BRANCHING LAWS AIM LECTURES JULY 2003 BY PETER TRAPA

NOTES BY DAVID A. VOGAN, JR.

July 30, 2003

ABSTRACT. A basic problem in the representation theory of a compact Lie group is to calculate the restriction of an irreducible representation of K to a closed subgroup. Most classical results on this problem concern connected groups. We'll recall some of those, and talk about some of the modifications needed to work with disconnected groups.

1. WEIGHT MULTIPLICITIES

Suppose $\mathfrak{g}$ is a complex semisimple Lie algebra and $\mathfrak{h}$ a Cartan subalgebra. If we fix a Borel subalgebra $\mathfrak{b} = \mathfrak{h} + \mathfrak{n}$ of $\mathfrak{g}$, then the irreducible finite-dimensional representations of $\mathfrak{g}$ are parametrized by dominant integral weights $\mathfrak{h}^*$. Write V^λ for the representation of highest weight λ. The restriction of V^λ to $\mathfrak{h}$ is a direct sum of weight spaces:

$$V^\lambda|_{\mathfrak{h}} = \bigoplus_{\mu in \mathfrak{h}^*} m_{\lambda\mu}\mathbb{C}_\mu, \tag{1.1}$$

with $m_{\lambda\mu}$ a non-negative integer (the weight multiplicity). The most basic classical branching problem is to calculate these multiplicities. One solution is Kostant's multiplicity formula:

$$m_{\lambda\mu} = \sum_{w\in W} (-1)^w P_{\mathfrak{n}}(w(\lambda+\rho) - (\mu+\rho)). \qquad (1.2)(a)$$

Here is what the terms on the right mean. $(-1)^w$ is the determinant of the action of w on $\mathfrak{h}^*$ (which is ± 1). The function

$$P_{\mathfrak{n}}\colon \mathfrak{h}^* \to \mathbb{N} \qquad (1.2)(b)$$

is Kostant's partition function: its value at γ is the number of distinct expressions

$$\gamma = \sum_{\alpha\in\Delta(\mathfrak{n},\mathfrak{h})} n_\alpha \alpha$$

Typeset by $\mathcal{AMS}$-TeX

of γ as a non-negative integer combination of positive roots. There are many proofs of Kostant's formula, using for example the BGG resolution of V^λ, or the smoothness of the flag variety. A very algebraic reference is [Hump].

Knowledge of weight multiplicities leads directly to solutions of many other branching problems. An example of a problem that it does *not* solve directly is this: calculate the representation of W on the zero weight space of V^λ. The heart of the difficulty is that the group W (or the normalizer of H in G) is disconnected.

Before moving to other topics, notice that two basic ingredients in Kostant's formula are the Cartan-Weyl parametrization of representations by dominant weights, and the action of W on weights.

2. K multiplicities.

Suppose G is a real reductive group with maximal compact subgroup K and complexified Lie algebra $\mathfrak{g}$. If X is a $(\mathfrak{g}, K)$-module of finite length, then by analogy with (1.1) we have

$$X|_K = \bigoplus_{\lambda\in\widehat{K}} m_X^\lambda E^\lambda, \tag{2.1}$$

with m_X^λ a non-negative integer. The fundamental branching problem in this setting is computation of m_X^λ. Here is a way to do that. Suppose for definiteness that X is irreducible. Write $\mathcal{P}ar$ for the (finite) set of Langlands parameters for representations of infinitesimal character equal to that of X. To each $\delta \in \mathcal{P}ar$ there is associated a "standard Harish-Chandra module"

$$X(\delta) = \mathrm{Ind}_{MAN}^G(X_{ds}^M \otimes e^\nu \otimes 1). \qquad (2.2)(a)$$

Here X_{ds}^M is a Harish-Chandra module for M in the limits of the discrete series, and $\nu \in \mathfrak{a}^*$. Each such standard module has a unique Langlands subquotient $\overline{X}(\delta)$. (Uniqueness of the Langlands subquotient depends on choosing exactly the right parameter set; it is not true for arbitrary representations of the form $(2.2)(a)$.) The irreducible modules $\overline{X}(\delta)$ are inequivalent, and every irreducible Harish-Chandra module of the same infinitesimal character as X is equivalent to one of them.

