in X can be ϵ-refined to a λ-chain. Then the following hold.*

1) Given δ such that $\lambda < \delta < \epsilon$, if Y is any compact, chain-connected metric space such that

$$d_{GH}(X,Y) < \frac{\min\{\epsilon-\delta, \delta-\lambda\}}{4},$$

then there is a surjective homomorphism $\theta : \pi_\delta(Y) \to \pi_\epsilon(X)$.

*2) Assume, in addition to the assumptions in part 1, that **1)** the map $X_{t_1} \to X_{t_2}$ is surjective for every $\lambda \leq t_1 < t_2 \leq \epsilon$, and **2)** Y has the same property. Given $\lambda < \delta < \delta' < \epsilon' < \epsilon$, if*

$$d_{GH}(X,Y) < \frac{1}{4}\min\{\delta-\lambda, \delta'-\delta, \epsilon'-\delta', \epsilon-\epsilon'\},$$

then each homomorphism $\theta_1 : \pi_\delta(X) \to \pi_{\delta'}(Y)$, $\theta_2 : \pi_{\epsilon'}(Y) \to \pi_\epsilon(X)$ is well-defined and we have the following commutative diagram:

$$\begin{array}{ccc} \pi_{\epsilon'}(Y) & \xrightarrow{\theta_2} & \pi_\epsilon(X) \\ \uparrow \Phi_{\epsilon'\delta'} & & \uparrow \Phi_{\epsilon\delta} \\ \pi_{\delta'}(Y) & \xleftarrow{\theta_1} & \pi_\delta(X). \end{array}$$

Remark Note the difference between the surjectivity assumptions in parts 1 and 2. In part 1, we are only assuming that each map $X_t \to X_\epsilon$ is surjective for every $t \in [\lambda, \epsilon]$. In part 2, we are assuming that, for *every* $\lambda \leq t_1 < t_2 \leq \epsilon$, the map $X_{t_1} \to X_{t_2}$ is surjective. These surjectivity conditions are always satisfied for geodesic spaces.

Proof Let Y be such that $d_{GH}(X,Y) < \tau := \frac{1}{4}\min\{\epsilon-\delta, \delta-\lambda\}$. Fix a base point $* \in X$, and choose a base point, $\star \in Y$ such that $d(*, \star) < \tau$. Since X and Y are chain-connected, the ϵ-groups of X (resp. δ-groups of Y) at any two base points are isomorphic. Thus, if we show that there is a surjective homomorphism from $\pi_\delta(Y,\star)$ onto $\pi_\epsilon(X,*)$, then the same conclusion will hold for any choice of base points.

Let $[\gamma]_\delta \in \pi_\delta(Y,\star)$ be given, and denote γ by $\gamma = \{\star = y_0, \ldots, y_n = \star\}$. For each $i = 0, 1, \ldots, n$, choose $x_i \in X$ such that $d(x_i, y_i) < \tau$, being sure to choose $x_0 = x_n = *$. Let $\tilde{\gamma} = \{* = x_0, \ldots, x_n = *\}$. The fact that $\tilde{\gamma}$ is an ϵ-loop follows because

$$d(x_{i-1}, x_i) \leq d(x_{i-1}, y_{i-1}) + d(y_{i-1}, y_i) + d(y_i, x_i) < 2\tau + \delta < \epsilon.$$

Suppose $* = \bar{x}_0, \bar{x}_1, \ldots, \bar{x}_n = *$ is any other choice of points in X such that $d(y_i, \bar{x}_i) < \tau$ for each i, and let $\tilde{\gamma}'$ be the ϵ-loop, $\{* = \bar{x}_0, \ldots, \bar{x}_n = *\}$. Note that for each i we have

$$d(\bar{x}_i, x_i) \leq d(\bar{x}_i, y_i) + d(y_i, x_i) < 2\tau < \epsilon,$$

and for each $i = 1, \ldots, n$, we have

$$\begin{aligned} d(\bar{x}_{i-1}, x_i) &\leq d(\bar{x}_{i-1}, x_{i-1}) + d(x_{i-1}, x_i) \\ &\leq d(\bar{x}_{i-1}, y_{i-1}) + d(y_{i-1}, x_{i-1}) + d(x_{i-1}, x_i) \\ &\leq 2\tau + 2\tau + \delta < \epsilon. \end{aligned}$$

Thus, we can insert $\bar{x}_1$ into $\tilde{\gamma}$ between $* = x_0$ and x_1 and then remove x_1 to obtain the chain $\{* = \bar{x}_0, \bar{x}_1, x_2, \ldots, x_n = *\}$. We can then insert $\bar{x}_2$ into this chain between $\bar{x}_1$ and x_2 and then remove x_2 to obtain the chain $\{* = \bar{x}_0, \bar{x}_1, \bar{x}_2, \ldots, x_n = *\}$. Continuing in this way, we can successively insert each $\bar{x}_i$ and remove the corresponding x_i, thus transforming, via ϵ-homotopy, $\tilde{\gamma}$ to $\tilde{\gamma}'$. Hence, it does not matter which points, $* = x_0, x_1, \ldots, x_n = *$, we choose in X satisfying $d(y_i, x_i) < \tau$, $x_0 = x_n = *$; for any such choice of points, the resulting loops belong to the same ϵ-homotopy class, $[\tilde{\gamma}]_\epsilon$.

Note that, thus far, we have defined a map, $\gamma \mapsto [\tilde{\gamma}]_\epsilon$, taking δ-loops at $\star$ to homotopy classes of ϵ-loops at $*$, and we have shown that this map is well-defined regardless of what points we choose for $\tilde{\gamma}$ satisfying the required conditions. So, we define $\theta : \pi_\delta(Y) \to \pi_\epsilon(X)$ by $\theta([\gamma]_\delta) = [\tilde{\gamma}]_\epsilon$, where - if $\gamma = \{\star = y_0, \ldots, y_n = \star\}$ - then $\tilde{\gamma} = \{x_0, \ldots, x_n\}$ is any choice of points in X such that $x_0 = x_n = *$ and $d(x_i, y_i) < \tau$ for $i = 1, \ldots, n-1$. We still need to show, however, that θ is well-defined, or if $\gamma' \sim_\delta \gamma$ in Y, and if $\tilde{\gamma}'$ and $\tilde{\gamma}$ are chosen as before, corresponding to γ' and γ, respectively, then $\tilde{\gamma}' \sim_\epsilon \tilde{\gamma}$. We prove this, first, in the case where γ' and γ differ by just a basic move. Assume, first, that γ' differs from γ by addition of a single point, say $\gamma = \{\star = y_0, \ldots, y_{i-1}, y_i, \ldots, y_n = \star\}$ and $\gamma' = \{\star = y_0, \ldots, y_{i-1}, y, y_i, \ldots, y = \star\}$. As before, we choose, for each $i = 0, 1, \ldots, n$, $x_i \in X$ so that $d(y_i, x_i) < \tau$ and $x_0 = y_0 = *$, and we let $\tilde{\gamma} = \{* = x_0, \ldots, x_n = *\}$. Choose $x \in X$ so that $d(y, x) < \tau$. Since the homotopy class of the resulting loop in X is independent of which points we choose, we can set $\tilde{\gamma}' = \{* = x_0, \ldots, x_{i-1}, x, x_i, \ldots, x_n = *\}$. This is an ϵ-loop because

$$d(x_{i-1}, x) \leq d(x_{i-1}, y_{i-1}) + d(y_{i-1}, y) + d(y, x) < \delta + 2\tau < \epsilon$$

$$d(x, x_i) \leq d(x, y) + d(y, y_i) + d(y_i, x_i) < \delta + 2\tau < \epsilon.$$

Moreover, we clearly have $\tilde{\gamma}' \sim_\epsilon \tilde{\gamma}$, since the former is obtained by adding a point to the latter.

