THE CAT(0) GEOMETRY OF COXETER GROUPS AND ARTIN GROUPS

KATHERINE M. GOLDMAN

ABSTRACT. Our goal is to discuss CAT(0) spaces and complexes of groups, with heavy emphasis on Coxeter groups and Artin groups as the main motivation and examples.

These notes are an ongoing work in progress and will be continually updated throughout the minicourse, and possibly after.

LAST UPDATED. May 22, 2023

CONTENTS

1. A brief introduction

The main focus of this minicourse is to present results about Coxeter groups and Artin groups attained through the lens of CAT(0) geometry. This is not solely a course on these groups, nor is it comprehensive. Much of the focus will be on the theory behind CAT(0) spaces and complexes of groups, two tools which have been invaluable in the study of these classes of groups. Not only will these prove to be useful in the study of Coxeter groups and Artin groups, but the groups themselves provide interesting and non-trivial examples of CAT(0) spaces and complexes of groups. We're going to jump around these topics rather than group them all together in order to try to illustrate these connections.

Let's start with the definition of Coxeter groups and Artin groups, since these are certainly the easiest things we'll define for a while. There are many equivalent definitions of a Coxeter group. The simplest is to start with a finite set I, then form symbols $S = \{s_i\}_{i\in I}$, which give rise to a group with the presentation

$$W = \langle\ r_i \mid (r_i r_j)^{m_{ij}} = 1\ \rangle,$$

where $1 \leq m_{ij} \leq \infty$, $m_{ij} = m_{ji}$, and $1 = m_{ij}$ if and only if $i = j$. If $m_{ij} = \infty$ then we exclude the corresponding relation. We will call the pair (W, S) a *Coxeter system.* (The choice of generators is important!)

Typically the m_{ij} are arranged in either a *Coxeter matrix* or *Coxeter diagram.* A Coxeter matrix M is just a symmetric integral matrix $(m_{ij})_{i,j\in I}$ with 1s along the diagonal, and where each off-diagonal entry is between 2 and ∞ (inclusive). A Coxeter diagram Γ is a simplicial graph (no loops or double edges) with vertex set I whose edges $\{i, j\}$ are labeled with an integer $3 \leq m_{ij} \leq \infty$. If $m_{ij} = 3$ we often omit the label. Two vertices not joined by an edge represent $m_{ij} = 2$, and we always take $m_{ii} = 1$.

Given a Coxeter matrix M or diagram Γ, we often denote the corresponding Coxeter group by W_M or W_Γ, in which case the distinguished generators S are to be taken implicitly. These various expressions of Coxeter groups have their own utility, but typically the diagrams will be most useful for visualizing the structure of the so-called "special subgroups", which become important later.

The prototypical example of a Coxeter group is the symmetric group on n letters, S_n. The Coxeter generators are given by the transpositions $(i, i+1)$. For an infinite example, see the group of reflection symmetries of the equilateral triangle tiling of $\mathbb{R}^2$ (aka the "3-3-3 triangle group").

Artin groups appear very similar to Coxeter groups on a surface level. Start with your favorite Coxeter matrix M or Coxeter diagram Γ, but then form a group with the presentation

$$A = A_M = A_\Gamma \coloneqq \langle\ s_i \mid \underbrace{s_i s_j s_i \ldots}_{m_{ij}\text{ letters}} = \underbrace{s_j s_i s_j \ldots}_{m_{ij}\text{ letters}},\ i \neq j\ \rangle.$$

As with Coxeter groups, we call the pair (A, S) an *Artin system.*

One can easily verify that by adding in the relations $s_i^2 = 1$, we attain the corresponding Coxeter group W_Γ. We let $\pi = \pi_\Gamma = \pi_M : A_\Gamma \to W_\Gamma$ denote the homomorphism sending s_i to r_i. (In fact, there is a set-theoretic section $\sigma : W \to A$ sending r_i to s_i, but the existence of this section is non-trivial.)

The prototypical example of an Artin group is the Braid group on n-strands Br_n, which has diagram

The corresponding Artin group is S_n.

While we know a great deal about Coxeter groups, and the presentation of Artin groups are very similar to that of Coxeter groups, surprisingly, we know nearly nothing about Artin groups. The word problem, conjugacy problem, and almost anything else you might think to ask about Artin groups are unknown, except in special cases. This can be attributed to the fact that Coxeter groups in general possess deep connections to well-studied geometry, while the geometry of Artin groups remains mysterious in general. Much of this course will focus on said geometry, which, as we've gestured toward, relies heavily on complexes of groups and CAT(0) spaces.

2. Constant curvature Riemannian manifolds

We begin the more formal part of the survey with some prerequisite material on the simply connected Riemannian manifolds of constant curvature κ. We won't need details regarding the Riemannian structure, we'll mostly just cover some of the interesting theorems. This section mostly follows Bridson and Haefliger, Ch I.2

Definition 1. We let M^n_κ denote the (unique!) simply connected Riemannian n-manifold of constant curvature κ. When $n = 2$, we will sometimes call M^2_κ the κ-model plane.

There are three broad classes of these manifolds, which we can use to give simple descriptions of these spaces solely as metric spaces. Namely, these can be grouped into those manifolds with κ negative, positive, and equal to 0. In fact, each manifold in these respective classes are simply scalings of the spaces M^n_{-1}, M^n_1, and M^n_0, respectively. So it suffices to describe these spaces. We'll discuss the geodesics, angles, and hyperplanes, as these become important pretty soon (particularly the geodesics).

2.1. Euclidean space ($\kappa = 0$).

We let $\mathbb{E}^n$ denote a fixed n-dimensional real vector space on basis $\{e_i\}$ equipped with the standard Euclidean inner product

$$(x, y) \coloneqq \sum_{i=1}^{n} x_i y_i,$$

where $x = x_1e_1 + \cdots + x_ne_n$ and $y = y_1e_1 + \cdots + y_ne_n$. This of course comes with the standard norm $\|x\|^2 = (x, x)$ and Euclidean distance $d_E(x, y) = \|x - y\|$. Endowed with this metric, $\mathbb{E}^n$ is (isometric to) our model space M^n_0.

There isn't much that you probably don't know about $\mathbb{E}^n$. The geodesics are affine lines, angles are usual angles induced by the inner product, and hyperplanes are affine subspaces of codimension 1.

2.2. The sphere ($\kappa > 0$).

The n-sphere $\mathbb{S}^n$ is the usual set of unit vectors in $\mathbb{R}^{n+1}$, namely

$$\left\{ x = (x_1, \dots, x_{n+1}) \in \mathbb{R}^{n+1} : (x, x) = 1 \right\}.$$

However the usual metric is of course not induced from $\mathbb{R}^{n+1}$. Rather, it's the function $d_S : \mathbb{S}^n \times \mathbb{S}^n \to [0, \pi]$ uniquely defined by

$$\cos d_S(A, B) = (A, B).$$

(Check BH for details on why this is a metric.) Then the sphere with this metric gives us the model space M_1^n.

We can view $\mathbb{S}^n \subseteq \mathbb{R}^{n+1}$ to get a quick description of the geodesics and hyperplanes. The geodesics are the intersection of $\mathbb{S}^n$ with 2-dimensional linear subspaces of $\mathbb{R}^{n+1}$, and the hyperplanes are intersection of $\mathbb{S}^n$ with codimension-1 linear subspaces of $\mathbb{R}^{n+1}$.

Angles are slightly trickier to define. Let $\gamma_i = V_i \cap \mathbb{S}^n$ with $V_i \subseteq \mathbb{R}^{n+1}$ a 2-dimensional linear subspace for $i = 1, 2$. Then $V_1 \cap V_2 = \mathbb{R}u$ for some unit vector u. Pick a $u_i \in V_i$ orthogonal to u. Then the angle between the geodesics γ_1 and γ_2 is the $\alpha \in [0, \pi]$ satisfying

$$\cos \alpha = (\pm u_1, \pm u_2),$$

where the sign gives the different supplementary angles.

2.3. **Hyperbolic space ($\kappa < 0$).** In Riemannian geometry, one usually defines hyperbolic space via the disk/ball/Poincaré model, but for our purposes (and to give a more unified presentation), we're going to use the hyperboloid/Minkowski model.

We let $\mathbb{E}^{n,1}$ denote a fixed n-dimensional real vector space on basis $\{v_i\}$ endowed with the symmetric bilinear form given by

$$\langle x, y \rangle = \sum_{i=1}^{n} x_i y_i - x_{n+1} y_{n+1},$$

the standard non-degenerate bilinear form of type $(n, 1)$ (hence the naming convention). Then we define hyperbolic n-space to be the set

$$\mathbb{H}^n = \left\{ x \in \mathbb{E}^{n,1} : \langle x, x \rangle = -1, x_{n+1} > 0 \right\}.$$

Why: The equation $\langle x, x \rangle = -1$ defines a hyperboloid of two sheets in $\mathbb{R}^{n+1}$, so requiring $x_{n+1} > 0$ gives us a single sheet.

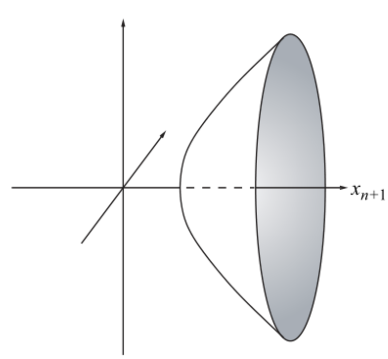

Now the metric can be defined similarly to the spherical case. Namely, we let $d_H : \mathbb{H}^n \times \mathbb{H}^n \to [0, \infty)$ denote the function uniquely defined by

$$\cosh d_H(A, B) = -\langle A, B \rangle.$$

(Again, see BH for details.) Then $\mathbb{H}^n$ with this metric gives us the model space M_{-1}^n.

Geodesics are also defined similarly to the spherical case. If V is a 2-dimensional vector subspace of $\mathbb{E}^{n,1}$ (and intersects $\mathbb{H}^n$ non-trivially), then $V \cap \mathbb{H}^n$ is a geodesic, and all geodesics of $\mathbb{H}^n$ arise in this way. Angles and hyperplanes are defined completely analogously as well.

2.4. **The model spaces.** Now, we can define the spaces M^n_κ in terms of the above metric spaces. In fact, they're just scalings of these spaces.

Definition 2. For $\kappa \in \mathbb{R}$, we define M^n_κ to be

(1) Euclidean space $\mathbb{E}^n$ with the above metric if $\kappa = 0$
(2) Hyperbolic space $\mathbb{H}^n$ with the metric $\frac{1}{\sqrt{-\kappa}} d_H$ if $\kappa < 0$
(3) The sphere $\mathbb{S}^n$ with the metric $\frac{1}{\sqrt{\kappa}} d_S$ if $\kappa > 0$.

The following are easy consequences of the definitions.

Proposition 3.

(1) M^n_κ is a geodesic metric space.
(2) If $\kappa \leq 0$ then M^n_κ is uniquely geodesic and all balls are convex.
(3) If $\kappa > 0$ then there is a unique geodesic segment joining $x, y \in M^n_\kappa$ if and only if $d(x,y) < \pi/\sqrt{k}$, and closed balls of radius $< \pi/(2\sqrt{\kappa})$ are convex.

For a number of reasons it becomes useful to make the following definition.

