Computing Discrete Series Multiplicities in the Connected Case
Atlas of Lie Groups AIM Workshop III
25–29 July 2005

John R. Stembridge ⟨jrs@umich.edu⟩

Department of Mathematics
University of Michigan
Ann Arbor, Michigan 48109–1043

Contents

1. The Blattner Multiplicity Formula

The main question we are addressing in these lecture notes is whether the computation of discrete series multiplicities via the Blattner formula (or relatives thereof) is practical.

The short answer is: mostly yes, but it depends on exactly how far you want to go in computing the multiplicities in a given series.

A. (Forgive my choice of) Notation.

Let G be a complex, connected, semisimple (or reductive) Lie group, T a maximal torus in G, and K the identity component of the fixed points of an involution in G. Equivalently, K is the complexification of a maximal compact subgroup of the identity component of a real form of G. Of course we assume $T \subset K$.

We let Λ_G denote the lattice of T-characters (the weight lattice), $\Phi_G \subset \Lambda_G$ the root system, $\Phi_G^\vee$ the co-roots, $\Lambda_G^\vee$ the co-weights, and $W(G)$ the Weyl group. Fix a choice of positive roots, say Φ_G^+, and let ρ_G denote half the sum of Φ_G^+.

The root system of K and its Weyl group $W(K)$ are determined by a $\mathbb{Z}/2\mathbb{Z}$-grading of Φ_G; i.e., a map $\varepsilon : \Phi_G \to \mathbb{Z}/2\mathbb{Z}$ such that $\varepsilon(\alpha+\beta) = \varepsilon(\alpha)+\varepsilon(\beta)$ whenever $\alpha, \beta, \alpha+\beta$ are G-roots. In these terms,

$$\Phi_K := \{\alpha \in \Phi_G : \varepsilon(\alpha) = 0\}$$

is the root system of K (also known as the set of "compact" roots), and

$$\Psi := \{\alpha \in \Phi_G^+ : \varepsilon(\alpha) = 1\}$$

is the set of (positive) "non-compact" roots. The positive roots of G induce a compatible positive system within Φ_K; namely, $\Phi_K^+ := \Phi_K \cap \Phi_G^+$.

We will primarily be focused on the weight lattice of G and the root system and Weyl group of K, so from now on, we will use the abbreviations

$$\Lambda = \Lambda_G, \quad \Phi = \Phi_K, \quad W = W(K), \quad \rho = \rho_K.$$

We associate to (G, K) the partition function

$$\mathcal{P}_\Psi(\gamma) := \left|\left\{m : \Psi \to \mathbb{N} : \gamma = \textstyle\sum_\beta m(\beta)\beta\right\}\right| \quad (\gamma \in \Lambda).$$

Thus, $\mathcal{P}_\Psi(\gamma)$ counts the number of ways to write γ as an unordered sum of roots from Ψ. Next, for each pair $\mu, \nu \in \Lambda$ we set

$$B_\Psi(\mu, \nu) := \sum_{w \in W} \operatorname{sgn}(w)\mathcal{P}_\Psi(w(\mu + \rho) - (\nu + \rho)). \tag{1}$$

Blattner Multiplicity Formula (Hecht-Schmid [**HS**]). *If $\mu \in \Lambda$ is K-dominant and $\lambda \in \Lambda$ is dominant and regular for G, then $B_\Psi(\mu, \lambda + \rho_G - 2\rho)$ is the multiplicity of μ in the discrete series for G indexed by λ.*

In the above formula, note that $\nu = \lambda + \rho_G - 2\rho$ is the lowest K-type in the series; it occurs with multiplicity one.

Since it seems to pose no significant additional obstacles, we will relax the assumptions on λ (or equivalently, ν), and assume only that $\nu \in \Lambda$ is K-dominant. It is not hard to see that this holds whenever λ is dominant and regular for G, but the converse fails.

Thus we seek to compute $B_\Psi(\mu, \nu)$ for all K-dominant $\mu, \nu \in \Lambda$.

2. The Classification of $\mathbb{Z}/2\mathbb{Z}$-Gradings

The complexity of the Blattner Formula varies greatly depending on the geometry of the positive non-compact roots Ψ; this geometry is slightly easier to understand if we analyze the structure of $\mathbb{Z}/2\mathbb{Z}$-gradings of Φ_G.

