Exploring the Mukhin–Varchenko conjecture

Tom Chappell

Advisor: Professor Ole Warnaar

School of Mathematics and Physics
The University of Queensland
February 13, 2013

Contents

Introduction 1

The hypergeometric function ${}_2F_1(\alpha,\beta,\gamma;z)$ was studied extensively in the eighteenth century by Euler and Gauss. Beyond this first study it has continued to interest mathematicians and inspired developments throughout the dicipline. The hypergeometric function has three common descriptions given by a power series, a solution to the hypergeometric differential equation, and an integral formula.

The hypergeometric differential equation is

$$z(1-z)\frac{d^2F}{dz^2}+(\gamma-(\alpha+\beta+1)z)\frac{dF}{dz}-\alpha\beta F=0.$$

The regular solution to this equation is the hypergeometric function. Euler was able to show that, provided $0<\mathrm{Re}(\beta)<\mathrm{Re}(\gamma)$, this solution is also described by the following integral,

$${}_2F_1(\alpha,\beta,\gamma;z)=\frac{\Gamma(\gamma)}{\Gamma(\beta)\Gamma(\gamma-\beta)}\int_0^1 t^{\beta-1}(1-t)^{\gamma-\beta-1}(1-zt)^{-\alpha}\,dt.$$

When z is zero the hypergeometric function is one. As such, taking z to be zero and setting $\gamma=\alpha+\beta$ in the above formula results in the Euler beta integral. This integral is evaluated as a ratio of gamma functions.

$$\int_0^1 t^{\alpha-1}(1-t)^{\beta-1}\,dt=\frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}\qquad \mathrm{Re}(\alpha),\mathrm{Re}(\beta)>0.$$

As we will later show, gamma function evaluations such as this are a more general phenomenom.

The Knizhnik–Zamolodchikov equations are partial differential equations which arise from the representation theory of Lie algebras. They may be though of as a generalisation of the hypergeometric differential equation. From the Knizhnik–Zamolodchikov equations arise phase integrals similar to the Euler beta integral. When normalised, these phase integrals have the following description.

For a simple Lie algebra $\mathfrak{g}$ and highest weight representations V_μ, V_ν and V_λ with $\lambda=\mu+\nu-\sum_{i=1}^n k_i\alpha_i$, the phase integral is given by

$$\int\left[\prod_{i=1}^k t_i^{-(\mu,\alpha_{t_i})}(1-t_i)^{-(\nu,\alpha_{t_i})}\prod_{1\le i<j\le k}|t_i-t_j|^{(\alpha_{t_i},\alpha_{t_j})}\right]^\gamma dt_1\ldots dt_k.$$

In 1988 Mukhin and Varchenko [11] conjectured that if the representation V_λ occurs with multiplicity one in the tensor product decomposition of $V_\mu \otimes V_\nu$, then this integral has an evaluation as a ratio of gamma functions. Various subsequent results [10, 16–18] have supported this conjecture.

To write down the integral, let alone to evaluate it, requires knowledge of the k_i values. For any particular μ, ν and λ, determining these values is a simple calculation. However, previous results towards the conjecture suggest that a fruitful approach is to consider infinite families of μ, ν and λ where the conjecture applies. Once such an infinite family has been indentified, an explicit description of the set of k_i values which occur is required. This is the aspect of the problem which we have focused on. The goal is to use combinatorial tools to find infinite families with λ occuring with multiplicity one and then to provide a description of the corresponding k_i values.

Beyond this introduction, the document is divided into three key sections. The first provides an overview of the relevant background material required to state and understand the conjecture. In the second section we state the Mukhin–Varchenko conjecture more explicitly and examine a few previous evaluations of the integral. Finally we explore several situations where the conditions of the conjecture hold and describe explicitly the k_i values there.

Fundamentals 2

This chapter provides an overview of the key background material required to understand the later sections of the thesis. All the material found here (and more) may be found in the excellent books by Erdmann and Wildon [2], Fulton and Harris [6], and Humphereys [8]. A second book by Fulton [5] is an excellent reference to Young tableaux and the Littlewood–Richardson rule.

Each section of this introductory material presents and explores a key concept related to the later parts of the thesis. In the first section we introduce Lie algebras, the object upon which the whole field is based. Following from Lie algebras is the notion of their representations. The third section is particularly important as it introduces root systems and weight lattices which form the language of most of the thesis. To connect representations and root systems we then discuss highest weight modules. Following directly from such modules is the Weyl character formula and two helpful algorithms for decomposing tensor products.

2.1 Lie algebras.

A Lie algebra is a vector space L together with a binary operation $L \times L \to L$ called the Lie bracket. This is notated with square brackets, $(x, y) \mapsto [x, y]$, and must satisfy

$$\begin{aligned} [\cdot,\cdot] &\text{ is bilinear} \\ [x,x] &= 0 \text{ for all } x \in L \\ [x,[y,z]] + [y,[z,x]] + [z,[x,y]] &= 0 \text{ for all } x,y,z \in L. \end{aligned}$$

The third property above is called the Jacobi identity.

By applying bilinearity and the second property to $x + y$ we get

$$0 = [x + y, x + y] = [x, x] + [x, y] + [y, x] + [y, y] = [x, y] + [y, x]$$

and hence the bracket is antisymmetric, $[x, y] = -[y, x]$. Some authors replace the second property with this antisymmetry as an alternative definition.

As a canonical example of a Lie algebra, given any vector space V of dimension n, the set of linear transformations $V \to V$ with Lie bracket $[x, y] = xy - yx$ forms a Lie algebra. This is called the general linear Lie algebra over V and is denoted $\mathfrak{gl}(V)$. Often it is helpful to fix a basis for V in which case we identify $\mathfrak{gl}(V)$ with the space of $n \times n$ matrices over F, the field of V. Then we write $\mathfrak{gl}_n(F)$.

Any subalgebra of $\mathfrak{gl}_n(F)$ is called a linear Lie algebra. The most prominent example of these is the special linear Lie algebra $\mathfrak{sl}_n(F)$. This corresponds to the set of matrices in $\mathfrak{gl}_n(F)$ with trace zero. We are usually interested in the field of complex numbers so we often write just $\mathfrak{gl}_n$ and $\mathfrak{sl}_n$ instead of $\mathfrak{gl}_n(\mathbb{C})$ and $\mathfrak{sl}_n(\mathbb{C})$ respectively.

Two substructures play a key role in understanding Lie algebras; subalgebras and ideals. A subalgebra of a Lie algebra L is a vector subspace which is a Lie algebra itself under the restriction of the Lie bracket. An ideal of a Lie algebra L is a subalgebra $I \subseteq L$ satisfying $[x, y] \in I$ for all $x \in L$, $y \in I$. Unlike rings, there is no distinction between left and right ideals as the Lie bracket satisfies antisymmetry.

Using the Lie bracket we may define the product of two Lie subalgebras I and J,

$$[I, J] = \mathrm{Span}\{[x, y] : x \in I, y \in J\}.$$

With this notation the definition of an ideal may be restated as a subalgebra I of L satisfying $[I, L] = I$. Frequently we use this product notation to define new ideals.

The simplest nontrivial Lie algebras are the abelian ones. A Lie algebra is called abelian if $[L, L] = 0$ or equivalently if $[x, y] = 0$ for all $x, y \in L$. This nomenclature may be understood intuitively by considering $\mathfrak{gl}_n$. Here the Lie bracket is $[x, y] = xy - yx$ and hence,

$$[x, y] = 0 \Leftrightarrow xy = yx.$$

Closely related are the 'nearly abelian' solvable Lie algebras. These have a terminating chain of ideals where each consecutive quotient is abelian. Alternatively we define the derived algebra of L to be $[L, L]$ and the derived series by,

$$\begin{aligned} L^{(0)} &= L \\ L^{(n+1)} &= [L^{(n)}, L^{(n)}]. \end{aligned}$$

Then a Lie algebra is solvable if its derived series terminates, that is $L^{(n)} = 0$ for some $n > 0$.

There is a similar definition for nilpotent Lie algebras. Define,

$$\begin{aligned} L^{0} &= L \\ L^{n+1} &= [L, L^{n}]. \end{aligned}$$

Then a Lie algebra is nilpotent if this series terminates.

The definition of a solvable Lie algebra is very similar to that of a solvable group. Unsurprisingly we also have a similar definition to simple groups where ideals play the role of normal subgroups. With this is mind, a Lie algebra is simple if it is not abelian and has no nontrivial ideals. That is, the only ideals of L are 0 and L itself.

Slightly more complicated are the semisimple Lie algebras. If a Lie algebra has no nontrivial *solvable* ideals it is called semisimple. Clearly every simple Lie algebra is semisimple and hence the name. As a convention, in later sections we will use L to denote a general Lie algebra and $\mathfrak{g}$ for a semisimple Lie algebra.

In this document, and throughout much of the literature, the focus is on the simple Lie algebras. The remainder of this section is a justification of this seemingly narrow focus with two key lemmas about the structure of Lie algebras.

Define the radical of L, $\operatorname{Rad} L$, to be a maximal solvable ideal of L. That is, for any solvable ideal I of L there must hold $I \subseteq \operatorname{Rad} L$. Since sums of solvable ideals are solvable, the radical is unique.

The first lemma below demonstrates that any Lie algebra has a solvable ideal with a semisimple quotient. The second lemma then describes the structure of semisimple Lie algebras. Although not required anywhere else in the thesis we give the proof of these two lemmas for interest's sake. The first proof uses only the elementary definitions already described. The second relies of some more advanced machinery so we only give a sketch of the result.

Lemma 2.1.1. If L is a Lie algebra then $L/\operatorname{Rad} L$ is semisimple.

Proof. First we show that for an ideal I of L, if I and L/I are both solvable, then so too is L. Let $\phi : L \to L/I$ be the canonical homomorphism, $\phi(a) = a + I$. Then we have $\phi(L^{(0)}) = \phi(L) = L/I = (L/I)^{(0)}$. Furthermore,

$$\begin{aligned}
\phi(L^{(k+1)}) &= \phi([L^{(k)}, L^{(k)}]) \\
&= [\phi(L^{(k)}), \phi(L^{(k)})] \\
&= [(L/I)^{(k)}, (L/I)^{(k)}] \\
&= (L/I)^{(k+1)}.
\end{aligned}$$

That is, the homomorphism preserves the derived series. Then noting that I is not necessarily an ideal of $L^{(k)}$ we have,

$$\begin{aligned}
(L/I)^{(k)} &= \phi(L)^{(k)} \\
&= \phi(L^{(k)}) \\
&= (L^{(k)} + I)/I.
\end{aligned}$$

Since L/I is solvable, $(L/I)^{(n)} = 0$ for some n. Using the above equivalence there thus holds $L^{(n)} \subseteq I$. But I is also solvable, so $I^{(m)} = 0$ and hence $(L^{(n)})^{(m)} = L^{(n+m)} = 0$. Therefore L is solvable.

Now suppose that J is a solvable ideal of $L/\operatorname{Rad} L$. Then there exists some ideal I with $\operatorname{Rad} L \subseteq I \subseteq L$ and $\phi(I) = J$. From the result just proved this I would be solvable. But $\operatorname{Rad} L$ is maximal and hence $I \subseteq \operatorname{Rad} L$ and so $J = 0$. Thus $L/\operatorname{Rad} L$ has no nontrivial solvable ideals and as such is semisimple. □

Lemma 2.1.2. Every semisimple Lie algebra is a direct sum of simple Lie algebras.

Proof. This result uses Cartan's criterion for semisimplicity. The Killing form is a product on a Lie algebra. The product is described somewhat in Section 2.4 where it arises to form a root system. Cartan's criterion states that a Lie algebra is semisimple if and only if the Killing form is nondegenerate.

From here the proof of the lemma follows by considering orthogonal complements with respects to this Killing form. If $\mathfrak{g}$ is semisimple but not simple it has some nontrivial ideal I. Then $I^{\perp}$ is also an ideal of $\mathfrak{g}$ and since the Killing form is nondegenerate, $\mathfrak{g} = I \oplus I^{\perp}$. Restricting the form then tells us that I and $I^{\perp}$ are also semisimple. By repeating this process we get the result. □

2.2 Representations

Representations and modules are methods to make a vector space resemble an algebraic object. In particular we are interested in representations of Lie algebras. The hope is that understanding the structure of a Lie algebra's representations will aid in understanding the Lie algebra itself.

Suppose V is a vector space over the same field as a Lie algebra L. Then a representation of L is a homomorphism $\varphi : L \to \mathfrak{gl}(V)$. Here homomorphism means that the Lie bracket is preserved, that is,

$$\varphi([x,y]) = [\varphi(x), \varphi(y)] = \varphi(x)\varphi(y) - \varphi(y)\varphi(x).$$

A module of L is a vector space V over the same field as L together with an action $L \times V \to V$ denoted $x \cdot v$. This action must satisfy bilinearity,

$$(\lambda x + \mu y) \cdot (\sigma v + \tau w) = \lambda\sigma(x \cdot v) + \mu\tau(y \cdot w) + \lambda\tau(x \cdot w) + \mu\sigma(y \cdot v)$$

and be consistent with the Lie bracket,

$$[x,y] \cdot v = x \cdot (y \cdot v) - y \cdot (x \cdot v)$$

for all elements $x, y \in L$, $v, w \in V$ and λ, μ, σ and $\tau \in F$. We say that V is an L-module.

Modules and representations are effectively two equivalent ways to describe the same object. Given a module, define the function $\phi_x(y)$ to be the homomorphism of V induced by the action of x, $\phi_x(y) = x \cdot y$. Then the representation corresponding to the module is given by $x \mapsto \phi_x$.

As a first example we give the trivial representation. This representation maps every element of the Lie algebra to the zero element of $\mathfrak{gl}(V)$. It automatically satisfies the bilinearity and homomorphism conditions. However this representation is not very interesting.

Second is the vector, standard or defining representation (these are all names for the same object). This is the inclusion map of any linear Lie algebra into $\mathfrak{gl}_n(\mathbb{C})$. For the classical Lie algebras this is their usual description as a set of matrices, hence the name defining representation.

Of more interest is the adjoint representation. This representation is very useful as it is closely linked to the Lie bracket. Given any $x \in L$ we define a map $\mathrm{ad}_x : L \to L$ by,

$$\mathrm{ad}_x(y) = [x,y].$$

Then the adjoint representation maps each $x \in L$ to ad_x. This makes L itself an L-module.

Given two representations of L, we may combine them to form a new representation with the direct sum. Here the homomorphisms simply act on each component separately.

Submodules of a module are vector subspaces which are modules in their own right. Two subspaces which are always submodules are the trivial cases: the zero space and the whole module. A module is said to be irreducible if it has no nontrivial submodules. Conversely, V is a reducible module if it can be decomposed into a direct sum of two submodules, $V = V_1 \oplus V_2$. Continuing this process, a module is completely reducible if it may be decomposed as a direct sum of irreducible modules. It is a theorem of Weyl that every finite dimensional module of a Lie algebra is completely reducible with a unique decomposition.

More complicated than the direct sum is the notion of the tensor product of two representations of a Lie algebra. Consider representations $\phi : L \to \mathfrak{gl}(V)$ and $\psi : L \to \mathfrak{gl}(W)$ of the Lie algebra L into the vector spaces V and W. Then the tensor product of these representations is the map $(\phi \otimes \psi) : L \to \mathfrak{gl}(V \otimes W)$ which sends each element g of L to a linear transformation of $V \otimes W$. Denoting this transformation by $(\phi \otimes \psi)_g$ it is then described by,

$$(\phi \otimes \psi)_g(u \otimes v) = \phi_g(u) \otimes v + u \otimes \psi_g(v).$$

We give the same definition again in the language of modules. Given L-modules V and W, the tensor product $V \otimes W$ is an L-module under the action,

$$x \cdot (v \otimes w) = (x \cdot v) \otimes w + v \otimes (x \cdot w).$$

Later in this chapter we study the structure of these tensor products in greater detail. Particularly we are interested in decomposing the tensor products into a direct sum of irreducible modules.

2.3 Root systems

Root systems arise naturally in the study of irreducible representations of simple Lie algebras. Before developing this connection, it is helpful to introduce the concept abstractly so as to develop a strong geometric intuition about these structures.

Reflections form the basic building blocks of a root system. For any vector α in some Euclidean space E we denote the reflection through the hyperplane orthogonal to α with s_α. This reflection s_α acting on a vector β is described explicitly by the formula,

$$s_\alpha(\beta) = \beta - \frac{2(\beta, \alpha)}{(\alpha, \alpha)}\alpha.$$

The quantity $2(\beta, \alpha)/(\alpha, \alpha)$ occurs so frequently in this theory that we abbreviate it by $\langle \beta, \alpha \rangle$. In contrast to the standard Euclidean inner product (β, α) this is not symmetric.

Root systems are, in essence, a set of vectors with a certain degree of reflective symmetry. As a precursor to the definition, consider the example below in Figure 2.1. This root system is known as B_2. Notice that these vectors are invariant

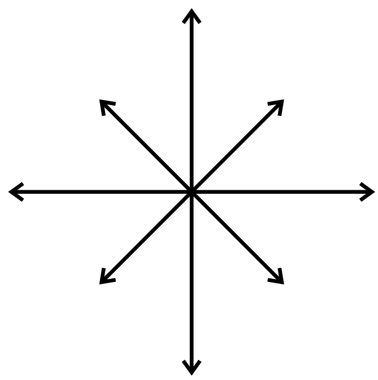

Figure 2.1: The root system B_2.

under any of their reflections. Also observe that the only colinear vectors are negatives of each other. These two properties together with restriction on the angles between the roots characterise a root system. With this in mind we define a root system.

Definition 2.3.1. A root system is a set of vectors Φ contained in some Euclidean space E that satisfy the following axioms.

- Φ is finite and spans E.
- For all $\alpha \in \Phi$, the only multiples of α in the root system are α and $-\alpha$.

◇ $s_\alpha(\Phi) = \Phi$ for all $\alpha \in \Phi$.

◇ $\langle \alpha, \beta \rangle \in \mathbb{Z}$ for all $\alpha, \beta \in \Phi$.

