

Structure constants of Kac–Moody Lie algebras

Bill Casselman

To Nolan Wallach, wishing him many more years of achievement

Abstract This paper outlines an algorithm for computing structure constants of Kac–Moody Lie algebras. In contrast to the methods currently used for finite-dimensional Lie algebras, which rely on the additive structure of the roots, it reduces to computations in the extended Weyl group first defined by Jacques Tits in about 1966. The new algorithm has some theoretical interest, and its basis is a mathematical result generalizing a theorem of Tits about the finite-dimensional case. The explicit algorithm seems to be new, however, even in the finite-dimensional case. I include towards the end some remarks about repetitive patterns of structure constants, which I expect to play an important role in understanding the associated groups. That neither the idea of Tits nor the phenomenon of repetition has already been exploited I take as an indication of how little we know about Kac–Moody structures.

Keywords: Structure constants • Kac–Moody • Lie algebras

Mathematics Subject Classification: 17B05, 17B22, 17B45, 17B67

1 Introduction

This paper is based on a lecture I gave at a conference in San Diego in honor of the achievements of Nolan Wallach.

Suppose Δ to be a finite set. In this paper a *Cartan matrix* indexed by Δ will be an arbitrary integral matrix $C = (c_{\alpha,\beta})$ ($\alpha, \beta \in \Delta$) satisfying these conditions:

B. Casselman
Mathematics Department, University of British Columbia, Vancouver, BC, Canada
e-mail: cass@math.ubc.ca

- (C1) $c_{\alpha,\alpha} = 2$;
 (C2) for $\alpha \neq \beta$, $c_{\alpha,\beta} \leq 0$;
 (C3) $c_{\alpha,\beta} = 0$ if and only if $c_{\beta,\alpha} = 0$.

More commonly such a matrix is called a *generalized Cartan matrix*, whereas a Cartan matrix is taken to be one with the additional requirement that

- (CD) there exist a diagonal matrix D such that CD is positive definite.

But times change, and “generalized” is now closer to “usual”. In this paper, if condition (CD) is satisfied I will call C positive definite. For purely technical reasons I will assume further that

- (C4) C is irreducible;
 (C5) C is invertible.

The first condition means that Δ cannot be expressed as $\Delta_1 \cup \Delta_2$ with $c_{\alpha,\beta} = 0$ for $\alpha \in \Delta_1$, $\beta \in \Delta_2$. This condition is no significant restriction on results, since one can work with the summands of the root systems corresponding to the Δ_i . The second condition excludes the affine Cartan matrices. This will restrict results, but the missing cases can be easily dealt with separately.

Let

$$L = \mathbb{Z}^\Delta, \quad L^\vee = \text{Hom}(L, \mathbb{Z}).$$

Thus Δ embeds into L as a basis. For every subset Θ of Δ let L_Θ be the span of Θ in L , so that in particular $L_\Delta = L$. Since C is integral and invertible, there exists a unique maximal linearly independent subset Δ^\vee of L^\vee and a bijection $\alpha \mapsto \alpha^\vee$ with

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

Let $L_{\Delta^\vee}^\vee$ be the submodule of L^\vee spanned by Δ^\vee . The linear transformation

$$s_\alpha: \lambda \mapsto \lambda - \langle \lambda, \alpha^\vee \rangle \alpha$$

is a (possibly skew) reflection, taking L to itself. Its contragredient s_{α^\vee} takes L^\vee and $L_{\Delta^\vee}^\vee$ to themselves. These reflections generate the *Weyl group* W associated to the Cartan matrix. The W -orbit of Δ is the set of *real roots* $\Sigma_{\mathbb{R}}$. It is the disjoint union of two halves, positive and negative. The positive (resp. negative) roots $\Sigma_{\mathbb{R}}^+$ (resp. $\Sigma_{\mathbb{R}}^-$) are integral linear combinations of elements in Δ with non-negative (resp. non-positive) coefficients.

Let \mathcal{D} be the region of all v in $L \otimes \mathbb{R}$ such that $\langle v, \alpha^\vee \rangle \geq 0$ for all $\alpha \in \Delta$, and define the *closed Tits cone* $\bar{\mathcal{C}}$ to be the closure of $W(\mathcal{D})$. It is convex, and has a non-empty interior, on which the group W acts discretely with fundamental domain \mathcal{D} . I define $\Sigma_{\mathbb{I}}$, the set of *imaginary roots*, to be the union of

$$\Sigma_{\mathbb{I}}^+ = (L_\Delta \cap \bar{\mathcal{C}}) - \{0\}$$




Figure 1 Real and imaginary roots for a system of rank 2

and its opposite $\Sigma_{\mathbb{I}}^-$. It, too, is W -stable, as is each half $\Sigma_{\mathbb{I}}^\pm$ and the set Σ of all roots, the union of real and imaginary ones. (For all this, refer to Proposition 5.2, Theorem 5.4, Lemma 5.8, and Exercise 5.12 in Chapter 5 of [Kac85]. Figure 1 shows how things look in a simple case.) The terminology is perhaps motivated by the fact that many root systems possess a natural quadratic form with respect to which real roots have real lengths and imaginary ones have imaginary ones.

Let

$$\mathfrak{h} = L^\vee \otimes \mathbb{C},$$

which may be identified with an Abelian Lie algebra, and let

$$\mathfrak{h}^\vee = L \otimes \mathbb{C}$$

be its complex linear dual. For each α in Δ let h_α be the image of α^\vee in \mathfrak{h} . In other words, α^\vee may be identified with a linear injection of \mathbb{C} into \mathfrak{h} taking x to xh_α . Let \mathfrak{g} be the Kac–Moody Lie algebra associated to these data. One is given in the construction of \mathfrak{g} an embedding of \mathfrak{h} into it. The adjoint action of \mathfrak{h} on \mathfrak{g} is semi-simple and locally finite, breaking it up into the sum of \mathfrak{h} and a number of root spaces \mathfrak{g}_λ , with $\lambda: \mathfrak{h} \rightarrow \mathbb{C}$ lying inside the image of L in \mathfrak{h}^\vee . The fundamental fact

about these data is that $\mathfrak{g}_\lambda \neq 0$ if and only if λ is a root in the sense spoken of earlier. This is complemented by the assertion that for each α in Δ , \mathfrak{g}_α has dimension one. It is also true that \mathfrak{g}_λ has dimension one for every real root λ .

The validity of these assertions depends on assumption (C5). Without it, there still exists a Kac–Moody algebra defined by C , but its description is a bit more subtle, as explained in the opening chapter of [Kac85], and elaborated on in Chapter 6. The corresponding Lie algebras turn out to be extensions of $\mathfrak{g} \otimes \mathbb{C}[t^{\pm 1}]$, with \mathfrak{g} finite-dimensional, and computations in it reduce easily to computations in \mathfrak{g} itself. This is why (C5) is not very restrictive.

There are two basic problems this paper will deal with:

Problem 1. To specify a good choice of basis elements e_λ of \mathfrak{g}_λ for all real roots λ ;

Problem 2. To find, for every pair of real roots λ, μ such that $\lambda + \mu$ is also a real root, the structure constant $N_{\lambda, \mu}$ such that $[e_\lambda, e_\mu] = N_{\lambda, \mu} e_{\lambda + \mu}$.

The method I will use to solve these problems originates in [Tits66a], which sketches what happens when \mathfrak{g} has finite dimension. The extension to Kac–Moody algebras is not quite trivial, but neither is it particularly tricky. I hope that it will be useful in exploring what happens for bracket computations involving imaginary roots, which is one of the great mysteries of Kac–Moody algebras, although I say nothing about this problem here.

When \mathfrak{g} is finite-dimensional, there are already in the literature two practical approaches to constructing the e_λ and computing the $N_{\lambda, \mu}$. One can be found in [Car72] and is explained in more detail in [CMT04]. The construction uses induction on roots, going from λ to $\lambda + \alpha$ ($\alpha \in \Delta$) according to a certain rule. Calculation of the structure constants amounts to transliterations of the Jacobi identity. The other method, introduced in [FK80], exhibits an interesting extension of groups to interpret the constants as related to cohomology. This technique was applied originally only to simply laced systems, but it has since been extended to the rest of the finite-dimensional \mathfrak{g} in [Ryl00] by ‘folding’. As far as I can tell these methods cease to be valid for arbitrary Kac–Moody algebras. In any event, in the cases for which they do work they do not seem to be a great deal faster than a program based on the method to be explained here. To the extent to which they are faster, they are ‘hard-wired’, incorporating for each system special short-cuts that do not apply in general.

Acknowledgments. I thank the organizers of the conference in honor of Nolan Wallach for inviting me to contribute to it. I also wish to thank my colleague Julia Gordon, who—rather casually—started me off on this project. I am extremely happy to have had this opportunity to try to understand and elaborate on Jacques Tits’ work on structure constants. I first saw that volume of the *Publications de l’IHES* for sale in Schoenhof’s when I was a very ignorant graduate student, and the memory of my puzzled thought, “What exactly is a structure constant?” has remained with me ever since. And finally I wish to thank the referee, who complained much about lacunae in an earlier draft but nonetheless continued to read and criticize carefully.

2 Chevalley bases

In this section I will start looking at Problem 1, that of constructing a good basis of \mathfrak{g} . The beginning is straightforward—for each α in Δ choose an arbitrary $e_\alpha \neq 0$ in \mathfrak{g}_α . It doesn't make any difference what choice you make, because all choices will be conjugate with respect to an automorphism of the Lie algebra \mathfrak{g} . For each α in Δ , $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$ will be the one-dimensional subspace spanned by h_α . Following [Tits66a], then choose $e_{-\alpha}$ so

$$[e_\alpha, e_{-\alpha}] = -h_\alpha .$$

The usual convention imposes a plus sign on the right hand side of this equation, but Tits' change of sign is extraordinarily convenient, in fact obligatory if the symmetries exploited later on in this paper are to remain comprehensible. In my opinion, Tits' choice should have become the conventional one long ago.

