Quasiconvexity in Word-Hyperbolic Coxeter Groups

Adam Simon Levine
Mathematics Department
Harvard University

April 3, 2005

Abstract

In Chapter 1, I outline many of the basic properties of word-hyperbolic groups and give an introduction to CAT(κ) spaces. In Chapter 2, I prove a number of important theorems about Coxeter groups, including the Exchange Condition and Tits's solution to the word problem. I also discuss Moussong's theorem on which Coxeter groups are word-hyperbolic. Finally, in Chapter 3, I present some original research showing that certain subgroups of certain word-hyperbolic Coxeter groups are quasiconvex.

Acknowledgements

I am deeply grateful to the following people for their assistance with this thesis:

To my advisor, Prof. Dylan Thurston, for guiding me through the thesis-writing process, always taking the time to answer my questions, and constantly challenging me to learn more.

To Profs. Patrick Bahls and Kim Whittlesey, for supervising my summer research at the University of Illinois at Urbana-Champaign and sparking my interest in geometric group theory.

To Mark Lipson, for his meticulous proofreading and his pages of detailed suggestions and corrections.

To Peter Green and David Escott, for their helpful comments.

To my roommates, for putting up with my math puns for four years.

To Amy Pasternack, for her late-night cookie deliveries to the Science Center computer lab.

To my brother Ben, for being a good sport during my lengthy explanations of Coxeter groups, and for being a dear friend.

To my parents, for their helpful advice on how to make my thesis comprehensible to laypeople, their valiant efforts to explain to friends and relatives exactly what I study, and their constant love and support.

And finally, to my high school calculus teacher, Mr. Robert Arrigo, for instilling in me a lifelong love of mathematics.

Contents

List of Figures

Introduction

Given a finite set of letters, consider the set of words — not English words, but simply finite sequences of letters — that can be written with those letters. We define two operations:

- Whenever a letter appears twice consecutively, we may delete both occurrences. Likewise, we may insert two consecutive, identical letters into a word at any point.
- For each pair of letters (say, a and b), choose one of the following rules: $ab = ba$, $aba = bab$, $abab = baba$, and so on. That is, if we see a sequence of the form $aba\cdots$, we may replace it with $bab\cdots$ (where both sequences are of the specified length), and vice versa. These rules need not be the same for different pairs of letters; we could say, for instance, that $ab = ba$, $aca = cac$, and $bcbc = cbcb$. We may also choose not to impose any such rule on particular pairs of letters.

Two words are said to be *equivalent* if it is possible to get from one to the other using these two operations. The set of all classes of equivalent words is called a *Coxeter group* (named for the geometer H.S.M. Coxeter, who studied this construction starting in the 1930s). For instance, if the rules are $ab = ba$, $aca = cac$, and $bcbc = cbcb$ as above, then the words $abac$, $aabc$, $baac$, and bc are all equivalent, so they represent the same element of the Coxeter group. They are not equivalent, however, to the word cb, which thus represents a different element of the Coxeter group.[1]

As an example, let the set of letters (also called *generators*) consist of a and b, and suppose that $aba = bab$. The elements of the Coxeter group can be written as a, b, ab, ba, aba (which is the same as bab), and the so-called *empty word*, a word of zero letters, often denoted by 1. Any longer word can be reduced to one of these six possibilities. As another example, the group generated by a, b, and c, with the rules $aba = bab$, $bcb = cbc$, and $ac = ca$, has 24 distinct elements. On the other hand, if the rules are $aba = bab$, $bcb = cbc$, and $aca = cac$, then the group contains infinitely many distinct elements, for the words abc, $abcabc$, $abcabcabc$, $abcabcabcabc$, etc. all represent different elements.

Many Coxeter groups occur naturally as so-called *reflection groups*, sets of transformations of a space that can be realized as the composition of a given set of reflections. For instance, consider two lines in the plane, labelled L_a and L_b, that intersect at an angle of $60°$.

[1] Formally, in order to call this set a group, we need to define a law of composition. However, as this Introduction is intended for a general audience, we shall use this colloquial description. A rigorous definition of Coxeter groups will appear in Chapter 2.

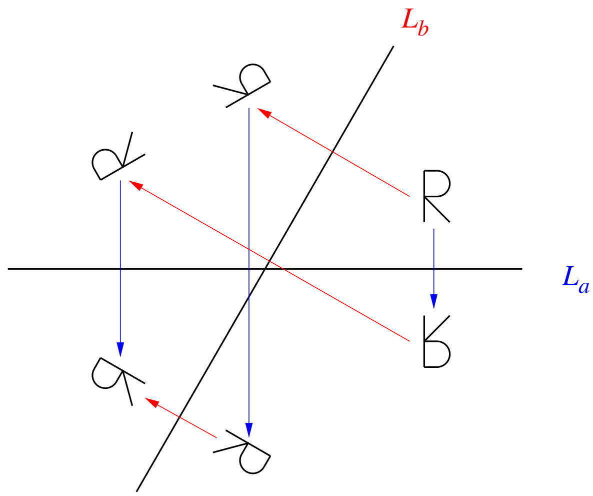

Figure 1: Reflections through lines meeting at 60°.

Take any figure, such as the upper right letter R shown in Figure 1, and reflect it across the line L_a to obtain the upside-down R in the lower-right corner. Next, reflect the new figure across L_b to obtain the R in the upper-left corner, and then reflect that down across L_a to the lower-left backward R. Next, starting with the original figure again, reflect first across L_b, then across L_a, then across L_b again, and note that the final image is exactly the same. In other words, if a and b represent reflections across L_a and L_b, respectively, then we obtain the rule $aba = bab$. Also, performing the same reflection twice in a row leaves every point fixed, so we may insert or delete pairs of consecutive, identical reflections into any sequence without consequence. Thus, we can say that the set of transformations of the plane that are compositions of the reflections a and b is a Coxeter group. More generally, if L_a and L_b intersect instead at an angle of $180°/n$, where n is an integer greater than 1, we obtain a Coxeter group with the rule

$$\underbrace{aba\cdots}_{n \text{ letters}} = \underbrace{bab\cdots}_{n \text{ letters}}.$$

One way to visualize a Coxeter group geometrically is the *Cayley graph*. This is a graph (a network of vertices connected by edges) with one vertex for every distinct element of the group and with edges labelled with the different letters, one of each type at each vertex. Words written with the generating letters correspond to paths emanating from a fixed starting point, and two words represent the same element of the group if and only if the corresponding paths end at the same point.

For example, the Cayley graph for the Coxeter group generated by a and b with the rule that $ab = ba$ is a square with its sides alternately labelled a and b. If $aba = bab$, the graph is a hexagon; if $abab = baba$, an octagon; and so on. Assume that the polygons are regular; that is, all the sides and all the angles are equal.[2] With more than two letters, the graph

[2]Formally, the angle measures are a property not of the graph itself but of the embedding of the graph into a larger space, in this case $\mathbb{R}^n$. To avoid having to introduce the concept of an abstractly defined graph,

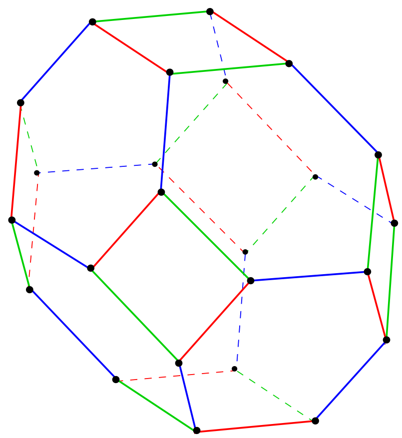

Figure 2: Cayley graph for Coxeter group generated by a, b, and c, subject to the rules $aba = bab$, $bcb = cbc$, and $ac = ca$.

can be formed by gluing together polygons of the appropriate types in such a way that there is one polygon of each type at every vertex. Two such graphs are shown in Figures 2 and 3.

The Cayley graph gives a useful criterion for determining whether a Coxeter group is finite or infinite. Let us consider the Coxeter group generated by a, b, and c, subject to the following rules:

$$\underbrace{aba\cdots}_{k\text{ letters}} = \underbrace{bab\cdots}_{k\text{ letters}}, \qquad \underbrace{bcb\cdots}_{l\text{ letters}} = \underbrace{cbc\cdots}_{l\text{ letters}}, \qquad \underbrace{aca\cdots}_{m\text{ letters}} = \underbrace{cac\cdots}_{m\text{ letters}}.$$

This group is called the (k, l, m) *triangle group.* The Cayley graph consists of polygons with $2k$, $2l$, and $2m$ sides, arranged with one polygon of each type coming together at every vertex. Note that the Cayley graph for the $(2, 3, 3)$ triangle group (Figure 2) is a bounded polyhedron, so the group is finite. On the other hand, the graph for the $(3, 3, 3)$ group (Figure 3) is a tessellation, or tiling, of the plane that extends infinitely in all directions, so the group is infinite, as seen above. For an arbitrary triangle group, the algebraic question of whether or not the group is finite thus turns into a geometric question of whether we can construct a polyhedron with the appropriate faces. This type of analysis is typical of the field of mathematics known as *geometric group theory*, in which one studies groups by looking at their Cayley graphs and other geometric objects.

One of the requirements for building a polyhedron with given faces is that the sum of the angles that come together at any vertex must be strictly less than 360°. (For example, in a cube, the sum is 270°; in a tetrahedron, it is 180°.) The formula for the sum of the interior angles of an n-sided polygon is $(n-2)180°$. If all the angle measures are equal, then each angle measures $\left(\frac{n-2}{n}\right)180°$. If we take $n = 2k$, $n = 2l$, and $n = 2m$ respectively, the criterion

we assume that the graph is embedded in $\mathbb{R}^n$. The condition that the polygons be regular will be used formally in Section 2.6.

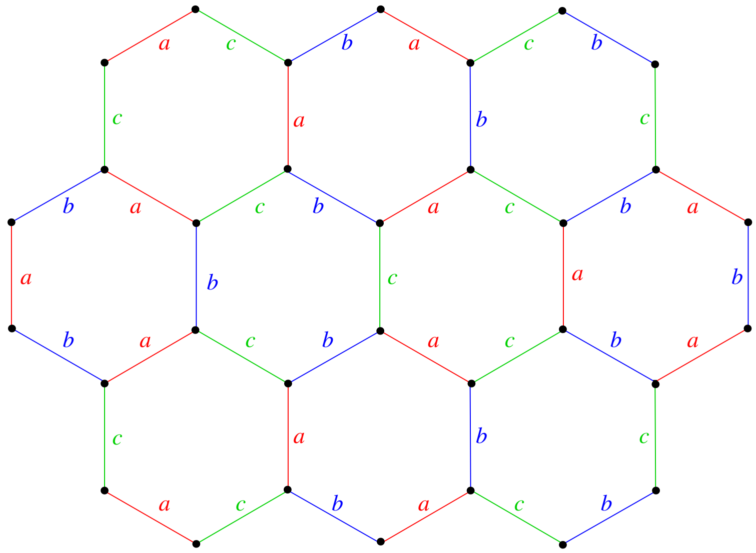

Figure 3: Cayley graph for Coxeter group generated by a, b, and c, subject to the rules $aba = bab$, $bcb = cbc$, and $aca = cac$.

for the group to be finite thus becomes the formula:

$$\left(\frac{2k-2}{2k}\right)180^\circ + \left(\frac{2l-2}{2l}\right)180^\circ + \left(\frac{2m-2}{2m}\right)180^\circ < 360^\circ.$$

After dividing both sides by 180° and rearranging terms, we can simplify this inequality to:

$$\frac{1}{k} + \frac{1}{l} + \frac{1}{m} > 1.$$

In the case where $aba = bab$, $bcb = cbc$, and $ac = ca$, discussed above, we have $\frac{1}{3} + \frac{1}{3} + \frac{1}{2} = \frac{7}{6} > 1$, so the group is finite. The Cayley graph is the 24-vertex polyhedron known as a truncated octahedron (Figure 2). On the other hand, if $aba = bab$, $bcb = cbc$, and $aca = cac$, then the formula gives $\frac{1}{3} + \frac{1}{3} + \frac{1}{3} = 1$, so the group must be infinite. In the latter example, the angle sum at each vertex is exactly 360° ($= 120^\circ + 120^\circ + 120^\circ$), as seen in the "honeycomb" Cayley graph (Figure 3).

Now consider the Coxeter group in which $ab = ba$, $bcbc = cbcb$, and $acacac = cacaca$, i.e., the $(2, 4, 6)$ triangle group. Since $\frac{1}{2} + \frac{1}{4} + \frac{1}{6} = \frac{11}{12} < 1$, the group must be infinite. In the Cayley graph, a square, an octagon (8 sides, 135° angles), and a dodecagon (12 sides, 150° angles) should come together at every vertex. However, if we try to draw this in a flat plane, we find that there is not enough room at each vertex for all three regular polygons.

Instead, we need to use the *hyperbolic plane* (denoted $\mathbb{H}^2$), one of the most fundamental constructions in non-Euclidean geometry. One way to define $\mathbb{H}^2$ is as the graph of the function $z = \sqrt{1 + x^2 + y^2}$, called a *hyperboloid*. (See Figure 4.) Given any two points in $\mathbb{H}^2$, the *hyperbolic line* between them is defined as the intersection of $\mathbb{H}^2$ with the unique

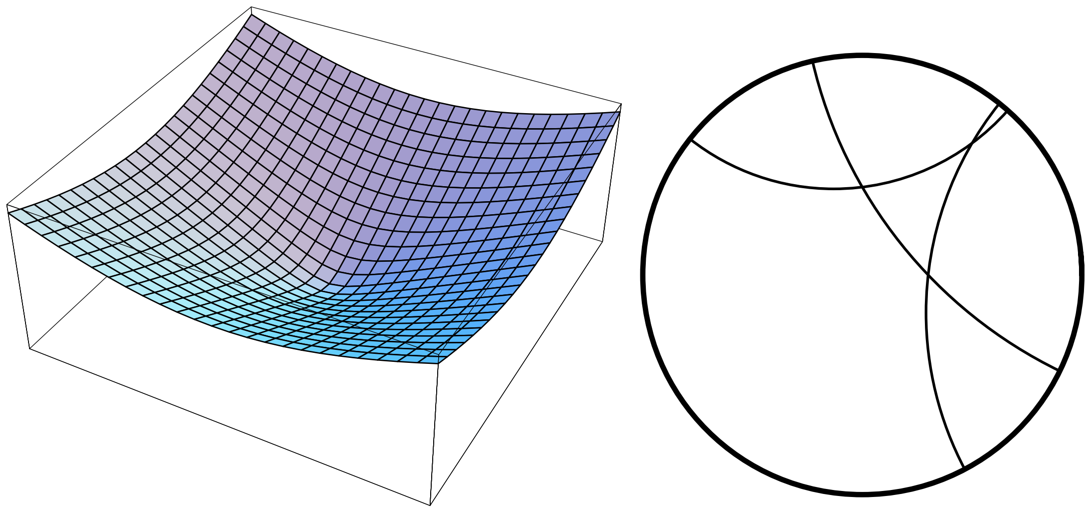

Figure 4: The hyperboloid (left) and Poincaré (right) models of the hyperbolic plane.

plane containing both points and the origin. Although these paths are not straight in the conventional sense of the word, they still satisfy many of the fundamental properties of ordinary lines in classical geometry, so we are justified in using the term "line."

Another way to visualize the hyperbolic plane (which can be taken as the definition if one prefers) is the so-called *Poincaré disk*. Let D be a circular disk. Given points p and q in D, there is a unique circle that passes through both p and q and is perpendicular to the boundary circle of D at both points of intersection. We declare that path to be the hyperbolic line joining p and q. In order for this to work, we need to use a special metric, or definition of distance, on D. The *hyperbolic distance* between two points is not the same as the distance on the page between the points lie on the page but rather is given by a complicated function. Using that function as the definition of distance, the hyperbolic line just described is in fact the shortest path between two points p and q. One property of the hyperbolic metric is that points that lie near the boundary may be very far apart, even if they appear on paper to be close together. For example, in Figure 5, the line segments all have the same length according to the hyperbolic metric, even though that certainly does not appear on paper to be true. A helpful analogy is that the Poincaré disk is to the hyperboloid as a flat map of the earth is to a globe. Because we are representing a curved surface using a flat one, some distortion of distances is inevitable, but the advantage of being able to draw complicated figures makes up for this difficulty.

In either model, the fact that the lines are "curved" yields a crucial result: the sum of the interior angles of any triangle is less than 180°. In small triangles, the sides are not severely curved, so the angles are not far from what they would be if the sides were "straight." In sufficiently large triangles, though, the angles can be made arbitrarily small. Moreover, by gluing together triangles, we find that the same is true for other polygons. For example, we can find quadrilaterals in which the sides all have the same (hyperbolic) length and in which the angles all have the same value, less than 90°. We may do the same with octagons

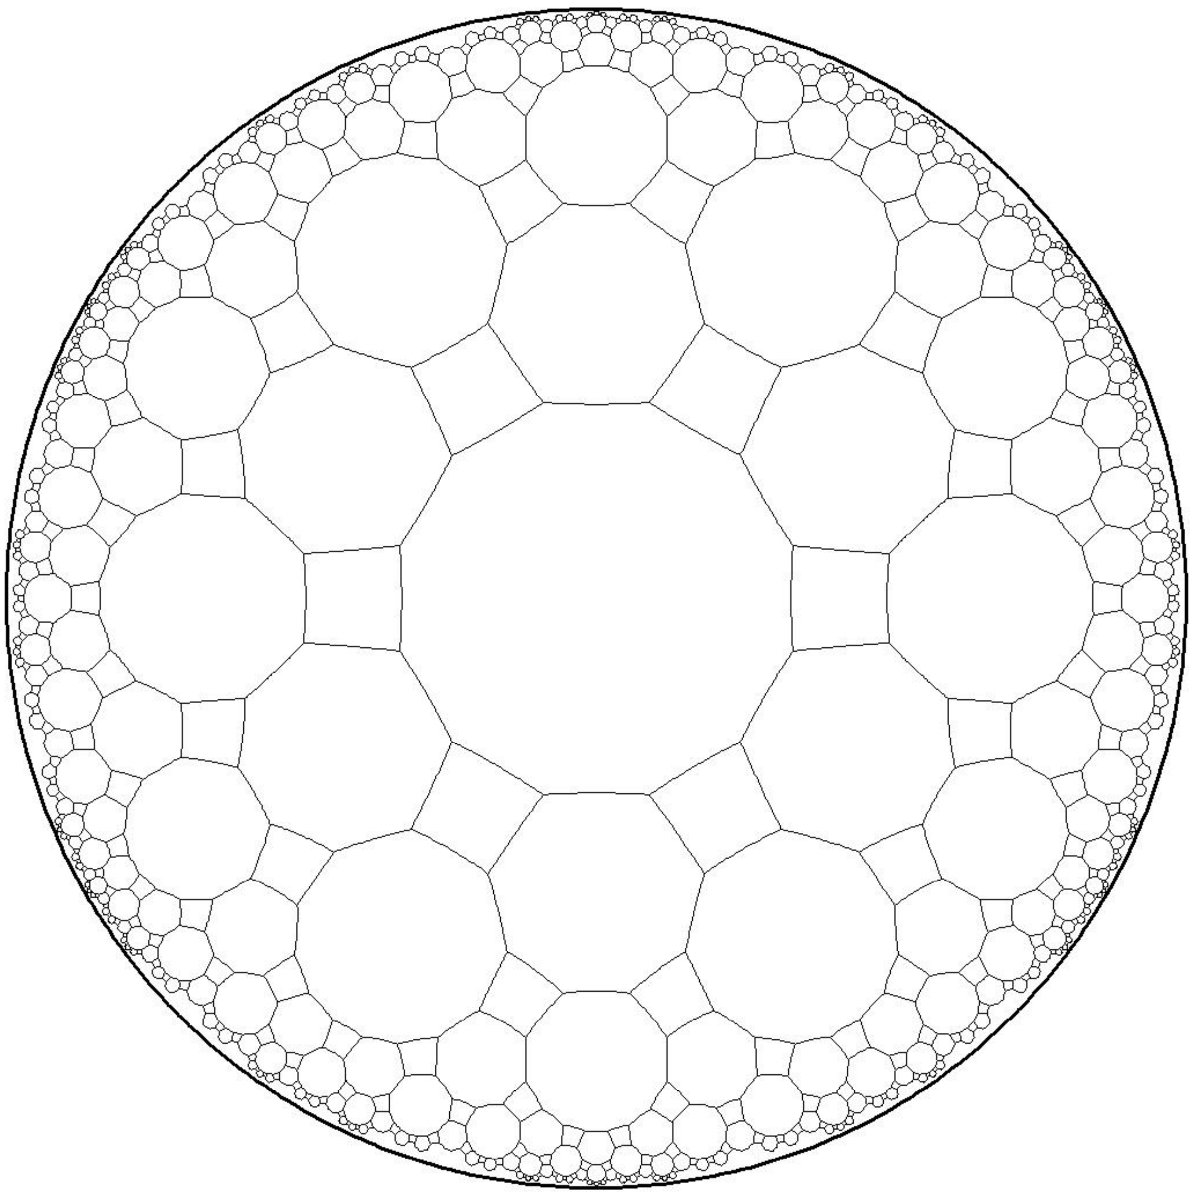

Figure 5: Cayley graph for the $(2, 4, 6)$ triangle group, embedded in the Poincaré model of the hyperbolic plane $\mathbb{H}^2$.

and dodecagons. If we choose the sizes correctly, we can fit these polygons together to form a tessellation of the hyperbolic plane with one of each type of polygon coming together at every vertex, shown in Figure 5. Notice that the sides are all slightly curved and that the angles of any polygon are all the same. According to the hyperbolic metric, the polygons are all regular, since (to reiterate) all the edges have the same length. Thus, by allowing for "curved lines," we find that it is possible to put together a tessellation of $\mathbb{H}^2$ that would be impossible in the ordinary (Euclidean) plane. The vertices and edges of this tessellation make up the Cayley graph for the $(2, 4, 6)$ triangle group.[3]

The $(2, 4, 6)$ triangle group is an example of a *word-hyperbolic group*. Since its introduction by Mikhael Gromov in the 1980s, this concept has had an enormous impact on the field of geometric group theory. In Chapter 1, I will prove some of the fundamental results concerning word-hyperbolic groups, such as the fact that they are always finitely presentable. I will discuss so-called *quasiconvex* subgroups of word-hyperbolic groups and show that these subgroups are themselves word-hyperbolic. Finally, I will discuss spaces of non-positive curvature, which are extremely important for the study of word-hyperbolic groups.

Chapter 2 deals with Coxeter groups. I will show how any Coxeter group can act on a vector space by reflections, as in the example discussed above, and use this action to prove some of the important algebraic properties of Coxeter groups. I will then prove a theorem of Jacques Tits [19] that solves the Word Problem in Coxeter groups. One of the goals is to provide a more rigorous understanding of the construction of the Cayley graph, which will then allow us to discuss Gabor Moussong's dissertation [16] about which Coxeter groups are word-hyperbolic.

In Chapter 3, I will discuss some original research that I did at the University of Illinois at Urbana-Champaign during the summer of 2004. I show that certain subgroups of certain word-hyperbolic Coxeter groups are quasiconvex, making explicit use of many of the theorems presented in the first two chapters.

I have tried to make this Introduction as readable as possible to a general audience. The rest of the thesis presumes more mathematical background, but it should still be quite accessible to anyone who has taken introductory courses in linear algebra, group theory, and topology (topics including metric spaces, fundamental groups, and simplicial and CW complexes). Additionally, some familiarity with the geometry of the hyperbolic plane is useful for Chapter 1.

[3] M.C. Escher's woodcut *Circular Limit I* (Figure 6) is based on the $(2, 4, 6)$ triangle group; note the similarity with Figure 5. Figuring out the precise connection is left to the reader as an exercise; see page 41 for the solution. Escher wrote, "I had an enthusiastic letter from Coxeter about my coloured fish, which I sent him. Three pages of explanation of what I actually did... It's a pity that I understand nothing, absolutely nothing of it" [1, pp. 100–101].