Definition 4. For $\kappa \in \mathbb{R}$, we denote the *diameter of* M^n_κ by D_κ. Explicitly,

$$D_\kappa = \begin{cases} \frac{\pi}{\sqrt{\kappa}} & \kappa > 0 \\ \infty & \kappa \leq 0 \end{cases}$$

Then our previous proposition can be rephrased more succinctly as "M^n_κ is D_κ-uniquely geodesic."

3. Simplices in M_κ

Before moving on further, it will be useful to discuss the definition and some basic properties of simplices in the model spaces.

Definition 5. Let $\kappa \in \mathbb{R}$ and $0 \leq n \leq m$ be integers. An *n-plane* in M^m_κ is a subspace which is isometric to M^n_κ. We say that $n+1$ points of M^m_κ are in *general position* if there is no $(n-1)$-plane containing them. Then a *(geodesic) n-simplex* $\sigma \subseteq M^m_\kappa$ is the convex hull of $n+1$ points in general position. These points are called the *vertices* of σ. **Note:** If $\kappa > 0$, we require the points lie in an open ball of radius $D_\kappa/2$.

A subset $T \subseteq S$ is called a *face* if it's the convex hull of a collection of vertices of S. It's called a proper face if $T \neq S$. (Note that faces are also simplices.) The interior of S is the set of points which don't lie in a proper face.

It will become useful later to determine a criteria for such simplices to exist. While there are criteria for each of the model spaces, the simplest and most useful for us will be to examine the spherical case. The following comes from Davis, §6.8.

Let $\sigma \subseteq \mathbb{S}^n \subseteq \mathbb{R}^{n+1}$ be an n-simplex. There are $n+1$ codim-1 faces, which we index by $\{\sigma_0, \ldots, \sigma_n\}$. Let $u_i \in \mathbb{R}^{n+1}$ be the "inward-pointing" unit normal to σ_i.

This means σ and u_i lie on the same side of the hyperplane $(u_i, x) = 0$. Thus by definition,

$$\sigma = \bigcap_{i=0}^{n} \{ x \in \mathbb{S}^n : (u_i, x) \geq 0 \} .$$

The dihedral angle θ_{ij} of σ between σ_i and σ_j is then

$$\theta_{ij} = \pi - \cos^{-1}\big((u_i, u_j)\big) = \cos^{-1}(-(u_i, u_j)).$$

Proposition 6. *Suppose we are given numbers $0 < \theta_{ij} = \theta_{ji} < \pi$ when $0 \leq i < j \leq n$ and $\theta_{ii} = \pi$. Then there exists a simplex $\sigma \subseteq \mathbb{S}^n$ with dihedral angles θ_{ij} if and only if the matrix $(-\cos\theta_{ij})$ is positive definite.*

Proof. If σ is a simplex then $-\cos(\theta_{ij}) = (u_i, u_j)$ and thus the matrix is positive definite.

Conversely, suppose $A = (-\cos\theta_{ij})$ is positive definite. Then since A is also symmetric, we can find a non-singular U such that $U^tU = A$ (e.g. the square root of A). Let $u_0, \ldots, u_n$ be the column vectors of U. Note these are linearly independent. Since the diagonal entries of A are 1, the u_i are unit vectors. The half-spaces $(x, u_i) \geq 0$ have non-empty intersection (as a consequence of linear independence). Thus their intersection defines a spherical simplex with dihedral angles θ_{ij}. □

4. Geometric reflection groups

To make the standard geometry of Coxeter groups immediately clear, let's talk about geometric reflection groups while these constant curvature spaces are still in our head. By a geometric reflection group, we mean a discrete subgroup of $\mathrm{Isom}(M^n_\kappa)$ for some n and some κ which is generated by reflections. But what is a reflection in M^n_κ?

Definition 7. A reflection $r \in \mathrm{Isom}(M^n_\kappa)$ is an isometry of order 2 ($r^2 = \mathrm{id}$) whose fixed point set, denoted M^r, is a hyperplane (as defined above). We sometimes call M^r the *wall* of r. We then call $W = \langle S \rangle \leq \mathrm{Isom}(M^n_\kappa)$ a *(geometric) reflection group* if S is a (necessarily finite) set of reflections and W is discrete in $\mathrm{Isom}(M^n_\kappa)$.

Note that each of these isometry groups are Lie groups[1] (hence the notion of a discrete subgroup is well-defined).

Example 8. The group generated by reflections about the faces of a right-angled n-gon in $\mathbb{H}^2$ is a geometric reflection group.

We have the following interesting theorem, which we won't prove (yet, but might later?).

Theorem 9. *Every geometric reflection group is a Coxeter group.*

The converse is not true in general. However, it is known exactly when a given Coxeter group gives a geometric reflection group, and on what type of space (negative, positive, or zero curvature). So we're not overwhelmed at first, let's focus just on the case of reflection groups in $\mathbb{S}^n$, since this turns out to be the most important kind in general anyway.

[1]In fact, they're each generated by reflections. However, of course the isometry groups are not themselves discrete, and thus are not referred to as "reflection groups". (And there are certainly isometries which aren't reflections.)

Since this space is compact, it's easy to see that any reflection group on it will be finite. Less obvious is the following

Proposition 10. *If W is a finite Coxeter group, then it admits a free, faithful action as a geometric reflection group on $\mathbb{S}^n$ for some n.*

In other words, a Coxeter group is finite if and only if it's a geometric reflection group on $\mathbb{S}^n$. We will prove this proposition, since it gives us a chance to introduce one of the main players in the study of Coxeter groups.

Definition 11. Let $M = (m_{ij})_{i,j \in I}$ be a Coxeter matrix with (W, S) the corresponding Coxeter system. Let $V = \mathbb{R}^I$ with basis $\{e_i\}_{i \in I}$. Define a symmetric bilinear form $B = B_M$ on V by

$$B(e_i, e_j) = -\cos(\pi/m_{ij}).$$

We then define *linear reflections* $r_i \in O(B_M) \subseteq \mathrm{GL}(V)$ by

$$r_i(x) = x - 2B(x, e_i)e_i.$$

Then, let $r : S \to \mathrm{GL}(V)$ be the map $r(s_i) = r_i$.

Proposition 12. *The map $r : S \to \mathrm{GL}(V)$ extends to an injective homomorphism $\rho : W \to \mathrm{GL}(V)$ mapping $s_i \mapsto r_i$. Moreover, $\rho(W)$ is a discrete subgroup of* $\mathrm{GL}(V)$.

This is called the *canonical representation* of W. It and its dual (which we might discuss later) are one of the central objects we'll discuss, in particular for finite Coxeter groups. Proving this theorem is definitely outside the scope of these notes. (It's not that hard, but it takes a lot more machinery.)

How do we use this to show finite Coxeter groups act on $\mathbb{S}^n$ as geometric reflection groups? It is a known fact that any finite-dimensional representation of a finite group preserves a positive definite form. In particular, this form gives a Euclidean metric. Since it preserves the positive definite form, it also preserves its set of unit vectors, and thus the unit sphere. Since it's finite, it's necessarily a discrete group.

There are a couple basic reasons why we care about this representation. One of the main ones is just that we can better visualize the structure of W. Another is that it allows us to define some basic building blocks that we'll rely on very heavily later. And one more won't reveal itself until probably the end of the course.

Let's describe the first two at the same time now.

Since ρ is a faithful representation of W, we'll often conflate w and $\rho(w)$ for $w \in W$. In particular, for $w \in W$, we let V^w denote the fixed set of $\rho(w)$.

Definition 13. Let (W, S) be a Coxeter system. The *set of reflections* for W is defined to be $R = \left\{ w^{-1}sw : s \in S, w \in W \right\}$, the conjugates of the generators S. Sometimes the elements of S are called the *basic reflections*.

This terminology comes from the fact that $\rho(W)$ contains elements which are not reflections, and in fact, the reflections are precisely the conjugates of the $\rho(s_i)$.

There are two perspectives we'll use now, which somewhat go hand in hand. First, the more classical viewpoint.

Definition 14. Consider a finite Coxeter group W acting on an $n+1$ dimensional vector space V with the above representation. We fix a simplicial cone $\mathcal{C}_0$, the intersection of the half-spaces

$$\{(x, e_i) \geq 0\}.$$

We form a simplicial decomposition of the sphere based on this cone. Namely, we examine the translates of $\mathcal{C} := \mathbb{S}^n \cap \mathcal{C}_0$. This is a spherical simplex (check!), which we call the *fundamental chamber of* W. It is part of a collection of so-called (open) "chambers" of $\mathbb{S}^n$; these are the connected components of $\mathbb{S}^n \setminus \bigcup_{r\in R} V^r$. The closure of any given chamber (referred to as simply a "chamber") has a decomposition into *facets*, corresponding to the intersection of said chamber with various walls V^r. Moreover, this induces a simplicial structure on the sphere, which we sometimes call the *Coxeter complex* $C = C(W,S)$ of W. We call a 0-dimensional facet a *vertex*, 1-dimensional an *edge*, 2-dim a *face*. It is clear from the definition of chamber and the W-action that W preserves this cell structure, and in particular preserves the set of chambers. We will later show that in fact W acts simply transitively on the set of chambers.

Example with Kaleidotile

So, this gives a classical simplicial complex. But there's another viewpoint, which seems identical at first, but actually turns out to be incredibly useful.

Definition 15. Let $x \in V$ be a generic point for the action of W (i.e., has trivial stabilizer, i.e., doesn't intersect any walls). Consider $\Sigma = \Sigma(W,S)$, the convex hull of the orbit Wx of x under W. Since W is finite, this gives us a polytope in V. The natural cell structure on this polytope is a cellulation of the ball whose boundary $\partial\Sigma$ is a cellulation of the sphere. We call Σ a *Coxeter polytope.*

Proposition 16. *The boundary $\partial\Sigma$ is dual to the Coxeter complex C.*

Kaleidotile again

Later, we will put a specific metric on Σ, but we note that the choice of base point does not change the combinatorial type of Σ.

4.1. **Tangent on Artin groups.** While we're discussing the canonical representation, now would actually be a great time to bring Artin groups into the picture as well. It turns out that in addition to the similar definition, their geometry (or at least, the geometry which is currently believed to be the most worth studying) comes from a similar place, although perhaps unexpectedly. The motivating fact is the following.

Theorem 17. *Let Γ be a Coxeter diagram whose corresponding Coxeter group W_Γ is finite. Let $\rho : W \to \mathrm{GL}(V)$ be the canonical representation. Consider the diagonal action of W on $V \times V$ via ρ, i.e.,*

$$w \cdot (v_1, v_2) := (wv_1, wv_2).$$

Define

$$X = V \times V - \bigcup_{r\in R} V^r \times V^r,$$
$$Y = X/W.$$

Then

$$A_\Gamma = \pi_1(Y),$$
$$\ker(A_\Gamma \to W_\Gamma) = \pi_1(X).$$

For various reasons, including this theorem, the Artin groups A_Γ corresponding to the finite W_Γ are usually called *spherical-type*, and the finite W_Γ are called *spherical.*

We probably won't prove this theorem just for the sake of time. It's not incredibly elucidating, it's mostly just giving the construction of the loops in the π_1 then proving they work. However this result belies the following conjecture, probably one of the more famous problems regarding Artin groups:

Theorem 18. *When A is a spherical-type Artin group, Y is a $K(A, 1)$.*

Eventually, we'll relate Y even more explicitly to the canonical representation (actually, to $C(W, S)$, cell structure and all). In addition, we'll generalize this to arbitrary Coxeter and Artin groups[2], which gives the famous "$K(\pi, 1)$ conjecture," one of the most prominent unsolved problems in Artin groups today. This generalization is a bit more complicated than I care to get into now, so let's move on for now before we get bogged down in the group theory side of things.

