Limit Roots of Lorentzian Coxeter Systems

Nic Brody

December 7, 2015

Abstract

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 discrete group generated by such transformations. 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 reflection group.

A Coxeter group is a generalization of a finite reflection group, whose have 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. To each group is an associated symmetric bilinear form which it preserves, and the signature of the bilinear form contains valuable information about W; if 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 V, but if we projectivize, $\mathbb{P}V$ contains limit points, which have the appearance of a fractal.

Contents

Introduction

In this thesis, we aim to give an exposition of the nascent theory of limit roots. We assume no prior knowledge of Coxeter groups, and introduce all relevant definitions and concepts. We take detours along the way to explore properties of Coxeter groups which may not be explicitly involved in questions about limit roots, but will at the very least help the reader gain an appreciation for the wide-ranging and diverse theory of Coxeter groups. The thesis is structured as follows. The first part of the paper is simply background. We provide many definitions and results that will be necessary to understand the main points of interest. In the second part, we define Coxeter groups and study them in general. In the third part, we study the specific case of Lorentzian Coxeter groups.

Part 1 of the paper begins with a short introduction to the basics of geometric group theory to give the reader a sense of the connections between the group theory and geometry. We build up to the Schwarz-Milnor Lemma, which tells us that for every finitely generated group, there is an essentially unique geometric space it can act on in a certain way. Moreover, it gives us a way to find this class of spaces. In the next section, 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.

After setting the stage in this way, we begin to study reflection groups in Part II. We first focus on the finite case since it motivates the general definition, which may seem rather obscure at first glance. In this section 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 previous theory 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 paper: how are the roots of a Lorentzian Coxeter group distributed among our vector space? This problem has caught the attention of a number of mathematicians recently; in the last five years, Dyer, Hohlweg, Ripoll, Labbé, and Chen have studied this question.

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 paper. We prove the basic theorems that will be crucial to understanding Coxeter groups, but nothing in this section lies within the domain of Coxeter theory.

1 Geometric Group Theory

We begin with some basic geometric group theory. As mentioned, the connections between geometry and group theory are the focus of this paper, so we take some time to introduce some relevant concepts. Specifically, we introduce geometric group actions and prove the Schwarz-Milnor lemma.

1.1 Group Actions

Group actions come up in a variety of settings; historically, groups were studied to understand symmetries of objects, and not as objects themselves. Here, we define some concepts for a group acting a metric space.

Definition 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 equal to $\{y \in X \mid d(x, y) \leq r\}$. For a subset S of X, we can define the *(open) ball of radius r about S* to be $B_r(S) = \cup_{x \in S} B_r(x)$. We 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. 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 $x, y \in [a, b]$. Finally, we 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 $\gamma(d(x, y)) = y$.

Example 1.2. A subset of $\mathbb{R}^n$ with the subspace metric is a geodesic metric space if and only if it is convex. For example, $X = \mathbb{R} \setminus \{0\}$ is not, since there is no curve $\gamma \colon [0, 2] \to X$ with $\gamma(0) = -1$ and $\gamma(2) = 1$. Some of the time, we can remedy this by endowing a connected subset of $\mathbb{R}^n$ with the *induced intrinsic metric*, in which $d'(x, y) = \inf\{\ell(\gamma) \mid \gamma \text{ is a curve from } x \text{ to } y\}$, where $\ell(\gamma)$ is the length of the curve. However, this process does not turn $\mathbb{R}^n \setminus \{0\}$ into a geodesic metric space, for example.

Definition 1.3 (Group Actions)**.** A *group action* of a group G on a mathematical structure X is a homomorphism ϕ from G to $\mathrm{Aut}(X)$, the group of structure-preserving maps from X to itself. For example, if X is a topological space, $\mathrm{Aut}(X)$ consists of homeomorphisms from X to X; if X is a vector space, $\mathrm{Aut}(X)$ consists of invertible linear transformations from X to X (more commonly known as $GL(X)$). 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 1.4 (Quotients)**.** Suppose a group G acts on a topological space X. We can 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 topologize with the final topology from q. That is, $U \subseteq X/G$ is open if and only if $q^{-1}(U)$ is open in X.

Definition 1.5 (Geometric group action)**.** Let (X, d) be a proper, geodesic metric space and G a finitely generated group. We say the group G acts on the space X by *isometries* if for every $g \in G$ and $x, y \in X$, we have $d(g.x, g.y) = d(x, y)$. We say a group G acts *cocompactly* if X/G, the quotient space of X induced by the action, is compact as a topological space. We call an action *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).

If G acts on X by isometries, cocompactly, and properly discontinuously, we say G acts on X *geometrically.*

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

It will also be useful to note that as long as X is a proper metric space (and G is a finitely generated group?), the following condition is equivalent to proper discontinuity: For every pair $x, y \in X$, there exists an $r > 0$ so that the set $\{g \in G \mid g.B_r(x) \cap B_r(y) \neq \varnothing\}$ contains finitely many group elements.

To see that our definition implies this, note that the closure of any open ball is compact. For the other direction, choose a compact set K (maybe try to cover it with open balls [how do I determine radii?], take finite subcover and hope that any pair satisfies this condition... probably need finite generation)

Example 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 if $U_x = B_{1/4}(x)$ and $U_y = B_{1/4}(y)$, then $|\{n \in \mathbb{Z} \mid n.U_x \cap U_y \neq \varnothing\}| \leq 1$.

1.2 Cayley Graph

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

Definition 1.7 (Cayley Graph)**.** If S is a finite generating set of a group G not containing the identity and symmetric ($S = S^{-1}$), we can build a graph called the *Cayley graph* using group elements as vertices, and directing an edge from g to gs whenever $s \in S$. That is, let $\Gamma_{(G,S)} = (G, E)$, where $E = \{(g, gs) \mid s \in S\}$. Then, $\Gamma_{(G,S)}$ is a graph which is locally finite since S is finite, connected since S generates G, simple since the identity is not in S, and undirected since S is symmetric.

Definition 1.8 (Metrizing Γ)**.** We can turn $\Gamma_{(G,S)}$ into 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 $i = 1, \ldots, n$; n is the length of the path, and we say γ is a path from $\gamma(0)$ to $\gamma(n)$. If the graph Γ is connected and undirected, we can set $d_S(x, y)$ (or just $d(x, y)$) to be 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 finding elements of S so that $xs_1 \ldots s_k = y$, or equivalently, finding minimal length expressions for $x^{-1}y$. This is a metric space since $d(x, y) = 0$ if and only if $x = y$, $d(x, y) = d(y, x)$ since Γ is undirected, and the triangle inequality is satisfied since a path from x to y followed by a path from y to z is 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.

It is not hard to check that, by identifying each edge in the graph with an interval $[0, 1]$, Γ becomes a proper geodesic metric space (proper since Γ is locally finite, geodesics exist because distances are defined in terms of paths).

We wonder if we can make G act geometrically on Γ; we try left multiplication (recall the vertices are labeled by group elements). It is clear that $d(x, y) = d(g.x, g.y)$ for each g. It also happens that $\Gamma/G \cong S^1$, so this action is compact. And the only $g \in G$ for which $g.B_{1/4}(x)$ intersects $B_{1/4}(y)$ is $g = yx^{-1}$, so indeed this is a geometric action.

Surprisingly, we will see shortly 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 we need a somewhat coarse identification of metric spaces. For example, any finite group acts geometrically on a one point metric space. This correspondence can only see the large-scale properties of either category, so we will choose our equivalence relation acoordingly.

Definition 1.9 (Quasi-isometry)**.** A function $f\colon (X, d_X) \to (Y, d_Y)$ is called a quasi-isometry if there are constants $A \geq 1$ and $B, C \geq 0$ so that for every pair of points x_1, x_2 in X, we have

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

and for every point $y \in Y$, there is a point $x \in X$ so that $d_Y(f(x), y) \leq C$.

In light of Definition 1.1, the final condition can be restated as $\overline{B_C(f(X))} = Y$.

Remark. It is equivalent to define quasi-isometries using four constants $A_1, A_2 \geq 1$ and $B_1, B_2 \geq 0$ and requiring $\frac{1}{A_1} d_X(x_1, x_2) - B_1 \leq d_Y(f(x_1), f(x_2)) \leq A_2 d_X(x_1, x_2) + B_2$. Of course, any map satisfying the first definition satisfies the second, and to show the second implies the first, 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 a symmetric relation.

Proposition 1.10 (Equivalence relation)**.** *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. [MAYBE STILL GIVE EXPLICIT PROOFS] □

At this point, the reader should contemplate what equivalence classes look like here. In turns out 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 $X \times [0,1]$. It is a somewhat coarse equivalence relation, but it turns out to be exactly what we need.

Example 1.11 (Infinite generating set)**.** Returning to our 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 geometric, this Cayley graph is not quasi-isometric 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 we see that the demand that our 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; not requiring S to be finite may give a different quasi-isometry class.

1.4 Schwarz-Milnor lemma

The following fact was observed by Efremovich in 1953, by Albert Schwarz in 1955, and John Milnor in 1968. [Give more history]

Theorem 1.12 (Schwarz-Milnor lemma)**.** *Suppose G (with symmetric generating set S) acts geometrically on a proper geodesic space (X, d_X). Then, for every $x \in X$, the map $f_x\colon \Gamma_{(G,S)} \to X$ defined by $f_x(g) = g.x$ is a quasi-isometry.*

We postpone the proof to the end of the section, since we will need a bit of preparation.