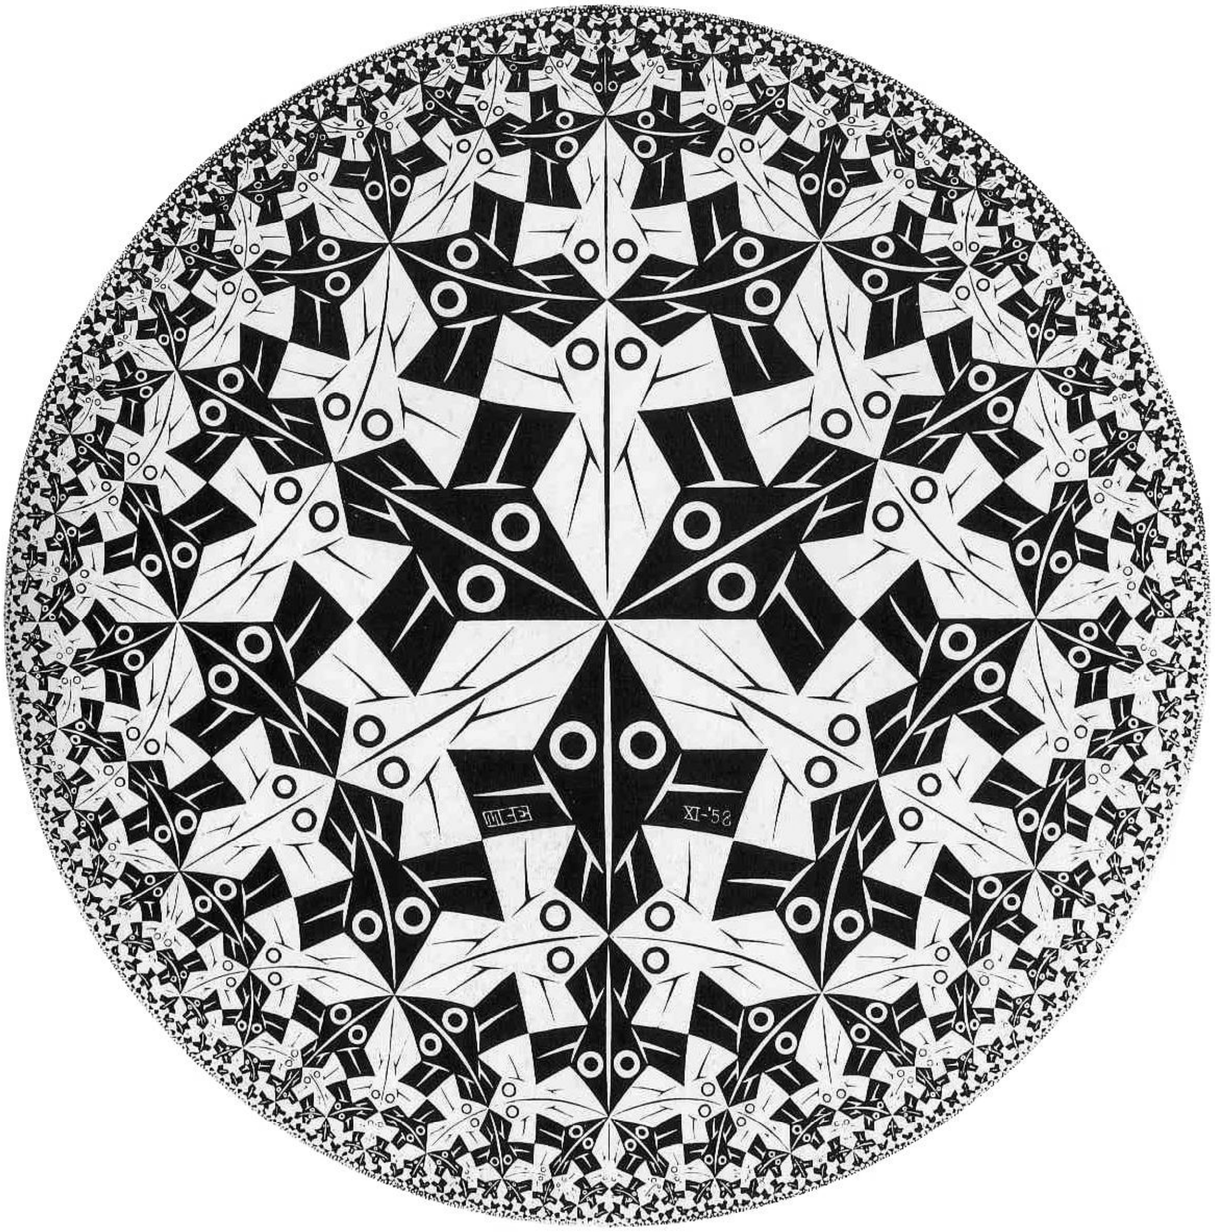

Figure 6: M.C. Escher's woodcut *Circular Limit I* [1, p. 319].

Chapter 1

The Theory of Word-Hyperbolic Groups

In geometric group theory, one studies abstractly presented groups by giving them a geometric structure or, more generally, by considering their action on various metric spaces. While many of the ideas involved (such as the Cayley graph) date back as far as the nineteenth century, the field really took off with the development of word-hyperbolic groups by Gromov and others in the 1980s. Roughly speaking, a group is word-hyperbolic if it acts nicely on a metric space that possesses certain properties of the hyperbolic plane. In the first two sections, we will prove some of the basic properties of word-hyperbolic groups and try to get a feel for the kinds of proof that one sees in geometric group theory. In particular, we will see in Section 1.3 that all word-hyperbolic groups are finitely presentable. Section 1.4 deals with quasiconvex subgroups, which are important to the original research presented in Chapter 3. Finally, in Section 1.5, we will look at CAT(κ) complexes, which are an extremely important source of examples of metric spaces on which word-hyperbolic groups can act.

1.1 δ-Hyperbolic Metric Spaces and Quasi-Isometries

We begin with an observation about the hyperbolic plane $\mathbb{H}^2$. Let Δ be a triangle with sides E_1, E_2, and E_3, and let p be a point of E_1 other than a vertex. Let r be the distance between p and $E_2 \cup E_3$, and let C be the circle of radius r centered at p, to which either E_2 or E_3 is tangent. The intersection of the interiors of C and Δ is then the interior of a semicircle. Note two important facts about the hyperbolic plane:

- The area of any triangle is equal to π minus the sum of the measures of the angles (in radians); in particular, it is strictly less than π.

- The area of a circle of radius r is at least πr^2.

Therefore, we have $\frac{1}{2}\pi r^2 \leq \pi$, so $r \leq \sqrt{2}$. In other words, every point of E_1 lies within $\sqrt{2}$ of some point in $E_2 \cup E_3$. (It is actually possible to obtain a better bound on r, but the existence of a bound is sufficient for our purposes.)

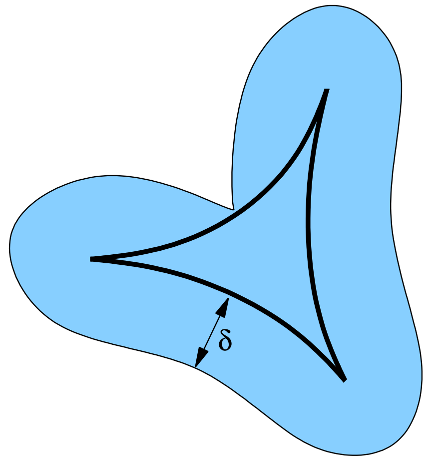

Figure 7: A δ-slim triangle.

Let (X, d) be a metric space. A *geodesic segment* is an isometric embedding of an interval of the real line into X, i.e., a continuous, injective map $f : [a, b] \hookrightarrow X$ such that $d(f(t), f(t')) = |t - t'|$ for all $t, t' \in [a, b]$. The space X is called a *geodesic space* if every pair of points can be joined by a geodesic. Unless otherwise noted, we will assume that our geodesic spaces are *proper*, i.e., that any closed ball is compact. A *geodesic triangle* is the union of three geodesic segments, every pair of which share an endpoint.

The idea of a hyperbolic metric space can be formulated in several ways. We shall use the following definition, known as the *slim triangles condition*:

Definition 1.1. Let (X, d) be a geodesic space, and let $\delta \geq 0$. A geodesic triangle in X is called *δ-slim* if each edge is contained within the δ-neighborhood of the union of the other two edges. The space X is called *δ-hyperbolic* if every geodesic triangle in X is δ-slim.

Thus, the hyperbolic plane $\mathbb{H}^2$ is δ-hyperbolic for $\delta = \sqrt{2}$. When $n \geq 2$, Euclidean space $\mathbb{E}^n$ is not δ-hyperbolic for any δ since the midpoint of one edge of a sufficiently large triangle will lie outside the δ-neighborhood of the other two edges. Finally, any tree (simply connected graph) is 0-hyperbolic since each edge of a triangle is actually contained in the union of the other two edges. Obviously, if a space is δ-hyperbolic, then it is δ'-hyperbolic for all $\delta' \geq \delta$.

A word on notation: a space is often simply called *hyperbolic* if it is δ-hyperbolic for some $\delta \geq 0$. (The word hyperbolic can have other meanings, however, so in many contexts the term δ-hyperbolic is preferable.) The same convention will apply for many other properties that will be discussed below.

Many of the concepts in hyperbolic geometry are what Bridson and Haefliger [3] call "quasifications" of previously existing geometric notions: quasi-isometry, quasiconvexity, and so on. Essentially, the quasification generalizes the original term up to linear approximation of distances. These concepts are highly compatible with the definition of hyperbolicity. Here is the first one:

Definition 1.2. Let (X, d_Y) and (Y, d_Y) be metric spaces, and let $\lambda \geq 1, \epsilon \geq 0$. A map $f : X \to Y$ (not necessarily continuous) is called a (λ, ϵ)-*quasi-isometric embedding* if for any $x, x' \in X$, we have

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

If, moreover, there exists a (λ', ϵ')-quasi-isometric embedding $g : Y \to X$ and a constant $\rho \geq 0$ such that for all $x \in X$ and $y \in Y$, we have $d_X(x, gf(x)) \leq \rho$ and $d_Y(y, fg(y)) \leq \rho$, the map f is called a *quasi-isometry*, and g is a *quasi-inverse* for f.

It is easy to check that the composition of two quasi-isometric embeddings is a quasi-isometric embedding and that the composition of two quasi-isometries is a quasi-isometry. Therefore, quasi-isometry is an equivalence relation between metric spaces. Colloquially, two spaces that are quasi-isometric look the same when viewed from a great distance. For instance, the canonical inclusion $\mathbb{Z} \hookrightarrow \mathbb{R}$ is a quasi-isometry whose quasi-inverse is "rounding to the nearest integer." On the other hand, an isometric embedding $\mathbb{R} \hookrightarrow \mathbb{R}^2$ does not have a quasi-inverse.

Let $C \geq 0$. Given a metric space Y, a subspace $Z \subset Y$ is called C-*quasi-dense* in Y if every point in Y lies within distance C of some point of Z.[1] The following proposition gives a good description of quasi-isometries:

Proposition 1.3. *Let X and Y be metric spaces. A (λ, ϵ)-quasi-isometric embedding $f : X \to Y$ is a quasi-isometry if and only if its image is C-quasi-dense in Y for some $C \geq 0$.*

Proof. Suppose the image of f is C-quasi-dense. For each $y \in Y$, if $y \in \operatorname{im} f$, then let $a_y = y$; otherwise, choose $a_y \in \operatorname{im} f$ such that $d(y, a_y) \leq C$. If $a_y = f(x_y)$, define $g : Y \to X$ by $g(y) = x_y$. (Of course, this map is not uniquely defined.) Let $y, y' \in Y$. By the triangle inequality,

$$d(y, y') - 2C \leq d(a_y, a_{y'}) \leq d(y, y') + 2C.$$

By the definition of a quasi-isometric embedding,

$$\frac{1}{\lambda} d_Y(a_y, a_{y'}) - \epsilon \leq d_X(x_y, x_{y'}) \leq \lambda d_Y(a_y, a_{y'}) + \epsilon.$$

Combining these two facts, we obtain:

$$\frac{1}{\lambda} d_Y(y, y') - \frac{2C}{\lambda} - \epsilon \leq d_X(g(y), g(y')) \leq \lambda d_Y(y, y') + 2C\lambda + \epsilon,$$

so g is a (λ', ϵ')-quasi-isometric embedding, where $\lambda' = \lambda$ and $\epsilon' = 2C\lambda + \epsilon$. To check that f and g are quasi-inverses, note that $d_Y(y, fg(y)) = d_Y(y, a_y) \leq C$ and that $d_X(x, gf(x)) \leq \lambda d_Y(f(x), fgf(x)) + \epsilon \leq \lambda C + \epsilon$, so let $\rho = \lambda C + \epsilon$.

The converse follows trivially from the requirement that $d(y, fg(y)) \leq \rho$ for all $y \in Y$. $\square$

[1] With many quasifications, taking the minimal value of the constant yields the original term. For instance, a $(1, 0)$-quasi-isometric is an isometric embedding, and (as we will see in Section 1.4) 0-quasiconvex means convex. However, this is not the case for quasi-density. If Z is 0-quasi-dense in Y, then $Y = Z$. The statement "Z is dense in Y" is properly quasified as "Z is C-quasi-dense in Y for all $C > 0$."

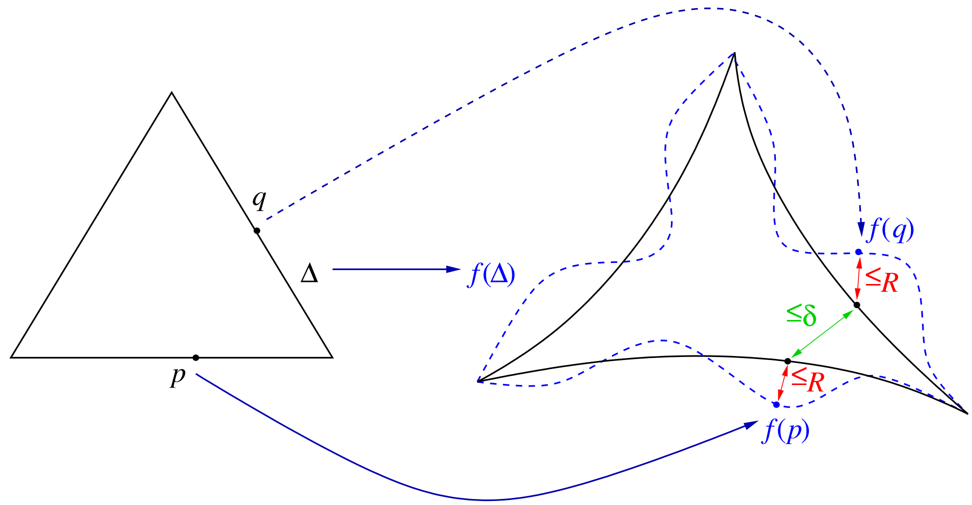

Figure 8: Schematic of the proof of Theorem 1.4.

The key fact is that hyperbolicity is an invariant of quasi-isometry. Specifically:

Theorem 1.4. *Let $\lambda \geq 1$, $\epsilon \geq 0$, and $\delta \geq 0$. There exists a constant $\delta' \geq 0$, determined solely by λ, ϵ, and δ, such that the following holds:*

Let X and Y be geodesic spaces, and let $f : X \to Y$ be a (λ, ϵ)-quasi-isometric embedding. If Y is δ-hyperbolic, then X is δ'-hyperbolic.

To prove this theorem, we will need to consider the image of a geodesic triangle in X. Analogous to the definition of a geodesic, a (λ, ϵ)-*quasi-geodesic* is defined as a (λ, ϵ)-quasi-isometric embedding of an interval in $\mathbb{R}$. The map f then sends each edge of the triangle to a quasi-geodesic in Y. The following lemma then puts a control on the behavior of these quasi-geodesics:

Lemma 1.5 (Stability of quasi-geodesics). *Let $\lambda \geq 1$, $\epsilon \geq 0$, and $\delta \geq 0$. There exists a constant $R \geq 0$, determined solely by λ, ϵ, and δ, such that the following holds:*

Let Y be a δ-hyperbolic geodesic space, and let $g : [a, b] \to Y$ be a (λ, ϵ)-quasi-geodesic. Say that $g(a) = p$ and $g(b) = q$, and let $[p, q]$ be a geodesic segment joining p and q. Then $[p, q]$ lies in the R-neighborhood of the image of g, and the image of g lies in the R-neighborhood of $[p, q]$.

The proof of this lemma requires some very involved chasing of δ's, λ's, and ϵ's, so we omit it. Essentially, one first "tames" an arbitrary quasi-geodesic by showing that it lies close to a continuous quasi-geodesic with the same endpoints and then shows that the latter must satisfy the lemma using some general properties of continuous curves in hyperbolic spaces. Complete proofs can be found in Bridson and Haefliger [3, pp. 403–405] or Ghys and de la Harpe [8, pp. 82-87].

Proof of Theorem 1.4. Let $\Delta \subset X$ be a geodesic triangle with edges E_1, E_2 and E_3, and let $p \in E_1$ (without loss of generality). The image of Δ under f is the union of three (λ, ϵ)-quasi-geodesics, g_1, g_2, and g_3, which agree at their endpoints. (See Figure 8). Let $\Delta' \subset Y$ be a geodesic triangle with the same vertices as $f(\Delta)$; denote the edges of Δ' by E'_1, E'_2, and E'_3. By Lemma 1.5, there exists a point $q \in E_1$ such that $d_Y(f(p), q) \leq R$. As Δ' is δ-slim, there exists a point $q' \in E'_2 \cup E'_3$ such that $d_Y(q, q') \leq \delta$. Lemma 1.5 then gives a point $p' \in E_2 \cup E_3$ such that $d_Y(f(p'), q') \leq R$. By the triangle inequality, we have $d_Y(f(p), f(p')) \leq \delta + 2R$. As f is a (λ, ϵ)-quasi-isometric embedding, we have $d_X(p, p') \leq \lambda d_X(f(p), f(p')) + \lambda\epsilon$. We thus see that the triangle Δ is δ'-slim, where $\delta' = \lambda(\delta + 2R + \epsilon)$. This value does not depend on our choice of Δ, so X is δ'-hyperbolic. $\square$

1.2 Word-Hyperbolic Groups

In this section, we will show how to turn a group into a metric space to which the results of the previous section are applicable.

Let G be a group and A a finite generating set for G. We may assume if $a \in A$ and its inverse a^{-1} are distinct elements of G, then only one of them, say a, lies in A. Each element $g \in G$ can be written as a product of elements of A and their inverses: $g = a_1 \cdots a_r$, where $a_i \in A^{\pm 1}$. Define the *length function* $\ell_A : G \to \mathbb{Z}$ as follows: For any word $g \in G$, let $\ell_A(g)$ be the minimum value of r for which g can be written as a product of r elements of $A^{\pm 1}$. If $g = a_1 \cdots a_r$, then g^{-1} can be written as $a_r^{-1} \cdots a_1^{-1}$, so $\ell_A(g^{-1}) \leq \ell_A(g)$; interchanging the roles of g and g^{-1}, we see that $\ell_A(g^{-1}) = \ell_A(g)$. Also, $\ell_A(g) = 0$ if and only if g is the identity element of A.

We may then define a function $d_A : G \times G \to \mathbb{Z}_{\geq 0}$ by $d_A(g, g') = \ell_A(g^{-1}g')$. The previous paragraph shows that this function is symmetric and that $d_A(g, g') = 0$ if and only if $g = g'$. Additionally, if $d(g, g') = k$ and $d(g', g'') = l$, note that $g^{-1}g'' = g^{-1}g'g'^{-1}g''$, which has a representation of $k + l$ letters, so $d(g, g'') \leq k + l$. Thus, the triangle inequality holds, so d_A is a metric on G, the *word metric with respect to A*. This metric is invariant under left multiplication: for any $g, g', h \in G$, we have $d(g, g') = d(hg, hg')$ since $(hg)^{-1}(hg') = g^{-1}h^{-1}hg' = g^{-1}g'$.

Since the word metric only takes integer values, the metric space (G, d_A) is discrete. As the previous section's results apply only to geodesic spaces, it is useful to consider a larger space, the *Cayley graph* $\mathcal{C}_A(G)$. This is a graph with vertex set $\{v_g \mid g \in G\}$ and an edge joining v_g and v'_g if and only if $d_A(g, g') = 1$. Without loss of generality, we then have $g^{-1}g' = a$ for some $a \in A$; the edge joining v_g and $v_{g'}$ is then labelled a and is oriented *from v_g to $v_{g'}$*. Note that if a has order 2 in G, then $g'^{-1}g = a^{-1} = a$, so the edge is simultaneously given the reverse orientation. Every vertex then has exactly one oriented edge coming in and one oriented edge going out for each element of A.[2]

[2]Many authors define the edge set of $\mathcal{C}_A(G)$ slightly differently: there is an oriented edge from v_g to v_{ga} for every $g \in G, a \in A$. The practical difference between these two definitions occurs when a generator a has order 2 in G. Using the other definition, the vertices v_g and v_{ga} are then joined by two distinct

Fix a vertex v of $\mathcal{C}_A(G)$ as a base point. Edge-paths in $\mathcal{C}_A(G)$ starting at the vertex v correspond bijectively with words in the letters $A^{\pm 1}$. (Traversing an edge labelled a in the direction opposite its orientation corresponds to the letter a^{-1}.) Two words represent the same element of G if and only if the corresponding paths end at the same vertex. We metrize $\mathcal{C}_A(G)$ by giving each edge length 1 and defining the distance between two points to be the length of the shortest edge-path joining them. The restriction of this metric to the vertices then agrees with the word metric d_A on G; in other words, the inclusion $G \hookrightarrow \mathcal{C}_A(G)$ is an isometric embedding. Every point of $\mathcal{C}_A(G)$ lies within distance $\frac{1}{2}$ of a vertex, so this inclusion is a quasi-isometry between G and $\mathcal{C}_A(G)$.

The Cayley graph is a geodesic metric space, motivating the following definition:

Definition 1.6. A group G is called *word-hyperbolic* (or simply *hyperbolic*) with respect to a finite generating set A if the Cayley graph $\mathcal{C}_A(G)$ is δ-hyperbolic for some $\delta \geq 0$.

The dependence on the choice of generating set disappears in light of Theorem 1.4 and the following lemma:

Lemma 1.7. *Let G be a group, and let A and B be two finite generating sets for G. Then the Cayley graphs $\mathcal{C}_A(G)$ and $\mathcal{C}_B(G)$ are quasi-isometric.*

Proof. It suffices to show that the metric spaces (G, d_A) and (G, d_B) are quasi-isometric. Each generator $a \in A$ can be written as a reduced word in the letters of B, and vice versa. Let λ be the maximum length of all these words. Any element $w \in G$ with a reduced expression $w = a_1 \cdots a_k$ in the letters of A can then be written as a word of length at most λk in the letters of B, so $\ell_B(w) \leq \lambda \ell_A(w)$, and likewise $\ell_A(w) \leq \lambda \ell_B(w)$. For any $g, g' \in G$, we then have

$$\frac{1}{\lambda} d_B(g, g') \leq d_A(g, g') \leq \lambda d_B(g, g'),$$

meaning that the identity map is a $(\lambda, 0)$-quasi-isometry between (G, d_A) and (G, d_B). $\square$

In other words, the geometric structure of a finitely generated group is well defined up to quasi-isometry. In particular, this allows us to state unambiguously whether or not a group is word-hyperbolic, since the choice of generating set does not matter. The constant of hyperbolicity, δ, could vary, but the existence or nonexistence of such a constant does not.

