Title
Limit Roots of Lorentzian Coxeter Systems

Author
Brody, Nicholas R.

Publication Date
2016

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Santa Barbara

Limit Roots of Lorentzian Coxeter Systems

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

Master of Arts

in

Mathematics

by

Nicolas R. Brody

Committee in Charge:

Professor Jon McCammond, Chair

Professor Darren Long

Professor Stephen Bigelow

June 2016

The Thesis of
Nicolas R. Brody is approved:

Professor Darren Long

Professor Stephen Bigelow

Professor Jon McCammond, Committee Chairperson

May 2016

Limit Roots of Lorentzian Coxeter Systems

Acknowledgements

My time at UC Santa Barbara has been inundated with wonderful friends, teachers, mentors, and mathematicians. I am especially indebted to Jon McCammond, Jordan Schettler, and Maribel Cachadia.

A very special thanks to Jean-Phillippe Labb for sharing his code which was used to produce Figures 7.1, 7.3, and 7.4.

This thesis is dedicated to my three roommates: Viki, Luna, and Layla. I hope we can all be roommates again soon.

Abstract

Limit Roots of Lorentzian Coxeter Systems

Nicolas R. Brody

A reflection in a real vector space equipped with a positive definite symmetric bilinear form is any automorphism that sends some nonzero vector v to its negative and pointwise fixes its orthogonal complement, and a finite reflection group is a group generated by such transformations that acts properly on the sphere of unit vectors. We note two important classes of groups which occur as finite reflection groups: for a 2-dimensional vector space, we recover precisely the finite dihedral groups as reflection groups, and permuting basis vectors in an n-dimensional vector space gives a way of viewing a symmetric group as a reflection group.

A Coxeter group is a generalization of a finite reflection group, whose rich geometric and algebraic properties interact in surprising ways. Any finite rank Coxeter group W acts faithfully on a finite dimensional real vector space V. Each such group has an associated symmetric bilinear form that it preserves, and the signature of this bilinear form contains valuable information about W. When it has type $(n, 1)$, we call such a group Lorentzian, and there is a natural action of such a group on a hyperbolic space. Inspired by a conjecture of Dyer in 2011, Hohlweg, Labbé and Ripoll have studied the set of reflection vectors in Lorentzian Coxeter groups. We summarize their results here. The reflection vectors form an infinite discrete subset of the vector space V, but the projective version of the reflection vectors has limit points. Understanding these limit points is the

primary goal of this text. They lie within the light cone of the Lorentz space, and have intimate connections with the boundary of the group.

Contents

Introduction

This thesis is an exposition of the nascent theory of limit roots. No prior knowledge of Coxeter groups is assumed, and all relevant definitions and concepts are introduced. Detours are taken along the way to explore properties of the wide-ranging and diverse theory of Coxeter groups, even if these properties are not directly related to limit roots. The thesis is structured as follows. The first part provides many background definitions and results that are necessary to understand the main points of interest, the second part defines Coxeter groups and studies them in general, and the third part focuses on the specific case of Lorentzian Coxeter groups. A more detailed description of the content is as follows.

Part I begins with a short introduction to some basics of geometric group theory to give the reader a sense of the connections between the group theory and the geometry. We build up to the Schwarz-Milnor Lemma, which says that for every finitely generated group, there is an essentially unique geometric space it can act on in a certain way. Moreover, this lemma provides a way to find this space. In the next chapter, we then recall some linear algebraic facts and definitions that will be useful throughout. In particular, we consider affine linear algebra and bilinear forms. We prove that every isometry of a quadratic space is a product of reflections, and that bilinear forms are determined by the signs of their "eigenvalues". We then turn our attention to hyperbolic geometry so that we can later consider group actions on hyperbolic space. We construct a few models of hyperbolic geometry and deduce some first properties, which will prove important when we study actions of Coxeter groups on such spaces.

After setting the stage in this way, we begin to study reflection groups in Part II. We first focus on finite reflection groups since these motivate the general definition of Coxeter groups, whose definition looks rather obscure at first glance. In this chapter we first encounter root systems, and we think about things very geometrically until we find the algebraic conditions for a finite reflection group. This enables us to define Coxeter groups in general, and then we do not waste much time describing a representation as a reflection group. We show that this contains the theory of finite reflection groups as a special case, and then we uncover the various properties that determine which type of space a given Coxeter group can act on.

In Part III, upon identifying Lorentzian space as the most tractable class of Coxeter groups which are not completely well-understood, we develop language in order to ask the question that motivates the entire thesis: how are the roots of a Lorentzian Coxeter group distributed among the vector space? This problem has caught the attention of a number of mathematicians recently. In the last five years, Dyer, Hohlweg, Ripoll, Preaux, Labbé, and Chen have studied this question. By considering the roots projectively, it is possible to find some accumulation points of roots, which turn out to lie within the light cone. Some alternative notions of "limits" of a Lorentzian Coxeter group follow this, one combinatorial and one arising from the group action. Some results on when these notions coincide are provided.

Part I

Geometric Preliminaries

The primary aim in Part I is to introduce the concepts and definitions which will be used in the remainder of the thesis. It contains some basic theorems that are crucial to understanding Coxeter groups, but nothing in this section lies within the domain of Coxeter theory.

Chapter 1 provides some basic concepts in the field of geometric group theory. Geometric group actions and Cayley graphs are defined, and the Schwarz-Milnor Lemma, which provides a correspondence between groups and metric spaces, is proved.

Chapter 2 sets notation and contains definitions in linear algebra, including affine linear algebra and bilinear forms. The Cartan-Dieudonne theorem and Sylvester's Law of Inertia are included in this chapter.

Chapter 3 develops the theory of hyperbolic geometry, and discusses some models of hyperbolic space. Also present is the theory of δ-hyperbolic spaces which are connected to the group theory in Chapter 8.

1. Geometric Group Theory

Historically, groups were studied to understand symmetries of objects, and not as objects in their own right. The modern abstract approach to group theory, though very powerful and beautiful itself, has divorced group theory from this underlying geometry. The burgeoning field of geometric group theory is amending this disconnect. As mentioned, the connections between geometry and group theory are the focus of this thesis, so time is taken to introduce some relevant concepts. Specifically, geometric group actions are introduced and the Schwarz-Milnor lemma is proven.

1.1 Group Actions

Group actions come up in a variety of settings. Here, some concepts for a group acting a metric space are defined.

Definition 1.1.1 (Geometric Definitions)**.** Let (X, d) be a metric space. The *(open) ball of radius r about $x \in X$* is the set $B_r(x) = \{y \in X \mid d(x, y) < r\}$. Its closure is the *closed ball of radius r about $x \in X$*, and is $\overline{B}_r(x) = \{y \in X \mid d(x, y) \leq r\}$. For a subset S of X, define the *(open) ball of radius r about S* to be $B_r(S) = \cup_{x \in S} B_r(x)$. Say X is *proper* if for every $x \in X$ and $r \in \mathbb{R}$, the closed ball about x of radius r is compact.

Say that z is *between* x and y in a metric space (X, d) if the triangle inequality $d(x, y) = d(x, z) + d(z, y)$ is an equality. A (closed) *interval* $[x, y]$ in a metric space is the set of points between x and y, i.e. $[x, y] = \{z \in X \mid d(x, z) + d(z, y) = d(x, y)\}$. This notation extends to (x, y) as expected.

Say a subset $K \subseteq X$ is *convex* if whenever $x, y \in K$ and $z \in [x, y]$, then $z \in K$. Given a subset Y of a metric space (X, d), the *convex hull of* Y, denoted $\mathsf{conv}(Y)$, is the intersection of all convex subsets of X containing Y.

There is an induced partial order on $[x, y]$ such that $z \leq z'$ if and only if $z \in [x, z']$, or equivalently, $z' \in [z, y]$. This ordering is not total in general; note that the interval between antipodal points on a sphere is the entire space. However, this notion is used only in the setting of a vector space, in which case there is an order isomorphism $[0, 1] \to [x, y] = \{x + t(y - x) \mid t \in [0, 1]\}$.

A *curve* in X is a continuous map $\gamma\colon [a, b] \to X$. A curve is a *geodesic* if γ is an isometric embedding; that is, if $d(x, y) = d(\gamma(x), \gamma(y))$ for every $x, y \in [a, b]$. Say (X, d) is a *geodesic metric space* if for every pair $x, y \in X$, there is a geodesic $\gamma\colon [0, d(x, y)] \to X$ so that $\gamma(0) = x$ and $\gamma(d(x, y)) = y$. A *geometry* is a proper geodesic metric space.

Example 1.1.2. A subset of $\mathbb{R}^n$ with the metric of $\mathbb{R}^n$ is a geodesic metric space if and only if it is convex. Some of the time, this can be remedied by endowing a connected subset of $\mathbb{R}^n$ with the *induced intrinsic metric*, in which $d'(x, y) = \inf\{\mathsf{length}(\gamma) \mid \gamma \text{ is a curve from } x \text{ to } y\}$. However, this process does not turn $\mathbb{R}^n \setminus \{0\}$ into a geodesic metric space since there is no geodesic connecting V to $-v$.

Definition 1.1.3 (Group Actions)**.** A *group action* of a group G on a mathematical structure X is a homomorphism $\phi\colon G \to \mathsf{Aut}(X)$, where $\mathsf{Aut}(X)$ is the group of structure-preserving maps from X to itself. For example, if X is a topological space, $\mathsf{Aut}(X)$

consists of homeomorphisms from X to X; if S is a set, $\mathsf{Aut}(S)$ consists of bijections from S to S; if V is a vector space, $\mathsf{Aut}(V)$ consists of invertible linear transformations from V to V ($\mathsf{Aut}(V)$ is more commonly known as $\mathsf{GL}(V)$ in this case. A group action in this sense is typically called a *representation*). Writing $g.x$ for $\phi(g)(x)$, an action satisfies (i) $1_G.x = x$ and (ii) $g.(h.x) = (gh).x$, and these properties characterize group actions.

Definition 1.1.4 (Quotients)**.** Suppose a group G acts on a topological space X. Partition X into its G-orbits and denote by X/G the set of equivalence classes. Let $q\colon X \to X/G$ by $q(x) = [x]$, and give X/G the quotient topology. That is, $U \subseteq X/G$ is open if and only if $q^{-1}(U)$ is open in X. Call X/G the *quotient of the space X by the group G*.

Definition 1.1.5 (Geometric group action)**.** Let (X, d) be a proper geodesic metric space and let G be a finitely generated group. The group G acts on the space X by *isometries* if whenever $g \in G$ and $x, y \in X$, then $d(g.x, g.y) = d(x, y)$. The group G acts *cocompactly* if the quotient space X/G is compact. An action is *properly discontinuous* if for every compact set $K \subseteq X$, there are only finitely many $g \in G$ so that $gK \cap K \neq \varnothing$ (so in particular, the stabilizer of any point is finite). When the action of G on X is properly discontinuous, cocompact, and by isometries, it is called a *geometric* action.

Remark. In a sense, to demand that the action is cocompact guarantees that the space is not too big for the group. Conversely, the group is not too large because the action is properly discontinuous. These two properties force geometric group actions to balance out in a nice way.

Note that as long as X is a proper metric space and the group action is isometric, the following condition is equivalent to proper discontinuity: For every $x \in X$, there exists an $r > 0$ so that the set $\{g \in G \mid g.B_r(x) \cap B_r(x) \neq \varnothing\}$ contains finitely many group elements.

To see that this implies the definition, take $x \in X$ and any $r > 0$; then since X is proper, $\overline{B_r(x)}$ is compact, and hence there are only finitely many group elements so that $g.\overline{B_r(x)} \cap \overline{B_r(x)} \neq \varnothing$. For the other direction, let K be a compact set. By assumption, for each $x \in K$, there is some r_x so that $g.B_{r_x}(x)$ intersects $B_{r_x}(x)$ finitely many times. The set $\{B_{r_x/2}(x) \mid x \in K\}$ is an open cover for K, so it admits a finite subcover, say by $x_1, \ldots, x_n$. If any set $\{g \in G \mid g.B_{r_i/2}(x_i) \cap B_{r_j/2}(x_j)\}$ is infinite, then there is an injective sequence $\{g_k\}$ in this set. But then $\{g_1^{-1}g_k \mid k \in \mathbb{N}\}$ is an infinite subset of G with $g_1^{-1}g_k.B_{r_i}(x_i) \cap B_{r_i}(x_i)$ for each k, a contradiction.

Example 1.1.6 (A geometric action)**.** Consider the action of $\mathbb{Z}$ on $\mathbb{R}$ by translation; that is, $n.x = n + x$ for $n \in \mathbb{Z}$ and $x \in \mathbb{R}$. Indeed, we have $|n.x - n.y| = |x - y|$, $\mathbb{R}/\mathbb{Z} \cong S^1$, and for any $x \in \mathbb{R}$, $|\{n \in \mathbb{Z} \mid n.B_{1/4}(x) \cap B_{1/4}(x) \neq \varnothing\}| = 1$.

1.2 Cayley Graph

For a finitely generated group G, build a directed graph with vertices in bijection with the group elements and the edges leaving each vertex corresponding to a fixed finite generating set. Then the group G can act on this graph by left multiplication, and this action is geometric.

Definition 1.2.1 (Cayley Graph)**.** If S is a subset of a group G, build a directed graph called the *(right) Cayley graph* using group elements as vertices, and directing an edge from g to gs whenever $s \in S$. That is, let $\Gamma_{(G,S)}$ have vertex set G and edge set $E = \{(g, gs) \mid g \in G, s \in S\}$.

If S is finite, then $\Gamma_{(G,S)}$ is locally finite, meaning each vertex is incident to only finitely many edges. If S is a generating set, then $\Gamma_{(G,S)}$ is connected. If S is symmetric, meaning $S = S^{-1}$, then $\Gamma_{(G,S)}$ may be viewed as undirected, since whenever (g, gs) is an edge, so is $(gs, gss^{-1} = g)$. This pair of directed edges becomes a single undirected edge. If the identity is not in S, then $\Gamma_{(G,S)}$ has no loops.

Definition 1.2.2 (Metrizing Γ)**.** Henceforth, assume S is a finite symmetric generating set not containing 1. The Cayley graph $\Gamma_{(G,S)}$ becomes a metric space using the graph distance. A *path* in a graph $\Gamma = (V, E)$ is a function $\gamma \colon \{0, 1, \ldots, n\} \to \Gamma$ so that $(\gamma(i-1), \gamma(i)) \in E$ for $i = 1, \ldots, n$. In this case n is *the length of the path*, and γ is *a path from* $\gamma(0)$ *to* $\gamma(n)$. Since the graph $\Gamma_{(G,S)}$ is connected and undirected, $d_S(x, y)$ (or just $d(x, y)$) is the minimal length of a path from x to y (a path exists since the graph is connected). Since paths correspond to sequences of right multiplications, this is the same as finding elements of S so that $xs_1 \ldots s_k = y$, or equivalently, finding minimal length expressions for $x^{-1}y$. This defines a metric on the vertices of $\Gamma_{(G,S)}$ since $d(x, y) = 0$ if and only if $x = y$, $d(x, y) = d(y, x)$ since the graph is undirected, and the triangle inequality is satisfied since a path from x to y followed by a path from y to z gives a path from x to z of length $d(x, y) + d(y, z)$, and so $d(x, z)$ is bounded above by this number. This metric is often called the *word metric* on G.

By metrically identifying each edge in the graph with an interval $[0, 1]$, the Cayley graph becomes a proper geodesic metric space. The distance between two vertices has already been defined, and it is clear how to calculate the distance between two points on the same edge. So suppose $x \in (g, gs)$, and $y \in G$. Then $d(x, y) = \min\{d(x, g) + d(g, y), d(x, gs) + d(gs, y)\}$. Then the graph is proper as a metric space since the graph is locally finite, and geodesics exist because distances are defined as the length of an already specified path.

The group G acts geometrically on its Cayley graph by extending the action of left multiplication on the vertices to the edges using the metric in the following way. Let $g \in G$, $\varepsilon \in (0, 1)$, $s \in S$. Write $g + \varepsilon s$ for the point on the edge (g, gs) which is ε away from g. Then $h.(g + \varepsilon s) = hg + \varepsilon s \in (hg, hgs)$. It is clear that $d(x, y) = d(g.x, g.y)$ for each $g \in G$ and $x, y \in \Gamma_{(G,S)}$. It also happens that $\Gamma_{(G,S)}/G$ is a wedge of $|S|$ circles, so this action is compact. And the only group element $g \in G$ for which $g.B_{1/4}(x)$ intersects $B_{1/4}(x)$ is the identity, so indeed this is a geometric action.

Surprisingly, this is pretty much the only type of geometric group action to be found.

1.3 Quasi-isometries

In order to make the desired correspondence between groups and metric spaces, it should be reasonably clear that a somewhat coarse identification of metric spaces is necessary. For example, the trivial action of any finite group is geometric on any compact metric space. This correspondence can only see the large-scale properties of either category, so the equivalence relation is chosen accordingly.

Definition 1.3.1 (Quasi-isometry)**.** A function $f\colon (X, d_X) \to (Y, d_Y)$ is called a *quasi-isometric embedding* if there are constants $A \geq 1$ and $B \geq 0$ so that whenever x_1, x_2 in X, then

$$\frac{1}{A} d_X(x_1, x_2) - B \leq d_Y(f(x_1), f(x_2)) \leq A d_X(x_1, x_2) + B.$$

A function $f\colon (X, d_X) \to (Y, d_Y)$ is called *quasi-onto* if there exists a $C \geq 0$ such that for every point $y \in Y$, there is a point $x \in X$ so that $d_Y(f(x), y) \leq C$.

If f is a quasi-isometric embedding which is quasi-onto, f is called a *quasi-isometry*.

In light of Definition 1.1.1, the quasi-onto property can be restated as $\overline{B_C(f(X))} = Y$.

Remark. It is equivalent to define quasi-isometric embeddings allowing the use of four constants $A_1, A_2 \geq 1$ and $B_1, B_2 \geq 0$ and requiring instead the similar condition $\frac{1}{A_1} d_X(x_1, x_2) - B_1 \leq d_Y(f(x_1), f(x_2)) \leq A_2 d_X(x_1, x_2) + B_2$. Of course, any map satisfying the official definition satisfies this new one, and to show the new one implies the official one, just take $A = \max\{A_1, A_2\}$ and $B = \max\{B_1, B_2\}$.

Metric spaces X and Y are *quasi-isometric* if there exists a quasi-isometry between them. This choice of language makes it sound like an equivalence relation.

Proposition 1.3.2. *Quasi-isometry is an equivalence relation.*

Proof. Of course the identity map is a quasi-isometry from a space to itself. If $f\colon X \to Y$ is a quasi-isometry, we look for a function $g\colon Y \to X$ which is a quasi-isometry. Given $y \in Y$, choose an $x \in X$ so that $d_Y(f(x), y) \leq C$, and set $g(y) = x$. Verifying that g is a quasi-isometry is straightforward, as is verifying that a composition of quasi-isometries is again a quasi-isometry. $\square$

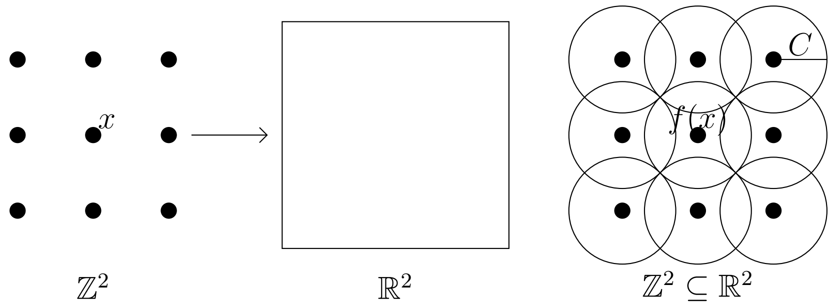

Figure 1.1: A quasi-isometry from $\mathbb{Z}^2$ to $\mathbb{R}^2$.

At this point, the reader should contemplate what equivalence classes look like here. It isn't hard to see that any bounded metric space is quasi-isometric to a point, $\mathbb{R}^n$ is quasi-isometric to $\mathbb{Z}^n$, and X is quasi-isometric to $X \times [0,1]$.

The inclusion map is an isometric embedding of $\mathbb{Z}^n \to \mathbb{R}^n$, so we may take $A = 1$ and $B = 0$. It is "quasi-onto" with $C = \sqrt{n}$. An inverse map can be found by mapping points in $\prod_{i=1}^{n}[x_i - \frac{1}{2}, x_i + \frac{1}{2})$ to $(x_1, \ldots, x_n) \in \mathbb{Z}^n$. This time take $C = 0$ because the map is surjective, but take $A = 1$ and $B = \sqrt{n}$.

Example 1.3.3 (Infinite generating set)**.** Returning to the previous example, the Cayley graph of $\mathbb{Z}$ with respect to $S = \{\pm 1\}$ is isometric (hence quasi-isometric) to $\mathbb{R}$. However, the Cayley graph of $\mathbb{Z}$ with respect to the generating set $S = \mathbb{Z}$ tells a different tale. Although the action of $\mathbb{Z}$ on $\Gamma_{(\mathbb{Z},\mathbb{Z})}$ is properly discontinuous, cocompact, and by isometries, this Cayley graph is not quasi-isometric to $\mathbb{Z}$ or to $\mathbb{R}$; $d_{\mathbb{Z}}(0, n) = 1$ for every n, but $d_{\{\pm 1\}}(0, n) = n$. A bounding constant would have to be larger than every natural number, so there is no quasi-isometry. Here it is evident that the demand that the metric space be proper is in fact necessary. This was not a valid action because the closed unit ball at 0 in $\Gamma_{(\mathbb{Z},\mathbb{Z})}$ is actually the entire graph, which is not compact. The Cayley graph for (G, S) is proper if and only if S is finite. Allowing an infinite set S would produce a different quasi-isometry class.

1.4 Schwarz-Milnor Lemma

The following theorem has been attributed to V. A. Efremovich in 1953, his student Albert Schwarz in 1955, and John Milnor in 1968. In the author's opinion, this could very well be called the fundamental theorem of geometric group theory. The concept of Gromov hyperbolic groups is contingent upon this theorem, for example.

Theorem 1.4.1 (Schwarz-Milnor lemma)**.** *Suppose a finitely generated group G acts geometrically on a proper geodesic space (X, d_X). Then for any choice of $x \in X$ and finite generating set S, the map $f_x \colon G \to X$ defined by $f_x(g) = g.x$ is a quasi-isometry.*

The proof requires a bit more preparation, so it is postponed to the end of the section.

Definition 1.4.2 (Diameter)**.** Let (X, d) be a metric space. The *diameter of X* is $diam(X) = \inf\{R \in \mathbb{R} \mid d(x, y) \leq R \text{ for all } x, y \in \mathbb{R}\}$.

Lemma 1.4.3. *If (X, d_X) is a compact metric space, its diameter is finite.*

Proof. Let $z \in X$. Since X is compact, the open cover $\cup_{r>0} B_r(z)$ has a finite subcover. Since these sets form a chain, we can choose r_0 to be the maximum r in the finite subcover, and deduce that $\overline{B_{r_0}(z)} = X$. Now notice that $R = 2r_0 \geq d_X(x, z) + d_X(z, y) \geq d_X(x, y)$ for every $x, y \in X$. □

Remark. Since the function $d_X(x, y)$ is continuous, it is true that $\overline{B_{\mathrm{diam}(X)}(x)} = X$ for every $x \in X$. Moreover, if $R > \mathrm{diam}(X)$, then $B_R(x) = X$.

Lemma 1.4.4 (Cayley graphs). *If S and S' are finite symmetric generating sets of a group G not containing the identity, $\Gamma_{(G,S)}$ and $\Gamma_{(G,S')}$ are quasi-isometric.*

Proof. It suffices to find $A \geq 1$ so that $d_S(g,h) \leq A \cdot d_{S'}(g,h)$. We take

$$A = \max\{d_S(1_G, s') \mid s' \in S'\}.$$

Note that $d_{S'}(g,h)$ is equivalent to finding the minimal length of a path from the g to h in elements of S'. Since each s' can be replaced by a path in elements of S of length at most A, we obtain $d_S(g,h) \leq Ad_{S'}(g,h)$. Symmetrically, we can find A' so that $d_{S'}(g,h) \leq A'd_S(g,h)$. $\square$

This means that, up to quasi-isometry, the generating set may be chosen as late in the game as desired.

Proof of Theorem 1.4.1. Since the action of G on both $\Gamma_{(G,S)}$ and X is by isometries, it suffices to check that the quasi-isometry inequalities hold when one of g, h in $d(g,h)$ is the identity. We aim to find $A \in [0,1], B \geq 0, C \geq 1$ so that $Ad_S(1,g) - B \leq d_X(x, g.x) \leq Cd_S(1,g)$, and $r > 0$ so that $\overline{B_r(f_x(\Gamma))} = X$.

