August 2023

Collapsibility and Z-Compactifications of CAT(0) Cube Complexes

Daniel L. Gulbrandsen
University of Wisconsin-Milwaukee

COLLAPSIBILITY AND Z-COMPACTIFICATIONS OF CAT(0) CUBE COMPLEXES

by

Daniel L. Gulbrandsen

A Dissertation Submitted in

Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

in Mathematics

at

The University of Wisconsin–Milwaukee

August 2023

ABSTRACT

COLLAPSIBILITY AND Z-COMPACTIFICATIONS OF CAT(0) CUBE COMPLEXES

by

Daniel L. Gulbrandsen

The University of Wisconsin–Milwaukee, 2023
Under the Supervision of Professor Craig R. Guilbault

We extend the notion of collapsibility to non-compact complexes and prove collapsibility of locally-finite CAT(0) cube complexes. Namely, we construct such a cube complex X out of nested convex compact subcomplexes $\{C_i\}_{i=0}^{\infty}$ with the properties that $X = \cup_{i=0}^{\infty} C_i$ and C_i collapses to C_{i-1} for all $i \geq 1$.

We then define bonding maps r_i between the compacta C_i and construct an inverse sequence yielding the inverse limit space $\varprojlim\{C_i, r_i\}$. This will provide a new way of Z-compactifying X. In particular, the process will yield a new Z-boundary, called the cubical boundary.

I dedicate this work to my Parents who have always supported me.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I would like to express my gratitude to the department of mathematical sciences at Universiy of Wisconsin, Milwaukee, and to the many colleagues who have become close friends of mine. The environment at UWM is as welcoming and conducive for doing mathematics as any I have found.

I am specifically indebted to the topology group at UWM. This work would not have even seen the light of day without their involvement. Most helpful were the *top secret* topology meetings organized by the students.

I would like to thank my advisor Craig R. Guilbault for their patience, and for their enthusiasm for doing mathematics.

1 INTRODUCTION

Collapses were first employed by Whitehead in [Whi50] in an attempt to replace methods in homotopy theory with more combinatorial ones. This more combinatorial approach, first applied to simplicial complexes, shortly thereafter to more general polyhedral complexes, and then to CW complexes, became known as *simple homotopy theory*. See [Coh73] for general reference. Under this approach two spaces are *simple homotopy equivalent* if one can be transformed into the other via a sequence of combinatorial moves called collapses and expansions, and a space is called *collapsible* if there exists a sequence of collapses that transform it to a point.

Classically, the definition of collapsible was applied to compact spaces only (spaces built out of finitely many simplices). It is somewhat surprising that very little has been done in extending the definition to non-compact spaces in the literature. In this paper, we provide such an extension, see Definition 3.8. We then prove the following, which appears on page 67.

Theorem 3.21. *All locally-finite CAT(0) cube complexes (finite or infinite) are collapsible.*

Even if one begins with a collapsible cube complex X, there is no guarantee that a chosen sequence of collapses and expansions will collapse X to a point. In the simplicial setting, it was shown in [LN19] just how wrong things can go. In that article, the authors characterize the ways in which an n-simplex can collapse to a subcomplex from which no further collapses may be made. In the cubical setting, *Bing's house with two rooms* is an example of a square complex that is not collapsible, even though it is contractible. However, upon crossing this complex with an interval, it becomes collapsible.

Our approach makes heavy use the rich combinatorial structure enjoyed by CAT(0) cube complexes and avoids any obstruction to collapsibility one might incur otherwise. More specifically, given a pointed locally-finite CAT(0) cube complex (X, v_0), we describe (Theorem 3.17) a sequence of nested convex compact subcomplexes $\{C_i\}_{i=0}^{\infty}$ such that

1. $\bigcup_{i=1}^{\infty} C_i = X$; and

2. C_{i+1} collapses to C_i via finitely many elementary collapses for all $i \geq 0$.

Throughout this exposition, we will rely heavily on the theory of hyperplanes in CAT(0) cube complexes, defined in Definition 2.29. These are particular subsets which are well-suited to combinatorial methods. To illustrate their usefulness, in [Sag95], Sageev used hyperplanes to generalize results in Bass-Serre theory to splittings of groups, for groups acting on CAT(0) cube complexes. And, in [AGM12], Agol, Groves, and Manning identified certain hyperplanes in *non-positively curved* cubed complexes to help resolve the Virtual Haken Conjecture.

We will use properties of hyperplanes to construct particularly nice maps $\phi_i : C_i \times [0,1] \to C_{i-1}$. At time $t = 1$, these maps become retractions. Then, in chapter 3, we use the compacta $\{C_i\}$ to construct a compactification of X. Using the r_i's as bonding maps, we assemble the collection $\{C_i\}$ into an inverse sequence

$$\{v_0\} = C_0 \xleftarrow{r_1} C_1 \xleftarrow{r_2} C_2 \xleftarrow{r_3} \cdots$$

In Lemma 4.11, we prove that X embeds in the product space ΠC_i. We then show that the closure of this embedding is, in fact, the inverse limit space. Let $\overline{X} = \varprojlim\{C_i, r_i\}$. Then $\overline{X}$ gives a compactification with desirable properties. In particular, on page 79, we prove the following.

Theorem 4.12. *For a locally-finite CAT(0) cube complex X, The space $\overline{X}$ gives a Z-compactification for X, the boundary of which $\overline{X} \setminus X$ is a Z-set.*

We call $\overline{X}$ the *cubical compactification* of X, and we call the remainder space $\partial_{\square} X = \overline{X} \setminus X$ the *cubical boundary* of X.

The notion of a Z-set was first introduced by R Anderson in [And67]. Then, in [BM91] and [Bes96], the authors popularized their use in geometric group theory. In particular, Bestvina in [Bes96] introduced *Z-structures*, which combine properties of groups and properties of Z-sets for groups acting geometrically.

In the case where X admits a geometric group action, $\partial_{\square} X$ need not satisfy all the conditions for being a Z-structure. Bestvina's *nullity condition* will not hold in general. The cubical boundary does however meet the conditions for a *weak Z-structure*. For more information on weak Z-structures see [Gui13] and [Gui16].

An immediate consequence of Theorem 4.12 is that $\partial_{\square} X$ is *shape equivalent* to the visual boundary of X.

A number of questions concerning the cubical compactification and boundary remain. Some of these are listed in chapter 4.

2 PRELIMINARIES

Our main objects of study will be *CAT(0) cube complexes*, which are certain types of polyhedral complexes with rich combinatorial structure. However, we begin with a brief description of the more familiar setting of simplicial complexes. This is not a waste of time though as many of the arguments and techniques that follow will rely on the theory of simplicial complexes. Then we define cube complexes and discuss their geometry. This will lead us to CAT(0) cube complexes. After introducing some key examples, we end this chapter with a series of lemmas that will be needed in the following chapters.

2.1 Simplicial Complexes

The study of simplicial complexes is pervasive in mathematics. In particular, they are familiar objects within the studies of algebraic topology and combinatorics. For the purposes of this paper, we will use simplicial complexes to determine geometric properties of cube complexes: Simplicial complexes provide a combinatorial means of telling whether or not a cube complex is non-positively curved. Additionally, simplicial complexes will be used to determine when subcomplexes are convex. These ideas will be examined in detail in the following sections. Readers already familiar with simplicial complexes would not suffer greatly by skipping to section 1.2 *Cube Complexes* and returning to read the lemmas as needed.

We prefer the definition of an *abstract simplicial complex*.

Definition 2.1 (Simplicial Complex, n-Simplex)**.** We define a **simplicial complex** K as a collection of finite sets that is closed under taking subsets. An **n-simplex** (or simply **simplex** when n is obvious) is the power set of an $n+1$-element set minus the empty set. We say the dimension of an n-simplex is n.

For an example of a simplex, the 2-simplex σ can be described by the following collection of subsets of the set $\{a, b, c\}$: $\sigma = \{\{a\}, \{b\}, \{c\}, \{a, b\}, \{b, c\}, \{a, c\}, \{a, b, c\}\}$. A 0-simplex will also be referred to as a vertex, and a 1-simplex will often be referred to as an edge. For a simplicial complex K, we call the union of n-simplices in K the n-**skeleton** of K and denote it by K^n. The 0-skeleton and 1-skeleton of a simplicial complex are called its vertex set and edge set, respectively. We call the vertex set of a simplex a **spanning set** for that simplex or say that the vertices span the simplex. For instance, the vertices a, b, c span σ from above. We say the dimension of a simplicial complex K is n if K contains n-simplices but does contain any $n+1$-simplices. Note that if K does not contain any $n+1$-simplices, then it cannot contain an m-simplex for any $m > n$.

Definition 2.2 ((Simplicial) Subcomplex, Full Subcomplex)**.** A subcollection L of K is a **subcomplex** if L is itself closed under taking subsets. We say a subcomplex L is **full** in K if whenever L contains the spanning set of a simplex in K, L contains the simplex.

For instance, K^n is a subcomplex of K. However, it is not necessarily a full subcomplex. In fact, unless n equals the dimension of K, K^n is guaranteed not to be a full subcomplex. This is because it will contain vertices that span a simplex that is of dimension larger than n.

The question of whether or not a subcomplex of a simplicial complex is full will be of paramount interest to us. Knowing when certain simplical subcomplexes are full will provide a certain test for determining whether or not a subcomplex of a cube complex is convex.

Now, let σ be an n-simplex. A *face*, f, of σ is a k-simplex ($k \leq n$) spanned by a $k+1$-element subset of the $n+1$-element set of vertices that span σ. If f is a k-simplex for some $k < n$, we call f a *proper face* of σ. In the 2-simplex spanned by vertices a, b, c, the collection $\{\{a\}, \{b\}, \{a, b\}\}$ is a proper face, while the entire 2-simplex is a face which is not proper.

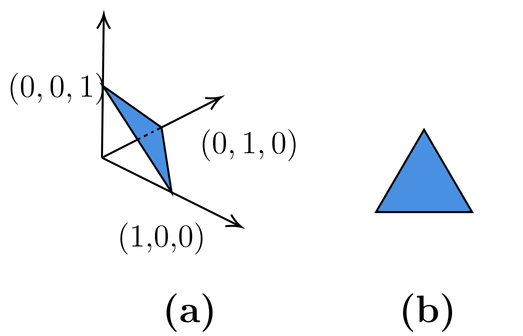

Figure 1: Geometric realization of a 2-simplex

Although we have defined simplicial complexes abstractly, we will at times make reference to the *geometric realization* of a simplex. For an n-simplex σ, one forms the geometric realization of σ by taking the convex hull of points represented by the standard basis vectors in $\mathbb{R}^{n+1}$. For example, the geometric realization of a 1-simplex is the convex hull of points $(1, 0)$ and $(0, 1)$ in $\mathbb{R}^2$. In figure 1, (a) shows the geometric realization of a 2-simplex. Figure (b) shows the same 2-simplex, but without the distracting coordinate axes. We let σ denote both a simplex and its geometric realization.

Given the geometric realization of σ, there is a natural way of identifying the faces of σ. Simply take the convex hull of some subset of points represented by basis vectors in the geometric realization. For instance, the convex hull of points $(0, 0, 1)$ and $(1, 0, 0)$ in Figure 1 is a face. Then, to define the geometric realization of a simplicial complex K, let $N = |K^0|$ and identify the vertices of K with the vertices of an $N-1$ simplex σ in $\mathbb{R}^N$. The **geometric realization** of K is the union of faces f of σ such that the vertex set in K corresponding to the vertices of f spans a simplex in K. We will also write K to denote the geometric realization of K.

We now state a few definitions and lemmas concerning simplicial complexes that will be used further below.

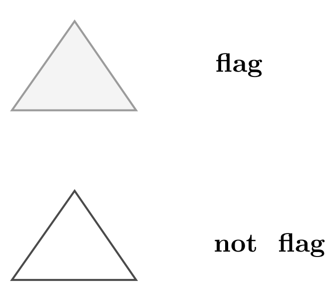

Figure 2: A 2-simplex and a 1-dimensional simplicial complex

Definition 2.3 (Flag Complex)**.** A simplicial complex K is called a **flag complex** if whenever K contains the 1-skeleton of a simplex σ, K contains σ. We will often write K is flag to mean it is a flag complex.

Clearly, a full subcomplex L of a flag complex K is itself a flag complex. Consider a collection of 1-simpices $\sigma_1, \ldots, \sigma_n$ in L which constitute the 1-skeleton of a simplex σ. Since K is flag, σ is in K. Since L contains the vertex set of σ and is full, σ is in L. Thus L is flag. The following example shows that the converse is false.

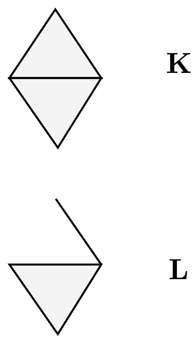

Figure 3: $L \leq K$ is flag but not full

Example 2.4 *Figure 3 shows the geometric realization of a flag subcomplex L of K that is not a full subcomplex. Note that, in this example, K is flag.*

If K is a simplicial complex, then the **cone** of K is defined as the union of simplices in K, an additional vertex p, called a **cone point**, and simplices that see p appended to each simplex in K. For instance, if $K = \{\{a\}, \{b\}, \{a, b\}\}$, then the cone of K is $\{\{a\}, \{b\}, \{p\}, \{a, b\}, \{a, p\}, \{b, p\}, \{a, b, p\}\}$. We denote the cone of a simplicial complex K with cone point p as $K * \{p\}$. Note that, if f is a face of K, then $f * \{p\}$ is a face of $K * \{p\}$. Moreover, this face is proper if f is a proper face in K.

Another useful fact about full subcomplexes is the following:

Lemma 2.5. *Let K and L be subcomplexes of a simplicial complex S. If K is full in L and L is full in S, then K is full in S.*

Proof. Let $v_1, \ldots, v_n$ be vertices in K that span a simplex σ in S. Since $K \leq L$, $v_1, \ldots, v_n \in L$. Hence $\sigma \in L$. Since K is full in L, $\sigma \in K$. $\square$

Lemma 2.6. *A simplicial complex K is flag if and only if $K * \{p\}$ is flag.*

Proof. Assume $K * \{p\}$ is flag, and suppose that K contains the 1-skeleton of an n-simplex, σ. Then σ contains $n+1$ vertices, say $v_1, \ldots, v_{n+1}$. As $K * \{p\}$ is flag, σ belongs to $K * \{p\}$. But since each of $v_1, \ldots, v_{n+1}$ belong to K, and the only additional simplices that $K * \{p\}$ contributes necessarily contain the vertex p, it must be the case that σ belongs to K.

Conversely, suppose K is flag and let τ be an n-simplex whose 1-skeleton is contained in $K*\{p\}$ and with p a vertex in τ. The vertices that span τ can then be written as $v_1, \ldots, v_n, p$. Observe that the 1-skeleton of the $n-1$-simplex spanned by vertices $v_1, \ldots, v_n$ belongs to K, as these 1-simplices do not contain p. Since K is flag, this $n-1$-simplex then belongs to K. By definition of cone of a simplex, it follows that τ belongs to $K * \{p\}$. □

Let K be a subcomplex of X where X contains a copy of $K * \{p\}$. Then, Lemma 2.6 describes when $K * \{p\}$ is a flag subcomplex. In the following lemma, we show that a similar statement holds to determine when $K * \{p\}$ is a full subcomplex. We do however require that X be a flag complex.

Lemma 2.7. *Let X be a flag simplicial complex and K a subcomplex. If there is a vertex p in X such that $K * \{p\}$ is a subcomplex of X, then $K * \{p\}$ is full if and only if K is full.*

Proof. Suppose K is full. Let $v_1, \ldots, v_n$ be vertices in K such that $v_1, \ldots, v_n, p$ span a simplex σ in X. Then $v_1, \ldots, v_n, p$ are pairwise adjacent in X. Since K is full, the 1-skeleton spanned by $v_1, \ldots, v_n$ is contained in K. Since K is flag (by Lemma 2.6), the simplex, say τ, spanned by $v_1, \ldots, v_n$ is in K and $\sigma = \tau * \{p\}$. Thus $\sigma \in K * \{p\}$ by definition.

Conversely, suppose $K * \{p\}$ is full and that K contains vertices $v_1, \ldots, v_n$ that span a simplex σ in X. By supposition, $\sigma \in K * \{p\}$ and it then follows from the definition of a simplicial cone that $\sigma \in K$. □

Definition 2.8 (Simplicial Map, Isomorphism)**.** Let K and K' be simplicial complexes. A map f from the vertex set of K to the vertex set of K' is called **simplicial** if f sends every simplex in K to a simplex in K'. We say two simplicial complexes K and K' are **isomorphic** if there is a bijective simplicial map f between them with simplicial inverse, and we call f a **simplicial isomorphism**.

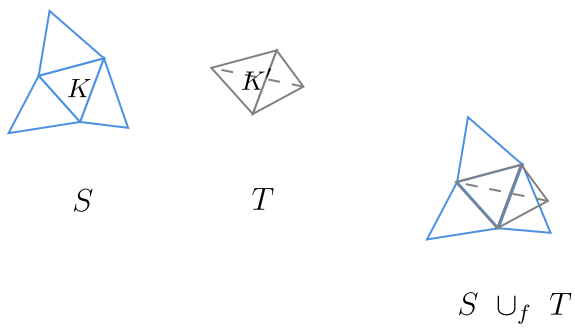

Figure 4: Simplicial complexes S and T attched along isomorphic subcomplexes to obtain the simplicial adjunction space $S \cup_f T$

Now, let S and T be simplicial complexes that contain isomorphic subcomplexes K and K', respectively. We describe a way of attaching S and T together along their isomorphic subcomplexes. Let $f : K \to K'$ be a simplicial isomorphism. The **simplicial adjunction space** $S \cup_K T$ is defined as the union of simplices in S with the union of simplices in T where the simplices in K are identified with the simplices in K' under f.

The next lemma provides a condition for determining when the simplicial adjunction space is a flag complex. The following lemma after that provides a stronger result which gives a way of determining when the simplical adjunction space is a full subcomplex of some ambient simplicial complex.

Lemma 2.9. *Let S and T be simplical complexes with $K \leq S, T$ where K is a full subcomplex of both S and T and $S \cap T = K$. Then the simplical complex $S \cup_K T$ obtained by attaching S and T along K is a flag complex if and only if each of S and T are flag complexes.*

Proof. First, suppose that $S \cup_K T$ is flag and let $v_1, \ldots, v_{n+1}$ be a collection of vertices in S that pairwise span 1-simplices in S. We seek to show that the simplex σ spanned by $v_1, \ldots, v_{n+1}$ is contained in S. If $\sigma \notin S$, then $\sigma \notin S \cup_K T$, as $S \cup_K T$ is simply the union of simplices in S with simplices in T. This contradicts the fact that $S \cup_K T$ is flag. Hence $\sigma \in S$ and we conclude that S is flag. The same argument will show that T is flag as well.

Now suppose that each of S and T are flag, and let $v_1, \dots, v_{n+1}$ be a collection of vertices in $S \cup_K T$ that pairwise span 1-simplices in $S \cup_K T$. If each v_i is contained in S, then the simplex spanned by $v_1, \dots, v_{n+1}$ is contained in S and thus $S \cup_K T$. The desired result follows similarly if each v_i is contained in T.

Now, suppose $v_1, \dots, v_{n+1}$ are not all contained in S or in T. Let $v_i \in S$, $v_i \notin T$, and $v_j \in T$, $v_j \notin S$, and let e be the edge they span. Since K is full in both S and T, e must belong to either S or T, implying that v_i and v_j either both belong to S or both belong to T. Thus $v_1, \dots, v_n$ either all belong to S or to T. $\square$

The next lemma may seem oddly specific, but it will be a key ingredient later when proving convexity of certain subcomplexes.

Lemma 2.10. *Let $K, S \leq X$ with $K \leq S$ and such that*

$$Y = S \bigcup_K K \star p$$

is a subcomplex of X with the cone point p not adjacent to any vertices in S except those in K. If K and S are full subcomplexes of X, then Y is full as well.

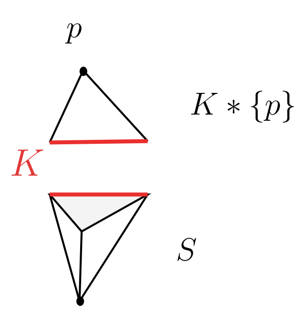

Figure 5

proof. Since S and K are full, they are flag. Furthermore, Lemma 2.6 implies $K \star p$ is flag. Lemma 2.9 then implies that Y is flag.

Suppose Y is not full and let n be the smallest positive integer for which there is a set of vertices $v_1, \dots, v_n$ in Y that span an $n-1$ simplex σ in X, but σ is not contained in Y. One of $v_1, \dots, v_n$ must be p. Suppose $p = v_i$ for some i. It follows then that the collection $v_1, \dots, v_n$ is contained in $K \star p$ (since p is assumed to be adjacent no vertices of S except for those in K). By fullness, σ is in $K \star p$. $\square$

As a special case of Lemma 2.10, if $S = K$, then $Y = K * \{p\}$ and we are in the case of Lemma 2.7.

We point out that, in general, a full subcomplex attached to a full subcomplex along a common full subcomplex need not be full. Consider two 2-simplices attached along a common edge. As a subcomplex of a 3-simplex, this is not full.

2.2 Cube Complexes

By a cube we mean a finite product $[-1, 1]^n$, which we may also refer to as an n-cube if we want to call attention to its dimension. 0-cubes will also be called vertices and 1-cubes will also be called edges. When it is useful to think of cubes as geometric objects, we will endow them with one of the following metrics.

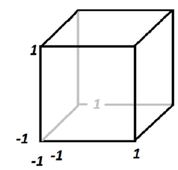

Figure 6: A 3-cube

1. (Taxi-Cab Metric) the n-cube is embedded in $\mathbb{R}^n$ and given the subspace metric induced by the l_1 norm.

2. (Euclidean Metric) again the n-cube is embedded in $\mathbb{R}^n$ and given the subspace metric induced by the Euclidean norm.

By default, open sets in a cube are defined using the Eucliedean metric.

For most of what follows, it will be more relevant to think of cubes as combinatorial objects. To this end, we set a framework similar to that of simplicial complexes.

Definition 2.11 (Faces, Facets)**.** A **face** of a cube is obtained by fixing some number, possibly zero, of its coordinates with a 1 or a -1. By forgetting the coordinates fixed with ± 1, one sees that a face is itself a cube.

A face of an n-cube that is of dimension strictly less than n is said to be a **proper face**. A proper face that is exactly one dimension less than the cube to which it belongs is called a **facet** of that cube. The vertices of a cube are then obtained by restricting all of the coordinates to 1 or to -1.

