UV home.

Associated Varieties and Unipotent Representations

David A. Vogan, Jr.*
Department of Mathematics
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

1. Introduction.

Suppose $G_{\mathbf{R}}$ is a semisimple Lie group. The philosophy of coadjoint orbits, as propounded by Kirillov and Kostant, suggests that unitary representations of $G_{\mathbf{R}}$ are closely related to the orbits of $G_{\mathbf{R}}$ on the dual $\mathfrak{g}_{\mathbf{R}}^*$ of the Lie algebra $\mathfrak{g}_{\mathbf{R}}$ of $G_{\mathbf{R}}$. One knows how to attach representations to semisimple orbits, but the methods used (which rely on the existence of nice "polarizing subalgebras" of $\mathfrak{g}$) cannot be applied to most nilpotent orbits.

One notion that *is* available for all representations is that of "associated variety." Let $K_{\mathbf{R}}$ be a maximal compact subgroup of $G_{\mathbf{R}}$, and K its complexification. Write $\mathfrak{g}$ for the complexification of $\mathfrak{g}_{\mathbf{R}}$. Attached to any admissible representation of $G_{\mathbf{R}}$ (for example, to any irreducible unitary representation) is a Harish-Chandra module X, which carries an algebraic action of K and a Lie algebra representation of $\mathfrak{g}$. If the original representation has finite length, as we assume from now on, then X is finitely generated as a $U(\mathfrak{g})$-module. Choose a finite-dimensional K-invariant generating subspace X_0 of X, and set

$$X_n = U_n(\mathfrak{g}) \cdot X_0 \tag{1.1)(a}$$

Here as usual $U_n(\mathfrak{g})$ is the (finite-dimensional) subspace of $U(\mathfrak{g})$ spanned by products of at most n elements of $\mathfrak{g}$. This defines a K-invariant increasing filtration on X, which is compatible with the standard filtration of $U(\mathfrak{g})$ in the sense that

$$U_p(\mathfrak{g}) \cdot X_q \subset X_{p+q} \tag{1.1)(b}$$

It follows that the associated graded space $\operatorname{gr} X$ is a module over $\operatorname{gr} U(\mathfrak{g})$. By the Poincaré-Birkhoff-Witt theorem, this last ring is naturally isomorphic to the symmetric algebra $S(\mathfrak{g})$. Because the filtration of X is K-invariant, K acts on $\operatorname{gr} X$ as well. Because of the compatibility of the K and $\mathfrak{g}$ actions on X, the action of the Lie algebra $\mathfrak{k}$ also preserves the filtration of X. It follows that the ideal generated by $\mathfrak{k}$ in $S(\mathfrak{g})$ annihilates $\operatorname{gr} X$. Consequently $\operatorname{gr} X$ may be regarded as an $S(\mathfrak{g}/\mathfrak{k})$-module, equipped with a compatible action of K. Condition (1.1)(a) guarantees that $\operatorname{gr} X$ is generated by X_0; in particular, it is finitely generated.

* Supported in part by NSF grant DMS-8805665.

Recall that the *associated variety* $\mathcal{V}(M)$ of a module M over a commutative ring R is defined to be the set of all prime ideals containing $\operatorname{Ann} M$. (When R is a finitely generated algebra over $\mathbb{C}$, $\mathcal{V}(M)$ is a closed subvariety of the affine algebraic variety $\operatorname{Spec} R$.) In the setting of the preceding paragraph, it is easy to check that the associated variety of $\operatorname{gr} X$ is independent of the choice of X_0; we will recall the argument in section 2. It is called the *associated variety of* X, and written $\mathcal{V}(X)$.

If V is any complex vector space, then the symmetric algebra $S(V)$ may be regarded as the algebra of polynomial functions on the dual vector space V^*. Evaluation at λ in V^* defines a homomorphism from $S(V)$ to $\mathbb{C}$, and therefore a maximal ideal in $S(V)$. All maximal ideals in $S(V)$ arise in this way, so the maximal spectrum of $S(V)$ (the closed points of $\operatorname{Spec} S(V)$) may be identified with V^*. Using these identifications (and speaking a little loosely), we may therefore write

$$\mathcal{V}(X) \subset \mathfrak{g}^* \tag{1.2)(a}$$

Explicitly,

$$\mathcal{V}(X) = \{\, \lambda \in \mathfrak{g}^* \mid p(\lambda) = 0 \text{ whenever } p \in \operatorname{Ann}(\operatorname{gr} X) \,\}. \tag{1.2)(b}$$

Because the elements of $\mathfrak{k}$ annihilate $\operatorname{gr} X$, they must be zero at elements of $\mathcal{V}(X)$. Consequently

$$\mathcal{V}(X) \subset (\mathfrak{g}/\mathfrak{k})^*. \tag{1.2)(c}$$

Because of the compatibility of the K action and the module structure on $\operatorname{gr} X$, $\mathcal{V}(X)$ is a K-invariant subvariety of $\mathfrak{g}^*$ (or of $(\mathfrak{g}/\mathfrak{k})^*$). It turns out in fact that $\mathcal{V}(X)$ is a union of a finite number of nilpotent orbits of K (Corollary 5.23).

We have therefore attached to any admissible representation of $G_{\mathbf{R}}$ of finite length a finite union of nilpotent K-orbits on $(\mathfrak{g}/\mathfrak{k})^*$. Our original intention was to relate representations to nilpotent orbits of $G_{\mathbf{R}}$ on $\mathfrak{g}_{\mathbf{R}}^*$. However, Sekiguchi has shown that there is a close formal relationship between these two kinds of orbits. We will recall his results in section 6, and related results of Schwartz in section 7. For now, it is sufficient to say that one can hope to pursue the philosophy of coadjoint orbits by investigating the relationship between a Harish-Chandra module and its associated variety. (A more precise formulation of this statement may be found in section 8.) Several natural questions arise at once.

1. Is $\mathcal{V}(X)$ the closure of a single orbit of K?
2. Is the closure of every nilpotent K-orbit on $(\mathfrak{g}/\mathfrak{k})^*$ the associated variety of an irreducible unitary representation of $G_{\mathbf{R}}$?
3. Is X uniquely determined by $\mathcal{V}(X)$?
4. Can the structure of X (global character, multiplicities of representations of K) be read off from $\mathcal{V}(X)$?

In a sense these questions are easy: the answer to each is no. The purpose of this paper is to establish some positive results along the lines suggested by the questions. Our inspiration comes from the (closely related) theory of characteristic varieties of D-modules. (We will have no need to be precise about what D is, or to define characteristic varieties; the expert reader can supply appropriate hypotheses.) There Kashiwara-Kawai, Gabber, and others have established deep and powerful results on what the characteristic variety of an irreducible regular holonomic D-module can look like, the extra structure it carries, and the extent to which a D-module can be recovered from its characteristic variety. For example, the characteristic variety of a simple D-module need not be irreducible. But Kashiwara and Kawai have shown that the irreducible components form a single equivalence class under the relation "intersect in codimension 1"; that is, they cannot be too far apart. We prove a weaker related result for associated varieties in section 4. Here is part of it.

Theorem 1.3. *Suppose X is an irreducible Harish-Chandra module, and $\mathcal{O}_\theta$ a K-orbit of maximal dimension in $\mathcal{V}(X)$. Suppose that the complement of $\mathcal{O}_\theta$ has codimension at least two in $\overline{\mathcal{O}_\theta}$. Then $\mathcal{V}(X) = \overline{\mathcal{O}_\theta}$.*

The proof uses a filtration of X different from (1.1)(a), but still satisfying (1.1)(b). Although it is entirely elementary, it is suggested by the non-commutative localization used by Gabber and Kashiwara-Kawai.

The codimension condition is satisfied in many interesting cases (for example, for representations attached to most "non-induced" orbits). In this paper, however, we will apply Theorem 1.3 only to prove the theorem of Borho-Brylinski and Joseph that the associated variety of a primitive ideal is irreducible (Corollary 4.7).

We turn now to the other questions listed above. In one way or another, they ask how X is related to $\mathcal{V}(X)$. To understand that, we must first understand how $\operatorname{gr} X$ is related to $\mathcal{V}(X)$. Very roughly speaking, a finitely generated module over a finitely generated $\mathbb{C}$-algebra looks like the space of sections of a vector bundle over its associated variety. That is, $\operatorname{gr} X$ is approximately the space of sections of a (K-equivariant) vector bundle on $\mathcal{V}(X)$. Now we have seen that $\mathcal{V}(X)$ is approximately a homogeneous space K/H. A K-equivariant vector bundle on K/H is (by passage to the isotropy action at the identity coset) the same thing as a representation of H. The corresponding space of sections is then an induced representation of K.

This suggests that we should be able to attach to X not only the variety $\mathcal{V}(X)$, but also representations of appropriate subgroups of K. This can be done without much difficulty (Definition 2.12). The resulting structure is analogous to the "characteristic cycle" for D-modules; the dimensions of the representations are the analogues of the multiplicities in the cycle. Under appropriate hypotheses on the annihilator of X, it turns out to be possible to place very strong constraints on the possibilities for the representations (Theorem 8.7 below). This is analogous to the fact that the smooth part of the characteristic variety of a D-module carries a natural local system. In some cases (considered for example in [Schwartz]) these constraints cannot be satisfied at all; we get in this way a partial understanding of the negative answer to the second question above. In general we find that, as a representation of K, X must be (approximately) induced from a very special representation of a very special subgroup. This provides some information about the fourth question above. In particular, we formulate in section 12 a precise conjecture (together with strong evidence) about the restrictions to K of a large class of unipotent representations.

The theory of associated varieties of primitive ideals is perhaps a little more familiar to some readers, so we recall briefly how it is related to these ideas. More details may be found in section 4. Suppose I is any two- sided ideal in $U(\mathfrak{g})$. Then the quotient ring $U(\mathfrak{g})/I$ is a finitely generated $U(\mathfrak{g})$-module, so we may define an associated graded $S(\mathfrak{g})$-module $\operatorname{gr} U(\mathfrak{g})/I \simeq S(\mathfrak{g})/\operatorname{gr} I$. The annihilator of this module is obviously $\operatorname{gr} I$, so its associated variety (generally written $\mathcal{V}(I)$) is

$$\mathcal{V}(I) = \{ \lambda \in \mathfrak{g}^* \mid p(\lambda) = 0 \text{ whenever } p \in \operatorname{gr} I) \}. \tag{1.4}$$

(Notice that this notation is inconsistent with that of (1.2): we should really call this $\mathcal{V}(U(\mathfrak{g})/I)$.) Because I is a two-sided ideal, the quotient ring $U(\mathfrak{g})/I)$ inherits the action of the adjoint group G_{ad} of $\mathfrak{g}$. It follows that $\mathcal{V}(I)$ is a G_{ad}-invariant subvariety of $\mathfrak{g}^*$.

Now suppose that I is the annihilator in $U(\mathfrak{g})$ of a Harish-Chandra module X. It is immediate from the definitions that

$$\operatorname{gr} \operatorname{Ann}(X) \subset \operatorname{Ann}(\operatorname{gr} X). \qquad (1.5)(a)$$

Comparing (1.2) with (1.4) therefore gives

$$\mathcal{V}(X) \subset \mathcal{V}(\operatorname{Ann}(X)) \cap (\mathfrak{g}/\mathfrak{k})^*. \qquad (1.5)(b)$$

It turns out that this containment is almost (but not quite) an equality; a precise statement appears in Theorem 8.4.

2. Associated varieties: elementary properties.

We continue now the discussion of associated varieties begun in the introduction. The results of this section are all easy and well-known, although in a few cases it is difficult to find good references. We have therefore included more proofs than the experts will need.

It is convenient and instructive to work in a slightly greater degree of generality. We continue to assume that $\mathfrak{g}$ is a complex reductive Lie algebra. Let K be an algebraic group equipped with an action Ad on $\mathfrak{g}$ (by automorphisms). We assume also that we are given an injective map on Lie algebras

$$i : \mathfrak{k} \to \mathfrak{g}$$

compatible with the differential of Ad. (We will use i to regard $\mathfrak{k}$ as a subalgebra of $\mathfrak{g}$, and generally drop it from the notation. The reader may wonder why we do not simply require K to be a subgroup of some algebraic group G with Lie algebra $\mathfrak{g}$. The reason is that in the setting of the introduction, this will not be possible if $G_{\mathbf{R}}$ is not a linear group.)

A module X for $\mathfrak{g}$ is called a $(\mathfrak{g}, K)$-*module* if it is equipped with an algebraic representation π of K satisfying the two conditions

$$\begin{aligned} \pi(k)(u \cdot x) &= (Ad(k)u) \cdot \pi(k)x \qquad && (k \in K,\, u \in U(\mathfrak{g}),\, x \in X) \\ d\pi(Z)x &= Z \cdot x && (Z \in \mathfrak{k},\, x \in X). \end{aligned} \qquad (2.1(a)$$

An increasing filtration of X indexed by $\mathbf{Z}$ is called *compatible* if it satisfies

$$\begin{aligned} U_p(\mathfrak{g}) \cdot X_q &\subset X_{p+q} \\ \pi(K)X_n &\subset X_n \end{aligned} \qquad (2.1)(b)$$

The first condition allows one to define on $\operatorname{gr} X = \sum_{n \in \mathbf{Z}} X_n / X_{n-1}$ the structure of a graded $S(\mathfrak{g})$-module. (The nth summand will sometimes be called $\operatorname{gr}_n(X)$.) The second condition provides a graded algebraic action (still called π) of K on $\operatorname{gr} X$. These two structures satisfy

$$\begin{aligned} \pi(k)(p \cdot m) &= (Ad(k)p) \cdot \pi(k)m && (k \in K,\, p \in S(\mathfrak{g}),\, m \in \operatorname{gr} X) \\ Z \cdot m &= 0 && (Z \in \mathfrak{k},\, m \in \operatorname{gr} X). \end{aligned} \tag{2.1)(c}$$

An $S(\mathfrak{g})$-module carrying a representation of K satisfying (2.1)(c) is called an $(S(\mathfrak{g}), K)$-*module.* A compatible filtration of X is called *good* if $\cap_{n \in \mathbf{Z}} X_n = 0$, $\cup_{n \in \mathbf{Z}} X_n = X$, and $\operatorname{gr} X$ is a finitely generated $S(\mathfrak{g})$-module. This amounts to four conditions on the filtration:

$$\begin{aligned} X_{-n} &= 0 && \text{(all } n \text{ sufficiently large);} \\ \cup_{n \in \mathbf{Z}} X_n &= X; \\ \dim X_n &< \infty; \\ U_p(\mathfrak{g}) \cdot X_q &= X_{p+q} && \text{(all } q \text{ sufficiently large, all } p \geq 0). \end{aligned} \tag{2.1)(d}$$

The existence of a good filtration evidently implies that X is finitely generated (by X_q for large enough q, say.) Conversely, if X is finitely generated, then we can construct a good filtration of X as in (1.1). The first problem is that X will have many different good filtrations; in order to extract well-defined invariants from the $(S(\mathfrak{g}), K)$-module structure on $\operatorname{gr} X$, we must investigate the dependence of this structure on the filtration. Here and at many points below we will therefore need to consider several different filtrations at the same time, and the subscript notation for them becomes inconvenient. In these cases we may say that $\mathcal{F}$ is a filtration of X, and write $\mathcal{F}_n(X)$ and $\operatorname{gr}(X, \mathcal{F})$ instead of X_n and $\operatorname{gr} X$.

Proposition 2.2. *Suppose X is a $(\mathfrak{g}, K)$-module, and $\mathcal{F}$ and $\mathcal{G}$ are good filtrations of X.*

a) There are integers s and t so that for every integer p,

$$\mathcal{G}_{p-s}(X) \subset \mathcal{F}_p(X) \subset \mathcal{G}_{p+t}(X).$$

b) There are finite filtrations

$$0 = \operatorname{gr}(X, \mathcal{F})_{-1} \subset \operatorname{gr}(X, \mathcal{F})_0 \subset \cdots \subset \operatorname{gr}(X, \mathcal{F})_{s+t} = \operatorname{gr}(X, \mathcal{F})$$

$$0 = \operatorname{gr}(X, \mathcal{G})_{-1} \subset \operatorname{gr}(X, \mathcal{G})_0 \subset \cdots \subset \operatorname{gr}(X, \mathcal{G})_{s+t} = \operatorname{gr}(X, \mathcal{G})$$

by graded $(S(\mathfrak{g}), K)$-submodules, with the property that the corresponding subquotients are isomorphic:

$$\operatorname{gr}(X, \mathcal{F})_j / \operatorname{gr}(X, \mathcal{F})_{j-1} \cong \operatorname{gr}(X, \mathcal{G})_j / \operatorname{gr}(X, \mathcal{G})_{j-1} \qquad (0 \leq j \leq s+t).$$

(The isomorphism shifts the grading of the jth subquotient by $j - s$.)

Proof. The assertion in (a) is a consequence of the properties of good filtrations listed in (2.1)(d). We first choose n so large that $\mathcal{F}_p(X) = U_{p-n}(\mathfrak{g}) \cdot \mathcal{F}_n(X)$, for all $p \geq n$. Next, we choose t so large that $\mathcal{F}_p(X) \subset \mathcal{G}_{p+t}(X)$ for all $p \leq n$; there are essentially only finitely many values of p to consider. For $p \geq n$, we have

$$\mathcal{F}_p(X) = U_{p-n}(\mathfrak{g}) \cdot \mathcal{F}_n(X) \subset U_{p-n}(\mathfrak{g}) \cdot \mathcal{G}_{n+t}(X) \subset \mathcal{G}_{p+t}(X).$$

(The last step uses the compatibility condition (2.1)(b) on $\mathcal{G}$.) This establishes the second containment in (a); the first follows by reversing the roles of $\mathcal{F}$ and $\mathcal{G}$.

For (b), we will construct the required filtration very explicitly. We want to define a graded submodule $\mathrm{gr}(X, \mathcal{F})_j$. We define its component in degree n to be the image of $\mathcal{F}_n \cap \mathcal{G}_{n+j-s}$ in $\mathrm{gr}_n(X, \mathcal{F})$. That is,

$$\mathrm{gr}_n(X, \mathcal{F})_j = (\mathcal{F}_n \cap \mathcal{G}_{n+j-s})/(\mathcal{F}_{n-1} \cap \mathcal{G}_{n+j-s}).$$

That this actually defines an $(S(\mathfrak{g}), K)$-submodule of $\mathrm{gr}(X, \mathcal{F})$ follows from the compatibility of the filtration $\mathcal{G}$. By (a), the submodule is zero if $j < 0$, and is all of $\mathrm{gr}(X, \mathcal{F})$ if $j \geq s + t$. Computing the subquotient modules is easy: the nth graded piece of $\mathrm{gr}(X, \mathcal{F})_j/\mathrm{gr}(X, \mathcal{F})_{j-1}$ is

$$(\mathcal{F}_n \cap \mathcal{G}_{n+j-s})/(\mathcal{F}_{n-1} \cap \mathcal{G}_{n+j-s} + \mathcal{F}_n \cap \mathcal{G}_{n+j-s-1})$$

Except for the shift from n to $n + j - s$, this formula is symmetric in $\mathcal{F}$ and $\mathcal{G}$. This provides the isomorphism in (b). Q.E.D.

Recall that a map d from an abelian category to an abelian semigroup is called *additive* if whenever

$$0 \to A \to B \to C \to 0$$

is a short exact sequence, we have $d(B) = d(A) + d(C)$. If d is such a map and A is an object in the category with a finite filtration $0 = A_{-1} \subset A_0 \subset \cdots \subset A_m = A$, then $d(A) = \sum_{j=0}^{m} d(A_j/A_{j-1})$ If B is another object admitting a filtration with subquotients isomorphic to those of A, then $d(A) = d(B)$. Proposition 2.2 therefore guarantees that any additive map on the category of finitely generated $(S(\mathfrak{g}), K)$-modules will give a well-defined (in fact additive) map on the category of $(\mathfrak{g}, K)$-modules.

Here are two important examples. Any finitely generated $S(\mathfrak{g})$-module A has a *Krull dimension* $\dim A$ which is a non-negative integer. (It is often convenient to say that the zero module has Krull dimension -1.) For a short exact sequence as above, we have $\dim B = \max(\dim A, \dim B)$. This map is additive if we make the integers into a semigroup by defining the "sum" of two integers to be their maximum. We therefore get a corresponding invariant for a finitely generated $U(\mathfrak{g})$-modules; it is the *Gelfand-Kirillov dimension*, usually written $\operatorname{Dim} X$.

Next, recall from the introduction that the associated variety $\mathcal{V}(A)$ for an $S(\mathfrak{g})$-module is the set of prime ideals containing the annihilator of A. We can make the collection of sets of prime ideals in $S(\mathfrak{g})$ into a semigroup with the union operation; then $\mathcal{V}$ is an additive map. (The reason is that given a short exact sequence as above, we must have

$$(\operatorname{Ann} A)(\operatorname{Ann} C) \subset \operatorname{Ann} B \subset (\operatorname{Ann} A) \cap (\operatorname{Ann} C).$$

Any prime ideal containing $\mathrm{Ann}\, B$ must contain the first term, and therefore must contain either $\mathrm{Ann}\, A$ or $\mathrm{Ann}\, C$. Conversely, any prime ideal containing either $\mathrm{Ann}\, A$ or $\mathrm{Ann}\, C$ contains their intersection, and therefore contains $\mathrm{Ann}\, B$.) We can therefore define the associated variety $\mathcal{V}(X)$ for a finitely generated $U(\mathfrak{g})$-module, as was claimed in the introduction. By the Nullstellensatz, this amounts to the assertion that the radical $\sqrt{\mathrm{Ann}(\mathrm{gr}\, X)}$ is independent of the choice of good filtration.

We want now to introduce some refinements of these invariants. We begin by recalling a little commutative algebra. Suppose M is a module for the commutative ring R. In addition to the associated variety, there are two other important sets of prime ideals attached to M. The *support* $\mathrm{Supp}\, M$ is defined to be the set of all prime ideals P in R for which the localization M_P is non-zero. The set of *associated primes* $\mathrm{Ass}\, M$ consists of those prime ideals which are annihilators of elements of M. It is easy to see that

$$\mathcal{V}(M) \supset \mathrm{Supp}\, M \supset \mathrm{Ass}\, M.$$

If M is finitely generated, then $\mathcal{V}(M) = \mathrm{Supp}\, M$ (see for example [Matsumura], pp.25-26). If in addition R is Noetherian, then $\mathrm{Ass}\, M$ is a finite set including the minimal elements of $\mathcal{V}(M)$ ([Matsumura], Theorem 6.5). It follows that $\mathcal{V}(M)$ is the Zariski closure of $\mathrm{Ass}\, M$.

Suppose now that R is Noetherian and M is finitely generated. Let $P_1, \ldots, P_r$ be the set of minimal elements in $\mathcal{V}(M)$; that is, the set of minimal primes containing the annihilator of M. (Each P_i corresponds to an irreducible component $\mathcal{V}(P_i)$ of $\mathcal{V}(M)$ regarded as an algebraic variety.) We are going to define the characteristic cycle of M to be a certain formal sum of these prime ideals (with positive integer multiplicities). Roughly speaking, the coefficient of $\mathcal{V}(P_i)$ will measure how many copies of R/P_i are contained in M.

Theorem 6.4 of [Matsumura] guarantees that we can find a finite filtration of M by R-submodules so that each subquotient M_j/M_{j-1} is isomorphic to R/Q_j, with Q_j in $\mathcal{V}(M)$ a prime ideal. We will define the coefficient of $\mathcal{V}(P_i)$ to be the number of values of j for which $P_i = Q_j$. We have to check that this definition is independent of the choice of the filtration of M. This requires a little care. The set of prime ideals occurring among the Q_j *does* depend on the choice of filtration; but the minimal elements have well-defined multiplicities. To see this, we can compare each of two filtrations of M to a common refinement of them, using formal arguments and the following easy lemma.

Lemma 2.3. *Suppose P is a prime ideal in a commutative ring R. Regard $A = R/P$ as an R-module, and fix a finite filtration of A with subquotients A_j/A_{j-1} of the form R/Q_j, with Q_j a prime ideal in R. Then $Q_1 = P$, and every other Q_j properly contains P.*

We leave the remaining details to the reader. We have now shown that the following definition makes sense.

Definition 2.4. Suppose R is a commutative Noetherian ring and M is a finitely generated R-module. Let $P_1, \ldots, P_r$ be the minimal prime ideals containing the annihilator of M. The *characteristic cycle* of M is the formal sum

$$\mathrm{Ch}(M) = \sum_{i=1}^{r} m(P_i, M) P_i,$$

where $m(P_i, M)$ is a positive integer defined as follows. Choose a finite filtration of M so that each subquotient M_j/M_{j-1} is of the form R/Q_j, with Q_j a prime ideal in R. Then $m(P_i, M)$ is the number of values of j for which $P_i = Q_j$.

Sometimes it is useful to define $m(Q, M)$ when Q is a prime ideal not among the $\{P_i\}$. If Q contains the annihilator of M but not minimally, we set $m(Q, M) = \infty$. If Q does not contain the annihilator of M, then $m(Q, M) = 0$. In terms of the local ring R_P at a prime ideal P, and the stalk $M_P = M \otimes_R R_P$ of M at a prime ideal P, it is not difficult to check that in every case

$$m(P, M) = \text{ length of } M_P \text{ as an } R_P\text{-module.}$$

Finally, we want to see that Ch is an additive map. To do that, we have to define a semigroup structure on its range. We take the range to be the set of finite formal sums $\sum_{i=1}^{r} m_i P_i$ (of prime ideals with positive integer coefficients) subject to the condition that there should be no containments among the P_i. To add two such expressions, we first throw away terms for ideals properly contained in other ideals, then add coefficients. Now the additivity is elementary.

As a consequence, we can define the *characteristic cycle* $\mathrm{Ch}(X)$ of a $(\mathfrak{g}, K)$-module; it consists of a positive integer weight attached to each irreducible component of $\mathcal{V}(X)$. Obviously this invariant refines the associated variety. Since the Gelfand-Kirillov dimension is just the dimension of $\mathcal{V}(X)$, the characteristic cycle contains that information as well.

One weakness of the characteristic cycle is that it contains no information about the action of K. In order to remedy this, we need some other ways to calculate the multiplicities in the characteristic cycle. Although our main results can be formulated strictly in terms of modules and ideals, it is very convenient to use the language of sheaves of modules along the way. For this we refer to [Hartshorne], chapter II, or [Shafarevich], Chapter VI. In particular, we use (for a commutative ring R) the equivalence between the category of R-modules and the category of quasi-coherent sheaves of modules on $\operatorname{Spec} R$.

Proposition 2.5. *Suppose P is a prime ideal in the commutative Noetherian ring R, and M is a finitely generated R-module annihilated by P. Then there is an element f of R, not belonging to P, with the property that the localization M_f is a free $(R/P)_f$-module. (That is, M is free on the open set $f \neq 0$ in $\operatorname{Spec} R/P$). Its rank is the multiplicity $m(P, M)$ of P in the characteristic cycle of M.*

The first assertion may be found for example in [Shafarevich], Proposition VI.3.1. The second is an immediate consequence of Definition 2.4 and the exactness of localization.

Corollary 2.6. *In the setting of Proposition 2.5, suppose Q is any prime ideal containing P but not containing f. Then the stalk M_Q of M at Q is a free R_Q/PR_Q-module of rank equal to $m(P, M)$. In particular, this multiplicity is equal to the dimension of the vector space M_Q/QM_Q over the quotient field R_Q/QR_Q of R/Q.*

Corollary 2.7. *In the setting of Proposition 2.5, suppose $\mathfrak{m}$ is any maximal ideal containing P but not containing f. Then*

$$m(P, M) = \dim M/\mathfrak{m}M;$$

the dimension is taken over the field $R/\mathfrak{m}$.

Corollary 2.8. *Suppose R is a finitely generated commutative algebra over $\mathbb{C}$, and M is a finitely generated R-module. Choose a finite filtration of M so that each subquotient is annihilated by some prime ideal in $\mathcal{V}(M)$. Fix one of the minimal primes P in $\mathcal{V}(M)$. Then there is an open dense subset U of $\operatorname{Spec} R/P$ (which we regard as a subset of $\operatorname{Spec} R$) so that if $\mathfrak{m}$ is any maximal ideal in U, the multiplicity of P is*

$$m(P,M) = \sum_j \dim_{\mathbb{C}} M_j/(\mathfrak{m}M_j + M_{j-1}).$$

It is tempting to try to omit the filtration of M in the last corollary; the formula would still make sense, and it would give the right answer for $M = R/Q$ whenever Q is prime. The difficulty is that the resulting function on modules is not additive. A simple example is $R = \mathbb{C}[x]$, $M = R/\langle x^2 \rangle$. Then $\mathcal{V}(M)$ consists of just the maximal ideal $\mathfrak{m} = \langle x \rangle$ in R, which has multiplicity two. Nevertheless, one calculates immediately that $\dim M/\mathfrak{m}M = 1$.

Despite the possibility of such problems, we are going to need to calculate multiplicities using filtrations satisfying somewhat weaker conditions than the one in Corollary 2.8. This can be done using the following elementary extension of Proposition 2.5.

Proposition 2.9. *Suppose P is a prime ideal in the commutative Noetherian ring R, and M is a finitely generated R-module. The following conditions are equivalent.*

a) There is an ideal I of R, not contained in P, such that M is annihilated by IP.

b) There is an element g of R, not contained in P, such that $g \cdot M$ is annihilated by P.

c) There is an element h of R, not contained in P, such that M_h is annihilated by P.

d) There is an element f of R, not belonging to P, with the property that M_f is annihilated by P, and defines a free $(R/P)_f$-module. (That is, M is free on the open set $f \neq 0$ in $\operatorname{Spec} R/P$).

In the setting (d), the rank of M_f is the multiplicity $m(P,M)$ of P in the characteristic cycle of M.

When the conditions (a)-(d) in the proposition are satisfied, we say that M is *generically reduced along P*. (Analogously, we might say that M is *reduced along P* when P annihilates M.) Notice that these are conditions only on the annihilator $\mathfrak{a}$ of M. We may also say that the ideal $\mathfrak{a}$ is *generically reduced along P* if the R-module $R/\mathfrak{a}$ is. This notion in the case of ideals is considered further in Lemmas 10.16 and 10.17.

