On the Representation Theory of Semisimple Lie Groups

by

Faisal Al-Faisal

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Pure Mathematics

Waterloo, Ontario, Canada, 2010

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

This thesis is an expository account of three central theorems in the representation theory of semisimple Lie groups, namely the theorems of Borel–Weil–Bott, Casselman–Osborne and Kostant. The first of these realizes all the irreducible holomorphic representations of a complex semisimple Lie group G in the cohomology of certain sheaves of equivariant line bundles over the flag variety of G. The latter two theorems describe the Lie algebra cohomology of a maximal nilpotent subalgebra of $\mathfrak{g}$ with coefficients in an irreducible $\mathfrak{g}$-module. Applications to geometry and representation theory are given.

Also included is a brief overview of Schmid's far-reaching generalization of the Borel–Weil–Bott theorem to the setting of unitary representations of real semisimple Lie groups on (possibly infinite-dimensional) Hilbert spaces.

Acknowledgements

I would like to express my sincere thanks to my family and friends for their support, to my advisor Professor Wentang Kuo for his patience and insight, to my readers Professors Spiro Karigiannis and Kevin Purbhoo for their helpful comments, and to the Department of Pure Mathematics at the University of Waterloo for being my home away from home. Special thanks are also due to Professors Brian Forrest, Kathryn Hare and Nico Spronk for their guidance, help and inspiration all through my time at Waterloo; it has been truly appreciated.

Lastly, I dedicate this thesis to my parents, whose love and encouragement got me to where I am today.

Contents

Introduction

Let G be a complex semisimple Lie group and fix a Borel subgroup $B \subset G$. Then the irreducible holomorphic representations of G lie in one-to-one correspondence with a certain subset—the set of so-called dominant weights—of the character group $\widehat{B}$. This correspondence identifies an irreducible representation with its highest weight: If V is an irreducible representation of highest weight λ, then there is a unique B-invariant line in V on which B acts by multiplication by λ. In the reverse direction, one can start with a dominant weight λ and attempt to induce it to a representation of G. Taking our cue from finite groups[1], we define

$$\operatorname{ind}_H^G(\lambda) = \{f \in \operatorname{Hol}(G, \mathbb{C})\colon f(gb) = \lambda(b)^{-1} f(g) \text{ for all } b \in B \text{ and } g \in G\}$$

and we let G act on this space by the left regular representation, viz.

$$(gf)(g') = f(g^{-1}g'),$$

where $g, g' \in G$ and $f \in \operatorname{ind}_H^G(\lambda)$. This construction makes sense for any $\lambda \in \widehat{B}$. However, $\operatorname{ind}_H^G(\lambda)$ is nonzero if and only if λ is dominant, in which case it is an irreducible representation of G of highest weight λ. This is the content of the Borel–Weil theorem (Theorem 2.3.1).

We can rephrase this induction procedure in more geometric terms. We begin by noting that each $\lambda \in \widehat{B}$ gives rise to a holomorphic line bundle

$$\mathcal{L}_\lambda = G \times_\lambda \mathbb{C} = G \times \mathbb{C}/\{(gb, v) \sim (g, \lambda(b)v)\}$$

[1]If $B \subset G$ are finite and V is a B-module, then $\mathbb{C}[G] \otimes_{\mathbb{C}[B]} V \cong \{f\colon G \to V\colon f(gb) = b^{-1}f(g)\}$.

over the projective variety G/B. There is a natural G-action on $\mathcal{L}_\lambda$ that lies over the G-action on G/B, i.e., $\mathcal{L}_\lambda$ is a G-equivariant line bundle. It follows that there is a linear G-action on the sheaf cohomology spaces $H^q(G/B, \mathcal{L}_\lambda)$. It is then easy to see (cf. Proposition 2.2.5) that we have an isomorphism of representations

$$H^0(G/B, \mathcal{L}_\lambda) \cong \mathrm{ind}_H^G(\lambda).$$

In other words, the space of global holomorphic sections of $\mathcal{L}_\lambda$ can be identified with $\mathrm{ind}_H^G(\lambda)$. In this setting the Borel–Weil theorem becomes the statement that $\mathcal{L}_\lambda$ has global holomorphic sections only when λ is dominant, in which case $H^0(G/B, \mathcal{L}_\lambda)$ is an irreducible representation of G of highest weight λ. Thus the representation theory of G tells us something about the geometry of G/B. Conversely, we obtain geometric realizations of all the irreducible holomorphic representations of G.

In view of the above description of $H^0(G/B, \mathcal{L}_\lambda)$ it is natural to wonder if similar descriptions exist for the higher cohomology of $\mathcal{L}_\lambda$. There is indeed a remarkable description of $H^*(G/B, \mathcal{L}_\lambda)$. By applying Kodaira's vanishing theorem, Borel and Hirzebruch [10] found that $H^q(G/B, \mathcal{L}_\lambda) = 0$ for all $q > 0$ if λ is dominant. Combined with the Borel–Weil theorem, this says that if λ is dominant, then $H^*(G/B, \mathcal{L}_\lambda)$ is concentrated in a single degree, namely 0, at which it is an irreducible representation of G of highest weight λ. In a follow-up paper [11] Borel and Hirzebruch take up the matter of nondominant $\lambda \in \widehat{B}$. They use Hirzebruch's Riemann–Roch formula to show that the Euler characteristic $\chi(G/B, \mathcal{L}_\lambda) = \sum_{i \geq 0} \dim H^i(G/B, \mathcal{L}_\lambda)$ either vanishes or else is equal to (plus or minus) the dimension of an irreducible representation of G. This led them to conjecture that, in general, $H^*(G/B, \mathcal{L}_\lambda)$ either vanishes entirely or else is concentrated in a single degree, at which it is an irreducible representation of G. This conjecture was proved by Bott in [13] and now goes by the name of the Borel–Weil–Bott theorem (Theorem 2.4.1). Bott's theorem includes concrete descriptions of the degree q_λ at which $H^q(G/B, \mathcal{L}_\lambda)$ is nonzero and of the highest weight of the resulting irreducible representation $H^{q_\lambda}(G/B, \mathcal{L}_\lambda)$. It should be remarked that the geometric content, i.e. the description of the cohomology of $\mathcal{L}_\lambda$, is what makes this theorem interesting; we have gained no new representation theoretic information: all the irreducible representations already occur in degree 0. Some applications are given in Section 2.5.

In [13] one can also find a reformulation of Bott's result in terms of Lie algebra cohomology. Bott's basic observation was that one can use the Dolbeault complex to compute $H^*(G/B, \mathcal{L}_\lambda)$. First one notes that $G/B = K/T$, where K is a maximal compact subgroup of G and $T = K \cap B$ is a maximal torus in K. Thus G/B endows the a priori real manifold K/T with a complex structure, and the antiholomorphic tangent bundle of K/T is then seen to be modeled on a maximal nilpotent subalgebra $\mathfrak{n}^-$ of the Lie algebra $\mathfrak{g}$ of G; here $\mathfrak{n}^-$ can be identified as the nilradical of the Borel subalgebra opposite to the Borel subalgebra $\mathfrak{b}$ corresponding to B. The Dolbeault complex, once pulled back to K from K/T, can then be shown to coincide with the complex $\mathrm{Hom}_T(\bigwedge^\bullet \mathfrak{n}, C^\infty(K))$, where $\mathfrak{n}$ is the nilradical of $\mathfrak{b}$ (see Section 3.5.1). This latter complex is a subcomplex of the complex that computes the Lie algebra cohomology $H^*(\mathfrak{n}, C^\infty(K))$. The Peter–Weyl theorem applied to $C^\infty(K)$ can then be used to relate the sheaf cohomology $H^*(G/B, \mathcal{L}_\lambda)$ to the Lie algebra cohomology $H^*(\mathfrak{n}, V)$ with coefficients in an irreducible $\mathfrak{g}$-module V. Specifically, one obtains an identification of K-modules

$$H^q(G/B, \mathcal{L}_\lambda) = \bigoplus V^* \otimes \mathrm{Hom}_{\mathfrak{h}}(\mathbb{C}, H^q(\mathfrak{n}, V)),$$

where the direct sum runs over all the irreducible (smooth) K-modules V and $\mathfrak{h} \subset \mathfrak{b}$ is a Cartan subalgebra of $\mathfrak{g}$, which acts on $\mathbb{C}$ via $-\lambda$ and on $H^q(\mathfrak{n}, V)$ in a certain, natural way (see Proposition 3.2.7).

Thus Bott's theorem can be used to describe—and can be deduced from a description of—the $\mathfrak{h}$-module structure on $H^q(\mathfrak{n}, V)$ for V an arbitrary irreducible $\mathfrak{g}$-module. This description was also obtained in a more algebraic manner by Kostant in [48], and is now called Kostant's theorem (Theorem 3.4.1). This string of ideas is taken up in detail in Chapters 2 and 3 and forms the bulk of this thesis.

In Chapter 4 our attention is turned to the unitary representation theory of real semisimple Lie groups. The ultimate goal is to be able to understand all the irreducible unitary representations of any such group G up to unitary equivalence. If G is compact, then it has a complexification $G_{\mathbb{C}}$, which is a complex semisimple Lie group, and the irreducible unitary representations of G—which turn out to always be finite-dimensional—lie in bijection with the irreducible holomorphic representations of $G_{\mathbb{C}}$. Thus the above discussion applies here, and our aforementioned goal is essentially attained. On the other hand, if G is noncompact, then the situation is more complicated and the goal remains out of reach. There is however a certain class of irreducible unitary representations, the so-called discrete series representations, which behave in many ways as though they are irreducible unitary representations of a compact group and as such have a similar description.

In some precise sense, the various Cartan subgroups of a real semisimple group G make contributions to the set of irreducible unitary representations. More specifically, there are finitely many conjugacy classes of Cartan subgroups, and each conjugacy class makes a certain contribution. There is at most one conjugacy class in G consisting of compact Cartan subgroups, and when such a conjugacy class exists, its contributed representations are precisely the discrete series representations. Contrast this with the case when G is compact, where all the irreducible unitary representations are induced from one-dimensional representations of a maximal torus $T \subset G$—a compact Cartan subgroup of G, which is conjugate to any other such. In other words, if G is compact, then all its irreducible unitary representations belong to the discrete series.

Motivated in part by the analogy between discrete series representations and representations of compact groups, Langlands conjectured in [50] an analogue of the Borel–Weil–Bott theorem for the discrete series in general (cf. Remark 4.4.3(i)). To be precise, let G be a real semisimple Lie group that contains a compact Cartan subgroup H, and consider the L^2 cohomology spaces $H^p_{(2)}(G/H, \mathcal{L}_\lambda)$, where $\mathcal{L}_\lambda = G \times_\lambda \mathbb{C}$ is the line bundle over G/H coming from a one-dimensional representation λ of H, just as before. Then Langlands' conjecture was that $H^*_{(2)}(G/H, \mathcal{L}_\lambda)$ either vanishes entirely or else is concentrated in a single degree at which it is a (specific) irreducible unitary representation in the discrete series of G. Langlands' conjecture was proved by Schmid in [55] (see Theorem 4.4.2 for the precise statement). Schmid's proof relies on, amongst other things, a masterful blending of the ideas used in the proofs of the Borel–Weil–Bott theorem and Kostant's theorem. As such, an outline of this proof as given in Section 4.5—the final section of this thesis—seems like a fitting finale.

Chapter 1

Preliminaries

This thesis makes heavy use of the structure theory and representation theory of semisimple Lie groups, as well as of various ideas coming from algebraic geometry and topology, the most prominent of which is the cohomology of sheaves. In this preliminary chapter we recall (mostly without proof) some basic facts concerning the former. The geometric and topological material shall not be reviewed; instead, we refer the reader to Griffiths and Harris [28] and Taylor [67] for the relevant background material. Good references for the Lie- and representation-theoretic material are Humphreys [40] and Knapp [44], [45].

1.1 Complex Semisimple Lie Groups

The general linear group $\mathrm{GL}(n,\mathbb{C})$ may be viewed as an affine variety by identifying it with the Zariski-closed subset $\{(x_{ij},z)\colon z\det(x_{ij})=1\}$ of $\mathbb{A}^{n^2}\times\mathbb{A}$. A **complex linear algebraic group** is then, by definition, a Zariski-closed subgroup G of $\mathrm{GL}(n,\mathbb{C})$. In particular, G is both an affine variety over $\mathbb{C}$ and a group whose group operations are regular morphisms, i.e. G is a bona fide algebraic group. The **radical** of G, denoted by $R(G)$, is the largest normal, connected, solvable subgroup of G; it is uniquely defined and closed in G, and contains all the normal, connected, solvable subgroups of G. The subset $R_u(G)$ of $R(G)$ consisting of all its unipotent elements (i.e. those which have 1 as their only eigenvalue) is called the **unipotent radical** of G; it is normal in G, and may be characterized as the largest connected normal unipotent subgroup of G. We say that G is **semisimple** (resp. **reductive**) if $R(G)$ (resp. $R_u(G)$) is trivial and G is connected.

A complex linear algebraic group G is in particular a complex Lie group. It therefore has a complex Lie algebra $\mathfrak{g}$. Recall that $\mathfrak{g}$ is said to be **semisimple** if it has no abelian ideals. It is then evident that a complex linear algebraic group G is semisimple if and only if its Lie algebra $\mathfrak{g}$ is semisimple.

On the other hand, suppose that G is a complex Lie group whose Lie algebra $\mathfrak{g}$ is semisimple. If G is also connected, then we say that G is a **complex semisimple Lie group**. It is a theorem that such a G admits, in a unique way, the structure of a complex linear algebraic group.

Our main example of a complex semisimple Lie group shall be the special linear group

$$\mathrm{SL}(n,\mathbb{C}) = \{A \in \mathrm{GL}(n,\mathbb{C})\colon \det A = 1\} \qquad (n > 1).$$

Observe that it is defined by the vanishing of the polynomial $\det(x_{ij}) - 1$, hence is Zariski-closed. Its Lie algebra is

$$\mathfrak{sl}(n,\mathbb{C}) = \{X \in \mathfrak{gl}(n,\mathbb{C})\colon \operatorname{trace} X = 0\},$$

which has no nontrivial ideals, let alone abelian ones. It is also not difficult to see that $\mathrm{SL}(n,\mathbb{C})$ is connected.

1.2 The Levi Decomposition

Every complex linear algebraic group G has a **Levi decomposition**

$$G = LR_u(G),$$

where L is a reductive subgroup (called a **Levi factor**) that intersects $R_u(G)$ trivially. The Levi factor decomposes further as

$$L = SC$$

where S is a semisimple subgroup (called the **semisimple part** of L), C is a torus (see next section), and SC is the centralizer of C in G.

The corresponding direct sum decompositions at the Lie algebra level are

$$\mathfrak{g} = \mathfrak{l} \oplus \mathfrak{r} = [\mathfrak{l}, \mathfrak{l}] \oplus Z(\mathfrak{l}) \oplus \mathfrak{r}.$$

Here $\mathfrak{l}$ is a reductive subalgebra and $\mathfrak{r}$ is the **nilradical** (largest nilpotent ideal) of $\mathfrak{g}$.

1.3 Maximal Tori and Borel Subgroups

An **(algebraic) torus** is a complex linear algebraic group S that is isomorphic to a finite direct product of $\mathrm{GL}(1, \mathbb{C})$'s:

$$S = \mathrm{GL}(1, \mathbb{C}) \times \cdots \times \mathrm{GL}(1, \mathbb{C}) = \mathbb{C}^\times \times \cdots \times \mathbb{C}^\times.$$

A **maximal torus** in G is a closed subgroup which is isomorphic to a torus whose dimension is maximal amongst the dimension of tori isomorphic to subgroups of G.

Assume now that G is a (connected) complex semisimple Lie group. Then G contains maximal tori, and any two of these are conjugate; their common dimension is called the **rank** of G. A **Borel subgroup** of G is a maximal connected solvable subgroup. These too are all conjugate, and their union is the whole of G. Any maximal torus is contained in a Borel subgroup. Conversely, every Borel subgroup contains a maximal torus, which will then be a maximal torus in G. If H is a maximal torus contained in a Borel subgroup B, then H intersects $U = R_u(B)$ trivially, and so we have a Levi decomposition $B = TU$.

At the Lie algebra level, a maximal solvable subalgebra of $\mathfrak{g}$ is called a **Borel subalgebra**, and a **Cartan subalgebra** is a maximal abelian subalgebra such that $\mathrm{ad}(\mathfrak{h})$ consists of diagonalizable operators (which are in fact *simultaneously* diagonalizable, because $\mathfrak{h}$ is abelian). One can then show that the Lie algebra of a maximal torus (resp. Borel subgroup) is a Cartan subalgebra (resp. Borel subalgebra) of $\mathfrak{g}$. Conversely, every Cartan subalgebra (resp. Borel subalgebra) of $\mathfrak{g}$ is the Lie algebra of a maximal torus (resp. Borel subgroup) of G.

To each Borel subgroup $B = HU$ of G there corresponds a (unique) **opposite** Borel subgroup $B^- = HU^-$ characterized by the fact that $B \cap B^- = H$. The Lie algebra of G then decomposes as

$$\mathfrak{g} = \mathfrak{u}^- \oplus \mathfrak{h} \oplus \mathfrak{u},$$

where $\mathfrak{u}^-$, $\mathfrak{h}$ and $\mathfrak{u}$ are the Lie subalgebras corresponding to U^-, H and U, respectively. This is known as the **triangular decomposition** of $\mathfrak{g}$.

Example 1.3.1. In $G = \mathrm{SL}(n, \mathbb{C})$ we may take B to be the upper triangular subgroup, B^- the lower triangular subgroup and T the diagonal subgroup. Then U (resp. U^-) is the upper (resp. lower) triangular subgroup with 1's down the diagonal. At the Lie algebra level, we have that $\mathfrak{b}$ (resp. $\mathfrak{b}^-$) is the upper (resp. lower) triangular subalgebra, $\mathfrak{u}$ (resp. $\mathfrak{u}^-$) is the strictly upper (resp. strictly lower) triangular subalgebra and $\mathfrak{h}$ is the diagonal subalgebra. ▲

1.4 The Jordan–Chevalley Decomposition

Every element x of a complex linear algebraic group $G \subseteq \mathrm{GL}(n, \mathbb{C})$ can be uniquely written as a product $x = x_{ss}x_u = x_u x_{ss}$, where x_{ss} is semisimple (diagonalizable) and x_u is unipotent. This is called the **Jordan–Chevalley decomposition** of x. The Lie algebra analogue is: every $X \in \mathfrak{g}$ can be uniquely written as a sum $X = X_{ss} + X_n$, where X_{ss} is semisimple and X_n is nilpotent, and $[X_{ss}, X_n] = 0$.

1.5 The Weyl Group

Assume that G is complex semisimple, and let H be a maximal torus of G. Then the **Weyl group** of G is, by definition, the *finite* group $W = N_G(H)/H$, where $N_G(H)$ is the normalizer of H in G. As any two maximal tori are conjugate in G, this definition is independent (up to isomorphism) of the choice of H. We shall have more to say about the Weyl group shortly. For now we just remark that that W acts naturally on H via conjugation: if $n_w \in N_G(H)$ is a representative of $w \in W$, then we define $w \cdot h = n_w h n_w^{-1}$ (for $h \in H$).

Example 1.5.1. If $G = \mathrm{SL}(n, \mathbb{C})$ then we take H to be the subgroup of diagonal matrices, in which case it can easily be shown that $N_G(H)$ consists of the matrices in G with exactly one nonzero entry in each row and each column. Consequently W can be identified with the permutation matrices in G, i.e. $W \cong S_n$. ▲

1.6 Representations

1.6.1 Main Definitions

Let V be a *nonzero* finite-dimensional complex vector space. By a **(rational) representation** of a complex linear algebraic group G on V we mean a regular (i.e., algebraic) homomorphism $\varphi\colon G \to \mathrm{GL}(V)$. Similarly, we define a **(holomorphic) representation** of a complex Lie group G on V to be a holomorphic homomorphism $\varphi\colon G \to \mathrm{GL}(V)$. In either case V shall be called a ***G*-module**; the adjectives "rational" and "holomorphic" might be included as well, if such emphasis is necessary. As is usual in such circumstances, reference to the map φ is typically omitted, and one writes gv (for $g \in G$ and $v \in V$) instead of the more correct $\varphi(g)v$.

An **equivariant** map between G-modules V and W is a linear map $T\colon V \to W$ that is also G-linear, i.e. $T(gv) = g(Tv)$ for all $g \in G$ and $v \in V$. Two G-modules V and W are said to be **isomorphic** if there is an invertible equivariant map $T\colon V \to W$ (in which case the inverse map is automatically equivariant).

A subspace U of a G-module V is said to ***G*-invariant** (or a **submodule** of V) if $gU \subseteq U$ for all $g \in G$, in which case (if $U \neq \{0\}$) it can be viewed as a G-module in its own right. If V has no nontrivial G-invariant subspaces, i.e. besides $\{0\}$ and V, then we say that V is **irreducible**. The terminology "simple G-module" is also common, and then V is said to be "semisimple" if it is the direct sum of simple submodules; we shall not be using this terminology.

1.6.2 Lie Algebra Representations

A holomorphic representation $\varphi\colon G \to \mathrm{GL}(V)$ of a complex Lie group G gives rise, by differentiation, to a complex-linear representation $\varphi\colon \mathfrak{g} \to \mathfrak{gl}(V)$ (which we denote by the same letter) of the Lie algebra $\mathfrak{g}$ of G. This passage preserves invariant subspaces and, as a result, takes irreducible G-modules to irreducible $\mathfrak{g}$-modules.

In the other direction, starting with a representation $\varphi\colon \mathfrak{g} \to \mathfrak{gl}(V)$, the best we can hope for is to exponentiate this to a representation of some covering group of G. The general fact is that the representations of $\mathfrak{g}$ lie in one-to-one correspondence with the representations of the universal covering group $\widetilde{G}$ of G (and, under this correspondence, irreducible $\mathfrak{g}$-modules correspond to irreducible $\widetilde{G}$-modules). If G is simply connected, then $\widetilde{G} = G$ and consequently G-modules and $\mathfrak{g}$-modules lie in bijection. In general, $G = \widetilde{G}/\pi_1(G)$, where the fundamental group $\pi_1(G)$ of G is (isomorphic to) a subgroup of the centre of $\widetilde{G}$. Thus the $\mathfrak{g}$-modules which exponentiate to G are those which, when exponentiated to $\widetilde{G}$, are trivial on $\pi_1(G) \subset Z(\widetilde{G})$.

1.6.3 Representations of Complex Semisimple Lie Groups

Assume now that G is a (connected) complex semisimple Lie group. Then we can either think of G as a linear algebraic group or a complex Lie group, and therefore we can speak of rational and holomorphic G-modules. Fortunately, these turn out to be the same thing: every rational G-module is holomorphic (this is obvious) and, conversely, every holomorphic G-module is rational (this lies deeper).

As far as the representation theory of G goes, the following theorem is fundamental. It provides some justification for all the attention we shall devote to the determination of the irreducible representations of G.

Theorem 1.6.1. *Every G-invariant subspace of a G-module has a G-invariant complement. Thus every G-module decomposes into a direct sum of irreducible submodules.* ∎

There are multiple proofs of this theorem; perhaps the easiest one is by means of Weyl's unitary trick, wherein one exploits the invariant integral of a compact real form of G.

1.6.4 Weight Space Decompositions

Assume that G is a complex semisimple Lie group, fix a maximal torus $H \subset G$, and let $\varphi\colon G \to \mathrm{GL}(V)$ be a representation of G. We need the following fact.

Theorem 1.6.2. *If $x = x_{ss}x_u$ is the Jordan–Chevalley decomposition (see Section 1.4) of $x \in G$, then $\varphi(x)_{ss} = \varphi(x_{ss})$ and $\varphi(x)_u = \varphi(x_u)$. That is, the Jordan–Chevalley decomposition factors through representations.* ■

It follows that $\varphi(H)$ consists of simultaneously diagonalizable operators. Thus V decomposes into a direct sum of simultaneous eigenspaces

$$V = \bigoplus V_\lambda.$$

Here the direct sum runs over λ in the **character group** of H, which is

$$\widehat{H} = \{\text{holomorphic homomorphisms } \lambda\colon H \to \mathbb{C}^\times\},$$

and

$$V_\lambda = \{v \in V : \varphi(h)v = \lambda(h)v \text{ for all } h \in H\}.$$

Because V is finite-dimensional, we have $V_\lambda = 0$ for all but finitely many values of λ. Those values of λ for which $V_\lambda \neq 0$ are called the **weights** of V, and V_λ is called the **weight space** in V corresponding to λ. The **multiplicity** of λ in V is defined to be $\dim V_\lambda$.

The action of the Weyl group $W = N_G(H)/H$ on H (see Section 1.5) gives rise to an action on $\widehat{H}$, namely via $(w \cdot \lambda)(h) = \lambda(w^{-1} \cdot h)$. This action clearly permutes the weights of V. More precisely, if $n_w \in N_G(H)$ is a representative of $w \in W$, then $n_w \cdot V_\lambda = V_{w\cdot\lambda}$. It follows that all the weights in the same W-orbit have the same multiplicity.

In general, if the group G is not semisimple, we can still apply the above to the semisimple part S of a Levi factor. Thus, if V is a G-module, its **weights** are defined to be its weights as an S-module. Here V is viewed as an S-module by restriction.

1.7 The Root Space Decomposition

Throughout the following subsections we assume that G is a complex semisimple Lie group, and we fix a maximal torus $H \subset G$.

1.7.1 The Setup

If we apply the ideas of Section 1.6.4 to the adjoint representation $\mathrm{Ad}\colon G \to \mathrm{GL}(\mathfrak{g})$, then the resulting weight space decomposition

$$\mathfrak{g} = \bigoplus \mathfrak{g}_\alpha$$

is very special. Let $\mathfrak{h}$ denote the Lie algebra of H and let the linear functional $\alpha\colon \mathfrak{h} \to \mathbb{C}$ denote the differential of α. Then

$$\begin{aligned} \mathfrak{g}_\alpha &= \{X \in \mathfrak{g}\colon \mathrm{Ad}(h)X = \alpha(h)X \text{ for all } h \in H\} \\ &= \{X \in \mathfrak{g}\colon [Z, X] = \alpha(Z)X \text{ for all } Z \in \mathfrak{h}\}, \end{aligned}$$

so that in particular

$$\mathfrak{g}_0 = \{X \in \mathfrak{g}\colon [Z, X] = 0 \text{ for all } Z \in \mathfrak{h}\} = \mathfrak{h},$$

because $\mathfrak{h}$, being a Cartan subalgebra of $\mathfrak{g}$, is maximal abelian. Thus if we let Φ denote the set of *nonzero* differentials of weights of $\mathfrak{g}$ under the adjoint representation, i.e. $\Phi = \{\alpha \in \mathfrak{h}^*\colon \alpha \neq 0, \mathfrak{g}_\alpha \neq 0\}$, we obtain the direct sum decomposition

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

This is called the **root space decomposition** of $\mathfrak{g}$. The weights in Φ are called the **roots** of $\mathfrak{g}$ and the weight spaces $\mathfrak{g}_\alpha$ are called the **root spaces** of $\mathfrak{g}$. The space $\mathfrak{g}_0 = \mathfrak{h}$ is *not* called a root space, by convention. Because any two maximal tori are conjugate, the root space decomposition is essentially independent of the choice of maximal torus H.

Example 1.7.1. Let $G = \mathrm{SL}(n, \mathbb{C})$, so that $\mathfrak{h}$ is the diagonal subalgebra of $\mathfrak{sl}(n, \mathbb{C})$. Let $e_i \in \mathfrak{h}^*$ be defined by

$$e_i \begin{pmatrix} h_1 & & \\ & \ddots & \\ & & h_n \end{pmatrix} = h_i$$

and let E_{ij} be the $n \times n$ matrix with an entry of 1 in the (i, j)th position and zeros elsewhere. Then the roots of G are $\{e_i - e_j : i \neq j\}$, and the corresponding root space are $\mathfrak{g}_{e_i - e_j} = \mathbb{C}E_{ij}$. ▲

Here are some important facts:

(i) Each root $\alpha \in \Phi$ has multiplicity equal to 1, that is, $\dim \mathfrak{g}_\alpha = 1$.

(ii) If $\alpha \in \Phi$ then $-\alpha \in \Phi$.

(iii) If $\alpha, \beta \in \Phi$ and $\beta \neq -\alpha$, then

$$[\mathfrak{g}_\alpha, \mathfrak{g}_\beta] = \begin{cases} \mathfrak{g}_{\alpha+\beta} & \text{if } \alpha + \beta \in \Phi \\ 0 & \text{otherwise.} \end{cases}$$

On the other hand, $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}] \subset \mathfrak{h}$.

Property (iii) is a special case of a more general phenomenon. Namely, let V be a G-module with weight space decomposition $V = \bigoplus_\lambda V_\lambda$, and let $\varphi \colon \mathfrak{g} \to \mathfrak{gl}(V)$ be the corresponding representation of $\mathfrak{g}$. Then one easily shows that

$$\varphi(\mathfrak{g}_\alpha)V_\lambda \subset V_{\lambda+\alpha} \quad \text{for all } \alpha \in \Phi.$$

In particular,

$$\varphi(\mathfrak{g}_\alpha)V_\lambda = 0 \quad \text{if } \alpha + \lambda \text{ is not a weight of } V. \tag{1.7.1}$$

These two observations form the figurative tip of the iceberg as far as the interaction of roots and weights goes.

We can say a bit more about the inclusion $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}] \subset \mathfrak{h}$. If $X_{\pm\alpha} \in \mathfrak{g}_{\pm\alpha}$ are nonzero, then one can show that $[X_\alpha, X_{-\alpha}] \in \mathfrak{h}$ is nonzero as well, and thus spans $[\mathfrak{g}_\alpha, \mathfrak{g}_{-\alpha}]$. By convention, we choose a normalization of $X_{\pm\alpha}$ so that $Z_\alpha = [X_\alpha, X_{-\alpha}] \neq 0$ satisfies $\alpha(Z_\alpha) = 2$ for all $\alpha \in \Phi$. The triple $\{X_{\pm\alpha}, Z_\alpha\}$ spans a subalgebra of $\mathfrak{g}$ isomorphic to $\mathfrak{sl}(2, \mathbb{C})$. This fact has several remarkable ramifications. For one, because the interaction of the roots and weights exerts some control over the representation theory of $\mathfrak{g}$, this allows the representation theory of $\mathfrak{sl}(2, \mathbb{C})$ to exert some control over that of $\mathfrak{g}$.

1.7.2 Positive and Negative Roots

There are several (equivalent) ways to split the set of roots Φ into a disjoint union of **positive roots** Φ^+ and **negative roots** Φ^- such that $-\Phi^+ = \Phi^-$. Here is one such way. Fix a linear functional $f\colon \mathrm{span}_{\mathbb{R}}\Phi \to \mathbb{R}$ whose kernel does not intersect Φ. Then put

$$\Phi^+ = \{\alpha \in \Phi\colon f(\alpha) > 0\} \quad \text{and} \quad \Phi^- = \{\alpha \in \Phi\colon f(\alpha) < 0\}.$$

We simply write $\alpha > 0$ (resp. $\alpha < 0$) if $\alpha \in \Phi^+$ (resp. $\alpha \in \Phi^-$). The different choices of positive roots, i.e. the different choices of linear functionals f, are related by a certain action of the Weyl group, and hence are ultimately immaterial.

It is easy to see that if $\alpha, \beta > 0$ and $\alpha + \beta \in \Phi$, then $\alpha + \beta > 0$; the analogous assertion holds for negative roots. Furthermore, we can choose a special subset $\Delta = \{\alpha_1, \ldots, \alpha_r\}$ of Φ^+, called the set of **simple roots**, such that every $\alpha > 0$ (resp. $\alpha < 0$) can be written uniquely in the form $\alpha = \sum_{i=1}^r n_i \alpha_i$ with the n_i nonnegative (resp. nonpositive) integers. The integer $r = |\Delta|$ is equal to the rank of G.

If we put $\mathfrak{u}^\pm = \bigoplus_{\alpha \in \Phi^\pm} \mathfrak{g}_\alpha$ and $\mathfrak{b}^\pm = \mathfrak{h} \oplus \mathfrak{u}^\pm$, then $\mathfrak{b}^\pm$ can be seen to be (opposite) Borel subalgebras of $\mathfrak{g}$ with $\mathfrak{u}^\pm$ their nilradicals. It can be shown that every Borel subalgebra containing $\mathfrak{h}$ arises in this fashion, for some choice of positive (or, equivalently, negative) roots, and then it follows that every Borel subalgebra of $\mathfrak{g}$ is conjugate to one of these.

Example 1.7.2. Let $G = \mathrm{SL}(n, \mathbb{C})$. In the notation of Example 1.7.1, we can take $\Phi^+ = \{e_i - e_j\colon i < j\}$, $\Phi^- = \{e_i - e_j\colon i > j\}$ and $\Delta = \{e_i - e_{i+1}\colon 1 \le i \le n-1\}$. The resulting subalgebras $\mathfrak{b}^\pm$ and $\mathfrak{u}^\pm$ are those mentioned in Example 1.3.1. ▲

1.7.3 The Killing Form

The Lie algebra $\mathfrak{g}$ has a canonical symmetric bilinear form, called the **Killing form**, which is given by

$$(X, Y) = \mathrm{trace}\,(\mathrm{ad}(X) \circ \mathrm{ad}(Y)).$$

Because $\mathfrak{g}$ is semisimple, the Killing form is nondegenerate.[1] Two root spaces $\mathfrak{g}_\alpha$ and $\mathfrak{g}_\beta$ are orthogonal with respect to the Killing form unless $\alpha + \beta = 0$; the Cartan subalgebra $\mathfrak{h}$ is orthogonal to $\mathfrak{g}_\alpha$ for all $\alpha \in \Phi^+$. Thus the restriction of the Killing form to $\mathfrak{h}$ remains nondegenerate, hence yields an isomorphism $\lambda \to H_\lambda$ of $\mathfrak{h}^*$ onto $\mathfrak{h}$. We use this to define a bilinear form on $\mathfrak{h}^*$ by $(\lambda, \mu) = (H_\lambda, H_\mu)$.

The restriction of this bilinear form to the r-dimensional real subspace $E = \mathrm{span}_{\mathbb{R}}\Phi$ of $\mathfrak{h}^*$ is an inner product: it is real-valued and positive-definite. We can thus endow E with the structure of a real Euclidean space.

1.7.4 Weyl Chambers

For $\alpha \in \Phi$, let $s_\alpha \colon E \to E$ be the reflection in the hyperplane $E_\alpha = \{\lambda \in E \colon (\lambda, \alpha) = 0\}$ orthogonal to α. Then

$$s_\alpha(\lambda) = \lambda - (\lambda, \alpha^\vee)\alpha,$$

where $\alpha^\vee = 2\alpha/(\alpha, \alpha)$ is the **coroot** corresponding to α. We have that $H_{\alpha^\vee} = Z_\alpha$, where Z_α is as in Section 1.7.1. That is, the triple $\{X_{\pm\alpha}, H_{\alpha^\vee}\}$ spans a subalgebra of $\mathfrak{g}$ isomorphic to $\mathfrak{sl}(2, \mathbb{C})$. This observation is the key to establishing many special properties of Φ, the most important of which is that Φ forms what is known as an **(abstract) root system**; we shall not go into any more details here.

