

Computation in Coxeter groups II. Constructing minimal roots

Bill Casselman
Mathematics Department
University of British Columbia
Canada
cass@math.ubc.ca

Abstract. *In the recent paper (Casselman, 2001) I described how a number of ideas due to Fokko du Cloux and myself could be incorporated into a reasonably efficient program to carry out multiplication in arbitrary Coxeter groups. At the end of that paper I discussed how this algorithm could be used to build the reflection table of minimal roots, which could in turn form the basis of a much more efficient multiplication algorithm. In this paper, following a suggestion of Robert Howlett, I explain how results due to Brigitte Brink can be used to construct the minimal root reflection table directly and more efficiently.*

Introduction

A **Coxeter system** is a pair (W, S) where W is a group with a set of generators S and relations

$$(st)^{m_{s,t}} = 1$$

for pairs s, t in S , with the conditions that $m_{s,s} = 1$ and $2 \leq m_{s,t} \leq \infty$ for $s \neq t$. In this paper I shall always assume S to be finite. Equivalent to these are the relations

$$s^2 = 1 \\ sts \dots = tst \dots \quad (m_{s,t} \text{ terms on each side})$$

The second type are called *braid relations*. Often the group alone is referred to, but almost always when one refers to a Coxeter group one has in mind a particular set S . The relations involved here are very simple, but these systems play an extremely important role in mathematics, one far from visible in the bare definition. Among other things, the Weyl groups of Kac-Moody Lie algebras possess a natural structure as a Coxeter system. Much work has been done on these structures, particularly on finite and affine Weyl groups, but many phenomena involving them remain unexplained, and it is likely that computer explorations will be even more significant in the future than they have been so far. These groups become extremely complex as the size of S grows, however, and computer programs dealing with them must be extremely efficient to be useful. Conventional symbolic algebra packages normally fail to handle the difficulties satisfactorily. This paper is the second of a series in which I describe what may be even in the long term the most efficient algorithms to do basic computations in arbitrary Coxeter groups, at least on serial machines. These algorithms depend strongly on mathematical results of Brigitte Brink, Robert Howlett, and Fokko du Cloux.

The most basic problems are those one encounters immediately: (1) how to represent elements of the group and (2) how to multiply two elements if they are given in whatever format has been chosen? These two problems are of course strongly related, and are by no means simple matters. For familiar examples such as S_n there are many satisfactory methods, and for the Weyl groups of Kac-Moody algebras one can use the representation of the group on the root lattice to reduce many problems to integer arithmetic. But arbitrary groups, even sometimes the exceptional finite groups H_3 and H_4 , are computationally more demanding. Even when working with Weyl groups one should be prepared to deal with very

large integers in arithmetic involving the root lattice. Solving computational problems in this domain leads to mathematically interesting questions as well as practical ones.

How to represent elements? The best combinatorial solutions—certainly, that of greatest theoretical interest—seem to be those in which elements of the group are represented as products of elements of S . More explicitly, impose an ordering on S . Suppose w to be equal to the product $s_1 s_2 \dots s_n$ as a product of generators in S , and for each i let w_i be the partial product $s_1 \dots s_i$. This expression for w is called its **InverseShortLex normal form** if for each i the element s_i is least in S such that $\ell(w_i s_i) < \ell(w_i)$. It is the **shortest** representation of w least in **lexicographical order** when read **in reverse**.

The multiplication problem can now be posed: *Given the normal form $\text{NF}(w)$ of w , what is $\text{NF}(sw)$?* The first general solution was implicit in work published (Tits, 1968). The main result of Tits' paper was that if one starts with a word in the generators and (a) applies the braid relations or (b) reduces ss to 1, one eventually gets all reduced words equivalent to the given one. The point is that one does not have to allow insertions of ss , which would increase the length of an expression and render the decision process potentially non-terminating. Tits' method of computing amounted essentially to listing all words of length no more than a given one that could be derived from it by elementary substitutions involving the basic Coxeter relations. The normal form could then be singled out from the list. Tits' algorithm had the virtue of being entirely combinatorial in nature, working only with words in S . Implementing it in an explicit program is an interesting exercise, but largely a futile one since it is extremely inefficient, requiring both time and storage space roughly exponentially proportional, or worse, to the length n of the normal form of w .

A second solution was proposed by Fokko du Cloux around 1990, but explicitly described only for finite Coxeter groups. It was this that I worked out in more detail in (Casselman, 2001). It is mildly recursive in nature, of unknown theoretical complexity but in practice not unreasonably slow. It also has the tremendous virtue of requiring little machine memory. If one wants to be able to carry out just a few multiplications, it is likely that this method cannot be improved on. But if one wants to carry out a great many multiplications for a given Coxeter group, then it is probably best to build an auxiliary data structure that will make multiplication much simpler. A structure perfectly suited to this purpose is a table describing how the generators of S act on the **minimal roots** of (Brink and Howlett, 1993). As will be recalled later on, it can be used in a multiplication algorithm of complexity essentially linear in the length of w , allowing multiplication to take place without backtracking in a left-to-right scan of the normal form of w .

In the paper (Casselman, 2001) I proposed that du Cloux's algorithm could be used as a kind of bootstrap to build this reflection table, but as I mentioned at the end of that paper it is not difficult to find groups where the number of minimal roots is rather small, hence the table itself also fairly small, but the task of building it in this way extremely tedious. This is unsatisfactory.

After I explained this method to Robert Howlett, he and I experimented with other tools, and after a short while he suggested that the results of Brigitte Brink's thesis seemed to be applicable. In this thesis, most of which has been summarized in her published articles, she describes rather explicitly what the set of minimal roots looks like for any Coxeter group. However, it was not at all apparent at first glance, at least to me, that what she does is computationally practical. Howlett showed that my first impression was false by producing a program in the symbolic algebra package *Magma* that builds the minimal root reflection table of an arbitrary Coxeter group. This implementation was much slower than subsequent programs written by Fokko du Cloux and myself, but already dealt easily with a previously difficult group described in my paper. All of these techniques depend on crucial results of Brink's thesis, but not on her more or less explicit lists, interesting though they may be.

This paper will explain in detail an elaboration of Howlett’s idea of how to apply the results of Brink’s thesis to practical programming in order to construct the minimal roots. The resulting program will not be the fastest possible, but it will be fairly fast. Howlett’s *Magma* program took a little over an hour to handle it, he tells me, whereas a program in *C* written by Fokko du Cloux took 10 seconds on a comparable machine. The version based on the technique to be given here lies in between the two—a version written in Java took a bit less than a minute on the same machine as that used by du Cloux. The principal drawback of the algorithm to be presented here as the basis of multiplication is that it claims an amount of machine memory proportional to the number of minimal roots. The number of minimal roots can grow quite fast with the size of S , so this is not a negligible objection. It thus exemplifies the usual programming trade-off between time and space. Familiar Coxeter groups are deceptive in this regard—for classical series of finite and affine groups the number of minimal roots is roughly proportional to $|S|^2$, whereas one of the first really difficult groups Howlett’s program dealt with was an exotic Coxeter group with $|S| = 22$ and several hundred thousand minimal roots. Groups for which the set of minimal roots is large are so complicated that there is not much exploration one can expect to do with them anyway, and in practice the number of minimal roots is only a relatively minor impediment.

I begin this paper recalling the basic properties of minimal roots; explain how they are to be used in multiplication; summarize results of Brigitte Brink’s thesis in the manner most useful for my purposes; and finally explain how to apply these results in practice to construct the minimal root reflection table. I should say that what I am going to come up with is somewhat different from Howlett’s program, and also quite different from a closely related program of Fokko du Cloux. It uses a few more advanced features of Brink’s thesis than Howlett’s program. It is less sophisticated than du Cloux’s program, but has the advantage of being possibly clearer in conception.

Once Howlett had suggested using the results of Brink’s thesis, the path to a practical program was in principle not too difficult. On the other hand, this path did not appear obvious to me until after a great deal of exploration, and it will surely be useful to place here a record of the outcome. This paper is largely self-contained; the order of exposition and its occasional geometric emphasis are new. My principal ambition in writing it has been to awaken interest in the work of Brink, Howlett, and du Cloux. This whole area of research is one of great charm, but it has apparently and unfortunately awakened little interest at large. This is likely due to the fact that it lies somewhere in the great no man’s land between pure mathematics and practical computation—*terra incognita* on most maps, even in the 21st century. I must also record that both my programs and this exposition benefited enormously from conversations with the late and very much lamented Fokko du Cloux. As I learned long ago to my own chagrin, competing directly in programming with him could be an embarrassing experience, but working alongside him was always, on the contrary, at once enjoyable and educational. With his death went a uniquely talented mathematician.

§§1–3 will recall some of the basic geometry of Coxeter groups, and §4 will recall how minimal roots can be used in multiplication. In §5 I start discussing the algorithm itself, with details presented in §§6–9.

1. The geometry of roots

The algorithms to be described are intimately related to geometric realizations of Coxeter groups. Following Brink and Howlett, I shall work here exclusively with the **standard realization**. To each element s of S is associated an element α_s of a basis Δ of a real vector space \hat{V} . An inner product is defined on \hat{V} according to the formula

$$\alpha_s \bullet \alpha_t = -\cos(\pi/m_{s,t}) = -(\zeta_{2m} + \zeta_{2m}^{-1})/2 \quad (m = m_{s,t}, \zeta_{2m} = e^{2\pi i/2m}).$$

In particular, $\alpha_s \bullet \alpha_s = 1$ for all s . The angle π/∞ is to be interpreted as 0.

In the dual V of \widehat{V} vectors α^\vee are then defined by the conditions

$$\langle \alpha, \beta^\vee \rangle = 2(\alpha \bullet \beta).$$

The α^\vee (for α in Δ^\vee) may not be linearly independent. This is so when the matrix $\alpha \bullet \beta$ is singular, as for example when it is

$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

and $\alpha^\vee = -\beta^\vee$. For each s in S reflections in V are defined by the formulas

$$sv = v - \langle \alpha_s, v \rangle \alpha_s^\vee.$$

Dual reflections in \widehat{V} are defined by the formulas

$$\langle s\lambda, v \rangle = \langle \lambda, sv \rangle$$

so that

$$s\lambda = \lambda - \langle \lambda, \alpha_s^\vee \rangle \alpha_s = \lambda - 2(\lambda \bullet \alpha_s) \alpha_s.$$

These reflections preserve the inner product. Thus, assuming $s \neq t$, we have $\alpha_s \bullet \alpha_t = -1$ whenever s and t generate an infinite group, $\alpha_s \bullet \alpha_t = 0$ when s and t commute, and $-1 < \alpha_s \bullet \alpha_t \leq -1/2$ otherwise. The maps taking the elements s to the corresponding reflections induce homomorphisms of W into both $\text{GL}(V)$ and $\text{GL}(\widehat{V})$.

The **Coxeter graph** associated to a Coxeter system has as nodes the elements of S , and an edge linking s and t if $m_{s,t}$, which is called the **degree** of the edge, is more than 2. In other words, the edges link s and t with $\alpha_s \bullet \alpha_t \neq 0$. In this case—i.e. when $\alpha \bullet \beta \neq 0$ —I write $\alpha \sim \beta$. A **simple link** in the graph is one with degree 3, a **multiple link** one of higher degree. The formula for reflection amounts to the specification that if $\lambda = \sum \lambda_\gamma \gamma$ and $s = s_\alpha$

$$(s\lambda)_\beta = \begin{cases} -\lambda_\alpha + \sum_{\alpha \sim \gamma} [-2\alpha \bullet \gamma] \lambda_\gamma & \text{if } \beta = \alpha \\ \lambda_\beta & \text{otherwise} \end{cases}$$

In calculations, it is frequently best to work with $\langle \alpha, \beta^\vee \rangle = 2\alpha \bullet \beta$ rather than the dot-product itself, since it is at once simpler and more efficient to do so.

The representation on V is faithful. Let

$$C = \{v \in V \mid \langle \alpha_s, v \rangle > 0 \text{ for all } s \in S\}.$$

The faces of this simplicial cone are the closures of the relatively open simplices

$$C_T = \{v \in V \mid \langle \alpha_s, v \rangle = 0 \text{ for } s \in T, \langle \alpha_s, v \rangle > 0 \text{ for } s \notin T\}$$

where $T \subseteq S$. Thus C itself is C_\emptyset , and a face of codimension one is some $C_{\{s\}}$.