Unsurprisingly the elements of Φ are called roots. The dimension of E is said to be the rank of the root system. The second or fourth requirements above are occasionally omitted in the definition of a root system. Then a root system satisfying the second is called reduced and a root system satisfying the fourth is called crystallographic. We will not need to make these distinctions.

To describe root systems we will usually embed them in $\mathbb{R}^n$. It is sometimes convenient to have E only a subspace of $\mathbb{R}^n$. With this in mind, to clarify some later notation, we will use r for the rank of the root system and n for the dimension of a particular embedding.

A subset $\Delta \subseteq \Phi$ is a base of Φ if Δ is a basis for E and every element of Φ may be expressed in this basis with either strictly positive or strictly negative coordinates. Elements of a base Δ are called simple roots. Given a particular choice of simple roots, the subset Φ^+ of Φ containing all roots with positive coordinates is called the positive roots. Throughout this exposition we use $\alpha_1, \ldots, \alpha_r$ to label the simple roots of a rank r system.

Two further examples of root systems are given in Figure 2.2 with a particular choice of simple roots shown. Note that B_3 is three dimensional. Indeed there exist root systems of any dimension.

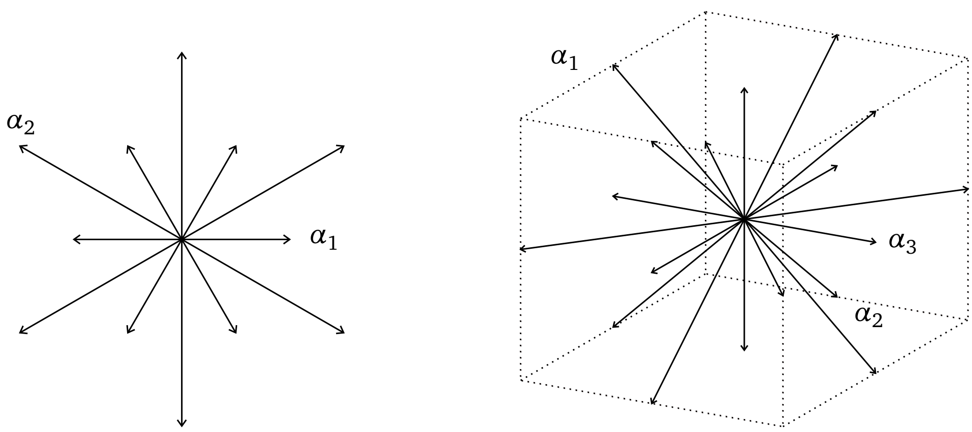

Figure 2.2: The root systems G_2 and B_3.

An important definition is that of the coroots. When a vector α has length $\sqrt{2}$ the two products (β, α) and $\langle \beta, \alpha \rangle$ will coincide. With this in mind, for a root α, the coroot $\alpha^\vee$ is a rescaling of α about the length $\sqrt{2}$.

$$\alpha^\vee = \frac{2\alpha}{(\alpha, \alpha)}.$$

This also gives $\langle \beta, \alpha \rangle = (\beta, \alpha^\vee)$.

As a consequence of the axioms of a root system, the set of reflections $\{s_\alpha : \alpha \in \Phi\}$ forms a group under composition called the Weyl group. This group is generated by the reflections corresponding to simple roots which are called simple reflections. Throughout this document we will use $\sigma_1, \ldots, \sigma_r$ to denote the simple reflections corresponding to the simple roots $\alpha_1, \ldots, \alpha_r$. The reflections give a natural action of the Weyl group on the Euclidean space in which the root system is embedded. Fundamental domains of this action are referred to as Weyl chambers and subspaces fixed by any reflection are called the walls of the Weyl chambers. Occasionally we will study similar actions with the reflections about a different point.

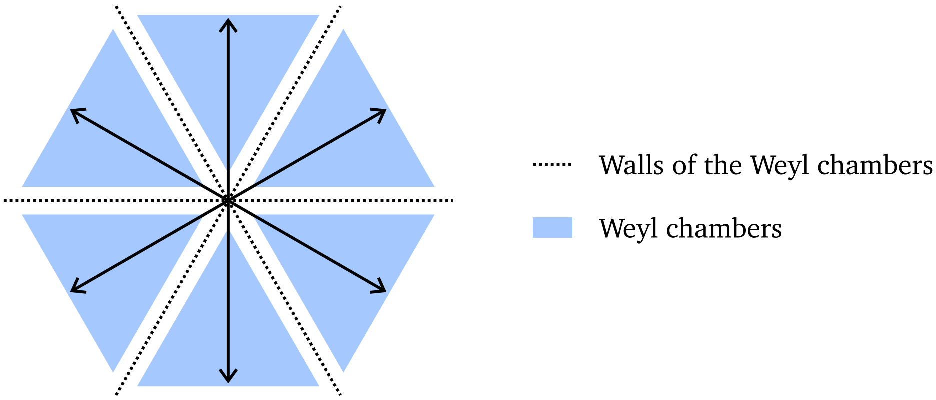

Figure 2.3: The root system A_2 with Weyl chambers and their walls highlighted.

As mentioned before, the fourth axiom of a root system is referred to as the crystallographic condition. This condition strongly restricts the possible angles θ that may occur between roots. Using the standard formula $(\alpha, \beta) = \|\alpha\| \, \|\beta\| \cos(\theta)$ we may rewrite the angle between roots as $\langle \alpha, \beta \rangle \langle \beta, \alpha \rangle = 4\cos^2(\theta)$. Table 2.1 lists the possible angles that may occur such that the roots match the crystallographic condition. Between simple roots the angles will be in the range $\pi/2 \leq \theta < \pi$.

Table 2.1: Angles between roots and corresponding Dynkin edges.

$\langle \alpha, \beta \rangle \langle \beta, \alpha \rangle$	Possible angles	Dynkin edge
0	$\pi/2$	α ○ ○ β
1	$\pi/3, 2\pi/3$	α ○—○ β
2	$\pi/4, 3\pi/4$	α ○⇒○ β
3	$\pi/6, 5\pi/6$	α ○⭆○ β
4	$0, \pi$	

The information in a root system is contained in the angles between the roots. In particular, to describe the system we need to know the angles between any pair of simple roots. This information is frequently encoded in two ways, by the Dynkin diagram and by the Cartan matrix.

The Dynkin diagram is a graph with a vertex for each simple root. Edges between these vertices indicate the angle between the two corresponding simple roots. An angle of $\pi/2$ is indicated with no edge, $2\pi/3$ by one edge, $3\pi/4$ by two edges and $5\pi/6$ by three edges. If the roots are of different lengths, then an arrow is drawn pointing towards the simple root of smaller length.

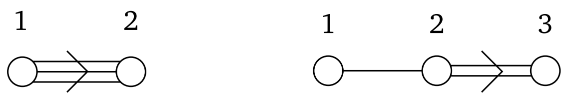

Figure 2.4: The Dynkin diagrams of the root systems G_2 and B_3.

The Cartan matrix contains the same information as the Dynkin diagram. This is the matrix given by $C = \left(\langle \alpha_i, \alpha_j \rangle\right)_{1 \leq i,j \leq r}$. Recall that r is the rank of the root system, so the elements of the Cartan matrix cover every pair of simple roots.

$$\begin{pmatrix} 2 & -1 & 0 \\ -1 & 2 & -2 \\ 0 & -1 & 2 \end{pmatrix}$$

Figure 2.5: The Cartan matrix for the root system B_3.

The Cartan matrix, Dynkin diagram and standard embedding of the classical root systems can be found in Appendix A.

New root systems may be constructed from old ones through the direct sum of the Euclidean spaces which contain them. Conversely, if a root system can be divided into two sets which are pairwise orthogonal, it may be split up into two smaller systems. In the Dynkin diagram this occurs as disconnected components of the graph and in the Cartan matrix as block matrices along the diagonal. A root system which cannot be split in this manner is called irreducible. These irreducible root systems turn out to be in one to one correspondence with the simple Lie algebras.

An important result is the classification of irreducible root systems. They fall into four infinite families, $\mathrm{A}_n, \mathrm{B}_n, \mathrm{C}_n$ and D_n and five exceptional cases, $\mathrm{E}_6, \mathrm{E}_7, \mathrm{E}_8, \mathrm{F}_4$ and G_2.

Of equal interest to representation theory are the integral weights. These are the points μ in E for which $\langle \mu, \alpha \rangle \in \mathbb{Z}$ for all $\alpha \in \Phi$. Many sets of weights form a basis of E with a particularly useful one given by the fundamental weights. These are the r weights Λ_i defined for a choice of simple system by $\langle \Lambda_i, \alpha_j \rangle = \delta_{i,j}$. Here

$\delta_{i,j}$ is the Kronecker delta,

$$\delta_{i,j} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

The set of all weights is an integer lattice over the fundamental weights.

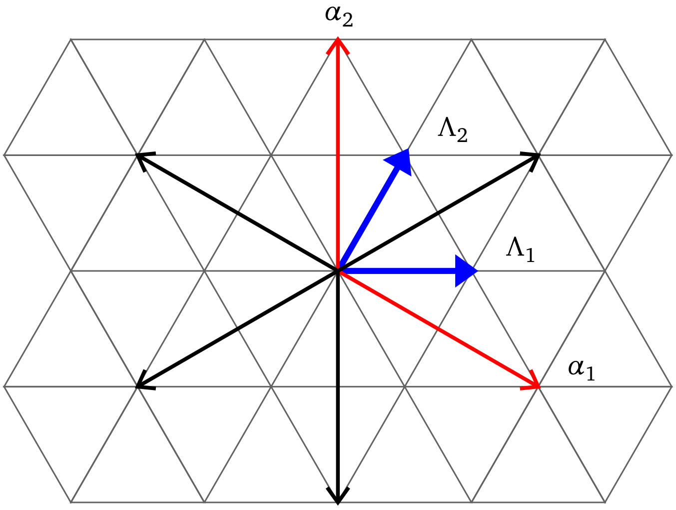

Figure 2.6: A section of the weight lattice for the root system A_2.

A useful element of the weight lattice is the Weyl vector ρ, half the sum of the positive roots or equivalently the sum of the fundamental weights.

$$\rho = \frac{1}{2} \sum_{\alpha \in \Phi^+} \alpha = \sum_{i=1}^{r} \Lambda_i$$

The fundamental weights are in the walls of a Weyl chamber since by definition they are orthogonal to all but one simple root. We also highlight a single Weyl chamber that is 'maximal' in the direction of the fundamental weights. Writing a weight μ as,

$$\mu = \sum_{i=1}^{r} \mu_i \Lambda_i,$$

the dominant chamber is the Weyl chamber formed from the μ with $\mu_i > 0$ for all i.

The fundamental weights form a basis for E. Writing μ over the fundamental weights as before we have

$$\mu = \sum_{i=1}^{n} \mu_i \Lambda_i$$

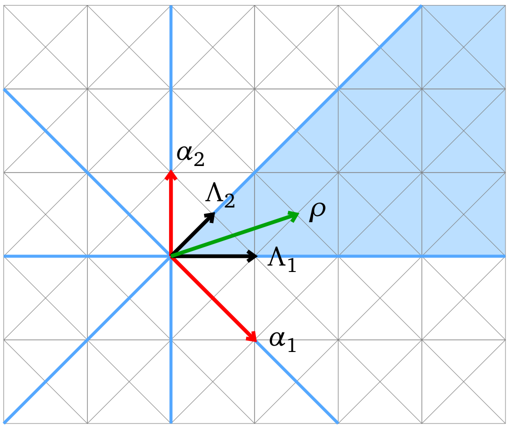

Figure 2.7: The weight lattice for the root system B_2 with the Weyl vector, dominant chamber and walls of the Weyl chambers shown.

so that by $\langle \Lambda_i, \alpha_j \rangle = \delta_{i,j}$ there holds

$$\langle \mu, \alpha_j \rangle = \mu_j.$$

That is $\langle \mu, \alpha_j \rangle$ gives the coordinate of μ in Λ_j. In particular the simple roots can be expressed as a sum over the fundamental weights with coefficients given by the Cartan matrix. Similarly, the fundamental weights may be expressed in terms of the simple roots using the rows of C^{-1}.

The simple reflections, which generate the Weyl group, are particularly easy to describe when they act on points in the weight lattice. Since $\langle \Lambda_i, \alpha_j \rangle = \delta_{i,j}$, the reflection of a fundamental weight leaves it unchanged unless $i = j$. In that case $s_{\alpha_i}(\Lambda_i) = \Lambda_i - \langle \Lambda_i, \alpha_i \rangle \alpha_i = \Lambda_i - \alpha_i$. So for a general weight $\mu = \sum_{i=1}^{n} \mu_i \Lambda_i$ and any $\alpha_i \in \Delta$ we have,

$$s_{\alpha_i}(\mu) = \mu - \mu_i \alpha_i.$$

2.4 Highest weight modules

The reason for introducing root systems and weight lattices in the previous section is that they arise in the study of representations of simple Lie algebras. In particular we are interested in representations indexed by a unique highest weight vector. This class of representation includes the irreducible finite-dimensional representations. In highest weight modules, roots and weights completely describe the module. Weights will give the vector space structure of the representation and roots will tell us how the Lie algebra acts on the module.

Roots and weights of a semisimple Lie algebra

To develop root systems in this context we need to introduce a very important object, the Cartan subalgebra. Given a lie algebra L and a subalgebra $H \subseteq L$, we define the idealiser of H to be $I(H)$,

$$I(H) = \{x \in L : [y, x] \in H \; \forall y \in H\}.$$

With this definition $I(H)$ is the largest subalgebra of L containing H as an ideal, hence the terminology idealiser.

We say that a subalgebra $\mathfrak{h}$ of L is a Cartan subalgebra if $\mathfrak{h}$ is nilpotent and $\mathfrak{h} = I(\mathfrak{h})$.

For semisimple Lie algebras we often talk instead about a maximal toral subalgebra. For a semisimple Lie algebra $\mathfrak{g}$, a subalgebra is toral if it consists only of elements which act diagonally. A maximal toral subalgebra $\mathfrak{h}$ is a toral subalgebra not contained in any other. That is $\mathfrak{h} \subset H \subseteq \mathfrak{g}$ implies that H is not toral. For semisimple Lie algebras, Cartan subalgebras and maximal toral subalgebras are equivalent.

Importantly Cartan subalgebras exist and are unique in the sense that any two differ only by an automorphism of the Lie algebra. Furthermore for semisimple Lie algebras they are abelian, $[\mathfrak{h}, \mathfrak{h}] = 0$.

Let $\mathfrak{g}$ be a semisimple Lie algebra. Suppose that V is a $\mathfrak{g}$-module with a corresponding action of $\mathfrak{g}$ on V. Consider the restriction of this action to an action of $\mathfrak{h}$ on V. This makes V into an $\mathfrak{h}$-module.

By definition, the elements of $\mathfrak{h}$ all act diagonally on V, so for each element of $\mathfrak{h}$ we may decompose V into a direct sum of eigenspaces. Now $\mathfrak{h}$ is abelian so the elements of $\mathfrak{h}$ commute and therefore are simultaneously diagonalisable, i.e., have common eigenspaces. Hence it makes sense to talk about eigenspaces for the action of $\mathfrak{h}$ on V.

In contrast, the eigenvalues need not be common among elements of $\mathfrak{h}$. Instead we introduce an 'eigenvalue function'. This is a function λ in the dual space of the Cartan subalgebra, $\lambda \in \mathfrak{h}^*$. The function λ is associated to a particular eigenspace and maps each element of $\mathfrak{h}$ to its eigenvalue. For any $\lambda \in \mathfrak{h}^*$ we define the weight space,

$$V(\lambda) = \{v \in V : h \cdot v = \lambda(h)v\}.$$

When $V(\lambda) \neq 0$ we call λ a weight of the module. The dimension of the weight space is referred to as the multiplicity of the weight.

Since V decomposes into weight spaces, once we know the weights we know the space V. To complete the module, what remains is to determine the action of the Lie algebra. This is where the roots come in to play.

Recall from Section 2.2 the adjoint representation

$$\mathrm{ad}_x : L \to L$$

which is defined by

$$\mathrm{ad}_x(y) = [x, y].$$

This gives L itself as an L-module.

Following the weight construction by restricting the adjoint representation to an $\mathfrak{h}$-module we get a decomposition of $\mathfrak{g}$ into weight spaces. One of these weight spaces, with weight 0 is $\mathfrak{h}$ itself. For this particular representation the nonzero weights are called roots and it gives rise to the Cartan decomposition,

$$\mathfrak{g} = \mathfrak{h} \oplus \bigoplus_{\alpha \in \Phi} \mathfrak{g}(\alpha).$$

The roots identified through this decomposition form a root system in the sense of Section 2.3.

As alluded to previously, the roots describe the action of the module. From the Cartan decomposition, we know that $\mathfrak{g}$ is generated by the root spaces and the Cartan subalgebra itself. Consider the weight space decomposition of a $\mathfrak{g}$-module V. From the definition of weight space we know that $\mathfrak{h}$ maps each weight space to itself. The question remains; how do elements from the root spaces act on the weight spaces? This is solved by what Fulton and Harris [6] refer to as the 'fundamental calculation'. For x, v and h, elements of the root space, weight space and Cartan subalgebra respectively we have,

$$\begin{aligned} [h, x] \cdot v &= h \cdot (x \cdot v) - x \cdot (h \cdot v) \\ \Rightarrow h \cdot (x \cdot v) &= x \cdot (h \cdot v) + [h, x] \cdot v. \end{aligned}$$

Now since v is in a weight space, $h \cdot v = \lambda(h)v$. Similarly, as x is in a root space and $[h, x]$ is the adjoint action of h on x we have $[h, x] = \alpha(h)x$. Together these give

$$h \cdot x \cdot v = (\lambda(h) + \alpha(h))\, x \cdot v.$$

So elements of the root space $\mathfrak{g}(\alpha)$ with root α map the weight space $V(\lambda)$ into the weight space $V(\lambda+\alpha)$. This describes the action of the module.

Before proceeding to discuss modules generated by a weight, it is interesting to identify the Euclidean space in which the root system is embedded. This is $\mathfrak{h}^*_{\mathbb{R}}$, the subspace of $\mathfrak{h}^*$ formed from real linear combinations of the roots.