Note that a cube $c = [-1, 1]^n$ is a face of itself.

As a quick aside, we enumerate the number of k-dimensional faces of an n-cube ($k \leq n$), which can themselves be realized as k-cubes. This is accomplished by counting the number of ways of fixing some subset of the n coordinates with 1's or -1's. Thus, for an n-cube, one finds that there are $2^{n-k}\binom{n}{k}$ k-dimensional faces. For example, the faces of a 3-cube consist of eight 0-cubes, twelve 1-cubes, six 2-cubes, and one 3-cube.

For a cube c, we say its vertices **span** c or are a **spanning set** for c.

Definition 2.12 (Cube Complex)**.** Let $\mathscr{C}$ denote a disjoint collection of cubes of varying dimensions and all their faces, and let X denote the quotient space with quotient map $q : \mathscr{C} \to X$, where q identifies various faces of cubes in $\mathscr{C}$ using euclidean isometries satisfying to the two following conditions:

1. for a cube c in $\mathscr{C}$, none of its faces are identified by q; and

2. for two distinct cubes c, d, the intersection $q(c) \cap q(d)$ is either a common face of each or the empty set

The space X is called a **cube complex**.

The topology on a cube complex X is assumed to be the quotient topology, unless otherwise stated. As such, a subset $U \subseteq X$ is open if $q^{-1}(U \cap q(c))$ is open in c for all $c \in \mathscr{C}$.

We will often abuse notation and identify $q(c)$ as simply c. For c, a cube in X, all the faces of c will also be cubes in X. We write $c \in X$ to mean that c is a cube in X. We call a cube **maximal** if it is not a proper face of any cube in X.

Condition 1 excludes the familiar construction of a torus whereby one identifies opposite edges of a 2-cube. However, upon subdividing the 2-cube and reparametrizing, the torus an be given a cube complex structure. Condition 2 is self-explanatory.

The following definition is a well-known equivalence relation on cube complexes.

Definition 2.13 (Cubical Isomorphism)**.** We say two cube complexes X and Y are **isomorphic** if there exists a bijection $X^0 \to Y^0$ so that if some collection of vertices span a cube in X then the images of those vertices span a cube in Y. We write $X \cong Y$ to mean that X and Y are isomorphic.

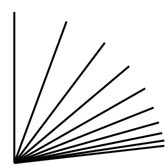

Figure 7: A cube complex that is not locally-finite

Let X be a cube complex. We say the dimension of X is n if X has cubes of dimension n but does not have cubes of dimension greater than n. If no such n exists, we say X is infinite-dimensional. Note that cube complexes are not necessarily finite-dimensional. A cube complex is **locally-finite** if every vertex meets only finitely many cubes. It is worth noting that there exist finite dimensional cube complexes that are not locally-finite. There also exist locally-finite cube complexes that are not finite-dimensional. In Figure 7, one is to imagine a space obtained by attaching countably many 1-cubes to a vertex. The resulting cube complex is 1-dimensional but not locally-finite. This is called a **Hedgehog space**. At the other extreme, consider the following example.

Example 2.14 Begin with a vertex and call it v_0. Attach an edge to v_0 and call it e. At the vertex of e that is not v_0, attach a square. Then, continue this process inductively. At the nth step attaching an $n+1$-cube to a facet of the n-cube that was attached in the previous step where this facet is disjoint from every cube that was added in the steps preceding the previous one. See the figure on the right for an example of such a space. This space is locally-finite, but not finite dimensional.

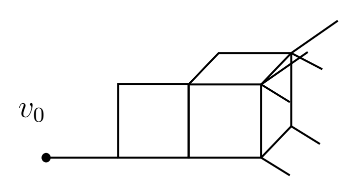

Figure 8: A cube complex that is locally-finite but not finite dimensional

One could also define cube complexes more analogously to simplicial complexes, although describing cube complexes this way is not as elegant. We first describe an n-cube combinatorially, and from there, we generalize to cube complexes.

Let c be an n-cube. Then c has 2^n vertices, and the faces of c have 2^k vertices for various $m \leq n$. Let $V = \{v_1, \ldots, v_{2^n}\}$ be the vertex set of c. The combinatorial description of c consists of a collection of subsets of $\mathcal{P}(V)$ corresponding to the vertex sets of faces of c. Note that while the faces of c are spanned by 2^k element subsets of V, not every subset of cardinality 2^k is included in the combinatorial description of c ($2^{n-k}\binom{n}{k}$ are included). Clearly then, this collection is not closed under taking subsets (as was the case with simplicial complexes).

A cube complex is then the union of its cubes described combinatorially. Two cubes, d and e are then connected along a common face f whenever the vertex set of f and all the subsets corresponding to the faces of f appear in both d and e. Without going into more detail, we will assume that the reader is comfortable with regarding cube complexes in this way, and we expect that if they were given a cube with vertices $v_1, \ldots, v_{2^n}$, they would be able to write the subsets of $\mathcal{P}(\{v_1, \ldots, v_{2^n}\})$ corresponding to its faces. For instance, a 1-cube with vertices labeled v_1 and v_2 could be described as $\{\{v_1\}, \{v_2\}, \{v_1, v_2\}\}$.

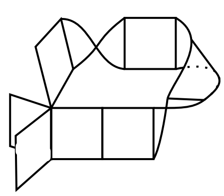

Figure 9: A square complex

A certain class of cube complexes have proven to provide accessible examples exhibiting a wide variety of interesting behavior. These are *square complexes*. Let X be a cube complex such that every maximal cube in X is a 2-cube. Such a space X is called a **square complex**. The figure on the left shows an example.

Like simplicial complexes, cube complexes enjoy a stratification based on the dimension of certain subcomplexes called skeleta. We let X^k denote the union of cubes in a cube complex X of dimension less than or equal to k, and call this the **k-skeleton** of X. We will often make reference to X^0, the 0-skeleton of a cube complex, which is also called the vertex set of X.

Throughout this paper, we will predominantly be working with *pointed* cube complexes. We will always choose the base point to be among the vertex set of X, and we let (X, v_0) denote the pointed cube complex with preferred basepoint v_0.

2.2.1 Carrier, Core, Cubical Interior, and Full Subcomplexes

Here, we identify various subsets and subcomplexes that will play recurring roles throughout this exposition.

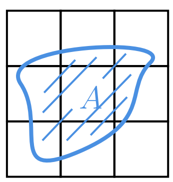

$core(A)$ = the center square

$\mathcal{C}(A)$ = the entire complex

Figure 10

Definition 2.15 (Subcomplex, Carrier, Core)**.** Let $\mathscr{C}' \subseteq \mathscr{C}$ such that if $c \in \mathscr{C}'$ then all the faces of c are also contained in $\mathscr{C}'$. We call the cube complex L, with quotient map $q' : \mathscr{C}' \to L$, a **subcomplex** of X if the map q' is the restriction of q to $\mathscr{C}'$. To identify L as a subcomplex of X, we will use the notation $L \leq X$.

Let $A \subset X$ be a subset. The **core** of A, denoted $core(A)$ is the largest subcomplex of X contained in A. The smallest subcomplex that contains A is called the **carrier** of A and is denoted $\mathcal{C}(A)$.

Lemma 2.16. *Let A, B be subsets of a cube complex X. Then $core(A \cap B) = core(A) \cap core(B)$.*

Proof. Let c be a cube in $core(A \cap B)$. Then $c \subseteq A$ and $c \subseteq B$. Thus $c \in core(A) \cap core(B)$.

Conversely, let c be a cube in $core(A) \cap core(B)$. Then $c \subseteq A \cap B$, and the desired result follows. $\square$

Definition 2.17 (Cubical Interior)**.** For a cube $c = [-1, 1]^n$, the **cubical interior** of c, denoted $\mathring{c}$, is the product $(-1, 1)^n$ if $n \geq 1$, and $\mathring{c} = \{1\}$ if $n = 0$.

By contrast, for an arbitrary subset $S \subseteq X$, we will let $int(S)$ denote its topological interior. Here we mean X with the quotient topology induced by the map q. Note that if c is a cube in X which is a proper face of some other cube in X, then $int(c) = \emptyset$, while $\mathring{c}$ is certainly not empty. At the other extreme, if X is a single n-cube c together with its faces, then $int(c) = X$.

Example 2.18. Let c be a maximal cube in X. Then $X - \mathring{c} \leq X$ is a subcomplex consisting of all the cubes of X except the cube c. Note however that the proper faces of c still belong to $X - \mathring{c}$. For example, consider the case where X equals a 2-cube c. Written combinatorially,

$$c = \{\{v_1\}, \{v_2\}, \{v_3\}, \{v_4\}, \{v_1, v_2\}, \{v_2, v_3\}, \{v_3, v_4\}, \{v_4, v_1\}, \{v_1, v_2, v_3, v_4\}\}$$

Then $X - \mathring{c}$ is the subcomplex

$$X - \mathring{c} = \{\{v_1\}, \{v_2\}, \{v_3\}, \{v_4\}, \{v_1, v_2\}, \{v_2, v_3\}, \{v_3, v_4\}, \{v_4, v_1\}\}$$

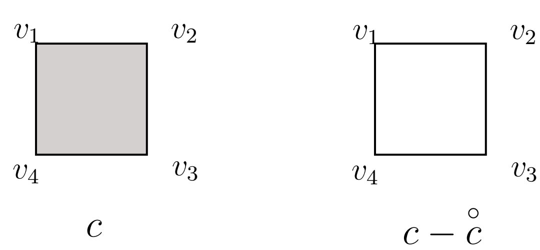

In general, one can also describe $c - \mathring{c}$ combinatorially. Let $v_1, \ldots, v_{2^n}$ be the vertices that span c. Then $c - \mathring{c}$ can be realized as the subcomplex of c obtained by removing $\{v_1, \ldots, v_{2^n}\}$ but retaining all other subsets corresponding to proper faces of c.

We say a subcomplex K of a cube complex X is **full** if, for any collection of vertices from K the set of cubes spanned by this collection in X are also cubes in K. The previous example yielded a subcomplex that is not full. This is because the subcomplex $X - \overset{\circ}{c}$ contains the vertex set of the cube c, but does not contain c itself. Note that this is essentially the same definition of full we gave for simplicial complexes. One need only change cubes to simplices.

Lemma 2.19. *Let X be a cube complex. Every set of vertices in X determines a unique full subcomplex.*

Proof. Let $V = \{v_1, v_2, \ldots\}$ be a set of vertices in X, and let K be the minimal full subcomplex containing this V. Let L denote the full subcomplex obtained by including every cube from X spanned by vertices in V. Clearly, $K \leq L$. Thus, we need only show that every cube in L is also in K. Let c be a cube in L spanned by vertices $\{u_1, \ldots, u_{2^n}\}$. As K contains this vertex set and is a full subcomplex, c is in L. $\square$

2.3 Geometry and Combinatorics of Cube Complexes

In this section, we define the metrics we will use on cube complexes. Following a brief discussion of the metric structure of cube complexes, we will turn our attention to more combinatorial methods to investigate the geometry of cube complexes. We then introduce the various concepts of convexity that will show up throughout the remainder. We finish the section by introducing the reader to some useful examples.

2.3.1 Curvature and Gromov's Link Condition

Before we begin working in the setting of CAT(0) cube complexes, we provide a brief discussion on non-positive curvature from a classical viewpoint. A geodesic metric space X is *non-positively curved*, if triangles in X are no thicker than their comparison triangles in Euclidean space. Given a geodesic metric space X, take three points $A, B, C \in X$ and form the geodesic triangle T with these points as vertices. Let the sides opposite A, B, and C have lengths a, b, and c, respectively. The comparison triangle is a triangle T' in the Euclidean plane with side lengths a, b, and c. To say that T is no thicker than T' means that between any two points on distinct sides in T, the cord joining them has length less than or equal to the cord joining corresponding points in T'. A non-positively curved space X is CAT(0) if it is simply-connected.

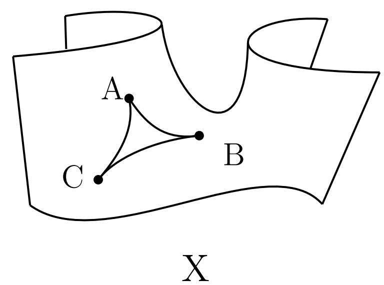

Figure 11

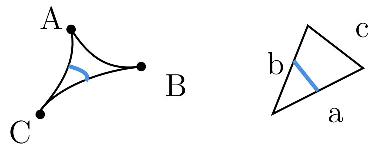

Figure 12

We will seldom make use of triangles in cube complexes, preferring to rely on more combinatorial methods that exist for CAT(0) cube complexes. One reason for this preference is that angles in CAT(0) spaces can be hard to work with in general. With that said, angles and triangles will ocasionally be mentioned. However, the angles we consider will usually arise as intersections of edges and hyperplanes. In all these cases, if the intersections are nontrivial then the resulting angles will all be right angles ($\pi/2$).

Now, let X be a cube complex. The distance between two points x and y in X is realized by one the of the following (path-length) metrics on X.

1. (Taxi-Cab Metric) With each cube endowed with the taxi-cab metric, we define $\rho_1(x, y)$ to be the infimum length of piecewise geodesic paths between x and y.

2. (Piecewise Euclidean Metric) We define $\rho_2(x, y)$ as the infimum length of piecewise geodesic paths between x and y, where each cube is endowed with the standard metric.

As a special case of (1), for $x, y \in X^0$, some authors refer to $\rho_1(x, y)$ as the **combinatorial distance** between x and y. In this case, the minimum length path between vertices x and y can be realized by a path that lives in the 1-skeleton of X. Such a path can be expressed as a sequence of vertices $x = v_1,\ v_2,\ \ldots,\ v_n = y$ such that v_i is on γ, and v_i and v_{i+1} are opposing endpoints of a 1-cube in X. A sequence of vertices realizing the combinatorial distance between two vertices is called a **combinatorial geodesic** and is denoted by $[v_1, \ldots, v_n]$. The following definition generalizes combinatorial geodesics to paths that can begin and end in the 1-skeleton.

Cube complexes are not necessarily connected. When points x and y belong to distinct components of X, we say the distance between them is infinite. Thus, the path-length metrics that we consider on X allow infinite values.

Definition 2.20. By a **combinatorial edge path** in X, we mean a sequence of edges $e_1, \ldots, e_n$ such that e_i and e_{i+1} are adjacent. A combinatorial edge path joins vertices $v_1, \ldots, v_{n+1}$ where e_i is spanned by v_i and v_{i+1}. We define an **edge path** as a combinatorial edge path in which one is allowed to append a portion of an edge to v_0 and to v_{n+1} such that the edges these portions belong to do not appear among $e_1, \ldots, e_n$. When an edge path realizes the distance (in the ρ_1 metric) between two points x and y in X^1, we call it a **geodesic edge path** between x and y.

One can think of a geodesic edge path between vertices as the *geometric realization* of the combinatorial geodesic between them.

The next definition describes the role simplicial complexes will play in the story of cube complexes.

Definition 2.21 (Link)**.** Let X be a cube complex. For a vertex $v \in X^0$, the **link** of v in X, denoted $lk_X(v)$ (or simply $lk(v)$ when X is understood), is a simplical complex obtained in the following manner: Each 1-cube that contains v as a face contributes a vertex to $lk(v)$. A collection of vertices in $lk(v)$ span a simplex if their corresponding 1-cubes belong to a common cube in X.

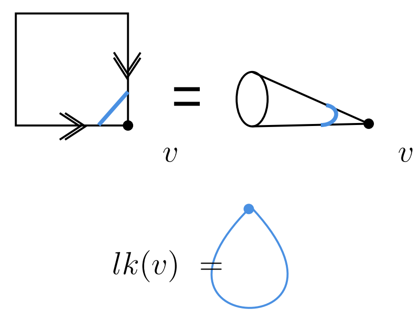

Figure 13: Not a cube complex; Link of v is not simplicial

With links of vertices, we can explain the need for condition (1) in Definition 2.12. If we were to allow identifying faces of the same cube, the link could fail to be a simplicial complex. See Figure 13. Also excluded from the class of simplicial complexes are bigons, two vertices attached by distinct edges.

Definition 2.22 (CAT(0) Cube Complex)**.** A cube complex is **non-positively curved** if it is connected and the link of every vertex is a flag complex. A non-positively curved cube complex is called a **CAT(0) cube complex** if it is simply connected.

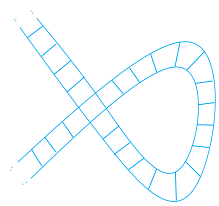

Figure 14: A non-positively curved cube complex that is not CAT(0)

In the image below, (a) shows a 2-dimensional cube complex where the link of each vertex is 1-dimensional, and (b) shows a 3-dimensional complex in which each the link of each vertex is 2-dimensional. Only the complex in (b) is non-positively curved. In fact, it is CAT(0).

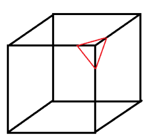

(a) Link is not flag

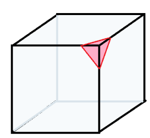

(b) Link is flag

Figure 15

It is a theorem of Gromov that the notion of non-positive curvature for cube complexes given combinatorially in terms of links agrees with the classical notion of non-positive curvature [Gro87]. This theorem was extended to infinite-dimensional cube complexes by Leary in [Lea13]. As an important consequence, when a cube complex is CAT(0) it is a unique geodesic space (with the piecewise Euclidean metric). As a consequence of this (and from the fact that CAT(0) $\implies$ simply-connected), it follows that CAT(0) cube complexes are contractible. For a CAT(0) cube complex, the piecewise Euclidean metric is often referred to as the CAT(0) metric.

Observe that, even if X is CAT(0), geodesic edge paths are not necessariy unique. In a 2-cube, for instance, there are two geodesic edge paths between diametrically opposed vertices. To further illustrate this issue, the figure on the right shows a CAT(0) cube complex with points p and q. The unique CAT(0) geodesic (in the CAT(0) metric) between p and q is highlighted in blue. One of the geodesic edge paths is highlighted in red. There are six more!

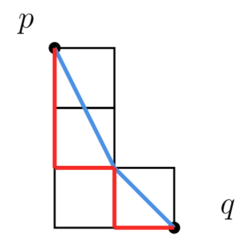

Figure 16

2.3.2 Convexity in Cube Complexes

For a connected, simply-conneected geodesic space, a subset A is called *convex* if every geodesic segment with endpoints in A is contained entirely in A. We will refer to subsets that satisfy this definition of convexity as **metrically convex**. As stated above, when endowed with the piecewise Euclidean metric, a CAT(0) cube complex has unique geodesics, and when we mention metrically convex subsets, we mean subsets that are convex with respect to the piecewise Euclidean metric.

For subcomplexes of a CAT(0) cube complex, a more combinatorial description of convexity is available.

Definition 2.23 (Combinatorially Convex)**.** Let X be a CAT(0) cube complex. We say a subcomplex K is **combinatorially convex** if K is connected and $lk_K(v)$ is a full subcomplex of $lk_X(v)$ for every vertex $v \in K$.

In the appendices of [Lea13] it is proved that if X is a CAT(0) cube complex, then a connected subcomplex of X is metrically convex if and only if it is combinatorially convex. We shall henceforth refer to a subcomplex in a CAT(0) cube complex that satisfies either (and hence both) of the definitions of convexity as simply **convex**.

It is not difficult to see that if K is a convex subcomplex of a CAT(0) cube complex X, then K is full. Consider a collection of vertices $v_1, \ldots, v_{2^n}$ in K that span an n-cube in X. By convexity, the diagonals of this cube belong to K. Clearly then, the entire cube belongs to K. It is also clear from the classical definition that triangles in K are no thicker than their comparison triangles in Euclidean space. Furthermore, as subcomplexes are closed and thus complete in X, Proposition II.2.4 in [BH99] implies that K is a deformation retract of X. In particular then, since X is simply-connected, K is simply-connected. It follows that a convex subcomplex is CAT(0). However, it is not always true that a full subcomplex is convex (consider two adjacent squares in a 3-cube).

We also allow the possibility that K can be a CAT(0) subcomplex and not necessarily convex in X. This simply means that the links of vertices in K are all flag but not all are necessarily full subcomplexes of their respective links in X.

The following remark follows from the discussion section 2.2 of [Hag07] and will be of great use to us as we will primarily work with geodesic edge paths rather than geodesic paths.

Remark. Let X be a CAT(0) cube complex and $K \leq X$ be a convex subcomplex. Then, for any pair of points in K^1, every geodesic edge path between them is contained in K.

According to Remark 2.3.2, a convex subcomplex K of a CAT(0) cube complex X is convex in the traditional sense with respect to either the ρ_1 or ρ_2 metrics. However, when discussing convexity of subsets of X we strictly mean convex with respect to the piecewise-Euclidean metric.

We will return to a discussion of convexity later in this chapter. For now, we finish by proving transitivity of convexity for subcomplexes of a CAT(0) cube complex. While this may seem obvious from the perspective of metric convexity, we felt the need to supply it as we predominantly consider convexity from a combinatorial perspective. The following result follows from Lemma 1.1 proved in section 1.1

Corollary 2.24. *Let X be a CAT(0) cube complex and let $K, L \leq X$ with L a convex subcomplex of X and K a convex subcomplex of L. Then K is also convex in X.*

Proof. Let v be a vertex in K. By convexity, $lk_K(v)$ is full in $lk_L(v)$ and $lk_L(v)$ is full in $lk_X(v)$. Lemma 2.5 then implies that $lk_K(v)$ is full in $lk_X(v)$. □

2.3.3 Some Simple Examples

The reader is encouraged to keep the following examples in mind throughout the remainder of this exposition. The majority of the following examples are non-compact.

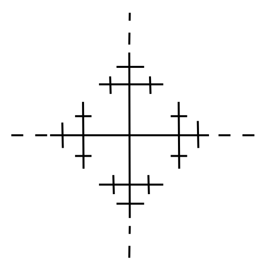

Figure 17

Example 2.25 (Trees)**.** Perhaps the most prototypical examples of CAT(0) cube complexes are trees. In fact, CAT(0) cube complexes are really a generalization of trees.

Given a tree, T, one realizes a cubical structure on T by parametrizing each edge as the interval $[-1, 1]$. We call this the standard cubical structure on T.

Figure 17 shows a tree where each vertex meets four 1-cubes. The link of any vertex is a discrete set with four points.

Example 2.26 (Euclidean n-Space)**.** The Euclidean line, $\mathbb{E}$ is a tree and follows the same prescription as the previous example. $\mathbb{E}^2$ has a standard tessellation by squares. We of course choose squares to be of the form $[-1, 1]^2$. This results in a square complex, where each vertex meets four squares in such a way that its link is a four-sided polygon. See Figure 18. More generally, $\mathbb{E}^n$ has a similar tesselation by n-cubes that sees 2^n n-cubes attached together at each vertex so that all vertex links are homeomorphic to $n-1$-spheres. This is the stand cubical structure on $\mathbb{E}^n$.