Lemma 1.13 (Diameter)**.** *If (X, d_X) is a compact metric space, there is an R so that $d_X(x, y) \leq R$ for every $x, y \in X$. The infimum of such R is called the* diameter *of X, denoted diam(X).*

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, we actually have $\overline{B_{\mathrm{diam}(X)}(x)} = X$ for every $x \in X$. Moreover, if $R > \mathrm{diam}(X)$, we have $B_R(x) = X$.

Lemma 1.14 (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 $d_S(g,h) \leq A d_{S'}(g,h)$. 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 $d_S(g,h) \leq A d_{S'}(g,h)$. Symmetrically, we can find A' so that $d_{S'}(g,h) \leq A' d_S(g,h)$. □

This tells us that we may choose a generating set as late in the game as we want when we only care about quasi-isometry classes.

Theorem 1.15 (Schwarz-Milnor lemma)**.** *Suppose a finitely generated group G acts geometrically on a proper geodesic space (X, d_X). Then for every $x \in X$ and symmetric generating set S not containing 1_G, the map $f_x \colon \Gamma_{(G,S)} \to X$ defined by $f_x(g) = g.x$ is a quasi-isometry.*

Proof. 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 $A d_S(1,g) - B \leq d_X(x, g.x) \leq C d_S(1,g)$, 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 1.13, choose $r > \operatorname{diam}(X/G)$. Let $K = \overline{B_r(x)}$ and choose $S = \{1 \neq g \in G \mid gK \cap K = \varnothing\}$ (by proper discontinuity, this is finite; let $B = |S|$). Note if $y \in gK \cap K$, then so is $g^{-1}y$, so S is symmetric. We wish to show that S generates G.

We chose r large enough so that the G-translates of K cover X (for example, note the image of K in X/G is the whole space, so its preimage in the quotient map is all of X). So let $g \in G$, and we will show that we can write $g = s_1 \dots s_k$ for some elements in S. [NEED THIS FOR THE NEXT PART]

Let $C = \max\{d_X(x, s.x) \mid s \in S\}$. Then $d_X(x, s.x) \leq C d_S(1,g) = C$ 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 $q(x) = x_0$. The definition of the quotient map tells us that $G.x = q^{-1}(x_0)$. We calculate that

$$X = q^{-1}(X/G) = q^{-1}\left(B_C(x_0)\right) = \bigcup_{y \in q^{-1}(x_0)} B_C(y) = B_C(q^{-1}(x_0)) = B_C(G.x)$$

□

As a corollary, we see that any two spaces which G can act on must be quasi-isometric. We conclude that all we really need to do is look at a Cayley graph. This incredible correspondence between geometry and group theory is a true gem.

2 Linear Algebra

In this section, we will define many of the geometric objects that will be of interest to us. We talk about affine linear algebra, bilinear forms, and reflections.

Definition 2.1 (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, the set of nonnegative 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, we have

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

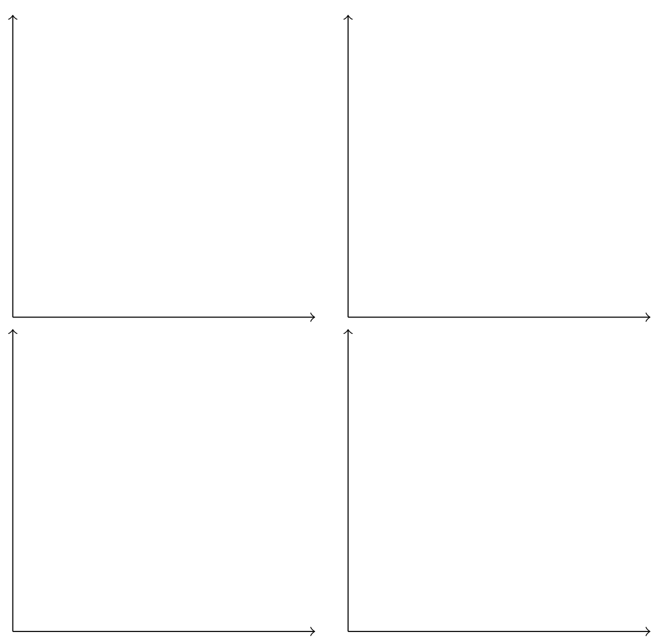

If A and B are (not necessarily finite) subsets of V, their *sum* is

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

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 $-\Phi \stackrel{\text{def}}{=} \{-v \mid v \in \Phi\}$.

2.1 Affine Linear Algebra

We work in the setting of a finite dimensional real vector space V. We will reserve an unspecified use of the word "subspace" to denote a *linear* subspace, and will always say *affine subspace* when we mean to refer to a subspace that does not necessarily contain the origin.

Proposition 2.2. *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 2.3. *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 $\sum_{i=1}^{k} c_i = 1$. □

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.

We also see that an m-dimensional affine subspace contained in another m-dimensional affine subspace implies equality.

Proposition 2.4. *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 $\min\{\dim(A_1), \dim(A_2)\}$. □

Theorem 2.5 (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 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 $1 \leq i \leq n-1$. Since we choose $u \notin W_n = -v + A_n$, we also have $v + u \notin A_n$, so 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. □

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 group 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 we 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 we no longer have an origin.

2.2 Bilinear Forms

We shall now acquaint ourselves with bilinear forms, working up to two main results. The first one will be a crucial fact we use many times, while the second will not be explicitly useful to us, but it will help motivate further work.

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 matrices relative a chosen basis, an endomorphism is a map $V \to V$ and a bilinear form is $V \times V \to \mathbb{R}$, when V is a real vector space. This distinction turns out to be quite consequential.

Definition 2.6 (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 $q(v) \stackrel{\text{def}}{=} B(v, v)$. We say a bilinear form B is *positive-semidefinite* if $q(v) = B(v, v) \geq 0$ for all $v \in V$. We 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, but we will stop saying this now) 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

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

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;

The *general linear group of* V, denoted $GL(V)$, 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:

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

Definition 2.7 (Similarity)**.** When we wish to express an endomorphism $T\colon V \to V$ of a finite dimensional vector space as a matrix A, we must first select an ordered basis $\mathcal{B} = \{v_1, \ldots, v_n\}$, and 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 $\mathcal{C} = \{w_1, \ldots, w_n\}$, 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, we call A, B *similar* if there is an invertible matrix S so that $A = S^{-1}BS$.

Definition 2.8 (Congruence)**.** The situation is different in the setting of bilinear forms $B\colon V \times V \to \mathbb{R}$. Instead, choosing a basis $\mathcal{B} = \{v_1, \ldots, v_n\}$ allows us to construct a matrix A so that $A_{i,j} = B(v_i, v_j)$, and consequently $B(u,w) = u^tAw$, thinking of u and w now being written as column vectors in the coordinates of $\mathcal{B}$. Now we observe that an automorphism S (with matrix T) changes $(S(u))^tA(S(w)) = u^t(T^tAT)w$, and so a change of basis amounts to multiplying on the left by the transpose of the invertible matrix on the right. We call such pairs of matrices A and T^tAT *congruent.*

Definition 2.9 (B_v)**.** For a given bilinear form, each $v \in V$ gives rise to a 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.

Lemma 2.10. *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 us to define the *signature* of a bilinear form.

Theorem 2.11 (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 2.10 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 $v_1, \ldots, v_n$ of $\ker(B_v)$ for which $B|_{v^\perp \times v^\perp}$ is diagonal. Observe that B is diagonal also with respect to $\{v_1, \ldots, v_n, v\}$.

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

2.3 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 $\dim(V)$ reflections. Our main interest is that such an expression of any length exists, but we will prove the strong version.

Theorem 2.12 (Cartan-Dieudonné). *Let B be a bilinear form on an n-dimensional space V over $\mathbb{R}$. If $f \in O_B(V)$, then there exist B-reflections $s_1, \ldots, s_k \in O_B(V)$, with $k \leq n$, so that $s_1 \ldots s_k = f$.*

Proof. We proceed by induction on $n = \dim(V)$. The only B-isometries if $n = 1$ are $f(v) = v$ and $f(v) = -v$. The $n = 2$ case was done above [NOT ABOVE ANYMORE, MAYBE JUST PROVE HERE AND REFER TO LATER]). 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 $f(v) \neq v$; if $f(v) = \alpha v$, then $\alpha = -1$, and so $s_v f$ restricts to an isometry on $v^\perp$, so $s_v f$ admits an expression $s_v f = s_1 \ldots s_k$, where $k \leq n-1$, hence $f = s_1 \ldots s_k s_v$. Otherwise, $H = \text{span}\{v, f(v)\}$ is 2-dimensional, so $f|_H$ is a product of at most 2 reflections, and $f|_{H^\perp}$ is a product of at most $\dim(H^\perp) = n-2$ reflections.