Proof. We will show that (a) $\Rightarrow$ (b) $\Rightarrow$ (c) $\Rightarrow$ (d) $\Rightarrow$ (a). Assume (a), and choose g in I not in P; then (b) is immediate. Next, suppose g is as in (b); then (c) follows with $h = g$. Suppose (c) holds. Apply Proposition 2.5 to the prime ideal PR_h in R_h, obtaining an element $f_0 = h^{-n} f_1$ of R_h; here f_1 is in R. It is easy to verify that the element $f = f_1 h$ has the properties we require (since R_f is naturally isomorphic to $(R_h)_{f_0}$.

Suppose f is as in (d). Choose generators $(m_1, \ldots, m_r)$ for M and $(p_1, \ldots, p_s)$ for P. The assumption that P annihilates M_f means that for each i,j there is a positive integer $N(i,j)$ so that $f^{N(i,j)} \cdot (p_i \cdot m_j) = 0$. Taking N to be the maximum of the $N(i,j)$, we get $f^N \cdot (P \cdot M) = 0$. We can take I to be the ideal generated by f^N in (a).

Finally, we prove the assertion about multiplicities. Fix f as in (d), and let S be the submodule of M annihilated by P. As we just saw, the action of f^N maps M into S. It follows that P does not contain the annihilator of M/S, so $m(P, M/S) = 0$, and consequently $m(P, M) = m(P, S)$. By the exactness of localization, $M_f = S_f$. Proposition 2.5 applied to S therefore gives the result. Q.E.D.

Corollary 2.10. *Suppose R is a finitely generated commutative algebra over* C, *and M is a finitely generated R-module. Fix one of the minimal primes P in $\mathcal{V}(M)$. Choose a finite filtration of M so that each subquotient is generically reduced along P. Then there is an open dense subset U of* $\operatorname{Spec} R/P$ *(which we regard as a subset of* $\operatorname{Spec} R$*) so that if* $\mathfrak{m}$ *is any maximal ideal in U, the multiplicity of P is*

$$m(P, M) = \sum_j \dim_{\mathbb{C}} M_j/(\mathfrak{m}M_j + M_{j-1}).$$

We are going to define a refinement of the characteristic cycle that reflects some of the action of K. We will make no attempt to define it in very great generality. Recall that a *virtual character* of an algebraic group H is a finite formal integer combination of irreducible (finite-dimensional algebraic) representations of H. The set of all virtual characters is a free abelian group with basis the set of irreducible representations of H. A character is called *genuine* if the coefficients are non-negative; that is, if it is the character of a (finite-dimensional reducible algebraic) representation π of H. We write $\Theta(\pi)$ for the character of π. The connection with more classical terminology in representation theory is that we may identify $\Theta(\pi)$ with the function on H (also denoted $\Theta(\pi)$) sending h to $\operatorname{tr} \pi(h)$. Passing to differences, we can attach such a function to any virtual character. This identifies the lattice of virtual characters of H with a certain space of conjugation-invariant regular functions on H. Because the unipotent radical U of H acts by unipotent operators in any algebraic representation, these functions must be constant on U. Consequently they descend to H/U. The theory of characters is really therefore a theory about reductive groups.

In order to define the refined characteristic cycle, we need one more fact.

Lemma 2.11. *Suppose M is a finitely generated $(S(\mathfrak{g}), K)$-module. Then there is a finite filtration of M by $(S(\mathfrak{g}), K)$-submodules with the property that every subquotient is generically reduced along every minimal prime in $\mathcal{V}(M)$.*

Proof. If M is zero, $\mathcal{V}(M)$ is empty and there is nothing to prove. Otherwise we proceed by induction on the dimension of $\mathcal{V}(M)$, and then by induction on the sum of the multiplicities of the components of largest dimension. So fix a minimal prime ideal P_1 in $\mathcal{V}(M)$, with R/P_1 of maximal dimension. The action of K on $S(\mathfrak{g})$ evidently permutes the prime ideals in $\mathcal{V}(M)$, and therefore the (finitely many) minimal primes. List the minimal primes in $\mathcal{V}(M)$ and the K orbit of P_1 as $P_1, \ldots, P_s$. Define M_1 to be the submodule of M annihilated by the intersection J of all of these ideals. Because J is K-invariant, M_1 is an $(S(\mathfrak{g}), K)$-submodule. Since P_1 is an associated prime of M, P_1 has non-zero multiplicity in M_1; so the inductive hypothesis applies to the quotient M/M_1. So we need only prove that M_1 is generically reduced along every minimal prime P in $\mathcal{V}(M)$. If P is not one of the P_i, then P does not contain J; so condition (a) of Proposition 2.9 may be satisfied (in a trivial way) by taking I equal to J. If P is one of the P_i, take I equal to the intersection of the remaining ones. Then IP is contained in J, which annihilates M_1; so condition (a) of Proposition 2.9 is again satisfied. Q.E.D.

Definition 2.12. Suppose M is a finitely generated $(S(\mathfrak{g}), K)$-module, and P is a minimal prime ideal in $\mathcal{V}(M)$ (so that $\mathcal{V}(P)$ is a component of $\mathcal{V}(M)$). Assume that for some $\lambda \in \mathfrak{g}^*$ the K-orbit of λ contains a dense open subset of $\mathcal{V}(P)$. Write $K(\lambda)$ for the isotropy group of the action of K at λ, and $\mathfrak{m}(\lambda)$ for the maximal ideal in $S(\mathfrak{g})$ corresponding to λ. We want to attach to M a finite-dimensional representation of $K(\lambda)$. If N is any $(S(\mathfrak{g}), K)$-module, then $K(\lambda)$ acts algebraically on the finite-dimensional vector space $N/\mathfrak{m}(\lambda)N$. Choose a finite filtration of M by $(S(\mathfrak{g}), K)$-submodules M_i, so that the subquotients are generically reduced along P (Lemma 2.11). The *character* of M at λ is the (genuine) virtual character

$$\chi(\lambda, M) = \sum_j M_j/(\mathfrak{m}(\lambda)M_j + M_{j-1})$$

of $K(\lambda)$. (We will explain in a moment why this is independent of choices.) By Corollary 2.10,

$$\dim \chi(\lambda, M) = m(P, M).$$

There is no difficulty in extending this definition to the setting of a coherent sheaf of modules with K-action over an algebraic variety Z on which K acts. If $Z = K/H$ is a homogeneous space, then such a sheaf M must be the sheaf of sections of a vector bundle on which K acts. Such a vector bundle in turn is given by a representation π of H. In this case $\chi(eH, M)$ is the virtual representation of H represented by π. The definition (like that of the characteristic cycle) gives precision to the idea that a module is more or less the space of sections of a vector bundle over its support.

To see that $\chi(\lambda, M)$ is well-defined, one can imitate the proof of Proposition 2.2. At various points in the argument one encounters short exact sequences

$$0 \to A \to B \to C \to 0.$$

These give rise to right exact sequences

$$A/\mathfrak{m}(\lambda)A \to B/\mathfrak{m}(\lambda)B \to C/\mathfrak{m}(\lambda)C \to 0$$

which must be shown to be exact. But the last formula in Definition 2.12 will guarantee the additivity of dimensions in this sequence, and the exactness follows. (This argument is the reason we needed Corollary 2.10.) Once we know that $\chi(\lambda, M)$ is well-defined, its additivity is obvious.

If the element λ of Proposition 2.2 is replaced by some $k \cdot \lambda$, with $k \in K$, then $K(\lambda)$ and $\chi(\lambda, M)$ are replaced by their conjugates under k. In this sense the choice of λ is immaterial. We could in fact even manage without the hypothesis that some K orbit meet $\mathcal{V}(P)$ in a dense set: in any case there is a dense open set in $\mathcal{V}(P)$ on which the isotropy groups of K belong to a single conjugacy class, and (on a slightly smaller open set) the same definition as above will yield a well-defined conjugacy class of genuine virtual representations of this class of subgroups. We will have no need for this generality, however.

The next theorem summarizes some of what we have established in this section.

Theorem 2.13. *Suppose X is a finitely generated $(\mathfrak{g}, K)$-module, and $\mathcal{V}(X) \subset \mathfrak{g}^*$ is its associated variety. Assume that K has a finite number of orbits on $\mathcal{V}(X)$. List the maximal orbits (those not contained in the closures of others) as $\mathcal{O}_1, \ldots, \mathcal{O}_s$, and choose a representative $\lambda_i \in \mathcal{O}_i$ for each orbit. Write K_i for the isotropy group of K at λ_i. Then there is attached to X a non-zero genuine virtual representation $\chi_i = \chi(\lambda_i, X)$ of K_i.*

In light of the remarks after Definition 2.12, this theorem says that X has something to do with sections of the vector bundles $K \times_{K_i} \chi_i$. We will make this more precise in Theorem 4.2.

3. Associated varieties: microlocal properties

In this section we will consider properties of associated varieties that are suggested by the "microlocalization" techniques of Sato-Kashiwara-Kawai and Gabber (see [Ginsburg], [Gabber], [Springer], and [SKK]). First we frame some of the definitions of section 2 more generally, beginning with a filtered algebra A (with a unit) over $\mathbf{C}$. (Often A will be the universal enveloping algebra $U(\mathfrak{g})$ or a quotient of it.) This means that A is equipped with an increasing filtration by subspaces indexed by $\mathbf{Z}$:

$$\cdots \subset A_{-1} \subset A_0 \subset A_1 \subset \cdots, \qquad A_p A_q \subset A_{p+q}. \tag{3.1)(a}$$

(Occasionally it will be convenient to index a filtration by $b\mathbf{Z}$ for some fraction b: this causes no difficulties.) We can then define an associated graded ring

$$R = \operatorname{gr} A = \sum_{n \in \mathbf{Z}} R^n, \qquad R^n = A_n / A_{n-1}. \tag{3.1)(b}$$

We will assume for convenience that R is commutative, although some of the preliminary formalism requires much less. Of course if A is $U(\mathfrak{g})$, then R is $S(\mathfrak{g})$. Suppose X is an A-module. A *compatible filtration* on X is an increasing family of subspaces of X indexed by $\mathbf{Z}$ (or, occasionally, by $a + b\mathbf{Z}$), satisfying

$$A_p \cdot X_q \subset X_{p+q}. \tag{3.1)(c}$$

In this case we can define an associated graded module

$$M = \operatorname{gr} X = \sum_{n \in \mathbf{Z}} M^n, \qquad M^n = X_n / X_{n-1}, \tag{3.1)(d}$$

a graded R-module. When we want the notation to allow for several filtrations, we write $\mathcal{F}_n(X)$ and $\operatorname{gr}(X, \mathcal{F})$ instead of X_n and $\operatorname{gr} X$. We will need our filtrations to be *exhaustive*:

$$\bigcup_n A_n = A, \qquad \bigcup_n X_n = X. \tag{3.1)(e}$$

(The dual requirements that $\bigcap_n A_n = 0$ and $\bigcap_n X_n = 0$ will appear eventually, but they are not needed for the general formal development.)

The notation of (3.1) will be in force throughout this section.

The notion of good filtration is a little subtle in general (see [Ginsburg], Proposition 1.2.2). We will need it only in the special case when $A_{-1} = 0$. Then it is given by the obvious analogue of (2.1)(d), and Proposition 2.2 carries over immediately.

The theory of microlocalization constructs certain (noncommutative) filtered localizations A_S of A related to graded localizations of R. Tensoring with these localizations gives a localization theory for A-modules; that is, a way of mapping an A-module X into larger, "smoother"objects X_S. A compatible filtration $\mathcal{F}$ on X induces one $\mathcal{F}_S$ on X_S as well. We can then pull $\mathcal{F}_S$ back to X, getting a new and "smoother" compatible filtration $\mathcal{F}(S)$ on X. From the existence of these improved filtrations, one can hope to deduce interesting results about X. One problem with this program is that it appears to require an understanding of microlocalization. What we propose to do is construct the new filtrations directly, without explicit use of A_S. The price is of course a loss in conceptual power; the present paragraph is intended to alleviate that loss.

What we cannot avoid is commutative localization, and we recall now a little about that. Recall that R is a graded commutative ring. A *closed cone* in $\operatorname{Spec} R$ is any subvariety defined by a homogeneous ideal in R, and an *open cone* is the complement of a closed cone. If f is in R, write $D(f) \subset \operatorname{Spec} R$ for the set of prime ideals not containing f. An open cone is a union of such sets, for various homogeneous elements f. Let us make this explicit in the case of $S(\mathfrak{g})$. Suppose U is an open cone in $\mathfrak{g}^*$. The complement of U is a closed cone, which is therefore the set of simultaneous zeros of a finite set S of homogeneous polynomials. The usual "identification" of $\operatorname{Spec} S(\mathfrak{g})$ with $\mathfrak{g}^*$ identifies $D(f)$ with the subset of $\mathfrak{g}^*$ at which the polynomial f does not vanish. Hence

$$\begin{aligned} U &= \{ \lambda \in \mathfrak{g}^* \mid f(\lambda) \neq 0 \text{ for some } f \text{ in } S \} \\ &= \bigcup_{f \in S} D(f). \end{aligned}$$

In general, whenever S is a subset of a commutative ring R, we will write

$$U_S = \bigcup_{f \in S} D(f) \tag{3.2)(a}$$

for the set of prime ideals not containing S, and

$$Z_S = \{ P \in \operatorname{Spec} R \mid S \subset P \} \tag{3.2)(b}$$

for its complement. Of course Z_S is just the variety of S, a closed subset of $\operatorname{Spec} R$.

Suppose M is an R-module. Recall (say from [Hartshorne]) that M defines a sheaf of modules (which we denote by $\tilde{M}$) on $\operatorname{Spec} R$. By definition, $\tilde{M}(D(f))$ is the localization M_f of M at f. A typical element of this localization is of the form $f^{-n}m$, with n a non-negative integer and $m \in M$. Two such elements $f^{-n}m$ and $f^{-n'}m'$ are equal if $f^{N-n}m$ is equal to $f^{N-n'}m'$ in M for all large N. If f is homogeneous of degree p, we can therefore grade M_f by defining

$$M_f^n = \sum_{k \in \mathbb{N}} f^{-k} M^{n+kp}. \tag{3.2)(c}$$

Notice that this grading may have negative terms even if the one on M does not.

Next, we want to compute the value of the sheaf $\tilde{M}$ on the open cone U_S defined by a set S of homogeneous elements. If f and g belong to S, then there is a natural map $\phi_{f,g}$ from M_f to M_{fg} (sending $f^{-n}m$ to $(fg)^{-n}(g^n m)$). The module $\tilde{M}(U_S)$ is contained in the direct product over S of the various localizations M_f; it is defined to be

$$\tilde{M}(U_S) = \left\{ m = (m_f) \in \prod_{f \in S} M_f \mid \phi_{f,g}(m_f) = \phi_{g,f}(m_g), \text{ all } f,\, g \in S \right\}. \tag{3.2)(d}$$

The maps $\phi_{f,g}$ preserve degrees in the gradings of the previous paragraph. If R is noetherian (so that S may be taken to be finite) then it follows that $\tilde{M}(U_S)$ is spanned by elements (m_f) in which all coordinates have the same degree:

$$\tilde{M}(U_S)^n = \{ m = (m_f) \in \prod_{f \in S} M_f^n \mid \phi_{f,g}(m_f) = \phi_{g,f}(m_g), \text{ all } f,\, g \in S \}. \tag{3.2)(e}$$

In this way $\tilde{M}(U_S)$ acquires a grading; again it may have negative terms. As an immediate consequence of this description, we get the following lemma.

Lemma 3.3. *Suppose R is a graded commutative ring, S is a homogeneous subset of R, and U_S is the open cone (3.2)(a). Then the kernel of the natural map $M \to \tilde{M}(U_S)$ is*

$$\{ m \in M \mid \text{ for all } f \in S \text{ there is some } N = N(f,m) \text{ so that } f^N m = 0 \}.$$

The support of this kernel is contained in the closed cone Z_S defined by S.

Definition 3.4. In the setting of (3.1), suppose $\mathcal{F}$ is a compatible filtration of the A-module X, and S is a homogeneous subset of $R = \operatorname{gr} A$. List the elements of S as $\{f_i\}_{i \in I}$; say f_i has degree p_i. Choose a set $\Sigma = \{\phi_i\}_{i \in I}$ of representatives of S in A:

$$\phi_i \in A_{p_i}, \qquad \operatorname{gr} \phi_i = f_i.$$

Suppose $J = (i_1, \ldots, i_N) \in I^N$ is an ordered N-tuple of elements of I. Define

$$p_J = \sum_{j=1}^{N} p_{i_j}, \qquad f_J = \prod_{j=1}^{N} f_{i_j} \in R^{p_J}, \qquad \phi_J = \prod_{j=1}^{N} \phi_{i_j} \in A_{p_J}.$$

Notice that ϕ_J depends on the order in which the product is taken.

The *S-localization of* $\mathcal{F}$ is a new filtration $\mathcal{F}(S)$ on X, defined as follows:

$$\mathcal{F}(S)_n(X) = \{ x \in X \mid \text{ for all } N \text{ sufficiently large, and all } J \in I^N,\ \phi_J \cdot x \in \mathcal{F}_{n+p_J}(X) \}.$$

When the set S and the filtration $\mathcal{F}$ are understood, it will be convenient to write $X_{[n]}$ in place of $\mathcal{F}(S)_n(X)$. We will then write $[\mathrm{gr}]\,X$ for the associated graded object.

In order to work conveniently with this definition, we need to extend its notation somewhat.

Definition 3.5 In the setting of (3.1)(a) and (3.1)(b), fix a set $\Sigma = \{\phi_i\}_{i\in I}$ of elements of A; say $\phi_i \in A_{p_i}$. Define ϕ_J, p_J as in Definition 3.4. By a *product of type* (N, r, Σ) we will mean (roughly) a product π of elements of A of degrees adding up to r, with at least N of the factors in Σ. (When Σ is understood, we may simply say a product of type (N, r).) A little more precisely, we mean that there are elements $(b_0, \ldots, b_N)$ of A, with $b_j \in A_{q_j}$, and a sequence $J = (i_j) \in I^N$, so that

$$\pi = b_0\phi_{i_1}b_1 \cdots b_{N-1}\phi_{i_N}b_N, \qquad r = p_J + \sum_{j=0}^{N} q_j.$$

Notice that a product of type (N, r) belongs to A_r.

Lemma 3.6. *In the setting of Definitions 3.4 and 3.5, $X_{[n]} = \mathcal{F}(S)_n(X)$ is a vector subspace of X containing $X_n = \mathcal{F}_n(X)$. Suppose that x is an element of $X_{[n]}$. Then there is a positive integer N_0 (depending on x) having the following property: if π is a product of type (N, r) and $N \geq N_0$, then $\pi \cdot x \in X_{n+r}$.*

If we could rearrange the product π to put all the terms in Σ on the right , then the result would be obvious. By Definition 3.4, the ϕ_J factor would take x into X_{n+p_J} (if N is large enough); and the remaining factor (which lies in A_{r-p_J}) would take X_{n+p_J} into X_{n+r}. The difficulty is that A is not commutative, so we are not completely free to rearrange the product. What saves us is the commutativity of $R = \mathrm{gr}\, A$. This says that we can rearrange products up to error terms of lower order: if $\phi \in A_p$ and $b \in A_q$, then there is a $c \in A_{p+q-1}$ such that $\phi b = b\phi + c$. Repeating this argument, we see that any product π of type (N, r) may be rewritten as

$$\pi = b\phi_J + (\text{sum of products of type } (N-1, r-1)) \tag{3.7}$$

with $b \in A_{r-p_J}$. (What is critical for this is that $\mathrm{gr}\, A$ be commutative, or at least that the elements $\mathrm{gr}\,\phi_i$ be central in it.)

Lemma 3.8. *In the setting of Definitions 3.4 and 3.5, fix integers N, s, and t, with $0 \leq s, t \leq N$, and $s + t \leq N + 1$. Then any product π of type (N, r) is equal to a sum of terms of two forms. The first form is $c\phi_J$, with $J \in I^s$ and c a product of type $(N - s - m, r - p_J - m)$ (with $0 \leq m \leq t - 1$). The second form is just products of type $(N - t, r - t)$.*

Notice that a term of the first form in the lemma is a special kind of product of type $(N-m, r-m)$. The simplest case of the lemma is $N=1, s=t=1$. In that case it says that if $\phi \in A_p, b \in A_q$, and $b' \in A_{q'}$ (so that the r of the lemma is $p+q+q'$), then $b'\phi b = c\phi + c'$, with $c \in A_{q+q'}$, and $c' \in A_{p+q+q'-1}$. The term $c\phi$ is of the first form in the lemma (with $m = 0 = t-1$) and the term c' is of the second form. Of course we may take $c = b'b, c' = b'(\phi b - b\phi)$.

Proof of Lemma 3.8. We proceed by induction on N. If $t=0$, then the result is trivial (since π is already of the second form allowed in the conclusion). If $t=1$, then the conclusion of the lemma follows from (3.7). So suppose $t > 1$. It follows that N is positive, and that $s \leq N-1$. Therefore we can find $b \in A_q, \phi \in A_p \cap \Sigma$, and a product π' of type $(N-1, r-p-q)$, so that $\pi = b\phi\pi'$. We apply the inductive hypothesis to $(\pi', N-1, s, t-1)$, and multiply the resulting expansion of π by $b\phi$. The only difficulty arises from terms π'' of type $((N-1)-(t-1), r-p-q-t+1)$ in the expansion of π'. After multiplication by $b\phi$, such a term is of type $(N-t+1, r-t+1)$. Applying the inductive hypothesis to $(b\phi\pi'', N-t+1, s, 1)$ gives the conclusion of the lemma. Q.E.D.

Proof of Lemma 3.6. The first assertion of the lemma is obvious. For the rest, choose s so large that for all all $J \in I^s$, $\phi_J \cdot x \in X_{n+p_J}$ (Definition 3.4). Since the filtration of X is exhaustive, there is a positive integer t such that $x \in X_{n+t}$. Set $N_0 = s+t-1$. Suppose $N \geq N_0$, and π is a product of type (N, r). We must show that $\pi \cdot x \in X_{n+r}$. By Lemma 3.8, we may replace π by a sum of terms of two forms. If $c\phi_J$ is a term of the first form, then $\phi_J \cdot x \in X_{n+p_J}$ and $c \in A_{r-p_J-m}$. It follows that $c\phi_J \cdot x \in X_{n+r-m}$ (with $m \geq 0$). If π' is of the second form, then it belongs to A_{r-t}. Since $x \in X_{n+t}$, it follows that $\pi' \cdot x \in X_{(n+t)+(r-t)}$. Q.E.D.

Corollary 3.9. *In the setting of Definition 3.4, the localized filtration $\mathcal{F}(S)$ may be described as*

$$\mathcal{F}(S)_n(X) = \{\, x \in X \mid \textit{ for all } N \textit{ large, and all } \pi \textit{ of type } (N,r),\ \pi \cdot x \in \mathcal{F}_{n+r}(X)\,\}.$$

It is a compatible exhaustive filtration of X, depending only on the set S (and not on the choice of representatives Σ in A).

Proof. Write $X_{[n]}$ for $\mathcal{F}(S)_n(X)$. The description of $X_{[n]}$ is immediate from Lemma 3.6. Since $X_n \subset X_{[n]}$, the localized filtration is exhaustive (cf. (3.1)(e)). To check the compatibility, suppose $x \in X_{[n]}$ and $b \in A_q$. We want to show that $b \cdot x \in X_{[n+q]}$. We use the description just established for the filtration. If π is a product of type (N, r), then $\pi \cdot b$ is a product of type $(N, r+q)$. It follows that if N is large enough, $\pi \cdot (b \cdot x) \in X_{n+r+q}$. That is, $b \cdot x$ satisfies our new criterion for belonging to $X_{[n+q]}$.

To see that $\mathcal{F}(S)$ is independent of the choice of Σ, suppose $\Sigma' = \{\phi_i'\}_{i \in I}$ is another such set. We can write $\phi_i' = \phi_i + v_i$, with $v_i \in A_{p_i - 1}$. If $J \in I^N$, then clearly ϕ'_J is equal to ϕ_J plus a sum of products of various types $(N-s, p_J - s)$, with $1 \leq s \leq N$; these are obtained from ϕ_J by replacing s of the ϕ_i factors by v_i. Now it follows from Lemma 3.6 that if $x \in X_{[n]}$ and N is large enough, then $\phi'_J \cdot x \in X_{[n+p_J]}$. This shows that the filtration defined using Σ' contains the one defined using Σ. Q.E.D.

A similar argument shows that $\mathcal{F}(S)$ is unaffected by adding to S any finite homogeneous subset of the ideal generated by S. A consequence is

Corollary 3.10. *In the setting of Definition 3.4, suppose S is finite, and that S' is another finite homogeneous set generating the same ideal in R. Then $\mathcal{F}(S) = \mathcal{F}(S')$.*

With a little more effort, one sees that only the radical of the ideal matters. Since we will not use this fact, we omit a detailed proof.

Proposition 3.11. *In the setting of Definition 3.4, write U_S for the open cone in $\operatorname{Spec} R$ whose complement is defined by S. Then there is a natural map of R-modules*

$$\sigma(S) : \operatorname{gr}(X, \mathcal{F}(S)) \to \tilde{M}(U_S)$$

giving rise to a commutative diagram

$$\begin{array}{ccc} \operatorname{gr}(X,\mathcal{F}) & \longrightarrow & \operatorname{gr}(X,\mathcal{F}(S)) \\ \downarrow & & \downarrow \\ M & \longrightarrow & \tilde{M}(U_S). \end{array}$$

Here the first vertical arrow is the isomorphism defining M, the upper horizontal arrow comes from the inclusion $\mathcal{F}_n(X) \subset \mathcal{F}(S)_n(X)$ of Lemma 3.6, and the lower horizontal arrow is restriction of sections (cf. Lemma 3.3).

If S is finite, then $\sigma(S)$ is injective.

Proof. Write $\sigma_n : X_n \to M^n$ for the quotient map defining M. We may construct $\sigma(S)$ as a family of (symbol) maps

$$\sigma(S)_n : \mathcal{F}(S)_n(X) \to \tilde{M}(U_S)^n \tag{3.12)(a}$$

trivial on $\mathcal{F}(S)_{[n-1]}(X)$. As usual it is convenient to write $X_{[n]}$ for $\mathcal{F}(S)_n(X)$. Then we write $\sigma_{[n]}$ for $\sigma(S)_n$:

$$\sigma_{[n]} : X_{[n]} \to \tilde{M}(U_S)^n. \tag{3.12)(a)'\quad(}$$

Fix an element $x \in X_{[n]}$. Choose N_0 as in Lemma 3.6, and fix $N > N_0$. We want to define an element

$$\sigma_{[n]}(x) = m = (m_i)_{i\in I} \in \tilde{M}(U_S) \tag{3.12)(b}$$

(cf. (3.2)(d)), with $m_i \in M^n_{f_i}$. By Lemma 3.6, $(\phi_i)^N \cdot x \in X_{n+Np_i}$. Write

$$m_i' = \sigma_{n+Np_i}((\phi_i)^N \cdot x), \qquad m_i = (f_i)^{-N} m_i'. \tag{3.12)(c}$$

The first problem is to show that (m_i) actually belongs to $\tilde{M}(U_S)$. By (3.2)(d), this is equivalent to showing that for $i, j \in I$ we have $f_j^N \cdot m_i' = f_i^N \cdot m_j'$ in the localized module $M_{f_i f_j}$. Of course it suffices to prove the equality in M. By inspection of the definitions, we see that it amounts to

$$(\phi_j)^N(\phi_i)^N \cdot x = (\phi_i)^N(\phi_j)^N \cdot x \pmod{X_{n+Np_i+Np_j-1}}. \tag{*}$$

Now since $\operatorname{gr} A$ is commutative, it is easy to check that $(\phi_j)^N(\phi_i)^N$ is equal to $(\phi_i)^N(\phi_j)^N$ modulo products of type $(2N-2, Np_i + Np_j - 1)$. Lemma 3.6 guarantees that such products map x into $X_{n+Np_i+Np_j-1}$, proving $(*)$. A similar argument shows that $\sigma(S)$ respects the action of R. If $x \in X_n$, then m_i' is equal to $(f_i)^N \cdot (x + X_{n-1})$. It follows that m is the natural image of $\sigma(x)$ in $\tilde{M}(U_S)$, proving the commutativity of the diagram in the proposition.

It remains to prove (assuming S is finite) the injectivity of $\sigma(S)$; that is, that the kernel of $\sigma_{[n]}$ is precisely $X_{[n-1]}$. So suppose $x \in X_{[n]}$, and $\sigma_{[n]}(x) = 0$. This means that for every $i \in I$, the element m_i of M_{f_i} must be zero. By the remarks before (3.2)(c), this means that $f_i^N \cdot m_i' = 0$ (as an element of M) for all large N. By the definition of m_i', this means that there is a positive integer N_i such that

$$(\phi_i)^{N_i} \cdot x \in X_{n+N_ip_i-1}. \qquad (**)$$

We want to deduce from this that $x \in X_{[n-1]}$. We use the criterion of Definition 3.4. Suppose N is larger than the cardinality of S times the maximum of the various N_i, and also larger than N_0; and that $J \in I^N$. We claim that $\phi_J \cdot x \in X_{n+p_J-1}$. This will complete the proof. Clearly there is an i that occurs at least N_i times among the ϕ_{i_j}. Write J' for what is left after N_i i's are removed from J (so that $p_J = p_{J'} + N_ip_i$). Then ϕ_J is equal to $\phi_{J'}\phi_i^{N_i}$ plus a sum of products of type $(N-1, p_J - 1)$. Since $N-1$ is at least N_0, the second kind of product maps x into X_{n+p_J-1}. By $(**)$, the first term has this property as well. Q.E.D.

To see that the localized filtration $\mathcal{F}(S)$ still captures much of the structure of X, we need a simple definition and a condition on X.

Definition 3.13 In the setting of (3.1), set

$$A_{-\infty} = \bigcap_n A_n, \qquad X_{-\infty} = \bigcap_n X_n.$$

When we wish to emphasize the filtration, we may write $\mathcal{F}_{-\infty}(X)$. In the simplified notation at the end of Definition 3.4, we write $X_{[-\infty]}$ for $\mathcal{F}(S)_{-\infty}(X)$.