The Kazhdan-Lusztig conjectures (which are theorems at least for connected linear groups) allow us to write

$$\overline{X}(\gamma) = \sum_{\delta\in\mathcal{P}ar} m_{\gamma\delta} X(\delta), \qquad (2.2)(b)$$

with $m_{\gamma\delta}$ an integer given by a reasonable[1] algorithm (see the talk by Fokko du Cloux).

The problem of computing $X|_K$ is reduced by $(2.2)(b)$ to the problem of computing each $X(\delta)|_K$. Because of $(2.2)(a)$,

$$X(\delta)|_K = \mathrm{Ind}_{M\cap K}^K(X_{ds}^M|_{M\cap K}).$$

Because of Frobenius reciprocity, this calculation in turn can be broken into two parts:

(1) compute $X_{ds}^M|_{M\cap K}$; and
(2) compute $V^\lambda|_{M\cap K}$ for every irreducible representation $\lambda \in \widehat{K}$.

[1] The definition of "reasonable" shifted rather drastically on the occasion of the publication of the original article of Kazhdan and Lusztig.

The first of these problems can be approached using the Blattner formula, which is closely analogous to (and as computable as) the Kostant multiplicity formula (1.2). There are complications arising from the disconnectedness of M, but never mind. We will therefore concentrate on the second, which is entirely a branching problem for compact groups. Here are two illuminating examples.

Example 2.3. Suppose $G = U(p, q)$, so that $K = U(p) \times U(q)$. The choices for M are parametrized by a non-negative integer a less than or equal to both p and q. Then

$$M \cap K = U(p-a) \times U(1)^a_\Delta \times U(q-a).$$

This is the subgroup of $[U(p-a) \times U(1)^a] \times [U(1)^a \times U(q-a)]$ obtained by taking the diagonal copy of $U(1)$ in each of the a copies of $U(1) \times U(1)$. Branching to these subgroups is easily handled by mild generalizations of Kostant's multiplicity formula.

Example 2.4. Suppose $G = O(p, q)$, so that $K = O(p) \times O(q)$. Some of the possibilities for M (there are more) are parametrized by a non-negative integer a less than or equal to both p and q. Then

$$M \cap K = O(p-a) \times O(1)^a_\Delta \times O(q-a).$$

Formally this example looks very similar to the preceding one, but technically it is very different: the subgroup $M \cap K$ has (usually) 2^{a+2} connected components, so that one cannot expect to get a branching law using only Lie algebra ideas.

Example 2.5. Suppose G is the split real linear group of type E_8, so that $K = \mathrm{Spin}(16)/\{1, \epsilon\}$. If MAN is the Borel subgroup of G, then $M = M \cap K \simeq (\mathbb{Z}/2\mathbb{Z})^8$. At most 16 of the 256 elements of M belong to a common maximal torus of K, so again Lie algebra methods cannot provide much help with the branching problem.

3. Branching for connected compact groups.

In this section we will recall Kostant's general branching theorem for connected compact Lie groups. We can work entirely with Lie algebras; so fix a complex reductive Lie algebra $\mathfrak{g}$, and let $\mathfrak{k}$ be a subalgebra reductive in $\mathfrak{g}$. (The connection with compact groups appears when $\mathfrak{g}$ and $\mathfrak{k}$ are complexified Lie algebras for compact connected groups $G \supset K$.) There is an $\mathrm{ad}(\mathfrak{k})$-invariant complement $\mathfrak{p}$ for $\mathfrak{k}$:

$$\mathfrak{g} = \mathfrak{k} \oplus \mathfrak{p}. \qquad (3.1)(a)$$

Choose a Cartan subalgebra $\mathfrak{t} \subset \mathfrak{k}$, and extend it to a Cartan subalgebra $\mathfrak{h}$ of $\mathfrak{g}$:

$$\mathfrak{h} = \mathfrak{t} \oplus (\mathfrak{h} \cap \mathfrak{p}). \qquad (3.1)(b)$$