The hyperplanes $\{E_\alpha\}_{\alpha \in \Phi^+}$ partition E into open convex cones called **Weyl chambers**. More precisely, the Weyl chambers are the connected components of the set $E \setminus \cup_{\alpha \in \Phi^+} E_\alpha$. Any such chamber is of the form

$$\{\lambda \in E \colon \varepsilon(\alpha)(\lambda, \alpha^\vee) > 0 \text{ for all } \alpha \in \Phi^+\} \tag{1.7.2}$$

[1] In fact, semisimple Lie algebras are characterized by the fact that their Killing form is nondegenerate: this is *Cartan's criterion for semisimplicity.*

where ε is some function $\Phi^+ \to \{\pm 1\}$, though not every such function defines a nonempty set. There is a distinguished Weyl chamber, called the **fundamental** (or **dominant**) **chamber**, which is given by

$$\{\lambda \in E\colon (\lambda, \alpha^\vee) > 0 \text{ for all } \alpha \in \Phi^+\} = \{\lambda \in E\colon (\lambda, \alpha^\vee) > 0 \text{ for all } \alpha \in \Delta\}.$$

Also important is its closure

$$\{\lambda \in E\colon (\lambda, \alpha^\vee) \geq 0 \text{ for all } \alpha \in \Delta\} = \{\lambda \in E\colon (\lambda, \alpha^\vee) \geq 0 \text{ for all } \alpha \in \Phi^+\}$$

which is called, naturally enough, the **closed fundamental chamber**.

An element $\lambda \in E$ is said to be **regular** if $(\lambda, \alpha^\vee) \neq 0$ for all $\alpha \in \Phi$, otherwise it is said to be **singular**.

The picture one ought to have in mind is the following one. The Weyl chambers divide the space E into open convex cones. The closure of any one of these chambers simply consists of the chamber and its "walls" (just change the > 0 to ≥ 0 in (1.7.2)). Then an element $\lambda \in E$ is regular if it does not lie on any of these walls.

1.7.5 The Weyl Group Revisited

The subgroup of $\mathrm{GL}(E)$ (actually of $\mathrm{O}(E)$) generated by the reflections $\{s_\alpha\}_{\alpha\in\Phi}$ is isomorphic to the Weyl group $W = N_G(H)/H$ of G. One can show that W permutes the Weyl chambers *freely* and *transitively*: if C and C' are Weyl chambers, then there exists a unique $w \in W$ such that $wC = C'$. It follows that possible sets of simple roots (and positive roots) are also permuted freely and transitively by W.

It suffices to take the **simple reflections** s_α with $\alpha \in \Delta$ as a generating set for W. Hence each $w \in W$ can be expressed as a product of simple reflections. The minimum number of simple reflections needed for such an expression is called the **length** of w and is denoted by $\ell(w)$. Note that $\ell(w) = \ell(w^{-1})$. It is not difficult to show that

$$\ell(w) = |\{\alpha \in \Phi^+ \colon w\alpha < 0\}|,$$

and from this we deduce that there is a unique **longest element** w_0 in W whose length is

$$\ell(w) = |\Phi^+| = \dim \mathfrak{u}^+ = \dim \mathfrak{g}/\mathfrak{b}^-.$$

This element w_0 is distinguished by the fact that $w_0\Phi^+ = \Phi^-$. Additionally, we have $w_0^2 = 1$ and any $w \in W$ can be expressed as $w = w'w_0$, where $\ell(w) = \ell(w_0) - \ell(w')$.

1.8 Cartan–Weyl Theory

In this section we recall the Cartan–Weyl parametrization of irreducible G-modules by means of the theorem of the highest weight. To this end, fix a complex semisimple Lie group G, a maximal torus $H \subset G$, a choice of positive roots $\Phi^+ \subset \Phi$ (in which case the negative roots are $\Phi^- = -\Phi^+$) and a choice of simple roots $\Delta \subset \Phi^+$.

1.8.1 Integrality and Dominance

Denote by L the kernel of the exponential map $\exp\colon \mathfrak{h} \twoheadrightarrow H$ (then $H = \exp(\mathfrak{h}) \cong \mathfrak{h}/L$). The **weight lattice** in $\mathfrak{h}^*$ is the lattice

$$\Lambda = \{\lambda \in \mathfrak{h}^* \colon \lambda(L) \subset 2\pi i\mathbb{Z}\}.$$

Linear functionals $\lambda \in \mathfrak{h}^*$ which belong to Λ are said to be **analytically integral**.

We have a group isomorphism $\lambda \to e^\lambda$ of Λ onto the character group $\widehat{H}$, where e^λ is defined explicitly by

$$e^\lambda(\exp Z) = \exp(\lambda(Z)) \quad \text{for } Z \in \mathfrak{h}.$$

Thus if $V = \bigoplus_\lambda V_\lambda$ is the weight space decomposition of a G-module V, then each weight λ lies in Λ. In particular, $\Phi \subset \Lambda$.

One can show that every $\lambda \in \Lambda$ satisfies the following integrality property:

$$(\lambda, \alpha^\vee) \in \mathbb{Z} \quad \text{for all } \alpha \in \Phi. \tag{1.8.1}$$

An immediate consequence is that $\Lambda \subset E = \mathrm{span}_{\mathbb{R}}\Phi$. One easily checks that the action of the Weyl group on E leaves Λ invariant.

Remark 1.8.1. In general, if a linear functional $\lambda \in \mathfrak{h}^*$ satisfies (1.8.1) then it is said to be **algebraically integral**. The set of algebraically integral weights forms a lattice Λ' in $\mathfrak{h}^*$ which contains the weight lattice Λ. The quotient Λ'/Λ is isomorphic to $\pi_1(G)$. So if (and only if) G is simply connected, then every algebraically integral weight is analytically integral. ▲

We say that $\lambda \in \Lambda$ is **dominant** if it belongs to the closed fundamental Weyl chamber. That is, $\lambda \in \Lambda$ is dominant if and only if

$$(\lambda, \alpha^\vee) \geq 0$$

for all $\alpha \in \Phi^+$ (or, equivalently, for all $\alpha \in \Delta$).

The fact that the Weyl group shuffles the the Weyl chambers can be used to show that, for every $\lambda \in \Lambda$, there is a *unique* $w \in W$ such that $w\lambda$ is dominant.

1.8.2 The Theorem of the Highest Weight

Let $\|\cdot\|$ denote the norm on E induced by the inner product $(\cdot,\cdot)$.

Theorem 1.8.2 (Theorem of the Highest Weight)**.** *Let V be an irreducible G-module. Then there exists exactly one weight λ of V with the following properties:*

(i) *The weight space V_λ is one-dimensional.*

(ii) *Every weight μ of V can be expressed as $\lambda - \sum_{i=1}^{r} n_i\alpha_i$, where the n_i's are in $\mathbb{Z}_{\geq 0}$ and the α_i's are the simple roots.*

(iii) *Every weight μ of V satisfies the inequality $\|\mu\| \leq \|\lambda\|$. Equality holds if and only if $\mu = w\lambda$ for some $w \in W$ (such weights μ are called **extremal** in V).*

(iv) *$\lambda + \alpha$ is not a weight of V for any $\alpha \in \Phi^+$. Hence the positive root spaces $\mathfrak{g}_\alpha$ all annihilate V_λ (see (1.7.1)), and the members of V_λ are the only vectors with this property.*

The weight λ is called the ***highest weight*** *of V. It determines the representation V up to isomorphism: if W is another irreducible G-module with highest weight μ, then $V \cong W$ if and only if $\lambda = \mu$. The highest weight λ is dominant and every dominant member of Λ arises as the highest weight of some irreducible G-module.* ■

Thus the irreducible G-modules are parametrized by the set of dominant weights in Λ. The choice of highest weight in an irreducible representation depends on the choice of the fundamental Weyl chamber. Thus the possible choices lie in the same W-orbit, i.e. they are the extremal weights.

In proving the above theorem one finds that the last assertion—showing that every dominant weight in Λ arises as the highest weight of an irreducible G-module—requires the most effort. We shall give several proofs of this assertion in later chapters. Our method will be to construct the desired G-module in some geometric fashion. And on this note, we now turn to more geometric matters.

1.9 Homogeneous Spaces

1.9.1 The Case of a Complex Lie Group

Let G be a complex Lie group. If X is a complex manifold, then by a **G-action** on X we shall mean a holomorphic map $G \times X \to X$ that is a group action. A space X endowed with a G-action shall be referred to as a **G-space**. A **homogeneous space** is a G-space with a transitive G-action.

The principal example of a homogeneous space is the quotient space G/H of G by a (closed) complex Lie subgroup H. This is a complex manifold and the action of G on it (by left translation) is holomorphic. Conversely, if X is a homogeneous space, then the isotropy subgroup $G_x = \{g \in G \colon gx = x\}$ at a point $x \in X$ is a complex Lie subgroup of G, and the map $g \to gx$ of G onto X induces an equivariant (i.e. action-respecting) isomorphism $G/G_x \cong X$ of complex manifolds.

1.9.2 The Case of a Complex Linear Algebraic Group

Let G be a complex linear algebraic group (or indeed a general algebraic group). In analogy with the preceding section, we would like to define a homogeneous space in this setting to be the quotient G/H of G by a closed subgroup H. However, there is the subtle issue of endowing G/H with the appropriate structure: it turns out that we must permit G/H to be a non-affine variety. The next result allows us to equip G/H with the structure a quasiprojective variety.

Theorem 1.9.1 (Chevalley)**.** *Let G be an algebraic group and $H \subset G$ a closed subgroup. Then there is a rational representation $\varphi\colon G \to \mathrm{GL}(V)$ and a one-dimensional subspace L of V such that $H = \{x \in G \colon \varphi(x)L = L\}$.* ■

This theorem implies that if we think of G as acting on $\mathbb{P}(V)$ by means of the representation φ, then H will be the isotropy subgroup of the point L. Hence we can identify the quotient space G/H with the G-orbit in $\mathbb{P}(V)$ of L. By general principles, this orbit is open in its closure, i.e. it is a quasiprojective variety, and we endow G/H with this structure. It follows that if G/H is complete, then it is a *projective variety*, which—as far as this thesis is concerned!—is highly desirable.

1.10 The (Generalized) Flag Variety

Let G be a complex semisimple Lie group. We may (and shall) view G as a complex linear algebraic group. A closed subgroup P of G is called **parabolic** if the quotient space G/P is complete, hence projective. Because G/P carries a transitive G-action, it is in fact a *smooth* projective variety, and as such can also be viewed as a compact complex manifold.

Theorem 1.10.1. *A closed subgroup is parabolic if and only if it contains a Borel subgroup.*[2] *In particular, a connected subgroup B of G is a Borel subgroup if and only if G/B is a projective variety.* ■

[2] It has been claimed that this property explains why “parabolic subgroups” are so-named: a parabolic subgroup is *para-Borelic*. I do not have any evidence to support this claim.

The proof of this theorem is not difficult, but it requires a slight digression (by way of Borel's fixed point theorem), so we shall omit it.

If B is a Borel subgroup of G, then the projective variety G/B is called the **flag variety** of G. Since any two Borel subgroups are conjugate, the flag variety is well-defined up to isomorphism. An alternative description may be obtained as follows. Let G act on its Lie algebra $\mathfrak{g}$ (via Ad). Since any two Borel subalgebras are conjugate, this action is transitive on the set X of Borel subalgebras of $\mathfrak{g}$. The isotropy subgroup at a fixed point $\mathfrak{b} \in X$ is just the centralizer $Z_G(\mathfrak{b})$ of $\mathfrak{b}$ in G, and thus we have a (set-theoretic) identification $G/Z_G(\mathfrak{b}) \cong X$. The Lie algebra of $Z_G(\mathfrak{b})$ in $\mathfrak{g}$ is easily seen to be $\mathfrak{b}$ itself. This means that $Z_G(\mathfrak{b})$ (or rather its connected component—but it can be shown that $Z_G(\mathfrak{b})$ is connected, *because G is complex*) is a Borel subgroup of G. We have thus identified X with the flag variety of G.

This alternative description of the flag variety is often useful. Had we started with it as a definition, we would have been able to give an easy proof of the fact that the flag variety is a projective variety. Indeed, let d be the dimension of any (hence all) Borel subalgebras of $\mathfrak{g}$. Then X can be identified with a subspace of the Grassmannian $\mathrm{Gr}(d, \mathfrak{g})$ of d-planes in $\mathfrak{g}$. The Grassmannian is a well-known example of a projective variety (the *Plücker map* $W \to \bigwedge^d W$ is a closed embedding of $\mathrm{Gr}(d, \mathfrak{g})$ into $\mathbb{P}(\bigwedge^d \mathfrak{g})$), and the fact that Borel subalgebras are solvable translates into a system of homogeneous polynomial equations that cut out X in $\mathrm{Gr}(d, \mathfrak{g})$, whence X is projective. The details of this argument can be found in [66, §19.7].

The variety G/P, where P is a parabolic subgroup of G, is called a **generalized flag variety**. The flag variety G/B is distinguished amongst the generalized flag varieties G/P by being the "largest." More precisely, the inclusion $B \subset P$ induces a surjective map $G/B \to G/P$, forcing $\dim_{\mathbb{C}} G/P \leq \dim_{\mathbb{C}} G/B$. As an aside: it can be shown that the map $G/B \to G/P$ is a fibration, and that the fibre P/B is itself a projective variety. Indeed, if $P = SCU_P$ and $B = HU$ are Levi decompositions, then $U_P \subset U$ and $CU_P \subset B$, so that $P/B = S/(S \cap B)$ and $S \cap B$ is a Borel subgroup of the semisimple group S. We shall explicitly identify P/B for a special class of parabolic subgroups in the next section.

Example 1.10.2. Let $G = \mathrm{SL}(n, \mathbb{C})$ and B the subgroup of upper triangular matrices.

Then B fixes a complete flag $\{0\} = V_0 \subset V_1 \subset \cdots \subset V_n = \mathbb{C}^n$ of subspaces V_i in $\mathbb{C}$—here the word complete means that $\dim V_i = i$, so that there are no dimensional gaps in the chain.[3] The group G acts transitively on the set of such flags, with B being the isotropy subgroup of the aforementioned flag $V_\bullet$. Hence G/B, the flag variety of $\mathrm{SL}(n, \mathbb{C})$, may be identified with the variety of complete flags in $\mathbb{C}^n$, which is typically denoted by $\mathcal{F}l(\mathbb{C}^n)$. In particular, the flag variety of $\mathrm{SL}(2, \mathbb{C})$ is the projective line $\mathbb{P}^1$.

Now let P be a parabolic subgroup containing B. Arguing as in the preceding paragraph, it is easy to see that G/P will be the variety $\mathcal{F}l(d_1, \ldots, d_k; \mathbb{C}^n)$ of partial flags $0 = V_0 \subset V_1 \subset \cdots \subset V_k = \mathbb{C}^n$ in $\mathbb{C}^n$ of a fixed signature $(d_1, d_1, \ldots, d_k)$, where $d_i = \dim V_i$ and $d_i < d_{i+1}$. In particular, the Grassmannian $\mathrm{Gr}(d, \mathbb{C}^n) = \mathcal{F}l(0, d, n; \mathbb{C}^n)$ is a generalized flag variety. ▲

1.11 Parabolic Subgroups and Parabolic Subalgebras

Let G be a complex semisimple Lie group. We shall now examine more closely the structure of parabolic subgroups. We first make the obvious definition: a **parabolic subalgebra** is the Lie algebra of a parabolic subgroup of G; it can be equivalently defined as a Lie subalgebra of $\mathfrak{g}$ that contains a Borel subalgebra.

Fix a Borel subalgebra $\mathfrak{b}$ in $\mathfrak{g}$. Then $\mathfrak{b} = \mathfrak{h} \oplus \bigoplus_{\alpha \in \Phi^-} \mathfrak{g}_\alpha$, where $\mathfrak{h}$ is a Cartan subalgebra of $\mathfrak{g}$ and Φ^- is a choice of negative (say) roots. It is now a simple matter to determine the parabolic subalgebras containing $\mathfrak{b}$. They are all of the form

$$\mathfrak{p} = \mathfrak{b} \oplus \bigoplus_{\alpha \in \Phi_P^+} \mathfrak{g}_\alpha,$$

where Φ_P^+ is a subset of Φ^+ that is closed under the addition of roots. Write $\Delta = \{\alpha_1, \ldots, \alpha_r\}$ for the simple roots in Φ^+. It must be the case that Φ_P^+ is generated by a subset of Δ, say $\{\alpha_i\}_{i \in I}$ for some index set $I \subseteq \{1, \ldots, r\}$. Hence the parabolic subalgebras containing $\mathfrak{b}$ lie in bijection with the 2^r subsets of $\{1, \ldots, r\}$. We write $\mathfrak{p}_I$ for

[3] Procesi [53, p.514] suggests the following explanation for the word "flag." In two dimensions, a complete flag consists of a point, a line and a plane. Think of the plane as an actual flag, the line as its pole, and the point as its base!

the parabolic subalgebra corresponding to $I \subseteq \{1, \ldots, r\}$. In particular, if $I = \{i\}$ is a singleton, then we simply write $\mathfrak{p}_{\alpha_i}$ for $\mathfrak{p}_I$, and call the resulting subalgebras the **minimal parabolic subalgebras** of $\mathfrak{g}$.

Analogous remarks hold at the Lie group level. For example, we write P_I for the parabolic subgroup corresponding to the set $I \subseteq \{1, \ldots, r\}$, etc. In this way we obtain also all the parabolic subgroups containing the Borel subgroup $B = Z_G(\mathfrak{b})$.

Let $P = SCU_P$ be the Levi decomposition of parabolic subgroup $P \supseteq B$. Then the semisimple part S has Lie algebra

$$\mathfrak{s} = \bigoplus_{\alpha \in \Phi_P^+} (\mathbb{C}H_{\alpha^\vee} \oplus \mathfrak{g}_\alpha \oplus \mathfrak{g}_{-\alpha}),$$

the torus C has Lie algebra

$$\mathfrak{c} = \bigoplus_{\alpha \notin I} \mathbb{C}H_{\alpha^\vee},$$

and the unipotent radical U_P has Lie algebra

$$\mathfrak{u}_P = \bigoplus_{\alpha \in \Phi^+ - \Phi_P^+} \mathfrak{g}_{-\alpha}.$$

In particular, if $P_I = P_\alpha$ is a minimal parabolic subgroup corresponding to a simple root α, then P_α/B is the flag variety of $\mathfrak{s} = \mathbb{C}H_{\alpha^\vee} \oplus \mathfrak{g}_\alpha \oplus \mathfrak{g}_{-\alpha} = \mathfrak{sl}(2, \mathbb{C})$ (see the remarks preceding Example 1.10.2), that is, $P_\alpha/B = \mathbb{P}^1$.

1.12 The Bruhat Decomposition

Let G be a complex semisimple Lie group, fix a maximal torus $H \subset G$ and a Borel subgroup B containing H, and let $W = N_G(H)/H$ be the Weyl group. In its weakest form, the **Bruhat decomposition** refers to the decomposition of G as a disjoint union $G = \bigsqcup_{w \in W} BwB$ of double cosets of B.

We shall require a more refined version of this decomposition. To this end, let U denote the unipotent radical of B. Pick for each $w \in W$ a coset representative $n_w \in N_G(H)$. Then as $H \cap U = \{1\}$ and n_w normalizes H, we have

$$BwB = Bn_wB = Un_wB.$$

Note that $Un_wB = Un_{w'}B$ if and only if $w = w'$ in W. Then the following disjoint union

$$G = \bigsqcup_{w \in W} Un_wB$$

is also called the **Bruhat decomposition** of G. We have $\dim_{\mathbb{C}} Un_wB = \ell(w) + \dim_{\mathbb{C}} B$ for all $w \in W$. If $B^- = HU^-$ is the opposite Borel subgroup, then the multiplication map $U^- \times B \to G$ is an isomorphism of $U^- \times B$ onto the dense open subset U^-B of G. We call U^-B (and anything like it) the **big cell** in G.

The Bruhat decomposition gives a cellular decomposition of the flag variety G/B. Let $\pi \colon G \to G/B$ be the canonical map, and set $X_w = \pi(Un_wB)$. Then

$$G/B = \bigsqcup_{w \in W} X_w$$

and each X_w is an affine variety (called a **Bruhat cell**) isomorphic to $\mathbb{A}^{\ell(w)}$. The Bruhat cell X_{w_0} corresponding to the longest element $w_0 \in W$ is called the **big cell** in G/B. It is the image of a big cell in G under π.

A similar description is also available for generalized flag varieties G/P, though we shall not require it. One thing we conclude from this is that the odd betti numbers of G/B (and also of G/P) are all zero. In particular, (generalized) flag varieties are simply connected.

1.13 Compact Real Forms

Let G be a complex semisimple Lie group. By a **compact real form** of G we mean a compact Lie subgroup K with Lie algebra $\mathfrak{k}$ such that $\mathfrak{g} = \mathfrak{k} \oplus i\mathfrak{k}$. Such subgroups exist, are automatically connected and maximal amongst the compact subgroups of G, and any two of them are conjugate. Conversely, if K is a compact semisimple Lie group, then there is a unique (up to isomorphism) complex semisimple Lie group G which contains K as a maximal compact subgroup. Then K is also a real form of G, in the sense that we have a splitting $\mathfrak{g} = \mathfrak{k} \oplus i\mathfrak{k}$.

Let K and G be as above, and fix a maximal (analytic) torus $T \subset K$ (this is isomorphic to a direct product of circle groups $\mathrm{SO}(2)$, and not $\mathrm{GL}(1, \mathbb{C})$'s). Then the complexified Lie algebra $\mathfrak{h} = \mathfrak{t} \oplus i\mathfrak{t}$ of T is a Cartan subalgebra of $\mathfrak{g}$. Let H denote the corresponding maximal (algebraic) torus in G, and fix a Borel subgroup B containing H. Then $G = KB$ and $B \cap H = T$, whence we obtain another description of the flag variety as

$$G/B = K/T.$$

We state one last important fact concerning the two groups K and G. First, it is clear that holomorphic representations of G restrict to yield smooth representations of K. In the opposite direction, it is possible to analytically continue smooth representations of K to holomorphic representations of G. These processes of restriction and analytic continuation are inverses of each other and take irreducible to irreducibles.

Example 1.13.1. If $G = \mathrm{SL}(n, \mathbb{C})$ then we can take $K = \mathrm{SU}(n)$. ▲

1.14 The Infinitesimal Character

Let G be a complex semisimple Lie group; in this section we are primarily interested in the Lie algebra $\mathfrak{g}$ of G. Let $\mathcal{U}(\mathfrak{g})$ and $\mathcal{Z}(\mathfrak{g})$ denote, respectively, the universal enveloping algebra of $\mathfrak{g}$ and its centre. By applying the Poincaré–Birkhoff–Witt theorem to the triangular decomposition $\mathfrak{g} = \mathfrak{u}^- \oplus \mathfrak{h} \oplus \mathfrak{u}$, we find that

$$\begin{aligned} \mathcal{U}(\mathfrak{g}) &= \mathcal{U}(\mathfrak{u}^-) \otimes \mathcal{U}(\mathfrak{h}) \otimes \mathcal{U}(\mathfrak{u}) \\ &= \mathcal{U}(\mathfrak{h}) \oplus (\mathfrak{u}^- \mathcal{U}(\mathfrak{g}) + \mathcal{U}(\mathfrak{g})\mathfrak{u}). \end{aligned}$$

The **auxiliary Harish-Chandra map** γ is the projection of $\mathcal{Z}(\mathfrak{g}) \subset \mathcal{U}(\mathfrak{g})$ onto the first factor $\mathcal{U}(\mathfrak{h}) = \mathbb{Z}[\mathfrak{h}]$ above. Then one can show that

$$z - \gamma(z) \in \mathcal{U}(\mathfrak{g})\mathfrak{u} \quad \text{for all } z \in \mathcal{Z}(\mathfrak{g}).$$

It follows that γ is an algebra homomorphism. Indeed, γ is linear by definition, and we observe that, for $z, z' \in \mathcal{Z}(\mathfrak{g})$,

$$zz' - \gamma(z)\gamma(z') = z(z' - \gamma(z')) + \gamma(z')(z - \gamma(z))$$

belongs to $\mathcal{U}(\mathfrak{g})\mathfrak{u}$, whence $\gamma(zz') = \gamma(z)\gamma(z')$.

Let V be a (finite-dimensional) irreducible $\mathfrak{g}$-module. Then $\mathcal{Z}(\mathfrak{g})$ acts on V by scalars: there is a homomorphism $\chi_V \colon Z(\mathfrak{g}) \to \mathbb{C}$, called the **infinitesimal character** of V, such that $zv = \chi_V(z)v$ for all $z \in \mathcal{Z}(\mathfrak{g})$ and all $v \in V$. If λ is the highest weight of V, then it is one of the main properties of the auxiliary Harish-Chandra map that $\chi_V(z) = \lambda(\gamma(z))$ for all $z \in \mathcal{Z}(\mathfrak{g})$.

For general $\lambda \in \mathfrak{h}^*$ one defines a homomorphism $\chi_\lambda \colon \mathcal{Z}(\mathfrak{g}) \to \mathbb{C}$ by setting

$$\chi_\lambda(z) = (\lambda - \rho)(\gamma(z)),$$

where $\rho = \frac{1}{2}\sum_{\alpha>0} \alpha \in E$. Then one proves that $\chi_\lambda = \chi_\mu$ (for $\lambda, \mu \in \mathfrak{h}^*$) if and only if $\lambda = w\mu$ for some $w \in W$.

If λ is the highest weight of the irreducible $\mathfrak{g}$-module V, then the infinitesimal character χ_V of V is equal to $\chi_{\lambda+\rho}$. It thus might seem that this ρ-shift adds an unnecessary complication. Nonetheless, we have introduced it because it is standard and, ultimately, more convenient.

Chapter 2

The Borel–Weil–Bott Theorem

2.1 Introduction

In this chapter we present the celebrated Borel–Weil theorem (Theorem 2.3.1). This theorem provides a geometric realization of all the irreducible representations of a complex semisimple Lie group G in terms of spaces of sections of equivariant line bundles over the flag variety of G. Proved in 1954 by Borel and Weil (see also Remark 2.3.2(i)), it may be regarded as the first theorem of geometric representation theory proper, a subject which has since seen—and still sees—a rich and beautiful development.

In this vein, we give Bott's equally celebrated generalization of the Borel–Weil theorem (Theorems 2.4.1 and 2.4.6). Bott's theorem computes the cohomology of certain equivariant vector bundles over the generalized flag variety G/P in terms of the representation theory of G. It was, for a time, the only method by which a description of these cohomology spaces could be obtained (see Section 2.5.1).

Two further applications are also given: the local rigidity of the complex structure on generalized flag varieties is established (Section 2.5.2), and Weyl's formula for the dimension of an arbitrary irreducible G-module is deduced with the aid of the Hirzebruch–Riemann–Roch formula (Section 2.5.3).

Finally, in Section 2.6, we say a few words about the issue of generalizing the main theorems of this chapter to the setting of semisimple algebraic groups over an arbitrary algebraically closed field k.

2.2 Equivariant Vector Bundles

Let G be a complex Lie group, H a closed subgroup of G, and $X = G/H$ the associated homogeneous space. We wish to exhibit a one-to-one correspondence between (holomorphic) representations of H and G-equivariant vector bundles over G/H.[1] Recall that a **G-equivariant vector bundle** over a G-space Y is a holomorphic vector bundle $\pi\colon E \to Y$ with a G-action on the total space E that is compatible with the G-action on the base space Y, i.e.

- the projection map is G-equivariant: $\pi(gv) = g\pi(v)$ for all $g \in G$ and $v \in E$;
- translation between fibres is linear: for each $g \in G$ and $y \in Y$, the map $E_y \to E_{gy}$ given by $v \mapsto gv$ is linear (and hence is an isomorphism of vector spaces).

The fact that the map $g \mapsto gv$ takes E_y to E_{gy} is a consequence of the G-equivariance of π. The inverse map is of course given by the action of g^{-1}, viz. $E_{gy} \to E_{g^{-1}(gy)} = E_y$.

Given a G-equivariant vector bundle E over G/H, it follows from the above definition that the fibre E_H at the point $H \in G/H$ is a representation of H. Conversely, given a representation $\varphi\colon H \to \mathrm{GL}(V)$ of H, consider the fibre product

$$G \times_\varphi V = (G \times V)/\sim,$$

where $\sim$ is the equivalence relation defined by $(g, v) \sim (gh^{-1}, \varphi(h)v)$ for $g \in G, h \in H$ and $v \in V$. With the quotient topology (and structure sheaf), $G \times_\varphi V$ becomes a complex manifold. Note that G acts on $G \times_\varphi V$ by left-multiplication of the first factor.

Proposition 2.2.1. *The G-space $G \times_\varphi V$ is a G-equivariant vector bundle (with fibre isomorphic to V) over G/H. The projection map $\pi\colon G \times_\varphi V \to G/H$ is given by $\pi([g, v]) = gH$.*

[1]Throughout this chapter, unless stated otherwise, we shall assume that our vector bundles are finite-dimensional, complex and holomorphic.

Proof. It is clear that π is well-defined. Let $q\colon G \to G/H$ denote the quotient map and fix $gH \in G/H$. Then, as a consequence of the Implicit Function Theorem, we can find an open neighborhood U of gH in G/H and a holomorphic map $\alpha\colon U \to G$ such that $q \circ \alpha = \mathrm{id}$ and the map $U \times H \to q^{-1}(U)$ defined by $(u,h) \to \alpha(u)h$ is biholomorphic [67, Lemma 14.3.3]. This induces a biholomorphic map $U \times V = U \times H \times_\varphi V \to \pi^{-1}(U)$ that is fibre-wise linear. Thus $\pi\colon G \times_\varphi V \to G/H$ is a holomorphic vector bundle; it is obviously G-equivariant. Finally, the fibre at the point gH is simply $\{[gh,v]\colon h \in H\} = \{[g,\varphi(h)v]\colon h \in H\} = \{[g,v]\colon v \in V\}$, which is isomorphic to V. ■

Observe that by restricting the action of H on $G \times_\varphi V$ to the fibre $\{[1,v]\colon v \in V\}$ lying over $H \in G/H$ we recover the original representation φ. We have thus established our desired bijection:

$$\begin{array}{ccc} \{\text{holomorphic representations of } H\} & \longleftrightarrow & \{G\text{-equivariant vector bundles over } G/H\} \\ \varphi\colon H \to \mathrm{GL}(V) & \longmapsto & G \times_\varphi V \to G/H \end{array}$$

This bijection is in fact an equivalence of categories.

Our goal now is to use this correspondence to induce representations from H to G. We will do this by exploiting the G-action on the equivariant vector bundle corresponding to the given representation of H. A G-action will often produce a representation, provided we can find a vector space (usually a function space) associated to the action. In the present setting there is at least one obvious candidate, namely the space of sections. Let $\varphi\colon H \to \mathrm{GL}(V)$ be a representation of H and let $E \to G/H$ denote the corresponding G-equivariant vector bundle (namely $E = G \times_\varphi V$). Then G acts linearly on the vector space $\Gamma(E)$ of sections of E as follows:

$$(gs)(x) = g\,s(g^{-1}x), \qquad \text{where } g \in G, s \in \Gamma(E) \text{ and } x \in G/H. \tag{2.2.1}$$

More generally, let $\mathcal{V}_\varphi$ denote the sheaf of holomorphic sections of E and consider the sheaf cohomology spaces $H^q(G/H, \mathcal{V}_\varphi)$ $(q \geq 0)$. Then $H^0(G/H, \mathcal{V}_\varphi) = \Gamma(E)$ and, by naturality, G acts linearly on every $H^q(G/H, \mathcal{V}_\varphi)$ because the sheaf $\mathcal{V}_\varphi$ is G-equivariant. One way to explicitly realize this action is to first define it on the Čech cocycle space $C^q(\mathcal{U}, \mathcal{V}_\varphi)$

corresponding to an open cover $\mathcal{U}$ of G/H, and then argue carefully as to how this action can be carried over to $H^q(G/H, \mathcal{V}_\varphi)$. This is done in full detail in [1, §4.1].

Although this action of G does produce a homomorphism $G \to \mathrm{GL}(H^q(G/H, \mathcal{V}_\varphi))$, to call this homomorphism a "representation" of G (as per our definition), we need to know that

- $H^q(G/H, \mathcal{V}_\varphi)$ is finite-dimensional, and that
- the homomorphism $G \to \mathrm{GL}(H^q(G/H, \mathcal{V}_\varphi))$ is holomorphic.

One way to deal with the finite-dimensionality issue is to recall the following (special case of a) theorem of Cartan and Serre. (See [67, Theorem 11.10.2].)

Theorem 2.2.2 (Cartan–Serre). *If $\mathcal{V}$ is a holomorphic vector bundle over a compact complex manifold X, then $H^q(X, \mathcal{V})$ is finite-dimensional for all $q \geq 0$.* ■

Thus, if G/H is compact, the Cartan–Serre theorem will ensure that $H^q(G/H, \mathcal{V}_\varphi)$ is finite-dimensional. It turns out that the compactness of G/H also guarantees the holomorphicity of the representation of G on $H^q(G/H, \mathcal{V}_\varphi)$. Indeed, the action of G on $H^q(G/H, \mathcal{V}_\varphi)$ is built up from translations, which in particular implies that the resulting homomorphism $G \to \mathrm{GL}(H^q(G/H, \mathcal{V}_\varphi))$ is continuous, hence holomorphic, being a map between complex Lie groups.

To summarize:

Proposition 2.2.3. *If G/H is compact, then the homomorphism $G \to \mathrm{GL}(H^q(G/H, \mathcal{V}_\varphi))$ obtained above is a holomorphic representation of G.* ■

Remark 2.2.4. There remains the potential issue of having $H^q(G/H, \mathcal{V}_\varphi) = 0$ (and in fact this turns out to be the case more often than not!), in which case it is not a "representation" according to our definition. This is only a technical nuisance, and therefore it should be implicitly understood that whenever we do speak of a representation on $H^q(G/H, \mathcal{V}_\varphi)$, we are assuming that $H^q(G/H, \mathcal{V}_\varphi)$ is nonzero. ▲

The final result of this section is a description of $H^0(G/H, \mathcal{V}_\varphi)$ as a subspace of $\mathrm{Hol}(G, V)$, the vector space of holomorphic functions from G to V. Let $s\colon G/H \to \mathcal{V}_\varphi$ be a holomorphic section in $H^0(G/H, \mathcal{V}_\varphi)$. Then there is a function $f\colon G \to V$ such that $s(gH) = [g, f(g)]$ for all $g \in G$. This function f must be holomorphic because s is. Additionally, as $[g, f(g)] = [gh, \varphi(h)^{-1} f(g)]$ for $g \in G$ and $h \in H$, the function f must satisfy

$$f(gh) = \varphi(h)^{-1} f(g) \qquad \text{for all } g \in G \text{ and } h \in H. \tag{2.2.2}$$

Conversely, if $f \in \mathrm{Hol}(G, V)$ satisfies (2.2.2), then the map $s\colon G/H \to \mathcal{V}_\varphi$ defined by $s(gH) = [g, f(g)]$ is easily seen to be a holomorphic section of $\mathcal{V}_\varphi$. To summarize:

Proposition 2.2.5. *We have an identification of representations:*

$$H^0(G/H, \mathcal{V}_\varphi) = \{f \in \mathrm{Hol}(G, V)\colon f(gh) = \varphi(h)^{-1} f(g) \textit{ for all } g \in G \textit{ and } h \in H\},$$

where G acts on the space on the left as in (2.2.1) *and on the space on the right via the left regular representation, viz.* $(gf)(g') = f(g^{-1}g')$ *for all* $g, g' \in G$.