Notice that $X_{-\infty}$ is automatically an A-submodule of X. (This uses only the compatibility of the filtration and the assumption that the filtration of A is exhaustive.) Combining Proposition 3.11 with Lemma 3.3, we get

Corollary 3.14. *In the setting of Definition 3.4, suppose S is finite. Then the restriction of the original filtration $\mathcal{F}$ to $\mathcal{F}(S)_{-\infty}(X) = X_{[-\infty]}$ defines an injection of* $\mathrm{gr}(X_{[-\infty]})$ *into a submodule of M supported on the closed cone Z_S defined by S (cf. (3.2)(b).*

Suppose in particular that $A_{-1} = 0$, R is Noetherian, and $\mathcal{F}$ is a good filtration of X. Then the characteristic variety of $X_{[-\infty]}$ is contained in Z_S.

When X is irreducible and Z_S does not contain its characteristic variety (or even under various weaker conditions) Corollary 3.14 will allow us to deduce that $X_{[-\infty]}$ must be zero, and therefore that the localized filtration "sees" all of the structure of X. The reason we cannot easily get a great deal of information from this is that the localized module $\tilde{M}(U_S)$ need not be finitely generated (over R), and therefore the localized filtration need not be good. In the next section we will see how it is sometimes possible to circumvent this problem in the case of $(\mathfrak{g}, K)$-modules, using the extra rigidity imposed by the action of K.

4. Microlocalization for $(\mathfrak{g}, K)$-modules

In this section we will apply the results of section 3 to associated varieties of $(\mathfrak{g}, K)$-modules. The main problem will be to find conditions under which the module $\tilde{M}(U_S)$ of (3.2) is finitely generated over $S(\mathfrak{g})$. Before considering this, we give a simpler application. Recall from after Corollary 2.10 the notion of virtual characters of an algebraic group H. Suppose Θ is such a virtual character, and π is a finite-dimensional representation of H. We can find a finite set of distinct irreducible representations $\{ \rho_i \}$ of H and integers m_i so that $\Theta = \sum m_i \rho_i$, and $\pi = \sum n_i \rho_i$ (as a virtual representation). Define the *quotient multiplicity of* Θ *in* π to be

$$[\Theta : \pi]_{H,\mathrm{quo}} = \sum m_i \dim \mathrm{Hom}_H(\pi, \rho_i). \tag{4.1)(a}$$

(We may drop the H if this causes no confusion.) Similarly, the *submodule multiplicity of* Θ *in* π is

$$[\Theta' : \pi]_{H,\mathrm{sub}} = \sum m_i \dim \mathrm{Hom}_H(\rho_i, \pi). \tag{4.1)(b}$$

Finally, the *multiplicity of* Θ *in* π is

$$[\Theta : \pi]_H = \sum m_i n_i. \tag{4.1)(c}$$

If π is completely reducible (for example, if H is reductive) then the three definitions coincide; we may drop the subscripts sub and quo from the notation in that case. These definitions may be extended to the case when π is any rational representation (possibly infinite-dimensional); in that case we need to require either that Θ be genuine, or that π have finite multiplicities, to avoid the appearance of $\infty - \infty$. If Θ is the character of some representation τ, then

$$[\Theta : \pi] \geq [\Theta : \pi]_{\mathrm{quo}} \geq \dim \mathrm{Hom}_H(\pi, \tau) \tag{4.1)(d}$$

with equality whenever π and τ are completely reducible. Of course there is an analogous inequality for submodule multiplicities.

Theorem 4.2. *Suppose X is an irreducible $(\mathfrak{g}, K)$-module (cf. (2.1)), $\lambda \in \mathfrak{g}^*$, and that the K-orbit*

$$\mathcal{O} = K \cdot \lambda \simeq K/K(\lambda)$$

of λ is dense in some irreducible component of the associated variety $\mathcal{V}(X)$. Define a genuine virtual representation $\chi(\lambda, X)$ of $K(\lambda)$ as in section 2. If τ is any representation of K, then

$$\dim \mathrm{Hom}_K(\tau, X) \leq [\chi(\lambda, X) : \tau]_{K(\lambda),\mathrm{quo}}.$$

Because of Frobenius reciprocity, this theorem can be regarded as a precise version of the remarks after Theorem 2.13. (The proof should make this clearer.) We are most interested in the case when K is reductive and τ is irreducible. Then the left side is the multiplicity of τ as a K-type of X. A somewhat surprising feature of the result is that one needs only a single component of $\mathcal{V}(X)$ to control all of the K-types of X.

Proof. Fix a good filtration $\mathcal{F}$ of X, and write $M = \operatorname{gr} X$ as in sections 2 and 3. Necessarily $\mathcal{O}$ is open in $\mathcal{V}(X)$, so its complement in $\mathcal{V}(X)$ is a closed cone in $\mathfrak{g}^*$. Choose a finite homogeneous set S in $S(\mathfrak{g})$ defining this complement: in the notation of (3.2)(b),

$$\mathcal{V}(X) = \mathcal{O} \cup Z_S,$$

a disjoint union. We may assume that the ideal generated by S (or even the linear span of S) is K-invariant. The open cone U_S (cf. (3.2)(a)) meets $\mathcal{V}(X)$ precisely in $\mathcal{O}$. Let $\mathcal{F}(S)$ be the localized filtration of X (Definition 3.4); this will be K-invariant by Corollary 3.10. (The proof of that corollary becomes almost trivial when S and S' have the same linear span, which is the only case we need.) Then $X_{[-\infty]}$ is a $(\mathfrak{g}, K)$-submodule of X, with characteristic variety contained in Z_S (Corollary 3.14). Since $\mathcal{V}(X)$ is not contained in Z_S, $X_{[-\infty]}$ is a proper submodule; so it is zero by the irreducibility of X. Proposition 3.11 gives an embedding $[\operatorname{gr}] X \hookrightarrow \tilde{M}(U_S)$, and it follows that

$$\dim \operatorname{Hom}_K(\tau, X) \leq \dim \operatorname{Hom}_K(\tau, [\operatorname{gr}] X) \leq \dim \operatorname{Hom}_K(\tau, \tilde{M}(U_S)). \tag{4.3)(a}$$

Now choose a finite filtration of M as in Lemma 2.11. The functor taking a module N to $\tilde{N}(U_S)$ is left exact, as is clear from (3.2) and the exactness of localization. Consequently

$$\dim \operatorname{Hom}_K(\tau, \tilde{M}(U_S)) \leq \sum_i \dim \operatorname{Hom}_K(\tau, (\widetilde{M_i/M_{i-1}})(U_S)). \tag{4.3)(b}$$

But U_S meets $\mathcal{V}(X)$ only in the orbit $\mathcal{O}$, which may be identified with the homogeneous space $K/K(\lambda)$. The assumption that M_i/M_{i-1} is generically reduced along $\mathcal{O}$ implies by K-invariance that it is reduced everywhere on $\mathcal{O}$; so the restriction of the sheaf $\widetilde{M_i/M_{i-1}}$ to U_S may be identified with a K-equivariant sheaf of modules on the homogeneous space $K/K(\lambda)$. Such a sheaf of modules is necessarily the equivariant vector bundle induced by the (fiber) representation of $K(\lambda)$ on $E_i = M_i/(\mathfrak{m}(\lambda)M_i + M_{i-1})$. By Frobenius reciprocity,

$$\dim \operatorname{Hom}_K(\tau, (\widetilde{M_i/M_{i-1}})(U_S)) = \dim \operatorname{Hom}_{K(\lambda)}(\tau, E_i). \tag{4.3)(c}$$

By (4.1)(d), this is bounded above by the quotient multiplicity of E_i in τ.

Now the sum of the various E_i is $\chi(\lambda, X)$ (Definition 2.12); so the inequalities in (4.3) give the conclusion of the theorem. Q.E.D.

In order to decide when a localized filtration is good, we need to know when the R-module $\tilde{M}(U_S)$ is finitely generated.

Theorem 4.4 ([Grothendieck], Proposition 5.11.1). *Suppose R is a finitely generated commutative algebra over a field k, M is a finitely generated R-module, and U is an open* *set in* Spec R. *Write* Z *for the complement of* U *in* Spec R. *Then the* R*-module* $\tilde{M}(U)$ *is finitely generated if and only if for every prime ideal* $P \in U \cup \operatorname{Ass} M$, *the closure* Y *of* P *in* Spec R *satisfies*

$$Y \cap Z \text{ has codimension at least 2 in } Y.$$

Needless to say, Grothendieck actually proves a much more general result.

To check the condition in this theorem, we will use the following easy lemma.

Lemma 4.5. *Suppose* M *is a finitely generated* $(S(\mathfrak{g}), K)$*-module. Then* K *(through its action on* $S(\mathfrak{g})$*) permutes the set of associated primes of* M*; and the identity component* K_0 *preserves each associated prime.*

Proof. The first claim is obvious. Since $\operatorname{Ass} M$ is finite, the stabilizer in K of each associated prime must be a closed subgroup of K of finite index. Such a subgroup contains the identity component. Q.E.D.

Theorem 4.6. *In the setting of (2.1), suppose* X *is a finitely generated* $(\mathfrak{g}, K)$*-module. Fix* $\lambda \in \mathfrak{g}^*$ *belonging to* $\mathcal{V}(X)$, *and assume that*

i) the closure of $\mathcal{O} = K \cdot \lambda$ *contains an irreducible component of* $\mathcal{V}(X)$*; and*

ii) if X' *is a non-zero* $(\mathfrak{g}, K)$*-submodule of* X, *then* $\mathcal{V}(X') \supset \mathcal{O}$.

Define $\partial\mathcal{O}$ *to be the complement of* $\mathcal{O}$ *in its closure* $\overline{\mathcal{O}}$. *Then there are two (non-exclusive) possibilities: either*

a) $\mathcal{V}(X)$ *is equal to* $\overline{\mathcal{O}}$*; or*

b) $\partial\mathcal{O}$ *has codimension one in* $\overline{\mathcal{O}}$.

Proof. Choose a good filtration $\mathcal{F}$ of X, and write $M = \operatorname{gr} X$ as usual. Define Z to be the complement of $\mathcal{O}$ in $\mathcal{V}(X)$. Then Z is a closed cone in $\mathfrak{g}^*$, and $Z \cap \overline{\mathcal{O}} = \partial\mathcal{O}$. Choose a finite homogeneous subset S of $S(\mathfrak{g})$ defining Z, and construct the localized filtration $\mathcal{F}(S)$ as in section 3. By Corollary 3.9, this is a compatible exhaustive filtration of X. By Corollary 3.14 and hypothesis (ii) of the theorem, $X_{[-\infty]} = 0$.

Suppose for the rest of the proof that conclusion (b) of the theorem fails; that is, that $Z \cap \overline{\mathcal{O}}$ has codimension at least two in $\overline{\mathcal{O}}$. Theorem 4.4 implies that $\tilde{M}(U_S)$ is finitely generated. To see that, suppose P is an associated prime of M not belonging to Z. By Lemma 4.5, the associated variety Y of P meets $\mathcal{O}$ in a K_0-invariant subset Y_0. Consequently Y_0 is open in $\mathcal{O}$, so $Z \cap Y \subset Z \cap \overline{\mathcal{O}}$ has codimension at least two in Y.

Now Proposition 3.11 implies that $\operatorname{gr}(X, \mathcal{F}(S))$ is finitely generated. By the definition at (2.1), it follows that $\mathcal{F}(S)$ is a good filtration of X. Now it is easy to check that the associated primes of $\tilde{M}(U)$ must belong to U; so in our case the associated primes of $\operatorname{gr}(X, \mathcal{F}(S))$ must correspond to the connected components of $\mathcal{O}$. Since the associated primes are dense in the characteristic variety, we get $\mathcal{V}(X) = \overline{\mathcal{O}}$. Q.E.D.

Corollary 4.7 ([Borho-Brylinski], [Joseph]). *Suppose* $\mathfrak{g}$ *is a reductive Lie algebra, and* $I \subset U(\mathfrak{g})$ *is a primitive ideal. Then the associated variety* $\mathcal{V}(I)$ *(defined using the Poincaré-Birkhoff-Witt theorem to be the variety in* $\mathfrak{g}^*$ *corresponding to the ideal* $\operatorname{gr} I$ *in* $S(\mathfrak{g})$*) is the closure of a single nilpotent coadjoint orbit in* $\mathfrak{g}^*$.

Proof. The expert will recognize this as almost an immediate consequence of Theorem 4.6; but then the expert already knew the result. For the benefit of other readers, we will give a more complete argument, including sketches of proofs for some standard intermediate results. To begin, we need some auxiliary definitions.

Suppose G is a connected reductive algebraic group with Lie algebra $\mathfrak{g}$. The adjoint action Ad of G on $\mathfrak{g}$ extends to an algebraic action (still denoted Ad) of G on $U(\mathfrak{g})$ by algebra automorphisms. The differential of this action, denoted ad, sends an element $Y \in \mathfrak{g}$ to the derivation

$$(\operatorname{ad} Y)(u) = Yu - uY \qquad (u \in U(\mathfrak{g})). \tag{4.8)(a}$$

Because I is a two-sided ideal, $\operatorname{ad}(\mathfrak{g})$ preserves I. Since G is connected, $\operatorname{Ad}(G)$ preserves I as well. We therefore get an action (still called Ad) of G on $U(\mathfrak{g})/I$ by algebra automorphisms. The differential of this action is given by a formula like (4.8)(a).

On the other hand, the Lie algebra $\mathfrak{g} \times \mathfrak{g}$ acts on $U(\mathfrak{g})$ by

$$(Y_1, Y_2) \cdot u = Y_1 u - uY_2 \qquad (u \in U(\mathfrak{g}), Y_i \in \mathfrak{g}. \tag{4.8)(b}$$

The defining relations of $U(\mathfrak{g})$ show that this action is a Lie algebra representation. Again the fact that I is a two-sided ideal implies that the action factors to $U(\mathfrak{g})/I$. The diagonal embedding of $\mathfrak{g}$ in $\mathfrak{g} \times \mathfrak{g}$ now provides the structure considered in (2.1), and (4.8)(a) and (b) make $U(\mathfrak{g})/I$ into a $(\mathfrak{g} \times \mathfrak{g}, G)$-module. We express this by saying that $U(\mathfrak{g})/I$ is a *Harish-Chandra bimodule.* We sometimes write $\mathfrak{g}_\Delta \subset \mathfrak{g} \times \mathfrak{g}$ for the diagonal subalgebra, which will play the role of $\mathfrak{k}$.

Restriction of linear functionals to the first factor provides a G-equivariant isomorphism

$$(\mathfrak{g} \times \mathfrak{g}/\mathfrak{g}_\Delta)^* \to \mathfrak{g}^* \tag{4.8)(c}$$

which we use to identify $\mathcal{V}(U(\mathfrak{g})/I)$ (defined by (1.2)) with the associated variety $\mathcal{V}(I)$ defined in the statement of the corollary. (The main point is that the standard filtration of $U(\mathfrak{g})$ defines by passage to the quotient a good filtration of the $(\mathfrak{g} \times \mathfrak{g}, G)$-module $U(\mathfrak{g})/I$.)

Write

$$\mathfrak{Z}(\mathfrak{g}) = \text{ center of } U(\mathfrak{g}); \tag{4.9}$$

this is the subalgebra of $U(\mathfrak{g})$ on which the adjoint action of G is trivial. Because I is primitive, it contains a maximal ideal $\mathcal{I}$ in $\mathfrak{Z}(\mathfrak{g})$. It follows at once that the $(\mathfrak{g} \times \mathfrak{g}, G)$-module $U(\mathfrak{g})/I$ is annihilated by the maximal ideal

$$\mathfrak{Z}(\mathfrak{g}) \otimes \mathcal{I} + \mathcal{I} \otimes \mathfrak{Z}(\mathfrak{g}) \subset \mathfrak{Z}(\mathfrak{g} \times \mathfrak{g}).$$

Consequently the Harish-Chandra bimodule $U(\mathfrak{g})/I$ is finitely generated and annihilated by an ideal of finite codimension in the center of the enveloping algebra. By one of Harish-Chandra's basic finiteness theorems, it follows that $U(\mathfrak{g})/I$ has finite length as a bimodule. In particular, it satisfies ACC and DCC on two-sided ideals. Following [Duflo], we can therefore choose a minimal two-sided ideal J properly containing I. Because I is prime, J is unique. It follows that every non-zero submodule X' of $U(\mathfrak{g})/I$ must contain J/I. In particular, we must have

$$\mathcal{V}(X') \supset \mathcal{V}(J/I). \tag{4.10}$$

We will use this in a moment to deduce hypothesis (ii) in Theorem 4.6.

The finiteness under the center of the enveloping algebra also guarantees that (for any subquotient X of $U(\mathfrak{g})/I$) $\mathcal{V}(X)$ consists of nilpotent elements (see Corollary 5.4 below). Since G has only a finite number of nilpotent coadjoint orbits (Theorem 5.8 below) $\mathcal{V}(X)$ must be a finite union of these, of dimensions bounded above by $\operatorname{Dim} X$. Now the additivity of associated varieties implies that

$$\mathcal{V}(U(\mathfrak{g})/I) = \mathcal{V}(U(\mathfrak{g})/J) \cup \mathcal{V}(J/I).$$

On the other hand, J properly contains the prime ideal I. By a theorem in [Borho-Kraft], it follows that $\operatorname{Dim} U(\mathfrak{g})/I > \operatorname{Dim} U(\mathfrak{g})/J$. Assembling these observations, we find that $\mathcal{V}(J/I)$ must contain a G-orbit $\mathcal{O}$ of dimension equal to the Gelfand-Kirillov dimension of $U(\mathfrak{g})/I$, and that the closure of $\mathcal{O}$ will in that case be an irreducible component of $\mathcal{V}(U(\mathfrak{g})/I)$. In conjunction with (4.10), this establishes the hypotheses for Theorem 4.6.

To complete the argument, recall that any coadjoint orbit of a Lie group carries a natural symplectic structure, and is therefore even-dimensional. The boundary of $\mathcal{O}$ (as a finite union of nilpotent coadjoint orbits) must therefore have even codimension in the closure of $\mathcal{O}$. This rules out conclusion (b) of Theorem 4.6; and (a) is what we wished to show. Q.E.D.

As a final application, we return to the problem of K-multiplicities.

Theorem 4.11. *In the setting of (2.1), suppose X is a finitely generated $(\mathfrak{g}, K)$-module. Fix $\lambda \in \mathfrak{g}^*$ belonging to $\mathcal{V}(X)$, and define*

$$\mathcal{O} = K \cdot \lambda = K/K(\lambda), \qquad \partial\mathcal{O} = \overline{\mathcal{O}} - \mathcal{O}.$$

Assume that

i) $\overline{\mathcal{O}}$ contains an irreducible component of $\mathcal{V}(X)$;
ii) if X' is a non-zero $(\mathfrak{g}, K)$-submodule of X, then $\mathcal{V}(X') \supset \mathcal{O}$; and
iii) $\partial\mathcal{O}$ has codimension at least two in $\overline{\mathcal{O}}$.

Define a genuine virtual representation $\chi(\lambda, X)$ of $K(\lambda)$ as in section 2, and choose a completely reducible representation $V(\lambda, X)$ of $K(\lambda)$ representing $\chi(\lambda, X)$. Then there is a finitely generated $(S(\mathfrak{g}), K)$-module Q supported on $\partial\mathcal{O}$, with the property that

$$X = Ind_{K(\lambda)}^{K}(V(\lambda, X)) - Q$$

as virtual representations of K. If K is reductive, this means that the multiplicity in X of any irreducible representation τ of K is

$$\dim \operatorname{Hom}_K(\tau, X) = \dim \operatorname{Hom}_{K(\lambda)}(\tau\mid_{K(\lambda)}, V(\lambda, X)) - \dim \operatorname{Hom}_K(\tau, Q).$$

This follows from the argument given for Theorem 4.2, in conjunction with Theorem 4.4. Because the details are a little delicate, and we will use this result (in section 12) only as evidence for a conjecture, we omit the details.

5. Associated varieties for $(\mathfrak{g}, K)$-modules

In this section we recall those basic facts about associated varieties for $(\mathfrak{g}, K)$-modules that depend on the structure theory of reductive Lie groups. The main result is Corollary 5.23; Corollary 5.20 will also be crucial for the proof of Theorem 8.7.

For this section, we will assume that

$$\mathfrak{g} \text{ is an algebraic Lie algebra.} \tag{5.1)(a}$$

As in (4.8) and (4.9), it is often convenient to fix a connected algebraic group G with Lie algebra $\mathfrak{g}$. Then G acts by Ad on the algebra $S(\mathfrak{g})$ of polynomial functions on $\mathfrak{g}^*$; we write

$$S(\mathfrak{g})^G \tag{5.1)(b}$$

for the algebra of Ad(G)-invariant polynomials, and

$$\mathfrak{Z}(\mathfrak{g}) = U(\mathfrak{g})^G \tag{5.1)(c}$$

for the algebra of G-invariants in $U(\mathfrak{g})$. (If G is connected, $\mathfrak{Z}(\mathfrak{g})$ is the center of $U(\mathfrak{g})$.)

Lemma 5.2. *Filter $\mathfrak{Z}(\mathfrak{g})$ by the restriction of the standard filtration of $U(\mathfrak{g})$. Then*

$$\operatorname{gr} \mathfrak{Z}(\mathfrak{g}) \simeq S(\mathfrak{g})^G.$$

Proof. Obviously the symbol maps

$$\sigma_n : U_n(\mathfrak{g})/U_{n-1}(\mathfrak{g}) \to S^n(\mathfrak{g})$$

of the Poincaré-Birkhoff-Witt theorem restrict to an inclusion

$$\sigma : \operatorname{gr} \mathfrak{Z}(\mathfrak{g}) \hookrightarrow S(\mathfrak{g})^G.$$

We must prove that σ is surjective. The symmetrization map β provides a degree-preserving Ad(G)-equivariant map from $S(\mathfrak{g})$ to $U(\mathfrak{g})$. When restricted to homogeneous polynomials, β is a one-sided inverse for the symbol maps. That is, if p is a homogeneous polynomial of degree n, then $\sigma_n(\beta(p)) = p$. Since β respects the adjoint action, it maps $S(\mathfrak{g})^G$ into $\mathfrak{Z}(\mathfrak{g})$. The surjectivity of σ follows immediately. Q.E.D.

Corollary 5.3. *Suppose $\mathcal{I}$ is a proper ideal of finite codimension in $\mathfrak{Z}(\mathfrak{g})$. Then $\operatorname{gr}\mathcal{I}$ is a proper graded ideal of finite codimension in $S(\mathfrak{g})^G$. Its radical is the ideal $S^+(\mathfrak{g})^G$ of invariant polynomials without constant term.*

Proof. Only the last assertion requires comment. The radical will be the intersection of the maximal ideals containing $\operatorname{gr}\mathcal{I}$. These must form a cone (since $\operatorname{gr}\mathcal{I}$ is graded) and a finite set (since $\operatorname{gr}\mathcal{I}$ has finite codimension). A set of maximal ideals in an $\mathbb{N}$-graded $\mathbb{C}$-algebra with these two properties corresponds to maximal ideals in the degree zero subalgebra. In our case this subalgebra is reduced to the constants, so the only maximal ideal involved is $S^+(\mathfrak{g})^G$. Q.E.D.

Corollary 5.4. *Suppose X is a $\mathfrak{g}$-module of finite length. Then the associated variety $\mathcal{V}(X)$ is contained in the cone $\mathcal{N}^*$ defined by $S^+(\mathfrak{g})^G$:*

$$\mathcal{N}^* = \{\, \lambda \in \mathfrak{g}^* \mid p(\lambda) = 0, \text{ all } p \in S^+(\mathfrak{g})^G \,\}.$$

Proof. Any irreducible $\mathfrak{g}$-module is annihilated by a maximal ideal in $\mathfrak{Z}(\mathfrak{g})$; so any $\mathfrak{g}$-module of finite length is annihilated by the product of a finite number of maximal ideals. Such a product is of finite codimension in $\mathfrak{Z}(\mathfrak{g})$. Now apply Corollary 5.3 and (1.2). Q.E.D.

Because of Corollary 5.4, we turn our attention now to the cone $\mathcal{N}^*$. Write Ad^* for the coadjoint action of G on $\mathfrak{g}^*$. For any $\lambda \in \mathfrak{g}^*$, define

$$G(\lambda) = \{\, g \in G \mid \mathrm{Ad}^*(g)(\lambda) = \lambda \,\}, \qquad \mathfrak{g}(\lambda) = \mathrm{Lie}(G(\lambda). \tag{5.5)(a}$$

More generally, if H is a Lie group acting by automorphisms on $\mathfrak{g}$, we will define $H(\lambda)$ analogously; this notation will be applied particularly with $H = K$ in the setting of (2.1), and with $H = G_{\mathbf{R}}$ (a real Lie group with complexified Lie algebra $\mathfrak{g}$) in the setting of the introduction. It is an elementary exercise to verify that

$$\mathfrak{g}(\lambda) = \{\, x \in \mathfrak{g} \mid \text{ for all } y \in \mathfrak{g}, \lambda([x,y]) = 0 \,\}. \tag{5.5)(b}$$

We say that λ is *nilpotent* if its restriction to the subalgebra $\mathfrak{g}(\lambda)$ is zero; that is, if

$$\lambda([x,\mathfrak{g}]) = 0 \Rightarrow \lambda(x) = 0. \tag{5.5)(c}$$

When G is reductive, we are going to see that $\mathcal{N}^*$ is precisely the cone of nilpotent elements (and that this definition of nilpotent agrees with more familiar ones). Here is a preliminary step.

Lemma 5.6. *In the setting of (5.5), identify the tangent space $T_\lambda(G \cdot \lambda)$ at with a subspace of the ambient vector space $\mathfrak{g}^*$. Then*

$$T_\lambda(G \cdot \lambda) = \{\, \mu \in \mathfrak{g}^* \mid \mu\,|_{\mathfrak{g}(\lambda)} = 0 \,\}.$$

Proof. By general results on homogeneous spaces, the tangent space in question may be obtained by applying the differentiated (coadjoint) action of G to λ. A typical element is therefore of the form $\mathrm{ad}^*(x)(\lambda)$. We check that such an element satisfies the condition in the lemma. Fix y in $\mathfrak{g}(\lambda)$. Then

$$\mathrm{ad}^*(x)(\lambda)(y) = -\lambda(\mathrm{ad}(x)(y)) = \lambda([y,x]) = 0,$$

the last equality coming from (5.5)(b). On the other hand, both spaces in the lemma have dimension equal to the codimension of $\mathfrak{g}(\lambda)$ in $\mathfrak{g}$. The containment we have just proved therefore implies their equality. Q.E.D.

Theorem 5.7. *Suppose G is a complex connected algebraic group with Lie algebra $\mathfrak{g}$, and $\lambda \in \mathfrak{g}^*$. The following conditions are equivalent.*

a) λ is nilpotent; that is, $\lambda \mid_{\mathfrak{g}(\lambda)} = 0$.
b) $\lambda \in T_\lambda(G \cdot \lambda)$
c) There is an element $x \in \mathfrak{g}$ such that $\mathrm{ad}^*(x)\lambda = \lambda$
d) For all non-zero complex numbers t, $t\lambda \in G \cdot \lambda$.
e) For infinitely many complex numbers t, $t\lambda \in G \cdot \lambda$.

If in addition G is reductive, these are also equivalent to

f) There is a Borel subalgebra $\mathfrak{b}$ of $\mathfrak{g}$ such that $\lambda \mid_{\mathfrak{b}} = 0$
g) $0 \in \overline{G \cdot \lambda}$
h) $\lambda \in \mathcal{N}^*$; that is, every G-invariant polynomial without constant term vanishes at λ.

Proof. We first prove the equivalence of (a)–(e). The equivalence of (a) and (b) is Lemma 5.6, and that of (b) and (c) is formal (see the proof of Lemma 5.6). Exponentiating (c) gives

$$\mathrm{Ad}^*(\exp(sx))\lambda = e^s \cdot \lambda,$$

which implies (d); and (e) follows from (d) . Conversely, assume (e). The set of complex numbers for which the condition in (e) holds is automatically a subgroup of $\mathbb{C}^\times$. It therefore contains a sequence converging to 1, and (c) follows.

Next, we show that (c) implies (f). By the definition of ad^*, (c) is equivalent to

$$\lambda \mid_{\mathrm{im}(1+\mathrm{ad}\, x)} = 0.$$

The image in question (call it $\mathfrak{s}$) contains all of the generalized eigenspaces of $\mathrm{ad}\, x$ except that for the eigenvalue -1. On the other hand, the sum of all the generalized eigenspaces of $\mathrm{ad}\, x$ corresponding to eigenvalues with non-negative real part is always a parabolic subalgebra of $\mathfrak{g}$. (Here we are using for the first time the assumption that $\mathfrak{g}$ is reductive.) It follows that $\mathfrak{s}$ contains a parabolic subalgebra, and hence a Borel subalgebra. This is (f).

To see that (f) implies (g), suppose B is the Borel subgroup of G corresponding to $\mathfrak{b}$. There is an element h of B such that $\mathrm{Ad}^*(h)$ has only real eigenvalues strictly smaller than one on $(\mathfrak{g}/\mathfrak{b})^*$. It follows that $\lim_{n\to\infty} \mathrm{Ad}(h^n)(\lambda) = 0$, which implies (g). That (g) implies (h) is easy (since G-invariant polynomials are constant on orbit closures).

To complete the proof, it is enough to show that (h) implies (e). This seems to be deeper than the rest of the argument. We will use the following result of Kostant.

Theorem 5.8 ([Kostant]). *If G is reductive, the cone $\mathcal{N}^*$ (Corollary 5.4) is the union of a finite number of orbits of G.*

Assume now that $\lambda \in \mathcal{N}^*$. Since $\mathcal{N}^*$ is a cone, all multiples of λ belong to it as well. By Theorem 5.8, we can find infinitely many in a single G-orbit. If one of these is $t\lambda$, then multiplying by t^{-1} gives infinitely many multiples of λ in the orbit of λ, as required. Q.E.D.

Although we will make every effort to avoid doing so, it is occasionally convenient to use an identification of $\mathfrak{g}^*$ with $\mathfrak{g}$.

Lemma 5.9. *Suppose G is a reductive Lie group with Lie algebra $\mathfrak{g}$. Then there is a non-degenerate symmetric G-invariant bilinear form $\langle , \rangle$ on $\mathfrak{g}$. If $\mathfrak{p}$ is any parabolic subalgebra of $\mathfrak{g}$, with nil radical $\mathfrak{n}$, then*

$$\mathfrak{p}^\perp = \textit{ radical of } \langle , \rangle \mid_{\mathfrak{p}} = \mathfrak{n}.$$

This is standard and easy: one can add to the Killing form any non-degenerate symmetric form on the center. Such a form gives a G-equivariant identification

$$\mathfrak{g}^* \simeq \mathfrak{g}, \qquad \lambda \mapsto x_\lambda. \tag{5.10}$$

Because G is an algebraic group, there is a notion of semisimple and nilpotent elements in $\mathfrak{g}$ (and a Jordan decomposition). In any algebraic Lie algebra, an element x is nilpotent if and only if it belongs to the nil radical of some Borel subalgebra. (This is not quite the definition, but it follows immediately from the fact every element belongs to a Borel subalgebra.)