Any $\mathbb{Z}/2\mathbb{Z}$-grading is completely determined by its values on the simple roots, and conversely, any choice of parities for the simple roots extends to a grading. Thus there are 2^n gradings of Φ_G, where $n = \operatorname{rk} T = \operatorname{rk} \Phi_G$.

Claim (presumably well-known). *If G is simple and $\widetilde{\alpha}_G = c_1\alpha_1 + \cdots + c_n\alpha_n$ is the highest G-root (where the α_i's are the simple roots of G), then each $\mathbb{Z}/2\mathbb{Z}$-grading of G may be classified as belonging to one of the following three types:*

(i) The trivial case. Every root is compact, $G = K$, and $\Psi = \varnothing$.

(ii) The minuscule (or Hermitian symmetric) case. The rank of Φ_K is $n-1$, and the set of simple roots for K is some $W(G)$-conjugate of $\{\alpha_1, \ldots, \alpha_n\} - \{\alpha_i\}$ for some index i such that $c_i = 1$. Furthermore, there is a co-weight ω (a $W(G)$-conjugate of the i-th fundamental co-weight, and hence minuscule[1]) such that $\Phi = \{\alpha \in \Phi_G : \langle \alpha, \omega \rangle = 0\}$ and $\Psi = \{\alpha \in \Phi_G^+ : \langle \alpha, \omega \rangle = \pm 1\}$.

(iii) The classical case. The rank of Φ_K is n, and the set of simple roots for K is some $W(G)$-conjugate of $\{\alpha_1, \ldots, \alpha_n, -\widetilde{\alpha}_G\} - \{\alpha_i\}$ for some index i such that $c_i = 2$. Furthermore, there is a co-weight ω (a $W(G)$-conjugate of the i-th fundamental co-weight) such that $\Phi = \{\alpha \in \Phi_G : \langle \alpha, \omega \rangle = 0 \text{ or } \pm 2\}$ and $\Psi = \{\alpha \in \Phi_G^+ : \langle \alpha, \omega \rangle = \pm 1\}$.

Note that in case (ii), the Dynkin diagram of Φ_K is obtained by deleting the i-th node of the Dynkin diagram of Φ_G, whereas in case (iii), it is obtained by deleting the i-th node of the *extended* Dynkin diagram of Φ_G.

Sketch of Proof. The $\mathbb{Z}/2\mathbb{Z}$-gradings of Φ_G may be indexed by $\Lambda_G^\vee / 2\Lambda_G^\vee$, the grading corresponding to $\omega \in \Lambda_G^\vee$ being given by

$$\varepsilon(\beta) := \langle \beta, \omega \rangle \pmod 2.$$

The root system Φ_K is determined by the choice of ω, and is constant up to isomorphism on each $W(G)$-orbit in $\Lambda_G^\vee / 2\Lambda_G^\vee$, or equivalently, on the $\widehat{W}(G)$-orbits in $(1/2)\Lambda_G^\vee$, where $\widehat{W}(G) = \Lambda_G^\vee \rtimes W$ denotes the extended affine Weyl group associated to Φ_G. Furthermore, the stabilizer of $\omega/2$ in $\widehat{W}(G)$ is a finite Weyl group (in fact, a conjugate of some parabolic subgroup of the "ordinary" affine Weyl group associated to Φ_G) whose root system is Φ_K.

Certainly every $\widehat{W}(G)$-orbit hits the fundamental alcove

$$A_0 = \{\lambda : 0 \leqslant \langle \alpha, \lambda \rangle \leqslant 1 \text{ for all } \alpha \in \Phi_G\},$$

so we may assume $\omega/2 \in A_0$; thus, ω is dominant and $\langle \widetilde{\alpha}_G, \omega \rangle \leqslant 2$. Writing $\omega = a_1\omega_1 + \cdots + a_n\omega_n$, where $a_i \in \mathbb{N}$ and ω_i denotes the i-th fundamental co-weight, we obtain $a_1c_1 + \cdots + a_nc_n \leqslant 2$. Renumbering if necessary, we must have either

1. $\omega = 0$, or
2. $c_1 = 1$ and $\omega = \omega_1$, or
3. $c_1 = c_2 = 1$ and $\omega = \omega_1 + \omega_2$, or
4. $c_1 = 1$ and $\omega = 2\omega_1$, or
5. $c_1 = 2$ and $\omega = \omega_1$.