In the second case, assume γ' is obtained by removing a point from γ, say

$$\gamma = \{\star = y_0, \ldots, y_{i-1}, y_i, y_{i+1}, \ldots, y_n = \star\}$$

$$\gamma' = \{\star = y_0, \ldots, y_{i-1}, y_{i+1}, \ldots, y = \star\}.$$

We choose, for each $i = 0, 1, \ldots, n$, $x_i \in X$ so that $d(y_i, x_i) < \tau$ and $x_0 = y_0 = *$, and we let $\tilde{\gamma} = \{* = x_0, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_n\}$ and $\tilde{\gamma}' = \{* = x_0, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n = *\}$. Then $\tilde{\gamma}'$ is an ϵ-loop because

$$d(x_{i-1}, x_{i+1}) \leq d(x_{i-1}, y_{i-1}) + d(y_{i-1}, y_{i+1}) + d(y_{i+1}, x_{i+1}) < \delta + 2\tau < \epsilon.$$

So, $\tilde{\gamma}' \sim_\epsilon \tilde{\gamma}$, showing that if γ' and γ differ by a basic move, then γ' and γ map to the same homotopy class of ϵ-loops in X.

Now, for the general case, if $\gamma' \in [\gamma]_\delta$, there is a sequence of δ-loops, $\gamma = \gamma_1, \ldots, \gamma_k = \gamma'$, such that γ_i differs from γ_{i-1} by a basic move for each $i = 2, \ldots, k$. But, by the previous two cases, each γ_i will map to the same homotopy class of ϵ-loops at $* \in X$, implying that γ and γ' will, also. Thus, θ is well-defined.

To see that θ is a homomorphism, suppose $[\gamma_1]_\delta$, $[\gamma_2]_\delta \in \pi_\delta(Y)$ are given, and denote γ_1, γ_2, respectively, by $\gamma_1 = \{\star = y_0, \ldots, y_n = \star\}$ and $\gamma_2 = \{\star = z_0, \ldots, z_m = \star\}$. For each $i = 1, \ldots, n-1$, choose $x_i \in X$ so that $d(y_i, x_i) < \tau$, and, for each $j = 1, \ldots, m-1$, choose $x'_j \in X$ so that $d(z_j, x'_j) < \tau$. Let $\tilde{\gamma}_1 = \{* = x_0, x_1, \ldots, x_n = *\}$ and $\tilde{\gamma}_2 = \{* = x'_0, x'_1, \ldots, x'_m = *\}$. Then, by construction, we have $\theta([\gamma_1]_\delta) = [\tilde{\gamma}_1]_\epsilon$ and $\theta([\gamma_2]_\delta) = [\tilde{\gamma}_2]_\epsilon$. Also by construction, we have $\theta([\gamma_1\gamma_2]_\delta) = [\tilde{\gamma}_1\tilde{\gamma}_2]_\epsilon$. Thus,

$$\theta([\gamma_1]_\delta[\gamma_2]_\delta) = \theta([\gamma_1\gamma_2]_\delta) = [\tilde{\gamma}_1\tilde{\gamma}_2]_\epsilon = [\tilde{\gamma}_1]_\epsilon[\tilde{\gamma}_2]_\epsilon = \theta([\gamma_1]_\delta)\theta([\gamma_2]_\delta).$$

Finally, we show that θ is surjective. Note that we have not used the refinability of X, yet, so this homomorphism is well-defined whether X is ϵ-surjective from below or not. The surjectivity of θ is the only property of θ that requires refinability. Let $[\tilde{\gamma}]_\epsilon \in \pi_\epsilon(X)$ be given. By assumption, there is a λ-loop, $\tilde{\gamma}'$, in $[\tilde{\gamma}]_\epsilon$. Denote this loop by $\tilde{\gamma}' = \{* = x_0, \ldots, x_n = *\}$, and choose, for each $i = 1, \ldots, n-1$, $y_i \in Y$ such that $d(x_i, y_i) < \tau$. Let $\gamma = \{\star = y_0, y_1, \ldots, y_{n-1}, y_n = *\}$. This is a δ-chain, since

$$d(y_{i-1}, y_i) \leq d(y_{i-1}, x_{i-1}) + d(x_{i-1}, x_i) + d(x_i, y_i) < 2\tau + \lambda < \delta.$$

Moreover, by definition of θ, we have $\theta([\gamma]_\delta) = [\tilde{\gamma}]_\epsilon$. Thus, θ is surjective.

The proof of part 2 now follows mostly from part 1. By the stated assumptions, every δ'-chain in Y can be δ'-refined to a λ-chain. Moreover, we have $\lambda < \delta < \delta'$ and

$$d_{GH}(X, Y) < \frac{1}{4}\min\{\delta - \lambda, \delta' - \delta, \epsilon' - \delta', \epsilon - \epsilon'\} \leq \frac{1}{4}\min\{\delta - \lambda, \delta' - \delta\},$$

but these are precisely the conditions needed to apply part 1 to conclude that we have a surjective homomorphism $\theta_1 : \pi_\delta(X) \to \pi_{\delta'}(Y)$. Likewise, by assumption, every ϵ-chain in X can be ϵ-refined to a δ'-chain, and

$$d_{GH}(X, Y) < \frac{1}{4}\min\{\delta - \lambda, \delta' - \delta, \epsilon' - \delta', \epsilon - \epsilon'\} \leq \frac{1}{4}\min\{\epsilon' - \delta', \epsilon - \epsilon'\},$$

we can apply part 1, again, to get the surjective homomorphism $\theta_2 : \pi_{\epsilon'}(Y) \to \pi_\epsilon(X)$. All that remains is to show that the diagram above commutes, but this really just follows from how we define θ. Let

$$\tau = \frac{1}{4}\min\{\delta - \lambda, \delta' - \delta, \epsilon' - \delta', \epsilon - \epsilon'\}.$$

Given $[\gamma]_\delta \in \pi_\delta(X)$, with $\gamma = \{* = x_0, \ldots, x_n = *\}$, we choose $y_i \in Y$ so that $d(x_i, y_i) < \tau$ for each $i = 1, \ldots, n-1$, and we set $\tilde{\gamma} = \{\star = y_0, y_1, \ldots, y_{n-1}, y_n = \star\}$. Then, by construction, $\tilde{\gamma}$ is a δ'-loop at $\star \in Y$, and $\theta_1([\gamma]_\delta) = [\tilde{\gamma}]_{\delta'}$. So, $\Phi_{\epsilon'\delta'}(\theta_1([\gamma]_\delta)) = [\tilde{\gamma}]_{\epsilon'}$. Now, by construction, to determine $\theta_2([\tilde{\gamma}]_{\epsilon'})$, we take any ϵ'-loop in $[\tilde{\gamma}]_\epsilon$, say $\{\star = z_0, \ldots, z_m = \star\}$, and choose any set of points $\{* = \bar{x}_0, \bar{x}_1, \ldots, \bar{x}_n = *\}$ such that $d(z_i, \bar{x}_i) < \tau$ for $i = 1, \ldots, m-1$. But one such ϵ'-loop is $\tilde{\gamma} = \{\star = y_0, \ldots, y_n = \star\}$, and, so, we can simply choose the points $\{* = x_0, x_1, \ldots, x_{n-1}, x_n = *\}$. That is, we can just choose γ. Thus,

$$\theta_2([\tilde{\gamma}]_{\epsilon'}) = [\gamma]_\epsilon$$

$$\Rightarrow \Phi_{\epsilon\delta}([\gamma]_\delta) = [\gamma]_\epsilon = \theta_2([\tilde{\gamma}]_{\epsilon'}) = (\theta_2 \circ \Phi_{\epsilon'\delta'} \circ \theta_1)([\gamma]_\delta). \quad \blacksquare$$

Note that we did not need the refinability condition in showing that the above diagram commuted. The only place we used refinability was in showing that the homomorphisms were surjective. Even without refinability, however, we still get the above commutative diagram.