For example, any free group is hyperbolic, since its Cayley graph is a tree. Any finite group is hyperbolic, since we can take δ to be equal to the diameter of the Cayley graph. On the other hand, the free abelian group $\mathbb{Z}^n$ is quasi-isometric to $\mathbb{R}^n$ and therefore is not hyperbolic for $n \geq 2$. Indeed, using a certain group action, one can show that no hyperbolic group contains any subgroup isomorphic to $\mathbb{Z}^n$. See Ghys and de la Harpe [8, p. 157].

More generally, one may learn a great deal about the geometry of a group up to quasi-isometry by considering the group's action on a geodesic space. In general, suppose that a

edges, one with each orientation. In our definition, which follows Coornaert *et al.* [4], there is there is only one (doubly-oriented) edge between these vertices. In any case, the two versions of the Cayley graph are quasi-isometric and essentially function identically.

group G acts on a topological space X. The action of G on X is called *proper* (or *properly discontinuous*) if for every compact subspace $K \subset X$, the set of elements $g \in G$ such that $K \cap g.K \neq \varnothing$ is finite. It is called *cocompact* if there exists a compact subspace $K \subset X$ such that $G.K = X$, i.e., if the translates of K cover X. (If the action is cocompact, then the orbit space X/G is compact.) If X is a metric space and each element of G is an isometry of X, then the action is *by isometries*. These three properties prove to be extremely powerful when the space X is a geodesic space, in view of the following theorem [3, p. 140]:

Theorem 1.8 (Svarc-Milnor Lemma). *Let G be a group and X a proper geodesic space. If G acts properly and cocompactly by isometries on X, then G is finitely generated, and for any $x_0 \in X$, the map $G \to X$ defined by $g \mapsto g.x_0$ is a quasi-isometry.*

Proof. Fix $x_0 \in X$. As the action of G on X is cocompact, there exists a compact set $C \subset X$ whose orbit covers X; by translation, we may assume that $x_0 \in C$. Any compact subspace of a metric space is bounded, so there exists $R > 0$ such that C is contained in the closed ball $B(x_0, R/3)$; obviously, the orbit of this ball also covers X. Let $A \subset G$ be the set of elements $g \in G$ for which the intersection $g.B(x_0, R) \cap B(x_0, R)$ is nonempty. As G acts properly on X and $B(x_0, R)$ is compact, the set A is finite.

Let $g \in G$, and choose a geodesic $c : [0, a] \to X$ from x_0 to $g.x_0$. By division with remainder, we have $(n-1)R/3 \leq a < nR/3$ for some $n \geq 1$. Let $t_i = iR/3$ for $0 \leq i < n$, and let $t_n = a$; then $d(c(t_{i-1}), c(t_i)) \leq R/3$. For each i, there exists $g_i \in G$ such that $d(g_i.x_0, c(t_i)) \leq R/3$. In particular, choose $g_0 = 1$ and $g_n = g$. By the triangle inequality, $d(g_{i-1}.x_0, g_i.x_0) \leq R$. As G acts by isometries, we have $d(x_0, g_{i-1}^{-1}g_i.x_0) \leq R$, so the element $a_i = g_{i-1}^{-1}g_i$ is contained in A. Therefore,

$$g = g_0(g_0^{-1}g_1)(g_1^{-1}g_2)\cdots(g_{n-1}^{-1}g_n) = a_1 \cdots a_n,$$

so A generates G.

As g can be written as a product of n elements of A, we have $d_A(1, g) \leq n$. By the construction of n, we have $(n-1)R/3 \leq d(x_0, g.x_0)$, so

$$d_A(1, g) \leq \frac{3d(x_0, g.x_0)}{R} + 1.$$

At the same time, since A is finite, let $D = \max\{d(x_0, a.x_0) \mid a \in A\}$. By taking an A-reduced expression for g and using the triangle inequality, we see that $d(x_0, g(x_0)) \leq D d_A(1, g)$. This proves that the map $g \mapsto g.x_0$ is a $(\lambda, 1)$-quasi-isometric embedding of G into X, where $\lambda = \max\{1, 3/R\}$. Moreover, since the orbit of x_0 is $R/3$-quasi-dense in X, the map is a quasi-isometry. ∎

The action of a group G on itself by left multiplication extends naturally to an action on $\mathcal{C}_A(G)$ by isometries. It is easy to check that this action is proper and cocompact. Therefore:

Corollary 1.9. *A group is word-hyperbolic if and only if it acts properly, cocompactly, and by isometries on a proper geodesic space that is δ-hyperbolic for some $\delta \geq 0$.* ∎

The three criteria of Theorem 1.8 — that a group action be proper, cocompact, and by isometries — appear constantly in geometric group theory. Often, one seeks to show that a word-hyperbolic group acts on a space that has properties more specific than merely being δ-hyperbolic. We shall see an example of this approach in the next section.

1.3 The Rips Complex and Finite Presentations

Let G be a group. A *group presentation* for G consists of a generating set A and a subset R of the free group $F(A)$ such that G is isomorphic to the quotient of $F(A)$ by the normal closure of R (the smallest normal subgroup of $F(A)$ containing R): $G \approx F(A)/N(R)$. If so, we write $G = \langle A \mid R \rangle$. In other words, G is the quotient obtained from $F(A)$ by imposing the relations $r = 1$ for each $r \in R$. Every group has a presentation, since we can take A to be the entire group and R to consist of elements of the form abc^{-1} for every pair $a, b \in G$. Obviously, many different presentations are possible for a given group; the goal of combinatorial group theory is to determine the properties of a group given a presentation.

A presentation is called *finite* if both A and R are finite sets; a group that admits a finite presentation is called called *finitely presentable.* Determining whether a finitely generated group is finitely presentable is an extremely important question in combinatorial group theory. Not every finitely generated group is finitely presentable, as the following example shows.

Example 1.10. Let F_2 be a free group of rank 2 with basis $\{a_1, a_2\}$, and let G be the direct product $F_2 \times F_2$. Let $h : G \to \mathbb{Z}$ be the homomorphism that sends each of the generators $(a_i, 1)$ and $(1, a_i)$ to $1 \in \mathbb{Z}$. Denote the kernel of this map by SB_2 (after John Stallings and Robert Bieri, who investigated a family of groups including this one). It is not hard to show that SB_2 is generated by

$$\{(a_1 a_2^{-1}, 1), (1, a_1 a_2^{-1}), (a_1, a_1^{-1}), (a_2, a_2^{-1})\}$$

and that the inclusion $\mathrm{SB}_2 \hookrightarrow G$ is an isometric embedding. However, using group homology, one can show that SB_2 is not finitely presentable. (See Bridson and Haefliger [3, p. 483].)

The elements $(a_1 a_2^{-1}, 1), (1, a_1 a_2^{-1}) \in \mathrm{SB}_2$ commute and each have infinite order, so they generate a subgroup that is isomorphic to $\mathbb{Z}^2$. Therefore, SB_2 is not hyperbolic, by the remark in the previous section.

Indeed, we have the following theorem:

Theorem 1.11. *Every word-hyperbolic group is finitely presentable.*

To prove this fact, we will make use of the following construction:

Definition 1.12. Let (X, d) be a metric space, and let $r > 0$. The *Rips complex* $P_R(X)$ is the simplicial complex in which every $(n+1)$-element subset $Y \subset X$ of diameter at most r spans an n-simplex.

A key fact about this complex is the following theorem of Rips, whose proof closely follows the one in Bridson-Haefliger [3, pp. 468–470].

Lemma 1.13. *Let Y be a δ-hyperbolic geodesic space, and let X be an r-quasi-dense subspace. For $R \geq 4\delta + 6r$, the Rips complex $P_R(X)$ is contractible.*

Proof. By Whitehead's theorem (a fundamental result in homotopy theory), a CW complex is contractible if and only if all its homotopy groups are trivial [10, p. 346]. As the sphere S^n is compact, the image of any map $S^n \to P_R(X)$ lies in a finite subcomplex of $P_R(X)$. Therefore, it suffices to show that every finite subcomplex L can be contracted to a point in $P_R(X)$. Let $V \subset X$ be the set of points that are vertices of L, and fix a basepoint $x_0 \in X$.

First, suppose that for every vertex $v \in V$, we have $d(x_0, v) \leq \frac{R}{2}$. (Bear in mind that the metric is on X, not L or $P_R(X)$.) Then the points in V span a simplex σ in $P_R(X)$, in which L is contained, so we can contract L to a point.

Now, assume that v is a vertex of L such that $d(x_0, v) > \frac{R}{2}$. Choose v such that this distance is maximal. We will deform L by a homotopy that fixes all the other vertices and moves v to a vertex v' that is closer to x_0 by a definite distance. Performing a finite number of these homotopies then yields the previous case.

Consider a geodesic $[x_0, v]$ in Y, and let y be a point on this geodesic such that $d(v, y) = \frac{R}{2}$. Choose a point $v' \in X$ such that $d(y, v') \leq r$. (The points y and v' are not necessarily distinct.) Let $\rho = d(v, v')$. By the triangle inequality, $\rho \leq \frac{R}{2} + r$ and $\rho \geq \frac{R}{2} - r \geq 2\delta + 2r$. We then have:

$$\begin{aligned} d(x_0, v') &\leq d(x_0, y) + d(y', y) \\ &\leq d(x_0, v) - \frac{R}{2} + r \\ &\leq d(x_0, v) - 2\delta - 2r, \end{aligned}$$

so v' is closer to x_0 than v is.

We shall now show that if another vertex u of L with $d(u, v) \leq R$, then $d(u, v') \leq R$ as well. Let Δ be a geodesic triangle with vertices x_0, u, and v, where the edge joining x_0 and v is the geodesic from the previous paragraph. This triangle is δ-slim, so y lies within the δ-neighborhood of the union of the other two sides, $[x_0, u] \cup [u, v]$. We consider two cases, illustrated in Figure 9:

1. There exists a point $u' \in [x_0, u]$ such that $d(y, u') \leq \delta$. Using the assumption that $d(x_0, v)$ is maximal and the triangle inequality, we have:

$$\begin{aligned} d(x_0, u') + d(u', u) &= d(x_0, u) \\ &\leq d(x_0, v) \\ &\leq d(x_0, u') + d(u', y) + d(y, v') + d(v', v) \end{aligned}$$

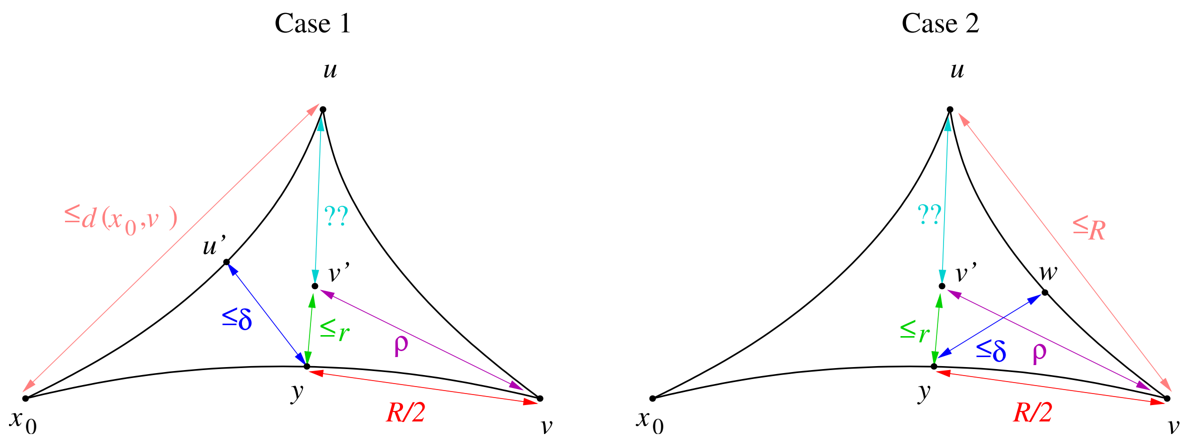

Figure 9: The two cases in the proof of Lemma 1.13.

Subtracting $d(x_0, u')$ from both sides, and noting that $d(u', y) \leq \delta$ and $d(y, v') \leq r$, we obtain $d(u, u') \leq \rho + \delta + r$. Then:

$$\begin{aligned} d(u, v') &\leq d(u, u') + d(u', y) + d(y, v') \\ &\leq \mu + 2\delta + 2r \\ &\leq \frac{R}{2} + 2\delta + 3r \\ &\leq R. \end{aligned}$$

2. There exists a point $w \in [u, v]$ such that $d(y, w) \leq \delta$. By the triangle inequality,

$$\rho = d(v, v') \leq d(v, w) + d(w, y) + d(y, v') \leq d(v, w) + \delta + r.$$

Then $d(u, w) = d(u, v) - d(v, w) \leq R - \rho + \delta + r$, so:

$$\begin{aligned} d(u, v') &\leq d(u, w) + d(w, y) + d(y, v') \\ &\leq R - \rho + 2\delta + 2r \\ &\leq R. \end{aligned}$$

The *star* of a vertex in a simplicial complex is the union of all closed simplices that contain that vertex. We have thus shown that every vertex in the star of v is also in the star of v'. Let $L' \subset P_R(X)$ be the subcomplex obtained from L by replacing each simplex of the form $[v, x_1, \ldots, x_k]$ by $[v', x_1, \ldots, x_k]$. For each such pair, there is an natural homotopy $\Delta^k \times I \to P_R(X)$ between the characteristic maps of $[v, x_1, \ldots, x_k]$ and $[v', x_1, \ldots, x_k]$. These homotopies, combined with the identity maps on the simplices not in the star of v (or v'), give a homotopy between the inclusions $L \hookrightarrow P_R(X)$ and $L' \hookrightarrow P_R(X)$. By applying this procedure finitely many times, we obtain a complex whose vertices all lie within the $\frac{R}{2}$-neighborhood of x_0, which can then be contracted as in the first case. Thus, $P_R(X)$ is contractible for $R \geq 4\delta + 6r$. $\square$

Now let G be a group with finite generating set A. Give G the word metric with respect to A, and consider the Rips complex $P_R(G) = P_R(G, A)$. Note that this complex is finite-dimensional and locally finite, since G contains only finitely many elements of a given length. If G is word-hyperbolic, say that the Cayley graph $\mathcal{C}_A(G)$ is δ-hyperbolic. The image of G in $\mathcal{C}_A(G)$ is $\frac{1}{2}$-quasi-dense, since each edge of $\mathcal{C}_A(G)$ has length 1, so $P_R(G)$ is contractible for $R \geq 4\delta + 3$.

As the metric $d = d_S$ is invariant under left multiplication, the diameter of the subset $Y = \{g_0, \ldots, g_n\}$ is equal to that of $hY = \{hg_0, \ldots, hg_n\}$. Hence, left multiplication by $h \in G$ determines a simplicial automorphism of $P_R(G)$. This defines a action of G on $P_R(G)$ by simplicial automorphisms: $\Phi : G \to \mathrm{Aut}(P_R(G))$.

Lemma 1.14. *The action of G on $P_R(G)$ by left multiplication is faithful, proper, and cocompact, and the stabilizer of any point is finite.*

Proof. First of all, note that G acts freely (i.e., without fixed points) and transitively on the vertices of $P = P_R(G)$, since the vertices are simply the elements of G. A nontrivial element $g \in G$ does not fix any vertices, so $\Phi(g)$ is not the identity. This shows Φ is faithful.

Let σ, σ' be simplices. If $g.\sigma \cap \sigma'$ is nonempty, then $\Phi(g)$ sends some subset of the vertices of σ to vertices of σ'. Since G acts freely on the vertices of P, there are only finitely many such elements g. Any compact set $C \subset P$ intersects finitely many simplices. Therefore, the set

$$\{g \in G \mid g.C \cap C \neq \varnothing\}$$

is finite, meaning that Φ is proper.

As $P_R(G)$ is locally finite, the star K of the vertex corresponding to the identity element of G is a finite subcomplex of $P_R(G)$, hence compact. The translates $g.K$ are simply the stars of the other vertices. As G acts transitively on the vertices, every cell of $P_R(G)$ lies within some translate of K, so the translates cover $P_R(G)$: $G.K = P_R(G)$. Thus, the action is cocompact.

Let $p \in P$, and let σ be the simplex of lowest dimension containing p. As the action of G on P is simplicial, any element $g \in G$ that fixes p must permute the vertices of σ. Since G acts freely on the vertices of P, there are only finitely many such elements g. Thus, the stabilizer of p is finite. $\square$

Partial proof of Theorem 1.11. For an arbitrary hyperbolic group G, one can use the contractibility of the Rips complex to find an explicit finite presentation for G; see Ghys and de la Harpe [8, pp. 75–77]. Here we present a much more interesting topological proof that holds when G is torsion-free.

Note that if the action $g \in G$ has a fixed point p, then g is an element of the finite stabilizer subgroup $\mathrm{Stab}(p)$, so g is a torsion element. Therefore, if G is torsion-free, it must act freely on P: for any simplex σ of P and any nontrivial element $g \in G$, we must have $g.\sigma \cap \sigma = \varnothing$. Let P'' be the second barycentric subdivision of P. The action of G on P extends naturally to a simplicial action of G on P'' that satisfies all the properties of Lemma 1.14. Additionally, the translates of the star of any point are all disjoint. Thus, the action is

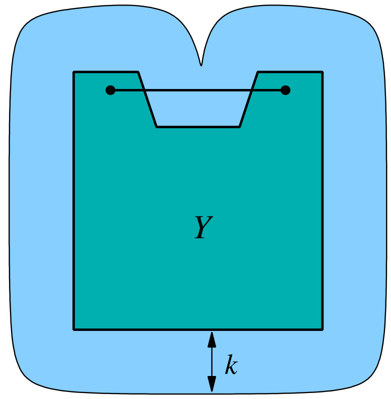

Figure 10: A k-quasiconvex subspace Y.

what Hatcher [10, p. 72] calls a *covering space action*, which implies that the quotient map $P'' \to P''/G$ is a normal covering space. In particular, for R sufficiently large, P'' is the universal cover of P''/G, so $\pi_1(P''/G) \approx G$.

The orbit space P''/G is a simplicial complex, and the map $P'' \to P''/G$ is simplicial. (This is not the case for the non-subdivided complex P.) As the action of G is cocompact, P''/G is a compact, finite-dimensional complex and therefore is finite. Note that for any CW complex X, the fundamental group $\pi_1(X)$ is determined by the 2-skeleton X^2. In particular, $\pi_1(X)$ is isomorphic to the quotient of $\pi_1(X^1)$ by the normal subgroup generated by the images of the attaching maps of the 2-cells. In our case, the complex P''/G contains only finitely many simplices, so its fundamental group G is finitely presentable. $\square$

When G is torsion-free, the quotient space P''/G is in fact an Eilenberg-Maclane space $K(G,1)$, since its fundamental group is G and its universal cover is contractible. The existence of a finite $K(G,1)$ implies a number of important facts about the group cohomology of G. For instance, all but finitely many of the rational cohomology groups $H^k(G;\mathbb{Q}) = H^k(K(G,1);\mathbb{Q})$ are zero. This fact is actually true for all word-hyperbolic groups, although proving it in the general case requires a more complicated argument using spectral sequences. (See Ghys and de la Harpe [8, pp. 74–75].)

1.4 Quasiconvexity

We now introduce another one of the "quasifications" mentioned earlier. Recall that a subspace Y of a geodesic space X is called *convex* if any geodesic joining two points of Y lies entirely within Y. We generalize this notion as follows:

Definition 1.15. Let $k \geq 0$. A subspace $Y \subset X$ is called k-*quasiconvex* if any geodesic joining two points of Y lies entirely within the k-neighborhood of Y.

A few examples in $\mathbb{R}^2$ should help illustrate this concept. Any bounded subspace Y is quasiconvex, since for sufficiently large k, the k-neighborhood of Y contains a circular

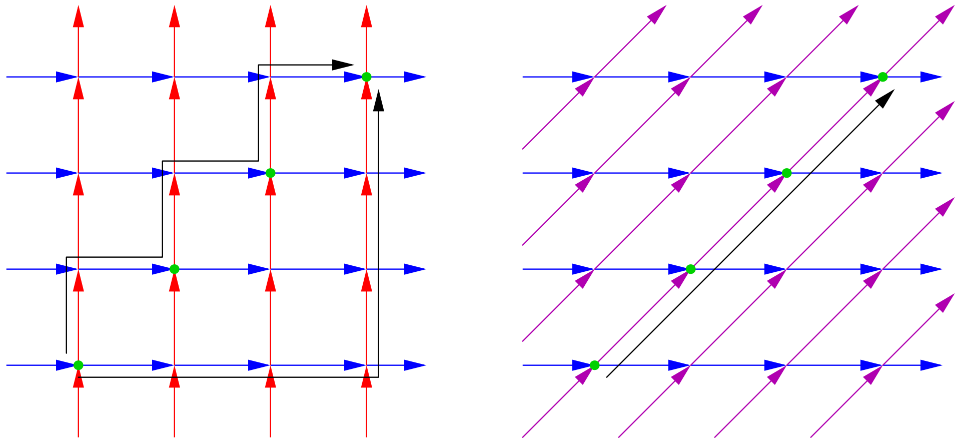

Figure 11: The Cayley graphs $\mathcal{C}_A(\mathbb{Z}^2)$ (left) and $\mathcal{C}_B(\mathbb{Z}^2)$ (right) for Example 1.16.

disk that contains Y. For a non-bounded example, note that the graph of $y = \sin x$ is 2-quasiconvex, since any segment joining two points on that graph lies within the region $\mathbb{R} \times [-1, 1]$, and any point in that region lies within the 2-neighborhood of the curve. On the other hand, the union of the two coordinate axes is not quasiconvex. For any $k > 0$, a segment joining the points $(2k, 0)$ and $(0, 2k)$ contains the point $(k\sqrt{2}, k\sqrt{2})$, which lies outside the k-neighborhood of the axes.

Now let G be a group with finite generating set A. Naturally, we define a subgroup H to be *quasiconvex with respect to* A if the set of vertices of $\mathcal{C}_A(G)$ corresponding to H is quasiconvex in $\mathcal{C}_A(G)$. Unfortunately, this definition, in general, depends on the choice of generating set, as the following example shows:

Example 1.16. Let H be the subgroup of $\mathbb{Z}^2$ generated by $(1, 1)$. Let $A = \{(1, 0), (0, 1)\}$ be the standard basis for $\mathbb{Z}^2$, and consider the Cayley graph $\mathcal{C}_A(\mathbb{Z}^2)$. One of the geodesics joining $(0, 0)$ and (k, k) is a path that goes k units to the right and then k units up. The point $(k, 0)$, which lies on this geodesic, is then distance k from the subspace H. Since k can be made arbitrarily large, H is not quasiconvex. On the other hand, if $B = \{(1, 0), (1, 1)\}$, then H is in fact convex in $\mathcal{C}_B(\mathbb{Z}^2)$. (See Figure 11.)

