ISSN 1038-751X

Report 93-38 October 1993

Miscellaneous facts about Coxeter groups

Robert B. Howlett

Lectures given at the ANU Group Actions Workshop

School of Mathematics and Statistics
The University of Sydney
NSW 2006
Australia

The definition of "Coxeter group" (to be presented later) seems rather technical at first sight, and one might think, therefore, that the study of Coxeter groups is a very specialized subject of little relevance to the rest of mathematics. The truth is, however, that Coxeter groups are incredibly ubiquitous, and all mathematicians, whatever their specialities, are likely to encounter Coxeter groups at one time or another. The basic reason for this is the fact that

Finite Coxeter groups ≡ Finite Euclidean reflection groups.

It is obvious that Euclidean space is important in mathematics, as is symmetry, and when these ingredients are combined the result may well be a Euclidean reflection group.

Other Coxeter groups—the infinite ones—are less important than finite Euclidean reflection groups, but important still. Euclidean geometry is not the only important geometry. There is a large body of current mathematical research concerned with infinite Coxeter groups and their applications.

Coxeter's name appears in this context because of his investigations of discrete groups of linear and affine transformations of Euclidean space.

Let us start by observing that the product of two reflections r, s of Euclidean space is a rotation through the angle 2θ, where θ is the angle between the mirrors corresponding to r and s. In the 2-dimensional case the mirrors are just lines through the origin, and the above fact is familiar from undergraduate courses. In three dimensions the result is still true, although since the mirrors are now planes through the origin, the notion of the angle between the mirrors perhaps needs clarification. In fact, we simply define the angle between two planes through O to be the angle between the two lines through O which are the normals to the planes. In general, in n-dimensional Euclidean space, we can use the formula $a \cdot b = \|a\|\,\|b\| \cos\theta$ to define the angle θ between any two vectors a and b, and then define the angle between two hyperplanes (that is, $(n-1)$-dimensional subspaces) to be the angle between the normals. The reflection in a hyperplane is the linear transformation which fixes each vector on the hyperplane and takes each vector normal to the hyperplane to its negative, and the product of the reflections corresponding to hyperplanes at an angle θ is a rotation through 2θ about the ($(n-2)$-dimensional) axis which is the intersection of the two hyperplanes. (Note that in these lectures we will not deal with reflections in hyperplanes that do not pass through the origin; thus translations, and other transformations which do not fix the origin, will not arise.)

The following formula makes it easy to do calculations involving reflections:

Proposition. *Let a be a vector in Euclidean n-space and let ρ_a be the reflection in the hyperplane orthogonal to a. Then $\rho_a(v) = v - \frac{2(a\cdot v)}{(a\cdot a)}a$ for all $v \in V$.*

Proof. Since reflections are linear transformations, and since the given formula is clearly linear in v, it is enough to check the validity of the formula on a spanning set. If v lies in the reflecting hyperplane then $a \cdot v = 0$ and the formula gives $\rho_a(v) = v$, as it should. Putting $v = a$ the formula gives $\rho_a(a) = a - 2a = -a$, which is also right. Since the whole space is spanned by the vectors in the hyperplane together with the normal vector a, the formula must be valid for all v. □

Clearly, a rotation can lie in a finite group of transformations only if it is a rotation through an angle which is a rational multiple of π; otherwise the rotation itself will have infinite order. And if a and b are vectors with $\mathrm{angle}(a,b) = \frac{k}{n}\pi$, where $\gcd(k,n) = 1$, then the rotation $\rho_a\rho_b$ has order n. Furthermore, it is easily seen that ρ_a and ρ_b generate a dihedral group of order $2n$. Indeed, since $\rho_a^2 = \rho_b^2 = 1$, this follows from a general group-theoretic fact:

Proposition. Let r, s be elements of order 2 (in any group) and suppose that rs has order n. Then the subgroup generated by r and s is dihedral of order $2n$.

It is easy to check the truth of this by simply enumerating the elements of the subgroup generated by r and s. Since r and s are self-inverse, the only group elements obtainable as products of r's and s's are 1, r, s, rs, sr, rsr, srs, …, two products of each length greater than 0. But the relation $(rs)^n = 1$ gives $rsrs\ldots = srsr\ldots$, where there are n factors on each side; so the two products of length n give the same group element. And this is easily seen to be the element of maximal length in the subgroup, since pre- or post-multiplying it be either r or s will result in cancellation, giving one or other of the elements of length $n-1$.

Examples (1) If $n = 3$ then $\langle r, s\rangle$ consists of the six elements 1, r, s, rs, sr, $rsr = srs$. To realize this geometrically, let $r = \rho_a$ and $s = \rho_b$, where a and b are unit vectors such that $a \cdot b = \cos\frac{n-1}{n}\pi = -\frac{1}{2}$. The mirrors corresponding to r and s intersect at an angle of $\frac{2\pi}{3}$, and are axes of symmetry of an appropriately positioned equilateral triangle. It can be seen that $\langle r, s\rangle$ is the group of symmetries of this triangle. (For each $n \geq 2$, the dihedral group of order n is the group of symmetries of a regular n-sided polygon.)

(2) If $n = 2$ then $rs = sr$, and $\langle r, s\rangle = \{1, r, s, rs\}$ is the direct product $\langle r\rangle \times \langle s\rangle$. Note that this corresponds to the case of perpendicular hyperplanes; thus, one might say "perpendicular reflections commute."

(3) If n is infinite then the two expressions of length l give different group elements for all l; this is the *infinite dihedral group*. Note that, unlike our previous two examples, we cannot in this case let $r = \rho_a$ and $s = \rho_b$, where the unit vectors a and b span two-dimensional Euclidean space and satisfy $a \cdot b = \cos\frac{n-1}{n}\pi$. For, interpreting this as $a \cdot b = \lim_{n\to\infty}\cos\frac{n-1}{n}\pi$, we find that $a \cdot b = -1$, giving $(a+b)\cdot(a+b) = 0$, which is impossible in Euclidean space. However, we do obtain in this way a realization of the infinite dihedral group as a reflection group on a non-Euclidean two-dimensional space.

Since the transformation group generated by a set of reflections $\{\,\rho_a \mid a \in S\,\}$ preserves the subspace spanned by the set of vectors S and fixes every vector in the orthogonal complement $S^\perp$, there is no loss of generality in ignoring $S^\perp$ altogether, and assuming that S spans the whole space. The dihedral groups are thus essentially 2-dimensional reflection groups. It is also fairly easy to see that the dihedral groups are the only finite reflection groups on 2-dimensional Euclidean space.

The 3-dimensional case is more complicated and more interesting. If a, b, c are linearly independent vectors with the property that the corresponding reflections ρ_a, ρ_b, ρ_c lie in a finite group, then $\operatorname{angle}(a, b)$, $\operatorname{angle}(a, c)$ and $\operatorname{angle}(b, c)$ must all be rational multiples of π, and this is not easy to achieve. It can be done by making $\operatorname{angle}(a, b)$ an arbitrary multiple of π, and then making c perpendicular to both a and b. In this case $\langle\rho_a, \rho_b\rangle$ is a dihedral group, and since ρ_c commutes with both ρ_a and ρ_b it follows that $\langle\rho_a, \rho_b, \rho_c\rangle$ is the direct product of the dihedral group $\langle\rho_a, \rho_b\rangle$ and the cyclic group of order 2 generated by ρ_c. Apart from these direct products, there are only three 3-dimensional Euclidean reflection groups. These are the groups of symmetries of a regular tetrahedron, a cube, and a regular dodecahedron. Note that a regular octahedron, being in some sense the dual of a cube, has the same group of symmetries as a cube. Similarly, a regular icosahedron has the same group of symmetries as a regular dodecahedron.

Tetrahedron. Given a regular tetrahedron centred at the origin, there is a uniquely determined "dual" tetrahedron which is congruent to the given one, and has the property that each edge of the given tetrahedron is perpendicularly bisected by an edge of the dual. The vertices of the two tetrahedra together give the vertices of a cube. Let a and c be the position vectors relative to the origin of the midpoints of a pair of parallel (but not opposite) edges e_1 and e_2. Then a and c are perpendicular. The face opposite to that determined by e_1 and e_2 has two edges which are not parallel to e_1 and e_2; let b be the position vector of the midpoint of one of these. Then b makes an angle of $\frac{2}{3}\pi$ with a, since a rotation of the cube through a third of a revolution about one of its diagonals can be seen to take a to b. Similarly, the angle between b and c is also $\frac{2}{3}\pi$. Summarizing:

$$\begin{aligned}
\operatorname{angle}(a, b) &= \tfrac{2}{3}\pi, &\quad& \rho_a\rho_b \text{ has order } 3,\\
\operatorname{angle}(b, c) &= \tfrac{2}{3}\pi, && \rho_b\rho_c \text{ has order } 3,\\
\operatorname{angle}(a, c) &= \tfrac{1}{2}\pi, && \rho_a\rho_c \text{ has order } 2.
\end{aligned}$$

The reflecting plane corresponding to ρ_a perpendicularly bisects one edge of the tetrahedron we started with, and the opposite edge lies in this plane. Thus ρ_a is a symmetry of the tetrahedron; the same is true for ρ_b and ρ_c. Furthermore, it is also easy to see that these three reflections generate the group of all symmetries of the tetrahedron (which is just $\mathrm{Sym}(4)$).

Cube. Using the same cube as in the previous example, let b and c be as before, but now let a be the position vector of the midpoint of that face which is perpendicular to e_2 but does not include the edge corresponding to b. Then

$$\begin{aligned} \text{angle}(a,b) &= \tfrac{3}{4}\pi, && \rho_a\rho_b \text{ has order } 4,\\ \text{angle}(b,c) &= \tfrac{2}{3}\pi, && \rho_b\rho_c \text{ has order } 3,\\ \text{angle}(a,c) &= \tfrac{1}{2}\pi, && \rho_a\rho_c \text{ has order } 2. \end{aligned}$$

The reflecting plane corresponding to a is parallel to and midway between a pair of opposite faces of the cube, and there are two other similar reflections which together with ρ_a generate a group isomorphic to $C_2 \times C_2 \times C_2$. The group generated by ρ_b and ρ_c is isomorphic to Sym(3), and acts on the $C_2 \times C_2 \times C_2$ by permuting the three generating reflections. The group $\text{Sym}(3) \ltimes (C_2 \times C_2 \times C_2)$ generated by ρ_a, ρ_b and ρ_c is the full symmetry group of the cube. (This extends naturally to n dimensions: the symmetry group of the "hypercube" is a Euclidean reflection group isomorphic to $\text{Sym}(n) \ltimes (C_2)^n$.)

Dodecahedron. This time we can choose a, b and c to be the position vectors of A, B and C, midpoints of three suitable edges of the regular dodecahedron, satisfying

$$\begin{aligned} \text{angle}(a,b) &= \tfrac{4}{5}\pi, && \rho_a\rho_b \text{ has order } 5,\\ \text{angle}(b,c) &= \tfrac{2}{3}\pi, && \rho_b\rho_c \text{ has order } 3,\\ \text{angle}(a,c) &= \tfrac{1}{2}\pi, && \rho_a\rho_c \text{ has order } 2. \end{aligned}$$

Each of the twelve pentagonal faces has five lines of symmetry bisecting it, and each such line determines a plane through O, the reflection in which is a symmetry of the dodecahedron. Each of these planes bisect four faces, and so we obtain a set $\mathcal{H}$ of $5 \times 12/4 = 15$ planes corresponding to fifteen reflection symmetries. Each plane in $\mathcal{H}$ passes through a uniquely determined pair of opposite edges—note that there are thirty edges altogether—and perpendicularly bisects another pair of opposite edges. Furthermore, the line through O normal to the plane bisects a third pair of opposite edges. In particular the reflection in the plane orthogonal to the position vector of the midpoint of an edge is indeed a symmetry of the dodecahedron.