Cases 1 and 4 both imply $\omega \in 2\Lambda_G^\vee$, and yield a trivial grading (type (i)).

[1] Recall that a co-weight ω is minuscule if $\langle \alpha, \omega \rangle \in \{0, \pm 1\}$ for all $\alpha \in \Phi_G$.

Case 3 reduces to Case 2: if $c_1 = c_2 = 1$, then ω_1 and ω_2 are both minuscule, and so is $\omega_1 - \omega_2$ (exercise!). Hence, the $\widehat{W}(G)$-orbit of $\omega/2$ is generated by some (dominant) half-minuscule weight as in Case 2, and we have a grading of type (ii). In particular, note that if ω is dominant minuscule, then $\omega/2$ is the midpoint of the vertices 0 and ω of A_0, and thus has a $\widehat{W}(G)$-stabilizer that is a maximal parabolic subgroup of $W(G)$, leading to the description in (ii).

In Case 5, $\omega/2$ is a vertex of A_0, so it has a stabilizer that is obtained by deleting a node from the extended diagram as described in (iii). □

3. Generalized Partition Functions and Multiplicities

In order to analyze the problem of computing $B_\Psi(\mu,\nu)$, we find it helpful to throw away as much structure as possible. This makes it easier to identify the essential features.

We will start by assuming that Φ is a (crystallographic) root system with positive roots Φ^+, co-roots $\Phi^\vee$, and Weyl group W. As usual, ρ denotes half the sum of the positive roots. Second, we assume that Λ is a lattice that contains Φ and is dual-compatible with it in the sense that $\langle \lambda, \alpha^\vee \rangle \in \mathbb{Z}$ for all $\lambda \in \Lambda$ and all co-roots $\alpha^\vee$. Third, we assume merely Ψ is a finite subset or multisubset (i.e., subset with repetitions) of Λ.

These given, we introduce a graded partition function $\mathcal{P}_\Psi(\gamma;t)$ via the formal expansion

$$\prod_{\lambda\in\Psi} \frac{1}{1-te^{\lambda}} = \sum_{\gamma\in\Lambda} \mathcal{P}_\Psi(\gamma;t)e^{\gamma}. \tag{2}$$

Thus the coefficient of t^l in $\mathcal{P}_\Psi(\gamma;t)$ is the number of ways partition γ into an unordered sum of l weights $\lambda \in \Psi$. By analogy with (1), we define

$$B_\Psi(\mu,\nu;t) := \sum_{w\in W} \operatorname{sgn}(w)\mathcal{P}_\Psi(w(\mu+\rho)-(\nu+\rho);t) \qquad (\mu,\nu\in\Lambda). \tag{3}$$

We now consider the problem of computing $B_\Psi(\mu,\nu;t)$ for all dominant $\mu,\nu \in \Lambda$.

To translate back to the discrete series context of §1, note that Φ, Φ^+, $\Phi^\vee$, W, and ρ may be viewed as being inherited from K, whereas Λ is inherited from the weight lattice of G. We have thrown away most of the structure provided by G, including the notions of $\mathbb{Z}/2\mathbb{Z}$-grading and compact vs. non-compact roots, but vestiges remain in the virtually unconstrained choice of Ψ. Note also that we now have only one notion of dominance, the one implicit in the choice of Φ^+.

It is not surprising at this point that we have thrown away too much structure; we will restore what is necessary or convenient in several subsequent steps.

Remarks. (a) At this level of generality, $\mathcal{P}_\Psi(\gamma;t)$ and $B_\Psi(\mu,\nu;t)$ belong to $\mathbb{Z}[[t]]$; they need not be polynomials. In fact, $\mathcal{P}_\Psi(\gamma;t)$ (and hence also $B_\Psi(\mu,\nu;t)$) is a polynomial for all $\gamma \in \Lambda$ if and only if Ψ is *pointed*; i.e., the cone generated by Ψ contains no 1-dimensional subspace, or equivalently, there is some $\delta \in \Lambda^\vee$ such that $\langle \lambda, \delta \rangle > 0$ for all $\lambda \in \Psi$.