The union of the domains $w\overline{C}$ (w in W) is a convex conical region called the **Tits cone** \mathcal{C} of the realization. The group W acts discretely on \mathcal{C} , and the closure $\overline{\mathcal{C}}$ of the open simplicial cone is a fundamental domain for this action. Each open face of $w\overline{C}$ is the transform of a unique C_T , and is said to be **labeled** by T .

A transform $w\alpha_s$ of an element of Δ is called a **root** of this realization. (This terminology differs somewhat from that in the theory of Kac-Moody algebras, where it would be called

a **real** root.) Points in the root hyperplane $\lambda = 0$ are fixed by the **root reflection** s_λ where $s_{w\alpha} = ws_\alpha w^{-1}$, the conjugate of an element in S . For any root λ , define

$$\begin{aligned}\mathcal{C}_{\lambda=0} &= \{v \in \mathcal{C} \mid \langle \lambda, v \rangle = 0\} \\ \mathcal{C}_{\lambda \geq 0} &= \{v \in \mathcal{C} \mid \langle \lambda, v \rangle \geq 0\} \\ \mathcal{C}_{\lambda \leq 0} &= \{v \in \mathcal{C} \mid \langle \lambda, v \rangle \leq 0\}.\end{aligned}$$

The collection of root hyperplanes is locally finite in \mathcal{C} .

A root λ is called **positive** if $\lambda > 0$ on \mathcal{C} , **negative** if $\lambda < 0$ on \mathcal{C} . Every root is either positive or negative, which means that no root hyperplane ever intersects \mathcal{C} . A root λ is positive if and only if $\lambda = \sum_{\alpha \in \Delta} \lambda_\alpha \alpha$ with all $\lambda_\alpha \geq 0$.

1.1. Proposition. *Suppose λ and μ to be distinct positive roots. The following are equivalent:*

- (a) $\lambda \bullet \mu = \cos(\pi k / m_{s,t})$ with $0 < k < m_{s,t}$ for some s, t in S with $m_{s,t} < \infty$;
- (b) $|\lambda \bullet \mu| < 1$;
- (c) the group generated by the root reflections s_λ and s_μ is finite;
- (d) the hyperplanes $\lambda = 0$ and $\mu = 0$ intersect in the interior of \mathcal{C} ;
- (e) there exists a point of \mathcal{C} in the region where $\lambda < 0$ and $\mu < 0$;
- (f) there exists a point of V in the region where $\lambda > 0$ and $\mu > 0$ whose orbit under the group generated by s_λ and s_μ contains a point in the region where $\lambda < 0$ and $\mu < 0$.

PROOF. This is implicit in the very early article (Vinberg, 1971). That (a) implies (b) is trivial; that (b) implies (c) is a consequence of the simplest part of his discussion of pairs of reflections. If G is any finite subgroup of W then the G -orbit of any point in the interior of \mathcal{C} will be finite, and since \mathcal{C} is convex the centre of mass of the orbit will be in the interior of \mathcal{C} as well. This shows that (c) implies (d). Suppose $L = \{\lambda = 0\} \cap \{\mu = 0\}$ to contain a point in the interior of \mathcal{C} . The whole of \mathcal{C} is tiled by simplices wC_T . Points in the open cones wC are fixed only by 1 in W , and those in some $wC_{\{s\}}$ are fixed only by 1 and a single reflection. Therefore the linear space L must contain some wC_T with $|T| = 2$. This means that s and t can be simultaneously conjugated into W_T , and shows that (d) implies (a).

That (d) implies (e) is trivial.

To show that (e) implies (f), let P be a point in \mathcal{C} where $\lambda < 0$ and $\mu < 0$. We may assume that P is in the interior of \mathcal{C} . If Q is a point of \mathcal{C} , the segment from P to Q intersects only a finite number of root hyperplanes, which means that there are only a finite number of root hyperplanes separating P from Q . But if κ is any positive root and $\langle \kappa, P \rangle < 0$ then P lies on the side of $\kappa = 0$ away from \mathcal{C} , while $s_\kappa P$ lies on the same side. So a finite number of applications of s_λ and s_μ to P will bring it to the region where $\lambda > 0$, $\mu > 0$.

It remains to show that (f) implies (c). In effect, this hypothesis means that the group has a longest element. This is a standard argument, which I'll not repeat here. **QED**

Following Brink and Howlett, I say that the positive root λ **dominates** the positive root μ if the region $\mathcal{C}_{\lambda \geq 0}$ contains the region $\mathcal{C}_{\mu \geq 0}$. Another way of putting this is to require that $\lambda > 0$ in every chamber wC on which $\mu > 0$.

1.2. Proposition. *Suppose λ and μ to be two distinct positive roots. If λ dominates μ then s_λ and s_μ generate an infinite group. Conversely, if s_λ and s_μ generate an infinite group then one of three possibilities holds:*

- (a) λ dominates μ ;
- (b) μ dominates λ ;
- (c) $s_\mu\lambda$ dominates μ and $s_\lambda\mu$ dominates λ .

PROOF. For any root λ , let $\{\lambda = 0\}^{\text{int}}$ be the intersection of the hyperplane $\lambda = 0$ with the interior of \mathcal{C} . According to Proposition 1.1, that the pair s_λ and s_μ generate an infinite group means neither more nor less than that

$$\{\lambda = 0\}^{\text{int}} \cap \{\mu = 0\}^{\text{int}} = \emptyset .$$

Suppose that λ dominates μ , so that $\mathcal{C}_{\mu \geq 0}$ is contained in $\mathcal{C}_{\lambda \geq 0}$. Since λ and μ are distinct, the hyperplanes $\lambda = 0$ and $\mu = 0$ are distinct. The intersection of the two hyperplanes has codimension one in either of them. If P lies in the intersection $\{\lambda = 0\}^{\text{int}} \cap \{\mu = 0\}^{\text{int}}$ then there will exist points in the interior of \mathcal{C} where $\mu = 0$ and λ is negative. This contradicts the dominance of μ by λ . From this we conclude that if λ dominates μ then s_λ and s_μ generate an infinite group.

Suppose that s_λ and s_μ generate an infinite group. Since

$$\{\lambda = 0\}^{\text{int}} \cap \{\mu = 0\}^{\text{int}} = \emptyset ,$$


the sign of μ on $\{\lambda = 0\}^{\text{int}}$ is a non-zero constant, as is that of λ on $\{\mu = 0\}^{\text{int}}$.

Suppose that $\lambda < 0$ on $\{\mu = 0\}^{\text{int}}$. I claim that in this case μ dominates λ . If not, then there exists a point P in \mathcal{C} where $\lambda \geq 0$ but $\mu < 0$. If Q lies in the interior of \mathcal{C} the open segment QP must lie entirely in the convex set where $\lambda \geq 0$ as well as the interior of \mathcal{C} . It crosses from the region where $\mu > 0$ to that where $\mu < 0$, and must therefore contain a point where $\mu = 0$ but $\lambda \geq 0$, a contradiction.

Similarly, if $\lambda > 0$ on $\{\mu = 0\}^{\text{int}}$ and $\mu < 0$ on $\{\lambda = 0\}^{\text{int}}$ then λ dominates μ .

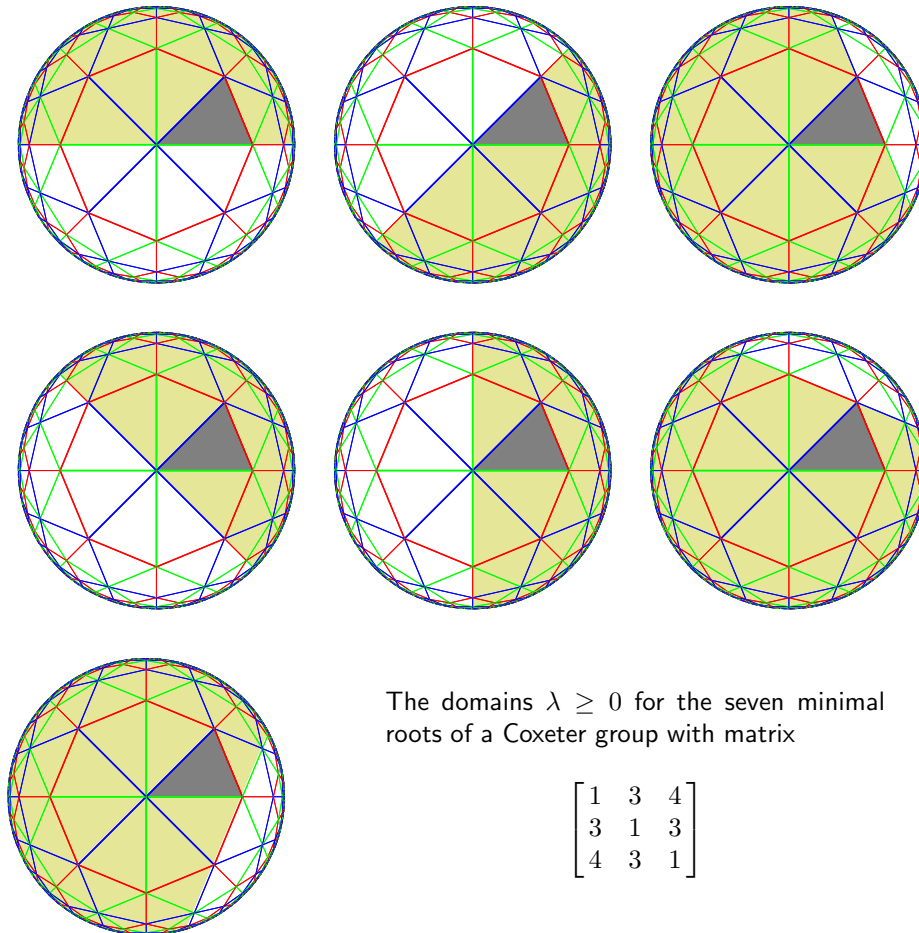
Finally, suppose that $\lambda > 0$ on $\{\mu = 0\}^{\text{int}}$, and that $\mu > 0$ on $\{\lambda = 0\}^{\text{int}}$. I claim that $\mu > 0$ on $s_\lambda\mathcal{C}$. If not, then $\mu < 0$ on $s_\lambda\mathcal{C}$. If P lies in $s_\lambda\mathcal{C}$ and Q lies on $\{\lambda = 0\}^{\text{int}}$ then on the one hand $\lambda < 0$ on $[P, Q)$ while on the other $\langle \mu, P \rangle < 0$ while $\langle \mu, Q \rangle > 0$ so $\mu = 0$ for some point of (P, Q) , contradicting the assumption that $\lambda > 0$ on $\{\mu = 0\}^{\text{int}}$. But since $\mu > 0$ on $s_\lambda\mathcal{C}$, $s_\lambda\mu > 0$ on \mathcal{C} and $s_\lambda\mu$ is a positive root. But $\lambda < 0$ on $s_\lambda\{\mu = 0\}$ so if we replace μ by $s_\lambda\mu$ we are back in the first case. Similarly for $s_\mu\lambda$. **QED**

The three cases are shown in these figures:



2. Minimal roots

A **minimal root** is a positive root that dominates only itself. It is an important and remarkable fact about arbitrary Coxeter groups, first proved by Brink and Howlett, that the set of minimal roots is always finite. The significance of this result does not appear in familiar cases, since all roots are minimal when W is finite, and the minimal roots for affine Coxeter groups are rather simple in nature.



In the rest of this section I recall from (Brink and Howlett, 1993) some basic facts about minimal roots. Proofs and statements will often be the same as theirs, and I'll refer to them when this happens. But in some places what I take from them is only implicit in their paper, and in other places I have given both statements and proofs a geometrical flavour, somewhat as in (Casselman, 1995).