An immediate corollary of part 1 of this lemma is the following.

Corollary 6.3.4 *If $\{X_n\}$ is a sequence of compact, chain-connected metric spaces converging to a compact metric space, X, and if there is some $\lambda < \epsilon$ such that every ϵ-chain in X can be ϵ-refined to a λ-chain, then, for any δ such that $\lambda < \delta < \epsilon$, there is a natural number $N(\delta)$ such that for all $n \geq N(\delta)$ there is a surjective homomorphism $\theta : \pi_\delta(X_n) \to \pi_\epsilon(X)$.*

The following example shows why we need refinability for the surjectivity of this homomorphism to hold.

Example 6.3.5 *Let X_n be the subspace of the geodesic circle of circumference 1 formed by removing a gap of length $\frac{1}{4}(1 + \frac{1}{n})$. These spaces converge to the subspace of the same geodesic circle with a gap of length $\frac{1}{4}$ removed. Let $\epsilon = \frac{1}{3}$. Then every ϵ-chain in X can be ϵ-refined to a λ-chain for all $\lambda > \frac{1}{4}$, but there are ϵ-chains that cannot be ϵ-refined to $\frac{1}{4}$-chains. Moreover, $\pi_{1/3}(X) \cong \mathbb{Z}$, while $\pi_{1/4}(X_n)$ is trivial for all n. So, there is no surjective homomorphism from $\pi_{1/4}(X_n)$ onto $\pi_{1/3}(X)$, because we cannot refine all ϵ-chains in X to be finer than $\frac{1}{4}$-chains. On the other hand, for $\frac{1}{4} < \delta < \frac{1}{3}$, $\pi_\delta(X_n) \cong \mathbb{Z}$ for large enough n, so we do get a surjective homomorphism $\theta : \pi_\delta(X_n) \to \pi_{1/3}(X)$ for large enough n. The difference, of course, is that we can refine ϵ-chains in X to δ-chains for such δ.* ∎

Recall that, in the previous section, we showed that the limit of the ϵ-covers of a given sequence, when it exists, need not be ϵ-connected. The following corollary of our precompactness theorem, Theorem 6.3.1, not only provides sufficient conditions for the cardinalities, $|\mathscr{G}_\epsilon|$, to be bounded but it also provides conditions under which the limit of any converging subsequence of ϵ-covers is ϵ-connected.

Corollary 6.3.6 *Let $\mathscr{X}$ be a precompact collection of compact, pointed, chain-connected metric spaces, and let $\mathscr{X}_\epsilon$ be the corresponding pointed collection of ϵ-covers. Assume, for some $\epsilon > 0$, that the following hold.*

1) *The spaces in $\mathscr{X}$ have a **uniform** ϵ-BMR property, meaning that there exist $\lambda < \epsilon$ and $N \in \mathbb{N}$ such that every two-point ϵ-chain in any $X \in \mathscr{X}$ can be ϵ-refined to a λ-chain of N or fewer points.*
2) *There exists some non-empty interval, $(\epsilon - \tau, \epsilon]$, such that for any $t \in (\epsilon - \tau, \epsilon]$ and any $X \in \mathscr{X}$, the cover X_t is proper.*

Then $\mathscr{X}_\epsilon$ is precompact. Moreover, the limit, Y, of any convergent sequence $\{(X^n_\epsilon, \tilde{*}_n)\} \subset \mathscr{X}_\epsilon$ is ϵ-connected.

Proof As in the proof of Theorem 6.3.1, we need to show that any subsequence $\{(X^n_\epsilon, \tilde{*}_n)\} \subset \mathscr{X}_\epsilon$ has a convergent subsequence. Reasoning as in that proof, we may assume that the corresponding sequence of base spaces, $\{(X_n, *_n)\}$, converges. We may assume without loss of generality that $\lambda > \max\{\epsilon - \tau, \epsilon/2\}$. Fix δ so that $\lambda < \delta < \epsilon$. Then, since $\{(X_n, *_n)\}$ is Cauchy, we can choose $K \in \mathbb{N}$ so that $m, n \geq K$ implies that

$$d_{GH}\big((X_n, *_n), (X_m, *_m)\big) < \frac{1}{4}\min\{\epsilon - \delta, \delta - \lambda\}.$$

By the close homomorphism lemma, there is a surjective homomorphism $\theta_n : \pi_\delta(X_K) \to \pi_\epsilon(X_n)$ for any $n \geq K$. (We're assuming, of course, that these are the groups based at $*_K$ and $*_n$, respectively.)

Let M be the smallest natural number such that $M\epsilon \geq \sup_n\{diam_\epsilon(X_n)\}$. Fix $n \geq K$, and let d^n_ϵ denote the metric on X^n_ϵ. Furthermore, let $\mathscr{G}^n_\epsilon$ and $\tilde{X}^n_\epsilon$ be the ϵ-generating set of $\pi_\epsilon(X_n)$ and the standard fundamental domain of X_n, respectively. Then, if $[\alpha]_\epsilon \in \tilde{X}^n_\epsilon$, we have $L([\alpha]_\epsilon) \leq M\epsilon$. Let $[\gamma]_\epsilon \in \mathscr{G}^n_\epsilon$ be given, so that $dist([\gamma]_\epsilon \tilde{X}^n_\epsilon, \tilde{X}^n_\epsilon) < \epsilon$. Since these sets are compact, there are $[\alpha]_\epsilon, [\beta]_\epsilon \in \tilde{X}^n_\epsilon$ such that

$$\epsilon > dist([\gamma]_\epsilon \tilde{X}^n_\epsilon, \tilde{X}^n_\epsilon) = d^n_\epsilon([\gamma]_\epsilon[\beta]_\epsilon, [\alpha]_\epsilon).$$

It follows that

$$\begin{aligned} d^n_\epsilon(\tilde{*}_n,[\gamma]_\epsilon) &\leq d^n_\epsilon(\tilde{*}_n,[\alpha]_\epsilon)+d^n_\epsilon([\alpha]_\epsilon,[\gamma]_\epsilon[\beta]_\epsilon)+d^n_\epsilon([\gamma]_\epsilon[\beta]_\epsilon,[\gamma]_\epsilon) \\ &< 2M\epsilon+\epsilon. \end{aligned}$$

Thus, $L([\gamma]_\epsilon) < (2M+1)\epsilon$, and this holds for any $[\gamma]_\epsilon \in \mathscr{G}^n_\epsilon$. So, by Lemma 2.7.3, given any $[\gamma]_\epsilon \in \mathscr{G}^n_\epsilon$, there is some loop in $[\gamma]_\epsilon$ consisting of, at most, $4M+3$, points. By assumption, we can, then, ϵ-refine $[\gamma]_\epsilon$ to a λ-loop consisting of fewer than $(4M+2)N$ points. (This is a true overestimate, but it will make the exposition clearer than computing the exact value; all we need is an upper bound.) For simplicity, we will just assume that γ is a λ-loop with $(4M+2)N$ or fewer points.