The points A, B and C can be chosen as follows. Let P be the centre of one of the faces, let Q be the midpoint of one of the edges of this face and let R be one of the vertices on this edge. The plane containing the triangle QRO is in the set $\mathcal{H}$, and so its normal through O passes through points A and A' which are midpoints of opposite edges. Choose A to be the one which is on the same side of the plane QRO as the point P. Similarly, choose B on the same side of PRO as Q with that BO normal to PRO, and C on the same side of PQO as R with CO normal to PQO. It can be shown that the reflections ρ_a, ρ_b and ρ_c thus determined—that is, the reflections in the planes QRO, PRO and PQO—generate the full group of symmetries of the dodecahedron.

(This group is in fact isomorphic to $\text{Alt}(5) \times C_2$. As we have seen above, each plane in $\mathcal{H}$ determines three pairs of opposite edges; these three pairs are mutually perpendicular to each other, and the fifteen pairs of opposite edges split into five such sets of three. The symmetry group permutes these five sets and also contains the transformation $-I$, which fixes each of the sets by taking each edge to its opposite.)

Returning to the example of the tetrahedron, recall that we found three elements $r = \rho_a$, $s = \rho_b$ and $t = \rho_c$ satisfying the relations

$$r^2 = s^2 = t^2 = 1,\ rsr = srs,\ sts = tst,\ rt = tr,$$

and generating a group isomorphic to Sym(4). We can easily show that the abstract group W defined by these generators and relations is in fact isomorphic to Sym(4). Note first that $r \mapsto (1,2)$, $s \mapsto (2,3)$ and $t \mapsto (3,4)$ certainly defines a surjective homomorphism $W \to \text{Sym}(4)$, as these three permutations generate Sym(4) and satisfy the required relations. Thus it remains to show that the order of W is at most 24. Now the subgroup $P = \langle r, s\rangle$ contains at most six elements, namely 1, r, s, rs, sr and $rsr = srs$. We proceed to list the right cosets of P in W. Firstly there is P itself, which is fixed by right multiplication by r and s. Right multiplication by t takes P to Pt. Right multiplying by r fixes Pt, since $(Pt)r = Prt = Pt$, and right multiplication by t takes Pt back to P. Right multiplication by s takes Pt to Pts, which is fixed by t

(since $Ptst = Psts = Pts$), and taken back to Pt by s. Right multiplying by r takes Pts to $Ptsr$, but no further cosets can be found, since s and t both fix $Ptsr$ (since $(Ptsr)s = Ptrsr = Prtsr = Ptsr$ and $(Ptsr)t = Ptstr = Pstsr = Ptsr$), while r takes $Ptsr$ back to Pts. These calculations have shown that the (nonempty) set $P \cup Pt \cup Pts \cup Ptsr$ is invariant under right multiplication by the generators of W, and therefore equals the whole of W.

Similar calculations can be used to show that the group of symmetries of the cube (of order 48) is defined abstractly by three generators r, s and t and the defining relations

$$r^2 = s^2 = t^2 = 1, \ rsrs = srsr, \ sts = tst, \ rt = tr,$$

while the group of symmetries of the dodecahedron (order 120) is defined by generators r, s and t with defining relations

$$r^2 = s^2 = t^2 = 1, \ rsrsr = srsrs, \ sts = tst, \ rt = tr.$$

Definition. A *Coxeter group* is a group with a presentation of the following form:

$$W = \mathrm{gp}\langle \{ r_a \mid a \in \Pi \} \mid r_a^2 = 1 \text{ for all } a \in \Pi, \, (r_a r_b)^{n_{ab}} = 1 \text{ for all } a, \, b \in \Pi \text{ with } a \neq b \rangle$$

where the set Π indexing the generators can be of any cardinality (although in these lectures it will always be assumed to be finite), and the parameters n_{ab} satisfy $n_{ab} = n_{ba} \in \mathbb{Z} \cup \{\infty\}$ and $n_{ab} \geq 2$ in all cases. A "relation" $(r_a r_b)^\infty = 1$ is to be interpreted as vacuous.

(Note that allowing $n_{ab} = 1$ would not yield any extra groups, since the relation $(r_a r_b)^1 = 1$ would just give $r_a = r_b$, allowing the generator s to be dropped.)

Lecture 2

The following table describes the classification of finite Coxeter groups:

Type	Diagram	Rank	Exponents	Order
A_n		n	$1, 2, 3, \ldots, n$	$(n+1)!$
$B_n = C_n$	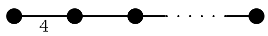	n	$1, 3, 5, \ldots, 2n-1$	$2^n n!$
D_n	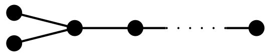	n	$1, 3, 5, \ldots, 2n-3, n-1$	$2^{n-1} n!$
$I_2(p)$	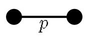	2	$1, p-1$	$2p$
H_3	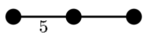	3	$1, 5, 9$	120
H_4	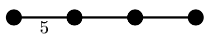	4	$1, 11, 19, 29$	14400
F_4	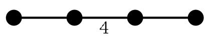	4	$1, 5, 7, 11$	1152
E_6	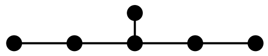	6	$1, 4, 5, 7, 8, 11$	$2^7 \cdot 3^4 \cdot 5$
E_7	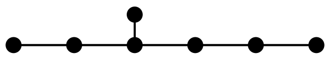	7	$1, 5, 7, 9, 11, 13, 17$	$2^{10} \cdot 3^4 \cdot 5 \cdot 7$
E_8	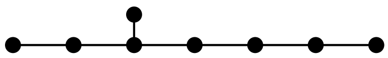	8	$1, 7, 11, 13, 17, 19, 23, 29$	$2^{14} \cdot 3^5 \cdot 5^2 \cdot 7$

and arbitrary direct products of groups from this list.

Notes.

1. The diagram is simply an abbreviated notation for the Coxeter presentation described at the end of Lecture 1. Each vertex of the diagram represents an element of Π. The parameter n_{ab} is 2 for nonadjacent vertices, 3 for vertices joined by an unlabelled edge, and is the label on the edge in the other cases.

The direct product of two Coxeter groups is a Coxeter group whose diagram is simply the union of the diagrams of the two factors (with no edges joining these component pieces, since if r_a and r_b are generators from different factors then $r_a r_b = r_b r_a$, whence $n_{ab} = 2$.) The connected components of the diagram of a Coxeter group W are diagrams of Coxeter groups which we will refer to as the *irreducible components* of W; clearly W is the direct product of these components.

2. The rank is the number of generators in the presentation—that is, the cardinality of Π—and is also the dimension of a Euclidean space on which the group acts as a group generated by reflections.
3. The classification almost coincides with the classification of complex semisimple Lie algebras. Indeed, the Weyl group of such a Lie algebra is a Euclidean reflection group, and all the groups on the above list can occur, with the exceptions of H_3, H_4 and $I_2(p)$ for $p = 5$ and all $p \geq 7$. The direct sum of two complex semisimple Lie algebras gives a complex semisimple Lie algebra, the Weyl group of which is the direct product of the Weyl groups of the factors. So a finite Coxeter group corresponds to a Lie algebra provided its irreducible components are all of types A_n, B_n, D_n, E_n, F_4 or $I_2(6)$ (this last case corresponding to the Lie algebra of type G_2).
4. For $n > 2$ the Lie algebras B_n and C_n are not isomorphic, but their Weyl groups are; hence the "$B_n = C_n$" on the above list.
5. All of the irreducible finite Coxeter groups can be described as the symmetry groups of various regular or nearly regular objects in n-dimensional space (where n is the rank). For instance, the group of type A_n, which is just $\mathrm{Sym}(n+1)$, is the symmetry group of a regular n-simplex (the n-dimensional generalization of the 2-dimensional equilateral triangle and the 3-dimensional regular tetrahedron). The group of type B_n is the symmetry group of the hypercube (measure polytope) or its dual, the cross polytope (the generalization of the octahedron: its $2n$ vertices lie on the n axes of a rectangular coordinate system at a fixed distance from the origin). The 2^{n-1} points with coordinates $(\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_n)$ where $\varepsilon_i = \pm 1$ for all i and $\varepsilon_1 \varepsilon_2 \cdots \varepsilon_n = 1$ form the vertices of a polytope whose symmetry group is the Coxeter group of type D_n; this group is isomorphic to the group of all monomial matrices whose nonzero entries are all ± 1, with an even number of -1's. Types H_4 and F_4 are symmetry groups of regular 4-dimensional polytopes, while E_6, E_7 and E_8 correspond to certain semiregular polytopes.
6. The exponents of a rank n Coxeter group W are n numbers m_1, m_2, $\ldots$, m_n, possessing many remarkable properties. Let $c = r_{a_1} r_{a_2} \cdots r_{a_n} \in W$, where $\Pi = \{a_1, a_2, \ldots, a_n\}$. That is, c is the product, in some order, of all the generators of W in the Coxeter presentation. Obviously c depends on the choice of the ordering of Π, but since the Coxeter diagram of W is a graph with no circuits, and since the generators corresponding to any pair of non-adjacent vertices commute with each other, a relatively easy graph-theoretic argument shows that different orderings give conjugate elements c. These elements are called the *Coxeter elements* of W. They have order $h = 2N/n$, where N is the number of reflections in W and n is the rank. The eigenvalues of c as a transformation of n-dimensional space are $\omega^{m_1}, \omega^{m_2}, \ldots, \omega^{m_n}$, where $\omega = e^{2\pi i/h}$; this is the reason for the name "exponents". It turns out that the numbers $m_1 + 1$, $m_2 + 1$, $\ldots$, $m_n + 1$ are the degrees of certain homogeneous polynomials in n variables, known as the "basic polynomial invariants of W". The product $(m_1+1)(m_2+1)\cdots(m_n+1)$ equals the order of W, and the sum $m_1 + m_2 + \cdots + m_n$ equals N, the number of reflections. (Proving these facts is beyond the scope of these lectures!)

From this point on we assume that W is an arbitrary finitely generated Coxeter group, not necessarily finite. Let V be a vector space over R with basis Π (in bijective correspondence with the distinguished generating set of W). Define a bilinear symmetric scalar product $(u, v) \mapsto u \cdot v$ on V by

$$a \cdot b = \begin{cases} 1 & \text{if } a = b \\ -\cos(\pi/n_{ab}) & \text{if } a \neq b \end{cases}$$

for all a, $b \in \Pi$. Note that despite the use of the "dot" notation we are not claiming that this scalar product is positive definite: indeed, $v \cdot v < 0$ does sometimes occur. Thus V is not necessarily Euclidean. However, if a, $b \in \Pi$ have $n_{ab} \neq \infty$, then the restriction of the scalar product to the subspace spanned by a and b is positive definite. This is clear since it is certainly possible to find unit vectors in the Euclidean plane whose inner product is $-\cos(\pi/n_{ab})$, for this just says that the angle between them is $(n_{ab} - 1)\pi/n_{ab}$. So in the case $n_{ab} < \infty$ it is natural to say that the angle between the basis vectors a, $b \in \Pi$ is $(n_{ab}-1)\pi/n_{ab}$. When $n_{ab} = \infty$ (giving $a \cdot b = -1$) the angle between a and b is undefined.