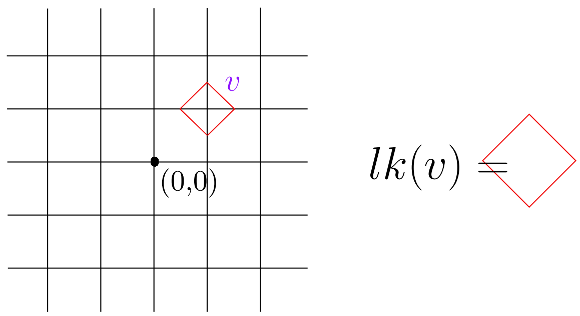

Figure 18

Of course, $\mathbb{E}^n$ already has a product structure. In particular, the standard cubical structure on $\mathbb{E}^2$ can be realized by giving $\mathbb{E}$ its standard cubical structure, and then using the product structure of $\mathbb{E} \times \mathbb{E}$ to define cubes in $\mathbb{E}^2$. One can then define $\mathbb{E}^n$ inductively as $\mathbb{E}^{n-1} \times \mathbb{E}$. The precise manner in which one should form products of cube complexes will be described in the section on products and adjunction spaces.

Example 2.27 (A More Hyperbolic Cube Complex)**.** We call the following construction the 5-plane. Begin with the order-4 pentagonal tiling of the hyperbolic plane, whereby each vertex meets four pentagons. Place a vertex at the centerpoint of each pentagon and at the centerpoint of each edge. These new vertices together with the original vertices give the vertex set of a square complex where some vertices meet four squares (the vertices on the edges, including the original vertices), and some vertices meet five squares (the centerpoints of pentagons). See Figure 19.

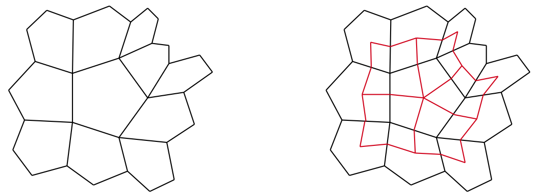

Figure 19

Reparametrize so that the four squares that meet at the valent-4 vertices give a single 2-cube. Performing this reparametrization at all such vertices yields a square complex where each vertex meets five squares. The link of each vertex is a pentagon.

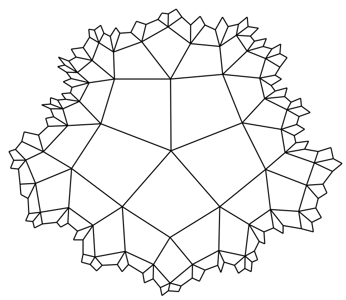

Figure 20

Although locally points live in cubes with euclicean geometry, the 5-plane with the path-length metric is quasi-isometric to the hyperbolic plane.

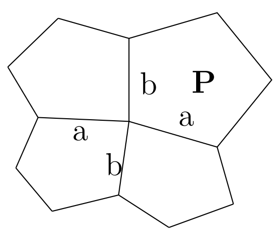

Figure 21: The polygon P is reflected about its edges

Alternatively, one could construct the 5-plane by beginning with a pentagon P and reflecting P about its edges in such a way that if a and b are two consecutive edges, then the result of reflecting about a then b is the same as reflecting about b then a. Performing the same identifications as before will yield the 5-plane.

Notice that, in this alternative construction, we began a pentagon P, the boundary of which is a simplicial complex consisting of five edges. It is no coincidence that the links of vertices in the 5-plane are isomorphic to the boundary of P. This exemplifies a more general construction known as the Davis complex, where one begins with a simplicial graph and then attaches simplices to every pairwise connected collection of vertices. The resulting flag complex is called a nerve. One then creates a CAT(0) cube complex where links of vertices are isomorphic copies of the nerve. Background on and applications of the Davis complex can be found in [Dav08].

2.3.4 Hyperplanes and Halfspaces

It is well-known that CAT(0) cube complexes possess a rich combinatorial structure which one can access by considering a collection of subsets called *hyperplanes*. The collection of hyperplanes in a CAT(0) cube complex will facilitate many of the forthcoming arguments in this paper. We begin this subsection by explaining how to build hyperplanes out of *midcubes*. We then recall well-known established properties of hyperplanes in CAT(0) cube complexes. We finish by using hyperplanes to define local coordinates and state a few other useful facts.

For a cube complex X, let c be an n-cube in X. A **midcube** of c is a subset, m, obtained by restricting precisely one coordinate of c to 0. So, for $c = J^n$, a midcube of c has the form $m = \{0\} \times J^{n-1}$. One immediately sees that, for $n \geq 1$, an n-cube will have n distinct midcubes. A 0-cube has no midcubes. By "forgetting" the coordinate that is fixed with a zero, one also sees that a midcube of an n-cube can be viewed as an $(n-1)$-cube in its own right. We emphasize however that a midcube is never a subcomplex of X.

The **midpoint** of an n-cube c is the point $(0, 0, \ldots, 0)$. This definition extends to midpoints of faces of cubes in the natural way.

The following is a way of obtaining a new cube complex from an existing one by subdividing cubes into "smaller" ones. The new subcomplex will be useful when we want to endow certain subsets with a cubical structure.

Definition 2.28 (Barycentric Cubing)**.** Given an n-cube, c, we will inductively describe a process of building the 1-skeleton of a cube complex. Subdivide each edge of c into two segments each containing the midpoint as one of its endpoints. Then, inductively attach a segment between the midpoint of each face f of c and the midpoint of each facet of f. Reparametrizing each segment to be of the form $J = [-1, 1]$ gives the 1-skeleton of a cube complex, and filling in each cube in the obvious way produces a cube complex c' called the **barycentric cubing** of the cube c.

One can now extend the definition to an entire cube complex X by taking the barycentric cubing of each cube in X. We denote the barycentric cubing of X as X'.

Note that the barycentric cubing of an n-cube, c, will be a cube complex with 2^n n-cubes. Each such n-cube will meet at the midpoint of c.

Now, let c and d be distinct cubes in X whose intersection is nonempty. Then their intersection will be a common face, say $f = c \cap d$. If f is a vertex, then the midcubes of c and the midcubes of d do not intersect. If f is an edge, then c and d will contain 2-cubes as faces that intersect along f. Each of these two cubes will contain midcubes that intersect in the midcube (midpoint) of f. More generally, for cubes c and d with $c \cap d = f$, there will be a number of midcubes of c and of d that intersect in f. Some of these midcubes will intersect in midcubes of f and some not. The following definition considers those midcubes of intersecting cubes whose intersection is a midcube of the common face.

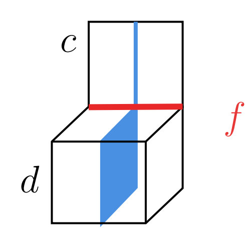

Figure 22

Definition 2.29 (Hyperplane). A **hyperplane** in X is constructed by beginning with a midcube m. Then, for each midcube n such that $m \cap n$ is a common face of both m and n, we include it in the hyperplane. We then repeat this process on each midcube as it is included in the hyperplane until we have exhausted all possible midcubes to add in this manner. We will express hyperplanes using lowercase letters with hats, such as $\hat{h}$. The collection of all hyperplanes will be denoted by $\hat{\mathcal{H}}_X$ (or simply $\hat{\mathcal{H}}$).

An example of a hyperplane in a CAT(0) cube complex is shown below. Carriers of hyper-

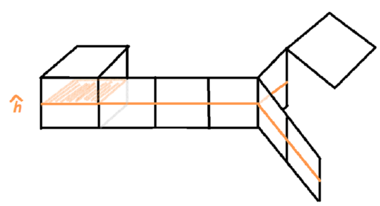

Figure 23: A hyperplane $\hat{h}$

planes will play an important role in many of the forthcoming arguments. They will prove especially useful in arguing that locally-finite CAT(0) cube complexes are collapsible. In the figure beneath, the carrier of $\hat{h}$ is highlighted.

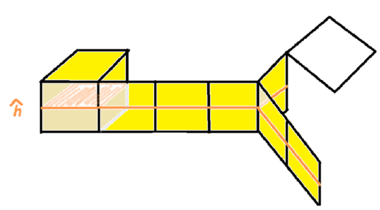

Figure 24: A hyperplane carrier $\mathcal{C}(\hat{h})$

In the case that X is a CAT(0) cube complex, hyperplanes in X enjoy a number of useful properties, which we list here.

Properties of Hyperplanes. Let X be a CAT(0) cube complex and let $\hat{h}$ be a hyperplane in X. Then

1. $\hat{h}$ is a convex subset and $\mathcal{C}(\hat{h})$ is a convex subcomplex of X.

2. $X \setminus \hat{h}$ consists of two components, each of which is a convex subset.

3. For a vertex $v \in X$, $\rho_1(v, \hat{h})$ equals 1 plus twice the number of hyperplanes separating v and $\hat{h}$.

4. If two hyperplanes $\hat{h}$ and $\hat{k}$ intersect, then for a point $x \in \hat{h}$, $\rho_2(x, \hat{k})$ is realized by a geodesic that lives in $\hat{h}$.

5. **Helly's Property:** If a collection of hyperplanes pairwise intersect, then the total intersection of the collection is nonempty

6. The carrier of $\hat{h}$ enjoys a product structure: $\mathcal{C}(\hat{h}) \cong \hat{h} \times [-1, 1]$, where $\hat{h}$ is identified with $\hat{h} \times \{0\}$.

7. For each edge e, there exists a unique hyperplane $\hat{h}$ such that $e \cap \hat{h}$ is a midpoint of e.

The closure of either of the two components of $X \setminus \hat{h}$ is called a **halfspace**. It follows from the definition that hyperplanes are disjoint from X^0. Thus, for a pointed CAT(0) cube complex (X, v_0), the two halfspaces corresponding to $\hat{h}$ are distinguished by either containing v_0 or not. We will denote the halfspace that contains v_0 by h, while the halfspace that does not contain v_0 will be denoted by h^*.

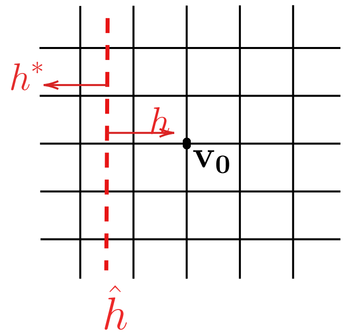

Figure 25: Example of halfspaces h and h^*

Proofs of properties (1) and (2) can be found in [Far16] Property (3) is a restatement of Lemma 2.17 in [CFI16]. Property (4) follows from (3) and Lemma 2.18 of [CFI16].

A restatement and proof of Helly's property (5) are provided by Lemma 4.14 in [Sag95]. Note that, since hyperplanes are convex, if a collection of hyperplanes intersect nontrivially, then the intersection is convex. In particular the intersection is connected.

Proofs of properties (6) and (7) can also be found in [Sag95].

The following example shows that Helly's property can fail if the flag condition on vertex links is relaxed.

Example 2.30. Let X be the cube complex that has vertices of alternating valences three and six, where every valent-3 vertex link is a triangle and every valent-6 vertex link is a hexegon. Links of 3-valent vertices are not lag. Thus, this space is not nonpositively curved and, in this space, there will exist triples of pairwise intersecting hyperplanes that do not intersect in total. The image on the left shows three such hyperplanes.

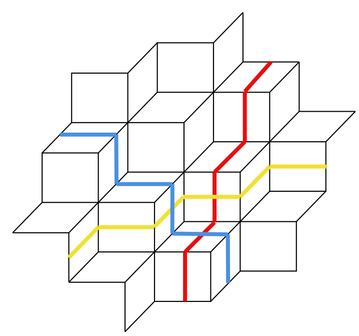

Figure 26: A cube complex where Helly's property is violated

Lemma 2.31. *Let (X, v_0) be a pointed locally-finite CAT(0) cube complex, and let v be a vertex in X. Then there are only finitely many hyperplanes separating v_0 and v and, every geodesic edge paths connecting v_0 with v will cross each of the hyperplanes separating them.*

Proof. That only finitely many hyperplanes separate v_0 and v follows from property (3) of hyperplanes.

Now, let $\hat{h}_{k_1}, \ldots \hat{h}_{k_n}$ be the hyperplanes separating v from v_0. As hyperplanes separate the space into two convex (path) components, any geodesic between v and v_0 will have to cross each of $\hat{h}_{k_1}, \ldots \hat{h}_{k_n}$ exactly once. By (6), it is easy to see that any edge path that crosses a hyperplane more than once can be shortened. So a geodesic edge path intersects a hyperplane once or not at all.

Let γ be a geodesic edge path between v_0 and v. Suppose γ crosses a hyperplane $\hat{h}$ that is not among $\hat{h}_{k_1}, \ldots, \hat{h}_{k_n}$. Since $\hat{h}$ does not separate v_0 and v, both v and v_0 belong to h. In particular then, they belong to $core(h)$. Since $core(h)$ is a convex subcomplex, γ must live entirely in $core(h)$. But then γ cannot cross $\hat{h}$. □

Let X be a CAT(0) cube complex with hyperplanes $\hat{\mathcal{H}}$. For locally-finite cube complexes consisting of countably many components, the set $\hat{\mathcal{H}}$ will be at most countably-infinite. In the case of a connected cube complex, it will be convenient to enumerate the hyperplanes as follows: we let $\hat{\mathcal{H}} = \{\hat{h}_1, \hat{h}_2, \ldots\}$ so that the map $i \mapsto \rho_1(v_0, \hat{h}_i)$ is monotone non-decreasing. We call an enumeration on $\hat{\mathcal{H}}$ that satisfies this property **distance-respecting**. Note that, in order to achieve a distance-respecting enumeration, X must be locally-finite. Otherwise, for some fixed distance d, there will be infinitely many hyperplanes a distance d from v_0 and we would be unable enumerate the hyperplanes whose distance from v_0 is greater than d. X must also be connected for the same reason.

Remark. Given a hyperplane $\hat{h}$ in X, we can use the barycentric cubing construction (Definition 2.28) to endow it with a cubical structure. Observe that, for each edge that $\hat{h}$ meets, it does so in its midpoint. Even more, if $\hat{h}$ intersects a cube, then $\hat{h}$ will contain the midpoint of that cube. We thus define the vertex set of $\hat{h}$ as the set of vertices in X' that are midpoints of cubes (in X) intersecting $\hat{h}$. Then $\hat{h}$ is naturally identified with the full subcomplex of X' corresponding to this vertex set.

It is unfortunate that, in general, geodesic edge paths between points in a CAT(0) cube complex are not unique. However, hyperplanes with a distance-respecting enumeration provide a way of choosing canonical geodesic edge paths when one endpoint is the base vertex.

Let X be a CAT(0) cube complex and let e be an edge in X. By property (7) of hyperplanes, e intersects a unique hyperplane, $\hat{h}$, of X. Suppose e belongs to a geodesic edge path, γ. Then, in traversing e, γ will intersect $\hat{h}$. It is therefore the case that each edge on a γ determines a unique hyperplane. Furthermore, the collection of all edges determines the entire collection of hyperplanes crossed by γ.

Let (X, v_0) be a pointed locally-finite CAT(0) cube complex and give $\hat{\mathcal{H}} = \{\hat{h}_i\}_{i=1}^{\infty}$ a distance-respecting enumeration. Let $[v_0, \ldots, v_n]$ be a combinatorial geodesic between vertices v_0 and v_n in X, and let γ be the corresponding geodesic edge path. Then γ is a sequence of edges $e_1, \ldots, e_n$ such that v_0 is an endpoint of e_1, v_n is an end point of e_n, and e_i shares an endpoint with e_{i+1} for $i = 1, \ldots, n-1$. Let $\hat{h}_{k_i}$ denote the hyperplane that intersects e_i. Then $\hat{h}_{k_1}, \ldots, \hat{h}_{k_n}$ gives a *sequence* of hyperplanes in X that γ crosses in the order written when travelling along γ from v_0 to v_n.

Since γ corresponds to a unique sequence of edges, it follows that the sequence $\hat{h}_{k_1}, \ldots, \hat{h}_{k_n}$ uniquely determines γ, so long as v_0 is undertood to be the starting point. We can therefore write $[\![\hat{h}_{k_1}, \ldots, \hat{h}_{k_n}]\!]$ to represent the same combinatorial geodesic as $[v_0, \ldots, v_n]$.

In addition to being used to distinguish the use of hyperplanes to express a combinatorial geodesic from the use of vertices, the double brackets ($[\![\ ,\ ]\!]$) also serve to caution the reader that this notation is only well-defined with the standard assumption that v_0 is the beginning vertex of the corresponding combinatorial geodesic.

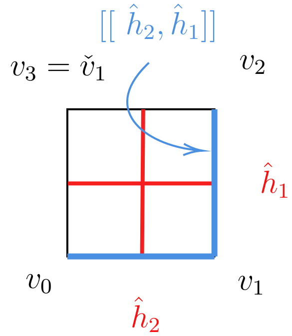

Figure 27

Example 2.32 Consider a square $c = J^2$ with vertex set $\{v_0, v_1, v_2, v_3\}$, where v_0, and v_2 are not adjacent (they are diametrically opposed). c has two hyperplanes $\hat{h}_1$ and $\hat{h}_2$, which are enumerated as in the accompanying figure. Then $[\![\hat{h}_1, \hat{h}_2]\!]$ and $[\![\hat{h}_2, \hat{h}_1]\!]$ describe two geodesic edge paths between v_0 and v_2. Note that when written with vertices, the effect on the combinatorial geodesic of switching the order in which $\hat{h}_1$ and $\hat{h}_2$ appear is to swap v_1 and v_3. In this example, $[\![\hat{h}_1, \hat{h}_2]\!]$ and $[\![\hat{h}_2, \hat{h}_1]\!]$ correspond to $[v_0, v_3, v_2]$ and $[v_0, v_1, v_2]$, respectively.

More generally, consider a geodesic edge path $\gamma = [v_1, \ldots, v_n]$ in which there are three consecutive vertices, v_{i-1}, v_i, and v_{i+1}, which belong to a common square c. Let $\check{v}_i$ denote the remaining vertex of c. Then exchanging $\check{v}_i$ for v_i results in a different geodesic edge path $\gamma' = [v_1, \ldots, \check{v}_i, \ldots, v_n]$. Following [Sag95], we refer to such a move as a **corner move**. Geometrically, the only difference between γ and γ' is that they travel through different edges of c. Otherwise, they traverse the same edges in the same order.

As Example 1.6 shows, an effect of performing a corner move on a geodesic edge path in the square c is to switch the order in which the two hyperplanes that intersect c appear in $[\![\hat{h}_{k_1}, \ldots, \hat{h}_{k_n}]\!]$. This fact is the essence of why we are able to choose "canonical" geodesic edge paths from v_0.

In [Sag95] corner moves provide the following result, which appears there as Theorem 4.6.

Lemma 2.33. *If u and v are vertices in a CAT(0) cube complex X, and α and β are two geodesic edge paths from u to v, then there exists a finite sequence of geodesic edge paths α_i from u to v such that $\alpha_1 = \alpha$, $\alpha_n = \beta$, and α_i and α_{i+1} differ by a corner move.*

We will use corner moves to rearrange the order in which geodesic edge paths cross hyperplanes. This technique will be used to prove Lemma 2.38.

Consider two hyperplanes $\hat{h}_1$ and $\hat{h}_2$ that intersect, and let x be a point on $\hat{h}_1$. By work done in [CFI16], the distance between x and $\hat{h}_2$ can be measured entirely in $\hat{h}_1$. More precisely, there exists a geodesic from x to $\hat{h}_2$ that lives entirely in $\hat{h}_1$.

Let (X, v_0) be a pointed locally-finite CAT(0) cube complex and suppose that $\hat{\mathcal{H}} = \{\hat{h}_i\}_{i=1}^{\infty}$ is given a distance-preserving enumeration. Consider a cube $c = [-1,1]^n$ in X. We will use the enumeration on $\hat{\mathcal{H}}$ and the orientations of each hyperplane to label local coordinates of points in c. Firstly, c has n midcubes, say $m_1, \ldots, m_n$, where the ordering is inherited from the ordering on $\hat{\mathcal{H}}$, that is, if m_i and m_j are midcubes of c that live on hyperplanes $\hat{h}_{k_i}$ and $\hat{h}_{k_j}$, respectively, then $i < j$ if and only if $k_i < k_j$.

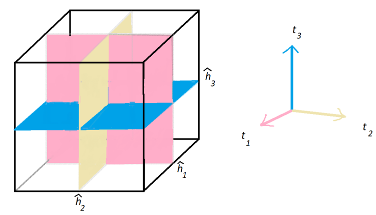

Figure 28: Local coordinates in a 3-cube

Then, let x be a point in c. To describe x with coordinates, observe that in $c = m_i \times [-1,1]$, the factor $[-1,1]$ gives the signed distance a point of c is from m_i. So, x can be given coordinates $(t_1, \ldots, t_n)$ where each t_i gives the signed distance from x to m_i. Finally, we orient the interval in the ith coordinate of c so that $m_i \times \{1\}$ is contained in h_{m_i}, the same halfspace as v_0. This establishes the signs of the coordinates of points in c. Figure 28 shows a schematic for this in a 3-cube.

Remark (A partition of $\hat{\mathcal{H}}$)**.** Let $\hat{h}$ be a hyperplane in a CAT(0) cube complex, X, and let γ be a geodesic edge path from v_0 to $\hat{h}$. Then γ is composed of edges which have length 2 and an additional segment on an edge that intersects $\hat{h}$ and has length 1. Summing all these lengths together shows that $\rho_2(v_0, \hat{h})$ is an odd integer. Then, for a locally-finite CAT(0) cube complex, the family $\hat{\mathcal{H}}$ can be subdivided into finite subsets corresponding to the hyperplanes a distance $2i-1$ from v_0, for $i = 1, 2, \ldots$. Let $F_i \subseteq \hat{\mathcal{H}}$ be the collection of hyperplanes a distance $2i-1$ from v_0, and let f_i denote the cardinality of F_i. If we let $n_i = f_1 + \cdots + f_i$, then $\hat{h}_{n_i}$ is the last hyperplane appearing in the enumeration of $\hat{\mathcal{H}}$ that is a distance $2i-1$ from v_0.