Now, recalling the proof of surjectivity in Lemma 6.3.3, we can find an element of $\pi_\delta(X_K)$ mapping to $[\gamma]_\epsilon \in \mathscr{G}^n_\epsilon \subset \pi_\epsilon(X_n)$ by choosing a δ-loop at $*_K \in X_K$, say $\tilde{\gamma}$, with the same number of points as this λ-loop, γ. It follows that $L(\tilde{\gamma}) < (4M+2)N\delta \Rightarrow L([\tilde{\gamma}]_\delta) < (4M+2)N\delta$. In other words, each $[\gamma]_\epsilon \in \mathscr{G}^n_\epsilon$ is the image, under θ_n, of an element in $\pi_\delta(X_K)$ lying in the ball $C(\tilde{*}_K,R)\cap\pi_\delta(X_K)$, where $R \geq (4M+2)N\delta$. But, since X^K_δ is proper by hypothesis, the number of elements of $\pi_\delta(X_K)$ lying in this ball is finite. Hence, we have $|\mathscr{G}^n_\epsilon| \leq |\pi_\delta(X_K)\cap C(\tilde{*}_K,R)|$, and this holds for all $n \geq K$. This shows that the cardinalities, $\{|\mathscr{G}^n_\epsilon|\}$, are uniformly bounded in n, which, in turn, by Theorem 6.3.1, shows that $\{(X^n_\epsilon,\tilde{*}_n)\}$ is precompact. Thus, it contains a convergent subsequence, and this proves that $\mathscr{X}_\epsilon$ is precompact.

Next, suppose that $(X^n_\epsilon,\tilde{*}_n) \to (Y,\bar{*})$. We will show that Y is ϵ-connected. As before, we may just assume that the sequence of base spaces, $\{(X_n,*_n)\}$, converges, also. Now, given $[\alpha]_\epsilon \in B(\tilde{*}_n,M\epsilon) \subset X^n_\epsilon$, for some sufficiently large $M \in \mathbb{N}$, there is a λ-chain in $[\alpha]_\epsilon$ by assumption 1 above. Specifically, the fact that $L([\alpha]_\epsilon) < M\epsilon$ implies that there is an ϵ-chain in $[\alpha]_\epsilon$ consisting of at most $2M+1$ points. We can ϵ-refine this chain to a λ-chain consisting of fewer than $(2M+1)N$ points. Lifting this chain to $\tilde{*}_n \in X^n_\epsilon$ gives us a λ-chain from $\tilde{*}_n$ to $[\alpha]_\epsilon$ of length less than $(2M+1)N\lambda$.

So, let $y \in Y$ be given, and let M be the smallest natural number such that

$$\frac{(2M-1)\epsilon+\lambda}{2} > d(\bar{*},y).$$

Then choose R so that $R > (2M+1)N\lambda$. Our assumption that $\lambda > \frac{\epsilon}{2}$ implies that $(2M+1)N\lambda > M\epsilon$, which, in turn, implies that $d(\bar{*},y) < R$. Choose n so that

$$d_{GH}\big(C(\tilde{*}_n,R),C(\bar{*},R)\big) < \frac{\epsilon-\lambda}{4}$$

and $d(\tilde{*}_n,\bar{*}) < (\epsilon-\lambda)/4$. Let $[\alpha]_\epsilon \in C(\tilde{*}_n,R)$ be such that $d([\alpha]_\epsilon,y) < (\epsilon-\lambda)/4$. We then have

$$d(\tilde{*}_n,[\alpha]_\epsilon) \leq d(\tilde{*}_n,\bar{*})+d(\bar{*},y)+d(y,[\alpha]_\epsilon) < 2\frac{\epsilon-\lambda}{4}+\frac{(2M-1)\epsilon+\lambda}{2}$$

$$\Rightarrow d(\tilde{*}_n,[\alpha]_\epsilon) < \frac{\epsilon-\lambda}{2}+\frac{(2M-1)\epsilon+\lambda}{2} = M\epsilon.$$

Thus, $[\alpha]_\epsilon$ lies in $C(\tilde{*}_n,R)$, as does a λ-chain from $\tilde{*}_n$ to $[\alpha]_\epsilon$. Denote this λ-chain by $\{\tilde{*}_n = z_0,z_1,\ldots,z_l = [\alpha]_\epsilon\}$. For each $i = 0,1,\ldots,l$, choose $y_i \in C(\bar{*},R)$ so that $d(z_i,y_i) < (\epsilon-\lambda)/4$, making sure to choose $y_0 = \bar{*}$ and $y_l = y$. Then $\{\tilde{*}_n = y_0,\ldots,y_l = y\}$ is an ϵ-chain, since

$$d(y_{i-1},y_i) \leq d(y_{i-1},z_{i-1})+d(z_{i-1},z_i)+d(z_i,y_i) < 2\frac{\epsilon-\lambda}{4}+\lambda = \frac{\epsilon+\lambda}{2} < \epsilon.$$

This shows that Y is ϵ-connected. ∎

We have already observed two facts: the ϵ-covers of a proper geodesic space are proper for every ϵ, and, for any $\epsilon > 0$, all geodesic spaces satisfy a uniform ϵ-BMR property, since every two-point ϵ-chain can be ϵ-refined to an $\frac{\epsilon}{2}$-chain consisting of 3 points by adding the midpoint. Thus, an immediate corollary of the previous result and the results from the previous section is the following.

Corollary 6.3.7 *Let $\mathscr{X}$ be a pointed, precompact collection of compact geodesic spaces, and, for any $\epsilon > 0$, let $\mathscr{X}_\epsilon$ be the corresponding pointed collection of ϵ-covers. Then $\mathscr{X}_\epsilon$ is precompact. Moreover, if $(Y, \bar{*})$ is the limit of any sequence $\{(X^n_\epsilon, \tilde{*}_n)\}$, where the corresponding sequence of base spaces, $\{(X_n, *_n)\}$, converges to $(X, *)$, then Y is ϵ-intrinsic (geodesic, in fact) and is isometric to a quotient of X_ϵ by a subgroup of $\pi_\epsilon(X)$.*

6.4 Convergence of Critical Values

In this section, we will utilize the previous results and the close homomorphism lemma to investigate the behavior of the critical values under convergence. We begin, however, with a useful result concerning *non*-critical values.

Lemma 6.4.1 *Let $\{X_n\}$ be a sequence of compact, chain-connected metric spaces converging to a compact metric space, X. Assume that there is a non-empty positive interval, (a, b), such that, for every n, X_n does not have any refinement critical values in this interval, or, equivalently, for each n and any $a < t_1 < t_2 < b$, the map $X_{t_1} \to X_{t_2}$ is surjective. Then X does not have any refinement critical values in the interval (a, b).*

Proof Let $a < \delta < \epsilon < b$ be given. Let $\alpha = \{x, y\}$ be any two-point ϵ-chain in X. Then $\lambda := d(x, y) < \epsilon$. Choose δ' so that $a < \delta' < \delta$. Now, fix n large enough that

$$d_{GH}(X_n, X) < \tau := \min\Big\{\frac{\epsilon - \lambda}{8}, \frac{\delta - \delta'}{2}\Big\},$$

and then choose $z, w \in X_n$ so that $d(x, z), d(w, y) < \tau$. Then

$$d(z, w) \leq d(z, x) + d(x, y) + d(y, w) < \frac{\epsilon - \lambda}{4} + \lambda.$$

Let $\epsilon' = \lambda + (\epsilon - \lambda)/4$, so that $\alpha' := \{z, w\}$ is an ϵ'-chain, and note that $\epsilon' < \epsilon$. By hypothesis, α' can be ϵ'-refined to a δ'-chain. Let $H' = \{\alpha' = \beta'_0, \beta'_1, \dots, \beta'_k\}$ be an ϵ'-homotopy taking α' to a δ'-chain. Denote each β'_i by $\beta'_i = \{z = u^i_0, \dots, u^i_{m_i} = w\}$. We will construct a homotopy, $H = \{\alpha = \beta_0, \beta_1, \dots, \beta_k\}$ in X inductively.

Clearly, we let $\beta_0 = \alpha$. Since $\alpha' = \beta'_0$ consists of only two points, β'_1 must be obtained by adding a point to β'_0, say $\beta'_1 = \{z, u, w\}$. Choose $v \in X$ so that $d(u, v) < \tau$, and then set $\beta_1 = \{x, v, y\}$. Then β_1 is an ϵ-chain, since

$$d(x, v) \leq d(x, z) + d(z, u) + d(u, v) < 2\tau + \lambda + \frac{\epsilon - \lambda}{4} < \lambda + \frac{\epsilon - \lambda}{2} < \epsilon$$

and likewise for $d(v, y)$. Thus, $\alpha = \beta_0 \sim_\epsilon \beta_1$.