For each $a \in \Pi$ define $\rho_a\colon V \to V$ by

$$\rho_a(v) = v - 2(v \cdot a)a$$

for all $v \in V$. Clearly, ρ_a is a linear transformation.

Proposition. *The transformations ρ_a just defined satisfy the relations $\rho_a^2 = 1$ and $(\rho_a\rho_b)^{n_{ab}} = 1$ for all a, $b \in \Pi$. They also preserve the scalar product.*

Proof. For all $a \in \Pi$ and u, $v \in V$ we have

$$\begin{aligned}\rho_a(v)\cdot\rho_a(u) &= (v - 2(v\cdot a)a)\cdot(u - 2(u\cdot a)a)\\ &= v\cdot u - 2(v\cdot a)(a\cdot u) - 2(u\cdot a)(v\cdot a) + 4(v\cdot a)(u\cdot a)(a\cdot a)\\ &= v\cdot u\end{aligned}$$

since $a \cdot a = 1$. Thus ρ_a preserves the scalar product.

Trivially, $\rho_a(a) = -a$, and hence $\rho_a^2(v) = \rho_a(v - 2(v\cdot a)a) = \rho_a(v) + 2(v\cdot a)a = v$ for all $v \in V$. Thus the relation $\rho_a^2 = 1$ holds.

Let a, $b \in \Pi$ and assume that $n = n_{ab} \neq \infty$ (since the relation $(\rho_a\rho_b)^{n_{ab}} = 1$ is vacuously true if $n_{ab} = \infty$). The subspace U of V spanned by a, b and v is ρ_a-invariant (as, indeed, is every subspace containing a), and the matrix of ρ_a on U is

$$R_a = \begin{pmatrix} -1 & -2(b\cdot a) & -2(v\cdot a)\\ 0 & 1 & 0\\ 0 & 0 & 1\end{pmatrix}$$

(in the sense that the diagram

$$\begin{tikzcd} \mathbb{R}^3 \arrow[r, "\phi"] \arrow[d, "R_a"'] & U \arrow[d, "\rho_a"] \\ \mathbb{R}^3 \arrow[r, "\phi"] & U \end{tikzcd}$$

commutes, where ϕ is the map which takes the (column) vector $(\lambda,\mu,\nu)^T \in \mathbb{R}^3$ to $\lambda a + \mu b + \nu v \in U$, and the mapping from $\mathbb{R}^3$ to $\mathbb{R}^3$ is left multiplication by R_a). Similarly, U is also ρ_b-invariant, and the matrix of ρ_b on U is

$$R_b = \begin{pmatrix} 1 & 0 & 0\\ -2(a\cdot b) & -1 & -2(v\cdot b)\\ 0 & 0 & 1\end{pmatrix}.$$

Computing R_aR_b we find that the matrix of $\rho_a\rho_b$ on U is $\begin{pmatrix} M & \mathbf{x}\\ \mathbf{0} & 1\end{pmatrix}$, where $\mathbf{0}$ is the two-component zero row vector, $\mathbf{x}$ is a two-component column whose entries are immaterial for the present argument, and M is the 2×2 matrix describing the action of $\rho_a\rho_b$ on the space spanned by a and b. Now since a and b span a Euclidean plane on which ρ_a and ρ_b act as reflections in the lines perpendicular to a and b, and since the angle between a and b is $\pi - \theta$, where $\theta = \pi/n$, it follows that on this plane $\rho_a\rho_b$ is a rotation through an angle of 2θ. Consequently $M^n = 1$. Actually,

$$M = \begin{pmatrix} -1 + 4\cos^2\theta & -2\cos\theta\\ 2\cos\theta & -1\end{pmatrix} = \frac{1}{\sin\theta}\begin{pmatrix}\sin 3\theta & -\sin 2\theta\\ \sin 2\theta & -\sin\theta\end{pmatrix}$$

(which is not the familiar matrix for a rotation through 2θ since it is not computed relative to an orthonormal basis), and induction can be used to show that

$$M^k = \frac{1}{\sin\theta}\begin{pmatrix}\sin(2k+1)\theta & -\sin 2k\theta\\ \sin 2k\theta & -\sin(2k-1)\theta\end{pmatrix}$$

for all integers k. Although this calculation was not needed for the present proof, we will use it later, and it does confirm that $M^n = 1$.

The eigenvalues of M are $\cos\theta \pm i\sin\theta$, and so $M - 1$ is nonsingular. So since

$$0 = M^n - 1 = (M-1)(M^{n-1} + M^{n-2} + \cdots + M + 1)$$

it follows that $M^{n-1} + M^{n-2} + \cdots + M + 1 = 0$. But the matrix of $(\rho_a\rho_b)^n$ on the space U is

$$\begin{pmatrix} M & \mathbf{x} \\ \mathbf{0} & 1 \end{pmatrix}^n = \begin{pmatrix} M^n & \mathbf{x}' \\ \mathbf{0} & 1 \end{pmatrix}$$

where $\mathbf{x}' = (M^{n-1} + M^{n-2} + \cdots + M + 1)\mathbf{x}$ (as a short calculation shows). Thus $\mathbf{x}' = \mathbf{0}$, which means that $(\rho_a\rho_b)^n$ fixes the vector v. Since v was an arbitrary vector in V we have shown that $(\rho_a\rho_b)^n = 1$, as required. $\square$

Lecture 3

The Proposition proved in Lecture 2 showed that the transformations ρ_a of the space V satisfy the defining relations of the group W. Thus there is a homomorphism ρ from W to the group of all nonsingular linear transformations of V such that $\rho(r_a) = \rho_a$ for all $a \in \Pi$. In other words, there is an action of W on the space V defined by

$$r_a v = v - 2(v \cdot a)a$$

for all $a \in \Pi$ and $v \in V$. This action turns out to be faithful, and invariably provides the best way to investigate W (more powerful than use of the building, or Coxeter complex, of W, since this topological object can be embedded in V).

Definition. The *root system* of W is the set $\Phi = \{\, wa \mid w \in W,\ a \in \Pi \,\}$. Elements of Π are called *simple roots*, and the elements $\{\, r_a \mid a \in \Pi \,\}$ (the distinguished generators for W) are called the *simple reflections*.

Note that when an element of Φ is expressed as a linear combination of the basis vectors in Π, the coefficients are not necessarily rational integers. Consequently the above definition of "root system" conflicts somewhat with the definition appropriate in the study of Lie algebras. Furthermore, for us all roots (elements of Φ) have length 1, which is not the case for root systems of Lie algebras. But a root system in the Lie algebra sense becomes a root system in our sense if each root is replaced by a unit vector in the same direction. And our definition gives a root system to every Coxeter group, not merely those which are Weyl groups of Lie algebras.

Properties of the root system provide the most powerful tools for investigating W.

Examples (1) Type A_3, 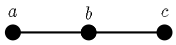

.

We have $\Pi = \{a, b, c\}$ with $a \cdot b = -\frac{1}{2} = b \cdot c$ and $a \cdot c = 0$. Using the formula $r_a v = v - 2(v \cdot a)a$ we find that

$$r_a a = -a, \qquad r_a b = b + a, \qquad r_a c = c.$$

Similarly

$$r_b a = a + b, \qquad r_b b = -b, \qquad r_b c = c + b,$$
$$r_c a = a, \qquad r_c b = b + c, \qquad r_c c = -c.$$

We can quickly generate all the roots in the root system by applying r_a, r_b and r_c to all vectors which are already known to be roots, using the simple roots to start the process. For example, our calculations above have already shown that $a + b$ and $b + c$ are roots, and so $r_a(b + c) = r_a b + r_a c = a + b + c$ must also be a root. Furthermore, since the negative of every simple root is a root (for example, $-a = r_a a$), it follows that the negative of every root is a root. So the following twelve vectors are all roots:

$$a,\ b,\ c,\ a+b,\ b+c,\ a+b+c,\ -a,\ -b,\ -c,\ -a-b,\ -b-c \text{ and } -a-b-c.$$

It is easily shown that this set is closed under the action of r_a, r_b and r_c; hence it is the entire root system. Quite naturally the roots with the minus signs are called *negative roots*, the others *positive roots*.

(2) Type D_4, 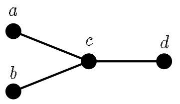

Here we have $\Pi = \{a, b, c, d\}$, with $a \cdot c = b \cdot c = d \cdot c = -\frac{1}{2}$ and $a \cdot b = b \cdot d = d \cdot a = 0$. Hence the actions of the simple reflections on the basis vectors are as follows:

$$\begin{array}{llll} r_a(a) = -a & r_a(b) = b & r_a(c) = c + a & r_a(d) = d \\ r_b(a) = a & r_a(b) = -b & r_b(c) = c + b & r_b(d) = d \\ r_c(a) = a + c & r_c(b) = b + c & r_c(c) = -c & r_c(d) = d + c \\ r_d(a) = a & r_d(b) = b & r_d(c) = c + d & r_d(d) = -d. \end{array}$$

This already shows that $a + c$, $b + c$ and $d + c$ are roots. Now $r_a(b + c) = r_a(b) + r_a(c) = b + a + c$ must be a root, and similarly $b + d + c$ and $a + d + c$ are roots also. Continuing in this way we find that $r_a(b + c + d) = a + b + c + d$ is a root, and so is $r_c(a + b + c + d) = a + b + 2c + d$. And, as always, the negative of a root is necessarily a root. Applying the simple reflections to the roots found so far turns out to not yield any more. Hence we have found all the roots; there are twenty-four altogether. They are

$$a,\ b,\ c,\ d,\ a + c,\ b + c,\ d + c,\ a + b + c,\ b + d + c,\ d + a + c,\ a + b + c + d,\ a + b + 2c + d$$

and their negatives. Observe once again the natural division of the root system into a set of positive roots and a set of negative roots: when a root is expressed as a linear combination of the simple roots, the coefficients are either all nonnegative or all nonpositive. This is an important general fact, which will be proved below.

(3) Type $I_2(p)$, 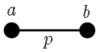

The group in this case is simply the dihedral group of order $2p$, and the vector space V can be identified with the Euclidean plane. The simple roots a and b are unit vectors, and the angle between them is $(p-1)\pi/p$. Since $r_a r_b$ acts as a rotation through $2\pi/p$, applying powers of $r_a r_b$ to a and $-b$ yields $2p$ distinct roots. Specifically, for each i from 0 to $2p - 1$, let v_i be the unit vector such that the anticlockwise angle from a to v_i is $i\pi/p$; these vectors v_i are the roots. The following diagram illustrates this when $p = 7$:

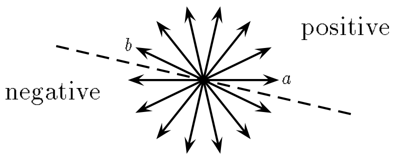

Observe that the roots on the positive side of the dashed line are positive combinations of a and b, since they lie "in between" a and b (in an angular sense). Those on the other side of the dashed line are the negatives of the positive roots, hence negative combinations of a and b.