The Euclidean inner product arises from the Killing form, a product defined over $\mathfrak{g}$. This is a symmetric bilinear form $\langle\cdot,\cdot\rangle$ which is again related to the adjoint representation,

$$\begin{aligned}\langle\cdot,\cdot\rangle &: \mathfrak{g}\times\mathfrak{g}\to\mathbb{C}\\ \langle x,y\rangle &= \mathrm{tr}(\mathrm{ad}_x\mathrm{ad}_y).\end{aligned}$$

This form is nondegenerate, even when restricted to $\mathfrak{h}$ (Cartan's criterion). So we use it to define a bijection between $\mathfrak{h}$ and $\mathfrak{h}^*$ by $x\mapsto f_x(y)=\langle x,y\rangle$. We then expand the Killing form to a product on $\mathfrak{h}^*$ by defining $(f_x,f_y)=\langle x,y\rangle$. Over $\mathfrak{h}^*_{\mathbb{R}}$ this is a Euclidean inner product which gives rise to the Euclidean space containing the root system.

Highest weight modules

Once a particular base of the root system has been chosen we inherit from the dominant chamber a 'positive' direction in the space of weights. This positive direction corresponds to greater coefficients in the fundamental weights. A weight μ in a module is called highest if it is furtherest in this direction amongst all the weights. Equivalently such a weight is annihilated by the action of every root space corresponding to a positive root. That is, $V(\mu+\alpha)=0$ for all $\alpha\in\Phi^+$.

We are interested here in representations of $\mathfrak{g}$ generated by a particular highest weight vector μ. This class of representations includes the finite dimensional irreducible representations. For these the weights are integral. That is, they are weights in the sense of Section 2.3. For the remainder of this section we assume all weights are integral.

For any integral highest weight vector μ in the dominant Weyl chamber (including its walls) there is a unique irreducible highest weight module. We notate this by V_μ. The weights of this module are all those points in the weight lattice that may be obtained by subtracting simple roots from μ with the condition that the resulting set of weights must be invariant under the reflections of the Weyl group. It is perhaps simpler to describe these weights as being contained in a convex hull formed by the orbit of the highest weight μ.

The Weyl orbit of a weight μ together with the weights between adjacent points on this orbit is called a shell. To restate, the convex hull is then exactly the shell which contains the highest weight. The multiplicity of the weights is invariant among each shell.

Given a root system and the set of weights obtained for a highest weight module we almost know the entire module. Each simple root describes the action of the generators of the Lie algebra on all of the weight spaces. What we do not yet know is the multiplicities of the weights, that is, the dimension of the weight spaces.

Before proceeding with these multiplicities, we give an example.

Example 2.4.1. Consider the irreducible module of $\mathfrak{sl}_3$ with highest weight $\mu = 3\Lambda_1 + 2\Lambda_2$. Here the root system is two-dimensional so we may draw the weight lattice in Figure 2.8. Each slide demonstrates a stage of the calculation. In the first slide we have drawn the highest weight μ.

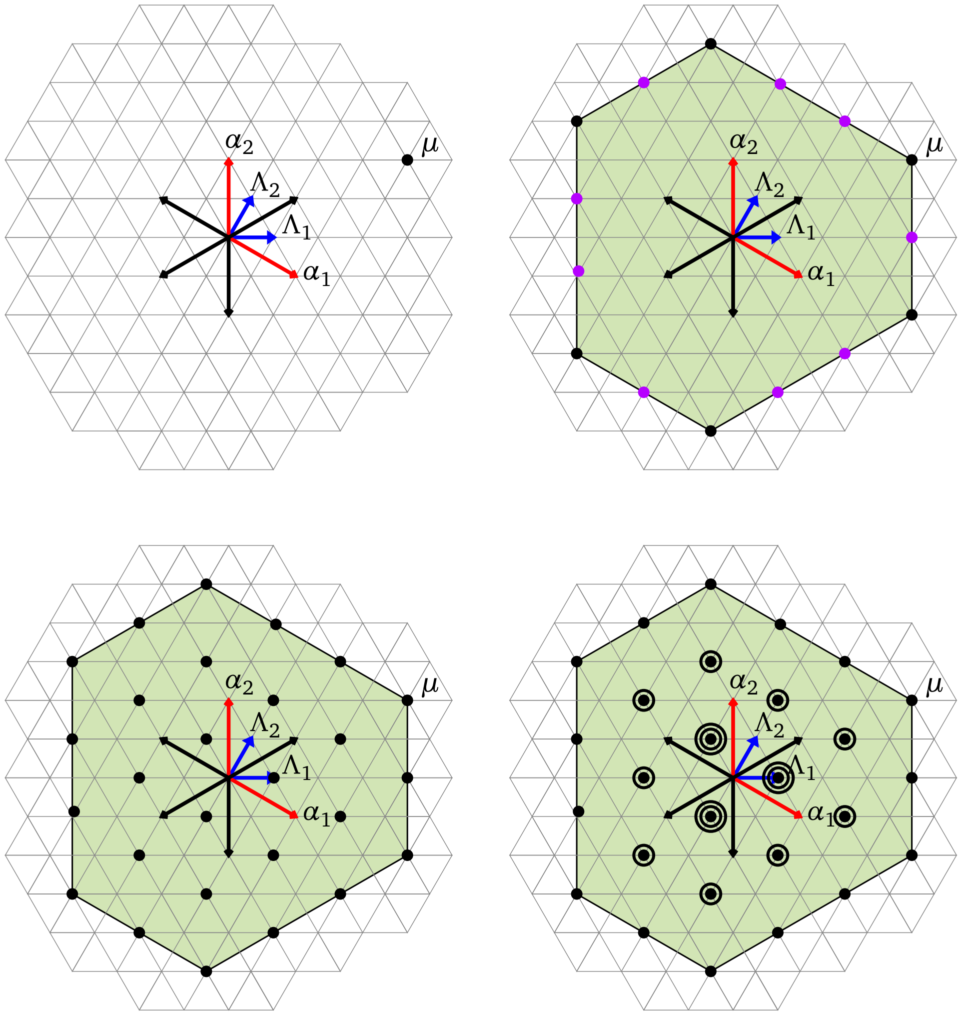

Figure 2.8: The highest weight module V_μ.

Then we consider the orbit of μ under the Weyl group. This orbit is shown in black in the second slide with the intermediate weights that complete the shell shown in purple. To complete the full set of weights we then add all points obtained by root subtractions (slide 3). Finally, although we do not yet in this exposition have a method for finding these, the weight multiplicities are indicated in the last slide. In these weight diagrams circles around the individual weights are used to indicate their multiplicity.

There are various ways to determine the weight multiplicities. Given a complete description of the modules corresponding to the fundamental weights and a technique for finding tensor product decompositions (Section 2.6) , it is possible to recursively describe the weight multiplicities.

When decomposing the tensor product of V_μ and V_ν, the highest weight module $V_{\mu+\nu}$ will always occur. The direct sum of two modules presents itself in the weight lattice as an overlapping (with multiplicity) of the weights in the modules. Hence by subtracting the weights of each of the smaller representations that occur in the tensor product $V_\mu \otimes V_\nu$, we will eventually be left with just the weights of $V_{\mu+\nu}$.

This process has been abstracted to create Freudenthal's formula [6]. As an alternative to this procedure, the Weyl character formula of Section 2.5 may also be used to determine the weight multiplicities.

Weights in a tensor product

Given vector spaces V and W, the tensor product $V \otimes W$ is spanned by the vectors $v_i \otimes w_j$ where v_i and w_j range over the basis vectors of V and W respectively. Suppose that v is an eigenvector of V with eigenvalue λ_v and similarly w is an eigenvector of W with eigenvalue λ_w. Then $v \otimes w$ will be an eigenvector of $V \otimes W$ with corresponding eigenvalue $\lambda_v + \lambda_w$. Thus the weights of the tensor product of two module will be the sum set of their individual weights.

For example in A_2, the modules V_{Λ_1} and V_{Λ_2} have weights

$$\{\Lambda_1, \Lambda_2 - \Lambda_1, -\Lambda_2\} \text{ and } \{\Lambda_2, \Lambda_1 - \Lambda_2, -\Lambda_1\}.$$

So their tensor product $V_{\Lambda_1} \otimes V_{\Lambda_2}$ will have weights

$$\{\Lambda_1 + \Lambda_2, 2\Lambda_1 - \Lambda_2, 0, 2\Lambda_2 - \Lambda_1, 0, \Lambda_2 - 2\Lambda_1, 0, \Lambda_1 - 2\Lambda_2, -\Lambda_1 - \Lambda_2\}.$$

This tells us that in the direct sum decomposition of $V_\mu \otimes V_\nu$, the highest weight module $V_{\mu+\nu}$ will always occur. Furthermore it can be shown that any highest weight module which occurs will be labelled by some λ of the form,

$$\lambda = \mu + \nu - \sum_{i=1}^{n} k_i \alpha_i.$$

2.5 Weyl characters

A highest weight module contains a set of weights. As seen in Section 2.4, these weights give the structure of the module. The formal Weyl character, χ, presents this information as a function attached to each module. The Weyl character is defined by,

$$\chi(V_\lambda) = \sum_\mu K_{\lambda\mu} e^\mu$$

where the sum ranges over all the weights μ of the highest weight module V_λ. Here $K_{\lambda\mu}$ is the multiplicity of the weight μ in the irreducible module of highest weight λ.

A useful property of the Weyl character is that direct sums and tensor products of modules correspond to sums and products of characters.

$$\chi_{V_\mu \oplus V_\nu} = \chi_{V_\mu} + \chi_{V_\nu}$$
$$\chi_{V_\mu \otimes V_\nu} = \chi_{V_\mu} \cdot \chi_{V_\nu}$$

One disadvantage of the Weyl characters is that finding them requires summing over the weights of a representation with multiplicity. Calculating these is not always an easy task. Fortunately there exists a remarkable formula that circumvents this, the Weyl character formula. This formula uses explicitly only the highest weight vector and the Weyl group of the Lie algebra. The function $l(w)$ is used to denote the smallest length of an element of the Weyl group expressed as a product of the simple reflections:

$$\chi(V_\lambda) = \frac{\sum\limits_{w \in W} (-1)^{l(w)} e^{w(\lambda+\rho)-\rho}}{\prod\limits_{\alpha>0} (1 - e^{-\alpha})}.$$

As an example, the characters for modules of $\mathfrak{sl}_n$ are the Schur functions. These are described with greater detail in Section 2.6 when we introduce the Littlewood–Richardson rule.

2.6 Tensor product algorithms

One of the most fundamental tools used in this research is the decomposition of the tensor product of two highest weight modules into a direct sum of irreducible modules. Two algorithms provide neat methods to do this. The first, Klimyk's rule, is applicable to all root systems and is described geometrically in the weight lattice. For $\mathfrak{sl}_n$ there is a very elegant description called the Littlewood–Richardson rule that uses tableaux instead of weights.

Klimyk's rule

An explicit formula for determining the direct sum decomposition of two highest weight modules is given by Klimyk's rule. This rule requires interpreting representations labelled by non-dominant weights. We define a new action of the Weyl group upon the weight lattice by $w \cdot \lambda = w(\lambda + \rho) - \rho$ for each element w of the Weyl group. This action corresponds to the typical reflections of the Weyl group centered on the negative Weyl vector. In such a scenario the Weyl chambers and walls are redefined to match this new action.

Representations labelled by a non-dominant weight are interpreted with the shifted action describe above. For any weight μ, either μ is in the wall of a Weyl chamber under this action or there is a unique element w in the Weyl group such that $w \cdot \mu$ is in the dominant chamber. In the first case, $V_\mu = 0$ and in the second $V_\mu = (-1)^{l(w)} V(w \cdot \mu)$.

With these two concepts, Klimyk's rule may be expressed as,

$$V_\mu \otimes V_\nu = \bigoplus_{\nu'} K_{\nu\nu'} V(\mu + \nu')$$

where ν' ranges over the weights of ν and again $K_{\nu\nu'}$ is the multiplicity of ν' in ν.

Example 2.6.1. For these examples it is easiest to study representations of $\mathfrak{sl}_3$. Consider the module $V_\nu \otimes V_\nu$ for $\nu = \Lambda_1 + \Lambda_2$. The weights of ν, over which ν' ranges are shown on the weight lattice in Figure 2.9. Adding ν to each of the

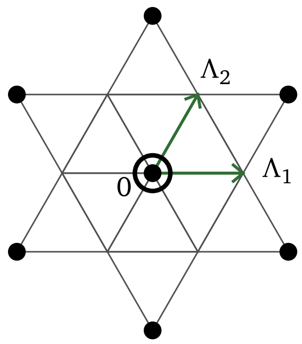

Figure 2.9: Weights of the module in example 1.

weights results in the points shown in Figure 2.10. Note that lines are drawn through $-\rho$ parallel to the fundamental weights. The action of the Weyl group described above corresponds to reflection about these lines. In this example two of the weights occur on the walls of the chamber and hence do not contribute while all others are in the positive chamber. Thus the decomposition is

$$V_{\Lambda_1+\Lambda_2} \otimes V_{\Lambda_1+\Lambda_2} = V_0 \oplus 2V_{\Lambda_1+\Lambda_2} \oplus V_{3\Lambda_1} \oplus V_{3\Lambda_2} \oplus V_{2\Lambda_1+2\Lambda_2}.$$

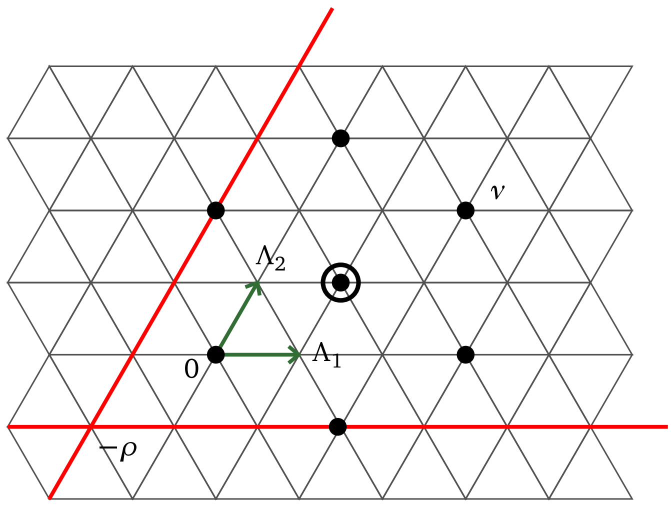

Figure 2.10: Diagram depicting Klimyk's rule.

Example 2.6.2. Let again $\mathfrak{g} = \mathfrak{sl}_3$ and take the tensor product $V_\mu \otimes V_\nu$ where $\mu = \Lambda_1 + 2\Lambda_2$ and $\nu = 2\Lambda_1 + \Lambda_2$. The weights of ν are shown on the lattice

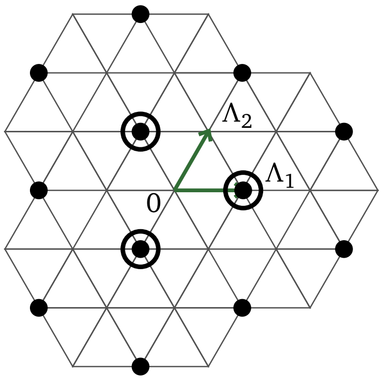

Figure 2.11: Weights of the module in example 2.

in Figure 2.11. Adding μ to each of the weights results in the points shown in Figure 2.12. Observe here that three of the weights occur in the walls of the chamber and hence do not contribute. There is a single weight outside the positive chamber. Acting on this weight by σ_1 with length one places it in the dominant chamber. This weight then cancels with one of the two weights already there. Thus the decomposition is

$$V_\mu \otimes V_\nu = V_0 \oplus 2V_{\Lambda_1+\Lambda_2} \oplus V_{3\Lambda_1} \oplus V_{3\Lambda_2} \oplus 2V_{2\Lambda_1+2\Lambda_2} \oplus V_{4\Lambda_1+\Lambda_2} \oplus V_{\Lambda_1+4\Lambda_2} \oplus V_{\mu+\nu}.$$

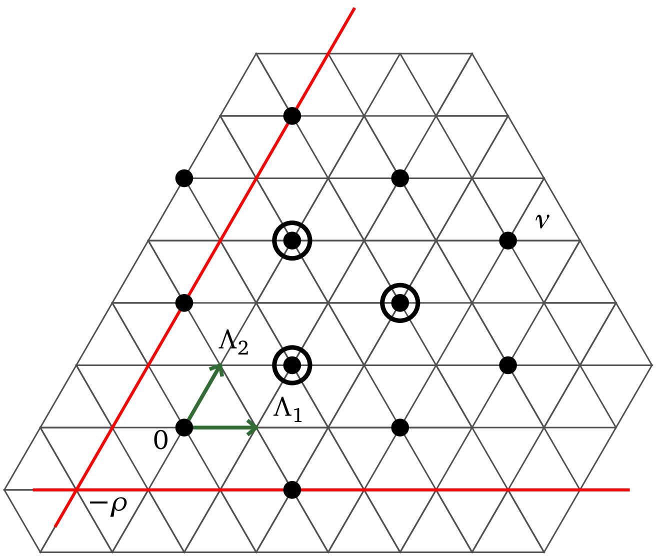

Figure 2.12: Diagram depicting Klimyk's rule.

Littlewood–Richardson rule

For representations of the Lie algebra $\mathfrak{sl}_n$, the tensor product decomposition has a particularly elegant description given by the Littlewood–Richardson rule. This rule uses representations labelled by partitions rather than weights. The notation between the two is identical but hopefully clear in context.

A partition $\lambda = (\lambda_1, \lambda_2, \dots)$ is a weakly deceasing sequence of nonnegative integers such that only finitely many λ_i are strictly positive. It is helpful to identify a partition with its diagram given by $\{(x, y) \in \mathbb{Z}^2 : x \leq \lambda_y\}$. We think of this as stacking rows of boxes of length λ_i and traditionally the diagram is drawn upside-down. For example, Figure 2.13 (left) corresponds to the partition $(6, 5, 5, 3, 1)$.