5. Curvature for metric spaces

There are many ways to extend the notion of curvature to arbitrary metric spaces. For example, Gromov's definition of δ-hyperbolicity (among his other work) essentially begat the field of geometric group theory. There have been many other generalizations of curvature as well. Typically, the best way to describe curvature is to examine how fast geodesics diverge. We will focus on a broadly-adopted characteristic typically credited to Cartan, Alexandrov, and Topongonov. This definition can be roughly interpreted as putting an "upper bound" on the curvature of a space, which, like in the case of Riemannian manifolds of non-positive curvature, has been successful in producing a wide variety of results (which we get into later). The way this is accomplished is by comparing the "thinness" of triangles of the given metric space to the triangles in a given model plane. We go over these details now.

Before we can do anything, we need the proper notion of a geodesic in a metric space (X, d). Like in Riemannian manifolds, these are the local isometries from $\mathbb{R}$ to X. More explicitly, they're the locally injective maps $\gamma : \mathbb{R} \to X$ so that

$$d(\gamma(a), \gamma(b)) = |b - a|$$

for all $a, b \in \mathbb{R}$. We call a metric space a *geodesic metric space* if any two points can be joined by a (not necessarily unique) geodesic. Similarly, we call a metric space an ℓ-geodesic metric space if every two points of distance no greater than ℓ can be joined by a geodesic.

A geodesic segment is a restriction of a geodesic γ to an interval $[a, b]$, and a geodesic ray is the restriction of a geodesic to an interval $[a, \infty)$ or $(\infty, b]$. Often we conflate the map with its image when we talk about geodesics.

Throughout the rest of the notes, we will fix $\kappa \in \mathbb{R}$, and let (X, d) denote a D_κ-geodesic metric space, and if we just say "metric space" we probably mean "D_κ-geodesic metric space"

[2]Fun fact: it took some time to generalize this properly after these ideas were first posed in terms of the finite case.

5.1. **Comparison triangles.** Triangles in a metric space are what one would probably expect; three geodesic segments with shared endpoints. More rigorously, a triangle T is a collection of geodesic segments $\gamma_i : [a_i, b_i] \to X$ with

$$\gamma_1(b_1) = \gamma_2(a_2)$$
$$\gamma_2(b_2) = \gamma_3(a_3)$$
$$\gamma_3(b_3) = \gamma_1(a_1).$$

These are the *vertices* of the triangle. Often we'll make reference to the triangle being the union of the images of the geodesic segments, namely

$$T = \bigcup \gamma_i([a_i, b_i]) \subseteq X$$

Since X is a geodesic metric space, given any three points in X we can form a triangle whose vertices are the given points. Note this isn't necessarily unique, but the side lengths will always be the same.

We let $\ell_i = |b_i - a_i|$ denote the side lengths of a triangle. In order to compare the thinness of triangles, we'd hope that a triangle in a given model plane would exist. Thankfully, we know exactly when they do!

Lemma 19 (BH Lemma 2.14)**.** *Let $\kappa \in \mathbb{R}$ and let $p, q, r \in X$ with $d(p, q) + d(q, r) + d(r, p) < 2D_\kappa$. Then there are three points $\tilde{p}, \tilde{q}, \tilde{r} \in M_\kappa^2$ so that*

$$d(\tilde{p}, \tilde{q}) = d(p, q)$$
$$d(\tilde{q}, \tilde{r}) = d(q, r)$$
$$d(\tilde{r}, \tilde{p}) = d(r, p).$$

We won't prove this here, see BH for proof.

As a consequence, we can make the following definition.

Definition 20. Let $\kappa \in \mathbb{R}$, and let $T = \bigcup \gamma_i([a_i, b_i])$ be a triangle in X with perimeter $< 2D_\kappa$ and vertices p, q, r. Let $\tilde{p}, \tilde{q}, \tilde{r}$ denote the corresponding points from the Lemma. Then we define the *κ-comparison triangle $\tilde{T}$ for T* to be the triangle in M_κ^2 whose sides are given by $[\tilde{p}, \tilde{q}]$, $[\tilde{q}, \tilde{r}]$, and $[\tilde{r}, \tilde{p}]$.

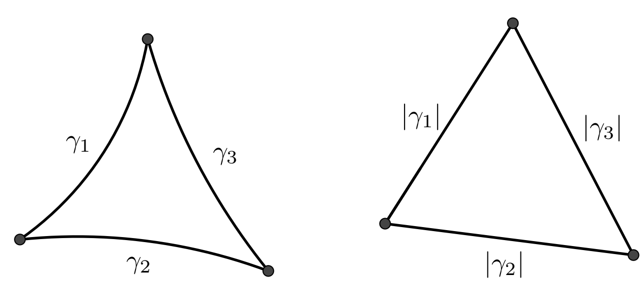

FIGURE 1. Comparison triangles

Now, we need a notion of comparing the thinness of triangles between these two spaces. This is given by the following.

Definition 21. Let $T = \bigcup \gamma_i([a_i, b_i])$ be a triangle in X with perimeter $< 2D_\kappa$. Let $\tilde{T} = \bigcup \tilde{\gamma}_i([a_i, b_i])$ the corresponding comparison triangle in M^2_κ. Choose points a in $\gamma_i\big((a_i, b_i)\big)$ and b in $\gamma_j\big((a_j, b_j)\big)$. Form "comparison points" $\tilde{a}, \tilde{b}$ in M^2_κ by choosing the (unique) points $\tilde{a}$ on $\tilde{\gamma}_i\big((a_i, b_i)\big)$ and $\tilde{b}$ on $\tilde{\gamma}_j\big((a_j, b_j)\big)$ which satisfy

$$d(\tilde{p}, \tilde{a}) = d(p, a)$$
$$d(\tilde{p}, \tilde{b}) = d(p, b).$$

where p is the shared vertex of γ_i and γ_j. We say that T is κ-thin if

$$d(a, b) \leq d(\tilde{a}, \tilde{b}).$$

for all choices of $a, b \in T$. (See Figure 2.)

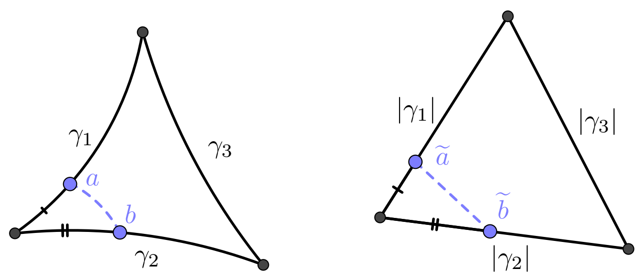

FIGURE 2. The CAT(κ) inequality

Definition 22. If $\kappa \leq 0$, then X is called a CAT(κ) space if X is a geodesic metric space and all of the triangles in X are κ-thin.

If $\kappa > 0$, then X is called a CAT(κ) space if X is a D_κ-geodesic metric space and all of the triangles in X are κ-thin.

Note: X doesn't have to be complete! Exercise: find a non-complete CAT(0) space. Harder: find one that isn't contractible (or even simply connected).

Before we give examples, we want to make use of the following relaxation of the CAT(κ) condition.

Definition 23. We say X has curvature $\leq \kappa$ if it is locally CAT(κ); that is, every point has a neighborhood which is CAT(κ) under the induced metric. If $\kappa = 0$ here, then we often say X is *non-positively curved*.

Now some examples

Examples 24.

(1) Normed linear spaces are CAT(0) under the metric induced by the norm. (Any triangle lives in a 2d subspace)
(2) Convex subsets of $\mathbb{E}^n$ are CAT(0).
(3) The product of CAT(0) spaces is CAT(0)
(4) A metric simplicial graph is CAT(κ) if and only if every locally embedded loop of length $< 2D_\kappa$ is embedded. In particular, a tree is CAT(κ) for all $\kappa \in \mathbb{R}$.

Where do we get more interesting examples? Well, in GGT, they typically come from so-called M_κ-polyhedral complexes, the higher dimensional analogue of Example 24.4, but we'll discuss those in a minute. First, let's prove some things about general CAT(κ) spaces.

6. Properties of CAT(κ) spaces

First, an interesting result relating our notion of curvature bounded above for metric spaces, and the traditional curvature of manifolds.

Theorem 25 (BH Ch. II.1 Thm 1A.6)**.** *A smooth Riemannian manifold M has curvature $\leq \kappa$ as a metric space if and only if the sectional curvature of M is $\leq \kappa$.*

The proof is quite long and technical, so we won't cover it here. But, it does illustrate that "curvature $\leq \kappa$" for metric spaces is good terminology.

One of the really nice things about CAT(κ) spaces compared to, say δ-hyperbolic spaces, is that there is a good notion of the "angle" between geodesics.

Definition 26. Let X be a geodesic metric space, and let $c_i : [0, a_i] \to X$, $i = 1, 2$ be geodesic paths with $c_1(0) = c_2(0) =: p$. For $t_i \in (0, a_i]$, let Δ_{t_1,t_2} be the geodesic triangle with vertices p, $c_1(t_1)$, and $c_2(t_2)$. Let $\widetilde{\Delta}_{t_1,t_2}$ be the corresponding comparison triangle in M_0^2. Let $\widetilde{\angle}_p(c_1(t_1), c_2(t_2))$ be the usual Euclidean angle of $\widetilde{\Delta}_{t_1,t_2}$ at the vertex corresponding to p. The *(Alexandrov or upper) angle* between the paths c_1 and c_2 is the number $\angle(c_1, c_2) \in [0, \pi]$ defined by

$$\begin{aligned}\angle(c_1, c_2) &= \limsup_{t_1,t_2\to 0} \widetilde{\angle}_p(c_1(t_1), c_2(t_2)) \\ &= \lim_{\varepsilon\to 0} \sup_{0<t,t'<\varepsilon} \widetilde{\angle}_p(c_1(t_1), c_2(t_2)).\end{aligned}$$

If X is uniquely geodesic and $p \neq x$, $p \neq y$, then the angle between the segments $[p, x]$ and $[p, y]$ may be denoted $\angle_p(x, y)$.

Proposition 27. *With the above set up,*

$$\cos(\widetilde{\angle}_p(c_1(t_1), c_2(t_2))) = \frac{1}{2t_1t_2}\big(t_1^2 + t_2^2 - d_X(c_1(t_1), c_2(t_2))^2\big).$$

Proof. Exercise. □

This gives us another interesting/useful criteria for determining when a space is CAT(κ).

Proposition 28. *X is a* CAT(κ) *space if and only if for every non-degenerate geodesic triangle Δ in X, the upper angle between any two sides is no greater than the κ-comparison triangle $\widetilde{\Delta}$.*

In particular, angles exist for any pair of geodesic segments in a CAT(κ) space.

Proof. Exercise; good for working out definition. □

Using this we can more easily derive the following.

Theorem 29.

(1) *If X is a* CAT(κ) *space, then it is a* CAT(κ') *space for all $\kappa' \geq \kappa$.*
(2) *If X is a* CAT(κ') *space for every $\kappa' > \kappa$, then it is a* CAT(κ) *space.*

Proof of (1). Suppose X is CAT(κ) and let $\kappa' \geq \kappa$. In particular, X is D_κ-geodesic. Then since D_κ is decreasing in κ, we see X is $D_{\kappa'}$-geodesic.