There is a natural restriction map

$$\mathrm{res}\colon \mathfrak{h}^* \to \mathfrak{t}^*, \qquad \lambda \mapsto \overline{\lambda}. \qquad (3.1)(c)$$

Define

$$\mathfrak{l} = \mathfrak{g}^{\mathfrak{t}}, \qquad (3.1)(d)$$

a Levi subalgebra of $\mathfrak{g}$ (corresponding to the roots of $\mathfrak{h}$ in $\mathfrak{g}$ that restrict to zero on $\mathfrak{t}$). Choose a parabolic subalgebra $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ of $\mathfrak{g}$ in such a way if α and β are any two roots of $\mathfrak{h}$ in $\mathfrak{g}$, then

$$\overline{\alpha} = \overline{\beta} \quad \text{and} \quad \alpha \in \Delta(\mathfrak{u}, \mathfrak{h}) \Rightarrow \beta \in \Delta(\mathfrak{u}, \mathfrak{h}). \tag{3.1)(e}$$

This is certainly possible to arrange. Finally, choose

$$\Delta^+(\mathfrak{g}, \mathfrak{h}) \supset \Delta(\mathfrak{u}, \mathfrak{h}). \tag{3.1)(f}$$

We need a little notation about the Weyl group. For $w \in W(\mathfrak{g}, \mathfrak{h}) = W$, define[2]

$$\Delta_w = \{\alpha \in \Delta^+ \mid w^{-1}\alpha < 0\}. \tag{3.2)(a}$$

Now define

$$W' = \{w \in W \mid \Delta_w \subset \Delta(\mathfrak{u}, \mathfrak{h}). \tag{3.2)(b}$$

The set W' is the set of (minimal length) coset representatives for $W(\mathfrak{l}, \mathfrak{h})$ in W; that is, multiplication defines a bijection

$$W(\mathfrak{g}, \mathfrak{h}) = W(\mathfrak{l}, \mathfrak{h}) \times W'. \tag{3.2)(c}$$

Theorem 3.3 (Kostant)**.** *Suppose V^λ is a finite-dimensional irreducible representation of $\mathfrak{g}$ of highest weight $\lambda \in \mathfrak{h}^*$, and V^ν is a finite-dimensional irreducible representation of $\mathfrak{k}$ of highest weight $\nu \in \mathfrak{t}^*$. Then the multiplicity of V^ν in $V^\lambda|_{\mathfrak{k}}$ is*

$$\sum_{w \in W'} (-1)^w d_\lambda(w) P_{\mathfrak{u}\cap\mathfrak{p}}(\overline{w(\lambda + \rho) - \rho} - \nu).$$

Here $d_\lambda(w)$ is the dimension of the irreducible representation of $\mathfrak{l}$ of highest weight $w(\lambda + \rho) - \rho$; and the partition function

$$P_{\mathfrak{u}\cap\mathfrak{p}} \colon \mathfrak{t}^* \to \mathbb{N}$$

counts expressions in terms of the weights of $\mathfrak{t}$ on $\mathfrak{u} \cap \mathfrak{p}$.

Another approach to this branching problem is due to Patera, Sharp, Moody, and McKay. The idea is to begin with the weight multiplicity formula

$$V^\lambda = \sum_{\mu \in \mathfrak{h}^* \text{dominant}} m_\mu (W \cdot \mu),$$

then to compute the branching

$$(W \cdot \mu)|_{\mathfrak{t}} = \sum_{\mu' \in \mathfrak{t}^* \text{dominant}} m_{\mu'} (W_{\mathfrak{k}}) \cdot \mu',$$

and then to deduce multiplicities of representations of $\mathfrak{k}$.