For example, if $\mathfrak{g} = \mathfrak{sl}_2$, we get

$$e_+ = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad e_- = \begin{bmatrix} 0 & 0 \\ -1 & 0 \end{bmatrix}, \quad h = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} .$$

The data of \mathfrak{h} together with the e_α make up what I will call a *frame* for \mathfrak{g} . (It is called in French an *épinglage*, frequently translated awkwardly into English as *pinning*. The reference in French is to the way butterflies are mounted. To those of us who have been parents of young children, the proposed English term has other, less fortunate, connotations.)

From now on, assume a frame to have been fixed. With this assumption, we are given also for each α in Δ a well-defined embedding

$$\varphi_\alpha: \mathfrak{sl}_2 \longrightarrow \mathfrak{g} \quad | \quad e_+ \longmapsto e_\alpha, \quad h \longmapsto h_\alpha, \quad e_- \longmapsto e_{-\alpha} .$$

Let $\mathfrak{sl}_2^{[\alpha]}$ be its image.

The map taking each $e_{\pm\alpha}$ to $e_{\mp\alpha}$, $h \mapsto -h$ for h in \mathfrak{h} extends to an involutory automorphism θ of all of \mathfrak{g} called the *canonical involution*.

For x in \mathfrak{g}_λ and y in \mathfrak{g}_μ , the bracket $[x, y]$ lies in $\mathfrak{g}_{\lambda+\mu}$ (which may be 0). If $\mu = -\lambda$, this means it will lie in \mathfrak{h} . For each real root λ the subspace $[\mathfrak{g}_\lambda, \mathfrak{g}_{-\lambda}]$ has dimension one in \mathfrak{h} , and is not contained in the kernel of λ . Hence there exists a unique h_λ in it with $\langle \lambda, h_\lambda \rangle = 2$.

Lemma 2.1. *For each real root λ there exists up to sign a unique e_λ in \mathfrak{g}_λ such that*

$$[e_\lambda, \theta(e_\lambda)] = -h_\lambda .$$

Proof. Given $e \neq 0$ in \mathfrak{g}_λ , we know that $\theta(e)$ lies in $\mathfrak{g}_{-\lambda}$ and hence

$$[e, \theta(e)] = ah_\lambda$$

for some constant $a \neq 0$. But then

$$[ce, \theta(ce)] = c^2ah_\lambda$$

so we choose ce with $c^2a = -1$. □

We may thus assemble a basis of \mathfrak{g} by choosing for each root λ one of the two choices this gives us. This is called a *Chevalley basis*. Part of the solution to Problem 1 above is to choose a Chevalley basis. It is unique only up to signs, and there is apparently no canonical choice.

Suppose we are given a Chevalley basis. Suppose further we are given real roots λ, μ such that $\lambda + \mu$ is also a real root. We have

$$[e_\lambda, e_\mu] = N_{\lambda, \mu}e_{\lambda+\mu}$$

in which the absolute value $|N_{\lambda, \mu}|$ is now uniquely determined. It is known never to be 0. We should not be too surprised to learn that it has a relatively simple expression. I will next explain the formula for it.

If λ and μ are roots with λ real, the λ -string through μ is the intersection of $\mu + \mathbb{Z}\lambda$ with Σ . It is always finite. Define $p_{\lambda, \mu}$ to be the maximum p such that $\mu - p\lambda$ is in the string, and $q_{\lambda, \mu}$ to be the maximum q such that $\mu + q\lambda$ lies in it. The string is the full segment $[\mu - p\lambda, \mu + q\lambda]$. The following identities are easy to verify:

$$\begin{aligned} p_{\mu, \lambda} &= p_{-\mu, -\lambda} \\ p_{\mu+\lambda, \lambda} &= p_{\mu, \lambda} + 1 \\ p_{\mu, -\lambda} &= n - p_{\mu, \lambda} \text{ if } n \text{ is the length of the string.} \end{aligned}$$

The following was first proved by Chevalley for finite-dimensional Lie algebras, probably by Tits more generally. See also [Mor87]. Following [Tits66a], I shall in effect prove it later on.

Theorem 2.2 (Chevalley's Theorem). *If λ, μ , and $\lambda + \mu$ are all real roots, then*

$$|N_{\lambda, \mu}| = p_{\lambda, \mu} + 1.$$

The constants $p_{\lambda, \mu}$ are simple to evaluate, so the whole problem of computing structure constants comes down to computing a sign.

This result makes it possible to assign a \mathbb{Z} -structure to \mathfrak{g} , and was classically the basis for Chevalley's construction of split group schemes defined over \mathbb{Z} . Incidentally, the usual proof of Chevalley's theorem is a case-by-case examination. This is somewhat unsatisfactory, and moreover it will not be possible when working with arbitrary Kac–Moody algebras. One is forced to come up with a proof that is illuminating even in the finite-dimensional cases.

3 The root graph

The algorithm we shall eventually see will require that we be able to list as many roots as we want. In fact, in principle we can construct even an infinite number all at the same time, as I will explain later.

Every real root is $w\alpha$ for some w in W , α in Δ . The natural way to construct real roots is therefore to start with all the α in Δ and apply the elements of W to them. How does this go? Keep in mind that the group W is generated by the elementary reflections s_α . We maintain two lists of positive roots, one the current list to be processed—the *waiting list*—and the other that of roots that have been processed—*serviced customers*. Processing a root means (a) removing it from the waiting list and (b) applying all elementary reflections to it to see whether we get new roots or roots we have already seen. When we see a new one, we put it on the waiting list, if the new root has height below some bound we have initially set. (I recall that the height of $\lambda = \sum \lambda_\alpha \alpha$ is $\sum |\lambda_\alpha|$.) Of course one may as well restrict oneself to the task of finding only positive roots.

It is not hard to use a look-up table to follow this method. It is guaranteed to give us eventually as many roots as we need, but it is not at all clear how long it will take. The problem that comes to mind is that when we apply elementary reflections we might expect *a priori* to go down in height to get a new root, and then down in height again, and so on. This is reminiscent of solving the word problem in group theory. In this case it ought not to be necessary—we would like to know that we can proceed by adding to the waiting list only reflections of greater height. This is easy to carry out, since $s_\alpha \lambda$ has height more than λ if and only if $\langle \lambda, \alpha^\vee \rangle < 0$. But seeing why this procedure will work is not quite trivial.

Define the *depth* of a positive root λ to be $\ell(w)$ for the w of least length such that $w^{-1}\lambda < 0$. Equivalently this is $\ell(w) + 1$ where w is of least length such that $\lambda = w\alpha$ with $\alpha \in \Delta$. For example, the depth of every α in Δ is 1. The fundamental fact about depth is Proposition 2.3 of [BH93], which is simple to state but not quite so simple to prove:

Proposition 3.1 (Brink and Howlett). *The depth of $s_\alpha \lambda$ is greater than that of λ if and only if $\langle \lambda, \alpha^\vee \rangle < 0$.*

This is also part of Lemma 3.3 in [Cas06].

One consequence is that the height of a root is greater than or equal to its depth, so if we have found all roots of depth $\leq n$ we have also found all roots of height $\leq n$. Another consequence is one we'll need in a later section:

Corollary 3.2. *If λ is a positive real root not in Δ , then there exists α in Δ with $\langle \lambda, \alpha^\vee \rangle > 0$.*

Here is another natural consequence:

Corollary 3.3. *If $\Theta \subseteq \Delta$, then the real roots in L_Θ are in the W_Θ -orbit of Θ .*

Here W_Θ is the subgroup of W generated by the s_α with α in Θ .

Proof. Arguing by induction on depth, it has to be shown that if λ is a positive real root in the linear span of Θ , there exists $\alpha \in \Theta$ such that $\langle \lambda, \alpha^\vee \rangle > 0$. But since λ is positive and in L_Θ , $\langle \lambda, \beta^\vee \rangle \leq 0$ for all β not in Θ , so that the claim follows from the previous Corollary. \square

In finding positive roots, we are in effect building what I call the *root graph*, from the bottom up. It is an important structure. With one exception the nodes of this graph are the positive roots. The exception is that for technical convenience I add a ‘dummy’ node at the bottom (Figures 2 and 3).

What about edges? They are oriented. There is an edge labeled s_α from the dummy node to α , for each α in Δ . There is an edge from a root λ to another μ if $\mu = s_\alpha \lambda$ and the height of μ is greater than the height of λ . It is labeled by s_α . Since

$$s_\alpha \lambda = \lambda - \langle \lambda, \alpha^\vee \rangle \alpha$$

this will happen if and only if $\langle \lambda, \alpha^\vee \rangle < 0$.




Figure 2 The roots and root graph of G_2




Figure 3 A different version of the root graph of G_2




Figure 4 The root graphs for A_3 and C_3




Figure 5 The bottom of an infinite root graph

In drawing the root graph, the loops are redundant, since the total number of edges coming in or out is a constant. (This is the reason for adding the ‘dummy’.) Furthermore, in the diagrams that follow the arrows can be replaced by simple lines, because the orientation is always upward (Figs. 4 and 5).