2.1. Lemma. (B & H 2.2(ii)) *Suppose λ and μ to be positive roots with λ dominating μ . If $w\mu$ is positive, then $w\lambda$ is also positive, and dominates $w\mu$.*

PROOF. Under the hypothesis, $\mathcal{C}_{\mu \geq 0}$ is contained in $\mathcal{C}_{\lambda \geq 0}$, hence

$$w\mathcal{C}_{\mu \geq 0} = \mathcal{C}_{w\mu \geq 0} \subseteq w\mathcal{C}_{\lambda \geq 0} = \mathcal{C}_{w\lambda \geq 0}.$$

If $w\mu$ is positive, the first contains C , and therefore $w\lambda > 0$ on C as well and $w\lambda$ is positive.. **QED**

It is well known that s in S permutes the complement of α_s in the set of all positive roots, or in other words that if s is in S and λ a positive root, then either $\lambda = \alpha_s$ and $s\lambda < 0$, or $s\lambda > 0$. For the minimal roots, there is a similar range of options:

2.2. Proposition. *If λ is a minimal root and s in S then exactly one of these three options holds:*

- $\lambda = \alpha_s$ and $s\lambda < 0$;
- $s\lambda$ is again minimal;
- $s\lambda$ dominates α_s .

PROOF. If $\lambda \neq \alpha_s > 0$, then $s\lambda > 0$. If $s_\alpha\lambda$ is not minimal, say it dominates $\beta > 0$ with $s\lambda \neq \beta$. If $\beta \neq \alpha$, then $s_\alpha\beta > 0$, and by the previous Lemma $\lambda = s_\alpha s_\alpha\lambda$ dominates $s_\alpha\beta$. Since λ is minimal, this can only happen if $\lambda = s_\alpha\beta$, a contradiction. **QED**

The **minimal root reflection table** is a matrix ρ of size $N \times |S|$, where N is the number of minimal roots. The entries are either minimal roots, or **virtual** minimal roots I arbitrarily label as \ominus and \oplus . These virtual roots are taken to themselves by all reflections. I set $\rho(s, \lambda) = \mu$ if $s\lambda = \mu$ is minimal, $\rho(s, \lambda) = \ominus$ if $\lambda = \alpha_s$, and $\rho(s, \lambda) = \oplus$ if $s\lambda$ is a non-minimal positive root. This table is one of the fundamental data structures of a Coxeter group. Here is what it looks like for the group pictured above, with the minimal roots numbered in the order in which the images are displayed:

		$s_\alpha\lambda$		
	α	0	1	2
λ	0	\ominus	4	6
1	1	3	\ominus	5
2	2	6	5	\ominus
3	3	1	3	\oplus
4	4	4	0	\oplus
5	5	\oplus	2	1
6	6	2	\oplus	0

This table can be calculated from the pictures by inspection.

3. The root graph

In this section I'll lay out the basic facts needed to construct the minimal roots and build this table for an arbitrary Coxeter group. This process will take place by induction on what might be called, roughly, the complexity of a root relative to those in Δ .

The **depth** $\delta(\lambda)$ of a positive root λ is the length of the smallest w such that $w^{-1}\lambda < 0$. The roots with depth 1 are therefore precisely those in Δ , and in all cases


$$\delta(\lambda) - 1 \leq \delta(s_\alpha\lambda) \leq \delta(\lambda) + 1 \quad (\alpha \in \Delta).$$

Equivalently, the depth n is one less than the length of a shortest gallery $C = C_0, C_1, \dots, C_n$ with C_n in the region $\lambda < 0$. Recall that a gallery is a chain of chambers with successive chambers sharing a common wall. The depth is also the smallest n with

$$\lambda = s_1 s_2 \dots s_{n-1} \alpha$$

for some $\alpha \in \Delta$. A partial order is thus induced on the set of all positive roots: $\lambda \preceq \mu$ if $\mu = w\lambda$ where $\delta(\mu) = \ell(w) + \delta(\lambda)$. This order gives rise to the **root graph**, whose nodes are positive roots, with edges $\lambda \rightarrow s\lambda$ whenever $\lambda \prec s\lambda$.

3.1. Lemma. *If λ and μ are distinct positive roots and $s_\mu\lambda$ dominates μ then $s_\mu\lambda > 0$ and $\delta(\lambda) < \delta(s_\mu\lambda)$.*



PROOF. If (C_i) is a gallery of least length from C to the region where $s_\mu\lambda < 0$, there will be some pair C_i, C_{i+1} sharing a wall on the hyperplane $\mu = 0$. The gallery reflected at $\mu = 0$ will be shorter. Hence $\delta(\lambda) < \delta(s_\mu\lambda)$. **QED**

The minimal roots lie at the bottom of the root graph, in the sense that:

3.2. Proposition. (B & H 2.2(iii)) *If λ and μ are both positive roots with $\lambda \preceq \mu$ and μ is minimal, so is λ .*

PROOF. This reduces to the case $\mu = s_\alpha\lambda$. If μ is minimal but $\mu \neq \alpha$, then either $s_\alpha\mu = \lambda$ is minimal, or it dominates α . But by the preceding Lemma, in the second case $\delta(\mu) < \delta(s_\alpha\mu)$, a contradiction. **QED**

3.3. Lemma. (B & H 2.3) *If $\lambda > 0$ and $\lambda \bullet \alpha > 0$ then $s_\alpha\lambda \prec \lambda$.*

PROOF. Suppose $w^{-1}\lambda < 0$ with $\ell(w) = \delta(\lambda)$.

If $w^{-1}\alpha < 0$ then we can write

$$\begin{aligned} w &= s_\alpha s_2 \dots s_n \\ \lambda &= s_\alpha s_2 \dots s_{n-1} \alpha_n \quad (\alpha_n = \alpha_{s_s}) \\ s_\alpha\lambda &= s_2 \dots s_{n-1} \alpha_n \\ \delta(s_\alpha\lambda) &= n - 1 < \delta(\lambda) = n. \end{aligned}$$

Now assume that $w^{-1}\alpha > 0$. Let w have the reduced expression $w = s_1 \dots s_n$, and let $u = ws_n = s_1 \dots s_{n-1}$. Thus $w\alpha_n < 0$. I claim that $u^{-1}s_\alpha\lambda < 0$, which implies that $\delta(s_\alpha\lambda) = n - 1$.

We can calculate that

$$\begin{aligned} w^{-1}s_\alpha\lambda &= w^{-1}\lambda - 2(\lambda \bullet \alpha)w^{-1}\alpha \\ u^{-1}s_\alpha\lambda &= s_n w^{-1}s_\alpha\lambda \\ &= s_n(w^{-1}\lambda - 2(\lambda \bullet \alpha)w^{-1}\alpha) \\ &= u^{-1}\lambda - 2(\lambda \bullet \alpha)s_n w^{-1}\alpha. \end{aligned}$$

Since the depth of λ is n , $\mu = u^{-1}\lambda$ will be positive. But then $\mu > 0$, $s_n\mu < 0$ implies that $\mu = \alpha_n$. So we have

$$u^{-1}s_\alpha\lambda = \alpha_n - 2(\lambda \bullet \alpha)s_n w^{-1}\alpha.$$

The root

$$\alpha_n - 2(\lambda \bullet \alpha) s_n w^{-1} \alpha$$

is either positive or negative. In the second term, the positive root $w^{-1} \alpha$ cannot be α_n , since then $\alpha = w \alpha_n$, contradicting $w \alpha_n < 0$. Therefore $s_n w^{-1} \alpha$ is a positive root also not α_n , and the second term is a negative root not equal to a multiple of $-\alpha_n$. The whole sum therefore has to be negative. **QED**

3.4. Corollary. (B & H 1.7) *Whenever λ is a positive root and s in S with $\lambda \neq \alpha_s$ there are three possibilities: (a) $s\lambda \prec \lambda$, (b) $s\lambda = \lambda$, or (c) $\lambda \prec s\lambda$, depending on whether $\lambda \bullet \alpha_s > 0$, $\lambda \bullet \alpha_s = 0$, or $\lambda \bullet \alpha_s < 0$, respectively.*

3.5. Corollary. (B & H 1.8) *If $\lambda = \sum \lambda_\alpha \alpha \preceq \mu = \sum \mu_\alpha \alpha$ then $\lambda_\alpha \leq \mu_\alpha$.*

In constructing the list of minimal roots and the minimal root reflection table, the roots will be dealt with by going up the root graph. We shall find ourselves in the situation where λ is a minimal root and we know all $s\lambda \prec \lambda$. We are to calculate an unknown $s\lambda$. We have to decide whether $s_\alpha \lambda$ is again minimal or not. It is easy to decide whether $\lambda = \alpha$, and in the remaining cases the possibilities, according to Proposition 2.2, are that it is minimal or that $s_\alpha \lambda$ dominates α . How do we tell which?

3.6. Proposition. *Suppose α to lie in Δ , λ to be a minimal root distinct from α . If $\lambda \prec s_\alpha \lambda$ then $s_\alpha \lambda$ is again minimal if and only if s_α and s_λ generate a finite group.*

Or, equivalently, $s_\alpha \lambda$ is not minimal if and only if s_α and s_λ generate an infinite group.

PROOF. Let $\mu = s_\alpha \lambda$.


Suppose that μ is not minimal. Then it dominates α , and s_α and s_μ generate an infinite group. But this is also the group generated by s_λ and s_α .

Conversely, suppose s_μ and s_α generate an infinite group. Among the three possibilities of Proposition 1.2 only the one in which μ dominates α is possible, since $\lambda \prec \mu$. **QED**

The minimal roots are minimal in the sense that they fill the bottom of the root graph. They are saturated in the sense that any root below a minimal root is still minimal. An edge $\lambda \prec s\lambda$ takes us out of the set of minimal roots if $s\lambda$ dominates α_s . Here is the bottom of the root graph for the Coxeter group with matrix

$$\begin{bmatrix} 1 & 3 & 4 \\ 3 & 1 & 3 \\ 4 & 3 & 1 \end{bmatrix}.$$

The single negative virtual root lies at the very bottom, the minimal roots are those at depth 1 and 2. Self-links are shown by coloured disks.



These results will soon be used to outline an algorithm to construct the minimal root reflection table. Processing roots will take place in order of increasing depth, applying this:

Minimality Criterion: *If $\lambda \neq \alpha$ is a minimal root and $\lambda \prec s_\alpha \lambda$, then $s_\alpha \lambda$ is no longer a minimal root if and only if $\lambda \bullet \alpha_s \leq -1$.*

Details follow, after I explain in the next section how to use the minimal root reflection table in multiplication.

4. Minimal roots and multiplication

Suppose w to have the normal form $w = s_1 \dots s_n$. Then according to du Cloux's strong form of the exchange principle (explained first in (du Cloux, 1990) and more generally in (Casselmann, 2001)) the normal form of sw will be obtained from that of w by either **insertion** or **deletion**:


$$\text{NF}(sw) = s_1 \dots s_i t s_{i+1} \dots s_n$$

or

$$\text{NF}(sw) = s_1 \dots s_{i-1} s_{i+1} \dots s_n .$$

How do we find which of the two is concerned, and where the insertion or deletion occurs? If insertion, which t is to be inserted?


The algorithm explained in (Casselmann, 2001) handles these problems in geometric terms. I define the **InverseShortLex tree** to be the directed graph whose nodes are elements of W , with a link from x to y labeled by s if $\text{NF}(y) = \text{NF}(x) \bullet s$. Its root is the identity element. Each $y \neq 1$ in W is the target of exactly one link, labeled by the s in S least with $ys < y$. Paths in this tree starting at 1 match the normal forms of elements of W .



The InverseShortLex tree for \tilde{A}_2

This tree gives rise to a geometrical figure in which we put a link from xC to yC if there is an edge in the tree from x to y . The multiplication problem can now be formulated: *given an InverseShortLex path from C to wC , how can we find the InverseShortLex path from C to swC ?* As explained in (Casselmann, 2001), the InverseShortLex figure is very close to

being symmetric with respect to the reflection s ; only links near the hyperplane $\alpha_s = 0$ are not preserved upon reflection. The chambers principally affected are what I call **exchange sites** of the `InverseShortLex` tree— wC is an exchange site (for the simple reflection s) if the chamber wC has a face on the hyperplane $\alpha_s = 0$, labeled by t in S which is less than the labelling of the `InverseShortLex` link entering wC .



Some of the exchange nodes for reflection by $\langle 1 \rangle$ in \tilde{A}_2 . Links modified under reflection are colored. The fundamental chamber is at the bottom.