Now, proceeding inductively, suppose for some $1 \leq i \leq k - 1$, we have constructed an ordered sequence of ϵ-chains from x to y, $\{\alpha = \beta_0, \beta_1, \dots, \beta_i\}$, such that each β_l has the same number of points as β'_l, β_{l+1} differs from β_l by a basic move for each $l = 0, \dots, i - 1$, and - if

$\beta'_l = \{z = u^l_0, u^l_1, \ldots, u^l_{m_l} = w\}$ - then $\beta_l = \{x = v^l_0, v^l_1, \ldots, v^l_{m_l} = y\}$ is such that $d(v^l_j, u^l_j) < \tau$ for every $0 \le l \le i$, $0 \le j \le m_l$. If β'_{i+1} is obtained by removing a point from β'_i - say

$$\beta'_i = \{z = u^i_0, \ldots, u^i_{j-1}, u^i_j, u^i_{j+1}, \ldots, u^i_{m_i} = w\}$$

$$\beta'_{i+1} = \{z = u^i_0, \ldots, u^i_{j-1}, u^i_{j+1}, \ldots, u^i_{m_i} = w\},$$

- then we let β_{i+1} be the chain β_i with that same corresponding point removed. Thus, if $\beta_i = \{x = v^i_0, \ldots, v^i_{j-1}, v^i_j, v^i_{j+1}, \ldots, v^i_{m_i} = y\}$, then we have $\beta_{i+1} = \{x = v^i_0, \ldots, v^i_{j-1}, v^i_{j+1}, \ldots, v^i_{m_i} = y\}$. Clearly β_{i+1} differs from β_i from a basic move, and β_{i+1} is an ϵ-chain since

$$\begin{aligned} d(v^i_{j-1}, v^i_{j+1}) &\le d(v^i_{j-1}, u^i_{j-1}) + d(u^i_{j-1}, u^i_{j+1}) + d(u^i_{j+1}, v^i_{j+1}) \\ &< 2\tau + \lambda + \frac{\epsilon - \lambda}{4} \\ &< \lambda + \frac{\epsilon - \lambda}{2} < \epsilon. \end{aligned}$$

If β'_{i+1} is obtained by adding a point to β'_i - say

$$\beta'_i = \{z = u^i_0, \ldots, u^i_j, u^i_{j+1}, \ldots, u^i_{m_i} = w\}$$

$$\beta'_{i+1} = \{z = u^i_0, \ldots, u^i_j, u, u^i_{j+1}, \ldots, u^i_{m_i} = y\},$$

then we choose a point $v \in X$ such that $d(u, v) < \tau$, and we define β_{i+1} to be the chain β_i with the point v added in to the same ordering position as u was added into β_i. So, if $\beta_i = \{x = v^i_0, \ldots, v^i_j, v^i_{j+1}, \ldots, v^i_{m_i} = y\}$, then $\beta_{i+1} = \{x = v^i_0, \ldots, v^i_j, v, v^i_{j+1}, \ldots, v^i_{m_i} = y\}$. As before, β_{i+1} differs from β_i by a basic move, and it is an ϵ-chain, since

$$\begin{aligned} d(v^i_j, v) &\le d(v^i_j, u^i_j) + d(u^i_j, u) + d(u, v) \\ &< 2\tau + \lambda + \frac{\epsilon - \lambda}{4} \\ &< \lambda + \frac{\epsilon - \lambda}{2} < \epsilon, \end{aligned}$$

and likewise for $d(v, v^i_{j+1})$.

This inductive process stops once we construct β_k, and this gives us an ϵ-homotopy, H, taking α to a chain β_k. But β'_k is a δ'-chain, and if $\beta'_k = \{z = u^k_0, \ldots, u^k_{m_k} = w\}$ and $\beta_k = \{x = v^k_0, \ldots, v^k_{m_k}\}$, then we have $d(v^k_j, u^k_j) < \tau$ for any $0 \le j \le m_k$. Thus, for any $1 \le j \le m_k$, we have

$$\begin{aligned} d(v^k_{j-1}, v^k_j) &\le d(v^k_{j-1}, u^k_{j-1}) + d(u^k_{j-1}, u^k_j) + d(u^k_j, v^k_j) \\ &< 2\tau + \delta' \\ &< \delta. \end{aligned}$$

Thus, β_k is a δ-chain. This shows that any two-point ϵ-chain in X can be ϵ-refined to a δ-chain, or, equivalently, that the map $X_\delta \to X_\epsilon$ is surjective. Since $a < \delta < \epsilon < b$ were arbitrary, the result follows. ∎

An immediate consequence of the previous lemma is an analog of the well-known result that the set of compact geodesic spaces is a closed subset of the set of compact metric spaces with respect to the Gromov-Hausdorff metric. Recall that a space, X, is refinable if $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$ is surjective for every $0 < \delta < \epsilon$.

Lemma 6.4.2 *Let $\mathscr{R}$ denote the collection of all compact, chain-connected, refinable metric spaces. Then $\mathscr{R}$ is closed in the Gromov-Hausdorff metric.*

Theorem 6.4.3 *Let $\{X_n\}$ be a sequence of compact, chain-connected metric spaces converging to a compact metric space, X. If a non-empty interval, (a, b), is a non-critical interval of X_n for every n, then (a, b) is a non-critical interval of X.*

Proof Our hypothesis means that the homomorphism $\Phi^n_{t_2t_1} : \pi_{t_1}(X_n) \to \pi_{t_2}(X_n)$ is an isomorphism for every n and every $a < t_1 < t_2 < b$. Along with the previous lemma, it also shows that, for every $a < t_1 < t_2 < b$, the homomorphism $\Phi_{t_2t_1} : \pi_{t_1}(X) \to \pi_{t_2}(X)$ is surjective, and this holds for each X_n, as well. Let $a < \delta < \epsilon < b$ be given. Then we need only show that $\Phi_{\epsilon\delta} : \pi_\delta(X) \to \pi_\epsilon(X)$ is injective.

Choose λ, δ', and ϵ' such that $a < \lambda < \delta' < \delta < \epsilon < \epsilon' \leq b$. Then choose n large enough that

$$d_{GH}(X_n, X) < \frac{1}{4} \min\{\delta' - \lambda, \delta - \delta', \epsilon - \delta, \epsilon' - \epsilon\}.$$

By part 2 of the close homomorphism lemma, we have surjective homomorphisms, $\theta_1 : \pi_{\delta'}(X_n) \to \pi_\delta(X)$ and $\theta_2 : \pi_\epsilon(X) \to \pi_{\epsilon'}(X_n)$, satisfying

$$\Phi^{X_n}_{\epsilon'\delta'} = \theta_2 \circ \Phi^X_{\epsilon\delta} \circ \theta_1.$$

But the map $\Phi^{X_n}_{\epsilon'\delta'}$ is an isomorphism by hypothesis, and all of the maps in the above commutativity relation are surjective. If $\Phi^X_{\epsilon\delta} : \pi_\delta(X) \to \pi_\epsilon(X)$ were not injective, then the map $\theta_2 \circ \Phi^X_{\epsilon\delta} \circ \theta_1 = \Phi^{X_n}_{\epsilon'\delta'}$ would not be either. This would be a contradiction. Hence, $\Phi^X_{\epsilon\delta}$ is an isomorphism. Since δ and ϵ were arbitrary, this shows that every map $\varphi_{\epsilon\delta} : X_\delta \to X_\epsilon$, for any $a < \delta < \epsilon < b$, is a bijection. Hence, (a, b) is a non-critical interval for X. ∎

The contrapositive of this last result immediately gives us an interesting and, perhaps, more enlightening interpretation. Essentially, critical values cannot suddenly appear out of nowhere in the limit; they must come from critical values of the spaces in the sequence. This is formally stated in the following corollary.