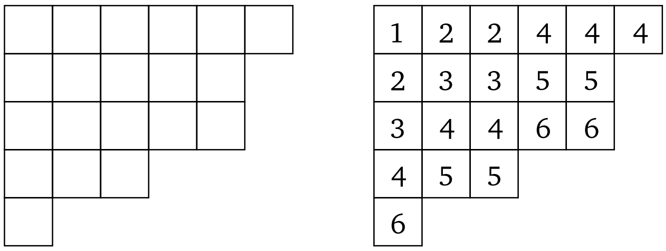

Figure 2.13: Examples of a partition and a semistandard tableau.

A tableau is a partition together with a filling of the boxes with positive integers. When using these objects to label representations of $\mathfrak{sl}_n$ the boxes may only be filled with the integers $1, \dots, n$. A tableau is called semistandard if the contents are weakly increasing across the rows and strictly increasing down the columns. Figure 2.13 (right) is a semistandard tableau. The reverse lattice word condition is the requirement that, as a tableau is read right to left, top to bottom, every partial word formed contains at least as many i as $i + 1$ for all i.

Given two partitions λ and μ we say that μ is contained in λ if $\mu_i \leq \lambda_i$ for all i. This is notated by $\mu \subseteq \lambda$. Assuming $\mu \subseteq \lambda$, the skew tableau λ/μ is given by the set subtraction of the diagram of μ from that of λ.

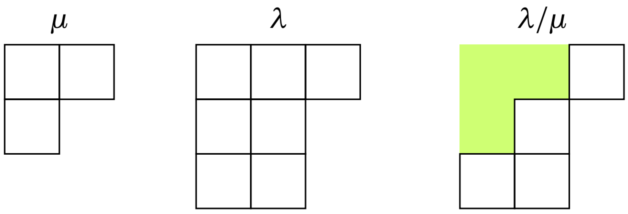

Figure 2.14: An example of a skew partition.

For a tableau, we define the contents to be $1^{n(1)}2^{n(2)}3^{n(3)}\dots$ where $n(i)$ is the number of i's in the filling.

Irreducible representations of $\mathfrak{sl}_n$ are commonly labelled by partitions through the use of Schur functions. These Schur functions are the specialisation of the Weyl character for $\mathfrak{sl}_n$. Before exploring the Schur functions more, we should state the Littlewood–Richardson rule.

If s_μ and s_ν are Schur functions labelled by partitions μ and ν, then the product $s_\mu \cdot s_\nu$ is given by

$$s_\mu \cdot s_\nu = \sum_\lambda c^\lambda_{\mu\nu} s_\lambda$$

such that $\mu \subseteq \lambda$ and hence λ/μ is a skew partition. The coefficient $c^\lambda_{\mu\nu}$ is the number of semistandard fillings of λ/μ which satisfy the reverse lattice word condition and the contents of the filling correspond to ν. By this we mean $n(i) = \nu_i$. Since the product of Weyl characters corresponds to the tensor product of representations, this rule gives the tensor product decomposition for highest weight modules of $\mathfrak{sl}_n$.

The Schur functions are symmetric homogeneous polynomials which form a basis for the ring of symmetric polynomials. They are usually defined with the following determinant formula,

$$s_\lambda = \frac{\det\limits_{1\le i,j\le n}\left(x_i^{\lambda_j+n-j}\right)}{\prod\limits_{1\le i<j\le n}(x_i - x_j)}.$$

Here λ is a partition. If we write $x_i = e^{\varepsilon_i}$ and assume $x_1 \dots x_n = 1$, this description of the Schur functions matches the Weyl character formula.

There is a simple bijection between the labelling of representations by partitions and the labelling by highest dominant weights. Suppose that s_λ is a Schur function labelled by the partition λ. Then the weight corresponding to the partition λ is $(\lambda_1 - \lambda_2, \lambda_2 - \lambda_3, \dots)$. This gives the coordinates of the highest weight vector.

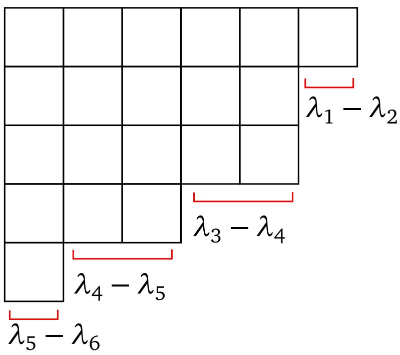

Figure 2.15: The bijection between partitions and weights.

An alternative combinatorial definition of the Schur functions may be given by a sum over tableaux. For a semistandard tableau T of shape λ, form the monomial

$$x^T = \prod_{i=1}^{n} x_i^{n(i)}.$$

Then the Schur function is the sum of these monomials over all semistandard tableaux of shape λ,

$$s_\lambda = \sum_T x^T.$$

As an example, the Schur function in three variables labelled by the partition $(2,1)$ is given by summing over the tableaux in Figure 2.16. This gives $s^3_{(2,1)} = x_1^2x_2 + x_1^2x_3 + x_1x_2^2 + 2(x_1x_2x_3) + x_1x_3^2 + x_2^2x_3 + x_2x_3^2$.

$$\begin{ytableau} 1 & 1 \\ 2 \end{ytableau} \quad \begin{ytableau} 1 & 1 \\ 3 \end{ytableau} \quad \begin{ytableau} 1 & 2 \\ 2 \end{ytableau} \quad \begin{ytableau} 1 & 2 \\ 3 \end{ytableau} \quad \begin{ytableau} 1 & 3 \\ 2 \end{ytableau} \quad \begin{ytableau} 1 & 3 \\ 3 \end{ytableau} \quad \begin{ytableau} 2 & 2 \\ 3 \end{ytableau} \quad \begin{ytableau} 2 & 3 \\ 3 \end{ytableau}$$

Figure 2.16: An example of a schur function.

When we use partitions to label representations of $\mathfrak{sl}_n$ we have to restrict the number of rows. If the partition has more than n rows it can have no semistandard filling with the numbers $1, \ldots, n$. Hence such partitions cannot label representations of $\mathfrak{sl}_n$ and so we ignore them in the Littlewood–Richardson rule. Additionally if a column has height exactly n then there is only a single possible way to fill this with the numbers $1, \ldots, n$ that is strictly increasing. So these columns do not contain any additional information and hence we ignore them also. This is equivalent to making the identification $x_1 \ldots x_n = 1$ in the corresponding Schur function. Since the Schur functions are homogeneous, if we make this identification no information is lost.

The Littlewood–Richardson rule completely describes the tensor product multiplicities and hence completely describes the tensor product decomposition. In our work we often want to view this rule as an algorithm for giving the direct sum decomposition. In this case the rule is interpreted as stacking boxes with content ν on top of the partition μ. Then the decomposition is simply given by all such stackings that are semistandard and satisfy the reverse lattice word condition.

Example 2.6.3. To demonstrate the Littlewood–Richardson rule, we echo the examples from the previous section on Klimyk's formula. The weight $\Lambda_1 + \Lambda_2$ corresponds to the partition $\mu = \nu = (2,1)$. Thus we must find all valid skew tableau λ/μ which have contents $1^2 2$. These are given in Figure 2.17.

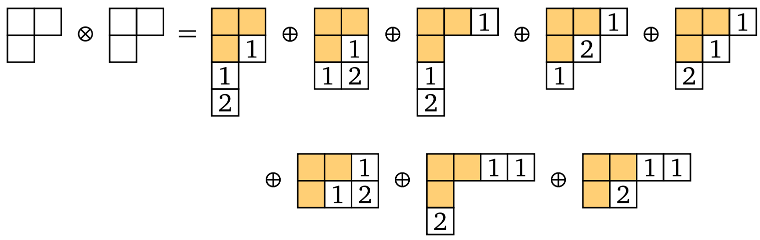

Figure 2.17: A first example of the Littlewood–Richardson rule.

If we restrict to $\mathfrak{sl}_3$ then the decomposition in this example is exactly that obtained through Klimyk's formula.

Example 2.6.4. The second example from before may also be repeated. The weight $\mu = \Lambda_1 + 2\Lambda_2$ corresponds to the partition $(3,2)$ and $\nu = 2\Lambda_1 + \Lambda_2$ corresponds to $(3,1)$. Thus we must find all valid skew tableau λ/μ which have contents $1^3 2$. These are given in Figure 2.18.

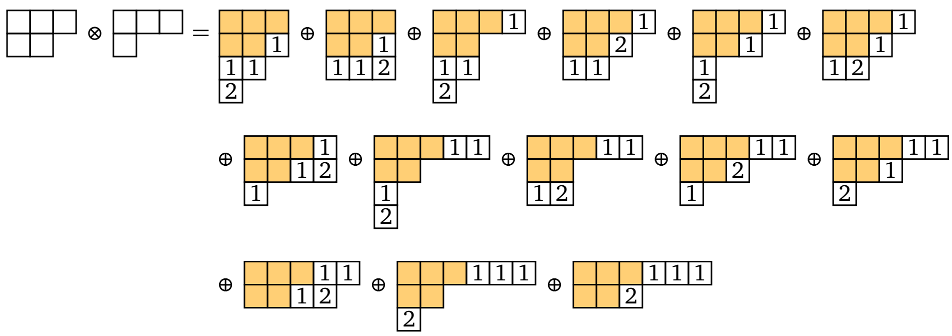

Figure 2.18: A second example of the Littlewood–Richardson rule.

If we restrict to $\mathfrak{sl}_3$ then the decomposition in this example is exactly that obtained through Klimyk's formula.

Research problem 3

The central question studied in this thesis is about evaluations of integrals defined for highest weights in a tensor product decomposition. In particular Mukhin and Varchenko made in their paper [11] three conjectures about a family of integrals. These conjectures remain open. Here we state the integral, the first conjecture and provide a brief overview of current progress on the problem.

3.1 The phase integral

Consider two highest weight modules of a simple Lie algebra $\mathfrak{g}$, labelled by μ and ν. Suppose that $\mu+\nu-\sum_{i=1}^{n} k_i\alpha_i$ is a particular highest weight in the direct sum decomposition of the tensor product $V_\mu \otimes V_\lambda$. The phase integral is obtained by attaching k_i integration variables to each simple root α_i and has the form,

$$\int \left[\prod_{i=1}^{k} t_i^{-(\mu,\alpha_{t_i})}(1-t_i)^{-(\nu,\alpha_{t_i})} \prod_{1\le i<j\le k} |t_i-t_j|^{(\alpha_{t_i},\alpha_{t_j})} \right]^{\gamma} dt_1 \dots dt_k.$$

Here t_i are the integration variables, α_{t_i} is the root to which t_i is attached and γ is an arbitrary complex number subject to convergence conditions. It is convenient to think of this attachment of integration variables in the Dynkin diagram.

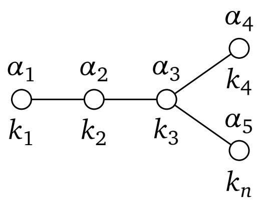

Figure 3.1: The Dynkin diagram of D_5 with k_i paired to each simple root.

We have not specified the domain of integration as it is not known in general. Identifying this domain is one of the key requirements for the evaluation, although not one which we studied extensively. The original conjecture had a constraint on this domain but it has since been shown that that was too restrictive.

3.2 Conjectures about the phase integral

Mukhin and Varchenko [11] made three conjectures about the phase integral. We are most interested in the first of these,

Conjecture 1. If the weight $\mu + \nu - \sum_{i=1}^{n} k_i \alpha_i$ occurs with multiplicity one in the tensor product decomposition then the phase integral can be evaluated as a ratio of gamma functions.

3.3 The gamma function

We are interested in evaluating integrals as a ratio of Euler gamma functions. The gamma function is a continuous and convex extension of the integer factorial to the complex numbers. It is defined by,

$$\Gamma(z+1) = \int_0^\infty t^z e^{-t}\, dt \qquad \mathrm{Re}(z) > -1.$$

For this function we have,

$$\Gamma(1) = \int_0^\infty e^{-t}\, dt = 1$$

and,

$$\begin{aligned} \Gamma(z+1) &= \int_0^\infty t^z e^{-t}\, dt \\ &= z \int t^{z-1} e^{-t}\, dt \\ &= z\Gamma(z). \end{aligned}$$

That is, the gamma function satisfies the recurrence relation $\Gamma(z+1) = z\Gamma(z)$ and $\Gamma(1) = 1$. On integers this then gives $\Gamma(n+1) = n!$.

3.4 Euler beta integral

When the phase integral from before has a single integration variable it reduces to the Euler beta integral.

$$\int_0^1 t^{\alpha-1}(1-t)^{\beta-1}\, dt = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)} \qquad \mathrm{Re}(\alpha), \mathrm{Re}(\beta) > 0.$$

This integral indeed has an evaluation as a ratio of gamma functions. A proof of this result follows.

$$\Gamma(\alpha)\Gamma(\beta) = \int_0^\infty \int_0^\infty e^{-u-v} e^{\alpha-1} v^{\beta-1}\, du dv.$$

Introducing the new integration variables z and t with $u = zt, v = z(A-t)$ and Jacobian determinant Az results in,

$$\begin{aligned}\Gamma(\alpha)\Gamma(\beta) &= \int_0^\infty \int_0^A e^{-zt-z(A-t)}(zt)^{\alpha-1}(z(A-t))^{\beta-1}Az\, dtdz \\ &= A\int_0^\infty e^{-Az}z^{\alpha+\beta-1}\, dz \int_0^A t^{\alpha-1}(A-t)^{\beta-1}\, dt \\ &= A^{1-\alpha-\beta}\Gamma(\alpha+\beta)\int_0^A t^{\alpha-1}(A-t)^{\beta-1}\, dt.\end{aligned}$$

Setting $A = 1$ we arrive at,

$$\int_0^1 t^{\alpha-1}(1-t)^{\beta-1}\, dt = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}.$$

The use of A here instead of simply starting with the interval $[0,1]$ is to foreshadow a technique used in Mukhin and Varchenko's paper.

3.5 Selberg's integral

The second more complicated example that we give in this section is the Selberg integral. When k is one, this is precisely the Euler beta integral from before. A proof of this integral may be found in Selberg's original paper [12] or with a different approach in work by Anderson [1].

$$\int_{0<t_1<\cdots<t_k<1} \prod_{i=1}^k t_i^{\alpha-1}(1-t_i)^{\beta-1} \prod_{1\le i<j\le k} |t_i - t_j|^{2\gamma}\, dt_1 \dots dt_k = \prod_{i=0}^{k-1} \frac{\Gamma(\alpha+i\gamma)\Gamma(\beta+i\gamma)\Gamma((i+1)\gamma)}{\Gamma(\alpha+\beta+(i+k-1)\gamma)\Gamma(\gamma)}$$

3.6 State of the conjecture

This section gives an overview of progress so far on the conjecture. Remarkably-few cases of the conjecture have been evaluated and even the simplest of these is difficult.

Original evidence

In the same paper where Mukhin and Varchenko made their conjecture they observed that, for the simple Lie algebra $\mathfrak{sl}_2$, the phase integral corresponds to Selberg's integral. As seen earlier this integral evaluates as a ratio of gamma functions and supports the conjecture. Note that every tensor product of irreducible modules in $\mathfrak{sl}_2$ is multiplicity free. This is apparent from the Littlewood–Richardson rule.

The authors then provided an evaluation of the integral for tensor products with the standard representation of $\mathfrak{sl}_n$. Any irreducible module that occurs in the tensor product decomposition of $V_\lambda \otimes V_{\Lambda_1}$ will occur without multiplicity. To see this, consider the Littlewood–Richardson rule when one partition is a single box. Their method was to integrate each variable with the Euler beta integral and develop a recursive formula for the evaluation.

Over the integration domain $0 < t_1 < t_2 < \cdots < t_k < 1$ the integral for this case has the form,

$$\int_{0<t_1<\cdots<t_k<1} t_1^{\alpha_1}(1-t_1)^{\beta_1} \prod_{j=2}^{n} t_j^{\alpha_j}(t_j - t_{j-1})^{\beta_j}\, dt_1 \ldots dt_k.$$

Then using the evaluation of the beta function with $A = t_{j-1}$ we may recursively integrate each variable, starting from $j = n$ and decreasing.

Each result in the following sections may be though of as a generalisation of one of Mukhin and Varchenko's original twoevaluations.

Extensions using Macdonald polynomials

Using the theory of Macdonald polynomials, Warnaar [17,18] was able to evaluate the phase integral for A_{n-1} in more general situations.

The Pieri rule is a method for computing the decomposition of tensor products with powers of the standard representation of $\mathfrak{sl}_n$. In the language of Littlewood–Richardson this may be stated as a tensor product with a single row of any length. Elements that occur in the decomposition are all the skew tableau which are horizontal strips. By this we mean that no two boxes share a column.

The first of Warnaar's results was to evaluate the phase integral for the Pieri rule. That is, for any n, the integral corresponding to a module in the decomposition of $V_\lambda \otimes V_{m\Lambda_1}$. Here he was able to give an explicit description of the k_i values which could occur in these decompositions. This allowed the integral to be evaluated in each of these cases.

His second major result was a complete evaluation of an integral for some more general representations in the Lie algebra $\mathfrak{sl}_3$. While this integral is extremely similar to the phase integral, the exact modules and tensor products to which it corresponds are unknown.

The standard representation in other root systems

Mimachi and Takamura [10] evaluated a particular phase integral for each of the root systems $\mathrm{B}_n, \mathrm{C}_n$ and D_n. The integral corresponded to the trivial representation occurring in the direct sum decomposition of the tensor square of the standard representation, $V_{\Lambda_1} \otimes V_{\Lambda_1}$. They obtained these results by extracting a two dimensional Selberg integral and using this to express the formula recursively.

Further results 4

This chapter contains various attempts to progress the state of the Mukhin–Varchenko conjecture.

4.1 Multiplicity-free tensor products

The Mukhin–Varchenko conjecture [11] is concerned with triples of weights μ, ν and λ with $\lambda = \mu + \nu - \sum k_i \alpha_i$ under the further condition that λ must occur with multiplicity one in the tensor product decomposition of $V_\mu \otimes V_\nu$. A more restrictive question is to ask for pairs μ and ν for which *every* irreducible module in this decomposition has multiplicity one. These pairs are called multiplicity-free tensor products and have been completely classified by Stembridge in two papers [14,15]. However Stembridge's papers do not contain an explicit description of the weights λ which occur in these multiplicity-free tensor products. These weights would give the values k_i.