(b) In the discrete series context, it is obvious that Ψ is always pointed, being a subset of the positive roots of G.

(c) There is a general theory of lattice points in polytopes (e.g., see §4.6 of [**EC1**]). In our context, it implies that $\mathcal{P}_\Psi(\gamma;t)$ and $B_\Psi(\mu,\nu;t)$ are always rational functions of t with poles at roots of unity.

(d) There are some reasonably sophisticated algorithms for counting lattice points in polytopes. They run in polynomial time in fixed dimensions, and might be useful.

(e) In some sense, we can always (1) eliminate t, and (2) reduce to the pointed case. Indeed, if we replace $\Lambda \to \mathbb{Z}\theta \oplus \Lambda$, $\Psi \to \theta + \Psi$ and set $t = e^\theta$, where θ is some neutral (i.e., W-fixed) element, then it is easy to check that

$$B_\Psi(\mu,\nu;t) = \sum_{m \geqslant 0} B_{\theta+\Psi}(\mu + m\theta, \nu; 1)t^m,$$

and $\theta + \Psi$ is necessarily pointed. However, this slices up the data in ways that may not be convenient.

Allowing Ψ to run wild over the lattice Λ and introducing the parameter t may seem gratuitous at this point, but we have a few reasons for doing this.

First, alternating sums of partition functions are ubiquitous in representation theory. Aside from the Blattner Formula, we would be negligent not to mention that when $\Psi = \Phi^+$, $B_\Psi(\mu,\nu;1)$ becomes Kostant's formula for the multiplicity of the weight ν in an irreducible representation of highest (dominant) weight μ. It seems worthwhile to assemble a large class of such alternating sums in one family.

Second, as we shall discuss in more detail later (see the last remark in §4), the discrete series cases of type (ii) (i.e., Hermitian symmetric) may be viewed as starting from a possibly *non-pointed* case (Remark (b) notwithstanding), having been converted to a pointed case implicitly via the transformation discussed in Remark (e). Understanding these cases as disguised versions of non-pointed cases turns out to be a useful point of view.

Third (and most important), the recurrence we present in §4 requires t.

Fourth, we find it intriguing (not to mention theoretically useful) that in the discrete series cases, the polynomials $B_\Psi(\mu,\nu;t)$ seem to always have nonnegative coefficients, assuming μ and ν are both dominant.

Research Topic. *Identify sufficient conditions on Ψ so that $B_\Psi(\mu,\nu;t)$ has nonnegative coefficients for all dominant $\mu,\nu \in \Lambda$.*

Let $\mathfrak{k}$ denote a complex semisimple Lie algebra with root system Φ and $\mathfrak{b}$ a Borel subalgebra of $\mathfrak{k}$ whose root spaces are the ones indexed by our chosen positive roots.

False Conjecture. *If $\sum_{\lambda\in\Psi} e^\lambda$ is the character of a $\mathfrak{b}$-submodule of a $\mathfrak{k}$-module of finite dimension, then $B_\Psi(\mu,\nu;t)$ should have $\geqslant 0$ coefficients for all dominant $\mu,\nu \in \Lambda$.*

It is easy to see that the hypothesis of the conjecture holds for all of the discrete series cases, and it is not hard to prove the conjecture for $\mathfrak{k} = sl(2)$. It is also true when $\Psi = \Phi^+$.

Indeed, it is known that $B_{\Phi^+}(\mu,\nu;t)$ (after suitable renormalization) is a Kazhdan-Lusztig polynomial for the associated affine Weyl group.

However, DAV has pointed out that a counterexample is obtained by deleting the lowest weight from the third symmetric power of the defining representation of $sl(3)$.

Nevertheless, we believe that there should be a natural (large) class of $\mathfrak{b}$-modules such that nonnegativity holds whenever $\sum_{\lambda\in\Psi} e^\lambda$ is the character of a member of the class.

In an email, DAV indicated that under the hypothesis of the False Conjecture, work of Griffiths should imply the nonnegativity of the coefficients for all dominant μ and all ν such that $\nu+2\rho-2\eta$ is dominant and regular, where η denotes half the sum of Ψ.