The relationship between this notion and multiplication is simple. Suppose $\text{NF}(w) = s_1 \dots s_n$ and for each i let w_i be the partial product $s_1 \dots s_i$. Suppose that $w_k C$ is the last exchange site in the gallery $C = C_0, \dots, C_n = w_n C$, and suppose that t labels the wall of $w_k C$ on $\alpha_s = 0$. Then either the gallery crosses $\alpha = 0$ there or it does not. In the first case, $t = s_{k+1}$ and

$$\text{NF}(sw) = s_1 \dots s_k s_{k+2} \dots s_n .$$

In the second,

$$\text{NF}(sw) = s_1 \dots s_k t s_{k+1} \dots s_n .$$

The question is now: *How do we recognize InverseShortLex exchange sites?* This is where roots and, even more pleasantly, minimal roots come in.

4.1. Lemma. *Suppose $\text{NF}(w) = s_1 \dots s_n$. The chamber wC is an InverseShortLex exchange site for α if and only if $w^{-1}\alpha = \beta$ with β in Δ and $s_\beta < s_n$.*

PROOF. The wall of wC labeled by β is embedded in the hyperplane $w\beta = 0$. **QED**

This suggests the following rough idea for an algorithm to calculate $\text{NF}(sw)$: Read the string $s_1 \dots s_n$ from left to right. As we read, keep track of an element t in S , an index k , the element $x = s_1 \dots s_{i-1}$ so far read, and a root λ . The index k is that of the last exchange, the element t is the label of the last exchange, $\lambda = x^{-1}\alpha_s$. The element x is implicit; we don't really have to keep track of it explicitly. At the start, $k = 0$, $t = s$, $x = 1$, $\lambda = \alpha_s$. As we read s_i , x changes to xs_i , λ changes to $s_i\lambda$. If $\lambda = \beta$ with $s_\beta < s_i$, t changes to s_β and

k to i . When we are about to read s_i , $sx = s_1 \dots s_k t s_{k+1} \dots s_{i-1}$. When we are through scanning, we have the data determining the last exchange.

So far, minimal roots have not come into play. But in fact, in calculating the reflected root $s\lambda$ we only need to apply the minimal root reflection table. If λ is minimal and $s\lambda = \ominus$, then the chambers $w_{i-1}C_{i-1}$ and $C_i = w_{i-1}s_i C_{i-1}$ lie on opposite sides of the hyperplane $\alpha = 0$, so we know we are looking at a deletion, and we need know no more. If $s\lambda$ is minimal and $s\lambda = \oplus$ then we know the gallery has crossed a hyperplane $\lambda = 0$ which separates wC from $\alpha = 0$. The exchange we last recorded will be the last we ever record, because the gallery cannot recross that hyperplane to reach $\alpha_s = 0$.

Pseudocode. The procedure has arguments s and the string $s_1 \dots s_n$ representing w . It uses variables

- t : an element of S ; initially s
- i : in the range $[0, n]$; initially 0
- k : also in the range $[0, n]$; initially 0
- λ : an extended minimal root—either a minimal root or \oplus or \ominus ; initially α_s
- x : an element of W ; initially 1

BEGINNING OF LOOP. At this point $x = s_1 \dots s_{i-1}$, $sx = s_1 \dots s_k t s_{k+1} s_{i-1}$, $\lambda = x^{-1}\alpha_s$.

Read the next generator s_i . Then x becomes $x = xs_i$. Set $\lambda = s_i\lambda$ from the minimal root reflection table.

If $\lambda = \beta$ in Δ with $s_\beta < s_i$, then set $t = s_\beta$, $k = i$.

If $\lambda = \ominus$, this means $k = i - 1$ and $s_i = t$.

Set $\text{NF}(sw) = s_1 \dots s_{i-1} s_{i+1} \dots s_n$ and stop.

If $\lambda = \oplus$, set $\text{NF}(sw) = s_1 \dots s_k t s_{k+1} s_{k+2} \dots s_n$ and stop.

Advance i to $i + 1$. If $i > n$, stop with $\text{NF}(sw) = s_1 \dots s_k t s_{k+1} s_{k+2} \dots s_n$. Else loop.

We now know how to use the minimal root reflection table to multiply, and it remains to see how to construct it.

5. A rough guide to constructing the minimal roots

A straightforward procedure to make a list of all minimal roots, along with a description of how reflections by the elements of S affect them, is not difficult to sketch.

The data associated to a minimal root λ , at least in this first approach, are to be:

- an array of coordinates;
- its index (in the list of minimal roots);
- a list of reflections $s\lambda$.

There are two natural choices of coordinates. Those I use in this note are the ones in the expression of a root as a linear combination of the roots in Δ :

$$\lambda = \sum_{\alpha \in \Delta} \lambda_\alpha \alpha .$$

The other reasonable choice is the array $(\lambda \bullet \alpha)$, which du Cloux has discovered to be a choice somewhat more efficient in programming, if also more difficult to work with. Indices are assigned in the order in which the minimal roots are first calculated; the roots in Δ are assigned the first $|S|$ integers. Furthermore, in the preliminary version of the algorithm we shall maintain a **look-up table** of some kind, enabling us to find a root's data given its coordinates, if it has already been encountered.

A minimal root is to be dealt with in two stages: (1) **installing** it and (2) **finishing** it. In the **installation** of a root

- ◇ its coordinate array is calculated;
- ◇ it is assigned an index;
- ◇ it is entered in the look-up table;
- ◇ a few reflections known at this time are registered;
- ◇ it is put on the waiting list of items to be finished.

The algorithm begins by installing the elementary roots α , which includes recording the reflections $s_\alpha\alpha = \ominus$.

Items are taken off the waiting list in the order in which they are put on—the list is, in programmers' terminology, a FIFO list (**first in, first out**) or **queue**. Roots are processed in order of depth, so that if $\delta(\lambda) < \delta(\mu)$ the root λ is installed before μ .

When a root λ is to be *finished*, it is taken off the queue and all of the reflections $s\lambda$ not already known are found. At this moment, all roots with depth less than that of λ have been finished, so that in particular we have recorded all reflections $\mu = s\lambda \prec \lambda$. We must calculate the $s\lambda$ not yet determined—we must decide whether $s\lambda = \lambda$, $s\lambda$ is a minimal root above λ in the root graph, or $s\lambda = \oplus$. We can tell easily whether $s\lambda = \lambda$. Otherwise we know that $\lambda \prec \mu = s\lambda$. There are three possibilities: (a) The root μ is minimal, and we have already installed it. It will be registered in our look-up table. We just add to the record that $s\lambda = \mu$ and $s\mu = \lambda$. (b) The root μ is not minimal. The criterion is that s_α and s_μ generate an infinite group, or equivalently that $\lambda \bullet \alpha_s \leq -1$. We can tell this when we calculate μ . In this case we set $s\lambda = \oplus$. (c) Otherwise $-1 < \lambda \bullet \alpha_s < 0$ and the root μ is a new minimal root not yet installed. We do so by assigning it an index, putting it in the look-up table, recording the reflection $s\lambda = \mu$, and putting μ in the process queue. After all $s\lambda$ have been found, we go on to the next root in the queue, as long as there is something there.

There are a couple of important simplifications that can be made to this procedure.

Eliminating the look-up table. Maintaining the look-up table is somewhat cumbersome. It was du Cloux who first observed that one can dispense with it. The only place where it is used is in case (a) above, where we are recording that $s_\alpha\lambda = \mu$ with μ already installed, and there is a better way to deal with this.

A root in the linear span of a pair of roots in Δ is called **dihedral**. Suppose for the moment that s and t are distinct elements of S . A root λ is a linear combination of α_s and α_t if and only if the orbit of λ with respect to the subgroup $W_{s,t}$ generated by s and t contains one of α_s or α_t . Installing the dihedral roots is a simple task, and can (and should) be done independently of all other installations, as we shall see later.

To any Coxeter system (W, S) is associated a partial ordering, sometimes called the **weak right Bruhat order**—the nodes are elements of W and there is an oriented edge $w \rightarrow ws$ labeled by s if $\ell(ws) = \ell(w) + 1$. The identity is the unique minimal node, and if the group is finite the unique highest node is the element of greatest length in W . Finite paths in this graph starting at 1 correspond precisely to reduced expressions of elements of W assembled from the labels of the edges followed. For a finite dihedral group $W_{s,t}$ this graph is particularly simple—there are two paths $st\dots$ and $ts\dots$ of length m from 1 to the longest element $w_{s,t}$ corresponding to the two reduced expressions for $w_{s,t}$ assured by the Coxeter relation $(st)^m = 1$. The Bruhat order is intimately related to the geometry of roots and chambers.

The first of the following two lemmas is well known:

5.1. Lemma. *For $T \subseteq S$, every element of the subgroup W_T permutes all the positive roots that are not in the linear span of the α_s with s in T .*

The next one has for me a mildly paradoxical flavour.

5.2. Lemma. *Suppose s, t to be a pair of elements of S , μ to be a positive root not in the span of α_s and α_t . Assume that $\alpha_s \bullet \mu > 0$ and $\alpha_t \bullet \mu > 0$. Then $W_{s,t}$ is finite, and the map $w \mapsto w\mu$ is a bijection between $W_{s,t}$ and $W_{s,t}\mu$ inducing an isomorphism of the weak Bruhat order with inverted precedence in the root graph.*

That is to say that $\ell(w) < \ell(sw)$ for s in T if and only if $sw\mu \prec w\mu$. Equivalently, if ν is the root of least depth in the $W_{s,t}$ -orbit then $w \mapsto w\nu$ is an order-preserving bijection between $W_{s,t}$ and $W_{s,t}\mu$.

PROOF. Let $T = \{s, t\}$. By assumption, all roots in the W_T -orbit of μ are positive roots. The region $\alpha_s \bullet v > 0, \alpha_t \bullet v > 0$ is a fundamental domain for $W_{s,t}$. If v is any vector in this region, w in W , then $\ell(xw) = \ell(w) + 1$ for x in T if and only if $wv \bullet \alpha_x > 0$. If ν is a root in the W_T -orbit of least depth then $\alpha_x \bullet \nu < 0$ for x in T . Proposition 1.1 implies that W_T is finite, and the rest follows easily. **QED**

How can this be applied in our algorithm to find all minimal roots? Suppose we are in the process of installing a minimal root μ . We may assume it not to be dihedral, because we can deal separately with all the dihedral roots in a particularly simple fashion. So we are looking at $\mu = s\lambda$ with $\lambda \prec \mu$. Two questions arise: (a) For t in S , how do we tell whether $t\mu \prec \mu$ or not? (b) If so, then $t\mu$ has already been finished. How do we identify it?

If $W_{s,t}$ is infinite, then by Proposition 5.2 the problem doesn't occur. Otherwise, let $w_{s,t}$ be the longest element in $W_{s,t}$. Calculate the element ν of least depth in $W_{s,t}\mu$, by finding alternately $t\lambda, st\lambda$, etc. until the chain starts to ascend in depth. The Lemma tells us that $t\mu \prec \mu$ if and only if the difference in depth between μ and ν is exactly $m_{s,t}$. In these circumstances the element $t\mu$ will then be $(tw_{s,t})\nu$, which can be calculated from the reflection tables of elements of depth less than λ . In other words, as soon as μ is installed we can calculate all the **descents** $t\mu \prec \mu$ —without doing any new calculations!

In the algorithm above we can therefore eliminate the look-up table. The procedure sketched above requires, however, that we know whether the reflection of a root has less depth or not, so we add to the data of a root λ its **descent set**, the set of s with $s\lambda \prec \lambda$. When we install a root, before we put it in the queue, we calculate all of its descents. (An alternative would be to store the depth of a root as part of its data.)

As has already been mentioned, because the result above is applicable only to roots not in the orbit of α_s and α_t under $W_{s,t}$, all the dihedral roots have to be handled somewhat specially—they should all be installed immediately after the roots in Δ , and those of depth greater than 2 have to be fed into the queue at the right moment.

One huge benefit of this technique is that it requires arithmetic only to decide whether $s_\alpha\lambda = \lambda$ or not.

Noting inherited locks. The following was observed by Brink and Howlett, and used as an important part of their proof that the set of minimal roots is finite:

5.3. Proposition. *If $s\lambda = \oplus$ and $\lambda \prec \mu$ then $s\mu = \oplus$ as well.*

PROOF. Let $\alpha = \alpha_s$. If λ is a minimal root, then according to the Minimality Criterion $s\lambda = \oplus$ if and only if $\lambda \bullet \alpha_s \leq -1$.

