Boundaries of K-types in discrete series

by

Markéta Havlíčková

S.B., Massachusetts Institute of Technology, June 2002

Submitted to the Department of Mathematics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2008

Author ....
Department of Mathematics
May 12, 2008

Certified by ....
Michael J. Hopkins
Professor of Mathematics, Harvard University
Thesis Supervisor

Certified by ....
David A. Vogan
Professor of Mathematics
Thesis Supervisor

Accepted by ....
David S. Jerison
Chairman, Department Committee on Graduate Students

Boundaries of K-types in discrete series

by

Markéta Havlíčková

Submitted to the Department of Mathematics
on May 12, 2008, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Abstract: A fundamental problem about irreducible representations of a reductive Lie group G is understanding their restriction to a maximal compact subgroup K. In certain important cases, known as the discrete series, we have a formula that gives the multiplicity of any given irreducible K-representation (or K-type) as an alternating sum. It is not immediately clear from this formula which K-types, indexed by their highest weights, have non-zero multiplicity. Evidence suggests that the collection is very close to a set of lattice points in a noncompact convex polyhedron. In this paper we shall describe a recursive algorithm for finding the boundary facets of this polyhedron.

Thesis Supervisor: Michael J. Hopkins
Title: Professor of Mathematics, Harvard University

Thesis Supervisor: David A. Vogan
Title: Professor of Mathematics

Acknowledgments

I would like to thank Bill Barker and Roger Howe for introducing me to Lie groups. Dan Stroock for never filtering his advice. George Lusztig for asking the right questions. Richard Borcherds for seeing through to the heart of every problem. David Vogan for spending endless hours showing me the world of Representation theory; for all his suggestions, care, and support. Mike Hopkins for telling me about Algebraic Topology and listening to Lie groups; for teaching me how to think about Mathematics; and, most of all, for always being my friend.

Contents

List of Figures

Chapter 1

Introduction

One extremely important aspect of understanding the representations of a reductive group G is understanding how they behave under restriction to a maximal subgroup K of G. It is a question analogous to finding the weights of the irreducible K-representation E_μ of highest weight μ. For compact groups, the Kostant multiplicity formula calculates the multiplicity of any given weight in E_μ. We also know a simple geometric rule for determining which weights actually occur: those lying in the convex hull of the Weyl group orbit of μ. This rule provides an extremely simple and quick way to get the most basic picture of E_μ. At present, there is no parallel to that for reductive groups and K-types.

A particularly interesting class of representations, which appears in automorphic forms and mathematical physics, is the discrete series. These are the irreducible subspaces of the regular representation $L^2(G)$. In their case, we have an explicit formula that can be used to determine the multiplicity of any given irreducible K-representation in X. It is not at all obvious which multiplicities will come out to be nonzero. There are infinitely many distinct ones, and they seem to be essentially all weights inside a convex polyhedron, subject to the obvious condition that any two differ by a sum of roots in the Lie algebra of G.

The problem of finding the boundary facets, or "edges", of this polyhedron was suggested by David Vogan, and the first results about it are due to Mark Sepanski. In this thesis, we shall describe a recursive algorithm for finding these facets, with induction on the rank of G.

We shall begin by describing the fundamental ideas behind the algorithm in this chapter, demonstrating most techniques on the well-understood case of compact groups. Section 1.1 will be spent on compact groups entirely, reviewing what we know about their representations. In Section 1.2 we shall introduce reductive groups and give some examples of the discrete series representations. In Section 1.3 we shall go back to compact groups, getting some idea on what "edges" are and how we can look for them. We shall use these ideas in Section 1.4 to sketch the process of finding edges in discrete series. We shall also state the main results of Mark Sepanski, as well as the main theorem in this paper.

In Chapter 2 we shall set up the stage by giving precise definitions of the basic objects we need to work with, as well as state some results which will come in handy later. The actual process of finding edges in a discrete series X has two main steps. The first is to find a nonzero map to X from a certain cohomologically induced module: we shall do this in Chapter 3. In Chapter 4, we will discuss how this map restricts the K-types of X. The main tool there will be the step algebra, which was introduced by Mickelsson in [10] and acts on the K-highest weights of any G-representation. Finally, in Chapter 5 we shall describe the recursive algorithm for finding edges, show how it works on some examples, and finish with some suggestions for future research.

1.1 An analogy

Let H be a compact torus with Lie algebra $\mathfrak{h}_0$, $\mathfrak{h}_0^*$ the dual of $\mathfrak{h}_0$, and $\hat{H}$ the lattice of analytically integral weights in $i\mathfrak{h}_0^*$. The irreducible representations of H are one-dimensional, one for each $\nu \in \hat{H}$. We shall denote the space corresponding to ν by $F(\nu)$, and its character by e^ν.

Let K be a compact connected Lie group with Lie algebra $\mathfrak{k}_0$ and H its maximal compact torus. We shall generally work with the complexified Lie algebras of $\mathfrak{k}_0$ and $\mathfrak{h}_0$: $\mathfrak{k} = \mathfrak{k}_0 \otimes_{\mathbb{R}} \mathbb{C}$ and $\mathfrak{h} = \mathfrak{h}_0 \otimes_{\mathbb{R}} \mathbb{C}$. We denote by $U(\mathfrak{k})$ the universal enveloping algebra of $\mathfrak{k}$. For any Ad_H-stable subalgebra $\mathfrak{g}' \subset \mathfrak{g}$ we let $\Delta(\mathfrak{g}')$ be the set of roots in $\mathfrak{g}'$.

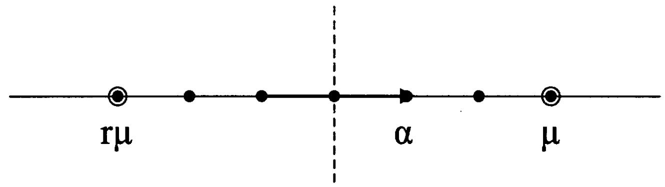

Figure 1-1: $\mathbf{E}_{3\alpha}$ for SU(2)

Let $\langle\ ,\ \rangle$ be the pairing between $\mathfrak{h}$ and $\mathfrak{h}^*$, and for any root α of $\mathfrak{k}$, we denote $\check{\alpha}$ the associated coroot in $\mathfrak{h}$.

Choose a set of positive roots Δ_+^K in $\mathfrak{k}$, and let ρ_c be the half sum: $\rho_c = \frac{1}{2}\sum_{\alpha\in\Delta_+^K}\alpha$. Let $\mathfrak{k}_+ \subset \mathfrak{k}$ be the nilpotent subalgebra with $\Delta(\mathfrak{k}_+) = \Delta_+^K$. Finally, let W_K be the Weyl group of K, generated by reflections about the simple roots in $\mathfrak{k}$. We have an associated length function $l : W_K \to \mathbb{Z}_{\geq 0}$, which for any element $w \in W_K$ returns the minimal number of simple reflections needed to write w. We will always denote the longest Weyl group element by w_0, of length $l(w_0) = S$.

The irreducible representations of K, which we shall refer to as K-types, are indexed by the dominant weights in $\hat{H}$. Each dominant $\mu \in \hat{H}$ corresponds to a representation E_μ with highest weight μ. As a representation of $\mathfrak{h}$, E_μ splits into a direct sum $E_\mu = \bigoplus_{\nu\in\hat{H}} E_\mu(\nu)$ where $E_\mu(\nu)$ is the weight ν subspace of E_μ. Let $m_\nu(E_\mu)$ be the dimension of $E_\mu(\nu)$; then $E_\mu(\nu)|_H$ is a direct sum of $m_\nu(E_\mu)$ copies of $F(\nu)$. We can therefore write $E_\mu = \bigoplus_{\nu\in\hat{H}} m_\nu(E_\mu)F(\nu)$; in other words, the character of E_μ, restricted to H, is $\chi_\mu = \sum_{\nu\in\hat{H}} m_\nu(E_\mu)e^\nu$.

The Weyl character formula gives us χ_μ as a ratio of two trigonometric polynomials:

$$\chi_\mu = \sum_{w\in W_K} \frac{(-1)^{l(w)}e^{w(\mu+\rho_c)}}{\Delta_K} \tag{1.1}$$

Here Δ_K is the Weyl denominator: $\Delta_K = \sum w \in W_K(-1)^{l(w)}e^{w\rho_c}$.

Given any particular weight ν, one can compute the multiplicity of $F(\nu)$ in E_μ by carrying out finitely many steps of the long division in the formula. If one wishes only to know which weights actually occur, there is a much simpler way to find out: the weights that have nonzero multiplicity are those congruent to μ modulo the root lattice of $\mathfrak{k}$, and lying in the convex hull of the set of points $\{w\mu\}_{w\in W_K}$ (see for example Chapter 14 in [1]).

Example 1. The most basic example is for the Lie algebra $\mathfrak{su}(2)$, with positive root α. A representation E_μ has weights $\{\mu, \mu-\alpha, \ldots, -\mu\}$, all occurring with multiplicity one. These lie inside the one-dimensional convex set with vertices $\{\mu, -\mu\}$. Note that $-\mu = r\mu$ where r is the reflection about α. A representation with highest weight $\mu = 3\alpha$ is shown in Figure 1-1.

Example 2. For a more illustrative example, let $\mathfrak{k}$ be the Lie algebra $\mathfrak{su}(3)$, with simple roots α_1 and α_2, and third positive root $\alpha_3 = \alpha_1 + \alpha_2$. Let r_1 and r_2 be reflections about α_1 and α_2, respectively, and let $r_{i_1 i_2 \cdots i_j} = r_{i_1} r_{i_2} \cdots r_{i_j}$. The Weyl group of $\mathfrak{k}$ is $W_K = \{1, r_1, r_2, r_{12}, r_{21}, r_{121} = r_{212}\}$. The representation E_μ with highest weight $\mu = 3\alpha_1 + 4\alpha_2$ is shown in Figure 1-2.

The "convex polygon" rule tells us exactly which weight multiplicities are nonzero. More importantly, it is an extremely simple way of getting the most basic picture of the representation, avoiding any calculations at all. It is precisely this type of description that we shall seek in the case of reductive groups, replacing weights with irreducible representations of a maximal compact subgroup.

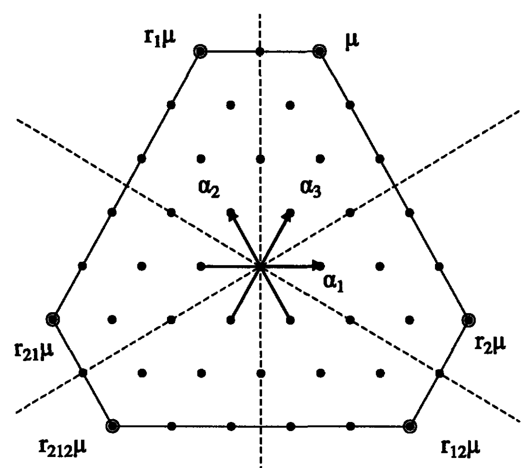

Figure 1-2: $E_{3\alpha_1+4\alpha_2}$ for SU(3)

1.2 Representations of reductive groups

A reductive group G is a slight generalization of a semisimple one. Its Lie algebra $\mathfrak{g}_0$ is a sum of a semisimple algebra and a commutative one. The precise conditions on the Lie group G are a little more technical and we shall come back to it in Chapter 2;

for the present, a real semisimple group is a good one to keep in mind. As usual, $\mathfrak{g}$ will denote the complexification $\mathfrak{g} = \mathfrak{g}_0 \otimes_{\mathbb{R}} \mathbb{C}$. Lastly, let K be a maximal compact subgroup of G, which is unique up to conjugation.