Corollary 6.4.4 *Let $\{X_n\}$ be a sequence of compact, chain-connected metric spaces converging to a compact metric space, X. If ϵ is a critical value of X, then there is a subsequence, $\{X_{n_k}\}$, and, for each k, a critical value $\epsilon_k \in Cr(X_{n_k})$ such that $\epsilon_k \to \epsilon$ in $\mathbb{R}$.*

Since critical values cannot suddenly appear in the limit, this brings up the other obvious question: can critical values suddenly *disappear* in the limit? It turns out that refinement critical values can suddenly disappear in the limit, as the following example shows.

Example 6.4.5 *Consider the space given in Example 4.1.2 (see Figure 4.1). There, the diagonal length, d, is greater than the gap length above it, which is what makes the upper gap length, l_1, an upper non-surjective critical value. As we saw in that example, the chain $\{x, y\}$, for all ϵ greater than but sufficiently close to l_1, is an ϵ-chain that cannot be ϵ-refined to an l_1-chain. Now, consider a sequence of such spaces, $\{X_n\}$, where the upper gap length remains fixed while the lower gap length decreases to a limiting value in such a way that the diagonal length, d, strictly decreases to a limiting value of l_1, the upper gap length. (One must also choose the height appropriately to make the dimensions work out this way, or else this phenomenon might* *not occur; that is just a simple matter of computation, however.) For each n, since the diagonal length is still greater than the upper gap length, l_1 is a refinement critical value of X_n by the same reasoning as before. However, in the limit space, where the diagonal length is exactly equal to the upper gap length, then l_1 is no longer an upper non-surjective critical value. In fact, for every $\epsilon > l_1$, $\{x, y\}$ is an ϵ-chain, and it can be ϵ-refined to a finer chain by first adding in u between x and y, and then adding in v between u and y, giving us the chain $\{x, u, v, y\}$. The dimensions can be arranged so that the lower gap length, in the limit space, has length strictly less than l_1, meaning that the chain $\{x, u, v, y\}$ is an l_1-chain. Thus, the refinement critical value, l_1, which is in the critical spectrum of each X_n, has disappeared upon passing to the limiting space.* ∎

However, as we will show in the next theorem, homotopy critical values *cannot* suddenly disappear in the limit when the spaces in question are sufficiently refinable, or at least refinable around the critical values. In other words, assuming sufficient refinability conditions, if $X_n \to X$ and there is a homotopy critical value $\epsilon_n \in Cr(X_n)$ for each n so that $\epsilon_n \to \epsilon$, then ϵ is a homotopy critical value of X. As has been mentioned several times, now, this should be further motivation for the study of the critical spectrum in the context of refinable spaces.

Theorem 6.4.6 *Let $\{X_n\}$ be a sequence of compact, chain-connected metric spaces converging to a compact metric space, X. Suppose, for each n, there is a homotopy critical value, $\epsilon_n \in Cr(X_n)$, such that $\epsilon_n \to \epsilon$ in $\mathbb{R}$, where $\epsilon > 0$. If there is some non-empty, positive interval, (a, b), that contains ϵ and is such that (a, b) contains no refinement critical values of X_n for every n, then ϵ is a homotopy critical value of X.*

Proof We first make the following observation: if there are values $\delta < \delta'$ such that the map $\varphi_{\delta'\delta} : X_\delta \to X_{\delta'}$ (or, equivalently, $\Phi_{\delta'\delta} : \pi_\delta(X) \to \pi_{\delta'}(X)$) is not injective, then there is a homotopy critical value of X between δ and δ'. To see why this is true, note that the assumption that $\varphi_{\delta'\delta}$ is not injective implies that there is a non-trivial δ-loop in X, say γ, that is δ'-null. Let $\delta^* = \inf\{t > \delta : \gamma \text{ is } t\text{-null}\}$. Then we have $\delta \leq \delta^* < \delta'$, where the last inequality holds because γ is δ'-null and every δ'-homotopy is a $(\delta' - t)$-homotopy for sufficiently small t. Note, also, that δ^* cannot be in the set above (i.e. γ cannot be δ^*-null), for if γ were δ^*-null, then it would be $(\delta^* - t)$-null for sufficiently small t, contradicting that δ^* is the infimum of this set. Thus, δ^* is the largest value for which γ is non-trivial, and for all $t > 0$, γ is $(\delta^* + t)$-null. In other words, the map $X_{\delta^*} \to X_{\delta^*+t}$ is non-injective for all $t > 0$. This makes δ^* critical.

Also, note that, by Lemma 6.4.1, the fact that (a, b) contains no refinement critical values of X_n for any n implies that this interval has the same property with respect to X. So, for any $a < t_1 < t_2 < b$, the map $X_{t_1} \to X_{t_2}$ is surjective.

We will prove the result by using the first observation to show that, for all sufficiently large $m \in \mathbb{N}$, there is a critical value of X in the interval $(\epsilon - 1/m, \epsilon + 1/m)$. Such a sequence of critical values must converge to ϵ, and the result will then follow from the fact that the critical spectrum contains all of its positive limit points.

So, fix a natural number m large enough that $(\epsilon - \frac{1}{m}, \epsilon + \frac{1}{m}) \subset (a, b)$ and $\epsilon - \frac{1}{m} > a$. Choose n large enough so that $|\epsilon_n - \epsilon| < \frac{1}{4m}$ and

$$d_{GH}(X_n, X) < \frac{1}{16m}.$$

We will assume that $\epsilon_n \geq \epsilon$, so that $\epsilon \leq \epsilon_n < \epsilon + \frac{1}{4m}$; the case $\epsilon_n < \epsilon$ is handled in exactly the same way. Since ϵ_n is a homotopy critical value of X_n, we can find δ arbitrarily close to ϵ_n so that the map between $X^n_{\epsilon_n}$ and X^n_δ is not injective. Thus, there are two subcases to consider. Suppose, first, that there are values, δ, *greater than* but arbitrarily close to ϵ_n such that the map $X^n_{\epsilon_n} \to X^n_\delta$ is not injective. Then, in particular, there is some δ such that $\epsilon \leq \epsilon_n < \delta < \epsilon + \frac{1}{4m}$ and the map $X^n_{\epsilon_n} \to X^n_\delta$ is not injective. Since the map $X^n_\epsilon \to X^n_{\epsilon_n}$ is surjective, it further follows that the map $X^n_\epsilon \to X^n_\delta$ is not injective. Finally, this implies that the map $X^n_\epsilon \to X^n_{\epsilon+1/4m}$ is not injective. Now, we clearly have the inequality

$$\epsilon - \frac{1}{m} < \epsilon - \frac{1}{4m} < \epsilon < \epsilon + \frac{1}{4m} < \epsilon + \frac{1}{2m},$$

and all of these values lie in (a, b). We also have the following:

$$d_{GH}(X_n, X) < \frac{1}{4}\frac{1}{4m} < \frac{1}{4}\frac{3}{4m} = \frac{1}{4}\left(\left(\epsilon - \frac{1}{4m}\right) - \left(\epsilon - \frac{1}{m}\right)\right),$$

$$d_{GH}(X_n, X) < \frac{1}{4}\frac{1}{4m} = \frac{1}{4}\left(\epsilon - \left(\epsilon - \frac{1}{4m}\right)\right) = \frac{1}{4}\left(\left(\epsilon + \frac{1}{4m}\right) - \epsilon\right),$$

$$d_{GH}(X_n, X) < \frac{1}{4}\frac{1}{4m} = \frac{1}{4}\left(\left(\epsilon + \frac{1}{2m}\right) - \left(\epsilon + \frac{1}{4m}\right)\right).$$

Thus, the close homomorphism lemma implies that we have surjective homomorphisms

$$\theta_1 : \pi_{\epsilon-1/4m}(X) \to \pi_\epsilon(X_n) \quad \text{and} \quad \theta_2 : \pi_{\epsilon+1/4m}(X_n) \to \pi_{\epsilon+1/2m}(X),$$

and these satisfy the commutativity relation

$$\Phi^X_{\epsilon+1/2m,\epsilon-1/4m} = \theta_2 \circ \Phi^{X_n}_{\epsilon+1/4m,\epsilon} \circ \theta_1.$$

But all three of the homomorphisms on the right are surjective, and the central one is not injective. Hence, the map on the right is not injective, from which it follows that $\Phi^X_{\epsilon+1/2m,\epsilon-1/4m}$ is not injective. Thus, there is some critical value of X in the interval $(\epsilon - 1/m, \epsilon + 1/m)$.