Thus, for arbitrary groups, quasiconvexity does not at first glance seem to be a very meaningful property. On the other hand, for hyperbolic groups, the situation is much better:

Lemma 1.17. *Let $f : X \to Y$ be a (λ, ϵ)-quasi-isometric embedding between geodesic spaces, and suppose Y is δ-hyperbolic. Let $Z \subset X$ be a k-quasiconvex subspace. Then the image of Z is quasiconvex in Y.*

Proof. Let $y_1, y_2 \in Y$ be points in the image of Z: $y_i = f(x_i)$ for $x_i \in Z$. Any geodesic $[x_1, x_2]$ in X lies within the k-neighborhood of Z, so $f([x_1, x_2])$ lies within the $(\lambda k + \epsilon)$-neighborhood of $f(Z)$. Note that $f([x_1, x_2])$ is a (λ, ϵ)-quasi-geodesic in Y, so by Lemma 1.5, any geodesic $[y_1, y_2]$ in Y lies within the R-neighborhood of $f([x_1, x_2])$, where R is a constant determined solely by δ, λ, and ϵ. Therefore, the subspace $f(Z)$ is $(\lambda k + \epsilon + R)$-quasiconvex in Y. $\square$

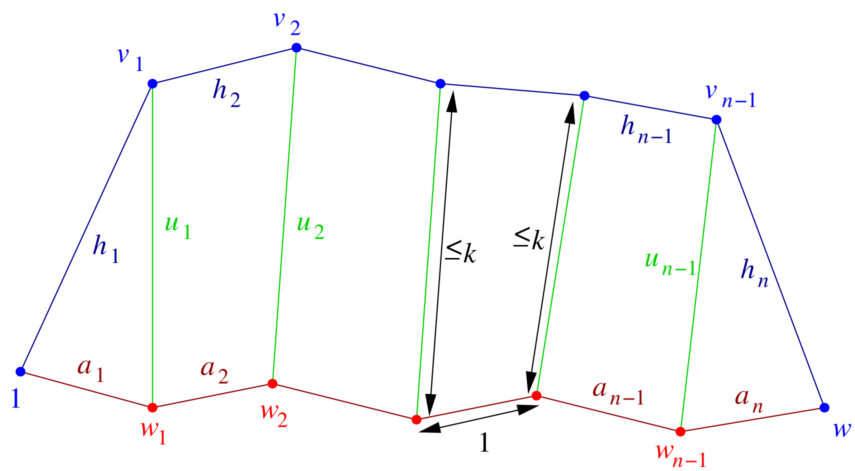

Figure 12: Schematic of the proof of Theorem 1.18.

As an immediate consequence, we see that if G is a hyperbolic group and H a subgroup that is quasiconvex with respect to a given generating set A, then H is quasiconvex with respect to all finite generating sets, by Lemma 1.7. Thus, in the hyperbolic case, we can say unambiguously whether or not a subgroup is quasiconvex. As a result, we may expect that quasiconvexity is connected to intrinsic algebraic properties that do not depend on the choice of generators. This is indeed the case.

Theorem 1.18. *Let G be a hyperbolic group and H a quasiconvex subgroup. Then H is finitely generated, the inclusion $H \hookrightarrow G$ is a quasi-isometric embedding, and H is word-hyperbolic.*

Proof. Fix a finite generating set A for G, and suppose that H is k-quasiconvex in $\mathcal{C}_A(G)$. Let $w \in H$, and let $w = a_1 \cdots a_n$ be a reduced expression for w in the letters of A. This word corresponds to a geodesic from 1 to w in $\mathcal{C}_A(G)$. For $j = 1, \ldots, n$, the truncation $w_j = a_1 \cdots a_j$ corresponds to a point on the geodesic. As H is k-quasiconvex, there exists $v_j \in H$ such that $d_A(w_j, v_j) \leq k$, so $u_j = w_j^{-1} v_j$ is a word of length at most k. In particular, take $u_n = 1$ and $v_n = w_n = w$. Also, let $u_0 = 1$.) Let $h_j = v_{j-1}^{-1} v_j$, so that $v_j = h_1 \cdots h_j$. In particular, $w = h_1 \cdots h_n$. (See Figure 12.)

The word $u_{j-1} a_j u_j^{-1}$ corresponds to a path joining v_{j-1} and v_j, so the relation $h_j = u_{j-1} a_j u_j^{-1}$ holds in G. Note that $\ell_A(h_j) \leq 2k+1$. In other words, any element of H can be written as a product of elements of H of length at most $2k+1$. Therefore, H is generated by the set

$$B = \{h \in H \mid \ell_A(h) \leq 2k+1\},$$

which is finite since A is finite.

Moreover, as $w = h_1 \cdots h_n$, we have $\ell_B(w) \leq n = \ell_A(w)$. At the same time, we have $\ell_A(w) \leq (2k+1)\ell_B(w)$. It follows that the inclusion $H \hookrightarrow G$ is a $(2k+1, 0)$-quasi-isometric embedding. $\square$

A word-hyperbolic group is called *locally quasiconvex* if every finitely generated subgroup is quasiconvex. Local quasiconvexity is an extremely strong condition. For example, applying Theorems 1.11 and 1.18, we see that any quasiconvex subgroup of a hyperbolic group is finitely presentable. Therefore, any locally quasiconvex group is *coherent*, meaning that every finitely generated subgroup is finitely presentable. The property of coherence has long been of interest to group theorists. Additionally, the intersection of two quasiconvex subgroups is again quasiconvex. If G is locally quasiconvex, it therefore satisfies satisfies *Howson's property*: the intersection of any two finitely generated subgroups is finitely generated. For more information on Howson's property, see Kapovich [12].

Rips [17] gives an example of a finitely generated subgroup of a hyperbolic group that is not finitely presented, hence not hyperbolic. Therefore, not every subgroup of a hyperbolic group is quasiconvex. Non-quasiconvex subgroups of hyperbolic groups are quite difficult to find, though; not many examples are known. At the same time, proving that a group is locally quasiconvex is also quite difficult. Thus, the subject of local quasiconvexity is a source of many open questions.

1.5 Spaces of Non-Positive Curvature

In this section, we will discuss an extremely important concept in non-Euclidean geometry that is closely related to δ-hyperbolic spaces: spaces of non-positive curvature.

The three best-known types of planar geometry are Euclidean, hyperbolic, and spherical. Accordingly, for $\kappa \in \{-1, 0, 1\}$, define the *model spaces* M_κ as follows: M_0 is the Euclidean plane $\mathbb{E}^2$, M_{-1} is the hyperbolic plane $\mathbb{H}^2$, and M_1 is the unit 2-sphere S^2. More generally, for other values $\kappa \in \mathbb{R}$, we may define M_κ by scaling the metric on $M_{\mathrm{sign}(\kappa)}$ by a factor of $1/\sqrt{|\kappa|}$. (For our purposes, though, we will only need to consider $\kappa \in \{-1, 0, 1\}$.) Denote the metric on M_κ by d_κ. In Riemannian geometry, each model space M_κ is known as the 2-manifold of constant sectional curvature κ. We wish to extend this notion of curvature to more general geodesic spaces.

Let (X, d_X) be a geodesic metric space, and let $\Delta = \Delta \subset X$ be a geodesic triangle with vertices x_1, x_2, x_3. In each model space M_κ, there exists a *comparison triangle* $\overline{\Delta}_\kappa = \Delta(\overline{x}_1, \overline{x}_2, \overline{x}_3) \subset M_\kappa$, unique up to congruence, such that $d_X(x_i, x_j) = d_\kappa(\overline{x}_i, \overline{x}_j)$ for each pair $i \neq j \in \{1, 2, 3\}$. (For $\kappa > 0$, we must assume that the distances are sufficiently small, since M_κ has finite diameter.) Each point $p \in \Delta$ corresponds to a unique point $\overline{p} \in \overline{\Delta}_\kappa$; if p lies in the edge $[x_i, x_j]$, then $\overline{p}$ lies in $[\overline{x}_i, \overline{x}_j]$, and we have $d_X(p, x_i) = d_\kappa(\overline{p}, \overline{x}_i)$ and $d_X(p, x_j) = d_\kappa(\overline{p}, \overline{x}_j)$.

Definition 1.19. A geodesic triangle $\Delta \subset X$ satisfies the CAT(κ) *inequality* if for any $p, q \in \Delta$ and any comparison points $\overline{p}, \overline{q} \in \overline{\Delta}_\kappa$, we have $d_X(p, q) \leq d_\kappa(\overline{p}, \overline{q})$. If every geodesic triangle in X satisfies the CAT(κ) inequality, we say that X is (globally) CAT(κ). If every point in X has a neighborhood that is CAT(κ), we say that X has *curvature* $\leq \kappa$.

In some sense, a space that is globally CAT(κ) is "more hyperbolic" or "more negatively curved" than the model space M_κ. Note that if $\kappa' > \kappa$, then CAT(κ) implies CAT(κ'), since M_κ is itself CAT(κ').

The abbreviation CAT, coined by Gromov [9, p. 106], stands for the Comparison[3] of A.D. Aleksandrov and A. Toponogov, who formulated the above definition of having curvature $\leq \kappa$ and other equivalent conditions. The idea of CAT(κ) spaces, particularly when $\kappa \leq 0$, is essential for the modern field of hyperbolic geometry, which in turn is connected closely to topology and differential geometry.

For our purposes, the importance of CAT(κ) spaces is the following lemma:

Proposition 1.20. *If a geodesic space X is* CAT(κ) *for some $\kappa < 0$, then X is δ-hyperbolic, where $\delta = \sqrt{-2/\kappa}$.*

Proof. Let $\Delta = \Delta(x, y, z)$ be a geodesic triangle in X, and let $\overline{\Delta}_\kappa = \Delta(\overline{x}, \overline{y}, \overline{z})$ be a comparison triangle for Δ in M_κ. Let $p \in \Delta$, suppose that $p \in [x, y]$, and let $\overline{p} \in [\overline{x}, \overline{y}]$ be a comparison point for p. Since the hyperbolic plane is $\sqrt{2}$-hyperbolic, the triangle Δ' is δ-slim, where $\delta = \sqrt{-2/\kappa}$. Then there exists a point $\overline{q} \in [\overline{x}, \overline{z}] \cup [\overline{y}, \overline{z}]$ such that $d_\kappa(\overline{p}, \overline{q}) \leq \delta$. The point $\overline{q}$ corresponds to some point $q \in [x, z] \cup [y, z]$, so by the CAT(κ) inequality, we obtain $d_X(p, q) \leq \delta$. Thus, X is δ-hyperbolic. □

Applying Corollary 1.9, we obtain:

Corollary 1.21. *If a group G acts properly, cocompactly, and by isometries on a* CAT(κ) *geodesic space X, where $\kappa < 0$, then G is hyperbolic.* □

Thus, one of the best ways to show that a group is hyperbolic is to show that it acts properly and cocompactly by isometries on a CAT(-1) space. We shall use this technique in Section 2.6. However, unlike in the case of Corollary 1.9, it is not known whether the converse to Corollary 1.21 is true for $\kappa = 0$, let alone $\kappa = -1$.

As in topology, it is easiest to work with spaces that are constructed out of smaller pieces — namely, cell complexes. Recall that a *convex polyhedron* in S^n, $\mathbb{E}^n$, or $\mathbb{H}^n$ is a bounded region B that is the intersection of a finite number of closed half-spaces. A *face* of B is any intersection of B and some of the bounding hyperplanes; each face is itself a convex polyhedron of dimension less than that of B. The *codimension* of a face is the dimension of B minus the dimension of the face.

A CW complex X is called *piecewise spherical* (PS), *piecewise Euclidean* (PE), or *piecewise hyperbolic* (PH) if the following conditions hold:

- For every closed cell B of X, there is a homeomorphism f_B of B onto a convex polyhedron in S^n, $\mathbb{E}^n$, or $\mathbb{H}^n$, respectively. The inverse images of the faces of this polyhedron are called faces of B.

- If B and C are closed cells, then $B \cap C$ is a face of both B and C, and the restriction of $f_C f_B^{-1}$ to $f_B(B \cap C)$ is an isometry between $f_B(B \cap C)$ and $f_C(B \cap C)$.

[3]Bridson and Haefliger [3, p. VII] claim that the C stands for E. Cartan, another pioneer of this area.

Each cell B of X can then be turned into a geodesic metric space by declaring f_B to be an isometry. We may then identify B with the image of f_B. In other words, the complex X is formed by gluing together polyhedra along identical faces using gluing maps that are isometries. To metrize the whole complex, define the distance d_X between two points to be the infimum of the lengths of the shortest path joining them (or infinity if there is no such path). Assuming X is path-connected, this infimum is always realized by a path that is the union of geodesic segments in the cells of X. Such a path is then a geodesic segment in X, making X into a geodesic space. (See Moussong [16] for more details.)

Any PS, PE, or PH complex has curvature $\leq \kappa$, where $\kappa = 1$, 0, or -1 respectively, since each point has a neighborhood (namely, the interior of the highest-dimensional cell containing that point) that satisfies CAT(κ). To determine whether the complex satisfies CAT(κ) globally, we need to consider the way the cells come together. For any cell C of X, we define a PS complex called the *link of C in X*, $\mathrm{lk}(C, X)$, as follows. Let k be the dimension of C, and let x be a point in the interior of C. In each closed, n-dimensional cell B containing C, let $\mathrm{lk}_x(C, B)$ be the intersection of B with a small $(n-k-1)$-sphere that lies in the $(n-k)$-plane orthogonal to C. We scale $\mathrm{lk}_x(C, B)$ up to be a convex polyhedron in an $(n-k-1)$-sphere of unit radius. It is easy to see that $\mathrm{lk}_x(C, B)$ does not depend on x, so we may define $\mathrm{lk}(C, B)$ accordingly. The link $\mathrm{lk}(C, X)$ is then defined by gluing together the cells $\mathrm{lk}(C, B)$ in the obvious manner to form a PS complex.

The link gives a measure of how much solid angle comes together at each cell. For instance, if C is a codimension-2 face of B (such as an edge of a 3-dimensional polyhedron, then $\mathrm{lk}(C, B)$ is an arc of angular length equal to the angle between the two codimension-1 faces that meet at B. If X consists of several such cells are glued together cyclically around C, then $\mathrm{lk}(X, C)$ is a PS closed curve whose length equals the sum of the dihedral angles.

A *closed geodesic* in a space Y is an isometric embedding $c : S_\ell \to Y$ of a circle of circumference ℓ into Y. The *girth* of Y is defined as the infimum of the lengths of closed geodesics in Y and denoted $\mathrm{g}(Y)$. If Y admits no closed geodesics, we set $\mathrm{g}(Y) = \infty$.

The criterion for X to be globally CAT(κ) is basically that when n-cells of X completely encircle a k-cell C, at least as much $(n-k-1)$-dimensional solid angle comes together at C as in an ordinary $(n-k-1)$-sphere. More formally, Gromov [9, p. 120] gives the following lemma:

Lemma 1.22. *Let X be a PS, PE, or PH complex, and let $\kappa = -1$, 0, or -1 respectively. The following are equivalent:*

1. *The complex X is globally* CAT(κ).
2. *For every cell C of X, the link* $\mathrm{lk}(C, X)$ *is* CAT(1).
3. *For every cell C of X, the girth of* $\mathrm{lk}(C, X)$ *is at least* 2π.

A complex satisfying (3) is sometimes said to satisfy the *link axiom*. We will make use of this property in Section 2.6.

Chapter 2

Coxeter Groups

In the Introduction, we gave a very colloquial definition of Coxeter groups that was meant to be comprehensible to someone not familiar with group theory. Here is a more formal definition:

Definition 2.1. Let W be a group with finite generating set S. A function $m : S \times S \to \{1, 2, \ldots, \infty\}$ is called a *Coxeter matrix* if $m(s, s) = 1$ and $m(s, t) = m(t, s) \neq 1$ when $s \neq t$. (We write $m_{s,t} = m(s, t)$.) The pair (W, S) is called a *Coxeter system* if W has a presentation

$$W = \langle S \mid (st)^{m_{s,t}} \ (s, t \in S,\ m_{s,t} \text{ finite})\rangle$$

for some Coxeter matrix m. The group W is then called a *Coxeter group*.

In other words, the group is defined by the relations $s^2 = 1$ and $(st)^{m_{s,t}} = 1$; the latter can also be written as

$$\underbrace{sts\cdots}_{m_{s,t}\text{ letters}} = \underbrace{tst\cdots}_{m_{s,t}\text{ letters}},$$

as in the colloquial description.

A Coxeter system is easily represented using a *Coxeter graph*, a labelled graph with vertices corresponding to the elements of S and with an edge labelled $m_{s,t}$ joining the vertices s and st if $m_{s,t} \geq 3$. When $m_{s,t} = 3$, we typically omit the label for convenience.

If S can be partitioned into disjoint, nonempty subsets $S_1, \cdots, S_k$ such that $m_{s,t} = 2$ whenever s and t are in different pieces of the partition, then S is isomorphic to the direct product $W_1 \times \cdots \times W_k$, where (W_i, S_i) is a Coxeter system generated by S_i with exponents inherited from (W, S). The connected components of the Coxeter graph for (W, S) are then the graphs for the systems (W_i, S_i). If W cannot be decomposed in this manner, it is called *irreducible*; in this case, the Coxeter graph for (W, S) is connected.

A Coxeter system is called *right-angled* if $m_{s,t}$ is either 2 or ∞ whenever $s \neq t$.

In his seminal paper [5], H.S.M. Coxeter showed that any *discrete reflection group* — that is, a group of isometries of $\mathbb{R}^n$ that is generated by reflections and under whose action the orbit of any point is discrete — can be presented in the form above, and he gave a complete classification of such groups. È.B. Vinberg [20] extended Coxeter's arguments to discrete reflection groups of hyperbolic space. The abstract definition is due to Nicolas Bourbaki, who proved many of the fundamental properties of Coxeter groups in his well-known work [2] on Lie groups and Lie algebras.

In keeping with the previous chapter, we will be especially interested in the geometry of Coxeter groups, as seen through the word metric and the Cayley graph. As described in the Introduction, one can intuitively visualize the Cayley graph by piecing together $2m_{s,t}$-gons whose sides are alternatively labelled s and t in an appropriate manner. (When $m_{s,t} = \infty$, we have infinite paths whose sides are labelled in this manner.) However, formally justifying this intuitive picture requires a fair amount of work, and that is the goal of the next few sections.

2.1 The Geometric Representation

Recall from Section 1.3 the definition of a group presentation. A group G can be given a presentation $G = \langle A \mid R \rangle$ if G is isomorphic to the quotient of the free group on A by the normal closure of R: $G \approx F(A)/N(R)$. Although it is a useful language for describing the elements of a group, a presentation actually carries extremely little information about the algebraic structure of the group. One problem that can often occur is that a group is "smaller" than its presentation makes it seem. For instance, one of the elements of the generating set A could in fact represent the identity in G, or two elements of A could represent the same element of G. Additionally, one of the relations in the group could say that $u^m = 1$ for some word u, while the order of the element u is actually a proper divisor of m. In this section, we will prove that these problems do not occur with Coxeter groups. Let (W, S) be a Coxeter system with coefficients $m_{s,t}$.

For clarity, we shall adopt a helpful notational convention used (in part) in Bourbaki [2] and Tits [19]. Abstract words will be written in boldface type. In particular, we write the generators themselves as $\mathbf{S} = \{\mathbf{s}, \mathbf{t}, \cdots\}$. Formally, words are the elements of the free monoid on $\mathbf{S}$, denoted $L(\mathbf{S})$. Elements of the free group $F(\mathbf{S})$ — reduced words — will also be written with boldface. The set of defining relations (for any group) is a subset of $F(\mathbf{S})$, so it is written as $\mathbf{R}$.

On the other hand, italic symbols represent elements of the group $W = F(\mathbf{S})/N(\mathbf{R})$. That is, if $\phi : F(\mathbf{S}) \to F(\mathbf{S})/N(\mathbf{R})$ is the quotient map, let $s = \phi(\mathbf{s})$ for each $\mathbf{s} \in \mathbf{S}$, and let $S = \phi(\mathbf{S})$. (We do not yet know, however, that the elements of S are all distinct, i.e., that $\phi|_{\mathbf{S}}$ is injective.) Also, define $\psi : L(\mathbf{S}) \to W$ to be the composition of ϕ with the canonical map $L(\mathbf{S}) \to F(\mathbf{S})$; this map sends every formal word to the element of W that it represents. The map ψ is surjective since each element of S is equal to its inverse in W.

First, we will check that none of the elements of S is the identity, i.e., that $\phi(\mathbf{s}) \neq \phi(\mathbf{1})$. Let $\epsilon : F(\mathbf{S}) \to \mathbb{Z}_2$ be the homomorphism sending each $\mathbf{s}_i$ to the nontrivial element of $\mathbb{Z}_2$. Each element of $\mathbf{R}$ is a word of even length and hence is sent to the identity in $\mathbb{Z}_2$. Therefore, the map ϵ induces a map $\overline{\epsilon} : W \to \mathbb{Z}_2$ sending each s_i to the nontrivial element, so $s_i \neq 1$. This resolves the first of the problems discussed above.

To resolve the other ambiguities, we will now see how an arbitrary Coxeter group can be said to act "by reflections" on a vector space. In the Introduction, we saw how the Coxeter group $\langle a, b \mid a^2, b^2, (ab)^m \rangle$ can be seen as the group of symmetries of $\mathbb{R}^2$ generated by reflections through two lines that meet at an angle of π/m. The angle between two unit vectors orthogonal to these lines is $\pi - \pi/m$, so the dot product of those vectors is $\cos(\pi - \pi/m) = -\cos(\pi/m)$. The vectors obviously form a basis for $\mathbb{R}^2$. We now try to generalize this notion for arbitrary Coxeter groups.

Let V be an n-dimensional real vector space with basis $(e_s)_{s\in S}$. Define a symmetric bilinear form B on V by

$$B(e_s, e_t) = \begin{cases} -\cos\frac{\pi}{m_{s,t}} & m_{s,t} < \infty \\ -1 & m_{s,t} = \infty. \end{cases}$$

Note that the e_s are unit vectors with respect to this form. The matrix of the form B with respect to the basis (e_s) is called the *cosine matrix* of W.

For each $s \in S$, define a linear map $\sigma_s : V \to V$ by

$$\sigma_s(v) = v - 2B(v, e_s)e_s;$$

these maps are called *reflections.* Let H_s be the orthogonal complement of e_s with respect to the form B: $H_s = e_s^\perp = \{v \in V \mid B(v, e_s) = 0\}$.

We may check that each reflection preserves the form B:

$$\begin{aligned} B(\sigma_s(v), \sigma_s(w)) &= B(v - 2B(v, e_s)e_s, w - 2B(w, e_s)e_s) \\ &= B(v, w) - 4B(w, e_s)B(v, e_s) + 4B(v, e_s)B(w, e_s)B(e_s, e_s) \\ &= B(v, w) \end{aligned}$$

Thus, $\sigma(s)$ fixes the subspace H_s pointwise and negates e_s.

Lemma 2.2. *The homomorphism $F(\mathbf{S}) \to GL(V)$ defined by $\mathbf{s} \mapsto \sigma_s$ induces a homomorphism $\sigma : W \to GL(V)$, known as the* geometric *or* canonical representation.

Proof. We must show that each of the relations in W holds in $GL(V)$, i.e., that $\sigma_s^2 = (\sigma_s\sigma_t)^{m_{s,t}} = \mathbf{1}_V$, the identity map in V. In fact, we will show more: the order of the product $\sigma_s\sigma_t$ is exactly $m_{s,t}$.

It is clear from the definitions that each reflection σ_s has order exactly 2. Next, choose $s \neq t$, and consider the subspace $V_{s,t}$ spanned by e_s and e_t. The reflections σ_s and σ_t stabilize $V_{s,t}$, and they fix the orthogonal complement $V_{s,t}^\perp$ pointwise.

Let $T : V_{s,t} \to V_{s,t}$ be the restriction of $\sigma_s \circ \sigma_t$ to $V_{s,t}$. By evaluating $T(e_s)$ and $T(e_t)$, we can easily compute that the matrix of T with respect to the basis (e_s, e_t) is

$$T = \begin{bmatrix} 4\cos^2\theta - 1 & -2\cos\theta \\ 2\cos\theta & -1 \end{bmatrix}$$

where $\theta = \pi/m_{s,t}$. We consider two cases:

- If $m_{s,t}$ is infinite, then this matrix becomes

$$T = \begin{bmatrix} 3 & -2 \\ 2 & -1 \end{bmatrix}.$$

  One can show by induction that

$$T^n = \begin{bmatrix} 2n+1 & -2n \\ 2n & -2n+1 \end{bmatrix}$$

  so T has infinite order, as does $\sigma_i \sigma_j$.

- If $m_{s,t}$ is finite, note that the characteristic polynomial of the matrix for T is $p_T(x) = x^2 - 2\cos(2\theta)x + 1$, and the eigenvalues are

$$x = \cos 2\theta \pm i \sin 2\theta = e^{\pm 2i\theta},$$

  which are $m_{s,t}^{\text{th}}$ roots of unity. Therefore, the order of the transformation T is $m_{s,t}$.