Suppose Δ is a geodesic triangle in X of perimeter $< 2D_\kappa$, and let Δ_κ and $\Delta_{\kappa'}$ be the M_κ^2 and $M_{\kappa'}^2$ comparison triangles, respectively. If we take $p, q \in \Delta$ and the respective comparison points p_κ, $p_{\kappa'}$ and q_κ, $q_{\kappa'}$ in Δ_κ and $\Delta_{\kappa'}$, we can somewhat intuitively see that

$$d_{M_\kappa}(p_\kappa, q_\kappa) \leq d_{M_{\kappa'}}(p_{\kappa'}, q_{\kappa'}),$$

or equivalently, that the angle at the vertices of Δ_κ is no greater than the corresponding angle of $\Delta_{\kappa'}$. Although the full rigorous argument is actually quite a bit more technical; see BH CH II.1 Lemma 1.13 for the full argument. □

Proof of (2). Let $x, y \in X$ with $d(x, y) < D_\kappa$. Then $d(x, y) < D_{\kappa'}$ for all $\kappa' > \kappa$ sufficiently close to κ. In particular, if X is $D_{\kappa'}$-geodesic for all $\kappa' > \kappa$ then X is D_κ-geodesic.

Let Δ be a geodesic triangle in X of perimeter $< 2D_\kappa$ with vertices p, q, r. Choose $\kappa' > \kappa$ sufficiently close so that Δ has perimeter $< 2D_{\kappa'}$. Let $\widetilde{\Delta}$ be the κ'-comparison triangle. Let a, b, c be the side lengths of Δ. Let γ be the angle between the sides of length a and b in X and let $\tilde{\gamma}$ be the corresponding angle in $\widetilde{\Delta}$. There are three cases to consider for κ but all are similar, so we'll just cover the case where $\kappa > 0$. In this case, the law of cosines in M_κ^2 applied to $\widetilde{\Delta}$ gives

$$\cos(c\sqrt{\kappa'}) = \cos(a\sqrt{\kappa'})\cos(b\sqrt{\kappa'}) + \sin(a\sqrt{\kappa'})\sin(b\sqrt{\kappa'})\cos(\tilde{\gamma}),$$

or rearranging, it gives

$$\cos(\tilde{\gamma}) = \frac{\cos(c\sqrt{\kappa'}) - \cos(a\sqrt{\kappa'})\cos(b\sqrt{\kappa'})}{\sin(a\sqrt{\kappa'})\sin(b\sqrt{\kappa'})}. \tag{1}$$

But by the previous characterization of CAT(κ) spaces, we know that $\gamma \leq \tilde{\gamma}$. Since (1) clearly varies continuously with κ', and $\gamma \leq \tilde{\gamma}$ for all $\kappa' > \kappa$ sufficiently close, it must hold for the κ-comparison triangle as well. Therefore X is CAT(κ). □

The following is one of the main reasons we care about CAT(κ) spaces.

Theorem 30. *Let X be a* CAT(κ) *space. Then for every pair of points $x, y \in X$ with $d(x, y) < D_\kappa$, there is a unique geodesic segment joining x and y. Moreover, this segment varies continuously with its endpoints.*

Proof (from BH). Consider $x, y \in X$ with $d(x, y) < D_\kappa$. Let $[x, y]$ and $[x, y]'$ be geodesic segments joining x to y, let $r \in [x, y]$, $r' \in [x, y]$ be such that

$$d(x, r) = d(x, r').$$

Let $[x, r]$ and $[r, y]$ be the segments whose union is $[x, y]$. Then any comparison triangle in M_κ^2 for the triangle with edges $[x, r]$, $[r, y]$, $[x, y]'$ is degenerate since $d(x, r) + d(r, y) = d(x, y) = \ell([x, y]')$. Thus the comparison points for r and r' are the same. Hence by the CAT(κ) inequality, $d(r, r') = 0$.

To show the geodesics vary continuously, we make note of a technical result from [BH]. Given any positive $\ell < D_\kappa$, there is a constant $C = C(\ell, \kappa)$ such that if $c, c' : [0, 1] \to M_\kappa^2$ are two "linearly parameterized" geodesic segments[3] and if $c(0) = c'(0)$, then $d(c(t), c'(t)) \leq C d(c(1), c'(1))$ for all t.

[3]This means there is a constant λ such that $d(c(t), c(t')) = \lambda|t - t'|$ for all t, t'.

Now, let x_n and y_n be sequences of points converging to x and y, resp. We may assume $d(x_n, y_n) < \ell$ and $d(x, y_n) < \ell$ for some positive $\ell < D_\kappa$. Let c, c_n, c'_n be linear parameterizations of the geodesics $[x, y], [x_n, y_n], [x, y_n]$, resp. By the $\mathrm{CAT}(\kappa)$ inequality,

$$\begin{aligned} d(c(t), c_n(t)) &\leq d(c(t), c'_n(t)) + d(c'_n(t), c_n(t)) \\ &\leq C(d(y, y_n) + d(x, x_n)). \end{aligned}$$

Hence $c_n \to c$ uniformly. □

Why is this important?

Corollary 31. *If X is a $\mathrm{CAT}(\kappa)$ space and B is a ball of radius $< D_\kappa$, then B is contractible. In particular, $\mathrm{CAT}(0)$ spaces are contractible.*

Proof. Let x be the point at which B is centered. The map $B \times [0, 1] \to X$ sending a point (y, t) to the point at distance $td(x, y)$ from y on the geodesic $[x, y]$ is a continuous retraction from B to x. □

This is a great criteria for showing a space is contractible, especially in the context of the material in the following couple sections. In fact, we have another great result regarding contractability of the universal cover of a space that we'll discuss later.

7. Polyhedral Cell Complexes

An M_κ-polyhedral cell complex is essentially a cell complex where each cell is given the metric of a polyhedra in M^n_κ. To make this more rigorous we need a couple definitions.

Definition 32. An M_κ-polyhedron P (sometimes called a convex polytope) is the non-empty intersection of finitely many halfspaces of M^n_κ. This endows P with a natural cell structure corresponding to the "facets" of the polyhedron. If $\kappa > 0$ then we require this intersection to be contained in a ball of radius $D_\kappa/2$. The dimension of P is the dimension of the smallest "subspace" containing it. If $x \in P$, then the "support" of x, denoted $\mathrm{supp}(x)$, is the smallest closed facet of P containing x.

Definition 33. Let $\{C_\lambda\}$ be a family of M_κ-polyhedra, and let X be their disjoint union. Let $\sim$ be an equivalence relation on X, and let $K = X/\sim$. Let $p : X \to K$ be the natural projection and $p_\lambda : C_\lambda \to K$ the restriction of p to C_λ. We call K an *M_κ-polyhedral complex* if

(1) p_λ is injective on the interior of C_λ, and
(2) If $x \in C_\lambda$ and $y \in C_\mu$ with $p_\lambda(x) = p_\mu(y)$, then there is an isometry $h : \mathrm{supp}(x) \to \mathrm{supp}(y)$ such that $p_\lambda(x_0) = p_\mu(h(x_0))$ for all $x_0 \in \mathrm{supp}(x)$

We denote the set of isometry classes of faces of K by $\mathsf{Shapes}(K)$

This is a technical way of saying "the gluing maps are isometries". Note that the condition (1) allows for self-gluings along the boundary of a given polyhedron.

Example 34. Any polyhedron in $\mathbb{E}^n$ is an M_0-polyhedral complex. Any convex tiling of $\mathbb{X}^2$ (where $X = \mathbb{S}, \mathbb{E}, \mathbb{H}$) is an M_κ-polyhedral complex. Cube complexes[4] are by definition M_0-polyhedral complexes.

[4] These are complexes where the cells are isometric to Euclidean cubes $[-1, 1]^n$ with no self-gluings. Sometimes we say "cubical complex" if the complex has self-gluings.

So when are these spaces CAT(κ)? First, we need to discuss the natural metric that a polyhedral complex is endowed with in order for this question to make sense. You could define it as the induced pseudometric from the projection $X \to K$, but there's a nicer more intuitive description in terms of geodesics.

Definition 35. A piecewise geodesic (segment) in a polyhedral complex K is a map $c : [a, b] \to K$ which has a subdivision $a = t_0 \leq t_1 \leq \cdots \leq t_k = b$ so that $c_i = c|_{[t_{i-1}, t_i]}$ is a geodesic within some C_{λ_i}, and for each $t \in [t_{i-1}, t_i]$ we have $c(t) = p_{\lambda_i}(c_i(t))$. If c is a piecewise geodesic, its length is

$$\ell(c) = \sum_{i=1}^{k} \ell(c_i).$$

We then have a well-defined pseudometric $d_K(a, b)$ given by taking the shortest length among the piecewise geodesics connecting a and b

In fact, we have the following

Theorem 36 (BH Ch I.7 Th 7.50)**.** *If K is an M_κ-polyhedral complex and $\mathsf{Shapes}(K)$ is finite, then (K, d_K) is a complete geodesic metric space.*

The condition that $\mathsf{Shapes}(K)$ be finite is entirely reasonable; in practice most polyhedral complexes have only finitely many isometry types of faces (e.g., if there is a cocompact group action).

Some of the primary reasons for studying these polyhedral complexes is that they allow us to more readily show they're CAT(κ), and that they arise naturally in the study of complexes of groups. We'll discuss complexes of groups in a minute, but we need some more combinatorial properties first.

7.1. **Cell complexes and posets.** The following material becomes important when we discuss complexes of groups, and the complex of groups for a Coxeter group in particular. I'll try to keep it brief so we're not so insanely overwhelmed with new material all at once.

Definition 37. Let X be a cell complex. The *barycentric subdivision* bX of X is the simplicial complex whose vertices are indexed by the cells of X, and where a set of vertices $V = \{c_0, \ldots, c_n\}$ span an n-simplex if and only if the cells c_i are "linearly" nested in X, i.e.,

$$c_0 \subseteq c_1 \subseteq \cdots \subseteq c_n.$$

If $\mathcal{P}$ is any poset, we define the *derived complex* $\mathcal{P}'$ of $\mathcal{P}$ to be the set of chains (linearly ordered subsets) of $\mathcal{P}$, itself ordered by inclusion. This is an "abstract simplicial complex" (meaning it's a poset of sets ordered by inclusion and closed under subsets). Every abstract simplicial complex has a *geometric realization* $|\mathcal{P}'|$; a simplicial complex whose set of simplices ordered by inclusion is order-isomorphic to $\mathcal{P}'$. More explicitly, $|\mathcal{P}'|$ is a simplicial complex whose vertices are the elements of $\mathcal{P}$, and a set of vertices $\{a_0, \ldots, a_n\}$ span an n-simplex if and only if it's linearly ordered. We denote this simplex by

$$[a_0, a_1, \ldots, a_n].$$

Proposition 38. *If X is a cell complex and $\mathcal{P}$ is its poset of cells ordered by inclusion, then bX is isomorphic to $|\mathcal{P}'|$.*

Proof. Exercise. □

Example 39. Let $X = [0,1]^2$. Label the cells as in Figure 3. Let $\mathcal{P}$ be the poset of faces of X. The barycentric subdivision of X is given in Figure 4. The

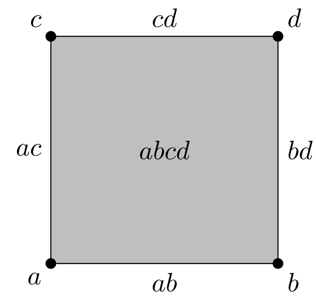

FIGURE 3. The cell complex X

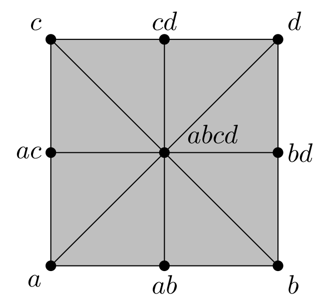

FIGURE 4. The barycentric subdivision bX of X

$(n-1)$-simplices of bX correspond directly to the chains of length n in $\mathcal{P}$.