Let $x \in X$. Since G acts on X cocompactly, the metric space X/G is compact. Using Lemma 1.4.3, choose $r > \operatorname{diam}(X/G)$. Let $K = \overline{B_r(x)}$ and choose S to be the nonidentity elements of G for which gK intersects K (by proper discontinuity, this is finite, so choose $B = |S|$). Note if $y \in gK \cap K$, then $g^{-1}y \in K \cap g^{-1}K$, so S is symmetric. We wish to show that S generates G.

The choice of r is large enough so that the G-translates of K cover X (note the image of K in X/G is the whole space, so its preimage in the quotient map is all of X). Let $g \in G$, and we will show that we can write $g = s_1 \dots s_k$ for some elements in S.

If $g \in S$, then we are done. Otherwise, fix a geodesic path $\gamma\colon [a,b] \to X$ with $\gamma(a) = x$ and $\gamma(b) = g.x$. Let $\varepsilon = \inf\{d(gK, K) \mid g \notin S \cup \{1\}\}$. Since K is closed and any element not in S has the property that K and gK are disjoint, ε is positive. Partition the path γ to obtain $x = x_0, x_1, \dots, x_k = g.x$, with k as small as possible with the property that $x_1 \in K$ and $d(x_i, x_{i+1}) < \varepsilon$ for $i \geq 1$.

Each x_i lies in some translate of K, by say g_i. This allows us to express the element g as $g = (g_1)(g_1^{-1}g_2)(g_2^{-1}g_3)\dots(g_{k-1}^{-1}g)$, so that each term here is an element of $S \cup \{1\}$. Since $x_1 \in K$, $g_1 \in S$. Because $x_i \in g_iK$ and $x_{i+1} \in g_{i+1}K$, we have $d(g_iK, g_{i+1}K) < \varepsilon$ and so $d(K, g_i^{-1}g_{i+1}K) < \varepsilon$. This implies that $g_i^{-1}g_{i+1}$ cannot be in the complement of $S \cup \{1\}$, or else we contradict $\varepsilon = \inf\{d(gK, K) \mid g \notin S \cup \{1\}\}$. Thus, the word length of g is at most k.

Let $C = \max\{d_X(x, s.x) \mid s \in S\}$. Then $d_X(x, s.x) \leq C d_S(1, s) = C$ for each $s \in S$, and so the triangle inequality yields

$$d_X(x, g.x) = d_X(x, s_1 \dots s_k.x) \leq \sum_{i=1}^{k} d_X(s_1 \dots s_{i-1}.x, s_1 \dots s_i.x)$$

$$= \sum_{i=1}^{k} d_X(x, s_i.x) \leq Ck = C d_S(1, g).$$

Now let $g, h \in G$. Then $d_X(f_x(g), f_x(h)) = d_X(g.x, h.x) = d_X(x, g^{-1}h.x)$.

Let $q\colon X \to X/G$ be the quotient map, and $q(x) = x'$. The definition of the quotient map tells us that $G.x = q^{-1}(x')$. We calculate that

$$X = q^{-1}(X/G) = q^{-1}\left(B_C(x')\right) = \bigcup_{y \in q^{-1}(x')} B_C(y) = B_C(q^{-1}(x')) = B_C(G.x),$$

and so f is quasi-onto, as desired. □

Corollary. *If a finitely generated group G acts geometrically on proper geodesic metric spaces (X, d_X) and (Y, d_Y), then X and Y are quasi-isometric.*

Thus, everything about geometric group actions can be determined by the Cayley graph. This incredible correspondence between geometry and group theory is a true gem.

2. Linear Algebra

This section defines the geometric objects of interest in the realms of affine linear algebra, bilinear forms, and reflections.

2.1 Linear Algebra

Definition 2.1.1 (Operations on subsets of a vector space)**.** Let $\Delta = \{v_1, \ldots, v_m\}$ be a finite subset of a vector space V. We call the set of linear combinations of Δ the *span of* Δ, and the set of nonnegative linear combinations of Δ the *cone of* Δ. The *affine subspace determined by* Δ is the set of linear combinations of Δ with coefficient sum 1, and the *convex hull of* Δ is the intersection of the cone and the affine subspace determined by Δ. In symbols,

$$\begin{aligned}
\mathsf{span}(\Delta) &\stackrel{\text{def}}{=} \{\alpha_1 v_1 + \cdots + \alpha_m v_m \mid \alpha_i \in \mathbb{R}\}, \\
\mathsf{cone}(\Delta) &\stackrel{\text{def}}{=} \{\alpha_1 v_1 + \cdots + \alpha_m v_m \mid \alpha_i \geq 0\}, \\
\mathsf{aff}(\Delta) &\stackrel{\text{def}}{=} \left\{\alpha_1 v_1 + \cdots + \alpha_m v_m \,\middle|\, \sum_{i=1}^{n} \alpha_i = 1\right\}, \text{ and finally,} \\
\mathsf{conv}(\Delta) &\stackrel{\text{def}}{=} \left\{\alpha_1 v_1 + \cdots + \alpha_m v_m \,\middle|\, \sum_{i=1}^{n} \alpha_i = 1, \alpha_i \geq 0\right\}.
\end{aligned}$$

If A and B are (not necessarily finite) subsets of V, their *sum* is

$$A + B \stackrel{\text{def}}{=} \{a + b \mid a \in A, b \in B\}.$$

In a classic case of abuse of notation, if $A = \{a\}$ we might write $a + B$ to denote $\{a\} + B$. If $\Phi \subseteq V$, the *negative of* Φ is $-\Phi \stackrel{\text{def}}{=} \{-v \mid v \in \Phi\}$.

2.2 Affine Linear Algebra

Let V be a finite dimensional real vector space. An unspecified use of the word "subspace" denotes a *linear* subspace. Any reference to an affine subspace as defined in Definition 2.1.1 includes the adjective *affine*.

Proposition 2.2.1. *Let A be a nonempty affine subspace of a vector space V (that is, A is the affine span of some subset of V). Then for any $a \in A$, the set $W = -a + A$ is a linear subspace of V.*

Proof. Consider a linear combination $w = \sum_{i=1}^{k} c_i(v_i - a)$ of vectors in $-a + A$. Observe that $w + a = \sum_{i=1}^{k} c_i v_i + \sum_{i=1}^{k}(-c_i)a + 1 \cdot a$, and so we have expressed $w + a$ as an affine combination of elements of A (since the sum of the coefficients is 1). Thus, $w + a \in A$, and $w \in -a + A$. □

Moreover, this property characterizes affine subspaces.

Proposition 2.2.2. *Any translate of a linear subspace is an affine subspace.*

Proof. To see this, take a vector a and a linear subspace W, and note that an affine combination in $a+W$ is a vector of the form $\sum_{i=1}^{k} c_i(a+w_i) = a + (\sum_{i=1}^{k} c_i w_i) \in a+W$, since $\sum_{i=1}^{k} c_i = 1$. $\square$

We call $-a+A$ *the linear subspace directing* A, and define the *dimension* of A to be $\dim(-a+A)$. We'll say $\dim(\varnothing) = -1$ as a convention.

Proposition 2.2.3. *The intersection of affine subspaces is again an affine subspace.*

Proof. Let V be a vector space, and A_1, A_2 affine subspaces. If $A_1 \cap A_2 = \varnothing$, we are done. Otherwise, pick $a \in A_1 \cap A_2$, so that $A_1 = a + W_1$ and $A_2 = a + W_2$ for some linear subspaces W_1, W_2. Then, $A_1 \cap A_2 = (a+W_1) \cap (a+W_2) = a + W_1 \cap W_2$, showing that $A_1 \cap A_2$ is an affine subspace, with dimension at most $\min\{\dim(A_1), \dim(A_2)\}$. $\square$

Theorem 2.2.4 (Proper subspaces are small)**.** *If $A_1, \ldots, A_n$ are proper affine subspaces of V, then $\bigcup_{i=1}^{n} A_i \neq V$.*

Proof. If V is the zero-dimensional vector space, the only proper affine subspace is the empty set, so the result holds. Suppose $\dim(V) \geq 1$.

Let $U_j = \bigcup_{i=1}^{j} A_i$, and let W_i be the linear subspace directing A_i for each i. We may assume that $A_n \not\subseteq U_{n-1}$, so there exists some $v \in A_n \setminus U_{n-1}$. Let $u \in V \setminus W_n$, and consider the affine line $L = \{v + tu \mid t \in k\}$. Since L contains v, $L \not\subseteq A_i$ for $1 \leq i \leq n-1$. We choose $u \notin W_n = -v + A_n$, so $v + u \notin A_n$, and thus L is not contained in A_n.

Now, $\dim(L \cap A_i) < 1$, and hence $|L \cap A_i| \leq 1$. It follows that $|L \cap U_n| \leq n < |\mathbb{R}|$, so $U_n \neq V$, as desired. $\square$

Remark. A more abstract approach to affine spaces is possible. Let V be a vector space, and A a set together with a transitive and free action of V on A, in which the action of $v \in V$ on $a \in A$ is denoted $v + a$. Here, "transitive" and "free" mean that for every $a, a' \in A$, there exists (transitive) a unique (free) vector v so that $v + a = a'$. Then call A together with the group action an affine space. Both of these approaches serve to create an analog of a vector space in which there is no longer a distinguished origin.

This sort of construction can be used in contexts other than vector spaces, and the more general concept is that of a *principal homogeneous space*, or *torsor*.

2.3 Bilinear Forms

This section discusses bilinear forms, working up to three main results. The first result, Sylvester's Law of Inertia, will be a crucial fact in this thesis. The second, equivalence of norms, is used precisely once in an important way. The third, the Cartan-Dieudonné theorem, is not explicitly useful for current goals, but it motivates certain definitions.

It is worth mentioning that diagonalizing a bilinear form is a different process from diagonalizing an endomorphism. Although both of these may be represented as square matrices relative a chosen basis, an endomorphism is a map $V \to V$ and a bilinear form is a map $V \times V \to \mathbb{R}$ for a real vector space V. This distinction turns out to be quite consequential, and the fact that they both look like squares of numbers is not quite as unifying as it may first appear. In the language of tensor analysis, an endomorphism is a $(1, 1)$-tensor, whereas a bilinear form is a $(0, 2)$-tensor.

Definition 2.3.1 (Bilinear Forms)**.** A *bilinear form* on a real vector space V is a function $B\colon V \times V \to \mathbb{R}$ which is linear in each coordinate, so that for each $v \in V$, the functions $B_v(w) \stackrel{\text{def}}{=} B(v,w)$ and $B^v(w) \stackrel{\text{def}}{=} B(w,v)$ are linear transformations $V \to \mathbb{R}$. A bilinear form is said to be *symmetric* if $B_v = B^v$ for every $v \in V$ (equivalently, $B(v,w) = B(w,v)$ for every $v, w \in V$). **We will only consider symmetric bilinear forms.**

Associated to any bilinear form B is a *quadratic form* $q\colon V \to \mathbb{R}$, which is defined by $q(v) \stackrel{\text{def}}{=} B(v,v)$. For this reason, the pair (V, B) may be referred to as a *quadratic space*. Usually, the term "quadratic space" refers to the pair (V, q). However, just as specifying a bilinear form gives rise to a quadratic form, specifying only the quadratic form q on a real vector space uniquely specifies a symmetric bilinear form by the formula $B(v,w) = \frac{1}{2}(q(v+w) - q(v) - q(w))$.

Say a bilinear form B is *positive-semidefinite* if $q(v) = B(v,v) \geq 0$ for all $v \in V$. Say the bilinear form B is *positive-definite* if it is positive semi-definite and $B(v,v) = 0$ if and only if $v = 0$. A symmetric positive-definite bilinear form is also called an *inner product*.

Let B be a (symmetric) bilinear form on a real vector space V, and v a vector in V. Then $v^\perp$, the *(B-) orthogonal complement of* v, is the kernel of B_v, or $v^\perp = \{w \in V \mid B(v,w) = 0\}$, and the *positive half space determined by* v is

$$\mathsf{Half}^+(v) \stackrel{\text{def}}{=} \{w \in V \mid B(v,w) > 0\}.$$

The *radical* of a bilinear form B on a real vector space V, denoted by $V^\perp$, is the set of vectors orthogonal to every other vector; that is, $\{v \in V \mid B(v,w) = 0\ \forall w \in V\}$.

The *general linear group of* V, denoted $\mathsf{GL}(V)$, is the group of invertible linear transformations from V to V, with the group operation of course being composition. The B*-orthogonal group* is the group of automorphisms which preserve the form:

$$O_B(V) \stackrel{\text{def}}{=} \{T \in \mathsf{GL}(V) \mid B(v,w) = B(T(v),T(w)) \text{ for every } v,w \in V\}\,.$$

Definition 2.3.2 (Reflections)**.** If (V,B) is a quadratic space and α is a vector with $B(\alpha,\alpha) \neq 0$, define $s_\alpha \colon V \to V$ by $s_\alpha(\lambda) = \lambda - 2\frac{B(\lambda,\alpha)}{B(\alpha,\alpha)}\alpha$. This has the expected properties in that the reflection in α fixes the orthogonal hyperplane and negates α. Moreover, s_α is a B-orthogonal map.

Definition 2.3.3 (Similarity)**.** When one wishes to express an endomorphism $T \colon V \to V$ of a finite dimensional vector space as a matrix A, one must first select an ordered basis $\mathcal{B} = \{v_1, \dots, v_n\}$. Then we can set entry $A_{i,j} = ((T(v_j))_{\mathcal{B}})_i$. In this setting, any automorphism S of V sets up a correspondence of $\mathcal{B}$ with another ordered basis $\mathcal{C} = \{w_1, \dots, w_n\}$, from which we may observe that $T \circ S(x) = S(y)$ for $T(x) = y$, and hence $S^{-1} \circ T \circ S(x) = y$. So the new matrix is obtained via conjugation by the invertible ‘matrix representing S. Consequently, call A, B *similar* if there is an invertible matrix S so that $A = S^{-1}BS$.

Definition 2.3.4 (Congruence)**.** The situation is different in the setting of bilinear forms $B \colon V \times V \to \mathbb{R}$. Instead, choosing a basis $\mathcal{B} = \{v_1, \dots, v_n\}$ allows one to construct a matrix A so that $A_{i,j} = B(v_i, v_j)$, and consequently $B(u,w) = u^t A w$, thinking of u and w now being written as column vectors in the coordinates of $\mathcal{B}$. Now observe that an automorphism S (with matrix T) changes $(S(u))^t A(S(w)) = u^t(T^t A T)w$, and so a change of basis amounts to multiplying on the left by the transpose of the invertible matrix on the right. Call such pairs of matrices A and $T^t AT$ *congruent*.

Definition 2.3.5 (B_v)**.** Recall the linear functional $B_v \colon V \to \mathbb{R}$ defined by $B_v(w) = B(v, w)$. When $B(v, v) \neq 0$, this linear functional is nonzero, is hence onto, and therefore $\ker(B_v)$ is codimension one, by the rank-nullity theorem.

2.3.1 The Type of a Bilinear Form

Lemma 2.3.6. *If $B \colon V \times V \to \mathbb{R}$ is a nonzero bilinear form, there is a vector $v \in V$ so that $B(v, v) \neq 0$.*

Proof. As B is not zero, there are vectors v and w so that $B(v, w) \neq 0$. We are done if either $B(v, v) \neq 0$ or if $B(w, w) \neq 0$. If not, $B(v + w, v + w) = 2B(v, w) \neq 0$, so $v + w$ qualifies. □

It turns out that any symmetric bilinear form on a real vector space can be diagonalized. This allows one to define the *signature* of a bilinear form in the forthcoming remark.

Theorem 2.3.7 (Sylvester's Law of Inertia)**.** *Let $B \colon V \times V \to \mathbb{R}$ be a symmetric bilinear form. Then there exists a basis $\{v_1, \ldots, v_n\}$ for V so that the matrix of B is given by*

$$\begin{pmatrix} I_p & & \\ & -I_q & \\ & & 0 \end{pmatrix}.$$

Proof. We induct on $n = \dim(V)$. If $n = 1$, B is always diagonal. Suppose the result holds for any n-dimensional vector space, and suppose $\dim(V) = n + 1$.

Note that if B is identically zero, it is already of this form. Suppose B is not identically zero, so we can apply Lemma 2.3.6 to obtain $v \in V$ with $B(v, v) \neq 0$. Now $\ker(B_v) = \{w \in V \mid B(v, w) = 0\}$ is n-dimensional, and there exists a basis $v_1, \dots, v_n$ of $\ker(B_v)$ for which $B|_{v^\perp \times v^\perp}$ is diagonal. Observe that B is diagonal also with respect to $\{v_1, \dots, v_n, v\}$.

Finally, having diagonalized, we can now modify our basis so that the final statement holds. First, reorder this basis so that $B(v_i, v_i)$ is positive for $1 \leq i \leq p$, negative for $p + 1 \leq i \leq p + q$, and 0 thereafter. Then set $w_i = v_i / \sqrt{|B(v_i, v_i)|}$ for $1 \leq i \leq p + q$, and observe that $B(w_i, w_j) = \begin{cases} 1 & \text{if } 1 \leq i = j \leq p \\ -1 & \text{if } p + 1 \leq i = j \leq p + q \\ 0 & \text{otherwise} \end{cases}$. □

Remark. Consequently, define the *signature* or *type* of B to be (p, q, r) where $p + q + r = n$. One may suppress r when it is zero.

2.3.2 Norms

Definition 2.3.8 (Norms)**.** When a bilinear form is positive-definite, it induces a *norm* $\|\cdot\| \colon V \to \mathbb{R}$ defined by $\|v\| = \sqrt{B(v, v)}$. Not all norms can be obtained in this way. Say two norms $\|\cdot\|_1$ and $\|\cdot\|_2$ are *equivalent* if there is a constant C so that $\frac{1}{C}\|v\|_1 \leq \|v\|_2 \leq C\|v\|_1$. Equivalent norms induce the same topology.

Theorem 2.3.9. *Any two norms on a finite-dimensional real vector space V are equivalent.*

Proof. Let $e_1, \dots, e_n$ be a basis for V. It is clear that equivalence of norms is an equivalence relation, so it suffices to show that an arbitrary norm $\|\cdot\|$ is equivalent to the specific norm $\|\cdot\|_1$, which we define by $\|\sum c_i e_i\|_1 = \sum |c_i|$. Now take $A = \max\{\|e_1\|, \dots, \|e_n\|\}$, and note that $\|\sum c_i e_i\| \leq \sum |c_i|\|e_i\| \leq A \cdot \sum |c_i| = A\|\sum c_i e_i\|_1$. Let B be the set of unit vectors of V with respect to $\|\cdot\|_1$; $\|\cdot\|$ is a continuous function on the compact set B, and so it attains a minimum (positive!) value m. So we have $m \leq \|\frac{v}{\|v\|_1}\|$, or in other words, $m\|v\|_1 \leq \|v\|$ for every v. Set $C = \max\{m, A\}$, and we have $\frac{1}{C}\|v\|_1 \leq \|v\| \leq C\|v\|_1$. $\square$

2.3.3 Cartan-Dieudonné Theorem

The Cartan-Dieudonné theorem guarantees an expression of any transformation in $O_B(V)$ in terms of a composition of at most $\dim(V) + 2$ reflections. The main interest is that such an expression of any length exists. In the case that B is positive definite, the length is at most $\dim(V)$.

Theorem 2.3.10 (Cartan-Dieudonné). *Let B be an inner product (positive-definite symmetric bilinear form) on an n-dimensional space V over $\mathbb{R}$. If $f \in O_B(V)$, then there exist B-reflections $s_1, \dots, s_k \in O_B(V)$, with $k \leq n$, so that $s_1 \dots s_k = f$.*

Proof. We proceed by induction on $n = \dim(V)$. The only B-isometries if $n = 1$ are $f(v) = v$ and $f(v) = -v$, which are compositions of length 0 and 1 respectively. Suppose the result is true for each dimension up to n. If f is the identity, it is the composition of 0 reflections. Otherwise, take some $v \in V$ so that $f(v) \neq v$, and let $\vec{m} = f(v) - v$ be the "move" vector of v. Then $s_{\vec{m}}(v) = f(v)$, and so $s_{\vec{m}}(f(v)) = s_{\vec{m}}(s_{\vec{m}}(v)) = v$. Thus $s_{\vec{m}}f$ fixes v, and hence $v^{\perp}$ is an $n-1$-dimensional invariant subspace, to which $s_{\vec{m}}f$ restricts to a $B|_{v^{\perp}}$ isometry. By induction, this is a composition of at most $n-1$ reflections $s_k \dots s_1$, and so $f = s_{\vec{m}}(s_{\vec{m}}f) = s_{\vec{m}}s_k \dots s_1$, as desired. $\square$

3. Hyperbolic Geometry

In the spherical and Euclidean cases, root systems of Coxeter groups are quite well-understood and in fact classified. The Coxeter groups of interest for present purposes relate to hyperbolic geometry.

3.1 Hyperbolic Space

Hyperbolic geometry holds fairly important historical value in mathematics. Euclidean geometry rests upon five axioms, one of which seems misfit: if ℓ is a line and p is a point not on ℓ, there is a unique line ℓ' through p which does not intersect ℓ. For a number of centuries, it was thought that this "parallel postulate" could be proven from the other four axioms. Instead, mathematicians were able to develop spherical geometry in which there is no such line ℓ', and hyperbolic geometry in which there are infinitely many such lines.

Definition 3.1.1 (Lorentz space)**.** Suppose V is an $(n+1)$-dimensional real vector space with a bilinear form B of type $(n, 1)$. Such a space with a bilinear form is called a *Lorentz space*. As before, choose a basis $\{e_1, \ldots, e_{n+1}\}$ for which $B(v, v) = q(v) = v_1^2 + \cdots + v_n^2 - v_{n+1}^2$. Call any such basis a *Lorentz basis*. With respect to this specific basis, call a vector *positive* if $v_{n+1} > 0$ and *negative* if $v_{n+1} < 0$. If one has a basis of space-like vectors, call a vector *positive* if it lies within the cone of the basis vectors, and negative if its negative is positive. Call $Q = q^{-1}(0)$ the *light cone* or the *isotropic cone*, and a

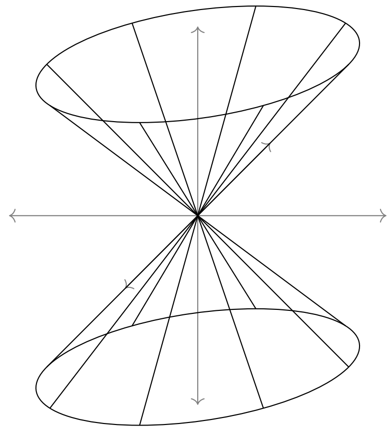

Figure 3.1: The light cone Q of hyperbolic 3-space

vector in Q is called *light-like*. A vector is *space-like* if it lies in $Q^+ = \{v \in V \mid q(v) > 0\}$ and *time-like* if it lies in $Q^- = \{v \in V \mid q(v) < 0\}$. Sometimes time-like vectors are referred to as having imaginary length, as one obtains by trying to define a norm by the usual formula. Of course in the case of a Lorentzian form, $\sqrt{B(v,v)}$ does not define a norm, and one cannot hope to define a topology arising from B. The sets Q^+ and Q^- are also called the *exterior* and *interior* of the light cone, respectively. Call a subspace of V space-like if every nonzero vector is space-like, time-like if it contains a time-like vector, and light-like otherwise.

If V is $(n+1)$-dimensional and B is a bilinear form on V of type $(n,1)$, impose an interesting geometry on a carefully chosen subset of V. By Sylvester's Law of Inertia (Theorem 2.3.7), according to some basis $(e_1, \dots, e_{n+1})$ of V, and writing $v = \sum v_i e_i$,then

$$B(v,w) = v_1 w_1 + \cdots + v_n w_n - v_{n+1} w_{n+1}.$$

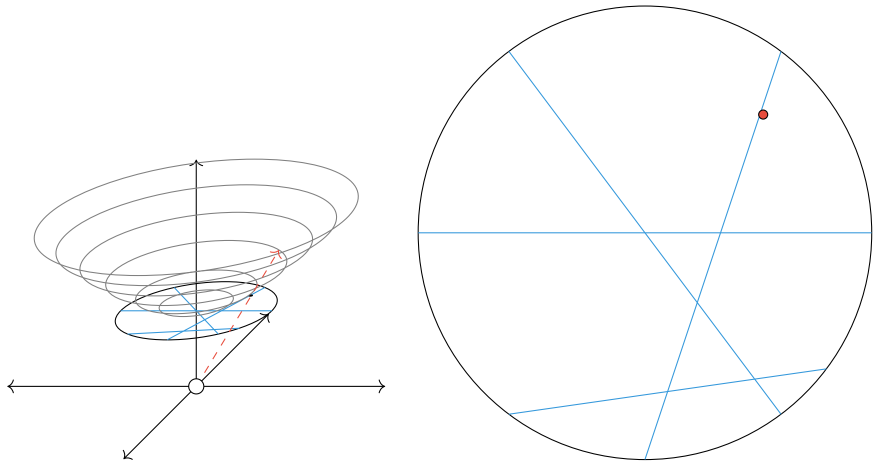

Figure 3.2: On the left in mesh is $\mathbb{H}$, the positive sheet of the -1 hyperboloid. By projecting this sheet onto the disc B_1 as the dashed vector suggests, one obtains the picture on the right. Intersections of hyperplanes with the projective disc are also pictured.