4. A Differential Recurrence

Continuing the setting of §3, notice that logarithmic differentiation of (2) yields

$$\frac{\partial}{\partial t}\prod_{\lambda\in\Psi}\frac{1}{1-te^\lambda}=\sum_{\lambda\in\Psi}\frac{e^\lambda}{1-te^\lambda}\cdot\prod_{\lambda\in\Psi}\frac{1}{1-te^\lambda}.$$

Using $'$ as an abbreviation for $\partial/\partial t$, it follows that

$$\mathcal{P}'_\Psi(\gamma;t)=\sum_{\lambda\in\Psi}\sum_{i\geqslant 1}t^{i-1}\mathcal{P}_\Psi(\gamma-i\lambda;t)\qquad(\gamma\in\Lambda).\tag{4}$$

Also, the constant term is easy: $\mathcal{P}_\Psi(\gamma;0)=\delta_{\gamma,0}$. Bearing in mind the linearity of this relation, we may immediately deduce the following identity from (3).

Proposition. *For all $\mu,\nu\in\Lambda$, we have*

$$B'_\Psi(\mu,\nu;t)=\sum_{\lambda\in\Psi}\sum_{i\geqslant 1}t^{i-1}B_\Psi(\mu,\nu+i\lambda;t).$$

Furthermore, $B_\Psi(\mu,\nu;0)=\begin{cases}\mathrm{sgn}(w) & \textit{if } w(\mu+\rho)=\nu+\rho \textit{ for some } w\in W,\\ 0 & \textit{otherwise.}\end{cases}$

Since the constant terms of $B_\Psi(\mu,\nu;t)$ and $\mathcal{P}_\Psi(\gamma;t)$ are easy to compute, it is clearly sufficient to determine their derivatives, and then integrate.

Remarks. (a) With no assumptions about Ψ, the above formula is not particularly useful for computing $B_\Psi(\mu,\nu;t)$. For example, if Ψ is not pointed, it may happen that $B_\Psi(\mu,\nu;t)$ is nonzero for all $\mu,\nu\in\Lambda$, and there is no entry point for starting a recurrence.

(b) We stole the basic idea for this (not yet a) recurrence from a 1995 paper of Broer [**B**], who used it in the special case $\Psi=\Phi^+$. This is not even the first time we have stolen this idea—we first used it in a 1998 paper on computational methods in representation theory.

(c) Even if $\nu\in\Lambda$ is dominant, terms of the form $\nu+i\lambda$ ($i\geqslant 1$, $\lambda\in\Psi$) may easily take us outside of the dominant chamber. In these cases, even in the context of the False Conjecture, there is no expectation that the coefficients of individual summands on the right side of the proposition should be nonnegative.

Assumption 1. *From now on, we assume that Ψ is pointed.*

Note that this is valid in the discrete series cases.

It follows that we may partially order Λ by defining

$$\mu \geqslant \nu \quad \Leftrightarrow \quad \mu - \nu \in \mathbb{N}\Psi.$$

In fact, yet another equivalent characterization of Ψ being pointed is that the above definition of $\geqslant$ is legitimate (no cycles).

It is clear that $\mathcal{P}_\Psi(\gamma; t) \neq 0$ if and only if $\gamma \geqslant 0$, and hence (see (3))

$$B_\Psi(\mu, \nu; t) \neq 0 \quad \Rightarrow \quad w(\mu + \rho) \geqslant (\nu + \rho) \text{ for some } w \in W. \tag{5}$$

It follows that the Proposition may be rewritten as

$$B'_\Psi(\mu, \nu; t) = \sum_{\nu + i\lambda \in S_\Psi(\mu,\nu)} t^{i-1} B_\Psi(\mu, \nu + i\lambda; t), \tag{6}$$

where $i \geqslant 1$ and $\lambda \in \Psi$ (as usual), and

$$S_\Psi(\mu, \nu) := \{\nu' \in \Lambda : w(\mu + \rho) \geqslant \nu' + \rho \geqslant \nu + \rho \text{ for some } w \in W\}.$$

Since W is finite and $\geqslant$ is locally finite, it follows that $S_\Psi(\mu, \nu)$ is finite. Furthermore, $S_\Psi(\mu, \nu)$ shrinks as ν moves higher in the partial order, so (6) qualifies as a valid recurrence for computing $B_\Psi(\mu, \nu; t)$.