[2] In the lecture, w was σ, and Δ_w was written W_σ. This notetaker was able to write that down, but finds himself unable to commit it to TEX.

4. Disconnected compact groups.

We mentioned at the end of section 1 that two of the basic ingredients in the Kostant multiplicity formula are the Cartan-Weyl parametrization of $\widehat{K}$ by weights, and the action of the Weyl group on the parameter space. Here we will begin to extend these ingredients to disconnected groups.

So assume that K is a compact Lie group, possibly disconnected. Write $\mathfrak{k}$ for the complexified Lie algebra of K, T_0 for a maximal torus in the identity component K_0, and

$$\mathfrak{b} = \mathfrak{t} + \mathfrak{n} \tag{4.1}$$

for a Borel subalgebra of $\mathfrak{k}$.

Definition 4.2. The *large Cartan subgroup* of K corresponding to $\mathfrak{b}$ is

$$T^+ = \{t \in K \mid \mathrm{Ad}(t)(\mathfrak{b} = \mathfrak{b}\} = N_K(\mathfrak{b}).$$

The *small Cartan subgroup* is

$$T^- = \{t \in K \mid \mathrm{Ad}(t)|\mathfrak{t} \text{ is trivial}\} = Z_K(T_0),$$

the centralizer in K of T_0. A general Cartan subgroup is any group T such that

$$T^- \subset T \subset T^+.$$

If T is any Cartan subgroup of K corresponding to $\mathfrak{b}$, then there is a well-defined character 2ρ of T,

$$2\rho(t) = \det \mathrm{Ad}(t)|_{\mathfrak{n}}, \qquad (4.3)(a)$$

the determinant of the adjoint action of T on $\mathfrak{n}$. Despite the suggestive notation, the character 2ρ need not have a square root. We need such a square root, so we introduce the group

$$\widetilde{T} = \ \rho \text{ cover of } T \ = \{(t, z) \in T \times \mathbb{C}^\times \mid 2\rho(t) = z^2\}. \qquad (4.3)(b)$$

Projection on the first factor defines a two-to-one group homomorphism

$$\pi \colon \widetilde{T} \to T, \qquad \pi(t, z) = t. \qquad (4.3)(c)$$

The kernel of π consists of the identity and $\epsilon = (1, -1)$. There is a short exact sequence

$$1 \longrightarrow \{1, \epsilon\} \longrightarrow \widetilde{T} \xrightarrow{\pi} T \longrightarrow 1. \qquad (4.3)(d)$$

Projection on the second factor defines a character that we call ρ:

$$\rho \colon \widetilde{T} \to \mathbb{C}^\times, \qquad \rho(t, z) = z. \qquad (4.3)(e)$$

The differential of ρ is equal to one half the differential of 2ρ, or one half the sum of the roots of $\mathfrak{t}$ in $\mathfrak{n}$. Furthermore ρ is a *genuine* character of $\widetilde{T}$, by which we mean that $\rho(\epsilon) = -1$.

Definition 4.4. Suppose that $\widetilde{T}$ is the ρ cover of T defined in (4.3). An irreducible representation χ of $\widetilde{T}$ is called *genuine* if $\chi(\epsilon) = -I$. (This is equivalent to assuming that χ does not factor through π to the quotient group T.) The differential of χ is a completely reducible representation of the Lie algebra $\mathfrak{t}$, and therefore a direct sum of various weights in $\mathfrak{t}^*$. The representation χ is called *dominant* if one (or, equivalentally, all) of these weights is dominant. It is called *regular* if one (or, equivalently, all) of these weights is non-zero on every coroot of $\mathfrak{t}$ in $\mathfrak{k}$.