Figure 6 The ISL root graph for A_3

Each root can be reached by one or more paths in the root graph, starting at the dummy node. But if we are to construct roots by traversing paths, we want to specify a unique path to every node. I choose the *ISL* (Inverse Short Lex) path, which is defined recursively (Figure 6). Assign an order to Δ . A path

$$\gamma = \lambda_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n} \lambda_n$$

is ISL if α_n is the least label on the edges leading up to λ_n , and similarly for all its initial segments.

4 Rank two systems

Chevalley’s formula for $|N_{\lambda,\mu}|$ indicates that the geometry of root configurations will likely be important in this business. In this section and the next I prove a few crucial properties.

I begin by recalling the original way of thinking about bases of roots. Suppose V to any real vector space, given a coordinate system (x_i) . Impose on V the associated *lexicographic order*: $x < y$ if $x_i = y_i$ for $i < m$ but $x_m < y_m$. This is a linear order, invariant under translation, and [Sat51] points out that any translation-invariant linear order that is continuous in some sense has to be one of these. In any free \mathbb{Z} -module contained in V there exists a least non-zero positive vector.

Suppose Δ to be a basis of a set Σ of roots, and take V to be $L_\Delta \otimes \mathbb{R}$. If we assign an order to Δ , we get a corresponding lexicographic order on V . *The positive roots in Σ are those > 0 with respect to the lexicographic order.* This should motivate the following discussion.

Lemma 4.1. *Suppose we are given an ordered set of real and linearly independent roots λ_i , and let Λ_k be the subset of the first k . There exists a subset of roots α_i in Δ and w in W such that $w\lambda_1 = \alpha_1$, and each $w\Lambda_k$ is contained in the non-negative integral span of the α_i for $i \leq k$.*

Proof. Extending the set of λ_i if necessary, we may assume they form a basis of V . Assign V the associated lexicographic order. Let V_k be the real span of Λ_k .

Define β_1 to be λ_1 , and define β_k to be the least positive root in V_k that is not in V_{k-1} (look at the linear order induced on the quotient to see that this exists). Let B be the set of all β_i , B_k the subset of the first k . The real span of B_k is V_k .

I claim that every real root λ in the real span of B_k positive with respect to this linear order (including the λ_i themselves) is in the non-negative integral span of B_k . For $k = 1$, there is nothing to prove. Proceed by induction on the smallest k such that λ is in the real span of Λ_k . Applying induction, we may assume λ is in $V_k - V_{k-1}$. If $\lambda = \beta_k$, no problem. Otherwise $\lambda > \beta_k$. According to Corollary 3.2 of § 3 there exists $i \leq k$ such that $\langle \lambda, \beta_i^\vee \rangle > 0$. Then $\lambda - \beta_i$ will again be a root, of smaller height. Repeat as required.

Now apply Theorem 2 of §5.9 in [MP95] to see that there exists w in W with $wB = \Delta$. □

Proposition 4.2. *If λ, μ are real roots, there exists w in W , α and β in Δ such that $w\lambda = \alpha$, $w\mu$ is a non-negative integral combination of α and β . The real root $w\mu$ is in $W_{\alpha,\beta}\{\alpha, \beta\}$.*

Proof. The last assertion follows from Corollary 3.3. □

The map from Δ to Δ^\vee has a natural extension to a map $\lambda \mapsto \lambda^\vee$ from $\Sigma_{\mathbb{R}}$ to $\Sigma_{\mathbb{R}}^\vee \subset L^\vee$. It is characterized by the formula $(w\lambda)^\vee = w\lambda^\vee$ for all real roots λ . If $\lambda = w(\alpha)$, then the reflection corresponding to λ is $s_\lambda = ws_\alpha w^{-1}$:

$$s_\lambda u = u - \langle u, \lambda^\vee \rangle \lambda .$$

For any pair of distinct real roots λ, μ let $L_{\lambda,\mu}$ be their integral span.

Proposition 4.3. *If λ and μ are any two distinct real roots with $\langle \lambda, \mu^\vee \rangle \neq 0$, there exists on $L_{\lambda,\mu}$ an inner product with respect to which s_λ and s_μ are orthogonal reflections, and $\lambda \cdot \lambda > 0$. It is unique up to a positive scalar, and $\kappa \cdot \kappa > 0$ for all real roots κ in $L_{\lambda,\mu}$.*

Proof. The reflection s_κ is orthogonal if and only if

$$2 \left(\frac{u \cdot \kappa}{\kappa \cdot \kappa} \right) = \langle u, \kappa^\vee \rangle$$

for all u in $L_{\lambda,\mu}$. This requires that

$$\begin{aligned} \langle \lambda, \mu^\vee \rangle &= 2 \left(\frac{\lambda \cdot \mu}{\mu \cdot \mu} \right) \\ \langle \mu, \lambda^\vee \rangle &= 2 \left(\frac{\lambda \cdot \mu}{\lambda \cdot \lambda} \right) . \end{aligned}$$

These equations tell us that $\lambda \cdot \mu$ and $\mu \cdot \mu$ are both determined by $\lambda \cdot \lambda$, so certainly there is a unique inner product defined uniquely on $L_{\lambda,\mu}$ by $\lambda \cdot \lambda$ and the requirement that reflections be orthogonal. Why are the norms of real roots then all positive?

According to Proposition 4.2, we may assume that $\lambda = \alpha$, $\mu = p\alpha + q\beta$ with α, β in Δ . In these circumstances $\langle \alpha, \beta^\vee \rangle \neq 0$, so there exists on $L_{\alpha, \beta}$ an essentially unique inner product with $\alpha \cdot \alpha > 0$. By the definition of a Cartan matrix, we also have $\beta \cdot \beta > 0$. The given inner product on $L_{\lambda, \mu}$ must be some scalar multiple of this one. But both λ and μ are Weyl transforms of α or β , so both also have positive norms. \square

Corollary 4.4. *Given two real roots λ, μ we have*

$$\langle \lambda, \mu^\vee \rangle \langle \mu, \lambda^\vee \rangle \geq 0$$

and one factor is 0 if and only if both are.

Proof. Since

$$\langle \lambda, \mu^\vee \rangle \langle \mu, \lambda^\vee \rangle = 4 \frac{(\lambda \cdot \mu)^2}{\|\lambda\|^2 \|\mu\|^2}. \quad \square$$

Since the transpose of a Cartan matrix is also a Cartan matrix, the previous result implies that there exists an essentially unique metric on Σ^\vee as well.

Corollary 4.5. *In an irreducible root system of rank two, the product $\|\lambda\| \|\lambda^\vee\|$ is constant.*

Proof. It is immediate from the defining formulas that

$$\|\alpha\| \|\alpha^\vee\| = \|\beta\| \|\beta^\vee\|. \quad (*)$$

Any other real root λ is of the form $w\alpha$ or $w\beta$, from which the claim follows. \square

Corollary 4.6. *The map*

$$v \longmapsto \frac{v^\vee}{\|v^\vee\|^2}$$

is linear on $\Sigma_{\mathbb{R}} \cap L_{\lambda, \mu}$.

5 Tits triples

In this section, suppose λ, μ , and ν to be what I call a **Tits triple**—a trio of real roots with $\lambda + \mu + \nu = 0$. In these circumstances, either $\langle \mu, \lambda^\vee \rangle$ or $\langle \nu, \lambda^\vee \rangle$ does not vanish, so by Corollary 4.3 there exists a well-defined inner product on $L_{\lambda, \mu} = L_{\lambda, \nu}$, with respect to which all real roots have positive norm.

In particular, the λ -string through μ contains at least the two real roots μ and $-\nu$. It is stable under s_λ , which reflects in the point where $\langle \mu, \lambda^\vee \rangle = 0$. By assumption, both μ and $-\nu$ are in this string.

Lemma 5.1. *Let $n = \langle \mu, \lambda^\vee \rangle$.*

- (a) $n < -1$ if and only if $\|\mu\| > \|\nu\|$;
- (b) $n = -1$ if and only if $\|\mu\| = \|\nu\|$;
- (c) $n > -1$ if and only if $\|\mu\| < \|\nu\|$.

Proof. If $n = -1$, then $s_\lambda \mu = -\nu$. If $n \geq 0$, then evaluate $(\mu + \lambda) \cdot (\mu + \lambda)$ to see that $\|\nu\|^2 > \|\mu\|^2$. In the remaining case, with $n \leq -2$, consider instead $s_\lambda \lambda = -\lambda$, $s_\lambda \nu$, $s_\lambda \mu$. □

There are a limited number of configurations of the λ -string through μ . If κ is a positive (resp. negative) imaginary root in the string, then $s_\lambda \kappa$ is also one, since the positive (resp. negative) imaginary roots are stable under W . Since the positive (resp. negative) imaginary roots are the intersection of a convex set with the roots, all intervening roots must also be positive (resp. negative) and imaginary. Therefore any real roots in the string must be at its ends. Since s_λ reflects in the hyperplane $\langle \mu, \lambda^\vee \rangle = 0$ and the real roots in the string can have at most two lengths, the previous result implies:

Lemma 5.2. *There do not exist three real roots κ in a λ -string with $\langle \kappa, \lambda^\vee \rangle \leq 0$.*

In other words, the following figures, with real roots dark, show all possibilities for strings containing a real root. The lengths in the string decrease until at most the half-way point, then increase (Figure 7).

All these are in fact possible, as you can see by perusing classical root diagrams or Figure 1.

Lemma 5.3. *The following are equivalent:*

- (a) $s_\lambda \mu = -\nu$;
- (b) $\langle \mu, \lambda^\vee \rangle = -1$;
- (c) $\|\lambda\| \geq \|\mu\|, \|\nu\|$.