To state the phase integral, let alone evaluate it, requires knowledge of the k_i values. With this in mind we would like to find explicit descriptions of these k_i values in situations where, according to the conjecture, the integral should have a gamma function evaluation. An excellent place to start is with the multiplicity-free tensor products.

4.2 Multiplicity-free tensor products in A_{n-1}

With the root system A_{n-1} we have the advantage of the Littlewood–Richardson rule to aid in finding the k_i values. Unfortunately while the rule itself gives the tensor product decomposition, it does not explicitly provide k_i values. To determine these we first reinterpret root subtraction as a shifting of boxes in partitions.

The multiplicity-free tensor products of partitions identified by Stembridge are,

- A row or column with any partition (Pieri's rule).
- Two rectangles.
- A fat hook with a two-line rectangle.

◇ A near rectangle with a rectangle.

Rows, columns and rectangles are partitions of those shapes. A fat hook is a partition with exactly two different row lengths (or equivalently column heights). A near rectangle is a partition from which the removal of one row or column results in a rectangle. We will describe the k_i values corresponding to all but the last of these cases.

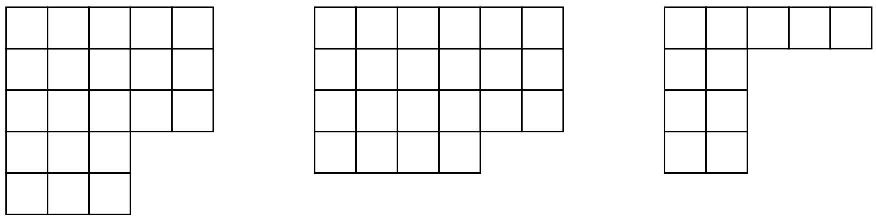

Figure 4.1: A fat hook and two near rectangles.

The challenge in describing the k_i values is twofold. First we describe the possible box shifts in the partition diagram that match the Littlewood–Richardson rule. This corresponds to the tensor product of the representations that the partitions label (Section 2.6). To apply these inequalities to representations of $\mathfrak{sl}_n$ we must then restrict the resulting partitions to height n and make the necessary corrections in the k_i values.

Root subtraction and box shifting

To assist in the description of k_i values for the multiplicity-free A_{n-1} tensor products, we reinterpret root subtraction as box shifting. This allows us to use the elegant description of the tensor product decomposition given by the Littlewood–Richardson rule.

Suppose that a box in a partition is moved down from one row to the next. This requires that there is a box to shift and a place to put the box that results in a valid partition. Moving a box from row i to row $i+1$ adds one to the lengths of the protrusions in row $i-1$ and row $i+1$. It also subtracts two from the length of the protrusion in row i. This is illustrated in Figure 4.2.

Using the bijection between partitions and weights this impact is described by the following transformations,

$$\begin{aligned} \mu_{i-1} &\mapsto \mu_{i-1} + 1 \\ \mu_i &\mapsto \mu_i - 2 \\ \mu_{i+1} &\mapsto \mu_{i+1} + 1. \end{aligned}$$

In the extreme cases, the transformations are slightly different as μ_0 and μ_{n+1} are not defined. Simply remove these and the interpretation is consistent. This

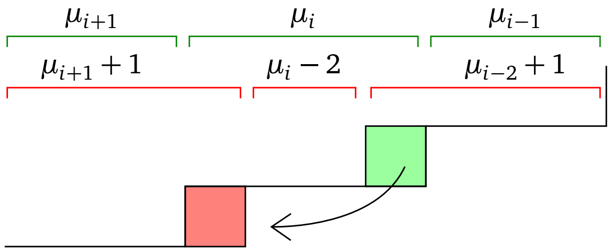

Figure 4.2: The effect of box shifting on the weight bijection.

removal matches the difference in the first and last rows of the Cartan matrix for A_{n-1}.

This transformation of weights corresponds precisely to the subtraction of the root α_i in the weight lattice. To state this explicitly, translation of a box from row i to row $i+1$ in a partition is equivalent to subtracting the root α_i from the corresponding weight. Similarly, movement from row i to row j with $i<j$ is the subtraction of the (not necessarily simple) root $(\alpha_i+\cdots+\alpha_{j-1})$.

It is worth noting that the root subtractions that leave the dominant chamber are precisely those box shifts that would not result in a valid partition.

Tensor products and box shifting

As seen earlier, in the language of weights the tensor product $V_\mu \otimes V_\nu$ decomposes into elements of the form V_λ with $\lambda = \mu+\nu-\sum_{i=1}^{n} k_i\alpha_i$. We are interested in describing these k_i values. The weight $\mu+\nu$ corresponds to the pairwise sum of μ and ν as partitions. Such a sum is demonstrated in Figure 4.3.

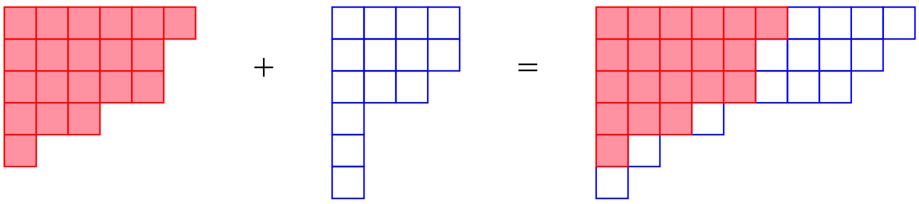

Figure 4.3: The sum of two partitions.

With the reinterpretation of root subtraction as box shifting, finding the k_i values becomes a problem of determining which box shifts from $\mu+\nu$ result in valid semistandard skew tableau that satisfy the reverse lattice word condition. An example of this for two small partitions is given in Figure 4.4.

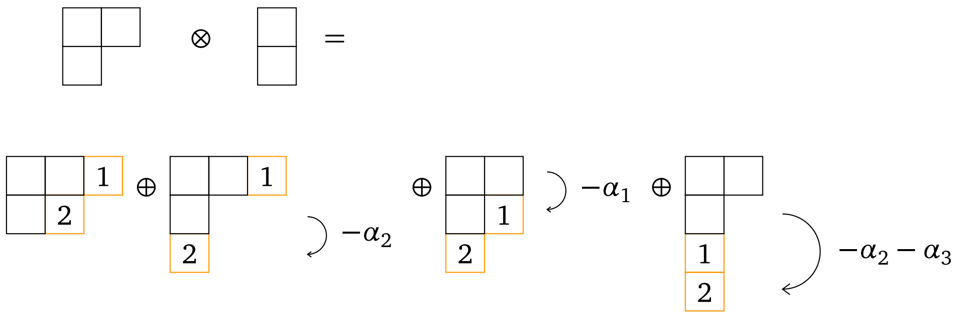

Figure 4.4: How k_i values occur in the tensor product of partitions.

Tensor product with a single row

This product is exactly the Pieri rule and the phase integral has been evaluated in this scenario [18]. Nonetheless, this is also the simplest case of a multiplicity-free tensor product and provides an opportunity to demonstrate the common techniques of this section.

Consider shifting boxes from $\mu + \nu$ where ν is (ν_1), a single row. Certainly no more boxes may be shifted than are in ν and by the semistandard condition no more boxes may be shifted than the width of μ. Thus $k_1 \leq \min\{\mu_1, \nu_1\}$. The number of boxes deposited in a row, say row i, is given by the difference between the number shifted in and the number shifted out, $k_i - k_{i+1}$. This difference may be no more that the number of available places in this row, $\mu_i - \mu_{i+1}$, otherwise the result either voids semistandardness or is not a tableau. This gives the set of inequalities,

$$\begin{gathered} k_1 \leq \min\{\mu_1, \nu_1\} \\ 0 \leq k_i - k_{i+1} \leq \mu_i - \mu_{i+1}. \end{gathered}$$

Note that the μ_i here label partitions. The value $\mu_i - \mu_{i+1}$ is then precisely the coefficient of Λ_i in the corresponding weight.

The restriction to $\mathfrak{sl}_n$ may be made by setting $k_n = 0$.

Tensor product with a single column

Let μ be an arbitrary partition and ν be a single column of height t, $\nu = (1^t)$. Then from the semistandard and reverse-lattice-word conditions λ will occur in the decomposition only if λ/μ is a vertical strip. A vertical strip is a skew tableau in which no two boxes share a row.

The construction here is more difficult to describe than the product with a row. This is because the bijection between partitions and weights is inherently connected to the horizontal structure of the partition, rather than the vertical.

To describe the k_i values, first consider moving all the boxes into a formal 'row zero'. Then move the boxes down the tableau accounting for the locations where a box may potentially be deposited.

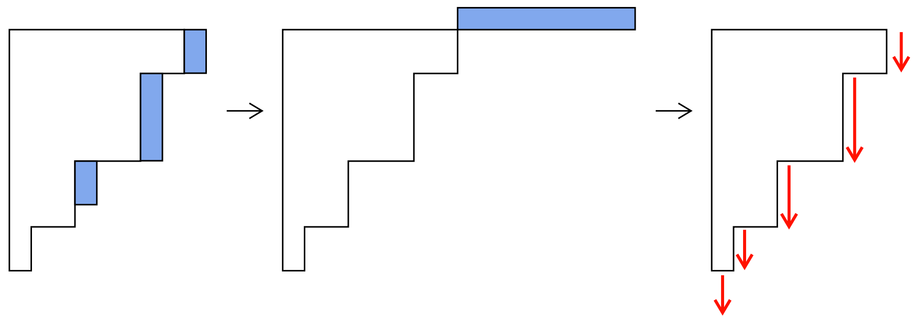

Figure 4.5: Describing the tensor product of a partition with a single column.

We use the values h_i to list the number of boxes *not* placed by row i. To match the artificial row, set $h_0 = t$. Then a box may be deposited at a row if either the partition has a shelf or a box had previously been deposited in the row below. Furthermore only one box may be deposited in each row. This gives the inequality $0 \leq h_i - h_{i+1} \leq \min\{1, (k_{i-1} - k_i) + (\mu_i - \mu_{i+1})\}$.

To convert from the h values to the desired k values we must account for the shift into the formal row zero. This is done by adding $i - t$ to each h_i.

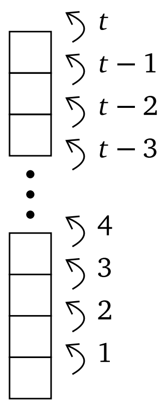

Figure 4.6: Moving the column to the artificial 'row 0'.

Together these two processes give the complete set of inequalities,

$$
\begin{gathered}
h_0 = t, h_{l(\mu)+t} = 0 \\
0 \leq h_i - h_{i+1} \leq \min\{1, (k_{i-1} - k_i) + (\mu_i - \mu_{i+1})\} \\
k_i = h_i - t + i \text{ for } 0 \leq i \leq t \\
k_i = h_i \text{ for } i > t.
\end{gathered}
$$

To restrict this description, simply remove all sets of k_i values for which $k_n \geq 0$.

Tensor product of two rectangles

Consider the tensor product of the representations labelled by $\mu = p\Lambda_r$ and $\nu = q\Lambda_s$. This corresponds to the product of two partitions which are rectangles of widths p and q, and heights r and s respectively. We may assume by the symmetry of the Littlewood–Richardson coefficients that $s \leq r$.

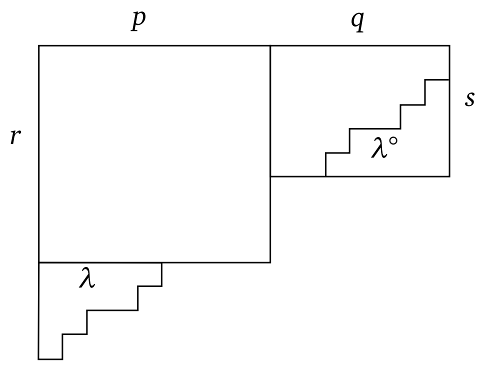

Figure 4.7: A diagram depicting $\mu + \nu$ in the product of two rectangles.

Firstly we describe the partitions λ which may occur. Any resulting partition is formed by removing a 180° rotated partition from the bottom right corner of ν in $\mu + \nu$ and placing it below μ.

This follows since the semistandard condition forces any resulting skew tableau to be strictly increasing vertically. As such no boxes may be shifted to the space below ν. Furthermore the remaining unshifted boxes are still a partition, so the section removed from ν is in the shape of a rotated tableau. The question is then; how these may be placed above μ?

Since this is a multiplicity-free product it should be no surprise that there is a unique way to do this. The reverse lattice word condition implies that each element in the removed section must be read after the box directly above it. Working through the columns of the removed section right to left and moving them to the top of μ gives the unique placement. This is illustrated in Figure 4.8.

Next we use the weight lattice interpretation to develop k_i values from this combinatorial description. To do this, choose the partition λ to be removed from ν ensuring that $\lambda_i \leq \min\{p, q\}$. As demonstrated in Figure 4.9, consider the total movement of boxes as the individual shifting of each row from $\mu + \nu$ to a row of same width in the final partition.

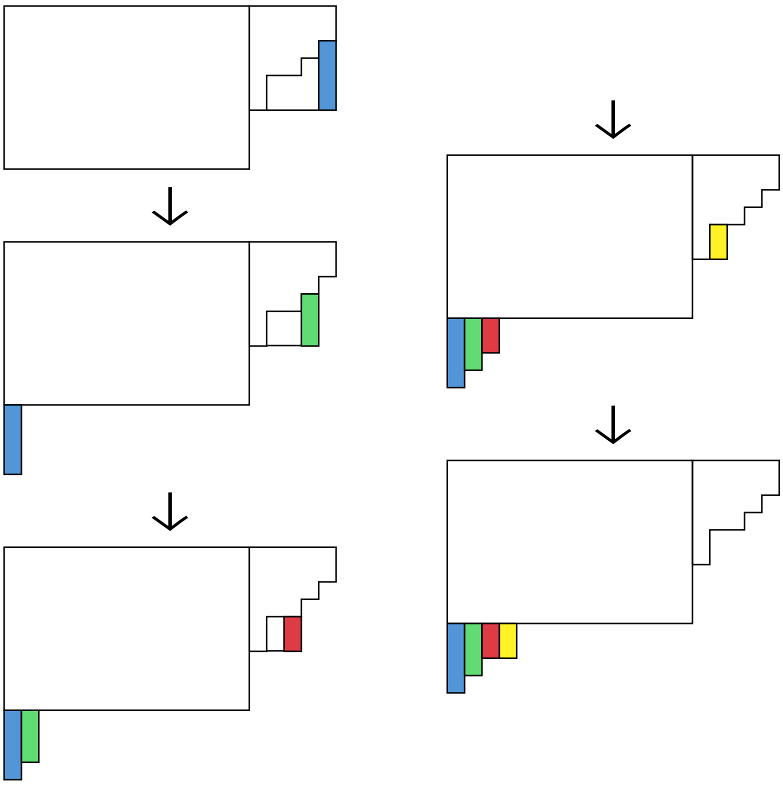

Figure 4.8: Positioning the cut out section of ν.

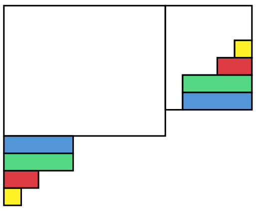

Figure 4.9: Determining the k_i values.

In terms of individual shifts this gives,

$$\begin{gathered}\lambda_1(\alpha_s + \cdots + \alpha_r)\\ \lambda_2(\alpha_{s-1} + \cdots + \alpha_{r+1})\\ \vdots\\ \lambda_s(\alpha_1 + \cdots + \alpha_{r+s-1})\end{gathered}$$

Collecting like terms we find the coefficients for the simple roots.

$$\begin{gathered}k_1 = k_{r+s-1} = \lambda_s\\ k_2 = k_{r+s-2} = \lambda_s + \lambda_{s-1}\\ \vdots\\ k_s = \cdots = k_r = \lambda_s + \cdots + \lambda_1\end{gathered}$$

Using the physical restrictions on the partition λ, this may be reduced to the set of inequalities,

$$\begin{gathered}k_1 = k_{q+s-1} \leq k_2 = k_{q+s-1} \leq \cdots \leq k_s = k_{s+1} = \cdots = k_q\\ k_i \leq (\min\{p,q\}) \cdot i\end{gathered}$$

These inequalities have a strong symmetry in them. This is perhaps most easily seen in Figure 4.10.

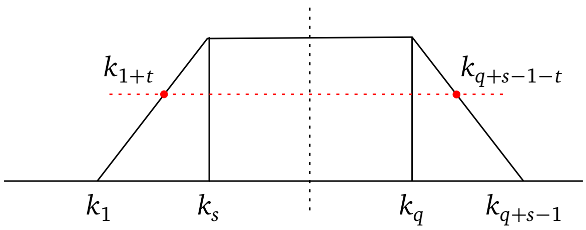

Figure 4.10: A fat hook and a two line rectangle.

As a partial verification, we check the limiting case where ν is a single row. Here the rectangle rule reduces to $k_1 = k_2 = \cdots = k_r \leq \min\{p,q\}$. This agrees with the description already established.

To end this explanation we give an example. Suppose we have the product of two squares, $\mu = (3,3,3)$ and $\nu = (2,2)$. Then the skew partitions which occur and the corresponding k_i values are shown in Figure 4.11.

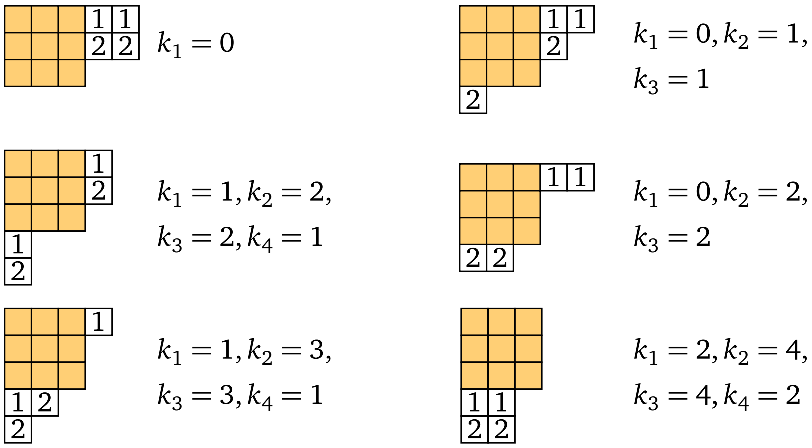

Figure 4.11: An example of the skew tableaux that occur in the product of two rectangles.