Theorem 4.5. *Suppose K is a compact Lie group and T^+ is a large Cartan subgroup. Let $\widetilde{T^+}$ be the ρ covering of T^+ (cf. (4.3)). Then there is a natural bijection between the set $\widehat{K}$ of irreducible representations of K, and the set of genuine dominant regular irreducible representations of $\widetilde{T^+}$.*

Describing this bijection is fairly easy. If V is an irreducible representation of K, then the subspace $V^{\mathfrak{n}}$ of highest weight vectors is an irreducible representation of T^+. This provides a bijection

$$\widehat{K} \leftrightarrow \text{dominant irreducible representations of } T^+.$$

Next, ρ is a genuine irreducible one-dimensional representation of $\widetilde{T^+}$. It follows that tensoring with ρ provides a bijection

$$\text{irreducible representations of } T^+ \leftrightarrow \text{genuine irreducible representations of } \widetilde{T^+}.$$

This bijection identifies dominant representations on the left with dominant regular representations on the right. Combining these two bijections gives the theorem.

You may wonder why we introduced the covering, which so far has served only to complicate the statement and proof of Theorem 4.5. The reason is that the covering will make possible a statement of the Weyl character formula, in which both numerator and denominator are genuine functions on $\widetilde{T^+}$. (A "genuine function" is one sent to its negative under translation by ϵ.) The actual lecture concluded with a precise and correct statement of the Weyl character formula on the small Cartan subgroup T^-. The notes will conclude instead with a statement of the Weyl character formula on T^+ more or less equivalent to the one described in [Vorange].

Definition 4.6. Suppose K is a compact Lie group, $\mathfrak{b}$ is a Borel subalgebra of $\mathfrak{k}$, and T^+ is the normalizer of $\mathfrak{b}$ in $\mathfrak{k}$. Then the homogeneous space K/T^+ is naturally identified with the flag variety $\mathcal{B}$ of all Borel subalgebras of $\mathfrak{k}$. An element of $k \in K$ is called *regular for* K if some (equivalently every? I'm not sure what to guess) fixed point of k on $\mathcal{B}$ is regular: that is, if the induced action of k on the tangent space at the fixed point does not have 1 as an eigenvalue.

An element $t \in T^+$ is called *regular for* K if the adjoint action $\mathrm{Ad}(t)|_{\mathfrak{n}}$ does not have 1 as an eigenvalue. (In order for this definition to be nicely consistent with the one above, one needs to be able to replace "some" by "every.") Regularity in T^+ is equivalent to the condition

$$\det(I - \mathrm{Ad}(t)|_{\mathfrak{n}}) \neq 0.$$

The *Weyl denominator* is the function

$$\Delta(\tilde{t}) = \rho(\tilde{t})^{-1} \det(I - \mathrm{Ad}(\pi(\tilde{t}))|_{\mathfrak{n}})$$

on $\widetilde{T^+}$. It is non-zero exactly on the preimages of regular elements of T^+.

The (unchecked) hope is that regular elements in T^+ are dense in T^+, and that regular elements of K are dense in K, and that every conjugacy class in K meets T^+. None of this should be very hard; the danger is that it's false.

We turn now to the Weyl character formula. Just as for connected K, the formula expresses the trace of a regular element $t \in T^+$ as a quotient of two terms, each of which depends on a preimage of t in $\widetilde{T^+}$. (Changing the chosen preimage changes the sign of each term, so the quotient depends only on t.) The denominator is the Weyl denominator of Definition 4.6. The numerator is (roughly speaking) a sum over the Weyl group of T_0 in K_0. What makes life complicated is that this Weyl group does not act on T^+. The solution (roughly speaking) is simply to throw away the terms that don't make sense.