Problem. *Given (dominant) $\mu, \nu \in \Lambda$,*

(a) *quickly test membership of ν' in $S_\Psi(\mu, \nu)$.*

(b) *quickly generate $S_\Psi(\mu, \nu)$.*

Part (a) suffices for generating the summands of (6); for each $\lambda \in \Psi$, we terminate the sum as soon as we reach an index i such that $\nu + i\lambda \notin S_\Psi(\mu, \nu)$. Note that subsequent iterations involving the computation of (say) $B_\Psi(\mu, \nu'; t)$ may hypothetically require testing membership of some ν'' in $S_\Psi(\mu, \nu')$. However, this occurs only when $\nu'' \geqslant \nu'$, in which case this is equivalent to membership of ν'' in $S_\Psi(\mu, \nu)$.

Barring some miraculous strategy that provides for extensive pruning of the recurrence tree, the Obvious Way to use (6) to compute $B_\Psi(\mu, \nu; t)$ will involve saving the values $B_\Psi(\mu, \nu'; t)$ for all $\nu' \in S_\Psi(\mu, \nu)$. In that case, a solution of (b) may be used as the basis for a cheap solution of (a). In any case, the space requirements of the Obvious Way will be proportional to $|S_\Psi(\mu, \nu)|$.

We have good ideas for solving the above Problem, but our current Maple code for computing discrete series multiplicities[2] instead takes an easier (and slower, less space efficient) approach based on the following.

[2]See ⟨`www.math.lsa.umich.edu/~jrs/data/blattner`⟩.

Assumption 2. *We assume Ψ is positively pointed; i.e., $\Psi \cup \Phi^+$ is pointed.*

Note that in the discrete series cases, $\Psi \cup \Phi^+$ is the set of positive G-roots (hence pointed), so this is an assumption we don't mind making.

Given this hypothesis, we can define a stronger partial ordering of Λ by setting

$$\mu \succcurlyeq \nu \quad \Leftrightarrow \quad \mu - \nu \in \mathbb{N}(\Psi \cup \Phi^+).$$

Dominant weights are $\succcurlyeq$-maximums within their W-orbits, so (5) implies

$$B_\Psi(\mu, \nu; t) \neq 0 \quad \Rightarrow \quad \mu \succcurlyeq \nu,$$

assuming that μ is dominant. It follows that

$$S_\Psi(\mu, \nu) \subseteq \{\nu' \in \Lambda : \mu \succcurlyeq \nu' \geqslant \nu\},$$

and thus a cruder version of (6) may be obtained by using this bound as a replacement for $S_\Psi(\mu, \nu)$. This is easier to implement, since we can generate the weights $\nu' \geqslant \nu$ while solving the recurrence, and truncate the sum whenever $\mu - \nu'$ crosses a wall of the cone generated by $\Psi \cup \Phi^+$. (Of course, this only approximates the relation $\mu \succcurlyeq \nu'$.)

In general, it seems to be hard to accurately estimate the number of lattice points in $S_\Psi(\mu, \nu)$ or $\{\nu' : \mu \succcurlyeq \nu' \geqslant \nu\}$. However, in the discrete series cases the cone $\mathbb{N}(\Psi \cup \Phi^+)$ is simplicial, being generated by the simple G-roots, so we have the easy bounds

$$|S_\Psi(\mu, \nu)| \ \leqslant \ |\{\nu' : \mu \succcurlyeq \nu' \geqslant \nu\}| \ \leqslant \ |\{\nu' : \mu \succcurlyeq \nu' \succcurlyeq \nu\}| \ = \ (1 + a_1) \cdots (1 + a_n),$$

where $a_1, \ldots, a_n$ are the coordinates of $\mu - \nu$ in terms of the simple G-roots. Our intuition is that these bounds are usually far from tight; however, if every simple G-root is noncompact, then $\mathbb{N}(\Psi \cup \Phi^+) = \mathbb{N}\Psi$, the partial orders $\succcurlyeq$ and $\geqslant$ coincide, and these allegedly non-tight bounds are equalities.

The following is useful for theoretical purposes.