Proof. All that remains is to prove that the actions are compatible, but this is trivial: if $s(g'H) = [g', f(g')]$, then $(gs)(g'H) = gs(g^{-1}g'H) = g[g^{-1}g', f(g^{-1}g')] = [g', f(g^{-1}g')] = [g', (gf)(g')]$. ■

Remarks 2.2.6.

(i) Suppose that V is a one-dimensional representation of H, in which case the induced vector bundle $\mathcal{L}$ on G/H is a line bundle. Using the above description of $H^0(G/H, \mathcal{L})$, we can give a simple proof (due to Serre [63, p.70]) that this space is finite-dimensional when G/H is compact. We begin by endowing $H^0(G/H, \mathcal{L})$ with the topology of uniform convergence on compact subsets, making it a Banach space. Recall that a uniformly bounded sequence of holomorphic functions has a subsequence that converges uniformly on compact subsets (by Montel's theorem). Thus $H^0(G/H, \mathcal{L})$ is a *locally compact* Banach space, hence is finite-dimensional. Observe also that the left regular representation of G on $H^0(G/H, \mathcal{L})$ is obviously holomorphic.

(ii) Suppose that G is a finite group. A representation $\varphi\colon H \to \mathrm{GL}(V)$ of a subgroup H, thought of as the $\mathbb{C}[H]$-module V, gives rise to a $\mathbb{C}[G]$-module $\mathbb{C}[G] \otimes_{\mathbb{C}[H]} V$. This latter module is regarded as being *induced* from the former, and as such is typically denoted by $\mathrm{ind}_H^G V$ or $\mathrm{ind}_H^G \varphi$. It may be identified, as a representation of G, with the function space $\{f\colon G \to \mathbb{C}\colon f(gh) = \varphi(h)^{-1} f(g)\}$ on which G acts via the left regular representation. Returning to complex Lie groups and holomorphic representations, we see that the process of forming the representation $H^0(G/H, \mathcal{V}_\varphi)$ of G may be thought of as the analogue of induction in this setting. For this reason, the space $H^0(G/H, \mathcal{V}_\varphi)$ is often denoted by $\mathrm{ind}_H^G V$ or $\mathrm{ind}_H^G \varphi$ in the literature. It is also common to see $H^q(G/H, \mathcal{V}_\varphi)$ referred to as $\mathrm{R}^q\mathrm{ind}_H^G V$ or $\mathrm{R}^q\mathrm{ind}_H^G \varphi$, the qth right derived functor of ind_H^G. ▲

In the next section, where we meet the Borel–Weil theorem, the objects of interest are one-dimensional representations of certain subgroups H and the resulting spaces of sections. The first remark above then serves to show that, insofar as the Borel–Weil theorem is concerned, we can avoid the technical issues associated with the action of G on the higher cohomology groups $H^q(X, \mathcal{V})$ $(q > 0)$.

2.3 The Borel–Weil Theorem

Assume now that G is semisimple. In light of the results of the previous section, we would like the space G/H to be compact. This is the case if and only if H is parabolic. So let P be a parabolic subgroup of G and consider the generalized flag variety G/P. To each representation $\varphi\colon P \to \mathrm{GL}(V)$ of P we associate the sheaf $\mathcal{V}_\varphi$ on G/P as before. Our goal is to determine when $H^q(G/P, \mathcal{V}_\varphi) \neq 0$; in which case we would also like to understand the induced representation $G \to \mathrm{GL}(H^q(G/P, \mathcal{V}_\varphi))$. This goal is realized in the Borel–Weil–Bott theorem (for irreducible φ). In this section we will describe part of this theorem.

We will be working with the flag variety $X = G/B^-$, where B^- is some fixed Borel subgroup of G. The choice of B^- is, of course, immaterial. However, once a choice is made, it corresponds to a choice of positive roots at the Lie algebra level. Let U^- denote the unipotent radical of B^-, so that $B^- = HU^-$ where H is a maximal (algebraic) torus. The minus signs designate that the Lie algebras of U^- and B^- are $\mathfrak{u}^- = \oplus_{\alpha<0}\mathfrak{g}_\alpha$ and $\mathfrak{b}^- = \mathfrak{h} \oplus \mathfrak{u}^-$, respectively, where $\mathfrak{h}$ is the Lie algebra of H (a Cartan subalgebra of $\mathfrak{g}$). Let $B = HU$ denote the Borel subgroup opposite B^-. Then the Lie algebras of U and B are $\mathfrak{u} = \oplus_{\alpha>0}\mathfrak{g}_\alpha$ and $\mathfrak{b} = \mathfrak{h} \oplus \mathfrak{u}$, respectively.

Let Λ denote the weight lattice in $\mathfrak{h}^*$, which we may identify with the character group $\widehat{H}$ of H (see Section 1.8.1). Because $U^- = [B^-, B^-]$, each character $\lambda \in \widehat{H}$ of H can be extended to a character of B^- by setting $\lambda(u) = 1$ for $u \in U^-$; conversely, every character of B^- restricts to a character of H. Thus we can identify $\widehat{H} = \Lambda$ with the character group of B^-, which, because B^- is solvable, is the set of irreducible holomorphic representations of B^-. Finally, let $\mathcal{L}_\lambda = G \times_\lambda \mathbb{C}$ denote the line bundle (or, where appropriate, its sheaf of holomorphic sections) over $X = G/B^-$ corresponding to the character $\lambda \in \widehat{H}$.

Theorem 2.3.1 (Borel–Weil)**.** *Assume the preceding notation and let $\lambda \in \Lambda$. If λ is dominant, then $H^0(X, \mathcal{L}_\lambda)$ is a (nonzero) irreducible representation of G of highest weight λ. If λ is not dominant, then $H^0(X, \mathcal{L}_\lambda) = 0$.*

Remarks 2.3.2.

(i) Borel and Weil proved this theorem in 1954; the details were written up by Serre and appear in [62]. Similar results were obtained around the same time by Harish-Chandra [34] and Tits [68].

(ii) If we had used the "positive" Borel subgroup B (instead of B^-), then the statement of the theorem would have been: If λ is dominant, then $H^0(X, \mathcal{L}_{-\lambda})$ is an irreducible representation of G of *lowest* weight $-\lambda$. Otherwise $H^0(X, \mathcal{L}_\lambda) = 0$. ▲

We will give two proofs of the Borel–Weil theorem in this section: an analytic one (relying on the Peter-Weyl theorem) and an algebraic one (relying on the Bruhat decomposition). But before doing so, let us take a look at a simple, but important, example.

Example 2.3.3. Let $G = \mathrm{SL}(2, \mathbb{C})$, $B = \{\left(\begin{smallmatrix} * & * \\ 0 & * \end{smallmatrix}\right)\}$, $B^- = \{\left(\begin{smallmatrix} * & 0 \\ * & * \end{smallmatrix}\right)\}$ and $X = G/B^- = \mathbb{P}^1$. For $n \in \mathbb{Z}$, let $\lambda_n \colon B^- \to \mathbb{C}^\times$ be the character defined by $\lambda_n(\left(\begin{smallmatrix} a & 0 \\ b & a^{-1} \end{smallmatrix}\right)) = a^n$. These are all the characters of B^-. Let $\mathcal{O}(n)$ denote the line bundle corresponding to λ_n. (The line bundle $G \times_{\lambda_n} \mathbb{C} \to \mathbb{P}^1$ obviously has degree n, so our usage of $\mathcal{O}(n)$ coincides with its usage in algebraic geometry.)

We wish to determine (more explicitly) the space

$$H^0(X, \mathcal{O}(n)) = \{f \in \mathrm{Hol}(G, \mathbb{C})\colon f(gb) = \lambda_n(b)^{-1} f(g) \text{ for } b \in B^- \text{ and } g \in G\}.$$

Using the Bruhat decomposition $G = B^- W B^- = B^- \sqcup B^- \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right) B^-$, we find that

$$H^0(X, \mathcal{O}(n)) \cong \{f \in \mathrm{Hol}(\mathbb{C}^2 - \{0\}, \mathbb{C})\colon f(tz) = t^n f(z) \text{ where } t \in \mathbb{C}^* \text{ and } z \in \mathbb{C}^2\}.$$

By Hartog's theorem [28, p.7], any $f \in H^0(X, \mathcal{O}(n))$ can be extended to a holomorphic function on all of $\mathbb{C}^2$ which will be homogeneous of degree n, i.e. $f(\lambda z) = \lambda^n z$ for $\lambda \in \mathbb{C}$ and $z \in \mathbb{C}^2$. This implies that there are no nontrivial sections of $\mathcal{O}(n)$ when $n < 0$; in other words,

$$H^0(X, \mathcal{O}(n)) = 0 \qquad \text{if } n < 0.$$

Now assume that $n \geq 0$. Then, by a power series expansion about $(0, 0)$, we see that a nontrivial $f \in H^0(X, \mathcal{O}(n))$ is a homogeneous *polynomial* of degree n; conversely, every such homogeneous polynomial gives rise to an element of $H^0(X, \mathcal{O}(n))$. Two polynomials that agree on $\mathbb{C}^2 - \{0\}$ must agree everywhere. If we let $\mathbb{C}[z_1, z_2]_n$ denote the space of degree n homogeneous polynomials in $\mathbb{C}[z_1, z_2]$, we see that

$$H^0(X, \mathcal{O}(n)) = \mathbb{C}[z_1, z_2]_n \qquad \text{if } n \geq 0.$$

By Proposition 2.2.5, the action of G is given by the left regular representation:

$$(gf)(z) = f(g^{-1} z) \qquad \text{for all } g \in G \text{ and } f \in \mathbb{C}[z_1, z_2]_n.$$

Anyone familiar with the representation theory of $\mathrm{SL}(2, \mathbb{C})$ will recognize that this representation is *irreducible* and that, moreover, every irreducible representation of $\mathrm{SL}(2, \mathbb{C})$ can be realized in this fashion, for an appropriate $n \geq 0$. Our findings are compatible with the Borel–Weil theorem. Indeed, λ_n is dominant if and only if $n \geq 0$. ▲

For the remainder of this section, let $\lambda \in \widehat{H}$ and set $H^0(\lambda) = H^0(X, \mathcal{L}_\lambda)$. According to Proposition 2.2.5,

$$H^0(\lambda) = \{f \in \mathrm{Hol}(G, \mathbb{C})\colon f(gb) = \lambda(b)^{-1} f(g) \text{ for all } b \in B^- \text{ and } g \in G\} \tag{2.3.1}$$

with G acting via the left regular representation.

First proof of the Borel–Weil theorem. Let $K \subset G$ be a compact real form of G. Then $T = B^- \cap K = H \cap K$ is a maximal torus in K and we may identify $X = G/B^-$ with K/T (see Section 1.13). Then each (smooth) character of T may be extended to a (holomorphic) character of B^- and, conversely, we may restrict characters from B^- to T. Consequently, we may view $\mathcal{L}_\lambda$ as either a smooth K-equivariant line bundle over K/T (given by $K \times_\lambda \mathbb{C}$, with λ restricted to T) or a holomorphic G-equivariant line bundle over G/B (given by $G \times_\lambda \mathbb{C}$, as before). Then, by mimicking the proof of Proposition 2.2.5, the space of global smooth sections of $\mathcal{L}_\lambda$ is

$$C^\infty(K/T, \mathcal{L}_\lambda) = \{f \in C^\infty(K)\colon f(kt) = \lambda(t)^{-1} f(k) \text{ for all } t \in T \text{ and } k \in K\}.$$

While K does act on this space via the left regular representation, $C^\infty(X, \mathcal{L}_\lambda)$ is not necessarily finite-dimensional, and hence will not in general be a representation of K (as per our definition of "representation"). At any rate, notice that any $f \in C^\infty(X, \mathcal{L}_\lambda)$ may be extended to a *smooth* section $\tilde{f}$ of the line bundle $\mathcal{L}_\lambda \to G/B^-$ by setting $\tilde{f}(kb) = \lambda(b)^{-1} f(k)$ for $kb \in KB^- = G$, and in the reverse direction, any smooth section of $\mathcal{L}_\lambda \to G/B^-$ restricts to one of $\mathcal{L}_\lambda \to K/T$. In this way we obtain two K-equivariant maps,

$$C^\infty(K/T, \mathcal{L}_\lambda) \to C^\infty(G/B^-, \mathcal{L}_\lambda) \quad \text{and} \quad C^\infty(G/B^-, \mathcal{L}_\lambda) \to C^\infty(K/T, \mathcal{L}_\lambda),$$

which are inverses to each other, as the reader may easily verify. Thus we can identify $H^0(\lambda)$, as a representation of K, with the subspace of $C^\infty(K/T, \mathcal{L}_\lambda)$ consisting of sections f whose extensions $\tilde{f}$ are holomorphic. The following lemma gives a helpful characterization of such sections.

Lemma 2.3.4. *Let $f \in C^\infty(K/T, \mathcal{L}_\lambda)$. Then $\tilde{f} \in C^\infty(G/B, \mathcal{L}_\lambda)$ is holomorphic if and only if $r_{\mathbb{C}}(Z)f = 0$ for all $Z \in \mathfrak{u}^- = \oplus_{\alpha<0}\mathfrak{g}_\alpha$, where $r_{\mathbb{C}}(Z)$ denotes infinitesimal right translation on K by Z.*

Recall the definition of $r_{\mathbb{C}}(Z)$: If $f \in C^\infty(K)$ and $X \in \mathfrak{k}$, then

$$(r(X)f)(k) = \frac{d}{dt} f(ke^{tX})|_{t=0}.$$

If $Z = X + iY \in \mathfrak{g} = \mathfrak{k} \oplus i\mathfrak{k}$, we let $r_{\mathbb{C}}(Z) = r(X) + ir(Y)$. We also define $r(Z)$ on $C^\infty(G)$ by

$$(r(Z)F)(g) = \frac{d}{dt} F(ge^{tZ})|_{t=0} \qquad (F \in C^\infty(G)).$$

Note that if $Z \in \mathfrak{k}$ then the two definitions of $r(Z)$, the first as an operator on $C^\infty(K)$ and the second as an operator on $C^\infty(G)$, are compatible.

Proof of Lemma 2.3.4. A function $F \in C^\infty(G)$ is holomorphic if and only if it satisfies the Cauchy–Riemann equations:

$$dF(iZ) = idF(Z) \qquad \text{for all } Z \in T_g(G) \text{ and } g \in G. \tag{2.3.2}$$

Let $l_g\colon G \to G$ denote left translation by g. Then $dl_g\colon \mathfrak{g} = T_e(G) \to T_g(G)$ is an isomorphism for all g, whence F satisfies (2.3.2) if and only if

$$dF(dl_g(iZ)) = idF(dl_g Z) \qquad \text{for all } Z \in \mathfrak{g} \text{ and } g \in G. \tag{2.3.3}$$

Note that

$$dF(dl_g(iZ)) = \frac{d}{dt} F(ge^{tiZ})|_{t=0} = (r(iZ)F)(g).$$

Similarly, $dF(l_g Z) = (r(Z)F)(g)$. We conclude that F is holomorphic if and only if

$$r(iZ)F = ir(Z)F \qquad \text{for all } Z \in \mathfrak{g}. \tag{2.3.4}$$

Now let f and $\tilde{f}$ be as in the statement of the lemma. Then (2.3.4) tells us that $\tilde{f}$ is holomorphic if and only if $r(iZ)\tilde{f} = ir(Z)\tilde{f}$ for all $Z \in \mathfrak{g}$. If Z is in $\mathfrak{b}^-$, then $e^{tZ} \in B^-$ so that $\tilde{f}(ge^{tZ}) = e^{-t\lambda(Z)}\tilde{f}(g)$. Consequently (2.3.4) holds for all such Z, since in this case both sides of the equation are $-i\lambda(Z)\tilde{f}(g)$. As $\mathfrak{g} = \mathfrak{b}^- \oplus \mathfrak{u}$, we conclude that $\tilde{f}$ is holomorphic if and only if (2.3.4) holds for all $Z \in \mathfrak{u}$. Let $Z \in \mathfrak{u}$. Then $\overline{Z} \in \mathfrak{u}^-$, where the bar is complex conjugation with respect to the decomposition $\mathfrak{g} = \mathfrak{k} \oplus i\mathfrak{k}$. Now, $r(iZ)\tilde{f} = ir(Z)\tilde{f} \iff r_{\mathbb{C}}(Z)f = r(Z)f$. But

$$r_{\mathbb{C}}(Z)\tilde{f} = r_{\mathbb{C}}(Z + \overline{Z})\tilde{f} - r_{\mathbb{C}}(\overline{Z})\tilde{f}$$

and, because $Z + \overline{Z} \in \mathfrak{k}$,

$$r(Z)\tilde{f} = r(Z + \overline{Z})\tilde{f} - r(\overline{Z})\tilde{f} = r(Z + \overline{Z})\tilde{f} = r_{\mathbb{C}}(Z + \overline{Z})\tilde{f}.$$

So $r_{\mathbb{C}}(Z)\tilde{f} = r(Z)f \iff r_{\mathbb{C}}(\overline{Z})\tilde{f} = 0$. It follows that $\tilde{f}$ is holomorphic if and only if $r_{\mathbb{C}}(\mathfrak{u}^-)\tilde{f} = 0$.

In particular, if $\tilde{f}$ is holomorphic, then as $f = \tilde{f}|_K$, we find that $r_{\mathbb{C}}(\mathfrak{u}^-)f = 0$. Conversely, suppose that $r_{\mathbb{C}}(\mathfrak{u}^-)f = 0$. We wish to show that $r_{\mathbb{C}}(\mathfrak{u}^-)\tilde{f} = 0$. We begin by noting that $r(Z)\tilde{f}(k) = r_{\mathbb{C}}(Z)\tilde{f}(k)$ for all $Z \in \mathfrak{g}$ and $k \in K$, by the reasoning in the preceding paragraph. Then, for $X \in \mathfrak{k}$, $k \in K$ and $b \in B^-$,

$$\begin{aligned}
(r(X)\tilde{f})(kb) &= \frac{d}{dt}\tilde{f}(kbe^{tX})|_{t=0} \\
&= \frac{d}{dt}\tilde{f}(ke^{t\mathrm{Ad}(b)X}b)|_{t=0} \\
&= \lambda(b)^{-1}\frac{d}{dt}\tilde{f}(ke^{t\mathrm{Ad}(b)}X)|_{t=0} \\
&= \lambda(b)^{-1}(r(\mathrm{Ad}(b)X)\tilde{f})(k) \\
&= \lambda(b)^{-1}(r_{\mathbb{C}}(\mathrm{Ad}(b)X)\tilde{f})(k).
\end{aligned}$$

So if $Z = X + iY \in \mathfrak{u}^-$, with $X, Y \in \mathfrak{k}$, then $\mathrm{Ad}(b)Z \in \mathfrak{u}^-$ and

$$\begin{aligned}
(r_{\mathbb{C}}(Z)\tilde{f})(kb) &= (r(X)\tilde{f})(kb) + i(r(Y)\tilde{f})(kb) \\
&= \lambda(b)^{-1}\left[(r_{\mathbb{C}}(\mathrm{Ad}(b)X)\tilde{f})(k) + (r_{\mathbb{C}}(i\mathrm{Ad}(b)Y)\tilde{f})(k)\right] \\
&= \lambda(b)^{-1}(r_{\mathbb{C}}(\mathrm{Ad}(b)Z)\tilde{f})(k) \\
&= \lambda(b)^{-1}(r_{\mathbb{C}}(\mathrm{Ad}(b)Z)f)(k) \\
&= 0,
\end{aligned}$$

as desired. ■

It follows that

$$H^0(\lambda) \cong \{f \in C^\infty(K)\colon r(\mathfrak{u}^-)f = 0 \text{ and } f(kt) = \lambda(t)^{-1}f(k) \text{ for all } t \in T\},$$

as a representation of K.

Now $C^\infty(K) \subset L^2(K)$, and the Peter–Weyl theorem yields the Hilbert space direct sum

$$L^2(K) = \widehat{\bigoplus_i} V_i \otimes V_i^*,$$

which runs over all the irreducible (unitary) representations of K. We identify $V_i \otimes V_i^*$ with the space of matrix coefficients of V_i. This decomposition affords a $K \times K$ action, with K acting on V_i by left translation and on V_i^* by right translation. The subspace of $C^\infty(K)$ (and hence of $L^2(K)$) corresponding to $H^0(\lambda)$ is finite-dimensional and is therefore contained in the *algebraic* direct sum $\bigoplus_i V_i \otimes V_i^*$. So we can write any given $f \in H^0(\lambda)$ as a finite sum $f = u_1 \otimes v_1^* + \cdots + u_n \otimes v_n^*$ of matrix coefficients of K. The stipulations $r(\mathfrak{u}^-)f = 0$ and $f(kt) = \lambda(t)^{-1} f(k)$ (for all $t \in T$) are then seen to be equivalent to the assertions that $\mathfrak{u}^- v_i^* = 0$ and that v_i^* belongs to the space of T-invariants $(V_i^* \otimes \mathbb{C})^T$ (with T acting by right translation on V_i^* and via λ on $\mathbb{C}$) for all i, respectively. We thus find that

$$\begin{aligned} H^0(\lambda) &\cong \bigoplus_i V_i \otimes \{v \in (V_i^* \otimes \mathbb{C}_\lambda)^T : \mathfrak{u}^- v = 0\} \\ &\cong \bigoplus_i V_i \otimes ((V_i/\mathfrak{u}^- V_i)^* \otimes \mathbb{C}_\lambda)^T \\ &\cong \bigoplus_i V_i \otimes \mathrm{Hom}_T(V_i/\mathfrak{u}^- V_i, \mathbb{C}_\lambda), \end{aligned}$$

as representations of K, where the subscript in $\mathbb{C}_\lambda$ indicates that T acts on $\mathbb{C}_\lambda = \mathbb{C}$ via λ. Now observe that $V_i/\mathfrak{u}^- V_i$ is the highest weight space of V_i, so if μ_i is the highest weight of V_i, then $\mathrm{Hom}_T(V_i/\mathfrak{u}^- V_i, \mathbb{C}_\lambda) = \mathrm{Hom}_T(\mathbb{C}_{\mu_i}, \mathbb{C}_\lambda)$, which is nonzero if and only if $\mu_i = \lambda$. Consequently, if λ is not dominant, then because no μ_i can be equal to λ, it follows that $H^0(\lambda) = 0$. On the other hand, if λ is dominant, then the above direct sum decomposition reduces to

$$H^0(\lambda) \cong V_i,$$

where V_i is the irreducible representation of K of highest weight λ. The theorem now follows by analytic continuation. ■

Second proof of the Borel–Weil theorem. Assume that $H^0(\lambda) \neq 0$. By the Lie–Kolchin theorem [40, §17.6], there exists a B^--invariant line ℓ in $H^0(\lambda)$ of weight μ (say). Thus, for $f \in \ell$, we have

$$f(b^{-1}g) = \mu(b)f(g) \qquad \text{for all } b \in B^- \text{ and } g \in G.$$

In particular, if w_0 denotes the longest element in the Weyl group of G (with respect to B^-), then

$$f(bw_0b') = \mu(b)^{-1}\lambda(b')^{-1}f(w_0) \qquad \text{for all } b, b' \in B^-.$$

As $B^-w_0B^-$ is open and dense in G, we see that f is completely determined by its value at w_0. It follows that ℓ is the unique B^--invariant line in $H^0(\lambda)$, proving that $H^0(\lambda)$ is irreducible and of highest weight μ. In fact $\mu = \lambda$; to see this, note that if $t \in H \subset B^-$ and $f \in \ell$ is nonzero (which is the case if and only if $f(w_0) \neq 0$) then

$$\mu(t)^{-1}f(w_0) = f(tw_0) = f(w_0w_0^{-1}tw_0) = \lambda(w_0^{-1}tw_0)^{-1}f(w_0) = \lambda(t)^{-1}f(w_0),$$

so that $\mu = \lambda$ on H, and hence on B. In particular, λ is the highest weight of an irreducible representation, so it must be dominant.

To finish the proof, we must show that if λ is dominant then $H^0(\lambda) \neq 0$. This amounts to constructing a nonzero holomorphic function $f_\lambda \colon G \to \mathbb{C}$ such that

$$f_\lambda(gtu) = \lambda(t)^{-1}f(g) \qquad \text{for all } g \in G,\ t \in H,\ u \in U^-.$$

We shall do this by exploiting the Bruhat decomposition $G = \sqcup_w BwB^-$ of G. The idea is to define f_λ on enough "large" cells to ensure the existence of a holomorphic extension to all of G via Hartog's theorem. We start by noting that the only cell of dimension equal to $\dim_{\mathbb{C}} G$ is $Bw_0B^- = w_0UB^-$, and that the only ones of codimension one are those of the form $Bw_0s_\alpha B^-$, where α is a simple root.[2] Each $Bw_0s_\alpha B^-$ is a closed subvariety of $w_0s_\alpha UB^-$, so that the complement of $w_0UB^- \cup \bigcup_{\alpha\in\Delta} w_0s_\alpha UB^-$ has codimension ≥ 2 in G; by multiplying through by w_0, we see that the same is true of the complement of $UB^- \cup \bigcup_{\alpha\in\Delta} s_\alpha UB^-$. Thus Hartog's theorem tells us that any holomorphic function defined on $UB^- \cup \bigcup_{\alpha\in\Delta} s_\alpha UB^-$ admits a unique holomorphic extension to G.

We begin our construction of f_λ by defining it on the big cell $UB^- = UHU^-$ by

$$f_\lambda(u_1tu_2) = \lambda(t)^{-1} \qquad (u_1 \in U,\ t \in H,\ u_2 \in U^-). \tag{2.3.5}$$

[2] Recall that the codimension of BwB^- is given by $\dim_{\mathbb{C}} G/B - \ell(w)$, where $\ell(w)$ is the length of w.

Since the multiplication map $U \times B^- \to UB^-$ is biholomorphic, we see that f_λ is holomorphic on UB^-. Next, fix a simple root $\alpha \in \Delta$. We want to extend the domain of f_λ to include $s_\alpha UB^-$. To do so, we must first recall some basic facts.

Let $\beta \in \Phi$ be an arbitrary root of G and let $x_\beta \colon \mathfrak{g}_\beta \to G$ be the exponential map restricted to $\mathfrak{g}_\beta$. Note that

$$tx_\beta(a)t^{-1} = x_\beta(\beta(t)a) \qquad \text{for all } t \in T \text{ and } a \in \mathfrak{g}_\beta. \tag{2.3.6}$$

The image of x_β is a closed subgroup of G, which we denote by U_β and call the *root subgroup* corresponding to β. The inclusion $\mathfrak{g}_\beta \oplus \mathfrak{g}_{-\beta} \oplus \mathbb{C}\beta^\vee \cong \mathfrak{sl}(2,\mathbb{C}) \subseteq \mathfrak{g}$ gives rise to a holomorphic homomorphism $\varphi_\beta \colon \mathrm{SL}(2,\mathbb{C}) \to G$ with the property that

$$\varphi_\beta \begin{pmatrix} 1 & a \\ 0 & 1 \end{pmatrix} = x_\beta(a) \quad \text{and} \quad \varphi_\beta \begin{pmatrix} 1 & 0 \\ a & 1 \end{pmatrix} = x_{-\beta}(a)$$

for all $a \in \mathbb{C}$. We then obtain the following elements of $N_G(T)$:

$$n_\beta(a) = x_\beta(a)x_{-\beta}(-a^{-1})x_\beta(a) = \varphi_\beta \begin{pmatrix} 0 & a \\ -a^{-1} & 0 \end{pmatrix} \qquad (a \in \mathbb{C}).$$

It is immediate that

$$\beta^\vee(a) = n_\beta(a)n_\beta(1)^{-1} = \varphi_\beta \begin{pmatrix} a & 0 \\ 0 & a^{-1} \end{pmatrix} \qquad (a \in \mathbb{C}^\times), \tag{2.3.7}$$

and that the image of $n_\beta(1)$ in $W = N_G(T)/T$ is none other than the simple reflection s_β.

Recall also that the subgroup U is generated by the root subgroups $\{U_\beta \colon \beta > 0\}$; moreover, multiplication (for any ordering of the roots) yields a biholomorphism

$$\prod_{\beta > 0} U_\beta \xrightarrow{\sim} U,$$

and if we set $V_\alpha = \langle U_\beta \colon \beta > 0, \beta \neq \alpha \rangle$ (for α a simple root, as above), we find that

$$U = V_\alpha U_\alpha \cong V_\alpha \times U_\alpha.$$

Then, as s_α normalizes V_α, we have that $s_\alpha UB^- = V_\alpha s_\alpha U_\alpha B^-$. Putting all this together, we see that the map

$$\begin{array}{rcl} V_\alpha \times \mathfrak{g}_\alpha \times H \times U^- & \longrightarrow & s_\alpha UB^- \\ (v,a,t,u) & \longmapsto & v s_\alpha x_\alpha(a) t u \end{array}$$

is a biholomorphism. We may suppose that $s_\alpha = n_\alpha(1)$. Then, for $a \neq 0$,

$$s_\alpha x_\alpha(a) = x_\alpha(-a^{-1})\alpha^\vee(-a^{-1})x_{-\alpha}(a^{-1})$$

so that

$$\begin{aligned} v s_\alpha x_\alpha(a) t u &= v x_\alpha(-a^{-1})\alpha^\vee(-a^{-1})x_{-\alpha}(a^{-1}) t u \\ &= v x_\alpha(-a^{-1})\alpha^\vee(-a^{-1}) t x_{-\alpha}(\alpha(t)a^{-1}) u \end{aligned}$$

by (2.3.6) applied to $\beta = -\alpha$. In particular, $v s_\alpha x_\alpha(a) t u \in UHU^-$ for all $v \in V_\alpha$, $t \in H$, $u \in U^-$ and $a \in \mathfrak{g}_\alpha \backslash \{0\}$, in which case

$$\begin{aligned} f_\lambda(v s_\alpha x_\alpha(a) t u) &= \lambda(\alpha^\vee(-a^{-1}))^{-1}\lambda(t)^{-1} \\ &= (-a)^{(\lambda,\alpha^\vee)}\lambda(t)^{-1} \end{aligned}$$

by (2.3.5). As λ is dominant, $(\lambda,\alpha^\vee) \geq 0$, so that f_λ defines a holomorphic function on $V_\alpha \times \mathfrak{g}_\alpha \backslash \{0\} \times H \times U$ and hence can be (uniquely) extended to a holomorphic function on $V_\alpha \times \mathfrak{g}_\alpha \times H \times U \cong s_\alpha UB^-$, as desired.

Finally, observe that if $\alpha, \beta \in \Delta$ then the definitions of f_λ on $s_\alpha UB^-$ and $s_\beta UB-$ coincide on $s_\alpha UB^- \cap s_\beta UB^-$, because they coincide on the dense subset $s_\alpha UB^- \cap s_\beta UB^- \cap UB^-$. Thus the extensions of f_λ obtained above glue together to yield a well-defined (and nonzero!) holomorphic function on $UB^- \cup \bigcup_{\alpha \in \Delta} s_\alpha UB^-$, which is what we were after. ■

Remark 2.3.5. We emphasize that the two proofs presented above did not use the fact that there exists an irreducible representation of highest weight any given dominant weight λ. Thus the Borel–Weil theorem may be viewed as giving the "existence" part of the theorem of the highest weight. On the other hand, one can prove by independent means (for example, using Verma modules) that there exists an irreducible representation of highest weight λ for any given dominant weight λ, and from this one can then deduce the Borel–Weil theorem in a perhaps more straightforward manner; see Humphreys [40, §31.4] or Knapp [44, Ch.V, §7]. ▲

2.4 The Borel–Weil–Bott Theorem

We shall continue using the notation of the previous section. Thus G denotes a complex semisimple Lie group, $B^- = HU^-$ a Borel subgroup of G (made up from negative roots), $X = G/B^-$ the flag variety of G, $\widehat{H}$ the character group of H (and hence of B^-) and, for $\lambda \in \Lambda = \widehat{H}$, we let $\mathcal{L}_\lambda$ denote the line bundle $G \times_\lambda \mathbb{C}$ on X (or its sheaf of holomorphic sections).

The main goal of this section is to prove a generalization, due to Bott [13], of the Borel–Weil theorem. To state this result, let

$$\rho = \frac{1}{2} \sum_{\alpha > 0} \alpha \in \Lambda \otimes_{\mathbb{Z}} \mathbb{Q}.$$

and, for $\lambda \in \Lambda$, let

$$\begin{aligned} q_\lambda &= |\{\alpha \in \Phi^+ \colon (\lambda + \rho, \alpha) < 0\}| \\ &= \min\{\ell(w) \colon w \in W \text{ such that } w(\lambda + \rho) \text{ is dominant}\}, \end{aligned}$$

where $(\cdot, \cdot)$ is the bilinear form on $\mathfrak{h}^*$ induced by the Killing form and W is the Weyl group of G.[3] One knows that if $\lambda+\rho$ is regular, then there is a unique $w \in W$ that makes $w(\lambda+\rho)$ dominant regular, in which case $w(\lambda + \rho) - \rho$ is a dominant weight and $q_\lambda = \ell(w)$.

Theorem 2.4.1 (Borel–Weil–Bott)**.** *Assume the preceding notation and let $\lambda \in \Lambda$.*

(i) *If $\lambda + \rho$ is singular, then $H^q(X, \mathcal{L}_\lambda) = 0$ for all $q \geq 0$.*

(ii) *If $\lambda + \rho$ is regular, then $H^q(X, \mathcal{L}_\lambda)$ is nonzero if and only if $q = q_\lambda$, in which case it is an irreducible representation of G of highest weight $w(\lambda + \rho) - \rho$, where w is the unique element of W that makes $w(\lambda + \rho)$ dominant.*

In particular, $H^0(X, \mathcal{L}_\lambda)$ is nonzero if and only if $\lambda + \rho$ is dominant regular, which is the case if and only if λ is dominant. In this case the element of the Weyl group mentioned in the theorem is $w = 1$, and we conclude:

[3] For an alternative description of q_λ, see Remark 2.4.7(vi).

Corollary 2.4.2 (Borel–Weil). *$H^0(X, \mathcal{L}_\lambda)$ is nonzero if and only if λ is dominant, in which case it is an irreducible representation of G of highest weight λ.* ■

We shall prove the Borel–Weil–Bott theorem by means of an inductive argument whose inductive step requires a direct verification of the theorem for $G = \mathrm{SL}(2, \mathbb{C})$. This verification is performed in the following example.