Recall Helly's property for hyerplanes: in a collection of pairwise intersecting hyperplanes, the total intersection is nonempty. There is a natural extension of Helly's property to the interiors of hyperplane carriers. Let $\mathcal{C}(\hat{h}_1), \ldots, \mathcal{C}(\hat{h}_n)$ be the carriers of the hyperplanes $\hat{h}_1, \ldots, \hat{h}_n$, respectively, and suppose $int(\mathcal{C}(\hat{h}_i)) \cap int(\mathcal{C}(\hat{h}_j)) \neq \emptyset$ for all i, j. It then follows that $\hat{h}_i \cap \hat{h}_j \neq \emptyset$ for all i, j. Helly's property for hyperplanes then implies that $\hat{h}_1 \cap \cdots \cap \hat{h}_n \neq \emptyset$. Clearly then, $\hat{h}_1 \cap \cdots \cap \hat{h}_n \subset \mathcal{C}(\hat{h}_1) \cap \cdots \cap \mathcal{C}(\hat{h}_n)$. Hence $\mathcal{C}(\hat{h}_1) \cap \cdots \cap \mathcal{C}(\hat{h}_n) \neq \emptyset$.

Lemma 2.34. *If $\hat{h}_1, \ldots, \hat{h}_n$ is a maximal collection of pairwise intersecting hyperplanes in a locally-finite CAT(0) cube complex X, then*

$$\mathcal{C}(\hat{h}_1) \cap \cdots \cap \mathcal{C}(\hat{h}_n) = c,$$

where c is a maximal cube in X. Moreover, every maximal cube c in X corresponds to a maximal collection of pairwise intersecting hyperplanes.

Proof. Let $\hat{h}_1, \ldots, \hat{h}_n$ be a maximal collection of pairwise intersecting hyperplanes, and let c be a cube in the total intersection $\mathcal{C}(\hat{h}_1) \cap \cdots \cap \mathcal{C}(\hat{h}_n)$. Then $\hat{h}_1 \cap \cdots \cap \hat{h}_n \cap c \neq \emptyset$. Since each hyperplane $\hat{h}_i$ contributes a midcube to c, we see that c has at least n midcubes and is therefore at least of dimension n. c cannot however contain any other midcubes because that would imply the existence of an additional hyperplane that intersects each $\hat{h}_i$ contradicting maximality. Therefore, c is of dimension n. If c is not maximal in X, then there is a cube $d = c \times J$ in X. d intersects the hyperplanes $\hat{h}_1, \ldots, \hat{h}_n$ and an additional hyperplane, say $\hat{h}$. But then $\hat{h}$ intersects each $\hat{h}_i$ contradicting maximality.

To see that $\mathcal{C}(\hat{h}_1) \cap \cdots \cap \mathcal{C}(\hat{h}_n) \subseteq c$, observe first that $\hat{h}_1 \cap \cdots \cap \hat{h}_n$ is a single point, namely, the center point of c. Well, if there were some cube d in $\mathcal{C}(\hat{h}_1) \cap \cdots \cap \mathcal{C}(\hat{h}_n)$ with $c \cap d = \emptyset$, then d would contain $\hat{h}_1 \cap \cdots \cap \hat{h}_n$ as well, a contradiction. Thus every cube in

$\mathcal{C}(\hat{h}_1) \cap \cdots \cap \mathcal{C}(\hat{h}_n)$ intersects c. It is not hard to see then that every cube in this intersection is a face of c.

Conversely, suppose $c = J^n$ is a maximal cube in X. For $i = 1, \ldots, n$, let m_i denote the midcubes of c, and let $\hat{h}_{k_i}$ denote the hyperplane containing m_i. Clearly, the hyperplanes $\hat{h}_i$, $i = 1, \ldots, n$ are pairwise intersecting. In fact it is immediate that the total intersection is nonempty:

$$\hat{h}_1 \cap \cdots \cap \hat{h}_n = m_1 \cap \cdots \cap m_n = (0, \ldots, 0),$$

where $(0, \ldots, 0)$ is the "center point" of c. Now, if there were some other hyperplane $\hat{h}$ that intersects $\hat{h}_i$ for $i = 1, \ldots, n$, then Helly's property for hyperplanes gives that the total intersection is nonempty. But then,

$$\hat{h}_1 \cap \cdots \cap \hat{h}_n \cap \hat{h} \subseteq \hat{h}_1 \cap \cdots \cap \hat{h}_n \subseteq c$$

Thus $\hat{h}$ intersects c. But hyperplanes intersect cubes in midcubes or not at all, and all the midcubes of c are already accounted for. Hence, the hyperplanes $\hat{h}_1, \ldots, \hat{h}_n$ give a maximal collection of pairwise intersecting hyperplanes. $\square$

Lemma 2.35. *Let X be a locally-finite CAT(0) cube complex, and let c be a maximal cube in X. If x is a point in X with $x \notin \mathring{c}$, then there exists a hyperplane $\hat{h}$ containing a midcube of c such that $\rho_1(x, \hat{h}) \geq 1$.*

Proof. Let $m_1, \ldots, m_n$ denote the midcubes of c. First, suppose $x \in c$, and let $[t_1, \ldots, t_n]$ be the cubical coordinates of x in c. Since x is not in the cubical interior of c, there must be some t_i with $t_i \in \{\pm 1\}$. Then $\rho_1(x, m_i) = 1$. But m_i belongs to a hyperplane, say $\hat{h}_i$. Therefore $\rho_1(x, \hat{h}_i) = 1$.

Now, suppose $x \notin c$. If there does not exist a hyperplane $\hat{h}$ containing a midcube of c with $\rho_1(x, \hat{h}) \geq 1$, then x is less than a distance 1 from all hyperplanes intersecting c. By maximality of c, this collection of pairwise intersecting hyperplanes will be a maximal collection. We thus have x belongs to the interiors of the carriers of all such hyperplanes.

Hence, x belongs to the total intersection of these carriers. By Helly's property for carriers, the total intersection of these carriers is c itself. Therefore $x \in c$, which is a contradiction. $\square$

2.3.5 Products and Adjunction Spaces

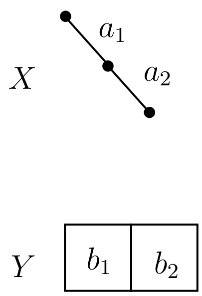

Figure 29: Two cube complexes

For cube complexes X and Y, with quotient maps $q : \mathscr{C} \to X$ and $q' : \mathscr{C}' \to Y$, the product $X \times Y$ is defined using the quotient map $q \times q' : \mathscr{C} \times \mathscr{C}' \to X \times Y$. Note that cubes in $\mathscr{C} \times \mathscr{C}'$ are of the form $c \times d$ for cubes c in $\mathscr{C}$ and d in $\mathscr{C}'$.

For cubes c_1 and c_2 in X that share a common face f, the cubes $c_1 \times d$ and $c_2 \times d$ share the common face $f \times d$. Similarly, cubes $c \times d_1$ and $c \times d_2$ have the face $c \times g$ in common, where d_1 and d_2 are cubes in Y that share a common face g. One can easily iterate this process to form finite products of cube complexes $X_1 \times X_2 \times \cdots \times X_n$.

In the figure on the left, X is a 1-dimensional cube complex and Y is a 2-dimensional cube complex. Figure 30 on the right shows their product $X \times Y$. Clearly, for finite-dimensional cube complexes X and Y, $X \times Y$ is a finite-dimensional cube complex with dimension equal to the sum of the dimensions of X and Y.

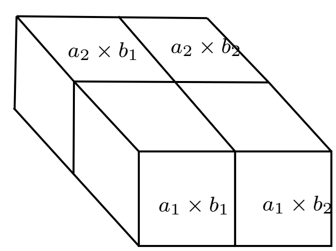

Figure 30: The product $X \times Y$

As a special case of a product, let K be a cube complex and let $J = [-1, 1]$ be a 1-cube with vertices $\{-1\}$ and $\{1\}$. Then $K \times J$ can be viewed as a cube complex whose cubes come in three flavors: $c \times J$, $c \times \{-1\}$, and $c \times \{1\}$. Note that each subset of the form $K \times \{t\}$ can be given a cube complex structure by removing the $\{t\}$ factor, and the resulting cube complex is isomorphic to K. Then, for all $t_1, t_2 \in [-1, 1]$ the map $N : \hat{h} \times \{t_1\} \to \hat{h} \times \{t_2\}$ defined by $N(x, t_1) = (x, t_2)$ composed with recognizing these spaces as cube complexes defines a natural isomorphism. As two special cases, let K^+ and K^- denote $K \times \{1\}$ and $K \times \{-1\}$, respectively. K^+ and K^- are special in part because they are actual subcomplexes of $K \times J$.

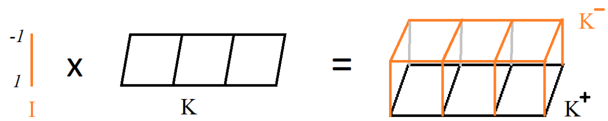

Figure 31: The product $K \times J$

For cubes of the form $c \times J$, $c \times \{0\}$ in fact gives a midcube of $c \times J$. Moreover, $K \times \{0\}$ is the union of midcubes of this form. Suppose K is a connected cube complex. Then each pair of intersecting cubes in K intersect in a common face, and since $K \cong K \times \{0\}$, it follows that $K \times \{0\}$ satisfies the conditions of Definition 2.29 and is therefore a hyperplane in $K \times J$. This hyperplane is disjoint from K and, the distance from $K \times \{0\}$ to K is clearly 1 (using either the ρ_1 or ρ_2 metrics). If K is not connected, then $K \times \{0\}$ gives a disjoint collection of hyperplanes in $K \times J$, one for each component of K. As a slight abuse of notation, we will sometimes identify K with K^+.

Recall property 4 of hyperplanes, which says that carriers have product neighborhoods. We previously wrote these as $\mathcal{C}(\hat{h}) \cong \hat{h} \times J = \hat{h} \times [-1, 1]$. This notation differs from the notation used in the definition of products of complexes. This is because $\hat{h}$ is not a subcomplex, while $\mathcal{C}(\hat{h})$ is. However, defining $\hat{h}^+$ and $\hat{h}^-$ as $\hat{h} \times \{1\}$ and $\hat{h} \times \{-1\}$, we can write carriers of hyperplanes in the following preferable way.

$$\mathcal{C}(\hat{h}) = \hat{h}^+ \times J.$$

The subcomplexes $\hat{h}^+$ and $\hat{h}^-$ each have an additional characterization due to the fact that $\hat{h}$ separates $\mathcal{C}(\hat{h})$ into two components. $\hat{h}^+$ consists of all cubes in $\mathcal{C}(\hat{h})$ that are contained in h, while $\hat{h}^-$ consists of the cubes in $\mathcal{C}(\hat{h})$ that are contained in h^*. Observe that the projection of $\hat{h} \times \{t\}$ onto the $\hat{h}$ coordinate is an isomorphism. Then $\hat{h} \times \{t\}$ and $\hat{h} \times \{s\}$ are isomorphic for all $s, t \in J$. In particular, with $\hat{h}$ identified with $\hat{h} \times \{0\}$, both $\hat{h}^+$ and $\hat{h}^-$ are isomorphic copies of $\hat{h}$. These isomorphisms induce an isomorphism between $\hat{h}^+ \times J$ and $\hat{h}^- \times J$.

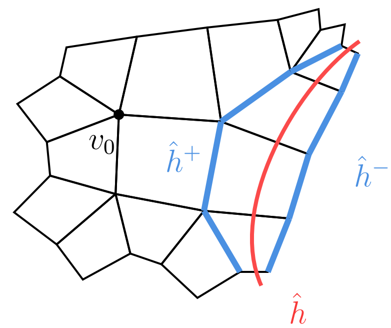

Figure 32: Subcomplexes $\hat{h}^+$ and $\hat{h}^-$

Lemma 2.36. *For a hyperplane $\hat{h}$ in a CAT(0) cube complex X,*

$$core(h) \cap \mathcal{C}(\hat{h}) = \hat{h}^+.$$

Proof. By definition $\hat{h}^+ \leq \mathcal{C}(X)$. Since $core(h)$ is the largest subcomplex (the union of all cubes) contained in h and $\hat{h}^+$ is contained in h, we have that $\hat{h}^+ \subseteq core(h) \cap \mathcal{C}(\hat{h})$.

Conversely, consider a cube in both $core(h)$ and $\mathcal{C}(\hat{h})$. This is a cube in $\mathcal{C}(\hat{h})$ contained in h. These are precisely the cubes in $\hat{h}^+$. Thus $core(h) \cap \mathcal{C}(\hat{h}) \subseteq \hat{h}^+$. $\square$

Of course the same argument used to prove Lemma 2.36 can be used to show that $core(h^*) \cap \mathcal{C}(\hat{h}) = \hat{h}^-$.

The next lemma identifies convex subcomplexes of $K \times J$, when K is itself a convex subcomplex a CAT(0) cube complex X.

Lemma 2.37. *Let K be a CAT(0) cube complex and $L \leq K \times J$. Then L is a convex subcomplex of $K \times J$ if and only if one of the following holds:*

a. *L is a convex subcomplex of $K^+ = K$*

b. *L is a convex subcomplex of K^-*

c. *$L = L' \times J$ for some convex subcomplex $L' \leq K$*

Proof. To begin, suppose $L \cap K^- = \emptyset$. Then it follows quickly that $L \leq K^+$. For, if L is not contained in K^+, then L contains a cube c not contained in K^+. By definition, either $c \in K^-$ or $c = d \times J$ for some cube $d \in K$. But if $c = d \times J$ for some cube $d \in K$, then c intersects K^- in the face $d \times \{-1\}$. In either case, $c \cap K^- \neq \emptyset$. A contradiction.

Now, suppose L is not contained in either K^+ or K^-. Then L must contain a cube of the form $d \times J$, for some $d \in K$. The desired result will follow if we show that all the maximal cubes of L are of the form $d \times J$, for some cube $d \in K$. So, let c be a maximal cube in L that shares a face with a cube of the form $e \times J \in L$, where $e \in K$. If c is not equal to $d \times J$ for some $d \in K$, then either $c \in K^+$ or $c \in K^-$. Suppose $c \in K^+$. Observe that $c \times J$ is a cube in $K \times J$ with c as a face. Let v be a vertex of $c \cap e$ and consider $lk_L(v)$. By convexity, $lk_L(v)$ is full in $lk_{K \times J}(v)$. Then, in $e \times J$, there is an edge $v \times J$ and this edge contributes a vertex to the simplex $lk_L(v)$. This vertex, together with the vertices provided by c give the vertex set of the simplex corresponding to $c \times J$. By fullness, the cube $c \times J$ must be L in order for the corresponding simplex to be in $lk_L(v)$. Thus all the maximal cubes in L are of the form $d \times J$, for some $d \in K$. By convexity of L, $L' = L^+$ is convex.

For cases (a) and (b), the converse follows directly from Corollary 2.24. Let $L = A \times J$ for some convex subcomplex $A \leq K$, and let v be a vertex in L. Either $v \in A$ or $v \in A^-$. Suppose $v \in A$. Then $lk_L(v) = lk_A(v) * \{p\}$, which is a full subcomplex of $lk_{K \times J}(v)$ by Lemma 2.7. The case for $v \in A^-$ is completely analogous. $\square$

Corollary 2.24 together with Lemma 2.37 guarantee that if X is CAT(0) and $K \times J \leq X$ is a subcomplex where K is convex, then convex subcomplexes of $K \times J$ of the form (a), (b), or (c) are also convex subcomplexes of X.

Lemma 2.38. *Let (X, v_0) be a pointed locally-finite CAT(0) cube complex, and let $[\![\hat{h}_{k_1}, \ldots, \hat{h}_{k_n}]\!]$ be a combinatorial geodesic between v_0 and some vertex v in X. Then there exists a permutation σ of the set $\{1, \ldots, n\}$ such that $[\![\hat{h}_{k_{\sigma(1)}}, \ldots, \hat{h}_{k_{\sigma(n)}}]\!]$ is a geodesic edge path from v_0 to v and $k_{\sigma(i)} < k_{\sigma(j)}$ for all $i < j$.*

Proof. We will show the existence of a geodesic edge path between v_0 and v which crosses hyperplanes in an order that respects the distance-respecting enumeration.

For $\gamma = [\![\hat{h}_{k_1}, \ldots, \hat{h}_{k_n}]\!]$, let j be the smallest integer such that $k_{j-1} > k_j$. We will first show that $\hat{h}_{k_{j-1}} \cap \hat{h}_{k_j} \neq \emptyset$. For, if they do not intersect then, from the fact that $\hat{h}_{k_{j-1}}$ separates X and by the distance-respecting enumeration, it must be the case that $\hat{h}_{k_j} \subset h_{k_{j-1}}$. In particular, $\hat{h}_{k_j}$ separates v_0 and $\hat{h}_{k_{j-1}}$. Clearly then, γ cannot cross $\hat{h}_{k_{j-1}}$ before crossing $\hat{h}_{k_j}$ when traveling from v_0 to v. Hence, $\hat{h}_{k_{j-1}} \cap \hat{h}_{k_j} \neq \emptyset$. It follows that $A = \mathcal{C}(\hat{h}_{k_{j-1}}) \cap \mathcal{C}(\hat{h}_{k_j})$ is nonempty and convex. Note that, since $\hat{h}_{k_{j-1}} \cap \hat{h}_{k_j} \neq \emptyset$, Lemma 2.37 implies that $A = A' \times J$ for some $A' \leq \hat{h}^+_{k_j}$.

Now, let $[v_0, \ldots, v_{n-1}, v] = [\![\hat{h}_{k_1}, \ldots, \hat{h}_{k_n}]\!]$. Then, in traveling from v_{j-2} to v_j, γ crosses $\hat{h}_{k_{j-1}}$ and then $\hat{h}_{k_j}$. Suppose v_{j-2}, v_{j-1}, and v_j do not belong to a common square. Then $[v_{j-2}, v_{j-1}, v_j]$ is a convex subcomplex. Note that the edge $[v_{j-2}, v_{j-1}]$ belongs to $\mathcal{C}(\hat{h}_{k_{j-1}})$, the edge $[v_{j-1}, v_j]$ belongs to $\mathcal{C}(\hat{h}_{k_j})$, and their intersection v_{j-1} belongs to A. In fact, $v_{j-1} \in A'$. Thus $v_{j-1} \times J \leq \mathcal{C}(\hat{h}_{k_{j-1}}) \cap \mathcal{C}(\hat{h}_{k_j})$. But $v_{j-1} \times J$ must be one of $[v_{j-2}, v_{j-1}]$ or $[v_{j-1}, v_j]$, a contradiction. Hence v_{j-2}, v_{j-1}, and v_j belong to a common square on which we can perform a corner move. This produces the transposition τ of the set $\{1, \ldots, n\}$ that transposes $j-1$ and j.

Finally, for each pair of consecutive integers $i-1, i$ with $k_{i-1} > k_i$ we can similarly produce a square that meets both $\hat{h}_{k_{i-1}}$ and $\hat{h}_{k_i}$. Performing corner moves on these squares produces corresponding transpositions. We need only repeat this performance finitely many times to produce finitely many transpositions that can be assembled into the desired permutation. $\square$

Lemmas 2.33 and 2.38 then imply that, given a geodesic edge path γ from v_0 to some vertex $v \in X$, there exist finitely many corner moves that can be used to transform γ into a geodesic edge path that crosses hyperplanes in an order that respects the distance-respecting enumeration.

Now, let X and Y be cube complexes with $X \cap Y = K$, a common subcomplex. Again, suppose X and Y have associated quotient maps $q : \mathscr{C} \to X$ and $q' : \mathscr{C}' \to Y$. Note that if $K \neq \emptyset$, then $\mathscr{C} \cap \mathscr{C}' \neq \emptyset$. Define $q \cup q' : \mathscr{C} \cup \mathscr{C}' \to X \cup Y$ by

$$q \cup q'(x) = \begin{cases} q(x), & x \in \mathscr{C} \\ q'(x), & x \in \mathscr{C}' \end{cases}$$

This defines a cube complex called the **adjunction** of X and Y along K, which we write as $X \cup_K Y$. Note that if X and Y are subcomplexes of the cube complex Z with $X \cap Y = K$, then $X \cup_K Y$ is a subcomplex of Z.

More generally, if $K \leq X$, $K' \leq Y$, and $f : K \to K'$ is an isomorphism, then one may define the adjunction space $X \cup_f Y$ as the cube complex formed by gluing X and Y together along their isomorphic subcomplexes K and K'. Cubes c in K are identified with $f(c)$.

Let $K \leq X$. Our primary interest in adjunction spaces will be in forming the space $X \cup_K (K \times J)$. This construction will provide the key ingredient to showing that all locally-finite CAT(0) cube complexes are collapsible.

Figure 33 shows an adjunction space $X \cup_K (K \times J)$. The subcomplex K is highlighted in red. The subcomplexes K^+ and K^- are highlighted in red and blue, respectively. Note that, in this example, the adjunction space of X and $K \times J$ along K results in a space that is not CAT(0), even though the subcomplex K is CAT(0). This is because the subcomplex K is not convex, as Corollary 2.39 shows. In proving the corollary, we will rely heavily on links of vertices. For a vertex $v \in K$, the effect of attaching $K \times J$ to X along K is to add a

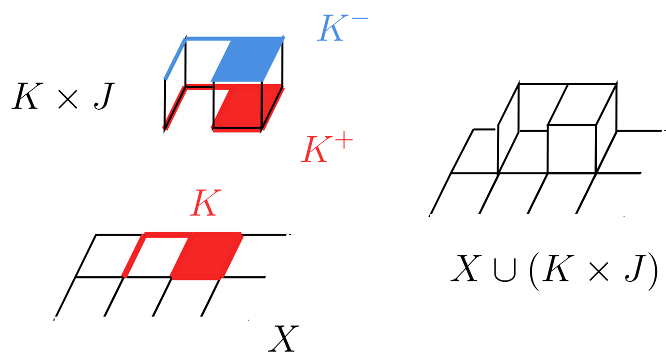

Figure 33: Forming the adjunction space $X \cup (K \times J)$

cone point, p, to $lk_K(v)$. That is, $lk_{K\times J}(v) = lk_K(v) * p$. The point p is contributed by the edge $v \times J$ in the product $K \times J$. This warrants its own remark.

Remark. In the space $X \cup_K K \times J$, for a vertex v in K, the link of v is equal to the cone $lk_K(v) * w$ attached to $lk_X(v)$ along $lk_K(v)$. The cone point w is not adjacent to any vertex in $lk_X(v) \setminus lk_K(v)$.

The following corollary follows from results in section 1.1.