Corollary 5.11 *In the setting of Theorem 5.7, suppose G is reductive. Then the isomorphism of (5.10) identifies $\mathcal{N}^*$ with the cone $\mathcal{N}$ of nilpotent elements in $\mathfrak{g}$.*

Proof. Suppose $\lambda \in \mathfrak{g}^*$ corresponds to $x_\lambda \in \mathfrak{g}$. If $\mathfrak{b}$ is a Borel subalgebra (or indeed any subspace of $\mathfrak{g}$) then $\lambda \mid_{\mathfrak{b}} = 0$ if and only if $x_\lambda \in \mathfrak{b}^\perp$. Now apply condition (f) of Theorem 5.7, Lemma 5.9, and the remarks preceding the corollary. Q.E.D.

Except under special additional hypotheses on the module X (as in the proof of Corollary 4.7, for example) there is no reason for the group G to act on an associated variety $\mathcal{V}(X)$. The basic finiteness result in Theorem 5.8 is therefore insufficient for us. We first recall the additional structure used in section 2.

Definition 5.12 A *pair* is a pair (G, K) of algebraic groups endowed with
i) an inclusion of Lie algebras $i : \mathfrak{k} \to \mathfrak{g}$; and
ii) an algebraic action Ad of K on $\mathfrak{g}$ by automorphisms, compatible with i.
Since the conditions (i) and (ii) refer only to $\mathfrak{g}$, we may also speak of the pair $(\mathfrak{g}, K)$. Define the *nilpotent cone for the pair* by

$$\begin{aligned}\mathcal{N}^*_{\mathfrak{k}} &= \{ \lambda \in \mathcal{N}^* \mid \lambda \mid_{\mathfrak{k}} = 0 \} \\ &= \mathcal{N}^* \cap (\mathfrak{g}/\mathfrak{k})^*.\end{aligned}$$

From Corollary 5.4 and (1.2)(c) we get at once

Corollary 5.13. *Suppose X is a $(\mathfrak{g}, K)$-module of finite length. Then the associated variety $\mathcal{V}(X)$ is a union of K-orbits contained in the cone $\mathcal{N}^*_{\mathfrak{k}}$ of Definition 5.12.*

We must therefore study the orbits of K on $\mathcal{N}_{\mathfrak{k}}^*$. Many of our results are most conveniently phrased in terms of the symplectic structure on a coadjoint orbit, which we therefore recall. For details the reader may consult for example [Guillemin-Sternberg]. Fix a G-orbit

$$\mathcal{O} \subset \mathfrak{g}^*. \tag{5.14)(a}$$

For each $\lambda \in \mathcal{O}$, define a skew-symmetric bilinear form ω_λ on $\mathfrak{g}$ by

$$\omega_\lambda(x, y) = \lambda([x, y]). \tag{5.14)(b}$$

The radical of ω_λ is $\mathfrak{g}(\lambda)$ (by (5.5)(b)), so ω_λ defines a non-degenerate symplectic form

$$\omega_\lambda \text{ on } \mathfrak{g}/\mathfrak{g}(\lambda) \simeq T_\lambda(\mathcal{O}). \tag{5.14)(c}$$

Evidently these forms fit together to define a smooth algebraic 2-form ω on $\mathcal{O}$. One checks fairly easily that ω is closed, so $\mathcal{O}$ is in a natural G-invariant way a complex symplectic manifold. (In particular, the dimension of $\mathcal{O}$ is even; recall that this was critical to the proof of Corollary 4.7.) Write $\mathrm{Sp}(\omega_\lambda)$ for the group of linear transformations of $T_\lambda(\mathcal{O})$ preserving the form ω_λ. Then the isotropy action at λ gives a natural homomorphism

$$G(\lambda) \to \mathrm{Sp}(\omega_\lambda). \tag{5.14)(d}$$

A somewhat different construction of ω_λ, emphasizing the "Poisson structure," is implicit in sections 10 and 11 (see (11.13)).

A submanifold Y of a symplectic manifold (Z, ω) is called *isotropic* if the symplectic form restricts to zero on each tangent space; that is, if

$$T_y(Y)^\perp \supset T_y(Y) \tag{5.15)(a}$$

for every $y \in Y$. (Here we identify $T_y(Y)$ with a subspace of $T_y(Z)$, and form the perpendicular with respect to the symplectic form ω_y. We say that Y is *coisotropic* if

$$T_y(Y)^\perp \subset T_y(Y), \tag{5.15)(b}$$

and *Lagrangian* if it is both isotropic and coisotropic; that is, if

$$T_y(Y)^\perp = T_y(Y). \tag{5.15)(c}$$

If Z is algebraic, we could make these definitions for arbitrary subvarieties, or even subschemes, using the Zariski tangent space. This leads to a conflict with standard terminology, however: a subvariety is usually called Lagrangian if its smooth locus is Lagrangian. Such a subvariety will satisfy (5.15)(b) but not (5.15)(c) at singular points y. We will therefore try to avoid the terminology in the singular case.

Proposition 5.16. *Suppose (G, K) is a pair (Definition 5.12), $\lambda \in (\mathfrak{g}/\mathfrak{k})^*$, and $\mathcal{O} = G \cdot \lambda$ is a coadjoint orbit of dimension $2n$. Define*

$$\mathcal{O}_{\mathfrak{k}} = \mathcal{O} \cap (\mathfrak{g}/\mathfrak{k})^*,$$

which we regard as a subscheme of the algebraic variety $\mathcal{O}$.

a) *The orbit $K \cdot \lambda$ is a smooth isotropic subvariety of $\mathcal{O}$ (cf. (5.15)(a)), contained in $\mathcal{O}_{\mathfrak{k}}$. In particular, $\dim K \cdot \lambda \leq n$.*

b) *The subscheme $\mathcal{O}_{\mathfrak{k}}$ of $\mathcal{O}$ is coisotropic, in the sense that its Zariski tangent space at each (closed) point λ satisfies (5.15)(b). In particular, the dimension of each such tangent space is at least n.*

In (b), we do not claim that $\mathcal{O}_{\mathfrak{k}}$ has dimension at least n; this scheme could be non-reduced, and so could have no smooth points. I know of no example where this happens, however.

Proof. The orbit $K \cdot \lambda$ is smooth because it is homogeneous. Its tangent space at λ is $\mathfrak{k}/\mathfrak{k}(\lambda)$. If x and y are elements of $\mathfrak{k}$ (representing tangent vectors) then $[x,y]$ also belongs to $\mathfrak{k}$. Since λ is assumed to vanish on $\mathfrak{k}$, it follows that

$$\omega_\lambda(x,y) = \lambda([x,y]) = 0.$$

This is (a).

For (b), we compute first that

$$T_\lambda(K \cdot \lambda)^\perp = \{\, y + \mathfrak{g}(\lambda) \in \mathfrak{g}/\mathfrak{g}(\lambda) \mid \omega_\lambda(y, \mathfrak{k}) = 0 \,\}$$

. Lemma 5.6 allows us to identify this with a subspace of $\mathfrak{g}^*$, namely

$$\begin{aligned} &\{\, \mu \in \mathfrak{g}^* \mid \mu\,|_{\mathfrak{k}+\mathfrak{g}(\lambda)} = 0 \,\} \\ =&\{\, \mu \in \mathfrak{g}^* \mid \mu\,|_{\mathfrak{g}(\lambda)} = 0 \,\} \cap \{\, \mu \in \mathfrak{g}^* \mid \mu\,|_{\mathfrak{k}} = 0 \,\} \\ =&T_\lambda(G \cdot \lambda) \cap T_\lambda((\mathfrak{g}/\mathfrak{k})^*). \end{aligned}$$

Now the Zariski tangent space of an intersection is the intersection of the tangent spaces. (Here it is essential to take scheme-theoretic intersection; this assertion would not be true in general if we considered only the underlying variety.) We have therefore shown that

$$T_\lambda(K \cdot \lambda)^\perp = T_\lambda\left(G \cdot \lambda \cap (\mathfrak{g}/\mathfrak{k})^*\right). \tag{5.17}$$

Now (b) follows from (a). Q.E.D.

Corollary 5.18. *In the setting of Proposition 5.16, the following are equivalent.*

a) $\dim K \cdot \lambda = n$.

b) *$K \cdot \lambda$ is a Lagrangian subvariety of $G \cdot \lambda$.*

c) *The intersection $G \cdot \lambda \cap (\mathfrak{g}/\mathfrak{k})^*$ is reduced at λ, and $K \cdot \lambda$ is open in it.*

Proposition 5.19 ([Kostant-Rallis]). *In the setting of Proposition 5.16, assume that $\mathfrak{k}$ is the algebra of fixed points of an involutive automorphism θ of $\mathfrak{g}$. Then the conditions of Corollary 5.18 are satisfied.*

Proof. Write $\mathfrak{p}$ for the -1 eigenspace of θ on $\mathfrak{g}$, and $\mathfrak{p}(\lambda)$ for its intersection with $\mathfrak{g}(\lambda)$. If x and y belong to $\mathfrak{p}$, then

$$\theta([x,y]) = [\theta(x), \theta(y)] = [-x, -y] = [x,y];$$

so $[x, y]$ belongs to $\mathfrak{k}$. The argument for Proposition 5.16(a) shows that $\mathfrak{p}/\mathfrak{p}(\lambda)$ is an isotropic subspace of $T_\lambda(G \cdot \lambda)$. Since $\mathfrak{g} = \mathfrak{k} + \mathfrak{p}$, $T_\lambda(G \cdot \lambda)$ is spanned by the two isotropic subspaces $\mathfrak{k}/\mathfrak{k}(\lambda)$ and $\mathfrak{p}/\mathfrak{p}(\lambda)$. By linear algebra, the sum must be direct and the subspaces of dimension n. Q.E.D.

Corollary 5.20 ([Kostant-Rallis]). *Suppose (G, K) is a pair (Definition 5.12), and $\lambda \in \mathfrak{g}^*$. Assume that $\mathfrak{k}$ is the algebra of fixed points of an involutive automorphism θ of $\mathfrak{g}$. Then the intersection $G \cdot \lambda \cap (\mathfrak{g}/\mathfrak{k})^*$ is a finite union of K-orbits. It is a smooth reduced Lagrangian subvariety of $G \cdot \lambda$.*

Only the finiteness of the number of K-orbits requires comment; and this follows from the fact that each is open in the intersection.

Definition 5.21 A *reductive symmetric pair* is a pair (G, K) of reductive algebraic groups as in Definition 5.12, endowed in addition with
iii) an involutive automorphism θ of $\mathfrak{g}$, commuting with $\mathrm{Ad}\, K$, with fixed point set $\mathfrak{k}$.
We say that (G, K) is of *Harish-Chandra class* if the automorphisms in $\mathrm{Ad}\, K$ are inner. (This is automatic if K is connected.) We may again speak of the pair $(\mathfrak{g}, K)$.

Corollary 5.22 ([Kostant-Rallis]). *Suppose $(\mathfrak{g}, K)$ is a reductive symmetric pair. Then the cone $\mathcal{N}^*_{\mathfrak{k}}$ of Definition 5.12 is a finite union of K-orbits.*

Corollary 5.23. *Suppose $(\mathfrak{g}, K)$ is a reductive symmetric pair, and X is a $(\mathfrak{g}, K)$-module of finite length. Then the associated variety $\mathcal{V}(X)$ is a finite union of K-orbits in $\mathcal{N}^*_{\mathfrak{k}}$.*

6. Connection with real nilpotent orbits.

In this section we recall results of Kostant-Rallis and Sekiguchi relating the nilpotent K orbits considered in section 5 with real nilpotent orbits. The ultimate goal, about which we will say more in section 7, is to make some philosophical connections between associated varieties and the method of coadjoint orbits.

Definition 6.1. A *real reductive Lie group* $G_{\mathbf{R}}$ is one with the following three properties: $G_{\mathbf{R}}$ has a finite number of connected components; the Lie algebra $\mathfrak{g}_{\mathbf{R}}$ is reductive; and the center of the derived group of the identity component is finite. Such a group has a maximal compact subgroup $K_{\mathbf{R}}$, unique up to conjugacy by $G_{\mathbf{R}}$, and a Cartan involution θ with fixed point group $K_{\mathbf{R}}$. The complexification K of $K_{\mathbf{R}}$ is a complex algebraic group, which acts on the complexification $\mathfrak{g}$ of $\mathfrak{g}_{\mathbf{R}}$. Consequently $(\mathfrak{g}, K)$ is a reductive symmetric pair. Conversely, it can be shown that every reductive symmetric pair arises in this manner from a real reductive group.

Given a real reductive group $G_{\mathbf{R}}$, we identify $\mathfrak{g}^*$ with the space of $\mathbf{R}$-linear maps from $\mathfrak{g}_{\mathbf{R}}$ to $\mathbf{C}$. This is a complex vector space containing as a real form the space $\mathfrak{g}^*_{\mathbf{R}}$ of real-valued linear functionals on $\mathfrak{g}_{\mathbf{R}}$. Taking the real part defines a restriction map

$$\mathrm{Re} : \mathfrak{g}^* \to \mathfrak{g}^*_{\mathbf{R}}$$

analogous to the restriction map from $\mathfrak{g}^*$ to $\mathfrak{k}^*$ used in section 5. The analogue of $(\mathfrak{g}/\mathfrak{k})^*$ is then the space $\mathfrak{g}_{i\mathbf{R}}^*$ of purely imaginary-valued linear functionals on $\mathfrak{g}_{\mathbf{R}}$. The analogue of $\mathcal{N}_{\mathfrak{k}}^*$ is the *imaginary nilpotent cone*

$$\mathcal{N}_{i\mathbf{R}}^* = \{\lambda \in \mathcal{N}^* \mid \operatorname{Re}\lambda = 0\} = \mathcal{N}^* \cap \mathfrak{g}_{i\mathbf{R}}^*.$$

Of course multiplication by i defines a $G_{\mathbf{R}}$-equivariant isomorphism from (for example) $\mathfrak{g}_{i\mathbf{R}}^*$ onto $\mathfrak{g}_{\mathbf{R}}^*$. We could therefore equally well consider the real nilpotent cone, and it is traditional to do this. The aesthetic advantages of $\mathcal{N}_{i\mathbf{R}}^*$ (such as the improved analogy with $\mathcal{N}_{\mathfrak{k}}^*$) were pointed out to me by H. Matumoto. For the moment, notice only that the $G_{\mathbf{R}}$-orbits on $\mathfrak{g}_{i\mathbf{R}}^*$ are a very natural setting for the method of coadjoint orbits.

Theorem 6.2 ([Sekiguchi]). *Suppose that $G_{\mathbf{R}}$ is a real reductive group (Definition 6.1) with maximal compact subgroup $K_{\mathbf{R}}$. Then there is a natural one-to-one correspondence between the (finite) set of $G_{\mathbf{R}}$-orbits on the imaginary nilpotent cone $\mathcal{N}_{i\mathbf{R}}^*$ and the K-orbits on $\mathcal{N}_{\mathfrak{k}}^*$ (Definition 5.12), implemented by Theorem 6.4 below. Suppose that in this correspondence the orbit of $\lambda_{i\mathbf{R}} \in \mathcal{N}_{i\mathbf{R}}^*$ corresponds to that of $\lambda_{\mathfrak{k}} \in \mathcal{N}_{\mathfrak{k}}^*$. Let G be any complex group with Lie algebra $\mathfrak{g}$.*

a) The G-orbits of $\lambda_{\mathfrak{k}}$ and $\lambda_{i\mathbf{R}}$ coincide.

b) $\dim_{\mathbf{R}} G_{\mathbf{R}} \cdot \lambda_{i\mathbf{R}} = 2 \cdot \dim_{\mathbf{C}} K \cdot \lambda_{\mathfrak{k}} = \dim_{\mathbf{C}} G \cdot \lambda_{\mathfrak{k}}$.

c) The maximal compact subgroups of the isotropy groups $K(\lambda_{\mathfrak{k}})$ and $G_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ are isomorphic (canonically, up to inner automorphism).

We recall the outline of Sekiguchi's argument (since we need most of it just to write down the correspondence). Unfortunately we must begin by choosing a non-degenerate symmetric real bilinear form on $\mathfrak{g}_{\mathbf{R}}$, which is invariant under $G_{\mathbf{R}}$ and θ, negative definite on $\mathfrak{k}_{\mathbf{R}}$, and positive definite on the -1-eigenspace $\mathfrak{p}_{\mathbf{R}}$ of θ. Such a form exists and is unique up to a positive real scalar on each simple factor of $\mathfrak{g}_{\mathbf{R}}$. We use it to identify $\mathcal{N}_{i\mathbf{R}}^*$ with $\mathcal{N}_{i\mathbf{R}}$ (the cone of purely imaginary nilpotent elements of the complexified Lie algebra) and $\mathcal{N}_{\mathfrak{k}}^*$ with $\mathcal{N}_{\mathfrak{k}}$ (the cone of nilpotent elements in the -1-eigenspace of θ). In this way the theorem becomes one about nilpotent Lie algebra elements. (The identification of elements made here depends on the choice of the form, but the identification of *orbits* does not. The reason is that multiplication by a positive real scalar on each simple factor sends a nilpotent element to a conjugate one.)

Write σ for the complex conjugation on $\mathfrak{g}$ defining the real form $G_{\mathbf{R}}$. Then σ commutes with θ, and the involution $\tau = \sigma\theta$ of $\mathfrak{g}$ is a Cartan involution for the complex Lie algebra. After the reduction of the previous paragraph, Theorem 6.2 describes a relationship between nilpotent elements in the -1-eigenspaces of σ and θ on $\mathfrak{g}$. We are going to define it by reduction to the case of $\mathfrak{sl}(2,\mathbf{R})$. It is convenient to consider the three (commuting) involutive automorphisms of $\mathfrak{sl}(2)$ analogous to θ, σ, and τ. Recall first that $\mathfrak{sl}(2)$ consists of the two by two complex matrices of trace zero. Then for $x \in \mathfrak{sl}(2)$, set

$$\theta_0 x = -{}^t x \qquad \sigma_0 x = \bar{x} \qquad \tau_0 x = -{}^t\bar{x}. \tag{6.3}$$

Then θ_0 is a complexified Cartan involution for the real form $\mathfrak{sl}(2,\mathbf{R})$; σ_0 is the complex conjugation for $\mathfrak{sl}(2,\mathbf{R})$; and τ_0 is a Cartan involution for $\mathfrak{sl}(2)$. (Equivalently, τ_0 is the complex conjugation for a compact real form of $\mathfrak{sl}(2)$. With this notation, we may reformulate the bijection in Sekiguchi's result as follows.

Theorem 6.4. *Suppose $G_{\mathbf{R}}$ is a real reductive group (Definition 6.1) with Cartan involution θ and maximal compact subgroup $K_{\mathbf{R}}$. Write σ for the complex conjugation on $\mathfrak{g}$ defining $\mathfrak{g}_{\mathbf{R}}$, and K for the complexification of $K_{\mathbf{R}}$. Then the following sets are in natural one-to-one correspondence.*

a) $G_{\mathbf{R}}$-orbits on the cone $\mathcal{N}_{i\mathbf{R}}$ of purely imaginary nilpotent elements in $\mathfrak{g}$.
b) $G_{\mathbf{R}}$-conjugacy classes of homomorphisms $\phi_{i\mathbf{R}}$ from $\mathfrak{sl}(2)$ to $\mathfrak{g}$, intertwining σ_0 with σ. (That is, we require $\sigma(\phi_{i\mathbf{R}}(x)) = \phi_{i\mathbf{R}}(\sigma_0(x))$.)
c) $K_{\mathbf{R}}$-conjugacy classes of homomorphisms $\phi_{i\mathbf{R},\mathfrak{k}}$ from $\mathfrak{sl}(2)$ to $\mathfrak{g}$, intertwining σ_0 with σ and θ_0 with θ.
d) K-conjugacy classes of homomorphisms $\phi_{\mathfrak{k}}$ from $\mathfrak{sl}(2)$ to $\mathfrak{g}$, intertwining θ_0 with θ.
e) K-orbits on the cone $\mathcal{N}_{\mathfrak{k}}$ of nilpotent elements in the -1-eigenspace of θ.

Here the bijections from (c) to (b) and (d) are given by the obvious inclusions; that from (b) to (a) sends (the conjugacy class of) $\phi_{i\mathbf{R}}$ to (the orbit of) $\phi_{i\mathbf{R}}\begin{pmatrix} 0 & i \\ 0 & 0 \end{pmatrix}$; and that from (d) to (e) sends $\phi_{\mathfrak{k}}$ to $\phi_{\mathfrak{k}}\begin{pmatrix} 1/2 & -i/2 \\ -i/2 & -1/2 \end{pmatrix}$.

Sketch of proof. Fix a nilpotent element $x_{i\mathbf{R}} \in \mathcal{N}_{i\mathbf{R}}$. By the Jacobson-Morozov theorem (applied to the real Lie algebra $\mathfrak{g}_{\mathbf{R}}$) we can find elements $y_{i\mathbf{R}} \in \mathcal{N}_{i\mathbf{R}}$ and $h_{i\mathbf{R}} \in \mathfrak{g}_{\mathbf{R}}$ satisfying

$$[h_{i\mathbf{R}}, x_{i\mathbf{R}}] = 2x_{i\mathbf{R}}, \quad [h_{i\mathbf{R}}, y_{i\mathbf{R}}] = -2y_{i\mathbf{R}}, \quad [x_{i\mathbf{R}}, y_{i\mathbf{R}}] = h_{i\mathbf{R}} \tag{6.5)(a}$$

$$\sigma x_{i\mathbf{R}} = -x_{i\mathbf{R}}, \quad \sigma y_{i\mathbf{R}} = -y_{i\mathbf{R}}, \quad \sigma h_{i\mathbf{R}} = h_{i\mathbf{R}}. \tag{6.5)(b}$$

The elements $y_{i\mathbf{R}}$ and $h_{i\mathbf{R}}$ are unique up to conjugation by the centralizer of $x_{i\mathbf{R}}$ in $G_{\mathbf{R}}$. The three elements span the complexification of a real subalgebra $\mathfrak{s}_{\mathbf{R}}$ of $\mathfrak{g}_{\mathbf{R}}$, which is obviously a homomorphic image of $\mathfrak{sl}(2,\mathbf{R})$. Explicitly, we can define a homomorphism $\phi_{i\mathbf{R}}$ by

$$\phi_{i\mathbf{R}}\begin{pmatrix} a & b \\ c & -a \end{pmatrix} = ah_{i\mathbf{R}} - ibx_{i\mathbf{R}} + icy_{i\mathbf{R}}. \tag{6.5)(c}$$

Because of (6.5)(b), the homomorphism $\phi_{i\mathbf{R}}$ intertwines the complex conjugation σ_0 for $\mathfrak{sl}(2,\mathbf{R})$ with σ. This gives the correspondence between (a) and (b) in Theorem 6.4.

To go from (b) to (c), we must conjugate $\phi_{i\mathbf{R}}$ by an element of $G_{\mathbf{R}}$ to make it intertwine Cartan involutions. The standard Cartan involution θ_0 for $\mathfrak{sl}(2,\mathbf{R})$ (negative transpose) is mapped to a certain automorphism θ_0' of the image. On the other hand, it is known that any Cartan involution of a semisimple subalgebra of $\mathfrak{g}_{\mathbf{R}}$ must extend to one on all of $\mathfrak{g}_{\mathbf{R}}$. Consequently θ_0' is the restriction of some Cartan involution θ' of $\mathfrak{g}_{\mathbf{R}}$. Clearly $\phi_{i\mathbf{R}}$ intertwines θ_0 with θ'. Now (by the uniqueness of Cartan involutions) θ' must differ from θ by conjugation by some element g of (the identity component of) $G_{\mathbf{R}}$. Write

$$\phi_{i\mathbf{R},\mathfrak{k}} = \mathrm{Ad}(g) \circ \phi_{i\mathbf{R}}. \tag{6.6}$$

Then $\phi_{i\mathbf{R},\mathfrak{k}}$ intertwines θ_0 with θ and σ_0 with σ. Sekiguchi shows ([Sekiguchi], Lemma 1.5) that this property (together with the specified $G_{\mathbf{R}}$-conjugacy class of $\phi_{i\mathbf{R}}$) determines $\phi_{i\mathbf{R},\mathfrak{k}}$ up to $K_{\mathbf{R}}$-conjugacy. This gives the correspondence from (b) to (c).

Next, we show how to go from (e) to (d). We begin with a nilpotent element $x_{\mathfrak{k}} \in \mathcal{N}_{\mathfrak{k}}$. As is shown in [Kostant-Rallis], we can find elements $y_{\mathfrak{k}} \in \mathcal{N}_{\mathfrak{k}}$ and $h_{\mathfrak{k}} \in \mathfrak{k}$ so that

$$[h_{\mathfrak{k}}, x_{\mathfrak{k}}] = 2x_{\mathfrak{k}}, \quad [h_{\mathfrak{k}}, y_{\mathfrak{k}}] = -2y_{\mathfrak{k}}, \quad [x_{\mathfrak{k}}, y_{\mathfrak{k}}] = h_{\mathfrak{k}}. \tag{6.7)(a}$$

$$\theta x_{\mathfrak{k}} = -x_{\mathfrak{k}}, \quad \theta y_{\mathfrak{k}} = -y_{\mathfrak{k}}, \quad \theta h_{\mathfrak{k}} = h_{\mathfrak{k}}. \tag{6.7)(b}$$

Again these three elements determine a homomorphism $\phi_{\mathfrak{k}}$ from $\mathfrak{sl}(2)$ to $\mathfrak{g}$, by the requirements

$$\phi_{\mathfrak{k}}\begin{pmatrix} 0 & i \\ -i & 0 \end{pmatrix} = h_{\mathfrak{k}}, \quad \phi_{\mathfrak{k}}\begin{pmatrix} 1/2 & -i/2 \\ -i/2 & -1/2 \end{pmatrix} = x_{\mathfrak{k}}, \quad \phi_{\mathfrak{k}}\begin{pmatrix} 1/2 & i/2 \\ i/2 & -1/2 \end{pmatrix} = y_{\mathfrak{k}}. \tag{6.7)(c}$$

The homomorphism $\phi_{\mathfrak{k}}$ intertwines θ_0 with θ. This establishes the correspondence from (e) to (d).

Suppose now that we are given $\phi_{\mathfrak{k}}$ as in (d) of Theorem 6.4. We wish to modify $\phi_{\mathfrak{k}}$ by conjugation by an element of K so that it intertwines σ_0 with σ. Of course it is equivalent to have τ_0 intertwined with τ. Now τ_0 is mapped by $\phi_{\mathfrak{k}}$ to a Cartan involution τ_0' of the image. Now we use a slight refinement of the result about extending Cartan involutions from subalgebras: that if the subalgebra is preserved by a fixed involutive automorphism θ, and the given Cartan involution on the subalgebra commutes with θ, then the extension may be chosen to commute with θ as well ([van Dijk], Proposition 2). We conclude that τ_0' may be extended to a Cartan involution τ' of $\mathfrak{g}$ commuting with θ. Now τ' and τ are two Cartan involutions commuting with θ. Consequently they differ by conjugation by some element k of (the identity component of) $K_{\mathbf{C}}$ ([Loos], Chapter IV, Theorem 2.1). Write

$$\phi_{i\mathbf{R},\mathfrak{k}} = \mathrm{Ad}(k) \circ \phi_{\mathfrak{k}}. \tag{6.8}$$

Then $\phi_{i\mathbf{R},\mathfrak{k}}$ intertwines θ_0 with θ and τ_0 with τ; so it also intertwines σ_0 with σ. This is the correspondence from (d) to (c); again we refer to [Sekiguchi] for the proof that it is well-defined. Q.E.D.

We turn now to the rest of the proof of Theorem 6.2. That $G \cdot \lambda_{\mathfrak{k}}$ is equal to $G \cdot \lambda_{i\mathbf{R}}$ is clear from the construction of the bijection in Theorem 6.4. We have already seen (Corollary 5.20) that $K \cdot \lambda_{\mathfrak{k}}$ may be regarded as a complex Lagrangian subvariety of $G \cdot \lambda_{\mathfrak{k}}$; this proves the second equality of dimensions in (b). Suppose now that we regard $G \cdot \lambda_{i\mathbf{R}}$ as a *real* symplectic manifold, by considering only the real part of $\omega_{\lambda_{i\mathbf{R}}}$ (cf. (5.14)). Then the assumption that $\lambda_{i\mathbf{R}}$ is imaginary-valued forces the real submanifold $G_{\mathbf{R}} \cdot \lambda_{i\mathbf{R}}$ to be isotropic; and the proof of Proposition 5.19 (with σ replacing θ) shows that it is actually Lagrangian. This gives the first equality of dimensions in (b).