Example 2.4.3. Let $G = \mathrm{SL}(2, \mathbb{C})$, $B^- = \{\left(\begin{smallmatrix} * & 0 \\ * & * \end{smallmatrix}\right)\}$, $H = \{\left(\begin{smallmatrix} * & 0 \\ 0 & * \end{smallmatrix}\right)\}$ and $X = G/B^- = \mathbb{P}^1$. In Example 2.3.3 we parameterized $\widehat{H}$ using $\mathbb{Z}$: to each $n \in \mathbb{Z}$ we associated a unique character λ_n, whose corresponding line bundle on X was $\mathcal{O}(n)$. We then found that

$$H^0(X, \mathcal{O}(n)) = \begin{cases} 0 & \text{if } n < 0 \\ \mathbb{C}[z_1, z_2]_n & \text{if } n \geq 0. \end{cases}$$

We want to now compute the higher cohomology groups. As $\dim_{\mathbb{C}} X = 1$, $H^q(X, \mathcal{O}(n))$ is potentially nonzero only if $q = 0$ or 1. The quickest way to compute $H^1(X, \mathcal{O}(n))$ is to use Serre duality, which tells us that $H^1(X, \mathcal{O}(n))$ and $H^0(X, \mathcal{O}(n)^* \otimes K)$ are dual and therefore isomorphic as representations of $G = \mathrm{SL}(2, \mathbb{C})$.[4] Here K is the canonical line bundle on $\mathbb{P}^1$, which is easily seen to be isomorphic to $\mathcal{O}(-2)$ (see also the remarks following Corollary 2.4.5). Thus, because $\mathcal{O}(n)^* \otimes \mathcal{O}(-2) = \mathcal{O}(-n) \otimes (-2) = \mathcal{O}(-n-2)$, we find that

$$H^1(X, \mathcal{O}(n)) \cong H^0(X, \mathcal{O}(-n-2)) = \begin{cases} 0 & \text{if } n > -2 \\ \mathbb{C}[z_1, z_2]_{-n-2} & \text{if } n \leq -2. \end{cases}$$

To compare this with the assertions of the Borel–Weil–Bott theorem, we first note that under the identification $\widehat{H} \cong \Lambda$, λ_n is given by $\lambda_n(\left(\begin{smallmatrix} x & 0 \\ 0 & -x \end{smallmatrix}\right)) = nx$ on $\mathfrak{h}$. In particular, $\Phi = \{\pm\lambda_2\}$, and λ_2 is the only positive root of G, whence $\rho = \frac{1}{2}\lambda_2 = \lambda_1$. Consequently, $\lambda_n + \rho = \lambda_{n+1}$ is singular if and only if $(\lambda_{n+1}, \lambda_2) = 0$. But $(\lambda_{n+1}, \lambda_2) = 2(n+1)(\lambda_1, \lambda_1)$, so that $\lambda_n + \rho$ is singular if and only if $n = -1$. We also see that q_{λ_n} is either 0 or 1,

[4] If V is an irreducible representation of G of highest weight λ, then the dual representation V^* is irreducible and of highest weight $-w_0\lambda$, where w_0 is the longest element of W. For $G = \mathrm{SL}(2, \mathbb{C})$, $-w_0\lambda = \lambda$.

depending on whether $n > -1$ or $n < -1$, respectively. So if we identify the Weyl group of G with $\mathbb{Z}/2\mathbb{Z} = \{1, \sigma\}$ as usual (so that $\sigma\lambda_n = \lambda_{-n}$), then the Borel–Weil–Bott theorem for G amounts to the following:

(i) $H^q(X, \mathcal{O}(-1)) = 0$ for all $q \geq 0$.

(ii) If $n \geq 0$, then $H^q(X, \mathcal{O}(n))$ is nonzero if and only if $q = 0$, in which case it is an irreducible representation of highest weight λ_n. On the other hand, if $n \leq -2$, then $H^q(X, \mathcal{O}(n))$ is nonzero if and only if $q = 1$, in which case it is an irreducible representation of highest weight $\sigma(\lambda_n + \rho) - \rho = \lambda_{-n-2}$.

This agrees with our findings above. ▲

The inductive argument we have in mind will also make use of the Leray spectral sequence of a fibration. The version we need is given below. (See Godement [26, Ch.2, Théorème 4.17.1] and Bott [13, §11] for the details.)

Proposition 2.4.4. *Let X and Y be compact and $\pi\colon X \to Y$ a locally trivial holomorphic fibration whose fibre F is compact and connected. Let V be a holomorphic vector bundle on X and denote by $\mathcal{V}$ (resp. $\mathcal{V}_F$) the sheaf of holomorphic sections of V (resp. $V|_F$). Then there is a spectral sequence E_r whose final term is associated to $H^*(X, \mathcal{V})$ and whose E_2 term is*

$$E_2^{q,p} = H^q(Y, \mathcal{W}^p),$$

where $\mathcal{W}^p$ is the sheaf of holomorphic sections of a vector bundle W^p over Y whose fibre at y is $H^p(\pi^{-1}(y), \mathcal{V}_{\pi^{-1}(y)}) \cong H^p(F, \mathcal{V}_F)$.[5] ■

Assume the notation in the preceding proposition and suppose further that X and Y are G-spaces, with G acting on Y transitively, and that π and V are G-equivariant. Then G_y, the isotropy subgroup of G at $y \in Y$, acts on the fibre $X_y \cong F$ and thus, by Proposition 2.2.3, $H^p(F, \mathcal{V}_F)$ becomes a holomorphic representation of G_y. We consequently obtain a G-equivariant vector bundle over $Y = G/G_y$, namely the fibre product of G and $H^p(F, \mathcal{V}_F)$

[5] For simplicity, we write $H^p(F, \mathcal{V}_F)$ for the sheaf $\mathcal{W}^p$, so that the E_2 term of the spectral sequence is $E_2^{q,p} = H^q(Y, H^p(F, \mathcal{V}_F))$.

(cf. Proposition 2.2.1). It follows that all the terms of the Leray spectral sequence are G-modules and that the associated differentials are G-equivariant. Moreover, the filtration of $H^k(X, \mathcal{V})$ is G-invariant and its corresponding graded module is isomorphic, as a G-module, to $\bigoplus_{p+q=k} E_\infty^{q,p}$. In particular:

Corollary 2.4.5. *If $H^p(F, \mathcal{V}_F) = 0$ for all p, except perhaps for $p = p'$, then the spectral sequence collapses and we obtain G-module isomorphisms*

$$H^k(X, \mathcal{V}) \cong \bigoplus_{p+q=k} E_2^{p,q} \cong H^{k-p'}(Y, H^{p'}(F, \mathcal{V}_F))$$

for each $k \geq 0$. ■

We need one last piece of information. Let P be a parabolic subgroup of G that contains B^- and set $Y = G/P$. The Lie algebra of P is of the form $\mathfrak{p} = \mathfrak{b}^- \oplus \bigoplus_{\alpha \in \Phi_P} \mathfrak{g}_\alpha$, where Φ_P is a subset of the *positive* roots Φ^+. Let $K_Y = \bigwedge^n T_Y^*$ denote the canonical bundle of Y. Here $n = \dim_{\mathbb{C}} Y = |\Phi^+ - \Phi_P|$ and T_Y^* is the cotangent bundle of Y. Since the latter is simply $(\mathfrak{g}/\mathfrak{p})^* = \mathfrak{p}^\perp = \bigoplus_{\alpha \in \Phi^+ - \Phi_P} \mathfrak{g}_\alpha$, we see that K_Y becomes a G-equivariant line bundle under the adjoint action of G, and hence corresponds to a one-dimensional representation of P. The weight associated to this representation is clearly $-2\rho_P$, where

$$\rho_P = \frac{1}{2} \sum_{\alpha \in \Phi^+ - \Phi_P} \alpha \in \Lambda \otimes_{\mathbb{Z}} \mathbb{Q}.$$

In particular, the canonical bundle of $X = G/B^-$ is $\mathcal{L}_{-2\rho}$.

With this in hand, we can now proceed to the

Proof of the Borel–Weil–Bott theorem. Fix a positive simple root $\alpha \in \Phi^+$ and let P_α denote the corresponding parabolic subgroup of G containing B^-. Set $Y_\alpha = G/P_\alpha$. Then the inclusion $B \subset P_\alpha$ induces a fibration $X \xrightarrow{\pi} Y_\alpha$ with fibre $P_\alpha/B \cong \mathbb{P}^1$, the flag variety of $\mathrm{SL}(2, \mathbb{C})$ (see end of Section 1.11). Proposition 2.4.4 provides us with a spectral sequence E_r for computing $H^*(X, \mathcal{L}_\lambda)$ with E_2 term

$$E_2^{q,p} = H^q(Y_\alpha, H^p(\mathbb{P}^1, \mathcal{L}_\lambda)).$$

Here it is to be understood that $H^p(\mathbb{P}^1, \mathcal{L}_\lambda)$ stands for the sheaf of holomorphic sections of a G-equivariant vector bundle on Y_α whose fibre at $y \in Y_\alpha$ is $H^p(\pi^{-1}(y), \mathcal{L}_\lambda|_{\pi^{-1}(y)})$. At any rate, we see that $H^p(\mathbb{P}^1, \mathcal{L}_\lambda) = 0$ if $p > 1$. The Borel–Weil–Bott theorem for $\mathrm{SL}(2, \mathbb{C})$ (cf. Example 2.4.3) tells us that the same is true in at least one of degrees $p = 0$ and $p = 1$. Corollary 2.4.5 then asserts that there is a p' such that

$$H^k(X, \mathcal{L}_\lambda) = H^{k-p'}(Y_\alpha, H^{p'}(\mathbb{P}^1, \mathcal{L}_\lambda)). \tag{2.4.1}$$

Let $s_\alpha \in W$ denote the simple reflection in α. The same argument applied to $s_\alpha(\lambda+\rho)-\rho$ yields a p'' such that

$$H^k(X, \mathcal{L}_{s_\alpha(\lambda+\rho)-\rho}) = H^{k-p''}(Y_\alpha, H^{p''}(\mathbb{P}^1, \mathcal{L}_{s_\alpha(\lambda+\rho)-\rho})). \tag{2.4.2}$$

Now recall that $(\rho, \alpha^\vee) = 2(\rho, \alpha)/(\alpha, \alpha) = 1$ (as is true for any simple α). Thus if we put $m = (\lambda, \alpha^\vee)$, we find that

$$s_\alpha(\lambda+\rho) - \rho = (\lambda + \rho - (\lambda+\rho, \alpha^\vee)\alpha) - \rho = \lambda - (m+1)\alpha.$$

By the remarks following Corollary 2.4.5 we know that the canonical bundle of Y_α is $\mathcal{L}_{-\alpha}$. So if $m \geq -1$, then

$$\mathcal{L}_{s_\alpha(\lambda+\rho)-\rho} \cong \mathcal{L}_\lambda \otimes (K_{Y_\alpha})^{\otimes m+1}. \tag{2.4.3}$$

Let's take a closer look at the restriction of the line bundle $\mathcal{L}_\lambda$ to the fibre $\pi^{-1}(y) \cong P_\alpha/B \cong \mathbb{P}^1$. The last isomorphism arises from the root homomorphism $x_\alpha \colon \mathrm{SL}(2, \mathbb{C}) \to G$. Recall from (2.3.7) that $x_\alpha(\mathrm{diag}(a, a^{-1})) = \alpha^\vee(a)$ for all $a \in \mathbb{C}^\times$. But then, if we think of λ as a character in $\widehat{H}$, we see that $\lambda \circ x_\alpha$ is a character $\mathbb{C}^* \to \mathbb{C}^*$ and hence must be of the form $(\lambda \circ x_\alpha)(a) = a^k$ for some integer k. This integer k is none other than $(\lambda, \alpha^\vee)$. Consequently, the restriction of $\mathcal{L}_\lambda$ to $\mathbb{P}^1$ is simply $\mathcal{O}((\lambda, \alpha^\vee)) = \mathcal{O}(m)$ in the notation of Example 2.3.3. Similarly, the restriction of the canonical bundle $K_{Y_\alpha} = \mathcal{L}_{-\alpha}$ to $\mathbb{P}^1$ is $\mathcal{O}((-\alpha, \alpha^\vee)) = \mathcal{O}(-2)$—the canonical bundle of $\mathbb{P}^1$, as expected. In this light, the right-hand side of (2.4.3) becomes

$$\mathcal{O}(m) \otimes \mathcal{O}(-2m-2) \cong \mathcal{O}(-m) \otimes \mathcal{O}(-2) \cong \mathcal{L}_\lambda^* \otimes \mathcal{O}(-2)$$

when restricted to $\pi^{-1}(y)$. So, by Serre duality,

$$H^0(\mathbb{P}^1, \mathcal{L}_{s_\alpha(\lambda+\rho)-\rho}) = H^1(\mathbb{P}^1, \mathcal{L}_\lambda) \quad \text{and} \quad H^1(\mathbb{P}^1, \mathcal{L}_{s_\alpha(\lambda+\rho)-\rho}) = H^0(\mathbb{P}^1, \mathcal{L}_\lambda) \tag{2.4.4}$$

as representations of G (cf. the footnote on page 43). Thus, if $m \geq -1$, then (2.4.3) yields

$$H^0(\mathbb{P}^1, \mathcal{L}_{s_\alpha(\lambda+\rho)-\rho}) = H^1(\mathbb{P}^1, \mathcal{L}_\lambda).$$

But if $m \geq -1$ we also have that $H^1(\mathbb{P}^1, \mathcal{L}_\lambda) = 0$ (by the Borel–Weil–Bott theorem for $\mathrm{SL}(2, \mathbb{C})$, for example). So, going back to (2.4.1) and (2.4.2) we see that $p' = 0$ and $p'' = 1$. Then, by combining (2.4.1), (2.4.2) and (2.4.4), we find that

$$H^k(X, \mathcal{L}_\lambda) \cong H^{k+1}(X, \mathcal{L}_{s_\alpha(\lambda+\rho)-\rho}) \qquad (k \geq 0) \tag{2.4.5}$$

if $(\lambda + \rho, \alpha) \geq 0$.

If $(\lambda+\rho, \alpha) = 0$, then $m = -1$ and both sides of (2.4.5) are zero by the Borel–Weil–Bott theorem for $\mathrm{SL}(2, \mathbb{C})$; as this is true for any simple root α, part (i) of the theorem follows.

We now turn to part (ii). We shall first prove that $H^k(X, \mathcal{L}_\lambda) = 0$ for all $k > 0$ when λ is dominant. Let w_0 denote the longest element of W, which we can write as a product of $n = \ell(w_0) = \dim_{\mathbb{C}} X$ simple reflections, say $w_0 = s_{\alpha_1} \cdots s_{\alpha_n}$. If λ is dominant, then $(s_{\alpha_{i+1}} \cdots s_{\alpha_n}(\lambda + \rho), \alpha_i) \geq 0$ for each i, hence by repeatedly applying the isomorphism (2.4.5) we obtain

$$H^k(X, \mathcal{L}_\lambda) \cong H^{k+\dim_{\mathbb{C}} X}(X, \mathcal{L}_{w_0(\lambda+\rho)-\rho}) \qquad (k \geq 0)$$

and the desired conclusion follows since $H^q(X, \mathcal{L}_{w_0(\lambda+\rho)-\rho})$ vanishes if $q > \dim_{\mathbb{C}} X$.

Finally, assume that $\lambda + \rho$ is regular and let w denote the unique element of W that makes $w(\lambda + \rho)$ dominant. Let $w = s_{\alpha_1} \cdots s_{\alpha_l}$ be the minimal expansion of w as a product of simple reflections. Here $l = \ell(w) = q_\lambda$. Thus, by iterating (2.4.5) as before, we obtain

$$H^k(X, \mathcal{L}_{w(\lambda+\rho)-\rho}) \cong H^{k+q_\lambda}(X, \mathcal{L}_\lambda) \qquad (k \geq 0).$$

Since $w(\lambda + \rho) - \rho$ is dominant, it follows from our conclusion in the previous paragraph that $H^{k+q_\lambda}(X, \mathcal{L}_\lambda) = 0$ if $k > 0$. And if $k = 0$, then $H^{q_\lambda}(X, \mathcal{L}_\lambda) \cong H^0(X, \mathcal{L}_{w(\lambda+\rho)-\rho})$ is irreducible and of highest weight $w(\lambda + \rho) - \rho$, by the Borel–Weil theorem. ∎

We shall now generalize the Borel–Weil–Bott theorem to G/P, where P is a parabolic subgroup of G that contains B^-.

Theorem 2.4.6 (Borel–Weil–Bott for G/P)**.** *Let $\mathcal{V}_\varphi = G \times_\varphi V$ be the G-equivariant vector bundle over $Y = G/P$ corresponding to an irreducible representation $\varphi\colon P \to \mathrm{GL}(V)$, and let $\lambda \in \Lambda$ be the highest weight of φ.*

(i) *If $\lambda + \rho$ is singular, then $H^q(Y, \mathcal{V}_\varphi) = 0$ for all $q \geq 0$.*

(ii) *If $\lambda + \rho$ is regular, then $H^q(Y, \mathcal{V}_\varphi)$ is nonzero if and only if $q = q_\lambda$, in which case it is an irreducible representation of G of highest weight $w(\lambda + \rho) - \rho$, where w is the unique element of W that makes $w(\lambda + \rho)$ dominant.*

Proof. We will prove that

$$H^q(Y, \mathcal{V}_\varphi) \cong H^q(X, \mathcal{L}_\lambda) \qquad \text{for all } q \geq 0.$$

The theorem will then follow by the Borel–Weil–Bott theorem for $X = G/B^-$.

Consider the fibration $X \xrightarrow{\pi} Y$ with fibre $F = P/B^-$. We wish to use the Leray spectral sequence (Proposition 2.4.4) to compute $H^q(X, \mathcal{L}_\lambda)$. To do so we must examine the P-module structure on $H^q(F, \mathcal{L}_\lambda)$, where $\mathcal{L}_\lambda$ is tacitly understood to be the restriction of $\mathcal{L}_\lambda \to X$ to F. Let $P = SCU_P$ be the Levi decomposition of P. Here S is semisimple, $C \subset H$ is an algebraic torus, $U_P \subset U^-$ is the unipotent radical of P, and SC is the centralizer of C. In particular, $CU_P \subset HU^- = B$ and therefore

$$F = P/B^- = S/S \cap B^-.$$

As $S \cap B^-$ is a Borel subgroup of S, the Borel–Weil–Bott theorem tells us that $H^q(F, \mathcal{L}_\lambda) = 0$ for $q > 0$ and that $H^0(F, \mathcal{L}_\lambda)$ is an irreducible representation of S of highest weight λ. We claim that $H^0(F, \mathcal{L}_\lambda)$ is in fact an irreducible representation of P of highest weight λ, and therefore must be isomorphic to V. To this end, we must examine the action of C on $H^0(F, \mathcal{L}_\lambda)$. Concretely, we have

$$H^0(F, \mathcal{L}_\lambda) = \{f \in \mathrm{Hol}(P, \mathbb{C})\colon f(pb) = \varphi(b)^{-1} f(p) \text{ for } p \in P \text{ and } b \in B^-\}$$

with P acting via the left regular representation (cf. Proposition 2.2.5). So if $a \in C$ and $p = scu \in P = SCU_P$, then

$$f(a^{-1}p) = f(a^{-1}scu) = f(a^{-1}sc) = f(sca^{-1}) = \varphi(a)f(p).$$

This means that C acts on $H^0(F, \mathcal{L}_\lambda)$ via φ, and this proves our claim. Corollary 2.4.5 then tells us that

$$H^q(X, \mathcal{L}_\lambda) \cong H^q(Y, H^0(F, \mathcal{L}_\lambda)) = H^q(Y, \mathcal{V}_\varphi),$$

as desired. ■

Remarks 2.4.7.

(i) In [10], Borel and Hirzebruch determine which line bundles $\mathcal{L}_\lambda$, $\lambda \in \Lambda$, over $Y = G/P$ are positive (in the sense of Kodaira). They find that $\mathcal{L}_\lambda$ is positive if $(\lambda, \alpha^\vee) > 0$ for all $\alpha \in \Phi^+ - \Phi_P$. In particular, if $P = B^-$ (so that $\Phi_P = \emptyset$) and λ is dominant, then $\mathcal{L}_\lambda \otimes K_X^{-1}$ is positive, in which case the Kodaira vanishing theorem implies that $H^q(X, \mathcal{L}_\lambda) = 0$ for $q > 0$. Combined with the Borel–Weil theorem, this yields the Borel–Weil–Bott theorem for dominant weights. For general $\lambda \in \Lambda$, Borel and Hirzebruch [11] employ the Hirzebruch–Riemann–Roch theorem to compute the Euler characteristic $\chi(X, \mathcal{L}_\lambda)$. They find the answer to be either zero or equal to (plus or minus) the dimension of an irreducible G-module. This led them to conjecture what is now known as the Borel–Weil–Bott theorem (Theorem 2.4.1).[6]

(ii) The ρ-shift in the Borel–Weil–Bott theorem is precisely what makes the statement compatible with Serre duality—this is especially apparent in the proof we have given.

(iii) The proof of Theorem 2.4.1 given above (which we have adapted from Schmid [59]) is close to Bott's original one. The most fundamental difference is that Bott used the Kodaira vanishing theorem to establish the vanishing of high degree cohomology. We have avoided this theorem by cleverly using the symmetry of the Weyl group; this line of reasoning appears to be due to Demazure (cf. [17] and [18]), although the germ of it was already present in Bott [13].

[6] Note that while the Borel–Hirzebruch paper appeared roughly seven months after Bott's [13], Bott was aware of it and in fact references it in his own paper.

(iv) The Borel–Weil theorem was used to finish the proof of the Borel–Weil–Bott theorem. Thus, strictly speaking, we cannot call the Borel–Weil theorem a corollary (as we did in Corollary 2.4.2). However, there are other proofs of the Borel–Weil–Bott theorem which do not rely on the Borel–Weil theorem. See, for example, Demazure [18] for an elegant "algebraic" proof that is much in the spirit of the proof given above. Another "algebraic" proof is presented in Chapter 3, where a theorem of Kostant will be used to reduce the Borel–Weil–Bott theorem to a calculation in Lie algebra cohomology which is independent of the Borel–Weil theorem (see Section 3.5.1).

(v) In Chapter 4 we will introduce a far-reaching generalization of the Borel–Weil–Bott theorem due to Schmid.

(vi) Here is a more geometric-looking reformulation of the Borel–Weil–Bott theorem. Assume for simplicity that G is simply connected. Let $\mathcal{L}$ be a G-equivariant line bundle over X. We know that $X = G/B^- = K/T$, where K is a compact real form of G and $T = K \cap B$ is a maximal torus in K. Thus any hermitian metric on $\mathcal{L}$ may be averaged over K to yield a K-invariant metric. The curvature form $\Theta_{\mathcal{L}}$ of this metric is then completely determined by its value at a point, and hence may be viewed as a hermitian form on the holomorphic tangent space of X at any point. One says that an equivariant line bundle $\mathcal{L}$ is *regular* if $\Theta_{\mathcal{L}}$ is nondegenerate; otherwise $\mathcal{L}$ is said to be *singular*. The canonical bundle K_X has a square root: there is line a bundle $K_X^{1/2}$, which is equivariant because G is simply connected, such that $K_X^{1/2} \otimes K_X^{1/2} = K_X$. The Borel–Weil–Bott theorem then states that $H^q(X, \mathcal{L}) = 0$ for all q if $\mathcal{L} \otimes K_X^{-1/2}$ is singular; and if $\mathcal{L} \otimes K_X^{-1/2}$ is regular, then $H^q(X, \mathcal{L}) = 0$ for all but one $q = q_{\mathcal{L}}$, in which case it is an irreducible G-module whose highest weight may be written down explicitly in terms of the first Chern class of $\mathcal{L}$. The first Chern class itself is simply

$$c_1(\mathcal{L}) = \frac{i}{2\pi} \sum_{\alpha>0} (\lambda, \alpha^\vee) dx_\alpha \wedge d\overline{x}_\alpha,$$

where λ is the element of $\Lambda = \mathrm{Pic}(X)$ associated to $\mathcal{L}$ and $(x_\alpha)_{\alpha>0}$ are the exponential coordinates near the identity of G. In particular, $c_1(\mathcal{L})$ and λ, thought of as elements in $\mathrm{Pic}(X) \cong H^1(X, \mathcal{O}_X^*)$, both represent the same element in $H^2(X, \mathbb{R})$. Finally, the number $q_{\mathcal{L}}$ is equal to the number of negative eigenvalues of the form $\Theta_{\mathcal{L} \otimes K_X^{-1/2}}$ on

$\mathcal{L} \otimes K_X^{-1/2}$. For the details, see Borel and Hirzebruch [10], [11], [12], Griffiths and Schmid [29], and Snow [64]. ▲

2.5 Applications

We retain the notation of the previous two sections. In particular, we denote by G a complex semisimple Lie group and by P a parabolic subgroup of G.

2.5.1 Cohomology of Certain Homogeneous Spaces

Historically, the main application of the Borel–Weil–Bott theorem was the determination of the cohomology groups $H^q(G/P, \mathcal{V})$, where $\mathcal{V}$ is induced from an irreducible P-module (for example, $\mathcal{V}$ = an equivariant line bundle). The generalized flag varieties $X = G/P$ form an important class of examples in algebraic geometry, and therefore one would naturally like to know, for example, the dimensions of the cohomology groups $H^q(X, \mathcal{V})$. Prior to Bott's paper [13], only partial results in this direction were known, such as those of Borel and Hirzebruch (cf. Remark 2.4.7(i)).

At any rate, to determine $\dim H^q(X, \mathcal{V})$ using the Borel–Weil–Bott theorem, it seems that one needs to be able to compute the dimension of any given irreducible G-module. The Weyl dimension formula works well in this regard. In fact, one can use the Borel–Weil–Bott theorem to deduce the Weyl dimension formula (and, more generally, the Weyl character formula); we shall take this matter up in Section 2.5.3.

To place X in a perhaps more distinguished context, we have the following classical theorem.

Theorem 2.5.1 (Wang [71]). *The generalized flag variety $X = G/P$ is a compact, simply connected, homogeneous Kähler manifold, and conversely, every such manifold arises in this fashion.* ■

It follows that many important algebro-geometric spaces, such as projective spaces and more generally Grassmannians, are of the form $X = G/P$. The Borel–Weil–Bott theorem then computes the cohomology groups $H^*(X, \mathcal{V})$ of all the G-equivariant vector bundles $\mathcal{V}$ over X which are induced from irreducible representations of P. While not every vector bundle over X is equivariant, one at least has the following remarkable result for line bundles.

Proposition 2.5.2. *If G is simply connected, then every holomorphic line bundle on X is G-equivariant. Thus we obtain a group isomorphism* $\mathrm{Pic}(X) \cong \Lambda$. ■

A proof of the first assertion can be found in the proof of Theorem 1 in Lurie [51]. The second assertion follows from the obvious isomorphism $\mathcal{L}_\lambda \otimes \mathcal{L}_\mu \cong \mathcal{L}_{\lambda+\mu}$.

Remark 2.5.3. If G is not simply connected, then the above theorem is no longer true. For example, take $G = \mathrm{PSL}(n+1, \mathbb{C})$ and choose P so that $G/P = \mathbb{P}^n$. Then the line bundle $\mathcal{O}(m)$ is G-equivariant if and only if $n+1$ divides m (cf. [22, Proposition 23.13]). For instance, if $\mathcal{O}(1)$ were G-equivariant, then the Borel–Weil theorem would imply that $H^0(G/P, \mathcal{O}(1)) = \mathbb{C}^{n+1}$ (see Example 2.5.4) is an irreducible representation of G, but G does not have an $(n+1)$-dimensional irreducible representation. ▲

Thus, for simply connected G, the Borel–Weil–Bott theorem computes the cohomology $H^*(X, \mathcal{L})$ of any line bundle $\mathcal{L}$ over X. If G is not simply connected, then $\mathcal{L}$ will at least be acted on equivariantly by the universal covering group $\widetilde{G}$ of G. This is because the flag varieties of G and $\widetilde{G}$ coincide, for they can both be identified with the variety of Borel subalgebras of $\mathrm{Lie}(G) = \mathrm{Lie}(\widetilde{G})$. So we can still say something about $H^*(X, \mathcal{L})$ in this case, though only in terms of the representation theory of $\widetilde{G}$ and not necessarily in that of G.

Example 2.5.4. Here we compute $H^q(\mathbb{P}^n, \mathcal{O}(m))$. Let $G = \mathrm{SL}(n+1, \mathbb{C})$ and let P be the parabolic subgroup consisting of the block lower triangular matrices

$$\left(\begin{array}{ccc|c} * & \cdots & * & 0 \\ \vdots & \ddots & \vdots & \vdots \\ * & \cdots & * & 0 \\ \hline * & \cdots & * & a \end{array}\right)$$

so that $Y = G/P = \mathbb{P}^n$. Let $\lambda_m \colon P \to \mathbb{C}^\times$ be the character which sends the above matrix to a^{-m}. The corresponding line bundle $\mathcal{L}_{\lambda_m}$ over Y is $\mathcal{O}(m)$.[7]

Now identify $\widehat{H}$ with Λ. At the Lie algebra level, we have $\mathfrak{h} = \{$diagonal subalgebra of $\mathfrak{sl}(n+1,\mathbb{C})\}$. The roots are $\{e_i - e_j \colon i \neq j\}$, where

$$e_i \begin{pmatrix} h_1 & & \\ & \ddots & \\ & & h_{n+1} \end{pmatrix} = h_i,$$

the positive roots are $\Phi^+ = \{e_i - e_j \colon i < j\}$, the simple roots are $\Delta = \{e_i - e_{i+1} \colon i \leq n\}$, and

$$\rho = \frac{1}{2} \sum_{i<j} e_i - e_j = \sum_{i=1}^{n+1} (n-i) e_i.$$

Observe that $\rho \in \Lambda$ (which we know must be the case because G is simply connected). As $e_1 + \cdots + e_{n+1} =$ trace, we have

$$\Lambda = \mathbb{Z}[e_1, \ldots, e_{n+1}]/\{e_1 + \cdots e_{n+1} = 0\} \cong \mathbb{Z}^{n+1}/(1, \ldots, 1).$$

Under this identification, λ_m corresponds to the weight $-me_{m+1} = (0, \ldots, 0, -m)$. Thus $\lambda_m + \rho = (n, n-1, \ldots, 2, 1, -m)$ is singular if and only if $(\lambda_m + \rho, e_i - e_j) = 0$ for some $i \neq j$, which is the case if and only if $n+m+1 = i$ for some i ($1 \leq i \leq n$). We can rewrite this last condition as $1 \leq n+m+1 \leq n$, and then the Borel–Weil–Bott theorem tells us that, in this case,

$$H^q(\mathbb{P}^n, \mathcal{O}(m)) = 0 \qquad \text{for all } q \geq 0.$$

Next suppose that $\lambda_m + \rho$ is regular. Then

$$\begin{aligned} q_{\lambda_m} &= |\{\alpha \in \Phi^+ \colon (\lambda_m + \rho, \alpha) < 0\}| \\ &= |\{i \in \{1, \ldots, n\} \colon n+m+1 < i\}|, \end{aligned}$$

and we have two cases to consider.

(i) If $n+m+1 > n$, then $q_{\lambda_m} = 0$ and $\lambda_m + \rho$ is dominant. The Borel–Weil–Bott theorem then asserts that $H^q(\mathbb{P}^n, \mathcal{O}(m))$ vanishes if $q > 0$, and is an irreducible representation of G of highest weight λ_m if $q = 0$.

[7] It is a well-known fact of algebraic geometry that $\mathrm{Pic}(\mathbb{P}^n) = \{\mathcal{O}(m)\}_{m \in \mathbb{Z}}$. This illustrates Proposition 2.5.2 for $G = \mathrm{SL}(n+1, \mathbb{C})$.

(ii) If $n + m + 1 < 1$, then $q_{\lambda_m} = n$. So if we let w denote the unique element of W that makes $w(\lambda_m + \rho)$ dominant, the Borel–Weil–Bott theorem then tells us that $H^q(\mathbb{P}^n, \mathcal{O}(m))$ vanishes unless $q = n$, in which case it is an irreducible representation of highest weight $w(\lambda_m + \rho) - \rho$. As $\ell(w) = q_{\lambda_m} = n$, we find that $w(\lambda_m + \rho) - \rho = \lambda_{-n-m-1}$. (This last assertion can also be verified using the Weyl dimension formula.)

In summary:

$$H^q(\mathbb{P}^n, \mathcal{O}(m)) = \begin{cases} V^m & \text{if } q = 0 \text{ and } m \geq 0 \\ V^{-n-m-1} & \text{if } q = n \text{ and } m \leq -n-1 \\ 0 & \text{otherwise,} \end{cases} \tag{2.5.1}$$

where V^k is the irreducible representation of G of highest weight λ_k, which we can identify with the space of degree k homogeneous polynomials in $\mathbb{C}[z_1, \ldots, z_{n+1}]$. ▲

Example 2.5.5. We can use the Borel–Weil–Bott theorem to verify that the line bundle $\mathcal{O}(m)$ over $\mathbb{P}^n$ satisfies Serre duality. Indeed, in the notation of the previous example, the the subset of Φ^+ corresponding to P is $\Phi_P = \{e_i - e_j : i < j \leq n\}$, so that

$$\rho_P = \frac{1}{2} \sum_{\alpha \in \Phi^+ - \Phi_P} \alpha = \frac{1}{2}(1, \ldots, 1, -n) = \frac{1}{2}(0, \ldots, 0, -n-1) = \frac{1}{2}\lambda_{n+1} \in \Lambda \otimes_{\mathbb{Z}} \mathbb{R}.$$

By the remarks following Corollary 2.4.5, the canonical bundle of $\mathbb{P}^n = G/P$ is

$$K = \mathcal{L}_{-2\rho_P} = \mathcal{L}_{\lambda_{-n-1}} = \mathcal{O}(-n-1).$$

Thus $\mathcal{O}(m)^* \otimes K = \mathcal{O}(-m-n-1)$, and by looking at (2.5.1), we see that Serre duality is (of course) satisfied. ▲

Example 2.5.6. Our findings in Example 2.5.4 allow us to compute $H^q(\mathbb{P}^1, \mathcal{V})$ for any vector bundle $\mathcal{V}$ over $\mathbb{P}^1 = \mathrm{SL}(2, \mathbb{C})/B$. Indeed, a famous theorem of Grothendieck [30] asserts that $\mathcal{V}$ is a direct sum of line bundles, say $\mathcal{V} = \mathcal{O}(a_1) \oplus \cdots \oplus \mathcal{O}(a_r)$. Then

$$H^q(\mathbb{P}^1, \mathcal{V}) = H^q(\mathbb{P}^1, \mathcal{O}(a_1)) \oplus \cdots \oplus H^q(\mathbb{P}^1, \mathcal{O}(a_r))$$

as G-modules, and the right side can be dealt with using (2.5.1). ▲

2.5.2 The Complex Structure on G/P

In his original paper [13], Bott applied the Borel–Weil–Bott theorem to show that the complex structure on the generalized flag variety G/P is locally rigid, i.e., it cannot be locally deformed. This result requires a familiarity with some aspects of the deformation theory of complex manifolds. We briefly recall the relevant facts. (More details can be found, for example, in Huybrechts [41, Ch.6].)