Letting $q(v) = B(v,v)$, consider $\mathbb{H}^n = q^{-1}(\{-1\}) \cap U^{n+1}$, where $U^{n+1} = \{(x_1, \ldots, x_{n+1}) \in V \mid x_{n+1} > 0\}$ is the upper-half space. Since q is a polynomial and hence differentiable map, and -1 is a regular value, this is a submanifold of V. Metrize $\mathbb{H}^n$ via the Riemannian metric associated to the polynomial q; namely, $ds^2 = dx_1^2 + \cdots + dx_n^2 - dx_{n+1}^2$.

Definition 3.1.2 (Projective Ball Model)**.** Let B_1 be the unit n-ball at height 1; that is, $B_1 = \{x \in \mathbb{R}^{n+1} \mid \sum_{i=1}^n x_i^2 < 1, x_{n+1} = 1\}$. Define $p\colon B_1 \to \mathbb{H}^n$ by setting $\{p(x)\} = \mathbb{R}x \cap \mathbb{H}^n$, so $p(x_1, \ldots, x_n, 1) = (x_1, \ldots, x_n, 1)/\sqrt{1 - \sum_{i=1}^n x_i^2}$. Define a metric on B_1 so that $d_p(x,y) = d_{\mathbb{H}}(p(x), p(y))$, and let $\mathbb{H}_p^n = (B_1, d_p)$, the *projective ball model* of hyperbolic space (the "p" subscript is for "projective"). See Figure 3.2.

Definition 3.1.3 (Conformal Ball Model)**.** This model is not used in the thesis, primarily because hyperplanes look nicer in the projective model and these are more important than angles in the context of Coxeter groups. However, the conformal model is more familiar to most readers, so, to contrast with the projective model, the construction of

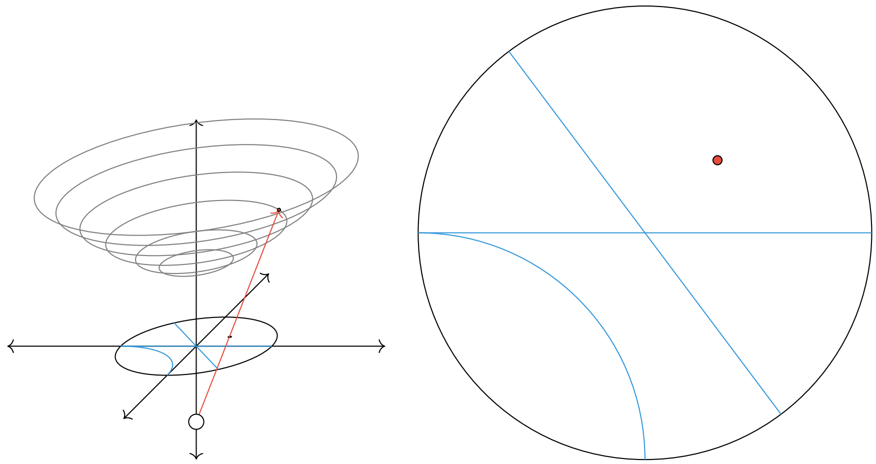

Figure 3.3: The conformal model construction. Here, hyperplanes can be constructed by intersecting a linear hyperplane with $\mathbb{H}^n$, and then projecting down to B_0. Note that the projection point is now $(0, \ldots, 0, -1)$ instead of the origin.

the conformal model is provided . See Figure 3.3. Let B_0 be the unit n-ball at the origin of $\mathbb{R}^{n+1}$; that is, $B_0 = \{x \in \mathbb{R}^{n+1} \mid \sum_{i=1}^{n} x_i^2 < 1, x_{n+1} = 0\}$. Define $c\colon \mathbb{H}^n \to B_0$ by setting $c(x_1, \ldots, x_{n+1}) = (x_1, \ldots, x_n, 0)/(1 + x_{n+1})$, and define $d_c(x, y) = d_{\mathbb{H}}(c^{-1}(x), c^{-1}(y))$. Let $\mathbb{H}^n_c = (B_0, d_c)$, the *conformal ball model* of hyperbolic space (the "c" subscript is for "conformal"). The map c can be viewed as the projection of the hyperboloid onto B_0 from $-e_1 = (0, \ldots, 0, -1)$.

Having constructed this metric space, it is of interest to note that it is in fact a geometry and thus admits the notion of a geometric group action. The further property that there is a unique geodesic between two points will allow for a unique triangle defined by a triple of noncollinear points.

Proposition 3.1.4. *Hyperbolic n-space is a proper geodesic metric space in which there is a unique geodesic between any two points.* □

3.2 Hyperbolic Isometries

Proposition 3.2.1. *For any $v \neq 0$, $v^\perp = \{w \in V \mid B(v, w) = 0\}$ is a codimension 1 linear subspace. If v is space-like, $v^\perp$ is time-like.*

Proof. If $v^\perp$ were not time-like, the restriction of B to $v^\perp$ would be positive-semidefinite, and then V would posses a basis with $B(v_i, v_j) \geq 0$ for each i, j, which is impossible by Theorem 2.3.7. □

Definition 3.2.2 (Lorentz transformations)**.** When V is a Lorentz space with bilinear form B, call an element of $O_B(V)$ a *Lorentz transformation*. It is a straightforward calculation to see that such a transformation must take a Lorentz basis to a Lorentz basis, and that this characterizes such linear maps. Call a Lorentz transformation *positive* if it takes some positive time-like vector to a positive time-like vector (by continuity, this implies that every positive time-like vector is sent to a positive time-like vector).

Theorem 3.2.3. *The group of isometries of $\mathbb{H}^n$ is isomorphic to the group of positive Lorentz transformations.*

3.3 Hyperbolic Reflections

A reflection in hyperbolic space is an isometry arising from a reflection in Lorentz space. Such reflections can be counterintuitive, in that moving the reflection vector closer to the light cone will move the reflecting hyperplane closer to the reflection vector. See Figure 3.4 to make sense of this statement. This vague qualitative behavior of the Lorentzian form is at the heart of the behavior of limit roots.

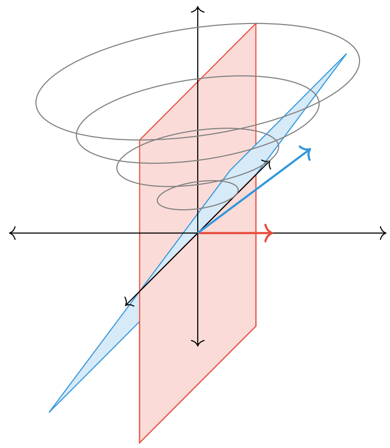

Figure 3.4: The parallel circles form the light cone. Moving from the horizontal to the diagonal vector sweeps out a family of hyperplanes from the vertical to the diagonal plane.

3.4 Hyperbolic Triangles

Trigonometry holds an important position within geometry. Understanding triangles is a prerequisite for understanding most objects one considers in geometry.

Definition 3.4.1. Suppose $x, y, z \in \mathbb{H}^n$ do not lie on a common line. Then the triangle T with vertices x_1, x_2, x_3 is the union $[x_1, x_2] \cup [x_2, x_3] \cup [x_3, x_1]$.

This definition is just the same as the corresponding one for Euclidean triangles. The properties of hyperbolic triangles are wildly different, however. In Euclidean space, given a circle of any radius, it is possible to construct a triangle containing that circle in its interior. This is not the case in hyperbolic space, as a consequence of the following result.

Proposition 3.4.2 (Triangles in $\mathbb{H}^n$ are thin)**.** *Let T be a triangle in $\mathbb{H}^n$ with vertices x_1, x_2, x_3. Then if $y \in [x_1, x_2]$, there is a point $z \in [x_2, x_3] \cup [x_3, x_1]$ so that $d(y, z) \leq \log(1 + \sqrt{2})$.*

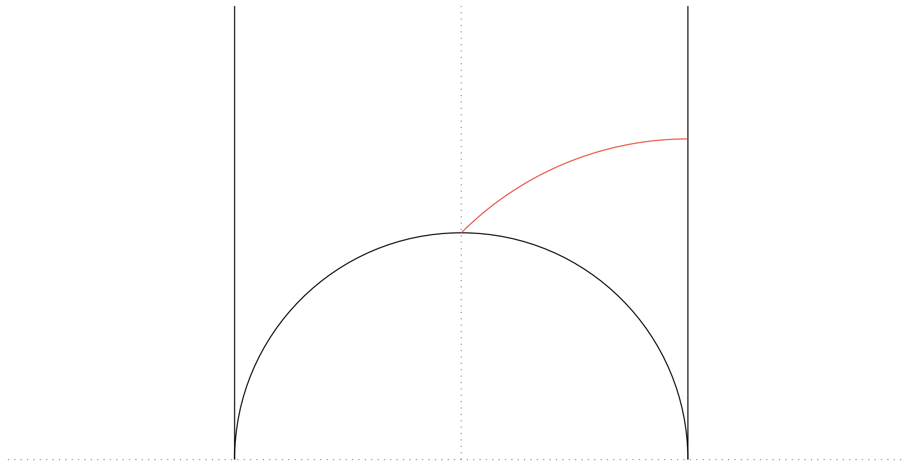

Figure 3.5: This figure illustrates the sketch of the proof of Proposition 3.4.2.

Sketch of proof. This sketch uses many concepts not introduced here. See [Rat06] for a more thorough development. Consider the ideal triangle T in the upper half plane with vertices -1, 1, and ∞. The shortest distance from i to the line $1 + i\mathbb{R}^+$ is $\log(1 + \sqrt{2})$, since the shortest geodesic intersects this line at $1+i\sqrt{2}$. Any other point on the geodesic between -1 and 1 is closer to one of the other sides than i is. Any triangle may be moved to the interior of T by an isometry. See Figure 3.5. $\square$

3.5 Hyperbolic Groups

Hyperbolicity is a central notion in geometric group theory. When a metric space possesses certain properties of hyperbolic geometry, it inherits many desirable characteristics.

Definition 3.5.1. Let (X, d) be a proper geodesic metric space (a geometry). Suppose $x_1, x_2, x_3 \in X$, and let $e_{1,2}$, $e_{2,3}$, and $e_{1,3}$ be geodesics between x_i and x_j. Let $T =$

$e_{1,2} \cup e_{2,3} \cup e_{1,3}$. Then say T is δ-*thin* if whenever $y \in e_{i,j}$, there is some $z \in T \setminus e_{i,j}$ with $d(y,z) < \delta$.

Say that (X,d) is δ-*hyperbolic* if every triangle in X is δ-thin.

Remark. By Proposition 3.4.2, hyperbolic space is $\log(1+\sqrt{2})$-hyperbolic. If a metric space X is δ-hyperbolic, it is also δ'-hyperbolic for any larger value δ'.

Definition 3.5.2. The Gromov product of y and z at x is

$$(y,z)_x \stackrel{\text{def}}{=} \frac{1}{2}(d(x,y) + d(x,z) - d(y,z)).$$

This concept affords an alternative characterization of a δ-hyperbolic space. (X,d) is δ-hyperbolic if $(x_2,x_3)_p \geq \min\{(x_1,x_2)_p,(x_1,x_3)_p\} - \delta$ whenever $p, x_1, x_2, x_3 \in X$.

Definition 3.5.3. Say a group is δ-*hyperbolic* if it acts geometrically on a δ-hyperbolic space, or equivalently, if its Cayley graph is δ-hyperbolic. A group is *hyperbolic* if it is δ-hyperbolic for some δ.

Example 3.5.4. The Cayley graph of a free group is a tree, which is easily seen to be 0-hyperbolic.

Part II

Coxeter Groups

Part II focuses on Coxeter groups, beginning with the theory of finite root systems in Chapter 4. A root system is a subset of a quadratic space satisfying certain axioms. Within any root system is a *simple system*, which, as the name suggests, is a simpler subset which contains all of the information of the root system. The simple system suggests a particularly nice group presentation of a finite reflection group.

The more general case of Coxeter groups is the subject of Chapter 5. The definition is purely algebraic, but Coxeter groups have a representation into the B-orthogonal group for some form B. With this representation, geometric properties of Coxeter groups can be considered. Another byproduct of this representation is a strong combinatorial theory of Coxeter groups, and these combinatorics yield an efficient solution of the word problem for a Coxeter group W.

In Chapter 6, the final chapter of Part II, a classification of finite root systems is provided, and a detailed view of the geometric group action is highlighted.

4. Finite Reflection Groups

The theory of finite reflection groups motivates the definition and study of general Coxeter groups. The geometric notions considered in the theory of finite reflection groups can be used to understand all Coxeter groups. Most of the proofs in this section are general enough such that once the corresponding objects are defined for general Coxeter groups, it is possible to apply these theorems in the more general case.

4.1 From Geometry to Algebra

The concept of a reflection in an inner product space is very geometric, but studying the collection of transformations that can be obtained via a sequence of reflections is algebraic.

Definition 4.1.1 (Reflections)**.** Let V be a finite dimensional vector space over $\mathbb{R}$ together with an inner product $\langle \cdot, \cdot \rangle$; that is, a positive-definite symmetric bilinear form. A linear map that sends some nonzero vector α to its negative and fixes its orthogonal complement $\alpha^{\perp} = \{\lambda \in V \mid \langle \alpha, \lambda \rangle = 0\}$ is called a *reflection of V in α*. In terms of a formula, this is the linear map $s_\alpha \colon V \to V$ such that $s_\alpha(\lambda) = \lambda - 2\frac{\langle \alpha, \lambda \rangle}{\langle \alpha, \alpha \rangle}\alpha$. Note that $s_{k\alpha} = s_\alpha$ for any $k \in \mathbb{R}^*$, so the formula is simplified by choosing $\langle \alpha, \alpha \rangle = 1$.

A finite subgroup of $\mathsf{GL}(V)$ is called a *finite reflection group* if it admits a generating set consisting of reflections. This requires in particular that the order of a product of any two reflections is finite.

Proposition 4.1.2 (Dihedral groups)**.** *If α, β are unit vectors in a vector space V, the subgroup of $\mathsf{GL}(V)$ generated by the reflections $\{s_\alpha, s_\beta\}$ is a (possibly infinite) dihedral group.*

Proof. The transformation $s_\alpha s_\beta$ acts as the identity map on $\mathsf{span}(\alpha, \beta)^\perp$, but is an order-preserving isometry which fixes the origin of $\mathsf{span}(\alpha, \beta)$. Therefore, this is a rotation. Depending on the rotation angle, this product has finite or infinite order. In either case, this generates a dihedral group. □

In other words, the product of any two reflections amounts to a rotation of a plane in a reflection group. To demand this group is finite requires the rotation angle is of finite order. Since specifying a finite reflection group amounts to specifying a generating set, any finite reflection group may be specified by a finite set of vectors.

Lemma 4.1.3 (Reflections are orthogonal)**.** *Reflections are elements of the orthogonal group of a vector space V. Thus a finite reflection group is a subgroup of $O(V)$.*

Proof. First, for formula lovers. Suppose without loss of generality that α is a unit vector.

$$\begin{aligned}\langle s_\alpha(v), s_\alpha(w)\rangle &= \langle v - 2\langle v, \alpha\rangle\alpha, w - 2\langle w, \alpha\rangle\alpha\rangle \\ &= \langle v, w\rangle - \langle 2\langle v, \alpha\rangle\alpha, w\rangle - \langle v, 2\langle w, \alpha\rangle\alpha\rangle + \langle 2\langle v, \alpha\rangle\alpha, 2\langle w, \alpha\rangle\alpha\rangle\end{aligned}$$

Now just observe that

$$\langle 2\langle v, \alpha\rangle\alpha, 2\langle w, \alpha\rangle\alpha\rangle = 4\langle v, \alpha\rangle\langle w, \alpha\rangle = \langle 2\langle v, \alpha\rangle\alpha, w\rangle + \langle v, 2\langle w, \alpha\rangle\alpha\rangle.$$

Alternately, extend $\{\alpha\}$ to an orthonormal basis with respect to the inner product and note that

$$s_\alpha = \begin{pmatrix} -1 & \\ & I_{\dim(V)-1} \end{pmatrix}$$

satisfies $s_\alpha^t = s_\alpha^{-1}$. □

Proposition 4.1.4 (Closure)**.** *Let W be a finite reflection group. Let $w \in W$ and suppose that $s_\alpha \in W$ for some nonzero vector α. Then the reflection in the vector $w(\alpha)$ is also in the group W.*

Proof. To see this, calculate that the w-conjugate $ws_\alpha w^{-1}$ is nothing more than $s_{w\alpha}$. Indeed, $ws_\alpha w^{-1}(w\alpha) = ws_\alpha(\alpha) = w(-\alpha) = -w\alpha$, so it sends $w(\alpha)$ to its negative. Now if $\langle \lambda, w(\alpha) \rangle = 0$, then by Lemma 4.1.3 we also have that $\langle w^{-1}\lambda, w^{-1}w(\alpha) \rangle = 0$. So $s_\alpha(w^{-1}\lambda) = w^{-1}\lambda$, and thus $(ws_\alpha w^{-1})(\lambda) = w(s_\alpha(w^{-1}\lambda)) = w(w^{-1}\lambda) = \lambda$, as required. □

4.2 Root Systems

Following the work in the previous section, particular collections of reflecting vectors are considered. Examined here are three types of sets of reflecting vectors and a way to pass between the three types (the full set of roots, the positive roots, and the simple roots). The way this leads to a group presentation of reflection groups is discussed, perhaps a bit informally. At this point, the general Coxeter group definition is sufficiently motivated.

Definition 4.2.1 (Root systems)**.** A finite subset $\Phi \subseteq V$ is called a *root system* when for each $\alpha \in \Phi$,

(i) $\mathbb{R}\alpha \cap \Phi = \{\alpha, -\alpha\}$,

(ii) $s_\alpha(\Phi) = \Phi$,

(iii) $\mathsf{span}(\Phi) = V$.

The first condition reflects the fact that $s_\alpha = s_{k\alpha}$ for any $k \in \mathbb{R}^*$ (cf. Definition 4.1.1), and the second is motivated by the fact that reflecting roots are closed under the group action (cf. Proposition 4.1.4).

There are two types of subsets of any root system which are of interest. Since the definition requires that Φ counts every reflection exactly twice, it is reasonable to take a filter out a single vector for each reflection. Any vector v in V which is not orthogonal to any vector in Φ (see Theorem 2.2.4) determines a set $\Phi^+ = \Phi \cap \mathsf{Half}^+(v)$ which is called a *positive root system*, and clearly Φ is the disjoint union of Φ^+ and $-\Phi^+$.

A linearly independent subset $\Delta \subseteq \Phi$ is called a *simple root system* if $\Phi^+ \subseteq \mathsf{cone}(\Delta)$ for some choice of Φ^+. The pair (Φ, Δ) is called a *based root system.*

Proposition 4.2.2 (Simple systems exist)**.** *Let Φ be a root system with a specified positive root system Φ^+. Then there is a linearly independent subset $\Delta \subseteq \Phi$ so that each positive root is a nonnegative linear combination of elements of Δ. That is, Δ is a simple system corresponding to Φ^+.*

Proof. Fix a positive root system Φ^+ in a root system Φ. Suppose Δ is a subset Φ^+ such that $\mathsf{cone}(\Delta) \supseteq \Phi^+$, and that no proper subset of Δ has this property.

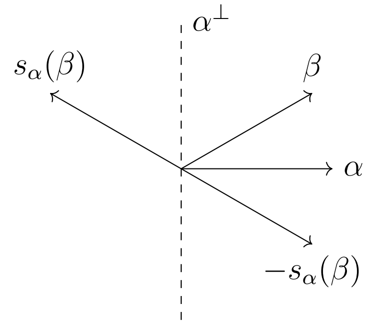

Figure 4.1: For the dimension two case, if α and β are vectors with $\langle \alpha, \beta \rangle > 0$, then neither $s_\alpha(\beta)$ or its negative can be a positive linear combination of α and β.

Then for any choice of distinct $\alpha, \beta \in \Delta$, $\langle \alpha, \beta \rangle \leq 0$. Suppose some pair had $\langle \alpha, \beta \rangle > 0$. Then $s_\alpha(\beta) = \beta - 2\frac{\langle \alpha, \beta \rangle}{\langle \alpha, \alpha \rangle}\alpha$, where the coefficient of α is positive. Suppose first that $s_\alpha(\beta)$ is a positive root. Then there is at least one expression $s_\alpha\beta = \sum_{\gamma \in \Delta} c_\gamma \gamma$ with nonnegative coefficients. If the β coefficient in this particular expression is less than 1, then $(1 - c_\beta)\beta = 2\frac{\langle \alpha, \beta \rangle}{\langle \alpha, \alpha \rangle}\alpha + \sum_{\gamma \neq \beta} c_\gamma \gamma$, and so β can be removed from Δ since it was already expressible as a positive combination of Δ. If the β coefficient is at least 1, then $0 = (c_\beta - 1)\beta + 2\frac{\langle \alpha, \beta \rangle}{\langle \alpha, \alpha \rangle}\alpha + \sum_{\gamma \neq \beta} c_\gamma \gamma$, which means that 0 is a nontrivial positive combination of elements of Δ. This is impossible, because the cone of the positive roots is disjoint from the cone of the negative roots.

However, this condition on the inner product forces Δ to be a basis for $\mathsf{span}(\Phi)$. If we had a dependence relation $\sum a_i \alpha_i = 0$ with the $\alpha_i \in \Delta$ we could move the vectors with negative coefficients to the other side to obtain a vector v which possesses two distinct expressions in terms of positive linear combinations of Δ. Since our bilinear form is positive definite, we have $0 \leq \langle v, v \rangle$, and by the condition on distinct elements of Δ, we have $\langle v, v \rangle \leq 0$. So $v = 0$ and Δ is linearly independent. $\square$

Remark. In the proof, one observes that the angle between any two roots in a simple system is obtuse.

Lemma 4.2.3 (Simple reflections on positive roots)**.** *Let Δ be a simple system contained in a root system $\Phi \subseteq V$. Let $\alpha \in \Delta$. Then $s_\alpha(\alpha)$ is a negative root, but for any other $\beta \in \Phi^+$, $s_\alpha(\beta) \in \Phi^+$.*

Proof. Recall that $s_\alpha(\alpha) = -\alpha$. Now suppose β is not α. Since $\beta = \sum_{\gamma\in\Delta} c_\gamma \gamma$ lies in $\mathsf{cone}(\Delta)$ and is not α, there is some simple root γ_0 so that the corresponding coefficient c_{γ_0} is positive. When we apply s_α to β, it can only change the α coefficient c_α, so $s_\alpha(\beta)$ still has positive c_{γ_0} coefficient, and thus $s_\alpha(\beta) \notin \Phi^-$ (because every root in Φ^- has non-positive coefficients). Thus, β remains positive under the action of s_α. $\square$

Theorem 4.2.4 (Correspondence)**.** *Let W be a finite reflection group associated to a root system Φ. Every positive system Φ^+ contains a unique simple system Δ, and each simple system Δ is contained in a unique positive system Φ^+. Moreover, any two positive (hence simple) systems are conjugate.*

Proof. Let Δ be a simple system in a root system Φ, and note that $\sum_{\alpha\in\Delta} \alpha$ is a vector which is not orthogonal to any vector in Φ. This gives a way to choose a positive root system containing Δ. We constructed a simple system within a positive root system in the previous proposition, and noted it was unique.

For the next part, fix two positive systems Φ_1^+ and Φ_2^+ containing simple systems Δ_1 and Δ_2 respectively. We want to show that there is a $w \in W$ with $w\Phi_1^+ = \Phi_2^+$. First note that by Lemma 4.2.3, if α is a simple root of Φ_1^+, then $s_\alpha(\Phi_1^+)$ sends α to $-\alpha$, but otherwise permutes the elements of Φ_1^+.

Now if $\Phi_1^+ \neq \Phi_2^+$, there must be some $\alpha \in \Delta_1$ so that $\alpha \notin \Phi_2^+$, or in other words, $-\alpha \in \Phi_2^+$. So $s_\alpha(\Phi_1^+)$ intersects Φ_2^+ in one more root. Since there are finitely many roots, applying such reflections finitely many times provides an element of w so that $w(\Phi_1^+) = \Phi_2^+$. □

Remark. The previous proof does not generalize to the infinite Coxeter group case, but the result is more motivational than useful for present purposes.

Definition 4.2.5. Fix a simple system Δ within a root system $\Phi \subseteq V$. Associated to this basis is a linear functional $\varphi\colon V \to \mathbb{R}$ which just evaluates each basis vector to 1. Thus, the fiber of φ over 1 is the affine span $\mathsf{aff}(\Delta)$. If β is a root (or really any vector in V), define the *height* of β to be $\varphi(\beta)$, which is just $\sum_{\alpha\in\Delta} c_\alpha$, where $\beta = \sum_{\alpha\in\Delta} c_\alpha \alpha$.