For the second subcase, suppose there are values, δ, *less than* but arbitrarily close to ϵ_n such that the map $X^n_\delta \to X^n_{\epsilon_n}$ is not injective. Then, in particular, there is some δ such that $\epsilon - 1/4m < \delta < \epsilon_n$ and the map $X^n_\delta \to X^n_{\epsilon_n}$ is not injective. This implies that the map $X^n_\delta \to X^n_{\epsilon+1/4m}$ is not injective, and this further implies the same conclusion for the map $X^n_{\epsilon-1/4m} \to X^n_{\epsilon+1/4m}$ (again, using surjectivity). We have the inequality

$$\epsilon - \frac{1}{m} < \epsilon - \frac{1}{2m} < \epsilon - \frac{1}{4m} < \epsilon + \frac{1}{4m} < \epsilon + \frac{1}{2m},$$

and all of these values lie in (a, b). Reasoning as before, it is evident that $d_{GH}(X_n, X)$ is less than $\frac{1}{4}$ times the difference between any consecutive numbers in this inequality. So, applying the close homomorphism lemma again, we obtain surjective homomorphisms

$$\theta_1 : \pi_{\epsilon-1/2m}(X) \to \pi_{\epsilon-1/4m}(X_n) \quad \text{and} \quad \theta_2 : \pi_{\epsilon+1/4m}(X_n) \to \pi_{\epsilon+1/2m}(X)$$

$$\Phi^X_{\epsilon+1/2m,\epsilon-1/2m} = \theta_2 \circ \Phi^{X_n}_{\epsilon+1/4m,\epsilon-1/4m} \circ \theta_1.$$

Again, since all three maps on the right are surjective, and since the central one is non-injective, the map $\Phi_{\epsilon+1/2m,\epsilon-1/2m}$ is non-injective. So, in this case, also, there is a critical value of X in the interval $(\epsilon - 1/m, \epsilon + 1/m)$. ∎

Obviously the previous result generalizes immediately to the case in which there is a subsequence, $\{X_{n_k}\}$, with homotopy critical values $\epsilon_{n_k} \in Cr(X_{n_k})$ converging to $\epsilon > 0$.

Combining this theorem with Corollary 6.4.4, we obtain the following.

Corollary 6.4.7 *Let $\{X_n\}$ be a sequence of compact, chain-connected, refinable spaces converging to a compact metric space, X. Then $\epsilon > 0$ is a critical value of X if and only if there is a subsequence, $\{X_{n_k}\}$, with critical values, $\epsilon_k \in Cr(X_{n_k})$, such that $\epsilon_k \to \epsilon$. Since geodesic spaces are refinable, this holds for any convergent sequence of compact geodesic spaces, as well.*

We will conclude with a general finiteness theorem for metric spaces and geodesic spaces.

Lemma 6.4.8 *Let $\{X_n\}$ be a sequence of compact, chain-connected, refinable metric spaces converging to a compact metric space, X, and assume that there is some $\tau > 0$ such that $\inf Cr(X_n) \geq \tau$ for all n. Then there exists $N \in \mathbb{N}$ such that for any $n, m \geq N$ and any $0 < \delta, \epsilon < \tau$, $\pi_\delta(X_m)$ is isomorphic to $\pi_\epsilon(X_n)$. In particular, there are only finitely many isomorphism classes among the groups $\pi_\epsilon(X_n)$ for $n \geq 1$ and $0 < \epsilon < \tau$. Moreover, there also exists $M \in \mathbb{N}$ so that, for any $0 < \delta, \epsilon < \tau$ and $n \geq M$, $\pi_\epsilon(X)$ is isomorphic to $\pi_\delta(X_n)$.*

Proof Since no X_n has any critical values less than τ, we know that for any $n \geq 1$ and any $0 < \delta, \epsilon \leq \tau$, the groups $\pi_\delta(X_n)$ and $\pi_\epsilon(X_n)$ are isomorphic. Since $\{X_n\}$ is Cauchy, there is some N such that $n \geq N$ implies that $d_{GH}(X_n, X_N) < \tau/40$. We have the inequality

$$\frac{\tau}{10} < \frac{\tau}{5} < \frac{3\tau}{10} < \frac{2\tau}{5} < \frac{\tau}{2}.$$

We further have, for any $n \geq N$,

$$d_{GH}(X_n, X_N) < \frac{1}{4}\frac{\tau}{10} = \frac{(k+1)\tau}{10} - \frac{k\tau}{10}$$

for each $k = 1, 2, 3, 4$. This, along with the refinability hypothesis, allows us to use the close homomorphism lemma to obtain surjective homomorphisms

$$\theta_1 : \pi_{\tau/5}(X_n) \to \pi_{3\tau/10}(X_N) \text{ and } \theta_2 : \pi_{2\tau/5}(X_N) \to \pi_{\tau/2}(X_n)$$

$$\Phi^{X_n}_{\tau/2,\tau/5} = \theta_2 \circ \Phi^{X_N}_{2\tau/5,3\tau/10} \circ \theta_1.$$

The maps $\Phi^{X_n}_{\tau/2,\tau/5}$ and $\Phi^{X_N}_{2\tau/5,3\tau/10}$ are isomorphisms. This implies that θ_1 and θ_2 must be, also. In fact, if θ_1 were not injective, then $\theta_2 \circ \Phi^{X_N}_{2\tau/5,3\tau/10} \circ \theta_1$ would not be either, contradicting that $\Phi^{X_n}_{\tau/2,\tau/5}$ is an isomorphism. So, θ_1 is an isomorphism, and it now follows from the composition relation above that θ_2 is, also. Therefore, $\pi_{\tau/2}(X_n)$ is isomorphic to $\pi_{2\tau/5}(X_N)$, which, in turn, is isomorphic to $\pi_{\tau/2}(X_N)$.

Now, this holds for any $n \geq N$. So, let $n, m \geq N$ and $0 < \delta, \epsilon < \tau$ be given. Then

$$\pi_\epsilon(X_n) \cong \pi_{\tau/2}(X_n) \cong \pi_{\tau/2}(X_N) \cong \pi_{\tau/2}(X_m) \cong \pi_\delta(X_m),$$

which proves the first claim. The second claim is an immediate consequence of this, since there are, at most, $N-1$ isomorphism classes among the groups $\pi_\epsilon(X_n)$ for $1 \leq n \leq N-1$, $0 < \epsilon < \tau$, and we've just shown that the groups $\pi_\epsilon(X_n)$, with $n \geq N$ and $0 < \epsilon < \tau$, are isomorphic.