If we let $\theta = \pi/p$ then the matrix of $r_a r_b$ on V relative to the basis a, b is the matrix M which appeared in the proof in Lecture 2. Using the formula for M^k given there we see that

$$(r_a r_b)^k a = \frac{\sin(2k+1)\theta}{\sin\theta} a + \frac{\sin 2k\theta}{\sin\theta} b$$

and

$$(r_a r_b)^k (-b) = \frac{\sin 2k\theta}{\sin\theta} a + \frac{\sin(2k-1)\theta}{\sin\theta} b,$$

so that in fact the root system is

$$\Phi = \left\{ \frac{\sin(m+1)\theta}{\sin\theta} a + \frac{\sin m\theta}{\sin\theta} b \;\middle|\; 0 \le m < 2p \right\}.$$

Observe that both coefficients are nonnegative when $0 \le m < p$, nonpositive when $p \le m < 2p$.

(4) The infinite dihedral group, 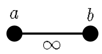

In this case we have $a \cdot b = -1$ and hence

$$\begin{array}{ll} r_a a = -a & r_b a = a + 2b \\ r_a b = a + 2b & r_b b = -b. \end{array}$$

A short calculation shows that the roots are precisely all vectors of the form $na + mb$ with $|n - m| = 1$ (as could have been predicted by letting $\theta \to 0$ in the $I_2(p)$ example).

(5) Type E_8, 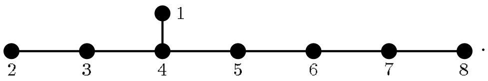

.

In this case it turns out that there are 120 positive roots and 120 negative roots. They are all integral linear combinations of the simple roots; determining them as combinations of the simple roots is somewhat tedious, but not too bad really, and gives information which can be useful occasionally. Alternatively, the root system can be described as follows. Let $e_1, e_2, \ldots, e_8$ be an orthogonal basis of eight-dimensional Euclidean space, with each e_i having length $1/\sqrt{2}$. Define

$$\Phi = \{\, \pm e_i \pm e_j \mid i \neq j \,\} \cup \{\, \tfrac{1}{2} \sum_{i=1}^{8} \varepsilon_i e_i \mid \varepsilon_i = \pm 1 \text{ and } \prod_{i=1}^{8} \varepsilon_i = 1 \,\}.$$

It is not hard to check that $u \cdot u = 1$ and $u \cdot v \in \{-\frac{1}{2}, 0, \frac{1}{2}\}$ for all $u, v \in \Phi$ with $u \neq \pm v$. Thus the angle between two vectors in Φ is always either $\pi/3$, $\pi/2$ or $2\pi/3$. Define $a_1 = \frac{1}{2}(e_1 - e_2 - e_3 - e_4 - e_5 - e_6 - e_7 + e_8)$ and $a_2 = e_1 + e_2$, and, for $3 \leq i \leq 8$, define $a_i = -e_{i-1} + e_{i-2}$. Then if we take $\Pi = \{\, a_i \mid 1 \leq i \leq 8 \,\}$ it can be checked that the inner products $a_i \cdot a_j$ are as they should be for consistency with the numbering chosen above for the vertices of the Coxeter diagram. (That is, $a_i \cdot a_j = -\frac{1}{2}$ if the vertices i and j are adjacent, $a_i \cdot a_j = 0$ otherwise.) It is also not hard to check that the reflections r_{a_i} do indeed preserve Φ. A root $\sum_{i=1}^{8} n_i a_i$ is positive if the largest i with $n_i \neq 0$ has $n_i > 0$. The "highest" root (that is, the root $\sum n_i a_i$ with $\sum n_i$ maximal) is the positive root which is perpendicular to all the simple roots but a_8, namely $e_7 + e_8 = 3a_1 + 2a_2 + 4a_3 + 6a_4 + 5a_5 + 4a_6 + 3a_7 + 2a_8$.

It is possible to explicitly describe the linear transformations $\{\, \rho(w) \mid w \in W \,\}$ (where W is the Coxeter group of type E_8) as matrices relative to the basis $e_1, \ldots, e_8$. Firstly, there are the $8!$ permutation matrices, and 2^7 diagonal matrices with diagonal entries ± 1 and determinant 1. These generate a group of order $2^7 8!$ which is a subgroup S of $\rho(W)$. The idea is to investigate the cosets of S. Now if $v_1, \ldots, v_8$ is any orthonormal basis of the space of eight component row vectors then there are $2^7 8!$ orthogonal matrices of the form xg where $x \in S$ and g is the matrix whose rows are $v_1, \ldots, v_8$. (These matrices are obtained from g by permuting the rows and multiplying an even number of rows by -1.) We proceed to describe a large number of orthonormal bases which give rise to elements of $\rho(W)$.

Let $\eta_1, \eta_2, \ldots, \eta_8$ be signs $\eta_i = \pm 1$ with $\prod_{i=1}^{8} \eta_i = 1$, and define

$$v_{ij} = \begin{cases} \frac{3}{4} & \text{if } i = j \\ -\frac{1}{4}\eta_i\eta_j & \text{if } i \neq j. \end{cases}$$

Let v_i be the row whose jth entry is v_{ij}. Then $v_1, \ldots, v_8$ form an orthonormal basis and give rise to a coset of S in $\rho(W)$. In fact this gives 64 cosets, corresponding to the 64 choices for the signs η_i, comprising a single (S, S) double coset in $\rho(W)$.

Choose a division of the set $\{1, 2, \ldots, 8\}$ into two subsets J and J' of four elements each—there are thirty-five ways of doing this—and let $\varepsilon = \pm 1$. Let $J = \{j_1, j_2, j_3, j_4\}$ and $J' = \{j_5, j_6, j_7, j_8\}$, and let α_{ij} be the (i, j)-entry of the matrix

$$X = \begin{pmatrix} 1 & \varepsilon & \varepsilon & \varepsilon & 0 & 0 & 0 & 0 \\ \varepsilon & 1 & -1 & -1 & 0 & 0 & 0 & 0 \\ \varepsilon & -1 & 1 & -1 & 0 & 0 & 0 & 0 \\ \varepsilon & -1 & -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \varepsilon & \varepsilon & \varepsilon \\ 0 & 0 & 0 & 0 & \varepsilon & 1 & -1 & -1 \\ 0 & 0 & 0 & 0 & \varepsilon & -1 & 1 & -1 \\ 0 & 0 & 0 & 0 & \varepsilon & -1 & -1 & 1 \end{pmatrix}$$

Let v_i be the vector whose kth entry is $\frac{1}{2}\alpha_{ij_k}$; then the matrix g whose rows are the v_i is in $\rho(W)$. Since there were 35 possible partions of $\{1, 2, \ldots, 8\}$ as $J \cup J'$ and two choices for ε, we have in fact obtained another 70 cosets of S, which again comprise one double coset. (The elements of this double coset are simply what you get by permuting the rows and columns of $\frac{1}{2}X$ and multiplying an even number of rows and/or columns by -1.)

We have now described $1+64+70=135$ cosets of S in $\rho(W)$, and in view of the assertion in Lecture 2 that $|W|=2^{14}\cdot 3^5\cdot 5^2\cdot 7$, this must be all of them. (And, incidentally, ρ must be an isomorphism.)

Lecture 4

In this lecture we will prove that the root system always splits into a set of positive roots and a set of negative roots—mixed positive and negative coefficients do not occur—and as a corollary we will deduce that ρ is always an isomorphism.

Since the group W is generated by the set of elements $\{\, r_a \mid a\in\Pi\,\}$ it follows that for each $w\in W$ there exist $a_1, a_2, \ldots, a_l \in \Pi$ with $w=r_{a_1}r_{a_2}\cdots r_{a_l}$. The least such nonnegative integer l is called the *length*, $l(w)$, of the element w. The identity element is the unique $w\in W$ of length 0.

The following lemma is a triviality:

Lemma. *If $w\in W$ and $a\in\Pi$ then $|l(w)-l(wr_a)|\le 1$. Moreover, if $w\neq 1$ then there exists $b\in\Pi$ with $l(wr_b)=l(w)-1$.*

Proof. If $w=r_{a_1}r_{a_2}\cdots r_{a_l}$ then $wr_a=r_{a_1}r_{a_2}\cdots r_{a_l}r_a$. Hence there is an expression for wr_a having length $l(w)+1$, and so $l(wr_a)\le l(w)+1$. But since we also have $w=(wr_a)r_a$, a symmetrical argument gives $l(w)\le l(wr_a)+1$, whence the first assertion follows. For the other part, let $w=r_{a_1}r_{a_2}\cdots r_{a_l}$ with the a_i in Π and $l=l(w)>0$, and let $b=a_l$. Then $wr_b=r_{a_1}r_{a_2}\cdots r_{a_{l-1}}$, and so $l(wr_b)\le l(w)-1$, and combined with the first part this gives $l(wr_b)=l(w)-1$. □

Define $\mathrm{PLC}(\Pi)=\{\,\sum_{a\in\Pi}\lambda_a a \mid \lambda_a\ge 0 \text{ for all } a\in\Pi\,\}$, the set of *positive linear combinations* of Π.

Theorem. *Let $w\in W$ and $a\in\Pi$. If $l(wr_a)\ge l(w)$ then $wa\in\mathrm{PLC}(\Pi)$.*

Proof. Choose $w\in W$ of minimal length such that the assertion of the theorem statement fails for some $a\in\Pi$, and choose such an a. Certainly w is not the identity element 1, since $1a=a$ is trivially a positive linear combination of Π (since $a\in\Pi$). So $l(w)\ge 1$, and we may choose $b\in\Pi$ such that $w_1=wr_b$ has length $l(w)-1$. If $l(w_1r_a)\ge l(w_1)$ then $l(w_1r_c)\ge l(w_1)$ holds for both $c=a$ and $c=b$ (since $w_1r_b=w$). Alternatively, if $l(w_1r_a)<l(w_1)$, we define $w_2=w_1r_a$, and note that $l(w_2r_c)\ge l(w_2)$ will hold for both $c=a$ and $c=b$ if $l(w_2r_b)\ge l(w_2)$. If this latter condition is not satisfied then we define $w_3=w_2r_b$. Continuing in this way we find, for some positive integer k, a sequence of elements $w_0=w$, w_1, w_2, $\ldots$, w_k with $l(w_i)=l(w)-i$ for all $i=0,1,2,\ldots,k$, and, when $i<k$,

$$w_{i+1}=\begin{cases} w_ir_a & \text{if } i \text{ is odd,}\\ w_ir_b & \text{if } i \text{ is even.}\end{cases}$$

Now since $0\le l(w_k)=l(w)-k$ we conclude that $l(w)$ is an upper bound for the possible values of k. Choosing k to be as large as possible, we deduce that $l(w_kr_c)\ge l(w_k)$ for both $c=a$ and $c=b$, for otherwise the process described above would allow a w_{k+1} to be found, contrary to the definition of k. By the minimality of our original counterexample it follows that w_ka and w_kb are both in $\mathrm{PLC}(\Pi)$.