As a vector space, the Lie algebra $\mathfrak{g}$ splits into a direct sum $\mathfrak{g} = \mathfrak{k} \oplus \mathfrak{p}$, where $\mathfrak{p}$ is the "noncompact" part of $\mathfrak{g}$. Let $\mathfrak{h}$ be a Cartan subalgebra of $\mathfrak{k}$. The main subject of this paper are discrete series, which only exist if G and K have the same rank. We shall therefore assume that $\mathfrak{h}$ is also a Cartan subalgebra of $\mathfrak{g}$. The root system of $\mathfrak{g}$ splits into roots lying in $\mathfrak{k}$ and roots lying in $\mathfrak{p}$; they are referred to as *compact*, resp. *noncompact*. Given a choice of positive roots in $\mathfrak{g}$, we denote by ρ_c and ρ_n the half-sums of positive compact, resp. noncompact roots. The half-sum of all roots is then $\rho = \rho_c + \rho_n$.

Example 1. Perhaps the most familiar example of a reductive group is $SL(2,\mathbb{R})$. We shall instead work with $SU(1,1)$, which is conjugate to $SL(2,\mathbb{R})$ inside $SL(2,\mathbb{C})$. Its complexified Lie algebra is $\mathfrak{su}(1,1) \simeq \mathfrak{sl}(2)$, with the standard basis $\{h, e, f\}$. Choosing $SU(1,1)$ instead of $SL(2,\mathbb{R})$ means that the Lie algebra of $K = H$ is $S(U(1)\times U(1)) \simeq SO(2)$, spanned by the diagonal element h, which makes restrictions to H easy to study. The single root β of $\mathfrak{su}(1,1)$ is noncompact; the subalgebra $\mathfrak{p}$ is spanned by the basis elements e and f. We will denote by Λ the fundamental weight corresponding to β; i.e. $\Lambda = \frac{1}{2}\beta$.

Example 2. Let $G = SU(2,1)$, with $K = S(U(2) \times U(1))$. The root system of $\mathfrak{g} = \mathfrak{su}(2,1)$ is of type A_2. There are two compact roots $\pm\alpha$, and four noncompact roots $\{\pm\beta, \pm\beta_2\}$, where $\beta_2 = \beta + \alpha$. We will label by Λ_α, resp. Λ_β the fundamental weights corresponding to α, resp. β; that is, $\Lambda_\alpha = \frac{1}{3}(2\alpha + \beta)$ and $\Lambda_\beta = \frac{1}{3}(\alpha + 2\beta)$.

Let $\tilde{X}$ be an irreducible unitary representation of G. To get the most basic picture of $\tilde{X}$, we need to understand its restriction to K. On a general vector $v \in \tilde{X}$, K acts in rather complicated ways, and the space $U(\mathfrak{k})\cdot v$ can have infinite dimension. Let X be the subspace of $\tilde{X}$ of all vectors where $U(\mathfrak{k})\cdot v$ is finite dimensional; then X is a sum of irreducible K-representations. By a result of Harish-Chandra, stated as Lemma 31 in [2], X is dense in $\tilde{X}$, and that every K-type occurs in X with finite multiplicity.

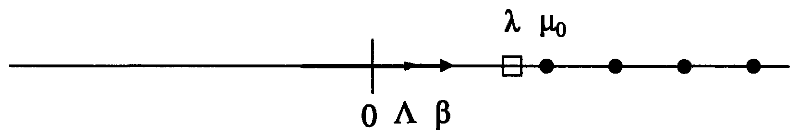

Figure 1-3: $\mathbf{X_{4\Lambda}}$ for $\mathbf{SU(1,1)}$

In such a case, the action of $\mathfrak{g}$ on X preserves the "finite under K" condition, so X is also a module for $\mathfrak{g}$. It contains all the information needed to recover the original representation $\tilde{X}$, and it is a much nicer module for K. We shall, therefore, restrict ourselves to studying X.

Example 1. Figure 1-3 shows a discrete series representation $X_{4\Lambda}$ of $SU(1,1)$. The label $\lambda = 4\Lambda$ is the *Harish-Chandra parameter* which indexes the discrete series; we shall talk more about it later. We have $K = H$ so K-type simply means a weight of H. The weights of this particular discrete series are $\{5\Lambda, 7\Lambda, 9\Lambda, \ldots\}$. There are infinitely many, each occurring with multiplicity one. Note that there is no notion of "highest K-type"; there is, however, a distinguished "lowest K-type", with weight 5Λ.

Example 2. A discrete series representation $X_{3\Lambda_\alpha+2\Lambda_\beta}$ for $G = SU(2,1)$ is shown on Figure 1-4. Each point on the figure represents the highest weight of a K-representation. For instance, the K-type $\mu_0 = 2\Lambda_\alpha + 4\Lambda_\beta$ corresponds to the three-dimensional representation whose weights are $\{\mu_0, \mu_0 - \alpha, \mu_0 - 2\alpha\}$. Again note that there is no notion of highest K-type. The K-type μ_0 is considered the "lowest", in that we get all the other K-types by adding positive roots, which in this case are $\{\alpha, \beta, \beta_2\}$.

Although some results in this paper are true in greater generality, we shall restrict ourselves to studying the discrete series representations. These are the irreducible subspaces of the regular representation $L^2(G)$. They are indexed by the aforementioned Harish-Chandra parameter λ, which is a regular element in the shifted lattice

$\hat{H}+\rho$, and is closely related to the lowest K-type. Any particular λ determines a set of positive roots in $\mathfrak{g}$:

$$\Delta_{+}=\{\alpha\in\mathfrak{g}\mid\langle\check{\alpha},\lambda\rangle>0\} \tag{1.2}$$

We shall label the positive noncompact positive roots by $\beta_1,\ldots,\beta_q$.

Let $X=X_\lambda$ be a discrete series with the Harish-Chandra parameter λ as above. The character of X is an invariant eigendistribution Θ on G. If we restrict Θ to K, we get a well defined distribution equal to a sum of irreducible characters of K with coefficients in $\mathbb{Z}$. The characters, restricted to H, are given by the appropriate Weyl character formula. The sum of these characters makes no sense over H, but the notation is a lot simpler if we ignore the above difficulties and express "$\Theta|_H$" as the formal sum

$$\Theta|_H=\sum_{n_1,\ldots,n_q\geq 0}\ \sum_{y\in W_K}\frac{(-1)^{l(y)}e^{y(\lambda+\rho_n+\sum n_j\beta_j)}}{\Delta_K} \tag{1.3}$$

The Blattner conjecture, proved by Hecht and Schmid in [3], says that the multiplicity of a particular K-type in X is exactly as predicted by Equation (1.3). Let us see what we can say about X just by looking at this formula.

The inner sum looks very much like the Weyl character formula for K. In fact, let us fix $n_1=n_2=\cdots=n_q=0$. Since λ is dominant regular for $\mathfrak{g}$, the weight $\lambda+\rho_n=\lambda+\rho-\rho_c$ is dominant regular for $\mathfrak{k}$: indeed, for every simple root $\alpha\in\Delta_{+}^{K}$ we have

$$\langle\check{\alpha},\lambda+\rho-\rho_c\rangle\geq 1+1-1=1 \tag{1.4}$$

Let $\nu_0=\lambda+\rho_n$. Since ν_0 is dominant regular for K, the inner summand in 1.3 is a

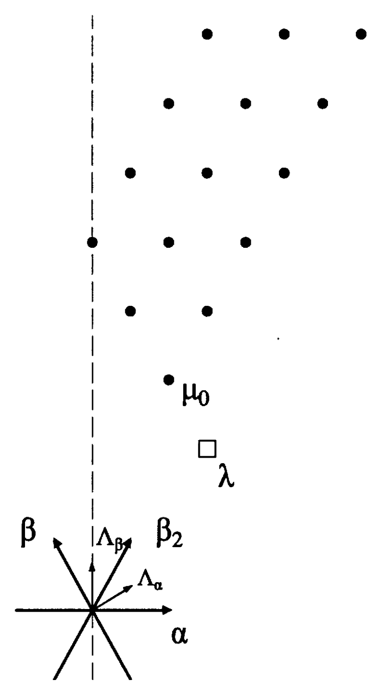

Figure 1-4: $\mathbf{X_{3\Lambda_\alpha+2\Lambda_\beta}}$ for $\mathbf{SU(2,1)}$

character for an irreducible K-representation with highest weight $\mu_0 = \nu_0 - \rho_c = \lambda + \rho - 2\rho_c$. This follows easily from the Weyl character formula (1.1).

Let us now choose n_1 through n_q not all zero, and let $\nu = \lambda + \rho_n + \sum n_j\beta_j$. If ν is dominant regular for $\mathfrak{k}$, then the inner summand is again the Weyl character formula for the irreducible representation with highest weight $\mu = \nu - \rho_c$.

Example 1. Let $G = SU(1,1)$ and $\lambda = 3\Lambda$. The Lie algebra $\mathfrak{g}$ has no compact roots and a single noncompact positive root β, so $\rho_c = 0$ and $\rho_n = \Lambda$. Recall that $K = H$, so K-types are weights. The lowest one is $\mu_0 = \lambda + \rho_n - \rho_c = 4\Lambda$.

Let us fix a value for $n_1 = n$; we get $\nu(n) = \lambda + n\beta = 2(n+2)\Lambda$. There is no question of dominance since K has no roots; the inner "sum" over the one trivial element in W_K gives us back the K-type $2(n+2)\Lambda$. What we get is exactly the discrete series shown in Figure 1-3.

Unfortunately, the weight $\nu = \lambda + \rho_n + \sum n_j\beta_j$ is not guaranteed to be dominant for $\mathfrak{k}$. Suppose that it is not. The inner sum runs over all elements of the Weyl group W_K, just like the Weyl character formula does. If ν is singular for $\mathfrak{k}$, then we get zero. If it is regular, then we get almost the character of the representation of K corresponding to the W_K-orbit of ν: the K-type with highest weight $\mu = w\nu - \rho_c$ where $w \in W_K$ is the element making $w\nu$ dominant. The only catch is the sign $(-1)^{l(w)}$: in the Weyl character formula, the summand corresponding to the highest weight occurs with a plus. If our w is odd length, then the highest-weight summand will occur with a minus. In other words, we will get $-E_\mu$ instead of E_μ.

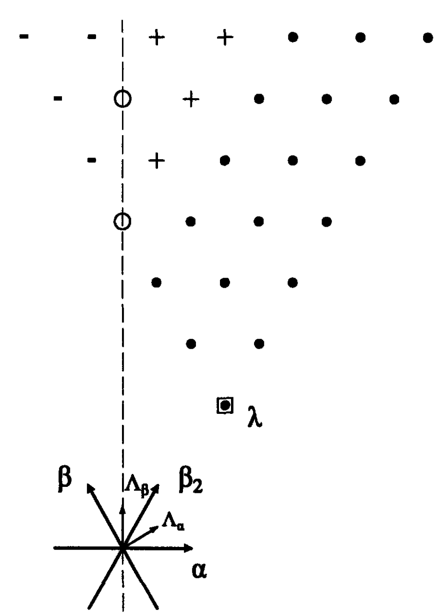

Figure 1-5: The cone for SU(2,1)

Example 2. Let $G = SU(2,1)$ and $\lambda = 3\Lambda_\alpha + 2\Lambda_\beta$. The half-sums are $\rho_n = \frac{1}{2}(\beta + \beta_2) = \frac{3}{2}\Lambda_\beta$ and $\rho_c = \frac{1}{2}\alpha = \Lambda_\alpha - \frac{1}{2}\Lambda_\beta$. The lowest K-type is therefore $\mu_0 = \lambda + \rho_n - \rho_c = 2\Lambda_\alpha + 4\Lambda_\beta$, in the discrete series which is shown in Figure 1-4.