Figure 7 Possible root strings containing a real root




Figure 8 An impossible real root configuration

The point of this result is that whenever $\lambda + \mu + \nu = 0$ we can cycle to obtain $s_\lambda \mu = -\nu$, since we can certainly cycle to get the third condition by taking λ of maximum length.

Proof. The equivalence of (a) and (b) is immediate.

Assume (b). Then $s_\lambda \mu = -\nu$, so $\|\mu\| = \|\nu\|$. But also $\langle \lambda, \mu^\vee \rangle \leq -1$, so by Lemma 5.1 we have $\|\lambda\| \geq \|\mu\|$. Thus (b) implies (c).

Assume (c). First of all, Lemma 5.1 implies that $\langle \lambda, \mu^\vee \rangle < 0$, for if not then $\|\nu\|^2 = \|\lambda + \mu\|^2 > \|\lambda\|^2$, a contradiction. This implies that $n = \langle \mu, \lambda^\vee \rangle < 0$ as well.

If $n = -1$ then $s_\lambda \mu = -\nu$ and $\|\mu\| = \|\nu\|$. If $n < -1$. But then by Lemma 5.1 we have $\|\mu\| \geq \|\mu + \lambda\| = \|\nu\|$. By symmetry, we also have $\|\nu\| \geq \|\mu\|$. But then $\|\mu\| = \|\nu\|$ and by Lemma 5.1 $s_\lambda \mu = -\nu$. □

One consequence is that we cannot have a trio of real roots with $\lambda + \mu + \nu = 0$, $\|\lambda\|, \|\mu\| > \|\nu\|$. The diagram of Figure 8 is impossible.

I shall now examine what happens when $\|\lambda\| \geq \|\mu\| = \|\nu\|$. There are two different cases.

Suppose first that $\|\lambda\| = \|\mu\| = \|\nu\|$. Then by Lemma 5.1

$$\begin{aligned}
 s_\lambda: \lambda &\mapsto -\lambda \\
 &\mu \mapsto -\nu \\
 &\nu \mapsto -\mu \\
 s_\mu: \lambda &\mapsto -\nu \\
 &\mu \mapsto -\mu \\
 &\nu \mapsto -\lambda
 \end{aligned}$$

so $s_\lambda s_\mu$ rotates (λ, μ, ν) to (μ, ν, λ) . Following this by s_μ maps

$$\begin{aligned}
 \lambda &\mapsto -\mu \\
 \mu &\mapsto -\lambda \\
 \nu &\mapsto -\nu.
 \end{aligned}$$

Thus

$$p_{\lambda,\mu} = p_{\mu,\nu} \text{ etc.}, \quad p_{\lambda,\mu} = p_{\mu,-\lambda} = p_{\mu,\lambda}.$$

Suppose next that $\|\lambda\| > \|\mu\| = \|\nu\|$. We may set these last to 1, $\|\lambda\|^2 = n > 1$. By Lemma 5.1 we must have $\langle \mu, \lambda^\vee \rangle = -1$. Since

$$\langle \lambda, \mu^\vee \rangle \|\mu\|^2 = \langle \mu, \lambda^\vee \rangle \|\lambda\|^2$$

we must also have $\langle \lambda, \mu^\vee \rangle = -n < -1$.

But then λ and $-\nu = \lambda + \mu$ must also lie in the initial half of the μ -string through λ , and therefore by Lemma 5.2 λ must be its beginning. Similarly, λ is at the beginning of its ν -string. Therefore

$$p_{\lambda,\mu} = p_{\lambda,\nu} = 0.$$

If ν were not at the beginning of its λ -string, there would exist a root $\nu - \lambda$ of squared-length

$$\|\nu - \lambda\|^2 = \|\nu\|^2 - 2\nu\dot{\lambda} + \|\lambda\|^2$$

greater than that of λ . Again by Lemma 5.2, this cannot happen. So ν is at the beginning of its λ -string. The same holds for μ . So

$$p_{\nu,\lambda} = p_{\mu,\lambda} = 0.$$

Since λ is at the beginning of its ν -string, $s_\nu\lambda = \lambda + n\nu$ is the end of that string, which is of length n . But $\lambda + \nu = -\mu$, so

$$p_{\mu,\nu} = p_{-\mu,-\nu} = n - 1$$

and also

$$p_{\nu,\mu} = p_{-\nu,-\mu} = n - 1.$$

These computations now imply:

Lemma 5.4 (Geometric Lemma). *We have*

$$\frac{p_{\lambda,\mu} + 1}{\|\nu\|^2} = \frac{p_{\mu,\nu} + 1}{\|\lambda\|^2} = \frac{p_{\nu,\lambda} + 1}{\|\mu\|^2}.$$

Lemma 5.5. *We have $p_{\mu,\lambda} = p_{\lambda,\mu}$.*

Proposition 5.6. *We have*

$$(p_{\lambda,\mu} + 1)\nu^\vee + (p_{\mu,\nu} + 1)\lambda^\vee + (p_{\nu,\lambda} + 1)\mu^\vee = 0.$$

Proof. Identify λ with $2\lambda^\vee/\|\lambda^\vee\|^2$, etc. Since $\lambda + \mu + \nu = 0$ we also have by Corollary 4.6

$$\frac{\lambda^\vee}{\|\lambda^\vee\|^2} + \frac{\mu^\vee}{\|\mu^\vee\|^2} + \frac{\nu^\vee}{\|\nu^\vee\|^2} = 0$$

But the product $\|\kappa\| \|\kappa^\vee\|$ does not depend on κ , so

$$\|\lambda\|^2 \lambda^\vee + \|\mu\|^2 \mu^\vee + \|\nu\|^2 \nu^\vee = 0.$$

Lemma 5.4 now implies that

$$(p_{\lambda,\mu} + 1)\nu^\vee + (p_{\mu,\nu} + 1)\lambda^\vee + (p_{\nu,\lambda} + 1)\mu^\vee = 0. \quad \square$$

6 Representations of $\mathrm{SL}(2)$

Representations of SL_2 play an important role in both Carter's and Tits' approaches to structure constants. In this section I recall briefly what is needed. Of course this is well-known material, but perhaps not in quite the form I wish to refer to.

Let

$$u = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad v = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

be the standard basis of \mathbb{C}^2 , on which $\mathrm{SL}_2(\mathbb{C})$ and \mathfrak{sl}_2 act. They also act on the symmetric space $S^n(\mathbb{C})$, with basis $u^k v^{n-k}$ for $0 \leq k \leq n$. Let π_n be this representation, which is of dimension $n + 1$.

$$\begin{aligned} \pi_n \left(\begin{bmatrix} e^x & 0 \\ 0 & e^{-x} \end{bmatrix} \right) : u^k v^{n-k} &\longmapsto e^{(2k-n)x} u^k v^{n-k} \\ \pi_n \left(\begin{bmatrix} 1 & x \\ 0 & 1 \end{bmatrix} \right) : u^k v^{n-k} &\longmapsto u^k (v + xu)^{n-k} \\ \pi_n \left(\begin{bmatrix} 1 & 0 \\ -x & 1 \end{bmatrix} \right) : u^k v^{n-k} &\longmapsto (u - xv)^k v^{n-k}. \end{aligned}$$

Now to translate these formulas into those for the action of \mathfrak{g} . Let $w_k = u^k v^{n-k}$.

$$\begin{aligned} \pi_n(h) : w_k &\longmapsto (2k - n) w_k \\ \pi_n(e_+) : w_k &\longmapsto (n - k) w_{k+1} \\ \pi_n(e_-) : w_k &\longmapsto -k w_{k-1}. \end{aligned}$$




Figure 9 The representation π_3

For many reasons, it is a good idea to use a different basis of the representation. Define *divided powers*

$$u^{[k]} = \frac{u^k}{k!},$$

and set $w_{[k]} = u^{[k]}v^{[n-k]}$ (Figure 9). Then

$$\begin{aligned} \pi_n(h) &: w_{[k]} \mapsto (2k - n) w_{[k]} \\ \pi_n(e_+) &: w_{[k]} \mapsto (k + 1)w_{[k+1]} \\ \pi_n(e_-) &: w_{[k]} \mapsto -(n - k + 1)w_{[k-1]}. \end{aligned}$$

If

$$\sigma = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix},$$

then we also have

$$\pi_n(\sigma): w_{[k]} \mapsto (-1)^k w_{[n-k]}.$$

The π_n exhaust the irreducible finite-dimensional representations of both $SL_2(\mathbb{C})$ and \mathfrak{sl}_2 , and every finite-dimensional representation of either is a direct sum of them.

Let $d = n + 1$, the dimension of π_n , assumed even. The weights of this with respect to h are

$$-d, \dots, -1, 1, \dots, d,$$

The formulas above imply immediately:

Proposition 6.1. *Suppose v to be a vector of weight -1 in this representation. Then $\pi(e_+)v$ and $\pi(\sigma)v$ are both of weight 1, and*

$$\pi(e_+)v = -(-1)^{d/2} \left(\frac{d}{2}\right) \pi(\sigma)v.$$

The sum of weight spaces \mathfrak{g}_μ for μ in a given λ -string is a representation of \mathfrak{sl}_2 . The formulas laid out in this section relate closely to the numbers $p_{\lambda,\mu}$ and $q_{\lambda,\mu}$. For one thing, as the picture above suggests and is easy to verify, they tell us that we can choose basis elements e_μ for each μ in the chain so that

$$\begin{aligned} [e_\lambda, e_\mu] &= (p_{\lambda,\mu} + 1)e_{\mu+\lambda} \\ [e_{-\lambda}, e_\mu] &= (-1)^{p_{\lambda,\mu}} (q_{\lambda,\mu} + 1)e_{\mu-\lambda}. \end{aligned}$$

The choice of basis for one chain unfortunately affects other chains as well, so this observation doesn't make the problem of structure constants trivial. But it is our first hint of a connection between structure constants and chains.