□

3 Hyperbolic Geometry

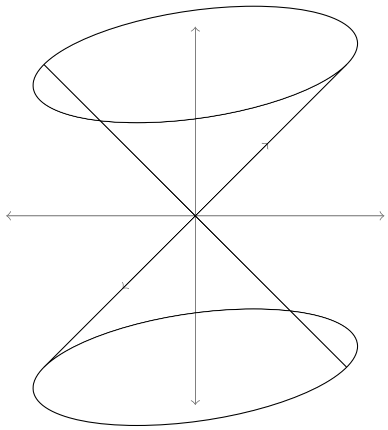

If V is $(n+1)$-dimensional and B is a bilinear form on V of type $(n,1)$, we can impose an interesting geometry on a carefully chosen subset of V. By Sylvester's Law of Inertia (Theorem 2.11), according to some basis $(e_1,\ldots,e_{n+1})$ of V, and writing $v=\sum v_i e_i$, we have

$$B(v,w)=v_1w_1+\cdots+v_nw_n-v_{n+1}w_{n+1}.$$

Letting $q(v)=B(v,v)$, we 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. In fact, we can even put a Riemannian metric on $\mathbb{H}^n$ in the following way. [Fill this in]

Definition 3.1 (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 $\mathbb{H}^n_p=(B_1,d_p)$, the *projective ball model* of hyperbolic space.

As we will prove in Section 3.1, the hyperplanes in the projective ball model can be viewed as intersections of linear subspaces of $\mathbb{R}^{n+1}$ with $\mathbb{H}^n_p$.

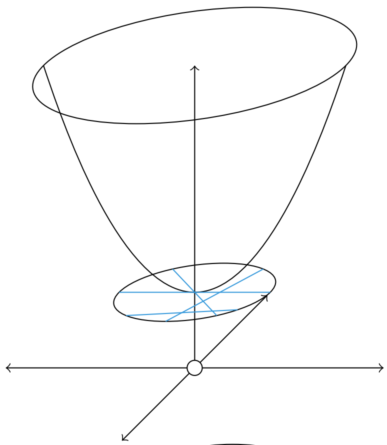

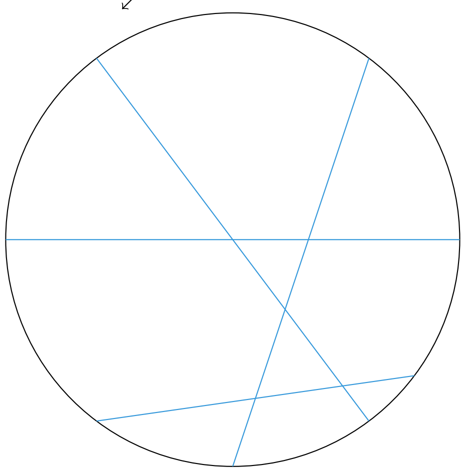

Definition 3.2 (Conformal Ball Model)**.** Let B_0 be the unit n-ball at the origin; that is, $B_0 = \{x \in \mathbb{R}^{n+1} \mid \sum_{i=1}^{n} x_i^2 < 1, x_{n+1} = 0\}$. Define $c\colon \mathbb{H}^n \to B_0$ by setting $c(x_1, \ldots, x_{n+1}) = (x_1, \ldots, x_n, 0)/(1 + x_{n+1})$, and define $d_c(x, y) = d_{\mathbb{H}}(c^{-1}(x), c^{-1}(y))$. Let $H_c^n = (B_0, d_c)$, the *conformal ball model* of hyperbolic space. This can be viewed as the projection of the hyperboloid onto B_0 from $-e_1 = (0, \ldots, 0, -1)$.

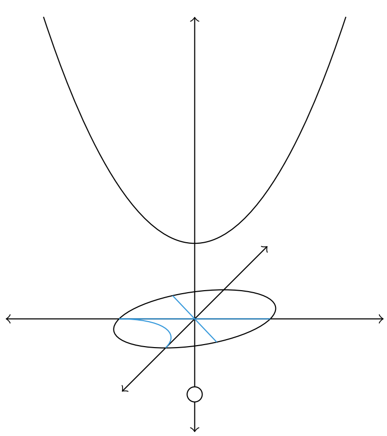

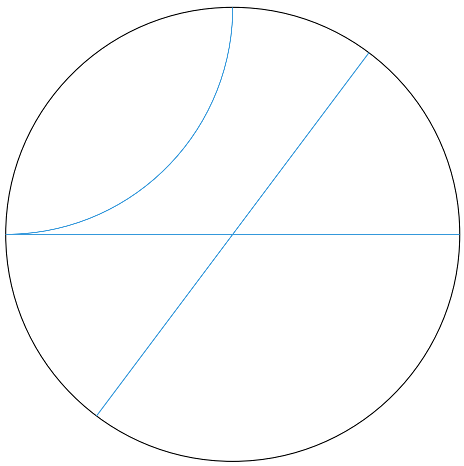

3.1 Hyperbolic Isometries

Definition 3.3 (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, we can choose a basis $\{e_1, \ldots, e_{n+1}\}$ for which $q(v) = v_1^2 + \cdots + v_n^2 - v_{n+1}^2$; we'll call any such basis a *Lorentz basis*. With respect to this specific basis, we call a vector *positive* if $v_{n+1} > 0$ and *negative* if $v_{n+1} < 0$. We call $Q = q^{-1}(0)$ the *light cone*, and a 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 $Q^- = \{v \in V \mid q(v) < 0\}$. These sets are also called the *exterior* and *interior* of the light cone, respectively. We 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.

Proposition 3.4. *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. Note that the linear functional B_v is nonzero, so its kernel is codimension 1. □

Definition 3.5 (Lorentz transformations)**.** When V is a Lorentz space with bilinear form B, we 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. We 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).

3.2 Reflections in Hyperbolic Space

It turns out that reflections in projective model are intersections of hyperplanes, and in the conformal model are intersections of spheres. [RATCLIFFE]

3.3 Hyperbolic Trigonometry

Define functions sinh : $\mathbb{R} \to \mathbb{R}$ and cosh : $\mathbb{R} \to \mathbb{R}$ by $\sinh(x) = \frac{e^x - e^{-x}}{2}$ and $\cosh(x) = \frac{e^x + e^{-x}}{2}$. Observe that sinh is odd, increasing, and surjective, while cosh is even, increasing for $x \geq 0$, and $\cosh(\mathbb{R}) = [1, \infty)$. [What do I need to say for this?]

Part II
Coxeter Groups

In Part II, we introduce the reader to Coxeter groups. We will begin with root systems and reflections [WHAT ELSE]

4 Finite Reflection Groups

Before considering the general case of an arbitrary Coxeter group, let us first explore what is meant by a finite reflection group. This will allow us to motivate the definition of a Coxeter group, and we will ultimately use geometric notions considered here to understand Coxeter groups. We will try to prove things in a general enough way that we can carry the facts to the Coxeter group case.

4.1 From Geometry to Algebra

Definition 4.1 (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 if $k \in \mathbb{R}^*$, we have $s_{k\alpha} = s_\alpha$, so the formula is simplified if we choose $\langle\alpha,\alpha\rangle = 1$.

A finite subgroup of $\mathrm{GL}(V)$ 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 4.2 (Dihedral groups)**.** *If α, β are unit vectors in a vector space V, the subgroup of $GL(V)$ generated by the reflections $\{s_\alpha, s_\beta\}$ is a (possibly infinite) dihedral group.*

Proof. If $\alpha = \pm\beta$, then $s_\alpha s_\beta(v) = v$ for any v. Otherwise, α and β are linearly independent, so let $H = \mathrm{span}\{\alpha,\beta\}$ be the plane spanned by them, and $H^\perp$ its orthogonal complement. Then if $v \in H^\perp$, $s_\alpha s_\beta(v) = s_\alpha(v) = v$. If instead $v \in H$, then $v = a\alpha + b\beta$, and $s_\alpha s_\beta(a\alpha + b\beta) = \left[(2a\langle\alpha,\beta\rangle + \frac{1}{2a})^2 + (b - \frac{4a^3+1}{4a^2})\right]\alpha + (b - 2a\langle\alpha,\beta\rangle)\beta$. [FIX THIS PROOF]

$s_\alpha s_\beta(\beta) = s_\alpha(-\beta) = 2\langle\alpha,\beta\rangle\alpha - \beta$

$s_\alpha s_\beta(\alpha) = s_\alpha(\alpha - 2\langle\alpha,\beta\rangle\beta) = -\alpha - 2\langle\alpha,\beta\rangle s_\alpha(\beta) = -\alpha - 2\langle\alpha,\beta\rangle(\beta - 2\langle\alpha,\beta\rangle\alpha) = (4\langle\alpha,\beta\rangle^2 - 1)\alpha - 2\langle\alpha,\beta\rangle\beta$. These equations don't seem to do much. Maybe just argue informally. Not sure how to just use geometry. □

In other words, the product of any two reflections amounts to a rotation of a plane in a reflection group; to demand this group be 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 4.3 (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, we 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 $s_\alpha^{tr} = s_\alpha^{-1}$. □

Proposition 4.4 (Closure)**.** *Let W be a finite reflection group. Suppose $s_\alpha \in W$, for some vector α, and $w \in W$. Then the reflection in the vector $w(\alpha)$ is also in the group W.*

Proof. To see this, we 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 4.3 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 sections, we look at particular collections of reflecting vectors. We will examine here three types of sets of reflecting vectors and a way to pass between the three types. We'll then discuss how this leads to a group presentation of reflection groups, perhaps a bit informally. At this point, we'll feel sufficiently motivated to define Coxeter groups.