For some choices of n_1 and n_2, the weight $\nu(n_1, n_2) = \lambda + \rho_n + n_1\beta + n_2\beta_2$ is dominant; for example if $n_1 = 0$ and we are simply adding multiples of β_2, which has positive product with α. However, if we set $n_2 = 0$ and start adding multiples of β, we remain dominant for $n_1 \leq 2$. The weight $\nu(3,0) = \lambda + \rho_n + 3\beta$ is singular and gives us no representation at all. Adding one more β gives us $\nu(4,0) = \lambda + \rho_n + 4\beta$, which lies in the negative Weyl chamber for K. It corresponds to a representation of highest weight $\mu = r\nu - \rho_c = \mu_0 + \alpha + 4\beta$, occurring with a minus sign.

We get the same K-type by choosing $n_1 = 3, n_2 = 1$, this time with a plus sign, as the weight $\nu(3,1) = \lambda + \rho_n + 3\beta + \beta_2$ is dominant. These two occurrences of E_μ are the only ones and exactly cancel each other out, leaving no copy of E_μ in X.

This cancellation happens for many of the weights. All of the possible ν lie in the cone starting at $\lambda + \rho_n$ going outwards in the directions β and β_2. These two roots are a basis for the two-dimensional space, so each ν comes from a unique choice of n_1 and n_2. Looking at Figure 1-5, we see that the left side of the cone runs off to the negative Weyl chamber. Each ν in that part of the cone, labeled by $-$ in the picture, corresponds to a representation of highest weight $\mu = r\nu - \rho_c$ occurring in the formula with a minus sign. This exactly cancels the $+1$ occurrence of E_μ, coming from $\nu' = \mu + \rho_c$ which is dominant, and labeled by $+$. The weights ν which give K-types in the discrete series all lie in a stripe going off to infinity in the direction of β_2, as shown in Figure 1-4.

The above example illustrates the procedure for finding K-types of X in general. To find out the multiplicity of E_μ in X, we have to look at $\nu_w = w(\mu + \rho_c)$ for all $w \in W_K$. There are finitely many ways to write ν_w as $\lambda + \rho_c + \sum n_j\beta_j$. Every such expression for ν_w is a contribution of $(-1)^{l(w)}$ to the multiplicity of μ in X. Summing up all the plus and minus signs tells us how many times E_μ occurs in X.

This is a rather tedious process, and it is not at all obvious which multiplicities will turn out to be nonzero. We shall try to answer that question here. The example

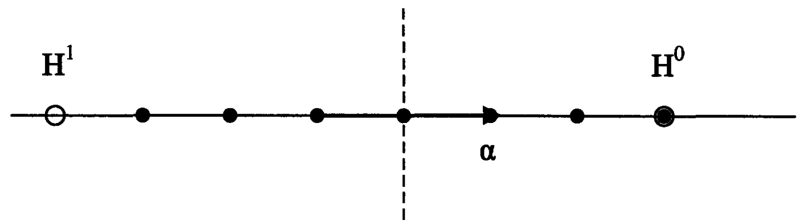

Figure 1-6: $\mathbf{H}^\bullet(\mathfrak{k}_+, \mathbf{E}_{3\alpha})$ for SU(2)

pictures are good indicators of what happens: the K-types of X seem to lie in a convex polytope, intersected with the lattice of noncompact roots shifted by ρ. The rest of this paper is dedicated to describing this polytope.

1.3 Analogy continued

Let us go back to the compact group K and one of its highest weight representations E_μ. Recall from Section 1.1 that the weights of E_μ lie in a convex set with vertices $\{w\mu\}_{w \in W_K}$. The vertices show up in the cohomology of $\mathfrak{k}_+$ in E_μ, in a way which motivates what happens in the case of reductive groups.

One way to draw this convex polygon is to find all the vertices and take their convex hull. We can also use a local description: each vertex comes with a cone where all the weights of E_μ live. This cone is given by half of the roots in $\mathfrak{k}$, which span a a nilpotent subalgebra $\bar{\mathfrak{n}}$. For reasons which will become clear later, it is more convenient to attach to each vertex the Borel subalgebra $\mathfrak{b} = \mathfrak{h} + \mathfrak{n}$ where $\mathfrak{n}$ is the opposite of $\bar{\mathfrak{n}}$, pointing in the "outward" directions where no weights of E_μ live.

Example 1. Let $K = SU(2)$, and pick a highest weight representation $X = E_\mu$. The vertices are at $\pm\mu$, each with all the weights of X on one side of it, and none on the other. The vertex at μ comes with the Borel subalgebra $\mathfrak{b} = \mathfrak{h} + \langle e \rangle$; we shall call it a $\mathfrak{b}$-*vertex*. The vertex at $-\mu$ is similarly associated with $r\mathfrak{b} = \mathfrak{h} + \langle f \rangle$.

There is nothing terribly revealing about such a description. In case of μ it simply says that there is no weight $\mu + \alpha$, which rephrases the fact that μ is the highest weight of E_μ. Let us restate it again: of all weight spaces in X, the highest one $F(\mu)$

has the special property that it is killed by the action of e. We write $F(\mu) = X^e$.

We could similarly say that $F(-\mu) = X^f$, but let us write all descriptions using e. All weight spaces in X are in the image of e, except $F(-\mu)$; i.e. $F(-\mu) = X/eX$.

Consider now the functor $(\ )^e$: it takes a $\mathfrak{k}$-module and returns the space of its highest weight vectors. It is a covariant right-exact functor, and its derived functors are referred to as the $\mathfrak{k}_+$-cohomology. The $\mathfrak{b}$-vertex is $F(\mu) = H^0(\mathfrak{k}_+, X)$. The $r\mathfrak{b}$ vertex is $F(r\mu) = F(-\mu)$, whose weight is very close to the degree one cohomology $H^1(\mathfrak{k}_+, X) = F(-\mu - \alpha)$. The cohomology for $E_{3\alpha}$ is shown in Figure 1-6.

The general result is given by Kostant's Theorem, found in Chapter IV of [8]:

Theorem 1.3.1. *Let $\mathfrak{k}$ be a compact Lie algebra, with $\mathfrak{k}_+$ a choice of positive roots, and E_μ the irreducible representation of $\mathfrak{k}$ with highest weight μ. For every element $w \in W_K$ we define the weight $\nu_w = w(\mu+\rho_c)-\rho_c$. The cohomology of $\mathfrak{k}_+$ in E_μ is*

$$H^i(\mathfrak{k}_+, E_\mu) = \bigoplus_{w\in W_K,\, l(w)=i} F(\nu_w) \tag{1.5}$$

Example 2. The Weyl group of $SU(2)$ has the trivial element of length zero, and the reflection r of length one; the cohomology is easily computed to be as stated in Example 1.

Let us consider $K = SU(3)$, and $X = E_\mu$. The half-sum of positive roots is $\rho_c = \alpha_1 + \alpha_2$. The Weyl group has the trivial element of length zero, $\{r_1, r_2\}$ of length one, $\{r_{12}, r_{21}\}$ of length two, and $r_{121} = r_{212}$ of length three (see Section 1.1 for notation). Let us denote by $H^w(\mathfrak{k}_+, E_\mu)$ the summand of $H^{l(w)}(\mathfrak{k}_+, E_\mu)$ corresponding to w. The weight of $H^w(\mathfrak{k}_+, E_\mu)$ is approximately $w^{-1}\mu$, except shifted by $(w^{-1}\rho_c - \rho_c)$. We can easily compute the cohomology:

$$\begin{array}{rclrcl}
H^1(\mathfrak{k}_+, E_\mu) & = & F(\mu) & H^{r_{12}}(\mathfrak{k}_+, E_\mu) & = & F(r_{21}\mu - \alpha_1 - 2\alpha_{@}) \\
H^{r_1}(\mathfrak{k}_+, E_\mu) & = & F(r_1\mu - \alpha_1) & H^{r_{21}}(\mathfrak{k}_+, E_\mu) & = & F(r_{12}\mu - 2\alpha_1 - \alpha_2) \\
H^{r_2}(\mathfrak{k}_+, E_\mu) & = & F(r_2\mu - \alpha_2) & H^{r_{121}}(\mathfrak{k}_+, E_\mu) & = & F(r_{121}\mu - 2\alpha_1 - 2\alpha_2)
\end{array}$$

Figure 1-7 shows these for the representation with highest weight $\mu = 3\alpha_1 + 4\alpha_2$.

Note again that for all $w \in W_K$, the vertex $w\mu$ is in fact a $w\mathfrak{k}_+$-vertex of X, and it lies close to $H^w(\mathfrak{k}_+, E_\mu)$,

We have seen that the cohomology of $\mathfrak{k}_+$ in E_μ gives us the vertices of E_μ and their orientation. We will take advantage of yet another useful attribute of Kostant's Theorem, namely the fact that it is "reversible", allowing us to recover the representation of K from any weight in the cohomology. Let X be a completely reducible representation of $\mathfrak{k}$, and suppose that a weight ν appears in $H^i(\mathfrak{k}_+, E_\mu)$. Let $w \in W_K$ be the element making $w(\nu + \rho_c)$ dominant. Then the representation E_μ occurs in X, where $\mu = w(\nu+\rho_c)-\rho_c$.

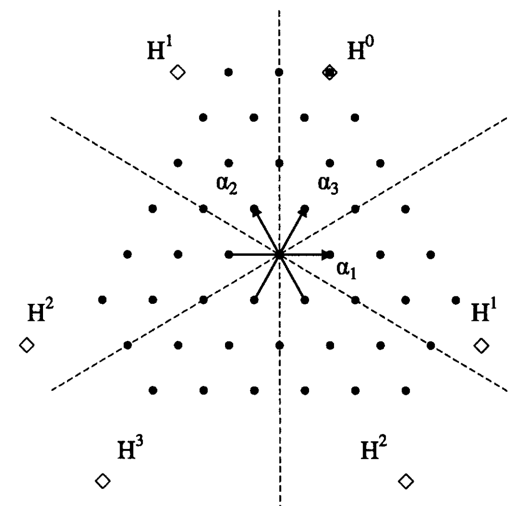

Figure 1-7: $\mathbf{H}^\bullet(\mathfrak{k}_+, \mathbf{E}_{3\alpha_1+4\alpha_2})$ for SU(3)

For example, if $i = 0$ then this simply says that if ν is a highest weight in X, then X contains the representation E_ν. Kostant's Theorem hands us a generalization of this to higher cohomology, which will come in very handy when we start working with reductive groups.

1.4 Edges

The K-types of a given discrete series X seem to be very close to a set of lattice points in a convex polyhedron. The polyhedron is noncompact, and as such it is not determined by its set of vertices. For example, the K-types polyhedron in Figure 1-4 has only two vertices, which do not specify the direction of the "infinite" boundaries of dimension one. This suggests that we may want to look for the lines themselves; or, in the general case, for facets of maximal dimension. We shall refer to these maximal facets as "edges". Let D be the dimension of the convex polyhedron. This number depends entirely on the group G; it satisfies $0 \leq D \leq \mathrm{rk}(G)$. If G is compact then $D = 0$; if G has no compact factors then $D = \mathrm{rk}(G)$. In any case, the facets we shall look for will be of dimension $D - 1$.