7 The extended Weyl group ...

Tits' beautiful idea is to reduce the computation of structure constants to computation in a certain extension of the Weyl group, the **extended** Weyl group, fitting into a short exact sequence

$$1 \rightarrow \{\pm 1\}^\Delta \rightarrow W_{\text{ext}} \rightarrow W \rightarrow 1.$$

If \mathfrak{g} has finite dimension, the group W_{ext} can be described succinctly as the normalizer of a maximal torus in the integral form of the simply connected split group with Lie algebra \mathfrak{g} . For arbitrary \mathfrak{g} it is a subgroup of a group G constructed in [KP85], a simply connected analogue for general Kac–Moody algebras.

I cannot resist remarking here about the literature in this field. That on Kac–Moody algebras is adequate. The primary reference is still [Kac85], but it needs updating and its exposition is dense. A useful supplement is [MP95]. For that matter, it is puzzling that there are a number of extremely basic questions about the algebras that have not yet been answered, such as whether or not they can be defined by Serre relations. This makes it very difficult to do explicit calculations with them. But the literature on the associated groups is far less satisfactory. The original paper [MT72] is very readable, but doesn't tell enough for practical purposes. There are a number of valuable expositions by Jacques Tits, such as [Tits87] and [Tits88], but many of these are difficult to obtain, and are in any event inconclusive. The paper [KP85] is exceptionally clear, but for full proofs one must refer back to earlier papers by the same authors that are not so clear. There is the book [Kum02], but it doesn't contain everything one wants. There are a number of expositions by Olivier Mathieu, but I have the impression that he stopped writing on this subject before he was through.

The good news is that we do not need to know a great deal about the Kac–Peterson group G .

(G1) There exists for each α in Δ a unique embedding

$$\varphi_\alpha: \text{SL}_2(\mathbb{C}) \hookrightarrow G$$

compatible with the map φ_α of Lie algebras. *The images $\text{SL}_2^{[\alpha]}(\mathbb{C})$ generate G .*

(G2) If π is any representation of \mathfrak{g} whose restriction to $\mathfrak{sl}_2^{[\alpha]}$ is a direct sum of irreducible representations of finite dimension, it lifts to a representation I will call π_α of $\text{SL}_2(\mathbb{C})$. *There exists a unique representation of G that I will also call π such that $\pi \circ \varphi_\alpha = \pi_\alpha$.* This is true in particular if $\pi = \text{ad}$.

(G3) Let ι be the embedding of \mathbb{C}^\times into $\mathrm{SL}_2(\mathbb{C})$:

$$\iota: x \mapsto \begin{bmatrix} x & 0 \\ 0 & 1/x \end{bmatrix}.$$

Abusing terminology slightly, for every α in Δ let $\alpha^\vee = \varphi_\alpha \circ \iota$. These give us a homomorphism

$$\varphi = \prod \alpha^\vee: (\mathbb{C}^\times)^\Delta \longrightarrow G.$$

It is an embedding, say with image H , which is its own centralizer in G . Let $H_{\mathbb{Z}}$ be the image of $\{\pm 1\}^\Delta$. (The notation is suggested by what happens for \mathfrak{g} finite-dimensional, in which case $H_{\mathbb{Z}}$ is the group of integral points in a maximal torus.)

(G4) Let

$$\omega(x) = \begin{bmatrix} 0 & x \\ -1/x & 0 \end{bmatrix}$$

The normalizer $N_G(H)$ is generated by H and the elements $\omega_\alpha(x) = \varphi_\alpha(\omega(x))$. The adjoint action of $\omega_\alpha(x)$ on \mathfrak{h} is the same as that of the Weyl reflection s_α . Let

$$\sigma_\alpha = \omega_\alpha(1).$$

(G5) The group W_{ext} is defined to be the subgroup of G generated by the σ_α . It fits into a short exact sequence:

$$1 \rightarrow H_{\mathbb{Z}} \rightarrow W_{\mathrm{ext}} \rightarrow W \rightarrow 1.$$

(G6) There exists a useful cross-section $w \mapsto \widehat{w}$ of the projection from W_{ext} to W . Define \widehat{s}_α to be σ_α . If $w = s_1 \dots s_n$ is an expression for w as a product of elementary reflections, the product

$$\widehat{w} = \widehat{s}_1 \dots \widehat{s}_n$$

depends only on w . Multiplication in W_{ext} is determined by the formulas

$$\widehat{s_\alpha w} = \begin{cases} \widehat{s_\alpha} \widehat{w} & \text{if } \ell(s_\alpha w) = 1 + \ell(w) \\ \alpha^\vee(-1) \widehat{s_\alpha} \widehat{w} & \text{otherwise.} \end{cases}$$

and

$$\widehat{w s_\alpha} = \begin{cases} \widehat{w s_\alpha} & \text{if } w s_\alpha > w \\ [w \alpha^\vee](-1) \widehat{w s_\alpha} & \text{otherwise.} \end{cases}$$

The properties of the cross section were first found by Tits and proved in [Tits66b] in the case of finite-dimensional \mathfrak{g} . Curiously, the explicit cocycle of the extension of W by $H_{\mathbb{Z}}$ was first exhibited in [LS87].

The generalization to arbitrary Kac–Moody algebras is by Kac and Peterson.

Why is W_{ext} relevant to structure constants? The connection between $N_G(H)$ and nilpotent elements of the Lie algebra can be seen already in SL_2 . In this group, the normalizer breaks up into two parts, the diagonal matrices and the inverse image M of the non-trivial element of the Weyl group, made up of the matrices $\omega(x)$ for $x \neq 0$. The matrix $\omega = \omega(x)$ satisfies

$$\omega^2 = \begin{bmatrix} 0 & x \\ -1/x & 0 \end{bmatrix}^2 = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}, \quad \text{hence} \quad \sigma_{\alpha}^{-1} = \alpha^{\vee}(-1)\sigma_{\alpha}.$$

According to the Bruhat decomposition, every element of SL_2 is either upper triangular or can be factored as $n_1 w(x) n_2$ with the n_i unipotent and upper triangular. Making this explicit:

$$\begin{aligned} \begin{bmatrix} 1 & 0 \\ x & 1 \end{bmatrix} &= \begin{bmatrix} 1 & 1/x \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -1/x \\ x & 0 \end{bmatrix} \begin{bmatrix} 1 & 1/x \\ 0 & 1 \end{bmatrix} \\ \begin{bmatrix} 0 & x \\ -1/x & 0 \end{bmatrix} &= \begin{bmatrix} 1 & x \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1/x & 1 \end{bmatrix} \begin{bmatrix} 1 & x \\ 0 & 1 \end{bmatrix}. \end{aligned}$$

Some easy calculating will verify further the following observation of Tits:

Lemma 7.1. *For any σ in the non-trivial coset of H in $N_G(H)$ in SL_2 there exist unique upper-triangular nilpotent e_+ and lower-triangular nilpotent e_- such that*

$$\sigma = \exp(e_+) \exp(e_-) \exp(e_+).$$

In this equation, any one of the three determines the other two.

Applying the involution θ we see that the roles of e_+ and e_- may be reversed. In other words, specifying an upper triangular nilpotent e_+ or a lower triangular nilpotent e_- is equivalent to specifying an element of M .

Suppose $\lambda = w\alpha$ to be a root of \mathfrak{g} . If ω in W_{ext} maps to w , define $SL_2^{[\lambda]}$ to be $\omega SL_2^{[\alpha]} \omega^{-1}$. This group is independent of the choice of ω , although its identification with SL_2 is not. Let $H_{\mathbb{Z}}^{[\lambda]}$ be the image of the diagonal matrices in $SL_2^{[\lambda]}(\mathbb{Z})$, and let $M_{\mathbb{Z}}^{[\lambda]}$ be the image of the matrices

$$\begin{bmatrix} 0 & \pm 1 \\ \mp 1 & 0 \end{bmatrix}$$

in $N_G(H)$. Each $M_{\mathbb{Z}}^{[\lambda]}$ has exactly two elements in it, corresponding to the possible choices of $e_{\pm\lambda}$. In other words, there is a well-defined map taking $\sigma \in M_{\mathbb{Z}}^{[\lambda]}$ to $e_{\lambda,\sigma}$ such that

$$\sigma = \exp(e_{\lambda,\sigma}) \exp(e_{-\lambda,\sigma}) \exp(e_{\lambda,\sigma}).$$

Here is the whole point:

—————
A choice of e_λ or $e_{-\lambda}$ is equivalent to a choice of σ in $M_{\mathbb{Z}}^{[\lambda]}$.
 —————

So now we can at last see the connection between structure constants and the extended Weyl group.

Lemma 7.2. *For ω in W_{ext} with image w in W we have*

$$\omega M_{\mathbb{Z}}^{[\lambda]} \omega^{-1} = M_{\mathbb{Z}}^{[w\lambda]} .$$

and furthermore for σ in $M_{\mathbb{Z}}^{[\lambda]}$

$$\text{Ad}(\omega)e_{\lambda,\sigma} = e_{w\lambda,\omega\sigma\omega^{-1}} .$$

Let me now solve Problem 1. We must choose for each real root $\lambda > 0$ an element σ_λ in $M_{\mathbb{Z}}^{[\lambda]}$. We have already chosen σ_α for α in Δ . Now follow the ISL root graph to set

$$\sigma_\mu = \sigma_\alpha \sigma_\lambda \sigma_\alpha^{-1}$$

when $\lambda \xrightarrow{\alpha} \mu$ is an edge in it. This technique, of defining objects associated to nodes of a graph by finding a spanning tree in it, is common in computer algorithms.