Here are the details. Recall that we are fixing $\mathfrak{b} = \mathfrak{t}+\mathfrak{n}$. Write $\mathfrak{n}^-$ for the opposite nil radical (which is automatically preserved by T^+). Suppose $w \in W(\mathfrak{k}, \mathfrak{t})$. Then w defines a second Borel subalgebra

$$\mathfrak{b}^w = \mathfrak{t} + \mathfrak{n}^w, \tag{4.7(a)}$$

corresponding to $w(\Delta^+(\mathfrak{k}, \mathfrak{t})$. This new Borel subalgebra defines a large Cartan subgroup

$$T^{+,w} = N_K(\mathfrak{b}^w). \tag{4.7(b)}$$

The identity component of $T^{+,w}$ is the same torus T_0 as for T^+, but the other components can be different. It's not difficult to see that

$$\mathfrak{n}^w = (\mathfrak{n} \cap \mathfrak{n}^w) \oplus (\mathfrak{n}^- \cap \mathfrak{n}^w), \tag{4.7(c)}$$

and therefore that

$$T^+ \cap T^{+,w} = \{t \in T^+ | \mathrm{Ad}(t) \text{ preserves } \mathfrak{n} \cap \mathfrak{n}^w\}. \tag{4.7(d)}$$

The representations of $T^+ \cap T^{+,w}$ on $\mathfrak{n}^- \cap \mathfrak{n}^w$ and on $\mathfrak{n} \cap \mathfrak{n}^{-,w}$ are contragredient.[3] It follows that on their common domain $T^+ \cap T^{+,w}$, the two characters 2ρ and $2\rho^w$ differ by the *square* of a determinant:

$$2\rho^w(t) = 2\rho(t) \cdot [\det(\mathrm{Ad}(t)|_{\mathfrak{n}^- \cap \mathfrak{n}^w})]^2 \qquad (t \in T^+ \cap T^{+,w}). \tag{4.7(e)}$$

This fact provides a natural isomorphism between the ρ covering of $T^+ \cap T^{+,w}$ and the ρ^w covering,

$$(t, z) \mapsto (t, z \cdot \det(\mathrm{Ad}(t)|_{\mathfrak{n}^- \cap \mathfrak{n}^w}) \qquad (t \in T^+ \cap T^{+,w}).$$

In this way we can regard ρ^w as a character of the ρ covering $(T^+ \cap T^{+,w})^\sim$.

We can now state the *Weyl denominator formula*

$$\Delta(\tilde{t}) = \sum_{\substack{w \in W(\mathfrak{k}, \mathfrak{t}) \\ t \in T^+ \cap T^{+,w}}} (-1)^w \rho^w(\tilde{t}). \tag{4.8}$$

Suppose now that χ is a genuine irreducible representation of T^+. Since it's the end of the day, I will merely assert confidently that there is a natural way to define a genuine representation χ^w of $(T^+ \cap T^{+,w})^\sim$. (The idea is to twist χ by conjugation by an element of K_0 defining the Weyl group element w. This conjugation will carry $T^+ \cap T^{+,w}$ into T^+.)

[3] This is a somewhat subtle point, even in the simplest case $\mathfrak{n}^w = \mathfrak{n}^-$. One way to see it is by introducing appropriate symplectic forms, and using the fact that the determinant of a symplectic linear map is one.

Theorem 4.9 (the Weyl character formula)**.** *Suppose T^+ is a large Cartan subgroup of a compact Lie group K, χ is a genuine dominant irreducible representation of $\widetilde{T^+}$, and V^χ is the corresponding irreducible representation of K. Suppose t is a regular element of T^+, and that $\tilde{t}$ is a preimage of t in the ρ covering $\widetilde{T^+}$. Then the character of V^χ at t is*

$$\Theta^\chi(t) = \left(\sum_{\substack{w \in W(\mathfrak{k},\mathfrak{t}) \\ t \in T^+ \cap T^{+,w}}} (-1)^w \chi^w(t) \right) / \Delta(\tilde{t}).$$

This formula could be proved using the Atiyah-Bott-Lefschetz fixed point formula.[4]

References

[Hump] J.E. Humphreys, *Introduction to Lie Algebras and Representation Theory*, Springer-Verlag, Berlin-Heidelberg-New York, 1972.

[Vorange] D. Vogan, *Unitary Representations of Reductive Lie Groups*, Annals of Mathematics Studies, Princeton University Press, Princeton, NJ, 1987.

[4]But let us not forget the possibility that this tissue of surmise and hypothesis has no basis in reality.