Corollary 2.39. *Let X be a CAT(0) cube complex and $K \leq X$ be a convex subcomplex. Then the cube complex $X \bigcup_K (K \times J)$ is CAT(0).*

Proof. Suppose X is CAT(0), and choose a vertex $v \in Z = X \bigcup_K (K \times J)$. If $v \notin K \times J$ then $lk_Z(v) = lk_X(v)$, which is flag since X is CAT(0). Suppose then that $v \in K$. Then

$$lk_Z(v) = lk_X(v) \bigcup_{lk_K(v)} lk_K(v) * p,$$

where the cone point p comes from the product line $v \times J$. Since K is a convex subcomplex of both X and $K \times J$, $lk_K(v)$ is full in both $lk_X(v)$ and $lk_K(v) * p$. Lemma 2.6 then implies that $lk_K(v) * p$ is flag, and $lk_X(v)$ is clearly flag. Therefore, Lemma 2.9 implies that $lk_Z(v)$ is flag.

If $v \in K^-$, then $lk_{K^-}(v)$ is flag as K^- is convex. It then follows that $lk_Z(v)$ is the cone over $lk_{K^-}(v)$ with cone point again being contributed by the edge $v \times J$. Thus, $lk_Z(v)$ is flag by Lemma 2.6.

Finally, Z is simply connected by Van Kampen's Theorem. $\square$

2.3.6 Subcomplexes and Hyperplanes

For the final section of this chapter, we include a discussion on some interactions between hyperplanes and subcomplexes. This will allow us to explicitly describe halfspaces and cores of halfspaces in subcomplexes.

Definition 2.40. In a CAT(0) cube complex (X, v_0), we call a hyperplane $\hat{h}$ **extremal** if either h or h^* does not contain any other hyperplanes of X (other than $\hat{h}$). For an extremal hyperplane, whichever corresponding halfspace does not contain any hyperplanes is called its **extremal side**.

Lemma 2.41. *In a CAT(0) cube complex, the hyperplane $\hat{h}$ is extremal if and only if one of h or h^* does not contain a maximal cube of X.*

Proof. Suppose $\hat{h}$ is extremal with extremal side h. If h contains a maximal cube c, then it must be the case that c is at least 2-dimensional and $c \cap \hat{h} = \emptyset$. ($c$ cannot be 1-dimensional for then its midpoint would be a hyperplane contained in h.) Let $\hat{h}_1, \hat{h}_2, \ldots, \hat{h}_n$ be the total collection of hyperplanes intersecting c. Since $\hat{h}$ is extremal, it must be the case that $\hat{h} \cap \hat{h}_i$ is nonempty for all i. Helly's property then implies that total intersection is nonempty. Now, by maximality of c, $\hat{h}_1 \cap \cdots \cap \hat{h}_n$ equals the center point of c. But $\hat{h}_1 \cap \cdots \cap \hat{h}_n \cap \hat{h} \subseteq \hat{h}_1 \cap \cdots \cap \hat{h}_n$ implying that $c \cap \hat{h} \neq \emptyset$, a contradiction.

Conversely, assume that h contains no maximal cubes but h contains a hyperplane, say $\hat{k}$. Then, let m be a midcube on $\hat{k}$. Let c be the maximal cube containing m. By assumption, c cannot be contained in h. But then, either $c \cap \hat{h} \neq \emptyset$ or $c \in h^*$. In the first case, it would follow that $\hat{k}$ intersects $\hat{h}$. In the second case, we would have that m is contained in both h and h^*. In either case, we arrive at a contradiction. $\square$

Note the following special case of Lemma 2.41: the hyperplane $\hat{h}$ is extremal with h not containing any hyperplanes except for $\hat{h}$ if and only if h does not contain a maximal cube. A similar statement of course exists for h^*.

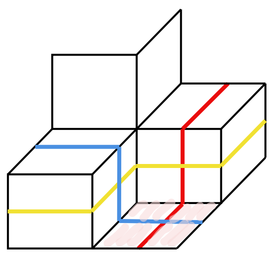

Figure 34

Example 2.42. Recall Example 2.3.4, which shows a cube complex X where Helly's property for hyperplanes fails. Figure 34 shows a finite subcomplex of X with three hyperplanes highlighted. The yellow hyperplane is extremal. However, Lemma 2.41 fails to hold in this case as there is a maximal cube contained in its extremal side.

One should compare the previous example with Lemma 2.34. The intersection of the carriers of the three hyperplanes shown in figure 26 consists of three edges joined at a singe vertex, which is not a maximal cube.

Recall Lemma 2.36, which says that $core(h) \cap \mathcal{C}(\hat{h}) = \hat{h}^+$. As $\mathcal{C}(\hat{h}) \cong \hat{h}^+ \times J$, it follows that

$$core(h) \bigcup_{\hat{h}^+} \mathcal{C}(\hat{h}) \cong core(h) \bigcup_{\hat{h}^+} \hat{h}^+ \times J \tag{2.1}$$

where $\hat{h}^+$ is identified with $\hat{h}^+ \times \{1\}$. Remark 2.3.5 then describes the links of vertices in $\hat{h}^+$. That is, for v a vertex in $\hat{h}^+$,

$$lk_X(v) = lk_{core(h)}(v) \bigcup_{lk_{\hat{h}^+}(v)} lk_{\hat{h}^+}(v) * \{p\}$$

Next, observe that $core(h)$ is a full subcomplex. Indeed, if $v_1, \ldots, v_{2^n}$ are vertices in $core(h)$ that span an n-cube in X, then this n-cube is contained entirely in h and hence is in $core(h)$.

Even more, h is a convex subcomplex. Note that all the vertices in h that are not contained in $\hat{h}^+$ have links equal to their links in X. Suppose $v \in \hat{h}^+$. Then

$$lk_X(v) = lk_{core(h)}(v) \bigcup_{lk_{\hat{h}^+}(v)} lk_{\hat{h}^+}(v) * \{p\},$$

where the cone point p is only adjacent to vertices in $lk_{\hat{h}^+}(v)$. Thus $lk_{core(h)}(v)$ is equal to $lk_X(v)$ with all the simplices containing p deleted. This is clearly a full subcomplex of $lk_X(v)$. Hence $core(h)$ is convex.

One can of course also demonstrate convexity of $core(h)$ via the more analytic definition.

Corollary 2.43. *Let $\hat{h}$ be an extremal hyperplane in a locally-finite CAT(0) cube complex X with h^* containing no hyperplanes except for $\hat{h}$. Then*

$$X = core(h) \bigcup_{\hat{h}^+} \mathcal{C}(\hat{h})$$

Proof. Let v be a vertex in X and suppose $v \notin \mathcal{C}(\hat{h})$. Let c be the maximal cube to which v belongs. By Lemma 2.41 $c \in h$. Hence, $c \in core(h)$. $\square$

Now, let (X, v_0) be a locally-finite CAT(0) cube complex and $K \leq X$ a connected subcomplex. Let $\hat{\mathcal{H}}_K$ denote the collection of hyperplanes in K. First, observe that, for each hyperplane $\hat{k} \in \hat{\mathcal{H}}_K$, there is a corresponding (unique) hyperplane $\hat{h}$ in X with $\hat{k} \subseteq \hat{h} \cap K$. Now, if K is convex then $\hat{h} \cap K$ is convex and it follows that $\hat{k} = \hat{h} \cap K$. However, even if K is CAT(0), it is not necessarily the case that hyperplanes of K are of the form $\hat{h} \cap K$, for some hyperplane $\hat{h}$ in X. Consider a 2-cube and let K be a subcomplex consisting of three consecutive edges. There will be a hyperplane of the 2-cube intersecting K in two components.

For a convex subcomplex $K \leq X$ then, it is clear that the hyperplanes of K are in a bijective correspondence with the set of hyperplanes of X that intersect K. Let $\hat{k} = \hat{h} \cap K$ be a hyperplane in K. If K contains the base vertex v_0 then we can define halfspaces k and k^* as the closures of the components of $K \setminus \hat{k}$, with $v_0 \in k$. One sees immediate that $k = h \cap K$ and $k^* = h^* \cap K$ and $\mathcal{C}(\hat{k}) = \mathcal{C}(\hat{h}) \cap K$.

Lemma 2.44. *Let (X, v_0) be a CAT(0) cube complex and K be a convex subcomplex that contains v_0. If $\hat{k} = \hat{h} \cap K$ is a hyperplane in K, then $core(k) = core(h) \cap K$ and $core(k^*) = core(h^*) \cap K$.*

Proof. Since $k = h \cap K$,

$$\begin{aligned} core(k) &= core(h \cap K) \\ &= core(h) \cap core(K) \\ &= core(h) \cap K. \end{aligned}$$

The same proof shows that $core(k^*) = core(h^*) \cap K$. □

Corollary 2.45. *Let (X, v_0) be a CAT(0) cube complex and K be a convex subcomplex that contains v_0. If $\hat{k} = \hat{h} \cap K$ is a hyperplane in K, then $\hat{k}^+ = \hat{h}^+ \cap K$.*

Proof. By Lemma 2.36, $\hat{k}^+ = core(k) \cap \mathcal{C}(\hat{k})$. By Lemma 2.44 we have $\hat{k}^+ = core(h) \cap \mathcal{C}(\hat{h}) \cap K$, and by another application of Lemma 2.36, we have $\hat{k}^+ = \hat{h}^+ \cap K$. □

We finish this section with a key definition and a handful of associated lemmas.

Definition 2.46 (Next to)**.** Let X be a cube complex and K a connected subcomplex. We say the hyperplane $\hat{h}$ is **next to** K if $\rho_1(K, \hat{h}) = 1$.

Remark. For a hyperplane $\hat{h}$ in a CAT(0) cube complex X, using the product structure of $\mathcal{C}(\hat{h}) = \hat{h}^+ \times J$, we can identify the locus of points a distance 1 from $\hat{h}$ (using either the taxicab or piecewise Euclidean metric). It is the subcomplex $\hat{h}^+ \cup \hat{h}^-$.

Example 2.47. Let $\hat{h}$ be a hyperplane in a locally-finite CAT(0) cube complex, and let $K = core(h)$. Then $\hat{h}^+ \leq K$ and it follows that no hyperplane of X could possibly separate K and $\hat{h}$. Thus, $\hat{h}$ is next to $K = core(h)$.

Similarly, $\hat{h}$ is next to $core(h^*)$, $\hat{h}^+$, and $\hat{h}^-$.

The following Lemma demonstrates a useful consequence when a hyperplane is next to a convex subcomplex of a CAT(0) cube complex.

Lemma 2.48. *Let K be a connected subcomplex of a locally-finite CAT(0) cube complex X. If $\hat{h}$ is next to K. Then, $\hat{h} \cap K = \emptyset$ and there does not exist a hyperplane separating K and $\hat{h}$.*

Proof. Suppose that $\rho_1(K, \hat{h}) = 1$. Clearly then $K \cap \hat{h} = \emptyset$. Suppose there is some hyperplane $\hat{k}$ that separates K and $\hat{h}$. Then any geodesic edge path γ between K and $\hat{h}$ that realizes the distance $\rho_1(K, \hat{h})$ will cross $\hat{k}$. But since $K \cap \hat{k} = \emptyset$, $\rho_1(K, \hat{k}) \geq 1$ implying that $\rho_1(K, \hat{h}) > 1$. A contradiction. □

In Lemma 2.22 in [CFI16] it is shown that if X is CAT(0) and $K \leq X$ is a convex subcomplex, then for any vertex $v \in X$, there is a unique point $p_K(v) \in K$ minimizing the combinatorial distance between v and K. Let L and K be subcomplexes of X. It follows that the distance between L and K can be measured between vertices. More explicitly, for $i = 1, 2$, there exist $u \in L$ and $v \in K$ such that $\rho_i(L, K)$ equals $\rho_i(u, v)$. And, if $\hat{h}$ is a hyperplane in X that is disjoint from K, then $\rho_i(K, \hat{h}) = \rho_i(v, \hat{h})$, for some vertex $v \in K$.

Corollary 2.49. *Let $K \leq X$ be a convex subcomplex, and let $\hat{h}$ be a hyperplane in X. If $\hat{h}$ is next to K, then either $K \cap \hat{h}^+ \neq \emptyset$ or $K \cap \hat{h}^- \neq \emptyset$.*

Proof. Let v be a vertex in K realizing the distance to $\hat{h}$. Then $\rho_2(v, \hat{h}) = 1$. By Remark 2.3.6 it follows that v belongs either to $\hat{h}^+$ or to $\hat{h}^-$. □

3 COLLAPSIBILITY OF CUBE COMPLEXES

In this chapter, we will define what it means for a cube complex to be collapsible. In particular, the definition we provide will apply to non-compact cube complexes. We begin with a couple of definitions which are analogous to the definitions one would first see when working in the setting of simplicial complexes.

Once the foundations are in place, we will describe an additional kind of collapse, called an *interval collapse* that is particulary well-suited for CAT(0) cube complexes. An single interval collapse allows one to combine a sequence of elementary cubical collapses into a singe move. Then, for a locally-finite CAT(0) cube complex X, we will describe a collection of nested convex finite subcomplexes, denoted $\{C_i\}$, with the property that C_0 is point and such that, for all i, there is an interval collapse from C_{i+1} to C_i. This process describes a collapsible CAT(0) cube complex $\cup C_i$. We will end the section by showing that, in fact, $X = \cup C_i$ thereby establishing the first main result of this paper, that all locally-finite CAT(0) cube complexes are collapsible.

3.1 Elementary Collapses and Collapsibility

First,we remind the reader of the classical definition of collapsibility. The definition is stated for simplicial complexes, but a version exists that applies to polyhedral complexes. In fact, an even more general definition that applies to CW complexes also exists, but we will not require such generality. One can find a more general treatment in [Coh73].

Definition 3.1 ((Simplicial) Free Face)**.** Let K be a simplicial complex containing simplices σ and τ. If $\sigma \leq \tau$ and no other simplices in K contains σ, then σ is called a **free face** in K.

Note that a free face σ of τ is of dimension exactly one less than τ.

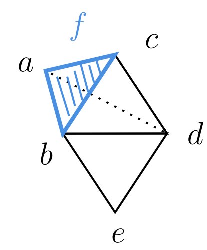

Figure 35: A simplicial free face, f.

Example 3.2. Figure 35 shows a 3-dimensional simplicial complex with vertex set $\{a, b, c, d, e\}$. The complex consists of a 2-simplex spanned by b, d, and e joined to a 3-simplex spanned by a, b, c, and d. The face f spanned by a, b, and c is a free face. In fact, every 2-dimensional face of the 3-simplex is a free face. However, the only edges that are free faces are the ones spanned by b and e and by d and e. These are all the free faces.

Definition 3.3 ((Simplicial) Collapsibility)**.** For a simplicial complex K with simplices σ and τ where $\sigma \leq \tau$ is a free face in K, the process of removing both σ and τ from K but retaining all other faces of τ such that the resulting space is a subcomplex of K is called an **elementary (simplicial) collapse**. If L denotes the resulting subcomplex, we denote this by $K \searrow^e L$.

We can similarly define an **elementary (simplicial) expansion** of L as a simplicial complex K such that $K \searrow^e L$. We may as well denote this as $L \,{}^e\!\nearrow K$.

A finite simplicial complex L is called **collapsible** if there exist subcomplexes $L_0, \ldots, L_n$ such that

1. L_0 is a vertex in L,

2. For $i < n$, $L_{i+1} \searrow^e L_i$, and

3. $\cup_{i=0}^n L_i = L$.

We now provide a couple of examples dealing with simplicial collapses.

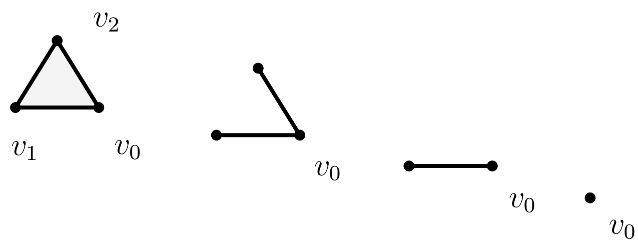

Figure 36: Simplicial collapses which show that the 2-simplex is collapsible

Example 3.4 (Collapsing a 2-Simplex)**.** Let σ be a 2-simplex spanned by vertices v_0, v_1, and v_2. $\{\{v_1\}, \{v_2\}, \{v_1, v_2\}\}$ is a free face in σ, and there is a corresponding elementary collapse taking σ onto a subcomplex consisting of two edges. Collapsing each of these edges, treating v_1 and v_2 as free faces, shows that σ is collapsible. See Figure 36.

It is easy to see that a collapsible simplicial complex is necessarily contractible. The following example shows that the converse of this statement can fail.

Example 3.5 (Dunce Hat)**.** Let σ be a 2-simplex. Identifying the edges of σ according to Figure 37 yields a CW complex, called the *dunce hat*. Regardless of how this space is triangulated, it will fail to have any free faces. Thus, the dunce hat is not collapsible. However, the identifications ensure that the boundary map is homotopic to the identity map on S^1. Hence the dunce hat is homotopic to a disc and therefore contractible.

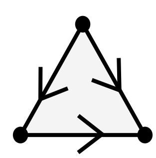

Figure 37: The dunce hat; a contractible but not collapsible space

The reader may be wondering why, in definition 3.3, the subcomplexes L_i are written in reverse order. Rather than starting with $L_0 = L$ and ending with L_n equal to a point, we have chosen to write the L_i in the reverse order to make the definition easier to extend to a definition that works even if infinitely many collapses are involved. One simply needs to adjust Definition 3.3 by allowing L to be an infinite subcomplex and also allowing countably many L_i with $\cup L_i = L$ and $L_{i+1} \searrow^e L_i$.

Definition 3.6. A cube in a cube complex X is a **free face** if it is a proper face of one and only one cube.

Note that a free face will necessarily be a facet of the cube for which it is a proper face of.

We first provide a purely combinatorial definition of a cubical collapse. More geometric interpretations will follow.

Definition 3.7. Let X be a cube complex. An **elementary collapse** in X is defined as the process of removing a free face and the cube to which it belongs. We denote this by $X \searrow^e (X - \mathring{c} - \mathring{f})$, where f is a free face and c the cube to which f is a proper face of. Note that the proper faces of f and the proper faces of c other than f still remain in the resulting complex.

In the above definition, the use of cubical interiors serves to emphasize the following geometric interpretation that exists for a cubical collapse. For a cube complex X, we could alternatively define an elementary collapse as the process of first removing $\mathring{c}$ from X, and then removing $\mathring{f}$. In fact, one could also define an elementary collapse as a deformation retract. We will provide the specifics of this later.

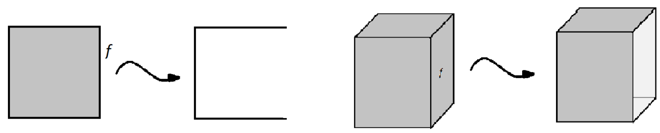

Figure 38: Elementary cubical collapses with free faces labeled f

While the process of an elementary collapse results in a subcomplex, the resulting space may no longer be CAT(0), even if the original space is CAT(0). For an example of this, consider an elementary collapse performed on a 3-cube. The resulting space will have vertices whose links are not flag complexes.

Definition 3.8. We say that L **collapses** to a subcomplex K, denoted $L \searrow K$, if there exists a, possibly infinite, sequence of subcomplexes D_i so that

1. $D_0 = K$
2. $D_{i+1} \searrow^e D_i$
3. $\bigcup_i D_i = L$

If $L \searrow K$ we say K **expands** to L, which we denote by $K \nearrow L$. We call L **collapsible** if L collapses to a vertex.

Note that, in the definition of collapsible, we allow the possibility of infinitely many collapses. This will allow us to prove collapsibility of non-compact spaces such as the cubified (Euclidean) plane, $\mathbb{E}^2$. It is essential then that the conditions of the definition can be demonstrated using elementary (cubical) expansions. Non-compact cube complexes such as $\mathbb{E}^2$ may not even have free faces.

3.2 Interval Expansions and Collapses

Next, we will show that certain adjunction spaces can be collapsed onto one of their summands. The following results will allow us to define maps which are equivalent in some sense to performing many elementary collapses at once.

Definition 3.9. For a subcomplex K of a cube complex X, we refer to the adjunction space

$$X \bigcup_K (K \times J)$$

as the **interval expansion of** X **along** K. As before, K is identified with $K \times \{1\}$, and we make the familiar identifications $K^+ = K = K \times \{1\}$ and $K^- = K \times \{-1\}$.

Definition 3.10. The **augmented dimension** of a finite cube complex is defined as the ordered pair (n, k), where the first coordinate is the dimension, n, of the complex, and the second coordinate represents the number of n-dimensional cubes.

Lemma 3.11. *Let X be a locally-finite CAT(0) cube complex, and $K \subset X$ be a finite subcomplex. If $Y = X \bigcup_K (K \times J)$, then $Y \searrow X$ via a finite sequence of elementary collapses.*

Proof. As above, we make the identifications $K = K^+ = K \times \{1\}$ and $K^- = K \times \{-1\}$. Then the faces contained in K^- will all necessarily belong to free faces of Y.

We prove the lemma by (transfinite) induction over the well-ordered (with lexicographic order) set $\mathbb{N} \times \mathbb{N}^+$, where an ordered pair (n, k) represents the augmented dimension of K. For the base case, K will be a singleton $\{v\}$, and $K \times J$ will be a single interval. In this case, Y will be a copy of X with this interval attached along $K = \{v\}$. The endpoint of this interval corresponding to -1 will be a free face. We can therefore accomplish the desired task with a single elementary collapse with respect to this free face.

For the inductive step, we suppose that for all L with augmented dimension less than (n, k) it is true that $X \bigcup_L (L \times J) \searrow X$. Let $K \subset X$ be a subcomplex with augmented dimension (n, k). There are two cases to consider: either $k = 1$, or $k > 1$. In the first case,

$K \times J$ will have a single $(n+1)$-cube, which is of the form $c \times J$ for some n-cube c of K. The face of this $(n+1)$-cube contained in K^-, call it f, will be a free face, on which we can perform an elementary collapse

$$X \bigcup_K (K \times J) \searrow^e X \bigcup_K (K \times J - \overset{\circ}{c} - \overset{\circ}{f}).$$

Note that

$$X \bigcup_K (K \times J - \overset{\circ}{c} - \overset{\circ}{f}) = X \bigcup_{K - \overset{\circ}{c}} ((K - \overset{\circ}{c}) \times J).$$

Since c was the only n-cube of K, $K - \overset{\circ}{c}$ is now a subcomplex of X with dimension strictly less than n. Thus the induction hypothesis applies and we get $X \bigcup_K ((K - \overset{\circ}{c}) \times J) \searrow X$ by a finite sequence of elementary collapses. Therefore $X \bigcup_K (K \times J) \searrow X$ via a finite sequence of elementary collapses.