An edge is therefore given by a hyperplane in $\mathfrak{h}^*$ placed at some weight, together with a preferred side where no K-types lie. A hyperplane corresponds to a linear functional ϕ on $\mathfrak{h}^*$, which is zero along the hyperplane directions, and negative on the side with no K-types. The functional ψ determines a parabolic subalgebra $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ of $\mathfrak{g}$:

$$\begin{aligned} \Delta(\mathfrak{l}) &= \{ \delta \in \mathfrak{g} \mid \phi(\delta) = 0 \} \\ \Delta(\mathfrak{u}) &= \{ \delta \in \mathfrak{g} \mid \phi(\delta) < 0 \} \end{aligned} \tag{1.6}$$

In other words, the reductive part $\mathfrak{l}$ runs along the edge directions, and $\mathfrak{u}$ points in the "outward" directions of no K-types. In Section 1.4, we defined D as the dimension of the minimal convex polyhedron containing the K-types of X. The edge will be a maximal one precisely when $\mathrm{rk}(\mathfrak{l}) = D - 1$.

Example 1. Let us first see what this description looks like in the case where vertices are the maximal facets. Let $G = SU(1,1)$ and X be the discrete series with Harish-Chandra parameter $\lambda = 3\lambda$, shown in Figure 1-3. There is exactly one edge, zero-dimensional, sitting at the K-type $\mu_0 = 4\lambda$. There are no "edge" directions, so $\mathfrak{l}$ is the Cartan subalgebra $\mathfrak{h} = \mathfrak{k}$. The "outward" direction is along $-\alpha$, which corresponds to the element $f \in \mathfrak{g}$. The edge is therefore given by the Borel subalgebra $\mathfrak{h} + \langle f \rangle$.

Example 2. Let us now consider our rank-two example where $G = SU(2,1)$, and X is the representation with Harish-Chandra parameter $\lambda = 3\Lambda_\alpha + 2\Lambda_\beta$, shown in Figure 1-4. There are two edges through μ_0: The first runs along the root β_2, which means that $\mathfrak{l} = \mathfrak{s}(\mathfrak{u}(1,1) \oplus \mathfrak{u}(1))$, spanned by $\mathfrak{h}$, e_{β_2}, and f_{β_2}. The outward directions are α and $-\beta$, so $\mathfrak{u}$ is spanned by e_α and f_β. The other edge runs along β, with outward directions $-\alpha$ and $-\beta_2$, giving the corresponding parabolic. The third edge of this discrete series runs through $\mu_0 + 2\beta$ along β_2; $\Delta(\mathfrak{u}) = \{-\alpha, \beta\}$.

Following the analogy of Section 1.1, we shall look for these edges using Lie algebra cohomology, only this time with respect to a nilpotent $\mathfrak{u}$ that is not necessarily maximal. We have the functor $(\ )^{\mathfrak{u}}$, which takes a $\mathfrak{g}$ module X and returns all vectors killed by $\mathfrak{u}$. The cohomology $H^i(\mathfrak{u}, X)$ is the i-th derived functor of $(\ )^{\mathfrak{u}}$. The reductive part $\mathfrak{l}$ of $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ preserves $X^{\mathfrak{u}} = H^0(\mathfrak{u}, X)$, and acts on the cohomology in higher degrees as well. We shall therefore think of each $H^i(\mathfrak{u}, X)$ as a representation of $\mathfrak{l}$. Similarly we can define the cohomology of X with respect to $\mathfrak{u} \cap \mathfrak{k}$: $H^i(\mathfrak{u} \cap \mathfrak{k}, X)$ will be a representation of $\mathfrak{l} \cap \mathfrak{k}$. In computing the cohomology, we shall always take a parabolic with all compact roots positive.

In Section 2.2 we will describe a restriction map

$$\rho : H^i(\mathfrak{u}, X) \to H^i(\mathfrak{u} \cap \mathfrak{k}, X) \tag{1.7}$$

Suppose that we find an $\mathfrak{l}$-representation $Z \subset H^i(\mathfrak{u}, X)$ which maps to an $\mathfrak{l} \cap \mathfrak{k}$-representation $\bar{Z}$ under this map. Let ν be the highest weight of $\bar{Z}$, and let $w \in W_K$ be the element making $w(\nu + \rho_c)$ dominant. A generalization of Kostant's Theorem to $\mathfrak{u} \cap \mathfrak{k}$ instead of $\mathfrak{k}_+$ will tell us that the representation E_μ occurs in X, where $\mu = w(\nu + \rho_c) - \rho_c$. This result comes from the occurrence of $\bar{Z}$ in $H^i(\mathfrak{u} \cap \mathfrak{k}, X)$, with no need to look at $H^i(\mathfrak{u}, X)$ at all. The fact that $\bar{Z}$ is in the image of ρ will tell us that not only does the K-type μ occur in X, but it lies on a boundary given by $w\mathfrak{q}$.

Example 1. Let $K = SL(2)$, X the discrete series with Harish Chandra parameter 3λ, and $\mathfrak{q} = \mathfrak{l} + \mathfrak{u} = \mathfrak{h} + \langle f \rangle$. In this case $\mathfrak{u} \cap \mathfrak{k}$ is zero, so *every* weight space of X lies in $H^0(\mathfrak{u} \cap \mathfrak{k}, X)$. Only the space $\mathbb{C}_{4\lambda}$ lies in $H^0(\mathfrak{u}, X)$, and this is indeed the K-type on the $\mathfrak{q}$-boundary of X.

The first step towards a result in this direction was taken by Mark Sepanski, who proved the following result in [12]: suppose that a K-type μ lies on a $\mathfrak{q}$-edge, for any parabolic $\mathfrak{q}$. Pick the shortest $w \in W_K$ such that the roots of $w(\mathfrak{u} \cap \mathfrak{k})$ are positive. The $\mathfrak{l} \cap \mathfrak{k}$-representation $\bar{Z}$ with highest weight $w^{-1}(\mu + \rho_c) - \rho_c$ occurs in $H^{l(w)}(w(\mathfrak{u} \cap \mathfrak{k}), X)$ by the generalized Kostant theorem. Sepanski proved that under these assumptions, $\bar{Z}$ is in the image of the restriction map in cohomology

$H^{l(w)}(w(\mathfrak{u}), X) \to H^{l(w)}(w(\mathfrak{u} \cap \mathfrak{k}), X)$.

In [13], Sepanski showed that for $G = SU(1, n)$ this is a one-to-one correspondence: every element in the image of the restriction map comes from a K-type on a boundary of X. The main result of this paper is the generalization of this statement to all simply-laced Lie groups:

Theorem 1.4.1. *Let G be a reductive Lie group with simply-laced Lie algebra $\mathfrak{g}$, and X a discrete series for G. Let $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ be a parabolic subalgebra of $\mathfrak{g}$, with $\mathfrak{u} \cap \mathfrak{k}$ positive. Let $\bar{Z}$ be an irreducible $L \cap K$-representation of highest weight ν that occurs in the image of an L-representation Z in the restriction map*

$$H^N(\mathfrak{u}, X) \xrightarrow{\rho} H^N(\mathfrak{u} \cap \mathfrak{k}, X) \tag{1.8}$$

in some degree N. Finally, let w be the Weyl group element which makes $w(\nu + \rho_c)$ dominant. Then there is a K-type $w(\nu+\rho_c)-\rho_c$ occurring in X, and it lies on an edge given by $w\mathfrak{q}$.

Chapter 2

Definitions and tools

In this chapter we shall describe the search process for edges. We will start by setting up all the definitions in Section 2.1. Section 2.2 contains the Hochschild-Serre spectral sequence, whose E_2 edge homomorphism is the restriction map ρ. In Section 2.3 we describe the cohomology of discrete series, as given by a theorem of Vogan. We shall finish in Section 2.4 by giving the complete definition of an edge and stating precisely the result of Mark Sepanski which says that every K-type on an edge gives a term in the image of ρ.

2.1 Definitions

2.1.1 Real reductive groups

Let $\mathfrak{g}$ be a real or complex Lie algebra. We say that $\mathfrak{g}$ is a *reductive Lie algebra* if it is fully reducible under $\mathrm{ad}_{\mathfrak{g}}$; that is, it is a direct sum of a semi-simple and a commutative Lie algebra. The commutative subalgebra is equal to the center of $\mathfrak{g}$, so we can write $\mathfrak{g} = [\mathfrak{g}, \mathfrak{g}] \oplus \mathrm{Z}(\mathfrak{g})$.

A complex algebraic Lie group is called *reductive* if its Lie algebra is reductive. A *real reductive Lie group* G is the real points of a complex reductive Lie group, or a finite cover thereof. A slightly more general definition can be found in Chapter IV of [8]. Here we shall list some relevant properties and further assumptions:

1. The real Lie algebra $\mathfrak{g}_0$ of G is reductive.

2. G contains a maximal compact subgroup K, which is unique up to conjugation.

3. There is a Lie algebra involution θ on $\mathfrak{g}_0$ whose $+1$ eigenspace is $\mathfrak{k}_0 = \mathrm{Lie}(K)$. Let $\mathfrak{p}_0$ be the -1 eigenspace. Then $\mathfrak{g}_0$ decomposes under θ into the direct sum $\mathfrak{g}_0 = \mathfrak{k}_0 \oplus \mathfrak{g}_0$.

4. We shall assume that G is connected.

5. We shall assume that $\mathrm{rk}(K) = \mathrm{rk}(G)$. Let H be a Cartan subgroup of K; it follows that H is also a Cartan subgroup of G. We denote its Lie algebra by $\mathfrak{h}_0$.

2.1.2 $(\mathfrak{g}, K)$-modules

The easiest kind of *reductive pair* $(\mathfrak{g}, K)$ is a Lie algebra $\mathfrak{g}$ together with a compact group K, coming from a reductive Lie group G as above. A more general definition can be found in [8], with $\mathfrak{g}$ any reductive Lie algebra and K a compact group with $\mathrm{Lie}(K) \subset \mathfrak{g}$. They are required to satisfy compatibility conditions which are automatic if the reductive group G is around.

In Chapter 1, we replaced a G-module $\tilde{X}$ with the maximal "nice" submodule with respect to K. The resulting module X was no longer a module for all of G, but had a well defined action of K and $\mathfrak{g}$, and behaved well under K. We shall replace $\tilde{X}$ by X from now on, without keeping G around at all. That is, we shall work in the category of $(\mathfrak{g}, K)$-modules:

Definition 2.1.1. A $(\mathfrak{g}, K)$-module X is a complex vector space with an action of $\mathfrak{g}$ and an action of K such that

1. The K representation is locally K-finite; that is, every vector $v \in X$ lies in a finite dimensional K-subspace.

2. The differentiated version of the K action is the restriction to $\mathfrak{k}$ of the $\mathfrak{g}$ action

The first condition says that X is a direct sum of K-types. The other condition ensures that the actions of $\mathfrak{g}$ and K are compatible, the way they would be if X came from an actual G-module.

2.1.3 The Hecke algebra

The *Hecke algebra* $R(\mathfrak{g}, K)$ is designed to take place of $U(\mathfrak{g})$, incorporating the action of K on $(\mathfrak{g}, K)$-modules. If $K = 1$ then $R(\mathfrak{g}, K)$ is in fact equal to $U(\mathfrak{g})$. If $\mathfrak{g} = \mathfrak{k}$, then $R(\mathfrak{k}, K) = R(K)$ is the K-finite distributions on K, with convolution as multiplication. Let $C(K)$ be the algebra of K-finite smooth functions on K $C(K) = \oplus_{\mu\in\hat{K}} E_\mu \otimes E_\mu^*$. If a metric on K is fixed, there is an isomorphism from $R(K)$ to the K-finite dual of $C(K)$.