By a **smooth family** of complex manifolds, we mean a triple $(\mathcal{X}, B, \pi)$, where $\mathcal{X}$ and B are complex manifolds and $\pi\colon \mathcal{X} \to B$ is a surjective smooth proper map. The fibres $\mathcal{X}_t = \pi^{-1}(t)$, $t \in B$, are then compact submanifolds of $\mathcal{X}$.

We have the following basic result.

Theorem 2.5.7 (Ehresmann). *Let $\pi\colon \mathcal{X} \to B$ be a smooth family of complex manifolds. If B is connected, then all the fibres $\mathcal{X}_t$ are diffeomorphic. Moreover, each $t_0 \in B$ has a neighborhood U such that $\mathcal{X}|_U = \pi^{-1}(U)$ is diffeomorphic to the product $\mathcal{X}_{t_0} \times B$.* ■

Remark 2.5.8. The theorem is no longer true if "diffeomorphic" is replaced by "isomorphic (as complex manifolds)". ▲

Let $\pi\colon \mathcal{X} \to B$ be a smooth family of complex manifolds, and assume that B is an open connected subset of $\mathbb{C}^n$ that contains 0. Set $X = \mathcal{X}_0$. We typically restrict our attention to the germ of 0, i.e. π is only considered over arbitrarily small neighborhoods of 0. In view of the above theorem, we can choose a diffeomorphism $\mathcal{X} \cong \mathcal{X}_{t_0} \times B$ (over some neighborhood of 0). Thus the family of fibres $\mathcal{X}_t$ (for "small t"), each of which is diffeomorphic to X, may be thought of as a smooth deformation of the complex structure on X. We say that the complex structure on X is **locally rigid** if, for each smooth deformation $\mathcal{X}_t$, there is a neighborhood U of 0 such that $\mathcal{X}_t$ is isomorphic to X, as a complex manifold, for all $t \in U$.

Theorem 2.5.9 (Frölicher–Nijenhuis [20]). *Let X be a compact complex manifold and denote by T_X the holomorphic tangent bundle of X. If $H^1(X, T_X) = 0$, then X is locally rigid.* ■

This theorem ought to be viewed as part of Kodaira–Spencer theory. In any case, we are now ready for Bott's result.

Theorem 2.5.10 (Bott). *Let G be a complex semisimple Lie group and P a parabolic subgroup of G, and set $X = G/P$. Then $H^q(X, T_X) = 0$ for $q \geq 1$.*

Proof. The short exact sequence of P-modules

$$0 \longrightarrow \mathfrak{p} \longrightarrow \mathfrak{g} \longrightarrow \mathfrak{g}/\mathfrak{p} \longrightarrow 0,$$

with P acting via the adjoint representation, induces a short exact sequence

$$0 \longrightarrow \mathcal{V} \longrightarrow \mathcal{E} \longrightarrow \mathcal{T} \longrightarrow 0$$

of G-equivariant vector bundles over X. The action of P on $\mathfrak{g}$ extends to an action of G, whence $\mathcal{E}$ is trivial and therefore $H^q(X, \mathcal{E}) = 0$ for $q \geq 1$. As $\mathcal{T} = T_X$, the associated long exact sequence in cohomology yields isomorphisms

$$H^q(X, T_X) \cong H^{q+1}(X, \mathcal{V}) \qquad \text{for each } q \geq 1.$$

So, to prove the theorem, it suffices to show that $H^q(X, \mathcal{V}) = 0$ for $q \geq 2$.

Let $\mathfrak{p} = V_n \supset V_{n-1} \supset \cdots \supset V_0 = 0$ be a filtration of $\mathfrak{p}$ by P-submodules V_i with irreducible quotients $W_{i+1} = V_{i+1}/V_i$ of highest weight λ_i. For each $i \in \{0, \ldots, n-1\}$, the short exact sequence of P-modules

$$0 \longrightarrow V_i \longrightarrow V_{i+1} \longrightarrow W_{i+1} \longrightarrow 0$$

induces a short exact sequence of G-equivariant vector bundles over X

$$0 \longrightarrow \mathcal{V}_i \longrightarrow \mathcal{V}_{i+1} \longrightarrow \mathcal{W}_{i+1} \longrightarrow 0$$

which in turn induces a long exact sequence of G-modules

$$\cdots \longrightarrow H^q(X, \mathcal{V}_i) \longrightarrow H^q(X, \mathcal{V}_{i+1}) \longrightarrow H^q(X, \mathcal{W}_{i+1}) \longrightarrow H^{q+1}(X, \mathcal{V}_i) \longrightarrow \cdots . \tag{2.5.2}$$

The Borel–Weil–Bott theorem (Theorem 2.4.6) asserts that $H^q(X, \mathcal{W}_{i+1}) = 0$ if $\lambda_i + \rho$ is singular or if $q \neq q_{\lambda_i}$. Since P is acting via the adjoint representation, each weight λ_i is either zero or else a root belonging to $\Phi^- \cup \Phi_P$.

Lemma 2.5.11 (Borel [13, Lemma 4.1]). *Let $\alpha \in \Phi$. If $\alpha + \rho$ is regular, then $q_\alpha < 2$.*

Proof. Recall that $q_\alpha = \{\beta \in \Phi^+ \colon (\alpha+\rho, \beta) < 0\}$ and that, if $\alpha+\rho$ is regular, there exists a unique $w \in W$ such that $w(\alpha+\rho)$ is dominant—and then $\ell(w) = q_\alpha$.

Assume without loss of generality that G is simple. Let $\alpha_1, \ldots, \alpha_l$ be the simple roots of G, and let $s_i = s_{\alpha_i} \in W$ denote the simple reflection in α_i. One knows that $(\alpha, \alpha_i^\vee) \in \{0, \pm 1, \pm 2, \pm 3\}$ with ± 3 occurring only if G is of type G_2.

If G is of type G_2, then there are two simple roots α_1 and α_2, with $(\alpha_1, \alpha_2^\vee) = -3$ and $(\alpha_2, \alpha_1^\vee) = -1$, so that $\Phi^+ = \{\alpha_1, \alpha_2, \alpha_1+\alpha_2, \alpha_1+2\alpha_2, \alpha_1+3\alpha_2, 2\alpha_1+3\alpha_2\}$, whence $\rho = 3\alpha_1 + 5\alpha_2$. It follows that $\alpha+\rho$ is regular if and only if $\alpha = -\alpha_1$ or $-\alpha_2$, and from here one can easily verify the lemma by explicit computation.

Assume now that G is not of type G_2. If $\alpha+\rho$ is regular then $(\alpha+\rho, \alpha_i^\vee) \neq 0$, or equivalently, $(\alpha, \alpha_i^\vee) \neq -1$, for all i. So we are left with the set $\{-2, 0, 1, 2\}$ of possible values for $(\alpha, \alpha_i^\vee)$. Order the simple roots so that $(\alpha, \alpha_i^\vee) = -2$ if $i \leq s$ and $(\alpha, \alpha_i^\vee) \geq 0$ if $i > s$. We claim that $s \leq 1$, i.e. that there is at most one simple root α_i such that $(\alpha, \alpha_i^\vee) = -2$. Suppose not. Recall that $(\alpha_i, \alpha_j^\vee) \leq 0$ if $i \neq j$. Hence $\alpha \neq -\alpha_1$. Clearly $\alpha \neq \alpha_1$, and as we are assuming that $(\alpha, \alpha_1) = -2 < 0$, it follows that $\alpha + \alpha_1$ is a root, and $(\alpha, \alpha) = 2(\alpha_1, \alpha_1)$, by a standard fact about root systems. Next observe that $(\alpha+\alpha_1, \alpha_2^\vee) = -2 + (\alpha_1, \alpha_2^\vee)$ because we are assuming $s \geq 2$. As G is not of type G_2, $(\alpha_1, \alpha_2^\vee) = 0$ and therefore $(\alpha_1, \alpha_2) = 0$. As before, we find that $\alpha \neq -\alpha_2$, and $(\alpha, \alpha) = 2(\alpha_2, \alpha_2)$. Now, $(\alpha, \alpha_1) = (\alpha, \alpha_1^\vee)(\alpha_1, \alpha_1)/2 = -(\alpha_1, \alpha_1)$, and similarly $(\alpha, \alpha_2) = -(\alpha_2, \alpha_2)$. Thus

$$(\alpha+\alpha_1+\alpha_2, \alpha+\alpha_1+\alpha_2) = (\alpha, \alpha) + (\alpha_1, \alpha_1) + (\alpha_2, \alpha_2) + 2(\alpha, \alpha_1) + 2(\alpha, \alpha_2) = 0,$$

so that $\alpha = -\alpha_1 - \alpha_2$. But then α cannot possibly be a root, for $-\alpha = \alpha_1 + \alpha_2$ is a sum of orthogonal simple roots. This contradiction proves that $s \leq 1$.

If $s < 1$, then α is dominant and $q_\alpha = 0$. So suppose that $s = 1$. If $\alpha = -\alpha_1$, then clearly $q_\alpha = 1$. If not, then $s_1(\alpha+\rho) = \alpha+\rho+\alpha_1$, so that $(s_1(\alpha+\rho), \alpha_i^\vee) = (\alpha, \alpha_i^\vee) + 1 + (\alpha_1, \alpha_i^\vee)$. If $i = 1$, this number is 1; if $i > 1$, this number is ≥ 0. Indeed, $(\alpha, \alpha_i^\vee) \geq 0$ because $s = 1$. On the other hand, $(\alpha_1, \alpha_i^\vee) \geq -1$, for if $(\alpha_1, \alpha_i^\vee) = -2$, then $\|\alpha_1\| = \sqrt{2}\|\alpha_i\|$ (where $\|\cdot\| = \sqrt{(\cdot,\cdot)}$). Also, $(\alpha, \alpha) = 2(\alpha_1, \alpha_1)$ because $\alpha \neq -\alpha_1$, so that $\|\alpha\| = \sqrt{2}\|\alpha_1\|$. But then we have three root lengths $\|\alpha\| > \|\alpha_1\| > \|\alpha_2\|$, which is impossible in an irreducible root system. The ultimate conclusion is that $s_1(\alpha+\rho)$ is dominant, whence $q_\alpha = 1$. This completes the proof of the lemma. ■

We conclude that $H^q(X, \mathcal{W}_{i+1}) = 0$ for $q \geq 2$ and $i = 0, \ldots, n-1$. Consequently, the long exact sequence (2.5.2) implies that $H^q(X, \mathcal{V}_i) = 0$ for $q \geq 2$ and $i = 0, \ldots, n$. In particular, $H^q(X, \mathcal{V}) = H^q(X, \mathcal{V}_n) = 0$ for $q \geq 2$, as desired. ■

By combining this theorem with the Frölicher–Nijenhuis theorem, we obtain:

Corollary 2.5.12. *The complex structure of G/P is locally rigid.* ■

In view of Theorem 2.5.1, this corollary can be rephrased as:

Corollary 2.5.13. *The complex structure of a compact, simply connected, homogeneous Kähler manifold is locally rigid.* ■

In particular, the complex structure of such spaces as $\mathbb{P}^n$, $\mathrm{Gr}(d, \mathbb{C}^n)$, $\mathcal{F}l(\mathbb{C}^n)$, etc. is locally rigid.

We also record the following interesting corollary of the proof of Theorem 2.5.10.

Proposition 2.5.14. *Let $\mathcal{V}_\varphi = G \times_\varphi V$ be the G-equivariant vector bundle on $X = G/P$ corresponding to a representation $\varphi\colon P \to \mathrm{GL}(V)$, and set $Q_\varphi = \{q_\lambda \colon \lambda$ is a highest weight of V and $\lambda + \rho$ is regular$\}$. Then $H^q(X, \mathcal{V}_\varphi) = 0$ if $q \notin Q_\varphi$.* ■

2.5.3 Weyl's Dimension Formula

We remarked in Section 2.5.1 that in order for the Borel–Weil–Bott theorem to be an effective tool for computing $\dim H^q(X, \mathcal{V})$, a knowledge of the dimensions of the irreducible G-modules seems necessary. We mentioned that the Weyl dimension formula proves useful in this regard. In this section we sketch an argument of how this formula can be deduced from the Borel–Weil–Bott theorem with the aid of the Hirzebruch–Riemann–Roch formula. The approach we have in mind is that of Borel–Hirzebruch [11, §22] and ultimately relies on the Weyl character formula. In fact, the Weyl character formula itself follows from a more refined study of the geometry of the flag variety; what is needed is a more sophisticated version of the Riemann–Roch formula: see Köck [47]. An alternative proof of the character formula will be given in Chapter 3.

Let V^λ denote the irreducible G-module of highest weight λ.

Theorem 2.5.15 (Weyl's Dimension Formula)**.**

$$\dim V^\lambda = \prod_{\alpha \in \Phi^+} \frac{(\lambda + \rho, \alpha)}{(\rho, \alpha)}.$$

Proof (sketch). Set $X = G/B^-$ and let $\lambda \in \Lambda$ be a dominant weight. We have

$$\chi(X, \mathcal{L}_\lambda) = \sum_{i \geq 0} (-1)^i \dim H^i(X, \mathcal{L}_\lambda)^i = \dim H^0(X, \mathcal{L}_\lambda) = \dim V^\lambda$$

by the Borel–Weil–Bott theorem. The Hirzebruch–Riemann–Roch formula then yields

$$\dim V^\lambda = \int_X \operatorname{ch}(\mathcal{L}_\lambda) \operatorname{td}(X),$$

where $\operatorname{ch}(\mathcal{L}_\lambda)$ is the Chern character of $\mathcal{L}_\lambda$ and $\operatorname{td}(X)$ is the Todd class of X. The Chern roots are the weights of the adjoint representation on $\mathfrak{g}/\mathfrak{b}^-$—that is, they are the positive roots Φ^+. Under the identification $\widehat{H} \cong \Lambda$: $\lambda \leftrightarrow e^\lambda$, the first Chern class of $\mathcal{L}_\lambda$ is $c_1(\mathcal{L}_\lambda) = e^\lambda$. Thus,

$$\dim V^\lambda = \int_X e^\lambda \prod_{\alpha \in \Phi^+} \frac{\alpha}{e^{\alpha/2} - e^{-\alpha/2}}.$$

By fibre integration, Borel and Hirzebruch show that this last expression is equal to the constant term of

$$\frac{\sum_{w \in W} \operatorname{sgn}(w) e^{w(\lambda+\rho)}}{\prod_{\alpha \in \Phi^+} e^{\alpha/2} - e^{-\alpha/2}},$$

which, by the Weyl character formula, is simply

$$\prod_{\alpha \in \Phi^+} \frac{(\lambda + \rho, \alpha)}{(\rho, \alpha)},$$

as desired. ■

Remark 2.5.16. It would be nice to have a simple geometric proof of the Weyl dimension formula that does not rely on the Weyl character formula. As far as I can ascertain, no such proof is known. ▲

2.6 Generalizing to Algebraic Groups over Arbitrary Fields

As was mentioned in Chapter 1, a complex semisimple Lie group is at the same time a semisimple linear algebraic group over $\mathbb{C}$. Thus it seems natural to wonder how, if at all, the main results of this chapter adapt to the setting of semisimple algebraic groups over a general field k. In this section we briefly delve into this line of inquiry. In order for our situation to not be completely hopeless, we shall assume that k is algebraically closed. It then turns out that theory has two distinct manifestations, depending on whether the characteristic of k is zero or $p > 0$.

Before proceeding, let us fix some notation. Let G be a semisimple algebraic group over k, $B \subset G$ a Borel subgroup, $T \subset B$ a maximal torus so that $B = TU$ where U is the unipotent radical of B, $W = N_G(T)/T$ the Weyl group, $X(T)$ the character group of T (which consists of all morphisms $T \to k^\times$) and $Y(T)$ the cocharacter group (which consists of all morphisms $k^\times \to T$). Denote by $(\cdot,\cdot)$ the pairing $X(T) \otimes_{\mathbb{Z}} Y(T) \to \mathbb{Z}$. Let $\Phi \subset X(T)$ denote the set of roots, i.e., the nonzero weights of T in Lie(G). Let $\alpha^\vee$ denote the coroot of $\alpha \in \Phi$. Fix a subset R^+ of positive roots, such that the weights of T in Lie(B) are negative, and let $\Delta \subset \Phi^+$ denote the base of simple roots. Let $X(T)_+$ denote the set of dominant weights, i.e. those $\lambda \in X(T)$ such that $(\lambda, \alpha^\vee) \geq 0$ for all $\alpha \in \Delta$ (or equivalently, for all $\alpha \in \Phi^+$). Finally, set $\rho = \sum_{\alpha \in R^+} \alpha \in \mathbb{Q} \otimes_{\mathbb{Z}} X(T)$. (The point is: most of the familiar objects from $k = \mathbb{C}$ also exist for general k. For the proper definitions, see Borel [9], Humphreys [40] or Springer [65].)

If $\operatorname{char} k = 0$, then the situation is relatively pleasant and the representation theory of G is well-understood. In particular, every rational G-module decomposes into a direct sum of irreducible submodules, the irreducible G-modules are classified by their highest weights, and the possible highest weights of irreducible G-modules lie in one-to-one correspondence with $X(T)_+$. Thus the situation is much in the same spirit of the Cartan–Weyl theory over $\mathbb{C}$. On the other hand, if $\operatorname{char} k > 0$, then there exists for every G a rational G-module which does not decompose into irreducibles. Remarkably, however, Chevalley has shown that the irreducible G-modules are still classified by their highest weights and that, moreover, the possible highest weights are in bijection with $X(T)_+$. We can therefore associate to each $\lambda \in X(T)_+$ an irreducible G-module $L(\lambda)$—and this makes sense in any characteristic.

A natural question at this point is whether there exists a Borel–Weil construction of $L(\lambda)$ ($\lambda \in X(T)_+$). Just as in the complex case, one can associate to each $\lambda \in X(T)$ a line bundle $\mathcal{L}_\lambda$ over the flag variety G/B. We let $H^i(\lambda)$ ($i \geq 0$) denote the corresponding cohomology groups. Then one has the following analogue of the Borel–Weil theorem:

Theorem 2.6.1 (Borel–Weil)**.** *$H^0(\lambda) \neq 0$ if and only if $\lambda \in X(T)_+$. If $\lambda \in X(T)_+$, then $L(\lambda)$ is the unique irreducible submodule of $H^0(\lambda)$; if $\operatorname{char} k = 0$ then $H^0(\lambda) = L(\lambda)$, but in positive characteristic it is possible that $H^0(\lambda) \neq L(\lambda)$.* ■

In fact, the second proof of the Borel–Weil theorem given in this chapter (see p.38) carries over, with minor modifications, to give the result in characteristic zero.

What about the Borel–Weil–Bott theorem?

Theorem 2.6.2 (Borel–Weil–Bott)**.** *Assume $\operatorname{char} k = 0$ and fix $\lambda \in X(T)$.*

(i) *If $\lambda + \rho$ is singular, then $H^i(\lambda) = 0$ for all $i \geq 0$.*

(ii) *If $\lambda + \rho$ is regular, then then $H^i(\lambda) = 0$ for all $i \neq \ell(w)$, where w is the unique element of W such that $w(\lambda + \rho) \in X(T)_+$, and $H^{\ell(w)}(\lambda) = H^0(w(\lambda + \rho) - \rho) = L(w(\lambda + \rho) - \rho)$.* ■

This can be proved using Kodaira's vanishing theorem (see also Kempf's vanishing theorem, below) just like in the complex case.

In positive characteristic the Borel–Weil–Bott theorem fails miserably. First, as we have already mentioned, there are nonzero G-modules $H^i(\lambda)$ which are not irreducible. Second, it is possible to have $H^i(\lambda) \neq 0$ in multiple degrees i. However, at least the following theorem holds in any characteristic.

Theorem 2.6.3 (Kempf's Vanishing Theorem)**.** *If $\lambda \in X(T)_+$ then $H^i(\lambda) = 0$ for all $i > 0$.* ■

In characteristic zero this is a consequence of Kodaira's vanishing theorem. The meat of this theorem is in the case of prime characteristic (where, as is well-known, there is no analogue of Kodaira's vanishing theorem)—this is the case that Kempf [43] handled and the reason the theorem bears his name.

There is much more that can be said here: a good reference for further reading is Jantzen [42].

Chapter 3

Lie Algebra Cohomology and the Theorems of Casselman–Osborne and Kostant

3.1 Introduction

Our focus in this chapter is not on Lie groups, but rather on (complex semisimple) Lie algebras. The first section introduces Lie algebra cohomology. This is a vast topic, and we have confined ourselves to the little part of it which is needed for stating and proving the two principal theorems of the chapter: the Casselman–Osborne theorem (Theorem 3.3.1) and Kostant's theorem (Theorem 3.4.1).

Fix a complex semisimple Lie algebra $\mathfrak{g}$, a Cartan subalgebra $\mathfrak{h} \subset \mathfrak{g}$ and a Borel subalgebra $\mathfrak{b} \supset \mathfrak{h}$, and let $\mathfrak{n}$ denote the nilradical of $\mathfrak{b}$. The Casselman–Osborne theorem describes the compatibility of two special actions, one by $\mathfrak{h}$ and the other by the centre $\mathcal{Z}(\mathfrak{g})$ of the universal enveloping algebra of $\mathfrak{g}$, on the Lie algebra cohomology space $H^p(\mathfrak{n}, V)$ with coefficients in an irreducible $\mathfrak{g}$-module V. A consequence is Kostant's theorem, which asserts that $H^p(\mathfrak{n}, V)$ decomposes, as an $\mathfrak{h}$-module, into one-dimensional weight spaces in a rather special, multiplicity-free manner.

To demonstrate the power of these results, we give as a first application another proof of the Borel–Weil–Bott theorem (Section 3.5.1). The relevant sheaf cohomology is computed by means of a Lie algebraic cochain complex, reducing the gist of the argument to a calculation of weight multiplicities, which Kostant's theorem handles readily.

Another application is Kostant's beautiful proof of the Weyl character formula, in which the formula is expressed as a ratio of Euler characteristics in Lie algebra cohomology (Section 3.5.2).

3.2 Lie Algebra Cohomology

3.2.1 The Definition of $H^p(\mathfrak{g}, V)$

Fix a (finite-dimensional) complex Lie algebra $\mathfrak{g}$ and a $\mathfrak{g}$-module (V, φ), which we do not assume to be finite-dimensional. For $p = 0, 1, 2, \ldots$, set

$$C^p(\mathfrak{g}, V) = \mathrm{Hom}_{\mathbb{C}}(\textstyle\bigwedge^p \mathfrak{g}, V).$$

Note that $C^0(\mathfrak{g}, V) = V$. We call $C^p(\mathfrak{g}, V)$ the space of V-valued p-cochains on $\mathfrak{g}$, or just p-**cochains** for short. We may think of a p-cochain $\omega \in C^p(\mathfrak{g}, V)$ as a p-linear alternating map $\omega\colon \mathfrak{g} \times \cdots \times \mathfrak{g} \to V$.

For $\omega \in C^p(\mathfrak{g}, V)$ and $X_1, \ldots, X_{p+1} \in \mathfrak{g}$ define

$$\begin{aligned} d\omega(X_1, \ldots, X_{p+1}) \;=\; & \sum_{j=1}^{p+1} (-1)^{j+1} \varphi(X_j) \omega(X_1, \ldots, \widehat{X_j}, \ldots, X_{p+1}) \\ & + \sum_{k<l} (-1)^{k+l} \omega([X_k, X_l], X_1, \ldots, \widehat{X_k}, \ldots, \widehat{X_l}, \ldots, X_{p+1}), \end{aligned}$$

where a hat over an argument means that argument is omitted. It can be easily shown that $d\omega \in C^{p+1}(\mathfrak{g}, V)$, and we therefore obtain a linear map

$$d\colon C^p(\mathfrak{g}, V) \to C^{p+1}(\mathfrak{g}, V),$$

called the **coboundary operator**.

We aim to prove that

$$C^0(\mathfrak{g},V) \xrightarrow{d} C^1(\mathfrak{g},V) \xrightarrow{d} \cdots$$

is a cochain complex. What we must show is that $d^2 = 0$. A direct verification is possible, but very messy. So we will take an alternative route.

We begin by observing that there is an action of $\mathfrak{g}$ on $C^p(\mathfrak{g},V) = \mathrm{Hom}(\bigwedge^p \mathfrak{g}, V)$, namely Hom of $\bigwedge^p \mathrm{ad}$ and φ. Explicitly, if this action is denoted by θ_X for $X \in \mathfrak{g}$, we have

$$(\theta_X \omega)(X_1,\dots,X_p) = \varphi(X)\omega(X_1,\dots,X_p) - \sum_{j=1}^{p} \omega(X_1,\dots,[X,X_j],\dots,X_p)$$

for $\omega \in C^p(\mathfrak{g},V)$. We call $\theta_X\omega$ the **Lie derivative** of ω relative to X (cf. Lemma 3.2.2 to see why).

We also define, for $X \in \mathfrak{g}$, the **interior product**

$$i_X \colon C^p(\mathfrak{g},V) \to C^{p-1}(\mathfrak{g},V)$$

by

$$(i_X\omega)(X_1,\dots,X_{p-1}) = \omega(X,X_1,\dots,X_{p-1}).$$

We now prove some technical lemmas.

Lemma 3.2.1. $[\theta_X, i_Y] = i_{[X,Y]}$ *for all* $X, Y \in \mathfrak{g}$.

Proof. Let $\omega \in C^p(\mathfrak{g},V)$ and $X_1,\dots,X_p \in \mathfrak{g}$. Then

$$(\theta_X i_Y \omega)(X_1,\dots,X_p) = \varphi(X)\omega(Y,X_1,\dots,X_p) - \sum_{j=1}^{p} \omega(Y,X_1,\dots,[X,X_j],\dots,X_p)$$

and

$$\begin{aligned}(i_Y\theta_X\omega)(X_1,\dots,X_p) &= \varphi(X)\omega(Y,X_1,\dots,X_p) - \sum_{j=1}^{p} \omega(Y,X_1,\dots,[X,X_j],\dots,X_p) \\ &\quad - \omega([X,Y],X_1,\dots,X_p).\end{aligned}$$

Subtracting the second equation from the first, we obtain

$$([\theta_X, i_Y]\omega)(X_1,\dots,X_p) = (i_{[X,Y]}\omega)(X_1,\dots,X_p),$$

as required. ■

Lemma 3.2.2 (Cartan's Formula)**.** *$\theta_X = i_X \circ d + d \circ i_X$ for all $X \in \mathfrak{g}$.*

Proof. Let $\omega \in C^p(\mathfrak{g}, V)$ and $X_0, \ldots, X_p \in \mathfrak{g}$. Then

$$\begin{aligned} i_{X_0} d\omega(X_1, \ldots, X_p) &= \sum_{j=0}^{p} (-1)^j \varphi(X_j)\omega(X_0, \ldots, \widehat{X_j}, \ldots, X_p) \\ &\quad + \sum_{k<l} (-1)^{l+k} \omega([X_k, X_l], X_0, \ldots, \widehat{X_k}, \ldots, \widehat{X_l}, \ldots, X_p). \end{aligned}$$

and

$$\begin{aligned} d i_{X_0} \omega(X_1, \ldots, X_p) &= \sum_{j=1}^{p} (-1)^{j+1} \varphi(X_j)\omega(X_0, X_1, \ldots, \widehat{X_j}, \ldots, X_p) \\ &\quad + \sum_{k<l} (-1)^{l+k} \omega(X_0, [X_k, X_l], X_1, \ldots, \widehat{X_k}, \ldots, \widehat{X_l}, \ldots, X_p). \end{aligned}$$

By adding these two equations together, we get

$$\begin{aligned} (i_{X_0} d\omega + d i_{X_0} \omega)(X_1, \ldots, X_p) &= \varphi(X_0)\omega(X_1, \ldots, X_p) \\ &\quad + \sum_{k=1}^{p} (-1)^k \omega([X_0, X_k], X_1, \ldots, \widehat{X_k}, \ldots, X_p). \end{aligned}$$

Using the fact that ω is alternating, the right-hand side reduces to $(\theta_{X_0}\omega)(X_1, \ldots, X_p)$, which is what we want. ∎

Lemma 3.2.3. *$d \circ \theta_X = \theta_X \circ d$ for all $X \in \mathfrak{g}$.*

Proof. Let $T = d \circ \theta_X - \theta_X \circ d$; we wish to show that $T = 0$ on $C^p(\mathfrak{g}, V)$. Note that if $\omega \in C^p(\mathfrak{g}, V)$ then $\omega(X_1, \ldots, X_p) = i_{X_p} \cdots i_{X_1} \omega$. Thus T is completely determined by its action on $C^0(\mathfrak{g}, V) = V$ and how it commutes with i_Y ($Y \in \mathfrak{g}$). If $v \in C^0(\mathfrak{g}, V)$, then

$$(d \circ \theta_X v)(Y) = (d(\varphi(X)v))(Y) = \varphi(Y)\varphi(X)v$$

and

$$(\theta_X dv)(Y) = \varphi(X)dv(Y) - dv([X, Y]) = \varphi(X)\varphi(Y)v - \varphi([X, Y])v.$$

Thus $Tv = 0$.

Next, let $Y \in \mathfrak{g}$. Then

$$i_Y d\theta_X = \theta_Y\theta_X - di_Y\theta_X = \theta_Y\theta_X - d\theta_X i_Y + di_{[X,Y]}$$

and

$$i_Y\theta_X d = \theta_X i_Y d - i_{[X,Y]}d = \theta_X\theta_Y - \theta_X di_Y - i_{[X,Y]}d$$

by the previous two lemmas. It follows that $i_Y T = -Ti_Y$ and consequently that $T = 0$ on $C^p(\mathfrak{g}, V)$, as desired. ∎

Finally:

Proposition 3.2.4. $d^2 = 0$. *Thus $(C^\bullet(\mathfrak{g}, V), d)$ is a cochain complex.*

Proof. We proceed as we did in the proof of the previous lemma. Fix $v \in V$. Then

$$\begin{aligned}(d^2v)(X,Y) &= \varphi(X)dv(Y) - \varphi(Y)dv(X) - dv([X,Y]) \\ &= \varphi(X)\varphi(Y)v - \varphi(Y)\varphi(X)v - \varphi([X,Y])v = 0.\end{aligned}$$

Next let $X \in \mathfrak{g}$. By combining the previous two lemmas, we find that

$$d^2 i_X + di_X d = d\theta_X = \theta_X d = i_X d^2 + di_X d$$

whence $d^2 i_X = i_X d^2$. Hence $d^2 = 0$. ∎

As usual, we now define $Z^p(\mathfrak{g}, V)$ to be the kernel of of d in $C^p(\mathfrak{g}, V)$ and $B^p(\mathfrak{g}, V)$ to be the image under d of $C^{p-1}(\mathfrak{g}, V)$; these two spaces are called the spaces of p-**cocycles** and p-**coboundaries**, respectively. As $d^2 = 0$, $B^p(\mathfrak{g}, V) \subseteq Z^p(\mathfrak{g}, V)$, and we may thus define the pth **Lie algebra cohomology space** of $\mathfrak{g}$ with coefficients in V as

$$H^p(\mathfrak{g}, V) = Z^p(\mathfrak{g}, V)/B^p(\mathfrak{g}, V).$$

This is a complex vector space.

Remarks 3.2.5.

(i) If $p > \dim \mathfrak{g}$, then $\bigwedge^p \mathfrak{g} = 0$ and consequently $H^p(\mathfrak{g}, V) = 0$ for all V.

(ii) $H^0(\mathfrak{g}, V) = Z^0(\mathfrak{g}, V) = \{v \in V : dv = 0\}$. By definition, $dv(X) = \varphi(X)v$ $(v \in V)$. Thus $H^0(\mathfrak{g}, V) = V^{\mathfrak{g}}$, the space of $\mathfrak{g}$-invariants in V. In particular, the induced action of $\mathfrak{g}$ on $H^0(\mathfrak{g}, V)$ is trivial.

(iii) In fact, the induced action $\mathfrak{g}$ on $H^p(\mathfrak{g}, V)$ is trivial for all p. Indeed we see from Cartan's formula (Lemma 3.2.2) that θ_X maps $Z^p(\mathfrak{g}, V)$ into $B^p(\mathfrak{g}, V)$. ▲

3.2.2 Functoriality

Suppose we are given a map $\alpha : U \to V$ of $\mathfrak{g}$-modules. Then we can extend α to a map $C^p(\mathfrak{g}, U) \to C^p(\mathfrak{g}, V)$ by defining

$$\alpha(\omega)(X_1, \ldots, X_p) = \alpha(\omega(X_1, \ldots, X_p)) \qquad (\omega \in C^p(\mathfrak{g}, U)).$$

It is then clear that $\alpha \circ d = d \circ \alpha$. Consequently, α descends to a map on cohomology, which we write as

$$\alpha^* : H^p(\mathfrak{g}, U) \to H^p(\mathfrak{g}, V) \qquad (p = 0, 1, 2, \ldots).$$

If $\beta : V \to W$ is another map of $\mathfrak{g}$-modules, then it is trivial to check that $(\beta \circ \alpha)^* = \beta^* \circ \alpha^*$.

In summary:

Proposition 3.2.6. *The assignment $V \rightsquigarrow H^p(\mathfrak{g}, V)$ is a covariant functor from the category of $\mathfrak{g}$-modules to the category of complex vector spaces.* ■

In fact, the functors $H^p(\mathfrak{g}, -)$ $(p = 0, 1, \ldots)$ form what is known as a δ-functor. The point—as far as we're concerned—is that every short exact sequence of $\mathfrak{g}$-modules

$$0 \longrightarrow U \longrightarrow V \longrightarrow W \longrightarrow 0$$

induces a long exact sequence

$$\begin{aligned} 0 \quad \longrightarrow \quad & H^0(\mathfrak{g}, U) \longrightarrow H^0(\mathfrak{g}, V) \longrightarrow H^0(\mathfrak{g}, W) \xrightarrow{\delta} H^1(\mathfrak{g}, U) \longrightarrow \cdots \\ \cdots \quad & \longrightarrow H^p(\mathfrak{g}, U) \longrightarrow H^p(\mathfrak{g}, V) \longrightarrow H^p(\mathfrak{g}, W) \xrightarrow{\delta} H^{p+1}(\mathfrak{g}, U) \longrightarrow \cdots \end{aligned}$$

in cohomology. The existence of the connecting maps δ follows from the exactness of

$$0 \longrightarrow C^p(\mathfrak{g}, U) \longrightarrow C^p(\mathfrak{g}, V) \longrightarrow C^p(\mathfrak{g}, W) \longrightarrow 0$$

and the usual diagram chase (see also Proposition 3.2.9). The other thing we need to check is the exactness of the sequence

$$0 \longrightarrow H^0(\mathfrak{g}, U) \longrightarrow H^0(\mathfrak{g}, V) \longrightarrow H^0(\mathfrak{g}, W).$$

In other words, we must show that $H^0(\mathfrak{g}, -)$ is left exact. But recall that $H^0(\mathfrak{g}, -)$ is the $\mathfrak{g}$-invariants functor (Remark 3.2.5(ii)), so this is obvious. In fact, one can show that $H^p(\mathfrak{g}, -)$ is the pth right derived functor of $H^0(\mathfrak{g}, -)$, and the δ-functoriality claim follows; indeed, $H^p(\mathfrak{g}, V) \cong \mathrm{Ext}^p_{\mathcal{U}(\mathfrak{g})}(\mathbb{C}, V)$, where $\mathbb{C}$ is regarded as the trivial $\mathcal{U}(\mathfrak{g})$-module.