Tensor product of a fat hook with a two-line rectangle

Consider the tensor product of the representations labelled by $\mu = p\Lambda_r + q\Lambda_s$ and $\nu = t\Lambda_2$. This corresponds to a fat hook and a two-line rectangle.

We restrict the problem to fat hooks that are 'big enough', that is $r \geq 2$ and $s - r \geq 2$. Note that the extreme cases not satisfying these inequalities are actually near-rectangles and should be multiplicity-free when tensored with more general ν. To ensure that the resulting shape is a partition boxes may only be placed in three locations, on the 'shelves' (Figure 4.13) of the fat hook. The semistandard condition prevents three boxes from sharing a single column so we may consider each shelf independently.

The reverse lattice word condition requires that the far right shelf has no fewer ones than twos and the far left shelf has no fewer twos than ones. The middle shelf is unrestricted by this condition except that it must accommodate the boxes

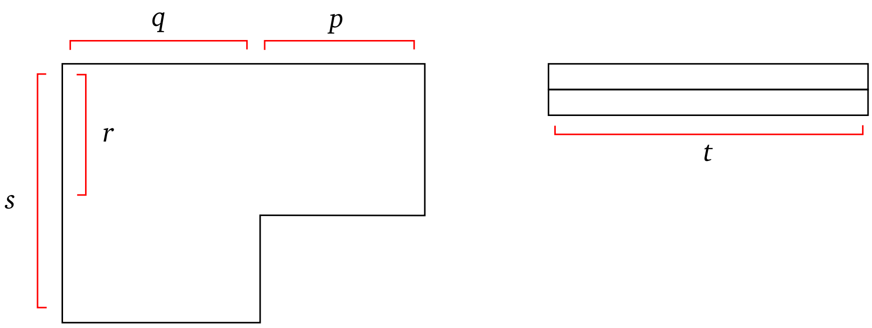

Figure 4.12: A fat hook and a two line rectangle.

from the other two. We can describe these observations in a combinatorial manner with a five-parameter description. This description consists of two stages.

In the first stage the parameters b_1 and b_2 describe the base layout. To do this, place height two rectangles of width b_1 and b_2 on shelves 2 and 3 respectively. This leaves a rectangle of width $t - b_1 - b_2$ on shelf 1.

The first stage covers every possible placement of the boxes containing a one. In the second stage we describe the further shift of boxes containing a two. Twos may be further moved from the first to second, second to third and first to third shelves without voiding semistandardness. List the number of these shifts with x_{12}, x_{23} and x_{13} respectively. The schematic in Figure 4.13 gives a visual representation of this construction.

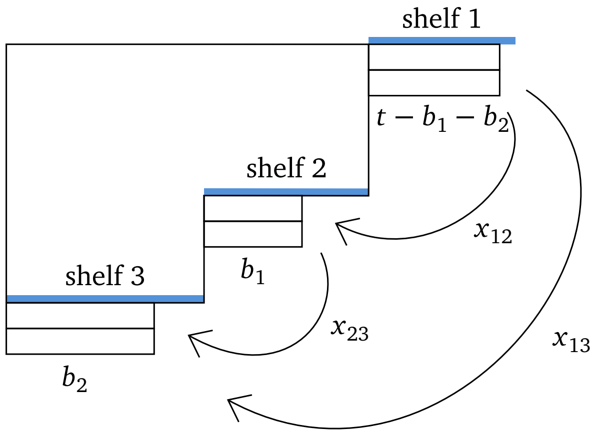

Figure 4.13: The five parameter description of the fat hook shifts.

By observing the geometry of the partition, a set of inequalities arises for these five parameters.

$$
\begin{gathered}
b_1 \leq p \\
b_2 \leq q \\
b_1 + b_2 \leq t \\
x_{12} + x_{23} \leq t - (b_1 + b_2) \\
x_{23} \leq b_1 \\
x_{12} \leq p - b_1 \\
x_{23} + x_{13} \leq q - b_2.
\end{gathered}
$$

To find the weight lattice interpretation and the k_i values we describe the shifts corresponding to each parameter.

$$\begin{aligned}
b_1 :& [\alpha_1 + \cdots + \alpha_r], [\alpha_2 + \cdots + \alpha_{r+1}] \\
b_2 :& [\alpha_1 + \cdots + \alpha_s], [\alpha_2 + \cdots + \alpha_{s+1}] \\
x_{12} :& [\alpha_2 + \cdots + \alpha_r] \\
x_{23} :& [\alpha_{r+2} + \cdots + \alpha_s] \\
x_{13} :& [\alpha_2 + \cdots + \alpha_s].
\end{aligned}$$

Collecting coefficients then gives the following expressions for the k_i,

$$\begin{aligned}
k_1 &= b_1 + b_2 \\
k_2, k_3, \ldots, k_r &= 2b_1 + 2x_{23} + x_{12} + x_{13} \\
k_{r+1} &= b_1 + 2b_2 + x_{13} \\
k_{r+2}, \ldots, k_s &= 2b_2 + x_{23} + x_{13} \\
k_{s+1} &= b_2.
\end{aligned}$$

Again to restrict this to representations of $\mathfrak{sl}_n$, simply remove all sets of k_i values where $k_n > 0$.

Before finishing the section we give an example of the fat hook construction. Let μ be the fat hook $(3,3,1,1)$ and ν be the two line rectangle $(3,3)$. There are six possible base configurations given by $(t - b_1 - b_2, b_1, b_2)$ equal to each of $(3,0,0), (2,1,0), (1,2,0), (2,0,1), (0,2,1), (1,1,1)$. We demonstrate the further shifts for two of these cases.

For $b_1 = 1$ and $b_2 = 0$ the possible shifts are given in Figure 4.14.

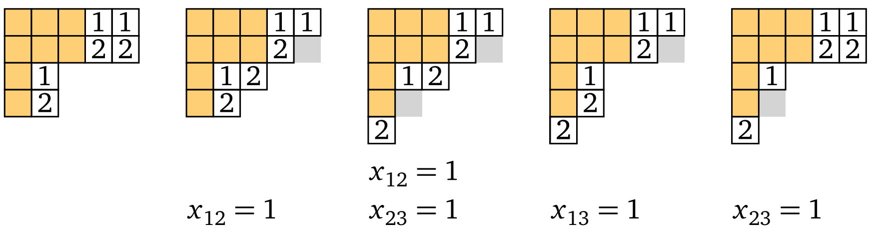

Figure 4.14: Shifts for one of the base layouts in the fat hook product example.

For $b_1 = 2$ and $b_2 = 0$ the possible shifts are given in Figure 4.15.

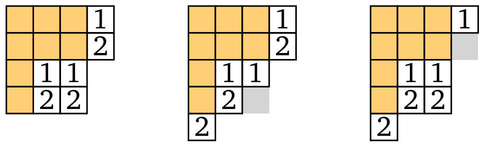

$x_{23} = 1$ $x_{13} = 1$

Figure 4.15: Shifts for another of the base layouts in the fat hook product example.

4.3 Lowest weights in a decomposition

In the Littlewood–Richardson rule, the subtraction of roots was interpreted as a shifting of boxes to the left. Intuitively, there should exist some 'leftmost' partition with a maximal number of root subtractions. Indeed this does exist, is unique and importantly, has multiplicity one.

Lemma 4.3.1 (Lowest weight in decomposition)**.** For $\mathfrak{sl}_n$ with n sufficiently large, given the tensor product decomposition of $V_\mu \otimes V_\nu$ into V_λ of the form $\lambda = \mu + \nu - \sum_{i=1}^n k_i \alpha_i$, there exists a unique λ with multiplicity one such that $k = \sum_{i=1}^n k_i$ is maximal. This is called the lowest weight in the decomposition.

Proof. Suppose μ and ν are partitions and that n is sufficiently large. Construct a new partition λ by appending the columns of ν to the columns of μ pairwise from the left. Since both μ and ν are partitions, λ is a valid partition. Consider a filling of λ/μ with weight ν. By the semistandard condition, each column of λ/μ must be filled with the integers $1, \dots, i$ where i is the height of the column. This completely determines the filling and hence $c^\lambda_{\mu\nu} = 1$. If there were any more boxes in the first column, a strictly increasing filling along this column could never be found. The same then holds recursively for the remaining columns and so the partition λ has a maximal number of left shifts. Hence the corresponding k is maximal. □

Such a lowest weight also exists without the requirement on n. However, in such scenarios we do now know whether this weight is always multiplicity free.

In the case where n is sufficiently large, there exists a simple explicit description of the weight λ with maximal k. The lowest weight in the decomposition of $V_{\Lambda_i} \otimes V_{\Lambda_j}$ is given by subtracting the root sum $\alpha_1 + 2\alpha_2 + \cdots + j(\alpha_j + \cdots + \alpha_i) + (j-1)\alpha_{j+1} + \cdots + 2\alpha_{i+j-2} + \alpha_{i+j-1}$ from $\Lambda_i + \Lambda_j$. This root subtraction corresponds to moving the first column from beside to below the second (Figure 4.16). Then for more general partitions, it suffices to use the above formula on pairs of columns, matched longest to shortest.

Unfortunately, the lowest weight phenomenon does not extend to other root systems. An explicit counterexample is given by the following decomposition in

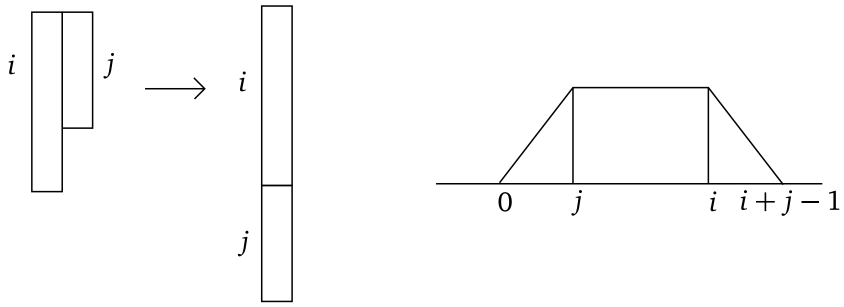

Figure 4.16: Values of k_i for the representation of lowest weight.

B_2,

$$V_{\Lambda_1} \otimes V_{\Lambda_1+2\Lambda_2} = V_{2\Lambda_1} \oplus V_{2\Lambda_2} \oplus V_{\Lambda_1+2\Lambda_2} \oplus V_{2\Lambda_1+2\Lambda_2} \oplus V_{4\Lambda_2}.$$

Here there are two weights with 'maximal' k. In any weight lattice, there is a set of parallel hyperplanes over which k is constant. In A_{n-1} roots are never parallel to these planes, however in other systems they may be. This is a possible explanation for the observations here.

4.4 Minuscule weights

In the previous sections we looked at generalising the known phase integral evaluations for $\mathfrak{sl}_n$ by considering other multiplicity-free tensor products. An important aspect of the Pieri rule is that, while one partition is restricted (either a single column or row), the second is completely free. We would like to study similar products in the other root systems. With this goal in mind, we look at the minuscule weights.

A very simple type of highest weight module would be one where the set of all weights in the module is a single orbit under the action of the Weyl group. For such modules, the highest weight is called minuscule. In various parts of the literature [6, 15] the more precise definition below is used.

Definition 4.4.1 (Minuscule)**.** A weight μ is minuscule if it satisfies

$$(\mu, \alpha^\vee) = \langle \mu, \alpha \rangle \leq 1 \text{ for all } \alpha \in \Phi^+.$$

This definition is equivalent to the weights being a single orbit under the Weyl group. A proof of this fact is given below.

Lemma 4.4.1. The highest weight module indexed by μ has all its weights in a single orbit under the Weyl group if and only if $\langle \mu, \alpha \rangle \leq 1$ for all $\alpha \in \Phi^+$.

Proof. The orbit of μ is given by $\{s_\alpha(\mu) : \alpha \in \Phi\}$. Since $s_\alpha = s_{-\alpha}$ this is equivalent to the set $\{s_\alpha(\mu) : \alpha \in \Phi^+\}$. From Section 2.3 the explicit form of these reflections is given by,

$$s_\alpha(\mu) = \mu - \langle \mu, \alpha \rangle \alpha.$$

Recall that any weight in the module is of the form $\mu - \sum_{i=1}^n k_i \alpha_i$ and that the orbit of μ forms an outer shell for these weights.

Suppose then that $\langle \mu, \alpha \rangle \leq 1$ for all $\alpha \in \Phi^+$. Then for any $\lambda = \mu - \sum_{i=1}^n k_i \alpha_i$, either λ is in the orbit of μ or λ is outside the outer shell. Thus the module is formed by a single orbit under the Weyl group.

For the reverse direction, suppose that $\langle \mu, \alpha \rangle > 1$ for some $\alpha \in \Phi^+$. Then $\mu - \alpha$ is not in the orbit of μ. But $\mu - \alpha$ is a weight in the highest weight module, so there must be at least two orbits. □

Another key property of the minuscule weights is that they are a subset of the fundamental weights. To see this, suppose μ is not fundamental. So we may infer that $\mu \geq \Lambda_i + \Lambda_j$ for some not necessarily distinct i and j. If $i = j$ then,

$$\begin{aligned} \langle \mu, \alpha_i \rangle &\geq 2\langle \Lambda_i, \alpha_i \rangle \\ &= 2. \end{aligned}$$

Thus the weight μ is not minuscule. Otherwise, since the angle between distinct simple roots is obtuse, we have $\alpha_i + \alpha_j \in \Phi^+$. Without loss of generality, assume that α_i is the shorter root. This then gives,

$$\begin{aligned} \langle \mu, \alpha_i + \alpha_j \rangle &= \frac{2(\Lambda_i + \Lambda_j, \alpha_i + \alpha_j)}{(\alpha_i + \alpha_j, \alpha_i + \alpha_j)} \\ &\geq \frac{2(\Lambda_i, \alpha_i)}{(\alpha_i, \alpha_i)} + \frac{2(\Lambda_j, \alpha_j)}{(\alpha_i, \alpha_i)} \\ &\geq \frac{2(\Lambda_i, \alpha_i)}{(\alpha_i, \alpha_i)} + \frac{2(\Lambda_j, \alpha_j)}{(\alpha_j, \alpha_j)} \\ &= \langle \Lambda_i, \alpha_i \rangle + \langle \Lambda_j, \alpha_j \rangle \\ &= 2 \end{aligned}$$

So in this case also the weight is not minuscule. Thus the minuscule weights are a subset of the fundamental weights.

We would like to know all the minuscule weights of the classical root systems. As an example we demonstrate the classification of the minuscule weights of B_n.

Example 4.4.1. For B_n the fundamental weights and positive weights are,

$$\begin{aligned} \Phi^+ &= \{\varepsilon_i \pm \varepsilon_j : 1 \leq i < j \leq n\} \cup \{\varepsilon_i : 1 \leq i \leq n\} \\ \Lambda_i &= \varepsilon_1 + \cdots + \varepsilon_i \quad \text{for } 1 \leq i < n \\ \Lambda_n &= \frac{1}{2}(\varepsilon_1 + \cdots + \varepsilon_n) \end{aligned}$$

Since the minuscule weights are a subset of the fundamental weights, it suffices to check each fundamental weight. For each $1 \leq i < n$ the weight Λ_i is not minuscule as,

$$\begin{aligned}\langle \Lambda_i, \varepsilon_1 \rangle &= \frac{2(\varepsilon_1 + \cdots + \varepsilon_i, \varepsilon_1)}{(\varepsilon_1, \varepsilon_1)} \\ &= 2.\end{aligned}$$

The weight Λ_n is minuscule. For positive roots of the form ϵ_i we have,

$$\begin{aligned}\langle \Lambda_n, \varepsilon_i \rangle &= \frac{2\left(\frac{1}{2}(\varepsilon_1 + \cdots + \varepsilon_n), \varepsilon_i\right)}{(\varepsilon_i, \varepsilon_i)} \\ &= 1.\end{aligned}$$

and it then follows that for weights of the form $\varepsilon_i - \varepsilon_j$ we have

$$\langle \Lambda_n, \varepsilon_i - \varepsilon_j \rangle = 0.$$

Thus the single fundamental weight of B_n is Λ_n.

Similar calculations for the other classical systems result in a classification of minuscule weights. These are given in Table 4.1 and also in Appendix A.

Root system	Minuscule weight(s)	Quasiminuscule
A_n	All fundamental weights are minuscule	
B_n	Λ_n	Λ_1
C_n	Λ_1	Λ_2
D_n	$\Lambda_1, \Lambda_{n-1}, \Lambda_n$	Λ_2

Table 4.1: The minuscule weights of the classical root systems.

In his second paper classifying the multiplicity-free tensor products Stembridge [15] defines the quasiminuscule weights. These satisfy the property that,

$$(\mu, \alpha^\vee) = \langle \mu, \alpha \rangle \leq 2 \text{ for all } \alpha \in \Phi^+.$$

All minuscule weights are quasiminuscule. Additional ones are listed in Table 4.1. The structure of modules generated by quasiminuscule weights are not as simple as with minuscule weights, but they do share some characteristics. We will hence also examine the Weyl orbits of some of these quasiminuscule weights.

Tensor products with minuscule weights

We would like to describe the k_i values corresponding to tensor products with the minuscule weights for $\mathrm{B}_n, \mathrm{C}_n$ and D_n. Since any weight $\mu = \sum_{i=1}^n \mu_i \Lambda_i$ in the orbit

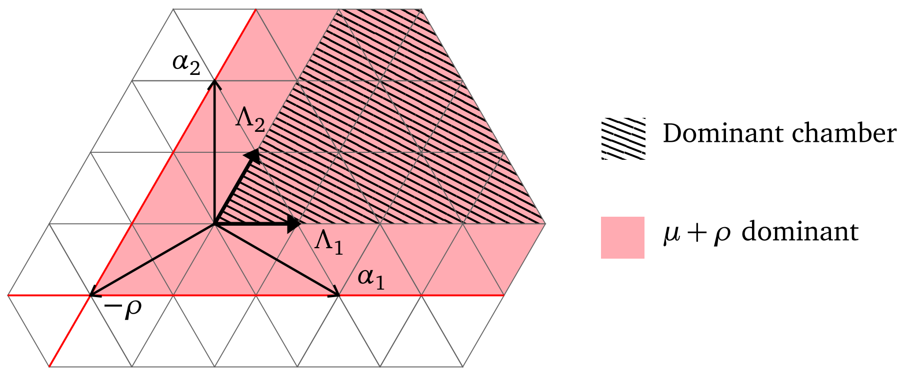

Figure 4.17: The positive chamber with respects to the negative Weyl vector.

of a minuscule weight satisfies $|\mu_i| \leq 1$, we must have $\mu+\rho$ a dominant weight. Colloquially, we think of this as μ being in the positive chamber with respects to the negative Weyl vector. This is illustrated in Figure 4.17. As a consequence of this, in Klimyk's rule no reflections can take place. In particular, this means that since all weights in the module have multiplicity one, the product is multiplicity free.