More generally, if the pair $(\mathfrak{g}, K)$ comes from a real reductive group G, then $R(\mathfrak{g}, K)$ is the algebra of left K-finite distributions on G with support in K. There is an isomorphism

$$R(\mathfrak{g}, K) = U(\mathfrak{g}) \otimes_{U(\mathfrak{k})} R(K) \tag{2.1}$$

Here $\mathfrak{k}$ acts on $U(\mathfrak{g})$ by right multiplication, and on $R(K)$ by the left regular representation. Equation (2.1) is also the definition of $R(\mathfrak{g}, K)$ in the most general case.

Hecke algebras do not have a unit, as the space $R(K)$ has no such thing. Instead, $R(K)$ has an *approximate identity*, which is a sequence of elements $\{s_1, s_2, \ldots\}$ with the following property: given an element $s \in R(K)$, there is an n such that $s_i s = s$ for all $i \geq n$.

A $R(\mathfrak{g}, K)$-module is called *approximately unital* if for every vector $v \in X$, there is an index $n(v)$ such that $s_i x = x$ for all $i \geq n(v)$. Every $(\mathfrak{g}, K)$-module X carries an action of $R(\mathfrak{g}, K)$, which satisfies this condition. The category of $(\mathfrak{g}, K)$-modules, defined in Section 2.1.2, is equivalent to the category of approximately unital $R(\mathfrak{g}, K)$-modules.

2.1.4 Cohomological functors

Suppose we have a subgroup $L \subset G$ and a subpair $(\mathfrak{l}, L \cap K)$ of $(\mathfrak{g}, K)$, where $\mathfrak{l} = \mathrm{Lie}(L) \otimes_{\mathbb{R}} \mathbb{C}$. Given a $(\mathfrak{l}, L \cap K)$-module V, we shall define two induction functors that make V into a $(\mathfrak{g}, K)$-module:

$$P^{\mathfrak{g},K}_{\mathfrak{l},L\cap K}(V) \;=\; R(\mathfrak{g},K)\otimes_{R(\mathfrak{l},L\cap K)} V \tag{2.2}$$

$$I^{\mathfrak{g},K}_{\mathfrak{l},L\cap K}(V) \;=\; \mathrm{Hom}_{R(\mathfrak{l},L\cap K)}(R(\mathfrak{g},K),V)_{K-\mathrm{fin}} \tag{2.3}$$

These are similar in spirit to the usual induction functors for universal enveloping algebras. The only unusual aspect of the definitions is due to the fact that $\mathrm{Hom}_{R(\mathfrak{l},L\cap K)}(R(\mathfrak{g},K),V)$ is not K-finite. In order to get a $(\mathfrak{g}, K)$-module, we need to take the space of K-finite vectors inside it.

Another difference is that these functors are not exact. The functor P is covariant and right exact, and sends projectives to projectives. The functor I is covariant and left exact, and sends injectives to injectives. Their derived functors are denoted by P_j and I^j, respectively. We shall see more of these later.

2.2 Hochschild-Serre spectral sequence

2.2.1 The cohomology $H^w(\mathfrak{k}_+, X)$

Let $w \in W_K$. In Chapter 1, we defined $H^w(\mathfrak{k}_+, E_\mu)$ to be the summand of $H^{l(w)}(\mathfrak{k}_+, E_\mu)$ whose weight was close to $w^{-1}(\mu)$. Let us make some more precise definitions here. We denote by C_w the set of weights

$$C_w = \{\nu \in \hat{H} \mid w(\nu + \rho_c) \text{ is dominant regular for } K\} \tag{2.4}$$

For any admissible G-representation X, $H^w(\mathfrak{k}_+, X)$ is the summand of $H^{l(w)}(\mathfrak{k}_+, X)$ with weights in C_w.

By Kostant's Theorem, the weights in $H^w(\mathfrak{k}_+, X)$ are in one-to-one correspondence to the highest weights of X. Given a highest weight vector $v \in X$, by abuse of notation we will say that $f \in H^w(\mathfrak{k}_+, X)$ "corresponds to v" if it is a nonzero element coming from the one-dimensional space $H^w(\mathfrak{k}_+, U(\mathfrak{k}) \cdot v) \subset H^w(\mathfrak{k}_+, X)$.

The following proposition will come in handy later:

Proposition 2.2.1. *Let $w \in W_K$ be an element of length i, and ν a weight in C_w. If β is one of the weights of $\mathfrak{p}$, then $\nu + \beta$ does not lie in $C_{w'}$ for any $w' \neq w$.*

Proof. Let $\Delta_+^K = \{\alpha_1, \ldots \alpha_m\}$, where the first i roots are those negative on $\nu + \rho_c$. That is,

$$\langle \check{\alpha}_j, \nu + \rho_c \rangle \leq -1 \quad \text{for } j \leq i, \quad \text{and} \tag{2.5}$$

$$\langle \check{\alpha}_j, \nu + \rho_c \rangle \geq 1 \quad \text{for } j > i \tag{2.6}$$

Since $\mathfrak{g}$ is a simply-laced Lie algebra, the product $\langle \check{\alpha}_j, \beta \rangle$ is at most 1 for any j. It follows that

$$\langle \check{\alpha}_j, \nu + \beta + \rho_c \rangle \leq 0 \quad \text{for } j \leq i, \quad \text{and} \tag{2.7}$$

$$\langle \check{\alpha}_j, \nu + \beta + \rho_c \rangle \geq 0 \quad \text{for } j > i \tag{2.8}$$

If equality holds for some α_j, then the weight $\nu + \beta + \rho_c$ is singular, and $\mu + \beta$ lies in no $C_{w'}$. Otherwise we must have that $\nu + \beta \in C_w$. □

2.2.2 Generalized Kostant's Theorem

Before we get to our K-types, we need more tools to work with. First is the generalized version of Kostant's Theorem, which was mentioned in Chapter 1.

For any $w \in W_K$, we define the set

$$\Delta^K(w) = \left\{\alpha \in \Delta_+^K |\, w^{-1}(\alpha) \in -\Delta_+^K\right\} \tag{2.9}$$

Given a θ-stable parabolic $\mathfrak{q}$ with $\mathfrak{u} \cap \mathfrak{k}$ positive, we let W_K^1 be the subset of W_K defined as

$$W_K^1 = \{w \in W_K \mid \Delta^K(w) \subset \Delta(\mathfrak{u} \cap \mathfrak{k})\} \tag{2.10}$$

The elements of W_K^1 are the minimal length representatives for the cosets of $W_{L \cap K}$ in W_K.

The following version of Kostant's Theorem is taken from Chapter IV in [8]:

Theorem 2.2.2. *With notation as above, let $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ be a θ-stable parabolic in $\mathfrak{k}$ with $\mathfrak{u} \subset \mathfrak{k}_+$, and let E_μ be the irreducible representation of K with highest weight μ. Let Y_ν denote the irreducible representation of L with highest weight ν. For $w \in W_K^1$, let $\nu_w = w(\mu + \rho_c) - \rho_c$. As a representation of L, the cohomology of $\mathfrak{u}$ in X is the direct sum*

$$H^i(\mathfrak{u}, X) = \bigoplus_{w \in W_K^1,\, l(w)=i} Y_{\nu_w} \tag{2.11}$$

This is very similar to the classical version, only now each summand is a representation of L, and the sum is only over those elements in W_K such that ν_w is actually a highest weight for L. Figure 2-1 shows the cohomology for $SU(3)$, with $\pm\alpha_1$ a root of $\mathfrak{l}$, α_2, α_3 the roots of $\mathfrak{u}$, and $\mu = 3\alpha_1 + 4\alpha_2$. The cohomology exists in degrees 0, 1 and 2; each is a single irreducible representation for $L = S(U(2) \times U(1)) \simeq SU(2) \oplus SO(2)$.

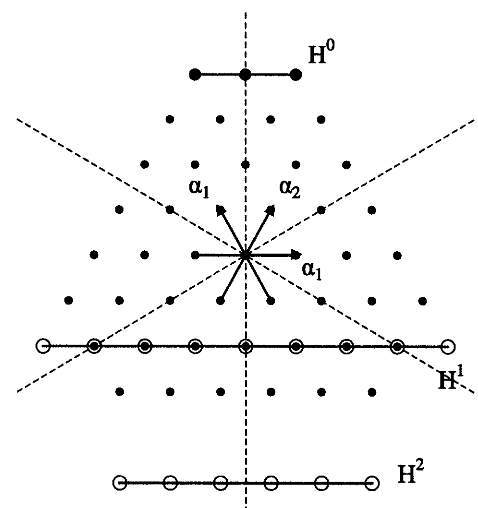

Figure 2-1: $\mathbf{H}^\bullet(\mathfrak{u}, \mathbf{X}_{3\alpha_1+4\alpha_2})$ for SU(3)

2.2.3 The spectral sequence

The next tool we need is the Hochschild-Serre spectral sequence, which, among other things, will hand to us the restriction map in cohomology. The sequence was introduced in [4].

Theorem 2.2.3. *Let $\mathfrak{q} \subset \mathfrak{g}$ be a θ-stable parabolic, $L \subset G$ a reductive subgroup, and X a $(\mathfrak{q}, L \cap K)$-module. There exists a spectral sequence converging to $H^{a+b}(\mathfrak{u}, X)$, with differential of bidegree $(r, 1-r)$ and E_1 term*

$$E_1^{a,b} = H^b(\mathfrak{u} \cap \mathfrak{k}, \bigwedge^a (\mathfrak{u}/\mathfrak{u} \cap \mathfrak{k})^* \otimes X) \tag{2.12}$$

The zeroth column in the sequence (2.12) is $E_1^{0,b} = H^b(\mathfrak{u} \cap \mathfrak{k}, X)$. The spectral sequence comes with a restriction map from the total cohomology $H^b(\mathfrak{u}, X)$ to this term:

$$H^b(\mathfrak{u}, X) \xrightarrow{\rho} H^b(\mathfrak{u} \cap \mathfrak{k}, X) \tag{2.13}$$

It is this restriction map which we shall use to find the K-types of our discrete series.

We shall need to understand what $L \cap K$-types can occur in the cohomology $H^i(\mathfrak{u}, X)$. To this end, let us analyze the E_1-terms in the spectral sequence (2.12). The easiest case is when $\mathfrak{u}\cap\mathfrak{k}$ acts on $\mathfrak{u}/\mathfrak{u}\cap\mathfrak{k}$ be zero; we shall indicate this by switching to the notation $\mathfrak{u}/\mathfrak{u}\cap\mathfrak{k} = \mathfrak{u}\cap\mathfrak{p}$. In this case we can take the term $\bigwedge^a(\mathfrak{u}/\mathfrak{u}\cap\mathfrak{k})^*$ out of the $\mathfrak{u}\cap\mathfrak{k}$ cohomology, leaving us with $E_1^{a,b} = H^b(\mathfrak{u}\cap\mathfrak{k}, X) \otimes \bigwedge^a(\mathfrak{u}\cap\mathfrak{p})^*$. This simplifies the situation enormously, for X is a sum of K-types whose $\mathfrak{u}\cap\mathfrak{k}$-cohomology we know by the generalized Kostant's Theorem 2.2.2; and of course we understand $\bigwedge^a(\mathfrak{u}/\mathfrak{u}\cap\mathfrak{k})^*$.

Unfortunately, in general $\mathfrak{u}\cap\mathfrak{k}$ does not act on $\bigwedge^a(\mathfrak{u}/\mathfrak{u}\cap\mathfrak{k})^*$ by zero. To deal with this case, we shall define a filtration on $\bigwedge(\mathfrak{u}/\mathfrak{u}\cap\mathfrak{k})^*$ whose associated graded *is* a trivial $\mathfrak{u}\cap\mathfrak{k}$ representation in every degree, and repeat the above analysis on the graded pieces.