Proposition 40. *As topological spaces, bX is homeomorphic to X. In fact, if X is an M_κ-polyhedral complex, then bX is an M_κ-simplicial complex and is canonically isometric to X (as metric spaces).*

Proof. Exercise. (Starting point: the vertices of bX are the barycenters (the "center of mass") of the corresponding cells of X.) □

Definition 41. Suppose X is a polyhedron[5] in M_κ^n. The dual X^* of X is the polyhedron whose proper m-faces (i.e., $0 < m < n$) are the $(n-m)$-faces of X, with adjacency preserved.

Proposition 42. *If X is a polyhedron and $\mathcal{P}$ is its set of cells ordered by inclusion, then $\mathcal{P}^{op}$ is the set of cells of X^* ordered by inclusion.*

Example 43. The dual of a cube is an octahedron, the dual of a tetrahedron is a tetrahedron, etc.

---

[5]Recall this is a non-empty intersection of finitely many halfspaces, equiv. the convex hull of finitely many points.

Definition 44. Let $\mathcal{P}$ be a poset, and let $a \in \mathcal{P}$. Then the *face of* a in $\mathcal{P}$ is

$$\mathcal{P}_{\leq a} = \{\, b \in \mathcal{P} : b \leq a \,\}\,,$$

and the *coface of* a in $\mathcal{P}$ is

$$\mathcal{P}_{\geq a} = \{\, b \in \mathcal{P} : b \geq a \,\}\,.$$

Replacing $\leq$ and $\geq$ with their strict counterparts gives the *lower link* and *upper link* of a in $\mathcal{P}$.

This terminology comes from the following observation: if X is a cell complex and $\mathcal{P}$ is its set of cells, then the face $\mathcal{P}_{\leq a}$ is the "face" of a in X, meaning, the subcomplex of X consisting of the closed cell a along with the cell structure coming from X. Moreover, if X has a dual, then the coface is simply the face of a in the dual of X.

8. Complexes of Groups

For the sake of simplicity (pun intended), these notes will focus on simple complexes of groups. The ideas for general complexes of groups are very similar, but much more technical to the point of obscuring the ideas, in my opinion. Plus, we only need simple complexes of groups for what we're doing here. For brevity, I'll probably just call them "complexes of groups" and omit the "simple" since all complexes of groups here will be simple.

A complex of groups is a higher-dimensional analogue of the famous notion of a graph of groups widely attributed to Serre. Essentially, it's built to give a way of reversing the process of quotienting a cell complex by a group action using only the data of the quotient and cell stabilizers. Where it becomes more interesting is when we involve metrics.

Anyway, enough chatter, here are the details.

Definition 45. A (simple) complex of groups $G(\mathcal{Q}) = (G_\sigma, \psi_{\tau\sigma})$ over a poset $(\mathcal{Q}, \leq)$ consists of the following data:

(1) for each $\sigma \in \mathcal{Q}$, a group G_σ, called the *local group at* σ, and
(2) for each $\tau < \sigma$, an injective homomorphism $\psi_{\tau\sigma} : G_\tau \to G_\sigma$ such that if $\tau < \sigma < \rho$, then

$$\psi_{\tau\rho} = \psi_{\sigma\rho}\psi_{\tau\sigma}.$$

Two simple complexes of groups $(G_\sigma, \psi_{\tau\sigma})$ and $(G'_\sigma, \psi'_{\tau\sigma})$ are *simply isomorphic* if there is a family of isomorphisms $\varphi_\sigma : G_\sigma \to G'_\sigma$ such that

$$\psi'_{\tau\sigma}\varphi_\tau = \varphi_\sigma\psi_{\tau\sigma}.$$

We can also define a morphism φ from a complex of groups $G(\mathcal{Q})$ to a group G; this is a collection of homomorphisms $\varphi_\sigma : G_\sigma \to G$ for $\sigma \in \mathcal{Q}$ so that

$$\varphi_\tau = \varphi_\sigma\psi_{\tau\sigma}$$

for each $\tau < \sigma$. We say that φ is *injective on the local groups* if each φ_σ is injective.

Example 46. Stallings' triangles of groups. 1-simplex of groups.

Example 47. Let (W, S) be a finite Coxeter system, (A, S) the associated Artin system, and let $\mathcal{S}$ denote the power set of S. For $T \subseteq S$ and $G = W$ or A, let G_T denote the subgroup of G generated by T. We can build a complex of groups $G(\mathcal{S})$ on

$\mathcal{S}$ by declaring the local groups to be the G_T, and the morphisms $\psi_{TU} : W_T \to W_U$ to be those induced from the inclusion $T \subseteq U$. Although, without more work, we don't necessarily know these maps are inclusions, hence we don't know if this actually determines a complex of groups. It turns out they are, but it will take some work to show this. These are the two motivating examples of complexes of groups, and we'll discuss them in great detail in the future.

So where does this come from? As mentioned before, it's based on the action of a group on a cell complex; we make that rigorous now.

Definition 48. Let X be a cell complex, and consider a cellular action $G \circlearrowleft X$ with a strict fundamental domain[6] Y. Let $\mathcal{Q}$ be the poset of cells of Y. We then associate a simple complex of groups $G(\mathcal{Q}) = (G_\sigma, \psi_{\tau\sigma})$ to this action as follows. We let G_σ be the stabilizer in G of the cell σ of Y, and $\psi_{\tau\sigma} : G_\tau \to G_\sigma$ is the inclusion homomorphism in G for cells $\tau \subseteq \sigma$. Then the inclusion maps $\varphi_\sigma : G_\sigma \to G$ define a morphism $\varphi : G(\mathcal{Q}) \to G$ which is injective on the local groups.

If a simple complex of groups $G(\mathcal{Q})$ arises in this fashion, then it is called *(strictly) developable.*

Why go through the trouble of all this abstract nonsense? One of the main reasons for defining complexes of groups is the following (this is the generalization of the development of a graph of groups in Bass-Serre theory).

Theorem 49 (The Basic Construction)**.** *Let Y be a cell complex, and let $\mathcal{Q}$ be its set of cells ordered by inclusion. Let $G(\mathcal{Q}) = (G_\sigma, \psi_{\tau\sigma})$ be a simple complex of groups over $\mathcal{Q}$. Let G be any group, and let $\varphi : G(\mathcal{Q}) \to G$ be a simple morphism which is injective on the local groups. Then there is a cell complex $D = D(Y, \varphi)$ called the* development of Y with respect to φ*, satisfying*

(1) Y is (isomorphic to) a subcomplex of D,
(2) There is a cellular action of G on D with strict fundamental domain Y

Note that this isn't quite the full generality of the development of an arbitrary complex of groups but this is all we'll need, and it's the most direct generalization of Bass-Serre theory. For more details, see BH.

Proof. Note that since φ is injective on the local groups, we identify each G_σ with its image $\varphi(G_\sigma) \subseteq G$. Note also that if $\tau < \sigma$ then $G_\tau \subseteq G_\sigma$ in G.

We define $D(Y, \varphi) = G \times Y/\sim$, where $\sim$ is the equivalence relation given by

$$(g, y) \sim (g', y') \qquad \Longleftrightarrow \qquad y = y' \text{ and } g^{-1}g' \in G_{\sigma(y)},$$

where $\sigma(y)$ is the smallest cell so that $y \in \sigma(y)$. We denote the equivalence classes by $[g, y]$. Then G acts by $g' \cdot [g, y] = [g'g, y]$. Note that Y may be identified with $\{1\} \times Y$. One then verifies this action has strict fundamental domain Y, and thus defines a complex of groups, which is then verified to be simply isomorphic to $G(\mathcal{Q})$. ∎

In fact we've shown the following

Corollary 50. *A simple complex of groups $G(\mathcal{Q})$ is strictly developable if and only if there is a group G and a morphism $\varphi : G(\mathcal{Q}) \to G$ which is injective on the local groups.*

[6]This is a subcomplex of X which intersects each orbit Gx in exactly one point.

Proof. If $G(\mathcal{Q})$ is developable, say from an action $G \circlearrowleft X$, then we saw that there is a morphism $G(\mathcal{Q}) \to G$ injective on the local groups. Conversely, the previous theorem shows that if there is such a morphism $\varphi : G(\mathcal{Q}) \to G$, then $G(\mathcal{Q})$ arises from the action of $G \circlearrowleft D(Y, \varphi)$, and is hence developable by definition. □

We note that in the more general cases of an arbitrary complex of groups, there is a weaker notion of developability, and it is possible for even a simple complex of groups to be developable but not strictly developable. (But we probably won't cover that here.) For the time being, we'll use "developable" and "strictly developable" as synonyms for brevity, and due to the fact that we won't be talking about general developability.

8.1. **General properties of complexes of groups.** The following are useful basic facts about complexes of groups.

Proposition 51. *Let Y be a cell complex and $\mathcal{Q}$ the set of cells of X. Let $G(\mathcal{Q})$ be a simple complex of groups, G an arbitrary group, and $\varphi : G(\mathcal{Q}) \to G$ a simple morphism which is injective on the local groups.*

1. *(Uniqueness) If X is a cell complex containing Y as a subcomplex and G acts on X with strict fundamental domain Y such that the associated complex of groups is isomorphic to $G(\mathcal{Q})$, then the identity map in Y extends to a G-equivariant isomorphism $D(Y, \varphi) \to X$.*
2. *(Functoriality) Let $\pi : G \to G'$ be a surjective homomorphism of groups. Let $\varphi' : G(\mathcal{Q}) \to G'$ be the morphism defined by $\varphi'_\sigma = \pi\varphi_\sigma$. If φ' is injective on the local groups, then the identity map on Y extends uniquely(!) to a G-equivariant morphism $p : D(Y, \varphi) \to D(Y, \varphi')$. The map p is a covering projection and $\ker(\pi)$ acts freely and transitively on the fibers of p.*

Proof.

1. The map is given by $[g, y] \mapsto gy$. (Check!)
2. We define $p : D(Y, \varphi) \to D(Y, \varphi')$ to be the map sending $x = [g, y]$ to $x' = [\pi(g), y]$. Clearly, this is G-equivariant. (Check uniqueness.) Then
$$p^{-1}(x') = \{\, [kx, y] : k \in \ker \pi \,\}\,.$$
Let $st_Y(y)$ denote the open star of y in Y and let σ be the smallest element of $st(y)$. Then, define $G_y = \varphi_\sigma(G_\sigma) \leq G$ and $G'_y = \varphi'_\sigma(G_\sigma) \leq G'$. Then the open star $st(x)$ of x in $D(Y, \varphi)$ is $gG_y st_Y(y)$ and the open star $st(x')$ of x' in $D(Y, \varphi')$ is $gG'_y st_Y(y)$ (check!). This shows p maps $st(x)$ homeomorphically onto $st(x')$ and $p^{-1}(st(x'))$ is the disjoint union of the open sets $k \cdot st(x)$ for $k \in \ker \pi$. Hence p is a Galois covering with Galois group $\ker \pi$. □

We collect some group-theoretic properties of the development here.