3.2.3 $\mathfrak{n}$-cohomology

Fix a complex Lie algebra $\mathfrak{b}$ and an ideal $\mathfrak{n}$ of $\mathfrak{b}$. (Our primary interest is when $\mathfrak{b}$ is a Borel subalgebra and $\mathfrak{n}$ its nilradical. Note that in other chapters we have denoted the nilradical by $\mathfrak{u}$; in the context of Lie algebra cohomology, the notation $\mathfrak{n}$ is more standard.) Let (V, φ) be a $\mathfrak{b}$-module, which we also view as an $\mathfrak{n}$-module by restriction. If $X \in \mathfrak{b}$ and $\omega \in C^p(\mathfrak{n}, V)$, we define $\theta_X \omega \in C^p(\mathfrak{n}, V)$ by

$$(\theta_X \omega)(X_1, \ldots, X_p) = \varphi(X)\omega(X_1, \ldots, X_p) - \sum_{j=1}^{p} \omega(X_1, \ldots, [X, X_j], \ldots, X_p).$$

This definition makes sense because $[\mathfrak{b}, \mathfrak{n}] \subseteq \mathfrak{n}$. We can thus view $C^p(\mathfrak{n}, V)$ as a $\mathfrak{b}$-module. From our work in the Section 3.2.1 (which trivially generalizes to this setting), we see that θ_X ($X \in \mathfrak{b}$) commutes with d and leaves $Z^p(\mathfrak{n}, V)$ and $B^p(\mathfrak{n}, V)$ invariant, whence it descends to a map on $H^p(\mathfrak{n}, V)$. As $\mathfrak{n}$ acts trivially on $H^p(\mathfrak{n}, V)$ (Remark 3.2.5(iii)), we conclude:

Proposition 3.2.7. *$H^p(\mathfrak{n}, V)$ is naturally a $\mathfrak{b}/\mathfrak{n}$-module.* ■

In view of our results in the previous section, we also have:

Corollary 3.2.8. *$H^p(\mathfrak{n}, -)$ is a functor from the category of $\mathfrak{b}$-modules to the category of $\mathfrak{b}/\mathfrak{n}$-modules.* ■

3.2.4 The Long Exact Sequence Revisited

Suppose now that $\mathfrak{g}$ is a complex semisimple Lie algebra with Cartan subalgebra $\mathfrak{h}$, and $\mathfrak{b} \supset \mathfrak{h}$ is a Borel subalgebra of $\mathfrak{g}$ with nilradical $\mathfrak{n}$. Then $\mathfrak{b} = \mathfrak{h} \oplus \mathfrak{n}$ and $\mathfrak{h} = \mathfrak{b}/\mathfrak{n}$, so, by what we have seen in the previous section, $H^p(\mathfrak{n}, V)$ is naturally an $\mathfrak{h}$-module for any $\mathfrak{b}$-module V. In fact, if V is a $\mathfrak{g}$-module, then by arguing along the same lines we see that $H^p(\mathfrak{n}, V)$ is an $\mathfrak{h}$-module.

In this case there is another important action on $H^p(\mathfrak{n}, V)$, this time by the centre $\mathcal{Z}(\mathfrak{g})$ of the universal enveloping algebra of $\mathfrak{g}$: if $z \in \mathcal{Z}(\mathfrak{g})$ and $\omega \in C^p(\mathfrak{n}, V)$, let

$$(z\omega)(X_1, \ldots, X_p) = z(\omega(X_1, \ldots, X_p)).$$

This action of $\mathcal{Z}(\mathfrak{g})$ on $C^p(\mathfrak{n}, V)$ clearly commutes with d and hence descends to an action on $H^p(\mathfrak{n}, V)$.

Proposition 3.2.9. *Let $0 \longrightarrow U \longrightarrow V \longrightarrow W \longrightarrow 0$ be a short exact sequence of $\mathfrak{g}$-modules. Then the maps in the long exact sequence*

$$\cdots \longrightarrow H^p(\mathfrak{n}, U) \longrightarrow H^p(\mathfrak{n}, V) \longrightarrow H^p(\mathfrak{n}, W) \xrightarrow{\delta} H^{p+1}(\mathfrak{n}, U) \longrightarrow \cdots$$

are maps of $\mathfrak{h}$-modules and $\mathcal{Z}(\mathfrak{g})$-modules.

Proof. The exactness of $0 \longrightarrow U \longrightarrow V \longrightarrow W \longrightarrow 0$ implies the exactness of $0 \longrightarrow C^p(\mathfrak{n}, U) \longrightarrow C^p(\mathfrak{n}, V) \longrightarrow C^p(\mathfrak{n}, W) \longrightarrow 0$ which in turn implies the assertions of the proposition for the maps that preserve degree. Thus we only have to worry about the connecting maps

$$H^p(\mathfrak{n}, W) \xrightarrow{\delta} H^{p+1}(\mathfrak{n}, U).$$

Let us recall how δ is defined. Fix a class $[\omega] \in H^p(\mathfrak{n}, W)$ with representative $\omega \in Z^p(\mathfrak{n}, W)$. As the natural map $C^p(\mathfrak{n}, V) \to C^p(\mathfrak{n}, W)$ is surjective, we can pick an element $\tau \in C^p(\mathfrak{n}, V)$ in the preimage of ω. Then $d\tau \in C^{p+1}(\mathfrak{n}, V)$; in fact, $d\tau \in C^{p+1}(\mathfrak{n}, U)$ (because its image under the natural map $C^{p+1}(\mathfrak{n}, V) \to C^{p+1}(\mathfrak{n}, W)$ is $d\omega = 0$), whence $d\tau \in Z^{p+1}(\mathfrak{n}, U)$ (because $d^2 = 0$). We define $\delta[\omega]$ to be $[d\tau]$.

Now if $z \in \mathcal{Z}(\mathfrak{g})$, then a representative of $z[\omega]$ is $z\omega$ and a preimage of $z\omega$ is $z\tau$. Then $\delta(z[\omega]) = \delta([z\omega]) = [dz\tau] = [zd\tau] = z(\delta[\omega])$. This proves that δ is a map of $\mathcal{Z}(\mathfrak{g})$-modules; that it is a map of $\mathfrak{h}$-modules is obvious. ■

3.2.5 Cohomology of Injective Modules

In this section we state two technical facts which are needed in the proof of the Casselman–Osborne theorem.

Proposition 3.2.10. *Let $\mathfrak{g}$ be a complex Lie algebra.*

(i) $\mathrm{Hom}_{\mathbb{C}}(\mathcal{U}(\mathfrak{g}), V)$ *is an injective* $\mathcal{U}(\mathfrak{g})$*-module for every* $\mathfrak{g}$*-module* V*.*

(ii) *Every* $\mathfrak{g}$*-module is a submodule of* $\mathrm{Hom}_{\mathbb{C}}(\mathcal{U}(\mathfrak{g}), V)$ *for some* V*.*

Proof. Both assertions follow from basic results in homological algebra; for example, if I is an injective S-module then one knows that $\mathrm{Hom}_S(R, I)$ is an injective R-module for any $R \supset S$—as V is an injective $\mathbb{C}$-module, (i) follows. Part (ii) is similarly a special case of a general homological theorem. ∎

Proposition 3.2.11. *Let $\mathfrak{g}$ be a complex Lie algebra, $\mathfrak{b}$ a subalgebra of $\mathfrak{g}$ and $\mathfrak{n}$ an ideal of $\mathfrak{b}$. If I is an injective* $\mathcal{U}(\mathfrak{g})$*-module then $H^p(\mathfrak{n}, I) = 0$ for $p > 0$.*

Proof. We have remarked that $H^p(\mathfrak{n}, -)$ is the pth right derived functor of the $\mathfrak{n}$-invariants functor. In particular, $H^p(\mathfrak{n}, -)$ is computed from injective resolutions. If I is an injective $\mathcal{U}(\mathfrak{g})$-module then, because $\mathcal{U}(\mathfrak{g})$ is a free $\mathcal{U}(\mathfrak{n})$-module (by the Poincaré–Birkhoff–Witt theorem), I is also an injective $\mathcal{U}(\mathfrak{n})$-module. Thus $0 \to I \to I \to 0$ is an injective resolution of I, and the proposition follows. ∎

3.3 The Casselman–Osborne Theorem

Let $\mathfrak{g}$ be a complex semisimple Lie algebra with Cartan subalgebra $\mathfrak{h}$ and let $\mathfrak{b} \supset \mathfrak{h}$ be a Borel subalgebra of $\mathfrak{g}$ with nilradical $\mathfrak{n}$. Then we have seen that $H^p(\mathfrak{n}, V)$ is naturally an $\mathfrak{h}$-module for any $\mathfrak{g}$-module V; we have also seen that $H^p(\mathfrak{n}, V)$ carries an action of $\mathcal{Z}(\mathfrak{g})$, the centre of the universal enveloping algebra $\mathcal{U}(\mathfrak{g})$ of $\mathfrak{g}$. The principal result of this section describes how these two actions are related.

Theorem 3.3.1 (Casselman–Osborne)**.** *Assume the preceding notation. If $z \in \mathcal{Z}(\mathfrak{g})$ and $\omega \in H^p(\mathfrak{n}, V)$ then $z\omega = \gamma(z)\omega$, where γ is the auxiliary Harish-Chandra map $\mathcal{Z}(\mathfrak{g}) \to \mathcal{U}(\mathfrak{h})$ (see Section 1.14).*

Remark 3.3.2. In the equation $z\omega = \gamma(z)\omega$, it is to be understood that the action on the left is that of $\mathcal{Z}(\mathfrak{g})$ while the action on the right is that of $\mathcal{U}(\mathfrak{h})$. ▲

Proof of Theorem 3.3.1. We proceed by induction on p. When $p = 0$, $H^p(\mathfrak{n}, V) = V^{\mathfrak{n}}$ as an $\mathfrak{h}$-module. So if $z \in \mathcal{Z}(\mathfrak{g})$, then $z - \gamma(z) \in \mathcal{U}(\mathfrak{g})\mathfrak{n}$, and it follows that $(z - \gamma(z))\omega = 0$ for all $\omega \in V^{\mathfrak{n}}$, as desired.

Next suppose that $p > 0$ and assume the result for degree $p-1$. By Proposition 3.2.10, we can find an injective $\mathcal{U}(g)$-module I that contains V as a submodule. This yields a short exact sequence of $\mathcal{U}(g)$-modules

$$0 \to V \to I \to Q \to 0$$

and hence a long exact sequence

$$\cdots \longrightarrow H^{p-1}(\mathfrak{n}, I) \longrightarrow H^{p-1}(\mathfrak{n}, Q) \xrightarrow{\delta} H^p(\mathfrak{n}, V) \longrightarrow H^p(\mathfrak{n}, I) \longrightarrow \cdots$$

as in Proposition 3.2.9. As I is injective, $H^p(\mathfrak{n}, I) = 0$ by Proposition 3.2.11. Thus, given $\omega \in H^p(\mathfrak{n}, V)$, we can find an $\widetilde{\omega} \in H^{p-1}(\mathfrak{n}, Q)$ such that $\delta(\widetilde{\omega}) = \omega$. By the inductive hypothesis, $z\widetilde{\omega} = \gamma(z)\widetilde{\omega}$ for all $z \in \mathcal{Z}(\mathfrak{g})$. Consequently, $z\omega = z\delta(\widetilde{\omega}) = \delta(z\widetilde{\omega}) = \delta(\gamma(z)\widetilde{\omega}) = \gamma(z)\delta(\widetilde{\omega}) = \gamma(z)\omega$, where we have used the fact that δ is a map of $\mathcal{Z}(\mathfrak{g})$- and $\mathcal{U}(\mathfrak{h})$-modules (see Proposition 3.2.9). ■

Remark 3.3.3. Theorem 3.3.1 was first proved by Casselman and Osborne in [16]; in fact, [16] contains a more general statement valid for reductive $\mathfrak{g}$ and parabolic subalgebras $\mathfrak{p} \supseteq \mathfrak{b}$. The proof presented above is taken from Knapp–Vogan [46, IV.10], where the proof of the more general statement can be found. We mention that the proof of this more general statement differs from the proof we have given only superficially. ▲

3.4 Kostant's Theorem

We continue using the notation of the previous section. Additionally, we fix a choice Φ^+ of positive roots and assume that $\mathfrak{n}$ (and hence $\mathfrak{b}$) are built from the positive root spaces. Also, letting W denote the Weyl group of the root system, we set

$$W_p = \{w \in W \colon \ell(w) = p\} \qquad (p = 0, 1, \ldots, \dim \mathfrak{n});$$

and for $w \in W$, we set

$$\Phi^+(w) = \{\alpha \in \Phi^+ \colon w^{-1}\alpha < 0\}$$

so that $\ell(w) = |\Phi^+(w)|$. Finally, we recall our old friend

$$\rho = \frac{1}{2} \sum_{\alpha \in \Phi^+} \alpha \in \Lambda \otimes_{\mathbb{Z}} \mathbb{Q}.$$

Our goal in this section is to prove:

Theorem 3.4.1 (Kostant's Theorem)**.** *Let V^λ be an irreducible $\mathfrak{g}$-module of highest weight λ. As an $\mathfrak{h}$-module,*

$$H^p(\mathfrak{n}, V^\lambda) = \bigoplus_{w \in W_p} \mathbb{C}_{w(\lambda+\rho)-\rho},$$

where $\mathbb{C}_{w(\lambda+\rho)-\rho}$ is the irreducible $\mathfrak{h}$-module of weight $w(\lambda + \rho) - \rho$.

Remark 3.4.2. Notice that $H^0(\mathfrak{n}, V^\lambda) = (V^\lambda)^{\mathfrak{n}}$ is the highest weight space of V^λ. As $W_0 = \{1\}$, Kostant's theorem implies that $(V^\lambda)^{\mathfrak{n}} = \mathbb{C}_\lambda$ as an $\mathfrak{h}$-module, as expected. ▲

The proof of Kostant's theorem will require several preliminary lemmas.

Lemma 3.4.3. *If $w \in W$, then*

$$\rho - w\rho = \sum_{\beta \in \Phi^+(w)} \beta.$$

Proof. Write

$$\rho = \frac{1}{2} \sum \{\alpha \in \Phi^+ \colon w^{-1}\alpha > 0\} + \frac{1}{2} \sum \{\alpha \in \Phi^+ \colon w^{-1}\alpha < 0\}$$

and

$$\begin{aligned} w\rho &= \frac{1}{2}w\sum\{\beta\in\Phi^+ \colon w^{-1}\beta>0\}+\frac{1}{2}w\sum\{\beta\in\Phi^+ \colon w^{-1}\beta<0\} \\ &= \frac{1}{2}\sum\{\beta\in\Phi^+ \colon w^{-1}\beta>0\}+\frac{1}{2}\sum\{\beta\in\Phi^- \colon w^{-1}\beta>0\} \\ &= \frac{1}{2}\sum\{\beta\in\Phi^+ \colon w^{-1}\beta>0\}-\frac{1}{2}\sum\{\beta\in\Phi^+ \colon w^{-1}\beta<0\}. \end{aligned}$$

Now subtract this last expression from that of ρ, and the lemma follows. ■

Lemma 3.4.4. *Let $\lambda \in \Lambda$ be dominant. Then, for all subsets $S \subseteq \Phi^+$ and $w \in W$,*

$$(w(\lambda+\rho)-\rho+\textstyle\sum_{\alpha\in S}\alpha, w(\lambda+\rho)-\rho+\sum_{\alpha\in S}\alpha) \geq (\lambda,\lambda),$$

with equality if and only if $S = \Phi^+$.

Proof. By the previous lemma we have

$$w(\lambda+\rho)-\rho+\sum_{\alpha\in S}\alpha = w\lambda+\sum_{\alpha\in S}\alpha-\sum_{\beta\in\Phi^+(w)}\beta.$$

Thus

$$\begin{aligned} & (w(\lambda+\rho)-\rho+\textstyle\sum_{\alpha\in S}\alpha, w(\lambda+\rho)-\rho+\sum_{\alpha\in S}\alpha)-(\lambda,\lambda) \\ =\ & (w(\lambda+\rho)-\rho+\textstyle\sum_{\alpha\in S}\alpha, w(\lambda+\rho)-\rho+\sum_{\alpha\in S}\alpha)-(w\lambda,w\lambda) \\ =\ & (\textstyle\sum_{\alpha\in S}\alpha-\sum_{\beta\in\Phi^+(w)}\beta, \sum_{\alpha\in S}\alpha-\sum_{\beta\in\Phi^+(w)}\beta)+2(w\lambda,\sum_{\alpha\in S}\alpha-\sum_{\beta\in\Phi^+(w)}\beta) \\ \geq\ & 2(w\lambda,\textstyle\sum_{\alpha\in S}\alpha-\sum_{\beta\in\Phi^+(w)}\beta). \end{aligned} \tag{$\clubsuit$}$$

Consider the expression $\sum_{\alpha\in S}\alpha-\sum_{\beta\in\Phi^+(w)}\beta$. By canceling off the roots that appear in the two sums, we are left with a sum of two types of roots. The first type is a root $\eta=\alpha$ that occurs in the first sum but not in the second. Then $w^{-1}\alpha>0$ and thus (recalling that λ is assumed to be dominant)

$$(w\lambda,\eta)=(\lambda,w^{-1}\eta)\geq 0 \qquad \text{and} \qquad (w\rho,\eta)=(\rho,w^{-1}\eta)>0. \tag{3.4.1}$$

The second type is a root $\eta=-\beta$ with β coming from the second sum. So $w^{-1}\eta>0$ and again the inequalities in (3.4.1) hold. It follows from the first inequality in (3.4.1) that the expression in (♣) is nonnegative, and this proves the inequality stated in the lemma. Note that the inequality in (♣) is an equality if and only if $\sum \eta = 0$, where the sum runs over the roots η mentioned above. But the second inequality in (3.4.1) implies that $\sum (w\rho, \eta) > 0$. So the sum $\sum \eta$ is 0 if and only if it is an empty sum—that is, if and only if $S = \Phi^+(w)$. ∎

Recall that $C^p(\mathfrak{n}, V^\lambda)$ is naturally an $\mathfrak{h}$-module.

Lemma 3.4.5. *The weight $w(\lambda + \rho) - \rho$ occurs in $C^p(\mathfrak{n}, V^\lambda)$ if and only if $p = \ell(w)$, in which case it occurs with multiplicity equal to* 1.

Proof. The weights μ occurring in $C^p(\mathfrak{n}, V^\lambda) = \mathrm{Hom}_{\mathbb{C}}(\bigwedge^p \mathfrak{n}, V^\lambda) = \bigwedge^p \mathfrak{n}^- \otimes V^\lambda$ are of the form

$$\mu = \eta - \sum_{\alpha \in S} \alpha,$$

where η is a weight of V^λ and S is a subset of Φ^+ of size p. Furthermore, the multiplicity of μ is equal to the number of such expressions, counted with the multiplicity of each occurring η.

In particular, if $w(\lambda + \rho) - \rho$ occurs in $C^p(\mathfrak{n}, V^\lambda)$ then there is a weight η of V^λ and a subset $S \subseteq \Phi^+$ of size p such that

$$w(\lambda + \rho) - \rho = \eta - \sum_{\alpha \in S} \alpha.$$

Lemma 3.4.4 then implies that $(\lambda, \lambda) \leq (\eta, \eta)$. But λ is the highest weight of V^λ, so $(\eta, \eta) \leq (\lambda, \lambda)$, and therefore $(\lambda, \lambda) = (\eta, \eta)$. Thus $S = \Phi^+(w)$, again by Lemma 3.4.4. It follows that $\sum_{\alpha \in S} \alpha = \rho - w\rho$ by Lemma 3.4.3, and the above expression reduces to $w\lambda = \eta$.

In conclusion, the weight $w(\lambda + \rho) - \rho$ occurs in $C^p(\mathfrak{n}, V^\lambda)$ if and only if $p = |S| = |\Phi^+(w)| = |\ell(w)|$, in which case it occurs with multiplicity equal to the multiplicity of $\eta = w\lambda$ in V^λ, which is 1. ∎

It follows that $\mathbb{C}_{w(\lambda+\rho)-\rho}$ occurs as a subspace of $C^p(\mathfrak{n}, V^\lambda)$ only in degree $p = \ell(w)$. The coboundary operator d maps $\mathbb{C}_{w(\lambda+\rho)-\rho}$ into $C^{\ell(w)+1}(\mathfrak{n}, V^\lambda)$, thereby annihilating it.

Furthermore, as the multiplicity of $\mathbb{C}_{w(\lambda+\rho)-\rho}$ in $C^{\ell(w)}(\mathfrak{n}, V^\lambda)$ is equal to 1, we see that $\mathbb{C}_{w(\lambda+\rho)}$ is not in the image of d applied to $C^{\ell(w)-1}(\mathfrak{n}, V^\lambda)$. Consequently, $\mathbb{C}_{w(\lambda+\rho)-\rho}$ occurs as a subspace of $H^p(\mathfrak{n}, V^\lambda)$ only if $p = \ell(w)$, in which case it occurs with multiplicity equal to 1.

Proof of Kostant's theorem. All that remains is to show that no other weights occur in $H^p(\mathfrak{n}, V^\lambda)$, that is, besides those of the form $w(\lambda+\rho)-\rho$ for $w \in W_p$. To this end, suppose that μ is a weight of $H^p(\mathfrak{n}, V^\lambda)$. Recall that the action of $\mathcal{Z}(\mathfrak{g})$ on $C^p(\mathfrak{n}, V^\lambda)$ is given by

$$(z\omega)(X) = z(\omega(X)) \qquad (z \in \mathcal{Z}(\mathfrak{g}), \omega \in C^p(\mathfrak{n}, V^\lambda), X \in \textstyle\bigwedge^p \mathfrak{n}).$$

As $\omega(X) \in V^\lambda$ and V^λ is irreducible, the right-hand side above is $\lambda(\gamma(z))\omega(X)$, where γ is the auxiliary Harish-Chandra character (see Section 1.14). We conclude that $z \in \mathcal{Z}(\mathfrak{g})$ acts on $H^p(\mathfrak{n}, V^\lambda)$ by the scalar $\lambda(\gamma(z))$.

Now let $\omega \in H^p(\mathfrak{n}, V^\lambda)$ be a weight vector of weight μ. Then $h \in \mathcal{U}(\mathfrak{h})$ acts on ω by the scalar $\mu(h)$.[1] Thus the Casselman–Osborne theorem (Theorem 3.3.1) implies that

$$\lambda(\gamma(z))\omega = z\omega = \gamma(z)\omega = \mu(\gamma(z))\omega \qquad \text{for all } z \in \mathcal{Z}(\mathfrak{g}).$$

So the scalars $\lambda(\gamma(z))$ and $\mu(\gamma(z))$ are equal for all $z \in \mathcal{Z}(\mathfrak{g})$, which implies the equality of the infinitesimal Harish-Chandra characters $\chi_{\lambda+\rho}$ and $\chi_{\mu+\rho}$. This in turn implies that $\lambda+\rho$ and $\mu+\rho$ lie in the same W-orbit, say $\mu+\rho = w(\lambda+\rho)$ for some $w \in W$. In fact, $w \in W_p$ by Lemma 3.4.5, and we're done. ∎

Remarks 3.4.6.

(i) A version of Theorem 3.4.1 appeared in Bott [13, §15]. Bott's proof was, of course, of a differential-geometric nature. The first algebraic proof of Theorem 3.4.1 was worked out—in a more general setting—by Kostant in [48]. Kostant acknowledges that his paper [48] (and its follow-up [49]) originated out of attempts to explain Bott's "strange" discovery [13, 15.3] that

$$\dim H^p(\mathfrak{n}, V^\lambda) = \dim H^{2p}(G/B, \mathbb{C}) = |W_p| \qquad \text{for all dominant } \lambda.$$

[1] Here $\mu \in \mathfrak{h}^*$ is extended to $\mathbb{Z}[\mathfrak{h}]$ in the obvious—and usual—manner: define $\mu(1) = 1$ and extend $\mathbb{Z}$-linearly.

In particular, $\dim H^p(\mathfrak{n}, V^\lambda)$ does not depend on λ. The equality

$$\dim H^{2p}(G/B, \mathbb{C}) = |W_p|$$

is easily established (using the Bruhat decomposition of G/B, for example). The point was to explain the equality

$$\dim H^p(\mathfrak{n}, V^\lambda) = |W_p|$$

in a more algebraic manner (as opposed to Bott's geometric one, via his generalized Borel–Weil theorem). This is of course easily accomplished by Kostant's theorem. Additionally, Kostant observed that his theorem was of a sufficiently robust character so as to yield a new, algebraic proof of the Borel–Weil–Bott theorem: see Section 3.5.1.

(ii) In [48] Kostant shows how to locate the one-dimensional subspace of $H^p(\mathfrak{n}, V^\lambda)$ of weight $w(\lambda+\rho)-\rho$ (here $\ell(w)=p$). Begin by fixing a nonzero root vector e_α for each $\alpha \in \Phi^+$, and let $v_{w\lambda}$ be a nonzero weight vector of weight $w\lambda$ (such a vector is unique up to a scalar). Define $\omega \in C^p(\mathfrak{n}, V^\lambda) = \mathrm{Hom}_{\mathbb{C}}(\bigwedge^p \mathfrak{n}, V^\lambda)$ by

$$\omega(e_{\alpha_1} \wedge \ldots \wedge e_{\alpha_p}) = \begin{cases} v_{w\lambda} & \text{if } \alpha_i = \beta_i \ (1 \le i \le p) \\ 0 & \text{if } \{\alpha_1, \ldots, \alpha_p\} \neq \{\beta_1, \ldots, \beta_p\}, \end{cases}$$

where $\beta_1, \ldots, \beta_p$ is an enumeration of $\Phi^+(w)$. Then $\omega \in B^p(\mathfrak{n}, V^\lambda)$ and $\mathbb{C}\omega$ is an $\mathfrak{h}$-invariant subspace of $H^p(\mathfrak{n}, V^\lambda)$ of weight

$$w\lambda - \sum_{\beta \in \Phi^+(w)} \beta = w\lambda - (\rho - w\rho) = w(\lambda+\rho) - \rho$$

by Lemma 3.4.3.

(iii) Kostant's proof of Theorem 3.4.1 relied on explicitly understanding the action of a special element $\Omega \in \mathcal{Z}(\mathfrak{g})$, namely the Casimir element, on $H^p(\mathfrak{n}, V^\lambda)$. In contrast, Casselman and Osborne [16] worked out the action of the whole of $\mathcal{Z}(\mathfrak{g})$ (Theorem 3.3.1), and this allowed them to give a more streamlined proof of Kostant's theorem.

(iv) Our proof of Kostant's theorem is from Knapp–Vogan [46, IV.9]; as was the case for the Casselman–Osborne theorem, one can find in [46] a more general statement whose proof differs only slightly from the proof we have given.

(v) For yet another proof of Kostant's theorem, see Aribaud [3]. ▲

3.5 Applications

We shall give two applications of Kostant's theorem (Theorem 3.4.1), both taken from Kostant's original paper [48]. We retain the notation of the previous section.

3.5.1 The Borel–Weil–Bott Theorem

In this section we will use Kostant's theorem to give another proof of the Borel–Weil–Bott theorem (Theorem 2.4.1).

Let G be a complex semisimple Lie group with Lie algebra $\mathfrak{g}$ and B the Borel subgroup corresponding to $\mathfrak{b}$. For convenience we also assume that G is simply connected. Let $X = G/B$ be the flag variety of G. We associate to each weight $\lambda \in \Lambda$, which may be viewed as a character of B, a G-equivariant line bundle $\mathcal{L}_\lambda = G \times_B \mathbb{C}_\lambda$ over G/B (see Section 2.2). The cohomology groups $H^p(X, \mathcal{L}_\lambda)$—which are naturally holomorphic representations of G—are then the subject of the Borel–Weil–Bott theorem.

Note that unlike in Chapter 2, our Borel subgroup B is now built from the positive root spaces, which means that the statement of the Borel–Weil–Bott theorem needs to be altered slightly. Before giving the statement, we recall a piece of notation: For $\lambda \in \Lambda$, let

$$q_\lambda = |\{\alpha \in \Phi^+ \colon (\lambda + \rho, \alpha) < 0\}|.$$

Theorem 3.5.1 (Borel–Weil–Bott)**.** *Assume the preceding notation and let $\lambda \in \Lambda$.*

(i) *If $\lambda + \rho$ is singular, then $H^q(X, \mathcal{L}_{-\lambda}) = 0$ for all $q \geq 0$.*

(ii) *If $\lambda + \rho$ is regular, then $H^q(X, \mathcal{L}_{-\lambda})$ is nonzero if and only if $q = q_\lambda$, in which case it is an irreducible representation of G of lowest weight $-(w(\lambda + \rho) - \rho)$, where w is the unique element of W that makes $w(\lambda + \rho)$ dominant. (In this case, $q_\lambda = \ell(w)$.)*

For the remainder of this section we let V^λ denote an irreducible representation of G of highest weight λ, and $\mathbb{C}_\lambda$ shall designate $\mathbb{C}$ with $\mathfrak{h}$ acting via λ.

In proving the Borel–Weil–Bott theorem, we shall proceed as we did in the first proof of the Borel–Weil theorem (see p.35). In particular, we shall prefer to work with the realization K/T of $X = G/B$, where K is a maximal compact subgroup of G and $T = K \cap B$ is a maximal torus in K. Our main tool is the next theorem, which may be thought of as a generalized Frobenius reciprocity.

Theorem 3.5.2 (Bott–Kostant Reciprocity)**.** *Let λ and μ be dominant weights. Then the multiplicity of $\mathbb{C}_{-\lambda}$ in $H^p(\mathfrak{n}, V^\mu)$ is equal to the multiplicity of $(V^\mu)^*$ in $H^p(X, \mathcal{L}_\lambda)$.*

The proof requires a little bit of differential geometry, which we now review.

Let $\Omega^p(X, \mathcal{L}_\lambda)$ denote the space of smooth, $\mathcal{L}_\lambda$-valued, type $(0,p)$ differential forms on $X = K/T$, and let $\overline{\partial}\colon \Omega^p(X, \mathcal{L}_\lambda) \to \Omega^{p+1}(X, \mathcal{L}_\lambda)$ be the associated Dolbeault operator. By definition, $\Omega^p(X, \mathcal{L}_\lambda)$ is the space of smooth sections of

$$\mathcal{L}_\lambda \otimes \textstyle\bigwedge^p (T^{(0,1)}X)^*,$$

where $T^{(0,1)}X$ is the antiholomorphic tangent bundle on X, and $\overline{\partial} = 1_{\mathcal{L}_\lambda} \otimes \overline{\partial}_X$, where $\overline{\partial}_X\colon \Omega^p(X) \to \Omega^{p+1}(X)$ is the usual Dolbeault operator, which is given by

$$\begin{aligned} \overline{\partial}_X \omega(X_1, \ldots, X_{p+1}) &= \sum_{j=1}^{p+1} (-1)^{j+1} X_j \omega(X_1, \ldots, \widehat{X_j}, \ldots, X_{p+1}) \\ &\quad + \sum_{k<l} (-1)^{k+l} \omega([X_k, X_l], X_1, \ldots, \widehat{X_l}, \ldots, \widehat{X_k}, \ldots, X_{p+1}). \end{aligned}$$

Then $\overline{\partial}^2 = 0$ and we have a complex $(\Omega^\bullet(X, \mathcal{L}_\lambda), \overline{\partial})$ which gives rise to the **Dolbeault cohomology spaces**

$$H^{0,p}(X, \mathcal{L}_\lambda) = \ker \overline{\partial} / \operatorname{im} \overline{\partial}.$$

The Dolbeault theorem then asserts that we have natural (hence equivariant, if we let K act in the obvious fashion on $\Omega^\bullet(X, \mathcal{L}_\lambda)$) isomorphisms

$$H^{0,p}(X, \mathcal{L}_\lambda) \cong H^p(X, \mathcal{L}_\lambda) \quad \text{for all } p \geq 0.$$

Now note that $(T^{(0,1)}X)^*$ is modeled on $\mathfrak{n}^-$. Thus we obtain, as usual, an identification

$$\Omega^p(X, \mathcal{L}_\lambda) \cong \{\omega\colon K \to {\textstyle\bigwedge}^p \mathfrak{n}^- \otimes \mathbb{C}_\lambda \colon \omega(gt) = \lambda(t)^{-1}\omega(g) \text{ for all } t \in T\}.$$

We can further identify the right-hand side (equivariantly) with

$$(C^\infty(K) \otimes \mathbb{C}_\lambda \otimes {\textstyle\bigwedge}^p \mathfrak{n}^-)^T \cong \mathrm{Hom}_T({\textstyle\bigwedge}^p \mathfrak{n}, C^\infty(K) \otimes \mathbb{C}_\lambda).$$

Here T acts on $C^\infty(K)$ via the *right* regular representation, on $\mathbb{C}_\lambda$ via λ and on $\bigwedge^p \mathfrak{n}^-$ via $\bigwedge^p \mathrm{Ad}^*$. An explicit isomorphism $\mathrm{Hom}_T(\bigwedge^p \mathfrak{n}, C^\infty(K) \otimes \mathbb{C}_\lambda) \to \Omega^p(X, \mathcal{L}_\lambda)$ sends T to the form ω_T given by

$$\omega_T(k)(X_1, \ldots, X_p) = T(X_1, \ldots, X_p)(k).$$

The details, which are straightforward, are omitted.

Finally, notice that the subspaces $\mathrm{Hom}_T(\bigwedge^\bullet \mathfrak{n}, C^\infty(K) \otimes \mathbb{C}_\lambda)$ sit inside the corresponding cocycle spaces $C^p(\bigwedge^\bullet \mathfrak{n}, C^\infty(K) \otimes \mathbb{C}_\lambda) = \mathrm{Hom}_{\mathbb{C}}(\bigwedge^\bullet \mathfrak{n}, C^\infty(K) \otimes \mathbb{C}_\lambda)$ of Lie algebra cohomology. In fact, the former form a subcomplex of the latter, and on this subcomplex the Lie algebra coboundary operator d coincides with the Dolbeault operator just by definition.

We are now ready to give the

Proof of Theorem 3.5.2. From the preceding discussion we immediately get $\mathfrak{h}$-module isomorphisms

$$(H^p(\mathfrak{n}, C^\infty(K) \otimes \mathbb{C}_{-\lambda}))^{\mathfrak{h}} \cong H^{0,p}(X, \mathcal{L}_{-\lambda}) \cong H^p(X, \mathcal{L}_{-\lambda}) \quad \text{for all } p \geq 0, \tag{3.5.1}$$

where the superscript $\mathfrak{h}$ denotes the subspace of $\mathfrak{h}$-invariants, as usual.