First we select in $\bigwedge(\mathfrak{u}/\mathfrak{u}\cap\mathfrak{k})$ a sequence $\{V_p\}_{p=0}^N$ of $L\cap K$-submodules such that

1. $V_0 = \bigwedge^0(\mathfrak{u}/\mathfrak{u}\cap\mathfrak{k}) = \mathbb{C}$
2. $V_p \subset \bigwedge^{r(p)}(\mathfrak{u}/\mathfrak{u}\cap\mathfrak{k})$; i.e. the members of V_p are homogeneous of some degree $r(p) \leq p$; also $p' \leq p$ implies $r(p') \leq r(p)$.
3. $(\mathfrak{u}\cap\mathfrak{k}) \cdot V_p \subset \langle V_0, \ldots, V_{p-1}\rangle$

This gives us a filtration of the spectral sequence (2.12), with graded E_1 terms

$$E_1^{a,b} = \sum_{r(p)=a} H^b(\mathfrak{u}\cap\mathfrak{k}, (V_p/V_{p-1})^* \otimes X) = \sum_{r(p)=a} H^b(\mathfrak{u}\cap\mathfrak{k}, X) \otimes (V_p/V_{p-1})^* \tag{2.14}$$

This tells us that the $L\cap K$-types in $E_1^{a,b}$ are a subset of the $L\cap K$-types in Equation (2.15). In other words, they are estimated by

$$(E_1^{a,b})^{\text{est}} = H^b(\mathfrak{u}\cap\mathfrak{k}, X) \otimes \bigwedge^a(\mathfrak{u}\cap\mathfrak{p})^* \tag{2.15}$$

2.3 The cohomology of discrete series

To analyze the image of the restriction map (2.13), we need to describe the cohomology of discrete series. This is done in Section 6 of [14]:

Given a G-regular weight λ, we define the number

$$l_{\mathfrak{q}}(\lambda) = |\{\alpha \in \mathfrak{u} \cap \mathfrak{k} \mid \langle \check{\alpha}, \lambda \rangle < 0\}| + |\{\alpha \in \mathfrak{u} \cap \mathfrak{p} \mid \langle \check{\alpha}, \lambda \rangle > 0\}| \tag{2.16}$$

Theorem 2.3.1. *Let X be a discrete series of $\mathfrak{g}$ with Harish-Chandra parameter λ, and $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ be a θ-stable parabolic subalgebra of $\mathfrak{g}$. Let T be the set*

$$T = \{w \in W_K^1 \mid l_{\mathfrak{q}}(\lambda_w) = i\} \tag{2.17}$$

Given an element $w \in W_K^1$, we let X_{λ_w} be the discrete series of $\mathfrak{l}$ with the parameter $\lambda_w = w(\lambda) - \rho_{\mathfrak{u}}$. As a representation of $\mathfrak{l}$, the i-th cohomology of $\mathfrak{u}$ in X is a sum of discrete series:

$$H^i(\mathfrak{u}, X) = \bigoplus_{w \in T} X_{\lambda_w} \tag{2.18}$$

The theorem says that $H^N(\mathfrak{u}, X)$ is a sum of discrete series for L. This will allow us to find the K-types of X by induction on the rank of our group. Indeed, let Z be one of the discrete series in $H^N(\mathfrak{u}, X)$. Recall that we defined D as the dimension of the minimal convex polyhedron containing the K-types of X (see Section 1.4). We are looking for facets of our polytope, which $\mathrm{rk}(L) = D-1 < \mathrm{rk}(G)$. We can therefore claim to understand the $L \cap K$-types of Z by induction. To see which ones give us walls of the polytope, we need to find out how these $L \cap K$-types behave under the restriction map (2.13).

More specifically, given an $L \cap K$-representation $\bar{Z} \subset Z$ of highest weight ν, we wish to know whether the restriction $\rho(\bar{Z})$ is nonzero. The first obvious condition is that $\nu \in C_w$ for $w \in W_K$ of length N. This is almost sufficient; to get the precise condition, we need to examine the Hochschild-Serre spectral sequence (2.12).

Pick an E_1 representative of $\bar{Z}$. It lives on the diagonal $a + b = i$. The restriction map is nonzero on $\bar{Z}$ precisely when this representative lies in in $E_1^{i,0} = H^i(\mathfrak{u} \cap \mathfrak{k}, X)$.

This means that $\nu \in C_w$ for some w of length i.

Suppose that $\rho(\bar{Z}) = 0$, which means that the representative is in some column $a > 0$. Expression (2.15) shows that there are distinct roots $\delta_1, \ldots, \delta_{r(a)}$ in $\mathfrak{u} \cap \mathfrak{p}$, such that $\nu' = \nu + \sum_{j=1}^{r(a)} \delta_j$ appears as a weight in $H^{w'}(\mathfrak{u} \cap \mathfrak{k}, X)$ for some w' of length $i - r(a)$. We are interested in $L \cap K$-types for which this cannot happen.

Definition 2.3.2. Let $\mathfrak{q}$ be a θ-stable parabolic and ν a weight lying in C_w for some $w \in W_K$ of length i. We shall say that ν is *safe* with respect to $\mathfrak{q}$ if for all $m > 0$ and all choices $\delta_1, \ldots \delta_m$ of distinct roots in $\mathfrak{u} \cap \mathfrak{p}$, the weight $\nu + \sum_{j=1}^{m} \delta_j$ does not lie in $C_{w'}$ for any $w' \in W_K$ of length $i - m$.

Based on the preceding discussion, Definition 2.3.2 is designed to make sure that $\bar{Z}$ is in the image of ρ. Note that if $i = 0$ then the condition is trivially satisfied; i.e. all dominant regular weights are safe. Proposition 2.2.1 proves that the condition is also satisfied when $i = 1$. Degrees higher than $i = 1$ need to be checked. In any case the unsafe weights must lie very close to the boundary of C_w.

Proposition 2.3.3. *Let X be a discrete series, and $\mathfrak{q}$ a θ-stable parabolic subalgebra of $\mathfrak{g}$. Let Z an $L \cap K$-representation occurring in $H^N(\mathfrak{u}, X)$ and $\bar{Z} \subset Z$ an irreducible $L \cap K$-representation of highest weight ν where $\nu \in C_w$ for $l(w) = N$. If ν is safe with respect to $\mathfrak{q}$, then the restriction map (2.13) is nonzero on $\bar{Z}$.*

The condition of being safe is sufficient but not necessary for the restriction map to be nonzero. To place an edge it is enough to give a single K-type lying on it. This means that we will find all edges except possibly those whose K-types all lie very close to the boundary of W_K. Most of these "bad" cases are eliminated by Sepanski's definition of an edge, but there may be others, for example if λ itself is close to the walls of W_K.

This issue is not at all relevant to the statement of Theorem 1.4.1. Its only effect is on convenience of finding the walls in practice. That is certainly worth worrying about, as it is interesting to understand why some $L \cap K$-types of Z occur in the image of ρ even though they *look like* they should not. We shall come back to this issue at the very end, in Section 5.3.

2.4 Edges II

Let us get back to the discrete series, and the convex polytope containing their K-types. Recall from Section 1.4 that an edge of this polytope is described by the means of a θ-stable parabolic subalgebra $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$. The edge directions are given by the reductive part $\mathfrak{l}$ and outward directions are given by the nilpotent $\mathfrak{u}$. The definition used in Chapter 1 was that $\mu \in \mathfrak{h}_0^*$ lies on a $\mathfrak{q}$-edge if it appears as a K-type of X and no K-type is of the form $\mu + \delta$ where δ is in the real positive span of the roots in $\mathfrak{u}$. There are some technical issues concerning K-types that are close to the walls of the fundamental chamber. The precise definition, as motivated by [12], is therefore more complicated.

The goal to keep in mind is Sepanski's result discussed at the end of Section 1.4. Let $w \in W_K$ be the shortest element making $w(\mathfrak{u} \cap \mathfrak{k})$ positive, and let $n = l(w)$. A K-type μ in X gives an $L \cap K$-representation $\bar{Z} \subset H^n(w(\mathfrak{u} \cap \mathfrak{k}), X)$ of highest weight $w(\mu + \rho_c) - \rho_c$. If μ lies on a $\mathfrak{q}$-edge, we want $\bar{Z}$ to lie in the image of the restriction map $H^n(w\mathfrak{u}, X) \to H^n(w(\mathfrak{u} \cap \mathfrak{k}), X)$.

We need to examine the Hochschild-Serre spectral sequence again. Recall that the E_1-terms of this sequence are estimated by

$$(E_1^{a,b})^{\text{est}} = H^b(\mathfrak{u} \cap \mathfrak{k}, X) \otimes \bigwedge^a (\mathfrak{u} \cap \mathfrak{p})^*$$

The differentials are of degree $(r, 1-r)$ and none for $r > 0$ can hit the $(0, n)$ square. It follows that $\bar{Z}$ survives the spectral sequence if and only if it doesn't hit anything on the $n+1$ diagonal, which computes $H^{n+1}(\mathfrak{u}, X)$.

The estimated E_1 terms on the diagonal are $\omega \otimes_{\mathbb{C}} p_1^* \wedge \ldots \wedge p_i^*$ where p_j are elements of $w(\mathfrak{u} \cap \mathfrak{p})$ and ω is a class in $H^{n+1-i}(\mathfrak{u} \cap \mathfrak{k}, X)$. The survival of $\bar{Z}$ is guaranteed if none of these terms can have weight ν. The definition of an edge ensures that K-types which *could* produce such troublesome weights in $H^{n+1-i}(\mathfrak{u} \cap \mathfrak{k}, X)$ are absent from X.

Recall that every weight of $H^{n+1-i}(\mathfrak{u} \cap \mathfrak{k}, X)$ lies in some $C_{w'}$ where $l(w') = n+1-i$. We only need to worry if $\nu + \delta_1 + \ldots + \delta_i$ actually lies in some such $C_{w'}$ for some choice of roots $\{\delta_j\}$ of $w(\mathfrak{u}\cap\mathfrak{p})$. The original condition appearing in [12] was somewhat stronger than necessary; what is handed to us by the Hochschild-Serre spectral sequence is the following definition:

Definition 2.4.1. Fix X a $(\mathfrak{g}, K)$ module, and $\mathfrak{q}$ a θ-stable parabolic. Let $w \in W_K$ be the shortest element which makes $w(\mathfrak{u}\cap\mathfrak{k})$ positive, and let $n = l(w)$. We say that $\mu \in i\mathfrak{h}_0^*$ lies on a $\mathfrak{q}$-edge if it satisfies two conditions:

1. The K-type μ appears in X.
2. Let $\nu = w(\mu+\rho_c)-\rho_c$, and $\delta_1, \ldots \delta_i$ be distinct roots in $w(\mathfrak{u}\cap\mathfrak{p})$. If the weight $\nu' = \nu + \sum_{j=1}^{i} \delta_j$ lies in $C_{w'}$ for some $w' \in W_K$ of length $n+1-i$, then the K-type $w(\nu' + \rho_c) - \rho_c$ does not appear in X

Having set up the precise definitions, let us restate the main result in [12], which guarantees that we will find all edges by looking at the restriction map in cohomology:

Theorem 2.4.2. *With the above setup, let μ be on a $\mathfrak{q}$-edge. Then the restriction map in cohomology*

$$H^{l(w)}(w\mathfrak{u}, X) \to H^{l(w)}(w(\mathfrak{u}\cap\mathfrak{k}), X) \tag{2.19}$$

is surjective on the $L \cap K$-types $w(\mu+\rho_c)-\rho_c$.