  Let $v = ae_s + be_t$ be a nonzero vector in $V_{s,t}$. We then have:

$$\begin{aligned} B(v,v) &= a^2 B(e_s,e_s) + 2abB(e_s,e_t) + b^2 B(e_t,e_t) \\ &= a^2 + b^2 - 2ab\cos\theta \\ &= (a - b\cos\theta)^2 + b^2 \sin^2\theta \end{aligned}$$

  Since $\theta \neq 0$ and $(a,b) \neq (0,0)$, at least one of the squared quantities is nonzero, so the sum is positive. Thus, the restriction of B to $V_{s,t}$ is positive definite, hence nondegenerate, so V is isomorphic to the direct sum $V_{s,t} \oplus V_{s,t}^{\perp}$. It follows that $\sigma_s \sigma_t = T \oplus \mathbf{1}_{V_{s,t}^{\perp}}$, so $\sigma_s \sigma_t$ has the same order as T, namely $m_{s,t}$.

□

Thus, we have an action of the group W on the space V, defined by $w.v = \sigma(w)(v)$ for $w \in W, v \in V$. (Remember that $(w_1 w_2).v = \sigma(w_1)(\sigma(w_2)(v))$; that is, we apply w_2 first, then w_1.)

Corollary 2.3. *The order of st in W is exactly $m_{s,t}$.*

Proof. Since $(st)^{m_{s,t}} = 1$ in W, the order of st is at most $m_{s,t}$. On the other hand, we have $\sigma(st) = \sigma_s \sigma_t$, which has order $m_{s,t}$, so so the order of st is at least $m_{s,t}$. □

One geometric consequence of Corollary 2.3 is that in the Cayley graph $\mathcal{C}_S(W)$, each path that is labelled $(st)^{m_{s,t}}$ consists of exactly $2m_{s,t}$ distinct edges. Moreover, the subgroup of W generated by s and t is isomorphic to the dihedral group $D_{m_{s,t}}$. Since the order of st is at least 2 for $s \neq t$, we see that $st \neq 1$, so $s \neq t$. We may therefore define the *rank* of the Coxeter system (W,S) as the cardinality of S. We have thus resolved all the difficulties posed at the beginning of this section.

2.2 The Exchange and Deletion Conditions

In this section, we will discuss the connection between the length function ℓ_S and the geometric representation, and we will use this to prove one of the most important properties of Coxeter groups, the so-called Exchange Condition. Our treatment largely follows Humphreys [11]; we will also briefly discuss another proof of the Exchange Condition given by Tits [19].

Let (W, S) be a Coxeter system. Recall that $\ell(w) = \ell_S(w)$ is defined as length of the shortest word in $L(\mathbf{S})$ that represents w: $\ell(w) = \min\{\ell(\mathbf{w}) \mid \mathbf{w} \in \psi^{-1}(w)\}$, where $\psi : L(\mathbf{S}) \to W$ is the canonical map and $\ell(\mathbf{w})$ is the number of letters in $\mathbf{w}$. Some general, easy-to-prove facts about the length function (true for all groups) are that $\ell(w) = \ell(w^{-1})$ and that $\ell(w) = 0$ if and only if w is the identity in W. Also, there is a triangle inequality:

$$\ell(w) - \ell(w') \leq \ell(ww') \leq \ell(w) + \ell(w').$$

Consider the geometric representation of W, defined in the previous section. We shall analyze the behavior of a certain set of unit vectors in V under the action of W. Let Φ be the set of vectors of the form $w.e_s$ for some $w \in W$, $s \in S$; Φ is called the *root system*, and the elements of Φ are called *roots*. The roots are all unit vectors with respect to the form B, since $B(w.e_s, w.e_s) = B(e_s, e_s) = 1$. Also, the root system is closed under negation, since $s.e_s = -e_s$.

Each pair of opposite roots $\pm w.e_s$ corresponds to a reflection in W that interchanges them, namely wsw^{-1}. To see this, note that for any $v \in V$, we have

$$\begin{aligned} wsw^{-1}.v &= w.(w^{-1}.v - 2B(w^{-1}.v, e_s)e_s) \\ &= v - 2B(w^{-1}.v, e_s)w.e_s \\ &= v - 2B(v, w.e_s)w.e_s \end{aligned}$$

which is the form for a reflection that negates the root $w.e_s$ and fixes its orthogonal complement pointwise. Thus, each root $\alpha = w.e_s \in \Phi$ determines a reflection $s_\alpha = wsw^{-1}$. Obviously, s_α and $s_{-\alpha}$ are identical. Conversely, if two roots α and β determine the same reflection, then $s_\alpha(\beta) = \beta - 2B(\beta, \alpha)\alpha = -\beta$, so $\beta = B(\beta, \alpha)\alpha$, and then $\beta = \pm\alpha$ since both are unit vectors. Thus, there is a natural bijection between the pairs of opposite roots and the reflections in W.[1] This correspondence has the following property:

Lemma 2.4. *If $\alpha, \beta \in \Phi$ and $\beta = w.\alpha$, then $ws_\alpha w^{-1} = s_\beta$.*

Proof. We have $ws_\alpha w^{-1}.\beta = ws_\alpha.\alpha = w. - \alpha = -\beta$, meaning that $ws_\alpha w^{-1}$ is the reflection that negates β, namely s_β. □

[1]Bourbaki [2] proves the results of this section by having W act on an abstractly defined set that is identical to Φ: the product set $\{\pm 1\} \times T$, where $T \subset W$ is the set of conjugates of the elements of S. He defines the action of W in a similarly fashion. While Bourbaki's proof is somewhat shorter than the one here, the advantage of our approach is that it allows for more geometric intuition.

Every root $\alpha \in \Phi$ can be uniquely as a linear combination of the e_i: $\alpha = \sum_{s \in S} c_s e_s$. If all the c_s are nonnegative, we say that α is *positive* and write $\alpha > 0$; if all the c_s are nonpositive, we say α is negative and write $\alpha < 0$. Let Π denote the set of positive roots; the set of negative roots is $-\Pi$.

The homomorphism $\bar{\epsilon} : W \to \mathbb{Z}_2$ shows that any two expressions for a particular $w \in W$ have the same length modulo 2. Taking $w' = s$ in the triangle inequality, we see that $\ell(ws)$ equals either $\ell(w) + 1$ or $\ell(w) - 1$. (Another formulation of this distinction is that $\ell(ws) = \ell(w) - 1$ if and only if w has a reduced expression whose last letter is $\mathbf{s}$.) The following lemma gives a geometric criterion in terms of the root system:

Lemma 2.5. *Let $w \in W$, $s \in S$. If $\ell(ws) = \ell(w) + 1$, then $w.e_s > 0$. If $\ell(ws) = \ell(w) - 1$, then $w.e_s < 0$.*

Proof. First, note that the second statement follows from the first. If $\ell(ws) = \ell(w) - 1$, then $\ell(wss) = \ell(ws) + 1$, so $ws.e_s > 0$, and therefore $w.e_s < 0$.

For the first statement, induct on $\ell(w)$. If $\ell(w) = 0$, then w is the identity, so obviously $w.e_s > 0$. If $\ell(w) = k > 0$, then there exists a reduced expression for w whose last letter is is $\mathbf{t} \neq \mathbf{s}$. That is, the set of reduced words

$$\mathbf{A} = \{\mathbf{w} \in \psi^{-1}(w) \mid \mathbf{w} = \mathbf{s}_1 \cdots \mathbf{s}_k,\ \mathbf{s}_k = \mathbf{s},\ \text{and } \mathbf{s}_{h+1}, \cdots, \mathbf{s}_k \in \{\mathbf{s}, \mathbf{t}\} \text{ for some } h < k\}$$

is nonempty. Let $\mathbf{w} = \mathbf{s}_1 \cdots \mathbf{s}_k$ be an element of $\mathbf{A}$ for which h is as small as possible. Write $\mathbf{w} = \mathbf{uv}$, where $\mathbf{u} = \mathbf{s}_1 \cdots \mathbf{s}_h$ and $\mathbf{v} = \mathbf{s}_{h+1} \cdots \mathbf{s}_k$, and set $u = \psi(\mathbf{u})$ and $v = \psi(\mathbf{v})$.

By the minimality of h, the word u does not have a representation ending in $\mathbf{s}$ or $\mathbf{t}$. Therefore, by the induction hypothesis, $u.e_s > 0$ and $u.e_t > 0$. It thus suffices to show that $v.e_s$ is a nonnegative linear combination of e_s and e_t, for $w.e_s$ is then a nonnegative linear combination of positive roots, hence itself a positive root.

The word $\mathbf{v}$ is a $(k-h)$-factor alternating product of $\mathbf{s}$ and $\mathbf{t}$ whose last letter is $\mathbf{t}$: either $\mathbf{v} = (\mathbf{st})^p$ (where $p \geq 1$) or $\mathbf{v} = \mathbf{t}(\mathbf{st})^p$ (where $p \geq 0$). We consider two cases, depending on $m_{s,t}$:

- If $m_{s,t} = \infty$, then since $B(e_s, e_t) = -1$,, we have $s.e_t = e_t + 2e_s$ and $t.e_s = e_s + 2e_t$. By induction, we see that $(st)^p.e_s = (2p+1)e_s + (2p)e_t$ and $t(st)^p.e_s = (2p+1)e_s + (2p+2)e_t$, so $v.e_s$ is a positive linear combination of e_s and e_t.

- If $m = m_{s,t}$ is finite, then we must have $\ell(\mathbf{v}) < m$, or else we could represent v with either a shorter word or a word ending in $\mathbf{s}$. Thus, $p < m/2$. As seen in the previous section, the restriction of B to $V_{s,t}$ is positive-definite, so we may identify $V_{s,t}$ with Euclidean space and the transformation st with a rotation through an angle of $2\pi/m$ in the direction from e_s to e_t. Therefore, $(st)^p$ rotates e_s through an angle of at most $\pi - 2\pi/m$. The angle between e_s and e_t is $\pi - \pi/m$, so the vector $(st)^p.e_s$ still lies within the positive cone spanned by e_s and e_t. Also, the angle between e_s and the line of reflection for t, L_t, is $\pi/2 - \pi/m$, so the angle between $(st)^p.e_s$ and L_t is at most $\pi/2$, and therefore $t(st)^p.e_s$ is also contained within the same cone. (See Figure 13.)

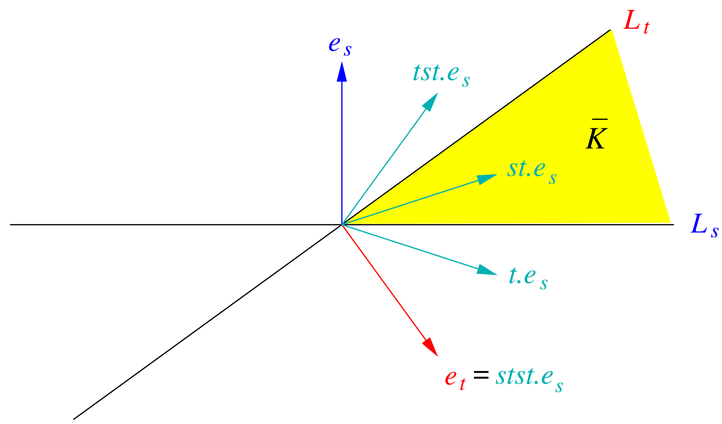

Figure 13: The positions of the roots $v.e_s$ in the proof of Lemma 2.5 when $m_{s,t} = 5$, along with the fundamental domain $\overline{K}$ defined at the end of the section.

In either case, we see that $v.e_s$ is equal to a nonnegative linear combination of e_s and e_t, and therefore $w.e_s$ is a positive root. $\square$

In particular, it follows that every root is either positive or negative. That is, $\Phi = \Pi \sqcup -\Pi$. Moreover, the lemma implies a crucial fact about the geometric representation:

Corollary 2.6. *The action of W on V is faithful. That is, the map $\sigma : W \to GL(V)$ is injective.*

Proof. If w is a nontrivial element of W, then let **s** be the last letter of a reduced expression for w. Then $\ell(ws) < \ell(w)$, so $w.e_s < 0$; in particular, $w.e_s \neq e_s$. Thus, $\sigma(w)$ is not the identity map on V. $\square$

The next lemma further clarifies the action of the generating reflections on the root system.

Lemma 2.7. *Each $s \in S$ sends e_s to its negative and permutes the other positive roots.*

Proof. By definition, $s.e_s = -e_s$. Let α be a positive root other than e_s; then

$$\alpha = \sum_{s \in S} c_s e_s,$$

where each $c_s \geq 0$. The root α cannot be a multiple of e_s, since all the roots are unit vectors, so $c_t > 0$ for some $t \neq s$. By definition, $s.\alpha = \alpha - B(\alpha, e_s)e_s$, so the coefficient of e_j in $s.\alpha$ is still strictly positive. Thus, $w.\alpha$ cannot be negative, so it is positive, and it is clearly distinct from e_i. Therefore, $s.(\Pi \setminus \{e_i\}) \subset \Pi \setminus \{e_i\}$. Applying s to both sides, we obtain the reverse inclusion, which shows that s permutes the elements of $\Pi \setminus \{e_i\}$. $\square$

In other words, each element of W of length 1 sends exactly one positive root to a negative root (and vice versa). By induction, one can prove a more general statement: the length of any element of W is equal to the number of positive roots that it sends to negative roots. Thus, the word-metric geometry on a Coxeter group is extremely closely connected to the intrinsic properties of the group, a statement that is not true for most groups given by presentations.

The preceding lemmas imply one of the most important properties of Coxeter groups:

Theorem 2.8 (Exchange Condition). *Let $w = s_1 \cdots s_r$ and suppose that for some $s \in S$, we have $\ell(ws) < \ell(w)$. Then there exists an index $q \in \{1, \ldots, r\}$ such that $ws = s_1 \cdots \widehat{s_q} \cdots s_r$, where the hat denotes omission.*

Proof. Since $e_s > 0$ and $w.e_s < 0$ by Lemma 2.5, there exists an index $q \leq r$ such that $(s_{q+1} \cdots s_r).e_j > 0$ but $(s_q \cdots s_r).e_s < 0$. That is, s_q sends the positive root $(s_{q+1} \cdots s_r).e_s$ to a negative one. By Lemma 2.7, the only such root is e_{s_q}, so $(s_{q+1} \cdots s_r).e_s = e_{s_q}$. Therefore, by Lemma 2.4, we have:

$$s_{q+1} \cdots s_r s s_r \cdots s_{q+1} = s_q.$$

The desired result follows immediately. $\square$

By taking inverses, we easily obtain a version of the Exchange Condition in which ws is replaced by sw, and so on.

An alternate form of the Exchange Condition that is sometimes more useful is the Deletion Condition:

Corollary 2.9 (Deletion Condition). *Let $w = s_1 \cdots s_r$, and suppose that $\ell(w) < r$. Then there exist indices $1 \leq p < q \leq r$ such that $w = s_1 \cdots \widehat{s_p} \cdots \widehat{s_q} \cdots s_r$.*

Proof. For each index j, let $w_q = s_1 \cdots s_q$. If for all q, we have $\ell(w_q) \geq \ell(w_{q-1})$, then by induction we must have $\ell(w_q) = q$ for each q, and in particular $\ell(w) = r$, which by assumption is not the case. Hence, for some q, we must have $\ell(w_q) < \ell(w_{q-1})$. By the Exchange Condition, there exists an index $p \in \{1, \ldots, q-1\}$ such that $w_q = s_1 \cdots \widehat{s_p} \cdots s_{q-1}$, and therefore $w = s_1 \cdots \widehat{s_p} \cdots \widehat{s_q} \cdots s_r$. $\square$

In other words, given a non-reduced word, a reduced word representing the same element of W can be obtained by deleting pairs of letters. It turns out that this property truly characterizes Coxeter groups, in view of the following theorem [11, pp. 16–18]:

Theorem 2.10. *Let W be a group with finite generating set S. Suppose that every element of S has order 2 and that the pair (W, S) satisfies the Deletion Condition. For $s, t \in S$, let $m_{s,t}$ be the order of st in W. Then (W, S) is a Coxeter system with exponents $m_{s,t}$.* $\square$

Tits [19] provides another way to think about the proof of the Exchange Condition. Let V^* be the dual vector space to V, with basis (e_s^*). In this space, the angle between two basis vectors e_s^*, e_t^* is $\pi/m_{s,t}$ rather than $\pi - \pi/m_{s,t}$. Define:

$$\begin{aligned} K^* &= \{\sum_{s\in S} c_s e_s^* \in V^* \mid c_s^* > 0\ \forall s \in S\} \\ \overline{K}^* &= \{\sum_{s\in S} c_s e_s^* \in V^* \mid c_s^* \geq 0\ \forall s \in S\}. \end{aligned}$$

Topologically, these can be seen as the interior and closure, respectively, of the same set. One can show that $\overline{K}^*$ is a fundamental domain for the induced action of W on V^* (the so-called *contragredient representation* of W) in the sense that if $w \in W$ fixes some $v^* \in \overline{K}^*$, then v^* lies on one of the hyperplanes that bound K^*. That is, $v^* \in \overline{K}^* \setminus K^*$. (See Bourbaki [2] for a proof.) Analagous to Lemma 2.5, Tits shows that if $\ell(sw) < \ell(w)$, then K^* and $w.K^*$ lie on opposite sides of the hyperplane $(c_s = 0)$ and uses this fact to prove the Exchange Condition.

Now return to the original vector space V. Each of the reflecting hyperplanes H_s separates V into two disjoint, open half-spaces. Let K_s be the one that contains the basis vector e_s, and let $\overline{K}_s$ be the closed half-space $K_s \cup H_s$. The dual to K^* (resp. $\overline{K}^*$) is the intersection of these open (resp. closed) half-spaces: $K = \bigcap_{s\in S} K_s$, $\overline{K} = \bigcap_{s\in S} \overline{K}_s$. (These are represented by the shaded area in Figure 13.) The closed cone $\overline{K}$ is then a fundamental domain for the action of W on V. We shall use this cone in Sections 2.3 and 2.4.

2.3 Discrete Reflection Groups

The geometric representation of Coxeter groups is easiest to visualize when the form B is positive-definite. In that case, the vector space V can be identified with $\mathbb{R}^n$ with the standard dot product (where $n = \mathrm{Card}(S)$). That is, we have a basis of vectors $e_s \in \mathbb{R}^n$ for which each product $e_s \cdot e_t$ is equal to $-\cos(\pi/m_{s,t})$. Therefore, the angle between e_s and e_t is $\pi - \pi/m_{s,t}$, and the hyperplanes $e_s^\perp$ and $e_t^\perp$ meet at a dihedral angle of $\pi/m_{s,t}$. (Note that all of the $m_{s,t}$ are finite, for otherwise the cosine matrix would have the matrix $\left[\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right]$, which is not positive definite, as a principal submatrix.)

Any reflection stabilizes the unit sphere $S^{n-1} \subset \mathbb{R}^n$, so the geometric representation induces an action of W on S^{n-1}. The fundamental domain C for this action is the intersection $S^{n-1} \cap \overline{K}$, an $(n-1)$-simplex in S^{n-1}. The generators of W are then reflections of S^{n-1} through the codimension-1 faces of C. The translates of C all have the same area since the reflections are isometries; these simplices form a tessellation of the sphere. As S^{n-1} is compact, it has finite volume, so the group W must be finite.

Coxeter [6] proved that the converse is also true: If W is a finite Coxeter group, then the form B is positive-definite. One way to prove this fact is to start with any positive-definite form β on V, for instance the standard dot product, and define a W-invariant, positive-definite form

$$\beta'(v, v') = \sum_{w\in W} \beta(w.v, w.v').$$

Using representation theory, one can then show that this form β' is a positive scalar multiple of B; therefore B is also positive-definite (For a full proof, see Humphreys [11, ch. 6].) It follows that every finite Coxeter group is naturally isomorphic to a *finite reflection group*, i.e., a finite subgroup of $\mathrm{GL}_n(\mathbb{R})$ that is generated by reflections.

Conversely, every finite reflection group is a Coxeter group. One way to prove this is to use a root system, as above, to show that the group satisfies the Deletion Condition; the result then follows from Theorem 2.10. We shall take a different, more general approach that is somewhat similar to that used by Coxeter in his original work [5], although in very different language.

Let $n \geq 2$, and let X be either the n-sphere S^n, Euclidean n-space $\mathbb{E}^n$, or hyperbolic n-space $\mathbb{H}^n$. In each of these spaces, any $(n-1)$-dimensional hyperplane P defines a unique reflection $\sigma_P : X \to X$, an isometry of X that fixes P pointwise and interchanges the two components of $X \setminus P$. Let $\mathrm{Isom}(X)$ be the group of isometries of X.

Definition 2.11. A group $W \subset \mathrm{Isom}(X)$ is called a *discrete reflection group* if it is generated by a finite set of reflections and the orbit of any point in X under the action of G is topologically discrete.

The fundamental domain C for the action of W on X is a convex region bounded by hyperplanes P_s (indexed by a finite set S) such that the reflections σ_{P_s} generate W. We may identify σ_{P_s} with s. If two planes P_s and P_t intersect, their intersection is spanned by a codimension-2 face of C. The composition st rotates C through an angle of 2θ around $P_s \cap P_t$, where θ is the dihedral angle between P_s and P_t. As C is a fundamental domain, we must then have $\theta = \pi/m_{s,t}$ for some integer $m_{s,t} \geq 2$. If P_s and P_t do not intersect, set $m_{s,t} = \infty$. Also, set $m_{s,s} = 1$ for each $s \in S$.

Theorem 2.12. *Let W be a discrete reflection group on X, as above. Then W is isomorphic to the Coxeter group $\langle S \mid R \rangle$, where $R = \{(st)^{m_{s,t}} \mid s, t \in S,\ m_{s,t} \text{ finite}\}$.*

Proof. Obviously, each reflection has order 2. The product st is a rotation of X through an angle of $2\pi/m_{s,t}$ around the $(n-2)$-plane containing $P_s \cap P_t$; this transformation has order $m_{s,t}$. Thus, the relations of a Coxeter group hold in W. We must now show that they are defining relations for the group.

The orbit of C gives a cell decomposition $\mathcal{T}$ of X. Let C' be an n-cell of $\mathcal{T}$. As C is a fundamental domain, there is a unique element $w \in W$ such that $w.C = C'$. For each $s \in S$, let $P'_s = w.P_s$. The reflection through P'_s is then equal to wsw^{-1}. Note that $wsw^{-1}.C' = ws.C$.

Let $\mathcal{T}^*$ be the dual cell structure to $\mathcal{T}$. The vertices of $\mathcal{T}^*$ correspond to the n-simplices of $\mathcal{T}$ and therefore to the elements of W. The 1-cells of $\mathcal{T}^*$ correspond to the $(n-1)$-cells of $\mathcal{T}$. In particular, for each $w \in W$ and $s \in S$, the $(n-1)$-cell $w.C \cap w.P_s$ corresponds to an edge joining the vertices w and ws. Thus, the 1-skeleton of $\mathcal{T}^*$ is isomorphic to the Cayley graph $\mathcal{C}_S(W)$.

By construction, the number of n-cells of $\mathcal{T}$ that come together at any codimension-2 face E of $\mathcal{T}$ is $2m_{s,t}$, where P_s and P_t are the hyperplanes that meet at the corresponding codimension-2 face of C. The 2-cell E^* in $\mathcal{T}^*$ that corresponds to E is therefore a $(2m_{s,t})$-gon whose sides are an edge-loop labelled $(st)^{m_{s,t}}$. Moreover, every such path in the 1-skeleton of $\mathcal{T}^*$ is filled in with such a 2-cell.