We have that $w=w_kv$, where v is expressible as an alternating product of r_a's and r_b's, ending in r_b and with k factors altogether. Obviously this means that $l(v)\le k$. Now combining a minimal length expression for w_k as a product of simple reflections with a minimal length expression for vr_a gives an expression for $w_kvr_a=wr_a$ of length $l(w_k)+l(vr_a)$. So in view of the hypothesis that $l(wr_a)\ge l(w)$ we conclude that

$$l(w_k)+l(vr_a)\ge l(wr_a)\ge l(w)=l(w_k)+k\ge l(w_k)+l(v),$$

and hence $l(vr_a)\ge l(v)=k$. Note also that v cannot have a minimal length expression in which the final factor is r_a, for if so vr_a would have a strictly shorter expression.

Let $n = n_{ab}$. Since r_a and r_b satisfy the defining relations of the dihedral group of order $2n$, it follows that every element of the subgroup generated by r_a and r_b has an expression of length less than $n+1$ as an alternating product of r_a's and r_b's. So, in particular, $k = l(v) \leq n$. Moreover if n is finite then the two alternating products of length n define the same element; so k cannot equal n, as v has no minimal expression ending with r_a. Now, as we have seen, the matrix of $(r_a r_b)^m$ on the space spanned by a and b (relative to the basis a, b) is

$$\begin{pmatrix} t_{2m+1} & -t_{2m} \\ t_{2m} & -t_{2m-1} \end{pmatrix}$$

where $t_i = \sin i\theta / \sin\theta$ (with $\theta = \pi/n$) if $n < \infty$, and $t_i = i$ if $n = \infty$. Observe that t_i is nonnegative for $0 \leq i \leq n$. But if k is even, $k = 2m$ say, then $va = (r_a r_b)^m a = t_{2m+1}a + t_{2m}b$, while if k is odd, $k = 2m+1$, then $va = (r_b r_a)^m r_b a = -(r_b r_a)^{m+1} a = t_{2m+2}b + t_{2m+1}a$. Thus the coefficients of a and b are t_k and t_{k+1}, in one order or the other, and these numbers are both nonnegative since $k < n$.

We have now shown that $va = \lambda a + \mu b$, with $\lambda \geq 0$ and $\mu \geq 0$, and hence

$$wa = w_k va = w_k(\lambda a + \mu b) = \lambda w_k a + \mu w_k b \in \mathrm{PLC}(\Pi)$$

since $w_k a$, $w_k b \in \mathrm{PLC}(\Pi)$. This contradicts our original choice of w and a as a counterexample to the statement of the theorem; so the proof is completed. □

The power of the theorem just proved may not be immediately apparent, but the proof of the next corollary makes it apparent: by use of the theorem, an important and highly non-obvious result is reduced to a triviality.

Corollary 1. *If $w \in W$ is such that $wv = v$ for all $v \in V$, then $w = 1$. That is, the kernel of the homomorphism $\rho\colon W \to \mathrm{GL}(V)$ is trivial.*

Proof. If $w \neq 1$ then $l(w) > 0$, and so we can write $w = w'r_a$ with $a \in \Pi$ and $l(w') = l(w) - 1$. Since $l(w'r_a) > l(w')$ the theorem yields $w'a \in \mathrm{PLC}(\Pi)$, and so

$$a = wa = (w'r_a)a = w'(r_a a) = w'(-a) = -w'(a) \in \mathrm{NLC}(\Pi),$$

where $\mathrm{NLC}(\Pi)$ is the set of negative linear combinations of Π. This contradicts the linear independence of Π, since $a \in \Pi$ cannot be expressed as a linear combination of Π in any way but the obvious one. □

We define $\Phi^+ = \Phi \cap \mathrm{PLC}(\Pi)$ and $\Phi^- = \{\, -v \mid v \in \Phi^+ \,\} = \Phi \cap \mathrm{NLC}(\Pi)$.

Corollary 2. $\Phi = \Phi^+ \cup \Phi^-$.

Proof. Let $v \in \Phi$. By the definition of Φ there exist $w \in W$ and $a \in \Pi$ with $v = wa$. Let $w' = wr_a$, noting that the lengths of w and w' differ by at most 1. Now if $l(w') \geq l(w)$ then $l(wr_a) \geq l(w)$ and the theorem yields that $v = wa \in \Phi^+$. On the other hand if $l(w) \geq l(w')$ then $l(w'r_a) \geq l(w')$ and the theorem yields that $w'a \in \Phi^+$, whence $v = w'r_a a = w'(-a) = -w'a \in \Phi^-$. □

The preceding proof shows, incidentally, that $l(wr_a) = l(w)$ is not a possibility, since the root v cannot simultaneously be positive and negative. So we have

Corollary 3. If $w \in W$ and $a \in \Pi$ then

$$l(wr_a) = \begin{cases} l(w) + 1 & \text{if } wa \in \Phi^+ \\ l(w) - 1 & \text{if } wa \in \Phi^-. \end{cases}$$

Observe that if $v \in \Phi$ and $a \in \Pi$ then when v and $r_a v$ are expressed as linear combinations of the basis vectors Π then the two are the same in all coefficients but the coefficient of a (because $v - r_a v = 2(v \cdot a)a$). So provided v is not a scalar multiple of a there is at least one nonzero coefficient in v which is the same as the corresponding coefficient in $r_a v$. So the roots v and $r_a v$ are either both positive or both negative. Since $v \cdot v = 1$ whenever v is a root it follows that the only scalar multiples of a which are roots are a and $-a$. Hence it follows that r_a permutes the positive roots other than a.

Definition. For $w \in W$ define $N(w) = \{\, v \in \Phi^+ \mid wv \in \Phi^- \,\}$

Our discussion above showed that if $a \in \Pi$ then $N(r_a) = \{a\}$. This combined with Corollary 3 above yields that if $l(wr_a) > l(w)$ then $N(wr_a) = r_a N(w) \cup \{a\}$, and, by induction, if $w = r_{a_1} r_{a_2} \cdots r_{a_l}$ with $l = l(w)$ and $a_i \in \Pi$ for all i, then $N(w) = \{a_l, r_{a_l} a_{l-1}, r_{a_l} r_{a_{l-1}} a_{l-2}, \ldots, r_{a_l} r_{a_{l-1}} \cdots r_{a_2} a_1\}$. In particular, $|N(w)| = l(w)$.

For each $J \subseteq \Pi$ we define W_J to be the subgroup of W generated by $\{ r_a \mid a \in J \}$. These subgroups are known as the *standard parabolic subgroups* of W. A *parabolic subgroup* is any conjugate of a W_J. To conclude this lecture we show that the parabolic subgroups are themselves Coxeter groups.

Proposition. *If $J \subseteq \Pi$ then W_J is isomorphic to the Coxeter group*

$$\widetilde{W}_J = \langle \{ s_a \mid a \in J \} \mid s_a^2 = 1 \text{ and } (s_a s_b)^{n_{ab}} = 1 \text{ for all } a,\, b \in J \} \rangle.$$

Before proving this let us look at what is being asserted in a special case. Let

$$W = \langle r, s, t, u \mid 1 = r^2 = s^2 = t^2 = u^2 = (rs)^3 = (rt)^8 = (ru)^2 = (st)^2 = (su)^6 = (rt)^5 \rangle,$$

(that is, W is the group corresponding to the diagram 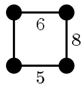

), and let $W_{\{r,s,t\}}$ be the parabolic subgroup of W generated by r, s and t. Let

$$\widetilde{W} = \langle r, s, t \mid 1 = r^2 = s^2 = t^2 = (rs)^3 = (rt)^8 == (st)^2 = (rt)^5 \rangle.$$

It is clear that there is a homomorphism from $\widetilde{W}$ onto $W_{\{r,s,t\}}$, since the defining relations of $\widetilde{W}$ are satisfied in W. It is not clear—although this is what the proposition asserts—that the map is an isomorphism. The relations in W which involve u could have some consequences which do not involve u, since it is conceivable that u could be eliminated from some pair of equations, yielding some relation that holds in $W_{\{r,s,t\}}$ and is not a consequence of the defining relations of $\widetilde{W}$. In general, for an arbitrary presentation of a group G, one would not expect to obtain a presentation for the subgroup generated by a subset of the generators of G by simply taking those defining relations which involve only these generators. But, remarkably, it does work for Coxeter groups, and the machinery we now have at hand makes it easy to prove.

Proof of the proposition. Let $\alpha\colon \widetilde{W}_J \to W_J$ be the epimorphism determined by $s_a \mapsto r_a$ for all $a \in J$. Let $\rho\colon W \to \mathrm{GL}(V)$ be the representation of W which we have been considering, and which is known (by Corollary 1) to be faithful. Let $\widetilde{\rho}\colon \widetilde{W}_J \to \mathrm{GL}(\widetilde{V})$ be the analogous representation of the Coxeter group $\widetilde{W}_J$, also faithful, for the same reason. The space $\widetilde{V}$ has a basis $\{ \tilde{a} \mid a \in J \}$, and there is a scalar product defined on $\widetilde{V}$ by $\tilde{a} \cdot \tilde{b} = -\cos(\pi/n_{ab})$ for all a, $b \in \Pi$. But this shows that if V_J is the subspace of V spanned by J then the linear isomorphism $\widetilde{V} \to V_J$ such that $\tilde{a} \mapsto a$ for all $a \in J$ in fact preserves the scalar product (since $\tilde{a} \cdot \tilde{b} = -\cos(\pi/n_{ab} = a \cdot b$ for all a, $b \in J$). Let $\phi\colon \mathrm{GL}(\widetilde{V}) \to \mathrm{GL}(V_J)$ be the isomorphism induced by this isomorphism $\widetilde{V} \to V_J$.

If a, $b \in J$ then $r_a b \in V_J$. So for all $a \in J$ we have that $r_a v \in V_J$ whenever $v \in V_J$. Hence the subspace V_J of V is W_J-invariant, and so there is a homomorphism $W_J \to \mathrm{GL}(V_J)$ which takes $w \in W_J$ to the restriction of $\rho(w)$. So we have a diagram

$$\begin{tikzcd}
\widetilde{W}_J \arrow[r, "\alpha"] \arrow[d, "\widetilde{\rho}"] & W_J \arrow[d, "\rho|_{V_J}"] \\
\mathrm{GL}(\widetilde{V}) \arrow[r, "\phi"] & \mathrm{GL}(V_J).
\end{tikzcd}$$

Check that it commutes! The result follows, since the left side and bottom are injective, and this forces α to be injective.

(To prove the commutativity of the diagram it is enough to look at the generators s_a ($a \in J$) of $\widetilde{W}_J$. Now $\widetilde{\rho}(s_a)$ is the transformation of $\widetilde{V}$ given by $\tilde{b} \mapsto \tilde{b} - 2(\tilde{b} \cdot \tilde{a})\tilde{a}$ for all $b \in J$, and the isomorphism ϕ takes this to the transformation of V_J given by $b \mapsto b - 2(\tilde{b} \cdot \tilde{a})a = b - 2(b \cdot a)a$. Since this is the restriction of $\rho(r_a)$, commutativity is proved.) $\square$

In this lecture we will prove the following remarkable property of Coxeter groups:

Theorem. *Suppose that W is any Coxeter group of finite rank, and H is a finite subgroup of W. Then H is a subgroup of some finite parabolic subgroup W_J of W.*

Notes.