On the other hand, arguing as in the proof of the Borel–Weil theorem (p.35), we see that the inclusion

$$C^\infty(K) \subset L^2(K) = \bigoplus_\mu V^\mu \otimes (V^\mu)^*$$

yields isomorphisms

$$(H^p(\mathfrak{n}, C^\infty(K) \otimes \mathbb{C}_\lambda))^{\mathfrak{h}} \cong \bigoplus_\mu (V^\mu)^* \otimes (H^p(\mathfrak{n}, V^\mu) \otimes \mathbb{C}_\lambda)^{\mathfrak{h}} \quad \text{for all } p \geq 0. \tag{3.5.2}$$

So by combining (3.5.1) and (3.5.2), we obtain

$$H^p(X,\mathcal{L}_\lambda) \cong \bigoplus_\mu (V^\mu)^* \otimes (H^p(\mathfrak{n},V^\mu)\otimes \mathbb{C}_\lambda)^{\mathfrak{h}} \quad \text{for all } p \geq 0. \tag{3.5.3}$$

But now observe that

$$\dim(H^p(\mathfrak{n},V^\mu)\otimes \mathbb{C}_\lambda)^{\mathfrak{h}} = \dim \operatorname{Hom}_{\mathfrak{h}}(\mathbb{C}_{-\lambda}, H^p(\mathfrak{n},V^\mu)) = \text{multiplicity of } \mathbb{C}_{-\lambda} \text{ in } H^p(\mathfrak{n},V^\mu),$$

which, according to Kostant's theorem (Theorem 3.4.1), is either 0 or 1. It follows from (3.5.3) that the multiplicity of $\mathbb{C}_{-\lambda}$ in $H^p(\mathfrak{n},V^\mu)$ is equal to the number of times $(V^\mu)^*$ appears in $H^p(X,\mathcal{L}_\lambda)$, as desired. ■

Remarks 3.5.3.

(i) With a little more effort it can be shown that

$$\operatorname{Hom}_K(V^\mu, H^p(X,\mathcal{L}_\lambda)) \cong \operatorname{Hom}_T(\mathbb{C}_{-\lambda}, H^p(\mathfrak{n},(V^\mu)^*)),$$

which is perhaps more reminiscent of Frobenius reciprocity.

(ii) Bott's original formulation [13, Theorem I] of Theorem 3.5.2 involved relative Lie algebra cohomology (see [13] for the definitions). He found that

$$\dim H^p(\mathfrak{b},\mathfrak{h},\operatorname{Hom}_{\mathbb{C}}((V^\mu)^*,\mathbb{C}_\lambda)) = \text{multiplicity of } (V^\mu)^* \text{ in } H^p(X,\mathcal{L}_\lambda).$$

Kostant observed that, in general,

$$\dim H^p(\mathfrak{n},V)^{\mathfrak{h}} = \dim H^p(\mathfrak{n},\mathfrak{h},V).$$

Taking $V = \operatorname{Hom}((V^\mu)^*,\mathbb{C}_\lambda)$ and noting that $H^p(\mathfrak{n},V) = H^p(\mathfrak{n},V^\mu)\otimes\mathbb{C}_\lambda$ (because $\mathfrak{n}$ acts trivially on $\mathbb{C}_\lambda$), one then finds

$$\dim H^p(\mathfrak{n},\mathfrak{h},\operatorname{Hom}((V^\mu)^*,\mathbb{C}_\lambda)) = \text{multiplicity of } \mathbb{C}_{-\lambda} \text{ in } H^p(\mathfrak{n},V^\mu).$$

This links Kostant's version [48, Proposition 6.3] of Theorem 3.5.2—the version we have quoted—to Bott's. ▲

At last, we come to the

Proof of the Borel–Weil–Bott theorem. Suppose that $\lambda+\rho$ is singular. Then the multiplicity of $\mathbb{C}_\lambda$ in $H^p(\mathfrak{n}, V^\mu)$ is zero for all $p \geq 0$ and all dominant μ, for otherwise Kostant's theorem would yield a $w \in W$ such that $\lambda = w(\mu+\rho)-\rho$, implying that $\lambda+\rho = w(\mu+\rho)$ is regular—a contradiction. We thus conclude that the multiplicity of $(V^\mu)^*$ is zero in $H^p(X, \mathcal{L}_{-\lambda})$ for all $p \geq 0$ and all dominant μ, by the Bott–Kostant reciprocity theorem. Hence $H^p(X, \mathcal{L}_{-\lambda}) = 0$ for all $p \geq 0$.

On the other hand, suppose that $\lambda+\rho$ is regular. Kostant's theorem asserts that the multiplicity of $\mathbb{C}_\lambda$ in $H^p(\mathfrak{n}, V^\mu)$ is zero for all $p \geq 0$ and all dominant μ unless $\lambda = w^{-1}(\mu+\rho)-\rho$ for some $w \in W_p$, in which case the multiplicity is 1. The equality $\lambda = w^{-1}(\mu+\rho)-\rho$ implies in particular that $w(\lambda+\rho) = \mu+\rho$ is dominant. There is a unique Weyl group element w for which the preceding statement is true, and for this w, $\ell(w) = q_\lambda$. Thus the multiplicity of $\mathbb{C}_\lambda$ in $H^p(\mathfrak{n}, V^\mu)$ is nonzero if and only if $\mu = w(\lambda+\rho)-\rho$ and $p = q_\lambda$. By appealing to the Bott–Kostant reciprocity theorem, we draw two conclusions. First, if $p \neq q_\lambda$, then the multiplicity of $(V^\mu)^*$ in $H^p(X, \mathcal{L}_{-\lambda})$ is zero for all dominant μ. Hence $H^p(X, \mathcal{L}_{-\lambda}) = 0$ in this case. The second conclusion is that $H^{q_\lambda}(X, \mathcal{L}_{-\lambda})$ is equivalent to $(V^{w(\lambda+\rho)-\rho})^*$, hence has lowest weight $-(w(\lambda+\rho)-\rho)$. ■

Remark 3.5.4. All the reasoning above can be reversed and one therefore obtains Kostant's theorem as a consequence of the Borel–Weil–Bott theorem. ▲

3.5.2 Weyl's Character Formula

Let $\{e^\mu : \mu \in \mathfrak{h}^*\}$ denote a $\mathbb{Z}$-module basis for the group ring $\mathbb{Z}[\mathfrak{h}^*]$, labeled so that $e^\mu e^\nu = e^{\mu+\nu}$. If V is a (finite-dimensional) $\mathfrak{b}$-module (so in particular a $\mathfrak{g}$-module), with weight space decomposition $V = \bigoplus_{\mu\in\mathfrak{h}^*} V_\mu$ (we adopt the convention that $V_\mu = 0$ if μ is not a weight of V), then the **formal character** of V is, by definition,

$$\operatorname{ch} V = \sum_{\mu\in\mathfrak{h}^*} (\dim V_\mu) e^\mu \in \mathbb{Z}[\mathfrak{h}^*]. \tag{3.5.4}$$

We also define the **$\mathfrak{n}$-Euler characteristic** of V by

$$\chi(\mathfrak{n}, V) = \sum_{p \geq 0} (-1)^p \operatorname{ch} H^p(\mathfrak{n}, V) = \sum_{p=0}^{\dim \mathfrak{n}} (-1)^p \operatorname{ch} H^p(\mathfrak{n}, V). \tag{3.5.5}$$

Lemma 3.5.5 (Euler–Poincaré Principle)**.** *Let V be a $\mathfrak{b}$-module. Then*

$$\chi(\mathfrak{n}, V) = \sum_{p=0}^{\dim \mathfrak{n}} (-1)^p \operatorname{ch} C^p(\mathfrak{n}, V).$$

Proof. The usual proof works here. Fix $\mu \in \mathfrak{h}^*$. From the fact that the coboundary operator d commutes with the action of $\mathfrak{h}$ on $C^p(\mathfrak{n}, V)$, we see that

$$H^p(\mathfrak{n}, V)_\mu = Z^p(\mathfrak{n}, V)_\mu / B^p(\mathfrak{n}, V)_\mu$$

whence

$$\dim H^p(\mathfrak{n}, V)_\mu = \dim Z^p(\mathfrak{n}, V)_\mu - \dim B^p(\mathfrak{n}, V)_\mu. \tag{3.5.6}$$

On the other hand, the short exact sequence

$$0 \longrightarrow Z^p(\mathfrak{n}, V)_\mu \longrightarrow C^p(\mathfrak{n}, V)_\mu \xrightarrow{d} B^{p+1}(\mathfrak{n}, V)_\mu \longrightarrow 0$$

(see the proof of Proposition 3.2.9) yields

$$\dim C^p(\mathfrak{n}, V)_\mu = \dim Z^p(\mathfrak{n}, V)_\mu + \dim B^{p+1}(\mathfrak{n}, V)_\mu. \tag{3.5.7}$$

By combining (3.5.6) and (3.5.7), we obtain

$$\sum_p (-1)^p \dim C^p(\mathfrak{n}, V)_\mu = \sum_p (-1)^p \dim H^p(\mathfrak{n}, V)_\mu,$$

and the lemma follows. ■

Now let V^λ be an irreducible $\mathfrak{g}$-module of highest weight λ. We wish to compute $\operatorname{ch} V^\lambda$.

Lemma 3.5.6. *We have*

$$\chi(\mathfrak{n}, V^\lambda) = \sum_{w \in W} \operatorname{sgn}(w) e^{w(\lambda+\rho)-\rho} = e^{-\rho} \sum_{w \in W} \operatorname{sgn}(w) e^{w(\lambda+\rho)}.$$

Proof. From (3.5.4) and Kostant's theorem (Theorem 3.4.1) we see that

$$\operatorname{ch} H^p(\mathfrak{n}, V^\lambda) = \sum_{w \in W^p} e^{w(\lambda+\rho)-\rho}.$$

The lemma now follows from (3.5.5) and the fact that $\operatorname{sgn}(w) = (-1)^p$ for $w \in W_p$. ■

Theorem 3.5.7 (Kostant)**.** *Let $\mathbb{C}$ denote the trivial $\mathfrak{g}$-module. Then*

$$\operatorname{ch} V^\lambda = \frac{\chi(\mathfrak{n}, V^\lambda)}{\chi(\mathfrak{n}, \mathbb{C})}$$

in $\mathbb{Z}[\mathfrak{h}^*]$.

Proof. In the proof of Lemma 3.4.5 we learned that the weights occuring in $C^p(\mathfrak{n}, V^\lambda)$ are of the form

$$\eta - \sum_{\alpha \in S} \alpha,$$

where η is a weight of V^λ and S is a subset of Φ^+ of size p. Thus, by summing over all such η and S, we find that

$$\operatorname{ch} C^p(\mathfrak{n}, V^\lambda) = \operatorname{ch} V^\lambda \sum_S e^{-\sum_{\alpha \in S} \alpha}$$

and then

$$\sum_p (-1)^p \operatorname{ch} C^p(\mathfrak{n}, V^\lambda) = \operatorname{ch} V^\lambda \sum_{S \subseteq \Phi^+} (-1)^{|S|} e^{-\sum_{\alpha \in S} \alpha}. \tag{3.5.8}$$

The left-hand side is $\chi(\mathfrak{n}, V^\lambda)$ by the Euler–Poincaré principle (Lemma 3.5.5). So it remains to deal with the sum

$$\sum_{S \subseteq \Phi^+} (-1)^{|S|} e^{-\sum_{\alpha \in S} \alpha} = \sum_p (-1)^p \sum_{|S|=p} e^{-\sum_{\alpha \in S} \alpha}.$$

The weight $-\sum_{\alpha \in S} \alpha$ occurs in $C^p(\mathfrak{n}, \mathbb{C})$ with multiplicity equal to the number of such S's (of size p) yielding the same sum. It follows that

$$\sum_{|S|=p} e^{-\sum_{\alpha \in S} \alpha} = \operatorname{ch} C^p(\mathfrak{n}, \mathbb{C}).$$

So, after another application of the Euler–Poincaré principle, (3.5.8) becomes

$$\chi(\mathfrak{n}, V^\lambda) = \chi(\mathfrak{n}, \mathbb{C}) \operatorname{ch} V^\lambda.$$

Hence $\chi(\mathfrak{n}, \mathbb{C})$ divides $\chi(\mathfrak{n}, V^\lambda)$ in $\mathbb{Z}[\mathfrak{h}^*]$, with $\operatorname{ch} V^\lambda$ as quotient. ■

Corollary 3.5.8 (Weyl's Character Formula)**.**

$$\operatorname{ch} V^\lambda = \frac{\sum_{w \in W} \operatorname{sgn}(w) e^{w(\lambda+\rho)}}{\sum_{w \in W} \operatorname{sgn}(w) e^{w\rho}}.$$

Proof. This follows immediately from the preceding theorem and Lemma 3.5.6. ■

Corollary 3.5.9 (Weyl's Dimension Formula)**.**

$$\dim V^\lambda = \prod_{\alpha \in \Phi^+} \frac{(\lambda + \rho, \alpha)}{(\rho, \alpha)}.$$

Proof. This is a well-known consequence of Weyl's character formula: see, e.g., [45, Theorem 5.84]. ■

Remark 3.5.10. Although Kostant was the first to explicitly show how Weyl's character formula could be distilled from an examination of $H^*(\mathfrak{n}, V^\lambda)$, the possibility of this had already been noted by Bott [13, p.247]. ▲

Chapter 4

Discrete Series Representations and Schmid's Theorem

4.1 Introduction

This chapter is a departure from the previous two in at least two significant ways. First, the groups under primary consideration here are the *real* semisimple groups—not the complex ones. Second, we shall study the unitary representation theory of these groups on (possibly infinite-dimensional) Hilbert spaces. In fact, the chief goal of this chapter is the construction of certain nontrivial irreducible unitary representations—the so-called *discrete series* representations—which are typically infinite-dimensional when the semisimple Lie group under consideration is noncompact.

It should be mentioned that, partly due to length constraints, this chapter does not contain much in the way of proofs. I have, however, attempted to provide references for all the main results. The reader will find that most of these references point to papers of Harish-Chandra, where the corresponding results were first proved. As Harish-Chandra's papers are rather demanding, it seems prudent to point out that a more streamlined account of several of these results can be found in the paper of Atiyah and Schmid [4].

We now turn to a brief outline of the chapter. The first two sections contain a summary of some of the main ideas of the theory of unitary representations of semisimple Lie groups. Most importantly, we define what it means for an irreducible unitary representation to belong to the discrete series (Definition 4.2.4), and we give Harish-Chandra's parametrization (Theorem 4.3.7) of these representations.

In Section 4.4 we present—without proof—Schmid's generalization of the Borel–Weil–Bott theorem (Theorem 4.4.2). This theorem realizes all the discrete series representations of a semisimple Lie group G in L^2 cohomology spaces of equivariant line bundles over the flag domain (the noncompact analogue of the flag variety) of G. A concise outline of its proof is given in Section 4.5.

In writing this chapter I have been greatly influenced by, and have closely followed, the papers of Schmid [54], [55], [56], [57], [59].

4.2 Preliminaries

Unless stated otherwise, G will always denote a real semisimple Lie group.[1] We are interested in studying the unitary representation theory of G. Thus an important role will be played by its unitary dual $\widehat{G}$. Recall that $\widehat{G}$ is the set of irreducible unitary representations of G modulo unitary equivalence.[2] We shall be nonchalant about identifying an irreducible unitary representation of G with its equivalence class in $\widehat{G}$. The Hilbert space on which a unitary representation π acts will be denoted by H_π.

Forgetting about semisimplicity for the moment, let us suppose that G is a compact group. Then a good deal of information is known about $\widehat{G}$. In particular, one knows that

[1] A **real semisimple Lie group** is a connected Lie group whose Lie algebra is semisimple, i.e. has no solvable ideals. For technical convenience, we shall always assume that our semisimple Lie groups have finite centre.

[2] A **unitary representation** of a Lie group G on a nonzero Hilbert space H is a homomorphism $\pi\colon G \to \mathcal{U}(H)$, where $\mathcal{U}(H)$ is the group of unitary operators on H, such that the function $G \to \mathbb{C}$ defined by $g \mapsto \langle \pi(g)u, v\rangle$ is continuous for all $u, v \in H$. Two unitary representations $\pi\colon G \to \mathcal{U}(H)$ and $\pi'\colon G \to \mathcal{U}(H')$ of G are said to be **unitarily equivalent** if there is a unitary operator $U\colon H \to H'$ such that $U\pi(g) = \pi(g)U$ for all $g \in G$. A unitary representation $\pi\colon G \to \mathcal{U}(H)$ of G is said to be **irreducible** if it has no *closed* invariant subspaces other than 0 and H.

each $\pi \in \widehat{G}$ is finite-dimensional and occurs as a direct summand (with multiplicity equal to its dimension) in the left regular representation of G on $L^2(G)$. More specifically, if we let M_π denote the subspace of $L^2(G)$ spanned by the matrix coefficients of π, then the Peter–Weyl theorem states that

$$L^2(G) = \bigoplus_{\pi \in \widehat{G}} M_\pi \cong \bigoplus_{\pi \in \widehat{G}} H_\pi^* \otimes H_\pi \cong \bigoplus_{\pi \in \widehat{G}} d_\pi H_\pi, \tag{4.2.1}$$

where $\bigoplus$ denotes the direct sum of Hilbert spaces and $d_\pi = \dim H_\pi$.[3] This is a decomposition of $L^2(G)$ into irreducible subrepresentations. The projection map

$$L^2(G) \to M_\pi$$

is given by

$$f \mapsto d_\pi(\chi_\pi * f)$$

where χ_π is the character of π and $*$ denotes the convolution[4] of functions. Thus each $f \in L^2(G)$ can be expressed as an L^2 sum

$$f = \sum_{\pi \in \widehat{G}} d_\pi(\chi_\pi * f). \tag{4.2.2}$$

If we let μ denote the measure on $\widehat{G}$ that assigns a mass of d_π to $\{\pi\} \subset \widehat{G}$, then we can rewrite the above expression as

$$f = \int_{\widehat{G}} \chi_\pi * f \, d\mu(\pi).$$

The measure μ is called the **Plancherel measure** of G. This form of (4.2.2) turns out to be more amenable to generalization. Before proceeding further, let us give a concrete example that will help illuminate the previous remarks.

Example 4.2.1. Let $G = \mathbb{T}$, the unit circle in $\mathbb{C}$. By Schur's Lemma, each π in $\widehat{\mathbb{T}}$ is one-dimensional and hence is (equivalent to) a continuous homomorphism $\mathbb{T} \to \mathcal{U}(\mathbb{C}) = \mathbb{T}$.

[3] The Peter–Weyl theorem gives us the first equality, which is comparatively the only difficult one to prove.

[4] By definition, $(\chi_\pi * f)(x) = \int_G \chi_\pi(xg^{-1}) f(g) dg$, $x \in G$, where dg is the (normalized) Haar measure of G.

These are all of the form $\chi_n\colon z \mapsto z^n$ for some $n \in \mathbb{Z}$. It follows that $n \leftrightarrow \chi_n$ gives us a one-to-one correspondence between $\widehat{\mathbb{T}}$ and $\mathbb{Z}$. Now, being one-dimensional, each χ_n is equal to its own character and therefore $M_\pi = \mathbb{C}\chi_n$. Thus the Peter–Weyl theorem (4.2.1) yields the decomposition

$$L^2(\mathbb{T}) = \bigoplus_{n\in\mathbb{Z}} \mathbb{C}\chi_n.$$

The projection of $L^2(\mathbb{T})$ onto $\mathbb{C}\chi_n$ is given by $f \mapsto \chi_n * f$. An easy computation shows us that $\chi_n * f = \hat{f}(n)\chi_n$, where $\hat{f}(n)$ is the nth Fourier coefficient of f. This means that the series expansion we obtained in (4.2.2) is simply

$$f(z) = \sum_{n\in\mathbb{Z}} \hat{f}(n) z^n, \quad z \in \mathbb{T},$$

which is none other than the Fourier expansion of an L^2 function on $\mathbb{T}$. Thus the remarks preceding this Example can be thought of as "Fourier analysis on compact groups." ▲

Let us now return to the setting of a semisimple Lie group G. In view of the above remarks, we are really only interested in the case of noncompact G. In this context, a typical irreducible unitary representation need not be finite-dimensional. In fact, if G is noncompact and simple, then G has no nontrivial finite-dimensional unitary representations—let alone irreducible ones. This is a significant complication. For example, it is not clear how one should define the character—a notion which is very useful in the compact setting—of an infinite-dimensional representation π. The naive definition $g \mapsto \operatorname{trace} \pi(g)$, $g \in G$, is of course inadequate: when π is infinite-dimensional, the unitary operator $\pi(g)$ has infinitely many eigenvalues on the unit circle and hence does not possess a well-defined trace. Another complication is that the left regular representation $L^2(G)$ generally does not contain all the irreducible unitary representations of G as subrepresentations; in fact, it need not contain any! This is the case, for example, if G is complex semisimple.

Fortunately all these obstacles can be surmounted. Instead of decomposing $L^2(G)$ as a direct sum of representations, one has a direct integral decomposition—the so-called **Plancherel decomposition** of G—

$$L^2(G) = \int_{\widehat{G}}^{\oplus} H_\pi^* \,\widehat{\otimes}\, H_\pi \, d\mu(\pi), \tag{4.2.3}$$

where $\widehat{\otimes}$ denotes the Hilbert space tensor product and μ is the **Plancherel measure** on $\widehat{G}$. Of course, the problem is to prove that such a measure μ actually exists. Moreover, having proved existence, one would like to be able to understand μ in some explicit fashion. We shall not concern ourselves with these problems too greatly; it suffices to say that they have satisfactory solutions, thanks to the work of Harish-Chandra. What is relevant to us now is the following. By isolating the subset of $\widehat{G}$ whose elements are assigned a finite positive mass by μ, we can effectively split the direct integral into a "discrete" part and a "continuous" part. The discrete part is in fact a direct sum. If G is compact, then the discrete part is all there is, for in this case the Plancherel measure assigns a positive mass of $d_\pi = \dim H_\pi$ to each singleton $\{\pi\}$ in $\widehat{G}$.

Proposition 4.2.2 (Godement–Harish-Chandra)**.** *Let G be a semisimple Lie group. For $\pi \in \widehat{G}$, the following conditions are equivalent.*

(a) *The Plancherel measure μ of $\widehat{G}$ assigns a positive mass to the singleton $\{\pi\}$.*

(b) *Some nonzero matrix coefficient of π is in $L^2(G)$.*

(c) *All the matrix coefficients of π are in $L^2(G)$.*

(d) *π is unitarily equivalent to a subrepresentation of $L^2(G)$.* ■

Remark 4.2.3. The equivalence of properties (a) through (c) is due to Godement [24], [25], who also showed that $\pi \in \widehat{G}_d$ satisfies an analogue of the Schur orthogonality relations. In particular,

$$\int_G |\langle \pi(g)u, v\rangle|^2 dg = d_\pi^{-1}\|u\|\|v\| \qquad \text{for all } u, v \in H_\pi,$$

where d_π is some positive constant (called the **formal degree** of π) that depends only on π. Harish-Chandra [33] established the equivalence of property (d) with the others, and in doing so showed that d_π is equal to the mass assigned to $\{\pi\}$ by the Plancherel measure, just as in the compact case. ▲

Definition 4.2.4. An irreducible unitary representation of G is said to belong to the **discrete series** of G if it satisfies the equivalent conditions in the previous Proposition. We denote by $\widehat{G}_d$ the subset of $\widehat{G}$ consisting of (the equivalence classes of) the discrete series representations.

Our goal in this chapter is to give a geometric construction of the discrete series representations of G. The construction we have in mind is much in the style of the Borel–Weil–Bott theorem—in fact, it will be seen to reduce exactly to the Borel–Weil–Bott theorem in case G is compact. Thus two natural questions arise:

- How is $\widehat{G}_d$ parameterized?
- What is the "noncompact" analogue of the flag variety?

4.3 Harish-Chandra's Parametrization of the Discrete Series

An important thing to note is that $\widehat{G}_d$ may very well be empty. The key result concerning the existence of discrete series representations is the following.

Theorem 4.3.1 (Harish-Chandra [37])**.** *For a semisimple Lie group G, the following conditions are equivalent.*

(a) $\widehat{G}_d \neq \emptyset$.

(b) *G has a compact Cartan subgroup.*[5]

(c) *If K is a maximal compact subgroup of G, then* $\operatorname{rank} G = \operatorname{rank} K$.[6] ∎

Examples 4.3.2.

(a) We know that $\widehat{G}_d = \widehat{G} \neq \emptyset$ if G is compact. This is consistent with the above theorem, as any maximal torus in a compact semisimple G is a compact Cartan subgroup.

[5] A subalgebra $\mathfrak{h}_0$ of a (real) Lie algebra $\mathfrak{g}_0$ is called a **Cartan subalgebra** if it is nilpotent and if $N_{\mathfrak{g}}(\mathfrak{h}) = \mathfrak{h}$; if $\mathfrak{g}_0$ is semisimple, then $\mathfrak{h}_0$ is a Cartan subalgebra if and only if its complexification $\mathbb{C} \otimes \mathfrak{h}_0$ is a Cartan subalgebra of the complex semisimple Lie algebra $\mathbb{C} \otimes \mathfrak{g}_0$ (see Section 1.3). A **Cartan subgroup** of a Lie group G is the centralizer $C_G(\mathfrak{h}_0)$ in G of a Cartan subalgebra $\mathfrak{h}_0$ in $\operatorname{Lie} G$.

[6] It can be shown that the Cartan subgroups of a connected Lie group G all have the same dimension. This common (real) dimension is called the **rank** of G.

(b) If G is a complex semisimple Lie group, then $\widehat{G}_d = \emptyset$. Indeed, if K is a compact real form (which in particular is a maximal compact subgroup) of G, then $\operatorname{rank} G = 2 \operatorname{rank} K$.

(c) If $G = \mathrm{SL}(n, \mathbb{R})$ (and $n > 1$ — for otherwise G is not semisimple), then $\widehat{G}_d \neq \emptyset$ if and only if $n = 2$. To prove this, note that $\mathrm{SO}(n)$ is a maximal compact subgroup of $\mathrm{SL}(n, \mathbb{R})$ of rank $\lfloor \frac{n}{2} \rfloor$ while the rank of $\mathrm{SL}(n, \mathbb{R})$ is $n - 1$. ▲

Let us elaborate on the relationship between compact Cartan subgroups and discrete series representations. A good place to start is with G compact. In this setting a Cartan subgroup of G is a maximal torus T. We also know that $\widehat{G}_d = \widehat{G}$ is parameterized by the dominant weights of G; what follows is a way to rephrase this. Let $W = N_G(T)/T$ denote the Weyl group of G, which we know to be independent of the choice of T. The action of W on T yields an action on the character group $\widehat{T} = \{\text{continuous homomorphisms } \chi \colon T \to \mathbb{T}\}$. Recalling that there is a unique dominant weight in the W-orbit of each $\chi \in \widehat{T}$, we see that the free orbits parameterize $\widehat{G}$. That is, to each $\chi \in \widehat{T}$ which is not fixed by any nontrivial element of W there corresponds a $\pi_\chi \in \widehat{G}$, and every element of $\widehat{G}$ is obtained this way. Moreover, two irreducible unitary representations π_χ and π'_χ of G are equivalent if and only if $\chi = w\chi'$ for some $w \in W$.

Now suppose that G is semisimple and let H be a compact Cartan subgroup of G. Although there may be multiple conjugacy classes of Cartan subgroups in G, there is only one conjugacy class of compact Cartan subgroups. Thus H is unique up to conjugacy. Let K be the unique maximal compact subgroup of G containing H. Let $\mathfrak{h}, \mathfrak{k}$ and $\mathfrak{g}$ denote the complexified Lie algebras of H, K and G, respectively. The Weyl groups of the root systems $(\mathfrak{g}, \mathfrak{h})$ and $(\mathfrak{k}, \mathfrak{h})$ are referred to as the *complex* Weyl group $W_{\mathbb{C}}$ and the *real* Weyl group $W_{\mathbb{R}}$, respectively. Of course $W_{\mathbb{R}} = W_{\mathbb{C}}$ if G is compact, but in general we can only say that $W_{\mathbb{R}} \subset W_{\mathbb{C}}$. In analogy with the compact case, both Weyl groups act on the character group $\widehat{H}$, and $\widehat{G}_d$ is parameterized by the orbits of these actions. The precise facts, due to Harish-Chandra, are as follows. If $\chi \in \widehat{H}$ is not fixed by any nontrivial element of $W_{\mathbb{C}}$, then one can associate to it an element $\pi_\chi \in \widehat{G}_d$. The real Weyl group acts on the subset of $\widehat{H}$ consisting of such χ, and two discrete series representations π_χ and π'_χ are equivalent if and only if they lie in the same $W_{\mathbb{R}}$-orbit.

Example 4.3.3. Let $G = \mathrm{SL}(2,\mathbb{R})$. Then $H = \mathrm{SO}(2) \cong \mathbb{T}$ is a Cartan subgroup of G and we know that $\widehat{H} = \{\chi_n \colon \chi_n(z) = z^n\} \cong \mathbb{Z}$. In the above notation, $H = K$—that is, H is maximal compact in G. Thus $W_{\mathbb{R}}$ is trivial. On the other hand, $W_{\mathbb{C}} = \langle\sigma\rangle \cong \mathbb{Z}/2\mathbb{Z}$, with $\sigma\chi_n = \chi_{-n}$ for all n. It follows that the only character of T which is fixed by a nontrivial element of $W_{\mathbb{C}}$ is χ_0. Consequently $\widehat{G}_d$ is parameterized by the nontrivial characters $\{\chi_n \colon n \neq 0\}$. ▲

It should be remarked that Harish-Chandra obtained this parametrization of $\widehat{G}_d$ without explicitly constructing the representation π_χ. He also obtained a result concerning the *global character* Θ_{π_χ} of π_χ. The global character is a useful generalization of the usual character to arbitrary (and possibly infinite-dimensional) unitary representations. To motivate its definition, let us suppose for the moment that G is a finite group. In this case we can identify $L^2(G)$ with the group algebra $\mathbb{C}G$ via the mapping $f \mapsto \sum_{g\in G} f(g)g$, $f \in L^2(G)$, and group representations of G can be extended to algebra representations of $L^2(G)$: if π is a representation of G, one defines

$$\pi(f) = \sum_{g\in G} f(g)\pi(g) \qquad (f \in L^2(G)).$$

Then the character Θ_π of π on $\mathbb{C}G$ is defined by $\Theta_\pi(f) = \operatorname{trace}\pi(f)$, $f \in L^2(G)$. The usual character χ_π of π can be extracted from Θ_π easily: we have $\chi_\pi(g) = \Theta_\pi(g)$, where on the right-hand side we are using g to denote the function on G which assumes the value 1 at g and 0 elsewhere.

The generalization to semisimple Lie groups is as follows. If π is an irreducible unitary representation of G, one defines

$$\pi(f) = \int_G f(g)\pi(g)dg$$

whenever f is a sufficiently well-behaved function on G so that the (operator-valued) integral makes sense. If f belongs to $C_c^\infty(G)$, then not only is the integral well-defined, but in fact the operator $\pi(f)$ is of trace class[7] on H_π. To see why this is so, one has to

[7] An operator T on a Hilbert space H is said to be of **trace class** if $\sum_i |\langle Te_i, e_i\rangle| < \infty$ for any (and hence all) orthonormal bases $\{e_i\}$ of H. In this case the number $\sum_i \langle Te_i, e_i\rangle$ is independent of the choice of basis and is called the **trace** of T.

consider what happens to π once it is restricted to the maximal compact subgroup K of G. One finds that, as a representation of K, H_π decomposes as

$$H_\pi = \bigoplus_{\tau \in \widehat{K}} m_\tau H_\tau,$$

with each irreducible unitary representation τ of K occurring with some multiplicity m_τ. The crucial fact is that, because π is irreducible, each m_τ is finite. This allows us, in some sense, to bring in K to help keep $\pi(f)$ well-behaved.

Theorem 4.3.4 (Harish-Chandra [31], [32])**.** *Let G be a semisimple Lie group and let $\pi \in \widehat{G}$. Then, for all $f \in C_c^\infty(G)$, the operator $\pi(f)$ is of trace class on H_π. The map $\Theta_\pi \colon C_c^\infty(G) \to \mathbb{C}$ sending f to $\operatorname{trace} \pi(f)$ is a conjugation-invariant distribution (in the sense of Schwarz) on G. Moreover, if π_1 and π_2 are unitary representations of G, then we have $\Theta_{\pi_1} = \Theta_{\pi_2}$ if and only if π_1 and π_2 are unitarily equivalent.* ■

Definition 4.3.5. If $\pi \in \widehat{G}$, then the distribution Θ_π is called the **global character** of π.

The following profoundly deep theorem of Harish-Chandra shows that Θ_π is a distribution of the best kind.

Theorem 4.3.6 (Harish-Chandra's Regularity Theorem [36])**.** *Let G be a semisimple Lie group. The global character Θ_π of $\pi \in \widehat{G}$ is locally an L^1 function on G, that is, there is a $\Theta_\pi \in L^1(G)$ such that $\Theta_\pi(f) = \int_G f(g)\Theta_\pi(g)dg$ for all $f \in C_c^\infty(G)$. This function Θ_π is conjugation-invariant. Moreover, there is a dense open subset of G (the set of **regular** elements), whose complement in G has measure zero, and on which all these functions Θ_π ($\pi \in \widehat{G}$) are real analytic.* ■

For example, if π is finite-dimensional, then

$$\Theta_\pi(f) = \int_G f(g)\chi_\pi(g)dg \qquad (f \in C_c^\infty(G)),$$

that is, the function Θ_π is the usual character χ_π of π.

We now turn our attention to the global characters of discrete series representations. If $\pi_\chi \in \widehat{G}_d$, let us write Θ_χ for Θ_{π_χ}. Recall that we are fixing a compact Cartan subgroup

$H \subset G$ and a maximal compact subgroup $K \supset H$. Recall also that $\mathfrak{h}, \mathfrak{k}$ and $\mathfrak{g}$ are the complexified Lie algebras of H, K and G, respectively. Let Φ denote the root system of $(\mathfrak{g}, \mathfrak{h})$ and fix a choice of positive roots $\Phi^+ \subset \Phi$. As usual we let Λ denote the weight lattice, which now lies in $i\mathfrak{h}_0^*$ where $\mathfrak{h}_0$ is the (real) Lie algebra of H, and we set $\rho = \frac{1}{2}\sum_{\alpha\in\Phi^+} \alpha \in \Lambda \otimes_{\mathbb{Z}} \mathbb{Q}$. We can (and do) identify Λ with the character group $\widehat{H}$ just as in the compact case.

Harish-Chandra's parametrization of $\widehat{G}_d$ is summarized as follows.