Chapter 3

A nonzero map

This chapter is dedicated to the first major step in proving Theorem 1.4.1, which is finding a nonzero map to our discrete series X from a certain cohomologically induced module denoted $\mathcal{L}Z$. We will define the construction of $\mathcal{L}Z$ in Section 3.1, and show that if $G = K$ it produces an irreducible K-representation. Another special case of the cohomological induction produces the discrete series. We shall describe it in Section 3.2, and discuss some consequences which will be important later. We shall spend sections 3.3 and 3.4 on describing the map from $\mathcal{L}Z$ to X, and proving that it is nonzero, thereby having a chance to restrict the K-types of X as soon as we find out more about $\mathcal{L}Z$ in Chapter 4. Unless stated otherwise, all definitions and results in this chapter come from [8], where all details can be found.

3.1 The map

In Section 2.1.4 we promised to come back to the cohomological functors P and I. We shall do so now, in a specific setting: let $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ be a θ-stable parabolic subalgebra of $\mathfrak{g}$, and Z be an $(\mathfrak{l}, L \cap K)$-module. First we make Z into a $(\mathfrak{q}, L \cap K)$-module $Z_{\mathfrak{q}}$ by letting $\mathfrak{u}$ act by zero. Then we define

$$\mathcal{L}Z = P^{\mathfrak{g},K}_{\mathfrak{q},L\cap K} Z_{\mathfrak{q}} = R(\mathfrak{g}, K) \otimes_{R(\mathfrak{q},L\cap K)} Z_{\mathfrak{q}} \tag{3.1}$$

Recall from Section 2.1.4 that the functor $\mathcal{L}$ is covariant and right exact. We denote its i-th derived functor by $\mathcal{L}_i$.

The functors $\mathcal{L}_i^K$ from $(\mathfrak{l}\cap\mathfrak{k}, L\cap K)$-modules to $(\mathfrak{k}, K)$-modules are defined similarly. That is, they are the derived functors of $\mathcal{L}^K$, which takes any $(\mathfrak{l}\cap\mathfrak{k}, L\cap K)$-module $\bar{Z}$ to

$$\mathcal{L}^K\bar{Z} = P^{\mathfrak{k},K}_{\mathfrak{q}\cap\mathfrak{k},L\cap K}\bar{Z}_{\mathfrak{q}} = R(\mathfrak{k}, K) \otimes_{R(\mathfrak{q}\cap\mathfrak{k},L\cap K)} \bar{Z}_{\mathfrak{q}} \tag{3.2}$$

The relevance of these cohomological functors is clear from the main result of this chapter:

Proposition 3.1.1. *Under the assumptions of Theorem 1.4.1, we get a $\mathfrak{g}$ map from $\mathcal{L}_N Z$ to X. This map will be nonzero on an irreducible K-representation of highest weight $w(\mu+\rho_c)-\rho_c$.*

Before we go on to proving this proposition, let us see what $\mathcal{L}_i^K\bar{Z}$ is if $\bar{Z}$ is an irreducible representation of $L\cap K$. The answer turns out to be rather simple: it is either zero, or an irreducible representation of K. To see this, let us start by stating Theorem 5.120 from [8], which relates the derived functors of $\mathcal{L}$ to $\mathfrak{u}$-cohomology:

Theorem 3.1.2. *Let X be a $(\mathfrak{g}, K)$-module, and Z an irreducible representation of L. There are two first-quadrant spectral sequences with differentials of bidegree $(r, 1-r)$, with respective E_2 terms*

$$E_2^{j,i} \;=\; \mathrm{Ext}^j_{(\mathfrak{l},L\cap K)}\left(Z, H^i(\mathfrak{u}, X)\right) \tag{3.3}$$

$$\text{and}\qquad E_2^{j,i} \;=\; \mathrm{Ext}^j_{(\mathfrak{g},K)}\left(\mathcal{L}_i Z, X\right) \tag{3.4}$$

and with a common abutment

$$E_2^{j,i} \Longrightarrow E_\infty^{i+j} = \mathrm{Ext}^{i+j}_{(\mathfrak{q},L\cap K)}\left(Z, X\right) \tag{3.5}$$

These two sequences give us the $j = 0$ column homomorphisms

$$E_\infty^i \xrightarrow{\pi_1} E_2^{0,i} = \mathrm{Hom}_{(\mathfrak{l}, L\cap K)}\left(Z, H^i(\mathfrak{u}, X)\right) \tag{3.6}$$

and

$$E_\infty^i \xrightarrow{\pi_2} E_2^{0,i} = \mathrm{Hom}_{(\mathfrak{g}, K)}\left(\mathcal{L}_i Z, X\right) \tag{3.7}$$

We get similar sequence if we restrict everything above to K. In the case of compact groups, the spectral sequences actually collapse, and the E^2-edge homomorphisms are isomorphisms (see Chapter V of [8] for proof):

$$\mathrm{Ext}^i_{(\mathfrak{q}\cap\mathfrak{k}, L\cap K)}\left(\bar{Z}_\mathfrak{q}, X\right) \xrightarrow{\sim} \mathrm{Hom}_{(\mathfrak{l}\cap\mathfrak{k}, L\cap K)}\left(\bar{Z}, H^i(\mathfrak{u}\cap\mathfrak{k}, X)\right) \tag{3.8}$$

$$\mathrm{Ext}^i_{(\mathfrak{q}\cap\mathfrak{k}, L\cap K)}\left(\bar{Z}_\mathfrak{q}, X\right) \xrightarrow{\sim} \mathrm{Hom}_{(\mathfrak{k}, K)}\left(\mathcal{L}_i^K \bar{Z}, X\right) \tag{3.9}$$

In other words, we have

$$\mathrm{Hom}_{(\mathfrak{k}, K)}\left(\mathcal{L}_i^K \bar{Z}, X\right) \simeq \mathrm{Hom}_{(\mathfrak{l}\cap\mathfrak{k}, L\cap K)}\left(\bar{Z}, H^i(\mathfrak{u}\cap\mathfrak{k}, X)\right) \tag{3.10}$$

Let us now take X to be the irreducible representation E_μ of K with highest weight μ. The right-hand side of Equation (3.10) is given to us by Kostant's Theorem: if $\bar{Z}$ is one of the summands in $H^i(\mathfrak{u}\cap\mathfrak{k}, X)$, then there is an inclusion $\bar{Z} \subset H^i(\mathfrak{u}\cap\mathfrak{k}, X)$; otherwise we get $\mathrm{Hom}_{(\mathfrak{l}\cap\mathfrak{k}, L\cap K)}\left(\bar{Z}, H^i(\mathfrak{u}\cap\mathfrak{k}, X)\right) = 0$. This tell us exactly what $\mathcal{L}_i^K \bar{Z}$ is:

Proposition 3.1.3. *Let $\bar{Z}$ be an irreducible representation of $L \cap K$ with highest weight ν. If $\nu + \rho_c$ is singular, then $\mathcal{L}_i^K \bar{Z} = 0$. If it is regular, let $w \in W_K$ be the element making $w(\nu + \rho_c)$ dominant, and let $\mu = w(\nu+\rho_c)-\rho_c$. Then*

$$\mathcal{L}_i^K \bar{Z} = \begin{cases} E_\mu & \text{for} \quad i = l(w) \\ 0 & \text{for} \quad i \neq l(w) \end{cases} \tag{3.11}$$

Proof. If $\nu + \rho_c$ is singular, then it cannot appear in any $\mathfrak{u}\cap\mathfrak{k}$ cohomology. Equation (3.10) show that $\mathrm{Hom}_{(\mathfrak{k}, K)}\left(\mathcal{L}_i^K \bar{Z}, E_\mu\right) = 0$ for all E_μ. The module $\mathcal{L}_i^K \bar{Z}$ is a direct sum of K-types and has none as its quotient, so it must be zero.

If $\nu + \rho_c$ is regular, then Kostant's Theorem tells us

$$\mathrm{Hom}_{(\mathfrak{l}\cap\mathfrak{k}, L\cap K)}\left(\bar{Z}, H^{w'}(\mathfrak{k}_+, E'_\mu)\right) = \begin{cases} \mathbb{C} & \text{for } w' = w,\ \mu' = \mu \\ 0 & \text{otherwise} \end{cases} \tag{3.12}$$

Translating this to the left-hand side of Equation (3.10) shows that $\bar{Z} \simeq E_\mu$. □

3.2 A few things about discrete series

One way to construct the discrete series X with Harish-Chandra parameter λ is using cohomological induction functors. This is stated as Theorem 11.178 in [8] using the functor I. Theorem 5.99 from [8] gives isomorphisms between P and I and lets us translate the discrete series result to P - or, more specifically, the functor $\mathcal{L}$ of previous section. Our parabolic will be the Borel subalgebra $\mathfrak{b} = \mathfrak{h} + \mathfrak{n}$ with all roots of $\mathfrak{n}$ negative on λ. Recall that $w_0 \in W_K$ is the longest Weyl group element, of length $l(w_0) = S$. Let $\mathbb{C}_{\nu_0}$ be the one-dimensional representation of H with weight $\nu_0 = w_0\lambda - \rho$. Then

$$\mathcal{L}_i\mathbb{C}_{\nu_0} = \begin{cases} X & \text{for} \quad i = S \\ 0 & \text{for} \quad i \neq S \end{cases} \tag{3.13}$$

We shall spend this section on some consequences of this construction. The first one, which shall become clear in Chapter 4 is that the lowest K-type of a module constructed in this manner is $\mathcal{L}^S_K\mathbb{C}_{\nu_0}$. Proposition 3.1.3 tells us exactly what this is, namely the representation E_{μ_0}, where

$$\mu_0 = w_0(\nu_0 + \rho_c) - \rho_c = w_0^2(\lambda + \rho) + (w_0\rho_c - \rho_c) = (\lambda + \rho) - 2\rho_c \tag{3.14}$$

This is indeed the lowest K-type as derived from the Blattner formula in Section 2.1.2

Next we shall use Theorem 3.1.2 to show that there are no higher Ext groups between discrete series. This was done for Ext^1 in [11]; here we shall prove the general case, which is an unpublished result of Zuckerman.

Proposition 3.2.1. *Let X and Y be discrete series of G with parameters λ_X and λ_Y, respectively. The space $\mathrm{Ext}^i_{(\mathfrak{g},K)}(Y,X)$ is nonzero exactly when $i=0$ and $\lambda_X=\lambda_Y$.*

Proof. Let $\mathfrak{b}_Y = \mathfrak{h}+\mathfrak{n}_Y$ be the Borel subalgebra of $\mathfrak{g}$ with all roots of $\mathfrak{n}_Y$ negative on λ_Y. We shall use the fact that $Y = \mathcal{L}_s\mathbb{C}_{\nu_0}$, where $\nu_0 = w_0\lambda_Y - \rho$, in the spectral sequence (3.4). The E_2 terms are $E_2^{b,a} = \mathrm{Ext}^b_{(\mathfrak{g},K)}(\mathcal{L}_a\mathbb{C}_{\nu_0}, X)$. As noted above, $\mathcal{L}_a\mathbb{C}_{\nu_0} = 0$ except when $a = S$. Therefore the sequence collapses, and we get $E_\infty^{b+S} = \mathrm{Ext}^b_{(\mathfrak{g},K)}(Y,X)$

Let us now examine the sister spectral sequence (3.3). The E_2 terms in it are $E_2^{d,c} = \mathrm{Ext}^d_{(\mathfrak{h},H)}(\mathbb{C}_{\nu_0}, H^c(\mathfrak{n}_Y,X))$ There are no higher extensions for representations of H, so again the spectral sequence collapses and we get $E_\infty^c = \mathrm{Hom}_{\mathfrak{h},H}(\mathbb{C}_{\nu_0}, H^c(\mathfrak{n}_Y,X))$

The final cohomology are to be the same, which gives us

$$E_\infty^i = \mathrm{Ext}^i_{(\mathfrak{g},K)}(Y,X) = \mathrm{Hom}_{\mathfrak{h},H}\left(\mathbb{C}_{\nu_0}, H^{i+S}(\mathfrak{n}_Y,X)\right) \tag{3.15}$$

Theorem 2.3.1 describes the cohomology $H^\bullet(\mathfrak{n}_Y,X)$ in detail. We shall only need the fact that $H^{i+S}(\mathfrak{n}_Y,X)$ it is a direct sum of irreducible representations of H whose weights are $\{w\lambda_X-\rho\}_{w\in T}$ for some (possibly empty) subset $T\subset W_K$. If $\mathrm{Ext}^i_{(\mathfrak{g},K)}(Y,X)$ is to be nonzero, then ν_0 must be one of these weights. In other words,

$$\nu_0 = w_0\lambda_X - \rho = w\lambda_Y - \rho \tag{3.16}$$

for some $w\in W_K$. Since λ_X and λ_Y are both dominant for K, this can only happen if $w=w_0$ and $\lambda_X=\lambda_Y$, in which case the map in question is the isomorphism $Y\simeq X$.