Next, assume K has $k > 1$ n-cubes. Choosing any $(n+1)$-cube of $K \times J$, say $d \times J$ for some n-cube d in K, it will have a free face in K^- on which we can perform a single elementary collapse. The result of this will be

$$X \bigcup_K (K \times J) \searrow X \bigcup_{K - \overset{\circ}{d}} ((K - \overset{\circ}{d}) \times J)$$

Note that $K - \overset{\circ}{d}$ is a finite subcomplex of X with the same dimension as K, but with $k - 1$ n-cubes. Thus the induction hypothesis implies that

$$X \bigcup_K (K \times J) \searrow X \bigcup_{K - \overset{\circ}{d}} ((K - \overset{\circ}{d}) \times J) \searrow X$$

via a finite sequence of elementary collapses. $\square$

In light of Lemma 3.11, we make the following definition.

Definition 3.12. If Y is the result of an interval expansion of X along K, we call X an **interval collapse of Y along K.**

From now on, if the spaces X, K, and Y are understood, we will sometimes simply write *interval expansion* and *interval collapse* without reference to X, K, or Y.

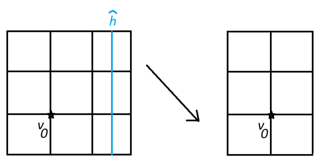

Figure 39: An interval collapse; $\hat{h}$ is extremal

Now, let (X, v_0) be a CAT(0) cube complex and $\hat{h}$ an extremal hyperplane in X with h^* extremal. By Corollary 2.43, and Equation (2.1) on page 47, we can write

$$X = core(h) \bigcup_{\hat{h}^+} (\hat{h}^+ \times J)$$

It then follows that $core(h)$ is an interval collapse of X along $\hat{h}^+$. Therefore, for an extremal hyperplane $\hat{h}$ with h^* extremal, $X \searrow core(h)$ via an interval collapse along $\hat{h}^+$.

More generally, let K be a subcomplex of a cube complex X with $\hat{h}$ next to K. Lemma 2.49 implies that either $K \cap \hat{h}^+ \neq \emptyset$ or $K \cap \hat{h}^- \neq \emptyset$. Assume, without loss of generality that $K \cap \hat{h}^+ \neq \emptyset$. The interval expansion of K along $K \cap \hat{h}^+$ has the form

$$K \bigcup_{K \cap \hat{h}^+} \Big((K \cap \hat{h}^+) \times J \Big).$$

The resulting complex can also be obtained by attaching a portion of $\mathcal{C}(\hat{h})$ to K along $K \cap \hat{h}^+$, and this is the essence of Lemma 3.13 below. Of course, K may only intersect $\hat{h}^+$ in a proper subcomplex, which explains the need to write $K \cap \hat{h}^+$. However, even in cases that $K \cap \hat{h}^+$ is a proper subset of $\hat{h}^+$, we will sometimes refer to $K \cup_{K \cap \hat{h}^+} ((K \cap \hat{h}^+) \times J)$ as the **interval expansion of K in the direction of $\hat{h}$.**

Lemma 3.13. *Let $Y \leq X$ be a subcomplex of (X, v_0) with $\hat{h}$ a hyperplane in X that is next to Y (with $Y \in h$). Then the interval expansion of Y in the direction of $\hat{h}$ is isomorphic to $Y \cup_K L$, where $K = Y \cap \hat{h}^+$ and L equals the collection of cubes in $\mathcal{C}(\hat{h})$ that intersect K in facets.*

Proof. Clearly, $L = \{c_j \mid c_j \in \mathcal{C}(\hat{h}) \; is \; maximal \; and \; c_j \cap K \; is \; a \; facet\}$. For each c_j in L, let $f_j = c_j \cap K$. Then $K = \cup f_j$, $c_j \cong f_j \times J$, and

$$\begin{aligned} Y \cup_K L &= Y \bigcup_K (\cup c_j) \\ &\cong Y \bigcup_K [\cup(f_j \times J)] \\ &= Y \bigcup_K [(\cup f_j) \times J] \\ &= Y \bigcup_K K \times J, \end{aligned} \tag{3.1}$$

□

When $\hat{h}$ is next to K, the interval expansion of K in the direction of $\hat{h}$ is nontrivial, meaning that K is a proper subset of the interval expansion. If $\hat{h}$ is not next to K and $\hat{h}^+ \cap K = \emptyset$, we say that the interval expansion is trivial. The case that $\hat{h}$ intersects K non-trivially, will not be considered.

Note that for a nontrivial interval expansion of K in the direction of the hyperplane $\hat{h}$, if L denotes the resulting complex, then L contains vertices that are at most a distance 2 away from K. Furthermore, for a vertex v in L with $\rho_1(v, K) = 2$, v will necessarily be a vertex in $(K \cap \hat{h}^+) \times \{-1\}$. Moreover, a geodesic edge path realizing $\rho_1(v, K)$ is given by the single edge $v \times J$.

Remark. Let (X, v_0) be a locally-finite CAT(0) cube complex, K a subcomplex, and L the interval expansion of K in the direction of $\hat{h}$. Let v be a vertex in $L - K$. Then the only additional hyperplane that any geodesic edge path connecting v with v_0 will cross is $\hat{h}$.

Remark 3.2 follows from the fact that, in X, hyperplanes are convex subsets that separate X into two components, one of which will contain v_0.

Corollary 2.39 on page 45 shows that the results of certain interval collapses and expansions preserve geometry in a strong sense. When Y and $K \times J$ are subcomplexes of a CAT(0) cube complex X with $K \leq Y$, Corollary 2.39 provides conditions necessary for performing an interval expansion of Y along K to be CAT(0). An even stronger statement would be that $Y \bigcup_K (K \times J)$ is a convex subcomplex of X whenever Y and K are convex subcomplexes. However, this is not always true. In Figure 40, take Y to be the subcomplex of $\mathbb{E}^3$ consisting of two 3-cubes attached along a common 2-cube. Although the subcomplex K highlighted in red is convex, the resulting adjunction space is not convex as a subcomplex of $\mathbb{R}^3$.

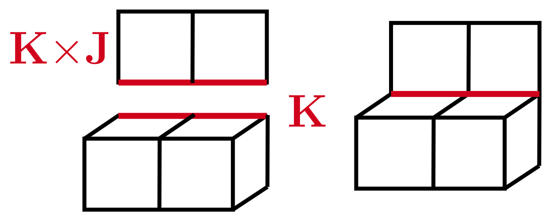

Figure 40: An interval expansion along the convex subcomplex K. The result is not convex.

In the space $Y \cup_K K \times J$, if we require K to be a connected subcomplex, then $K \times \{0\}$ in fact becomes a hyperplane of the resulting cube complex. Moreover, this hyperplane is extremal.

As further corollaries of Corollary 2.39, we have the following:

Corollary 3.14. *Let X be a CAT(0) cube complex, and Y a convex subcomplex. If $\hat{h}$ is a hyperplane next to Y, then the result of performing an interval expansion in the direction of $\hat{h}$ will be a CAT(0) subcomplex.*

Corollary 3.15. *Let X be a CAT(0) cube complex and let $\hat{h}$ be an extremal hyperplane in X with h^* not containing any hyperplanes other than $\hat{h}$. The result of performing the interval collapse in the direction of $\hat{h}$ will be a CAT(0) cube complex.*

In the following lemma, we show an instance where is convexity is preserved by an interval expansion.

Lemma 3.16. *Let X be a CAT(0) cube complex and let K be a convex subcomplex of X with $\hat{h}$ next to K. Then the interval expansion of X along K in the direction of $\hat{h}$ is a convex subcomplex of X.*

Proof. Let Y denote the interval expansion of X along K in the direction of $\hat{h}$. Then,

$$Y \cong K \bigcup_{\hat{h}^+ \cap K} \left((\hat{h}^+ \cap K) \times J \right).$$

Since $\hat{h}$ is next to K, $\mathcal{C}(\hat{h}) \cap K = \hat{h}^+ \cap K$ is convex because both K and $\hat{h}^+$ are convex. Now, let v be a vertex in Y. If $v \in K \setminus (\hat{h}^+ \cap K) \times [-1, 1]$, then the link of v in Y agrees with the link of v in K, and this link is a full subcomplex of $lk_X(v)$ because K is a convex of X. Suppose $v \in \hat{h}^+ \cap K$. Then

$$lk_Y(v) = lk_K(v) \bigcup_{lk_{\hat{h}^+ \cap K}(v)} lk_{\hat{h}^+ \cap K}(v) \star w,$$

where w is the cone point contributed by the edge $v \times J$ in $(\hat{h}^+ \cap K) \times J$. By Lemma 2.7 (page 8), $lk_{\hat{h}^+ \cap K}(v) \star w$ is full in $lk_X(v)$ because $lk_{\hat{h}^+ \cap K}(v)$ is full. Then, $lk_Y(v)$ is full by Lemma 2.10 (page 10). Finally, if $v \in (\hat{h}^+ \cap K) \times \{-1\}$, then $lk_Y(v) = lk_{\hat{h}^+ \cap K}(v) \star w$, which is full by Lemma 2.7 $\square$

3.3 Collapsibility of Locally-Finite CAT(0) Cube Complexes

In this section, we will describe a family of nested convex subcomplexes, $\{C_i\}$, of a locally-finite CAT(0) cube complex X. We will refer to the C_i's as *cubically expanded compacta.* We will show that, for all $i \geq 1$, C_i collapses onto C_i via finitely many elementary cubical collapses. We will conclude the chapter by showing that the family $\{C_i\}$ gives a compact exhaustion for X and thus proving collapsibility of locally-finite CAT(0) cube complexes. This result extends the results of [BL20], where collapsibility of finite square complexes is proven, and [AB19], where it is shown that compact CAT(0) cube complexes are collapsible. Concerning the scope of generality in article [AB19], math review MR4091542 is especially relevant.

Recall that we have enumerated the set $\hat{\mathcal{H}}_X$ so that it is distance-respecting. That is, the map $i \mapsto \rho_1(v_0, \hat{h}_i)$ is non-decreasing. Now, starting with $C_0 = \{v_0\}$, for $i \geq 1$, we inductively define C_i as the following subcomplex of X:

$$C_i = C_{i-1} \bigcup_{K_i} L_i$$

where $L_i \leq \mathcal{C}(\hat{h}_i)$ is the union of cubes c_j in $\mathcal{C}(\hat{h}_i)$ (along with their faces) such that c_j intersects C_{i-1} in a facet of c_j, and $K_i = C_{i-1} \cap \hat{h}_i^+$. If c_j is a cube in $\mathcal{C}(\hat{h})$ that intersects C_{i-1} in a facet, we let f_j denote the corresponding facet.

Of course, for this process to be non-trivial, meaning that C_{i-1} is a proper subset of C_i, it needs to be the case that $\mathcal{C}(\hat{h}_i) \cap C_{i-1}$ is nonempty. This will turn out to be the case, as we will show that the hyperplane $\hat{h}_i$ is next to C_{i-1}, for all i. Once we have established that each step in this construction is nontrivial, the requirements imposed on cubes in L_i will guarantee that the maximal cubes in L_i intersect $\hat{h}_i$.

To get our feet wet, let us investigate the first step in this construction. Clearly, $\hat{h}_1$ is next to C_0. Clearly then, $K_1 = C_0 \cap \hat{h}_1^+ = v_0$, and the only cube in $\mathcal{C}(\hat{h}_1)$ that intersects C_0 in a facet will be an edge that has v_0 as one of its endpoints. Thus

$$C_1 \cong C_0 \bigcup_{C_0} C_0 \times J$$

which is the interval expansion of C_0 in the direction of $\hat{h}_1$.

Now, for $i \geq 1$, with the assumption that $\hat{h}_i$ is next to C_{i-1}, Lemma 3.13 allows us to write

$$C_i \cong C_{i-1} \bigcup_{K_i} K_i \times J, \tag{3.2}$$

which is the interval expansion of C_{i-1} in the direction of $\hat{h}_i$. We could also call this the interval expansion of C_{i-1} along K_i. The C_i's thus have a dual nature. In their most innate sense, they are to be thought of as subcomplexes of X where one builds C_i from C_{i-1} by attaching a portion of $\mathcal{C}(\hat{h}_i)$ to C_{i-1}. This process results in a cube complex isomorphic to the interval expansion of C_{i-1} in the direction of $\hat{h}_i$.

Assume the interval expansion of C_{i-1} along $\hat{h}_i^+$ is nontrivial, and consider a vertex $v \in K_i = K_i^+$. Let j be the smallest integer for which $v \in C_j$. Since $v \notin C_{j-1}$, it must be the case that $v \in K_j^-$. Thus $K_j^- \cap K_i \neq \emptyset$. This will come up shortly when we prove that the cubically expanded compacta are convex. The links of vertices in C_i will be constructed from links in K_i and K_j^- for relevant i, j.

Equation (3.2) provides insight into the links of vertices of L_i. In particular, inside L_i are the subcomplexes $K_i^- = K_i \times \{-1\}$ and $K_i^+ = K_i \times \{1\}$. For each vertex $v \in K_i^+$, $v' = v \times \{-1\}$ is a vertex in K_i^- with $lk_{L_i}(v) \cong lk_{L_i}(v')$. From the product structure $L_i \cong K_i \times [-1, 1]$, it follows that $lk_{L_i}(v) \cong cone(lk_{K_i}(v))$ for all $v \in K_i$.

Note that Lemma 3.11 and Equation 3.1 together imply that C_i collapses to C_{i-1}.

The following theorem describes properties of the compacta C_i that distinguish them as useful subcomplexes.

Theorem 3.17. *Let (X, v_0) be a pointed locally-finite CAT(0) cube complex. For all $k = 0, 1, 2, \ldots$, the subcomplex*

$$C_k = C_{k-1} \bigcup_{K_k} L_k$$

satisfies the following:

1. C_k is convex.

2. The hyperplanes of C_k are in a one-to-one correspondence with the set $\hat{\mathcal{H}}_k = \{\hat{h}_i\}_{i=1}^k$ via the function

$$\hat{h}_i \mapsto \hat{h}_i \cap C_k$$

3. The vertex set of C_k is precisely the set of vertices in X all of whose geodesic edge paths to the basepoint cross only hyperplanes from the collection $\hat{\mathcal{H}}_k$.

4. For all $k \geq 1$, C_k collapses onto C_{k-1} via finitely many elementary cubical collapses.

Proof. Clearly, statements (1) - (3) are true for C_1. Now, suppose statements (1) - (3) are true for all $i < k$. We first argue that the hyperplane $\hat{h}_k$ is next to C_{k-1}. Let $x \in \hat{h}_k$ such that the combinatorial geodesic γ from v_0 to x realizes the combinatorial distance from v_0 to $\hat{h}_k$. Note x is then necessarily a vertex in $\hat{h}_k$. Choose a vertex v in $\hat{h}_k^-$ that is a distance 1 from x, and let γ be a geodesic edge path from v_0 to v with combinatorial geodesic $[\hat{h}_{k_1}, \ldots, \hat{h}_{k_n}]$. By Lemma 2.38 (page 43) we may assume $k_i < k_j$ for $i < j$. Suppose $\hat{h}_k = \hat{h}_{k_l}$, $1 \leq l \leq n$. Then $[\hat{h}_{k_1}, \ldots, \hat{h}_{k_l}]$ gives a combinatorial geodesic from v_0 to a vertex u contained in $\hat{h}_k^-$, and $[\hat{h}_{k_1} \ldots, \hat{h}_{l-1}]$ gives a combinatorial geodesic from v_0 to a vertex u' contained in $\hat{h}_k^+$. Clearly, $\rho_1(u', \hat{h}_k) = 1$. By the induction hypothesis, C_{k-1} satisfies (3) and therefore $u' \in C_{k-1}$. Thus

$\rho_1(C_{k-1}, \hat{h}_k) \leq 1$. (2) guarantees that $\rho_1(C_{k-1}, \hat{h}_k) \neq 0$. Thus $\rho_1(C_{k-1}, \hat{h}_k) = 1$ and we conclude that $\hat{h}_k$ is indeed next to C_{k-1}.

Then, since $\hat{h}_k$ is next to C_{k-1}, the induction hypothesis and Lemma 3.16 implies that C_k is convex.

Next, we show that C_k satisfies (2). By construction, $C_k \cap \hat{h}_k \neq \emptyset$. Also, for $i < k$, since $C_{k-1} \leq C_k$ and C_k is convex, $C_k \cap \hat{h}_i$ is nonempty and convex. We only need to ensure that there are no other hyperplanes of X that intersect C_k. Suppose there is some $l > k$ with $C_k \cap \hat{h}_l \neq \emptyset$. Since $C_{k-1} \cap \hat{h}_l = \emptyset$, it must be that $\hat{h}_l$ intersects $L_k \cong K_k \times J$. But for $\hat{h}_l$ to miss C_{k-1} and intersect L_k nontrivially, it must do so in the subset $K_k \times \{0\}$. But then $\hat{h}_l \cap L_k \subseteq \hat{h}_k$, from which it follows that $\hat{h}_l = \hat{h}_k$.

To show that C_k satisfies (3), let v be a vertex in C_k. Since $v_0 \in C_k$ and C_k is convex, Remark 2.3.2 (page 24) guarantees that any geodesic edge path between v_0 and v is contained entirely in C_k.

Conversely, suppose $v \in X$ is a vertex all of whose geodesic edge paths to the base point v_0 cross only hyperplanes from the collection $\hat{H}_k$. Let $\gamma = [\![\hat{h}_{a_1}, \ldots, \hat{h}_{a_n}]\!]$ be the geodesic edge path from v_0 to v that preserves the distance-respecting enumeration. If $v \notin C_k$, then $v \notin C_{k-1}$ and the induction hypothesis implies that any geodesic edge path (including γ) from v_0 to v must cross some hyperplane with index greater than $k-1$. Since the geodesic edge path γ is enumeration preserving, this hyperplane must be $\hat{h}_k$. Thus $a_n = k$ and $\gamma' = [\![\hat{h}_{a_1}, \ldots, \hat{h}_{a_{n-1}}]\!]$ is a geodesic edge path that crosses only hyperplanes from the set $\hat{\mathcal{H}}_{k-1}$. Letting v' denote the final vertex of γ', the induction hypothesis guarantees that $v' \in C_{k-1}$. Furthermore, v' and v are joined by an edge e. Now, e intersects $\hat{h}_k$ and thus is contained in $\mathcal{C}(\hat{h}_k)$. Therefore, v is contained on a cube e that meets C_{k-1} in a facet, v'. But this implies that $v \in L_k \leq C_k$.

Finally, that C_k collapses onto C_{k-1} via finitely many elementary cubical collapses for all $k \geq 1$ follows from Equation (3.2) and Lemma 3.11. $\square$

Example 3.18. Consider the standard cubulation of $\mathbb{R}^2$ with hyperplanes $\hat{\mathcal{H}} = \{\hat{h}_1, \hat{h}_2, \ldots\}$

Let v_0 be a chosen basepoint of $\mathbb{R}^2$ and enumerate $\hat{\mathcal{H}}$ in the following way: Observe that there are always exactly four hyperplanes a distance $2i - 1$ from the basepoint $(0, 0)$, and each of these four are distinguished by the direction (north, east, south, west) the geodesic connecting it to the basepoint must travel from the basepoint. For each of the four hyperplanes at equal distance from $(0, 0)$, we include them into the enumeration in the order north, east, south, and then west. This enumeration is distance-respecting. With $C_0 = \{v_0\}$,

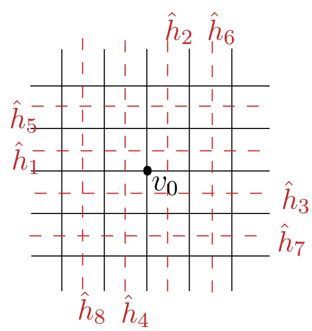

Figure 41: The Euclidean plane with a distance-respecting enumeration on hyperplanes

the following figure shows the first few corresponding cubically expanded compacta for the Euclidean plane. The hyperplane $\hat{h}_1 \cap C_i$ is shown in each compactum.

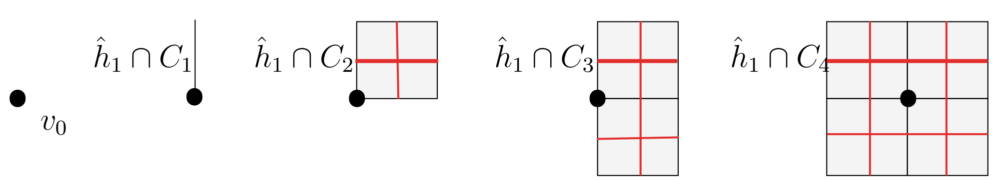

Figure 42: The cubically expanded compacta C_0 through C_4 for the Euclidean plane

Corollary 3.19. *Given a distance-respecting enumeration $\hat{\mathcal{H}} = \{\hat{h}_i\}_{i=1}^{\infty}$ of the hyperplanes of (X, v_0), let $\{C_i\}_{i=1}^{\infty}$ denote the corresponding cubically expanded compacta. Then C_i is precisely the full subcomplex of X whose vertex set consists of those vertices whose geodesic edge paths to v_0 intersect no hyperplanes except for $\hat{h}_1, \ldots, \hat{h}_i$.*

Proof. By Theorem 3.17, C_i is convex and thus a full subcomplex. Uniqueness follows from Lemma 2.19 on page 18. □

Corollary 3.20. *Given a distance-respecting enumeration of $\hat{\mathcal{H}}$, the hyperplanes of (X, v_0), let $\{C_i\}$ denote the corresponding cubically expanded compacta. Then*

$$\bigcup_i C_i = X.$$

Proof. Choose a vertex v in X, and consider the collection of geodesic edge paths between v and v_0. By local-finiteness, there will be finitely many such edge paths. Let $\hat{h}_{l_1}, \ldots, \hat{h}_{l_k}$ denote the (finite) collection of hyperplanes crossed by all possible geodesic edge baths between v and v_0. Here the ordering of subscripts is inherited from the enumeration on the set of all hyperplanes. By Theorem 3.17, C_{l_k} contains all vertices of X whose geodesic edge paths to v_0 only cross hyperplanes from the collection $\hat{h}_1, \hat{h}_2, \ldots, \hat{h}_{l_k}$. Clearly then v is contained in C_{l_k}. □

Corollary 3.20 is worth comparing to Theorem 4.2 in [BL14], where it is shown that a locally-finite CAT(0) simplicial complex is a monotone union of collapsible subcomplexes. However, their results fall short of proving collapsibility of non-compact simplicial complexes.