For Theorem 6.2(c), we need some additional notation based on (6.5). Write

$$G_{\mathbf{R}}(x_{i\mathbf{R}}) = \text{ centralizer in } G_{\mathbf{R}} \text{ of } x_{i\mathbf{R}}, \tag{6.9)(a}$$

$$G_{\mathbf{R}}(\phi_{i\mathbf{R}}) = \text{ centralizer in } G_{\mathbf{R}} \text{ of } \phi_{i\mathbf{R}}(\mathfrak{sl}(2)). \tag{6.9)(b}$$

The element $h_{i\mathbf{R}}$ has integral eigenvalues in its adjoint action on $\mathfrak{g}$, and preserves $\mathfrak{g}_{\mathbf{R}}(x_{i\mathbf{R}})$. Whenever $\mathfrak{s}$ is an $\mathrm{ad}(h_{i\mathbf{R}})$-stable subspace of $\mathfrak{g}$, we write

$$\mathfrak{s}[k;h_{i\mathbf{R}}] = \mathfrak{s}[k] = \{\, s \in \mathfrak{s} \mid [h_{i\mathbf{R}}, s] = ks \,\}. \qquad (6.9)(c)$$

By the representation theory of $\mathfrak{sl}(2)$), the eigenvalues of $\mathrm{ad}(h_{i\mathbf{R}})$ on $\mathfrak{g}_{\mathbf{R}}(x_{i\mathbf{R}})$ are non-negative, and $\mathfrak{g}_{\mathbf{R}}(\phi_{i\mathbf{R}})$ is precisely the zero eigenspace:

$$\mathfrak{g}_{\mathbf{R}}(\phi_{i\mathbf{R}}) = \mathfrak{g}_{\mathbf{R}}(x_{i\mathbf{R}})[0]. \qquad (6.9)(d)$$

Define

$$\mathfrak{u}_{\mathbf{R}}(x_{i\mathbf{R}}) = \sum_{k>0} \mathfrak{g}_{\mathbf{R}}(x_{i\mathbf{R}})[k], \qquad (6.9)(e)$$

the sum of the positive eigenspaces of $\mathrm{ad}(h_{i\mathbf{R}})$ on $\mathfrak{g}_{\mathbf{R}}(x_{i\mathbf{R}})$. Then $\mathfrak{g}_{\mathbf{R}}(x_{i\mathbf{R}})$ is evidently the semidirect product of the reductive subalgebra $\mathfrak{g}_{\mathbf{R}}(\phi_{i\mathbf{R}})$ by the nilpotent ideal $\mathfrak{u}_{\mathbf{R}}(x_{i\mathbf{R}})$. We want to get this decomposition on the group level. By putting $h_{i\mathbf{R}}$ in a Cartan subalgebra of $\mathfrak{g}_{\mathbf{R}}$, we see that $\mathfrak{u}_{\mathbf{R}}(x_{i\mathbf{R}})$ is contained in a maximal nilpotent subalgebra of $\mathfrak{g}_{\mathbf{R}}$. It follows that the corresponding connected subgroup

$$U_{\mathbf{R}}(x_{i\mathbf{R}}) = exp(\mathfrak{u}_{\mathbf{R}}(x_{i\mathbf{R}})) \qquad (6.9)(f)$$

is a simply connected unipotent Lie group.

Lemma 6.10(cf. [Barbasch-Vogan], section 2, and [Kostant]). *In the setting of (6.9), the centralizer $G_{\mathbf{R}}(x_{i\mathbf{R}})$ is the semidirect product of the reductive group $G_{\mathbf{R}}(\phi_{i\mathbf{R}})$ by the unipotent normal subgroup $U_{\mathbf{R}}(x_{i\mathbf{R}})$. In particular, suppose that the image of $\phi_{i\mathbf{R}}$ is preserved by θ. Then the restriction of θ to $G_{\mathbf{R}}(\phi_{i\mathbf{R}})$ is a Cartan involution, so*

$$K_{\mathbf{R}} \cap G_{\mathbf{R}}(\phi_{i\mathbf{R}})$$

is a maximal compact subgroup of $G_{\mathbf{R}}(\phi_{i\mathbf{R}})$ or of $G_{\mathbf{R}}(x_{i\mathbf{R}})$.

The main point in the proof is to show that, given $x_{i\mathbf{R}}$, any two choices of $y_{i\mathbf{R}}$ satisfying (6.5)(a)-(b) must be conjugate by a (unique) element of $U_{\mathbf{R}}(x_{i\mathbf{R}})$. We omit the argument.

We can define exactly parallel notation for K based on (6.7). In that case we are working with algebraic groups, and the analogue of Lemma 6.10 is precisely a Levi decomposition.

Lemma 6.11. *In the setting of (6.7) (and with notation analogous to (6.9)) the centralizer $K(x_{\mathfrak{k}})$ is the semidirect product of the reductive group $K(\phi_{\mathfrak{k}})$ by the unipotent normal subgroup $U(x_{\mathfrak{k}})$. In particular, suppose that the image of $\phi_{\mathfrak{k}}$ is preserved by the involution τ (defined before (6.3)). Then the restriction of τ to $K(\phi_{\mathfrak{k}})$ is a Cartan involution, so*

$$K_{\mathbf{R}} \cap K(\phi_{\mathfrak{k}})$$

is a maximal compact subgroup of $K(\phi_{\mathfrak{k}})$ or of $K(x_{\mathfrak{k}})$.

Corollary 6.12. *In the setting of Theorem 6.4, suppose ϕ is a homomorphism from $\mathfrak{sl}(2)$ to $\mathfrak{g}$, intertwining σ_0 with σ and θ_0 with θ. Define*

$$x_{i\mathbf{R}} = \phi\begin{pmatrix} 0 & i \\ 0 & 0 \end{pmatrix}, \qquad x_{\mathfrak{k}} = \phi\begin{pmatrix} 1/2 & -i/2 \\ -i/2 & -1/2 \end{pmatrix}.$$

Then the centralizer in $K_{\mathbf{R}}$ of the image of ϕ is a maximal compact subgroup of $G_{\mathbf{R}}(x_{i\mathbf{R}})$ and of $K(x_{\mathfrak{k}})$.

In light of Theorem 6.4, this establishes Theorem 6.2(c).

7. Admissible orbits

In this section we recall Duflo's notion of "admissible" (imaginary) coadjoint orbits. We then present a result of J. Schwartz (Theorem 7.14) describing which nilpotent $K_{\mathbf{C}}$ orbits on $\mathcal{N}_{\mathfrak{k}}^*$ correspond (via Theorem 6.2) to admissible imaginary orbits.

Suppose to begin that $G_{\mathbf{R}}$ is any real Lie group, and

$$\lambda_{i\mathbf{R}} \in \mathfrak{g}_{i\mathbf{R}}^* \tag{7.1)(a}$$

is a purely imaginary-valued linear functional. Write $\mathcal{O}_{i\mathbf{R}}$ for the $G_{\mathbf{R}}$-orbit of $\lambda_{i\mathbf{R}}$. As in (5.14) we can define an (imaginary-valued) symplectic form $\omega_{\lambda_{i\mathbf{R}}}$ on the tangent space

$$T_{\lambda_{i\mathbf{R}}}(\mathcal{O}_{i\mathbf{R}}) \simeq \mathfrak{g}_{\mathbf{R}}/\mathfrak{g}_{\mathbf{R}}(\lambda_{i\mathbf{R}}). \tag{7.1)(b}$$

Write $Sp(\omega_{\lambda_{i\mathbf{R}}})$ for the group of (real linear) symplectic linear transformations of this tangent space. Then the isotropy action gives a natural homomorphism

$$G_{\mathbf{R}}(\lambda_{i\mathbf{R}}) \xrightarrow{j} Sp(\omega_{\lambda_{i\mathbf{R}}}). \tag{7.1)(c}$$

On the other hand, the real symplectic group has a natural two-fold covering group, the *metaplectic group*:

$$1 \longrightarrow \{1, \epsilon\} \longrightarrow Mp(\omega_{\lambda_{i\mathbf{R}}}) \xrightarrow{p} Sp(\omega_{\lambda_{i\mathbf{R}}}) \longrightarrow 1. \tag{7.1)(d}$$

This covering may be pulled back via the homomorphism (7.1)(c) to give the *metaplectic double cover* of the isotropy group:

$$1 \longrightarrow \{1, \epsilon\} \longrightarrow \tilde{G}_{\mathbf{R}}(\lambda_{i\mathbf{R}}) \xrightarrow{p(\lambda_{i\mathbf{R}})} G_{\mathbf{R}}(\lambda_{i\mathbf{R}}). \tag{7.1)(e}$$

Explicitly, this covering group is defined by

$$\tilde{G}_{\mathbf{R}}(\lambda_{i\mathbf{R}}) = \{\, (g, m) \in G_{\mathbf{R}}(\lambda_{i\mathbf{R}}) \times Mp(\omega_{\lambda_{i\mathbf{R}}}) \mid j(g) = p(m) \,\}. \tag{7.1)(f}$$

Definition 7.2 ([Duflo]). In the setting of (7.1), a representation χ of $\tilde{G}_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ is called *genuine* if $\chi(\epsilon) = -I$. (In [Vogan], chapter 10, such representations are called *metaplectic.* Notice that if χ is irreducible, $\chi(\epsilon)$ is necessarily $+I$ or $-I$.) We say that χ is *admissible* if it is genuine, and the differential of χ is a multiple of $\lambda_{i\mathbf{R}}$; that is, if

$$\chi(\exp x) = \exp(\lambda_{i\mathbf{R}}(x)) \cdot I$$

for all $x \in \mathfrak{g}_{\mathbf{R}}(\lambda_{i\mathbf{R}})$. (Here the exponential map on the left is the one for $G_{\mathbf{R}}$, and the one on the right is for complex numbers.) Notice that if $G_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ has a finite number of connected components, an irreducible admissible representation of $\tilde{G}_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ is necessarily unitarizable.

If admissible representations exist, we say that $\lambda_{i\mathbf{R}}$ (or the orbit $\mathcal{O}_{i\mathbf{R}}$) is *admissible.* A pair $(\lambda_{i\mathbf{R}}, \chi)$ consisting of an element of $\mathfrak{g}^*_{i\mathbf{R}}$ and an irreducible admissible representation of $\tilde{G}_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ is called an *admissible* $G_{\mathbf{R}}$*-orbit datum.* Two such are called equivalent if they are conjugate by $G_{\mathbf{R}}$.

Admissible orbit data are the raw material of the orbit method. Here is a rough version of what one expects. (The words "nice" and "usually" below reflect my ignorance; I do not know a more precise statement that is correct.)

Desideratum 7.3. Suppose $G_{\mathbf{R}}$ is a nice type I Lie group, and $\mathcal{D} = (\lambda_{i\mathbf{R}}, \chi)$ is an admissible orbit datum for $G_{\mathbf{R}}$ (Definition 7.2). Attached to the equivalence class of $\mathcal{D}$ there should be a unitary representation $\pi(\mathcal{D})$ of $G_{\mathbf{R}}$. This representation should be a direct sum of a finite number (possibly zero) of irreducibles; and usually $\pi(\mathcal{D})$ itself should be irreducible.

For more information about what can be proved in this direction, the reader may consult for example [Duflo]; the case of reductive groups is discussed in [Vogan], chapter 10. Of course our primary concern here is with the case of nilpotent orbit data for reductive groups. The condition of admissibility is very simple in this case.

Observation 7.4. In the setting of Definition 7.2, assume that $\lambda_{i\mathbf{R}}$ is nilpotent (cf. (5.5)). Then a representation χ of $\tilde{G}_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ is admissible if and only if it is trivial on the identity component, and $\chi(\epsilon) = -I$. (Here ϵ is the non-trivial element of the kernel of the covering map $p(\lambda_{i\mathbf{R}})$.) In particular, $\lambda_{i\mathbf{R}}$ is admissible if and only if the preimage (under $p(\lambda_{i\mathbf{R}})$ of the identity component $G_{\mathbf{R}}(\lambda_{i\mathbf{R}})_0$ is disconnected.

Example 7.5. Suppose $G_{\mathbf{R}}$ is the group $SO(3)$. The Lie algebra $\mathfrak{g}_{\mathbf{R}}$ may be identified with skew-symmetric three by three real matrices. Fix a non-zero real number t, and let $\lambda_{i\mathbf{R}}(t)$ denote the linear functional whose value at a matrix x is it times the $(1,2)$ entry of x. Then the isotropy group $G_{\mathbf{R}}(\lambda_{i\mathbf{R}}(t))$ is $SO(2)$, embedded in $G_{\mathbf{R}}$ as the upper left two by two block. The characters of $SO(2)$ are parametrized by $\mathbf{Z}$; we can arrange the parametrization so that the nth character χ_n has differential $\lambda_{i\mathbf{R}}(n)$ (restricted to $\mathfrak{so}(2)$). The complexified isotropy action on $\mathfrak{g}/\mathfrak{g}(\lambda_{i\mathbf{R}})$ has the two weights corresponding to $+1$ and -1. The symplectic group for a two-dimensional vector space is just $SL(2, \mathbf{R})$. It follows that the isotropy representation maps $SO(2)$ isomorphically onto a maximal compact subgroup of $SL(2, \mathbf{R})$. Consequently the metaplectic double cover of $SO(2)$ is the connected double cover; its genuine representations are again one-dimensional characters χ_j parametrized by $j \in \mathbf{Z} + 1/2$. We have therefore found admissible orbit data $\mathcal{D}(j) = (\lambda_{i\mathbf{R}}(j), \chi_j)$ for each $j \in \mathbf{Z} + 1/2$. It turns out that the $\mathcal{D}(j)$ and $\mathcal{D}(-j)$ are conjugate, and that every admissible orbit datum except $\mathcal{T} = (0, \mathbf{C})$ is conjugate to some $\mathcal{D}(j)$. (Here $\mathcal{T}$ stands for "trivial"; this orbit datum exists for every $G_{\mathbf{R}}$, and $\pi(\mathcal{T})$ is the trivial representation.) There are various ways to attach representations to the other orbit data; the most natural all make $\pi(\mathcal{D}(j))$ the irreducible representation of $SO(3)$ of dimension $2j$. Notice in particular that the trivial representation is attached both to $\mathcal{T}$ and to $\mathcal{D}(1/2)$.

In this example, only the orbit datum $\mathcal{T}$ is nilpotent. Before we consider some interesting examples of nilpotent orbit data, it will be helpful to have a little more machinery.

Lemma 7.6. *Suppose G is a real Lie group having a finite number of connected components, K is a maximal compact subgroup of G, and ϵ is a central element of G of finite order m. Fix an mth root of unity ζ. Say that a representation of χ of G is admissible if it is trivial on the identity component of G, and $\chi(\epsilon) = \zeta I$; and define admissible representations of K analogously. Then restriction to K induces a bijection from admissible representations of G to admissible representations of K.*

This is an immediate consequence of Mostow's result that G is topologically the product of K with a vector space. Because of this result (and Observation 7.4), the question of admissibility for nilpotent orbit data can be studied on the level of maximal compact subgroups. Suppose V is a real vector space carrying a non-degenerate imaginary-valued symplectic form ω. We can choose a complex structure and positive-definite Hermitian form h on V so that ω is the imaginary part of h. Write $U(h)$ for the unitary group of h; this is a maximal compact subgroup of $Sp(\omega)$. The complex determinant defines a one-dimensional character

$$\det : U(h) \to \mathbf{C}^\times. \tag{7.7)(a}$$

Using this homomorphism, we can pull back the connected double cover of $\mathbf{C}^\times$ to a double cover of $U(h)$:

$$\tilde{U}(h) = \{\, (g, z) \in U(h) \times \mathbf{C}^\times \mid \det(g) = z^2 \,\} \tag{7.7)(b}$$

This is called the *square root of the determinant covering*, because projection on the second factor defines a homomorphism

$$\det{}^{1/2} : \tilde{U}(h) \to \mathbf{C}^\times. \tag{7.7)(c}$$

whose square is precisely det. It turns out that this (delightfully simple) covering is precisely the one induced by the (delightfully complicated) metaplectic covering $Mp(\omega)$. The next lemma shows how to compute with such coverings.

Lemma 7.8. *Suppose G is a real Lie group, and γ is a one-dimensional character of G. Let $\tilde{G}$ denote the square root of γ covering of G (cf. (7.7)), and ϵ the non-trivial element of the kernel of the covering map. Define an admissible representation χ of $\tilde{G}$ to be one trivial on the identity component, satisfying $\chi(\epsilon) = -1$. Define a γ-admissible representation χ_0 of G to be one whose differential is half the differential of γ:*

$$\chi_0)(\exp(x)) = \gamma(\exp(x/2)) \cdot I \qquad x \in \mathfrak{g})$$

Then there is a natural bijection between admissible representations of $\tilde{G}$ and γ-admissible representations of G.

Proof. We use the one-dimensional character $\gamma^{1/2}$ of $\tilde{G}$ (cf. (7.7)(c)). Tensoring with $\gamma^{1/2}$ sends admissible representations of $\tilde{G}$ to (the pullbacks to $\tilde{G}$ of) γ-admissible representations of G. Q.E.D.

Corollary 7.9 ([Schwartz]). *Suppose $G_{\mathbf{R}}$ is a real reductive Lie group, and $\lambda_{i\mathbf{R}} \in \mathcal{N}^*_{i\mathbf{R}}$ is a purely imaginary nilpotent linear functional. Fix a maximal compact subgroup H of the isotropy group $G_{\mathbf{R}}(\lambda_{i\mathbf{R}})$. Choose an H-invariant complex structure and hermitian form h on the symplectic vector space $\mathfrak{g}_{\mathbf{R}}/\mathfrak{g}_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ as in (7.7); this is possible since the compact subgroup $p(\lambda_{i\mathbf{R}})(H)$ of the symplectic group must be contained in some maximal compact subgroup. Now define γ to be the corresponding complex determinant character of H. Then the admissible representations of $\tilde{G}_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ (Definition 7.2) are in natural one-to-one correspondence with the γ-admissible representations of H (Lemma 7.8). In particular, $\lambda_{i\mathbf{R}}$ is admissible if and only if the restriction of γ to the identity component H_0 of H is the square of another character of H_0.*

This follows at once from Lemmas 7.6 and 7.8, and Observation 7.4. It is still short of complete information about admissibility, for the character γ is difficult to compute explicitly from this description. Nevertheless we can treat some illustrative examples.

Example 7.10 (see [Schwartz]). Suppose $G_{\mathbf{R}}$ is the symplectic group $Sp(2n, \mathbf{R})$ for the standard symplectic form ω on $\mathbf{R}^{2n} = \mathbf{R}^n \times \mathbf{R}^n$: in terms of the usual dot product, this is

$$\omega((x, y), (x', y')) = x \cdot y' - y \cdot x'.$$

Then $G_{\mathbf{R}}$ consists of $2n \times 2n$ matrices $\begin{pmatrix} A & B \\ C & D \end{pmatrix}$ such that A^tB and D^tC are symmetric, and $A^tD - B^tC = I$. Its Lie algebra consists of matrices $\begin{pmatrix} X & Y \\ Z & -{}^tX \end{pmatrix}$ with Y and Z symmetric.

Fix non-negative integers p, q, and r such that $p + q + r = n$. We are going to define a nilpotent linear functional $\lambda_{i\mathbf{R}}(p, q, r)$. With obvious notation for matrix entries, it is

$$\lambda_{i\mathbf{R}}(p, q, r) \begin{pmatrix} X & Y \\ Z & -{}^tX \end{pmatrix} = i \sum_{k=1}^{p} y_{kk} - i \sum_{k=p+1}^{p+q} y_{kk}.$$

Using the explicit description of the group given above, it is not too hard to compute the isotropy group $G_{\mathbf{R}}(\lambda_{i\mathbf{R}}(p, q, r))$ explicitly. The answer is most conveniently expressed in terms of the semidirect product decomposition (Lemma 6.10). The reductive factor is $O(p, q) \times Sp(2r, \mathbf{R})$ (embedded in $Sp(2n, \mathbf{R})$ in an "obvious" way which we encourage the reader to untangle). The unipotent normal subgroup is two-step nilpotent; its Lie algebra $\mathfrak{u}_{\mathbf{R}}(\lambda_{i\mathbf{R}}(p, q, r))$ consists of matrices of the form $\begin{pmatrix} A & 0 \\ C & -{}^tA \end{pmatrix}$, subject to the conditions

$$a_{kl} = 0 \text{ unless } k > p + q, l \leq p + q, C = {}^tC, \text{ and } c_{kl} = 0 \text{ if } k, l > p + q.$$

Consequently the maximal compact subgroup H of the isotropy group is $O(p)\times O(q)\times U(r)$.

To compute the character γ, we need to understand something about the symplectic vector space

$$V(p,q,r) = \mathfrak{g}_{\mathbf{R}}/\mathfrak{g}_{\mathbf{R}}(\lambda_{i\mathbf{R}}(p,q,r))$$

The dimension of $V(p,q,r)$ is easily computed from the description of the isotropy subgroup; it is $2r(p+q)+(p+q)(p+q+1)$. This formula suggests a description of $V(p,q,r)$ as a symplectic representation of $O(p,q)\times Sp(2r,\mathbf{R})$, which is not too difficult to verify:

$$V(p,q,r) \simeq \left(\mathbf{R}^{2r}\otimes \mathbf{R}^{p+q}\right)\oplus\left(S^2(\mathbf{R}^{p+q})\oplus S^2(\mathbf{R}^{p+q})^*\right).$$

Here the first summand carries the tensor product of the symplectic form on the first factor with the orthogonal form on the second; this is a symplectic form preserved by the product of the groups of the small forms. The second summand is the sum of a group representation and its dual; it therefore carries a natural group-invariant symplectic form. (The $Sp(2r,\mathbf{R})$ factor acts trivially on the second summand.)

We can now express the character γ of H (Corollary 7.9) in terms of the standard determinant characters of the factors $U(r)$, $O(p)$, and $O(q)$:

$$\gamma = (\det_{U(r)})^{p+q}\otimes(\det_{O(p)})^{n-1}\otimes(\det_{O(q)})^{n-1}.$$

On the identity component only the first factor is non-trivial. We conclude that $\lambda_{i\mathbf{R}}(p,q,r)$ is admissible if and only if either $p+q$ is even or $r=0$. In this case (Corollary 7.9 again) the number of irreducible admissible representations – that is, the number of inequivalent orbit data – is 4 if p and q are both non-zero, 2 if exactly one is non-zero, and 1 if both are zero. (This last case corresponds to the trivial orbit datum.)

It turns out that all the linear functionals $\lambda_{i\mathbf{R}}(p,q,r)$ are admissible for the metaplectic covering $Mp(2n,\mathbf{R})$. Schwartz gives many examples of nilpotent orbits for $Sp(2n,\mathbf{R})$ that are inadmissible for all coverings, however.

The next theorem relates the character γ of Corollary 7.9 to the corresponding K-orbit on $\mathcal{N}^*_{\mathfrak{k}}$ (Theorem 6.2).

Theorem 7.11 ([Schwartz]). *In the setting of Theorem 6.2, suppose that the element $\lambda_{\mathfrak{k}}\in\mathcal{N}^*_{\mathfrak{k}}$ corresponds to $\lambda_{i\mathbf{R}}\in\mathcal{N}^*_{i\mathbf{R}}$. Write $\gamma_{i\mathbf{R}}$ for the character of the maximal compact subgroup of $G_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ (Corollary 7.9). Define a character $\gamma_{\mathfrak{k}}$ of $K(\lambda_{\mathfrak{k}})$ by*

$$\gamma_{\mathfrak{k}}(k) = \det(\mathrm{Ad}(k))\,|_{(\mathfrak{k}/\mathfrak{k}(\lambda))^*}.$$

Then the restriction of $\gamma_{\mathfrak{k}}$ to the maximal compact subgroup is identified (by the isomorphism of Theorem 6.2) with $\gamma_{i\mathbf{R}}$.

The character $\gamma_{\mathfrak{k}}$ gives the action of $K(\lambda_{\mathfrak{k}})$ on top degree differential forms on $K\cdot\lambda_{\mathfrak{k}}$, at the point $\lambda_{\mathfrak{k}}$. Because of the symplectic structure, this is dual to the action on the top exterior power of the conormal bundle $T^*_{K\cdot\lambda_{\mathfrak{k}}}(G\cdot\lambda_{\mathfrak{k}})$ of $K\cdot\lambda_{\mathfrak{k}}$ in the (complex symplectic) variety $G\cdot\lambda_{\mathfrak{k}}$, at the point $\lambda_{\mathfrak{k}}$. By the proof of Proposition 5.19, this can be phrased as

$$\gamma_{\mathfrak{k}}(k) = \det(\mathrm{Ad}(k))\,|_{(\mathfrak{p}/\mathfrak{p}(\lambda))}. \tag{7.12}$$

Before giving the proof of Theorem 7.11, we record a consequence. It is convenient to make a definition parallel to Definition 7.2.

Definition 7.13. Suppose $(\mathfrak{g}, K)$ is a reductive symmetric pair (Definition 5.21), and $\lambda_{\mathfrak{k}} \in \mathcal{N}_{\mathfrak{k}}^*$. Define a character $\gamma_{\mathfrak{k}}$ of $K(\lambda_{\mathfrak{k}})$ as in Theorem 7.11 and (7.12). An algebraic representation χ of $K(\lambda_{\mathfrak{k}})$ is called *admissible* if its differential is half the differential of $\gamma_{\mathfrak{k}}$; that is, if

$$\chi(\exp(x)) = \gamma_{\mathfrak{k}}(\exp(x/2)) \cdot I$$

for all $x \in \mathfrak{k}(\lambda_{\mathfrak{k}})$. If admissible representations exist, we say that $\lambda_{\mathfrak{k}}$ (or the orbit $K \cdot \lambda_{\mathfrak{k}}$) is admissible. A pair $(\lambda_{\mathfrak{k}}, \chi)$ consisting of an element of $\mathcal{N}_{\mathfrak{k}}^*$ and an irreducible admissible representation of $K(\lambda_{\mathfrak{k}})$ is called a *nilpotent admissible K-orbit datum.* Two such are equivalent if they are conjugate by K.

Theorem 7.14 ([Schwartz]). *Suppose $G_{\mathbf{R}}$ is a real reductive Lie group, $K_{\mathbf{R}}$ is a maximal compact subgroup, and K is its complexification. Then there is a natural bijection between equivalence classes of nilpotent admissible $G_{\mathbf{R}}$-orbit data (Definition 7.2) and equivalence classes of nilpotent admissible K-orbit data (Definition 7.13).*

Proof. Fix a nilpotent linear functional $\lambda_{i\mathbf{R}} \in \mathcal{N}_{i\mathbf{R}}^*$, and let $\lambda_{\mathfrak{k}} \in \mathcal{N}_{\mathfrak{k}}^*$ be a corresponding element (Theorem 6.2). We want to associate to each admissible representation $\chi_{i\mathbf{R}}$ of $\tilde{G}_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ an admissible representation $K(\lambda_{\mathfrak{k}})$. Let H be a maximal compact subgroup of $G_{\mathbf{R}}(\lambda_{i\mathbf{R}})$. By Theorem 6.2, we may as well assume that H is also a maximal compact subgroup of $K(\lambda)$. Let γ be the character of H constructed in Corollary 7.9 (or, by Theorem 7.11, the restriction of $\gamma_{\mathfrak{k}}$ to H). By Lemma 7.8, restriction to the preimage of H and twisting by the "square root of γ" defines a bijection from admissible representations of $\tilde{G}_{\mathbf{R}}(\lambda_{i\mathbf{R}})$ to γ-admissible representations of H. An even simpler fact about algebraic groups (essentially Weyl's "unitarian trick") guarantees that restriction to H is a bijection from admissible representations of $K(\lambda_{\mathfrak{k}})$ to γ-admissible representations of H. (An admissible representation of $K(\lambda_{\mathfrak{k}})$ is automatically trivial on the unipotent radical, so this is really just a statement about reductive groups.) The theorem follows. Q.E.D.

Proof of Theorem 7.11. Fix a nondegenerate symmetric real bilinear form b on $\mathfrak{g}_{\mathbf{R}}$, invariant under $G_{\mathbf{R}}$ and θ, negative definite on $\mathfrak{k}_{\mathbf{R}}$, and positive definite on $\mathfrak{p}_{\mathbf{R}}$, as in the proof of Theorem 6.2. Then the bilinear form

$$b_\theta(u, v) = b(\theta u, v) = b(u, \theta v) \tag{7.15)(a}$$

is negative definite on $\mathfrak{g}_{\mathbf{R}}$. Define elements $x_{i\mathbf{R}}$ and $x_{\mathfrak{k}}$ of $\mathfrak{g}$ by the requirements

$$\lambda_{i\mathbf{R}}(y) = b(x_{i\mathbf{R}}, y), \qquad \lambda_{\mathfrak{k}}(y) = b(x_{\mathfrak{k}}, y) \tag{7.15)(b}$$

for all $y \in \mathfrak{g}$. By Theorem 6.4, we may as well assume that there is a homomorphism ϕ from $\mathfrak{sl}(2)$ to $\mathfrak{g}$, intertwining σ_0 with σ and θ_0 with θ, and satisfying

$$\phi\begin{pmatrix} 0 & i \\ 0 & 0 \end{pmatrix} = x_{i\mathbf{R}}, \qquad \phi\begin{pmatrix} 1/2 & -i/2 \\ -i/2 & -1/2 \end{pmatrix} = x_{\mathfrak{k}}. \tag{7.15)(c}$$

In terms of ϕ, we can give two descriptions of the compact group H:

$$H = \{ g \in G_{\mathbf{R}} \mid \mathrm{Ad}(g)(x_{i\mathbf{R}}) = x_{i\mathbf{R}}, \theta g = g \}; \tag{7.15)(d}$$

$$H = \{ k \in K \mid \mathrm{Ad}(k)(x_{\mathfrak{t}}) = x_{\mathfrak{t}}, \sigma k = k \}. \tag{7.15)(e}$$

The symplectic structure $\omega = \omega_{x_{i\mathbf{R}}}$ on $V = \mathfrak{g}_{\mathbf{R}}/\mathfrak{g}_{\mathbf{R}}(x_{i\mathbf{R}})$ is defined by

$$\omega(u,v) = b(x_{i\mathbf{R}}, [u,v]). \tag{7.16)(a}$$

In order to calculate the character $\gamma_{i\mathbf{R}}$, we must (according to Corollary 7.9) construct a complex structure J on V; that is, a linear transformation satisfying $J^2 = -I$. In addition, J must commute with the action of H, and satisfy

$$\omega(Ju,v) = \omega(Jv,u) \tag{7.16)(b}$$

$$(1/i)\omega(Ju,u) \geq 0. \tag{7.16)(c}$$

This makes V into a complex vector space, and H acts by complex-linear transformations. (Although we will not use it, the hermitian form h of (7.7) is

$$h(u,v) = (1/i)\omega(Ju,v) + \omega(u,v). \tag{7.16)(d}$$

The action of H is unitary for this hermitian form.) The character $\gamma_{i\mathbf{R}}$ is the determinant of the action of H on this complex vector space. Equivalently, it is the determinant of the action of H on the $+i$-eigenspace of J in $V_{\mathbf{C}} = \mathfrak{g}/\mathfrak{g}(x_{i\mathbf{R}})$.