For β in Δ , λ any root

$$s_\beta\lambda \bullet \alpha = (\lambda - 2(\lambda \bullet \beta)\beta) \bullet \alpha = (\lambda \bullet \alpha) - 2(\lambda \bullet \beta)(\beta \bullet \alpha)$$

which means that for $\beta \neq \alpha$ the dot product $\lambda \bullet \alpha$ does not increase under reflection by β , as long as $\lambda \prec s_\beta \lambda$ and therefore $\lambda \bullet \beta < 0$. Therefore once we have calculated that $s\lambda = \oplus$ this remains true for all minimal μ with $\lambda \prec \mu$. **QED**

This result is the crucial part of the proof by Brink and Howlett that the number of minimal roots is finite. By König's Lemma, if there an infinite number, then in the part of the root graph composed of minimal roots there would exist an infinite chain. For each λ in this chain, let \mathfrak{N}_λ be the subset of α in Δ with $s_\alpha \lambda = \oplus$. According to the Proposition, the set P_α is non-decreasing, so in an infinite chain of minimal roots it must be stationary with non-trivial complement. Let P be the intersection, and choose λ be such that $P_\lambda = P$ for $\lambda \preceq \mu$ in the chain. Beyond λ , reflection by α in P never occurs, so μ_α for all $\lambda \preceq \mu$. Furthermore, the set of possible values of $\mu \bullet \alpha$ for α in the complement of P lies in $(-1, 1)$, and must be finite. Therefore there must exist μ_1 and μ_2 such that (a) $\mu_{1,\alpha} = \mu_{2,\alpha}$ for α in P and (b) $\mu_1 \bullet \alpha = \mu_2 \bullet \alpha$ for α in its complement. Lemma 2.7 of B & H then implies that $\mu_1 \bullet \mu_2 = 1$, which means that one of them dominates the other, which is a contradiction.

I call an s with $s\lambda = \oplus$ a **lock** on λ at s . This Proposition tells us that locks are inherited as we go up the root graph and these inherited locks, like descents, can be immediately recorded upon installation.

In the simplest cases—when none of the Coxeter matrix entries $m_{s,t}$ are greater than 3—there is probably no better algorithm than the one sketched above, but incorporating the simple modifications concerning descents and locks. In fact, one can even apply it when the Coxeter group is the Weyl group of an integral root system, taking into account a different realization of the group. In this case, all doubled dot products $2(\lambda \bullet \alpha)$ will be integers. When $\lambda \prec s_\alpha \lambda$ this doubled dot-product will be a negative integer, and if it is anything other than -1 the root $s\lambda$ will no longer be minimal. What could be simpler?

If there are entries $m_{s,t} > 3$ and we are using the standard realization, however, we have to deal with non-rational cyclotomic arithmetic. The algorithm described still works, at least in principle, although there are a few annoyances that arise. All dot products in the range $(-1, 0)$ will be of the form $\cos(\pi k/m_{s,t})$ for $k < m_{s,t}/2$. So in order to tell whether $s\lambda$ dominates α_s we just have to compare the dot product $\lambda \bullet \alpha_s$ to numbers in this finite set. There is one potentially nasty problem, however. All we can assert *a priori* about the coordinates of λ is that they will be in the real cyclotomic field generated by $\cos(2\pi/M)$, where M is the lowest common multiple of the numbers $2m_{s,t}$. The amount of work involved would seem to grow disturbingly as M grows. I call this the problem of **mixed cyclotomy**.

There is one problem related to this that I might have mentioned earlier. It is difficult to tell whether real elements of a cyclotomic field are positive or not if they are represented as floating point real numbers. Thus one has to be careful. There are some obvious ways to get around this—for example, working very high precision—but they are inefficient. So at first sight it looks like we are faced with a dilemma. We want to work with exact cyclotomic arithmetic to avoid rounding errors, but will then have trouble detecting the sign of numbers and will also face intricate calculations. Brink's thesis tells us that both these apparent difficulties are spurious. In the next section I'll digress somewhat to explain how the technique introduced here can be used in another computation, but in the one after that I'll follow Brink in examining more carefully the coefficients of minimal roots.

6. Integral root systems

I'll explain in this section how the technique explained just now to avoid table look-up in constructing minimal roots of a Coxeter group can also be used to construct the root graph for the system of real roots of a Kac-Moody algebra, at least as much of it as is practical.

Let Δ be a basis of a free \mathbb{Z} -module L of finite rank, and suppose that to each α in Δ is associated a vector α^\vee in the dual lattice \widehat{L} , such that (a) $\langle \alpha, \alpha^\vee \rangle = 2$; (b) $\langle \alpha, \beta^\vee \rangle = 0$ if and only if $\langle \beta, \alpha^\vee \rangle = 0$; (c) $\langle \alpha, \beta^\vee \rangle \leq 0$ for $\alpha \neq \beta$. Then for each α the linear transformation

$$s_{\alpha^\vee}: \lambda \longmapsto \lambda - \langle \lambda, \alpha^\vee \rangle \alpha$$

is a reflection. Let

$$n_{\alpha, \beta} = \langle \alpha, \beta^\vee \rangle \langle \beta, \alpha^\vee \rangle.$$

It is non-negative. The group W generated by these reflections is a Coxeter group with

$$m_{\alpha, \beta} = \begin{cases} 2 & \text{if } n_{\alpha, \beta} = 0 \\ 3 & \text{if } n_{\alpha, \beta} = 1 \\ 4 & \text{if } n_{\alpha, \beta} = 2 \\ 6 & \text{if } n_{\alpha, \beta} = 3 \\ \infty & \text{if } n_{\alpha, \beta} \geq 4 \end{cases}$$

The transforms of elements of Δ by W are called the **roots** of the system, and they are the **real roots** of the Kac-Moody algebra determined by the Cartan matrix $(\langle \alpha, \beta^\vee \rangle)$. These roots are in bijection with those constructed earlier from the standard Cartan matrix $(-2 \cos \pi / m_{\alpha, \beta})$, but the corresponding linear representations of W won't be isomorphic unless the matrix $C = (\langle \alpha, \beta^\vee \rangle)$ is symmetrizable—that is to say DC is symmetric for some diagonal matrix D with positive diagonal entries. This happens rarely.

As earlier, the root graph has as its nodes the roots λ , and an oriented edge $\lambda \rightarrow s_\alpha \lambda$ if $\langle \lambda, \alpha^\vee \rangle < 0$ and a loop from λ to itself labelled by α if $s_\alpha \lambda = \lambda$ or equivalently $\langle \lambda, \alpha^\vee \rangle = 0$. This root graph is isomorphic to the one constructed from the standard Cartan matrix, but in constructing it one might want to keep track of explicit data such as the explicit expression $\lambda = \sum \lambda_\alpha \alpha$, so there is some point to describing its construction directly in terms of these coefficients.

The goal now is to find all roots up to a certain depth d . The processes of installation and finishing are pretty much the same, except that we have to include as part of the data attached to a root its depth as well as its index of coefficients, so we can avoid installing roots of depth $> d$. The important point is that we can use the same trick described before in order to tell for a new non-dihedral root $\mu = s\lambda$ with $t\lambda \prec \lambda$ whether $t\mu < \mu$ or not by calculating $t\lambda$, $st\lambda$, etc. until we get to a root of minimal depth in the $W_{s,t}$ -orbit of λ .

Incidentally, as the diagram for the root graph in §3 shows, root graphs exhibit an interesting repetitive structure. Howlett tells me that he has shown that it is always the set of paths through a finite automaton. One can in effect describe the whole root graph in finite, relatively simple terms.

7. Root coefficients

We want to know something about what possible values can occur for the coefficients of roots and, especially, minimal roots. Roots are produced by a sequence of reflections from the basic roots in Δ . As is the case in many algorithmic processes, although we understand each single step—here, reflection—quite well, the overall development is not so clear.

We begin by looking at dihedral roots.

Suppose S to possess exactly two elements s and t , corresponding to roots α and β . Let $m = m_{s,t}$. What are all the positive roots of the system? Let

$$z = \zeta_{2m} = e^{2\pi i/2m} = e^{\pi i/m} = \cos(\pi/m) + i \sin(\pi/m)$$


so that

$$-2\alpha \bullet \beta = 2 \cos(\pi/m) = (z + z^{-1}) = \text{(say) } c$$

and

$$\begin{aligned} s\lambda &= \lambda - 2(\lambda \bullet \alpha)\alpha \\ 2(p\alpha + q\beta) \bullet \alpha &= 2p - cq \\ s(p\alpha + q\beta) &= [-p + cq]\alpha + q\beta \\ t\lambda &= \lambda - 2(\lambda \bullet \beta)\beta \\ 2(p\alpha + q\beta) \bullet \beta &= 2q - cp \\ t(p\alpha + q\beta) &= p\alpha + [-q + cp]\beta. \end{aligned}$$

We can express this in a diagram



Suppose $m = \infty$. Then $z = 1$ and $c = z + z^{-1} = 2$. We get in succession (in tabular form) as transforms of α

$\lambda = c_\alpha \alpha + c_\beta \beta$	c_α	c_β
α	1	0
$t\alpha$	1	2
$st\alpha$	3	2
...
$(st)^n \alpha$	$2n + 1$	$2n$
$t(st)^n \alpha$	$2n + 1$	$2n + 2$
...		

and similarly for the transforms of β . Of these, only α and β are minimal roots.

If $m = 2$, the positive roots are α and β .

Now suppose $2 < m < \infty$. Then we get

λ	c_α	c_β
α	1	0
$t\alpha$	1	$z + z^{-1}$
$st\alpha$	$z^2 + z^{-2} + 1$	$z + z^{-1}$
...
$(st)^n \alpha$	C_{2n+1}	C_{2n}

$$t(st)^n\alpha \qquad C_{2n+1} \qquad C_{2n+2}$$

where the C_n are solutions of the difference equation $C_{n+2} = -C_n + (z + z^{-1})C_{n+1}$ with initial values $C_0 = 0, C_1 = 1$, giving

$$C_n = \frac{z^n - z^{-n}}{z - z^{-1}} = z^{n-1} + z^{n-3} + \dots + z^{-(n-3)} + z^{-(n-1)} = \frac{\sin(\pi n/m)}{\sin(\pi/m)}.$$

The first few are

$$0, 1, c, c^2 - 1, \dots$$

Note that $C_n > C_{n-1}$ as long as the real part of z^n is positive, which happens as long as $n < m/2$. Note also that $C_0 = C_m = 0$ and that more generally $C_n = C_{m-n}$. There are m roots in all for this dihedral system.

For $m = 3$ we get the table

λ	c_α	c_β
α	1	0
β	0	1
$t\alpha = s\beta$	1	1

For $m = 4$:

λ	c_α	c_β
α	1	0
β	0	$\frac{1}{2}$
$t\alpha = st\alpha$	$\frac{1}{2}$	$\frac{\sqrt{2}}{2}$
$s\beta = ts\beta$	$\frac{\sqrt{2}}{2}$	1

For $m = 5$, with $\theta = (1 + \sqrt{5})/2$:

λ	c_α	c_β
α	1	0
β	0	1
$t\alpha$	1	θ
$s\beta$	θ	1
$st\alpha = ts\beta$	θ	θ

For $m = 6$:

λ	c_α	c_β
α	1	0
β	0	$\frac{1}{2}$
$t\alpha$	$\frac{1}{2}$	$\frac{\sqrt{3}}{2}$
$s\beta$	$\frac{\sqrt{3}}{2}$	$\frac{1}{2}$
$st\alpha = tst\alpha$	$\frac{2}{3}$	$\frac{\sqrt{3}}{2}$
$ts\beta = sts\beta$	$\frac{\sqrt{3}}{2}$	1

For $m \geq 7$:

λ	c_α	c_β
α	1	0
β	0	1
$t\alpha$	1	c_m
$s\beta$	c_m	1
$st\alpha = tst\alpha$	$c_m^2 - 1$	c_m
$ts\beta = sts\beta$	c_m	$c_m^2 - 1$

Here $c_m = 2 \cos(\pi/m)$, and keep in mind that $c_m^2 - 1 > 2$ for $m \geq 7$. The table for $m > 6$ looks somewhat similar to the last, except that it is longer. All coefficients beyond these first 6 rows are greater than 2.