This observation also allows us to give a simple description of the k_i values. As no reflections occur, the highest weights in the tensor product decomposition described by Klimyk's rule will only be a translation of the weights in the orbit of μ. Since this shift does not change the 'relative' root subtractions, the k_i values are a subset of the number of roots subtracted from μ to give each weight in its orbit.

In the following sections we endeavour to describe these k_i values.

4.5 Minuscule modules over B_n

The root system B_n has a single minuscule weight Λ_n and a single quasiminuscule weight Λ_1. Both of these have all weights occurring with multiplicity one and exhibit uniform tensor product structures.

The quasiminuscule module V_{Λ_1}

To describe the orbit of Λ_1 under the Weyl group we construct a Cayley-like graph of the weights under the action. For every weight in the graph, draw an edge to depict the action of each generator of the Weyl group. Fortunately the simple reflections act trivially on all the fundamental weights except for,

$$\sigma_i(\Lambda_i) = \Lambda_i - \alpha_i.$$

Omitting these trivial actions from the diagram makes it much more manageable. Enumerating these edges and the coefficients of the corresponding fundamental weight precisely gives the list of k_i values.

All the Cayley graphs that we draw will arise from studying the orbit of a particular weight. With this in mind we will depict the edges as arrows to demonstrate how the calculations proceed and to show the direction of the root subtractions. Of course the reflections that these arrows describe have no inherent direction to them.

The orbit of Λ_1 is given below.

$$(\Lambda_1) \xrightarrow{\sigma_1} (\Lambda_2 - \Lambda_1) \xrightarrow{\sigma_2} (\Lambda_3 - \Lambda_2) \xrightarrow{\sigma_3} \dots \xrightarrow{\sigma_{n-2}} (\Lambda_{n-1} - \Lambda_{n-2}) \xrightarrow{\sigma_{n-1}} (2\Lambda_n - \Lambda_{n-1})$$
$$\xrightarrow{\sigma_n} (-[-2\Lambda_n - \Lambda_{n-1}]) \xrightarrow{\sigma_{n-1}} (-[\Lambda_{n-1} - \Lambda_{n-2}]) \xrightarrow{\sigma_{n-2}} \dots \xrightarrow{\sigma_1} (-\Lambda_1)$$

We see here $2n$ weights in the orbit. To complete the standard representation, which is $2n+1$ dimensional, we add the only permitted weight, 0.

The k_i values increase one after another until they are all one, then increase again in the reverse order until all are two. This may be described by the inequalities below. For each $0 \leq m \leq n$ we have,

$$k_i = 1 \text{ for } 1 \leq i \leq m$$
$$k_i = 0 \text{ for } m < i \leq n$$

or

$$k_i = 1 \text{ for } 1 \leq i \leq m$$
$$k_i = 2 \text{ for } m < i \leq n.$$

The minuscule module V_{Λ_n}

We wish to describe the weights in the module V_{Λ_n} of B_n. Since this module is minuscule these weights will be precisely the orbit of Λ_n under the Weyl group. This orbit is most convenient to describe recursively upon the rank of the root system. To explain this recursion we give a few low-dimensional examples of the construction and then describe the general pattern.

The weight lattice of B_2 with the orbit of Λ_n is drawn in Figure 4.18. Here there are only four weights in the orbit. Unfortunately the weight lattices of B_n for larger values of n are impractical to draw, so instead we show the Cayley like graphs of the orbits. This graph for B_2 is also given in Figure 4.18

The Cayley graphs depicting the orbit of Λ_n in B_3 and B_4 are shown below in Figures 4.19 and 4.20. Each graph is constructed by connecting two copies of the previous graph. This is the general pattern that continues as the rank increases, a result which we now prove.

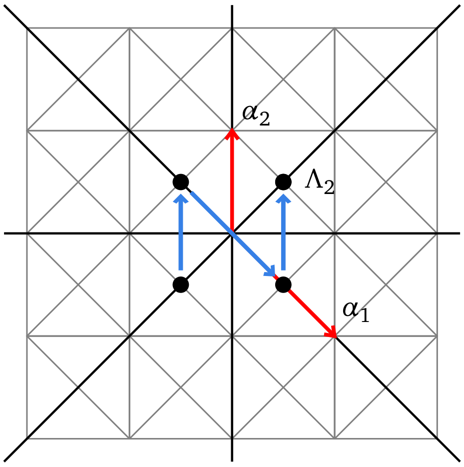

$$\begin{tikzcd}
\Lambda_2 \arrow[d, "\sigma_2"] \\
\Lambda_1 - \Lambda_2 \arrow[d, "\sigma_1"] \\
\Lambda_2 - \Lambda_1 \arrow[d, "\sigma_2"] \\
-\Lambda_2
\end{tikzcd}$$

Figure 4.18: The orbit of Λ_2 in B_2.

Lemma 4.5.1. The weight $-\Lambda_n$ occurs in V_{Λ_n} and has corresponding k_i values $k_i = i$.

Proof. This is simply a computation. Apply the generators $\sigma_n, \sigma_{n-1}, \dots, \sigma_1$ then $\sigma_n, \sigma_{n-1}, \dots, \sigma_2$ and so on until σ_n, σ_{n-1} and σ_n.

$$\begin{gathered}
(\Lambda_n) \xrightarrow{\sigma_n} (\Lambda_{n-1} - \Lambda_n) \xrightarrow{\sigma_{n-1}} (\Lambda_{n-2} - \Lambda_{n-1}) \xrightarrow{\sigma_{n-2}} \dots \xrightarrow{\sigma_2} (\Lambda_1 - \Lambda_2 + \Lambda_n) \\
\xrightarrow{\sigma_1} (-\Lambda_1 + \Lambda_n) \left[\xrightarrow{\sigma_n} \dots \xrightarrow{\sigma_2} (-\Lambda_2 + \Lambda_n) \right] \\
\left[\xrightarrow{\sigma_n} \dots \xrightarrow{\sigma_3} (-\Lambda_3 + \Lambda_n) \right] \dots \left[\xrightarrow{\sigma_n} \dots \xrightarrow{\sigma_{n-1}} (-\Lambda_{n-1} + \Lambda_n) \right] \xrightarrow{\sigma_n} (-\Lambda_n)
\end{gathered}$$

□

Lemma 4.5.2. The full collection of weights in the B_n module V_{Λ_n} is given by the following recursive construction.

- Take two copies of the Cayley graph for the orbit in B_{n-1}, call them X and Y.
- Add one to all indices.
- Recursively, each X and Y has a positive and negative half.
- Add Λ_1 to the negative half of X and the positive half of Y.
- Connect the negative weights in X to the positive weights in Y, following the action of σ_1.

$$\begin{tikzcd}
\Lambda_3 \arrow[d, "\sigma_3"] & \\
\Lambda_2 - \Lambda_3 \arrow[d, "\sigma_2"] & \\
\Lambda_1 - \Lambda_2 + \Lambda_3 \arrow[d, "\sigma_3"] \arrow[r, "\sigma_1"] & -\Lambda_1 + \Lambda_3 \arrow[d, "\sigma_3"] \\
\Lambda_1 - \Lambda_3 \arrow[r, "\sigma_1"] & -\Lambda_1 + \Lambda_2 - \Lambda_3 \arrow[d, "\sigma_2"] \\
 & -\Lambda_2 + \Lambda_3 \arrow[d, "\sigma_3"] \\
 & -\Lambda_3
\end{tikzcd}$$

Figure 4.19: The orbit of Λ_3 in B_3.

- ◇ Finally, call X the new positive half and Y the new negative half.

Proof. The proof is by induction on n. We use two hypotheses, that the orbit of Λ_n matches the construction described above and that half the weights at each stage have been acted on by σ_1, including $-\Lambda_n$.

As a base, the action for B_2 satisfies the hypotheses.

To begin the induction, consider the orbit of Λ_n under $\langle \sigma_2, \ldots, \sigma_n \rangle$. Following the Dynkin diagram, taking this subgroup of the Weyl group may be thought of as embedding B_{n-1} in B_n. So by the inductive hypothesis we know what this orbit looks like. The only difference (as reflected in the Cartan matrix) is that the generator σ_2 now adds Λ_1. However this does not change the orbit since $\sigma_2, \ldots, \sigma_n$ all act trivially on Λ_1. By the second inductive hypothesis this Λ_1 will occur in precisely half of the weights. We call these weight the negative half.

Notably $-\Lambda_n$ does not occur in the orbit of Λ_n under this restricted action. By the first lemma it does however occur in the module. Consider next its orbit under $\langle \sigma_2, \ldots, \sigma_n \rangle$. By linearity this will be precisely the negative of the previous orbit. Again add $-\Lambda_1$ to the positive half to account for the change in σ_2.

$$
\begin{tikzcd}
\Lambda_4 \arrow[d, "\sigma_4"] & & & \\
\Lambda_3 - \Lambda_3 \arrow[d, "\sigma_3"] & & & \\
\Lambda_2 - \Lambda_3 + \Lambda_4 \arrow[r, "\sigma_2"] \arrow[d, "\sigma_4"] & \Lambda_1 - \Lambda_2 + \Lambda_4 \arrow[r, "\sigma_1"] \arrow[d, "\sigma_4"] & -\Lambda_1 \Lambda_4 \arrow[d, "\sigma_4"] & \\
\Lambda_2 - \Lambda_4 \arrow[r, "\sigma_2"] & \Lambda_1 - \Lambda_2 + \Lambda_3 - \Lambda_4 \arrow[r, "\sigma_1"] \arrow[d, "\sigma_3"] & -\Lambda_1 + \Lambda_3 - \Lambda_4 \arrow[d, "\sigma_3"] & \\
 & \Lambda_1 - \Lambda_3 + \Lambda_4 \arrow[r, "\sigma_1"] \arrow[d, "\sigma_4"] & -\Lambda_1 + \Lambda_2 - \Lambda_3 + \Lambda_4 \arrow[r, "\sigma_2"] \arrow[d, "\sigma_4"] & -\Lambda_2 + \Lambda_4 \arrow[d, "\sigma_4"] \\
 & \Lambda_1 - \Lambda_4 \arrow[r, "\sigma_1"] & -\Lambda_1 + \Lambda_2 - \Lambda_4 \arrow[r, "\sigma_2"] & -\Lambda_2 + \Lambda_3 - \Lambda_4 \arrow[d, "\sigma_3"] \\
 & & & -\Lambda_3 + \Lambda_4 \arrow[d, "\sigma_4"] \\
 & & & -\Lambda_4
\end{tikzcd}
$$

Figure 4.20: The orbit of Λ_4 in B_4.

The only action on the weights mentioned so far that is not accounted for is that of s_1. This action gives a bijection between the negative halves of X and Y. To describe this bijection we need the $n-2$ case of the induction.

Structurally the negative half of X and the positive half of Y both resemble the $n-2$ case of the induction. If we remove Λ_1 and Λ_2 they are in fact identical. The bijection given by the action σ_1 is effectively the identity map here. All weights in the negative half of X have $+\Lambda_1$ and all in the positive half of Y have $-\Lambda_1$. From the $n-1$ case, the first half in the negative set of X will have $-\Lambda_2$ while the others will have no Λ_2. Similarly the first half in the positive set of Y will have no Λ_2 and the others will have $+\Lambda_2$.

Then the action of σ_1 forms a bijection by the calculations,

$$
\begin{aligned}
(\Lambda_1 - \Lambda_2 + \dots) &\xrightarrow{\sigma_1} (-\Lambda_1 + \dots) \\
(\Lambda_1 + \dots) &\xrightarrow{\sigma_1} (-\Lambda_1 + \Lambda_2 + \dots)\,.
\end{aligned}
$$

The various subsets here are more clearly illustrated with Figure 4.22. The

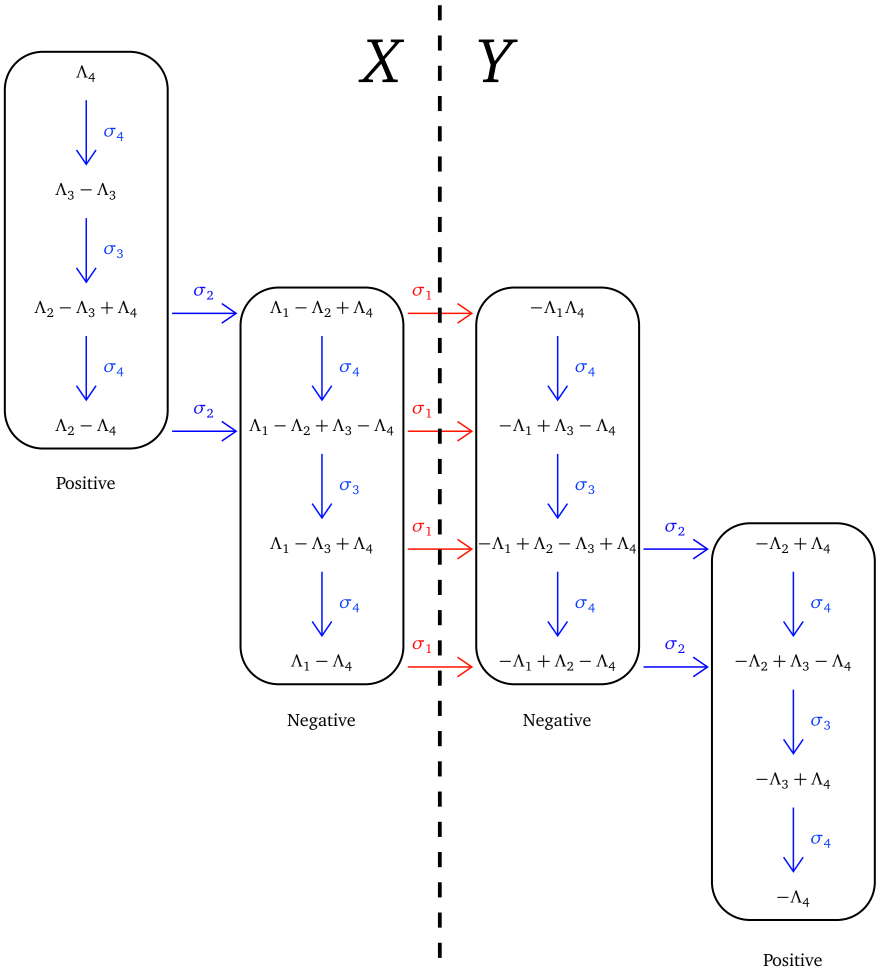

Figure 4.21: Sets used in the recursion.

bijection above then completely describes the action of σ_1. Thus we have found all the weights in the orbit and they follow the desired construction. $\square$

Every highest weight module which can occur in the direct sum decomposition will have k_i values corresponding to the number of roots subtracted from Λ_n to get each weight in the orbit. The following theorem describes these k_i values.

Theorem 4.5.3 (k_i values for V_{Λ_n})**.** Given the construction in the lemma above the k_i values are described by the two inequalities,

$$k_i \leq i$$
$$0 \leq k_{i+1} - k_i \leq 1.$$

Proof. The proof will be inductive. As a base case, the inequalities given precisely describe the k_i values in B_2.

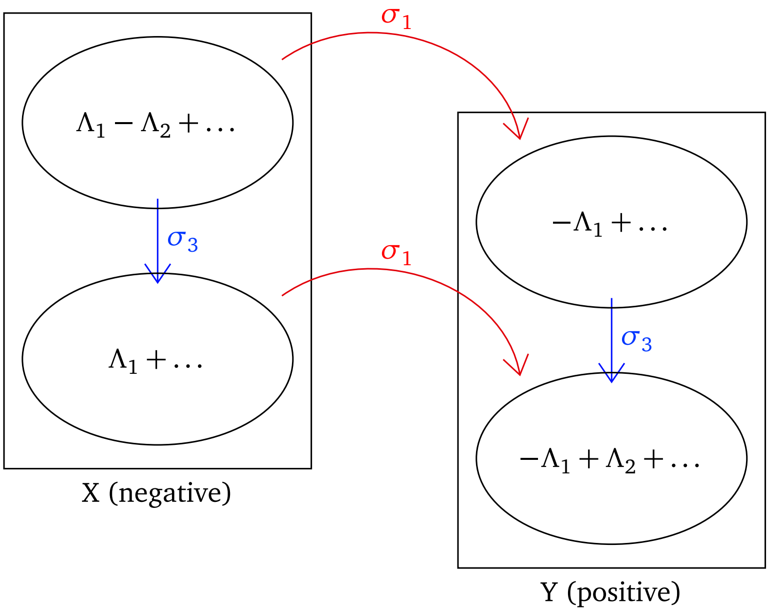

Figure 4.22: The bijection induced by σ_1.

Now from the recursive structure, the weights may be divided into two sets. Call the set X those which are attached to Λ_n under $\langle \sigma_{\alpha_2}, \dots, \sigma_{\alpha_n} \rangle$. The similar orbit of $-\Lambda_n$ is called Y and contains only the negatives of the weights in X.

We will use the $n-1$ case to describe the k_i values of the two sets. The weights in X have exactly the same number of roots subtracted as those in the B_{n-1} case. Symmetry allows us to describe the k_i values for Y. Since these weights are exactly the negatives of those in X, the B_{n-1} case tells us how many simple roots can be added to $-\Lambda_n$ to get each weight. From the first lemma we know exactly the coordinates of $-\Lambda_n$ and hence we may describe the rest.