1. This theorem is an exercise in Bourbaki; hence it is well known (perhaps). A proof is offered in the paper [1] of the lecturer and his student Brigitte Brink, but we made the mistake of not following Bourbaki's hints closely enough and wound up with an argument which is not valid if the scalar product on V is degenerate. The problem is easily avoided by working with the dual space.
2. The finite standard parabolic subgroups of W are easily found by inspection of the diagram of W and perusal of the list of finite Coxeter groups given in Lecture 2.
3. The theorem applies in particular to cyclic subgroups of W, and the conclusion is that the only elements of finite order are the elements of the finite parabolic subgroups.

We need a few preliminaries before starting the proof of the theorem. Firstly, we have

Lemma. *The group W is finite if and only if the root system Φ is finite.*

Proof. If $|W| < \infty$ then $|\Pi| < \infty$ and

$$|\Phi| = \big|\{\, wa \mid w \in W,\ a \in \Pi \,\}\big| \leq |W|\,|\Pi| < \infty.$$

Conversely, assume that $|\Phi| = M < \infty$. For each $w \in W$ define $\sigma_w\colon \Phi \to \Phi$ by $\sigma_w(v) = wv$ for all $v \in \Phi$. Then σ_w is a permutation of Φ, and furthermore $w \mapsto \sigma_w$ is a homomorphism $\sigma\colon W \to \mathrm{Sym}(\Phi)$, the symmetric group on Φ. Now if $w \in \ker\sigma$ then $wv = v$ for all $v \in \Phi$, which in turn implies that $wa = a$ for all $a \in \Pi$, and $wv = v$ for all $v \in V$ since Π spans V. But by Corollary 1 from Lecture 4 this implies that $w = 1$, and so we conclude that σ is injective. Hence

$$|W| \leq |\mathrm{Sym}(\Phi)| = M! < \infty$$

as required. $\square$

Let V^* be the space of all linear functions $V \to \mathbb{R}$, and for each $a \in \Pi$ define $\delta_a \in V^*$ by

$$\delta_a(b) = \begin{cases} 1 & \text{if } b = a \\ 0 & \text{otherwise} \end{cases} \qquad \text{for all } b \in \Pi.$$

Thus $\{\, \delta_a \mid a \in \Pi \,\}$ is the basis of V^* dual to the basis Π of V.

There is a natural right action of W on V^* as follows: if $w \in W$ and $\alpha \in V^*$ then $\alpha w \in V^*$ is defined by

$$(\alpha w)v = \alpha(wv) \qquad \text{for all } v \in V.$$

Note that if the scalar product on V were nondegenerate it would be possible to find elements $\hat{a} \in V$, one for each $a \in \Pi$, such that

$$\hat{a} \cdot b = \begin{cases} 1 & \text{if } b = a \\ 0 & \text{otherwise} \end{cases} \qquad \text{for all } b \in \Pi.$$

This would enable us to identify V^* with V via $\delta_a = \hat{a}$, and we would be able to avoid mention of the dual space. But the scalar product is not always nondegenerate.

If a, b, $c \in \Pi$ then

$$\begin{aligned}(\delta_b r_a)c &= \delta_b(c - 2(c\cdot a)a) \\ &= \delta_b c - 2(c\cdot a)\delta_b a \\ &= \begin{cases} \delta_b c & \text{if } b \neq a \\ (\delta_a - \sum_{d\in\Pi} 2(d\cdot a)\delta_d)c & \text{if } b = a, \end{cases}\end{aligned}$$

and since this holds for all $c \in \Pi$ we conclude that for all a, $b \in \Pi$,

$$\delta_b r_a = \begin{cases} \delta_b & \text{if } b \neq a \\ -\delta_a - \sum\limits_{d \neq a} 2(d \cdot a)\delta_d & \text{if } b = a. \end{cases}$$

For each $\alpha \in V^*$ let $\mathrm{Neg}(\alpha) = \{\, v \in \Phi^+ \mid \alpha(v) < 0 \,\}$, and define

$$C = \{\, \alpha \in V^* \mid \mathrm{Neg}(\alpha) \cap \Pi = \emptyset \,\}.$$

That is, C consists of all $\alpha \in V^*$ such that $\alpha(a) \geq 0$ for all simple roots α. If $\alpha \in C$ then clearly $\alpha(v) \geq 0$ for all $v \in \Phi^+$, since $\Phi^+ \subseteq \mathrm{PLC}(\Pi)$. Conversely, if $\alpha(v) \geq 0$ for all $v \in \Phi^+$ then *a fortiori* $\alpha(a) \geq 0$ for all $a \in \Pi$. So we have that

$$C = \{\, \alpha \in V^* \mid \alpha(v) \geq 0 \text{ for all } v \in \Phi^+ \,\} = \{\, \alpha \in V^* \mid \mathrm{Neg}(\alpha) = \emptyset \,\}.$$

Observe also that if $\alpha = \sum_{a \in \Pi} \lambda_a \delta_a \in V^*$ then $\alpha(b) = \lambda_b$ for each $b \in \Pi$. Thus

$$C = \{ \sum_{a \in \Pi} \lambda_a \delta_a \in V^* \mid \lambda_a \geq 0 \text{ for all } a \,\}.$$

Proposition. *Let $\alpha \in C$ and let $J = \{\, a \in \Pi \mid \alpha(a) = 0 \,\}$. Then*

$$\{\, w \in W \mid \alpha w = \alpha \,\} = W_J = \{\, w \in W \mid \alpha w \in C \,\}.$$

Proof. Let $\alpha = \sum_{a \in \Pi} \lambda_a \delta_a$. Since $\lambda_a = \alpha(a) = 0$ for $a \in J$, we have $\alpha = \sum_{b \in \Pi \setminus J} \lambda_b \delta_b$. Now if $a \in J$ we have $\delta_b r_a = \delta_b$ for all $b \notin J$, and so $\alpha r_a = \alpha$. By induction on the length of w we conclude that $\alpha w = \alpha$ for all $w \in W_J$. So $W_J \subseteq \{\, w \in W \mid \alpha w = \alpha \,\}$.

Since it is trivial that $\{\, w \in W \mid \alpha w = \alpha \,\} \subseteq \{\, w \in W \mid \alpha w \in C \,\}$, it remains only to prove that $\{\, w \in W \mid \alpha w \in C \,\} \subseteq W_J$. So let $w \in W$ and suppose that $\alpha w = \alpha' \in C$. If $w \neq 1$ we may choose $a \in \Pi$ with $l(w^{-1} r_a) = l(w^{-1}) - 1$, and write $w^{-1} r_a = w_1^{-1}$. Then $w_1^{-1}(a) \in \Phi^+$, and

$$0 \leq \alpha(a) = (\alpha' w^{-1})(a) = (\alpha' w_1^{-1} r_a)(a) = -\alpha'(w_1^{-1} a) \leq 0$$

since $\alpha' \in C$. It follows that $\alpha(a) = 0$, and so $a \in J$. Hence, by the first part of this proof, $\alpha r_a = \alpha$. Therefore

$$\alpha w_1 = \alpha r_a w_1 = \alpha w = \alpha',$$

and, furthermore, $W_J w_1 = W_J r_a w_1 = W_J w$. Now if $w_1 \neq 1$ we may repeat the argument and find $w_2 \in W$ with $l(w_2^{-1}) = l(w_1^{-1}) - 1$ such that $\alpha w_2 = \alpha'$ and $W_J w_2 = W_J w_1 = W_J w$. And then if $w_2 \neq 1$ we can repeat the argument again, and so construct a sequence w, w_1, w_2, w_3, ... of elements of W. Since the length decreases at each step the sequence must terminate; however, the process can be continued so long as $w_i \neq 1$. Thus $w_i = 1$ for some i, and

$$W_J w = W_J w_1 = W_J w_2 = \cdots = W_J w_i = W_J,$$

proving that $w \in W_J$, as required. □

In the case that $|W| < \infty$, every orbit of W on V^* contains an element of C, and so it follows that the stabilizer in W of an element of C is always a parabolic subgroup. This is no longer true in the infinite case; instead we have

Proposition. $\{\, \alpha w \bigm| \alpha \in C \text{ and } w \in W \,\} = \{\, \beta \in V^* \bigm| |\mathrm{Neg}(\beta)| < \infty \,\}$.

Proof. Let $\alpha \in C$ and $w \in W$. If $v \in \Phi^+$ then $(\alpha w)(v) = \alpha(wv)$ is nonnegative if $wv \in \Phi^+$, and so

$$\mathrm{Neg}(\alpha w) = \{\, v \in \Phi^+ \mid \alpha(wv) < 0 \,\} \subseteq \{\, v \in \Phi^+ \mid wv \in \Phi^- \,\} = N(w),$$

a finite set. So $\{\, \alpha w \mid \alpha \in C \text{ and } w \in W \,\} \subseteq \{\, \beta \in V^* \mid |\mathrm{Neg}(\beta)| < \infty \,\}$.

Conversely, let $\beta \in V^*$ be such that set $\mathrm{Neg}(\beta)$ is finite; we show that $\beta = \alpha w$ for some $\alpha \in C$ and $w \in W$. If $\mathrm{Neg}(\beta) = \emptyset$ then $\beta \in C$, and so we may take $\alpha = \beta$ and $w = 1$. Otherwise $\beta \notin C$, and so $\mathrm{Neg}(\beta) \cap \Pi \neq \emptyset$, and we may choose $a \in \Pi$ with $\beta(a) < 0$. Observe that $a \notin \mathrm{Neg}(\beta r_a)$, since $(\beta r_a)(a) = -\beta(a) > 0$. But r_a permutes $\Phi^+ \backslash \{a\}$, and if $v \in \Phi^+ \backslash \{a\}$ then $v \in \mathrm{Neg}(\beta r_a)$ if and only if $(\beta r_a)(v) = \beta(r_a v)$ is negative, which holds if and only $r_a v \in \mathrm{Neg}(\beta)$. So $\mathrm{Neg}(\beta r_a) = r_a(\mathrm{Neg}(\beta) \backslash \{a\})$ has fewer elements than $\mathrm{Neg}(\beta)$, and arguing by induction on the size of this set we deduce that $\beta r_a = \alpha w'$ for some $\alpha \in C$ and $w' \in W$. Thus $\beta = \alpha w$ where $w = w' r_a \in W$, as required. □

Proof of the theorem. Choose a counterexample to the theorem for which the rank of the Coxeter group W is as small as possible. Then W has a finite subgroup H which is not contained in any finite parabolic subgroup of W. In particular, the group W itself must be infinite. By the Lemma, Φ is infinite also.

Let $\alpha = \sum_{a \in \Pi} \lambda_a \delta_a$, where the coefficients λ_a are positive numbers (arbitrarily chosen). That is, $\alpha(a) > 0$ for all $a \in \Pi$. Then $\alpha \in C$—indeed, α is in the interior of C—since $\alpha(v)$ is strictly positive for all $v \in \Phi^+$. Note as a consequence that $\alpha(v) \neq 0$ for all $v \in \Phi$.