We can use these results, following Brink's thesis, to prove:

7.1. Proposition. *If a positive root is expressed as a sum $\sum c_\alpha \alpha$ of basic roots, then each c_α is a sum of products of the form $c_1 \dots c_k$ where each c_i is of the form*

$$\frac{\sin(\pi n/m)}{\sin(\pi/m)} = \frac{\zeta_{2m}^n - \zeta_{2m}^{-n}}{\zeta_{2m} - \zeta_{2m}^{-1}} = \zeta_{2m}^{n-1} + \zeta_{2m}^{n-3} + \dots + \zeta_{2m}^{-(n-1)} \quad (1 \leq n \leq m/2)$$

where m is one of the $m_{s,t}$.

For $n = 1$ this expression is 1, so we may also allow c_i to be a positive integer. For $n = 2$ it is $c_m = 2 \cos(\pi/m)$ and for $n = 3$ it is $c_m^2 - 1$, so this agrees with our calculations.

PROOF. By induction on the depth of a root. If the depth is 1 there is no problem. Otherwise suppose

$$\lambda = w\alpha = s_1 \dots s_{n-1} \alpha_t$$

with $\alpha = \alpha_t$ in Δ , $\delta(\lambda) = n$. Let $s = s_{n-1}$. Let x be the minimal element in $wW_{s,t}$, and let y be in $W_{s,t}$ with $w = xy$. Then $x\alpha_s > 0$, $x\alpha_t > 0$. We know from the discussion of dihedral roots what the coefficients of

$$y\alpha_t = p\alpha_s + q\alpha_t$$

are like, but then

$$w\alpha_t = px\alpha_s + qx\alpha_t.$$

We can apply an induction hypothesis to $x\alpha_s$ and $x\alpha_t$ since $\ell(x) < \ell(w)$. **QED**

7.2. Lemma. *The function $\sin(x)/\sin(ax)$ is strictly decreasing in $[0, \pi/2]$.*

PROOF. Its derivative is

$$\begin{aligned} \frac{(\sin x)' \sin(ax) - (\sin(ax))' \sin(x)}{\sin^2(ax)} &= \frac{\sin(x)}{\sin(ax)} \left(\frac{\cos x}{\sin x} - \frac{a \cos(ax)}{\sin(ax)} \right) \\ &= \frac{\sin(x)}{\sin(ax)} \left(\frac{1}{\tan x} - \frac{a}{\tan(ax)} \right) < 0 \end{aligned}$$

since

$$a \tan x > \tan ax$$

in $(0, \pi/2)$.

7.3. Corollary. *The smallest positive root coefficient is 1. Any root coefficient lying between 1 and 2 equals $2 \cos(\pi/m_{s,t})$ for some s, t .*

PROOF. Because of the Lemma, the minimum value of $\sin(x)/\sin(ax)$ in $[0, \pi/2]$ is $1/\sin(\pi a/2)$. If $a = 1/n$ with $n \geq 3$, this lies at or above 2. Therefore all possible values of $\sin(x)/\sin(x/n)$ for $0 \leq x \leq \pi/2$ in the range $(1, 2)$ are with $n = 2$, when this is just $2 \cos(x/2)$. The minimum value achieved is $2 \cos(\pi/4) = \sqrt{2}$. Any root coefficient will be a sum of terms

$$C = c_1 \dots c_k$$

where $k \geq 1$ and each c_i will be either a positive integer or one of the constants in the Proposition. If C itself is not 1 then we may assume that each $c_i > 1$. The smallest possible value for one of these constants is 1, and the second smallest possibility is $\sqrt{2}$. So if C is less than 2 then $k = 1$, and is an integer only if it is 1. **QED**

7.4. Corollary. *The smallest possible coefficient of a positive root, other than 1, is $\sqrt{2}$.*

8. The support of minimal roots

One of the principal results of Brink's thesis is that the apparent difficulties concerning mixed cyclotomy in the determination of the minimal roots do not actually arise—we need never do arithmetic in cyclotomic fields other than the ones generated by some $m_{s,t}$ -th roots of unity.

In order to see why this is so, we have to look carefully at how roots are calculated.

A root $\sum \lambda_\alpha \alpha$ may be considered as a function on the Coxeter graph: $s \mapsto \lambda_{\alpha_s}$. The **support** of a root $\lambda = \sum \lambda_\alpha \alpha$ is its support as a function—the set of s where $\lambda_{\alpha_s} \neq 0$. One of Brink's nice ideas is to track the computation of roots by looking at their support. All roots are constructed by applying some element of W to a basic root, and we are going to construct them by ascending the root graph.

Upon reflection by s , only the coefficient λ_α changes, and to

$$-\lambda_\alpha + \sum_{\lambda_\beta \neq 0, \alpha \sim \beta} [-2 \alpha \cdot \beta] \lambda_\beta .$$

According to Corollary 3.4, the support of a root cannot decrease as we go up the root graph. Under what circumstances it is extended? In going up the root graph from λ to $s\lambda$, the support of λ is extended if s is not already in its support. If s is not connected to that support in the Coxeter graph, then $s\lambda = \lambda$. This process gives a new root only when s is linked to the support of λ . Therefore:

8.1. Proposition. *The support of any root is connected.*

Suppose we apply a reflection s to a positive root λ whose support does not contain s . The formula for reflection is

$$s_\alpha \lambda = \lambda - 2(\lambda \cdot \alpha) \alpha$$

where

$$\alpha \cdot \lambda = \sum_{\beta} \lambda_\beta (\alpha \cdot \beta) = \lambda_\alpha + \sum_{\lambda_\beta \neq 0, \alpha \sim \beta} \lambda_\beta [\alpha \cdot \beta]$$

if $\lambda = \sum \lambda_\alpha \alpha$. How does the support of a root get extended? If λ is a minimal root with $\lambda_{\alpha_s} = 0$, then $s\lambda$ will also be a minimal root if and only if $-1 < \lambda \cdot \alpha$, or equivalently

$$\sum_{\lambda_\beta \neq 0, \beta \sim \alpha} [-2 \alpha \cdot \beta] \lambda_\beta < 2 .$$

But $\lambda_\beta \geq 1$ and also $-2\alpha \bullet \beta \geq 1$. The condition for minimality cannot hold here unless there is just one $\beta = \alpha_t$ in the sum. Furthermore, in that case the condition above amounts to requiring that

$$[-2\alpha \bullet \beta] \lambda_\beta < 2.$$

The first term is either 1 or $\geq \sqrt{2}$, as is the coefficient λ_β . All in all:

8.2. Proposition. *Suppose λ to be a minimal root, s an element of S not in the support Θ of λ . If $s\lambda$ is again minimal then there must be just a single link from s to Θ , say to t , and one of these holds:*

- (a) $m_{s,t} = 3$ and $\lambda_{\alpha_t} < 2$;
- (b) $3 < m_{s,t} < \infty$ and $\lambda_{\alpha_t} = 1$.

We can summarize this in a mnemonic diagram:

$$\begin{aligned} \bullet \quad 1 &\text{---} 0 & \mapsto & 1 &\text{---} 1 \\ \bullet \quad c &\text{---} 0 & \mapsto & c &\text{---} c \quad (1 < c < 2) \\ \bullet \quad 1 &\overset{m}{\rightleftarrows} 0 & \mapsto & 1 &\overset{m}{\rightleftarrows} c \quad (c = 2 \cos(\pi/m)) \end{aligned}$$

where --- means an edge of multiplicity 3 and \rightleftarrows one of higher multiplicity.

8.3. Corollary. *The support of a minimal root is a tree containing no links of infinite degree.*

9. Composition and decomposition

The most important of Brink's ideas is to understand that minimal roots can be assembled in a simple way from relatively simple ones.

The trick is to focus on the coefficients of a minimal root that are equal to the minimum possible value 1. This leads to a remarkable view of the structure of a minimal root, which I'll explain in the rest of this section. Suppose λ to be a minimal root with $\lambda_\alpha = 1$. Let $\alpha = \alpha_s$. We know that λ has been built up from a chain of minimal root predecessors

$$\lambda_1 \prec \dots \prec \lambda_n = \lambda.$$

Some one of these, say $\mu = \lambda_k$, will be the first to have a non-zero coefficient of α , and that coefficient must be 1, since by Corollary 3.5 coefficients can only grow as we move up the root graph. Unless $\mu = \lambda_1 = \alpha$, the root μ must have been constructed by extension from a root λ_{k-1} whose support did not contain s , and the earlier Proposition tells us that $\mu = s\lambda_{k-1}$, where s is linked to the support of λ_{k-1} at a single node t by a simple link in the Coxeter graph, and $\lambda_{k-1, \alpha_t} = \lambda_{k, \alpha_t} = 1$. We have $\lambda = w\mu$ where w is a product of generators other than s .

Since the support Θ of λ is a tree, the complement of s in Θ is the disjoint union of exactly as many components Θ_i^\bullet as there are edges in Θ linked to s , say ℓ . Reflections corresponding to nodes in different components will commute. Hence we can write $\lambda = w_1 \dots w_\ell \mu$ where each w_i is in $W_{\Theta_i^\bullet}$. All the w_i commute, so we can write the product in any order. Each of the $\lambda_i = w_i \alpha$ is itself a root with support in $\Theta_i = \Theta_i^\bullet \cup \{s\}$, minimal because it will be a predecessor of λ in the root graph. We have therefore proven:

9.1. Proposition. *Let λ be a minimal root with $\lambda_{\alpha_s} = 1$. Let Θ be the support of λ , the Θ_i^\bullet the connected components of $\Theta \setminus \{s\}$, $\Theta_i = \Theta_i^\bullet \cup \{s\}$. The restriction λ_i of λ to Θ_i is also a minimal root.*

This is called the **decomposition** of λ at s . Conversely:

9.2. Proposition. *Let Θ be a subset of S , s an element of Θ , the Θ_i^\bullet the components of $\Theta \setminus \{s\}$, $\Theta_i = \Theta_i^\bullet \cup \{s\}$. If for each i the minimal root λ_i has support on Θ_i with coefficient 1 at α_s , then the vector λ with support Θ whose restriction to each Θ_i agrees with λ_i is a minimal root.*

This is called **composition** or **amalgamation** at s .

PROOF. More difficult than the previous result. It follows from this result of Brink's:

9.3. Lemma. *If λ is any positive root with coefficient 1 at $\alpha = \alpha_s$, then $\lambda = w\alpha$, where w is a product of $t \neq s$ in S .*

PROOF. Induction on the depth of λ . Trivial if $\lambda = \alpha$. It remains to be shown that if $\delta(\lambda) > 1$ and $\lambda_\alpha = 1$ then we can find s such that $s\lambda \prec \lambda$ and $s\lambda$ also has coefficient 1 at α . There exists at any rate some s with $s\lambda \prec \lambda$. If this does not have coefficient 1 at α , then—as we have seen above— $s\lambda$ has coefficient 0 at α ; α has exactly one link to the support of $s\lambda$, say to β ; and the coefficient of $s\lambda$ at β is equal to 1. By induction, $s\lambda = w\beta$ where w is a product of t equal neither to s nor s_β . But then

$$\lambda = sw\beta = ws\alpha . \text{ QED}$$

These results have a natural generalization, allowing us to compose and decompose at collections of more than one node. The **unit support** of λ is the set of α with $\lambda_\alpha = 1$. Suppose λ to be a minimal root, Θ its support, $T \subset \Theta$ its unit support. We know that Θ is a connected tree. The complement $\Theta \setminus T$ can be expressed as the disjoint union of its connected components Θ_i^\bullet . For each of these, let Θ_i be its union with the elements of T attached to it. Each set Θ_i will then be connected, and no point of T will be in its interior. The set T will contain in addition some pairs of nodes of the Coxeter graph connected by links in it; let these be counted as additional sets Θ_i (but with empty interiors). Finally, there is one exceptional case where Θ has just one element, which is Θ_1 .

9.4. Lemma. *The set Θ is the union of the Θ_i , which overlap only in points of T .*

PROOF. We may suppose Θ to have more than one element in it. Suppose given t in Θ . If it is not in T , it belongs to a unique Θ_i^\bullet . Otherwise, it will be in T . If it is connected in the Coxeter diagram to one of the Θ_i^\bullet , it will be in Θ_i . Finally, we have the situation where the only elements of Θ it is linked to are in T . Each of the pairs forming such a link will be one of the Θ_i . QED

If T is any subset of S , a subset Θ is called **T-connected** if connected and none of its interior points belong to T . Each of the Θ_i defined above are T -connected, and are called the **T-components** of Θ . Following Brink, I call a minimal root **indecomposable** if its support is T -connected, or in other words if its unit support is contained in the boundary of its support.