Theorem 4.2.6 (Generated by simple system)**.** *Let W be a finite reflection group associated to a root system Φ. Suppose Δ is a simple system contained in Φ. Then, W is generated by the reflections S arising from roots in Δ.*

Proof. It is of course clear that W is generated by the reflections in the positive root system Φ^+ corresponding to Δ. Let W_S denote the subgroup of W generated by S. Choose some $\beta \in \Phi^+$, and consider the positive roots Φ_S^+ in the W_S-orbit of β. Among the roots in Φ_S^+, each of which lie in $\mathsf{cone}(\Delta)$, choose one (say γ) with minimal height as defined above. We'll show γ must be simple. We have $0 < \langle \gamma, \gamma \rangle = \sum_{\alpha\in\Delta} c_\alpha \langle \gamma, \alpha \rangle$, so there must be at least one $\alpha \in \Delta$ with $\langle \gamma, \alpha \rangle > 0$. If $\alpha \neq \gamma$, then by Lemma 4.2.3, $s_\alpha(\gamma)$ is positive, but we know the height of $s_\alpha(\gamma)$ is strictly less than the height of γ, a contradiction. So γ is in fact simple.

This says that if $\beta \in \Phi^+$, there is some $w \in W_S$ so that $w(\beta)$ is simple. This implies that the union of W_S orbits of Δ contains any $\beta \in \Phi^+$ as some $w^{-1}\alpha$. So if β is a negative root, then there is some $w \in W_S$ and $\alpha \in \Delta$ with $-\beta = w(\alpha)$. But then $\beta = ws_\alpha(\alpha)$, and so $W_S(\Delta) \supseteq \Phi^+$, and therefore $W_S = W$. $\square$

4.3 A Presentation

This section introduces a simple presentation for a given finite reflection group. To prove rigorously this is indeed a presentation requires too much development which would need to be repeated in the next section, so an argument relying on geometric intuition is given. The next chapter obtains a complete (but less memorable) proof as a byproduct of more general work in Chapter 5.

Proposition 4.3.1. *Let W be a finite reflection group on a vector space V with a simple system $\Delta = \{\alpha_1, \ldots, \alpha_n\}$, and let $S = \{s_1, \ldots, s_n\}$. Any pair α_i, α_j generates a finite dihedral group of order say $2m_{i,j}$. Then $W \cong \langle S \mid (s_i s_j)^{m_{i,j}} \rangle$.*

Sketch of proof. To see this define the *fundamental chamber* C to be the intersection of the positive half spaces determined by Δ. So $C = \bigcap_{\alpha \in \Delta} \mathsf{Half}^+(\alpha)$. This is an open simplicial cone. Call the hyperplanes $\alpha^\perp$ for $\alpha \in \Delta$ the *walls* of C.

Place a sphere S of radius $\frac{1}{2}$ in C so that it intersects each wall of C in exactly one point. This can be done because $\mathbb{R}^n$ is complete and because the chamber is a simplicial cone. Further, draw edges from the center of the sphere to each wall, labeled with the simple root corresponding to the wall of C. By reflecting the ball around the hyperplanes, one obtains the 1-skeleton of a polytope, the W-permutahedron, all of whose edges are length 1. Any subset of S corresponds to a unique face of the polytope containing the center of the ball. Taking a two-element subset determines a 2 dimensional face, and reading around the edge labels gives a word $(s_i s_j)^{m_{ij}}$

It is clear that all relations in the presentation hold. To complete the proof requires showing that any relation in W is a consequence of the ones given. This we will omit, but a full proof can be found in Section 1.9 of [Hum92].

$\square$

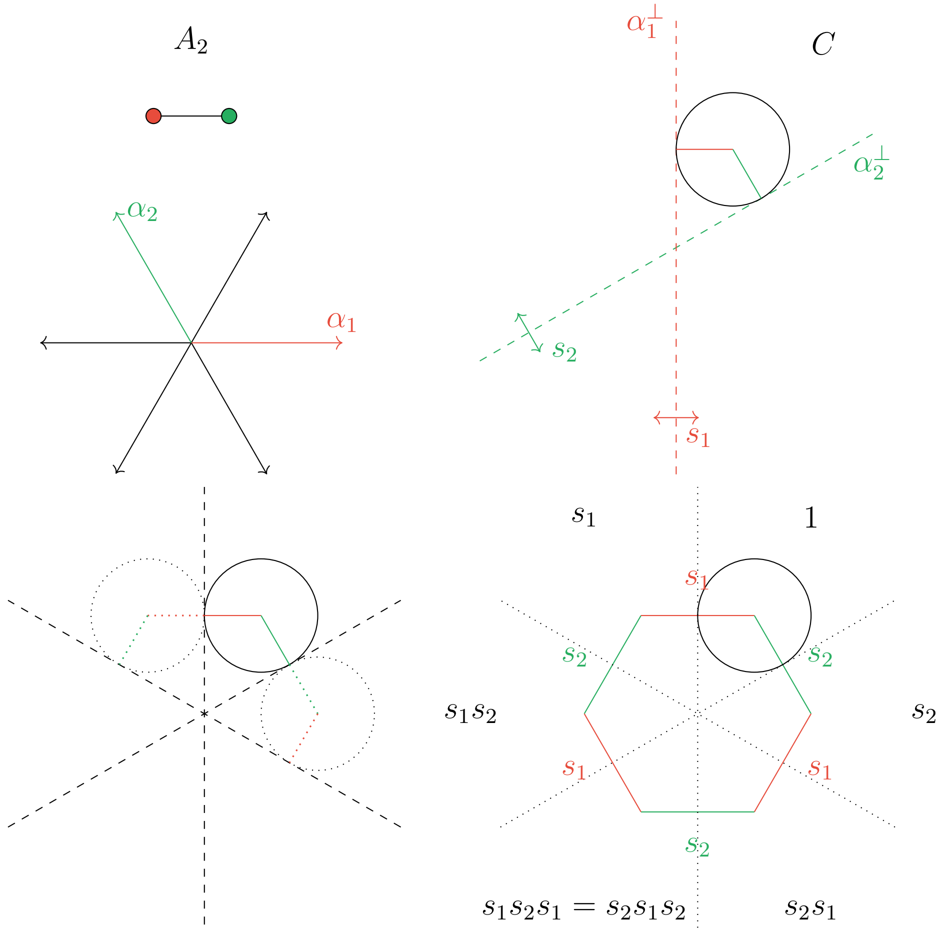

Figure 4.2: Beginning with the A_2 Coxeter diagram, one can construct the A_2 root system. By intersecting the positive half-spaces associated with the simple system, one obtains the fundamental chamber C. One can place a sphere of radius 1/2 within the closure D of C so that it is tangent to each hyperplane, and draw edges from the center of this sphere to the adjacent hyperplane, coloring each edge. The orbit of this configuration under the reflection group yields the Cayley graph, which is the 1-skeleton of the W-permutahedron.

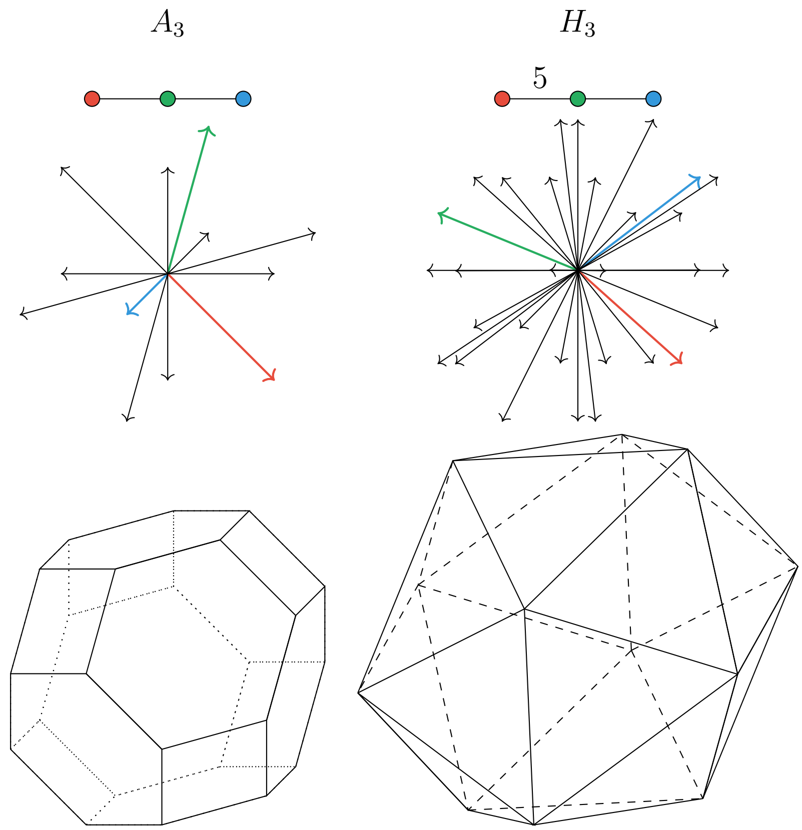

Figure 4.3: Pictured on the left is the A_3 root system followed by the A_3-permutahedron. On the right, the H_3 root system and the icosahedron, which is a regular polytope with symmetry group isomorphic to the H_3 reflection group.

5. Coxeter Groups

Every group generated by reflections which acts geometrically on an n-sphere has a presentation of a particularly simple type. By expanding the type of geometries allowed, every presentation of this type corresponds to a group generated by reflections acting geometrically on a geometry.

5.1 Defining Coxeter Groups

There are three ways to specify a Coxeter system. The first method is a group presentation. Associated to the group presentation is a matrix and a graph.

Definition 5.1.1 (Coxeter system)**.** A group W together with a generating set $S = \{s_1, \ldots, s_n\}$ is a (finitely generated) *Coxeter system* if

$$W = \langle s_1, \ldots, s_n \mid (s_i s_j)^{m_{i,j}} \rangle,$$

(i) $m_{i,j} \in \mathbb{N} \cup \infty$, where $m_{i,j} = \infty$ means that there is no relation on the product

(ii) $m_{i,i} = 1$

(iii) for $i \neq j$, $m_{i,j} = m_{j,i} \geq 2$.

We call $n = |S|$ the *rank* of the Coxeter system.

To emphasize how this generalizes finite reflection groups, this definition allows for groups with similar presentations but don't necessarily come from the same geometric process. Geometric information can be recovered, however.

This is a case where the presentation gives exactly what one would expect. In general, a group presentation can be misleading, but this situation has the desirable properties that $s_i = s_j$ implies $i = j$, and that the order of $s_i s_j$ is precisely $m_{i,j}$ rather than merely a divisor of $m_{i,j}$. One can see that the function $S \to \mathbb{Z}/2\mathbb{Z}$ which sends each generator in S to the nonidentity element in $\mathbb{Z}/2\mathbb{Z}$ has each relation in the kernel and hence extends to a homomorphism sgn onto $\mathbb{Z}/2\mathbb{Z}$, allowing us to conclude that the generators are order 2. The other properties are a consequence of the representation of W which is constructed in Section 5.2.

Two handy orthographic devices are now introduced, which can encode the information of a Coxeter group: one is a matrix and the other a graph. At first glance, these just provide a convenient notation schematic, but the graph-theoretic and linear-algebraic properties contain important mathematical information.

Definition 5.1.2 (Coxeter diagram)**.** For a Coxeter system (W, S), define the *Coxeter diagram* to be a graph on vertices $s_1, \ldots, s_n$, with an edge labeled $m_{i,j}$ between s_i and s_j if and only if $m_{i,j} \geq 3$. If an edge is unlabeled, implicitly it is labeled with a 3. It is clear that these two objects have all of the information of the Coxeter system, so one can in fact define a Coxeter system by either of these means.

Definition 5.1.3 (Coxeter matrix)**.** Given a Coxeter system (W, S), the *Coxeter matrix* is the $n \times n$ matrix M with $M_{ij} = m_{i,j}$. Define also the *Schläfli matrix* C with entry

$C_{i,j} = \cos(\pi - \pi/m_{i,j})$. A symmetric matrix corresponds quite directly to a symmetric bilinear form, as described in Section 2.3. Because the diagram and matrix encode identical information, adjectives typically reserved for bilinear forms are used to describe the graph and vice versa. For example, one might say that a graph is positive definite.

Example 5.1.4 (B_3/C_3)**.** To illustrate this correspondence, the following three datum contain the same information.

$$\langle s_1, s_2, s_3 \mid s_1^2 = s_2^2 = s_3^2 = (s_1s_2)^4 = (s_2s_1)^4 = (s_2s_3)^3 = (s_3s_2)^3 = (s_1s_3)^2 = (s_3s_1)^2\rangle$$

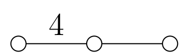

$$\begin{bmatrix} 1 & 4 & 2 \\ 4 & 1 & 3 \\ 2 & 3 & 1 \end{bmatrix} \leftrightarrow \begin{bmatrix} 1 & -\cos(\pi/4) & 0 \\ -\cos(\pi/4) & 1 & -1/2 \\ 0 & -1/2 & 1 \end{bmatrix}$$

Example 5.1.5. It is important to note that the definition is of a Coxeter *system* rather than a Coxeter *group*. The pair (W, S) is studied rather than the group W abstractly. The reason for this is that the a group W may admit multiple generating sets which differ in key properties. For example, the Coxeter groups defined by the diagrams in Figure 5.1 are isomorphic, but have very different Coxeter-theoretic properties. Thus, the appropriate category to study here is that of *marked groups*. A *marked group* is a pair consisting of a group G and a generating set S for G.

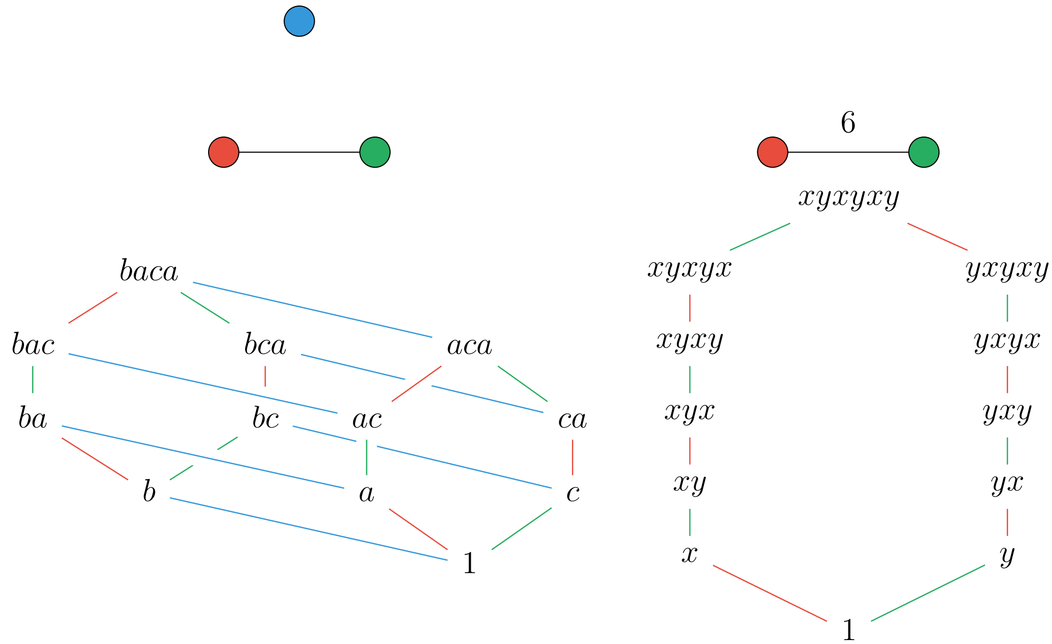

Figure 5.1: Two distinct Coxeter systems which are isomorphic as groups. Also evident in this figure is the property that a disconnected diagram yields a direct product of groups.

For more results on nonisomorphic diagrams giving the same group, see [BMMN02]. The theory of *diagram-twisting* developed in this paper allows one to generate examples of this phenomena.

5.2 The Geometric Representation

A representation of a group is a homomorphism into $\mathsf{GL}(V)$ for some vector space V. There is a uniform way of constructing such a representation for any Coxeter system. In the case of a finite Coxeter group, the geometric interpretation satisfies the definition of finite reflection groups.

Definition 5.2.1 (Bilinear form)**.** Let $(W, \{s_1, \ldots, s_n\})$ be a Coxeter system. Take V a real vector space with basis $\{\alpha_1, \ldots, \alpha_n\}$. Impose a geometry on V by constructing a bilinear form $B(\alpha_i, \alpha_j) = \cos(\pi - \frac{\pi}{m_{i,j}}) = -\cos(\pi/m_{i,j})$, interpreting $-\cos(\pi/\infty) = -1$. Later, any value less than -1 will be allowed in the latter case. Observe that, in particular, $B(\alpha_i, \alpha_i) = 1$ for each of the basis vectors.

Any vector v defines a linear functional $B_v \colon V \to \mathbb{R}$ so that $B_v(w) = B(v, w)$. In fact, the function $v \mapsto B_v$ is a linear map to the dual space $V \to V^*$. The *radical* of B is the subspace $B^\perp = \{v \in V \mid B_v(w) = 0 \ \forall w \in V\} = \ker(v \mapsto B_v)$. When this is trivial, say B is nondegenerate.

When v is not in the radical of B (that is, there is some w for which $B(v, w) \neq 0$), B_v is not the zero linear functional, so it is a surjection onto $\mathbb{R}$ and by the rank-nullity theorem, its kernel is a codimension one subspace of V. Accordingly, define the *hyperplane determined by* v to be $H_v = \ker(B_v)$.

Note that a reflection in this hyperplane is not always possible, because v may lie in H_v when $B(v, v) = 0$.

Definition 5.2.2 (Orthogonal group)**.** Define the *B-orthogonal group*, denoted $O_B(V)$, to be the set of linear transformations which preserve the bilinear form B; that is, $O_B(V) = \{T \in \mathsf{GL}(V) \mid B(T(v), T(w)) = B(v, w) \ \forall v, w \in V\}$. Suppose $B(\alpha, \alpha) \neq 0$, and define a *reflection in α with respect to B* to be the linear map

$$\sigma_\alpha(v) = v - 2\frac{B(\alpha, v)}{B(\alpha, \alpha)}\alpha,$$

and observe that σ_α fixes $v \in H_\alpha = \ker(B_\alpha)$ and $\sigma_\alpha(\alpha) = -\alpha$. The following calculation shows that $\sigma_\alpha \in O_B(V)$ for every α:

$$
\begin{aligned}
B(\sigma_\alpha(u), \sigma_\alpha(w)) =& B(u - 2\frac{B(\alpha, u)}{B(\alpha, \alpha)}\alpha, w - 2\frac{B(\alpha, w)}{B(\alpha, \alpha)}\alpha) \\
=& B(u, w) - 2\left(B(\frac{B(\alpha, u)}{B(\alpha, \alpha)}\alpha, w) + B(u, \frac{B(\alpha, w)}{B(\alpha, \alpha)}\alpha)\right) \\
&+ 4B\left(\frac{B(\alpha, u)}{B(\alpha, \alpha)}\alpha, \frac{B(\alpha, w)}{B(\alpha, \alpha)}\alpha\right) \\
=& B(u, w).
\end{aligned}
$$

Thus the group generated by $\{\sigma_\alpha \mid \alpha \in \Phi\}$ is actually a subgroup of $O_B(V)$.

The properties of this bilinear form unsurprisingly have geometric consequences. When it is positive definite, the action is in some sense (the sense of Section 1.1) best viewed on a sphere, say the unit ball $\{v \in V \mid B(v, v) = 1\}$ of V. If B is merely positive semidefinite, the group has a natural action on an $n-k$-dimensional Euclidean affine space, where k denotes the dimension of the radical of B. One can view the action (which is not always geometric) on a hyperbolic space if the form has type $(n-1, 1)$.

As it turns out, requiring $-\cos(\pi/\infty) = -1$ is a bit more restrictive than desired. When subgroups are studied in Section 7.1, any value satisfying $-\cos(\pi/\infty) \leq -1$ is allowed.

Definition 5.2.3 (More Geometric Objects)**.** A number of subsets of our vector space V constructed in the previous section are of interest. First of all, since Δ is a basis for the vector space, one can define a linear functional $\phi\colon V \to \mathbb{R}$ by mapping each basis vector to 1, so that $\phi(v)$ is the sum of the coordinates of v in Δ. The set $V_0 = \phi^{-1}(0)$ is a hyperplane in V, and $V_1 = \phi^{-1}(1)$ is an affine hyperplane in V (it is, in fact, $\mathsf{aff}(\Delta)$). Let

$\hat{} : V \setminus V_0 \to V_1$ by $\hat{v} = \frac{v}{\phi(v)}$. Note this map cannot extend to V_0, but it maps any other vector onto V_1 by the intersection of the line spanned by α with the affine hyperplane V_1.

One may want to consider the projectivization $\hat{\Phi}$ of Φ, but one needs to ensure $\Phi \cap V_0 = \varnothing$. Since $\Phi \subseteq \mathsf{cone}(\Delta) \cup - \mathsf{cone}(\Delta)$, it is enough to check that $\mathsf{conv}(\Delta) \cap V_0 = \varnothing$. This is clear: $\phi(\delta) = 1$ for any $\delta \in \mathsf{conv}(\Delta)$.

Definition 5.2.4 (Length function)**.** Suppose G is a group which is generated by a set S, and let $w \in G$. Then, the length of w is $\ell(w) = \min\{k \geq 0 \mid s_1 \dots s_k = w \text{ for some } s_i \in S\}$. Obviously, if $s_1 \dots s_k$ is an *expression* for w, then $\ell(w) \leq k$. This fact is used below

Proposition 5.2.5 (Technical lemma for length function)**.** *Take $w, w' \in W$. Then $\ell(w) = \ell(w^{-1})$, and $\ell(w) - \ell(w') \leq \ell(ww') \leq \ell(w) + \ell(w')$.*

Proof. Indeed, let $s_1 \dots s_k$ be a reduced expression for w. Then $s_k \dots s_1$ is an expression for w^{-1}. So $\ell(w^{-1}) \leq \ell(w)$. Applying this argument to w^{-1} yields $\ell(w) = \ell((w^{-1})^{-1}) \leq \ell(w^{-1})$, so the result follows.

Now let $s_1, \dots s_k$ be a reduced expression for w and $s_{k+1} \dots s_{k+q}$ a reduced expression for w'. Then $\ell(ww') \leq k + q = \ell(w) + \ell(w')$.

Finally, $\ell(w) = \ell(ww'(w')^{-1}) \leq \ell(ww') + \ell((w')^{-1}) = \ell(ww') + \ell(w')$, and so $\ell(w) - \ell(w') \leq \ell(ww')$. We used both of the previous parts in this calculation. $\square$

Corollary. *If $w \in W$ and $s \in S$, then $\ell(ws) = \ell(w) \pm 1$.*

Proof. By the second part of Proposition 5.2.5, and taking $w' = s$, it suffices to show that $\ell(ws) \neq \ell(w)$. But this is afforded to us by the sgn homomorphism from W to $\mathbb{Z}/2\mathbb{Z}$. This implies that any two expressions for the same group element must have the same parity, and so any expression for w and ws must have different parities. $\square$

In the previous corollary, whether $\ell(ws)$ is $\ell(w)+1$ or $\ell(w)-1$ has profound geometric and combinatorial consequences. Given an element $w \in W$, the set of $s \in S$ for which the length of ws is greater than the length of w is an important

The length function has a geometric rendition. This is introduced as a new function, but with some work it is shown that the new geometrically defined function coincides with the combinatorial length function. Having both of these descriptions allows for smooth passage between the geometric and algebraic theories.

Definition 5.2.6 (Geometric length)**.** Let (W, S) be a Coxeter system, together with quadratic space (V, B) and based root system (Φ, Δ). For $w \in W$, let $\mathsf{inv}(w) = \{\alpha \in \Phi^+ \mid w(\alpha) \in \Phi^-\}$, and let $n(w) = |\,\mathsf{inv}(w)|$. Recall that corresponding to each element s_i in S is a basis vector α_i in V. The set $\mathsf{inv}(w)$ can be computed as follows: take a reduced word $s_1 \ldots s_k = w$, and $\mathsf{inv}(w) = \{\alpha_k, s_k(\alpha_{k-1}), \ldots, s_k \ldots s_2(\alpha_1)\}$. Consequently, $n(w) = \ell(w)$.

Proposition 5.2.7 (Simple reflection on positive roots)**.** *Let $(W, S, V, B, \Phi, \Delta)$ be as above. Then if $s \in S$ and α a positive root, then $s(\alpha)$ is a negative root if and only if $\alpha = \alpha_s$ (the simple root corresponding to s). So s permutes the positive roots but sends α_s to its negative.*

Proof. Recall that $s(\alpha_s) = -\alpha_s$. Suppose α is not α_s; since $\alpha = \sum_{t\in S} c_t\alpha_t$ lies in $\mathsf{cone}(\Delta)$ and is not α_s, there is some s' so that the corresponding coefficient $c_{s'}$ is positive. When we apply s to α, it can only change the α_s coefficient c_s, so $s(\alpha)$ still has positive $c_{s'}$ coefficient, and thus $s(\alpha) \notin \Phi^-$. Thus, α remains positive under the action of s.