Let $\beta = \sum_{h \in H} \alpha h$. By the proposition proved above, $\mathrm{Neg}(\alpha h)$ is a finite set for each $h \in H$, and so the set $\Omega = \bigcup_{h \in H} \mathrm{Neg}(\alpha h)$ is finite also. If $v \in \Phi^+ \backslash \Omega$ then $v \in \Phi^+ \backslash \mathrm{Neg}(\alpha h)$ for all $h \in H$, and since $(\alpha h)(v) = \alpha(hv) \neq 0$ it follows that $(\alpha h)(v) > 0$ for all $h \in H$. Thus for all $v \in \Phi^+ \backslash \Omega$,

$$\beta(v) = \sum_{h \in H} (\alpha h)(v) > 0,$$

and so $v \notin \mathrm{Neg}(\beta)$. Thus $\mathrm{Neg}(\beta)$ is finite, and (by the proposition again), $\beta = \gamma w$ for some $\gamma \in C$ and $w \in W$. Furthermore, $\Phi^+ \backslash \Omega \neq \emptyset$ since Φ^+ is infinite; hence $\beta \neq 0$.

Note that, for all $k \in H$,

$$\beta k = \sum_{h \in H} \alpha(hk) = \sum_{g \in H} \alpha g = \beta$$

and therefore

$$\gamma(wkw^{-1}) = (\gamma w)kw^{-1} = \beta k w^{-1} = \beta w^{-1} = \gamma.$$

Hence $wkw^{-1} \in W_J$, where $J = \{\, a \in \Pi \mid \gamma(a) = 0 \,\}$. Note that if $J = \Pi$ then $\gamma = 0$, which is impossible since $\gamma w = \beta \neq 0$. So $H_1 = wHw^{-1} \leq W_J$, which is a Coxeter group of smaller rank than W. By minimality of our counterexample it follows that there is a $K \subseteq J$ such that W_K is finite and H_1 is contained in a conjugate of W_K. Hence $H = w^{-1} H_1 w$ is also contained in a conjugate of the finite parabolic subgroup W_K, contradicting our original hypothesis. □

Lecture 6

This lecture will be devoted to a brief discussion of the results of [1], in which it is shown that Coxeter groups of finite rank are automatic. The key concept need for the proof of this fact is a relation, which we call *dominance*, on the set of positive roots.

Definition. If u, $v \in \Phi^+$ with $u \neq v$ then u *dominates* v if $\{\, w \in W \mid wu \in \Phi^- \,\} \subseteq \{\, w \in W \mid wv \in \Phi^- \,\}$.

That is, u dominates v if every element of W which takes u to a negative root also takes v negative. If $u \in N(w)$, then $v \in N(w)$.

We define also $\Delta = \{\, u \in \Phi^+ \mid u \text{ dominates some } v \in \Phi^+ \backslash \{u\} \,\}$, and $\Delta' = \Phi^+ \backslash \Delta$. We remark that when W is finite the set Δ is empty.

I had dreamt up these concepts many years ago, but had no application for them before encountering the automaticity question. I had also privately conjectured the following theorem, but I needed Brigitte's help to prove it.

Theorem. For all Coxeter groups of finite rank, $|\Delta'| < \infty$.

The theorem says, in other words, that there are only finitely many positive roots with the property of not dominating anything. In my view the theorem is an important fact about Coxeter groups, and is the main result of [1], although the consequence, that Coxeter groups are automatic, seems to have attracted more interest.

Example

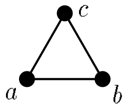

Since $a \cdot b = b \cdot c = c \cdot a = -\frac{1}{2}$ and $a \cdot a = b \cdot b = c \cdot c = 1$, we find that

$$\begin{array}{lll} r_a(a) = -a & r_b(a) = a + b & r_c(a) = a + c \\ r_a(b) = b + a & r_b(b) = -b & r_c(b) = b + c \\ r_a(c) = c + a & r_b(c) = c + b & r_c(c) = -c. \end{array}$$

Note that $r_a(a+b+c) = -a+(b+a)+(c+a) = a+b+c$, and since the same holds for r_b and r_c we conclude that if $\rho = a + b + c$ then $w\rho = \rho$ for all $w \in W$. We can now easily describe Φ^+. Firstly, a, b, $c \in \Phi^+$, of course. Applying simple reflections to these we find that $a + b$, $b + c$, $c + a \in \Phi^+$. Now applying simple reflections to these yields that $a + \rho$, $b + \rho$ and $c + \rho$ are positive roots. But since ρ is fixed by all elements of W it is clear what happens from here. We will find that

$$\Phi^+ = \{\, v + n\rho \mid n \text{ is a nonnegative integer and } \alpha \in \{a, b, c, a+b, b+c, c+a\} \,\}.$$

But now it is easy to see that $\Delta' = \{a, b, c, a+b, b+c, c+a\}$, for if $w \in W$ and $w(\alpha + n\rho)$ is negative, where $\alpha \in \{a, b, c, a+b, b+c, c+a\}$ and $n \geq 1$, then $w(\alpha)$ must be negative also, since otherwis $w(\alpha+n\rho) = w(\alpha)+n\rho$ would be a positive linear combination of Π. So in this example $|\Delta'| = 6$.

Before moving on to harder things, let us prove a result which is a triviality, but nevertheless illustrates the kind of use these ideas can be put to.

Proposition. *Let w_1, $w_2 \in W$, and suppose that $N(w_1) \cap N(w_2) \neq \emptyset$. Then $N(w_1) \cap N(w_2) \cap \Delta' \neq \emptyset$.*

Proof. Let $v \in N(w_1) \cap N(w_2)$. If $v \in \Delta$ then there exists $v' \in \Phi^+$ such that v dominates v'. Since $v \in N(w_i)$ we have $v' \in N(w_i)$, for each i. That is, $v' \in N(w_1) \cap N(w_2)$. Now if $v' \in \Delta$ we can repeat the argument, and we can continue to do so as long as we have not encountered an element of $N(w_1) \cap N(w_2) \cap \Delta'$. We obtain in this way a sequence of elements of $N(w_1) \cap N(w_2)$ which are necessarily all distinct from each other (as follows, for example, from the next lemma we shall prove). Since the sets $N(w_i)$ are finite, a point must be reached at which the sequence cannot be continued; that is, an element of $N(w_1) \cap N(w_2) \cap \Delta'$ must eventually be found. □

We define the *depth* of a positive root v, denoted by $\mathrm{dp}(v)$, to be the least l such that wv is negative for some $w \in W$ of length l. In fact, $\mathrm{dp}(v) = \frac{1}{2}(l(r_v) + 1)$.

Lemma. *If u dominates v then $\mathrm{dp}(u) > \mathrm{dp}(v)$.*

Proof. Choose $w \in W$ with $l(w) = \mathrm{dp}(u)$ and $wu \in \Phi^-$. Write $w = r_a w'$, where $a \in \Pi$ and $l(w') = l(w) - 1$. Then $w'u$ is positive and $r_a(w'u)$ is negative; so it follows that $w'u = a$. Now $wv \in \Phi^-$ since u dominates v, and therefore $r_a(w'v) \in \Phi^-$. But $w'v \neq w'u = a$ since $v \neq u$; hence $w'v$ is negative, and $\mathrm{dp}(v) \leq l(w') < \mathrm{dp}(u)$. □

If W is finite then there exists an element $w_0 \in W$ such that $w_0(\Phi^+) = \Phi^-$. It can be seen that $v \mapsto -w_0 v$ is a depth preserving and dominance reversing permutation of Φ^+; thus the lemma just proved shows (as we remarked earlier) that there can be no dominance in a finite root system.

Let $w = r_{a_1} r_{a_2} \cdots r_{a_n}$ where $a_i \in \Pi$ for all i. If $n = l(w)$, so that there is no shorter expression for w as a product of simple reflections, then we say that the expression $r_{a_1} r_{a_2} \cdots r_{a_n}$ is *reduced*.

To avoid second order subscripts we will in future write these expressions as $a_1 a_2 \cdots a_n$ instead of $r_{a_1} r_{a_2} \cdots r_{a_n}$.

Theorem. *There exists a finite state automaton which can recognize whether expressions are reduced.*

We describe a suitable automaton below. The theorem is part of what is needed to prove that W is automatic. Also needed are more finite state automata to deal with multiplying reduced expressions by fundamental reflections; these automata are slightly more complicated than the one we shall describe, but based on the same fundamental property, namely, the finiteness of Δ'.

A finite state automaton consists of a finite set $\mathcal{S}$ (the *states* of the automaton) together with a transition function $\tau\colon \mathcal{S} \times \Pi \to \mathcal{S}$. There is a distinguished element $S_0 \in \mathcal{S}$, called the *initial state*, and each state is either an *accept state* or a *reject state*. The idea is that, starting at the state S_0, the automaton reads an expression $a_1 a_2 \cdots a_n$ one symbol at a time, starting from the left, and moves from state to state by the rule $S_i = \tau(S_{i-1}, a_i)$. That is, the value of the transition function for the current state and the symbol being read gives determines the state to move to. When the end of the expression is reached the automaton reports "accept" or "reject" depending on the final state. The theorem says that such an automaton exists, having the property that the reduced expressions are accepted and non-reduced expressions rejected.

We start by describing an infinite state automaton which recognizes reduced expressions. The states of this will be subsets of Φ^+, all of which will be accept states, with the empty subset $\emptyset$ as the initial state, together with one reject state, which we will call $\underline{\mathrm{No}}$. That is

$$\mathcal{S} = \mathcal{P}(\Phi^+) \mathbin{\dot\cup} \{\underline{\mathrm{No}}\}$$

(using $\mathcal{P}(X)$ to denote the power set of the set X). Define the transition function $\tau\colon S \times \Pi \to S$ as follows: for all $a \in \Pi$,

$$(*) \qquad \begin{aligned} \tau(\underline{\mathrm{No}}, a) &= \underline{\mathrm{No}} \\ \tau(S, a) &= \begin{cases} \underline{\mathrm{No}} & \text{if } a \in S \\ r_a S \cup \{a\} & \text{if } a \notin S. \end{cases} \end{aligned}$$

Note that once the reject state is entered it is never left. This is reasonable, since an expression $a_1 a_2 \cdots a_n$ cannot be reduced if it has an initial segment $a_1 a_2 \cdots a_i$ which is not reduced.

Consider now an expression $a_1 a_2 \cdots a_n$, and for each i from 0 to n let w_i be the element of W corresponding to the expression $a_1 a_2 \cdots a_i$. Let S_i be the state of the automaton after reading a_i. We use induction on i to prove the following statement:

$$S_i = \begin{cases} \underline{\mathrm{No}} & \text{if } a_1 a_2 \cdots a_i \text{ is not reduced} \\ N(w_i) & \text{if } a_1 a_2 \cdots a_i \text{ is reduced.} \end{cases}$$

This is true when $i = 0$ since $N(1) = \emptyset = S_0$. Now suppose $i > 0$. If $a_1 a_2 \cdots a_{i-1}$ is not reduced then $a_1 a_2 \cdots a_i$ is not reduced, $S_{i-1} = \underline{\mathrm{No}}$ by the inductive hypothesis and $S_i = \tau(S_{i-1}, a_i) = \underline{\mathrm{No}}$ by $(*)$; so our statement holds in this case. So assume that $a_1 a_2 \cdots a_{i-1}$ is reduced. The inductive hypothesis gives $S_{i-1} = N(w_{i-1})$. If $a_i \in N(w_{i-1})$ then $l(w_i) < l(w_{i-1})$ and so $a_1 a_2 \cdots a_i$ is not reduced; furthermore $(*)$ gives $S_i = \tau(S_{i-1}, a_i) = \underline{\mathrm{No}}$ since $a_i \in S_{i-1}$. On the other hand, if $a_i \notin N(w_{i-1})$ then $a_1 a_2 \cdots a_i$ is reduced; moreover $N(w_i) = r_{a_i} N(w_{i-1}) \cup \{a_i\} = r_{a_i} S_{i-1} \cup \{a_i\} = \tau(S_{i-1}, a_i) = S_i$. This shows that the statement holds in all cases, and completes the induction.