The following useful theorem, proven in Bridson and Haefliger [3, p. 135], gives a geometric interpretation for the notion of a group presentation:

Lemma 2.13. *Let G be a group with generating set A, and let R be a subset of the kernel of the canonical map $\phi : F(A) \twoheadrightarrow G$. Let $\mathcal{C}_{A,R}(G)$ be the 2-complex formed from the Cayley graph $\mathcal{C}_A(G)$ by attaching a single 2-cell to each edge-loop labelled by a cyclic permutation of a reduced word $r \in R$. Then the complex $\mathcal{C}_{A,R}(G)$ is simply connected if and only if $\ker \phi$ is equal to the normal closure of R in $F(A)$, i.e., if $\langle A \mid R \rangle$ is a presentation for G.* □

In particular, in the complex $\mathcal{C}_{S,R}(W)$, each disk labelled s^2 has its two edges glued together to form a 2-sphere, which is simply connected. The $(st)^{m_{s,t}}$-labelled disks are then precisely the 2-cells of $\mathcal{T}^*$. Thus, $\mathcal{C}_{S,R}(W)$ is isomorphic the 2-skeleton of $\mathcal{T}^*$ with a 2-sphere adjoined to each edge. Recall that the fundamental group of a CW complex is determined by the complex's 2-skeleton. In particular, as $\mathcal{T}^*$ is a cell decomposition for X, we see that $\mathcal{C}_{S,R}(W)$ is simply connected. Therefore, by Lemma 2.13, we have $W \approx \langle S \mid R \rangle$. □

We may further obtain further results about discrete reflection groups depending on whether $X = S^n$, $\mathbb{E}^n$, or $\mathbb{H}^n$.

- When $X = S^n$, the group W is finite. Coxeter [5] showed that every convex polyhedron in S^n whose angles are all at most $\pi/2$ is a spherical n-simplex, so C has exactly $n+1$ faces. If we embed S^n as the unit sphere in $\mathbb{R}^{n+1}$, the $(n-1)$-dimensional hyperplanes in S^n correspond bijectively with the n-dimensional hyperplanes through the origin in $\mathbb{R}^{n+1}$, and reflections of S^n and of $\mathbb{R}^{n+1}$ correspond naturally. Thus, any finite reflection group (a finite subgroup of $GL_{n+1}(V)$ generated by reflections) can be realized as a discrete reflection group on S^n and therefore is a Coxeter group. Combining this result with the results above, we obtain the following:

  Proposition 2.14. *Let W be a group with finite generating set S of cardinality n. The following are equivalent:*

  1. *The pair (W, S) is a finite Coxeter system.*
  2. *The pair (W, S) is a Coxeter system whose cosine matrix is positive-definite.*
  3. *The group W embeds in $GL_{n-1}(\mathbb{R})$ as a finite reflection group.* □

- For $X = \mathbb{E}^n$, the group is called an *affine reflection group.* If C is unbounded, it is the product of $\mathbb{R}^k$ with a bounded, $(n-k)$-dimensional polyhedron whose dihedral angles are the same as those between the hyperplanes P_s, so by projecting orthogonally we may assume that C is bounded. Coxeter showed that every convex, Euclidean polyhedron whose angles are all at most $\pi/2$ is equal to a product of simplices and infinite lines in perpendicular dimensions. Any affine reflection group then splits accordingly into a direct product of affine reflection groups, since $m_{s,t} = 2$ for perpendicular hyperplanes.

The fundamental domain for an irreducible, affine reflection group is a simplex. In this case, any proper subset of the bounding hyperplanes meet at a single point and therefore generate a finite reflection group. In other words, every principal submatrix of the cosine matrix B of an irreducible affine reflection group is positive definite. It follows that B itself is positive-semidefinite but not positive-definite. (That is, for all $v \in V$, we have $B(v, v) \geq 0$, but there exists a nonzero $v \in V$ for which $B(v, v) = 0$.)

Accordingly, we define a Coxeter group to be *affine* if its cosine matrix is positive-semidefinite but not positive-definite. One can show that every irreducible, affine Coxeter group is an affine reflection group by considering a certain subspace of the vector space V on which the group acts in its geometric representation; see Humphreys [11, pp. 133–134] for details. In general, the only difference between the affine Coxeter groups and the affine reflection groups is that the product of an affine Coxeter group and a finite Coxeter group is affine because its cosine matrix is positive-semidefinite. On the other hand, any affine reflection group factors as the product of strictly affine groups.

It is quite easy to classify all the irreducible finite and affine Coxeter groups. One can check that all the groups listed in Figure 14 are either affine or finite, as the case may be, by looking at their cosine matrices. By process of elimination, one can then prove that this list is complete. The original result is due to Coxeter [5]; Humphreys [11, ch. 2] gives a very elegant proof.

- For $X = \mathbb{H}^n$, it is very difficult to obtain general results when the fundamental region is not a simplex, since the classification of hyperbolic polyhedra is extremely complicated. The hyperbolic reflection groups whose fundamental domains are simplices (allowing for noncompact simplices of finite volume with vertices at infinity) turn out to correspond to the irreducible Coxeter groups for which B has signature $(n-1, 1)$ and $B(v, v) < 0$ for any element $v \in K$, where K is the open cone of Section 2.2 [11, pp. 138–141]. Such groups are called (irreducible) *hyperbolic Coxeter groups*. (Do not confuse these with word-hyperbolic groups, discussed in Chapter 1; the latter will be discussed in Section 2.6.) Vinberg [20] gives more information on general hyperbolic Coxeter groups.

Example 2.15. Consider the *triangle groups* discussed in the Introduction: namely, the Coxeter groups of the form

$$W = \left\langle a, b, c \mid a^2, b^2, c^2, (ab)^k, (bc)^l, (ac)^m \right\rangle.$$

The fundamental domain for this group is a 2-simplex whose angles are π/k, π/l, and π/m. Whether this simplex is spherical, Euclidean, or hyperbolic depends on whether the sum

$$\frac{\pi}{k} + \frac{\pi}{l} + \frac{\pi}{m}$$

is greater than, equal to, or less than π, respectively.

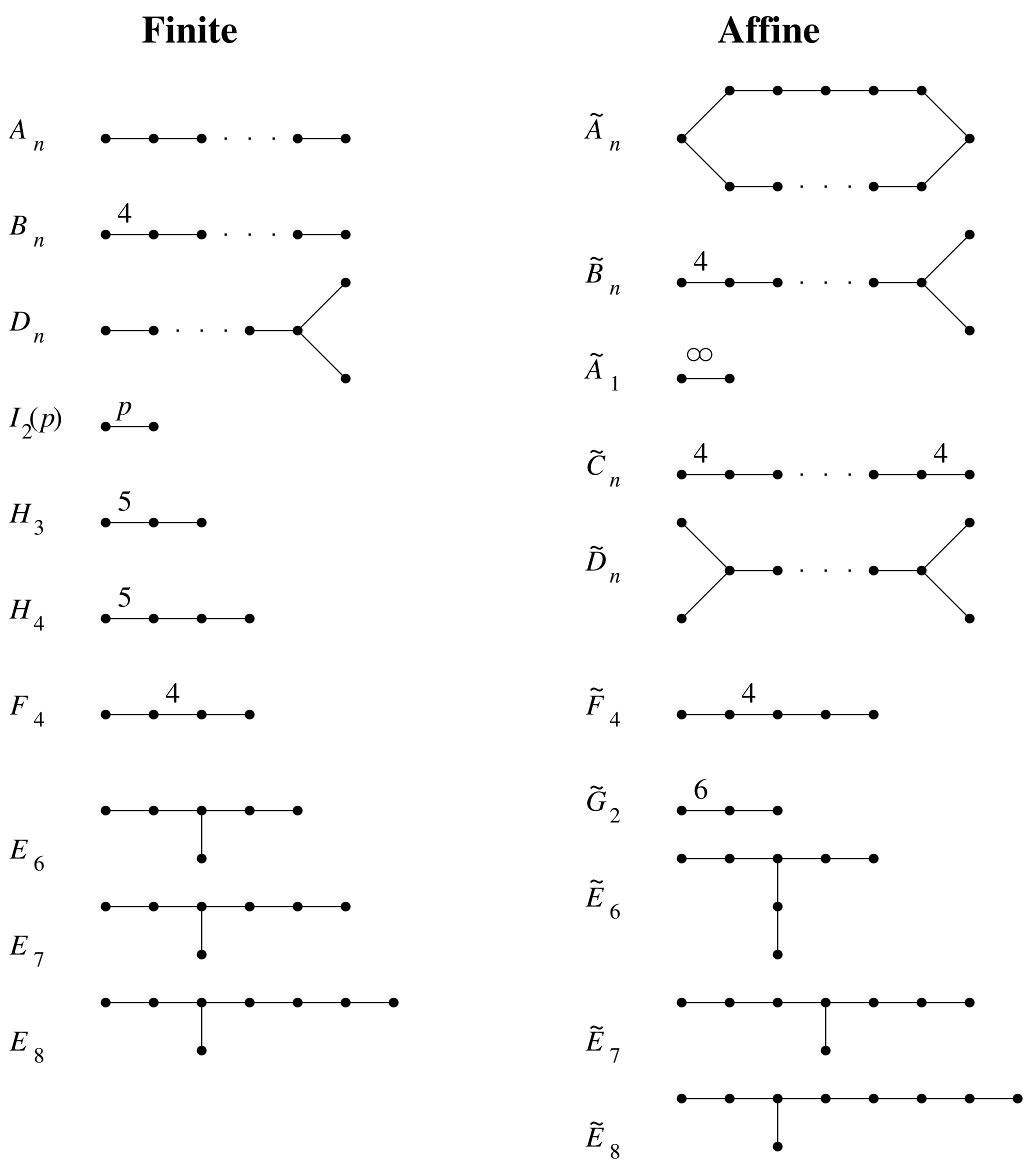

Figure 14: Coxeter diagrams for the irreducible finite and affine Coxeter groups.

We can now answer the question posed in the Introduction. Escher's work (Figure 6) is based on a tessellation of the hyperbolic plane by 30°-45°-90° triangles. As we have just seen, such a triangle is the fundamental domain for the $(2, 4, 6)$ triangle group. Each fish in the picture is composed of two triangles, with the fish's spine lying opposite the 90° angle of each. (In order to see the symmetries more easily, it is helpful to ignore the fish.) By taking the dual of such a tessellation, we obtain the Cayley graph shown in Figure 5.

2.4 The Word Problem in Coxeter Groups

As mentioned before, one of the serious challenges presented by group presentations is the word problem. Given a group presentation, the problem is to find an algorithm that determines in a predictable number of steps whether or not two given words in the generators represent the same element of the group. Novikov showed in 1955 that in general the word problem for groups is unsolvable. (See, for example, Lyndon and Schupp [14] or Stillwell [18] more about the solvability of word problems.)

Of course, in many instances the word problem can be solved. In general, this entails finding a canonical form in which every element of the group can be expressed uniquely and a procedure for obtaining this form from an arbitrary word in the generators of the group. For instance, in a free group, any word can be simplified to a unique reduced word in which there no consecutive pairs of inverse letters. In the free abelian group $\langle a, b \mid aba^{-1}b^{-1} \rangle$, each element can be expressed uniquely as $a^j b^k$ by transposing letters as needed. Note that it is sufficient to determine whether or not a given word represents the identity; given two arbitrary words w and w', we can consider the product $w^{-1}w'$.

Jacques Tits [19] solved the word problem in Coxeter groups in 1968. His algorithm for determining whether or not a word represents the identity is effective, in the sense that it always yields a definite result in a finite number of steps, although it is not very efficient in terms of computation time. In this section, we shall prove this theorem and discuss its consequences, closely following Tits's proof.

Let (W, S) be a Coxeter system. For each pair of distinct generators $s, t \in S$ for which $m_{s,t}$ is finite, let $\mathbf{w}_{s,t}$ denote the product of $m_{s,t}$ alternating factors of $\mathbf{s}$ and $\mathbf{t}$ whose last factor is $\mathbf{s}$. Clearly, the words $\mathbf{w}_{s,t}$ and $\mathbf{w}_{t,s}$ represent the same element $w_{s,t} \in W$.

Define an *elementary simplification* to be either of the following moves:

- Replacing a word of the form $\mathbf{xssy}$ with $\mathbf{xy}$.
- Replacing a word of the form $\mathbf{x}\mathbf{w}_{s,t}\mathbf{y}$ with $\mathbf{x}\mathbf{w}_{t,s}\mathbf{y}$.

A *simplification* is a finite sequence of elementary simplifications. For any word $\mathbf{x}$ in $L = L(\mathbf{S})$, define the *simplification set* $S(\mathbf{x}) \subset L$ as the set, obviously finite, of words that can be obtained from $\mathbf{x}$ by simplifications. Clearly, if $\mathbf{v} \in S(\mathbf{u})$, then $S(\mathbf{v}) \subset S(\mathbf{u})$. Also, if $\mathbf{v} \in S(\mathbf{u})$ and $\ell(u) = \ell(v)$, then all of the elementary simplifications that lead from $\mathbf{u}$ to $\mathbf{v}$ are of the second type. As these are reversible, we must have $\mathbf{u} \in S(\mathbf{v})$, and therefore $S(\mathbf{u}) = S(\mathbf{v})$.

Tits's key result says that one can always obtain a reduced word in W using these two operations. Specifically:

Theorem 2.16. *Let* $\mathbf{u}, \mathbf{v} \in L(\mathbf{S})$. *The words* $\mathbf{u}$ *and* $\mathbf{v}$ *represent the same element of* W *if and only if the intersection* $S(\mathbf{u}) \cap S(\mathbf{v})$ *is nonempty. In particular, if* $\mathbf{v}$ *is a reduced word, then this condition is equivalent to having* $\mathbf{v} \in S(\mathbf{u})$.

This theorem provides an algorithm, albeit extremely inefficient, for obtaining a reduced word from a given word. First, cancel any pairs of consecutive, identical letters. If there are none, then consider all of the finitely many words that can be obtained with substitutions of the form $sts\cdots = tst\cdots$. If any of these words contains a pair of consecutive, identical letters, cancel them and repeat; if not, stop. No matter the choice of the order of cancellations and substitutions, Theorem 2.16 implies that the end result is always a reduced word. The fact that order doesn't matter will be useful in Chapter 3.

Tits's algorithm is much more efficient in a right-angled Coxeter group. Here, the second type of elementary simplification consists of interchanging adjacent letters that commute. In reducing a word, it then suffices to look for pairs of identical letters that commute with all of the letters in between. This simplification of the search process provides a polynomial-time solution to the word problem in right-angled Coxeter groups. In particular, let $\mathbf{w}$ be a word of length r. The time to determine whether or not the i^{th} letter of $\mathbf{w}$ can be cancelled with the next occurrence of the same letter is a linear function of $r - i$, so the time to see if $\mathbf{w}$ contains a cancelling pair is a quadratic function of r, and the total time to reduce $\mathbf{w}$ completely is a cubic. We will explore this algorithm in more depth in Chapter 3.

The following lemma is in some sense a generalization of the Exchange Condition:

Lemma 2.17. *Let* $w \in W$, $s, t \in S$. *Suppose that* $\ell_S(ws) < \ell_S(w)$ *and* $\ell_S(wt) < \ell_S(w)$. *Then* $\ell_S(ww_{s,t}) = \ell_S(w) - m_{s,t}$.

Proof. Let $r = \ell_S(w)$. As in the proof of Lemma 2.5, we may find a reduced expression $\mathbf{w} = \mathbf{uv}$, where $u = \psi(\mathbf{u})$ satisfies $u.e_s > 0$ and $u.e_t > 0$, and $v = \psi(\mathbf{v})$ is an alternating product of s and t satisfying $v.e_s < 0$ and $v.e_t < 0$. The only element of the dihedral group generated by s and t that has reduced expressions ending in both s and t is the word of length $m_{s,t}$, namely $w_{s,t}$. Therefore, we have $ww_{s,t} = u$, which has length $\ell_S(w) - m_{s,t}$. □

Proof of Theorem 2.16. The elementary simplifications correspond to the relations in W, so any simplification of a word in L represents the same element of W as the original word. This proves the "if" statement. For the final statement, if $\mathbf{v}$ is reduced, then every element of $S(\mathbf{v})$ is obtained by a sequence of elementary simplifications of the second type, so $S(\mathbf{u}) \cap S(\mathbf{v})$ being nonempty implies that $\mathbf{v} \in S(\mathbf{u})$.

For the "only if" statement, give the set $L \times L$ a lexicographic order: $(\mathbf{x}, \mathbf{y}) < (\mathbf{x}', \mathbf{y}')$ if $\ell(\mathbf{x}) < \ell(\mathbf{x}')$ or if $\ell(\mathbf{x}) = \ell(\mathbf{x}')$ and $\ell(\mathbf{y}) < \ell(\mathbf{y}')$. This ordering is obviously a well-ordering. Let $\Sigma \subset L \times L$ be the set of all pairs of words $(\mathbf{x}, \mathbf{y})$ such that $\ell(\mathbf{x}) \geq \ell(\mathbf{y})$.

Suppose that the theorem is false, i.e., that the set

$$\Omega = \{(\mathbf{x}, \mathbf{y}) \in L \times L \mid S(\mathbf{x}) \cap S(\mathbf{y}) = \varnothing \text{ and } \psi(\mathbf{x}) = \psi(\mathbf{y})\}$$

is nonempty. Let $(\mathbf{u}, \mathbf{v})$ be a minimal element of $\Sigma \cap \Omega$ with respect to the lexicographic ordering. Let $u = \psi(\mathbf{u}) = \psi(\mathbf{v})$.

We claim that every element of $S(\mathbf{u})$ has the same length. For suppose $\mathbf{x}$ is an element of $S(\mathbf{u})$ with $\ell(\mathbf{x}) < \ell(\mathbf{u})$. Clearly $S(\mathbf{x}) \cap S(\mathbf{v}) = \varnothing$ and $\psi(\mathbf{x}) = \psi(\mathbf{u}) = \psi(\mathbf{v})$. Also, either $(\mathbf{x}, \mathbf{v})$ or $(\mathbf{v}, \mathbf{x})$ is in Σ, depending on whether or not $\ell(\mathbf{x}) \geq \ell(\mathbf{v})$. But this contradicts the minimality of $(\mathbf{u}, \mathbf{v})$ in $\Sigma \cap \Omega$.

Let $\mathbf{s}$ and $\mathbf{t}$ be the last letters of $\mathbf{u}$ and $\mathbf{v}$ respectively, and write $\mathbf{u} = \mathbf{u}'\mathbf{s}$ and $\mathbf{v} = \mathbf{v}'\mathbf{t}$. We claim that $\ell(\mathbf{u}) = \ell(\mathbf{v})$. For suppose that $\ell(\mathbf{u}) > \ell(\mathbf{v})$. We have seen that all the words in L that represent a given element of W have the same length modulo 2. In particular $\ell_S(u) \geq \ell_S(v) + 2$, so $\ell(\mathbf{vs}) \leq \ell_S(u) - 1 = \ell(\mathbf{u}')$. Note that $\psi(\mathbf{u}') = \psi(\mathbf{us}) = \psi(\mathbf{vs})$. By the minimality of $(\mathbf{u}, \mathbf{v})$, the intersection $S(\mathbf{u}') \cap S(\mathbf{vs})$ cannot be empty. Every element of $S(\mathbf{u}')$ has the same length; this fact follows from the corresponding result for $S(\mathbf{u})$. Hence, there exists $\mathbf{y} \in S(\mathbf{u}') \cap S(\mathbf{vs})$ such that $\ell(\mathbf{u}') = \ell(\mathbf{y}) \leq \ell(\mathbf{vs})$, which means that $\ell(\mathbf{vs}) = \ell(\mathbf{u}')$. Then $S(\mathbf{vs}) = S(\mathbf{u}')$, so $\mathbf{vs} \in S(\mathbf{u}')$, so $\mathbf{vss} \in S(\mathbf{u})$, so $\mathbf{v} \in S(\mathbf{u})$, which contradicts our initial assumption that $S(\mathbf{u}) \cap S(\mathbf{v}) = \varnothing$. Thus, $\ell(\mathbf{u}) = \ell(\mathbf{v})$.

By the minimality of $(\mathbf{u}, \mathbf{v})$, the words $\mathbf{u}$ and $\mathbf{v}$ must be reduced. Since $\psi(\mathbf{us}) = \psi(\mathbf{u}')$ and $\psi(\mathbf{vt}) = \psi(\mathbf{v}')$, we have $\ell_S(us) < \ell_S(u)$ and $\ell_S(ut) < \ell_S(u)$. Hence, by Lemma 2.17, there exists a word $\mathbf{x} \in L$ of length $\ell(\mathbf{u}) - m_{s,t}$ such that

$$\psi(\mathbf{x}\mathbf{w}_{s,t}) = \psi(\mathbf{x}\mathbf{w}_{t,s}) = \psi(\mathbf{u}) = \psi(\mathbf{v}).$$

The words $\mathbf{u}'$ and $\mathbf{x}\mathbf{w}_{s,t}\mathbf{s}^{-1}$ (i.e., $\mathbf{x}\mathbf{w}_{s,t}$ with its last letter deleted) have the same length, specifically $\ell(\mathbf{u}) - 1$, and both represent us, so by minimality the sets $S(\mathbf{u}')$ and $S(\mathbf{x}\mathbf{w}_{s,t}\mathbf{s}^{-1})$ must intersect and therefore be equal. Similarly, we have $S(\mathbf{v}') = S(\mathbf{x}\mathbf{w}_{t,s}\mathbf{t}^{-1})$. It then follows that

$$S(\mathbf{u}) = S(\mathbf{x}\mathbf{w}_{s,t}) = S(\mathbf{x}\mathbf{w}_{t,s}) = S(\mathbf{v}),$$

contradicting our initial assumption. $\square$

2.5 Parabolic Subgroups

Let (W, S) be a Coxeter system. In Section 2.1, we saw that the subgroup generated by two distinct elements $s, t \in S$ is the dihedral group $D_{m_{s,t}}$. We will now extend this notion to the subgroup of W generated by any subset of S.

For any word $\mathbf{w} \in L(\mathbf{S})$, define the *composition set* $T(\mathbf{w}) \subset \mathbf{S}$ as the set of letters that occur in $\mathbf{w}$. The following lemma is an immediate consequence of Tits's solution to the word problem:

Lemma 2.18. *Let $w \in W$, and let $\mathbf{w}_1, \mathbf{w}_2$ be two reduced expressions for w. Then $T(\mathbf{w}_1) = T(\mathbf{w}_2)$.*

Proof. By the second part of Theorem 2.16, we have $\mathbf{w}_2 \in S(\mathbf{w}_1)$, so there exists a sequence of elementary simplifications transforming $\mathbf{w}_1$ into $\mathbf{w}_2$. Since both words have minimal length in $\psi^{-1}(w)$, all these simplifications are of the second type. By construction, such moves preserve composition sets, so $T(\mathbf{w}_1) = T(\mathbf{w}_2)$. $\square$

Thus, the composition set $T(w)$ of an element of $w \in W$ is well-defined as the composition set of any reduced expression for w. Note that this is not the case for arbitrary groups. For instance, in the group $\langle a, b, c, d \mid abc^{-1}d^{-1}\rangle$, the reduced words ab and dc both represent the same element of the group.

For any subset $T \subset S$, define the *parabolic subgroup* $W_T \subset W$ to be the subgroup generated by the elements of T. Clearly, the subgroup W_T consists of those elements of w such that $T(w) \subset T$. One of the most important facts about Coxeter groups is that parabolic subgroups are extremely well-behaved.