Firstly we show that all the weights satisfy the inequalities. From the $n-1$ case, if a weight is in X it has k_i values given by,

$$k_1 = 0$$
$$k_i \leq i - 1$$
$$0 \leq k_{i+1} - k_i \leq 1 \text{ for } i = 2, \dots, n-1.$$

Since $k_1 = 0$ there holds $0 \leq k_2 - k_1 \leq 1$ and so all weights in X satisfy the proposed inequalities.

With the inductive case, any weight in Y may be described by $-\Lambda_n + \sum_{i=1}^{n} u_i \Lambda_i$ satisfying,

$$u_1 = 0$$
$$u_i \leq i - 1$$
$$0 \leq u_{i+1} - u_i \leq 1.$$

By the first lemma the k_i values are then given by $k_i = i - u_i$. Since $u_i \geq 0$, there must hold $k_i \leq i$. Also,

$$\begin{aligned} k_{i+1} - k_i &= (i + 1 - u_{i+1}) - (i - u_i) \\ &= 1 - (u_{i+1} - u_i) \end{aligned}$$

so $0 \leq k_{i+1} - k_i \leq 1$. Thus all the weights in Y satisfy the proposed inequalities.

For the second step we, show that every set of k_i values given by the inequalities correspond to some root.

If $k_1 = 0$ then,

$$\begin{aligned} k_i &= (k_i - k_{i-1}) + (k_{i-1} - k_{i-2}) + \cdots + (k_2 - k_1) \\ &\leq i - 1. \end{aligned}$$

This occurs as a weight in X by the inductive case.

If $k = 1$ consider $u_i = i - k_i$. From a similar calculation as before $0 \leq u_{i+1} - u_i \leq 1$. Since $u_1 = 0$ there holds $u_i \leq i - 1$. So these k_i values occur as a weight in Y. □

4.6 Minuscule modules over C_n

The root system C_n has only a single minuscule weight Λ_1.

The minuscule module V_{Λ_1}

This module closely resembles the quasiminuscule module of B_n. It has $2n$ weights which form a single chain as the Weyl orbit.

$$\begin{gathered} (\Lambda_1) \xrightarrow{\sigma_1} (\Lambda_2 - \Lambda_1) \xrightarrow{\sigma_2} (\Lambda_3 - \Lambda_2) \xrightarrow{\sigma_3} \dots \xrightarrow{\sigma_{n-1}} (\Lambda_n - \Lambda_{n-1}) \xrightarrow{\sigma_n} \\ (\Lambda_{n-1} - \Lambda_n) \xrightarrow{\sigma_{n-1}} (\Lambda_{n-2} - \Lambda_{n-1}) \xrightarrow{\sigma_{n-2}} \dots \xrightarrow{\sigma_1} (-\Lambda_1) \end{gathered}$$

The k_i values are also similar to those from B_n. They increase from 0 to 1 step by step for each i as i increases. Then they increase from 1 to 2 step by step for each $i < n$ as i decreases. This may be described with the following formula.

For each $1 \leq m \leq n$ we have a set of k_i values,

$$\begin{aligned} k_i &= 1 \text{ for } 1 \leq i \leq m \\ k_i &= 0 \text{ for } m + 1 \leq i \leq n \end{aligned}$$

and

$$\begin{aligned} k_i &= 2 \text{ for } 1 \leq i < m \\ k_i &= 1 \text{ for } m \leq i \leq n. \end{aligned}$$

4.7 Minuscule modules over D_n

The root system D_n has three minuscule weights, Λ_1, Λ_{n-1} and Λ_n.

The minuscule module V_{Λ_1}

Again this module shares much of its structure with the standard representations of B and C. The Weyl orbit is a single chain, until the reflections σ_{n-1} and σ_n have a nontrivial action. Then the orbit bifurcates before rejoining and continuing along the negative of the original chain. A calculation to verify this creates the Cayley graph of Figure 4.23.

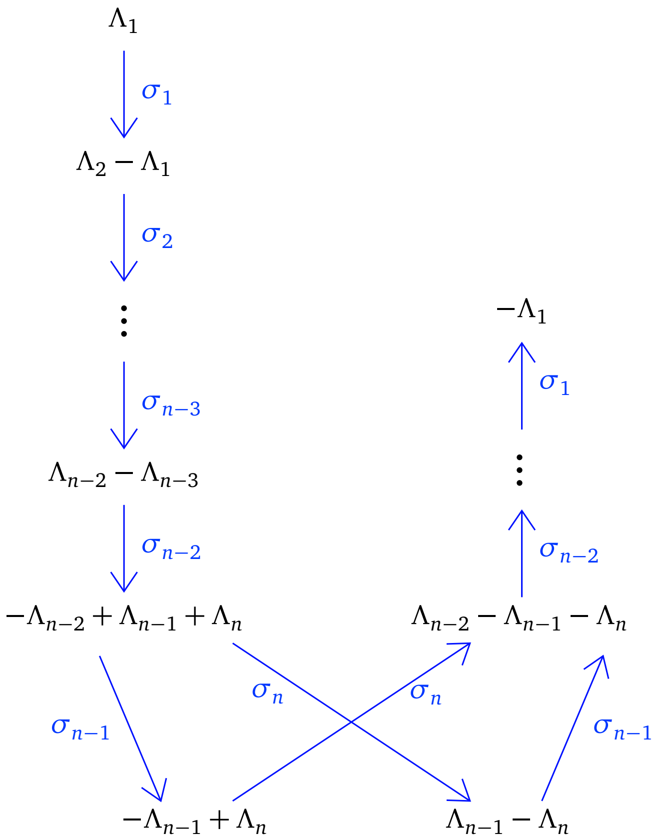

Figure 4.23: The weight orbit of Λ_1 in D_n.

We describe the two chains and the two symmetrical cases separately. For each m satisfying $1 \leq m \leq n-1$ we have the first chain,

$$k_i = 1 \text{ for } 1 \leq i \leq m$$
$$k_i = 0 \text{ for } i > m,$$

the two symmetrical cases,

$$k_i = 1 \text{ for } 1 \leq i \leq n-2$$
$$k_{n-1} = 1, k_n = 0$$

or

$$k_i = 1 \text{ for } 1 \leq i \leq n-2$$
$$k_{n-1} = 0, k_n = 1$$

and finally the last chain,

$$k_i = 1 \text{ for } 1 \leq i \leq m$$
$$k_i = 2 \text{ for } m < i \leq n-2$$
$$k_{n-1} = k_n = 1.$$

The minuscule modules $V_{\Lambda_{n-1}}$ and V_{Λ_n}

These modules are easiest to describe together. In D_4 both these weights have orbits identical to that of Λ_3 in B_3. These are shown in Figure 4.24.

$$\begin{tikzcd}
\Lambda_3 \arrow[d, "\sigma_3"] & & \Lambda_4 \arrow[d, "\sigma_4"] & \\
\Lambda_2 - \Lambda_3 \arrow[d, "\sigma_2"] & & \Lambda_2 - \Lambda_4 \arrow[d, "\sigma_2"] & \\
\Lambda_1 - \Lambda_2 + \Lambda_4 \arrow[r, "\sigma_1"] \arrow[d, "\sigma_4"] & -\Lambda_1 + \Lambda_4 \arrow[d, "\sigma_4"] & \Lambda_1 - \Lambda_2 + \Lambda_3 \arrow[r, "\sigma_1"] \arrow[d, "\sigma_3"] & -\Lambda_1 + \Lambda_3 \arrow[d, "\sigma_3"] \\
\Lambda_1 - \Lambda_4 \arrow[r, "\sigma_1"] & -\Lambda_1 + \Lambda_2 - \Lambda_4 \arrow[d, "\sigma_2"] & \Lambda_1 - \Lambda_3 \arrow[r, "\sigma_1"] & -\Lambda_1 + \Lambda_2 - \Lambda_3 \arrow[d, "\sigma_2"] \\
& -\Lambda_2 + \Lambda_3 \arrow[d, "\sigma_3"] & & -\Lambda_2 + \Lambda_4 \arrow[d, "\sigma_4"] \\
& -\Lambda_3 & & -\Lambda_4
\end{tikzcd}$$

Figure 4.24: The weight orbits of Λ_3 and Λ_4 in D_4.

In higher dimensions the orbit again echoes that of B_n. Specifically, the same recursive construction is observed. However instead of taking two copies of the previous orbit for X and Y we instead take the previous orbit of Λ_n as X and the previous orbit of Λ_{n-1} as Y. Continuing with the construction unchanged apart from this then results in the module for Λ_n.

Unlike in the B case, this module does not contain $-\Lambda_n$. It does however contain $-\Lambda_{n-1}$. Hence to get the orbit of $V_{\Lambda_{n-1}}$, simply take the negative of the orbit for Λ_n.

Classical root systems A

This appendix provides a summary of the relevant information about the classical root systems. The expressions given for the roots and weights correspond to the standard embedding in $\mathbb{R}^n$ [8].

A.1 The root system A_{n-1}.

The root system A_{n-1} corresponds to the semisimple Lie algebra $\mathfrak{sl}(n,\mathbb{C})$ of n-dimensional complex matrices with trace zero. Here it is easiest to describe the roots and weights if we embed the system in a vector space of dimension one higher than the rank of the system.

Dynkin diagram	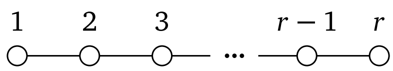
Cartan matrix	$\begin{pmatrix} 2 & -1 & 0 & & \\ -1 & 2 & -1 & 0 & \\ & \ddots & \ddots & \ddots & \\ & 0 & -1 & 2 & -1 \\ & & 0 & -1 & 2 \end{pmatrix}$
Roots	$\Phi = \{\varepsilon_i - \varepsilon_j : 1 \le i \ne j \le n\}$
Simple roots	$\alpha_i = \varepsilon_i - \varepsilon_{i+1}$ $(1 \le i \le n-1)$
Positive roots	$\Phi^+ = \{\varepsilon_i - \varepsilon_j : 1 \le i < j \le n\}$
Fundamental weights	$\Lambda_i = \varepsilon_1 + \cdots + \varepsilon_i - \frac{i}{n}(\varepsilon_1 + \cdots + \varepsilon_n)$
Minuscule weights	All fundamental weights are minuscule.

A.2 The root system B_n.

The root system B_n corresponds to the semisimple Lie algebra $\mathfrak{o}(2n+1,\mathbb{C})$ called the odd orthogonal algebra.

Dynkin diagram	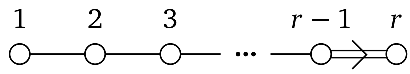
Cartan matrix	$\begin{pmatrix} 2 & -1 & 0 & & & & & \\ -1 & 2 & -1 & 0 & & & & \\ 0 & -1 & 2 & -1 & 0 & & & \\ & & \ddots & \ddots & \ddots & & & \\ & & 0 & -1 & 2 & -1 & 0 \\ & & & 0 & -1 & 2 & -2 \\ & & & & 0 & -1 & 2 \end{pmatrix}$
Roots	$\Phi = \{\pm(\varepsilon_i \pm \varepsilon_j) : 1 \leq i \neq j \leq n\} \cup \{\pm\varepsilon_i : 1 \leq i \leq n\}$
Simple roots	$\alpha_i = \varepsilon_i - \varepsilon_{i+1}$ $(1 \leq i \leq n-1)$, $\alpha_n = \varepsilon_n$
Positive roots	$\Phi^+ = \{\varepsilon_i \pm \varepsilon_j : 1 \leq i < j \leq n\} \cup \{\varepsilon_i : 1 \leq i \leq n\}$
Fundamental weights	$\Lambda_i = \varepsilon_1 + \cdots + \varepsilon_i$ $\Lambda_n = \frac{1}{2}(\varepsilon_1 + \cdots + \varepsilon_n)$
Minuscule weight	Λ_n
Quasiminuscule weight	Λ_1

A.3 The root system C_n.

The root system C_n corresponds to the semisimple Lie algebra $\mathfrak{sp}(2n, \mathbb{C})$ called the symplectic Lie algebra.

Dynkin diagram	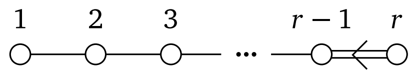
Cartan matrix	$\begin{pmatrix} 2 & -1 & 0 & & & & \\ -1 & 2 & -1 & 0 & & & \\ 0 & -1 & 2 & -1 & 0 & & \\ & & \ddots & \ddots & \ddots & & \\ & & 0 & -1 & 2 & -1 & 0 \\ & & & 0 & -1 & 2 & -1 \\ & & & & 0 & -2 & 2 \end{pmatrix}$
Roots	$\Phi = \{\pm(\varepsilon_i \pm \varepsilon_j) : 1 \leq i \neq j \leq n\} \cup \{\pm 2\varepsilon_i : 1 \leq i \leq n\}$
Simple roots	$\alpha_i = \varepsilon_i - \varepsilon_{i+1}$ $(1 \leq i \leq n-1)$, $\alpha_n = 2\varepsilon_n$
Positive roots	$\Phi^+ = \{\varepsilon_i \pm \varepsilon_j : 1 \leq i < j \leq n\} \cup \{2\varepsilon_i : 1 \leq i \leq n\}$
Fundamental weights	$\Lambda_i = \varepsilon_1 + \cdots + \varepsilon_i$
Minuscule weight	Λ_1
Quasiminuscule weight	Λ_2

A.4 The root system D_n.

The root system D_n corresponds to the semisimple Lie algebra $\mathfrak{o}(2n, \mathbb{C})$ called the even orthogonal algebra.

Dynkin diagram

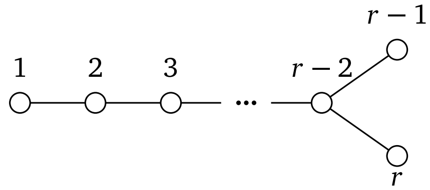

Cartan matrix

$$\begin{pmatrix} 2 & -1 & 0 & & & & & \\ -1 & 2 & -1 & 0 & & & & \\ 0 & -1 & 2 & -1 & 0 & & & \\ & & \ddots & \ddots & \ddots & & & \\ & & 0 & -1 & 2 & -1 & -1 \\ & & & 0 & -1 & 2 & 0 \\ & & & & -1 & 0 & 2 \end{pmatrix}$$

Roots	$\Phi = \{\pm(\varepsilon_i \pm \varepsilon_j) : 1 \leq i \neq j \leq n\}$
Simple roots	$\alpha_i = \varepsilon_i - \varepsilon_{i+1}$ $(1 \leq i \leq n-1)$, $\alpha_n = \varepsilon_{i-1} + \varepsilon_n$
Positive roots	$\Phi^+ = \{\varepsilon_i \pm \varepsilon_j : 1 \leq i < j \leq n\}$
Fundamental weights	$\Lambda_i = \varepsilon_1 + \cdots + \varepsilon_i$ $\Lambda_{n-1} = \frac{1}{2}(\varepsilon_1 + \cdots + \varepsilon_{n-1} - \varepsilon_n)$ $\Lambda_n = \frac{1}{2}(\varepsilon_1 + \cdots + \varepsilon_{n-1} + \varepsilon_n)$
Minuscule weights	$\Lambda_1, \Lambda_{n-1}, \Lambda_n$
Quasiminuscule weight	Λ_2

References

[1] Greg W. Anderson. A short proof of Selberg's generalized beta formula. *Forum Math.*, 3(4):415–417, 1991.

[2] Karin Erdmann and Mark J. Wildon. *Introduction to Lie algebras*. Springer Undergraduate Mathematics Series. Springer-Verlag London Ltd., London, 2006.

[3] Pavel I. Etingof, Igor B. Frenkel, and Alexander A. Kirillov, Jr. *Lectures on representation theory and Knizhnik-Zamolodchikov equations*, volume 58 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1998.

[4] Peter J. Forrester and S. Ole Warnaar. The importance of the Selberg integral. *Bull. Amer. Math. Soc. (N.S.)*, 45(4):489–534, 2008.

[5] William Fulton. *Young tableaux*, volume 35 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1997. With applications to representation theory and geometry.

[6] William Fulton and Joe Harris. *Representation theory*, volume 129 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991. A first course, Readings in Mathematics.

[7] Gert Heckman. Hypergeometric functions. Lecture notes, 2011. `http://www.math.ru.nl/~heckman/tsinghua.pdf`.

[8] James E. Humphreys. *Introduction to Lie algebras and representation theory*, volume 9 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1978. Second printing, revised.

[9] Alexander Kirillov, Jr. *An introduction to Lie groups and Lie algebras*, volume 113 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2008.

[10] Katsuhisa Mimachi and Takashi Takamuki. A generalization of the beta integral arising from the Knizhnik-Zamolodchikov equation for the vector representations of types B_n, C_n and D_n. *Kyushu J. Math.*, 59(1):117–126, 2005.

[11] Evgeny Mukhin and Alexander Varchenko. Remarks on critical points of phase functions and norms of Bethe vectors. In *Arrangements—Tokyo 1998*, volume 27 of *Adv. Stud. Pure Math.*, pages 239–246. Kinokuniya, Tokyo, 2000.

[12] Atle Selberg. Bemerkninger om et multipelt integral. *Norsk. Mat. Tidsskr*, 24:71–78, 1944.

[13] Gopala Krishna Srinivasan. The gamma function: an eclectic tour. *Amer. Math. Monthly*, 114(4):297–315, 2007.

[14] John R. Stembridge. Multiplicity-free products of Schur functions. *Ann. Comb.*, 5(2):113–121, 2001.

[15] John R. Stembridge. Multiplicity-free products and restrictions of Weyl characters. *Represent. Theory*, 7:404–439 (electronic), 2003.

[16] V. Tarasov and A. Varchenko. Selberg-type integrals associated with $\mathfrak{sl}_3$. *Lett. Math. Phys.*, 65(3):173–185, 2003.

[17] S. Ole Warnaar. A Selberg integral for the Lie algebra A_n. *Acta Math.*, 203(2):269–304, 2009.

[18] S. Ole Warnaar. The $\mathfrak{sl}_3$ Selberg integral. *Adv. Math.*, 224(2):499–524, 2010.

[19] Jerzy Weyman. Pieri's formulas for classical groups. In *Invariant theory (Denton, TX, 1986)*, volume 88 of *Contemp. Math.*, pages 177–184. Amer. Math. Soc., Providence, RI, 1989.