Finally, the refinability of the spaces, $\{X_n\}$, implies that X is refinable, as well. Moreover, X cannot have any critical values in the interval $(0, \tau)$, for if it did, Corollary 6.4.4 would imply that a subsequence of $\{X_n\}$ has critical values in $(0, \tau)$, also. Thus, we can apply the same argument as above (in fact, just replace X_N with X) to show that, for all sufficiently large n and $0 < \delta, \epsilon < \tau$, $\pi_\epsilon(X)$ is isomorphic to $\pi_\delta(X_n)$. ∎

Corollary 6.4.9 *If, in addition to the conditions of the previous lemma, the spaces are also geodesic and semi-locally simply connected, then there exists $N \in \mathbb{N}$ such that $\pi_1(X_n)$ is isomorphic to $\pi_1(X_m)$ for all $n, m \geq N$ and $\pi_1(X_n)$. In particular, there are only finitely many isomorphism classes among the fundamental groups, $\{\pi_1(X_n)\}_{n\geq 1}$.*

Proof In light of the previous lemma, we need only prove the following general result: the fundamental group of a compact, semi-locally simply connected geodesic space, X, is isomorphic to $\pi_\epsilon(X)$ for all sufficiently small ϵ. By Lemma 3.1.7, we know that we have a surjective homomorphism, $h_\epsilon : \pi_1(X) \to \pi_\epsilon(X)$ for every $\epsilon > 0$, namely the map taking every path loop homotopy class to the ϵ-class of strong ϵ-chains along that path. So, verifying the result is just a matter of showing that the semi-local simple connectivity of X implies that this map is injective for small ϵ.

Since X is compact and semi-locally simply connected, there is some $r > 0$ with the property that every path loop in X lying in a ball of radius r is nullhomotopic. Fix ϵ positive but small enough that $\frac{3\epsilon}{2} < r$. We will show that h_ϵ is injective.

Let γ be a path loop at the base point, $*$, along which there is a strong ϵ-chain that is ϵ-null (i.e. the path homotopy class of γ is in $\ker h_\epsilon$). We know that there exists an ultra ϵ-chain along γ, and we know that an ultra ϵ-chain along a path is also a strong ϵ-chain along that path. We further know that any two strong ϵ-chains along a path are ϵ-homotopic. Thus, we can assume without loss of generality that the strong ϵ-chain along γ is, in fact, an ultra ϵ-chain. Then, by Lemma 3.1.8, γ is path homotopic to a finite product of path loops of the form $\alpha\beta\alpha^{-1}$, where α is a path from $*$ to some point, x, and β is a loop at x lying in a ball of radius $3\epsilon/2$. So, each β lies in a ball of radius r, implying that it is nullhomotopic. Thus, each loop, $\alpha\beta\alpha^{-1}$, will be nullhomotopic, and γ will also, showing that h_ϵ is injective for ϵ such that $\frac{3\epsilon}{2} < r$. ∎

These results immediately yield the following finiteness theorem. The first part of the this theorem is one of the most general finiteness theorems proved to date, at least with regard to the types of spaces to which it applies.

Theorem 6.4.10 *Let $\mathscr{X}$ be a precompact collection of chain-connected, compact, refinable metric spaces with the property that there is some $\tau > 0$ such that $\inf Cr(X) \geq \tau$ for all $X \in \mathscr{X}$. Then there are only finitely many isomorphism classes among the groups, $\pi_\epsilon(X)$, for $X \in \mathscr{X}$ and $0 < \epsilon < \tau$. In particular, for any $0 < \epsilon < \tau$ and any convergent sequence $\{X_n\} \subset \mathscr{X}$ with $X = \lim_n X_n$, the groups $\pi_\epsilon(X_n)$ are eventually isomorphic to $\pi_\epsilon(X)$.*

If, in addition to the previous conditions, the spaces in $\mathscr{X}$ are geodesic and semi-locally simply connected (e.g. Riemannian manifolds), then there are only finitely isomorphism classes among the fundamental groups of the spaces in $\mathscr{X}$. Moreover, if $\{X_n\}$ is any convergent sequence in $\mathscr{X}$, then the fundamental groups, $\pi_1(X_n)$, are eventually isomorphic.

The condition that the critical spectra be uniformly bounded below can be strong, but it still allows for some broad applicability. In the first part of this theorem, where semi-local simple connectivity is not assumed, this uniform lower bound on the critical spectra does exclude, for example, "Hawaiian earring type" singularities, but it does not exclude the possibility of non-trivial path loops that are path homotopic to arbitrarily small loops (i.e. spaces that are not homotopically Hausdorff). However, the latter part of this last theorem could lead to other research threads in which one examines geometric properties of Riemannian manifolds and geodesic spaces (e.g. curvature, systoles, etc.) that might yield a uniform lower bound on the critical spectra.

Bibliography

[1] V. Berestovskii, C. Plaut, *Uniform Universal Covers of Uniform Spaces*, Topology and its Applications 154 (2007) 1748-1777.

[2] V. Berestovskii, C. Plaut, *Covering Group Theory for Topological Groups*, Topology and its Applications 114 (2001) 187 - 199.

[3] D. Burago, Y. Burago, S. Ivanov, *A Course in Metric Geometry (Graduate Studies in Mathematics, Vol. 33)*, American Mathematical Society (2001).

[4] J. Conant, C. Plaut, J. Wilkins, et al *Discrete Homotopy Theory and Critical Values of Metric Spaces*, preprint, 35 pages.

[5] P. de la Harpe, *Topics in Geometric Group Theory*, The University of Chicago Press, (2000).

[6] J. Milnor, *A Note on Curvature and the Fundamental Group*, Journal of Differential Geometry 2, (1968), 1-7.

[7] P. Petersen, *Riemannian Geometry (Second Edition)*, Springer (2006).

[8] C. Plaut, *Quotients of Uniform Spaces*, Topology and its Applications 153 (2006) 2430 - 244.

[9] C. Plaut, J. Wilkins, *Homotopy Critical Values and Gromov-Hausdorff Convergence*, preprint, 35 pages.

[10] E. Spanier, *Algebraic Topology*, Springer (December 1994).

[11] C. Sormani, *How Riemannian Manifolds Converge: A Survey*, to appear in Progress in Mathematics.

[12] C. Sormani, G. Wei, *Hausdorff Convergence and Universal Covers*, Transactions of the American Mathematical Society Vol. 353, No. 9 (2001) 3585-3602.

[13] C. Sormani, G. Wei, *The Covering Spectrum of a Compact Length Space*, J. Differential Geometry 67 (2004) 35 - 77.

[14] C. Sormani, G. Wei, *The Cut-Off Covering Spectrum*, Transactions of the American Mathematical Society Vol. 362, No. 5 (2010) 2339 - 2391.

[15] C. Sormani, G. Wei, *Universal Covers for Hausdorff Limits of Noncompact Spaces*, Transactions of the American Mathematical Society Vol. 356, No. 3 (2003) 1233 - 1270.

Vita

I attended undergraduate school at the Georgia Institute of Technology in my hometown of Atlanta, Georgia. I originally majored in mechanical engineering for three years before ultimately obtaining a degree in applied mathematics in 2003. I then attended Clemson University where I received a master's degree in Mathematics in May 2005. My master's thesis research focused on the metric structure of Lorentzian manifolds, including a simplified linear-algebraic version of the proof that a Riemannian manifold admits a Lorentzian metric if and only if it admits a smooth one-dimensional distribution. After two years at Clemson, I came to the University of Tennessee in Fall 2005, and I began my doctoral work with Dr. Conrad Plaut. Broadly speaking, my general research interest is in Riemannian geometry and the influence of geometric properties on the topology of the manifolds. My doctoral research focused on more general metric geometry, particularly singular and non-smooth metric spaces, which would include, for instance, Gromov-Hausdorff limits of Riemannian manifolds, fractals, and various continua such as solenoids. I also have research interests in analysis on fractals and, more generally, non-smooth metric spaces.