Proposition 52. *Let Y be a cell complex and $\mathcal{Q}$ its poset of faces. Suppose $G(\mathcal{Q})$ is a complex of groups with a morphism $\varphi : G(\mathcal{Q}) \to G$ which is injective on the local groups. Let $D = D(Y, \varphi)$ be the development, and let $G_0 \leq G$ be the subgroup generated by the $\varphi(G_\sigma) \leq G$ for the cells σ of Y. Let $D_0 = G_0 Y$, the G_0-orbit of the fundamental domain Y.*

1. *If $gY \cap Y \neq \varnothing$, then $g \in G_\sigma$ for some σ. In fact, if σ is a cell of y and $g \in G$, then the stabilizer in G of the cell $[g, \sigma]$ of D is $g^{-1} G_\sigma g$.*
2. *If $gD_0 \cap D_0 \neq \varnothing$ then $g \in G_0$, hence $gD_0 = D_0$.*
3. *D_0 is clopen in D.*

(4) *D_0 is connected if and only if Y is connected.*
(5) *D is connected if and only if Y is connected and $G = G_0$.*

Proof.

(1) This follows from the definition of the development, and the fact that Y is a strict fundamental domain.
(2) If $x \in gD_0 \cap D_0$, then there are points $y, y' \in Y$ and elements $g_0, g_0' \in G_0$ such that $x = g_0 y = (gg_0')y'$. But since Y is a strict fundamental domain, we have that $y = y'$, and hence $g_0^{-1} g g_0' \in G_\sigma$ where $y \in \sigma$.
(3) D_0 is a subcomplex, hence closed. Suppose $x \in \sigma$, a cell of D, with $x \notin D_0$. We claim that $\sigma \cap D_0 = \varnothing$, hence the complement of D_0 is a subcomplex of D, hence closed, hence D_0 is open. Choose $g \in G$ so that $g\sigma \subseteq Y$. If $\sigma \cap D_0$ were non-empty then $gD_0 \cap D_0$ would be non-empty, so by part (2), we would have $gD_0 = D_0$. However, $x \in gX_0$ and $x \notin D_0$ by assumption. Hence $\sigma \cap D_0 = \varnothing$, as desired.
(4) Suppose Y is connected. Each $g \in G_0$ is a product of elements $g_1, \dots, g_k$ from the various isotropy subgroups of Y. tk finish this proof
(5) Follows immediately from Part (4).

□

This inspires the following definition.

Definition 53. Let $G(\mathcal{Q})$ be a (simple) complex of groups. We define the *universal group* of $G(\mathcal{Q})$ to be

$$\widehat{G(\mathcal{Q})} \coloneqq \varinjlim_{\sigma \in \mathcal{Q}} G_\sigma,$$

the direct limit[7] of the system $(G_\sigma, \psi_{\tau\sigma})$ of groups and morphisms. The natural homomorphisms $\iota_\sigma : G_\sigma \to \widehat{G(\mathcal{Q})}$ give a canonical simple morphism $\iota : G(\mathcal{Q}) \to \widehat{G(\mathcal{Q})}$, where $\iota = (\iota_\sigma)$. (Note: in general ι_σ isn't injective!!)

Of course, this can also be defined in terms of universal properties.

Proposition 54. *Given any group G and any simple morphism $\varphi : G(\mathcal{Q}) \to G$, there is a unique homomorphism $\hat{\varphi} : \widehat{G(\mathcal{Q})} \to G$ such that $\varphi = \hat{\varphi} \circ \iota$ (meaning $\varphi_\sigma = \hat{\varphi} \circ \iota_\sigma$ for each $\sigma \in \mathcal{Q}$).*

Conversely, clearly every homomorphism $\hat{\varphi} : \widehat{G(\mathcal{Q})} \to G$ induces a simple morphism $\varphi : G(\mathcal{Q}) \to G$ by $\varphi_\sigma \coloneqq \hat{\varphi} \circ \iota_\sigma$. This shows the following

Proposition 55. *There is a simple morphism $\varphi : G(\mathcal{Q}) \to G$ which is injective on the local groups if and only if $\iota : G(\mathcal{Q}) \to \widehat{G(\mathcal{Q})}$ is injective on the local groups.*

In other words, $G(\mathcal{Q})$ is developable if and only if the homomorphisms $G_\sigma \to \widehat{G(\mathcal{Q})}$ are injective. For this reason we sometimes use the following terminology.

Definition 56. If $\mathcal{Q}$ is the poset of cells of a cell complex Y, then with the notations above, we let $\hat{D} = \hat{D}(\mathcal{G}(Q)) \coloneqq D(Y, \iota)$, which we sometimes refer to as the *universal development.*[8]

---

[7]Note: this is not the same as the usual definition over a directed set! Instead, it's the free product quotiented by the equivalence relation $a \sim b$ iff $b = \psi_{\tau\sigma}(a)$ for some $\tau < \sigma$.

[8]This terminology is non-standard, but we may refer to this by name later and this doesn't have a usual name in the literature.

The reason for this naming is as follows:

Proposition 57. *If Y is a connected, simply connected cell complex and $\mathcal{Q}$ is its poset of cells, then $\hat{D}(G(\mathcal{Q}))$ is connected and simply connected. Moreover, if $\varphi : G(\mathcal{Q}) \to G$ is a morphism which is injective on the local groups and the canonical homomorphism $\hat{\varphi} : \widehat{G(\mathcal{Q})} \to G$ is surjective, then $\pi_1(D(Y, \varphi)) \cong \ker(\hat{\varphi})$.*

So we can think of $\hat{D}$ as the "universal covering" of the development in a pretty literal sense.

To tie this together, let's give an example of how to rephrase the above results in order to say something useful about a very general situation.

Corollary 58. *Let G be a group acting on a simply connected cell complex X with strict fundamental domain Y. Then G is the direct limit of the isotropy subgroups of the cells of Y along its poset of cells.*

8.2. **The Complex of Groups $W(\mathcal{S})$.** Let's venture back to Coxeter land with what we've learned (and to provide another prominent example of a complex of groups). When W is finite, we defined a simplicial complex $C(W, S)$ in terms of the chambers of the action of W on $\mathbb{S}^n$, and we distinguished a fundamental chamber $\mathcal{C}$. In fact, we have the following classical result:

Proposition 59. *$\mathcal{C}$ is a strict fundamental domain for the action of W on $C(W, S)$.*

Proof. We'll present an outline of the proof, since some of it requires classical material on Coxeter groups that we haven't covered.

In order to show the Proposition, it suffices to show that W acts simply transitively on the set of all chambers. To this end, let Ω be the simplicial graph whose set of vertices is the set of chambers, with two vertices connected by an edge if and only if they're adjacent[9]. Note that Ω must be connected. Moreover, by the definition of Ω, two vertices are joined by an edge if and only if the corresponding chambers are adjacent; in other words, if and only if the intersection of the chambers is contained in a wall V^r. This means if A, B are two vertices, there is a unique reflection $r \in R$ swapping them. Hence W acts transitively on the vertices of Ω. This shows $\mathcal{C}$ is a fundamental domain.

To see this fundamental domain is strict, we need to show that for all $w \in W$, if $wx = y$ then $x = y$. We do this by induction on the word length $k = \ell_S(w)$ of w.

The case $k = 0$ is trivial ($w = 1$). Suppose $k \geq 1$. Let A be a chamber containing x. Then there is a wall V^r of A so that $w = rw'$ with $\ell_S(w') = \ell_S(w) - 1$. In particular, A and wA lie on opposite sides of V^r. Hence we must have $y, x \in V^r$, and $y = ry = r(wx) = r^2 w' x = w' x$. By inductive hypothesis, $y = x$. $\square$

Thus the action of W on $C(W, S)$ gives us a simple complex of groups, which we denote by $W(\mathcal{S}^{op}_{<S})$ (this notation will become more clear later). We can explicitly determine the data of the complex of groups by examining the isotropy subgroups of the fundamental chamber.

Previously, our fundamental chamber was arbitrary, but now we make a canonical choice of chamber: there is a unique chamber $\mathcal{C}$ of $C(W, S)$ whose walls are precisely V^s for $s \in S$; this chamber is the intersection of $\mathbb{S}^n$ with the half-spaces

$$\{ x_1 e_1 + \cdots + x_{n+1} e_{n+1} : x_i \geq 0 \},$$

[9]They intersect in a common codimension-1 facet.

where the e_i are the basis vectors from the canonical representation (i.e., the -1 eigenvectors of the reflections in S). In particular, there is an order-reversing isomorphism between the facets of the chamber and the proper[10] subsets of S, the latter of which we will denote $\mathcal{S}_{<S}$. In other words, $\mathcal{S}_{<S}^{op}$ is the poset of cells of $\mathcal{C}$. Now, one could build a complex of groups from this action, but it turns out for our purposes to be more beneficial to look at Σ instead.

As mentioned a while ago, $\partial\Sigma$ is the dual of C (when W is finite). Since $\mathcal{C}$ is a strict fundamental domain for $W \circlearrowleft C$, it follows that $L := \mathcal{C} \cap \Sigma$ is a strict fundamental domain for $W \circlearrowleft \partial\Sigma$. Since Σ is the cone on its boundary, it follows that the cone K on L is a strict fundamental domain for $W \circlearrowleft \Sigma$. Note that K is not a subcomplex of the cell structure we've put on Σ, it's just a subspace. However, it is a subcomplex of the barycentric subdivision $b\Sigma$: since $\partial\Sigma$ and C are dual, they have the same barycentric subdivision, and thus L, the intersection of $b\partial\Sigma$ with $\mathcal{C}$, is a subcomplex of this subdivision $b\partial\Sigma$. It follows that K is a subcomplex of $b\Sigma$. Later we will endow K with a cell structure inherited from the natural one on Σ. First, we need to describe K combinatorially.

The poset of cells of $\mathcal{C}$ is $\mathcal{S}_{<S}^{op}$. Thus its barycentric subdivision is $b\mathcal{C} = |(\mathcal{S}_{<S}^{op})'|$. Since $\partial\Sigma$ is dual to C it follows that the cells of $b\partial\Sigma$ corresponding to those of $b\mathcal{C}$ are $L = |(\mathcal{S}_{<S})'|$. The cone point comes from adding S back in, thus

$$K = |\mathcal{S}'|.$$

In other words, we want to build our complex of groups over $\mathcal{S}$.