Note that this proof is almost identical to that of Lemma 4.2.3 $\square$

We can extend the corollary of Proposition 5.2.5 as follows.

Proposition 5.2.8. *Let $w \in W$ and $s \in S$. Then $\ell(ws) = \ell(w) + 1$ if and only if $w(\alpha_s)$ is a positive root.*

Theorem 5.2.9. *Let (W, S) be a Coxeter system and V a real vector space with basis $\{\alpha_1, \dots, \alpha_n\}$ in bijection with S. Let B be the corresponding bilinear form. Let $\sigma \colon W \to O_B(V)$ by setting $\sigma(s_i) = \sigma_{\alpha_i}$, the B-reflection in α_i. Then the order of $\sigma(s_i s_j)$ in $O_B(V)$ is $m_{i,j}$. Thus σ is a representation of W in $O_B(V)$. Moreover, σ is faithful.*

Proof. Let $w \in \ker(\sigma)$. Then $w(\alpha_s) = \alpha_s$ for each $s \in S$, which by Prop 5.2.8 implies that $\ell(ws) = \ell(w) + 1$ for each $s \in S$. Suppose $\ell(w) = k \geq 1$. Take a reduced expression $w = s_1 \dots s_k$, then note$ws_k = s_1 \dots s_{k-1}$, a contradiction. So $\ell(w) = 0$, which implies that $w = 1$, and so σ is faithful. $\square$

5.3 Parabolic Subgroups

Let (W, S) be a Coxeter system, with corresponding diagram Γ. Two groups may be defined by a subset $I \subseteq S$. First, consider the subgroup of W which is generated by I, which is called a *standard parabolic subgroup* and denote it by W_I. Second, this corresponds to a subgraph Γ_I of Γ by including only the vertices in I. This diagram abstractly defines a Coxeter system (W_I', I). Then W_I and W_I' are isomorphic, and hence a parabolic subgroup is a Coxeter system.

Recall the faithful representation $\sigma \colon W \to \mathsf{GL}(V)$ from Section 5.2. Recall V was defined from the basis $\{\alpha_s\}_{s \in S}$. Consider the subspace V_I spanned by $\{\alpha_s\}_{s \in I}$, and restrict the representation σ to the parabolic subgroup W_I. On the other hand, since

(W_I', I) is a Coxeter system, there is a corresponding faithful geometric representation $\sigma'\colon W_I' \to \mathsf{GL}(V_I')$, and a canonical isomorphism $V_I \to V_I'$.

So the following commutative square

$$\begin{tikzcd} W_I \arrow[r] \arrow[d] & \mathsf{GL}(V_I) \arrow[d] \\ W_I' \arrow[r] & \mathsf{GL}(V_I') \end{tikzcd}$$

implies that $W_I \to W_I'$ is an isomorphism. Hence, for a subset I of S, the pair (W_I, I) is itself a Coxeter system.

Moreover, the map $I \mapsto W_I$ is an order isomorphism, which implies that S is in some sense a minimal generating set. That is, no element in S can be expressed as a product of other elements of S.

More generally, a *parabolic subgroup* is any subgroup of W which is conjugate to a standard parabolic subgroup. Similarly (wW_Iw^{-1}, wIw^{-1}) is a Coxeter system.

If Γ is a disconnected graph, partition its vertex set into disjoint components I and J. Then note that if $i \in I$ and $j \in J$, then $m_{i,j} = 2$ and so any two elements in W_I and W_J commute. Thus W_I and W_J centralize one another, so W_I and W_J are both normal subgroups of W. Since their product contains the generating set S and they intersect trivially, $W \cong W_I \times W_J$.

Accordingly, call a Coxeter system (W, S) *irreducible* if it corresponds to a connected graph. Note that Example 5.1.5 shows that irreducibility is a property of a Coxeter system and not a property of a Coxeter group.

5.4 Combinatorics

There is a fascinating combinatorial theory of Coxeter Groups. The necessary facts are outlined here, but a book by Björner and Brenti illuminates this theory in full.

There is a natural grading of elements of a Coxeter group by word length. There are a few partial orders that can be defined for Coxeter groups. The weak order is most important for this thesis. With respect to the weak order, a finite Coxeter group possesses a unique element of maximal length. In the infinite case, the lack of a longest element helps provide a way to understand the limit roots, the main object of this thesis.

The symmetric group is intimately connected to combinatorial objects such as parking functions and noncrossing partitions. By understanding the connections of these combinatorial objects to the symmetric group as a finite Coxeter group, it is possible to use other Coxeter groups to generalize these combinatorial objects.

Definition 5.4.1 (Words in Coxeter groups)**.** Given a set S, let S^* denote the free monoid generated by S (relevant terminology: S^* consists of *words* in the *alphabet* S). Embed $S^n \hookrightarrow S^*$ so that $S^* \cong \bigcup_{n\in\mathbb{N}} S^n$, which provides a notion of word length in S^*. Now if (W, S) is a Coxeter system, then since each element of S has order 2 in W and is hence self-inverse, every element of W is a word in the alphabet S. In other words, there is a surjection $\varepsilon\colon S^* \to W$ given by $\varepsilon(s_1 s_2 \ldots s_r) = s_1 s_2 \ldots s_r$. Define $\ell\colon W \to \mathbb{N}$ so that $\ell(w)$ is the minimal length of a word in S^* having image w under ε. Since ε is surjective, this is a nonempty subset of $\mathbb{N}$ and hence possesses a least element. In general, there will be many different words of length $\ell(w)$ in S^* that map to w. Any such word is called *reduced*. Beware that elements of S^* are often conflated with their image under ε.

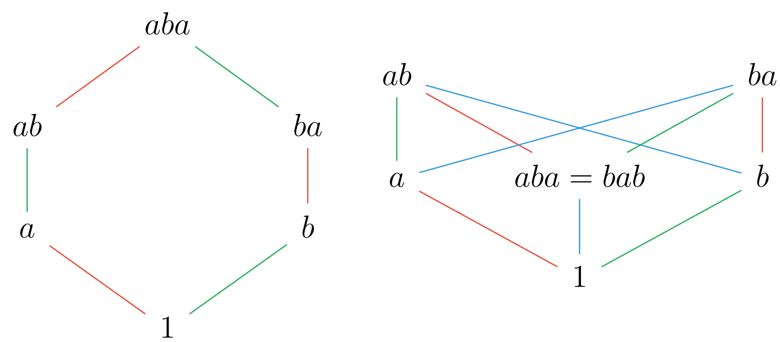

Figure 5.2: Hasse diagrams for A_2 under the right weak order (left) and the absolute order (right).

For example, in the A_2 Coxeter group generated by a, b, $\varepsilon(ba) = \varepsilon(aaba) = \varepsilon(abab) = \varepsilon((ab)^5) = \varepsilon(babb) = \varepsilon(ba) = ba$.

Let $R = \{wsw^{-1} \mid w \in W, s \in S\}$. Call R the set of *reflections* in W. There is an analogous map $\varepsilon_R \colon R^* \to W$ satisfying $\varepsilon(r_1 \dots r_k) = r_1 \dots r_k$. Define $\ell_R \colon W \to \mathbb{N}$ to be the minimal length word in R^* having image w under ε_R. Any word achieving this length is called *R-reduced*. The relationship between the sets S and R is similar to the relationship between simple and positive root systems, respectively.

There is a theory of *dual Coxeter systems*, where one considers pairs (W, R) which satisfy the corresponding properties of a Coxeter group and its full set of reflections. Define the *absolute order* on W, given by $u \leq v$ if and only if $\ell_R(u) \leq \ell_R(v)$ and a (hence any) reduced R-expression for u is a prefix for some reduced R-expression for v. Consider the Hasse diagrams in Figure 5.2, where W is the Coxeter group of type A_2. Then the length of the longest S-expression is equal to $|R|$, and the length of the longest R-expression is equal to $|S|$. Further numerology and connections of this sort motivates the terminology of *dual*.

Definition 5.4.2 (Lattices)**.** Let $(P, \leq)$ be a partially ordered set (or poset). If $S \subseteq P$, say x is an upper bound for S if $s \leq x$ for every $s \in S$. If x also has the property that if y is another upper bound for S, then $x \leq y$, say x is a *least upper bound*, or *join*, denoted $\bigvee S$. Dually, define the *greatest lower bound*, or *meet* of a set S, denoted $\bigwedge S$. It is clear that meets and joins are unique in case of existence, but in a general poset, these objects need not exist. If every nonempty finite subset S of P has a join, call $(P, \leq)$ a *join semi-lattice*, and similarly for *meet semi-lattice*. Call$(P, \leq)$ a *lattice* if it is both. The adjective *complete* means the corresponding operation can be done for arbitrary nonempty subsets.

The collection of subsets of a given set forms a lattice under inclusion, which is a key example to keep in mind when trying to recall the terminology and notation. In this case, the meet of a collection of subsets is their intersection $\cap$: the set of elements in which they "meet". Likewise, one can "join" a collection of subsets together by taking their union $\cup$.

Definition 5.4.3 (Weak order)**.** Let (W, S) be a Coxeter system, and let $u, w \in W$. Then say $u \leq_R w$ if and only if some (and in fact any) reduced expression for u is a prefix for a reduced expression for w. This is equivalent to $\ell(u) + \ell(u^{-1}w) = \ell(w)$. This just means if $s_1 \dots s_k$ is a reduced expression for u, it can be extended to a *reduced* expression $s_1 \dots s_k \dots s_n$ for w. Similarly define the left weak order by setting $u \leq_L w$ if and only if u is a *suffix* for a reduced expression for w.

Theorem 5.4.4. $(W, \leq_R)$ *is a complete meet semi-lattice.*

Sketch of proof. Let $x, y \in W$. Then since $1 \leq_R x$ and $1 \leq_R y$, the set L of lower bounds of $\{x, y\}$ is nonempty. An element z of maximal length in L is a meet.

To show this involves the *left descent sets* of x, y, namely those $s \in S$ for which $\ell(sx) < \ell(x)$. This has the property that whenever s is in the descent set for both u and w, then $u \leq_R w$ if and only if $su \leq_R sw$. This lemma allows one to compare another element of L with z and "strip away" simple reflections until the other element is the identity.

The fact that meets exist for pairs together with the fact that ℓ is an $\mathbb{N}$-valued rank function for $(W, \leq)$ allows for a descending chain argument that guarantees a meet for an arbitrary subset. $\square$

5.5 Deletion and Exchange Conditions

Because of the prominent role played by the generating set in the theory of Coxeter groups, it is important to understand the various ways to relate certain expressions of elements of the group W in terms of the set S. There are two criterion which pin down Coxeter groups in the following sense: A group W together with a generating set S of order two elements, (W, S) is a Coxeter system if and only if (W, S) satisfies the Deletion Property if and only if (W, S) satisfies the Exchange Property ($C \Leftrightarrow D \Leftrightarrow E$).

Definition 5.5.1 (Deletion Property)**.** Given an unreduced expression $w = s_1 \ldots s_k$, there exist $1 \leq i < j \leq k$ so that $w = s_1 \ldots \widehat{s_i} \ldots \widehat{s_j} \ldots s_k$.

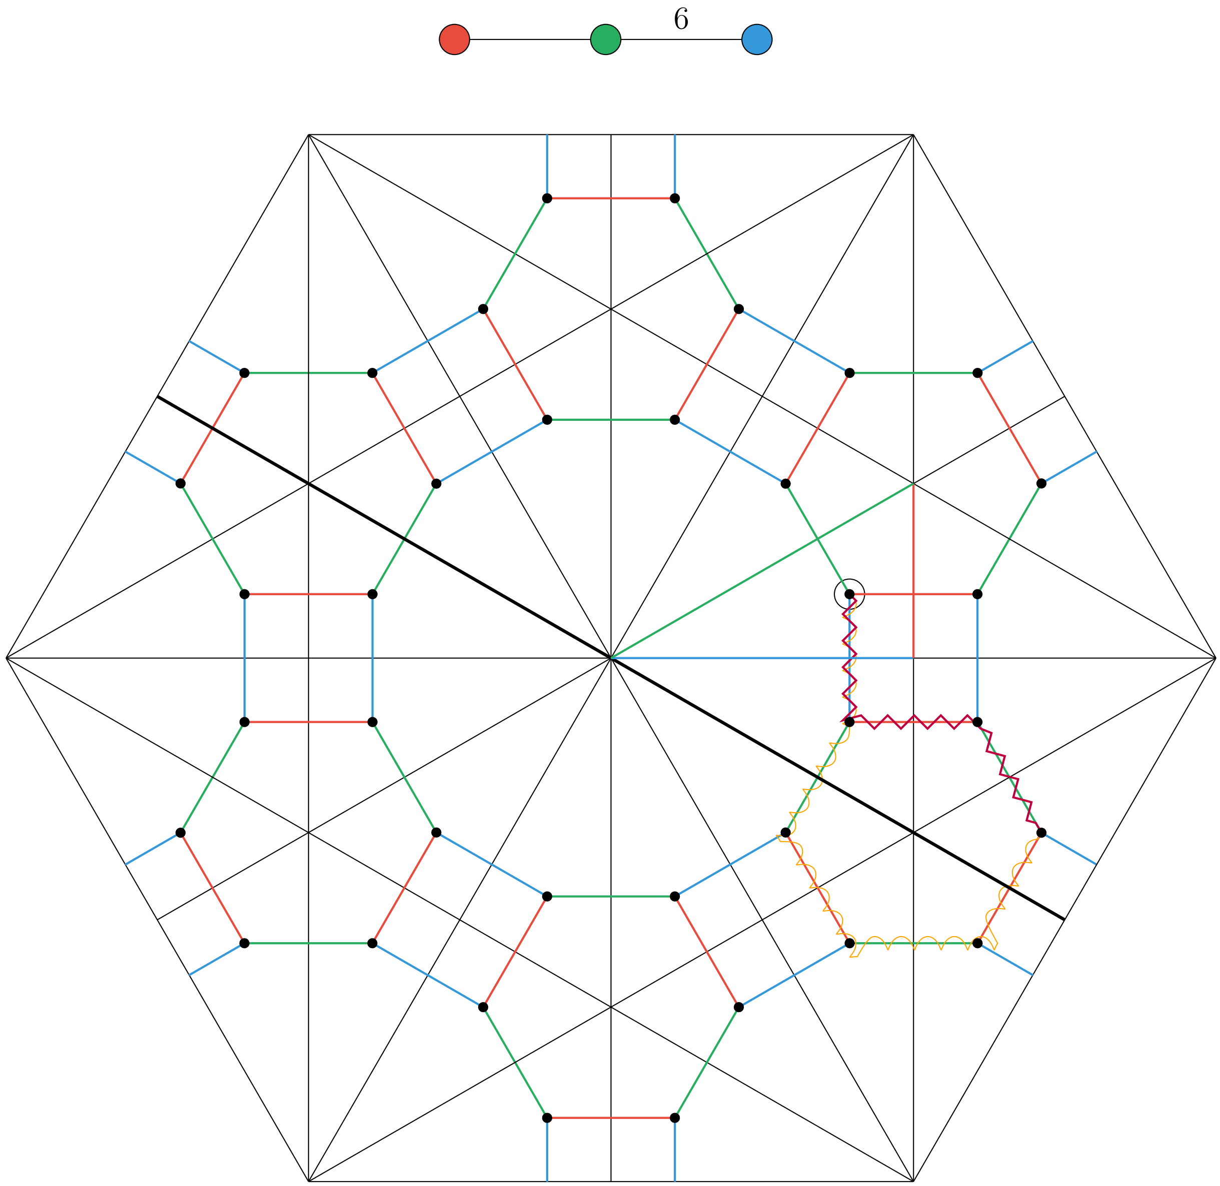

Figure 5.3: The deletion property in $\widetilde{G_2}$.

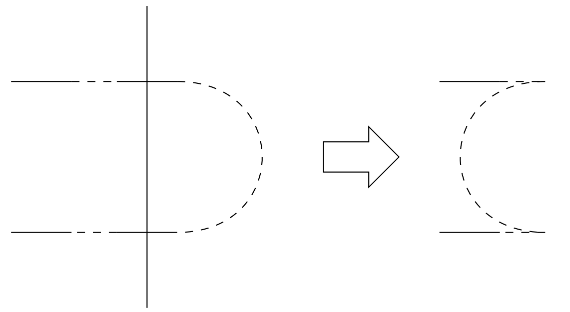

Remark. As the picture suggests, the deleted letters correspond to crossing over the same hyperplane twice. However, this doesn't necessarily mean that $s_i = s_j$.

Definition 5.5.2 (Exchange Property)**.** Take a reduced expression $w = s_1 s_2 \dots s_k$ and $s \in S$. Then $\ell(sw) \leq \ell(w)$ implies that sw has an expression $s_1 \dots \widehat{s_i} \dots s_k$, for some $1 \leq i \leq k$, where $\hat{}$ denotes a removed letter.

Proposition 5.5.3 (Equivalence)**.** *Whenever W is a group generated by a set S of involutions, (W, S) is a Coxeter system if and only if (W, S) satisfies the Deletion Property if and only if (W, S) satisfies the Exchange Property.*

Proof. This is a direct consequence of Theorems 3.2.16 and 3.2.17 in [Dav08]. □

Example 5.5.4. Figure 5.3 illustrates the deletion property in a portion of the Coxeter complex of the Coxeter group $\widetilde{G_2}$. The unreduced word $bgrgr$ becomes the reduced word $b\hat{g}rg\hat{r} = brg$ by avoiding the slightly thicker hyperplane. This hyperplane corresponds to the reflection bgb. The "two wrongs make a right" principle says the following. The action of brg says "apply g, apply r, apply b," leads us to reflect in the corresponding hyperplanes in that order, but reading the word left to right and traversing the Cayley graph in that order will end in the correct place. What is happening in this case is a reflection in b, then a reflection in *the image of r under b*, then a reflection in *the image* *of g under the image of r under b.* The fact that both of these work is related to the equation $(brgrb)(brb)b = brg$

5.6 The Word Problem

Let G be a group with presentation $\langle X \mid R\rangle$. How can one determine whether two words w_1, w_2 in the free group F_X on X represent the same group element in G?

For example, consider the presentation $\langle a, b \mid aba^{-1}b^{-1}\rangle$ for $\mathbb{Z}^2$. Set $w_1 = a^3b^{-2}aba$, and $w_2 = bab^{-1}aba^2$. In order to determine whether $w_1 = w_2$, one can use the relation to move the a's and b's around until there is a word of the form $a^m b^n$. One obtains $w_1 = a^5b^{-1}$ and $w_2 = a^4b$. Call $a^m b^n$ a *normal form* for the presentation $\langle X \mid R\rangle$. Each word in F_X can be arranged to be of this form, and no two words of this form represent the same element in G.

This is called *the word problem for G*. While this reads like a problem in algebra, it has surprising logical implications, and most solutions are both geometric and combinatorial in nature.

It turns out that this problem is *undecidable*, in the sense that there does not exist a single algorithm which will solve this problem for an arbitrary group. In fact, more is true: there exist groups for which there does not exist an algorithm that can solve the word problem. This was shown by Pyotr Novikov in 1955 and again by William Boone in 1958.

A rather salient property about the presentation of a Coxeter group is that the word problem is solvable in Coxeter groups. It can be shown that any word may be converted into a reduced expression by applying *braid-moves* (replacing $sts\ldots$ with $tst\ldots$) and *nil-*

moves (replacing ss with the empty word λ), and moreover, any reduced expressions are related by braid-moves. Thus, one can specify a normal form for each word in a Coxeter group, and then use finitely many braid-moves and nil-moves to convert an arbitrary word to its normal form. The point is one never has to apply a move of the form $\lambda \mapsto ss$.

Theorem 5.6.1. *Let (W, S) be a Coxeter system. If $s_1 \dots s_k$ and $s'_1 \dots s'_k$ are reduced expressions for the same group element w, then there is a sequence of braid-moves connecting them. Moreover, if an expression is not reduced, it may be reduced using only braid-moves and nil-moves.*

Proof. Let $\langle s_i, s_j \rangle^m$ denote the alternating product of length m, starting with s_i. Then a braid move corresponds to replacing $\langle s_i, s_j \rangle^{m_{i,j}}$ with $\langle s_j, s_i \rangle^{m_{i,j}}$.

We induct on k; if $\ell(w) = 0$, it is clear. So suppose it holds for $k-1$. If $s_1 = s'_1$, then the induction hypothesis allows us to relate $s_2 \dots s_k = s'_2 \dots s'_k$ using only braid-moves. So suppose $s_1 \neq s'_1$. Then since both s_1 and s'_1 are prefixes for w, $s_1 \leq w$ and $s'_1 \leq w$. It would be convenient to find some pair of expressions for w which begin with s_1 and s'_1 and are obviously related by a braid-move.

Note that the existence of joins in the weak order supplies us with just that. Since w is an upper bound for both s_1 and s'_1, the set U of upper bounds of $\{s_1, s'_1\}$ is nonempty and so the meet of U is a join for $\{s_1, s'_1\}$. Then w must lie above the join of these two elements, which is $\langle s_1, s'_1 \rangle^{m_{1,1'}}$ (this fact is not completely obvious, but we will omit the proof). So $\langle s_1, s'_1 \rangle^{m_{1,1'}}$ is a prefix for w. Thus, there exists some word β with $\langle s_1, s'_1 \rangle^m \beta = \langle s'_1, s_1 \rangle^m \beta = w$.

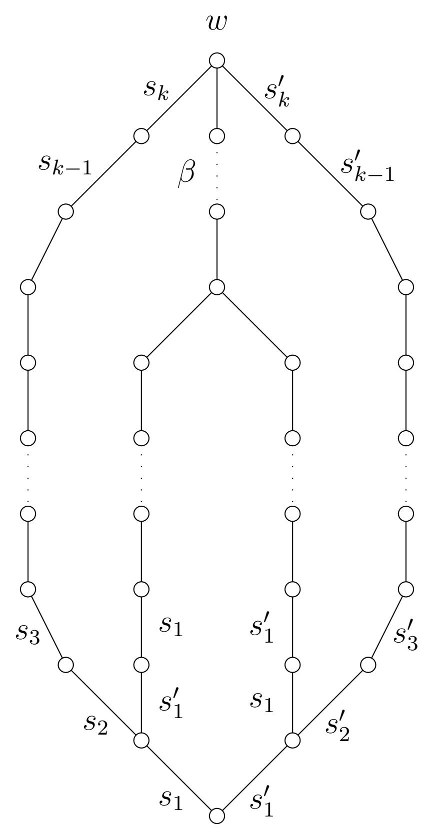

Figure 5.4: An illustration of the proof that the word problem is solvable. Sweep from left to right in the diagram, using the induction hypothesis, followed by a braid-move, and once again using the induction hypothesis to relate two reduced expressions using only braid-moves.

Now, we can use induction to relate $s_1 \ldots s_k$ to $\langle s_1, s'_1 \rangle^m \beta$ via braid-moves since they both begin with s_1, we can relate $\langle s_1, s'_1 \rangle^m \beta$ with $\langle s'_1, s_1 \rangle^m \beta$ by the obvious braid-move, and then again use induction to relate to $s'_1 \ldots s'_k$. $\square$

6. Spherical and Euclidean Coxeter Groups

The type of the bilinear form corresponding to a given Coxeter group is a very consequential piece of information. When the form is positive definite form, the corresponding group turns out to be finite. In fact, it is also true that every finite Coxeter group is positive definite. In this case, the representation acts by linear isometries of V, so the action restricts to a faithful geometric action on S_V, the unit sphere. For this reason, finite Coxeter groups are also called spherical.

When the form is positive semi-definite but not positive definite, the action of W restricts to an affine subspace of V in which W contains a translation, and so there is a geometric action on a Euclidean space. Call such Coxeter groups Euclidean.

Both of these types are well-studied and even classified. The remaining Coxeter groups are the wild ones, and the main objects of interest for us. There is a bit more control when the form has type $(n-1, 1)$, but this is not always the case. When this is the case, call the Coxeter system (W, S) *Lorentzian*.

6.1 Finite Reflection Groups: Classification

It is not too difficult to classify the Coxeter groups which are finite. This was done first by H.S.M. Coxeter in 1935 [Cox35]. This classification is very closely related to many other classifications in mathematics, being an instance of a so-called ADE classification (as are the classification of complex semisimple Lie algebras, and the classification of finite-type quivers).

Lemma 6.1.1. *The Coxeter graph of an irreducible finite reflection group is connected and acyclic, and any graph obtained by removing vertices or decreasing labels corresponds to a finite reflection group.*

Proof. Section 5.3 tells us it suffices to classify the case where the graph is connected, because otherwise the group is simply a direct product of the connected components.