Definition 4.5 (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) $s_\alpha(\Phi) = \Phi$.

The first condition reflects the fact that $s_\alpha = s_{k\alpha}$ for any $k \in \mathbb{R}^*$ (cf. Definition 4.1), and the second is motivated by the fact that reflecting roots are closed under the group action (cf. Proposition 4.4).

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 canonical choice of vector for each reflection. Any vector v in V which is not orthogonal to any vector in Φ (see Theorem 2.5) determines a set $\Phi^+ = \Phi \cap \text{Half}(v)$ which we call a *positive root system*, and clearly Φ is the disjoint union of Φ^+ and $-\Phi^+$.

We call $\Delta \subseteq \Phi$ a *simple root system* if Δ is a basis for span(Φ) and $\Phi^+ \subseteq \text{cone}(\Delta)$ for some choice of Φ^+. We say a root system Φ is *essential* if it spans the vector space.

Proposition 4.6 (Simple systems)**.** *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.*

Proof. Fix a positive root system Φ^+ in a root system Φ, and let Δ be the intersection of all subsets X of Φ for which we can write every element of Φ^+ as a nonnegative linear combination of elements of X. Taking $X = \Phi^+$ allows us to do so, so $\Delta \subseteq \Phi^+$. If there is some distinct pair of vectors $\alpha, \beta \in \Delta$ with $\langle \alpha, \beta \rangle > 0$, then neither $s_\alpha(\beta)$ nor $-s_\alpha(\beta)$ can lie in cone(α, β). This forces $\langle \alpha, \beta \rangle \leq 0$ for every distinct pair of roots.

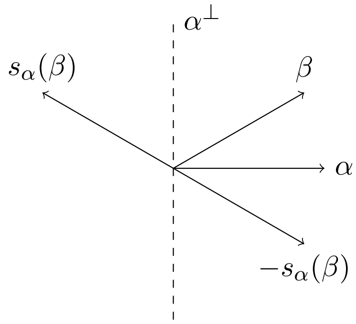

However, this condition on the inner product forces Δ to be a basis for span(Φ); indeed, if we had a dependence relation $\sum a_i \alpha_i = 0$ with the $\alpha_i \in \Delta$ we could move the vectors with negative coefficients to the other side to obtain a vector v which possesses two distinct expressions 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 indeed Δ is linearly independent. □

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

Theorem 4.7 (Correspondence)**.** *Let W be a finite reflection group associated to a root system Φ. Every positive system contains a unique simple system, and each simple system is contained in a unique positive system. 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^+; we want to show that there is a $w \in W$ with $w\Phi_1^+ = \Phi_2^+$. We first note that 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^+. □

Theorem 4.8 (Generated by simple system)**.** *Proof.* □

4.3 A Presentation

In this section, we convince the reader that we can write down a presentation for a given finite reflection group. To prove this rigorously requires too much development which would need to be repeated in the next section, so we rely on geometric intuition, and we'll get a more formal and complete (but less memorable) proof as a byproduct of our more general work in Section 5.

Let W be a finite reflection group on a vector space V with a simple system Δ. First we define the *(open) simplicial cone* C to be the intersection of the positive half spaces determined by Δ; so $C = \bigcap_{\alpha\in\Delta} \mathrm{Half}(\alpha)$.

"Drop ball into cone, draw edges from center to hyperplanes." By reflecting the ball around the hyperplanes, we obtain the 1-skeleton of a polytope, 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 $(s_is_j)^{m_{ij}}$

It is clear that all relations in the presentation hold; we need to show that these generate all of the relations in the group W. Suppose we have a relation $s_1 \dots s_r = 1$.

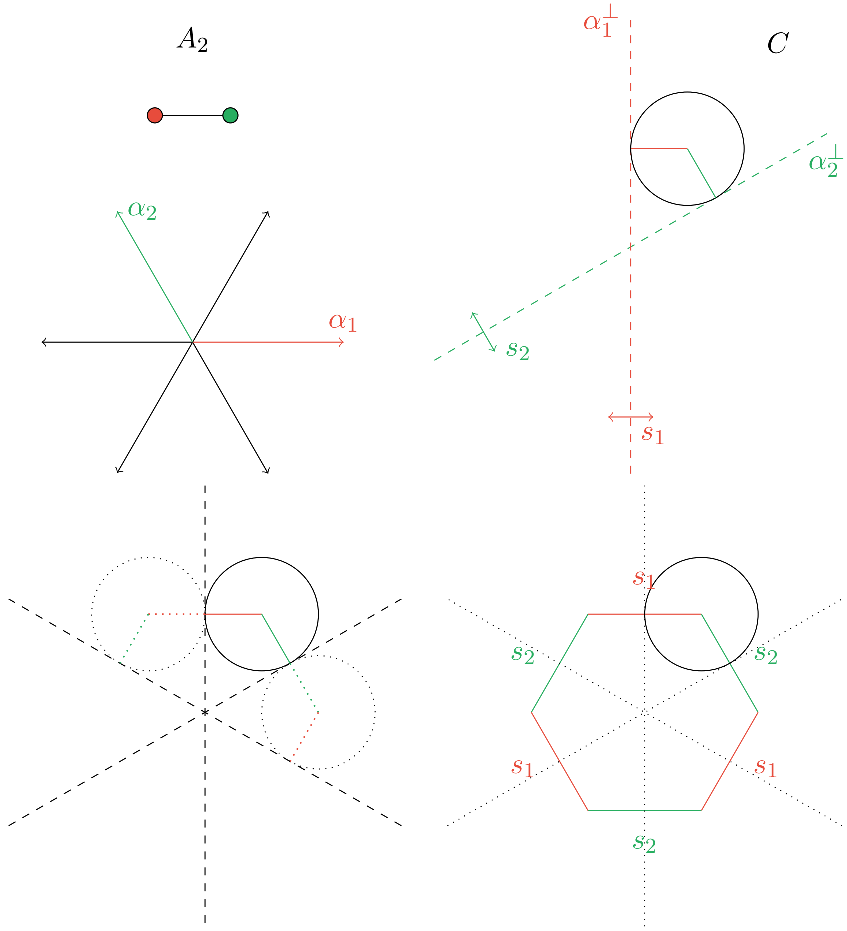

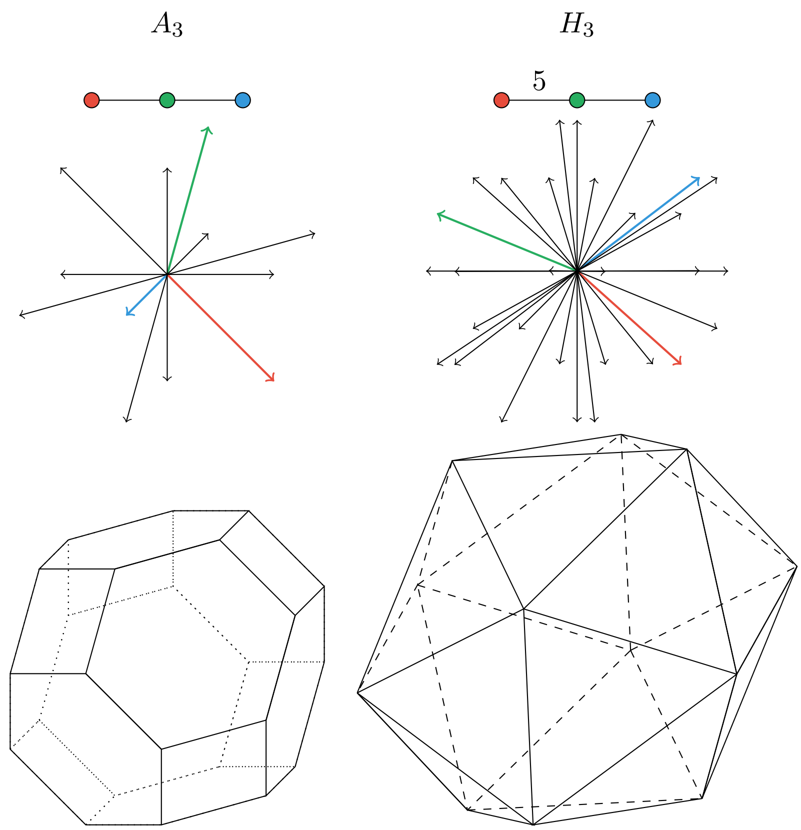

5 Coxeter Groups

We have demonstrated that a class of interesting groups possesses a certain presentation. After building up this theory, we run this process in reverse; we began with geometric considerations and made a deduction about its algebraic properties. Now we weaken the algebraic conditions and see what we can get geometrically.

5.1 Defining Coxeter Groups: From Algebra to Geometry

Definition 5.1 (Coxeter systems)**.** We say that a group W together with a generating set $S = \{s_1, \ldots, s_n\}$ is a *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 we impose 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. We will see shortly that we can recover geometric information.