First, we need to examine the structure of the local groups. Since $\partial\Sigma$ and C have a common barycentric subdivision, the cell structure of $b\Sigma$ comes (in part) from the hyperplanes V^r. In particular, the stabilizer of a vertex x of K is based on the collection $T(x)$ of $t \in S$ such that

$$x \in V^t.$$

(In fact, these hyperplanes completely determine x, as it's the unique vertex of K contained in their intersection.) Namely, the stabilizer of x is the reflection group generated by the orthogonal reflections about the hyperplanes V^t for $t \in T(x)$. Since this is a finite geometric reflection group, we know it's a Coxeter group, and we denote this by W_T. It's easy to verify that the stabilizer has presentation

$$W_T := \langle\, T \mid (t_i t_j)^{m_{ij}} = 1 \,\rangle, \tag{2}$$

where the m_{ij} are the same as those which define W. (Note: this is not a subgroup of W yet!) So, for $T \in \mathcal{S}$, we declare the local group to be W_T, as defined above. The maps ψ_{UT} are just the homomorphisms induced by the inclusions $U \subseteq T$. This data defines our complex of groups $W(\mathcal{S})$ with development $b\Sigma(W, S)$. Note that we have the following as a consequence of Corollary 50:

Proposition 60. *Let $M = (m_{ij})$ be a Coxeter matrix and (W, S) the associated Coxeter system. Moreover, suppose W is finite. Let $T \subseteq S$, and let*

$$W_T := \langle\, T \mid (t_i t_j)^{m_{ij}} = 1 \,\rangle. \tag{3}$$

Then the homomorphism $W_T \to W$ induced by the inclusion $T \subseteq S$ is injective.

---

[10] S isn't included because W acts freely on the sphere when it's finite. $\varnothing$ is included because this is the stabilizer of a top-dimensional cell.

(Note that this is actually a priori non-obvious.)

There's a nice way to describe the development combinatorially based on these definitions; at the risk of using confusing notation, we let

$$W\mathcal{S} = \bigcup_{T \in \mathcal{S}} W/W_T = \{\, wW_T : w \in W, T \in \mathcal{S} \,\} .$$

One can verify that $W\mathcal{S}$ is order-isomorphic to the set of cells of $\Sigma(W, S)$ and thus $|(W\mathcal{S})'|$ is isomorphic to $b\Sigma(W, S)$.

If you notice, there's really nothing stopping us from making these exact definitions on the complex-of-groups side work when W is infinite, or even for Artin groups. So, let's see what happens.

Definition 61. Suppose (G, S) is an arbitrary Coxeter or Artin system. For $T \subseteq S$, let G_T be defined as in (3) if G is a Coxeter group, or the corresponding presentation if G is an Artin group. Then, let

$$\mathcal{S} = \{\, T \subseteq S : W_T \text{ is finite} \,\} .$$

(We'll later see why we want these to be finite, and why we always want it based on whether the Coxeter group is finite.) Then, we define a complex of groups $G(\mathcal{S})$ whose local groups are G_T for $T \in \mathcal{S}$, and whose morphisms ψ_{UT} are the morphisms induced by the inclusions $U \subseteq T$. (These are injections by Proposition 60.)

However, now we come to an issue—is this complex of groups developable? In the finite case, $W(\mathcal{S})$ was developable by definition, since it arose from the action of W on the Coxeter polytope Σ. But for an arbitrary Coxeter or Artin group, we don't necessarily have a nice action on a space like this[11]. We have one way to determine if a complex of groups is developable so far, which is to construct a morphism which is injective on the local groups. However, a priori, we don't know if W_T injects into W! It turns out there's another way to show a complex of groups is developable using the notions of non-positive curvature we discussed previously.

9. Curvature for Complexes of Groups

We now want to introduce the notion of "curvature $\leq \kappa$" for a complex of groups, in a similar vein to metric spaces. One might think that a good notion would be to just require that the development support a metric with curvature $\leq \kappa$, but we want to define this notion for complexes of groups which are a priori not developable, so we need to do a bit more work. First, we'll need some more generalities about polyhedral complexes.

Definition 62. Let K be an M_κ-polyhedral complex. Fix $x \in K$. The *open star* of x in K, denoted $st(x)$, is the union of the interiors of all cells of K containing X (equiv,. the complement of the largest subcomplex of X not containing x). The *closed star* $St(x)$ is the smallest subcomplex of K containing $st(x)$.

The *link* of x in K, denoted $Lk(x, K)$ is the "space of directions" at x endowed with the subspace metric. More explicitly, let $\varepsilon(x)$ denote the distance from x to the boundary $\partial st(x)$ of the open star of x. Then the link $Lk(x, K)$ is (isometric to) the sphere centered at x of radius $\varepsilon(x)/2$. The link inherits a polyhedral cell structure from the intersection of the sphere with the cell structure of K.

[11]Actually, the dual of the canonical representation for an arbitrary Coxeter group gives us an action with a strict fundamental domain, but for the sake of illustration, we're going to use an alternate approach, since this still doesn't work for Artin groups.

These are both immensely useful notions, as we will see soon.

9.1. **The local development.** It turns out that developability is actually a global condition; there is a "local development" associated to any complex of groups. To be more precise, one can construct what would be the star of a point in the development, if the complex were developable.

Definition 63. Let $\mathcal{P}$ be a poset, and $G(\mathcal{P})$ a simple complex of groups over $\mathcal{P}$. Let $\sigma \in \mathcal{P}$. Define $G(\mathcal{P}_{<\sigma})$ to be the restriction of $G(\mathcal{P})$ to the subposet $\mathcal{P}_{<\sigma} := \{\, \rho \in \mathcal{P} : \rho < \sigma \,\}$. There is a canonical simple morphism $\psi_\sigma : G(\mathcal{P}_{<\sigma}) \to G_\sigma$ given by $(\psi_\sigma)_\rho = \psi_{\rho\sigma} : G_\rho \to G_\sigma$. In particular, ψ_σ is injective on the local groups, so $G(\mathcal{P}_{<\sigma})$ is developable, with development $D = D(\mathcal{P}_{<\sigma}, \psi_\sigma)$. Sometimes, $D = Lk_\sigma$ is called the *lower link* of σ. We then define

$$St(\sigma) = |\mathcal{P}_{\geq\sigma} * D|,$$
$$st(\sigma) = \operatorname{int}(St(\sigma)).$$

Sometimes, $\mathcal{P}_{\geq\sigma} = St(\sigma, \mathcal{P})$ is called the *combinatorial star* of σ. This $St(\sigma)$ is called the *local development of* $G(\mathcal{P})$ *at the vertex* σ.

Note that there is a natural action of G_σ on $St(\sigma)$ (i.e., the usual action on D and the trivial action on $\mathcal{P}_{\geq\sigma}$), and it has strict fundamental domain $St(\sigma, \mathcal{P})$. This gives us the information for a developable complex of groups, which we sometimes call $G(\sigma)$.

We can now define the correct notion of curvature:

Definition 64. Let $\kappa \in \mathbb{R}$. Let $\mathcal{Q}$ be a poset such that $|\mathcal{Q}'|$ is an M_κ-polyhedral complex, and let $G(\mathcal{Q})$ be a simple complex of groups over $\mathcal{Q}$. Then $G(\mathcal{Q})$ is said to be of curvature $\leq \kappa$ if the induced metric on the open star $st(\sigma)$ of σ in the local development $St(\sigma)$ is of curvature $\leq \kappa$ for each $\sigma \in \mathcal{Q}$.

Now, we have

Theorem 65. *Let $\mathcal{Q}$ be a poset such that $\Delta := |\mathcal{Q}'|$ is simply connected. Let $G(\mathcal{Q})$ be a simple complex of groups over $\mathcal{Q}$. Suppose $\kappa \leq 0$. If Δ supports a metric which makes it an M_κ-polyhedral complex with $\mathsf{Shapes}(\Delta)$ finite and $G(\mathcal{Q})$ has curvature $\leq \kappa$, then $G(\mathcal{Q})$ is strictly developable.*

The proof of this is well outside the scope of these notes. If you're interested, it's covered in detail in BH Ch III.$\mathcal{G}$.

So now all we have to do to check developability is to verify that the open stars are CAT(0)! But... we still don't have a very easy way to check if a space is CAT(0). So let's fix that!!

10. Curvature for M_κ-polyhedral complexes

We're going to go through a series of lemmas which will culminate in really nice characterizations of curvature $\leq \kappa$ for M_κ-polyhedral complexes. The main ideas mostly hinge on the following definition.

Definition 66. An M_κ-polyhedral complex K satisfies the *link condition* if for every vertex $v \in K$, the link $Lk(v, K)$ is a CAT(1) space.

Before we continue, we need a technical result, which, while important, has a very long and unenlightening proof, so we'll skip it.

Theorem 67 (BH Ch. II, Th 3.14)**.** *Let Y be a metric space. Then Y is* CAT(1) *if and only if the κ-cone*[12] *$C_\kappa Y$ is* CAT(κ)*.*

Our first important theorem is the following.

Theorem 68. *An M_κ-polyhedral complex K with* Shapes(K) *finite has curvature $\leq \kappa$ if and only if it satisfies the link condition.*

Proof. Let $v \in K$ be a vertex. Recall that the link of v is the $\varepsilon(v)$-sphere centered at v, so the $\varepsilon(v)$-ball is the cone on $Lk(v, K)$. In fact, it's isometric to the κ-cone $C_\kappa L(v, K)$ of the link (this follows from the definition of the metric on the κ-cone). Thus K satisfies the link condition if and only if each vertex has a neighborhood which is a CAT(κ) space. To finish the proof, use the following exercise. □

Exercise 69. For any $x \in K$, if v is a vertex of $supp(x)$ and $\eta > 0$ is sufficiently small, then $B(x, \eta)$ is isometric to $B(x', \eta)$ for some $x' \in B(v, \varepsilon(v))$.

So this theorem applied to the complex of groups setting gives us

Corollary 70. *Let $\kappa \in \mathbb{R}$. Let $\mathcal{Q}$ be a poset such that $|\mathcal{Q}'|$ is an M_κ-polyhedral complex, and let $G(\mathcal{Q})$ be a simple complex of groups over $\mathcal{Q}$. Then $G(\mathcal{Q})$ is of curvature $\leq \kappa$ if and only if the link of σ in the open star $st(\sigma)$ is* CAT(1)*.*

We didn't describe it before but it turns out this link has a nice description: it's actually just $|\mathcal{P}_{<\sigma} * Lk^\sigma|$ (see previous chapter for the terminology). So showing a complex of groups is nonpositively curved amounts to putting a CAT(1) metric on that complex. Although checking that a complex is CAT(1) is still no easy feat, it turns out that we can do this inductively, in a way. Namely, we have the following

Theorem 71 (BH Ch II Thm 5.4)**.** *Let K be an M_κ-polyhedral complex with* Shapes(K) *finite. If $\kappa \leq 0$ then TFAE*

(1) K is CAT(κ)*,*
(2) K is uniquely geodesic,
(3) K satisfies the link condition and contains no isometrically embedded circles, and
(4) K satisfies the link condition and is simply connected.

If $\kappa > 0$, then TFAE

(1) K is CAT(κ)*,*
(2) K is $(\pi/\sqrt{\kappa})$-uniquely geodesic, and
(3) K satisfies the link condition and contains no isometrically embedded circles of length less than $2\pi/\sqrt{\kappa}$.

In both cases, (1) $\iff$ (3) follows from our Theorem 68 and from Proposition 4.17 in [BH Ch. II]. We won't show the other directions since this is basically all we need in our settings. (Sometimes we'll use (4) $\implies$ (3) but the proof isn't very enlightening, so we'll skip it.)

So, in order to verify an arbitrary M_κ-polyhedral complex (with Shapes finite) is CAT(κ), it suffices to check that the links of the vertices are CAT(1), which we can do by verifying there are no "small circles", then checking that the links of the vertices in the link are CAT(1), and so on. While this is much easier than verifying the space is CAT(κ) by definition, in practice it could still be very hard. This is where the works of Gromov and Moussong come in.