Theorem 3.21. *All locally-finite CAT(0) cube complexes are collapsible.*

Proof. Let X be a CAT(0) cube complex and let $\{C_i\}$ denote the cubically expanded compacta corresponding to a distance-respecting enumeration on $\hat{\mathcal{H}}_X$. By Corollary 3.20, $\cup_i C_i = X$. By Theorem 3.17, each C_i collapses onto C_{i-1} via finitely many elementary cubical collapses. Finally, as $C_0 = \{v_0\}$, it follows that X is collapsible. □

4 COMPACTIFYING FROM COLLAPSING MAPS

4.1 Topological Collapses

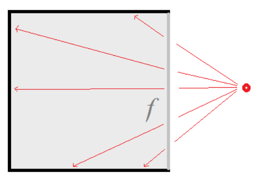

Figure 43: $c \searrow^e c - \mathring{c} - \mathring{f}$

One can realize a combinatorial elementary collapse topologically as the end result of a particularly nice deformation retraction. For a free face f of the cube c, we will define a map which, starting from a point directly opposite f from the center point of c, pushes radially outward toward the remaining facets of $c - \mathring{f}$ continuously mapping c onto $c - \mathring{c} - \mathring{f}$. To this end, let $c = J^n$ be an n-cube. Embed c in $\mathbb{E}^n$ by the map $(x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_n)$. We will treat the facet $\{1\} \times J^{n-1}$ as a free face and describe an elementary collapse on c by a homotopy. Let $b > 1$ and $x = (x_1, \ldots, x_n)$ be a point in c. Then there is a line between x and the point $(b, 0, \ldots, 0)$. This line will necessarily intersect $c - \mathring{c} - \mathring{f}$ in a unique point. Call this intersection point y and let l_x denote the line segment between x and y parametrized isometrically by $[0, 1]$ to start at x $(t = 0)$ and end at y $(t = 1)$. Define the map $\alpha : c \times [0, 1] \to c - \mathring{c} - \mathring{f}$ by letting $\alpha(x, t)$ be the position at time t along l_x. We leave it to the reader to verify that α is continuous. α thus defines a deformation retraction of c. Let $r(x) = \alpha(x, 1)$. Then r defines a retraction of c onto $c - \mathring{c} - \mathring{f}$.

If c belongs to a cube complex X, then the homotopy α extends to the identity homotopy on $X - \mathring{c} - \mathring{f}$.

If $c = [-1, 1]$ is an edge with endpoints $u = \{1\}$ and $v = \{-1\}$, then, with v as a free face, this deformation describes a contraction of the interval $[-1, 1]$ onto the endpoint u. The following example shows precisely what we mean by this in the case of a 2-cube.

Example 4.1. *Let $c = J^2$ be a 2-cube isometrically embedded in $\mathbb{E}^2$ by the map $(u_1, u_2) \mapsto (u_1, u_2)$, and let $b > 1$. We will treat the edge $f = \{1\} \times J$ as a free face. For each point $u = (u_1, u_2) \in c$, the line l_u through u and $(b, 0)$ has the form*

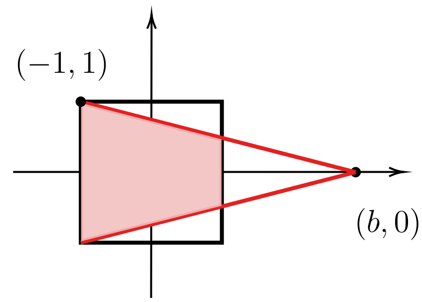

Figure 44

$$y = \frac{u_2}{u_1 - b}x + \frac{bu_2}{b - u_1}.$$

Where this line intersects $c - \overset{\circ}{c} - \overset{\circ}{f}$, call the intersection point Q_u. Q_u will belong to one of the three edges other than f. However, which one depends on the position of the point u. The three cases are as follows:

- If $\frac{u_1}{1+b} - \frac{b}{1+b} \leq u_2 \leq -\frac{u_1}{1+b} + \frac{b}{1+b}$ then l_u intersects the edge $\{-1\} \times J$.
- If $-1 \leq u_2 < \frac{u_1}{1+b} - \frac{b}{1+b}$ then l_u intersects edge $J \times \{-1\}$.
- If $-\frac{u_1}{1+b} + \frac{b}{1+b} < u_2 \leq 1$ then intersects edge $J \times \{1\}$.

Letting $\alpha : c \times [0,1] \to c$ defined by letting each point $u \in c$ travel along l_u ending at the point $Q_u \in c - \overset{\circ}{c} - \overset{\circ}{f}$ defines a deformation retraction realizing an elementary cubical collapse.

We will return to this alternative viewpoint of elementary cubical collapses as continuous deformations later when we describe CAT(0) cube complexes as inverse systems of compact subcomplexes. We now return our attention to interval collapses. We will define a different deformation retraction which has image equal to the result of doing a single interval collapse.

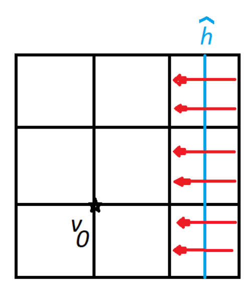

Figure 45: (Topological) Interval Collapse

Definition 4.2 (interval collapse (topological version)). Let K be a subcomplex of the cube complex A. The map

$$\phi : \left[A \bigcup_K (K \times J) \right] \times [0,1] \ \rightarrow \ A \bigcup_K (K \times J)$$

defined by

$$\phi((k,s),t) = (k, s(1-t)+t)$$

for $(k,s) \in K \times J$, $t \in [0,1]$ and

$$\phi(x,t) = x$$

for all $x \in A$ and $t \in [0,1]$, is called a **topological interval collapse**.

Lemma 4.3. *The map ϕ defined in Definition 4.2 is continuous.*

Proof. For points in $K \times J$, the coordinate functions are clearly continous. For points in A, ϕ is the identity at all times t and thus continuous. Then, since $K = K \times \{1\}$ and $\phi((k,1),t) = (k,1)$ for all $t \in [0,1]$, the map gluing lemma guarantees that ϕ is continuous on all of $A \cup_K K \times J$. □

Remark. Observe that at time $t = 1$ of a topological interval collapse ϕ, the point $(k,s) \in K \times J$ is mapped to $(k,1)$. All points in A are fixed. In particular, points in $K = K \times \{1\}$ are fixed. The homotopy ϕ is thus a deformation retraction, and the map $r = \phi(x,1)$ defines a retraction of $X = A \cup_K K \times J$ onto A.

Definition 4.4. Let $\varphi : X \times I \rightarrow X$ be a homotopy. We say that φ is **track-faithful** if $\varphi(\varphi(x,t),1) = \varphi(x,1)$ for all $x \in X$ and $t \in I$.

Lemma 4.5. *Let A be a cube complex and K a subcomplex. Then the topological interval collapse*

$$\phi : \left[A \bigcup_K (K \times J) \right] \times [0,1] \ \rightarrow \ A \bigcup_K (K \times J)$$

is track-faithful.

Proof. Let $(k,s) \in K \times J$ and $t \in [0,1]$. Then

$$\begin{aligned}\phi(\phi((k,s),t),1) &= \phi((k, s(1-t)+t),1) \\ &= (k, [s(1-t)+t](1-1)+1) \\ &= (k,1) \\ &= \phi((k,s),1).\end{aligned}$$

□

Remark. The deformation retractions described in section 3.1 are also track-faithful. See Example 4.1. We leave it to the reader to check the details.

Example 4.6. *Let X be a CAT(0) cube complex, and let $\hat{h}$ be an extremal hyperplane in X. By Corollary 2.43, $X = core(h) \cup_{\hat{h}^+} \mathcal{C}(\hat{h}) \cong core(h) \cup_{\hat{h}^+} \hat{h}^+ \times J$. There is thus a topological interval collapse ϕ which, at time $t = 1$, gives a retraction*

$$r : core(h) \cup_{\hat{h}^+} \mathcal{C}(\hat{h}) \to core(h).$$

Consider a point $(k,s) \in \mathcal{C}(\hat{h})$. At all times t, $\phi(x,t)$ is contained in the product line $x \times [-1,1]$ (with x identified with $x \times \{1\}$). Moreover, Lemma 4.5 guarantees that every point on the product line $x \times [-1,1]$ is mapped to x under ϕ at time 1 by track-faithfullness.

Note that, for a topological collapse ϕ of the space $A \cup_K K \times J$, the image of the retraction $r(x) = \phi(x,1)$ agrees with the result of doing a combinatorial interval collapse (see Definitions 3.9 and 3.12). Then, for a finite cube complex K, there is a finite sequence of elementary collapses of $K \times J$, the end result of which agrees with the image of $r : K \times J \to K$.

4.2 Z-Compactifications, Z-Sets, and Inverse Limits

In this section, we will touch briefly on topics related to Z-sets and inverse limits. For a more detailed exposition, see [Gui16] and [GM19].

Definition 4.7 (Homotopy Negligible, Z-Set)**.** A subset A of a space X is said to be **homotopy negligible** if there exists a homotopy $H : X \times [0,1] \to X$ such that $H_0 = id_X$ and $H_t(X) \subset X - A$ for all $t > 0$. If X is a locally-compact metric space and A is a closed homotopy negligible subspace, then we call A a **Z-set**.

Perhaps the prototypical examples of Z-sets are closed subsets of ∂M for M a manifold.

A **Z-compactification** of a space X is a compactification $\overline{X} = X \sqcup Z$ with the property that Z is a Z-set in $\overline{X}$. In this case, Z is also referred to as a Z-boundary for $\overline{X}$. An attractive feature of Z-compactifications is that they preserve the homotopy type of the space. In particular, if X is contractible, then so is $\overline{X}$.

Now, let $\{X_i\}_{i=0}^{\infty}$ be a monotone sequence of nested compacta and let $X = \cup_{i=0}^{\infty} X_i$. Suppose there exist retractions $r_i : X_i \to X_{i-1}$ for all $i \geq 1$. The situation is neatly described using an *inverse sequence*

$$X_0 \xleftarrow{r_1} X_1 \xleftarrow{r_2} X_2 \xleftarrow{r_3} \cdots \tag{4.1}$$

The following definition describes a way of using the maps r_i to assemble the compacta X_i into a subspace of the product $\prod_{i=1}^{\infty} X_i$.

Definition 4.8 (Inverse Limit)**.** For the inverse sequence in (4.1), its **inverse limit** is defined as

$$\varprojlim\{X_i, r_i\} = \left\{ (x_0, x_1, \ldots) \in \prod_{i=0}^{\infty} X_i \mid r_i(x_i) = x_{i-1} \; for \; all \; i > 0 \right\}$$

Where $\varprojlim\{X_i, r_i\}$ is topologized as a subspace of $\prod_{i=0}^{\infty} X_i$ with the product topology.

A key feature of inverse limits for us is the fact that if X_i is compact for all $i \geq 0$ then $\varprojlim\{X_i, r_i\}$ is compact.

The following remark is common knowledge among those familiar with inverse limits.

Remark. Consider again the inverse sequence in (4.1). Given any subsequence of natural numbers, $\{k_i\}_{i=0}^{\infty}$, there is a corresponding inverse sequence

$$X_{k_0} \xleftarrow{r'_1} X_{k_1} \xleftarrow{r'_2} X_{k_2} \xleftarrow{r'_3} \cdots$$

which is a subsequence of (4.1). Here, $r'_i = r_{k_{i-1}+1} \circ r_{k_{i-1}+2} \circ \cdots \circ r_{k_i}$. Moreover, $\varprojlim\{X_i, r_i\} \approx \varprojlim\{X_{k_i}, r'_i\}$.

4.3 The Cubical Compactification and Boundary

Recall the cubically expanded compacta $\{C_i\}_{i=0}^{\infty}$. Since $C_i \cong C_{i-1} \cup_{K_i} K_i \times J$, there is a topological interval collapse $\phi_i : C_i \times [0,1] \to C_i$. Let $r_i(x) = \phi(x,1)$. Then r_i defines a retraction of C_i onto C_{i-1}. Using these retractions as bonding maps, we assemble the collection $\{C_i\}$ into an inverse sequence:

$$\{v_0\} = C_0 \xleftarrow{r_1} C_1 \xleftarrow{r_2} C_2 \xleftarrow{r_3} \cdots$$

As each C_i is compact, the inverse limit, $\varprojlim\{C_i, r_i\}$, is compact.

In this section, we will show that a locally-finite CAT(0) cube complex $X = \cup C_i$ embeds in $\varprojlim\{C_i, r_i\}$ via a map α. It will be shown that $\overline{\alpha(X)}$ provides a compactification of $\alpha(X)$. We will then show that the remainder space $\overline{\alpha(X)} - \alpha(X)$ is a Z-set and thus $\overline{\alpha(X)}$ provides a Z-compactification of X. The following lemma will be used to obtain a well-defined map, α, which will serve as the embedding map of X into $\overline{\alpha(X)}$.

Before proceeding, we say a few words about the topology on X. Since $X = \cup_i^{\infty} C_i$, one can naturally associate to X the weak topology with respect to the compacta C_i, whereby a subset U of X is open if and only if $U \cap C_i$ is open in C_i for all i. Since X is locally-finite, this topology agrees with the topology induced by the path-length (Euclidean) metric on X. As contrast, consider the hedgehog space shown in Figure 7. This space is not locally-finite and, in this case, the weak topology does not agree with the path-length topology.

Lemma 4.9. *Let X be a locally-finite CAT(0) cube complex and endow $\hat{\mathcal{H}}$ with a distance-respecting enumeration. Let $\{C_i\}$ denote the corresponding cubically expanded compacta. Then for all i, there is a j ($j > i$) with $C_i \subset \mathrm{int}_X(C_j)$.*

Proof. Let S be the star neighborhood of C_i. By Corollary 3.20, there is some j with $S \subset C_j$. Then $C_i \subset \mathrm{int}(S) \subset \mathrm{int}(C_j)$. □

In light of Lemma 4.9 and Corollary 3.20, we can now write

$$X = \bigcup_{i=0}^{\infty} int(C_i)$$

Corollary 4.10. *For all i, there exists a k so that $r_j^{-1}(x) = x$ for all $x \in C_i$ and $j > k$.*

Proof. By Lemma 4.9, there exists a $k > i$ with $C_i \subset \mathrm{int}(C_k)$ Let $j > k$ and $x \in C_i$. Then $r_j(x) = x$ and, there cannot exist $y \in C_i$ with $y \neq x$ and $r_j(y) = x$, for otherwise, x would belong to $L_j \cong K_j \times J$ (See Remark 4.1). But $C_i \subseteq \mathrm{int}(C_k)$ implies that $C_i \subseteq \mathrm{int}(C_{j-1})$, and $\mathrm{int}(C_{j-1}) \cap L_j = \emptyset$. □

Recall that for $x \in X$, by Corollary 3.20, there is a j with $x \in C_j$. For $i < j$, define $r_{ij} : C_j \to C_i$ by $r_{ij}(x) = r_i \circ r_{i+1} \circ \cdots \circ r_j(x)$. We let $r_{ii} = r_i$ and we leave r_{ij} undefined for $j < i$. Observe that r_{ij} is a composition of continuous functions and is therefore continuous. Clearly, r_{ij} defines a retraction from C_j to C_i.

Now, consider the function $s_i : X \to C_i$ defined as the union of maps $\cup_{j \geq i} r_{ij}$. That s_i is well-defined follows from the fact that r_{ik} agrees with r_{il} on $C_k \cap C_l = C_{min\{k,l\}}$. Since r_j fixes points in C_i for $i < j$, we have that if $x \in C_i$ and $i < j$, then $r_{ij}(x) = r_i(x)$. Therefore,

$$s_i(x) = \begin{cases} r_i(x), & x \in C_i \\ r_{ij}(x), & j > i,\ x \in C_j,\ and\ x \notin C_{j-1} \end{cases}$$

To show that s_i is continuous, we observe that s_i is continuous when restricted $int(C_k)$, for $k = 0, 1, \ldots$ Indeed, if $k \leq i$ then the restriction of s_i to $int(C_k)$ agrees with the map r_i restricted to $int(C_k)$. And, if $k > i$ then the restriction agrees with r_{ik}. Then, since

$X = \cup_{k=0}^{\infty} int(C_k)$, it follows that s_i is continuous.

The following function will serve as the embedding of X into $\varprojlim\{C_i, r_i\}$. Define

$$\alpha : X \to \prod_{i=0}^{\infty} C_i$$

by $\alpha(x) = (s_0(x), s_1(x), s_2(x), \ldots)$. That α is well-defined, follows from the fact that each s_i is well-defined.

Next we show that, for $x \in X$, $\alpha(x)$ is the sequence whose coordinates are eventually the point x itself. The following claim makes this precise.

Claim: For $x \in X$, let $\alpha(x) = (x_0, x_1, \ldots)$. Then, there exists a k such that if $j \geq k$ then $x_j = x$. Furthermore, if $\overline{y} = (y_0, y_1, \ldots)$ is any point in $\alpha(X)$ such that there exists a k' with $y_j = x$ for all $j \geq k'$, then $\overline{y} = \alpha(x)$.

Proof of Claim. Let $x \in X$. We can take k to be such that $x \in C_{k-1}$ but $x \notin C_{k-2}$. For $j \geq k$, $s_j(x) = x$ proving the first part of the claim.

Suppose $\overline{y} = (y_0, y_1, \ldots) \in \alpha(X)$ such that there exists k' with $y_j = x$ for $j \geq k'$. Let $y \in X$ with $\alpha(y) = \overline{y}$. If $y \neq x$, then $\alpha(y)$ is the sequence whose coordinates are eventually y, which cannot be the same sequence as $\alpha(x)$. □

Note that for $k > j$ and $x \in C_j$, $s_k(x) = x$. Thus, for $x \in X$, $\alpha(x)$ is the sequence whose coordinates are eventually the point x itself.

Lemma 4.11. *The map α is an embedding.*

Proof. Throughout the proof the metric on X is assumed to be the Euclidean path-length metric.

That α is continuous follows from the fact that each s_i is continuous. To show that α is one-to-one, let $x, y \in X$ with $x \neq y$, $x \in C_i$, and $y \in C_j$. Then, for $k >$max$\{i, j\}$, $s_k(x) = x$ and $s_k(y) = y$. Since $x \neq y$, it follows that $\alpha(x) \neq \alpha(y)$.

Lastly, we will show that α is open onto $\alpha(X)$. To this end, let $x \in X$ and let $B_\epsilon(x)$ be an open metric ball in X centered at x and with radius $\epsilon > 0$. Let $\overline{x} \in \alpha(B_\epsilon(x))$. By Lemma 4.9, $B_\epsilon(x) \subseteq \text{int}(C_i)$ for some i. In particular then, $\overline{x} = (x_0, x_1, \ldots)$ is contained $\alpha(int(C_i))$. Thus, $x_j = x_i$ for all $j \geq i$ and we conclude that $\overline{x} = \alpha(x_i)$. Here we are implicitly using the fact that $r_j^{-1}(x_i) = x_i$ for all $j \geq i$ (see Corollary 4.10). In particular, $x_i \in B_\epsilon(x)$. Take $0 < \delta < \epsilon$ such that $B_\delta(x_i) \subseteq B_\epsilon(x)$. Then, let $U = B_\delta(x_i) \times \Pi_{j \neq i} C_j$ and $V = U \cap \alpha(X)$. Then V is an open set in the subspace topology with $\overline{x} \in V$. Clearly,

$$V = (B_\delta(x_i) \times \Pi_{j \neq i} C_j) \cap \alpha(X) \subseteq (B_\epsilon(x) \times \Pi_{j \neq i} C_j) \cap \alpha(X) \subseteq \alpha(B_\epsilon(x))$$

□

Corollary 4.12. $\overline{\alpha(X)} = \varprojlim\{C_i, r_i\}$.

Proof. Let $\overline{x} \in \overline{\alpha(X)}$. Suppose $\overline{x} = \alpha(x)$ for some $x \in X$ and let $\alpha(x) = (x_0, x_1, \ldots)$. Then there is a k such that $x_j = x$ for all $j \geq k$. It must be the case then that $x \in C_j$ for all $j \geq k$. Therefore $r_j(x_j) = r_j(x) = x = x_{j-1}$ for all $j > k$. And, if $j \leq k$, then

$$s_j(x) \;=\; r_j \circ \cdots \circ r_k(x) \;=\; r_j \circ \cdots \circ r_k(x_k) \;=\; r_j(x_j) \;=\; x_{j-1}.$$

Suppose $\overline{x} \notin \alpha(X)$. Let $\overline{x} = (x_0, x_1, \ldots)$. We seek to show that $r_i(x_i) = x_{i-1}$ or all $i > 0$. Note that $x_i \in C_i$ for all i and $\alpha(x_i) = (x_0, x_1, \ldots, x_{i-1}, x_i, x_i, x_i, \ldots)$. Clearly the sequence $\{\alpha(x_i)\}_{i=0}^\infty$ converges to $\overline{x}$. Now, each $\alpha(x_i)$ belongs to $\varprojlim\{C_i, r_i\}$ by the first part of the proof. Thus, $r_i(x_i) = x_{i-1}$ for all $i \geq 1$ as desired.

For $\overline{x} = (x_0, x_1, \ldots) \in \varprojlim\{C_i, r_i\}$, we again consider the sequence $\{\alpha(x_i)\}_{i=0}^\infty$. This sequence lives in $\alpha(X)$ and converges to $\overline{x}$. □

Let $R = \overline{\alpha(X)} - \alpha(X)$. We define $\overline{X} = X \sqcup R$ and let $\overline{\alpha} : \overline{X} \to \overline{\alpha(X)}$ be defined by

$$\overline{\alpha}(x) = \begin{cases} \alpha(x), & x \in X \\ x, & x \in R \end{cases}$$

We declare a set $U \subseteq \overline{X}$ to be open if and only if $\overline{\alpha}(U)$ is open in $\overline{\alpha(X)}$. Clearly, $\overline{X} \approx \overline{\alpha(X)}$.

Definition 4.13 (The Cubical Boundary)**.** For a locally-finite CAT(0) cube complex X, we denote the remainder space $\overline{X} - X$ as $\partial_{\square} X$ and call it the **cubical boundary** of X.

For a finite CAT(0) cube complex X we define $\partial_{\square} X$ as the empty set.