This is a case where the presentation gives exactly what one would expect. In general, a group presentation can be misleading, but in this situation, we have the desirable properties that $s_i = s_j$ implies $i = j$, and that the order of $s_i s_j$ is precisely $m_{i,j}$. One can see that the map which sends each generator in S to 1 in $\mathbb{Z}/2\mathbb{Z}$ has each relation in the kernel and hence extends to a homomorphism 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 we will construct in Section 5.4.

Definition 5.2 (Graphs and matrices)**.** Given a Coxeter system (W, S), we define the *Coxeter matrix* to be the $n \times n$ matrix M with $M_{ij} = m_{i,j}$. We 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 we do not label an edge, implicitly it is labeled with a 3. It is clear that these two objects have all of the information of the Coxeter system, so we can in fact define a Coxeter system by either of these means.

Define also the *Schlafli matrix* $C_{i,j} = -2\cos(\pi/m_{i,j})$. A symmetric matrix corresponds quite directly to a bilinear form, as described in Section **??**. Because the Dynkin diagram and matrix encode identical information, we will use adjectives typically reserved for bilinear forms to describe the graph and vice versa; for example, we might say that a graph is positive definite.

5.2 Combinatorics

There is a fascinating combinatorial theory of Coxeter Groups. We outline the necessary facts 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 three ways to organize this into a partial order, which we discuss in Section 10. Each of these partial orders has an interesting geometric interpretation. Two of them will come into play throughout the paper.

Definition 5.3 (Words in Coxeter groups)**.** Given a set S, let S^* denote the free monoid generated by S (relevant language: S^* consists of *words* in the *alphabet* S). We also embed $S^n \hookrightarrow S^*$ so that $S^* \cong \bigcup_{n \in \mathbb{N}} S^n$, which provides us with 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 in fact 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; we call any such word *reduced.* Beware that we will frequently conflate an element in S^* with its image under ε.

Let $R = \{wsw^{-1} \mid w \in W, s \in S\}$. We call R the set of reflections in W. We analogously have a map $\varepsilon_R \colon R^* \to W$ satisfying $\varepsilon(r_1 \ldots r_k) = r_1 \ldots r_k$; define $\ell_R \colon W \to \mathbb{N}$ to be the minimal length word in R^* having image w under ε_R. The sets S and R have a similar relationship to that of a simple and positive root system.

5.3 The Standard Representation

5.4 The Geometric Representation

We now aim to recover a geometric interpretation of our generalized class of groups, which in the case of a finite Coxeter group coincides with our definition of finite reflection groups.

Definition 5.4 (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 $B(\alpha_i, \alpha_j) = -\cos(\pi/m_{i,j})$, interpreting $-\cos(\pi/\infty) = -1$. Later, we will allow for any value less than -1. 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 $B_v(w) = B(v, w)$. The *radical* of B is the subspace $B^\perp = \{v \in V \mid B_v(w) = 0\ \forall w \in V\}$. When this is trivial, we 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 $H_v = \ker(B_v)$.

Definition 5.5 (Orthogonal group)**.** Define the B*-orthogonal group* $O_B(V) = \{T \in GL(V) \mid B(T(v), T(w)) = B(v, w)\ \forall v, w \in V\}$. 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 $\sigma_\alpha(\alpha) = -\alpha$. A quick calculation shows that $\sigma_\alpha \in O_B(V)$ for every α, and so the group generated by $\{\sigma_\alpha \mid \alpha \in \Phi\}$ is actually a subgroup of $O_B(V)$.

Indeed, write $v = v_0 + v_1$ and $w = w_0 + w_1$ with $v_0, w_0 \in \mathbb{R}\alpha$ and $v_1, w_1 \in H_\alpha$. Then

$$B(\sigma_\alpha(v), \sigma_\alpha(w)) = B(\sigma_\alpha(v_0 + v_1), \sigma_\alpha(w_0 + w_1))$$

$$= B(-v_0 + v_1, -w_0 + w_1) = B(-v_0, -w_0) + B(v_1, w_1) = B(v, w).$$

Proposition 5.6. *We set* $\sigma(s_i) = \sigma_{\alpha_i}$*. Then the order of* $\sigma(s_i s_j)$ *in* $O_B(V)$ *is* $m_{i,j}$*. Thus* σ *is a faithful representation of* W *in* $O_B(V)$*.*

Proof. □

The properties of this bilinear form unsurprisingly have geometric consequences. When it is positive definite, we obtain an action that in some sense (the sense of Section 1.1) is best viewed on a sphere, say the unit ball $\{v \in V \mid B(v, v) = 1\}$ of V. If we are 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. We can view the action 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 we'd like. When we study subgroups in Section 7.1, we will allow any value satisfying $-\cos(\pi/\infty) \leq -1$.

Definition 5.7 (More Geometric Objects)**.** We are interested in a number of subsets of our vector space V constructed in the previous section. First of all, since Δ is a basis for our vector space, we 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. Let $\hat{}: V \setminus V_0 \to V_1$ by $\hat{v} = \frac{v}{\phi(v)}$.

We will want to think about $\hat{\Phi}$, but we need to ensure $\Phi \cap V_0 = \varnothing$. Since $\Phi \subseteq \operatorname{cone}(\Delta) \cup -\operatorname{cone}(\Delta)$, we merely need to confirm that $\operatorname{conv}(\Delta) \cap V_0 = \varnothing$.

Definition 5.8 (Length function)**.** Suppose G is a group which is generated by a set S, and let $w \in G$. Define $\ell(w) = \min\{k \geq 0 \mid s_1 \dots s_k = w \text{ for some } s_i \in S\}$. We call any expression for w of length $\ell(w)$ a *reduced expression*. Obviously, if $s_1 \dots s_k$ is an *expression* for w, then $\ell(w) \leq k$. We'll use this fact below.

Proposition 5.9 (Technical lemma for length function)**.** *Take $w, w' \in W$. We have $\ell(w) = \ell(w^{-1})$, 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. ☐

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

Proof. By the second part of Proposition 5.9, it suffices to show that $\ell(ws) \neq \ell(w)$. (We need to use the alternating stuff or something about orientation). ☐

We will now introduce a geometric interpretation of the length function; it will require some work to show this new function is equivalent to the combinatorial length, but having both descriptions will allow us to relate the geometry to the algebra very smoothly.

Definition 5.10 (Geometric length)**.** Let (W, S) be a Coxeter system, with $R = W^{-1}SW$, the set of reflections. Define a function $n\colon W \to \mathbb{N}$ by

5.5 Deletion and Exchange

We can carry our definitions of reduced words in a group G with generating set S to a more general setting. We then