We begin now the construction of J. Define

$$E = \phi\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \qquad F = \phi\begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} = -\theta E. \tag{7.17)(a}$$

These are elements of $\mathfrak{g}_{\mathbf{R}}$, and $iE = x_{i\mathbf{R}}$. If T is a linear transformation of $\mathfrak{g}_{\mathbf{R}}$, write T^θ for its adjoint with respect to the negative definite form b_θ of (7.15)(a):

$$b_\theta(Tu,v) = b_\theta(u,T^\theta v).$$

Obviously θ is self-adjoint, and

$$\mathrm{ad}(x)^\theta = -\mathrm{ad}(\theta x). \tag{7.17)(b}$$

Our first approximation to J is the linear transformation

$$Q = \theta \circ \mathrm{ad}\, E = -\mathrm{ad}\, F \circ \theta. \tag{7.17)(c}$$

Obviously the kernel of Q is precisely $\mathfrak{g}_{\mathbf{R}}(x_{i\mathbf{R}})$, and it follows from (7.17)(b) that Q is skew-adjoint with respect to b_θ: $Q^\theta = -Q$. A first consequence is that Q defines an invertible linear transformation $\overline{Q}$ on V. A second is that $R = -Q^2$ is a non-negative self-adjoint linear transformation on $\mathfrak{g}_{\mathbf{R}}$. By (7.17)(c),

$$R = \operatorname{ad} F \circ \operatorname{ad} E. \tag{7.17)(d}$$

Clearly Q commutes with the action of H. By (7.16),

$$\omega(Qu, v) = b(iE, [\theta[E, u], v]) = -ib(\theta[E, u], [E, v]) = -ib_\theta([E, u], [E, v]). \tag{7.18)(a}$$

From the last expression it follows that

$$\omega(Qu, v) = \omega(Qv, u), \qquad (1/i)\omega(Qu, u) \geq 0. \tag{7.18)(b}$$

These assertions correspond to the requirements (7.16)(b) and (c) for J. The only difficulty is that $\overline{Q}^2$ is not $-I$, but only the negative operator $-\overline{R}$. We correct this using a square root in the usual way. Let S denote the non-negative self-adjoint square root of R. Then S defines an invertible linear transformation $\overline{S}$ on V, and we set

$$J = (\overline{S})^{-1}\overline{Q}. \tag{7.18)(c}$$

The correction factor commutes with all operators commuting with $\overline{Q}$, including $\overline{Q}$ itself and the adjoint action of H; so $J^2 = -I$, and J commutes with H. Now (7.18)(b) implies that $\omega(-Q^2u, v) = \omega(u, -Q^2v)$. Considering this equation on each eigenspace of $-Q^2$, we deduce that

$$\omega(Su, v) = \omega(u, Sv).$$

Now a trivial formal argument gives (7.16)(b) and (c) from (7.18)(b).

In light of the remarks after (7.16)(d), and the construction just given for J, we find that the character $\gamma_{i\mathbf{R}}$ of H may be described as the determinant of the adjoint action on the span of the positive eigenspaces of

$$(1/i)Q = -\theta \circ ad(x_{i\mathbf{R}}); \tag{7.19}$$

this operator may be taken to act on all of $\mathfrak{g}$, or just on the quotient $V_{\mathbf{C}} = \mathfrak{g}/\mathfrak{g}(x_{i\mathbf{R}})$. (It is a hermitian operator with respect to the hermitian form obtained from b_θ; this is another way to understand the fact that it is diagonalizable, with real eigenvalues.) Our next task is to relate these positive eigenspaces to $\mathfrak{k}/\mathfrak{k}(x_{\mathfrak{k}})$.

We begin by constructing a Cayley transform. Even though the homomorphism ϕ may not exponentiate to SL(2), its composition with ad does. Every element g of SL(2, C) therefore gives rise to an automorphism of $\mathfrak{g}$, which we write as $\operatorname{Ad}(\phi(g))$ (even though $\phi(g)$ by itself is undefined). These automorphisms commute with the action of H. Set

$$c = (1/\sqrt{2})\begin{pmatrix} i & -1 \\ 1 & -i \end{pmatrix}. \tag{7.20)(a}$$

By calculation in SL(2), we find that

$$\operatorname{Ad}(\phi(c))(x_{i\mathbf{R}}) = x_{\mathfrak{k}}. \tag{7.20)(b}$$

Define

$$T = \mathrm{Ad}(\phi(c)) \circ (1/i)Q \circ \mathrm{Ad}(\phi(c))^{-1}. \tag{7.20)(c}$$

Then T may be regarded as a linear transformation on $\mathfrak{g}$ or on $\mathfrak{g}/\mathfrak{g}(x_{\mathfrak{k}})$. It is diagonalizable with real eigenvalues, and the character $\gamma_{i\mathbf{R}}$ of H is the determinant of the adjoint action on the span of the positive eigenspaces. To calculate T, we need to know how θ acts on $\mathrm{Ad}(\phi(c))$. This can be computed from the fact that ϕ intertwines θ_0 and θ. Using (7.20)(b) and (7.19), we find (after a little calculation in SL(2))

$$T = -\theta \circ \mathrm{Ad}(\phi\begin{pmatrix} 0 & -i \\ -i & 0 \end{pmatrix}) \circ ad(x_{\mathfrak{k}}) = \mathrm{Ad}(\phi\begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}) \circ ad(x_{\mathfrak{k}}) \circ \theta. \tag{7.20)(d}$$

Write $W = \mathfrak{g}/\mathfrak{g}(x_{\mathfrak{k}})$, W^+ for the sum of the positive eigenspaces of T on W, and W^- for the sum of the negative eigenspaces. We have

$$W = W^+ \oplus W^- = (W \cap \mathfrak{p}) \oplus (W \cap \mathfrak{k}). \tag{7.21)(a}$$

The theorem we are trying to prove says that the determinants of the actions of H on W^+ and $W \cap \mathfrak{k}$ agree. Obviously it suffices to prove that

$$W^+ \simeq W \cap \mathfrak{p} \tag{7.21)(b}$$

as representations of H. In order to do this, we need another decomposition of W. Define $\zeta = \mathrm{Ad}(\phi(-I))$, an involutive automorphism of $\mathfrak{g}$. Then ζ commutes with θ, T, and the image of ϕ. It therefore lifts to a linear transformation of order 2 on W (still called ζ), commuting with everything else. Write

$$W = W_e \oplus W_o \tag{7.21)(c}$$

for the decomposition into the +1 ("even") and -1 ("odd") eigenspaces of ζ. We use analogous notation for other spaces and operators; thus for example W_o^+ is the sum of the positive eigenspaces of T_o. The desired isomorphism (7.21)(b) would follow from two separate isomorphisms

$$W_e^+ \simeq W_e \cap \mathfrak{p}, \tag{7.21)(d}$$

and

$$W_o^+ \simeq W_o \cap \mathfrak{p}. \tag{7.21)(e}$$

These seem to require rather different treatments, and we will prove them separately.

First we consider (7.21)(d). It follows from (7.20)(d) that T and θ anticommute on W_e:

$$T_e\theta_e = -\theta_e T_e. \tag{7.22)(a}$$

Consequently θ_e interchanges W_e^+ and W_e^-, and T_e interchanges $W_e \cap \mathfrak{k}$ and $W_e \cap \mathfrak{p}$. It follows immediately that all four representations of H are isomorphic:

$$W_e^+ \simeq W_e^- \simeq W_e \cap \mathfrak{k} \simeq W_e \cap \mathfrak{p}. \tag{7.22)(b}$$

The first isomorphism is given by θ_e; the second by restricting to W_e^- the projection $(I+\theta_e)/2$; and the third by T_e. In particular, we get (7.21)(d).

For (7.21)(e), we need to decompose W further using the action $\mathrm{ad}\circ\phi$ of $\mathfrak{sl}(2)$. Fix a non-negative integer N (the highest weight), and write $S(N)$ for the irreducible representation of $\mathfrak{sl}(2)$ of dimension $N+1$ (realized say on homogeneous polynomials of degree N in two variables). Define

$$\mathfrak{g}(N) = \text{ sum of all copies of } S(N) \text{ in } \mathfrak{g}, \tag{7.23)(a}$$

the corresponding isotypic subspace of $\mathfrak{g}$. We have

$$\mathfrak{g}_o = \sum_{N\ odd} \mathfrak{g}(N) \tag{7.23)(b}$$

Recall from (6.7)(c) the element

$$h_{\mathfrak{k}} = \phi\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \in \mathfrak{k}.$$

The eigenvalues of $\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$ on $S(N)$ are $N, N-2, \ldots, -N+2, -N$, each occurring once. Write $S(N)_m$ for the m-eigenspace, and

$$\mathfrak{g}(N)_m = \ m\text{-eigenspace of } h_{\mathfrak{k}} \text{ in } \mathfrak{g}(N). \tag{7.23)(c}$$

We have

$$\mathrm{ad}(x_{\mathfrak{k}}) : \mathfrak{g}(N)_m \to \mathfrak{g}(N)_{m+2}\ ; \tag{7.23)(d}$$

this map is an isomorphism unless $m = N$, in which case it is zero. Because it also interchanges $\mathfrak{k}$ and $\mathfrak{p}$, and commutes with the action of H, we deduce that

$$\mathfrak{g}(N)_m \cap \mathfrak{k} \simeq \mathfrak{g}(N)_{m+2} \cap \mathfrak{p} \qquad (m \neq N), \tag{7.23)(e}$$

and similarly with $\mathfrak{k}$ and $\mathfrak{p}$ exchanged. These facts pass at once to W. Writing $W(N)_m$ for the image of $\mathfrak{g}(N)_m$ in W, we find that

$$W = \sum_{N \neq m} W(N)_m, \qquad W(N)_m \simeq \mathfrak{g}(N)_m \quad (N \neq m). \tag{7.23)(f}$$

All of these spaces are preserved by θ, and

$$T : W(N)_m \to W(N)_{-m-2}. \tag{7.23)(g}$$

In particular, $W(N)$ is T-invariant, so we may speak of $W(N)^+$ and $W(N)^-$. The isomorphism (7.21)(e) will follow if we can show that

$$W(N)^+ \simeq W(N) \cap \mathfrak{p} \tag{7.23)(h}$$

for every odd non-negative integer N.

So fix such an integer $N = 2j+1$. We first compute the right side of (7.23)(h). Define

$$A(N) = W(N)_{-1} \cap \mathfrak{k}, \tag{7.24)(a}$$

a representation of H. The bilinear form b induces a natural isomorphism

$$(\mathfrak{g}(N)_{-1} \cap \mathfrak{p})^* \simeq \mathfrak{g}(N)_1 \cap \mathfrak{p}. \tag{7.24)(b}$$

Because of (7.23)(e), it follows that

$$A(N)^* \simeq W(N)_{-1} \cap \mathfrak{p}, \tag{7.24)(c}$$

In the same way we may compute all the spaces $W(N)_m \cap \mathfrak{k}$ and $W(N)_m \cap \mathfrak{p}$ in terms of $A(N)$ and $A(N)^*$:

$$W(N)_m \cap \mathfrak{k} \simeq A(N), \qquad W(N)_m \cap \mathfrak{p} \simeq A(N)^* \qquad (m \equiv -1 \pmod 4) \tag{7.24)(d}$$

$$W(N)_m \cap \mathfrak{k} \simeq A(N)^*, \qquad W(N)_m \cap \mathfrak{p} \simeq A(N) \qquad (m \equiv 1 \pmod 4). \tag{7.24)(e}$$

In particular,

$$\begin{aligned} W(N) \cap \mathfrak{p} &\simeq (A(N) \oplus A(N)^*)^j \oplus A(N)^* \qquad (N \equiv 1 \pmod 4) \\ &\simeq (A(N) \oplus A(N)^*)^j \oplus A(N) \qquad (N \equiv 3 \pmod 4) \end{aligned} \tag{7.24)(f}$$

Next, we compute the left side of (7.23)(h). Fix $m > -1$. Because of (7.23)(g), the sum of the positive eigenspaces of T on $W(N)_m \oplus W(N)_{-m-2}$ is isomorphic (by projection on the first summand) to $W(N)_m$. By (7.24), this space is isomorphic to $A(N) \oplus A(N)^*$. Consequently

$$W(N)^+ \simeq (A(N) \oplus A(N)^*)^j \oplus W(N)^+_{-1}. \tag{7.25)(a}$$

Recall now the formula (7.20)(d) for T. Comparing it with (7.24)(f) and (7.25)(a), we find that the desired isomorphism (7.23)(h) follows from the fact that

$$(-1)^j \mathrm{Ad}(\phi \begin{pmatrix} 0 & -i \\ -i & 0 \end{pmatrix}) \circ ad(x_{\mathfrak{k}})$$

acts by a positive scalar on $W(N)_{-1}$. (This is an assertion about finite-dimensional representations of $\mathfrak{sl}(2)$, and is easily verified by direct calculation.) This proves (7.23)(h), and therefore (7.21)(e), and therefore (7.21)(b), and therefore the theorem. Q.E.D.

The proof has some useful consequences. Because W_e and W_o are perpendicular with respect to the symplectic form $\omega_{x_{\mathfrak{k}}}$ on W, each of them separately is a symplectic vector space. This structure is preserved by H, so the determinant of H acting on W_e is one. The isomorphisms (7.22)(b) imply that this determinant is the square of the determinant on (say) $W_e \cap \mathfrak{k}$. This last determinant is therefore ± 1:

$$\det(\mathrm{Ad}(x)) \text{ on } (W_e \cap \mathfrak{p}) = \pm 1. \tag{7.26)(a}$$

In particular, this determinant is trivial on the identity component of H. On the odd part of W, the determinant characters of H on $A(N)$ and $A(N)^*$ differ by the square of a character. It therefore follows from (7.24)(f) that there is a character γ_1 of H such that (as characters of H)

$$\det \circ \mathrm{Ad} \text{ on } (W_o \cap \mathfrak{p}) = [\det \circ \mathrm{Ad} \text{ on } (W_{-1} \cap \mathfrak{p})][\gamma_1]^2. \tag{7.26)(b}$$

In light of (7.26)(a), we can drop the subscript o on the identity component of H:

$$\gamma_{\mathfrak{k}} \,|_{H_0} = [\det \circ \mathrm{Ad} \text{ on } (W_{-1} \cap \mathfrak{p})][\gamma_1]^2. \tag{7.26)(c}$$

Corollary 7.27 ([Schwartz]). *In the setting of (6.7), write H for a maximal compact subgroup of the centralizer $K(x_{\mathfrak{k}})$. Then the nilpotent element $x_{\mathfrak{k}}$ is admissible (Definition 7.13) if and only if the determinant of the action of H_0 on the -1-eigenspace of $h_{\mathfrak{k}}$ on $\mathfrak{p}$ is the square of another character of H_0.*

Recall that a nilpotent element is called *even* if the corresponding semisimple element $h_{\mathfrak{k}}$ (or, equivalently, $h_{i\mathbf{R}}$) has only even eigenvalues. This means that the representation $\mathrm{ad} \circ \phi$ of $\mathfrak{sl}(2)$ has only odd-dimensional irreducible constituents. In the classical groups, a nilpotent element is even if and only if all its Jordan block sizes have the same parity. For an even nilpotent the space W_{-1} is zero, so we deduce

Corollary 7.28 ([Schwartz]). *Every even nilpotent element is admissible. For such an element, the character of Theorem 7.11 takes values in $\{\pm 1\}$.*

8. Representations attached to admissible orbits.

Recall that one of our goals is a better understanding of what it means for a representation of a reductive group $G_{\mathbf{R}}$ to be attached to a nilpotent imaginary orbit $G_{\mathbf{R}} \cdot \lambda_{i\mathbf{R}}$. We have asserted in section 7 that one should actually try to attach a representation to a $G_{\mathbf{R}}$-admissible orbit datum $(\lambda_{i\mathbf{R}}, \chi_{i\mathbf{R}})$. In section 6, we saw that $\lambda_{i\mathbf{R}}$ corresponds to a nilpotent element $\lambda_{\mathfrak{k}} \in (\mathfrak{g}/\mathfrak{k})^*$. In section 7, we saw that $\chi_{i\mathbf{R}}$ corresponds to a representation $\chi_{\mathfrak{k}}$ of the isotropy group $K(\lambda_{\mathfrak{k}})$, having the following property. Write $\gamma_{\mathfrak{k}}$ for the determinant character of $K(\lambda_{\mathfrak{k}})$ on

$$T^*_{\lambda_{\mathfrak{k}}}(K \cdot \lambda_{\mathfrak{k}}) \simeq (\mathfrak{k}/\mathfrak{k}(\lambda_{\mathfrak{k}}))^* \simeq \mathfrak{p}/\mathfrak{p}(\lambda_{\mathfrak{k}}). \tag{8.1)(a}$$

Then we require

$$\chi_{\mathfrak{k}}(x^2) = \gamma_{\mathfrak{k}}(x) \cdot I \tag{8.1)(b}$$

for all x in the identity component of $K(\lambda_{\mathfrak{k}})$. On the other hand, suppose X is a $(\mathfrak{g}, K)$-module of finite length. By Theorem 2.13 (and Corollary 5.23) one can associate to X a certain finite set of K-orbits

$$K \cdot \lambda_1, \ldots, K \cdot \lambda_s \subset \mathcal{N}^*_{\mathfrak{k}} \tag{8.1)(c}$$

and (genuine virtual) representations χ_i of $K(\lambda_i)$. It is therefore natural to impose the following requirement on the (still undefined) process of attaching representations to nilpotent orbits.

Desideratum 8.2. Suppose $G_{\mathbf{R}}$ is a real reductive group, with complexified maximal compact subgroup K. Suppose $\mathcal{D} = (\lambda_{i\mathbf{R}}, \chi_{i\mathbf{R}})$ is a nilpotent admissible $G_{\mathbf{R}}$-orbit datum, and $(\lambda_{\mathfrak{k}}, \chi_{\mathfrak{k}})$ is a corresponding nilpotent admissible K-orbit datum. If X is a $(\mathfrak{g}, K)$-module attached to $\mathcal{D}$, then $(\lambda_{\mathfrak{k}}, \chi_{\mathfrak{k}})$ should be (up to K-conjugacy) one of the pairs (λ_i, χ_i) attached to X by Theorem 2.13 (cf. (8.1)(c)).

This is a rather weak requirement. In the case of the principal nilpotent orbit for a quasisplit group, it allows whole translation families of fundamental series representations. (According to Desideratum 7.3, there should be at most finitely many representations attached to an orbit datum.) On the other hand, if $G_{\mathbf{R}}$ is semisimple, this requirement alone correctly attaches only the trivial representation to the trivial orbit datum $(0, \mathbb{C})$. In general the requirement is stronger for smaller orbits. Since it is the representations attached to small orbits that are the most troublesome technically, it is worthwhile to pursue Desideratum 8.2. We are therefore led to the problem that is the main concern of this section.

Problem 8.3. Given an orbit $K \cdot \lambda \subset \mathcal{N}_{\mathfrak{k}}^*$, find conditions on a Harish-Chandra module X guaranteeing that

a) $K \cdot \lambda$ contains a component of $\mathcal{V}(X)$, and
b) the corresponding isotropy representation $\chi(\lambda, X)$ (Theorem 2.13) is admissible (Definition 7.13).

Towards part (a) of Problem 8.3, we will prove only the following result.

Theorem 8.4. *Suppose (G, K) is a reductive symmetric pair of Harish-Chandra class (Definition 5.21), and X is an irreducible $(\mathfrak{g}, K)$-module. Write*

$$J = \operatorname{Ann} X \subset U(\mathfrak{g}),$$

a primitive ideal in $U(\mathfrak{g})$. Let $\mathcal{O} \subset \mathcal{V}(J) \subset \mathfrak{g}^*$ be the dense nilpotent G-orbit (Corollary 4.7).

a) $\mathcal{V}(X) \subset \mathcal{V}(J) \cap (\mathfrak{g}/\mathfrak{k})^*$.
b) $\mathcal{O} \cap (\mathfrak{g}/\mathfrak{k})^*$ is the union of a finite number of K-orbits $\mathcal{O}_1, \ldots, \mathcal{O}_r$, each of which has dimension equal to half the dimension of $\mathcal{O}$.
c) Some of the $\mathcal{O}_i$ are contained in $\mathcal{V}(X)$; they are precisely the K-orbits of maximal dimension in $\mathcal{V}(X)$.

Proof. Part (a) is an immediate consequence of the definitions: elements of $\operatorname{gr} J$ (defined using the standard filtration of $U(\mathfrak{g})$) obviously annihilate $\operatorname{gr} X$, and therefore constrain $\mathcal{V}(X)$ according to (1.2)(a). Part (b) follows from Corollary 5.20. Now (c) is evidently equivalent to the assertion that the Gelfand-Kirillov dimension of X (section 2) is at least half that of $U(\mathfrak{g})/J$. Such a relationship holds for any faithful module for a primitive algebra; a proof of the equality in this special case appears in [Vogan1], Lemma 3.4. Q.E.D.

Part (a) of Problem 8.3 is therefore just slightly stronger than asking that

$$\mathcal{V}(\mathrm{Ann}(X)) = \overline{G\cdot\lambda}. \tag{8.5}$$

A great deal is known about how to check such a condition, but we will not review it here.

We turn now to part (b) of Problem 8.3. To formulate a result, we need to recall the *transpose anti-automorphism* $u \mapsto {}^t u$ of $U(\mathfrak{g})$. This is a linear map, characterized by the properties

$${}^t x = -x \quad (x \in \mathfrak{g}), \qquad {}^t(uv) = {}^t v\, {}^t u \quad (u, v \in U(\mathfrak{g}). \tag{8.6}$$

Theorem 8.7. *Suppose (G, K) is a reductive symmetric pair of Harish-Chandra class (Definition 5.21), and X is an irreducible $(\mathfrak{g}, K)$-module. Write*

$$J = \mathrm{Ann}\, X \subset U(\mathfrak{g}),$$

a primitive ideal in $U(\mathfrak{g})$. Let $\mathcal{O} \subset \mathcal{V}(J) \subset \mathfrak{g}^*$ be the dense nilpotent G-orbit (Corollary 4.7). Assume that

i) The $S(\mathfrak{g})$-module $S(\mathfrak{g})/\mathrm{gr}\, J$ is generically reduced along $\mathcal{O}$ (Proposition 2.9; this is automatic if $\mathrm{gr}\, J$ is prime, or more generally if $\mathcal{O}$ has multiplicity one in the characteristic cycle of $S(\mathfrak{g})/\mathrm{gr}\, J$).

ii) The ideal J is preserved by the transpose antiautomorphism of $U(\mathfrak{g})$ (cf. (8.6).

Fix $\lambda \in \mathcal{O} \cap (\mathfrak{g}/\mathfrak{k})^*$, and write H for the corresponding isotropy subgroup of K. Then the character $\chi(\lambda, X)$ of X at λ (Theorem 2.13) is admissible (Definition 7.13).

The proof will occupy the next three sections. We will conclude this section with some easy differential geometry intended to provide motivation for the proof. (Since the results will not be applied directly, we will feel free to omit any inconvenient proofs.) The problem is to understand what kind of natural condition can force χ to be admissible. Suppose for a moment that instead of admissibility (which specifies the differential of χ in a slightly complicated way) we were seeking conditions that force the differential of χ to be zero. The following well-known result provides such conditions; we formulate it in the smooth context for the sake of familiarity.

Proposition 8.8. *Suppose $H \subset K$ are Lie groups, and (χ, V) is a finite-dimensional representation of H. Write $\mathcal{V}$ for the corresponding vector bundle on K/H. Then the following conditions are equivalent.*

a) The differential of χ is zero.

b) $\mathcal{V}$ has a K-invariant structure of local coefficient system on K/H.

c) Attached to every vector field ξ on K/H there is a differential operator L_ξ on the space $\Gamma\mathcal{V}$ of sections of $\mathcal{V}$. This correspondence is complex-linear, and satisfies

i) $L_\xi(f\cdot\sigma) = (\partial_\xi f)\cdot\sigma + f\cdot(L_\xi\sigma) \quad (f \in C^\infty(K/H), \sigma \in \Gamma\mathcal{V})$.

ii) $L_{f\xi}(\sigma) = f\cdot(L_\xi\sigma) \quad (f \in C^\infty(K/H), \sigma \in \Gamma\mathcal{V})$.

iii) Suppose $x \in \mathfrak{k}$, and $\xi(x)$ is the corresponding vector field on K/H. Then $L_{\xi(x)}$ is the natural action of x on sections of $\mathcal{V}$.

In (c), condition (i) simply says that L_ξ is a first-order differential operator with symbol ξ (times an appropriate identity operator). This is a consequence of (ii) and (iii), but we include it to clarify the nature of L.

Condition (b) has been included only for motivation; we are not going to define local coefficient system carefully here. The expert reader will easily supply a proof of its equivalence with (a) or (c). We will prove only the equivalence of (a) and (c).

Proof. Suppose that (c) holds. Write $\mathfrak{m}$ for the ideal in $C^\infty(K/H)$ of functions vanishing at eH. There is a natural isomorphism

$$\Gamma\mathcal{V}/\mathfrak{m}\Gamma\mathcal{V} \simeq V;$$

this takes the restriction to $\mathfrak{h}$ of the natural action of $\mathfrak{k}$ to the differential of χ. If $x \in \mathfrak{h}$, then the vector field $\xi(x)$ vanishes at eH, and may therefore be written in the form

$$\xi(x) = \sum f_i \xi_i$$

with $f_i \in \mathfrak{m}$. If $v \in V$ is represented by a section $\sigma \in \Gamma\mathcal{V}$, then (by (iii)) $\chi(x)v$ is represented by

$$L_{\xi(x)}\sigma = \sum L_{f_i\xi_i}\sigma = \sum f_i L_{\xi_i}\sigma.$$

(Here we have used condition (ii).) This last expression belongs to $\mathfrak{m}\Gamma\mathcal{V}$, and therefore represents zero in V. Therefore the differential of χ is zero, which is (a).

The other direction is very easy, and we will be sketchy. Suppose (a) holds. Write H_0 for the identity component of H. Then K/H_0 is a covering of K/H, and the pullback $\mathcal{V}_0$ of $\mathcal{V}$ to this covering is the vector bundle corresponding to the representation $\chi\mid_{H_0}$ of H_0. By assumption this last bundle is trivial (in a K-invariant way): its sections are just functions on K/H_0 with values in V. It follows immediately that $\mathcal{V}_0$ has the structure required in (c). (L_ξ acts on vector-valued functions by acting on each coordinate separately.) But the structure in (c) is purely local; so its existence on the covering space K/H_0 immediately implies its existence on K/H. Q.E.D.

Next, we prove an analogue of Proposition 8.8 for the case of admissible representations of H.

Proposition 8.9. *Suppose $H \subset K$ are Lie groups, and (χ, V) is a finite-dimensional representation of H. Write $\mathcal{V}$ for the corresponding vector bundle on K/H. Define γ to be the determinant character of H acting on $(\mathfrak{k}/\mathfrak{h})^*$ (the cotangent space at eH to K/H. Fix a complex number k. Then the following conditions are equivalent.*

a) The differential of χ is equal to k times the differential of γ.

b) Attached to every vector field ξ on K/H there is a differential operator L_ξ on the space $\Gamma\mathcal{V}$ of sections of $\mathcal{V}$. This correspondence is complex-linear, and satisfies

i) $L_\xi(f\cdot\sigma) = (\partial_\xi f)\cdot\sigma + f\cdot(L_\xi\sigma)\quad (f \in C^\infty(K/H),\ \sigma \in \Gamma\mathcal{V})$.

ii) $L_{f\xi}(\sigma) = f\cdot(L_\xi\sigma) + k\cdot(\partial_\xi f)\cdot\sigma\quad (f \in C^\infty(K/H),\ \sigma \in \Gamma\mathcal{V})$.

iii) Suppose $x \in \mathfrak{k}$, and $\xi(x)$ is the corresponding vector field on K/H. Then $L_{\xi(x)}$ is the natural action of x on sections of $\mathcal{V}$.

I do not know a good name for the kind of structure considered in (b); of course the second term in (ii) means that it is *not* a connection. The L is intended to suggest

"Lie derivative;" the Lie derivative action of vector fields on volume forms on a manifold satisfies the conditions in (b) with $k = 1$.

Proof. Suppose that (b) holds. Write $\mathcal{T}$ for the tangent bundle of K/H, so that $\Gamma\mathcal{T}$ is the space of vector fields. Exactly as in the proof of Proposition 8.8, we have

$$\Gamma\mathcal{T}/\mathfrak{m}\Gamma\mathcal{T} \simeq T_{eH}(K/H). \tag{8.10)(a}$$

The commutator of two vector fields vanishing at eH again vanishes at eH. Consequently there is a well-defined action of vector fields vanishing at eH on the tangent space at eH:

$$A : \mathfrak{m}\Gamma\mathcal{T} \to \operatorname{End}(T_{eH}(K/H)), \quad A(\tau)(\xi(eH)) = [\tau, \xi](eH). \tag{8.10)(b}$$

If x belongs to $\mathfrak{h}$, then $\xi(x)$ vanishes at eH, and the action of $\xi(x)$ just described is just the adjoint action on $\mathfrak{k}/\mathfrak{h}$. Its trace is therefore the negative of the differential of γ at x:

$$\gamma(x) = -\operatorname{tr} A(\xi(x)) \quad (x \in \mathfrak{h}). \tag{8.10)(c}$$

To compute this trace, choose a set of vector fields so that $\xi_1(eH), \ldots, \xi_n(eH)$ is a basis of $T_{eH}(K/H)$. If τ is any vector field vanishing at eH, we can write (at least near eH)

$$\tau = \sum f_i \xi_i \tag{8.11)(a}$$

for some functions f_i vanishing at eH. Then

$$\begin{aligned} A(\tau)(\xi_j(eH)) &= \sum_i [f_i\xi_i, \xi_j](eH) \\ &= \sum_i (f_i\xi_i\xi_j - \xi_j f_i \xi_i)(eH) \\ &= \sum_i \left(f_i(\xi_i\xi_j - \xi_j\xi_i) - (\partial_{\xi_j} f_i)\xi_i\right)(eH) \\ &= \sum_i \left(f_i(eH)[\xi_i, \xi_j](eH) - (\partial_{\xi_j} f_i)(eH)\xi_i(eH)\right) \\ &= \sum_i (\partial_{\xi_j} f_i)(eH)\xi_i(eH) \end{aligned} \tag{8.11)(b}$$

Consequently

$$\operatorname{tr} A(\tau) = -\sum_i (\partial_{\xi_i} f_i)(eH). \tag{8.11)(c}$$

Now assume that (c) holds. We calculate the differential of χ as in the proof of Proposition 8.5, by identifying V with $\Gamma\mathcal{V}/\mathfrak{m}\Gamma\mathcal{V}$. Suppose that an element $v \in V$ is represented by a section σ, and that $x \in \mathfrak{h}$. Then

$$\chi(x)v \text{ is represented by } L_{\xi(x)}(\sigma). \tag{8.12)(a}$$

Write $\xi(x) = \sum_i f_i \xi_i$ as in (8.11)(a), with $f_i(eH) = 0$. Then condition (ii) shows that

$$L_{\xi(x)}(\sigma) = \sum_i f_i L_{\xi_i}(\sigma) + k \cdot (\partial_{\xi_i} f_i) \cdot \sigma. \tag{8.12)(b}$$

Evaluating at eH gives

$$\chi(x)v = k \cdot \left(\sum_i (\partial_{\xi_i} f_i)(eH) \right) \cdot v. \tag{8.12)(c}$$

By (8.11), the sum is $\gamma(x)$, proving (a).