We have thus far shown that $\overline{X}$ is a compactification of X. In the following paragraphs, we will show that $\partial_{\square} X$ is a Z-set. This will establish that $\overline{X}$ is in fact a Z-compactification of X. To this end, we construct the following ladder diagram.

$$\begin{tikzcd}
C_0 \times I \arrow[d, "H_0"'] & C_1 \times I \arrow[l, "r_1 \times id"'] \arrow[d, "H_1"'] & C_2 \times I \arrow[l, "r_2 \times id"'] \arrow[d, "H_2"'] & C_3 \times I \arrow[l, "r_3 \times id"'] \arrow[d, "H_3"'] & \cdots \arrow[l, "r_4 \times id"'] \\
C_0 & C_1 \arrow[l, "r_1"'] & C_2 \arrow[l, "r_2"'] & C_3 \arrow[l, "r_3"'] & \cdots \arrow[l, "r_4"']
\end{tikzcd}$$

Consider the top inverse sequence with bonding maps $r_i \times id : C_i \times I \to C_{i-1} \times I$. Since each bonding map is the identity on the I factor, we can write

$$\begin{aligned} \varprojlim\{C_i \times I, r_i \times id\} &= \varprojlim\{C_i, r_i\} \times \varprojlim\{I, id\} \\ &= \varprojlim\{C_i, r_i\} \times I \end{aligned}$$

Our goal is to define maps H_i so that the squares in the above diagram commute. In fact, by the way we will define the maps H_i, they will be contractions of the C_i to the basepoint v_0.

To get started, let H_0 be the constant map that maps $C_0 \times I$ to v_0. Then, inductively define

$$H_i(x,t) = \begin{cases} x, & 0 \le t \le 1/2^i \\ \phi_i(x, 2^i t - 1), & 1/2^i \le t \le 1/2^{i-1} \\ H_{i-1}(r_i(x), t), & 1/2^{i-1} \le t \le 1 \end{cases}$$

First, consider the square

$$\begin{tikzcd} C_0 \times I \arrow[d, "H_0"'] & C_1 \times I \arrow[l, "r_1 \times id"'] \arrow[d, "H_1"] \\ C_0 & C_1 \arrow[l, "r_1"'] \end{tikzcd}$$

To write down the map H_1 explicitly, it is

$$H_1(x,t) = \begin{cases} x, & 0 \le t \le 1/2 \\ \phi_1(x, 2t-1), & 1/2 \le t \le 1 \end{cases}$$

Let us show that $r_1 \circ H_1 = H_0 \circ (r_1 \times id)$. To this end, let $(x,t) \in C_1 \times I$ and suppose $t \le 1/2$. Then

$$\begin{aligned} r_1 \circ H_1(x,t) &= r_1(x) \\ &= v_0 \\ &= H_0(v_0, t) \\ &= H_0(r_1(x), t) \\ &= H_0 \circ (r_1 \times id)(x,t) \end{aligned}$$

If $t \ge 1/2$ then

$$\begin{aligned} r_1 \circ H_1(x,t) &= r_1(\phi(x, 2t-1) \\ &= r_1(x) \qquad (by\ the\ track - faithful\ property) \\ &= H_0(r_1(x), t) \\ &= H_0 \circ (r_1 \times id)(x,t) \end{aligned}$$

Now, let us consider the ith square of the above diagram for $i > 1$.

$$\begin{array}{ccc} C_{i-1} \times I & \xleftarrow{r_i \times id} & C_i \times I \\ \big\downarrow{\scriptstyle H_{i-1}} & & \big\downarrow{\scriptstyle H_i} \\ C_{i-1} & \xleftarrow{r_i} & C_i \end{array}$$

Our goal is to show that $r_i \circ H_i = H_{i-1} \circ (r_i \times id)$. So, let $(x, t) \in C_i \times I$. By track-faithfulness, $r_i(\phi_i(x, t)) = r_i(x)$ for all $t \in [0, 1]$. Thus

$$r_i \circ H_i(x,t) = \begin{cases} r_i(x), & 0 \le t \le 1/2^{i-1} \\ r_i(H_{i-1}(r_i(x), t)), & 1/2^{i-1} \le t \le 1 \end{cases}$$

Because r_i is a retraction fixing C_{i-1}, we can simplify further.

$$r_i \circ H_i(x,t) = \begin{cases} r_i(x), & 0 \le t \le 1/2^{i-1} \\ H_{i-1}(r_i(x), t), & 1/2^{i-1} \le t \le 1 \end{cases}$$

On the other hand, we have

$$H_{i-1} \circ (r_i \times id)(x,t) = \begin{cases} r_i(x), & 0 \le t \le 1/2^{i-1} \\ \phi_{i-1}(r_i(x), 2^{i-1}t - 1), & 1/2^{i-1} \le t \le 1/2^{i-2} \\ H_{i-2}(r_{i-1}(r_i(x)), t), & 1/2^{i-2} \le t \le 1 \end{cases}$$

Comparing this with the definition of $H_{i-1}(r_i(x), t)$, it is clear that $r_i \circ H_i(x, t) = H_{i-1} \circ (r_i \times id)(x, t)$. The preceeding paragraphs allow us to state the following result.

Theorem 4.14. *For a locally-finite CAT(0) cube complex X, the cubical boundary $\partial_\square X = \overline{X} - X$ is a Z-set in $\overline{X}$.*

Proof. We will identify X with $\alpha(X)$, where α is the embedding in Lemma 4.11.

Since $r_i \circ H_i(x, t) = H_{i-1} \circ (r_i \times id)(x, t)$ for all $i \ge 1$ there is a well-defined map

$$H : \varprojlim\{C_i, r_i\} \times I \to \varprojlim\{C_i, r_i\},$$

which is the restriction of the product map $H_1 \times H_2 \times \cdots$ to points in $\overline{X}$. For $t > 0$, choose j such that $1/2^j \le t < 1/2^{j-1}$ and let $\overline{x} \in \overline{X}$. Then

$$H(\overline{x}, t) = (H_0(x_0), H_1(x_1), \ldots)$$

and, for $i \geq j$,

$$H_i(x_i, t) = H_i(r_{i+1}(x_{i+1}), t) = H_{i+1}(x_{i+1}, t)$$

Thus, $H(\overline{x}, t) \in X$. $\square$

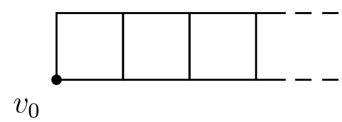

Figure 46: Square complex $r \times J$.

Example 4.15. Let r be a cubulated ray, and let X be the square complex $r \times J$. See figure 46. $\partial_{\square} X$ is shape equivalent to $\partial_\infty X$, which is a singleton.

Using bonding maps obtained from Example 4.1, we obtain a different compactification of X. Consider the point labeled p (see Figure 47). The path in red shows part of the preimage of p using bonding maps obtained from elementary cubical collapses. The rest of the preimage of p is contained in the "top" ray. These two pieces converge to different points in the compactification. This example shows that the identity map does not extend to a homeomorphism between this boundary and the cubical boundary.

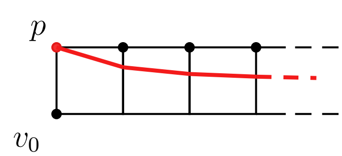

Figure 47

4.3.1 Well-Definedness of the Cubical Boundary

This section is devoted to showing that the cubical boundary is independent of the enumeration of hyperplanes of X so long as the chosen enumeration is distance-respecting.

The following corollary establishes that, given two distinct distance-respecting enumerations of $\hat{\mathcal{H}}$, the corresponding cubically-expanded compacta will agree at regular steps in the process. It follows immediately from Corollary 3.19.

Corollary 4.16. *Given two distance-respecting enumerations $\hat{\mathcal{H}}_A = \{\hat{a}_1, \hat{a}_2, \ldots\}$ and $\hat{\mathcal{H}}_B = \{\hat{b}_1, \hat{b}_2, \ldots\}$ of the hyperplanes of (X, v_0), let $\{A_i\}$ and $\{B_i\}$ be cubically expanded compacta corresponding to $\hat{\mathcal{H}}_A$ and $\hat{\mathcal{H}}_B$, respectively. For each $i = 1, 2, \ldots$, let $\hat{a}_{k_i}, \hat{b}_{k_i}$ denote the last hyperplanes appearing in their respective enumerations that are a distance $2i-1$ from v_0 (in the CAT(0) metric). We claim that $A_{k_i} = B_{k_i}$.*

Proof. A_{k_i} and B_{k_i} are both full subcomplexes of the same vertex set. □

From now on, for the two different sets of compacta $\{A_i\}$ and $\{B_i\}$ corresponding to distance-respecting enumerations $\hat{\mathcal{H}}_A$ and $\hat{\mathcal{H}}_B$ of the hyperplanes of X, we will let D_i denote the common compacta $A_{k_i} = B_{k_i}$ for which $\hat{a}_{k_i}, \hat{b}_{k_i}$ denote the last hyperplanes appearing in their respective enumerations that are a distance $2i - 1$ from v_0.

Definition 4.17 (Truncated Metric)**.** For a metric $\rho : X \times X \to \mathbb{R}$, we define the following function.

$$\overline{\rho}(x, y) = \min\{\rho(x, y), 1\}$$

Verifying that $\overline{\rho}$ defines a metric is left to the reader. We call $\overline{\rho}$ the **truncated metric** and refer to the distance between points x and y in the truncated metric as the **truncated distance** between them.

Lemma 4.18. *Let $\hat{h}_i \in \hat{\mathcal{H}}$, $x \in \mathcal{C}(\hat{h}_i)$, and let ϕ_i denote the interval collapse w.r.t. $\hat{h}_i$. Then the truncated distance between x and any hyperplane besides $\hat{h}_i$ remains unchanged throughout ϕ_i.*

Proof. This follows immediately from the fact that, in a cube, all points on the product line $\{x\} \times J$ are equidistant (in the truncated metric) from all hyperplanes except precisely one, the hyperplane $\hat{h}_i$. □

Remark. Let c be a cube that meets $\hat{h}_i$ in the midcube m_i. According to Lemma 4.18, the local coordinate corresponding to m_i is the only coordinate in which $\phi_i(x, t)$ takes different values for different times t.

Lemma 4.19. *For (X, v_0), a pointed locally finite CAT(0) cube complex, let $\hat{\mathcal{H}}_A$ and $\hat{\mathcal{H}}_B$ denote two distance-respecting enumerations of the hyperplanes of X. Let D_{i+1} and D_i be as above. Let r_A and r_B denote the compositions of retractions that takes D_{i+1} onto D_i in the enumerations $\hat{\mathcal{H}}_A$ and $\hat{\mathcal{H}}_B$, respectively. Then for all $x \in D_{i+1}$, $r_A(x) = r_B(x)$.*

Proof. If $x \in D_i$, then $r_A(x) = x = r_B(x)$.

Let $x \in D_{i+1} \setminus D_i$, and let $c = [-1, 1]^n$ be a maximal cube in D_{i+1} that contains $x = (x_1, \dots, x_n)$. Then $c \cap D_i$ is a face $f = [-1, 1]^m$ of c, where $m < n$. Without loss of generality, we identify f with the first m coordinates of c. We will show that $r_A(x) = (x_1, \dots, x_m, 1, \dots, 1)$. Each coordinate of x corresponds to a hyperplane, and the last $m - n$ coordinates correspond to hyperplanes that are collapsed via the maps ϕ_i. Thus r_A maps each such coordinate x_i $(i > m)$ to 1. Furthermore, since $f \in D_i$, the first m coordinates of x correspond to hyperplanes that still remain in D_i after collapsing D_{i+1} to D_i. Thus, these coordinates are fixed by r_A. This gives the desired result.

By the same argument, r_B maps x to $(x_1, \dots, x_m, 1, \dots, 1)$. Hence $r_A = r_B$. $\square$

Let $r_i : A_i \to A_{i-1}$ and $s_i : B_i \to B_{i-1}$ denote the retractions obtained from interval collapses (Definition 4.2). We have, a priori, two cubical compactifications of X given by $\varprojlim\{A_i, r_i\}$ and $\varprojlim\{B_i, s_i\}$. Let us also write $\overline{X}_A = X \sqcup R_A$ and $\overline{X}_B = X \sqcup R_B$, where $R_A = \varprojlim\{A_i, r_i\} \setminus \alpha_A(X)$ and $R_B = \varprojlim\{B_i, s_i\} \setminus \alpha_B(X)$. We will show that $\overline{X}_A$ and $\overline{X}_B$ are equivalent in a strong sense by demonstrating a homeomorphism between them that restricts to the identity map on X.

It follows from Lemma 4.19 that each inverse sequence

$$\{v_0\} = A_0 \xleftarrow{r_1} A_1 \xleftarrow{r_2} A_2 \xleftarrow{r_3} \cdots \qquad \& \qquad \{v_0\} = B_0 \xleftarrow{s_1} B_1 \xleftarrow{s_2} B_2 \xleftarrow{s_3} \cdots$$

contains

$$\{v_0\} = D_0 \xleftarrow{t_1} D_1 \xleftarrow{t_2} D_2 \xleftarrow{t_3} \cdots$$

as a subsequence, where the bonding maps t_i equal compositions r_{ij} for i, j integers corresponding to hyperplanes $\hat{h}_i$ and $\hat{h}_j$ which are the last hyperplanes appearing in the enumeration $\hat{\mathcal{H}}_A$ that are a distance $2k-1$ and $2k+1$, respectively, from the basepoint for some k. It is not hard to check that the map $\theta_A : \varprojlim\{A_i, r_i\} \to \varprojlim\{D_i, t_i\}$ defined by forgetting the coordinates corresponding to $A_i \neq D_i$ is a homeomorphism. We leave the details to the reader. The inverse of this map inserts coordinates in a way consistent with the bonding maps r_i. Define $\theta_B : \varprojlim\{B_i, s_i\} \to \varprojlim\{D_i, t_i\}$ similarly. Let $\alpha_A : X \to \Pi A_i$ and $\alpha_B : X \to \Pi B_i$ denote the embeddings from Lemma 4.11, and let $\overline{\alpha}_A : \overline{X}_A \to \overline{\alpha_A(X)}$ and $\overline{\alpha}_B : \overline{X}_B \to \overline{\alpha_B(X)}$ be the obvious extensions of these maps to $\overline{X}_A$ and $\overline{X}_B$, respectively.

Corollary 4.20. *For a locally-finite CAT(0) cube complex (X, v_0), let $\hat{\mathcal{H}}_A$, $\hat{\mathcal{H}}_B$, and $\{A_i\}$, $\{B_i\}$ be as above. Then the identity map extends to a homeomorphism between $\varprojlim\{A_i, r_i\}$ and $\varprojlim\{B_i, s_i\}$, where r_i and s_i are the retractions obtained via interval collapses.*

Proof. The map

$$\overline{\alpha}_B^{-1} \circ \theta_B^{-1} \circ \theta_A \circ \overline{\alpha}_A : \overline{X}_A \to \overline{X}_B$$

is a composition of homeomorphisms and thus a homeomorphism. For a point $x \in X$, let $\alpha_A(x) = (v_0, x_1, x_2, \ldots, x, x, \ldots)$ and let k_i denote the integer corresponding to the hyperplane $\hat{h}_{k_i}$ appearing last among all hyperplanes in $\hat{\mathcal{H}}_A$ a distance $2k_i - 1$ from the v_0. Then

$$\begin{aligned}
\overline{\alpha}_B^{-1} \circ \theta_B^{-1} \circ \theta_A \circ \overline{\alpha}_A(x) &= \overline{\alpha}_B^{-1} \circ \theta_B^{-1} \circ \theta_A \circ \alpha_A(x) \\
&= \overline{\alpha}_B^{-1} \circ \theta_B^{-1} \circ \theta_A(v_0, x_1, x_2, \ldots, x, x, \ldots) \\
&= \overline{\alpha}_B^{-1} \circ \theta_B^{-1}(v_0, x_{k_1}, x_{k_2}, \ldots, x, x, \ldots) \\
&= \alpha_B^{-1}(v_0, \ldots, s_{k_1}(x_{k_1}), x_{k_1}, \ldots, s_{k_2}(x_{k_2}), x_{k_2}, \ldots, x, x, \ldots) \; = \; x
\end{aligned}$$

where $\overline{\alpha}_B^{-1}$ recognizes $(v_0, \ldots, s_{k_1}(x_{k_1}), x_{k_1}, \ldots, s_{k_2}(x_{k_2}), x_{k_2}, \ldots, x, x, \ldots)$ as the image of a point in X. $\square$

4.3.2 Cubical Boundary and Visual Boundary

In [Gro87], Gromov introduced a new boundary of geodesic δ-hyperbolic metric spaces, which we will refer to as the *Gromov boundary.* For such a space, (X, d), the Gromov boundary is constructed by considering the family of geodesic rays emanating from a basepoint, o. An equivalence relation of placed on this family of rays whereby two rays γ_1 and γ_2 are considered equivalent if

$$d(\gamma_1(t), \gamma_2(t)) \leq K$$

for some bound K and all times t. The **Gromov boundary** of X is the collection of equivalence classes of geodesic rays emanating from o and is denoted $\partial_G X$. One can define a basis for the topology on $\partial_G X$ using the *gromov product*, which measures how, in a precise way, how long geodesic rays stay near each other before diverging. One can find the relevant definitions and details on the Gromov boundary in [KB02].

A closely related boundary exists for proper CAT(0) spaces, which is called the *visual boundary* and denoted $\partial_\infty X$, for an appropriate space X. As with $\partial_G X$, the visual boundary consists of equivalence classes of geodesic rays emanating from a basepoint o. However, the topology on $\partial_\infty X$ is instead given by the cone topology, a detailed description of which can be found chapter 2, section 8 of [BH99].

Of particular relevance, attaching the visual boundary to a CAT(0) cube complex in the appropriate manner provides a Z-compactification with the visual boundary being a Z-set. Although a CAT(0) cube complex may have different Z-boundaries, all these boundaries are *shape equivalent.* For details, see [BM91], [GM19], and [Gui16].

5 OPEN QUESTIONS

Question 1. Given a locally-finite CAT(0) cube complex X, are the cubical boundaries of X and its barycentric cubing X' the same?

Question 2. Are the locally-connected properties of the cubical boundary significantly different than the locally-connected properties of the visual boundary?

Question 3. In [Cor15], Cordes shows that the Morse boundary satisfies a certain uniqueness condition. Namely, if X and Y are quasi-isometric proper geodesic metric spaces then their Morse boundaries are homeomorphic. Taking X to be the cubified line and Y to be $X \times J$ immediately shows that this statement is false for the cubical boundary. However, the question of when two cubical boundaries are homeomorphic remains.

Question 4. What can be said of the other boundaries obtained by using bonding maps obtained from the deformation retractions described in Example 4.1 on page 69? These boundaries are not necessarily equivalent. In particular, the identity map will not always extend to a homeomorphism. But, we can ask whether these boundaries are homeomorphic.

Question 5. Can the process used to construct $\overline{X}$ be applied in more general settings. In particular, can the simply-connected assumption be dropped? If the "holes" of X are contained in a compact subcomplex A, then perhaps there is a way to make sense of cubically compactifying X by somehow ignoring A.

Alternatively, can the flag condition on links be relaxed? Example 2.3.4 on page 32 shows a cube complex that is not CAT(0), but for which cubically expanded compacta C_i can be defined, as well as corresponding bonding maps.

Question 6. What relationships exist between the cubical boundary and other boundaries? Such as the contracting boundary, Poisson boundary, etc.

REFERENCES

[AB19] Karim Adiprasito and Bruno Benedetti. Collapsibility of cat(0) spaces, 2019.

[AGM12] Ian Agol, Daniel Groves, and Jason Manning. The virtual haken conjecture, 2012.

[And67] R. David Anderson. On topological infinite deficiency. *Michigan Mathematical Journal*, 14:365–383, 1967.

[Bes96] Mladen Bestvina. Local homology properties of boundaries of groups. *Michigan Mathematical Journal*, 43(1):123 – 139, 1996.

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

[BL14] ore Baralić and Ioana-Claudia Lazar. A note on the combinatorial structure of finite and locally finite simplicial complexes of nonpositive curvature. 03 2014.

[BL20] Tudor Bînzar and Ioan Lazăr. The collapsibility of cat(0) square 2-complexes. *Publicationes Mathematicae Debrecen*, 97:1–25, 2020.

[BM91] Mladen Bestvina and Geoffrey Mess. The boundary of negatively curved groups. *Journal of the American Mathematical Society*, 4:469–481, 1991.

[CFI16] Indira Chatterji, Talia Fernós, and Alessandra Iozzi. The median class and superrigidity of actions on cat(0) cube complexes. *Journal of Topology*, 9(2):349–400, 2016.

[Coh73] Marshall M. Cohen. *A course in simple-homotopy theory.* Springer-Verlag, New York-Berlin, 1973. Graduate Texts in Mathematics, Vol. 10.

[Cor15] Matthew Cordes. Morse boundaries of proper geodesic metric spaces. *Groups, Geometry, and Dynamics*, 11, 02 2015.

[Dav08] Michael W. Davis. *The Geometry and Topology of Coxeter Groups. (LMS-32)*. Princeton University Press, 2008.

[Far16] Daniel Farley. A proof of sageev's theorem on hyperplanes in cat(0) cubical complexes. In Michael W. Davis, James Fowler, Jean-François Lafont, and Ian J. Leary, editors, *Topology and Geometric Group Theory*, pages 127–142, Cham, 2016. Springer International Publishing.

[GM19] Craig R. Guilbault and Molly A. Moran. Proper homotopy types and z-boundaries of spaces admitting geometric group actions. *Expositiones Mathematicae*, 37(3):292–313, 2019.

[Gro87] M. Gromov. Hyperbolic groups. In S. Gersten, editor, *Essays in Group Theory*, volume 8 of *MSRI Publications*, pages 75–263. Springer, New York, 1987.

[Gui13] Craig R. Guilbault. Weak z -structures for some classes of groups. *Algebraic & Geometric Topology*, 14:1123–1152, 2013.

[Gui16] Craig R. Guilbault. Ends, shapes, and boundaries in manifold topology and geometric group theory. In Michael W. Davis, James Fowler, Jean-François Lafont, and Ian J. Leary, editors, *Topology and Geometric Group Theory*, pages 45–125, Cham, 2016. Springer International Publishing.

[Hag07] Frédéric Haglund. Isometries of cat(0) cube complexes are semi-simple. *Annales mathématiques du Québec*, 2007.

[KB02] Ilya Kapovich and Nadia Benakli. *Boundaries of hyperbolic groups*, pages 39–93. 03 2002.

[Lea13] Ian J. Leary. A metric kan–thurston theorem. *Journal of Topology*, 6(1):251–284, 2013.

[LN19] Davide Lofano and Andrew Newman. The worst way to collapse a simplex. *Israel Journal of Mathematics*, 244:625 – 647, 2019.

[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proceedings of the London Mathematical Society*, s3-71(3):585–617, 1995.

[Whi50] John Henry Constantine Whitehead. Simple homotopy types. *American Journal of Mathematics*, 72:1, 1950.