Second, suppose the graph is acyclic, so it is a labeled tree. Let us see why this is. Suppose $s_1, \ldots, s_k$ are distinct labels of vertices forming a cycle in the graph (that is, with indices mod k, $s_j s_{j+1}$ has order at least three). Then I claim that the elements of the infinite sequence $\{(s_1 \ldots s_k)^i\}_{i=1}^{\infty}$ are all reduced words and hence distinct. For suppose $s_1 \ldots s_k s_1 \ldots s_k \ldots s_1 \ldots s_k$ is unreduced. Then as we showed in Section 5.6, we can perform a sequence of "braid-moves" and "nil-moves" to reduce it. Since no adjacent letters are equal in this word, there are no nil-moves to apply. Since all of the graph labels between adjacent elements are at least 3 by assumption, there are no braid-moves to apply. Thus, this word is already reduced (and in fact has a unique reduced expression).

The fact that removing vertices gives a finite subgroup follows from the forthcoming Section 5.3, which does not rely on this result.

Decreasing edge labels retains positive-definiteness due to some technical linear algebra related to Perron-Frobenius theory. It makes use of the fact that when $m < m'$, $\cos(\pi - \frac{\pi}{m}) > \cos(\pi - \frac{\pi}{m'})$. □

Theorem 6.1.2 (Classification of finite reflection groups)**.** *The diagrams listed in Figure 6.1 are all of the irreducible finite reflection groups.*

Proof. We show that these are the only possibilities, but do not explicitly calculate that these are indeed finite reflection groups. We rely on the fact that to have a subgraph corresponding to an infinite subgroup implies the group is not finite. The graph theoretic properties will be crucial. So suppose we have a tree. There is a unique rank 1 Coxeter group, $\mathbb{Z}/2\mathbb{Z}$. If the rank is 2, we can put any natural number m (but not ∞) as a label to obtain the diagram of type $I_2(m)$, and this gives a dihedral group.

Now suppose Γ has at least 3 vertices and, for now, suppose we have no labels. Considerations on degrees of vertices will be our first step.

Case 1: The degree of each vertex is at most 2. These are just the graphs of type A_n, each of which is a finite reflection group.

Case 2: The degree of some vertex is at least 4. In this case, we can take this vertex together with four of its neighbors to see that we contain a subgraph of type $\tilde{D}_4$, so no positive-definite graph can have a vertex of degree 4.

Case 3: Every vertex has degree at most 3.

Subcase i: More than one vertex has degree 3. Suppose x, y have degree 3. Then, since Γ is connected, there is a path from x to y. But this means we have a subgraph of type $\tilde{D}_n$, where $n-3$ is the length of the path.

Subcase ii: We have a unique vertex of degree 3, having legs of length $p \leq q \leq r$ (the vertex of degree 3 is counted in each of these leg lengths). If $p = 3$, then we contain a graph of type $\tilde{E}_6$ since $3 \leq q \leq r$. So p must be 2; if $q = 4$, then we contain a graph of type $\tilde{E}_7$. If $q = 3$ and $r = 6$, we contain a graph of type $\tilde{E}_8$, so we must have $r = 3$, $r = 4$, or $r = 5$, corresponding to the cases E_6, E_7, and E_8, respectively. If $q = 2$, then any value of r gives us a graph of type D_{r+2}.

This classifies the unlabeled connected Coxeter diagrams. So suppose we have two edges labeled. This gives a subgraph of type $\tilde{C}_n$, by decreasing these labels to 4 and removing edges not forming a path between these edges. So there can only be a single labeled edge. If there's a vertex of degree 3, we contain a graph of type $\tilde{B}_n$, so in fact we have a path with a single labeled edge.

If our label is at least 6, we contain a graph of type $\tilde{G}_2$, since we have at least three vertices. So suppose our label is 5. If our graph is type H_3 or H_4, it is finite. Any other graph with a label of 5 either contains 5 in an edge which is not on the end, so we contain a graph of type Z_4, or has at least 5 vertices, which contains the graph of type Z_5.

If our label is 4 and it is on the end of the path, we have type B_n/C_n. So suppose the labeled edge is not on the end. Then we either have type F_4, or contain $\tilde{F}_4$. This completes the classification. $\square$

The diagrams in Figure 6.2 and Figure 6.3 are used in the classification. The diagrams in Figure 6.3 have type $(3, 1)$ and $(4, 1)$ respectively.

6.2 Geometric Group Actions in Coxeter Groups

Recall Section 1.1 established a correspondence between groups and metric spaces. This section uses this correspondence to restrict the view of the representation constructed in Section 5.2 to a more appropriate domain.

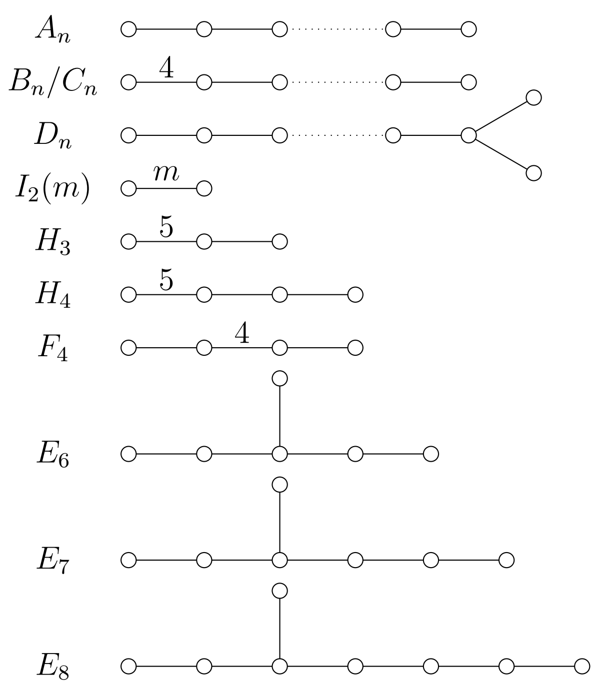

Figure 6.1: The positive definite graphs

The type of the bilinear form B determines the large scale geometry of a Coxeter system. The type of the form B determines what sort of geometric object is preserved under the B-orthogonal group. If B is positive definite, the unit sphere S^{n-1} of V is preserved. If B is positive semi-definite but not positive definite, the restriction to the positive-definite part of V is a sphere, but the entire radical is preserved. This is a cylinder $S^{n-2} \times \mathbb{R}$ (recall the radical is one-dimensional when (W, S) is irreducible). Finally, when the type is $(n-1, 1)$, a hyperboloid is preserved.

Definition 6.2.1 (Dual representation)**.** Given the geometric representation $\sigma\colon W \to O_B(V)$, where V is a vector space with basis $\{\alpha_s \mid s \in S\}$, can construct a dual representation $\sigma^*\colon W \to \mathsf{GL}(V^*)$, by demanding $(\sigma^*(w)f)(v) = f(\sigma(w^{-1})v)$, whenever $v \in V$, $w \in W$, and $f \in V^*$. The notation $\langle f, v\rangle \stackrel{\text{def}}{=} f(v)$ is thus useful, since then the re-

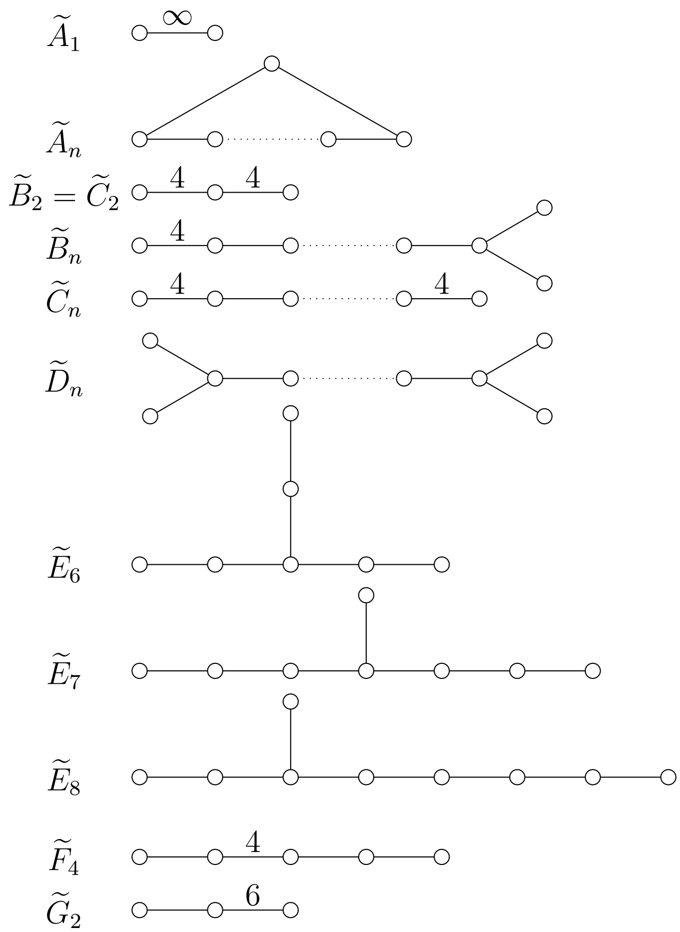

Figure 6.2: The affine Coxeter diagrams. Each of these graphs corresponds to a bilinear form which has a one-dimensional radical, and the restriction to any codimension one subspace not containing the radical is positive definite.

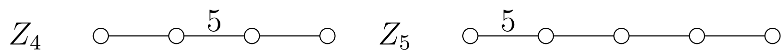

Figure 6.3: The two indefinite graphs encountered in the classification.

quired equation is $\langle \sigma^*(w)f, v\rangle = \langle f, \sigma(w^{-1})v\rangle$. This is meant to be reminiscent of the corresponding formula for inner products.

Definition 6.2.2 (Coxeter complex)**.** Let (W, S) be a Coxeter system, and consider the dual representation $\sigma^*\colon W \to \mathsf{GL}(V^*)$. The fundamental chamber is $C = \{f \in V^* \mid f(\alpha_s) > 0\ \forall s \in S\}$. The *Tits cone* is the W-orbit of the closed fundamental chamber $\overline{C}$, and the *Coxeter complex* $\Sigma(W, S)$ is the quotient of (the nonzero vectors of) the Tits cone by $\mathbb{R}^+$.

Example 6.2.3. Examine the case of the Coxeter group of type $I_2(m)$, which is the dihedral group D_{2m}. As this is a finite group, it can actually act geometrically on a one-point space. The quasi-isometry condition is a bit too coarse to keep track of finite or bounded information. However, since a one-point space is not quasi-isometric to $\mathbb{R}^2$, this can restrict to some smaller space. Upon taking a quotient by the positive reals, there is a geometric action on the unit sphere. This is not a coincidence but actually the general case for finite Coxeter groups: restrict to the unit sphere $S_V = \{v \in V \mid B(v, v) = 1\}$. Finite Coxeter groups are often called spherical Coxeter groups for this reason.

Since B is positive definite, it is possible to identify V and V^*, and so it is unnecessary to dualize.

Example 6.2.4 (Euclidean and $\widetilde{A_2}$)**.** In the case that B is positive semi-definite but not positive definite, the representation of W induces a geometric action on an affine Euclidean space. Suppose also that W is irreducible, and so the corresponding Coxeter diagram is connected.

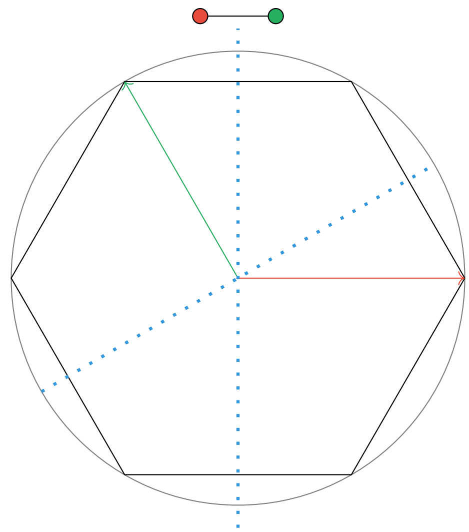

Figure 6.4: The Coxeter complex for $I_2(3) = A_2$. The red and green vectors are the simple roots, and their corresponding hyperplanes are shown dotted.

Let A be the matrix for B with respect to the basis $\{\alpha_s\}_{s\in S}$. Then a technical linear algebraic fact tells us that the radical of the form B coincides with the null space of A, and this null space is one-dimensional spanned by a vector with strictly positive coordinates $\lambda \in \mathsf{conv}(\{\alpha_w\}_{w\in S})$.

Focus on the following example.

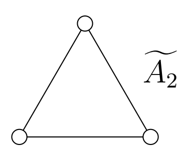

In this case, the simple roots are linearly independent, but the bilinear form is such that their 3 hyperplanes intersect in a common line spanned by $(1, 1, 1)$. What this means is that in the dual space, the corresponding linear functionals lie on a plane, and provide a tiling of equilateral triangles. However, in the original space, the generating reflections fix the cylinders with axis $\mathsf{span}(1, 1, 1)$. The roots all lie on one such cylinder.

The Schläfli matrix of the Coxeter group of type $\tilde{A}_2$ is $\begin{bmatrix} 1 & -1/2 & -1/2 \\ -1/2 & 1 & -1/2 \\ -1/2 & -1/2 & 1 \end{bmatrix}$. Denote the corresponding quadratic form by q, so $q(1,1,1) = 0$, and one can check that this vector spans the radical of the positive semi-definite form.

More directly, one can calculate the various orthogonal hyperplanes H_i for the simple roots α_i. These are $H_1 = \{v \in \mathbb{R}^3 \mid 2v_1 = v_2 + v_3\}$, $H_2 = \{v \in \mathbb{R}^3 \mid 2v_2 = v_1 + v_3\}$, and $H_3 = \{v \in \mathbb{R}^3 \mid 2v_3 = v_1 + v_2\}$. As stated, these planes all intersect along the line L spanned by $(1,1,1)$, and since each reflection s_i leaves fixed H_i, this line is fixed under the action of W. In fact, the cylinders centered around L are invariant under the action of W.

The matrices representing s_1, s_2, and s_3 can be calculated as $s_i(v) = v - 2B(v,\alpha_i)\alpha_i$. Thus

$$s_1 = \begin{pmatrix} -1 & 1 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, s_2 = \begin{pmatrix} 1 & 0 & 0 \\ 1 & -1 & 1 \\ 0 & 0 & 1 \end{pmatrix}, s_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & -1 \end{pmatrix}.$$

Pass to the dual representation. Here, the matrices are

$$s_1^* = \begin{pmatrix} -1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}, s_2^* = \begin{pmatrix} 1 & 1 & 0 \\ 0 & -1 & 0 \\ 0 & 1 & 1 \end{pmatrix}, s_3^* = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & -1 \end{pmatrix}.$$

Note that these preserve the affine hyperplane $v_1 + v_2 + v_3 = 1$. Indeed,

$$s_1^* \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} = \begin{pmatrix} -v_1 \\ v_1 + v_2 \\ v_1 + v_3 \end{pmatrix}, s_2^* \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} = \begin{pmatrix} v_1 + v_2 \\ -v_2 \\ v_2 + v_3 \end{pmatrix}, s_3^* \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} = \begin{pmatrix} v_1 + v_3 \\ v_2 + v_3 \\ -v_3 \end{pmatrix},$$

each of whose coordinates sum to $v_1 + v_2 + v_3$, which is again 1.

Example 6.2.5 (Rank 3 hyperbolic)**.** A rank 3 Coxeter group is described by the Coxeter matrix $\begin{bmatrix} 1 & a & b \\ a & 1 & c \\ b & c & 1 \end{bmatrix}$ corresponding to the graph below.

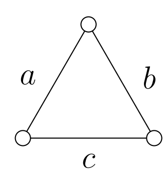

The geometry is determined by the number $\frac{1}{a} + \frac{1}{b} + \frac{1}{c}$. If this value is greater than 1, then the Coxeter group is finite. If the sum is exactly one, then the Coxeter group is Euclidean, and if the sum is less than one, we are in the Lorentzian case. One sees this by examining the corresponding bilinear form and checking its type.

A somewhat less formal but more geometric way to see why this is the case goes as follows. The fundamental chamber for the standard representation is a cone over a triangle with dihedral angles π/a, π/b, and π/c. Whether this is a spherical, Euclidean, or hyperbolic triangle depends on how the sum of these angles compares with π.

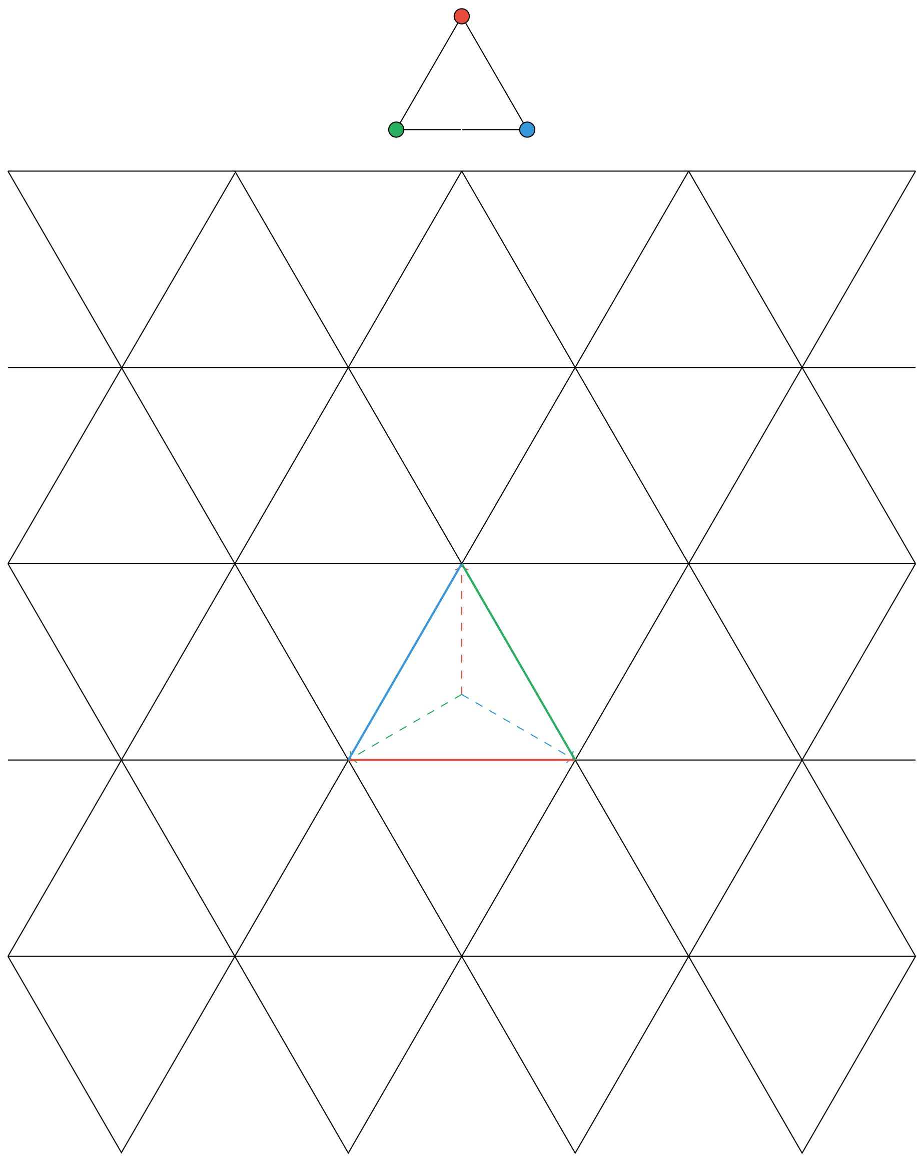

Figure 6.5: The $\widetilde{A}_2$ Coxeter group. The picture is of the affine span of the simple root system. The simple system is indicated by the dashed arrows. The intersection of the corresponding hyperplanes with the affine hyperplane V_1 is indicated by color. Dual to this collection of chambers is the Cayley graph.

Note that equality holds if and only if these numbers are $(2,3,6)$, as in $\widetilde{G_2}$, $(2,4,4)$ as in $\widetilde{B_2}$, or $(3,3,3)$, as in $\widetilde{A_2}$, corresponding to Euclidean Coxeter groups. The finite Coxeter groups were already classified.

The rank 3 examples lead us to a potential issue that can arise with hyperbolic Coxeter groups. Namely, if any of a, b, c is ∞, the corresponding fundamental domain is an unbounded region of the hyperbolic space it acts on, and so the action of such a Coxeter group is not geometric in the sense defined. Often, one can distinguish Coxeter groups of type $(n-1,1)$ by this feature.

The intersection of the simplicial cone and the hyperbolic space can be compact, or it can be an unbounded region with finite volume, or it may be neither of these cases. These distinction are explored in Part III

Part III

Lorentzian Coxeter Systems

If the bilinear form associated to a Coxeter system (W, S) is of Lorentzian type, the root system Φ is an infinite discrete subset of the vector space V. Although it does not have limit points itself, viewing the root system in a particular affine patch of the projective space provides more information about the root system. In particular, it is then possible to determine limiting directions of the roots. This is the subject of Chapter 7. These limit directions must be in the light cone, and hence the boundary of the hyperbolic space.

This fact might lead one to conjecture that the limit roots have some relationship to the the limit *set* of the group. Hyperbolic space has a natural notion of a boundary, and W inherits such a notion through its action on hyperbolic space. Chapter 8 studies this object. It is shown that the limit set is the same as the limit roots. In a sense, this theorem unites the dual theories of root systems and the hyperplanes in the hyperbolic space upon which the group acts.

As a note, a large part of this theory goes through without a problem without the Lorentzian hypothesis at all. However, to avoid highlighting when it comes in and out, it is a standing hypothesis in most definitions and theorems.

7. The Limit Roots

Let (W, S) be a Coxeter system, and Δ the associated simple system in (V, B). Let $\mathsf{Half}^+(\alpha) = \{v \in V \mid B(v, \alpha) \geq 0\}$, and define the *fundamental chamber* $C = \cap_{\alpha \in \Delta} \mathsf{Half}^+(\alpha)$. When (V, B) has type $(n, 1)$, consider $P = -C \cap \mathbb{H}$. In [Rat06], [Hum92], [Dav08], a Coxeter group is called hyperbolic if this set P has finite volume, and (W, S) is called compact hyperbolic if moreover P is compact. An ideal triangle is an example of a finite volume subset of $\mathbb{H}^2$ which is not compact. The hyperbolic Coxeter groups in this sense are classified.

To circumvent this dispute, we follow other authors by saying that a Coxeter group is *Lorentzian* when the associated bilinear form is of type $(n-1, 1)$. This is a more general notion than Humphreys' hyperbolic, as it does not require that the fundamental chamber C is contained in Q^-.

To each $v \in V$ with $B(v, v) \neq 0$, one can associate $s_v \in O_B(V)$ by $s_v(w) = w - 2\frac{B(v,w)}{B(v,v)}v$, which is called the *$B$-reflection associated with* v. So let $S = \{s_\alpha \mid \alpha \in \Delta\}$, and then $W = \langle S \rangle$ the subgroup of $O_B(V)$ generated by S. Finally, let $\Phi = W(\Delta)$ be the W-orbit of the simple system, and then (W, S) is a Coxeter system, and (Φ, Δ) is said to be a based root system in (V, B).

Note that (V, B) together with (Φ, Δ) actually determines (W, S), so a Lorentzian Coxeter system is uniquely specified by a based root system in a quadratic space.

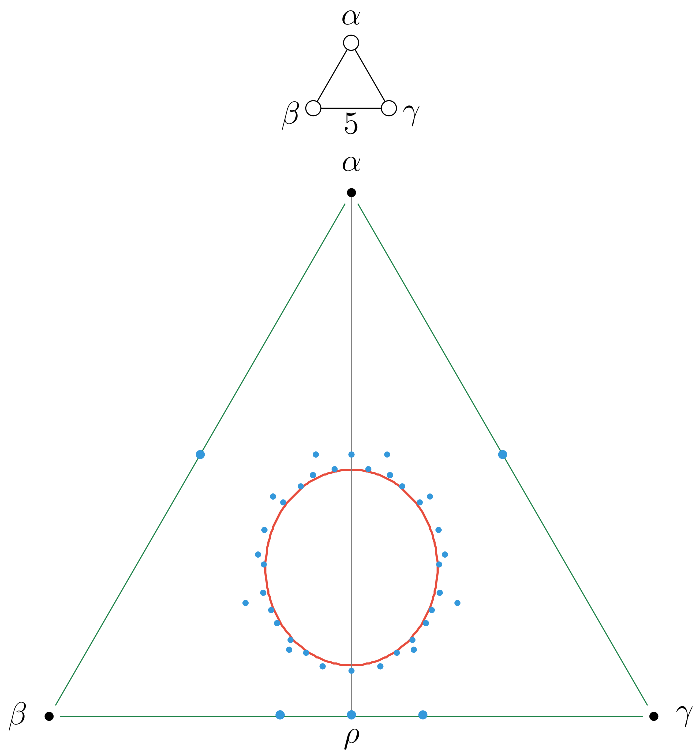

Figure 7.1: The gray line spanning between α and ρ intersects the 1-sphere $\widehat{Q}$ (in red) nontangentially, so the form is indefinite on this plane, with $B(\alpha, \rho) < -1$.