8 ... and structure constants

How will this new scheme relate to structure constants? Suppose we are given choices of the elements σ_λ . These determine nilpotent elements $e_{\pm\lambda,\sigma_\lambda}$ etc. in $\mathfrak{sl}_2^{[\lambda]}$. Suppose $\lambda + \mu + \nu = 0$. According to our definitions,

$$[e_{\lambda,\sigma_\lambda}, e_{\mu,\sigma_\mu}] = \varepsilon(\lambda, \mu, \nu, \sigma_\lambda, \sigma_\mu, \sigma_\nu)(p_{\lambda,\mu} + 1)e_{-\nu,\sigma_\nu} ,$$

for some ε factor. According to the generalization of Chevalley’s theorem this factor is ± 1 , but I won’t assume that.

The dependence on the roots λ, μ, ν is manifest but weak, in that σ_λ determines both e_λ and $e_{-\lambda}$. In fact we don’t need to take it into account at all. If $\lambda + \mu + \nu = 0$ then the only other linear combination $\pm\lambda \pm \mu \pm \nu$ that vanishes is $-\lambda - \mu - \nu$. Since $\theta(e_+) = e_-$ we know that $\varepsilon(-\lambda, -\mu, -\nu, \dots) = \varepsilon(\lambda, \mu, \nu, \dots)$, so we can just write $\varepsilon(\sigma_\lambda, \sigma_\mu, \sigma_\nu)$:

Proposition 8.1. *For real roots λ, μ, ν with $\lambda + \mu + \nu = 0$ we have*

$$[e_{\lambda, \sigma_\lambda}, e_{\mu, \sigma_\mu}] = \varepsilon(\sigma_\lambda, \sigma_\mu, \sigma_\nu)(p_{\lambda, \mu} + 1)e_{-\nu, \sigma_\nu}.$$

where ε depends only on $\sigma_\lambda, \sigma_\mu, \sigma_\nu$.

How do we compute $\varepsilon(\sigma_\lambda, \sigma_\mu, \sigma_\nu)$? Following §2.9 of [Tits66a], we get four basic rules.

Theorem 8.2. *Suppose that λ, μ, ν are real roots with $\lambda + \mu + \nu = 0$. Then*

- (a) $\varepsilon(\sigma_\mu, \sigma_\lambda, \sigma_\nu) = -\varepsilon(\sigma_\lambda, \sigma_\mu, \sigma_\nu)$.
- (b) $\varepsilon(\sigma_\lambda, \sigma_\mu, \sigma_\nu^{-1}) = -\varepsilon(\sigma_\lambda, \sigma_\mu, \sigma_\nu)$.
- (c) *Suppose $\|\lambda\| \geq \|\mu\|, \|\nu\|$. In this case $s_\lambda \mu = -\nu$ and $s_\nu = s_\lambda s_\mu s_\lambda$, so $\sigma_\lambda \sigma_\mu \sigma_\lambda^{-1}$ lies in $M_{\mathbb{Z}}^{[v]}$ and satisfies*

$$\varepsilon(\sigma_\lambda, \sigma_\mu, \sigma_\lambda \sigma_\mu \sigma_\lambda^{-1}) = (-1)^{p_{\lambda, \mu}}.$$

- (d) *The function ε is invariant under cyclic permutations:*

$$\varepsilon(\sigma_\lambda, \sigma_\mu, \sigma_\nu) = \varepsilon(\sigma_\mu, \sigma_\nu, \sigma_\lambda).$$

Proof. (a) follows from Lemma 5.5, since $[x, y] = -[y, x]$. (b) is elementary, since $e_{\nu, \sigma_\nu^{-1}} = -e_{\nu, \sigma_\nu}$.

It is the last two results that are significant. The first tells us that there is one case in which we can calculate the constant explicitly, and the second tells us that we can manipulate any case so as to fall in this one.

For (c): Proposition 6.1 tells that in this case

$$\begin{aligned} [e_{\lambda, \sigma_\lambda}, e_{\mu, \sigma_\mu}] &= (-1)^{p_{\lambda, \mu}}(p_{\lambda, \mu} + 1) \text{Ad}(\sigma_\lambda)e_{\mu, \sigma_\mu} \\ &= (-1)^{p_{\lambda, \mu}}(p_{\lambda, \mu} + 1)e_{s_\lambda \mu, \sigma_\lambda \sigma_\mu \sigma_\lambda^{-1}}. \end{aligned}$$

For (d): This reduces to the fact that the Jacobi identity has cyclic symmetry. For e_λ, e_μ , and e_ν it tells us that

$$\begin{aligned} 0 &= [[e_\lambda, e_\mu], e_\nu] + [[e_\mu, e_\nu], e_\lambda] + [[e_\nu, e_\lambda], e_\mu] \\ &= \varepsilon(\sigma_\lambda, \sigma_\mu, \sigma_\nu)(p_{\lambda, \mu} + 1)[e_{-\nu}, e_\nu] \\ &\quad + \varepsilon(\sigma_\mu, \sigma_\nu, \sigma_\lambda)(p_{\mu, \nu} + 1)[e_{-\lambda}, e_\lambda] \\ &\quad + \varepsilon(\sigma_\nu, \sigma_\lambda, \sigma_\mu)(p_{\nu, \lambda} + 1)[e_{-\mu}, e_\mu] \\ &= \varepsilon(\sigma_\lambda, \sigma_\mu, \sigma_\nu)(p_{\lambda, \mu} + 1)h_\nu \\ &\quad + \varepsilon(\sigma_\mu, \sigma_\nu, \sigma_\lambda)(p_{\mu, \nu} + 1)h_\lambda \\ &\quad + \varepsilon(\sigma_\nu, \sigma_\lambda, \sigma_\mu)(p_{\nu, \lambda} + 1)h_\mu. \end{aligned}$$

But according to Proposition 5.6,

$$(p_{\lambda, \mu} + 1)h_\nu + (p_{\mu, \nu} + 1)h_\lambda + (p_{\nu, \lambda} + 1)h_\mu = 0.$$

The vectors h_λ and h_μ are linearly independent, so the conclusion may be deduced from the following trivial observation: if u, v, w are vectors, two of them linearly independent, and

$$\begin{aligned} au + bv + cw &= 0 \\ u + v + w &= 0, \end{aligned}$$

then $a = b = c$. □

The relevance to computation should be evident. If $\lambda + \mu + \nu = 0$, then one of the three roots will at least weakly dominate in length, and by Theorem 8.2(d) we can cycle to get the condition assumed in Lemma 5.3.

Incidentally, from these rules follows Chevalley’s theorem:

Proposition 8.3. *The function ε always takes values ± 1 .*

At any rate, we are now faced with the algorithmic problem: Suppose $\lambda + \mu + \nu = 0$ with $\|\lambda\| \geq \|\mu\| = \|\nu\|$. Then $-\nu = s_\lambda \mu$ and $\sigma_\lambda \sigma_\mu \sigma_\lambda^{-1}$ lies in $M_{\mathbb{Z}}^{[\nu]}$. Is it equal to σ_ν or $\sigma_\nu^{-1} = \nu^\vee(-1)\sigma_\nu$? If $\sigma_\lambda \sigma_\mu \sigma_\lambda^{-1} = \nu^\vee(\pm 1)\sigma_\nu$ we can deduce now:

$$[e_{\lambda, \sigma_\lambda}, e_{\mu, \sigma_\mu}] = \pm(-1)^{p_{\lambda, \mu}}(p_{\lambda, \mu} + 1)e_{\nu, \sigma_\nu}.$$

9 The extended root graph

To each real root λ are associated two elements $\sigma_\lambda^{\pm 1}$ in $M_{\mathbb{Z}}^{[\lambda]}$. These are the nodes of the *extended root graph*, which is a two-fold covering of the root graph itself. Make edges from each node σ to $\sigma_\alpha \sigma \sigma_\alpha^{-1}$, and from σ to $\sigma_\alpha^{-1} \sigma \sigma_\alpha$ (for $\alpha \in \Delta$). Compute these as we assign the σ_λ . This can all be calculated by using the formulas for multiplication in W_{ext} (Figure 10).

This is what we need in order to compute the edges in the extended root graph, which we use to compute all $\sigma_\lambda \sigma_\mu \sigma_\lambda^{-1} / \sigma_\nu$ for Tits triples (λ, μ, ν) .

Let’s look at an example, the Lie algebra \mathfrak{sl}_3 .




Figure 10 The root system A_2

Set

$$\sigma_\gamma = \sigma_\alpha \sigma_\beta \sigma_\alpha^{-1} = s_\alpha s_\beta \alpha^\vee (-1) \hat{s}_\alpha \hat{s}_\beta \hat{s}_\alpha = \beta^\vee (-1) \hat{s}_\gamma.$$

Thus by definition $[e_\alpha, e_\beta] = e_\gamma$.

We can also use Tits' trick to calculate $[e_\beta, e_\alpha]$. We know

$$[e_\beta, e_\alpha] = e_{\gamma, \sigma_\beta \sigma_\alpha \sigma_\beta^{-1}}.$$

But

$$\sigma_\beta \sigma_\alpha \sigma_\beta^{-1} = s_\beta s_\alpha \beta^\vee (-1) \hat{s}_\beta \hat{s}_\alpha \hat{s}_\beta = \alpha^\vee (-1) \hat{s}_\gamma = \gamma^\vee (-1) \sigma_\gamma,$$

so $[e_\beta, e_\alpha] = -e_\gamma$. Of course this just confirms what we already knew.