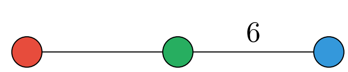

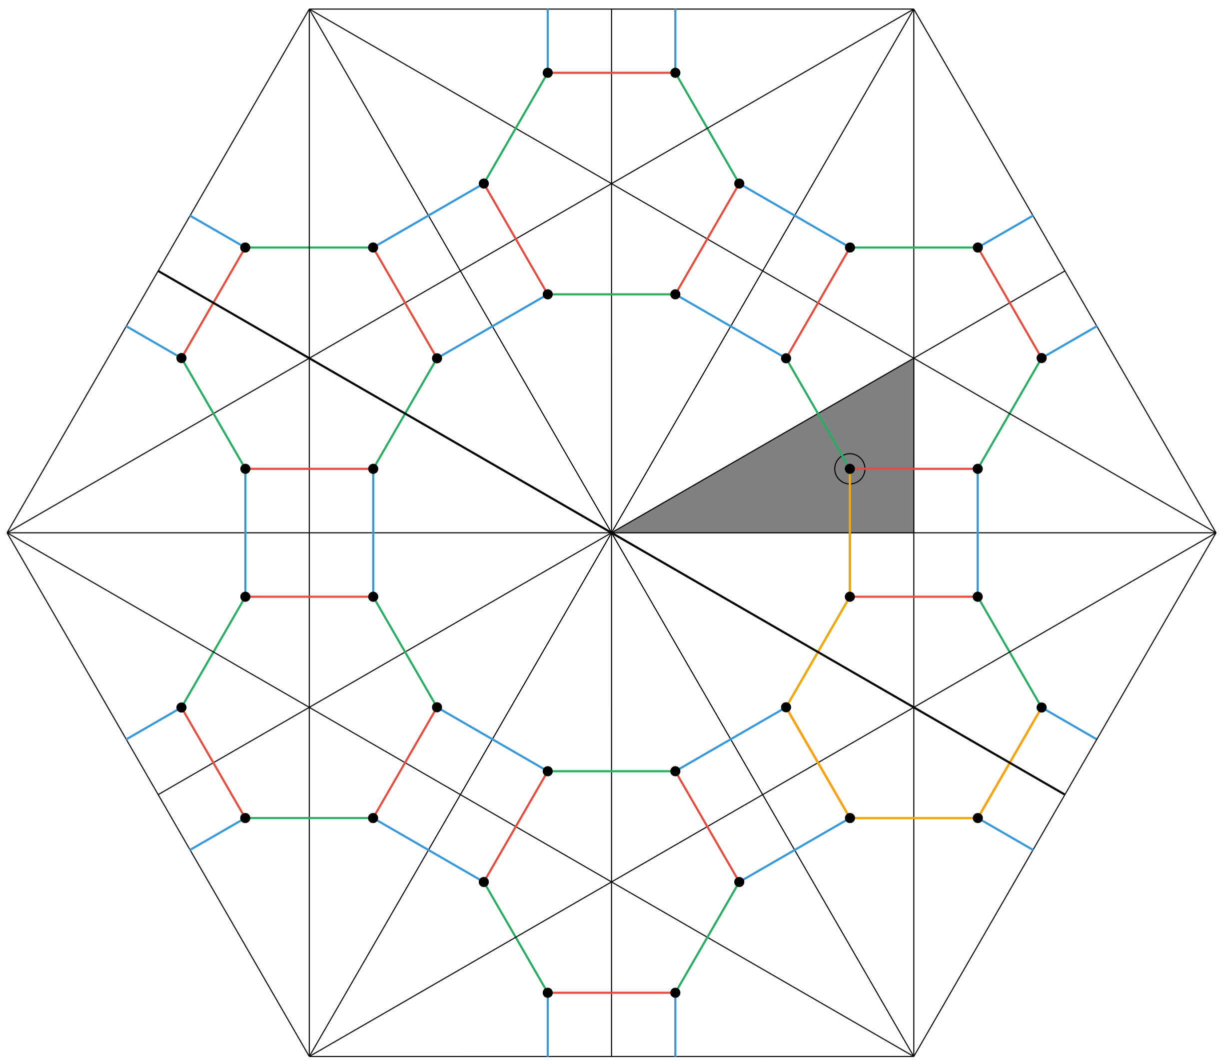

Proposition 5.11 (Deletion Property)**.** *Given an unreduced expression* $w = s_1 \dots s_k$*, there exist* $1 \leq i < j \leq k$ *so that* $w = s_1 \dots \hat{s_i} \dots \hat{s_j} \dots s_k$*.*

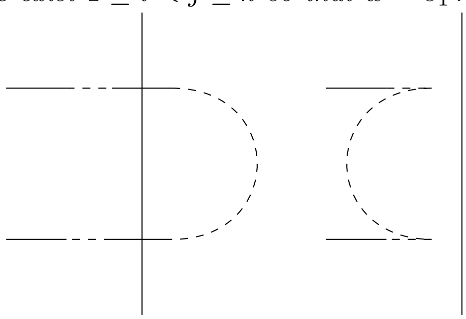

Proof. □

Proposition 5.12 (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* $s_1 \dots \hat{s_i} \dots s_k$*, for* *some $1 \leq i \leq k$, where ˆ denotes a removed letter.*

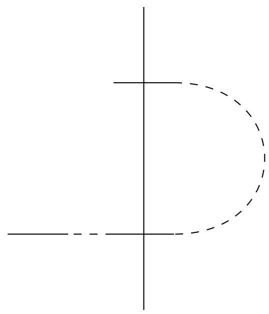

Proof. □

Amazingly, these properties turn out to characterize Coxeter groups in the following sense.

Proposition 5.13 (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. □

Theorem 5.14 (Deletion and Exchange Conditions)**.** *Proof.* □

Given a finite reflection group W on a Euclidean n-space V, let R denote the set of elements of W which are reflections. Each reflection is of the form s_α for some $\alpha \in V$; if for each element in R we pick such an α and its negative, the union of all such will give us a root system Φ (we arranged for the first property to hold by selecting two vectors for each reflection, and the fact that the second property holds is a consequence of Proposition 4.4). From this, we can select a positive root system and hence a simple system $\Delta = \{\alpha_1, \ldots, \alpha_n\}$. To simplify notation, let $s_i = s_{\alpha_i}$.

Theorem 5.15 (A presentation)**.** *With the above notation, we have*

$$W = \langle s_1, \ldots, s_n \mid (s_i s_j)^{m(i,j)} \rangle, \tag{1}$$

for some positive integers $m(i,j) \geq 2$ for $i \neq j$, and each $m(i,i) = 1$.

Proof. To prove this, we suppose $s_{i_1} \ldots s_{i_m}$ is the identity, and [follow Humphreys] □

We now introduce handy orthographic devices which, perhaps surprisingly, contain some important mathematical information. We can encode the information of a finite reflection group by means of either a graph or a matrix.

Definition 5.16 (Graphs and matrices)**.** Given a finite reflection group W, we can select a presentation as in Equation (1). We build a labeled graph $\Gamma = (S, E)$ with vertices corresponding to generators, and an edge between s_i, s_j if and only if $m_{i,j} > 2$, and labeled with this number. Such a graph contains precisely the same information as the presentation. Observe that vertices that are not connected with an edge are exactly those that commute.

5.6 Parabolic Subgroups

Let (W, S) be a Coxeter system. We can ask what happens when we look at a subset I of a simple system S. Let W_I be the group generated by the corresponding simple reflections. If we 1

(W_I, I) is again a Coxeter system.

if graph not connected we have $W \cong W_I \times W_J$

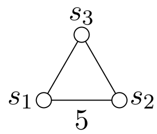

Consider the geometric representation (V, B) of the Coxeter group W associated to the graph above. The parabolic subgroup $W_{\{s_1,s_2\}}$ has as the longest word $s_1s_2s_1s_2s_1$; this corresponds to a root $s_1s_2(\alpha_1) = \frac{1+\sqrt{5}}{2}(\alpha_1 + \alpha_2)$ which we will call β. Then we have $B(\alpha_3, \beta) = \frac{1+\sqrt{5}}{2}(B(\alpha_3, \alpha_1) + B(\alpha_3, \alpha_2)) = -\frac{1+\sqrt{5}}{2} \lneq -1$. However, if we consider the geometric representation (V', B') coming from $W_{\{s_\beta,s_3\}}$, we get $B'(\alpha_3, \beta) = -1$. Thus we have shown that the geometric representation as we have defined it does not restrict to arbitrary reflection subgroups. As we alluded to before, we now have a motivation to allow for $-\cos(\pi/\infty) < -1$.

6 A Four-fold Split

The type of the bilinear form corresponding to a given Coxeter group is a very consequential datum. When we have a 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 we can actually restrict this action to a 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, we can restrict the action of W to an affine subspace of V in which W contains a translation, and so we have a geometric action on a Euclidean space; we 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. We have a bit more control when the form has type $(n - 1, 1)$, but this is not always the case.

6.1 Finite Reflection Groups: Classification

It is not too difficult to classify the Coxeter groups which are finite, and we would be remiss if we did not include this feat first achieved by –. This classification is very closely related to many other classifications in mathematics. The first attempts were – and –, but had these mistakes.

Lemma 6.1. *The Coxeter graph of an irreducible finite reflection group is connected and acyclic, and every subgraph corresponds to a finite reflection group.*

Proof. Section 5.6 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, we also get to suppose the graph is acyclic; so we are actually just looking at labeled trees. Let us see why this is. [Suppose $s_1, \ldots, s_k$ form a cycle, show the group must be infinite.]

The third condition also follows from Section 5.6 □

Theorem 6.2 (Classification of finite reflection groups)**.** *The diagrams listed in Figure 6.1 exhaust the irreducible finite reflection groups.*

Proof. [THIS NEEDS CLEANING UP] We show that these are the only possibilities, and leave the verification that these are indeed finite reflection groups to the reader. 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. If the rank is 2, we can put any natural number m 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 will exhaust the possibilities.

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 this vertex together with four of its neighbors to see that we contain a subgraph of type $\tilde{D}_4$.

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+?? is the length of the path.

Subcase ii: We have a unique vertex of degree 3, having legs of length $p \leq q \leq r$. If $p = 2$, then we contain a graph of type $\tilde{E}_6$ since $2 \leq q \leq r$. So p must be 1; if $q = 3$, then we contain a graph of type $\tilde{E}_7$. If $q = 2$ and $r = 5$, we contain a graph of type $\tilde{E}_8$, so we must have $r = 2$, $r = 3$, or $r = 4$, corresponding to the cases E_6, E_7, and E_8, respectively. If $q = 1$, then any value of r gives us a graph of type D_{r+3}.[What is a leg (find a good way to define these)]

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's 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.

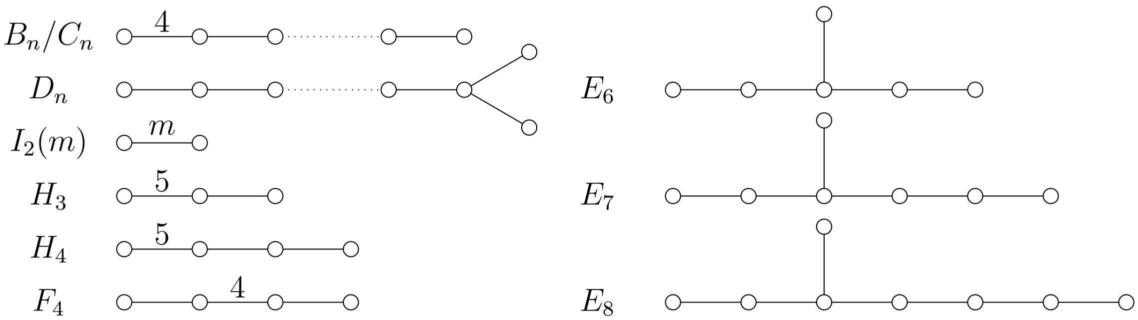

Figure 1: The positive definite graphs

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 $\tilde{F}_4$. $\square$

The following are used in the classification.