7.1 Reflection Subgroups and the Generalized Geometric Representation

The parabolic subgroups are the easiest to manage. That is, given a Coxeter system (W, S), one can consider subsets I of S, and examine the subgroup W_I of W generated by I. As in Section 5.3, the pair (W_I, I) is again a Coxeter system. The situation is less predictable when I is merely a subset of $R = \{w^{-1}sw \mid s \in S, w \in W\}$, as in Example 7.1.1.

Example 7.1.1. Consider the following Coxeter diagram and the standard representation associated to it.

The parabolic subgroup generated by s_β and s_γ has order 10, with longest word $s_\gamma s_\beta s_\gamma s_\beta s_\gamma = s_\beta s_\gamma s_\beta s_\gamma s_\beta$. The corresponding root is $\rho = s_\gamma s_\beta(\gamma) = s_\beta s_\gamma(\beta)$ (this satisfies the formula $s_{w(\alpha)} = w s_\alpha w^{-1}$ calculated previously).

We can calculate $\rho = \frac{1+\sqrt{5}}{2}(\beta + \gamma)$, and then see that

$$B(\alpha, \rho) = \frac{1+\sqrt{5}}{2}\left(B(\alpha,\beta) + B(\alpha,\gamma)\right) = -\frac{1+\sqrt{5}}{2} < -1.$$

The reflection subgroup generated by s_ρ and s_α is an infinite dihedral group by Lemma 7.2.1. If one considers this reflection subgroup abstractly and construct the corresponding standard representation, one obtains $B(\alpha, \rho) = -1$, which of course is different from the present situation. To amend this lack of "functoriality," one can allow for $B(\alpha_i, \alpha_j)$ to be any value less than or equal to -1 when this product has infinite order. When specifying a Coxeter group by a diagram, this thesis labels the infinite edges with the value $B(\alpha_i, \alpha_j)$, so that the diagram encodes all of the information of the representation.

7.2 Dihedral Subgroups

This section collects technical results about dihedral subgroups occurring in a Coxeter group. These results have important consequences in the developing theory of limit roots.

Lemma 7.2.1. *Suppose α and β are positive roots in a based root system (Φ, Δ).*

- *If $|B(\alpha, \beta)| < 1$, then $W_{\{\alpha,\beta\}}$ is finite.*

- *If $B(\alpha, \beta) \leq -1$, then $W_{\{\alpha,\beta\}}$ is an infinite dihedral group, and moreover, the roots $(s_\alpha s_\beta)^n(\alpha)$ are distinct for each n and lie in $\mathsf{cone}(\alpha, \beta)$. The symmetric statement in α and β of course holds as well.*

Sketch of proof. For the first part, the condition on $B(\alpha, \beta)$ implies that B is positive definite on the restriction to their span, and since the roots are discrete, this implies that the subgroup they generate is finite.

For the other statement, an explicit calculation shows that the α and β coefficients are strictly increasing in the sequence of roots $(s_\alpha s_\beta)^n(\alpha)$. □

Lemma 7.2.2. *The set of values in $(-1, 1)$ which occur as $B(\alpha, \rho)$ for some simple $\alpha \in \Delta$ and positive $\rho \in \Phi^+$ is finite (and nonempty). Thus, there is a minimum positive value κ occurring as $|B(\alpha, \rho)|$.*

Sketch of proof. The idea here is that, given a Coxeter group W, any value that occurs as $B(\alpha, \rho)$ with magnitude less than 1 will occur within a finite subgroup (and in fact generate one!). The collection of such subgroups up to conjugacy is finite since W is a finitely generated Coxeter group, and so there are only finitely many values of $B(\alpha, \rho)$ in $(-1, 1)$. Take κ to be the minimum. □

Remark. In the previous proof, it is not obvious that a finitely generated Coxeter group has only finitely many conjugacy classes of finite subgroups, but the hypothesis cannot be weakened much. It is easy to write down a non-finitely generated Coxeter system which has infinitely many conjugacy classes of finite subgroups. It is harder to find a finitely generated group with infinitely many conjugacy classes of finite subgroups, but [Bri00] provides such an example.

Of course, in the case of an infinite Coxeter group, the associated root system is also infinite. In this setting, the root system is a discrete subset of V, directional limits can be considered.

Definition 7.2.3 (Depth of a root)**.** Let (Φ, Δ) be a based root system in a quadratic space (V, B), with the corresponding Coxeter system (W, S). If ρ is a positive root, then define the *depth of* ρ to be $\mathsf{depth}(\rho) = 1 + \min\{\ell(w) \mid w(\rho) \in \Delta\}$. So depth 1 roots are the simple roots, depth 2 roots are the positive roots which can be obtained as a simple reflection of a simple root, and so on.

Since S is finite, there are only finitely many words in S^* of length at most n. Consequently, there are only finitely many roots of depth at most n. Thus, if ρ_n is an injective sequence of roots, $\mathsf{depth}(\rho_n) \to \infty$. To see that this implies that $\|\rho_n\| \to \infty$ a lower bound for the norm in terms of the depth can be obtained. Since any two norms on a finite dimensional vector space give the same topology, it is inconsequential how the norm is selected, so one may as well choose $\|\cdot\|$ to be the norm induced by declaring Δ to be an orthonormal basis.

The next few sections explore how the depth of a root relates to the Euclidean norm. This first requires a technical lemma regarding depth.

Lemma 7.2.4 (Depth lemma)**.** *Suppose (W, S) is a Coxeter system with corresponding based root system (Φ, Δ). Let $\alpha \in \Delta$ and $\beta \in \Phi^+ \setminus \{\alpha\}$. Then the depth of $s_\alpha(\beta)$ is*

$$\begin{cases} \mathsf{depth}(\beta) + 1 & \text{if } B(\alpha, \beta) < 0, \\ \mathsf{depth}(\beta) & \text{if } B(\alpha, \beta) = 0, \\ \mathsf{depth}(\beta) - 1 & \text{if } B(\alpha, \beta) > 0, \end{cases}$$

Proof. If $B(\alpha, \beta) = 0$, then $s_\alpha(\beta) = \beta$ and so $\mathsf{depth}(s_\alpha(\beta)) = \mathsf{depth}(\beta)$.

Now suppose $B(\alpha, \beta) > 0$. Choose a w with $w(\beta) < 0$ achieving $\ell(w) = \mathsf{depth}(\beta)$.

If $ws_\alpha < w$, then $ws_\alpha(s_\alpha(\beta)) = w(\beta) < 0$ and so $\mathsf{depth}(s_\alpha(\beta))$ is bounded above by $\ell(ws_\alpha) < \ell(w) = \mathsf{depth}(\beta)$. This has to be $\mathsf{depth}(\beta) - 1$ (since if it were smaller, we would have a u with $\ell(u) < \ell(w)$ for which $u(\beta) < 0$, a contradiction).

If instead $w < ws_\alpha$, then by Proposition 5.2.8, we have $w(\alpha) > 0$. Consider the root $ws_\alpha(\beta) = w(\beta) - 2B(\alpha, \beta)w(\alpha)$. Since $w(\beta) < 0$, $B(\alpha, \beta) > 0$, and $w(\alpha) > 0$, we have $ws_\alpha(\beta) < 0$ (and moreover, $ws_\alpha(\beta)$ is not simple because it is the sum of two negative roots). Now choose $t \in S$ so that $\ell(tw) < \ell(w)$. Then $tw(s_\alpha(\beta)) < 0$ since t permutes the negative roots which are not $-\alpha_t$. So $\mathsf{depth}(s_\alpha\beta) \leq \ell(tw) < \ell(w) = \mathsf{depth}(\beta)$.

For the final equation, suppose $B(\alpha, \beta) < 0$. Then $B(\alpha, s_\alpha(\beta)) > 0$, and so $\mathsf{depth}(\beta) = \mathsf{depth}(s_\alpha s_\alpha \beta) = \mathsf{depth}(s_\alpha(\beta)) - 1$, or just $\mathsf{depth}(s_\alpha(\beta)) = \mathsf{depth}(\beta) + 1$, as desired. □

Consequently, if ρ is a root of depth r, then there is a root π of depth $r - 1$ and a simple root $\alpha \in \Delta$ so that $\rho = s_\alpha(\pi)$ and moreover, $B(\alpha, \pi) < 0$.

Remark. Lemma 7.2.4 allows one to define a partial order on the set of positive roots. The covering relations of this order are defined by setting $\beta \lhd \gamma$ if there is some $s \in S$

with $s(\beta) = \gamma$ and $\mathsf{depth}(\gamma) = \mathsf{depth}(\beta) + 1$. As is clear from the definition, depth is a rank function for this ordering.

7.3 Roots Diverge

This section shows that infinite root systems are discrete, and that any bounded set has only finitely many roots.

Lemma 7.3.1. *There is a $\lambda > 0$ so that for any $\rho \in \phi^+$, $\|\rho\|^2 \geq 1 + \lambda(\mathsf{depth}(\rho) - 1)$.*

Proof. We induct on $\mathsf{depth}(\rho)$. Declaring the simple roots to be an orthonormal basis not only allows us to define a norm, but we obtain a corresponding inner product $\langle \cdot, \cdot \rangle$. Observe that for a simple root α and positive root π, we have $\langle \alpha, \pi \rangle \geq 0$, and $\|\alpha\| = 1$. With these observations, it is clear that the claim holds for a root of depth 1. Suppose the inequality holds for any root of depth $r - 1$, and let ρ be a root of depth r. Then by Lemma 7.2.4, there is a root π of depth $r - 1$ and a simple root α with $s_\alpha(\pi) = \rho$ and $B(\alpha, \pi) < 0$. We then obtain

$$\|\rho\|^2 = \|s_\alpha(\pi)\|^2 = \|\pi - 2B(\alpha, \pi)\alpha\|^2 \geq \|\pi\|^2 + 4B(\alpha, \pi)^2 - 4B(\alpha, \pi)\langle \alpha, \pi \rangle.$$

Now since $B(\alpha, \pi) < 0$ and $\langle \alpha, \pi \rangle > 0$, we have $\|\rho\|^2 \geq \|\pi\|^2 + 4B(\alpha, \pi)^2$. Applying the induction hypothesis yields $\|\pi\|^2 + 4B(\alpha, \pi)^2 \geq 1 + \lambda(\mathsf{depth}(\pi) - 1) + 4B(\alpha, \pi)^2 \geq 1 + \lambda(r - 2) + \lambda = 1 + \lambda(r - 1)$ □

Theorem 7.3.2. *The norm of any injective sequence of roots ρ_n goes to infinity.*

Proof. There are only finitely many roots of bounded depth. □

7.4 Projecting the Roots

Although the previous section implies there is no hope of finding a limit point of the set Φ in V, one can look for "directional limits." Since the concept of direction is captured in the projective space, one can use this and its topology to study limit roots. However, to keep things concrete, another approach of projecting the roots is used. Namely, the ray spanned by each root is intersected with an affine hyperplane, carefully chosen so that each such ray intersects it. This amounts to selecting an affine patch of the projectivization of V.

Example 7.4.1 ($\widetilde{A}_1$)**.** Consider the following representation of $\widetilde{A}_1$.

$$\alpha \, \overset{(-1.01)}{\underset{\infty}{\circ\!\!-\!\!\!-\!\!\!-\!\!\circ}} \, \beta$$

Note that $B(\alpha+\beta, \alpha+\beta) = B(\alpha, \alpha) + 2B(\alpha, \beta) + B(\beta, \beta) = 1 + 2(-1.01) + 1 = -.02 < 0$, so $\alpha + \beta \in Q^-$, the point being that Q^- is nonempty.

Definition 7.4.2 (Affine hyperplane)**.** Suppose (Φ, Δ) is a based root system in a quadratic space (V, B), and let $\varphi\colon V \to \mathbb{R}$ be the linear functional so that $\varphi(\alpha) = 1$ for each $\alpha \in \Delta$. Thus φ sums the coordinates of a vector written in the basis Δ. The kernel of φ is a hyperplane in V which does not intersect any simple root in Δ. The main object of interest will be the affine subspace determined by Δ, which can be viewed as $V_1 \overset{\text{def}}{=} \varphi^{-1}(\{1\}) = \mathsf{aff}(\Delta)$. While all simple roots lie in V_1, the other positive roots need not. These roots are, by their very definition, positive linear combinations of elements of Δ, and so for $\rho \in \Phi^+$, it is true that $\varphi(\rho) > 0$. So one can consider the projection of ρ onto V_1 defined by $\widehat{\rho} \overset{\text{def}}{=} \frac{\rho}{\varphi(\rho)}$. Note that for a positive root ρ, $\varphi(\rho) = |\rho|_1$, where $|\cdot|_1$ is

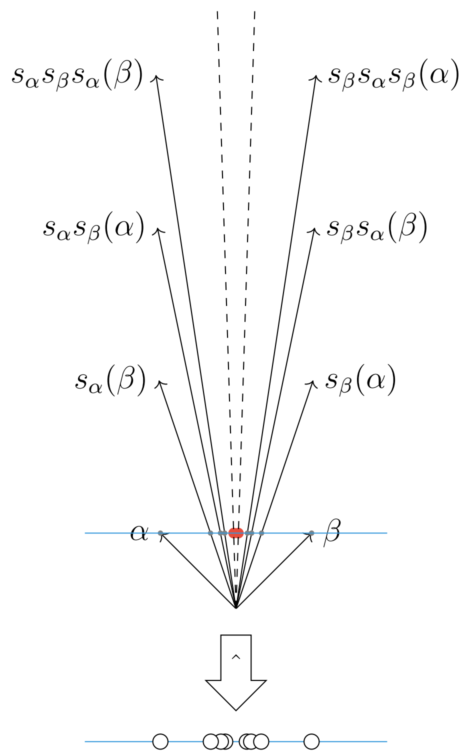

Figure 7.2: The roots in a Lorentzian representation of $\widetilde{A}_1$. There are two limit points when projected onto V_1, the affine span of Δ, and these limit points coincide with the projected light cone $\widehat{Q}$, which in this case consists of two points.

the ℓ_1 norm, defined by summing the absolute values of the coordinates in the basis Δ. Since $\mathbb{R}\alpha \cap \Phi = \{\pm\alpha\}$, the map between Φ^+ and $\widehat{\Phi}$ is a bijection and $|\Phi^+| = |\widehat{\Phi}|$.

In terms of Figure 7.2, the projectivization of the roots is the intersection of the line spanned by each positive root with the blue (horizontal) line, $V_1 = \mathsf{aff}\{\alpha, \beta\}$.

Proposition 7.4.3. *If (Φ, Δ) is a based root system in a quadratic space (V, B), then the set $\widehat{\Phi}$ of normalized roots lies within $\mathsf{conv}(\Delta)$.*

Proof. When ρ is a positive root, $\widehat{\rho}$ has positive coefficients as well. This just says that $\widehat{\rho} \in \mathsf{cone}(\Delta)$, and so in fact $\widehat{\rho}$ is in $\mathsf{cone}(\Delta) \cap \mathsf{aff}(\Delta) = \mathsf{conv}(\Delta)$. $\square$

Definition 7.4.4. If $\widehat{\Phi}$ is the set of normalized roots, let $\mathsf{Lim}(\Phi)$ be the limit set of $\widehat{\Phi}$.

Corollary. *Since the convex hull of a finite set is compact, $\mathsf{Lim}(\Phi)$ is empty if and only if Φ is finite if and only if B is positive-definite.*

7.5 Limit Roots are Light-like

The limit roots lie in the light cone $\widehat{Q} = \widehat{q^{-1}(0)}$. Since there is a natural way of viewing the interior of the projected isotropic cone Q as a hyperbolic space, there is ultimately a way of understanding the light cone as a sort of boundary of this hyperbolic space. This is explored in Section 8.1.

Theorem 7.5.1. *Let (Φ, Δ) be a based root system in a quadratic space (V, B), and recall that $q(v) = B(v, v)$. The set of limit roots $\mathsf{Lim}(\Phi)$ is contained within the set $\widehat{Q} = \{v \in V_1 \mid q(v) = 0\}$.*

Proof. Suppose ρ_n is an injective sequence of roots, and that $\widehat{\rho}_n$ converges. We'll first show that $\varphi(\rho_n) \to \infty$. For this, note that $\varphi(\rho_n) = |\rho_n|_1$, and we showed in Theorem 7.3.2 that $\|\rho_n\| \to \infty$. Invoke Theorem 2.3.9 to prove the claim. Now calculate

$$q(\widehat{\rho}_n) = q\left(\frac{\rho_n}{\varphi(\rho_n)}\right) = \frac{q(\rho_n)}{\varphi(\rho_n)^2} = \frac{1}{\varphi(\rho_n)^2}$$

where in the final equality, we have used the fact that each root is the image of a simple root under an element of $O_B(V)$, so every root ρ has $q(\rho) = 1$. Since we have shown that $\varphi(\rho_n) \to \infty$, we obtain $q(\widehat{\rho}_n) \to 0$, and so every limit root lies within $\widehat{Q}$. $\square$

So in fact $\mathsf{Lim}(\Phi) \subseteq \mathsf{conv}(\Delta) \cap \widehat{Q}$.

7.6 Group Action on Limit Roots

The Lorentzian space built out of the group presentation of course comes along with an action of the group on that space, which can be restricted to one sheet of the hyperboloid of two sheets. Upon projecting onto an affine hyperplane as in Section 7.4, this hyperboloid becomes the interior $\mathbb{H}_p$ of the set $\widehat{Q} = \widehat{q^{-1}(0)}$, and the action of W can extend to the boundary $\partial\mathbb{H}_p = \widehat{Q}$. Since the limit roots $\mathsf{Lim}(\Phi)$ lie on this boundary, this provides an action of W on $\mathsf{Lim}(\Phi)$.

Alternatively, one can try to let W act on V_1 by the formula $w \cdot \widehat{v} = \widehat{w \cdot v}$. Of course, this makes sense only when $w \cdot v$ does not lie in V_0. The maximal subset of V_1 upon which W can act cannot include V_0 or its W-orbit. Take $D = V_1 \setminus (\bigcup_{w \in W} wV_0)$, and the given formula provides an action of W on D.

Proposition 7.6.1. *If x is a limit root between ρ_1 and ρ_2, then wx is a limit root between $w\rho_1$ and $w\rho_2$.*

There is a geometric approach to understanding this action. If α is a projected root and $x \in D \cap Q$, the affine line in V_1 determined by α and x, $L(\alpha, x)$, may intersect Q tangentially, in which case s_α fixes x. This geometric condition is equivalent to $B(\alpha, x) = 0$. When this algebraic condition is not met, then $L(\alpha, x)$ intersects Q in a second point y, and $s_\alpha(x) = y$.

This suggests a partial order one can put on the set $\mathsf{conv}(\Delta) \cap Q^+$, which can be called the *visibility order*. Say that $\alpha < \beta$ if $B(\alpha, \beta) > 0$ and $\mathsf{Ray}(\alpha, \beta) = \{\alpha + t(\beta - \alpha) \mid t \geq 0\}$ intersects $\widehat{Q}$. This second geometric condition is equivalent to the inequality (which is satisfied with $+$ if and only if it is satisfied with $-$)

$$\frac{B(\alpha, \alpha - \beta) \pm \sqrt{B(\alpha, \beta)^2 - B(\alpha, \alpha)B(\beta, \beta)}}{B(\alpha - \beta, \alpha - \beta)} > 1,$$

or just

$$\sqrt{-\det \begin{pmatrix} B(\alpha, \alpha) & B(\alpha, \beta) \\ B(\beta, \alpha) & B(\beta, \beta) \end{pmatrix}} < B(\beta - \alpha, \beta).$$

For this expression to even make sense requires that the determinant is nonpositive, which is equivalent to the algebraic condition that the reflections associated to α and β generate an infinite dihedral group and the geometric condition that their plane spanned intersects Q.

Having established this order on the exterior of the projective light cone, the geometric interpretation is as follows. The orthogonal hyperplane determined by a vector α in $\mathsf{conv}(\Delta) \cap Q^+$ can be found geometrically by taking the affine span of points x in Q so that $L(\alpha, x)$ intersects Q tangentially. Thus $\alpha \leq \beta$ if and only if $\mathsf{Half}^+(\alpha) \subseteq \mathsf{Half}^+(\beta)$.

Conjecture 7.6.2. If λ is a limit root, there is a sequence $\alpha_1 = \rho_1 \triangleleft \rho_2 \triangleleft \rho_3 \triangleleft \dots$ with $\lim_{n\to\infty} \rho_n = \lambda$.

Remark. For some limit roots λ, it is possible to have a sequence $\rho_1 \not\leq \rho_2 \not\leq \dots$ which converges to λ. Suppose that λ is a limit root with an injective sequence $\{\lambda_n\}$ of limit roots. For each λ_n, choose ε_n such that $\varepsilon_n < d(\lambda_n, \lambda^\perp)$. Then since λ_n is a limit root, there is a projective root ρ_n with $d(\rho_n, \lambda_n) < \varepsilon$. Thus there is a sequence $\{\rho_n\} \to \lambda$ so that λ is not visible from any ρ_n.

Example 7.6.3. Figure 7.4 illustrates the geometric objects introduced in this section. It shows the projectivized root system for the Coxeter diagram in the figure, with roots of depth up to five. The following properties are visible. The page serves as the plane V_1. All of the action occurs in the convex hull of the simple system $\{\alpha, \beta, \gamma\}$, which is precisely because all of the positive roots are within $\mathsf{cone}(\Delta)$, and the normalization map makes it so that all roots have coefficient sum 1. In red is $\widehat{Q}$, the projectivized light cone. The plane spanned by $\{\alpha, \beta\}$ becomes a line in projective space, and because the parabolic subgroup generated by these roots is $I_2(5)$, there are precisely five roots along this line. Because the restriction of B to $\{\alpha, \beta\}$ is positive definite, the normalized light cone does not intersect this projective subspace. However, the restriction to the plane $\{\beta, \gamma\}$ is indefinite, and so this intersects $\widehat{Q}$ in two points. There are infinitely many

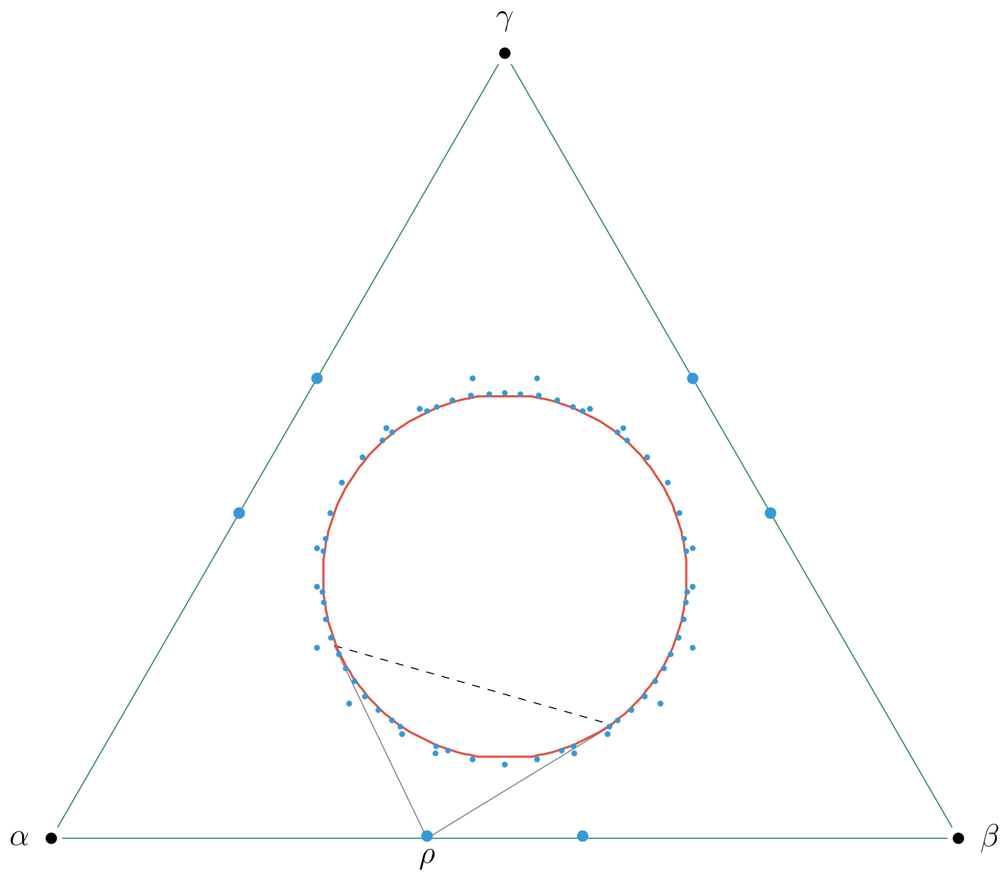

Figure 7.3: The up-set for the normalized root ρ with respect to the visibility order is pictured. The hyperplane associated to ρ is the dashed line.