10 Admissible triples

At this point I have explained how to construct the σ_λ and computed as much as we want of the extended root graph.

According to Proposition 8.2, we can reduce the computation of the factors ε to the case in which the arguments come from a Tits triple (λ, μ, ν) , and in this case it reduces more precisely to a calculation of a comparison of $\sigma_\lambda \sigma_\mu \sigma_\lambda^{-1} / \sigma_\nu$.

We therefore start by making a list of Tits triples, and the ISL tree can be used to do this. We start by dealing directly with all cases in which $\lambda = \alpha$ lies in Δ and $\mu > 0$, following from λ up the ISL tree. The cases where $\mu < 0$ can be dealt with at the same time, since if $\alpha + \mu = \nu$, then $-\nu + \alpha = -\mu$. Then we add the ones where λ is not in Δ . If (λ, μ, ν) is a Tits triple with $s_\alpha \lambda > 0$, then so is

$$s_\alpha(\lambda, \mu, \nu) = (s_\alpha \lambda, s_\alpha \mu, s_\alpha \nu).$$

Thus we can compile a complete list of Tits triples by going up the ISL tree.

We must then compare $\sigma_\lambda \sigma_\mu \sigma_\lambda^{-1}$ to σ_ν for all Tits triples (λ, μ, ν) . This computation may also be done inductively in the ISL tree, since

$$\sigma_\alpha \sigma_\lambda \sigma_\alpha^{-1} \cdot \sigma_\alpha \sigma_\mu \sigma_\alpha^{-1} \cdot \sigma_\alpha \sigma_\lambda^{-1} \sigma_\alpha^{-1} = \sigma_\alpha \cdot \sigma_\lambda \sigma_\mu \sigma_\lambda^{-1} \cdot \sigma_\alpha^{-1}.$$

To see exactly how this goes, fix for the moment assignments $\lambda \mapsto \sigma_\lambda$. For every triple (λ, μ, ν) with $s_\lambda \mu = \nu$ (not just Tits triples) define the factor $\tau_{\lambda, \mu, \nu}^\vee$ by the formula

$$\sigma_\lambda \sigma_\mu \sigma_\lambda^{-1} = \tau_{\lambda, \mu, \nu}^\vee \cdot \sigma_\nu.$$

The factor $\tau_{\lambda, \mu, \nu}^\vee$ will lie in $H_\nu(\mathbb{Z})$, hence will be either 1 or $\nu^\vee(-1)$. I will show in the next section how to compute these factors when λ lies in Δ . This will depend on something we haven't seen yet. But for the moment assume that the $\tau_{\alpha, \mu, \nu}^\vee$ have been calculated for all α in Δ and μ an arbitrary positive root. We can then proceed to calculate the τ -factors for all Tits triples by induction. Let

$$(\lambda_\bullet, \mu_\bullet, \nu_\bullet) = (s_\alpha \lambda, s_\alpha \mu, s_\alpha \nu).$$

Then

$$\begin{aligned}
\sigma_\lambda \sigma_\mu \sigma_\lambda^{-1} &= \tau_{\lambda, \mu, \nu}^\vee \cdot \sigma_\nu \\
\sigma_\alpha \sigma_\lambda \sigma_\alpha^{-1} \cdot \sigma_\alpha \sigma_\mu \sigma_\alpha^{-1} \cdot \sigma_\alpha \sigma_\lambda^{-1} \sigma_\alpha^{-1} &= \sigma_\alpha \cdot \sigma_\lambda \sigma_\mu \sigma_\lambda^{-1} \cdot \sigma_\alpha^{-1} \\
(\tau_{\alpha, \lambda, \lambda_\bullet}^\vee \cdot \sigma_{\lambda_\bullet}) (\tau_{\alpha, \mu, \mu_\bullet}^\vee \cdot \sigma_{\mu_\bullet}) (\tau_{\alpha, \lambda, \lambda_\bullet}^\vee \cdot \sigma_{\lambda_\bullet})^{-1} &= \sigma_\alpha \cdot \tau_{\lambda, \mu, \nu}^\vee \sigma_\nu \cdot \sigma_\alpha^{-1} \\
(\tau_{\alpha, \lambda, \lambda_\bullet}^\vee + s_\lambda \cdot \tau_{\alpha, \mu, \mu_\bullet}^\vee + s_\lambda \cdot s_\mu \cdot s_\lambda \cdot \tau_{\alpha, \lambda, \lambda_\bullet}^\vee) \cdot \sigma_{\lambda_\bullet} \sigma_{\mu_\bullet} \sigma_{\lambda_\bullet}^{-1} &= s_\alpha \tau_{\lambda, \mu, \nu}^\vee \cdot \sigma_\alpha \sigma_\nu \sigma_\alpha^{-1} \\
(\tau_{\alpha, \lambda, \lambda_\bullet}^\vee + s_\lambda \cdot \tau_{\alpha, \mu, \mu_\bullet}^\vee + s_\nu \cdot \tau_{\alpha, \lambda, \lambda_\bullet}^\vee) \cdot \sigma_{\lambda_\bullet} \sigma_{\mu_\bullet} \sigma_{\lambda_\bullet}^{-1} &= (s_\alpha \tau_{\lambda, \mu, \nu}^\vee + \tau_{\alpha, \nu, \nu_\bullet}^\vee) \cdot \sigma_{\nu_\bullet}
\end{aligned}$$

leading to:

Lemma 10.1. *If (λ, μ, ν) is a Tits triple and $\lambda \neq \alpha \in \Delta$, then so is*

$$(\lambda_\bullet, \mu_\bullet, \nu_\bullet) = (s_\alpha \lambda, s_\alpha \mu, s_\alpha \nu),$$

and

$$\tau_{\lambda_\bullet, \mu_\bullet, \nu_\bullet}^\vee = s_\alpha \tau_{\lambda, \mu, \nu}^\vee + \tau_{\alpha, \lambda, \lambda_\bullet}^\vee + s_\lambda \cdot \tau_{\alpha, \mu, \mu_\bullet}^\vee + s_\nu \cdot \tau_{\alpha, \lambda, \lambda_\bullet}^\vee + \tau_{\alpha, \nu, \nu_\bullet}^\vee.$$

I wish this formula were more enlightening. One must conclude that the relationship between the groups $\mathrm{SL}_2^{[\lambda]}$ and Tits' cross section is complicated.

The principal conclusion of these preliminary formulas is that for both the specification of the σ_λ and the calculation of the $\sigma_\lambda \sigma_\mu \sigma_\lambda^{-1}$ we are reduced to the single calculation: for α in Δ and $\lambda > 0$, given σ_λ how do we calculate $\sigma_\alpha \sigma_\lambda \sigma_\alpha^{-1}$? I must explain in detail not only how calculations are made, but also how elements of $N_H(\mathbb{Z})$ are interpreted in a computer program. I have already explained the basis of calculation in the extended Weyl group.

In understanding how efficient the computation of structure constants will be, we have to know roughly how many admissible triples there are. Following [Car72] and [CMT04], I assume the positive roots to be ordered, and I define a trio of roots λ, μ, ν to be *special* if $0 < \lambda < \mu$ and $\lambda + \mu = -\nu$ is again a root. If λ, μ, ν is any triple of roots with $\lambda + \mu + \nu = 0$, then (as Carter points out) exactly one of the following twelve triples is special:

$$\begin{aligned}
&(\lambda, \mu, \nu), (-\lambda, -\mu, -\nu), (\mu, \lambda, \nu), (-\mu, -\lambda, -\nu) \\
&(\nu, \lambda, \mu), (-\nu, -\lambda, -\mu), (\lambda, \nu, \mu), (-\lambda, -\nu, -\mu) \\
&(\mu, \nu, \lambda), (-\mu, -\nu, -\lambda), (\nu, \mu, \lambda), (-\nu, -\mu, -\lambda).
\end{aligned}$$

Hence there are at most 12 times as many Tits triples as special triples. How many special triples are there? This is independent of the ordering of Σ^+ , since it is one-half the number of pairs of positive roots λ, μ with $\lambda + \mu$ also a root. In [CMT04] it is asserted that for all classical systems the number is $O(n^3)$, where n is the rank of the system. Don Taylor has given me the following more precise table:

System Number of special triples

A_n	$n(n^2 - 1)/6$
B_n	$n(n - 1)(2n - 1)/3$
C_n	$n(n - 1)(2n - 1)/3$
D_n	$2n(n - 1)(n - 2)/3$
E_6	120
E_7	336
E_8	1120
F_4	68
G_2	5

11 Some details of computation

For this section I am going to simplify notation. Every element of the extended Weyl group may be represented uniquely as $t^\vee(-1) \cdot \widehat{w}$, where t^\vee is in $X_*(H) = L_{\Delta^\vee}$ and w is in W . I will express it as just $t^\vee \cdot \widehat{w}$, and of course it is only t^\vee modulo $2X_*(H)$ that counts. Also, I will refer to the group operation in L_{Δ^\vee} additively.