Since we will not use the converse, we omit a detailed proof. One approach is to pass to a covering as in the proof of Proposition 8.8, and relate $\mathcal{V}$ to the "kth symmetric power" of the volume form bundle. (This is straightforward if k is an integer.) We have already remarked that the Lie derivative action on volume forms satisfies the conditions in (b) with k=1, and one can proceed from there. Q.E.D.

With this proposition as a guide, we can now formulate some analogous algebraic results. We will start in a purely commutative setting (Proposition 9.9); then show how the commutative structure can arise from "Poisson algebras" (Proposition 10.7); and finally show (sometimes) how to get the necessary Poisson algebra structure from a $(\mathfrak{g}, K)$-module (section 11).

9. Proof of Theorem 8.7: commutative algebra.

Suppose C is a graded commutative $\mathbb{C}$-algebra. We are looking for a version of Proposition 8.9, in which this algebra will correspond to the smooth functions on K/H. Recall that a *derivation* of C is a linear endomorphism δ of C satisfying

$$\delta(cc') = c\delta(c') + \delta(c)c'$$

The derivations of C form a graded C-module $\mathrm{Der}(C)$: we say that δ has degree p if it carries C^n to C^{n+p-1}. If δ and δ' are derivations of degrees p and p', then the commutator $[\delta, \delta']$ is a derivation of degree $p + p' - 1$. For our purposes it will be convenient not to work with $\mathrm{Der}(C)$ directly.

Definition 9.1. A *module of derivations* of C is a graded C-module D endowed with a degree-preserving C-module map

$$\partial : D \to \mathrm{Der}(C), \quad \xi \mapsto \partial_\xi,$$

and a $\mathbb{C}$-bilinear skew-symmetric bracket

$$\{,\} : D \times D \to D$$

of degree -1. We require these to satisfy

$$\{\xi, c \cdot \xi'\} = c \cdot \{\xi, \xi'\} + \partial_\xi(c) \cdot \xi' \tag{9.1)(a}$$

$$\partial_{\{\xi,\xi'\}} = [\partial_\xi, \partial_{\xi'}] \tag{9.1)(b}$$

$$\{\xi, \{\xi', \xi''\}\} = \{\{\xi, \xi'\}, \xi''\} + \{\xi', \{\xi, \xi''\}\}. \tag{9.1)(c}$$

(The last condition is the Jacobi identity.) All three conditions are easily verified for the case $D = \mathrm{Der}(C)$. The module D will replace the vector fields in Proposition 8.9.

Definition 9.2. Suppose D is a module of derivations of C, and M is a graded C-module. A *Lie derivative* on M is a graded $\mathbf{C}$-linear map

$$L : D \to \mathrm{Hom}_{\mathbf{C}}(M, M), \quad \xi \mapsto L_\xi$$

satisfying

$$L_\xi(c \cdot m) = c \cdot L_\xi(m) + \partial_\xi(c) \cdot m. \tag{9.2)(a}$$

It is called *torsion-free* if L maps the bracket to commutator of operators:

$$L_{\{\xi,\xi'\}} = [L_\xi, L_{\xi'}]. \tag{9.2)(b}$$

(Notice that this condition makes M a representation of the Lie algebra D.) It is said to be of *k-form type* (for $k \in \mathbf{C}$) if

$$L_{c\cdot\xi}(m) = c \cdot L_\xi(m) + k \cdot \partial_\xi(c) \cdot m. \tag{9.2)(c}$$

(By "k-forms" we understand here not differential forms of degree k, but rather k-th powers of the volume form. Ultimately we will be concerned with the case $k = 1/2$, where the "half-form bundle" considered in distribution theory is of $1/2$-form type in the present sense. This is the origin of the terminology.) Of course such a module M will replace the space $\Gamma\mathcal{V}$ of sections of $\mathcal{V}$ in Proposition 8.9.

Notice that we can define a torsion-free Lie derivative on D itself by $L_\xi(\xi') = \{\xi, \xi'\}$. It is not in general of k-form type for any k, however. The map $L = \partial$ is a torsion-free Lie derivative of 0-form type for C (regarded as a module over itself).

We can now introduce the analogue of the adjoint action of $\mathfrak{h}$ on $\mathfrak{k}/\mathfrak{h}$, as in (8.10).

Lemma 9.3. *In the setting of Definition 9.1, suppose $\mathfrak{m}$ is a maximal ideal in C, corresponding to a homomorphism*

$$\lambda : C \to \mathbf{C}.$$

Then $\mathfrak{m}D$ is a Lie subalgebra of the Lie algebra $(D, \{,\})$. Write

$$A_\mathfrak{m} : \mathfrak{m}D \to \mathrm{End}(D/\mathfrak{m}D), \quad A_\mathfrak{m}(\tau)(\xi + \mathfrak{m}D) = \{\tau, \xi\} + \mathfrak{m}D$$

for the action induced by the adjoint action. Then every endomorphism $A_\mathfrak{m}(\tau)$ has finite rank; so we can define a one-dimensional representation γ of $\mathfrak{m}D$ by

$$\gamma_\mathfrak{m}(\tau) = -\mathrm{tr}\, A_\mathfrak{m}(\tau).$$

Explicitly, suppose $\{\xi_i\}$ *is a finite subset of* D, *and* $a_i \in \mathfrak{m}$. *Then*

$$\gamma_{\mathfrak{m}}\left(\sum a_i\xi_i\right) = \sum \lambda(\partial_{\xi_i}(a_i)).$$

Suppose that M *is a* C*-module and* L *is a Lie derivative on* M *(Definition 9.2). Then the submodule* $\mathfrak{m}M$ *is invariant under the action of* $\mathfrak{m}D$, *so there is a natural representation*

$$\chi_{\mathfrak{m}} : \mathfrak{m}D \to \mathrm{End}(M/\mathfrak{m}M), \quad \chi_{\mathfrak{m}}(\tau)(m + \mathfrak{m}M) = L_\xi(m) + \mathfrak{m}M$$

If L *is of* k*-form type (Definition 9.2), then*

$$\chi_{\mathfrak{m}}(\tau) = k \cdot \gamma_{\mathfrak{m}}(\tau) \cdot I \quad (\tau \in \mathfrak{m}D).$$

The proof is a formal translation of that of Proposition 8.9, and we omit it. (It is worth observing that $\mathfrak{m}^2D$ is a Lie ideal in $\mathfrak{m}D$, and that the representations $A_{\mathfrak{m}}$, $\gamma_{\mathfrak{m}}$, and $\chi_{\mathfrak{m}}$ all factor to $\mathfrak{m}D/\mathfrak{m}^2D$.)

We need to fit this structure together with group actions. So suppose K is a complex algebraic group, acting by degree-preserving algebraic automorphisms of the algebra C: we write

$$\mathrm{Ad} : K \to \mathrm{Aut}(C). \tag{9.4)(a}$$

When ambiguity might arise, we may write instead Ad_C.) The differential of the K action defines a Lie algebra homomorphism from $\mathfrak{k}$ to the (degree-preserving!) derivations of degree 1 of C:

$$\mathrm{ad} : \mathfrak{k} \to \mathrm{Der}^1(C). \tag{9.4)(b}$$

A degree-preserving algebraic action π of K on a C-module M is called *compatible* if it satisfies

$$\pi(g)(c \cdot m) = (\mathrm{Ad}(g)(c)) \cdot (\pi(g)(m)) \tag{9.4)(c}$$

(cf. (2.1)(c)). (Generally we will write such actions on modules with a dot ("module notation") rather than choose a name for the representation.) In this case we call M a *compatible* (C, K)*-module.*

Definition 9.5. In the setting of (9.4)(a), a module D of derivations of C is called *compatible* if we are given an action of K on D by Lie algebra automorphisms, and a K-equivariant Lie algebra homomorphism

$$i : \mathfrak{k} \to D^1$$

satisfying the following conditions:

a) D is a compatible (C, K)-module;
b) for all $x \in \mathfrak{k}$, we have $\partial_{i(x)} = ad_C(x)$; and
c) the differential of the action of K on D is the map $\mathrm{ad}_D \circ i$ (from $\mathfrak{k}$ to $\mathrm{End}(D)$).

Definition 9.6. Suppose D is a compatible module of derivations of C, and M is a compatible graded C-module. A Lie derivative L on M is called *compatible* if

a) $L_{g\cdot\xi}(g\cdot m) = g\cdot(L_\xi(m))$; and
b) the differential of the action of K on M is $L\circ i$.

Notice that the Lie derivatives defined on C and on D after Definition 9.2 are automatically compatible.

Suppose now that D is a compatible module of derivations of C, and that $\mathfrak{m}$ is a maximal ideal in C, corresponding to a homomorphism

$$\lambda : C \to \mathbf{C}. \tag{9.7)(a}$$

The group K acts on the set of such λ; define H to be the isotropy subgroup. Its Lie algebra is easily computed to be

$$\mathfrak{h} = \{\, x \in \mathfrak{k} \mid \mathrm{ad}(x)(\mathfrak{m}) \subset \mathfrak{m}\,\}. \tag{9.7)(b}$$

Because of (9.5)(b), this contains as a subalgebra

$$\mathfrak{h}_1 = \{\, x \in \mathfrak{k} \mid i(x) \in \mathfrak{m}D\,\}. \tag{9.7)(c}$$

The group H acts on $D/\mathfrak{m}D$ and (if M is a compatible (C,K)-module) on $M/\mathfrak{m}M$: we write

$$A_H : H \to \mathrm{End}(D/\mathfrak{m}D), \qquad \chi_H : H \to \mathrm{End}(M/\mathfrak{m}M). \tag{9.7)(d}$$

The definitions of compatibility have been arranged to guarantee that the differentials of these representations agree on $\mathfrak{h}_1$ with the Lie algebra representations of Lemma 9.3:

$$A_H(x) = A_{\mathfrak{m}}(i(x)), \quad \chi_H(x) = \chi_{\mathfrak{m}}(i(x)) \quad (x \in \mathfrak{h}_1) \tag{9.7)(e}$$

(In fact the representations $A_{\mathfrak{m}}$ and $\chi_{\mathfrak{m}}$ may be extended from $\mathfrak{m}D$ to the larger Lie algebra

$$\{\, \tau \in D \mid \partial_\tau(\mathfrak{m}) \subset \mathfrak{m}\,\},$$

which contains $i(x)$; and (9.7)(e) remains valid. The reason we have not done this is that the final — and most interesting — conclusion of Lemma 9.3 appears *not* to extend.)

Finally, we need something like the character γ for H. The Zariski tangent space to $\mathrm{Spec}\, C$ at $\mathfrak{m}$ is

$$T_{\mathfrak{m}}(\mathrm{Spec}\, C) = \mathrm{Hom}(\mathfrak{m}/\mathfrak{m}^2, C/\mathfrak{m}). \tag{9.8)(a}$$

Any derivation of C sends $\mathfrak{m}^2$ into $\mathfrak{m}$, and so defines an element of this tangent space; so we get a natural map

$$\partial_{\mathfrak{m}} : D/\mathfrak{m}D \to T_{\mathfrak{m}}(\mathrm{Spec}\, C). \tag{9.8)(b}$$

We can also define

$$i_{\mathfrak{m}} : \mathfrak{k} \to D/\mathfrak{m}D. \tag{9.8)(c}$$

By inspection of the definitions,

$$\mathfrak{h} = \ker \partial_{\mathfrak{m}} \circ i_{\mathfrak{m}}, \qquad \mathfrak{h}_1 = \ker i_{\mathfrak{m}}. \tag{9.8)(d}$$

(This suggests already the importance of the case when $\partial_{\mathfrak{m}}$ is an isomorphism.) Finally, we define a character γ_H of H by

$$\gamma_H(h) = \det A_H(h)^{-1} \tag{9.8)(e}$$

at least when $D/\mathfrak{m}D$ is finite-dimensional. As in (9.7)(e), this definition is obviously compatible with the one in Lemma 9.3:

$$\gamma_H(x) = \gamma_{\mathfrak{m}}(i(x)), \quad (x \in \mathfrak{h}_1). \tag{9.8)(f}$$

These considerations and Lemma 9.3 prove

Proposition 9.9. *Suppose K is an algebraic group of automorphisms of the graded commutative algebra C, D is a compatible module of derivations of C (Definition 9.6), and M is a compatible (C,K)-module of k-form type (Definition 9.2). Fix a maximal ideal $\mathfrak{m}$ in C, and let H denote its stabilizer in K. Define $\mathfrak{h}_1$ as in (9.7)(c). Let χ_H be the representation of H on $M/\mathfrak{m}M$, and let γ_H be the inverse of the determinant character of H on $D/\mathfrak{m}D$ (assuming this to be finite-dimensional). Then for $x \in \mathfrak{h}_1$, we have*

$$\chi_H(x) = k \cdot \gamma_H(x) \cdot I.$$

Suppose in addition that the map $\partial_{\mathfrak{m}}$ of (9.8)(b) is an isomorphism. Then $\mathfrak{h}_1$ is all of $\mathfrak{h}$, and γ_H is equal to the determinant of the action of H on the Zariski cotangent space $T^*_{\mathfrak{m}}(\operatorname{Spec} C)$.

10. Proof of Theorem 8.7: Poisson algebras.

Our next goal is to show how the structures required in Proposition 9.9 can arise from the theory of Poisson algebras.

Definition 10.1. A *graded Poisson algebra* is a graded commutative algebra R endowed with a Lie algebra structure

$$\{,\} : R \times R \to R$$

of degree -1, such that the bracket with each element of R is a derivation of the commutative algebra structure:

$$\{r, st\} = \{r, s\}t + s\{r, t\}.$$

An ideal $Q \subset R$ is called a *Poisson ideal* if $\{Q, R\} \subset Q$; in this case R/Q is again a Poisson algebra. An ideal I is called *integrable* if $\{I, I\} \subset I$; we will see in a moment how to exploit this weaker condition.

Example 10.2 Suppose G is a Lie group with Lie algebra $\mathfrak{g}$. Then the symmetric algebra $S(\mathfrak{g})$ carries a Poisson algebra structure, determined by the property

$$\{x, y\} = [x, y] \qquad (x, y \in \mathfrak{g}).$$

More generally

$$\{x, p\} = \mathrm{ad}(x)(p) \qquad (x \in \mathfrak{g}, p \in S(\mathfrak{g})).$$

Since ad is the differential of the adjoint action of G, it follows that any $\mathrm{Ad}(G)$-invariant subspace of $S(\mathfrak{g})$ is preserved by bracket with $\mathfrak{g}$; so any $\mathrm{Ad}(G)$-invariant ideal is preserved by Poisson bracket with anything. If X is an $\mathrm{Ad}(G)$-invariant subset of $\mathfrak{g}^*$, then the ideal of zeros $Q(X)$ is a Poisson ideal; so $S(\mathfrak{g})/Q(X)$ is a Poisson algebra. It will be graded if X is homogeneous — for example, if X is a nilpotent coadjoint orbit.

Integrable ideals are even easier to construct in this example. Suppose Q is a Poisson ideal in $S(\mathfrak{g})$, and $\mathfrak{k}$ is a subalgebra of $\mathfrak{g}$. Then $Q + \mathfrak{k}S(\mathfrak{g})$ is an integrable ideal. It will be graded if Q is.

Lemma 10.3. *Suppose R is a graded Poisson algebra, and I is a graded integrable ideal. Set*

$$C = R/I, \qquad D = I/I^2.$$

Then C is a commutative algebra, and D is a graded C-module. D has the structure of a module of derivations of C (Definition 9.1), by

$$\partial_{a+I^2}(b + I) = \{a, b\}, \qquad \{a + I^2, a' + I^2\} = \{a, a'\} + I^2 \qquad (a, a' \in I, b \in R).$$

We omit the straightforward verification.

Next, we consider the Poisson algebra structure that will give rise to Lie derivatives.

Definition 10.4. Suppose R is a graded Poisson algebra, and M is a graded R-module. A *first-order structure* on M is a graded (complex-linear) map of degree -1

$$a_1 : R \times M \to M$$

having the following property:

$$a_1(r, s \cdot m) + ra_1(s, m) = a_1(rs, m) + 1/2\{r, s\} \cdot m \qquad (r, s \in R, m \in M).$$

This structure is said to be *torsion-free* if whenever r and s both annihilate M, we have

$$a_1(\{r, s\}, m) = a_1(r, a_1(s, m)) - a_1(s, a_1(r, m)).$$

This definition (which comes from Gerstenhaber's theory of deformations) is certainly difficult to motivate directly. We will see later how first-order structures arise from $\mathfrak{g}$-modules. For the moment, we offer only a simpler analogy. If M is a complex vector space, we could define a *zeroth-order structure* on M to be a map a_0 of degree 0 from $R \times M$ to M, satisfying

$$a_0(r, a_0(s, m)) = a_0(rs, m).$$

This is nothing but an R-module structure. We could call it *torsion-free* if whenever $a_0(r, m)$ and $a_0(s, m)$ are both zero, then $a_0(\{r, s\}, m)$ is also zero. This just says that the annihilator of M is an integrable ideal. Definition 10.4 is in a technical sense a natural refinement of these notions (of R-module and integrable ideal).

Lemma 10.5. *Suppose R is a graded Poisson algebra, I is a graded integrable ideal, and M is a graded R-module annihilated by I. Define C, D, ∂, and $\{,\}$ as in Lemma 10.3. Then a first-order structure a_1 on M induces a Lie derivative*

$$L : D \to \mathrm{End}(M)$$

of half-form type (Definition 9.2), by the formula

$$L_{s+I^2}(m) = a_1(s, m) \qquad (s \in I, m \in M).$$

If a_1 is torsion-free, then so is L.

Proof. Again this is a routine verification from the definitions. Consider for example the "half-form type" condition. So suppose $r \in R$ and $s \in I$. Then by definition

$$L_{rs+I^2}(m) = a_1(rs, m) = ra_1(s, m) + a_1(r, sm) - 1/2\{r, s\}m.$$

Since s annihilates M, the second term on the right is zero; and the others are

$$(r + I)L_{s+I^2}(m) + 1/2\partial_{s+I^2}(r)m,$$

as required. We leave the other details to the reader. Q.E.D.

We need a way to understand the map $\partial_{\mathfrak{m}}$ of (9.8) in terms of the Poisson structure. Our goal is Proposition 10.9 below. Suppose R is a Poisson algebra and $\mathfrak{m}$ is a maximal ideal in R, corresponding to a homomorphism

$$\lambda : R \to \mathbb{C}, \qquad \ker \lambda = \mathfrak{m}. \tag{10.6)(a}$$

We define a (possibly degenerate) symplectic form $\phi_{\mathfrak{m}}$ (or ϕ_λ) on the Zariski *cotangent* space $\mathfrak{m}/\mathfrak{m}^2$ by

$$\phi_{\mathfrak{m}}(f + \mathfrak{m}^2, g + \mathfrak{m}^2) = \lambda(\{f, g\}) \qquad (f, g \in \mathfrak{m}). \tag{10.6)(b}$$

Of course a bilinear form on the cotangent space is the same as a linear map from the cotangent space to the tangent space; we write

$$\Phi_{\mathfrak{m}} : T^*_{\mathfrak{m}}(\mathrm{Spec}\, R) \to T_{\mathfrak{m}}(\mathrm{Spec}\, R). \tag{10.6)(c}$$

Then $\phi_{\mathfrak{m}}$ is non-degenerate if and only if $\Phi_{\mathfrak{m}}$ is an isomorphism.

Suppose now that $I \subset \mathfrak{m}$ is an ideal. Then I spans a subspace of $\mathfrak{m}/\mathfrak{m}^2$ that we call the *conormal space to R/I at* $\mathfrak{m}$:

$$T^*_{R/I,\mathfrak{m}}(\mathrm{Spec}\, R) = I/(\mathfrak{m}^2 \cap I). \tag{10.7)(a}$$

By definition, the Zariski tangent space $T_{\mathfrak{m}}(R/I)$ consists of linear functionals on $\mathfrak{m}/(\mathfrak{m}^2 + I)$. This may be identified with linear functionals on $\mathfrak{m}/\mathfrak{m}^2$ vanishing on the conormal space to R/I:

$$T_{\mathfrak{m}}(\operatorname{Spec} R/I) = \left(T^*_{R/I,\mathfrak{m}}(\operatorname{Spec} R)\right)^{\perp}. \tag{10.7)(b}$$

Suppose in addition that I is integrable. Then it is immediate from the definitions that

$$\phi_{\mathfrak{m}} = 0 \text{ on } T^*_{R/I,\mathfrak{m}}(\operatorname{Spec} R) \quad (I \text{ integrable}). \tag{10.7)(c}$$

Consequently $\Phi_{\mathfrak{m}}$ restricts to a map

$$\Phi_{\mathfrak{m},I} : I/(\mathfrak{m}^2 \cap I) \to T_{\mathfrak{m}}(\operatorname{Spec} R/I). \tag{10.7)(d}$$

It is clear from the definitions that $\partial_{\mathfrak{m}}$ is the composition of $\Phi_{\mathfrak{m},I}$ with the natural projection

$$I/\mathfrak{m}I \to I/(\mathfrak{m}^2 \cap I). \tag{10.7)(e}$$

Definition 10.8. Suppose R is a Poisson algebra (over $\mathbf{C}$ and that $\mathfrak{m}$ is a maximal ideal in R, with $R/\mathfrak{m} = \mathbf{C}$. We say that $\operatorname{Spec} R$ is *symplectic at* $\mathfrak{m}$ if

i) it is regular at $\mathfrak{m}$ (that is, if $R_{\mathfrak{m}}$ is a regular local ring); and
ii) the form $\phi_{\mathfrak{m}}$ is non-degenerate.

(It is not hard to see that the first condition is in fact a consequence of the second.) Now suppose $I \subset \mathfrak{m}$ is an integrable ideal. We say that R/I is *Lagrangian at* $\mathfrak{m}$ if

i) R is symplectic at $\mathfrak{m}$;
ii) R/I is regular at $\mathfrak{m}$; and
iii) the subspace $T^*_{R/I,\mathfrak{m}}(\operatorname{Spec} R)$ of $T^*_{\mathfrak{m}}(\operatorname{Spec} R)$ is maximally isotropic for the form $\phi_{\mathfrak{m}}$.

Again it is fairly easy to see that the second condition is a consequence of the other two.

Here is our result about $\partial_{\mathfrak{m}}$.

Proposition 10.9. *Suppose R is a Noetherian Poisson algebra, $I \subset R$ is an integrable ideal, and $\mathfrak{m} \supset I$ is a maximal ideal. Assume that R/I is Lagrangian at $\mathfrak{m}$. Then the map $\partial_{\mathfrak{m}}$ (Lemma 10.3 and (9.8)) is an isomorphism.*

Proof. We begin with the underlying linear algebra.

Lemma 10.10. *Suppose (V, ϕ) is a finite-dimensional vector space with a (possibly degenerate) symplectic form, and $S \subset V$ is an isotropic subspace. Write $S^{\perp} \subset V^*$ for the linear functionals vanishing on S. Then the natural map $S \to S^{\perp}$ defined by ϕ is*

i) surjective if and only if S is maximal isotropic, and
ii) injective if and only if S meets the radical of ϕ only in $\{0\}$.

We leave the elementary proof to the reader. As a consequence of Lemma 10.10 and the description of $\partial_{\mathfrak{m}}$ in (10.7), we get immediately the following result.

Lemma 10.11. *In the setting of Lemma 10.3, fix a maximal ideal $\mathfrak{m} \supset I$. Then $\partial_{\mathfrak{m}}$ (cf. (9.8)) is an isomorphism if and only if the following two conditions are satisfied:*

i) the subspace $T^*_{R/I,\mathfrak{m}}(\operatorname{Spec} R)$ is a maximal isotropic subspace;
ii) the symplectic form $\phi_\mathfrak{m}$ is non-degenerate; and
iii) the projection $I/\mathfrak{m}I \to I/(\mathfrak{m}^2 \cap I)$ is an isomorphism.

Conditions (i) and (ii) here are obvious consequences of the Lagrangian hypothesis in Proposition 10.9. For (iii), we use the following lemma.

Lemma 10.12. *Suppose R is a commutative Noetherian ring, $\mathfrak{m}$ is a maximal ideal in R, and $I \subset \mathfrak{m}$ is any ideal. Assume that R is regular at $\mathfrak{m}$. Then R/I is regular at $\mathfrak{m}$ if and only if*

$$I \cap \mathfrak{m}^2 = \mathfrak{m}I.$$

Proof. (We follow [Hartshorne], p.178, where the result is proved for varieties.) Write n for the dimension of the local ring $R_\mathfrak{m}$ (as a vector space over $R/\mathfrak{m}$). Then R is regular at $\mathfrak{m}$ if and only if the Zariski cotangent space $\mathfrak{m}/\mathfrak{m}^2$ has dimension n ([Atiyah-MacDonald], Theorem 11.22). For the cotangent space to R/I, we have the exact sequence

$$0 \to I/(I \cap \mathfrak{m}^2) \to \mathfrak{m}/\mathfrak{m}^2 \to \mathfrak{m}/(\mathfrak{m}^2 + I) \to 0; \qquad (10.13)(a)$$

that is,

$$0 \to I/(I \cap \mathfrak{m}^2) \to T^*_\mathfrak{m}(\operatorname{Spec} R) \to T^*_\mathfrak{m}(\operatorname{Spec} R/I) \to 0. \qquad (10.13)(b)$$

Suppose first that the condition in the lemma is satisfied. Write r for the dimension of $I/\mathfrak{m}I$. Then by Nakayama's lemma, I is generated by r elements near $\mathfrak{m}$ (that is, after inverting a finite number of elements not in $\mathfrak{m}$). It follows that $(R/I)_\mathfrak{m}$ has dimension at least $n-r$. On the other hand, (10.13)(b) shows that the cotangent space at $\mathfrak{m}$ to $\operatorname{Spec} R/I$ has dimension equal to $n-r$. This implies that the dimension of $(R/I)_\mathfrak{m}$ is at most $n-r$ ([Atiyah-MacDonald, Corollary 11.15), with equality if and only if R/I is regular at $\mathfrak{m}$.

Conversely, suppose that R/I is regular of dimension d at $\mathfrak{m}$. By (10.13)(b), $I/(I \cap \mathfrak{m}^2)$ has dimension $n-d$. Choose elements $x_1, \ldots, x_{n-d}$ of I whose images form a basis for this quotient, and define J to be the ideal they generate in R. Obviously $J \subset I$, and by construction

$$J/(J \cap \mathfrak{m}^2) = I/(I \cap \mathfrak{m}^2),$$

a space of dimension $n-d$. On the other hand, $J/\mathfrak{m}J$ has dimension at most $n-d$ (since J has $n-d$ generators); so the natural projection from $J/\mathfrak{m}J$ onto $J/(J \cap \mathfrak{m}^2)$ must be an isomorphism. So J satisfies the condition in the lemma. By the first half of the proof, $(R/J)_\mathfrak{m}$ is a regualr local ring of dimension d. On the other hand $(R/I)_\mathfrak{m}$ is a quotient of $(R/J)_\mathfrak{m}$. Since these are regular local rings (hence integral domains) of the same dimension, it follows that $I_\mathfrak{m} = J_\mathfrak{m}$. It follows at once that $I/\mathfrak{m}I = J/\mathfrak{m}J$; and we have already seen that the right-hand side is isomorphic to $I/(I \cap \mathfrak{m}^2)$. Q.E.D.

Proof of Proposition 10.9. Because of Lemma 10.12, the Lagrangian hypothesis in the proposition implies all three of the conditions given in Lemma 10.11 for $\partial_\mathfrak{m}$ to be an isomorphism. Q.E.D.

Finally, we can bring a group into the picture. Suppose the algebraic group K acts algebraically by graded automorphisms on the graded Poisson algebra R:

$$\operatorname{Ad} : K \to \operatorname{Aut}(R). \qquad (10.14)(a)$$

We say that this action is *Hamiltonian* if there is a K-equivariant Lie algebra homomorphism

$$i : \mathfrak{k} \to R^1 \tag{10.14)(b}$$

with the property that the differential of Ad at x is given by Poisson bracket with $i(x)$. An ideal I is called *K-integrable* if it is integrable and K-invariant, and

$$I \supset i(\mathfrak{k}). \tag{10.14)(c}$$

Proposition 10.15. *In the setting of Lemma 10.5, assume also that R carries a Hamiltonian action of K (cf. (10.14)); that M is a compatible (R, K)-module; that I is K-integrable; and that the first-order structure a_1 is K-equivariant. Then the structures C, D, ∂, $\{,\}$, and L of Lemmas 10.3 and 10.5 all carry compatible K-actions.*

We leave this to the reader.

The conditions appearing in Proposition 10.9 are phrased in terms of Zariski cotangent spaces. It will be convenient for us to relate such conditions to the notion of "generically reduced" (cf. Proposition 2.9). This is accomplished by Lemma 10.17. below.

Lemma 10.16. *Suppose P is a prime ideal in the commutative Noetherian ring R, and $Q \subset R$ is any other ideal. Then Q is generically reduced along P (Proposition 2.9) if and only if there is an element f of R, not belonging to P, such that Q_f is equal either to P_f or to zero. The second possibility occurs exactly when Q is not contained in P.*

Proof. Suppose first that Q is generically reduced. By Proposition 2.9(c), there is an f not in P so that $(R/Q)_f$ is a free $(R/P)_f$-module. In particular, $(R/Q)_f$ is annihilated by P, so it is a quotient of $(R/P)_f$. Any proper quotient of $(R/P)_f$ has a larger annihilator, and so cannot be free unless it is zero. The required statement now follows from the exactness of localization. The argument reverses trivially. Q.E.D.