9.5. Lemma. (Brink) (Decomposition) *The restriction of a minimal root to any of the T -components of its support is again a minimal root.*

PROOF. This follows by induction from the earlier decomposition. QED

As does this:

9.6. Proposition. (Brink) (Composition) *Let Θ be any connected, simply connected subset of S , T an arbitrary subset of Θ , Θ_i its T -components. If λ_i is for each i a minimal root with support Θ_i and unit support $T \cap \Theta_i$, then the vector λ with support Θ whose restriction to each Θ_i is λ_i will be a minimal root.*

This reduces the calculation of minimal roots to the calculation of indecomposable ones, together with a bit of accounting—i.e. joining indecomposable ones together in their unit supports. Finding the reflection table turns out not to be much more difficult.

It is useful to restate the criterion here:

- **(Indecomposability criterion)** Suppose λ to be a minimal root with support Θ . Then λ is indecomposable if and only if either (a) $\Theta = \{s\}$ or (b) every node of Θ in the unit support of λ is connected to exactly one other node in Θ .

10. Constructing the indecomposable minimal roots

So now the question we have to deal with is, *How are indecomposable roots to be found?* They are obtained by reflection from predecessors in the root graph. The predecessors may or may not be indecomposable. What can we say about the process of forming them? It takes place in simple steps, so we just have to understand the simple steps, going up the root graph one edge at a time. The most important question is this: *If λ is an indecomposable root and $\mu = s\lambda \prec \lambda$, what can we say about μ ?* Let Θ be the support of λ , and let $\alpha = \alpha_s$. There are several possibilities.

- Suppose that s is a boundary point of Θ with $\lambda_\alpha = 1$. Then according to Proposition 8.2 the root λ is obtained from μ by **extension** along a simple edge from a unit node t of μ . This is also a unit node of λ . By the indecomposability criterion this can happen only if $\mu = \alpha_t$.
- Suppose that s is in the boundary of Θ but that $\lambda_\alpha > 1$. Let t be the unique node that s is connected to. There are two possibilities: (a) the support of μ is one node less than that of λ or (b) it is the same. Case (a) also breaks down by Proposition 8.2 into two possibilities: (i) the node t is in the unit support of μ , hence of λ . Since λ is indecomposable $\mu = \alpha_t$, the link from s to t has multiplicity $m > 3$, and $\lambda_\alpha = 2 \cos(\pi/m)$. (ii) The link from s to t is simple, and $\lambda_\alpha = \lambda_{\alpha_t}$. In case (b) we can deduce at least that μ is also an indecomposable root.
- Suppose α lies in the interior of Θ . Then by Proposition 8.1 it cannot happen that $\mu_\alpha = 0$. Suppose $\mu_\alpha > 1$. Then μ is also indecomposable.
- Finally, suppose α lies in the interior of Θ but $\mu_\alpha = 1$. The support of μ is also Θ , so μ cannot be indecomposable. It will be the composite of two or more indecomposable roots μ_i adjoined to each other at α .

Running this backwards:

10.1. Proposition. *Every indecomposable root is constructed through a sequence of reflections $\mu \prec s\mu = \lambda$ falling into one of three possible types:*

- (1) **extension**, where the support of μ is less than that of λ ;
- (2) **promotion**, where μ is indecomposable and has the same support as λ ;
- (3) **fusion**, where μ is the composite of indecomposable roots μ_i each with α_s in its unit support.

We have already investigated extension in detail, but here I point out that if it is going to lead to an indecomposable root it has to be the process

$$\bullet \quad c \bullet \longrightarrow 0 \quad \longmapsto \quad c \bullet \longrightarrow c \quad (1 < c < 2).$$

There is nothing particularly subtle about promotion. In the next section we look at fusion, about which we can say something interesting.

11. Reflecting at junctions

This section is unavoidably technical, but also crucial towards understanding how the eventual program will work. Let me try to summarize how things go, before starting to plow.

Suppose we are given a minimal root λ with support Θ as well as an elementary reflection s . The basic question throughout this paper is whether $s\lambda$ is again a minimal root. Let $\alpha = \alpha_s$. There are three cases:

- (a) If s is in the support of λ , and the coefficient $\lambda_\alpha > 1$, we are looking at the process of promotion. This node will occur in the interior of one of the Brink components of Θ , so we are looking at promotion of this component. We might have to do some cyclotomic arithmetic to carry it out, but the structures we are looking at are not going to change much—at most we have to change a single root coefficient.
- (b) The generator s is not in Θ . If $s\lambda$ is again indecomposable, s must be linked to Θ by a single edge with certain restrictions we have already mentioned.
- (c) The node s is in the support of λ , but the coefficient $\lambda_\alpha = 1$. If this node is on the boundary of Θ , we are again looking at promotion, and again of a single component of Θ . Otherwise, this node has links to at least two nodes of Θ . This is the case I have called fusion, and it is the case we are going to examine in this section. The main consequence of the investigation is that mixed cyclotomy does not occur, which is extremely important to know when programming. Also, the code for fusion follows the treatment below very closely.

Suppose Θ to be a connected and simply connected subset of S , s an interior node of Θ , the Θ_i the connected components of $\Theta \setminus \{s\}$. Suppose that λ is a minimal root with support Θ and $\lambda_\alpha = 1$ where $\alpha = \alpha_s$. Let λ_i be the restriction of λ to Θ_i , and assume the λ_i are indecomposable. In this section we'll answer the following question:

- *In these circumstances, when is $s\lambda$ a minimal root?*

Each of the Θ_i contains a unique node t_i linked to s . Let $\beta_i = \alpha_{t_i}$. The sort of things we have to answer are: *How many components Θ_i are allowed? What are the possible values for λ_{β_i} ?*

The inner product $\lambda \bullet \alpha$ is

$$1 - \sum \lambda_{\beta_i} |\alpha \bullet \beta_i|.$$

Recall that $\lambda_{\beta_i} \geq 1$, $|\alpha \bullet \beta_i| \geq 1/2$. When the junction has more than 3 roots joining at it, this dot-product is ≤ -1 , and $s\lambda$ will not be minimal.

11.1. Proposition. *If $\Theta \setminus \{s\}$ has more than 3 components, $s\lambda = \oplus$.*

We may now assume there to be 2 or 3 links at α .

11.2. Lemma. *Suppose $c = 2 \cos(\pi/m_{s,t})$. The links $c \bullet \rightarrow 1$ and $1 \bullet \rightarrow 1$ cannot occur in any root.*

This introduces notation I hope to be self-explanatory. What I mean by this is that if λ is a root then there do not exist neighbouring nodes s and t in the Coxeter graph with (1) $m_{s,t} = 3$, $\lambda_{\alpha_s} = 1$, $1 < \lambda_{\alpha_t} < 2$; or (2) $m_{s,t} > 3$, $\lambda_{\alpha_s} = 1$, $\lambda_{\alpha_t} = 1$.

PROOF. Lemma 9.3 implies that $1 \bullet \rightarrow 1$ cannot occur, and then implies that if $c \bullet \rightarrow 1$ occurs, so does $1 \bullet \rightarrow 1$. **QED**

We now consider the possibilities for fusion.

Fusion from two links

The three *a priori* local possibilities for λ are (with s in the middle)

$$\begin{aligned} \text{(2.a)} \quad & \lambda_x \bullet \longrightarrow 1 \bullet \longrightarrow \lambda_y \\ \text{(2.b)} \quad & \lambda_x \bullet \longrightarrow 1 \bullet \longleftarrow \lambda_y \\ \text{(2.c)} \quad & \lambda_x \bullet \longleftarrow 1 \bullet \longleftarrow \lambda_y \end{aligned}$$

with λ_* .

Case (2.c) We can eliminate one of these immediately.

11.3. Proposition. *In the case (2.c) $s\lambda = \oplus$.*

PROOF. The inner product $\lambda \bullet \alpha$ is $1 - c_x \lambda_x / 2 - c_y \lambda_y / 2$. We know that each c_* is at least $\sqrt{2}$, and we know by the previous Lemma that $\lambda_* \geq \sqrt{2}$ as well. Thus the dot product is -1 or less. **QED**

Case (2.a) The diagram is $\lambda_x \bullet \longrightarrow 1 \bullet \longrightarrow \lambda_y$.

In this case, the condition that $s\lambda$ be minimal is that $1 - \lambda_x/2 - \lambda_y/2 > -1$, or $\lambda_x + \lambda_y < 4$. The case $\lambda_* = c$ is excluded by the previous Lemma, so either $\lambda_* = 1$ or $\lambda_* \geq 2$. It is not possible for both λ_* to be ≥ 2 , so one at least must be 1, say $\lambda_y = 1$.

We can do better: the inner product must be $-\cos(\pi k/m)$ for $k \leq m/2$. This gives us:

$m = 3$. The inner product is 0, with

$$\bullet \quad 1 \bullet \longrightarrow 1 \bullet \longrightarrow 1 \quad \longmapsto \quad 1 \bullet \longrightarrow 1 \bullet \longrightarrow 1$$

(which does not produce a new indecomposable root) or $-1/2$ with

$$\bullet \quad 2 \bullet \longrightarrow 1 \bullet \longrightarrow 1 \quad \longmapsto \quad 2 \bullet \longrightarrow 2 \bullet \longrightarrow 1 .$$

$m = 4$. We must have $1/2 - \lambda_x/2 = -\sqrt{2}/2$ or $\lambda_x = 1 + \sqrt{2}/2$. Impossible.

$m = 5$. $1 - \lambda_x = -c$ or $1 - c$, $\lambda_x = 1 + c$ or $\lambda = c$. The second is impossible, alas. So the only possibility is

$$\bullet \quad (1+c) \bullet \longrightarrow 1 \bullet \longrightarrow 1 \quad \longmapsto \quad (1+c) \bullet \longrightarrow (1+c) \bullet \longrightarrow 1 .$$

$m \geq 6$. Nothing possible.

Case (2.b) Here the diagram is

$$\lambda_x \bullet \longrightarrow 1 \bullet \longleftarrow \lambda_y ,$$

and the condition is

$$1 - \lambda_x/2 - (c/2)\lambda_y > -1, \quad \lambda_x + c\lambda_y < 4 .$$

$m = 4$. Only possibility is

$$\bullet \quad 1 \bullet \longrightarrow 1 \bullet \longleftarrow \sqrt{2} \quad \longmapsto \quad 1 \bullet \longrightarrow 2 \bullet \longleftarrow \sqrt{2} .$$

$m = 5$. Only possibility is

$$\bullet \quad 1 \bullet \rightarrow 1 \bullet \rightleftharpoons c \quad \mapsto \quad 1 \bullet \rightarrow (1+c) \bullet \rightleftharpoons c .$$

We can nicely summarize the results:

11.4. Proposition. *If λ is the composition of two indecomposable roots at s , then $s\lambda = \oplus$ unless one of the two components is $1 \bullet \rightarrow 1$, and in that case we are looking at one of the reflections expressed locally as*

$$\begin{array}{ll} 1 \bullet \rightarrow 1 \bullet \rightarrow 2 & \mapsto \quad 1 \bullet \rightarrow 2 \bullet \rightarrow 2 \\ 1 \bullet \rightarrow 1 \bullet \xrightarrow{4} \sqrt{2} & \mapsto \quad 1 \bullet \rightarrow 2 \bullet \xrightarrow{4} \sqrt{2} \\ 1 \bullet \rightarrow 1 \bullet \rightarrow (1+c) & \mapsto \quad 1 \bullet \rightarrow (1+c) \bullet \rightarrow (1+c) \\ 1 \bullet \rightarrow 1 \bullet \xrightarrow{5} c & \mapsto \quad 1 \bullet \rightarrow (1+c) \bullet \xrightarrow{5} c \end{array}$$


Note that in all these cases we start with λ a composition of two indecomposable roots, one of which is just $1 \bullet \rightarrow 1$. So it is trivial to reconstruct the reflection even when this is all embedded in a generic root.