Proposition 11.1. *Suppose α to be in Δ , $\lambda \neq \alpha > 0$. Then*

$$\widehat{s}_\alpha \widehat{s}_\lambda \widehat{s}_\alpha = \begin{cases} \widehat{s_{s_\alpha \lambda}} & \text{if } \langle \alpha, \lambda^\vee \rangle < 0 \\ \alpha^\vee \cdot \widehat{s}_\lambda & \langle \alpha, \lambda^\vee \rangle = 0 \\ (\alpha^\vee + s_\alpha s_\lambda \alpha^\vee) \cdot \widehat{s_{s_\alpha \lambda}} & \langle \alpha, \lambda^\vee \rangle > 0. \end{cases}$$

Proof. A preliminary calculation:

$$\begin{aligned} s_\lambda \alpha &= \alpha - \langle \alpha, \lambda^\vee \rangle \lambda \\ s_\alpha s_\lambda \alpha &= -\alpha - \langle \alpha, \lambda^\vee \rangle s_\alpha \lambda \\ &= -\alpha - \langle \alpha, \lambda^\vee \rangle (\lambda - \langle \lambda, \alpha^\vee \rangle \alpha) \\ &= -\alpha - \langle \alpha, \lambda^\vee \rangle \lambda + \langle \alpha, \lambda^\vee \rangle \langle \lambda, \alpha^\vee \rangle \alpha \\ &= -\langle \alpha, \lambda^\vee \rangle \lambda + (\langle \alpha, \lambda^\vee \rangle \langle \lambda, \alpha^\vee \rangle - 1) \alpha. \end{aligned}$$

Recall that by Corollary 4.4 the product $\langle \alpha, \lambda^\vee \rangle \langle \lambda, \alpha^\vee \rangle \geq 0$. Recall also that $w s_\alpha > w$ if and only if $w\alpha > 0$.

(a) $\langle \alpha, \lambda^\vee \rangle < 0$. Here $s_\lambda \alpha > 0$ and $s_\alpha s_\lambda \alpha > 0$ so $s_\lambda < s_\alpha s_\lambda < s_\alpha s_\lambda s_\alpha$, and $\ell(s_\alpha s_\lambda s_\alpha) = \ell(s_\lambda) + 2$, and

$$\widehat{s}_\alpha \widehat{s}_\lambda \widehat{s}_\alpha = \widehat{s_{s_\alpha s_\lambda s_\alpha}} = \widehat{s_{s_\alpha \lambda}}.$$

(b) $\langle \alpha, \lambda^\vee \rangle = 0$. Here $s_\alpha \lambda = \lambda$, and $\widehat{s_\lambda \hat{s}_\alpha} = \widehat{s_\lambda \hat{s}_\alpha}$, but $s_\alpha(s_\lambda s_\alpha) = s_\lambda$ so

$$\widehat{s_\alpha \hat{s}_\lambda \hat{s}_\alpha} = \widehat{s_\alpha \widehat{s_\lambda \hat{s}_\alpha}} = \alpha^\vee \cdot \widehat{s_\lambda}.$$

(c) $\langle \alpha, \lambda^\vee \rangle > 0$. Here one sees easily that $s_\lambda \alpha < 0$. But since $\lambda \neq \alpha$ we also have $s_\alpha s_\lambda \alpha < 0$ also. So $\ell(s_{s_\alpha \lambda}) = \ell(s_\lambda) - 2$.

$$\begin{aligned} \widehat{s_\alpha \hat{s}_\lambda} &= \alpha^\vee \cdot \widehat{s_\alpha s_\lambda} \\ \widehat{s_\alpha \hat{s}_\lambda \hat{s}_\alpha} &= \alpha^\vee \cdot \widehat{s_\alpha s_\lambda \hat{s}_\alpha} \\ &= \alpha^\vee \cdot \widehat{s_\alpha \lambda} \cdot \alpha^\vee \cdot 1 \\ &= (\alpha^\vee + s_\alpha s_\lambda \alpha^\vee) \cdot \widehat{s_\alpha \lambda}. \end{aligned}$$

□

Corollary 11.2. *Suppose $\sigma_\lambda = t_\lambda^\vee \cdot \widehat{s_\lambda}$. Then for $\alpha \neq \lambda$*

$$\sigma_\alpha \sigma_\lambda \sigma_\alpha^{-1} = \begin{cases} (s_\alpha t_\lambda^\vee + s_\alpha s_\lambda \alpha^\vee) \cdot \widehat{s_\alpha \lambda} & \langle \alpha, \lambda^\vee \rangle < 0 \\ (s_\alpha t_\lambda^\vee) \cdot \widehat{s_\alpha \lambda} & \langle \alpha, \lambda^\vee \rangle = 0 \\ (s_\alpha t_\lambda^\vee + \alpha^\vee) \widehat{s_\alpha \lambda} & \langle \alpha, \lambda^\vee \rangle > 0. \end{cases}$$

Keep in mind that the reflection associated to $s_\alpha \lambda$ is $s_\alpha s_\lambda s_\alpha$.

Proof. We start with

$$\sigma_\alpha \sigma_\lambda \sigma_\alpha^{-1} = \widehat{s_\alpha t_\lambda^\vee \hat{s}_\lambda \alpha^\vee \hat{s}_\alpha} = (\widehat{s_\alpha t_\lambda^\vee} + s_\alpha s_\lambda \alpha^\vee) \cdot \widehat{s_\alpha \hat{s}_\lambda \hat{s}_\alpha}.$$

Apply the Proposition. □

There are three cases, according to whether $\langle \alpha, \lambda^\vee \rangle$ is $<$, $=$, or $>$ 0. These correspond to how the length $\ell(s_\alpha s_\lambda s_\alpha)$ relates to $\ell(s_\lambda)$. So now finally we can compute the factors $\tau_{\alpha, \mu, s_\alpha \mu}^\vee$, comparing $\sigma_\alpha \sigma_\mu \sigma_\alpha^{-1}$ to $\sigma_{s_\alpha \mu}$.

Lemma 11.3. *Suppose $s_\alpha \mu = \nu$. If $\sigma_\alpha \sigma_\mu \sigma_\alpha^{-1} = t^\vee \widehat{s}_\nu$ and $\sigma_\nu = t_\nu^\vee \widehat{s}_\nu$, then*

$$\tau_{\alpha, \mu, \nu}^\vee = t^\vee + t_\nu^\vee.$$

12 Patterns in the computation

As a consequence of the main theorem of [BH93], Bob Howlett proved that the root graph of an arbitrary Coxeter group is described by a finite automaton. What this means is that paths in the root graph are the same as paths in a certain finite state machine. The machine, although finite, can be quite large. As Figure 11 indicates, there are 31 states in the machine producing paths in the root graph of the root system we have seen earlier in Figure 1. (In this diagram, nodes give rise to equivalent states if and only if the subsequent paths out of them are equivalent.




Figure 11 The finite state machine for the root graph of an infinite root system

Each state is noted by a unique shaded box. An unshaded one with the same label signifies another occurrence. The labeled boxes are all one needs to specify the finite state machine.)

One thing evident in this picture is the repetition of the weak Bruhat order of the finite groups W_θ . (Lemma 5.2 of [Cas06] explains this.) This as well as regularity should have some significance for the extended root graph as well, but I don't yet completely understand what is going on, and I will not discuss this topic here.

References

[BH93] Brigitte Brink and Robert Howlett, A finiteness property and an automatic structure for Coxeter groups, *Mathematische Annalen* **296** (1993), 179–190.
 [Cas06] Bill Casselman, Computations in Coxeter groups II. Constructing minimal roots, *Representation Theory* **12** (2008), 260–293.
 [Car72] Roger Carter, *Simple groups of Lie type*, Wiley, 1972.
 [Che55] Claude Chevalley, Sur certains groupes simples, *Tohoku Mathematics Journal* **48** (1955), 14–66.
 [CMT04] Arjeh Cohen, Scott Murray, and Don Taylor, Computing in groups of Lie type, *Mathematics of Computation* **73** (2004), 1477–1498.

- [FK80] Igor Frenkel and Victor Kac, Affine Lie algebras and dual resonance models, *Inventiones Mathematicae* **62** (1980), 23–66.
- [Kac85] Victor Kac, *Infinite Dimensional Lie Algebras* (second edition). Cambridge University Press, Cambridge, New York, 1985.
- [KP85] V. G. Kac et D. H. Peterson, ‘Defining relations of certain infinite dimensional groups’, in *Élie Cartan et les mathématiques d’aujourd’hui*, *Astérisque*, volume hors série (1985), 165–208.
- [Kum02] Shrawan Kumar, *Kac–Moody Groups, Their Flag Varieties, and Representation Theory*, Birkhäuser, Boston, 2002.
- [LS87] R. P. Langlands and D. Shelstad, On the definition of transfer factors, *Mathematische Annalen* **278** (1987), 219–271.
- [MP95] Robert Moody and Arturo Pianzola, *Lie Algebras with Triangular Decompositions*, J. Wiley, New York, 1995.
- [MT72] Robert Moody and K. L. Teo, Tits’ systems with crystallographic Weyl groups, *Journal of Algebra* **21** (1972), 178–190.
- [Mor87] Jun Morita, Commutator relations in Kac–Moody groups, *Proceedings of the Japanese Academy* **63** (1987), 21–22.
- [Ryl00] L. J. Rylands, Fast calculation of structure constants, preprint, 2000.
- [Sat51] Ichiro Satake, On a theorem of Cartan, *Journal of the Mathematical Society of Japan* **2** (1951), 284–305.
- [Tits66a] Jacques Tits, Sur les constants de structure et le théorème d’existence des algèbres de Lie semi-simple, *Publications de l’I. H. E. S.* **31** (1966), 21–58.
- [Tits66b] ———, Normalisateurs de tores I. Groupes de Coxeter étendus, *Journal of Algebra* **4** (1966), 96–116.
- [Tits68] ———, Le problème de mots dans les groupes de Coxeter, *Symposia Math.* **1** (1968), 175–185.
- [Tits87] ———, Uniqueness and presentation of Kac–Moody groups over fields, *Journal of Algebra* **105** (1987), 542–573.
- [Tits88] ———, Groupes associés aux algèbres de Kac–Moody, *Séminaire Bourbaki*, exposé 700 (1988).