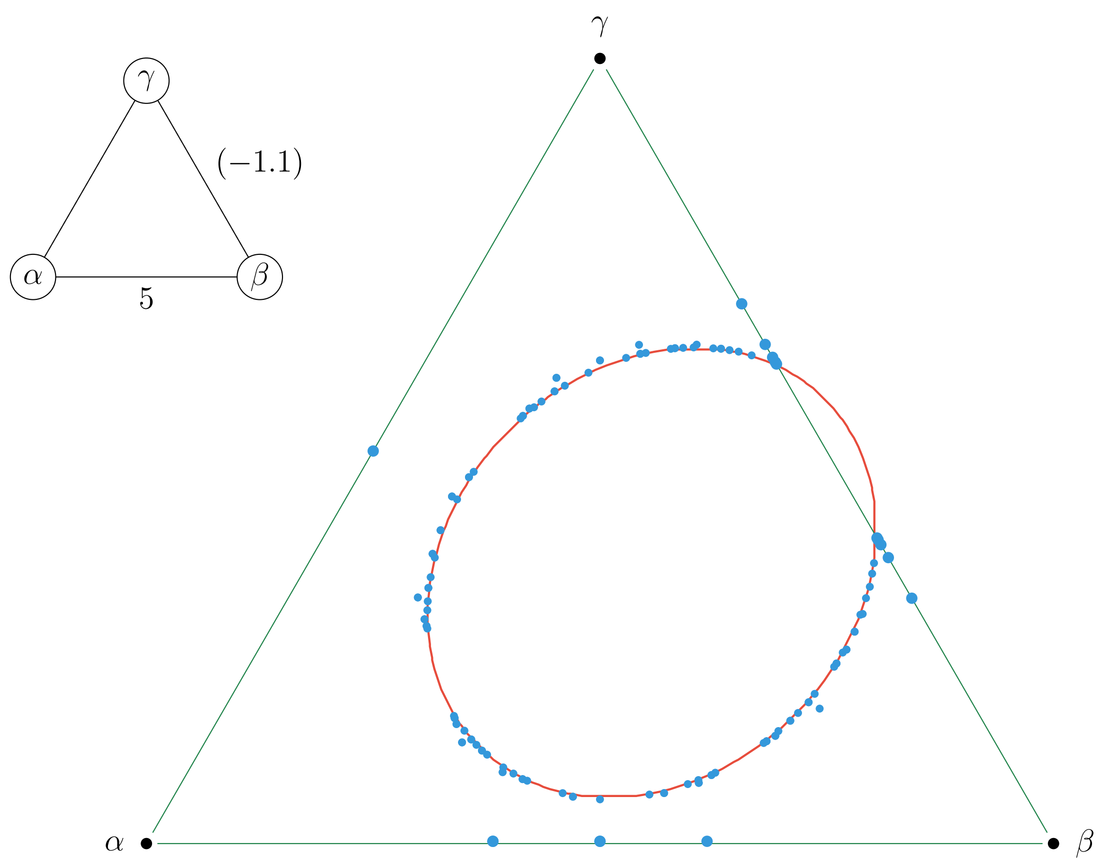

Figure 7.4: The projectivized root system for the generalized Coxeter diagram in the top left.

roots along this line, and two limit roots. Compare this with Figure 7.2. For the same reason, the reflection subgroup between β and the root between α and γ provides us with two additional limit roots. Finally, a dihedral group of order 10 can be seen between the root on the α-γ line and the root closest to γ on the β-γ line. In the rank 3 case, $\widehat{Q}$ is contained in the interior of the convex hull of Δ if and only if all labels are finite. With the standard representation when a label is infinite, $\widehat{Q}$ is tangential to the boundary of $\mathsf{conv}(\Delta)$, but if $B(\alpha_1, \alpha_2) < -1$ for some pair of simple roots, this line intersects the boundary in two points.

8. The Limit Set

Recall the setup in Chapter 3. A proper geodesic metric space X is δ-hyperbolic if every geodesic triangle is δ-thin, and a group is δ-hyperbolic if its Cayley graph has this property. The first section of this chapter develop the theory of hyperbolic groups before applying this theory to the present environment of Lorentzian Coxeter groups.

8.1 Boundary of a δ-Hyperbolic Space

Although there is some sense of a boundary for say $\mathbb{R}^n$, it does not have very strong "rigidity properties" that are quasi-isometry invariant. However, a quasi-isometry between δ-hyperbolic spaces induces a homeomorphism between their boundaries, which gives the boundary a very important role in hyperbolic geometry. Many of these results lie in the theory of "Mostow rigidity," in which the pioneering result is that the metric of a certain class of hyperbolic manifolds is determined completely by its fundamental group.

Definition 8.1.1 (Boundaries)**.** Let (X, d) be a (proper geodesic) δ-hyperbolic space, and $x \in X$. Recall that for $x, y, z \in X$, the Gromov product is $(y, z)_x = \frac{1}{2}(d(x, y) + d(x, z) - d(y, z))$. Say a sequence $\{x_n\}$ converges to ∞ if $\liminf_{i,j\to\infty}(x_i, x_j)_x = \infty$. Call two sequences $\{x_n\}$ and $\{y_n\}$ equivalent if $\liminf_{i,j\to\infty}(x_i, y_j)_x = \infty$. It can be shown that this relation is independent of the choice of x in the Gromov product.

The *(Gromov) boundary of* X, denoted ∂X, is the set of equivalence classes of sequences converging to ∞. If $p \in \partial X$, let $B_r(p) = \{q \in \partial X \mid \liminf_{i,j\to\infty}(p_i, q_j) \geq r$ for some representative sequences of $p, q\}$.

Taking this collection of sets as a basis for open sets determines a topology, and can be combined with the topology on X to obtain a compactification $\bar{X} = X \cup \partial X$ of X.

The boundary is a quasi-isometry invariant. A quasi-isometry from X to Y extends to a homeomorphism from ∂X to ∂Y.

An *end* of X is a connected component of the boundary. Another way to define the ends of a (hemicompact) topological space is as follows. Let $K_1 \subseteq K_2 \subseteq \ldots$ be a sequence of compact sets whose interiors cover X. Further choose nested connected components of their complements, $U_1 \supseteq U_2 \supseteq \ldots$ where U_i is a connected component of $X \setminus K_i$. Each such sequence of nonempty connected components is an end of X. It may appear that this is dependent on the choice of $K_1 \subseteq K_2 \subseteq \ldots$, but it can be shown that if $L_1 \subseteq L_2 \subseteq \ldots$ is another sequence with the same properties, then there is a natural bijection between the ends with respect to the K_i and the ends with respect to the L_i.

Applying these definitions to a δ-hyperbolic group W allows one to speak of *the boundary of* W. As an interesting side note, Bass-Serre theory answers some questions about what the ends of a (not necessarily hyperbolic) finitely generated group can look like.

Theorem 8.1.2 (Freudenthal-Hopf). *The number of ends of a finitely generated group G is 0, 1, 2, or ∞.*

Sketch of proof. See Theorem 11.27 in [Mei]. The idea is that if the Cayley graph of G has $k \geq 3$ ends, then it also has at least $2k-2$ ends, so it also has at least $2(2k-2)-2 = 4k-6$ ends, and thus it has infinitely many ends. This is done by using some group element g to translate the identity sufficiently far from itself such that a single end now splits into an additional $k-1$ ends.

□

Remark. This theorem was first proven in the 1930's, but was proven by Stallings using geometric group theory techniques in 1968.

The only case in which a group has more than 1 end is when G splits as an HNN extension over finite subgroups, or is a free product amalgamated over a finite group, but it is possible for a group to have all of 0, 1, 2 or infinitely many ends. A group has no end if and only if it is finite. The group $\mathbb{Z}$ has two ends, and any group with two ends has a finite-index subgroup isomorphic to $\mathbb{Z}$. When $n > 1$, $\mathbb{Z}^n$ acts geometrically on $\mathbb{R}^n$, whose boundary is connected and homeomorphic to an $n-1$-sphere. A free group of rank $n \geq 2$ acts on a $2n$-regular tree which has infinitely many ends.

Cannon's conjecture can also be formulated with the language developed.

Conjecture 8.1.3 (Cannon)**.** If G is a hyperbolic group such that ∂G is homeomorphic to a 2-sphere, then G acts geometrically on hyperbolic 3-space.

8.2 Infinite Reduced Words

This section develops a more combinatorial notion of a limit of W.

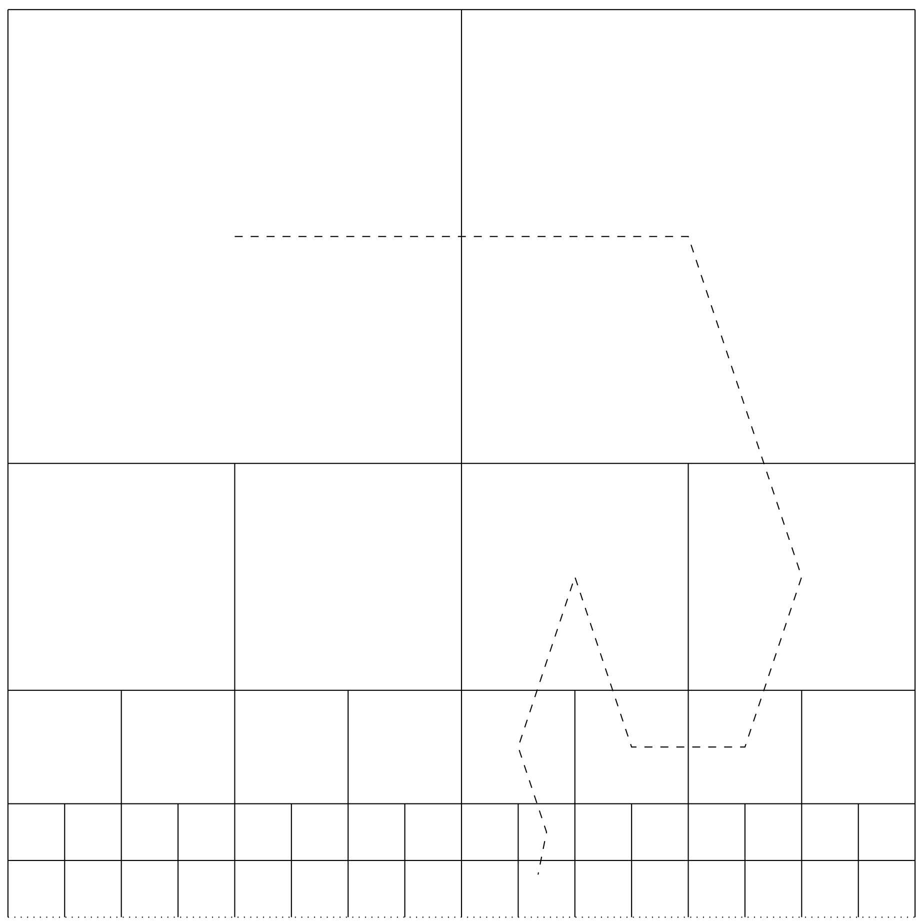

Figure 8.1: This can be viewed as a combinatorial version of the upper half model of hyperbolic space. A “random walk” is shown to tend towards the boundary of hyperbolic space. In a sense, a random walk has equal probability of moving to any of the five adjacent blocks. However, since two of these five adjacent blocks are below a given block, the generic move will be downwards.

Definition 8.2.1 (Infinite reduced words)**.** Let (W, S) be a Coxeter system. Call a sequence $\mathbf{w} = s_1 s_2 \cdots \in S^{\mathbb{N}}$ *reduced* if every initial segment is a reduced word in W; denote the collection of such sequences by $\mathcal{W}$. Recall the *(left) inversion set* $\mathsf{inv}(w) = \Phi^+ \cap w(\Phi^-)$ is the set of positive roots which are sent to negative roots by (left action of) $w \in W$. This can be extended to a reduced word $\mathbf{w}$ by $\mathsf{inv}(\mathbf{w}) = \{s_1 \ldots s_{k-1}(\alpha_{s_k})\}_{k \in \mathbb{N}} \subseteq \Phi^+$.

It is possible to endow the collection of reduced words with a pre-order by setting $\mathbf{w} \preceq \mathbf{w}'$ if every initial segment of $\mathbf{w}$ is below some initial segment of $\mathbf{w}'$ in the right weak order for W. Define an equivalence relation on $\mathcal{W}$ by setting $\mathbf{w} \simeq \mathbf{w}'$ when $\mathbf{w} \preceq \mathbf{w}'$ and $\mathbf{w}' \preceq \mathbf{w}$, then denote by W^∞ the corresponding quotient. This formal development makes precise the vague notion of infinite reduced words being equivalent via an infinite number of braid moves.

The weak order may be extended to $\overline{W} = W \cup W^\infty$ by setting $w \leq \mathbf{w}$ if a reduced expression for w is the initial segment of some infinite reduced word in the equivalence class of $\mathbf{w}$. Also, there is a left action of W on $\overline{W}$, given by preceding a sequence in W^∞ with a reduced word for the element of W, or just left multiplication in the case of an element of W.

Example 8.2.2. For example, in $\widetilde{A}_1$, the sequence $\mathbf{w}_1 = s_1 s_2 s_1 s_2 s_1 \ldots$ is reduced, and it is inequivalent to $\mathbf{w}_2 = s_2 s_1 s_2 s_1 s_2 \ldots$. The W action on W^∞ gives $s_1 \mathbf{w}_1 = s_2 \mathbf{w}_1 = \mathbf{w}_2$, and in fact, $w \in W$ swaps the elements of W^∞ if and only if $\ell(w)$ is odd. These two infinite reduced words correspond to distinct inversion sets. In the context of Figure 7.2, $\mathbf{w}_1$ corresponds to the roots to the left of Q, and $\mathbf{w}_2$ corresponds to the roots on the right.

Since the inversion sets $\mathsf{inv}(\mathbf{w}_1)$ and $\mathsf{inv}(\mathbf{w}_2)$ are infinite, they have limit points, which are limit roots. These limit roots are distinct when the representation is Lorentzian, but if the representation is the standard geometric one, they coincide. In a finite Coxeter group, there are no infinite reduced words.

8.3 Limit Set of W

There is yet another notion of boundary to be considered. A Lorentzian Coxeter system (W, S) has an action on projective hyperbolic space, and the orbit Wx of a point $x \in \mathbb{H}_p$ is an infinite subset of $\mathbb{H}_p$. When viewed within $\mathsf{conv}(\Delta)$, one can consider the set of limit points of the set $Wx \subseteq \mathbb{H}_p$.

Definition 8.3.1 (Limit sets)**.** Let (W, S) be a Lorentzian Coxeter system, and let $\mathbb{H}_p^n$ be the associated projective model of hyperbolic space. Let $x \in \mathbb{H}_p^n$. Then $\mathsf{Lim}(Wx)$ is the limit set of the orbit of x under W.

Proposition 8.3.2. *Notation as above, if $x, y \in \mathbb{H}_p^n$, then $\mathsf{Lim}(Wx) = \mathsf{Lim}(Wy)$. Thus there is a well-defined notion of $\mathsf{Lim}(W)$.*

Sketch of proof. Let $\| \cdot \|_\Delta$ denote the Euclidean norm obtained by declaring Δ an orthonormal basis. Since W acts discretely on $\mathbb{H}_p^n$ with respect to the hyperbolic metric d_H, any convergent sequence of points $\{w_i x\}$ in Wx must converge to a point on the boundary $\widehat{Q}$. We now observe that $d_{\mathbb{H}}(w_i x, w_i y) = d_{\mathbb{H}}(x, y)$ for every w_i. But if $w_i x$ goes to the boundary, then $\|w_i x - w_i y\| \to 0$. $\square$

No full proofs are found in this section, but results are outlined and proofs are sketched. The first result extracts a subset of $\mathsf{Lim}(\Phi)$ which is easier to calculate and understand, and the second result provides the precise set of limit roots for a subclass of Coxeter systems.

8.4 Dihedral Subgroups, Revisited

There is a much easier object to calculate if one wishes to find the set of limit roots. The set of vectors in Q that lie between some simple root and some (projective) positive root is in fact dense in $\mathsf{Lim}(\Phi)$. Consider the set

$$D = \bigcup_{\alpha \in \Delta} \bigcup_{\rho \in \widehat{\Phi}} ([\alpha, \rho] \cap \widehat{Q}).$$

Since $\mathsf{Lim}(\Phi)$ is closed, it suffices to show that D is dense in $\mathsf{Lim}(\Phi)$. So take $\lambda \in \mathsf{Lim}(\Phi)$. By assumption, there is a sequence of normalized roots $\{\rho_n\}_{n \in \mathbb{N}}$ converging to λ. The idea is that, for large enough N, these roots are close enough to Q such that some simple root α is "on the other side of Q" for $n \geq N$. This then allows us to find a sequence within D converging to λ by considering the intersection point of $[\alpha, \rho_n]$ with Q which is closer to λ. A proof that this is possible is in [HLR11], Section 4.2 and Section 4.3.

8.5 The Limit Roots and the Limit Set Coincide

Since there is already a lot known about the boundary of a hyperbolic group G, or the limit set of G, the final theorem in this thesis is perhaps the most informative. It shows that this new idea of studying limit roots is nothing more than studying the limit set of the Coxeter group. However, having both concepts allows the use of different tools in studying either objects.

Theorem 8.5.1. *Let (W, S) be a Lorentzian Coxeter system, with root system Φ. Then $\mathsf{Lim}(W) = \mathsf{Lim}(\Phi)$.*

Sketch of proof. The proof supplied in the paper [HPR] by Hohlweg, Preaux, and Ripoll relies on an object referred to as the *imaginary convex set*, whose main properties are established in the paper [DHR13] by Dyer, Hohlweg, and Ripoll. The imaginary convex set is the W-orbit of K, the vectors $v \in \mathsf{conv}(\Delta)$ for which $B(\alpha, v) \leq 0$ for every $\alpha \in \Delta$. The set K is a polytope with nonempty interior, and the closure of $W(K)$ is realized as the convex hull of the limit roots $\mathsf{Lim}(\Phi)$. It is shown in [DHR13] that $\mathsf{Lim}(\Phi)$ is the set of limit points of Wz, for $z \in W(K)$, and so we conclude that $\mathsf{Lim}(\Phi) = \mathsf{Lim}(W)$.

□

In some sense, this proof is unsatisfying. It relies on this general theory about the imaginary convex set developed in [DHR13] which does not say much about hyperbolic geometry. The authors state that they "do not know a direct proof of this statement using only tools from hyperbolic geometry." It seems that such a proof should be possible with all of the different geometric notions of boundaries of hyperbolic space that have been developed in the last few decades.

A strategy to pursue. Let λ be a limit root. By 7.6.2, there is an sequence $\{\rho_n\}$ of projective roots which is increasing with respect to the visibility order. Corollary 2.11 in the paper [HLR11] by Hohlweg, Labbé, and Ripoll says that for all $\varepsilon > 0$, there are only finitely many roots distance at least ε from $\widehat{Q}$. This implies that the interval $[\rho_n, \rho_{n+1}]$ is finite, so it is possible to refine this sequence to one in which (a) $\rho_1 \in \Delta$, and (b) $\rho_n \lhd \rho_{n+1}$ for each n (meaning that $\rho_n < \rho_{n+1}$, and that there is no other root $\rho \in (\rho_n, \rho_{n+1})$. From this, we (conjecturally) obtain a reduced word $\mathbf{w}$ such that $\mathsf{Lim}(\mathsf{inv}(\mathbf{w})) = \{\lambda\}$.

Now we change perspectives. With this reduced word $\mathbf{w}$, we consider the sequence $\{w_i x\}_{i\in\mathbb{N}} \subseteq \mathbb{H}$, and conjecturally claim that it converges to λ.

We also know that since the sequence ρ_n is strictly increasing with respect to the visibility order, we obtain a sequence of nested half-spaces in $\mathbb{H}$. The "two wrongs make a right" principle instructs that $\mathbf{w}$ follows a path within the nested half-spaces in $\mathbb{H}$. $\square$

Another strategy to pursue. Let $w \in W$ be an infinite order element. Then Perron-Frobenius theory implies that the matrix $\sigma(w)$ afforded by the representation has a largest eigenvalue strictly greater than 1. Since it has determinant 1, it must have some eigenvalue with modulus strictly less than 1. Moreover, these eigenspaces have dimension 1. Projectively, a 1-dimensional eigenspace is just a point. Since $w \in O_B(V)$, these eigenspaces must be contained in Q. Thus we have special points, say w^+ and w^- on the boundary of hyperbolic space which correspond to an eigenvalue with modulus larger than 1 and smaller than 1, respectively.

This induces so-called *North-South dynamics* on $\mathbb{H}^n$: if U^+ and U^- are open sets in $\partial\mathbb{H}^n$ containing q^+ and q^- respectively, then there is an m so that $w^m(\partial\mathbb{H}^n \setminus U^-) \subseteq U^+$. This implies that q^+ is an element of the limit set of W.

It seems feasible that, using similar techniques, one could find a reason that the same point q^+ is a limit root. $\square$

This concludes the thesis. There is much more known in the theory of limit roots than is outlined here, most of which can be found in the references. One such result is a characterization of which Coxeter systems have every point in $\widehat{Q}$ as a limit root. For the free Coxeter group of rank 4, the limit roots resemble an Appolonian gasket, and more general fractal properties are explored. The articles in the references also contain many more conjectures and directions for further study. Most of the unresolved conjectures are regarding the fractal properties that are especially evident in the free Coxeter group example.

Most of the results and conjectures among these papers are ultimately motivated by the figures computed by Jean-Philippe Labbé, who graciously shared his code which enabled the production of Figures 7.1, 7.3, and 7.4 in this thesis.

Bibliography

[BB06] Anders Bjorner and Francesco Brenti, *Combinatorics of coxeter groups*, vol. 231, Springer Science & Business Media, 2006.

[BH99] Martin R Bridson and André Haefliger, *Metric spaces of non-positive curvature*, vol. 319, Springer Science & Business Media, 1999.

[BM15] Noel Brady and Jon McCammond, *Factoring euclidean isometries*, International Journal of Algebra and Computation **25** (2015), no. 01n02, 325–347.

[BMMN02] Noel Brady, Jonathan P McCammond, Bernhard Mühlherr, and Walter D Neumann, *Rigidity of coxeter groups and artin groups*, Geometriae Dedicata **94** (2002), no. 1, 91–109.

[Bou08] Nicolas Bourbaki, *Lie groups and lie algebras: chapters 7-9*, vol. 7, Springer Science & Business Media, 2008.

[Bow06] Brian H Bowditch, *A course on geometric group theory.*

[Bri00] R. Martin Bridson, *Finiteness properties for subgroups of gl*$(n, \mathbb{Z})$, Mathematische Annalen **317** (2000), no. 4, 629–633.

[CFK+97] James W Cannon, William J Floyd, Richard Kenyon, Walter R Parry, et al., *Hyperbolic geometry*, Flavors of geometry **31** (1997), 59–115.

[CL15] Hao Chen and Jean-Philippe Labbé, *Lorentzian coxeter systems and boyd–maxwell ball packings*, Geometriae Dedicata **174** (2015), no. 1, 43–73.

[Cox35] HSM Coxeter, *The complete enumeration of finite groups of the form* $r_i^2 = (r_i r_j)^{k_{i,j}} = 1$, Journal of the London Mathematical Society **1** (1935), no. 1, 21–25.

[Dav08] Michael Davis, *The geometry and topology of coxeter groups*, vol. 32, Princeton University Press, 2008.

[DHR13] Matthew Dyer, Christophe Hohlweg, and Vivien Ripoll, *Imaginary cones and limit roots of infinite coxeter groups*, arXiv preprint arXiv:1303.6710 (2013).

[HL15] Christophe Hohlweg and Jean-Philippe Labbé, *On inversion sets and the weak order in coxeter groups*, arXiv preprint arXiv:1502.06926 (2015).

[HLR11] Christophe Hohlweg, Jean-Philippe Labbé, and Vivien Ripoll, *Asymptotical behaviour of roots of infinite coxeter groups*, arXiv preprint arXiv:1112.5415 (2011).

[HPR] Christophe Hohlweg, Jean-Philippe Praux, and Vivien Ripoll, *On the limit set of root systems of coxeter groups acting on lorentzian spaces.*

[Hum92] James E. Humphreys, *Reflection groups and coxeter groups*, Cambridge University Press, 1992.

[KB02] Ilya Kapovich and Nadia Benakli, *Boundaries of hyperbolic groups*, Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001) **296** (2002), 39–93.

[Kra09] Daan Krammer, *The conjugacy problem for coxeter groups*, Groups, Geometry, and Dynamics **3** (2009), no. 1, 71–171.

[LT15] Thomas Lam and Anne Thomas, *Infinite reduced words and the tits boundary of a coxeter group*, International Mathematics Research Notices **2015** (2015), no. 17, 7690–7733.

[Mei] John Meier, *Groups, graphs and trees: An introduction to the geometry of infinite groups, publisher = Cambridge University Press, year = 2008, isbn = 0521895456,*.

[Rat06] John Ratcliffe, *Foundations of hyperbolic manifolds*, Springer, 2006.