Theorem 2.19. *Let (W, S) be a Coxeter system, and let $T \subset S$.*

1. *The pair (W_T, T) is a Coxeter system with exponents inherited from (W, S).*
2. *The inclusion $(W_T, d_T) \hookrightarrow (W, d_S)$ is an isometric embedding, where d_T and d_S are the word metrics on W_T and W respectively.*
3. *The subgroup W_T is convex in W with respect to the generating set S.*

Proof. For (1), let T' be a set with the same cardinality as T, and choose a bijection $\phi : T' \to T$. Let (W', T') be the Coxeter system generated abstractly by T', with relations determined by the corresponding relations in (W, S). The map $\phi : T' \to T$ then induces a well-defined homomorphism $\overline{\phi} : W' \to W$ whose image is W_T. We must now check that this map is injective. Let $w \in W'$, and suppose that $\phi(w) = 1$. If $\mathbf{t}'_1 \cdots \mathbf{t}'_r$ is an expression for w, where $\mathbf{t}'_i \in \mathbf{T}'$, then $\phi(w)$ has an expression $\mathbf{t}_1 \cdots \mathbf{t}_r$, where $\mathbf{t}_i = \phi(\mathbf{t}'_i)$. By Theorem 2.16, the word $\mathbf{t}_1 \cdots \mathbf{t}_r$ can be reduced to the empty word using elementary simplifications involving only the letters in $T(\mathbf{t}_1 \cdots \mathbf{t}_r) \subset T$. We may then use the same simplifications, which are also valid in W', to reduce $\mathbf{t}'_1 \cdots \mathbf{t}'_r$ to the empty word, showing that $w = 1$ in T. Thus, $\overline{\phi}$ is injective. Identifying W' with its image $W_T \subset W$, we see that (W_T, T) is a Coxeter system. Both (2) and (3) follow immediately from the fact that $T(w) \subset T$ for any $w \in W_T$. $\square$

Theorem 2.19 is a very important property of Coxeter groups. Part (1) says that we can find a presentation for a parabolic subgroup simply by reading off the relations that involve only the generators of that subgroup. That is by no means true for arbitrary groups. For instance, in the group $G = \langle a, b, c \mid ab^{-2}, cb^{-3}\rangle$, the generators a and c are distinct, and there is no relation involving only them, so we might guess that the "parabolic subgroup" they generate is free of rank 2. However, since $a = b^2$ and $c = b^3$, the subgroup is in fact infinite cyclic, generated by b. A well-defined notion of parabolic subgroups is thus special to Coxeter groups.

One can easily construct the Cayley graph of a Coxeter system (W, S) piece by piece using the Cayley graphs of the parabolic subgroups. In particular, for each subset $T \subset S$, the graph $\mathcal{C}_S(W)$ contains one copy of $\mathcal{C}_T(W_T)$ for each coset in the quotient W/W_T. Since $W_T \cap W_{T'} = W_{T\cap T'}$, we form $\mathcal{C}_S(W)$ by gluing the copy of $\mathcal{C}_T(W_T)$ corresponding to wW_T to the copy of $\mathcal{C}_{T'}(W_{T'})$ corresponding to $wW_{T'}$ along the copy of $\mathcal{C}_{T\cap T'}(W_{T\cap T'})$ corresponding to $wW_{T\cap T'}$. This makes formal the idea of "gluing together $(2m_{s,t})$-gons" mentioned in the Introduction.

2.6 Word-Hyperbolic Coxeter Groups

We are finally ready to bring together the results of Chapters 1 and 2 and ask: When is a Coxeter group word-hyperbolic? G. Moussong answered this question in his doctoral dissertation [16] in 1988. We will not give a full proof of his results here, but we give a partial explanation of some of them. Moussong and M.W. Davis (his dissertation advisor) give a more readable exposition in [7].

First, a brief note on terminology: Although in Chapter 1, we used the terms "word-hyperbolic" and "hyperbolic" interchangeably, we saw in Section 2.3 that the term "hyperbolic Coxeter group" has a specific meaning different from Gromov's definition of hyperbolicity. Indeed, the former meaning implies the latter, for if (W, S) is a hyperbolic Coxeter system, then the Cayley graph embeds quasi-isometrically into $\mathbb{H}^n$, so W is word-hyperbolic. On the other hand, any finite Coxeter group is word-hyperbolic but not hyperbolic. To eliminate any ambiguity, we will always use "word-hyperbolic" in this section.

Here is Moussong's main theorem:

Theorem 2.20. *Let (W, S) be a Coxeter system. The group W is word-hyperbolic if and only if the following both hold:*

1. *There is no subset $T \subset S$ such that (W_T, T) is an affine Coxeter system of rank at least 3.*

2. *There is no pair of disjoint subsets $T_1, T_2 \subset S$ such that W_{T_1} and W_{T_2} are both infinite and commute with each other.*

Using the list of finite and affine Coxeter groups given in Figure 14, it is very easy to identify which Coxeter groups are word-hyperbolic.

The "only if" direction of this theorem is easy to prove. The Cayley graph of an affine Coxeter group is the 1-skeleton of the dualization of a tessellation of Euclidean space. Hence, such a group is quasi-isometric to $\mathbb{E}^n$ and thus not hyperbolic. Since the embedding $W_T \hookrightarrow W$ is an isometry, every parabolic subgroup of a word-hyperbolic Coxeter group is itself hyperbolic. Therefore, a Coxeter system that does not satisfy (1) is not hyperbolic. Likewise, if (W, S) does not satisfy (2), then it contains two commuting elements of infinite order, which generate a subgroup isomorphic to $\mathbb{Z} \times \mathbb{Z}$. Recall that no hyperbolic group contains any subgroup isomorphic to $\mathbb{Z} \times \mathbb{Z}$. Thus, (2) must hold in a hyperbolic group W.

For the "if" direction, the idea is to construct a CAT(0), piecewise Euclidean complex Σ (the *Davis-Moussong complex*) on which the Coxeter group acts properly and cocompactly by isometries. For groups satisfying the conditions of Theorem 2.20, the complex Σ can be also given a piecewise hyperbolic structure that is CAT(-1), and therefore the group itself is hyperbolic by Lemma 1.21.

The basic building block for the complex Σ is the following construction. First, let (W, S) be a finite Coxeter group of rank n, and let $\epsilon > 0$. As seen in Section 2.3, the vector space V used in the geometric representation of W can be identified with $\mathbb{R}^n$ with the standard Euclidean metric and dot product. Choose a point x_ϵ that lies at distance ϵ from each of the reflecting hyperplanes. Define the *Coxeter cell* $C_\epsilon(W) = C_\epsilon(W, S)$ as the convex hull of the orbit $W.x_\epsilon$ of x_ϵ. This is an n-dimensional, convex, Euclidean polyhedron, and it can be given a natural PE cell structure. It is easy to see that the 1-skeleton of $C_\epsilon(W, S)$ is the Cayley graph $\mathcal{C}_S(W)$, scaled so that each edge has length 2ϵ. More generally, for $0 \leq k \leq n$, the k-cells of $C_\epsilon(W, S)$ correspond naturally to the cosets of the parabolic subgroups of rank k, and a face corresponding to wW_T is isometrically isomorphic to $C_\epsilon(W_T, T)$.

Now let (W, S) be an arbitrary group. The complex Σ_ϵ is a piecewise Euclidean complex defined as follows: For every subset $T \subset S$ for which W_T is finite, Σ includes one cell isometric to $\mathcal{C}_\epsilon$ for each coset in the quotient W/W_T. The cells are glued together in the obvious way: a nonempty intersection of two cosets of parabolic subgroups is itself a coset of a parabolic subgroup, so we may glue the two Coxeter cells together along the corresponding face, which is itself a Coxeter cell. The vertices of Σ correspond to the cosets of the trivial subgroup, i.e., the elements of W, and the 1-cells correspond to cosets of the rank-1 parabolic subgroups generated by individual elements of S. Therefore, the 1-skeleton Σ_ϵ^1 is the Cayley graph $\mathcal{C}_S(W)$, scaled so that every edge has length 2ϵ.

Next, we define a piecewise hyperbolic complex $\Sigma_\epsilon^{\mathbb{H}}$ that is isomorphic (as a cell complex) to Σ_ϵ.

First, consider the case where (W, S) is a finite Coxeter system. The intersection of $C_\epsilon(W) \subset \mathbb{R}^n$ with the convex cone $\overline{K}$ is called a *block* and denoted $B_\epsilon(W)$. Note that the barycenter of every face of in the star of x lies in one or more of the hyperplanes that bound K. Therefore, the block $B_\epsilon(W)$ is a subcomplex of the barycentric subdivision of $C_\epsilon(W)$ and is a combinatorial n-cube. The faces of $B_\epsilon(W)$ that are contained in faces of $C_\epsilon(W)$ are called *outer faces*; those that are contained in the hyperplanes that bound $\overline{K}$ are called *inner faces*. The outer faces clearly meet the inner faces perpendicularly. (See Figure 15.)

Let $B_\epsilon^{\mathbb{H}}(W) \subset \mathbb{H}^n$ (the *hyperbolic block* for W) be a hyperbolic polyhedron whose edge lengths are the same as those of $B_\epsilon(W)$ and whose angles between two inner faces and between an inner face and an outer face are all the same as those in $B_\epsilon(W)$. The angles between the outer faces, however, are necessarily different. By gluing together copies of $B_\epsilon^{\mathbb{H}}(W)$ in the obvious way, we obtain a *hyperbolic Coxeter cell* $C_\epsilon^{\mathbb{H}}(W)$. In this case, since $\Sigma_\epsilon = C_\epsilon(W)$, we set $\Sigma_\epsilon^{\mathbb{H}} = C_\epsilon^{\mathbb{H}}(W)$. If (W, S) is now an arbitrary Coxeter system, then we form $\Sigma_\epsilon^{\mathbb{H}}$ from Σ_ϵ by replacing each cell $C_\epsilon(W_T)$ by $C_\epsilon^{\mathbb{H}}(W_T)$.

Given two polyhedral complexes X and X', a λ-*map* is a homeomorphism $X \to X'$ that preserves cell structure and is a $(\lambda, 0)$-quasi-isometry (where $\lambda \geq 1$). If such a map exists, X' is called a λ-*change* of X.

For the affine Coxeter groups considered in Section 2.3, note that the complex Σ_ϵ is $\mathcal{T}^*$, the dual to the tiling of $\mathbb{R}^n$ by the fundamental domain of W. For the hyperbolic Coxeter groups, the same is true for $\Sigma_\epsilon^{\mathbb{H}}$. Thus, these complexes somewhat generalize the way in which finite, affine, and hyperbolic Coxeter groups act on S^n, $\mathbb{E}^n$, and $\mathbb{H}^n$, respectively.

Lemma 2.21. *Let $\lambda > 1$. For ϵ sufficiently small, the complex $\Sigma_\epsilon^{\mathbb{H}}$ is a λ-change of Σ_ϵ.*

Proof. When W is finite, define a homeomorphism $f : C_\epsilon(W) \to C_\epsilon^{\mathbb{H}}(W)$ by proportional radial scaling. That is, for each $x \in C_\epsilon(W)$, draw a line segment L_x from the center O

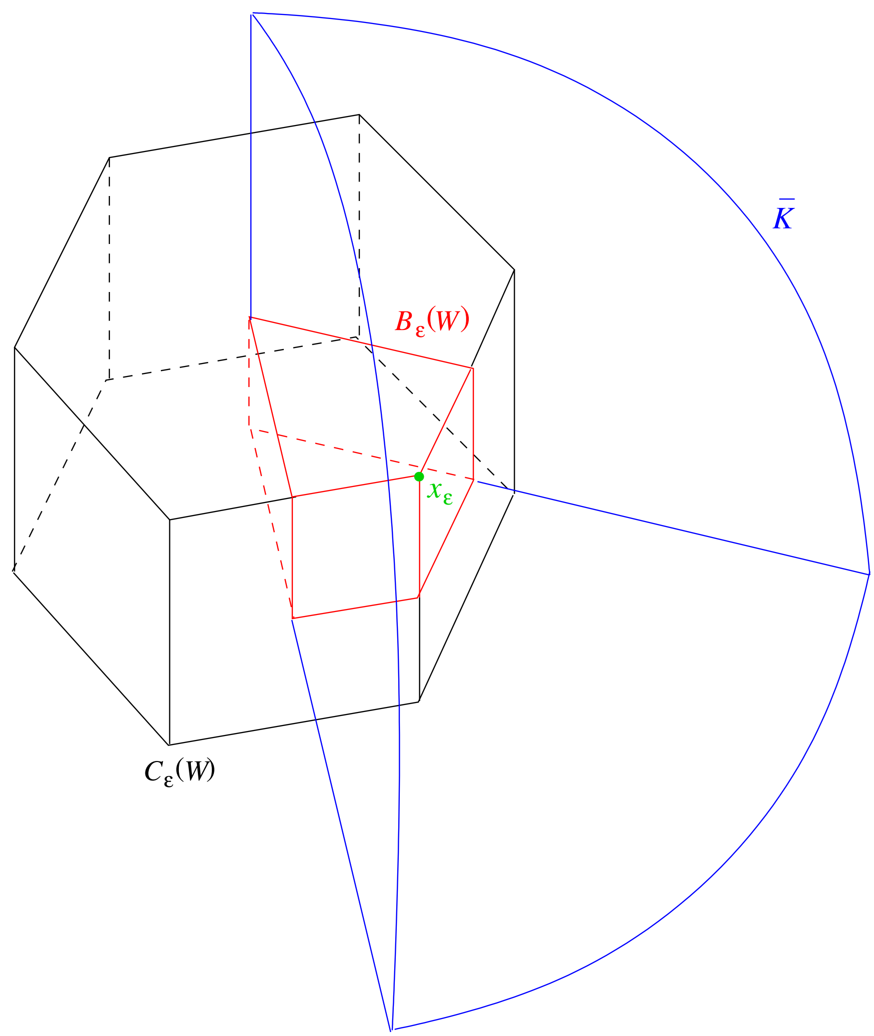

Figure 15: The Coxeter cell $C_\epsilon(W)$ and block $B_\epsilon(W)$ for the $(2, 2, 3)$ triangle group.

of $C_\epsilon(W)$ through x to the boundary of $C_\epsilon(W)$, and draw a corresponding segment L'_x in $C^{\mathbb{H}}_\epsilon(W)$. Define $f(x)$ to be the point on L'_x such that

$$\frac{d(O,x)}{|L_x|} = \frac{d(O',f(x))}{\ell\,|L'_x|}$$

where $\ell(L_x)$ denotes the length of L_x. Obviously, f preserves the cell structure of $C_\epsilon(X)$. Note that

$$\lim_{\epsilon\to 0}\inf_{x\in C_\epsilon(X)} \frac{\ell(L'_x)}{\ell(L_x)} = 1,$$

since small neighborhoods in hyperbolic space are approximately Euclidean. Therefore, for sufficiently small ϵ, the map f is a λ-map.

Now let (W,S) be an arbitrary Coxeter system. As S is finite, there are only finitely many isometric isomorphism classes of cells in Σ_ϵ, so we may choose ϵ sufficiently small that $C^{\mathbb{H}}_\epsilon(W_T)$ is a λ-change of $C_\epsilon(W_T)$ for every finite parabolic subgroup W_T. We then obtain a λ-map $\Sigma_\epsilon \to \Sigma^{\mathbb{H}}_\epsilon$ by gluing together the maps on the individual cells. □

Lemma 2.22. *The complexes Σ_ϵ and $\Sigma^{\mathbb{H}}_\epsilon$ are simply connected, and the group W acts properly and cocompactly by isometries on both.*

Proof. Let Σ be either Σ_ϵ or $\Sigma^{\mathbb{H}}_\epsilon$; the proof is identical for both.

To see that Σ is simply connected, note that the 2-skeleton of Σ is homeomorphic to the complex $\mathcal{C}_{S,R}(W)$ introduced in Lemma 2.13, where $R = \{(st)^{m_{s,t}} \mid s,t\in S,\ m_{st} \text{ finite}\}$. Since $\langle S \mid R\rangle$ is a presentation, $\mathcal{C}_{S,R}(W)$ is simply connected, as is Σ.

The action of W on itself by left multiplication extends naturally to an action of W on Σ by isometries. To see that this action is proper, note that any compact set $K \subset \Sigma$ is contained in a finite subcomplex Σ' of Σ. For any $w \in W$ for which $w.K \cap K \neq \varnothing$, we have $w.\Sigma' \cap \Sigma'$, so w permutes a subset of the vertices of Σ'. The set of such w is finite, as required. Finally, note that the star of any vertex is compact, and the orbit of the star covers all of Σ, so the action is cocompact. □

Finally, we must consider the global curvature of Σ_ϵ and $\Sigma^{\mathbb{H}}_\epsilon$.

Lemma 2.23. *Let (W,S) be a Coxeter system.*

1. *For any $\epsilon > 0$, the complex Σ_ϵ is globally* CAT(0)*.*
2. *If (W,S) satisfies the conditions of Theorem 2.20 and ϵ is sufficiently small, then the complex $\Sigma^{\mathbb{H}}_\epsilon$ is globally* CAT(−1)*.*

By Lemma 1.22, it suffices to show in each case that each complex satisfies the link axiom, i.e., that the link of any cell has girth at least 2π.

We consider the hyperbolic case first. Note the following lemma of Moussong [16, p. 17]:

Lemma 2.24. *Let X be a finite, piecewise spherical, simplicial complex. For any real number $\alpha < \mathrm{g}(X)$, there exists $\lambda > 1$ such that for any λ-change X' of X, we have $\mathrm{g}(X') \geq \alpha$.*

In particular, to prove (2), it suffices to show that the link of every cell in Σ_ϵ has girth strictly greater than 2π. If so, we can find $\alpha \geq 2\pi$ such that $\mathrm{g}(\mathrm{lk}(C, \Sigma_\epsilon)) > \alpha$ for every cell C of Σ_ϵ, since there are only finitely many isometric isomorphism classes of links. Let λ be the least of the values supplied for the different links by Lemma 2.24. By Lemma 2.21, there exists $\epsilon > 0$ such that $\Sigma_\epsilon^{\mathbb{H}}$ is a λ-change of Σ_ϵ. Each link $\mathrm{lk}(C_\epsilon^{\mathbb{H}}(W_T), \Sigma_\epsilon^H)$ is a λ-change of the corresponding link $\mathrm{lk}(C_\epsilon(W_T), \Sigma_\epsilon)$ and therefore has girth at least α, which proves that $\Sigma_\epsilon^{\mathbb{H}}$ is globally CAT(-1).

Thus, it suffices to consider the links of the cells in Σ_ϵ. Since the links depend only on angles, we may omit mention of ϵ. This proof is the main technical result of Moussong's thesis; it depends on and we shall not attempt to prove it here. However, we will prove an extremely simple case of the theorem.

The link of a vertex, $\mathrm{lk}(v, \Sigma)$, is generally easy to describe. It contains one vertex for each element of S. If $m_{s,t}$ is finite, then the vertices corresponding to s and t are connected by an edge of length $\pi - \pi/m_{s,t}$. For any subset $T \subset S$ of cardinality k, the vertices corresponding to T span a spherical $(k-1)$-simplex (whose shape is determined by its edge lengths) if and only if W_T is finite. Links of higher-dimensional cells are harder to describe, since they require knowing the various angles present in the Coxeter cells, but they can be determined algebraically using the cosine matrix for W.

Note that every edge in $\mathrm{lk}(v, \Sigma)$ (and, more generally, in $\mathrm{lk}(C, \Sigma)$ for any cell C) has length at least $\pi/2$. Let us restrict to the case of a closed geodesic c in that is an edge-loop in $\mathrm{lk}(v, \Sigma)$.

First, suppose that c consists of three edges of lengths $(\pi - \pi/k)$, $(\pi - \pi/l)$, and $(\pi - \pi/m)$, for a total circumference of $3\pi - \left(\frac{1}{k} + \frac{1}{l} + \frac{1}{m}\right)\pi$. If $\frac{1}{k} + \frac{1}{l} + \frac{1}{m} > 1$, then the letters corresponding to the vertices of this triangle generate a finite parabolic subgroup, so c is the boundary of a 2-simplex in $\mathrm{lk}(v, \Sigma)$, a contradiction. Thus, the length of c must be at least 2π. Moreover, this length is greater than 2π unless $\frac{1}{k} + \frac{1}{l} + \frac{1}{m} = 1$, in which case the letters generate an affine parabolic subgroup.

Next, note that every edge-loop consisting of four or more edges has length at least 2π. The only such loop whose length is exactly 2π is a square, each of whose sides has length $\pi/2$. Let $s_1, \cdots, s_4$ be the letters corresponding to the four vertices of this square; then $m_{s_i, s_{i+1}} = 2$ (indices modulo 4). Also, we must have $m_{s_1,s_3} = m_{s_2,s_4} = \infty$, since otherwise the square would be filled in by two 2-simplices. Therefore, the parabolic subgroup generated by $s_1, \cdots, s_4$ is the direct product of two infinite dihedral groups and is also an affine Coxeter group of rank 4. This completes the (extremely simplified) proof.

Chapter 3

Original Results

In this chapter, I will present the results of my research on quasiconvex subgroups of word-hyperbolic Coxeter groups.

Recall from Section 1.4 that the property of quasiconvexity ordinarily depends on one's choice of generating set; a subgroup can be quasiconvex with respect to one set of generators but not another. Only in hyperbolic groups is the choice of generators guaranteed to be irrelevant. However, a Coxeter group has a canonical presentation that is extremely closely related to the properties of the group, as the previous chapter showed. The length function with respect to the standard generators has a close geometric connection with the all-important geometric representation on the group, and the Deletion and Exchange Conditions and Tits's solution to the word problem make computation quite easy when using that presentation. Thus, even in the non-hyperbolic case, it is still meaningful to talk about quasiconvexity in Coxeter groups. In any case, the results in this chapter only concern hyperbolic Coxeter groups.

As mentioned above, one frequently considers quasiconvexity in showing that a hyperbolic group is *locally quasiconvex*, i.e., that every finitely generated subgroup is quasiconvex. This property implies the weaker but nevertheless important property of *coherence*: that every finitely generated subgroup is finitely presentable. McCammond and Wise [15] proved that a Coxeter group of rank r is coherent if every exponent m_{ij} is at least r and locally quasiconvex if the inequalities are strict. Kapovich and Schupp [13] proved several results saying that if the exponents m_{ij} in a Coxeter group (of arbitrary rank) are sufficiently large, then any subgroup generated by a sufficiently small number of generators is quasiconvex.

These results mostly concern Coxeter groups of *large type*, those in which every exponent $m_{s,t}$ is at least 3 when $s \neq t$. Such groups are examples of *small cancellation groups*, an often-studied class of groups in which, roughly speaking, the relations are "long words." (See Lyndon and Schupp [14, Ch. V] for an introduction to small cancellation theory.)

I am interested in the opposite type of Coxeter group, *right-angled* groups, in which $m_{s,t}$ is either 2 or ∞ when $s \neq t$. Since the results of small cancellation theory do not apply, a different approach is needed. Patrick Bahls, with whom I worked at the University of Illinois, has conjectured that every word-hyperbolic, right-angled Coxeter group is locally quasiconvex.

According to Theorem 2.19, the parabolic subgroups of Coxeter groups are always convex (with respect to the standard generating set). Therefore, consider the following generalization of the parabolic subgroups:

Definition 3.1. Let (W, S) be a Coxeter system. A finite subset $U \subset W$ is called a *peribolic subset* if no two elements of U share a letter of S in their reduced expressions, i.e., if $T(u) \cap T(u') = \varnothing$ for distinct elements $u, u' \in U$. A subgroup $H \leq W$ generated by a peribolic subset is called a *peribolic subgroup.*[1]

Like parabolic subgroups, peribolic subgroups are made possible by Lemma 2.18, which states that the composition set $T(u) \subset S$ is well-defined. In most groups, such a concept is not very meaningful, since the generators of H could be written in very different ways.

Bahls and I have proven the following result:

Theorem 3.2. *Let (W, S) be a word-hyperbolic, right-angled Coxeter system. Then every peribolic subgroup $H \leq W$ is quasiconvex.*

Proof. We must find a constant k such that for any $w \in H$ and any S-reduced word $\mathbf{x}$ representing w, any truncation $\mathbf{x}'$ of $\mathbf{x}$ represents an element $x' \in W$ such that $d_S(x', H) \leq k$.