[12] This is just the cone on Y with a special metric defined in Ch I.5, Definition 5.6 of BH

11. Combinatorial tools

The beauty of the following ideas come from the fact that they provide an (almost) exclusively combinatorial criteria for an M_κ-polyhedral complex to be CAT(κ). While they aren't 100% applicable all the time, they will see great use in our specific examples. Since they're combinatorial in nature, in order to explain the conditions, we'll again need more terminology.

Definition 72. Let L be a simplicial complex. We call L a *flag complex* (or "clique complex", or say "L satisfies the no-Δ condition") if for every finite set S of vertices of L, if the elements of S are pairwise joined by edges, then S spans a simplex of L.

Example 73. A circuit with 3 edges is not a flag complex, but a 2-simplex is (in fact any simplex is). Any simplicial graph is a flag complex if and only if all circuits have length ≥ 4.

There are two criteria we will discuss. The first, by Gromov, is more rigid, but has a beautiful application, which has seen widespread use in real-world subjects such as A.I. training and robot movement[13].

Definition 74. An M_1-polyhedral complex whose cells are all (isomorphic to) simplices will be called a *spherical simplicial complex*. A spherical simplicial complex is called *all-right* if each edge has length $\pi/2$.

Theorem 75. *Let L be a finite dimensional all-right spherical simplicial complex. Then L is* CAT(1) *if and only if it's a flag complex.*

tk fill in proof!!

The great corollary is thus:

Corollary 76 (Gromov's Link Condition)**.** *A finite dimensional cubed complex*[14] *is non-positively curved if and only if the link of each of its vertices is a flag complex.*

Proof. We showed that an M_0-polyhedral complex has non-positive curvature if and only if each vertex link is CAT(1). But in a cubed complex, the vertex links are all-right simplicial complexes. ☐

This is a great criteria, but really only useful for cube complexes (which isn't a bad thing of course, but lacks some generality we need). Thankfully, Moussong proved a wonderful generalization of this Theorem in his thesis [tk cite].

Definition 77. A spherical simplicial complex L is called a *metric flag complex* if every set $V = \{v_0, \ldots, v_k\}$ of vertices of L which is pairwise joined by an edge e_{ij} and there is a spherical k-simplex in $\mathbb{S}^k$ whose edge lengths are equal to the edge lengths $\ell(e_{ij})$, then V spans a k-simplex of L.

Lemma 78 (Moussong's Lemma)**.** *Suppose L is a spherical simplicial complex whose edges have length* at least $\pi/2$. *Then L is* CAT(1) *if and only if it's a metric flag complex.*

This lemma follows from the following lemma.

[13] Don't ask me why, no clue

[14] An M_0-polyhedral complex whose cells are all isomorphic to Euclidean cubes

Lemma 79 (Bowditch's Lemma, tk cite)**.** *Let X be a compact metric space of curvature ≤ 1. Suppose that every nonconstant closed rectifiable curve of length $< 2\pi$ is homotopically trivial through a homotopy which strictly decreases the length of the curves in the family. [tk wording] Then X is* CAT(1)*.*

This proof is unfortunately outside the scope of these notes. I'm told it uses the Birkoff curve-shortening process(?).

tk fill in proof of moussong

12. Σ is CAT(0)

After all this build up, we finally show that $\Sigma(W, S)$ is developable for an arbitrary Coxeter system (W, S), which as we've discussed amounts to putting a CAT(1) metric on the link of the local development. We establish a chain of lemmas to do this. First, we need to describe the cell structure on $K = |\mathcal{S}'|$ which gives the domain nonpositive curvature.

First, suppose (W, S) is a finite Coxeter system. We endow $K = K(W, S) = |\mathcal{S}'|$ with a natural cell structure from Σ and $\mathcal{C}$. Since Σ is the cone on $\partial\Sigma$, we consider the cone on $\mathcal{C}$. The set of cells of $\mathcal{C}$ is $\mathcal{S}^{op}_{<S}$, so the set of cells on the cone of $\mathcal{C}$ is $\mathcal{S}^{op}$. Since Σ is dual to C, it follows that the cells of Σ which are dual to (the cone on) $\mathcal{C}$ are indexed by $\mathcal{S}$. Now since K is the intersection of the cone on $\mathcal{C}$ and Σ, each face of K is an intersection of a cell of $\mathcal{C}$ with a cell of Σ. We are able to explicitly describe these faces combinatorially as follows.

Note that if $\mathcal{P}$ is the poset of cells of a cell complex X ordered by inclusion, then for any $a \in \mathcal{P}$, the set $\mathcal{P}_{\leq a}$ combinatorially represents the cell structure on a, as it includes all the cells of X which are subsets of a. Thus the faces of K are

$$F_{T_1,T_2} := \mathcal{S}_{\leq T_1} \cap \mathcal{S}^{op}_{\leq T_2}$$

For $T_1 \in \mathcal{S}$, $T_2 \in \mathcal{S}^{op}$. But since the second just has the order reversed, we can more simply write this as

$$F_{T_1,T_2} = \mathcal{S}_{\leq T_1} \cap \mathcal{S}_{\geq T_2}$$

for $T_1, T_2 \in \mathcal{S}$. To summarize, the vertices of K are the elements $T \in \mathcal{S}$, and a set V of vertices span a cell of K if and only if $V = \mathcal{S}_{\leq T_1} \cap \mathcal{S}_{\geq T_2}$ for some $T_1, T_2 \in \mathcal{S}$.

Proposition 80. *The cells of K are combinatorial cubes. Viewed as a subspace of Σ with the induced metric, this makes K into an M_0-polyhedral complex.*

Proof. Exercise. □

We now consider the correct metric and cell structure on K when (W, S) is arbitrary.

Definition 81 (The metric on K)**.** Let (W, S) be any Coxeter system, and let $K = |\mathcal{S}'|$, where

$$\mathcal{S} = \{\, T \subseteq S : W_T \text{ is finite} \,\}.$$

For $T_1 \subseteq T_2 \in \mathcal{S}$, let

$$F_{T_1,T_2} = \mathcal{S}_{\geq T_1} \cap \mathcal{S}_{\leq T_2}.$$

The cell structure and metric on K is as follows:

(1) The vertices of K are the elements of $\mathcal{S}$.

(2) A set V of vertices of K span a cell if and only if they're equal to F_{T_1,T_2} for some $T_1 \subseteq T_2 \in \mathcal{S}$.
(3) We endow the cell corresponding to F_{T_1,T_2} with the Euclidean metric described above for $K(W_T, T) \subseteq \Sigma(W_T, T)$.

Lemma 82. *The metric on K is well-defined.*

Proof. Exercise. $\square$

Lemma 83. *Let v be a vertex of K, say corresponding to $T \in \mathcal{S}$. Then under this cell structure, the link $Lk(v, st(v))$ in the local development of $W(\mathcal{S})$ is isomorphic to $|\mathcal{S}'_{>T}| * C(W_T, T)$. (The $*$ means the spherical join of the complexes as defined in BH Ch I.5 Def 5.13.)*

We'll provide the main argument of the proof, but the details should be verified by the reader to better understand what's going on.

Proof. The star $St(v)$ is by definition the join of the upper star $|\mathcal{S}'_{T\geq}|$ and lower link $Lk_v(\mathcal{S})$ as defined previously. Under the cubical structure, the cells containing v are thus

$$gF_{T_1,T_2} = \mathcal{S}_{\leq T_2} \cap g\mathcal{S}_{\geq T_1}$$

where $T_1 \subseteq T \subseteq T_2 \in \mathcal{S}$ and $g \in W_T$ (check!). But the decomposition of the star as a join means we can decompose these cells as a Cartesian product, namely,

$$\begin{aligned} gF_{T_1,T_2}\mathcal{S}_{\leq T_2} \cap g\mathcal{S}_{\geq T_1} \\ &= (\mathcal{S}_{\leq T_2} \cap \mathcal{S}_{\geq T}) \times (g\mathcal{S}_{\geq T_1} \cap \mathcal{S}_{\leq T}) \\ &= F_{T,T_2} \times gF_{T_1,T}. \end{aligned}$$

The set of F_{T,T_2} for $T_2 \in \mathcal{S}$ and $T \subseteq T_2$ is naturally identified with $\mathcal{S}_{\geq T}$, while the set of $F_{T_1,T}$ with $T_1 \subseteq T$ is naturally isomorphic to $\mathcal{S}^{op}_{\leq T}$ (check!). Then passing to the link, the cartesian product becomes a join[15], and we see that we're restricted to $T_2 \supsetneq T$ and $T_1 \subsetneq T$, so these become $\mathcal{S}_{>T}$ and $\mathcal{S}^{op}_{<T}$. Since T is spherical, the development of $\mathcal{S}^{op}_{<T}$ is naturally identified with $C(W_T, T)$ with the usual metric. $\square$

Lemma 84. *If $T \subseteq S$ is finite, then $C(W_T, T)$ is* CAT(1).

Proof. If T is finite, then $C(W_T, T)$ is a simplicial decomposition of the $|T|$-sphere, hence CAT(1). $\square$

Since the join of two CAT(1) spaces is again CAT(1), it remains to show that the upper link is CAT(1).

Lemma 85. *With the metric and cell structure induced from K, the upper link $|\mathcal{S}'_{>T}|$ has the following metric:*

Lemma 86. *Let $v_0 = [\varnothing]$ be the vertex corresponding to the trivial subset of $\mathcal{S}$. Then $Lk(v, st(v))$ is* CAT(1).

Proof. By Lemma 83, $Lk(v, st(v))$ is isomorphic to $K_0 \coloneqq |\mathcal{S}'_{>\varnothing}|$. Let $\sigma_T = |\mathcal{S}_{>\varnothing} \cap \mathcal{S}_{\leq T}|$. Each simplex of $Lk(v, st(v))$ is contained in some σ_T for $T \in \mathcal{S}$. (Check!) The piecewise spherical structure on K_0 induced from the metric on K assigns each simplex σ_T the dihedral angles $-\cos(\pi/m_{ij})$. $\square$

[15] Exercise: make this rigorous.

The following

Lemma 87. *$Lk(v, st(v))$ has edge lengths $\geq \pi/2$.*

tk include proof

Lemma 88. *$Lk(v, st(v))$ is a metric flag complex, hence* CAT(1).

tk include proof

Corollary 89. *$st(v)$ is* CAT(0)*, hence $W(\mathcal{S})$ is developable, and the development $\Sigma(W, S)$ is* CAT(0).

13. The Cartan-Hadamard Theorem

The main purpose of this section is to prove the following wonderful theorem (*cf.* non-positively curved Riemannian manifolds).

Theorem 90 (Cartan-Hadamard)**.** *Let X be a complete connected non-positively curved metric space of curvature $\leq \kappa$, where $\kappa \leq 0$. Then the universal cover $\widetilde{X}$ of X is* CAT(κ).

CAT(κ) spaces: the model spaces M_κ, CAT(κ) polyhedral complexes (including CAT(0) cube complexes), the link condition.

Complexes of groups: simple complexes of groups, the local development, global developability, the universal and fundamental groups, metrics of nonpositive curvature.

Coxeter groups: spherical (finite) linear reflection groups, the Davis-Moussong complex Σ for infinite Coxeter groups, the CAT(0) metric on Σ.

Artin groups: hyperplane arrangements, spherical-type Artin groups, the Deligne complex Φ for infinite-type Artin groups, right-angled Artin groups (including special cube complexes), type FC Artin groups and metrics on Φ, the $K(\pi, 1)$ conjecture.