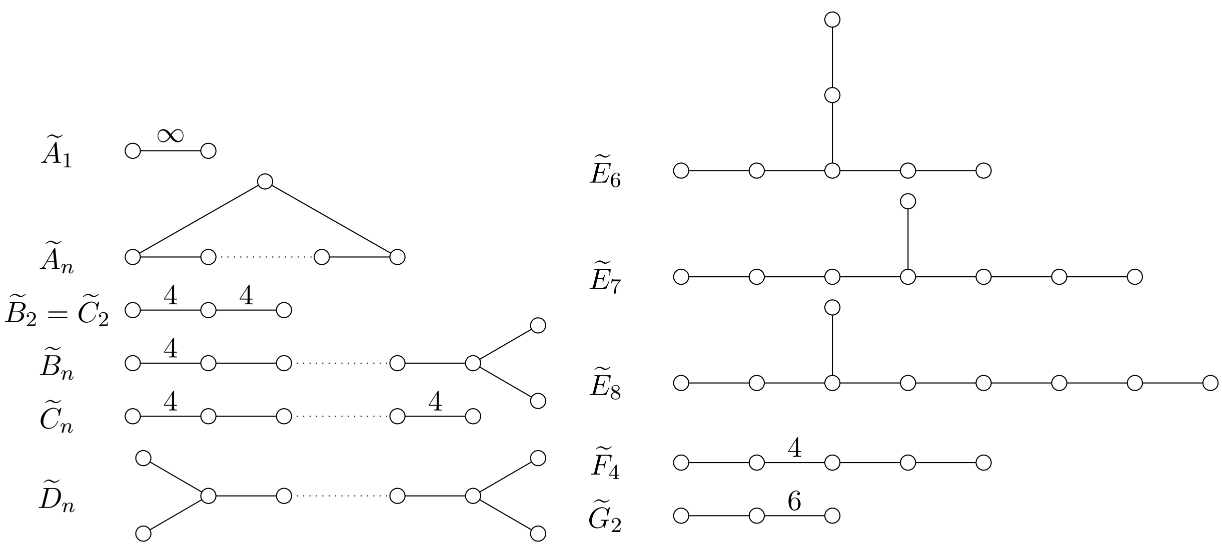

The following two Coxeter diagrams were also used in the calculation but are not affine; in fact, they have type $(3, 1)$ and $(4, 1)$, respectively.

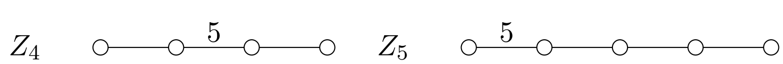

6.2 Affine Reflection Groups

In the case that B is semi-definite but not positive definite, we can view the representation of W as a geometric action on an affine Euclidean space. Suppose also that W is irreducible.

Let A be the matrix for B with respect to the basis $\{\alpha_s\}_{s\in S}$. Since W is irreducible, the graph is connected.

We have concluded that $V^\perp$ is spanned by some vector λ in $\text{conv}(\{\alpha_s\}_{s\in S})$

It may help the reader to focus on the following example. Consider the Coxeter group of type $\tilde{A}_2$ with Coxeter diagram as in Figure 6.2. The Schläfli matrix is $\begin{bmatrix} 1 & -1/2 & -1/2 \\ -1/2 & 1 & -1/2 \\ -1/2 & -1/2 & 1 \end{bmatrix}$, which has eigenvalues $(0, 3/2, 3/2)$.

(The simple roots are linearly independent, but the geometry is such that their 3 hyperplanes intersect in a common line $(\lambda, \lambda, \lambda)$. What this means is that in the dual space, the linear functionals lie on a plane, and we get a tiling of equilateral triangles. However, in the original space, the generating reflections fix a cylinder. We get roots decorating the cylinder)

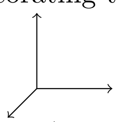

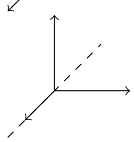

6.3 Geometric group actions in Coxeter Groups

Recall in Section 1.1 we established a correspondence between groups and metric spaces. We use this correspondence to restrict our view of the representation constructed in Section 5.4 to a more appropriate domain.

Example 6.3. We 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. Our quasi-isometry condition was a bit too coarse to keep track of finite information. However, the fact that a one-point space is not quasi-isometric to $\mathbb{R}^2$, we know we can restrict to some smaller space. Indeed, if we restrict to the unit circle, we have a geometric action. This is not a coincidence but actually the general case for finite Coxeter groups; we restrict to the unit sphere $S_V = \{v \in V \mid B(v, v) = 1\}$.

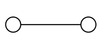

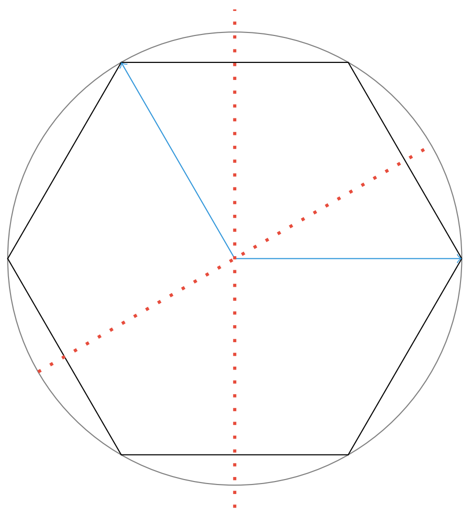

Example 6.4. We examine the case of the Coxeter group of type $\tilde{A}_2$. The Schläfli matrix is $\begin{bmatrix} 1 & -1/2 & -1/2 \\ -1/2 & 1 & -1/2 \\ -1/2 & -1/2 & 1 \end{bmatrix}$. Diagonalizing this matrix using the basis ?? yields $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ [Acts geometrically on affine hyperplane]

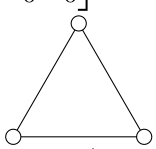

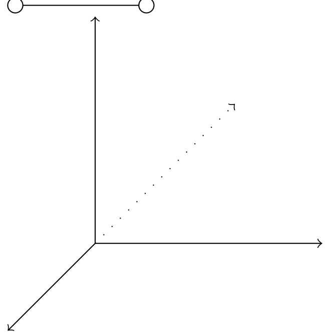

Example 6.5. Hyperbolic

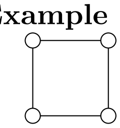

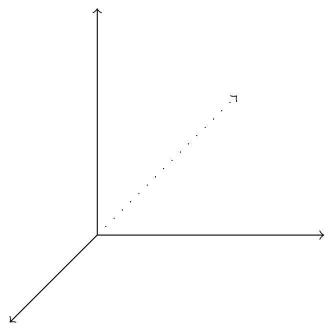

Part III
Lorentzian Coxeter Systems

7 Lorentzian Coxeter Systems

There is some dissonance in the literature as to what a hyperbolic Coxeter group is. For example, [DESCRIBE ALL DIFFERENT DEFNS OF HYPERBOLIC, EXPLAIN WHO MADE WHICH ONES ETC]

To circumvent this problem, we introduce a different word, and say that a Coxeter group is *Lorentzian* when the associated bilinear form is of type $(n-1,1)$. Note this is a more general notion than Humphreys' hyperbolic; indeed, we do not require that $B(v,v)<0$ for $v \in C$.

When we have a vector space V together with a bilinear form B on V, we will call the pair (V,B) a *quadratic space.*

Remark. Usually, the term "quadratic space" refers to a vector space together with a quadratic form. However, specifying a quadratic form Q on a real vector space uniquely specifies a symmetric bilinear form by the formula $B(v,w)=\frac{1}{2}(Q(v+w)-Q(v)-Q(w))$. Moreover, every symmetric bilinear form arises in this way from some quadratic form.

To each $v \in V$ with $B(v,v) \neq 0$, we can define $s_v \colon V \to V$ by $s_v(w) = w - 2\frac{B(v,w)}{B(v,v)}v$, which we call the *$B$-reflection associated with* v. Observe that $s_v \in O_B(V)$; indeed, $B(s_v(u), s_v(w)) = B(u-2\frac{B(v,u)}{B(v,v)}v, w-2\frac{B(v,w)}{B(v,v)}v) = B(u,w)-2\left(B(\frac{B(v,u)}{B(v,v)}v, w) + B(u, \frac{B(v,w)}{B(v,v)}v)\right)+ 4B(\frac{B(v,u)}{B(v,v)}v, \frac{B(v,w)}{B(v,v)}v) = B(u,w)$.

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 (Φ,Δ) together with (V,B) actually determines (W,S), so we can define a Lorentzian Coxeter system by a based root system in a quadratic space.

7.1 Subgroups and the Generalized Geometric Representation

The parabolic subgroups are the easiest to manage. Given a Coxeter system (W,S), we can consider subsets I of S, and examine the subgroup W_I of W generated by I. It turns out that (W_I, I) is a Coxeter system. The situation is less predictable when we take merely a subset of $R = \{w^{-1}sw \mid s \in S, w \in W\}$, as we'll see shortly.