However, to prove the theorem we need a finite state automaton, and the above automaton definitely has infinitely many states if W is infinite, since Φ^+ has infinitely many subsets in this case. The automaton could be modified by eliminating the states $S \subseteq \Pi$ which are not of the form $N(w)$ for some $w \in W$, but that would still leave infinitely many states since no two distinct elements $w \in W$ have the same $N(w)$. Instead, we modify the automaton by replacing the infinite set Φ^+ by the finite set Δ', defining

$$\mathcal{S} = \mathcal{P}(\Delta') \cup \{\underline{\mathrm{No}}\}.$$

The modified definition of the transition function is

$$\begin{aligned} \tau(\underline{\mathrm{No}}, a) &= \underline{\mathrm{No}} \\ \tau(S, a) &= \begin{cases} \underline{\mathrm{No}} & \text{if } a \in S \\ (r_a S \cup \{a\}) \cap \Delta' & \text{if } a \notin S. \end{cases} \end{aligned}$$

In other words, we simply throw away elements of Δ as they occur. This idea is based on the following lemma.

Lemma. *If $v \in \Delta'$ and $a \in \Pi$ are such that $r_a v \in \Delta$, then $r_a v$ dominates a.*

Proof. Suppose that $r_a v$ dominates $u \in \Phi^+$. Let $w \in W$ be such that $wv \in \Phi^-$. Then $(wr_a)(r_a v)$ is negative, and so $(wr_a)u$ is negative also (since $r_a v$ dominates u). So we have shown that $w(r_a u)$ is negative whenever wv is negative, and if $r_a u$ were positive this would say that v dominates $r_a u$. But v does not dominate anything, since $v \in \Delta'$, and so $r_a u$ must be negative. Thus $u \in N(r_a) = \{a\}$; so $u = a$, and $r_a v$ dominates a, as claimed. □

The lemma permits us to prove, by an induction very similar to the one we did before, that

$$S_i = \begin{cases} \underline{\text{No}} & \text{if } a_1 a_2 \cdots a_i \text{ is not reduced} \\ N(w_i) \cap \Delta' & \text{if } a_1 a_2 \cdots a_i \text{ is reduced.} \end{cases}$$

In the case that w_{i-1} is reduced but w_i is not, we know that $a_i \in N(w_{i-1})$, and so $a_i \in N(w_{i-1}) \cap \Delta' = S_{i-1}$, since the simple roots are certainly all in Δ'. So in this case $S_i = \tau(S_{i-1}, a_i) = \underline{\text{No}}$, which is right. If w_{i-1} and w_i are both reduced then

$$S_i = (r_{a_i} S_{i-1} \cup \{a_i\}) \cap \Delta' = (r_{a_i} S_{i-1} \cap \Delta') \cup \{a_i\},$$

while

$$N(w_i) \cap \Delta' = (r_{a_i} N(w_{i-1}) \cup \{a_i\}) \cap \Delta' = (r_{a_i} N(w_{i-1}) \cap \Delta') \cup \{a_i\}.$$

Obviously $S_i \subseteq N(w_i) \cap \Delta'$, since $S_{i-1} \subseteq N(w_{i-1})$ by the inductive hypothesis. If $v \in N(w_i) \cap \Delta'$ and $v \notin S_i$ then we must have $v \in r_{a_i} N(w_{i-1}) \cap \Delta'$ and $v \notin r_{a_i} S_{i-1}$. But the inductive hypothesis says that $S_{i-1} = N(w_{i-1}) \cap \Delta'$; so $v \notin r_{a_i} N(w_{i-1}) \cap r_{a_i} \Delta'$, and hence $v \notin r_{a_i} \Delta'$ since we know that $v \in r_{a_i} N(w_{i-1})$. Hence we have the situation described in the lemma: $v \in \Delta'$ and $r_{a_i} v \in \Delta$. By the lemma it follows that $r_{a_i} v$ dominates a_i. But $r_{a_i} v \in N(w_{i-1})$, and so we conclude that $a_i \in N(w_{i-1})$ also. But this contradicts the assumption that w_i is reduced. So $S_i = N(w_i) \cap \Delta'$, and the induction goes through.

The idea is perhaps better illustrated by the following table, in which the ith column gives the elements of $N(w_i)$, so long as $a_1 a_2 \cdots a_i$ is reduced. (We have abbreviated r_{a_i} to r_i.)

$$\begin{array}{cccccc}
a_1 & r_2(a_1) & r_3 r_2(a_1) & r_4 r_3 r_2(a_1) & \cdots & r_n \cdots r_5 r_4 r_3 r_2(a_1) \\
 & a_2 & r_3(a_2) & r_4 r_3(a_2) & \cdots & r_n \cdots r_5 r_4 r_3(a_2) \\
 & & a_3 & r_4(a_3) & \cdots & r_n \cdots r_5 r_4(a_3) \\
 & & & a_4 & \cdots & r_n \cdots r_5(a_4) \\
 & & & & \cdots & \vdots \\
 & & & & & a_n
\end{array}$$

The occurrence of any negative root in this would indicate that the expression $a_1 a_2 \cdots a_n$ is not reduced, and so we want the automaton to reject the expression in that case. If any of the roots in the table are in Δ then the automaton discards them, and all subsequent roots in the same row. But, by the lemma, if v is the first root in Δ in some row of the table, then v in fact dominates the simple root a_i at the bottom of its column. It follows that if a negative root subsequently appears in the row of v, then there is also a negative root, in the same column, in the row of a_i. So discarding roots in Δ does not inhibit the detection of negative roots.

To finish off, we give an outline of the proof of the main theorem (the finiteness of Δ'). The first step is to prove the following criterion for dominance:

Proposition. *If $u, v \in \Phi^+$ then u dominates v if and only if $\mathrm{dp}(u) > \mathrm{dp}(v)$ and $u \cdot v \geq 1$.*

Proof. Assume that u dominates v. By a previous lemma, $\mathrm{dp}(u) > \mathrm{dp}(v)$; hence have only to prove that $u \cdot v \geq 1$. Assume, for a contradiction, that $u \cdot v < 1$. It is easily shown that, under these circumstances, the two dimensional space P spanned by u and v is Euclidean. If the angle between u and v is not a rational multiple of π then the roots of the form $\lambda u + \mu v$ will be dense on the unit circle in P. This necessarily gives a root $\lambda u + \mu v$ for which λ and μ have opposite signs, and this contradicts our basic result about positive and negative roots. So the angle between u and v is a rational multiple of π, and the (dihedral) subgroup of W generated by r_u and r_v is finite. This contradicts the fact that there is no dominance in a finite root system.

Conversely, suppose that $u \cdot v = \lambda \geq 1$ and $\mathrm{dp}(u) > \mathrm{dp}(v)$. If we choose $w \in W$ of minimal length such that $wv \in \Pi$ then it is easily checked that $wu \in \Phi^+$ has depth greater than $1 = \mathrm{dp}(wv)$, and $wu \cdot wv = u \cdot v \geq 1$. Furthermore, u dominates v if and only if wu dominates wv. The upshot of this is that we may replace u and v by wu and wv, and assume that $v \in \Pi$. Now suppose that $x \in W$ has the property that xu is negative but xv is positive. Then $r_v u$ is a positive root, and

$$x(r_v u) = x(u - 2\lambda v) = xu + 2\lambda(-xv),$$

which is negative since $\lambda \geq 0$. Hence $N(x)$ contains u and $r_v u$, and hence all roots of the form $\alpha u + \beta r_v u$ with α, $\beta > 0$. But since $u \cdot r_v u = 1 - 2\lambda^2 \leq -1$ it can be shown that there are infinitely many such roots, and this contradicts the fact that $|N(x)| = l(x)$ is finite. □

Our proof of the theorem now proceeds along the following lines. If Δ' is not finite it will contain roots of arbitrarily large depth. So for any integer m there will exist a sequence v_1, v_2, $v_3, \ldots,$ v_m such that $v_i \in \Delta'$ for all i, the depth of v_i is i, and $v_i = r_{a_i} v_{i-1}$ for some $a_i \in \Pi$.

If $a \in \Pi$ and $v_i \cdot a \leq -1$ then $a_{i+1} \neq a$, since otherwise we would have $v_{i+1} \cdot a \geq 1$, which is impossible since v_{i+1} is not allowed to dominate a. It follows that $v_{i+1} \cdot a \leq -1$ also, and we deduce that $a_j \neq a$ for all $j > i$. Hence the coefficient of a in v_j (when v_j is expressed as a linear combination of the simple roots) remains fixed for $j > i$.

Fix a simple root $a \in \Pi$ and consider the sequence of numbers $v_i \cdot a$, for $i = 1,\ 2,\ \ldots,\ m$. These numbers are always less than 1, and once they fall below -1, then, as we have seen, the coefficient of a in v_i remains fixed. But while they remain in the interval $(-1, 1)$ they are constrained to lie in a finite set of values, consisting of numbers of the form $\cos(\frac{n}{k}\pi)$. This is because the angle between v_i and a, in the (Euclidean) subspace they span, has to be a rational multiple of π, since roots with mixed coefficients cannot occur. The denominator k is bounded since the subgroup generated by r_{v_i} and r_a is finite, and hence lies in a finite parabolic subgroup of W, and there are only finitely many conjugacy classes of these. If m is large enough then it follows from these considerations that there will be two distinct roots v_i and v_j in our sequence with the following property: for all $a \in \Pi$, either $v_i \cdot a$ and $v_j \cdot a$ are less than -1, and the coefficient of a in v_i is the same as its coefficient in v_j, or else $v_i \cdot a = v_j \cdot a$. The *coup de grace* is now delivered by the following lemma.

Lemma. *If $u = \sum_{a\in\Pi} \lambda_a a$ and $v = \sum_{a\in\Pi} \mu_a a$ are positive roots, and for all $a \in \Pi$ either $\lambda_a = \mu_a$ or $u \cdot a = v \cdot a$, then $u \cdot v = 1$.*

Proof. We have that $u - v = \sum \nu_a a$, with the sum ranging of those $a \in \Pi$ such that $u \cdot a = v \cdot a$. Hence

$$u \cdot (u - v) = v \cdot (u - v).$$

which rearranges to give $2(u \cdot v) = u \cdot u + v \cdot v$. Since $u \cdot u = v \cdot v = 1$, the result follows. □

The preceding lemma and proposition combine to show that either v_i dominates v_j or *vice versa.* But this is impossible, since v_i and v_j are both meant to be in Δ'. This places an upper bound on m (the length of our increasing sequence in Δ'), which in turn bounds the size of Δ' and completes the proof of the theorem.

A final remark: the bound that we obtain for the size of Δ' is an order of magnitude or two greater than its actual size in practice. It would be interesting to obtain more accurate information on the size of Δ'.

Reference

[1] Brink, B. and Howlett, R. "A finiteness property and an automatic structure for Coxeter groups" *Math. Ann.* **296** (1993) 179–190.