Let $U \subset W$ be a peribolic subset that generates H. For each $u \in U$, choose an S-reduced word $\mathbf{u} \in L(\mathbf{S})$ that represents u. Let m_U be the maximum length of the words $\mathbf{u}$; that is, $m_U = \max\{\ell_S(u) \mid u \in U\}$.

Any element $w \in H$ can be written as $w = u_1 \cdots u_r$, where $u_i \in U$. We may assume that r is the least integer for which such a decomposition exists; we say that $u_1 \cdots u_r$ is *U-reduced.* Accordingly, let $\mathbf{w} = \mathbf{u}_1 \cdots \mathbf{u}_r$. Say that $\mathbf{w} = \mathbf{s}_1 \cdots \mathbf{s}_N$, where $\mathbf{s}_\sigma \in \mathbf{S}$. As in the previously chapter, let s_σ denote the image of $\mathbf{s}_\sigma$ in W. (The use of σ as the index variable will be explained shortly.)

As seen in Section 2.4, it is possible to reduce $\mathbf{w}$ to an S-reduced word $\mathbf{w}_0$ by cancelling pairs of identical letters that commute with all the letters in between them. More precisely, suppose that $\mathbf{w}$ is not S-reduced. Then there exist indices $\sigma < \sigma'$ with $\mathbf{s}_\sigma = \mathbf{s}_{\sigma'}$ such that for every index τ between σ and σ', we have $s_\sigma s_\tau = s_\tau s_\sigma$. The word $\mathbf{s}_1 \cdots \widehat{\mathbf{s}_\sigma} \cdots \widehat{\mathbf{s}_{\sigma'}} \cdots \mathbf{s}_N$ then also represents w. Repeating this procedure a finite number of times, we obtain an S-reduced expression for w.

There are obviously many different ways to reduce $\mathbf{w}$ to a S-reduced word. (In light of the Cayley graph $C_S(W)$, we may also refer to an S-reduced word as a *geodesic.*) Not only are multiple geodesics possible, but there may be different sequences of cancellations (involving different pairings of letters) that lead to the same geodesic. For any sequence of cancellations that yields a geodesic, let $\Sigma \subset S \times \mathbb{N} \times \mathbb{N}$ be the set of triples (s, σ, σ') such that $\sigma < \sigma'$, $s_\sigma = s_{\sigma'} = s$, and the pair $(\mathbf{s}_\sigma, \mathbf{s}_{\sigma'})$ is one of the pairs of letters that cancel.

[1]The term "peribolic" is a pun coined by Bahls. The prefix "peri-" means around or near, so a peribolic subgroup is one that is nearly parabolic. For a detailed analysis of why one might find this name funny, see Sigmund Freud's *Jokes and their Relation to the Unconscious.*

Each element of Σ is called a *chord* (for reasons that will be made evident later). Note that no index $\sigma \in \{1, \ldots, N\}$ appears twice in the elements of Σ.

For any chord (s, σ, σ'), let $L(s, \sigma, \sigma')$ be the number of indices i such that $\sigma < i < \sigma'$ and $s_i = s$, i.e., the number of times that the letter $\mathbf{s}$ appears between $\mathbf{s}_\sigma$ and $\mathbf{s}_{\sigma'}$. Let $L(\Sigma)$ denote the sum

$$L(\Sigma) = \sum_{(s,\sigma,\sigma') \in \Sigma} L(s, \sigma, \sigma'),$$

a nonnegative integer. A sequence of cancellations is called *minimal* if $L(\Sigma)$ is minimal; for any word $\mathbf{w}$, there exists a minimal sequence, by the well-ordering of the natural numbers.

Note that $L(\Sigma)$ is not necessarily zero for a minimal sequence. For instance, if $st \neq ts$, then the only way to reduce the word $\mathbf{w} = \mathbf{stssts}$ to the empty word is to cancel the inner pair of $\mathbf{s}$'s first, then the $\mathbf{t}$'s, and finally the outer pair of $\mathbf{s}$'s; we then have $L(\Sigma) = 2$. On the other hand, if $st = ts$, then this sequence is no longer minimal, since we could instead cancel the first $\mathbf{s}$ with the second and the third with the fourth and obtain $L(\Sigma) = 0$.

For $1 \leq j \leq r$, let $\mathbf{w}_j = \mathbf{u}_1 \cdots \mathbf{u}_j$, and let N_j be the length of $\mathbf{w}_j$: that is, $\mathbf{w}_j = \mathbf{s}_1 \cdots \mathbf{s}_{N_j}$. Here is the key lemma:

Lemma 3.3. *Let (W, S) be a right-angled, word-hyperbolic Coxeter system, and let $H \leq W$ be a peribolic subgroup generated by U. There exists a constant B, determined solely by W, S, H, and U, such that the following is true:*

Let $\mathbf{w} = \mathbf{u}_1 \cdots \mathbf{u}_r$ be a U-reduced word, and let Σ be the chord set for a minimal sequence of cancellations for $\mathbf{w}$. Then for each $1 \leq j \leq r$, the number of chords $(s, \sigma, \sigma') \in \Sigma$ such that $\sigma \leq N_j < \sigma'$ is less than B.

We defer the proof of Lemma 3.3 and first show how it implies that H is quasiconvex.

Let $\widetilde{\mathbf{w}}$ be the S-reduced expression for w obtained from a minimal sequence of cancellations on $\mathbf{w}$, and let $\mathbf{x}$ be an arbitrary S-reduced expression for w. The paths in the Cayley graph $\mathcal{C}_S(W)$ corresponding to $\mathbf{x}$ and $\widetilde{\mathbf{w}}$ begin and end at the same point, so they can be viewed as a degenerate triangle in which one side is simply a point. Since $\mathcal{C}_S(W)$ is δ-hyperbolic for some constant $\delta \geq 0$, each of these paths is contained in the δ-neighborhood of the other. Therefore, if $\mathbf{x}'$ is any truncation of $\mathbf{x}$, there exists a truncation $\widetilde{\mathbf{w}}'$ of $\widetilde{\mathbf{w}}$ such that $d_S(x', \widetilde{w}') \leq \delta$ (where, as usual, $x' = \psi(\mathbf{x}')$ and $w' = \psi(\mathbf{w}')$.) (See Figure 16.)

We may write $\widetilde{\mathbf{w}}$ as a product $\widetilde{\mathbf{u}}_1 \cdots \widetilde{\mathbf{u}}_r$, where $\widetilde{u}_i$ is the word obtained from $\mathbf{u}_i$ in the reduction of $\mathbf{w}$ to $\widetilde{\mathbf{w}}$. For $1 \leq j \leq r$, let $\widetilde{\mathbf{w}}_j = \widetilde{\mathbf{u}}_1 \cdots \widetilde{\mathbf{u}}_j$ and $\widetilde{w}_j = \psi(\mathbf{w}_j)$. Choose j to be the least value for which the word $\widetilde{\mathbf{w}}'$ is a truncation of $\widetilde{w}_j$. Since the length of each word $\widetilde{\mathbf{u}}_i$ is at most m_U, we therefore have $d_S(\widetilde{w}', \widetilde{w}_j) \leq m_U$.

Let $\Sigma' \subset \Sigma$ be the set of all chords $(s, \sigma, \sigma') \in \Sigma$ such that $\sigma \leq N_j < \sigma'$. By Lemma 3.3, the cardinality of Σ' is at most B. Denote the elements of Σ' by $(s_1, \sigma_1, \sigma_1'), \ldots, (s_k, \sigma_k, \sigma_k')$, where $\sigma_1 < \cdots < \sigma_k$.

Let $\mathbf{y} = \mathbf{s}_{\sigma_k'} \mathbf{s}_{\sigma_{k-1}'} \cdots \mathbf{s}_{\sigma_1'}$. Using Tits's algorithm, we may reduce the word $\mathbf{w}_j\mathbf{y}$ to a geodesic by successively cancelling each pair of letters $(\mathbf{s}_{\sigma_i}, \mathbf{s}_{\sigma_i'})$ (as i ranges from k to 1), as well as all pairs $(\mathbf{s}_\tau, \mathbf{s}_{\tau'})$ that cancel in the reduction of $\mathbf{w}$ and for which $\tau < \tau' \leq N_j$.

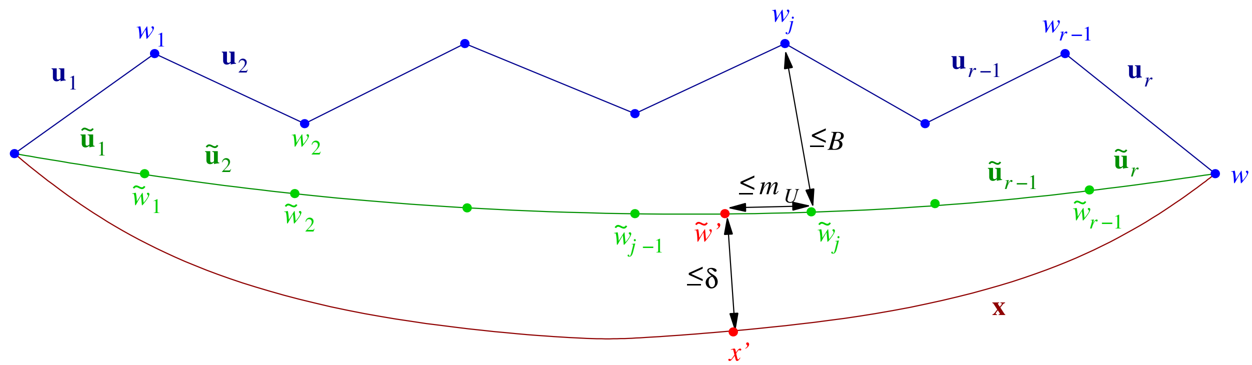

Figure 16: Schematic of the proof of Theorem 3.2.

The resulting geodesic is simply $\tilde{\mathbf{w}}_j$. Therefore, $w_j^{-1}\tilde{w}_j = y$, so $d_S(\tilde{w}_j, w_j) \le \ell_S(y) \le k \le B$. Note that $w_j \in H$.

By the triangle inequality, we have $d(x', w_j) \le K$, where $K = \delta + m_U + B$. The entire geodesic $\mathbf{x}$ thus lies within the $(K+1)$-neighborhood of the elements of H in the Cayley graph $\mathcal{C}_S(W)$. (The extra 1 is to account for points in the interiors of edges.) As K depends only on the choice of W, S, H, and U, we thus see that H is $(K+1)$-quasiconvex. $\square$

Note that the only part of this proof that depends on the construction of H — specifically, the fact that H is peribolic — is Lemma 3.3. In an unpublished paper, Bahls first developed the technique of bounding the cancellations of pairs in order to show that a different type of subgroup is quasiconvex: namely, subgroups generated by reflections. (For those subgroups, there is an easier proof.) Assuming that one can prove some analogue of Lemma 3.3, it may be possible to extend this type of argument to other types of Coxeter groups.

The proof of Lemma 3.3 is extremely detailed and would take many pages to present in its entirely. However, we will give a rough outline of the proof.

First of all, let us explain some of the nomenclature used above. We may represent the word $\mathbf{w}$ as a subdivision of a line segment into N intervals, or *cells*, corresponding to the letters $\mathbf{s}_i$. Each cell is labelled with the ordered pair $(s_\sigma, \sigma) \in S \times \mathbb{N}$. (Typically, we use a Roman letter for the element of S and the corresponding Greek letter for the index: thus, (s, σ), (t, τ), etc.) A cell (s, σ) is called an s-cell; by abuse of notation, we often refer to the cell (s, σ) as simply σ.

We may keep track of a sequence of cancellations by drawing an arc, or *chord*, joining each pair of cells that corresponds to a pair of cancelling letters. The resulting *cancellation diagram* is extremely useful for figuring out what commutation relations must hold in the group. For instance, if (t, τ) is a cell without a chord and lies between the endpoints of the chord (s, σ, σ'), then $st = ts$. Similarly, if the diagram contains chords (s, σ, σ') and (t, τ, τ') with $\sigma < \tau < \sigma' < \tau'$, then again we have $st = ts$, since one pair of letters must cancel before the other one. (Such chords are said to cross; this terminology makes sense if we draw all the chords below the base line segment.)

We may define a partial order on the chord set Σ by nesting: $(s, \sigma, \sigma') \prec (t, \tau, \tau')$ if and only if $\tau < \sigma < \sigma' < \tau'$. Many proofs involve induction on this partial ordering.

The cancellation diagram for a minimal sequence of cancellations is called a minimal diagram. A minimal diagram $\mathcal{D}$ satisfies several important properties. First of all, if (s, σ, σ') is a chord and (s, ψ) is an s-cell that lies between σ and σ', then ψ must contain an endpoint of a chord whose other endpoint is also between σ and σ'. Otherwise, we could replace the chord (s, σ, σ') with (s, σ, ψ) or (s, ψ, σ) and obtain a valid cancellation diagram, contradicting minimality of $\mathcal{D}$. Moreover, if $(s, \psi, \psi') \prec (s, \sigma, \sigma')$, then there exists a chord (t, τ, τ'), with $st \neq ts$, such that $(s, \psi, \psi') \prec (t, \tau, \tau') \prec (s, \sigma, \sigma')$. Otherwise, we could s-chords with (s, σ, ψ) and (s, ψ', σ') and obtain a valid diagram, again contradicting the minimality of $\mathcal{D}$.

When a word is given as a product $\mathbf{w} = \mathbf{u}_1 \cdots \mathbf{u}_r$, where the $\mathbf{u}_i$ are generators of a given subgroup (or their inverses), we may accordingly group the cells into *intervals* $I_1, \ldots, I_r$. An interval that represents either u or u^{-1} is called a u-interval. When there exists a chord (s, σ, σ') such that that σ is in interval I_i and σ' is in interval $I_{i'}$, the pair $(I_i, I_{i'})$ is called a *matched pair*. Since the $\mathbf{u}_i$ are reduced words, the endpoints of any chord must be contained in different intervals.

For peribolic subgroups, note that if $(I_i, I_{i'})$ is a matched pair joined by a chord (s, σ, σ'), then the intervals I_i and $I_{i'}$ both represent the same word $\mathbf{u}$ or its inverse, since no other element of U includes the letter s. Using this fact, it is possible to get an extremely good control on the behavior of the chords. For instance, one important lemma is that if σ is the m^{th} s-cell from the right end of I_i, then σ' is the m^{th} s-cell from the left end of $I_{i'}$.

The main technical result is as follows:

Proposition 3.4. *Let (W, S) be a right-angled, word-hyperbolic Coxeter system, and let H be a peribolic subgroup generated by U. Say that $\mathbf{w} = \mathbf{u}_1 \cdots \mathbf{u}_r$ as above, and consider a minimal cancellation diagram for $\mathbf{w}$. Let $s, t \in S$, and suppose that there are chords*

$$(s, \sigma_1, \sigma_1') \prec (t, \tau_1, \tau_1') \prec (s, \sigma_2, \sigma_2') \prec (t, \tau_2, \tau_2')$$

satisfying either of the following conditions:

1. *For some distinct $u, v \in U$, we have $s \in T(u), t \in T(v)$, so no two of the eight cells are contained in the same interval.*

2. *For some $u \in U$, we have $s, t \in T(u)$, and the four pairs of cells $\{\tau_2, \sigma_2\}$, $\{\tau_1, \sigma_1\}$, $\{\sigma_1', \tau_1'\}$, and $\{\sigma_2', \tau_2'\}$ are contained in separate u-intervals.*

Then $st = ts$.

Sketch of proof. Suppose that $st \neq ts$. The idea is to find two non-commuting letters $x, y \in S$ such that $sx = sx$, $tx = xt$, $sy = ys$, and $ty = yt$. The parabolic subgroup W_T generated by $T = \{s, t, x, y\}$ will then be the direct product of two copies of the infinite dihedral group. However, this is an affine Coxeter group of rank 4, which contradicts Theorem 2.20.

Given any set of indices $1 \leq i_1 < \cdots < i_k \leq r$, note that there must be at least one cell in the union $I_{i_1} \cup \cdots \cup I_{i_k}$ that either has no chord or has a chord whose other endpoint

Case 1:

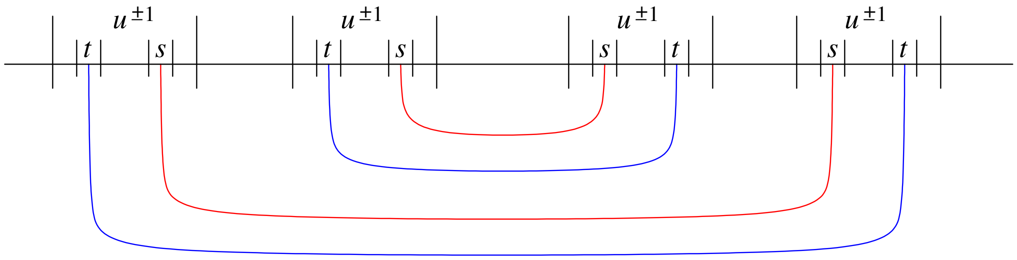

Case 2:

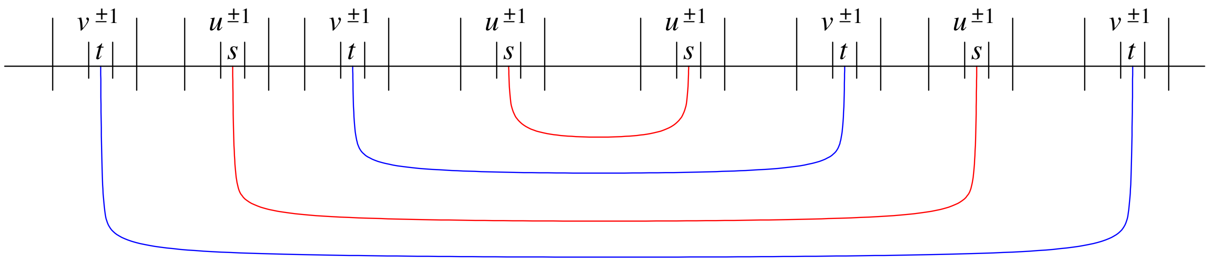

Figure 17: The two types of cancellation diagrams in Proposition 3.4.

is not in $I_{i_1} \cup \cdots \cup I_{i_k}$. Otherwise, the word $\mathbf{w}' = \mathbf{u}_1 \cdots \widehat{\mathbf{u}_{i_1}} \cdots \widehat{\mathbf{u}_{i_k}} \cdots \mathbf{u}_r$ would also be an expression for w, contradicting the fact that $\mathbf{w}$ is U-reduced.

Roughly speaking, we may use this technique to obtain a non-cancelling cell (x, ξ) that appears in the same interval as either σ_1 or σ_2 (varying in different cases). By a counting argument, there is another non-cancelling x-cell (x, ξ') to the right of ξ. As a result, there must be a cell (y, v) between ξ and ξ', with $xy \neq yx$, preventing the formation of a chord (x, ξ, ξ'). Figuring out the other commutation relations is an extremely tedious exercise in diagram-chasing, so I will not show the details here. $\square$

Proof of Lemma 3.3. In terms of cancellation diagrams, the lemma says the number of chords with one endpoint in $I_1 \cup \cdots \cup I_j$ and the other in $I_{j+1} \cup \cdots \cup I_r$ is universally bounded. We may prove this fact using repeated applications of the Pigeonhole Principle. For any $s \in S$, the s-chords of this type form a chain in the partial ordering on chords, since s-chords cannot cross: $(s, \sigma_1, \sigma_1') \prec \cdots \prec (s, \sigma_k, \sigma_k')$. By the minimality of the diagram, for $1 \leq i \leq k-1$, there must be a chord (t_i, τ_i, τ_i') such that $(s, \sigma_i, \sigma_i') \prec (t_i, \tau_i, \tau_i') \prec (s, \sigma_{i+1}, \sigma_{i+1}')$. If k is sufficiently large, then one of the letters t_i will repeat in one of the ways forbidden by Proposition 3.4, so k must be less than some constant k_0. Therefore, by the Pigeonhole Principle, the total number of chords joining $I_1 \cup \cdots \cup I_j$ and $I_{j+1} \cup \cdots \cup I_r$ is bounded by $B = k_0 \, |U| + 1$. $\square$

Poetic Conclusion

We ask, “What good can possibly come of
Turning a group into a metric space?”
When this idea was put forth by Gromov,
Group theory took off at a rapid pace.

Word-hyperbolic groups have much in store.
They all can be presented finitely,
And they have useful properties galore
When they are quasiconvex locally.

In Cox’ter groups, we learn lots from how long
The words are for a given element.
To tie it all together, there’s Moussong
Whose proof is anything but evident.

My thesis is complete, and sad to tell
I soon must bid fair Harvard a farewell.

Bibliography

[1] F.H. Bool, J.R. Kist, J.L. Locher, and F. Wierda, *M.C. Escher: His Life and Complete Graphic Work*, Abradale Press, 1992.

[2] N. Bourbaki, *Groupes et Algèbres de Lie, chapitres 4–6*, Masson, 1981.

[3] M. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Grundlehren der mathematischen Wissenschaften, vol. 319, Springer-Verlag, 1999.

[4] M. Coornaert, T. Delzant, and A. Papadopoulos, *Géometrie et Théorie des Groupes*, Lecture Notes in Mathematics, vol. 1441, Springer-Verlag, 1990.

[5] H.S.M. Coxeter, *Discrete groups generated by reflections*, Annals of Mathematics **35** (1934), 588–620.

[6] ———, *The complete enumeration of finite groups of the form* $R_i^2 = (R_iR_j)^{k_{ij}} = 1$, Journal of the London Mathematical Society **10** (1935), 21–25.

[7] M.W. Davis and G. Moussong, *Notes on nonpositively curved polyhedra*, Low Dimensional Topology (K. Böröczky, W. Neumann, and A. Stipicz, eds.), Bolyai Society Mathematical Studies, vol. 8, János Bolyai Mathematical Society, 1999, pp. 11–94.

[8] É. Ghys and P. de la Harpe, *Sur les Groupes Hyperboliques d'après Mikhael Gromov*, Progress in Mathematics, vol. 83, Birkhäuser, 1990.

[9] M. Gromov, *Hyperbolic groups*, Essays in Group Theory (S.M. Gersten, ed.), Mathematical Sciences Research Institute Publications, vol. 8, Springer-Verlag, 1987, pp. 75–263.

[10] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.

[11] J. Humphreys, *Reflection Groups and Coxeter Groups*, Cambridge Studies in Advanced Mathematics, vol. 29, Cambridge University Press, 1990.

[12] I. Kapovich, *Amalgamated products and the Howson property*, Canadian Mathematics Bulletin **40** (1997), no. 3, 330–340.

[13] I. Kapovich and P.E. Schupp, *Bounded rank subgroups of Coxeter groups, Artin groups, and one-relator groups with torsion*, Proceedings of the London Mathematical Society **88** (2004), no. 1, 88–113.

[14] R.C. Lyndon and P.E. Schupp, *Combinatorial Group Theory*, Ergebnisse der Mathematik und ihrer Grenzgebiete, vol. 89, Springer-Verlag, 1977.

[15] J. McCammond and D. Wise, *Locally quasiconvex small cancellation groups*, Preprint.

[16] G. Moussong, *Hyperbolic Coxeter groups*, Ph.D. thesis, The Ohio State University, 1988.

[17] E. Rips, *Subgroups of small cancellation groups*, Bulletin of the London Mathematical Society **14** (1982), 45–47.

[18] J. Stillwell, *Classical Topology and Combinatorial Group Theory*, Graduate Texts in Mathematics, vol. 72, Springer-Verlag, 1993.

[19] J. Tits, *Le problème des mots dans les groupes de Coxeter*, Symposia Mathematica (INDAM, Rome, 1967/68), vol. 1, Academic Press, 1969, pp. 175–185.

[20] È.B. Vinberg, *Discrete linear groups generated by reflections*, Mathematics of the USSR Izvestija **5** (1971), no. 5, 1083–1119.