Assumption 3. *It would be nice to assume that Ψ is upward-saturated; i.e., for all $\lambda \in \Psi$, $\alpha \in \Phi^+$ and $k \in \mathbb{Z}$ such that $\langle \lambda, \alpha^\vee \rangle \leqslant -k \leqslant 0$, we have $\lambda + k\alpha \in \Psi$.*

The above definition is sloppy if Ψ is a multiset, but can be fixed by inserting multiplicity obfuscations. Again, it is easily seen to be valid in the discrete series cases. If Φ and Ψ are the even and (positive) odd roots of a $\mathbb{Z}/2\mathbb{Z}$-grading of some larger root system, then the usual saturation property of root strings implies (in the context of the above assumption) that $\lambda + k\alpha$ is necessarily a root, and hence it must be in Ψ, since λ is odd and α is even. Similarly, this assumption also holds in the context of the False Conjecture.

Claim. *If Assumption 3 holds, then* $B_\Psi(\mu,\nu;t)\neq 0 \Rightarrow \mu \geqslant \nu$ *for dominant* $\mu,\nu\in\Lambda$.

Proof. If $B_\Psi(\mu,\nu;t)\neq 0$, then there is some $w\in W$ such that $w(\mu+\rho)-(\nu+\rho)=\lambda_1+\cdots+\lambda_l$ for suitable $\lambda_i\in\Psi$. If $w=1$, we are done. Otherwise, given that μ is dominant, $w(\mu+\rho)$ cannot be dominant, so there is some $\alpha\in\Phi^+$ such that

$$\langle\lambda_1+\cdots+\lambda_l,\alpha^\vee\rangle=\langle w(\mu+\rho)-(\nu+\rho),\alpha^\vee\rangle<\langle w(\mu+\rho),\alpha^\vee\rangle<0,$$

the first inequality being based on the fact that ν is dominant. Thus we may select integers k_i such that $\langle\lambda_i,\alpha^\vee\rangle\leqslant -k_i\leqslant 0$ and

$$\langle w(\mu+\rho),\alpha^\vee\rangle=-(k_1+\cdots+k_l),$$

in which case $\lambda_i+k_i\alpha\in\Psi$ by upper saturation, and

$$s_\alpha w(\mu+\rho)-(\nu+\rho)=(\lambda_1+k_1\alpha)+\cdots+(\lambda_l+k_l\alpha).$$

The result now follows by induction with respect to the length of w. □

Remark. Even in the discrete series cases, it can sometimes happen that for a fixed dominant $\mu\in\Lambda$, there are infinitely many dominant $\nu\in\Lambda$ such that $\mu\geqslant\nu$. (However, there can only be finitely many such ν that arise as the lowest K-type of some discrete series.) This is computationally annoying, since it would be nice to have a program that takes some dominant μ as input, and computes $B(\mu,\nu;t)$ for all dominant $\nu\leqslant\mu$. Unfortunately, this is not always a finite problem.

For example, let $\lambda\mapsto\bar\lambda$ denote orthogonal projection onto $\operatorname{Span}\Phi$, and consider the image $\bar\Psi$ of Ψ. If $\bar\Psi$ is not pointed, then there is a nontrivial relation $\sum a_i\bar\lambda_i=0$ ($a_i\in\mathbb{N}$, $\lambda_i\in\Psi$), in which case $\theta:=\sum a_i\lambda_i$ is a nonzero neutral element (hence anti-dominant), and $\mu\geqslant\mu-\theta\geqslant\mu-2\theta\geqslant\cdots$ is an infinite dominant chain for any dominant $\mu\in\Lambda$.

More specifically, in the Hermitian symmetric cases (see part (ii) of the Claim in §2), recall that there is some G-minuscule $\omega\in\Lambda^\vee$ such that

$$\Phi=\{\alpha\in\Phi_G:\langle\alpha,\omega\rangle=0\},\quad \pm\Psi=\{\beta\in\Phi_G:\langle\beta,\omega\rangle=\pm1\}.$$

If it happens that Ψ includes a full W-orbit (for example, ω is G-dominant and hence $\Psi=\{\beta\in\Phi_G:\langle\beta,\omega\rangle=1\}$), then $\bar\Psi$ includes a W-orbit in $\operatorname{Span}\Phi$, and hence cannot be pointed. In particular, the ugliness of the previous paragraph applies.