Fusion from three links

Assume the product of the reflection to be an indecomposable root. The conditions are $\lambda_x + \lambda_y + \lambda_z < 4$ if all links are simple, or $\lambda_x + \lambda_y + c_m \lambda_z < 4$ if one has degree $m > 3$. More of degree $m > 3$ cannot occur by simple calculation. If $m > 3$ then $c \geq \sqrt{2}$ so the second sum is at least 4, and cannot occur.

As for the first, we cannot have any $\lambda_* \geq 2$, so each can only be 1 or c . But the link $c \bullet \rightarrow 1$ cannot occur.

11.5. Proposition. *The only triple junction producing a minimal root by reflection at the junction is the case that all coefficients are 1. It reflects to a central node with value 2.*



The only triple junction giving rise to fusion.

Here is for us the principal result of Brink's thesis:

11.6. Theorem. *No indecomposable minimal root has more than one link of degree $m > 3$ in its support.*

This implies that mixed cyclotomy does not occur in computing the reflection $s\lambda$ of a minimal root λ .

PROOF. It is implicit in observations already made. It goes by induction on the size of the support. The Theorem is trivial for dihedral roots.

Suppose it to be true for all roots of size $< n$, and suppose we are looking at an indecomposable root of size n . It arose in one of these ways:

- Extension from an indecomposable root. If the number of multiple links is to rise, then we must have extended by a multiple root. But by Proposition 8.2 any extension by a multiple link is decomposable.
- Promotion. The support is not extended.
- Fusion of two links. In this case, according to Proposition 9.1 one of the roots being fused is a simple dihedral root, so there is no problem here, either.
- Fusion of three links. According to Proposition 9.2 the roots being fused are all simple dihedral roots. **QED**

Brink's thesis at this point goes on to build a rather explicit description of all possible indecomposable roots whose support is a given Coxeter graph. It is probably a poor idea to try to use these lists in a computer program, since it is easy enough to use basic properties of indecomposable roots to build these lists automatically.

There are a few important consequences of what we have seen so far that are not quite obvious. Keep in mind that the qualities of a root depends only on its support, not on how it is contained in a larger system. We know that a multiple link occurs no more than once in the support of an indecomposable root. If the degree is high enough, its possible location is also very much restricted.

11.7. Lemma. *Suppose $m > 6$, and suppose λ to be an indecomposable root with a multiple link in its support. Then any chain in the root graph to λ from Δ starts with one of the ends of the multiple link, which lies on the boundary of the support of λ .*

PROOF. Suppose $\mu \prec s\mu$ is the edge in the root graph in which the multiple link is added. This must be an extension $1 \iff 0 \mapsto 1 \iff c$. Let s be the left node here, t the right one. If μ is other than α_s , then because λ is indecomposable the reflection s must occur later in the chain, which can be easily seen to be an impossibility. A similar calculation shows that s remains at the end of the support of λ . **QED**

The following also follows from an easy calculation:

11.8. Proposition. *If λ is an indecomposable root whose construction starts with the node s and adds a link of degree $m \geq 7$ from s to t in the first step, then either (a) $\lambda_{\alpha_s} = 1$ and $\lambda_{\alpha_t} = nc_m$ for some integer n or (b) $\lambda_{\alpha_s} = c_m^2 - 1$ and $\lambda_{\alpha_t} = c_m$.*

11.9. Corollary. *If λ is an indecomposable root with an end $1 \iff nc_m$ or $(c_m^2 - 1) \iff c_m$ with $m \geq 7$, then all other coefficients of λ are integer multiples of c_m .*

11.10. Corollary. *Suppose λ to be an indecomposable non-dihedral root, and that it is not one of the roots occurring in the previous corollary. Then every coefficient of λ is of the form $a + bc_m$ with a and b integers.*

PROOF. If $m \leq 6$ the field generated by c_m is quadratic. **QED.**

Thus, with a bit of fiddling to handle the case $m \geq 7$ carefully, the three integers (a, b, m) are sufficient to describe any coefficient of a non-dihedral minimal root. This is a drastic version of the principle of 'no mixed cyclotomy'. Implementing this idea in an explicit program is straightforward enough, but the program to be sketched in the next section, which uses composition and decomposition in a more sophisticated fashion, seems to be more efficient.

12. The program

I am going to sketch here a program that lists all minimal roots and constructs the reflection table. It, too, is unavoidably technical, as details of efficient programming nearly always are.

The most important decision to be made is to answer the question, *What is a root?* Relying strongly on Brink's thesis, this program works with two kinds of roots, **indecomposable** and **composite**. The associated data structures are very different. Both have indices, descent sets, and a reflection list. But an indecomposable root has in addition an actual array of coefficients. Because mixed cyclotomy doesn't occur, these coefficients will all lie in one of the real cyclotomic rings generated by $\zeta_{2m} + \zeta_{2m}^{-1}$, where $m = m_{s,t}$ for some s, t . I call these rings **elementary**. An indecomposable root is also assigned its degree m , and when $m > 6$ its two special nodes, the ones spanning the unique link of degree m , are specified as well. A composite root has, instead, a list of its indecomposable components.

New roots are found by applying a reflection to a root already installed. There are several basic ways in which to do this: (1) **extension** of an indecomposable root to produce a new indecomposable root with a larger support; (2) **promotion** of an indecomposable root to produce a new indecomposable root with the same support; (3) **fusion** at a junction of several indecomposable roots to produce a new indecomposable root; (4) **composition** of an arbitrary root with a dihedral root attached at one end to its unit support; (5) **replacement** of one or more of the components of a composite root by a single indecomposable root.

It is only the first three that require actual arithmetic, and always in one of the elementary real cyclotomic rings. Composition is an extremely simple operation, while replacement amounts to reducing a composite reflection to one of the first three. That this can be done easily seems almost an accident depending on the limited number of ways fusion can take place. I'll not give details of this below, but I'll sketch an example here. Suppose we want to calculate $s_3\lambda$ where

$$\lambda: 1 \longrightarrow 1 \longrightarrow \mathbf{1} \iff \sqrt{2} \longrightarrow \sqrt{2}$$

and $s = s_3$ is the third node from the left. The root λ is the composite of three indecomposable roots

$$1 \longrightarrow 1, \quad 1 \longrightarrow \mathbf{1}, \quad \mathbf{1} \iff \sqrt{2} \longrightarrow \sqrt{2}$$

There are two components in the **star** of s , the components of λ containing it, and they do not exhaust λ . Instead, they make up the smaller root

$$\mu: 1 \longrightarrow \mathbf{1} \iff \sqrt{2} \longrightarrow \sqrt{2}$$

whose reflections we shall have already calculated when λ is finished. It is easy enough to locate μ , given the reflection tables already made, since it is equal to s_2 applied to the indecomposable root

$$\nu: \mathbf{1} \iff \sqrt{2} \longrightarrow \sqrt{2}$$

We look up $s_2\nu$ to get μ , then look up $s\mu$, which is the indecomposable root

$$1 \longrightarrow \mathbf{2} \iff \sqrt{2} \longrightarrow \sqrt{2}.$$

Finally, we replace the two components of μ by the single component $s\mu$ in the list of components of λ . All this can be carried out, with minor modifications, for each of the cases of fusion. In the pseudo-program below the details will be grossly telescoped.

The program:

First of all we install the basic roots, those in Δ .

Then we install specially the dihedral roots, those whose support consists of exactly two nodes of the Coxeter graphs. When we are through with these we have finished the basic roots. We put the dihedral roots of depth 2 in the queue, and after we have installed all of those we put those of depth 3 in the queue. In this way, we preserve the essential property of the queue that roots of depth d are put on before those of depth $d + 1$.

As for dihedral roots of depth 4 or more, we can finish them as soon as they are installed. Both of their coefficients are 2 or more, which means they cannot be extended, so the only reflections higher in the root graph are to dihedral roots with the same support. In other words, we never put them in the queue. Incidentally, these occur only for links with $m \geq 6$. One consequence is that in calculating reflections of roots taken off the queue, we have only to do arithmetic with algebraic integers of the form $a + bc_m$, where m is the degree of the root, $c_m = 2 \cos(\pi/m)$, a and b are integers.

At this point, all roots of depth 1 have been finished, all dihedral roots have been installed, and some have been finished. The queue contains only dihedral roots. While the queue is not empty, we remove roots from it to be finished.

Let λ be a root removed from the queue. All roots of less depth have been finished already, and in particular all descents $s\lambda \prec \lambda$ have been assigned. We run through all the s in S , and for each one of them where $s\lambda$ has not been assigned, we have to calculate $s\lambda$. Exactly what this involves depends on a number of things—the basic idea is to use composition and decomposition to avoid repeating actual arithmetic.

So now suppose we are looking at s and λ with $s\lambda$ not known. Let $\alpha = \alpha_s$. There are many cases to deal with.

- The node s lies in the support of λ .
 - It belongs to a unique component λ_i .
 - That component is λ itself, which is therefore indecomposable. We must calculate the dot product $\lambda \bullet \alpha$ to see if $s\lambda$ is minimal or not, and if it is install it.
 - The component λ_i is not all of λ , which is a composite root. We have calculated $s\lambda_i$ already, and $s\lambda$ is obtained by replacing λ_i with $s\lambda_i$.
 - It belongs to several components. It is therefore a junction, and we are looking at a possible fusion. Let μ be the star of s , the composite of those components whose supports contain s .
 - The root μ is the same as λ . We have to calculate $s\lambda$ explicitly. It may be λ , \oplus , or a new indecomposable root.
 - The root μ makes up only a part of λ . We have calculated $s\mu$, and obtain $s\lambda$ by replacing μ with $s\mu$. Exactly what we do to find μ and $s\mu$ depends on which case of fusion we are dealing with, as I mentioned in the example I looked at earlier. This can be decided by local information around s .
- The node s does not lie in the support of λ .
 - It is linked to more than one node in the support, or by a link of infinite degree. Here $s\lambda = \oplus$.
 - It is linked by a single link of finite degree.
 - The link is to a unit node of λ . Then $s\lambda$ is a new composite.
 - The link is not to a unit node.
 - The root λ is indecomposable. We have to calculate explicitly.

- The root λ is composite. If μ is the component it is linked to, we have already calculated $s\mu$. We replace μ by $s\mu$ in λ .

In summary, I should say that there are several advantages to following Brink's idea of expressing minimal roots as the amalgamation of indecomposable ones. For one thing, a given indecomposable root may occur in several other roots, so redundancy of storage is eliminated. For another, we only have to do actual arithmetic in promoting nodes of indecomposable roots, and this amounts to a huge saving in time.

References

1. Brigitte Brink, 'On root systems and automaticity of Coxeter groups', Ph. D. thesis, University of Sydney, 1994.
2. Brigitte Brink, 'The set of dominance-minimal roots', available as Report 94-43 from the School of mathematics and Statistics at the University of Sydney:
<http://www.maths.usyd.edu.au:8000/res/Algebra/Bri/dom-min-roots.html>
3. Brigitte Brink, 'The set of dominance-minimal roots', *Journal of Algebra* **206** (1998), 371-412. This publication, although it has the same title as the previous item, is very different from it. For purposes of computation it is decidedly less interesting, since it doesn't explicitly introduce the indecomposable minimal roots, and emphasizes hand-lists instead.
4. Brigitte Brink and Robert Howlett, 'A finiteness property and an automatic structure for Coxeter groups', *Math. Ann.* **296** (1993), 179-190.
5. Bill Casselman, 'Automata to perform basic calculations in Coxeter groups', in *Representations of groups, CMS Conference Proceedings 16*, A.M.S., 1994.
6. Bill Casselman, 'Computation in Coxeter groups—I. multiplication', preprint, 2001. *Electronic Journal of Combinatorics* **9(1)** (2002), #25.
7. Bill Casselman, Java code for finding minimal roots at
<http://www.math.ubc.ca/~cass/coxeter.tar.gz>
8. Fokko du Cloux, 'Un algorithme de forme normale pour les groupes de Coxeter', preprint, Centre de Mathématiques à l'École Polytechnique, 1990.
9. J. Tits, 'Le problème des mots dans les groupes de Coxeter', *Symposia Math.* **1** (1968), 175-185.
10. E. B. Vinberg, 'Discrete linear groups generated by reflections', *Math. USSR Izvestia* **5** (1971), 1083-1119.