Lemma 10.17. *Suppose R is a finitely generated commutative algebra over $\mathbb{C}$, P is a prime ideal in R, and $Q \subset P$ is any ideal. Write d for the dimension of R/P. Then Q is generically reduced along P if and only if there is an element f of R, not belonging to P, with the following property: for every maximal ideal $\mathfrak{m}$ of R containing P but not containing f, the Zariski cotangent space*

$$\mathfrak{m}/(\mathfrak{m}^2 + Q)$$

to $\operatorname{Spec} R/Q$ *at* $\mathfrak{m}$ *has dimension d. In this case we have*

$$T^*_{\mathfrak{m}}(\operatorname{Spec} R/Q) \simeq T^*_{\mathfrak{m}}(\operatorname{Spec} R/P)$$

for such $\mathfrak{m}$.

Proof. Suppose Q is generically reduced along P. By Lemma 10.16, there is an h not in P so that $(R/Q)_h = (R/P)_h$. Consequently the Zariski tangent spaces coincide away from the zeros of h. Choose g not in P vanishing on the singular locus of $\operatorname{Spec} R/P$; then $f = gh$ satisfies the requirement of the lemma.

Conversely, suppose such an f exists. Fix a maximal ideal $\mathfrak{m}$ containing P but not f, and consider the local ring $(R/Q)_{\mathfrak{m}}$. Since $Q \subset P$, the dimension of Q is at least d. The hypothesis therefore guarantees that $(R/Q)_{\mathfrak{m}}$ is a regular local ring ([Atiyah-MacDonald], Theorem 11.22) of dimension d. Just as in the proof of Lemma 10.12, it follows that $P_{\mathfrak{m}} = Q_{\mathfrak{m}}$. If we let g be a common denominator for the images in $Q_{\mathfrak{m}}$ of a finite set of generators of P, then $gP \subset Q$. Hence $P_g = Q_g$. The element g does not belong to $\mathfrak{m}$, so it does not belong to P either. By Lemma 10.16, Q is generically reduced along P. Q.E.D.

11. Proof of Theorem 8.7: representation theory.

We can now present the setting in which we will construct first-order structures, Lie derivatives of half-form type, and (eventually) admissible representations of isotropy groups. Suppose (G, K) is a reductive symmetric pair of Harish-Chandra class (Definition 5.21; we assume G is connected, although K need not be). Let X be a $(\mathfrak{g}, K)$-module of finite length. Define

$$J = \operatorname{Ann} X \subset U(\mathfrak{g}), \tag{11.1)(a}$$

a primitive ideal in $U(\mathfrak{g})$. Define

$$Q = \operatorname{gr} J \subset S(\mathfrak{g}), \tag{11.1)(b}$$

the associated graded ideal. Because J is $\mathrm{Ad}(G)$-invariant, Q is as well; so Q is a Poisson ideal (Example 10.2). Its associated variety is the union of the closures of several nilpotent G-orbits (just one if X is irreducible). Fix one of these $\mathcal{O}$, which we assume contains a component of $\mathcal{V}(J)$:

$$\mathcal{V}(J) = \mathcal{V}(Q) \supset \overline{\mathcal{O}}. \tag{11.1)(c}$$

Our Poisson algebra is

$$R = S(\mathfrak{g})/Q. \tag{11.1)(d}$$

Of course we have a Hamiltonian action Ad of K on R; the map i is induced by the inclusion of $\mathfrak{k}$ in $S(\mathfrak{g})$. Define

$$I = Q + \mathfrak{k}S(\mathfrak{g}). \tag{11.1)(e}$$

This is an integrable ideal by Example 10.2; it is obviously K-integrable by the definition (cf. (10.6)(c)). We can therefore introduce

$$C = R/I, D = I/I^2, \partial \tag{11.1)(f}$$

and so on as in Lemma 10.3. The set of maximal ideals in $\operatorname{Spec} C$ is the associated variety of I; that is (on the level of points)

$$\operatorname{Spec} C = \mathcal{V}(I) = \mathcal{V}(Q) \cap (\mathfrak{g}/\mathfrak{k})^* = \overline{\mathcal{O}} \cap (\mathfrak{g}/\mathfrak{k})^*. \tag{11.1)(g}$$

(Here we have used the obvious fact that the associated variety of $\mathfrak{k}S(\mathfrak{g})$ is $(\mathfrak{g}/\mathfrak{k})^*$.) To get the module M, choose any finite-dimensional K-stable generating subspace X_0 of X, and define a filtration as in (1.1). Set

$$M = \operatorname{gr} X; \tag{11.1)(h}$$

obviously M is a compatible (R, K)-module annihilated by I.

We can now state a crucial technical lemma.

Lemma 11.2. *In the setting of (11.1), assume that*

Q is generically reduced along $\mathcal{O}$

(cf. Proposition 2.9; the condition is automatically fulfilled if Q is prime, or more generally if $\mathcal{O}$ has multiplicity one in the characteristic cycle of $S(\mathfrak{g})/Q$). Fix a weight $\lambda \in \mathcal{O} \cap (\mathfrak{g}/\mathfrak{k})^*$, and write $\mathcal{O}_{\mathfrak{k}}$ for its orbit under K. Let $\mathfrak{m}$ denote the maximal ideal corresponding to evaluation at λ (in $S(\mathfrak{g})$, R, or C). Define H to be the stabilizer of λ in K.

a) The Poisson algebra R is symplectic at $\mathfrak{m}$ (Definition 10.8).
b) The ideal I and the C-module M are generically reduced along $\mathcal{O}_{\mathfrak{k}}$, so the representation χ_H of H on $M/\mathfrak{m}M$ (cf. (9.7)(d)) coincides with $\chi(\lambda, M)$ as defined in Definition 2.12.
c) The quotient R/I is Lagrangian at $\mathfrak{m}$ (Definition 10.8).
d) The Zariski tangent space $T_{\mathfrak{m}}(\operatorname{Spec} C)$ is naturally isomorphic to $\mathfrak{k}/\mathfrak{h}$; so the character γ_H of H defined in (9.8)(e) coincides with $\gamma_{\mathfrak{k}}$ as defined in Theorem 7.11.

We postpone the proof (which is a straightforward application of Proposition 5.20) to the end of the section.

In addition, we need a way to define a first-order structure on M.

Lemma 11.3. *In the setting of (11.1), assume that the ideal $J \subset U(\mathfrak{g})$ is preserved by the transpose antiautomorphism of (8.6). Then the R-module M admits a K-equivariant torsion-free first-order structure (Definition 10.4).*

Again we postpone the proof for a moment.

Proof of Theorem 8.7. We introduce the structure and notation of (11.1). By Lemma 11.2 and hypothesis (i), what we must show is that the differential of the representation χ_H of H (cf. (9.7)(d)) is half the differential of γ_H (cf. (8.17)(e)). By Proposition 9.9, it suffices to find a Lie derivative on M (as a C-module) of half-form type. By Lemma 10.5, it suffices to find a K-equivariant first-order structure on M as an R-module. By Lemma 11.3, the existence of such a structure follows from hypothesis (ii). Q.E.D.

Proof of Lemma 11.3. The idea of the proof is that the module structure on M captures X as an A-module "to order zero;" the first order structure captures slightly more of the A-module structure. The implementation we give of this idea is borrowed from [Gerstenhaber] and [BFFLS]; see also [Vogan3], section 3.

We begin with a simple fact about the transpose map. Recall the symmetrization map β from $S(\mathfrak{g})$ to $U(\mathfrak{g})$, and the symbol maps σ_n (see the proof of Lemma 5.2). If f is a homogeneous polynomial of degree n in $S(\mathfrak{g})$, then one checks easily that

$$^t(\beta(f)) = (-1)^n \beta(f). \tag{11.4}$$

Write

$$A = U(\mathfrak{g})/J, \quad J_n = J \cap U_n(\mathfrak{g}), \quad A_n = U_n(\mathfrak{g})/J_n \tag{11.5)(a}$$

Then A is a filtered algebra, and the adjoint action of G on $U(\mathfrak{g})$ factors to A. The isomorphism σ from $\operatorname{gr} A$ onto R is implemented by (surjective) symbol maps

$$\sigma_n : A_n \to R^n \tag{11.5)(b}$$

with kernel precisely A_{n-1}. These maps respect the adjoint action of G. Because G is reductive, we can choose G-equivariant cross-sections

$$\alpha_n : R^n \to A_n, \qquad \sigma_n \circ \alpha_n = \text{ identity}. \tag{11.5)(c}$$

These maps can be added over n, giving a filtered linear isomorphism

$$\alpha : R \to A. \tag{11.5)(d}$$

(When confusion may arise, we will write these maps as σ^A and α^A.) By the hypothesis on J, the transpose antiautomorphism factors to a filtered antiautomorphism of A. By (11.4), the associated graded (anti)automorphism of R acts by $(-1)^n$ on R^n. Now the transpose is involutive, so A is the direct sum of its $+1$ and -1 eigenspaces. In particular, there is a G-invariant complement for A_{n-1} in A_n on which the transpose acts by $(-1)^n$. If we use such a complement to define α, we get

$$^t\alpha_n(r) = (-1)^n \alpha_n(r). \tag{11.5)(e}$$

(It is easy to check that that this requirement determines α_n modulo $U_{n-2}(\mathfrak{g})$.)

Because α is a linear isomorphism, we may use it to pull back the algebra structure from A to R, obtaining a new multiplication that we write as $*$:

$$r * s = \alpha^{-1}(\alpha(r)\alpha(s)). \tag{11.6)(a}$$

This multiplication respects the filtration of R by degree, so we may write it as an infinite sum of bilinear maps m_n of degree $-n$. This means that

$$m_n : R^p \times R^q \to R^{p+q-n} \tag{11.6)(b}$$

and that

$$r * s = \sum_{n=0}^{\infty} m_n(r, s). \tag{11.6)(c}$$

There is no difficulty about convergence: for fixed r and s (say homogeneous of degrees p and q) we will have $m_n(r,s) = 0$ whenever $n > p + q$ (since the gradation of R has no negative terms). The sum in (11.6)(c) is therefore finite. By the definition of associated graded algebra,

$$m_0(r, s) = rs \tag{11.6)(d}$$

In order to prove the lemma, we will need to calculate m_1 as well. Notice first that if a and b belong to A, then

$$^t(ab+ba) = ab+ba, \qquad {}^t(ab-ba) = -(ab-ba). \tag{11.7)(a}$$

By (11.5)(e), it follows that

$$r*s+s*r \in R^{even}, \qquad r*s-s*r \in R^{odd} \tag{11.7)(b}$$

(with obvious notation). In terms of the expansion (11.6)(c), this says that

$$m_{2n+1}(r,s) = -m_{2n+1}(s,r), \qquad m_{2n}(r,s) = m_{2n}(s,r). \tag{11.7)(c}$$

Now if $u \in U_p(\mathfrak{g})$ and $v \in U_q(\mathfrak{g})$, then it is well-known that $uv - vu \in U_{p+q-1}(\mathfrak{g})$. If we write σ^U for the symbol maps in $U(\mathfrak{g})$, then

$$\sigma^U_{p+q-1}(uv-vu) = \{\sigma^U_p(u), \sigma^U_q(v)\}. \tag{11.7)(d}$$

(This is easily proved by induction on p and q; when $p = q = 1$, it amounts to the defining relation of $U(\mathfrak{g})$.) Of course the analogous formula relates commutators in A to the Poisson structure in R. Translated into a statement about the product $*$ on R, (11.7)(d) is

$$m_1(r,s) - m_1(s,r) = \{r,s\}.$$

Now use (11.7)(c) with $n = 0$ to get

$$m_1(r,s) = 1/2\{r,s\}. \tag{11.7)(e}$$

We now make analogous constructions relating X and M. By definition of M, there are K-equivariant symbol maps

$$\sigma_n : X_n \to M^n \tag{11.8)(a}$$

with kernel precisely X_{n-1}. Because K is reductive, we can choose K-equivariant cross-sections

$$\alpha_n : M^n \to X_n, \qquad \sigma_n \circ \alpha_n = \text{ identity}. \tag{11.8)(b}$$

These maps can be added over n, giving a filtered K-equivariant linear isomorphism

$$\alpha : M \to X. \tag{11.8)(c}$$

(Again we will sometimes write σ^X and α^X to avoid confusion.) This isomorphism can be used to pull the A-module structure on X back to an $(R,*)$-module structure on M, which we again denote with a $*$:

$$r * m = (\alpha^X)^{-1}\left(\alpha^A(r)\cdot\alpha^X(m)\right). \tag{11.8)(d}$$

Just as in the case of the product $*$ on R, this action respects the filtrations by degrees, so it may be expanded as a sum of K-equivariant bilinear maps a_n of degree $-n$:

$$a_n : R^p \times M^q \to M^{p+q-n}, \qquad r * m = \sum_{n=0}^{\infty} a_n(r,m). \tag{11.8)(e}$$

By the definition of $\operatorname{gr} M$,

$$a_0(r,m) = r \cdot m \tag{11.8)(f}$$

The map a_1 will turn out to be the first-order structure on M that we are seeking. We have provided no normalization of α "up to order $n-2$" analogous to (11.5)(e), so we cannot expect to find a closed formula for a_1 analogous to (11.7)(e). To get information about a_1, we use the fact that X is an A-module, and therefore that M is an $(R,*)$-module. This means that

$$r * (s * m) = (r * s) * m \qquad (r, s \in R, m \in M). \tag{11.9)(a}$$

Written in terms of the a_n and m_n, this becomes

$$\sum_p a_p\left(r, \sum_q a_q(s,m)\right) = \sum_p a_p\left(\sum_q m_q(r,s), m\right). \tag{11.9)(b}$$

Now suppose r, s, and m are homogeneous of degrees i, j, and k respectively. Then a typical term on either side is homogeneous of degree $i+j+k-p-q$. Collecting terms of the same degree, we obtain finally

$$\sum_{p+q=n} a_p(r, a_q(s,m)) = \sum_{p+q=n} a_p(m_q(r,s),m)). \tag{11.9)(c}$$

We examine these identities one at a time, beginning with $n=0$:

$$a_0(r, a_0(s,m)) = a_0(m_0(r,s),m).$$

Using (11.6)(d) and (11.8)(f), we can rewrite this as

$$r \cdot (s \cdot m) = (rs) \cdot m :$$

that is, M is an R-module. This is an important fact, but hardly new to us. For $n=1$, we get

$$a_0(r, a_1(s,m)) + a_1(r, a_0(s,m)) = a_0(m_1(r,s),m) + a_1(m_0(r,s),m).$$

Using (11.6)(d), (11.7)(e), and (11.8)(f), we get

$$r \cdot a_1(s,m) + a_1(r, s \cdot m) = 1/2\{r,s\} \cdot m + a_1(rs,m). \tag{11.9)(d}$$

This is precisely the requirement for a first-order structure.

To see that a_1 is torsion free, we consider (11.9)(c) with $n = 2$. This is

$$r \cdot a_2(s, m) + a_1(r, a_1(s, m)) + a_2(r, s \cdot m) = m_2(r, s) \cdot m + a_1(1/2\{r, s\}, m) + a_2(rs, m). \tag{11.10)(a}$$

This property is difficult to use, since we do not know m_2 explicitly. However, we do know that m_2 is symmetric (cf. (11.7)(c)). Skew-symmetrizing (11.10)(a) in the variables r and s therefore eliminates the first and third terms on the right, leaving

$$(r \cdot a_2(s, m) - s \cdot a_2(r, m)) + (a_1(r, s \cdot m) - a_1(s, r \cdot m)) + (a_2(r, s \cdot m) - a_2(s, r \cdot m)) = a_1(\{r, s\}, m). \tag{11.10)(b}$$

If r and s both annihilate M, then the first and third terms on the left vanish, leaving

$$a_1(r, s \cdot m) - a_1(s, r \cdot m) = a_1(\{r, s\}, m) \qquad (r, s \in \operatorname{Ann} M). \tag{11.10)(c}$$

This is the definition of torsion-free. Q.E.D.

Proof of Lemma 11.2. Write P for the prime ideal in R defined by $\mathcal{O}$. By the hypothesis and Lemma 10.17,

$$T^*_\lambda(\operatorname{Spec} R) \simeq T^*_\lambda(\operatorname{Spec} R/P) \simeq T^*_\lambda(\mathcal{O}). \tag{11.11)(a}$$

(The set of λ for which this holds is Zariski dense and G-invariant, and therefore includes all of $\mathcal{O}$.) We need the same fact for R/I. Choose g not in P so that $P_g = Q_g$ (Lemma 10.16); by the G-invariance of P and Q we may assume that $g \notin \mathfrak{m}$. Write P' for the prime ideal in R defined by the component of $\mathcal{O}_\mathfrak{k}$ containing λ, and consider the three ideals

$$P' \supset P + \mathfrak{k}R \supset I = Q + \mathfrak{k}R. \tag{11.11)(b}$$

By Proposition 5.20, the first two coincide after we localize at some element f not vanishing at λ. By hypothesis, P coincides with Q after localization at some element g, which may also be chosen not to vanish at λ. So all three coincide after localization at fg. It follows that I is generically reduced along $\mathcal{O}_\mathfrak{k}$. This is part (b) of the lemma (since I annihilates M). By Lemma 10.17,

$$T^*_\lambda(\operatorname{Spec} R/I) \simeq T^*_\lambda(\mathcal{O}_\mathfrak{k}). \tag{11.11)(c}$$

Part (d) of the lemma follows.

To continue, we need to understand the symplectic form ϕ_λ on $T^*_\lambda(\operatorname{Spec} R)$ (cf. (10.6)). Each function f in R defines an element

$$df = (f - \lambda(f)) + \mathfrak{m}^2 \in T^*_\lambda(\mathcal{O}). \tag{11.12)(a}$$

This function also defines a tangent vector $\partial_\lambda(f)$ by the requirement

$$\partial_\lambda(f)(dg) = \lambda(\{f, g\}). \tag{11.12)(b}$$

The map Φ_λ of (10.6) is

$$\Phi_\lambda(df) = \partial_\lambda(f). \tag{11.12)(c}$$

Suppose $x \in \mathfrak{g}$; write $i(x)$ for the restriction to $\mathcal{O}$ of the (linear) function x on $\mathfrak{g}^*$, and $\mathrm{ad}^*(x)$ for the vector field on $\mathcal{O}$ induced by the coadjoint action of G. Then (because the coadjoint action is Hamiltonian)

$$\partial_\lambda(i(x)) = \mathrm{ad}^*(x)(\lambda). \tag{11.12)(d}$$

Now every tangent vector to $\mathcal{O}$ at λ comes from $\mathfrak{g}$; so this implies that the map Φ_λ is surjective. Consequently ϕ_λ is non-degenerate, proving part (a) of the lemma.

Because Φ_λ is an isomorphism, we can use it to transfer the symplectic form ϕ_λ to a symplectic form ϕ_λ^* on the tangent space. Now we already have a non-degenerate symplectic form ω_λ on $T_\lambda(\mathcal{O})$ (cf. (5.14)(b)). Of course we would like to know that this it form coincides with ϕ_λ^*: that is, that

$$\omega_\lambda(\Phi_\lambda(df), \Phi_\lambda(dg)) = \phi_\lambda(df, dg). \tag{11.13)(a}$$

Because of (11.12)(c), this is equivalent to

$$\lambda(\{f, g\}) = \omega_\lambda(\partial_\lambda(f), \partial_\lambda(g)). \tag{11.13)(b}$$

Each side is a bilinear form vanishing on $\mathfrak{m}^2$ and on constants, so it suffices to prove (11.13)(b) for $f = i(x)$ and $g = i(y)$, with x and y in $\mathfrak{g}$. Then both sides are just $\lambda([x, y])$ (cf. Example 10.2 and (5.14)(b)).

Part (c) of the lemma follows from (11.13) and Corollary 5.20. Q.E.D.

12. K-multiplicities in unipotent representations.

In this section we present a conjectural description of the K-multiplicities in certain unipotent representations. Since unipotent representations have not been defined in general, the reader may prefer to regard this as a desideratum rather than a conjecture; but see the remarks after the statement.

Conjecture 12.1. Suppose $(\mathfrak{g}, K)$ is a reductive symmetric pair of Harish-Chandra class, and $\mathcal{O} \subset \mathfrak{g}^*$ is a nilpotent coadjoint orbit. Assume that $\partial\mathcal{O}$ has codimension at least four in $\overline{\mathcal{O}}$. Suppose X is an irreducible unipotent $(\mathfrak{g}, K)$-module attached to $\mathcal{O}$. Then there are

a) an element $\lambda \in \mathcal{O} \cap (\mathfrak{g}/\mathfrak{k})^*$, and

b) an admissible representation χ of the stabilizer $K(\lambda)$ (Definition 7.13)

so that

$$X \simeq Ind_{K(\lambda)}^K(\chi)$$

as a representation of K.

It is possible to replace the "unipotent $(\mathfrak{g}, K)$-modules attached to $\mathcal{O}$" in this conjecture by a precisely defined class. Here is one way to do that. One can find in [Vogan3], Definition 5.5 a definition of "unipotent Dixmier algebra attached to $\mathcal{O}$." (It is not known that algebras satisfying the requirements there exist, but it is conjectured that they do.) Such an algebra A is equipped with a map $U(\mathfrak{g}) \to A$, of which the kernel is a primitive ideal J_A satisfying $\mathcal{V}(\operatorname{gr} J_A) = \overline{\mathcal{O}}$. We can call J_A a *unipotent primitive ideal attached to* $\mathcal{O}$. A unipotent representation attached to $\mathcal{O}$ should have annihilator equal to some J_A. (It seems to be a bad idea to try to use this as a definition of unipotent representation — there are some non-unitary representations annihilated by a J_A, and they should probably not be called unipotent.) At any rate, the conjecture above should hold for any irreducible X with $\operatorname{Ann} X = J_A$ (always assuming the codimension condition on $\partial\mathcal{O}$).

Let us consider how close we are to proving this more precise conjecture. Of course the idea is to apply Theorem 8.7. The ideal J_A will always satisfy ${}^tJ_A = J_A$ (as a consequence of the definition of Dixmier algebra). It will *not* in general satisfy the first hypothesis of Theorem 8.7 (that $\operatorname{gr} J_A$ is generically reduced); but it will satisfy an analogous condition. This analogous condition (which is a little involved to restate here) allows the proof of Theorem 8.7 to be repeated almost without change. We therefore find an element $\lambda \in \mathcal{O} \cap (\mathfrak{g}/\mathfrak{k})^*$ so that $\chi = \chi(\lambda, X)$ is a non-zero admissible representation of $K(\lambda)$. Now Theorem 4.4 shows that as representations of K

$$X \simeq Ind_{K(\lambda)}^{K}(\chi) - E, \tag{12.2}$$

with an "error term" E related to $\partial\mathcal{O}$. (The "codimension at least four" condition in the conjecture gives the "codimension at least two" hypothesis in Theorem 4.4 because of Corollary 5.20.) The K-multiplicities in E are of a lower order of magnitude than those in X, because of the condition on the support of E. Therefore (12.2) says that Conjecture 12.1 is approximately correct. For the rest of this section, we will consider some examples offering various kinds of support or illumination for Conjecture 12.1.

Example 12.3. Suppose $G_{\mathbf{R}}$ is $SL(2,\mathbf{R})$, and $\mathcal{O}$ is the principal nilpotent orbit. This orbit does *not* satisfy the codimension condition of Conjecture 12.1; we want to see how the conclusion of the conjecture fails. The group K is isomorphic to $\mathbf{C}^\times$, so its representations are characters τ_n parametrized by the integers. There are two orbits of K on $\mathcal{O} \cap (\mathfrak{g}/\mathfrak{k})^*$; each has $K(\lambda) = \{\pm 1\}$. Thus $K(\lambda)$ has two admissible representations χ_0 and χ_1, both admissible; we have

$$Ind_{K(\lambda)}^{K}(\chi_a) = \sum_{n \equiv a \,(\mathrm{mod}\, 2)} \tau_n.$$

By contrast, Arthur attaches three (special) unipotent representations to $\mathcal{O}$ (the constituents of the unitary principal series with continuous parameter 0). Their K-characters are

$$\sum_{k=-\infty}^{\infty} \tau_{2k}, \quad \sum_{k=0}^{\infty} \tau_{2k+1}, \quad \sum_{k=0}^{\infty} \tau_{-2k-1}.$$

The first of these is given by a formula as in Conjecture 12.1, but the last two are not.

This example suggests that certain combinations of representations might obey the multiplicity formula in Conjecture 12.1 without the codimension hypothesis. This is false, as one sees by examining the principal orbit in $SU(2,1)$.

Example 12.4. Suppose $G_{\mathbb{R}}$ is the double cover of $SL(3,\mathbb{R})$, and $\mathcal{O}$ is the minimal non-zero nilpotent orbit. Then $\dim \mathcal{O} = 4$ (and $\partial\mathcal{O} = \{0\}$ so the codimension condition is satisfied). K is isomorphic to $SL(2,\mathbb{C})$, which has one irreducible representation τ_n of each dimension n. Up to K conjugacy there is only one possibility for λ, namely a highest weight vector in the representation (isomorphic to τ_5) of K on $\mathfrak{g}/\mathfrak{k})^*$. Consequently

$$K(\lambda) = \left\{ \begin{pmatrix} t & x \\ 0 & t^{-1} \end{pmatrix} \mid t^4 = 1 \right\}.$$

The admissible representations of $K(\lambda)$ are those trivial on the identity component. There are four irreducible ones, given by

$$\chi_a \begin{pmatrix} t & x \\ 0 & t^{-1} \end{pmatrix} = t^a \qquad (a = 0, 1, 2, 3).$$

The corresponding induced representations are

$$Ind^K_{K(\lambda)}(\chi_a) = \sum_{n \equiv a+1 \,(\mathrm{mod}\, 4)} \tau_n.$$

For $a = 0$, $a = 1$, and $a = 2$ there are unitary representations attached to $\mathcal{O}$ having these K-types (see [Torasso]). For $a = 3$ there is no $(\mathfrak{g}, K)$-module with exactly these K-types annihilated by a unipotent primitive ideal. (There *are* two irreducible $(\mathfrak{g}, K)$-modules with these K-types, but they are not unitary. Their primitive ideals are interchanged by the transpose antiautomorphism.)

The point of this example is to show that we cannot insist on the existence of unipotent representations attached to all admissible χ. A good explanation of why χ_3 is not allowed in this case would be extremely valuable.

References.

[Atiyah-MacDonald] M. Atiyah and I. MacDonald, *Introduction to Commutative Algebra.* Addison-Wesley, Reading, Massachusetts, 1969.

[Barbasch-Vogan] D. Barbasch and D. Vogan, "Unipotent representations of complex semisimple Lie groups," Ann. of Math. **121** (1985), 41-110.

[BFFLS] F. Bayen, M. Flato, C. Fronsdal, A. Lichnerowicz, and D. Sternheimer, "Deformation theory and quantization. I. Deformation of symplectic structures," Ann. Physics **111** (1978), 61-110.

[Borho-Brylinski] W. Borho and J.-L. Brylinski, "Differential operators on homogeneous spaces III," Inventiones math. **80** (1985), 1-68.

[Borho-Kraft] W. Borho and H. Kraft, "Über Bahnen und deren Deformationen bei linearen Aktionen reduktiver Gruppen," Comment. Math. Helvetici **54** (1979), 61-104.

[Duflo] M. Duflo, "Théorie de Mackey pour les groupes de Lie algèbriques," Acta Math. **149** (1982), 153-213.

[Gerstenhaber] M. Gerstenhaber, "On the deformation of rings and algebras," Ann. of Math. **79** (1964), 59-103.

[Ginsburg] V. Ginsburg, "Characteristic varieties and vanishing cycles," Inventiones math. **84** (1986), 327-402.

[Grothendieck] A. Grothendieck (rédigés avec la collaboration de J. Dieudonné) *Eléments de Géométrie Algébrique IV. Etude locale des schémas et des morphismes de schémas (Seconde Partie)*, Publications Mathématiques **24** (1965), Institut des Hautes Etudes Scientifiques, Le Bois-Marie, Bures-sur-Yvette, France.

[Guillemin-Sternberg] V. Guillemin and S. Sternberg, *Geometric Asymptotics*, Math Surveys **14**. American Mathematical Society, Providence, Rhode Island (1978).

[Hartshorne] R. Hartshorne, *Algebraic Geometry*. Springer-Verlag, New York-Heidelberg-Berlin, 1977.

[Joseph] A. Joseph, "On the associated variety of a primitive ideal,"J. Algebra **93** (1985), 509-523.

[Kostant] B. Kostant, "The principal three-dimensional subgroup and the Betti numbers of a complex simple Lie group," Amer. J. Math. **81** (1959), 973-1032.

[Kostant-Rallis] B. Kostant and S. Rallis, "Orbits and representations associated with symmetric spaces," Amer. J. Math. **93** (1971), 753-809.

[Loos] O. Loos, *Symmetric Spaces*, vol I. Benjamin, New York-Amsterdam, 1969.

[Matsumura] H. Matsumura, *Commutative Ring Theory*, translated by M. Reid. Cambridge University Press, Cambridge, 1986.

[Schwartz] J. Schwartz, "The determination of the admissible orbits in the real classical groups," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, Massachusetts, June, 1987.

[Sekiguchi] J. Sekiguchi, "Remarks on nilpotent orbits of a symmetric pair," J. Math. Soc. Japan **39** (1987), 127-138.

[Shafarevich] I. Shafarevich, *Basic Algebraic Geometry*, translated by K. Hirsch. Springer-Verlag, Berlin-Heidelberg-New York, 1977.

[SKK] M. Sato, M. Kashiwara, and T. Kawai, "Hyperfunctions and pseudodifferential equations," 265-529 in Lecture Notes in Mathematics **287**. Springer-Verlag, Berlin-Heidelberg-New York, 1973.

[Springer] T.A. Springer, "Microlocalisation algébrique," *Séminaire d'algébre Paul Dubreil et Marie-Paul Malliavin, 36ème année (Paris, 1983-1984)*, 299-316, Lecture Notes in Mathematics **1146**. Springer-Verlag, Berlin-New York, 1985.

[Torasso] P. Torasso, "Quantification geometrique, operateurs d'entrelacement et representations unitaires de $SL_3(\mathbf{R})$," Acta Math. **150** (1983), 153-242.

[van Dijk] G. van Dijk, "Invariant eigendistributions on the tangent space of a rank one semisimple symmetric space," Math. Ann. **268** (1984), 405-416.

[Vogan1] D. Vogan, "Gelfand-Kirillov dimension for Harish-Chandra modules," Inventiones math. **48** (1978), 75-98.

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

[Vogan3] D. Vogan, "Noncommutative algebras and unitary representations," in *The Mathematical Heritage of Hermann Weyl*, R. O. Wells, Jr., ed. Proceedings of Symposia in Pure Mathematics, volume 48. American Mathematical Society, Providence, Rhode Island, 1988.