Theorem 4.3.7 (Harish-Chandra [37]). *Let $H \subset K \subset G$ be as above. If $\lambda \in \Lambda$ is such that $\lambda + \rho$ is regular, then there exists a $\pi_{\lambda+\rho} \in \widehat{G}_d$ such that*

$$\Theta_{\lambda+\rho}|_H = (-1)^{\frac{1}{2}\dim G/K} \frac{\sum_{w\in W_{\mathbb{R}}} \operatorname{sgn}(w) e^{w(\lambda+\rho)}}{\prod_{\substack{\alpha\in\Phi, \\ (\lambda+\rho,\alpha)>0}} e^{\alpha/2} - e^{-\alpha/2}}.$$

Conversely, if $\pi \in \widehat{G}_d$ then π is unitarily equivalent to some such $\pi_{\lambda+\rho}$. Moreover, $\pi_{\lambda+\rho}$ is unitarily equivalent to $\pi_{\lambda'+\rho}$ if and only if $\lambda + \rho = w(\lambda' + \rho)$ for some $w \in W_{\mathbb{R}}$. ■

Remark 4.3.8. The above expression for $\Theta_{\lambda+\rho}|_H$ bears an uncanny resemblance to the Weyl character formula; in fact, if G is compact, then that is precisely what it reduces to (see p.100). Harish-Chandra [33] also obtained a formula for the formal degree of $\pi = \pi_{\lambda+\rho}$ (cf. Remark 4.2.3), namely

$$d_\pi = \prod_{\substack{\alpha\in\Phi \\ (\lambda+\rho,\alpha)>0}} \frac{(\lambda+\rho, \alpha)}{(\rho, \alpha)},$$

which of course reduces to the Weyl dimension formula in case G is compact. ▲

4.4 Schmid's Theorem

We continue using the notation of the previous section. In particular, G is real semisimple, H is a compact Cartan subgroup of G, and K is maximal compact subgroup of G containing H. Also, let $\mathfrak{h}_0, \mathfrak{k}_0$ and $\mathfrak{g}_0$ denote the Lie algebras of H, K and G, respectively, and let $\mathfrak{h}, \mathfrak{k}$ and $\mathfrak{g}$ denote their complexifications.

Our attention in this section shall be focused on the problem of explicitly constructing the representation $\pi_{\lambda+\rho} \in \widehat{G}_d$ of Theorem 4.3.7, where $\lambda \in \Lambda$ is such that $\lambda + \rho$ is regular.

As we have remarked, the construction we have in mind is entirely analogous to the one given by the Borel–Weil–Bott theorem (Theorem 2.4.1). Let us therefore briefly recall what this theorem states. Suppose that G is compact. Then H is a maximal torus, and the quotient space $X = G/H$ is a compact complex manifold (actually a smooth projective variety). Every $\lambda \in \Lambda = \widehat{H}$ gives rise to a holomorphic, G-equivariant line bundle $\mathcal{L}_\lambda$ over X, and the Borel–Weil–Bott theorem asserts that $H^q(X, \mathcal{L}_\lambda)$ is nonzero in at most one degree $q = q_\lambda$, at which $H^{q_\lambda}(X, \mathcal{L}_\lambda)$ is a specific irreducible representation.

Thus in order to generalize this to noncompact G, we need an appropriate analogue of the flag variety X. The obvious candidate is the homogeneous space $D = G/H$. A priori, D is only a smooth manifold, but it can be endowed with a G-invariant complex structure (in fact several) as follows. We begin by observing that $\mathfrak{g}_0$ gives rise to an involution (a "complex conjugation") on $\mathfrak{g} = \mathbb{C} \otimes \mathfrak{g}_0$ which we shall denote with a bar. As $\Phi \subset \Lambda \subset i\mathfrak{h}_0^*$, each root $\alpha \in \Phi$ assumes imaginary values on $\mathfrak{h}_0$, and thus $\overline{\mathfrak{g}_\alpha} = \mathfrak{g}_{-\alpha}$. In particular, we can choose a set Φ^+ of positive roots in Φ so that

$$\mathfrak{g} = \overline{\mathfrak{n}} \oplus \mathfrak{h} \oplus \mathfrak{n},$$

where

$$\mathfrak{n} = \oplus_{\alpha \in \Phi^+} \mathfrak{g}_{-\alpha}.$$

The complexified tangent space of D at eH can therefore be identified with $\overline{\mathfrak{n}} \oplus \mathfrak{n}$. It follows that D has a G-invariant almost complex structure. This structure is uniquely determined if we require that the $(1,0)$ part of the tangent space at eH coincides with $\overline{\mathfrak{n}}$; in general, the different choices of almost complex structures correspond to different choices of positive roots. The integrability of any of these almost complex structures amounts to the assertion that $[\mathfrak{g}_\alpha, \mathfrak{g}_\beta] \subset \mathfrak{g}_{\alpha+\beta}$ whenever α, β and $\alpha + \beta$ are in Φ, which we know holds. Thus these almost complex structures are in fact complex structures. In what follows we fix one such G-invariant complex structure on D.

Remark 4.4.1. If the group G is a real form of a complex semisimple $G_{\mathbb{C}}$, then there is an alternative way of realizing these G-invariant complex structures on $D = G/H$. Namely if X is the flag variety of $G_{\mathbb{C}}$, which we think of as the variety of Borel subalgebras of $\mathfrak{g} = \operatorname{Lie} G_{\mathbb{C}}$, then a theorem of Wolf [74] asserts that G acts on X with finitely many orbits.

Hence there must be open orbits. The isotropy subgroup of G at any Borel subalgebra $\mathfrak{b} \in X$ can be shown to coincide with H. So we can identify D with the G-orbit of $\mathfrak{b}$ in X, which must be open since $\dim_{\mathbb{R}} D = \dim_{\mathbb{R}} X$, and hence carries a G-invariant complex structure. For this reason spaces of the form G/H for G a real semisimple Lie group and H a compact Cartan subgroup are called **flag domains**.

In particular, if G is compact, then it can be regarded as a real form and maximal compact subgroup of a unique complex semisimple $G_{\mathbb{C}}$. One knows here that the action of G on X is transitive (any orbit is both open and, because G is compact, closed), which means that we can identify D with X (as we already know).

Much more information can be found in [74]. ▲

Each $\lambda \in \Lambda = \widehat{H}$ gives rise to a smooth G-equivariant line bundle $\mathcal{L}_\lambda = G \times_\lambda \mathbb{C}$ over D (see Section 2.2). Just as before, the space of sections of $\mathcal{L}_\lambda$ over an open set $U \subset D$ is naturally isomorphic to

$$\{f \in C^\infty(p^{-1}(U)) \colon f(gh) = \lambda(h)^{-1} f(g) \text{ for all } g \in p^{-1}(U) \text{ and } h \in H\}, \tag{4.4.1}$$

where p is the quotient map $G \to D = G/H$. According to Griffiths and Schmid [29, §2] (see also [54, pp. 5–6]), $\mathcal{L}_\lambda$ can be made into a holomorphic line bundle, and a section f in (4.4.1) is holomorphic if and only if $r_{\mathbb{C}}(Z)f = 0$ for all $Z \in \mathfrak{n}$ (compare Lemma 2.3.4).

As the sheaf of holomorphic sections of $\mathcal{L}_\lambda$ is G-equivariant, we get an action of G on the sheaf cohomology space $H^*(D, \mathcal{L}_\lambda)$ as before. However, unlike in the compact case, this cohomology space need not be finite-dimensional, which is fine because we are no longer expecting finite-dimensional representations, but there remains the issue of unitarity. There is at least a way to topologize $H^*(D, \mathcal{L}_\lambda)$ so that it becomes a Fréchet space and then one has many deep representation-theoretic results concerning the (continuous) action of G on $H^*(D, \mathcal{L}_\lambda)$. For our purposes, however, it will be more appropriate to consider the L^2 *cohomology* space $H^*_{(2)}(D, \mathcal{L}_\lambda)$ instead. The relevant definitions are as follows.

Let $\Omega^p_c(D, \mathcal{L}_\lambda)$ denote the space *compactly supported*, smooth, $\mathcal{L}_\lambda$-valued $(0,p)$ forms on D, and let $\overline{\partial} \colon \Omega^p_c(D, \mathcal{L}_\lambda) \to \Omega^{p+1}_c(D, \mathcal{L}_\lambda)$ be the restriction of the usual Dolbeault operator $\Omega^p(D, \mathcal{L}_\lambda) \to \Omega^{p+1}(D, \mathcal{L}_\lambda)$ to $\Omega^p_c(D, \mathcal{L}_\lambda)$ (it is obvious that its image lives in $\Omega^{p+1}_c(D, \mathcal{L}_\lambda)$).

Now put G-invariant hermitian metrics on D and $\mathcal{L}_\lambda$—the latter is unique up to a constant since H is compact. The natural action of G on $\Omega_c^p(D, \mathcal{L}_\lambda)$ commutes with $\overline{\partial}$, and, by integration over D, induces a G-invariant inner product on $\Omega_c^p(D, \mathcal{L}_\lambda)$. With respect to this inner product, $\overline{\partial}$ has a formal adjoint $\overline{\partial}^*\colon \Omega_c^p(D, \mathcal{L}_\lambda) \to \Omega_c^{p-1}(D, \mathcal{L}_\lambda)$ and, as usual, we define the Laplace–Beltrami operator by $\Delta = \overline{\partial}^*\overline{\partial} + \overline{\partial}\,\overline{\partial}^*$; this is an elliptic, G-invariant differential operator on $\Omega_c^p(D, \mathcal{L}_\lambda)$.

Let $L^p_{(2)}(D, \mathcal{L}_\lambda)$ denote the L^2-closure of $\Omega_c^p(D, \mathcal{L}_\lambda)$ with respect to the above inner product, that is, $L^p_{(2)}(D, \mathcal{L}_\lambda)$ is the subspace of square-integrable $\mathcal{L}_\lambda$-valued forms of type $(0, p)$. The operators $\overline{\partial}$, $\overline{\partial}^*$ and Δ admit closed extensions to $L^p_{(2)}(D, \mathcal{L}_\lambda)$. We denote the resulting (unbounded) operators by the same letters. Finally, we set

$$H^p_{(2)}(D, \mathcal{L}_\lambda) = \ker \Delta|_{L^p_{(2)}(D, \mathcal{L}_\lambda)}.$$

This is called the pth L^2 **cohomology** space of $\mathcal{L}_\lambda$; it is a Hilbert space and the natural action of G on $L^p_{(2)}(D, \mathcal{L}_\lambda)$ induces a unitary representation on $H^p_{(2)}(D, \mathcal{L}_\lambda)$.

As usual, we find that

$$H^p_{(2)}(D, \mathcal{L}_\lambda) = \ker \overline{\partial}|_{L^p_{(2)}(D, \mathcal{L}_\lambda)} \cap \ker \overline{\partial}^*|_{L^p_{(2)}(D, \mathcal{L}_\lambda)}.$$

More remarkably, we also have that

$$H^p_{(2)}(D, \mathcal{L}_\lambda) = \ker \Delta|_{L^p_{(2)}(D, \mathcal{L}_\lambda) \,\cap\, \Omega^p(D, \mathcal{L}_\lambda)}.$$

This is because the metric on D, being G-invariant, is complete, which, according to Andreotti and Vesentini [2, Proposition 7], implies that the largest and smallest closed extensions of Δ coincide.

To state Schmid's theorem we need to introduce one last piece of notation. To that end, we say that a root $\alpha \in \Phi$ is **compact** if the root space $\mathfrak{g}_\alpha$ is a subspace of $\mathfrak{k}$. If on the other hand $\mathfrak{g}_\alpha$ is a subspace of the orthogonal complement of $\mathfrak{k}$ (with respect to the Killing form), then α is said to be **noncompact**. We denote by Φ_c and Φ_n the sets of compact and noncompact roots, respectively. As $\dim \mathfrak{g}_\alpha = 1$ for all $\alpha \in \Phi$, we have $\Phi = \Phi_c \cup \Phi_n$. Also note that Φ_c is the root system of $(\mathfrak{k}, \mathfrak{h})$, hence is a subsystem of $(\mathfrak{g}, \mathfrak{h})$. Put

$$\Phi_c^+ = \Phi_c \cap \Phi^+ \quad \text{and} \quad \Phi_n^+ = \Phi_n \cap \Phi^+.$$

Then for $\lambda \in \Lambda$ we set

$$p_\lambda = |\{\alpha \in \Phi_c^+ \colon (\lambda + \rho, \alpha) < 0\}| + |\{\alpha \in \Phi_n^+ \colon (\lambda + \rho, \alpha) > 0\}|.$$

Observe that if G is compact then $p_\lambda = q_\lambda$ (in the notation of Chapter 2).

Theorem 4.4.2 (Schmid)**.** *Assume the preceding notation and let $\lambda \in \Lambda$.*

(i) *If $\lambda + \rho$ is singular, then $H^p_{(2)}(D, \mathcal{L}_\lambda) = 0$ for all $p \geq 0$.*

(ii) *If $\lambda + \rho$ is regular, then $H^p_{(2)}(D, \mathcal{L}_\lambda)$ is nonzero if and only if $p = p_\lambda$, in which case $H^{p_\lambda}_{(2)}(D, \mathcal{L}_\lambda) = \pi_{\lambda+\rho}$.* ■

We shall outline a proof of this theorem in the next section.

Remarks 4.4.3.

(i) Schmid's theorem was conjectured by Langlands in [50]. Langlands was inspired by a vanishing theorem of Griffiths for L^2 cohomology (see [29, §7]), which asserted that for certain λ, $H^p_{(2)}(D, \mathcal{L}_\lambda) = 0$ for all $p \neq p_\lambda$. By formally applying the Atiyah–Bott fixed point formula to the action of H on $H^{p_\lambda}_{(2)}(D, \mathcal{L}_\lambda)$, Langlands found that the value of the global character of $H^{p_\lambda}_{(2)}(D, \mathcal{L}_\lambda)$ at a regular element γ of H coincided with Harish-Chandra's formula for $\Theta_{\lambda+\rho}(\gamma)$ (see Theorem 4.3.7). This led him to conjecture that $H^{p_\lambda}_{(2)}(D, \mathcal{L}_\lambda)$ should be the representation $\pi_{\lambda+\rho}$.

(ii) The similarity between Schmid's theorem and the Borel–Weil–Bott theorem is obvious. In fact, the former reduces to the latter in case G is compact (see below). One particular difference is that, unlike in the compact setting, one needs to use cohomology in degree > 0 to realize the entire discrete series of an arbitrary semisimple Lie group G. That is, *there is no L^2 Borel–Weil theorem.* For example, if $G = \mathrm{SO}(4,1)_0$, then Dixmier [19] has shown that $\widehat{G}_d \neq \emptyset$ but $H^0_{(2)}(D, \mathcal{L}_\lambda) = 0$ for all $\lambda \in \Lambda$. ▲

In concluding this section we show that Schmid's theorem reduces to the Borel–Weil–Bott theorem when G is compact.

We begin by noting that, since Δ is elliptic, there is a natural map

$$H^p_{(2)}(D, \mathcal{L}_\lambda) \to H^p(D, \mathcal{L}_\lambda)$$

given by sending each $\omega \in H^p_{(2)}(D, \mathcal{L}_\lambda)$, which is in particular a $\overline{\partial}$-closed smooth Dolbeault form, to the Dolbeault cohomology class it determines. In general, this map is neither injective nor surjective. However, if D is compact (which is the case if and only if G is compact), then the Hodge theorem asserts that this map is an isomorphism. So assume that G is compact. Then most of Schmid's theorem reduces to the Borel–Weil–Bott theorem immediately; all that needs elaboration is the second half of assertion (ii). Thus let $\lambda \in \Lambda$ and assume that $\lambda + \rho$ is regular. Then $p_\lambda = q_\lambda$, and according to Theorem 4.3.7, the global character of the representation $\pi_{\lambda+\rho}$ satisfies

$$\Theta_{\lambda+\rho}|_H = \frac{\sum_{w \in W_{\mathbb{R}}} \operatorname{sgn}(w) e^{w(\lambda+\rho)}}{\prod_{\substack{\alpha \in \Phi, \\ (\lambda+\rho,\alpha)>0}} e^{\alpha/2} - e^{-\alpha/2}}. \tag{4.4.2}$$

Let w' denote the unique element of $W_{\mathbb{R}}$ such that $w'(\lambda+\rho)$ is dominant. Since the map $w \to ww'$ is a permutation of $W_{\mathbb{R}}$, the numerator of (4.4.2) is simply

$$\sum_{w \in W_{\mathbb{R}}} \operatorname{sgn}(ww') e^{w(\mu+\rho)} = \operatorname{sgn}(w') \sum_{w \in W_{\mathbb{R}}} \operatorname{sgn}(w) e^{w(\mu+\rho)},$$

where $\mu = w'(\lambda+\rho) - \rho$. The set $S = \{\alpha \in \Phi\colon (\lambda+\rho, \alpha) > 0\}$ determines a positive system in Φ. The element of $W_{\mathbb{R}}$ which transforms S to our fixed positive system Φ^+ is precisely w', that is, $w'S = \Phi^+$. The denominator of (4.4.2) thus reduces to

$$\prod_{\alpha \in \Phi^+} e^{(w')^{-1}\alpha/2} - e^{-(w')^{-1}\alpha/2} = (-1)^{\ell(w')} \prod_{\alpha \in \Phi^+} e^{\alpha/2} - e^{-\alpha/2} = \operatorname{sgn}(w') \prod_{\alpha \in \Phi^+} e^{\alpha/2} - e^{-\alpha/2}.$$

Since G is compact, $\pi_{\lambda+\rho}$ is finite-dimensional. This implies that $\Theta_{\lambda+\rho}$ is the usual character $\chi_{\pi_{\lambda+\rho}}$, which is completely determined by its values on any maximal torus in G (such as H). Putting all this together, we find that

$$\chi_{\pi_{\lambda+\rho}} = \frac{\sum_{w \in W_{\mathbb{R}}} \operatorname{sgn}(w) e^{w(\mu+\rho)}}{\prod_{\alpha \in \Phi^+} e^{\alpha/2} - e^{-\alpha/2}}.$$

The Weyl character formula now implies that the highest weight of $\pi_{\lambda+\rho}$ is μ, in accordance with the Borel–Weil–Bott theorem.

4.5 On the Proof of Schmid's Theorem

We retain the notation of the previous sections. Our goal in this section is to outline a proof of Theorem 4.4.2.

In [54], Schmid used the Plancherel decomposition of $L^2(G)$ to show that

$$H^p_{(2)}(D, \mathcal{L}_\lambda) = \int_{\widehat{G}}^{\oplus} H^*_\pi \,\widehat{\otimes}\, \mathcal{H}^p(\pi)_{-\lambda} \, d\mu(\pi),$$

where $\mathcal{H}^p(\pi)$ is a certain (generally infinite-dimensional) Hodge-theoretically defined H-module, which he called the *formal harmonic space* attached to $\pi \in \widehat{G}$, and $\mathcal{H}^p(\pi)_{-\lambda}$ is the $-\lambda$ weight space of $\mathcal{H}^p(\pi)$ under the action of H. Later [55, Theorem 3.1], Schmid obtained an H-module isomorphism between $\mathcal{H}^p(\pi)$ and the $\mathfrak{n}$-cohomology space $H^p(\mathfrak{n}, H^\infty_\pi)$. Here H^∞_π is the space of smooth vectors in H_π.[8] It follows that

$$H^p_{(2)}(D, \mathcal{L}_\lambda) = \int_{\widehat{G}}^{\oplus} H^*_\pi \,\widehat{\otimes}\, H^p(\mathfrak{n}, H^\infty_\pi)_{-\lambda} \, d\mu(\pi). \tag{4.5.1}$$

(Compare with equation (3.5.3) of Section 3.5.1.) Schmid then makes use of a lemma of Casselman and Osborne [16] which asserts that, if $H^p(\mathfrak{n}, H^\infty_\pi)_{-\lambda}$ is nonzero, then H^∞_π has infinitesimal character $\chi_{-\lambda-\rho}$.[9] However, one knows that there are only finitely many distinct representations in $\widehat{G}$ with the same infinitesimal character (see [35]). It follows that the integral in (4.5.1) is nonzero only on a finite set, in which case the occurring representations have positive Plancherel measure, hence belong to the discrete series by Theorem 4.2.3. That is, $H^p_{(2)}(D, \mathcal{L}_\lambda)$ is a finite (possibly empty) direct sum of discrete series representations, all having infinitesimal character $\chi_{-\lambda-\rho}$ (compare [55, Corollary 3.22]). According to a result of Atiyah and Schmid [4, Corollary 6.13], if χ_μ is the infinitesimal character of a discrete series representation, then μ must be regular.[10] This proves part (i) of Schmid's theorem, namely that $H^*_{(2)}(D, \mathcal{L}_\lambda) = 0$ if $\lambda + \rho$ is singular.

[8]By definition, $H^\infty_\pi = \{v \in H_\pi : \text{the map } g \mapsto \pi(g)v \text{ is smooth }\}$.

[9]This Casselman–Osborne result ought to be viewed in the context of generalizing Kostant's theorem (Theorem 3.4.1) to infinite-dimensional $\mathfrak{g}$-modules which possess an infinitesimal character (of which H^∞_π is an example).

[10]The proof of this result makes use of Atiyah's L^2 index theorem; the argument is outlined in [59, pp. 100–101].

Thus we may assume that $\lambda + \rho$ is regular. From the preceding discussion we obtain the following reciprocity result (which is valid for any $\lambda \in \Lambda$):

Proposition 4.5.1 (Schmid [55, Corollary 3.23])**.** *Let $\pi \in \widehat{G}_d$. Then the multiplicity of H_π^* in $H^p_{(2)}(D, \mathcal{L}_\lambda)$ is equal to* $\dim H^p(\mathfrak{n}, H_\pi^\infty)_{-\lambda}$. ■

So what is now required is an analysis of the spaces $H^p(\mathfrak{n}, H_\pi^\infty)_{-\lambda}$ for $\pi \in \widehat{G}_d$ such that H_π^∞ has infinitesimal character $\chi_{-\lambda-\rho}$. In terms of Harish-Chandra's parametrization (Theorem 4.3.7), these aforementioned discrete series representations are those whose global character is $\Theta_{w(-\lambda-\rho)}$ for some $w \in W_{\mathbb{R}}$. That is, there is only one of them (up to unitary equivalence), namely $\pi = \pi_{-\lambda-\rho}$. For such π, Schmid [55, Theorem 4.1] shows that $H^p(\mathfrak{n}, H_\pi^\infty)_{-\lambda}$ vanishes unless $p = p_\lambda$, in which case it is one-dimensional. Assertion (ii) of Schmid's theorem now follows from Proposition 4.5.1 together with the fact that $H_{\pi_{\lambda+\rho}} = (H_{\pi_{-\lambda-\rho}})^*$.

Remark 4.5.2. One does not have to use Harish-Chandra's parametrization of $\widehat{G}_d$ (Theorem 4.3.7) as we did. Indeed, in Atiyah–Schmid [4] one can find another approach to Schmid's theorem, which not only avoids Harish-Chandra's parametrization, but in fact proves it in the process. This approach is sketched in [59, §3]. ▲

Bibliography

[1] D. Akhiezer, *Lie Group Actions in Complex Analysis*, Aspects of Mathematics, E27, Friedr. Vieweg & Sohn, Braunschweig, 1995.

[2] A. Andreotti and E. Vesentini, *Carleman estimates for the Laplace-Beltrami equation on complex manifolds*, Inst. Hautes Études Sci. Publ. Math. (1965), no. 25, 81–130.

[3] F. Aribaud, *Une nouvelle démonstration d'un théorème de R. Bott et B. Kostant*, Bull. Soc. Math. France **95** (1967), 205–242.

[4] M. Atiyah and W. Schmid, *A geometric construction of the discrete series for semisimple Lie groups*, Invent. Math. **42** (1977), 1–62.

[5] ———, *Erratum: "A geometric construction of the discrete series for semisimple Lie groups" [Invent. Math.* **42** *(1977), 1–62]*, Invent. Math. **54** (1979), no. 2, 189–192.

[6] V. Bargmann, *Irreducible unitary representations of the Lorentz group*, Ann. of Math. (2) **48** (1947), 568–640.

[7] R. J. Baston and M. G. Eastwood, *The Penrose Transform: Its Interaction with Representation Theory*, Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York, 1989.

[8] I. Biswas and S. Subramanian, *On homogeneous vector bundles*, Bull. Sci. Math. **132** (2008), no. 5, 419–424.

[9] A. Borel, *Linear Algebraic Groups*, Second ed., Graduate Texts in Mathematics, vol. 126, Springer-Verlag, New York, 1991.

[10] A. Borel and F. Hirzebruch, *Characteristic classes and homogeneous spaces. I*, Amer. J. Math. **80** (1958), 458–538.

[11] ———, *Characteristic classes and homogeneous spaces. II*, Amer. J. Math. **81** (1959), 315–382.

[12] ———, *Characteristic classes and homogeneous spaces. III*, Amer. J. Math. **82** (1960), 491–504.

[13] R. Bott, *Homogeneous vector bundles*, Ann. of Math. (2) **66** (1957), 203–248.

[14] ———, *On induced representations*, The Mathematical Heritage of Hermann Weyl, Proc. Sympos. Pure Math., vol. 48, Amer. Math. Soc., Providence, RI, 1988, pp. 1–13.

[15] P. Cartier, *Remarks on "Lie algebra cohomology and the generalized Borel-Weil theorem", by B. Kostant*, Ann. of Math. (2) **74** (1961), 388–390.

[16] W. Casselman and M. S. Osborne, *The $\mathfrak{n}$-cohomology of representations with an infinitesimal character*, Compositio Math. **31** (1975), no. 2, 219–227.

[17] M. Demazure, *Une démonstration algébrique d'un théorème de Bott*, Invent. Math. **5** (1968), 349–356.

[18] ———, *A very simple proof of Bott's theorem*, Invent. Math. **33** (1976), no. 3, 271–272.

[19] J. Dixmier, *Représentations intégrables du groupe de De Sitter*, Bull. Soc. Math. France **89** (1961), 9–41.

[20] A. Frölicher and A. Nijenhuis, *A theorem on stability of complex structures*, Proc. Nat. Acad. Sci. U.S.A. **43** (1957), 239–241.

[21] W. Fulton, *Young Tableaux: With Applications to Representation Theory and Geometry*, London Mathematical Society Student Texts, vol. 35, Cambridge University Press, Cambridge, 1997.

[22] W. Fulton and J. Harris, *Representation Theory: A First Course*, Graduate Texts in Mathematics, vol. 129, Springer-Verlag, New York, 1991.

[23] I. M. Gelfand and M. A. Naimark, *Unitary representations of the Lorentz group*, Izvestiya Akad. Nauk SSSR. Ser. Mat. **11** (1947), 411–504.

[24] R. Godement, *Sur les relations d'orthogonalité de V. Bargmann. I. Résultats préliminaires*, C. R. Acad. Sci. Paris **225** (1947), 521–523.

[25] ———, *Sur les relations d'orthogonalité de V. Bargmann. II. Démonstration générale*, C. R. Acad. Sci. Paris **225** (1947), 657–659.

[26] ———, *Topologie Algébrique et Théorie des Faisceaux*, Hermann, Paris, 1958.

[27] R. Goodman and N. R. Wallach, *Representations and Invariants of the Classical Groups*, Encyclopedia of Mathematics and its Applications, vol. 68, Cambridge University Press, Cambridge, 1998.

[28] P. Griffiths and J. Harris, *Principles of Algebraic Geometry*, Wiley Classics Library, John Wiley & Sons Inc., New York, 1994.

[29] P. Griffiths and W. Schmid, *Locally homogeneous complex manifolds*, Acta Math. **123** (1969), 253–302.

[30] A. Grothendieck, *Sur la classification des fibrés holomorphes sur la sphère de Riemann*, Amer. J. Math. **79** (1957), 121–138.

[31] Harish-Chandra, *Representations of semisimple Lie groups. II*, Trans. Amer. Math. Soc. **76** (1954), 26–65.

[32] ———, *Representations of semisimple Lie groups. III*, Trans. Amer. Math. Soc. **76** (1954), 234–253.

[33] ———, *Integrable and square-integrable representations of a semisimple Lie group*, Proc. Nat. Acad. Sci. U.S.A. **41** (1955), 314–317.

[34] ———, *Representations of semisimple Lie groups. IV*, Amer. J. Math. **77** (1955), 743–777.

[35] ———, *The characters of semisimple Lie groups*, Trans. Amer. Math. Soc. **83** (1956), 98–163.

[36] ———, *Invariant eigendistributions on a semisimple Lie group*, Trans. Amer. Math. Soc. **119** (1965), 457–508.

[37] ———, *Discrete series for semisimple Lie groups. II. Explicit determination of the characters*, Acta Math. **116** (1966), 1–111.

[38] R. Hartshorne, *Algebraic Geometry*, Graduate Texts in Mathematics, vol. 52, Springer-Verlag, New York, 1977.

[39] R. A. Herb, *Harish-Chandra and his work*, Bull. Amer. Math. Soc. **25** (1991), 1–17.

[40] J. E. Humphreys, *Linear Algebraic Groups*, Graduate Texts in Mathematics, vol. 21, Springer-Verlag, New York, 1975.

[41] D. Huybrechts, *Complex Geometry: An Introduction*, Universitext, Springer-Verlag, Berlin, 2005.

[42] J. C. Jantzen, *Representations of Algebraic Groups*, Second ed., Mathematical Surveys and Monographs, vol. 107, American Mathematical Society, Providence, RI, 2003.

[43] G. R. Kempf, *Linear systems on homogeneous spaces*, Ann. of Math. **103** (1976), no. 3, 557–591.

[44] A. W. Knapp, *Representation Theory of Semisimple Groups: An Overview Based on Examples*, Princeton Landmarks in Mathematics, Princeton University Press, Princeton, NJ, 2001.

[45] ———, *Lie Groups Beyond an Introduction*, Second ed., Progress in Mathematics, vol. 140, Birkhäuser Boston Inc., Boston, MA, 2002.

[46] A. W. Knapp and D. A. Vogan, Jr., *Cohomological Induction and Unitary Representations*, Princeton Mathematical Series, vol. 45, Princeton University Press, Princeton, NJ, 1995.

[47] B. Köck, *The Grothendieck-Riemann-Roch theorem for group scheme actions*, Ann. Sci. École Norm. Sup. (4) **31** (1998), no. 3, 415–458.

[48] B. Kostant, *Lie algebra cohomology and the generalized Borel-Weil theorem*, Ann. of Math. **74** (1961), 329–387.

[49] ———, *Lie algebra cohomology and generalized Schubert cells*, Ann. of Math. **77** (1963), 72–144.

[50] R. P. Langlands, *Dimension of spaces of automorphic forms*, Algebraic Groups and Discontinuous Subgroups (Proc. Sympos. Pure Math., Boulder, Colo., 1965), Amer. Math. Soc., Providence, R.I., 1966, pp. 253–257.

[51] J. Lurie, *A proof of the Borel-Weil-Bott theorem*, `http://www.math.harvard.edu/~lurie/papers/bwb.pdf`.

[52] D. Montgomery, *Simply connected homogeneous spaces*, Proc. Amer. Math. Soc. **1** (1950), 467–469.

[53] C. Procesi, *Lie Groups: An Approach through Invariants and Representations*, Universitext, Springer, New York, 2007.

[54] W. Schmid, *On a conjecture of Langlands*, Ann. of Math. **93** (1971), 1–42.

[55] ———, *L^2-cohomology and the discrete series*, Ann. of Math. **103** (1976), 375–394.

[56] ———, *Representations of semi-simple Lie groups*, Representation Theory of Lie Groups: Proceedings of the SRC/LMS Research Symposium held in Oxford, June 28–July 15, 1977, London Mathematical Society Lecture Note Series, vol. 34, Cambridge University Press, Cambridge, 1979, pp. 185–235.

[57] ———, *Homogeneous complex manifolds and representations of semisimple Lie groups*, Representation theory and harmonic analysis on semisimple Lie groups, Math. Surveys Monogr., vol. 31, Amer. Math. Soc., Providence, RI, 1989, pp. 223–286.

[58] ———, *The mathematical legacy of the paper: "Homogeneous vector bundles" [Ann. of Math. (2)* **66** *(1957), 203–248] by R. Bott*, Raoul Bott: Collected Papers, Vol. 1, Contemp. Mathematicians, Birkhäuser Boston, Boston, MA, 1994, pp. 40–47.

[59] ———, *Discrete series*, Representation Theory and Automorphic Forms (Edinburgh, 1996), Proc. Sympos. Pure Math., vol. 61, Amer. Math. Soc., Providence, RI, 1997, pp. 83–113.

[60] ———, *Geometric methods in representation theory*, Poisson Geometry, Deformation Quantisation and Group Representations, Cambridge University Press, Cambridge, 2005.

[61] M. Sepanski, *Compact Lie Groups*, Graduate Texts in Mathematics, vol. 235, Springer, New York, 2007.

[62] J.-P. Serre, *Représentations linéaires et espaces homogènes kählériens des groupes de Lie compacts (d'après Armand Borel et André Weil)*, Séminaire Bourbaki, Vol. 2, Soc. Math. France, Paris, 1995, pp. Exp. No. 100, 447–454.

[63] ———, *Complex Semisimple Lie Algebras*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2001.

[64] D. Snow, *Homogeneous vector bundles*, `http://www.nd.edu/~snow/Papers/HomogVB.pdf`.

[65] T. A. Springer, *Linear Algebraic Groups*, Second ed., Progress in Mathematics, vol. 9, Birkhäuser Boston Inc., Boston, MA, 1998.

[66] P. Tauvel and R. W. T. Yu, *Lie Algebras and Algebraic Groups*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2005.

[67] J. L. Taylor, *Several Complex Variables with Connections to Algebraic Geometry and Lie Groups*, Graduate Studies in Mathematics, vol. 46, American Mathematical Society, Providence, RI, 2002.

[68] J. Tits, *Sur certaines classes d'espaces homogènes de groupes de Lie*, Acad. Roy. Belg. Cl. Sci. Mém. Coll. in 8° **29** (1955), no. 3, 268.

[69] D. A. Vogan, Jr., *Representations of Real Reductive Lie Groups*, Progress in Mathematics, vol. 15, Birkhäuser Boston, Mass., 1981.

[70] ———, *Unitary Representations of Reductive Lie Groups*, Annals of Mathematics Studies, vol. 118, Princeton University Press, Princeton, NJ, 1987.

[71] H.-C. Wang, *Closed manifolds with homogeneous complex structure*, Amer. J. Math. **76** (1954), 1–32.

[72] C. Weibel, *An Introduction to Homological Algebra*, Cambridge Studies in Advanced Mathematics, vol. 38, Cambridge University Press, Cambridge, 1994.

[73] F. L. Williams, *Frobenius reciprocity and Lie group representations on $\bar{\partial}$ cohomology spaces*, Enseign. Math. (2) **28** (1982), no. 1-2, 3–30.

[74] J. A. Wolf, *The action of a real semisimple group on a complex flag manifold. I. Orbit structure and holomorphic arc components*, Bull. Amer. Math. Soc. **75** (1969), 1121–1237.

[75] R. Zierau, *Representations in Dolbeault cohomology*, Representation Theory of Lie Groups, IAS/Park City Math. Ser., vol. 8, Amer. Math. Soc., Providence, RI, 2000, pp. 91–146.