Conversely, if Ψ is upper saturated (Assumption 3), one can show that if $(\Psi\cup\Phi^+)^-$ is not pointed, then $\bar\Psi$ is not pointed, so the only other possibility is that $(\Psi\cup\Phi^+)^-$ *is* pointed. In that case, we can find $\delta\in\operatorname{Span}\Phi^\vee$ such that $\langle\lambda,\delta\rangle>0$ for all $\lambda\in\Psi\cup\Phi^+$. However, if $\nu\in\Lambda$ is dominant, then so is $\bar\nu$, and hence $\bar\nu$ is in the $\mathbb{Q}^+$-span of Φ^+ (this is equivalent to the fact that the inverse of a Cartan matrix is nonnegative). Therefore, $\langle\nu,\delta\rangle=\langle\bar\nu,\delta\rangle\geqslant 0$ for all dominant $\nu\in\Lambda$. This proves that there can only finitely many dominant $\nu\leqslant\mu$, since each such ν must have the form $\nu=\mu-\sum a_i\lambda_i$ ($a_i\in\mathbb{N}$, $\lambda_i\in\Psi$) and there are only finitely many integer points $(a_1,a_2,\dots)$ in the simplex $\sum a_i\langle\lambda_i,\delta\rangle\leqslant\langle\mu,\delta\rangle$.

5. Alternatives and Other Remarks

The best feature of the algorithm discussed in the previous section is that the Weyl group W and the partition function $\mathcal{P}_\Psi$ have been (mostly) finessed away. The worst feature is that the space requirements could be huge, although we expect that this might be a problem only for discrete series involving G of type $\mathcal{E}_8$, in which case the root system Φ (assuming it is proper) is of type $\mathcal{D}_8$ or $\mathcal{A}_1 \oplus \mathcal{E}_7$.

A third significant aspect is that the algorithm is tailor-made for computing $B_\Psi(\mu, \nu; t)$ for μ fixed and varying ν. Whether this is a bug or a feature will depend on the application. For computing within a single discrete series, this is a bug, since this means fixing the Harish-Chandra parameter of the series (hence fixing ν), and computing $B_\Psi(\mu, \nu; 1)$ for (say) all dominant μ up to some height h. This means throwing away the table of values for $B_\Psi(\mu, \nu'; t)$ each time we change μ.

An alternative approach that we have not (yet) implemented would be to use similar techniques to compute the graded partition function $\mathcal{P}_\Psi(\gamma; t)$ via (4), and then compute $B_\Psi(\mu, \nu; t)$ directly from (3). The analogue of the Problem from §4 in this approach is:

Problem. *Given $\gamma, \mu, \nu \in \Lambda$ with μ, ν dominant,*

(a) *quickly test $\gamma \geqslant 0$,*
(b) *quickly visit all members of the W-orbit of $\mu + \rho$ that are $\geqslant \nu + \rho$.*

We expect that the two algorithms should have roughly equivalent time and space requirements for the computation of a single instance of $B_\Psi(\mu, \nu; t)$. On the other hand, the big advantage of this new approach is that we get to re-use and add to the table of values for $\mathcal{P}_\Psi(\gamma; t)$ during the computation of each $B_\Psi(\mu, \nu; t)$.

Finally, let us mention that if the space requirements do turn out to be a major obstacle, there is still the possibility of using sophisticated lattice-point counting methods to compute each value of the partition function $\mathcal{P}_\Psi(\gamma; 1)$ separately from scratch. These methods have minimal space requirements (good), but we cannot afford to save previous values (bad), so they are likely to be much slower if we accumulate the cost of computing all multiplicities up to some height.

References

[**B**] A. Broer, The sum of generalized exponents and Chevalley's restriction theorem for modules of covariants, *Indag. Math.* **6** (1995), 385–396.

[**HS**] H. Hecht and W. Schmid, A proof of Blattner's conjecture, *Invent. Math.* **31** (1975), 129–154.

[**EC1**] R. P. Stanley, "Enumerative Combinatorics, Vol. 1," Cambridge Univ. Press, Cambridge, 1997.