8 Limit Roots

Of course, in the case of an infinite Coxeter group, the associated root system is also infinite. In our setting, the root system is a discrete subset of V, but we can consider directional limits.

Definition 8.1 (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 we define the *depth of* ρ to be $\mathrm{dp}(\rho) = 1 + \min\{\ell(w) \mid w(\rho) \in \Delta\}$. 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. We set Φ_n^+ to be the positive roots of depth n.

8.1 Roots Diverge

In this section, we show that any bounded set contains only finitely many roots, and hence that the root system is a discrete set.

8.2 Projecting the Roots

Although there is no hope of finding a limit root

Definition 8.2 (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 $\alpha \in \Delta$ (so really what φ does is add up the coordinates of α in the basis Δ). The kernel of φ is a hyperplane in V which does not intersect any simple root in Δ. As we will shortly describe, the main object of interest will be the affine subspace determined by Δ, which we can view as $V_1 \overset{\text{def}}{=} \varphi^{-1}(\{1\}) = \mathrm{aff}(\Delta)$. While all simple roots lie in V_1, the other positive roots need not. However, these roots are (by their very definition) positive linear combinations of elements of Δ, and so for $\rho \in \Phi^+$, we have $\varphi(\rho) > 0$. So we can consider the projection of ρ onto V_1 defined by $\widehat{\rho} \overset{\text{def}}{=} \frac{\rho}{\varphi(\rho)}$.

Proposition 8.3. *If* (Φ, Δ) *is a based root system in a quadratic space* (V, B)*, then the set* $\widehat{\Phi}$ *of normalized roots lies within* $\mathrm{conv}(\Delta)$*.*

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

Corollary. *The set of limit roots is empty if and only if* B *is positive-definite.*

8.3 Random Walks in Hyperbolic Space

The following picture should give some intuition to the assertion that "random walks in hyperbolic space diverge". We will not explicitly need this fact, but it will help motivate the next section.

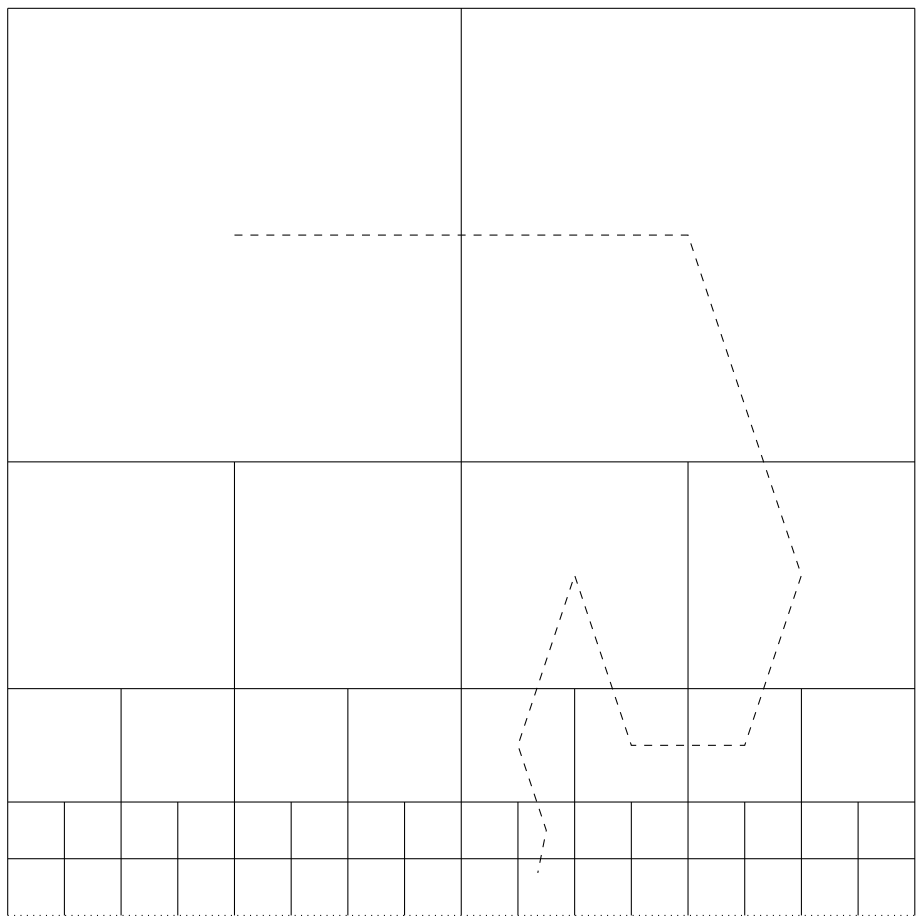

Interpret this picture as follows: a path in the upper half plane model of hyperbolic space is “twice as likely” to move down than up.

8.4 Infinite Reduced Words

Let (W, S) be a Coxeter system. Call a sequence $f\colon \mathbb{N} \to S$ *reduced* if every initial segment is a reduced word in W. For example, in $\widetilde{A}_1$, the sequence $s_1 s_2 s_1 s_2 s_1 \ldots$ is reduced. In a finite Coxeter group, there are no infinite reduced words.

OUTLINE: Show that the limit roots are contained in the light cone. The limit set of the orbit of any point is the same as the set of limit roots.

9 Examples

We consider examples of Lorentzian Coxeter systems, and aim to calculate their limit roots.

9.1 K_{n+1}

First, we consider K_{n+1}, the complete graph on $n+1$ vertices. For $n \geq 3$, this graph corresponds to a Lorentzian Coxeter system. The Schlafli matrix associated with this graph is $a_{ij} = \begin{cases} 1 & \text{if } i = j \\ -1/2 & \text{if } i \neq j \end{cases}$. In fact, the eigenvalues will be $\frac{3}{2}$ with multiplicity n, and $1 - \frac{n}{2}$ with multiplicity 1. Indeed, the vector $\vec{1}$ consisting of all ones is visibly an eigenvector with eigenvalue $1 - \frac{n}{2}$, and any permutation of the vector $(1, -1, 0, \ldots, 0)$ is an eigenvector with eigenvalue $\frac{3}{2}$. Any permutation of $(1, 1, 1, 0, \ldots, 0)$ will lie on the light cone.

9.2 (p,q,r)

Second, we look at (p, q, r) graphs. Let Γ be the graph formed by taking paths of length p, q, and r, and identifying an endpoint of each so that Γ is a tree. Thus, D_n is a $(2, 2, n-2)$ graph. Note that a (p, q, r) graph has rank $p + q + r - 2$, and that the ordering of p, q, and r is irrelevant so we may as well take $p \leq q \leq r$. Such a graph corresponds to a Lorentzian Coxeter system precisely when $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} < 1$.

To see this, we construct a representation of the Coxeter group corresponding to a (p, q, r) graph.

$$\left(\begin{array}{c|ccccc} & b_1 & \ldots & b_{p+1} & & \\ a & c_1 & \ldots & \ldots & c_{q+1} & \\ & d_1 & \ldots & \ldots & \ldots & d_{r+1} \end{array} \right)$$

Let $\lambda =$

9.3 $\tilde{E_8}$

Although $\tilde{E_8}$ may be viewed as a $(2, 3, 6)$ graph, we may more efficiently represent it in $\mathbb{R}^{9,1}$ using the following vectors. This is in hyperbolic space, but the span of the roots is affine.

$\{\pm(1 \mid 1^3 0^6)\} \cup \{(0 \mid 1^1 (-1)^1 0^7)\}$

Part IV
Combinatorics

10 Combinatorial Generalizations

We return our attention to the combinatorial aspects of Coxeter groups.

Definition 10.1 (Reflection length)**.** As before, let R be the set $\{wsw^{-1} \mid w \in W, s \in S\}$ of reflections in W. We define $\varepsilon_R \colon R^* \to W$ and $\ell_R \colon W \to \mathbb{N}$ analogously to ε and ℓ, so that $\ell_R(w)$ is the minimal length of a word in R^* having image w under ε_R.

Definition 10.2 (Partial orders)**.** For each of the following orders, we define the covering relations $\triangleleft$, and take the transitive closure to obtain a partial ordering $<$ on W. The Bruhat order turns out to be the most important, followed by the weak order.

Weak order We say $u \triangleleft_R w$ if $\ell(u) \leq \ell(w)$ and $u^{-1}w \in S$. Say $u \triangleleft_L w$ if instead $wu^{-1} \in S$.

Bruhat order We say $u \triangleleft w$ if $\ell(u) \leq \ell(w)$ and $u^{-1}w \in R$.

Absolute order We say $u \triangleleft_A w$ if $\ell_R(u) \leq \ell_R(w)$ and $u^{-1}w \in R$.

Proposition 10.3 (Symmetric orders)**.** *The Bruhat and absolute order are left-right symmetric.*

Proof. □

Definition 10.4. Upon considering these three orders at once, it seems natural to define a fourth order, in which the covering relations are $u \triangleleft' w$ if $\ell_R(u) \leq \ell_R(w)$ and $u^{-1}w \in S$. This order is fairly unstudied. We explore properties of $<'$ below.

Proposition 10.5 (Disconnected)**.** *The Hasse diagram of $<'$ is disconnected if [What??]*

Appendices

References