□

3.3 A map of spectral sequences

Vogan's theorem 2.3.1 hands us all L-representations Z that occur in $H^i(\mathfrak{u},X)$. Proposition 3.1.1 is a statement of when we can guarantee that $\mathcal{L}Z$ has a nonzero map to X. We know this already for $\mathcal{L}^K Z$ and compact groups: the answer is *always*, as given by the isomorphism (3.10). The trick in proving Proposition 3.1.1 lies in relating the spectral sequences of Theorem 3.1.2 to those for $\mathfrak{k}$.

Proposition 3.3.1. *Suppose $\bar{Z}$ is an $L \cap K$-representation that occurs in the image of Z under the restriction map in cohomology. The restriction map, together with the inclusion $\bar{Z} \subset Z$ of $L \cap K$-modules, give us the following commutative diagram:*

$$\begin{array}{ccc} \mathrm{Hom}_{(\mathfrak{l},L\cap K)}\left(Z, H^i(\mathfrak{u},X)\right) & \xrightarrow{\rho} & \mathrm{Hom}_{(\mathfrak{l}\cap\mathfrak{k},L\cap K)}\left(\bar{Z}, H^i(\mathfrak{u}\cap\mathfrak{k},X)\right) \\ \uparrow\pi_1 & & \uparrow\pi_1^K \\ \mathrm{Ext}^i_{(\mathfrak{q},L\cap K)}\left(Z_{\mathfrak{q}},X\right) & \xrightarrow{\rho} & \mathrm{Ext}^i_{(\mathfrak{q}\cap\mathfrak{k},L\cap K)}\left(\bar{Z}_{\mathfrak{q}},X\right) \\ \downarrow\pi_2 & & \downarrow\pi_2^K \\ \mathrm{Hom}_{(\mathfrak{g},K)}\left(\mathcal{L}_i Z,X\right) & \xrightarrow{\rho} & \mathrm{Hom}_{(\mathfrak{k},K)}\left(\mathcal{L}_i^K \bar{Z},X\right) \end{array} \tag{3.17}$$

Proof. The vertical maps π_1 and π_2 in diagram (3.17) come from the spectral sequences (3.3) and (3.4). The same is true for π_1^K and π_2^K, with everything happening inside of K.

To get the sequence (3.3), let us take a projective resolution $P_\bullet$ of Z in the category of $(\mathfrak{l}, L\cap K)$-modules, and an injective resolution $I^\bullet$ of X in the category of $(\mathfrak{q}, L\cap K)$-modules. We make the P_i into $(\mathfrak{q}, L\cap K)$-modules $P_{i,\mathfrak{q}}$ by letting $\mathfrak{u}$ act by zero.

The zeroth term of the first spectral sequence is

$$E_0^{i,j} = \mathrm{Hom}_{(\mathfrak{q},L\cap K)}(P_{i,\mathfrak{q}}, I^j) = \mathrm{Hom}_{(\mathfrak{l},L\cap K)}(P_i, I^{j,\mathfrak{u}}) \tag{3.18}$$

Let us take cohomology in rows first, using $E_0^{i,j} = \mathrm{Hom}_{(\mathfrak{q},L\cap K)}(P_{i.\mathfrak{q}}, I^j)$. Each I^j is injective, so

$$H^i\left[\mathrm{Hom}_{(\mathfrak{q},L\cap K)}(P_{\bullet,\mathfrak{q}}, I^j)\right] = \mathrm{Hom}_{(\mathfrak{q},L\cap K)}\left(H^i[P_{\bullet,\mathfrak{q}}], I^j\right) \tag{3.19}$$

The complex $(P_\bullet)_\mathfrak{q}$ has cohomology only in degree zero, so the sequence collapses into the zeroth column, with $E_1^{0,j} = \mathrm{Hom}_{(\mathfrak{q},L\cap K)}(Z, I^j)$. Taking the cohomology in rows now gives

$$E_\infty^j = E_2^{0,j} = \mathrm{Ext}^j_{(\mathfrak{q},L\cap K)}(Z,X) \tag{3.20}$$

Let us go back to the beginning and take cohomology in columns first, using $E_0^{i,j} = \mathrm{Hom}_{(\mathfrak{l},L\cap K)}(P_i, I^{j,\mathfrak{u}})$. The P_i are projective, so we get

$$H^j\left[\mathrm{Hom}_{(\mathfrak{l},L\cap K)}(P_i, I^{\bullet,\mathfrak{u}})\right] = \mathrm{Hom}_{(\mathfrak{l},L\cap K)}\left(P_i, H^j[I^{\bullet,\mathfrak{u}}]\right) \tag{3.21}$$

This gives us $E_1^{i,j} = \mathrm{Hom}_{(\mathfrak{l},L\cap K)}(P_i, H^j(\mathfrak{u},X))$. Taking cohomology in rows next, we get the E_2 term

$$E_2^{i,j} = \mathrm{Ext}^i_{(\mathfrak{l},L\cap K)}(Z, H^j(\mathfrak{u},X)) \tag{3.22}$$

In particular, the zeroth column term is $E_2^{0,j} = \mathrm{Hom}_{(\mathfrak{l},L\cap K)}(Z, H^j(\mathfrak{u},X)$, and comes with the restriction map π_1 from the total cohomology.

The resolution $P_\bullet$, resp. $I^\bullet$, remains projective, resp. injective, if we restrict them to the category of $(\mathfrak{l}\cap\mathfrak{k}, L\cap K)$, resp $(\mathfrak{q}\cap\mathfrak{k}, L\cap K)$-modules. There is an inclusion $(I^j)^{\mathfrak{u}} \to (I^j)^{\mathfrak{u}\cap\mathfrak{k}}$, making a natural map

$$\mathrm{Hom}_{(\mathfrak{l},L\cap K)}(P_i, (I^j)^{\mathfrak{u}}) \xrightarrow{\rho} \mathrm{Hom}_{(\mathfrak{l}\cap\mathfrak{k},L\cap K)}(P_i, (I^j)^{\mathfrak{u}\cap\mathfrak{k}}) \tag{3.23}$$

This gives us a map of spectral sequences, resulting in the commutative square

$$\begin{array}{ccc}
\mathrm{Hom}_{(\mathfrak{l},L\cap K)}\left(Z, H^i(\mathfrak{u},X)\right) & \xrightarrow{\rho} & \mathrm{Hom}_{(\mathfrak{l}\cap\mathfrak{k},L\cap K)}\left(Z, H^i(\mathfrak{u}\cap\mathfrak{k},X)\right) \\
\uparrow{\scriptstyle \pi_1} & & \uparrow{\scriptstyle \pi_1^K} \\
\mathrm{Ext}^i_{(\mathfrak{q},L\cap K)}\left(Z_\mathfrak{q}, X\right) & \xrightarrow{\rho} & \mathrm{Ext}^i_{(\mathfrak{q}\cap\mathfrak{k},L\cap K)}\left(Z_\mathfrak{q}, X\right)
\end{array} \tag{3.24}$$

Let us now look at the spectral sequence (3.4). First we take a projective resolution $P_\bullet$ of $Z_\mathfrak{q}$ in the category of $(\mathfrak{q}, L\cap K)$-modules, and an injective resolution $I^\bullet$ of X in the category of $(\mathfrak{g}, K)$-modules.

The zeroth term of the spectral sequence is

$$E_0^{i,j} = \mathrm{Hom}_{(\mathfrak{g},K)}(\mathcal{L}P_j, I^i) = Hom_{(\mathfrak{q},L\cap K)}(P_j, I^i) \tag{3.25}$$

Repeating the techniques used above, let us take cohomology in rows first, using $E_{i,j}^{0} = \mathrm{Hom}_{(\mathfrak{q},L\cap K)}(P_j, I^i)$. The spectral sequence collapses to the first column, with $E_1^{0,j} = \mathrm{Hom}_{(\mathfrak{q},L\cap K)}(P_j, X)$. The E_2-term becomes

$$E_\infty^j = E_2^{0,j} = \mathrm{Ext}^j_{(\mathfrak{q},L\cap K)}(Z, X) \tag{3.26}$$

If we take cohomology in columns first, using $E_0^{i,j} = \mathrm{Hom}_{(\mathfrak{g},K)}(\mathcal{L}P_j, I^i)$, we get $E_1^{i,j} = \mathrm{Hom}_{(\mathfrak{g},K)}(\mathcal{L}_j Z, I^i)$. This gives us the E_2 term

$$E_2^{i,j} = \mathrm{Ext}^i_{(\mathfrak{l},L\cap K)}(\mathcal{L}_j Z, X) \tag{3.27}$$

In particular, the zeroth column term is $E_2^{0,j} = \mathrm{Hom}_{(\mathfrak{g},K)}(L_j Z, X)$, and comes with the restriction map π_2 from the total cohomology.

Again the resolutions behave nicely if we restrict everything to inside $(\mathfrak{k}, K)$. There is a natural map

$$\mathrm{Hom}_{(\mathfrak{q},L\cap K)}(P_j, I^i) \xrightarrow{\rho} \mathrm{Hom}_{(\mathfrak{q}\cap\mathfrak{k},L\cap K)}(P_j, I^i) \tag{3.28}$$

This gives a map of spectral sequences, resulting in the commutative square

$$\begin{array}{ccc} \mathrm{Ext}^i_{(\mathfrak{q},L\cap K)}(Z_{\mathfrak{q}}, X) & \xrightarrow{\rho} & \mathrm{Ext}^i_{(\mathfrak{q}\cap\mathfrak{k},L\cap K)}(Z_{\mathfrak{q}}, X) \\ \downarrow\pi_2 & & \downarrow\pi_2^K \\ \mathrm{Hom}_{(\mathfrak{g},K)}(\mathcal{L}_i Z, X) & \xrightarrow{\rho} & \mathrm{Hom}_{(\mathfrak{k},K)}(\mathcal{L}_i^K Z, X) \end{array} \tag{3.29}$$

This finishes the proof in the case where $\bar{Z} = Z$. To prove the general case, note that both spectral sequences are natural in the first variable. The inclusion $\bar{Z} \subset Z$ of $(\mathfrak{l}\cap\mathfrak{k}, L\cap K)$-representations results in a corresponding map of spectral sequences, which we can compose with the maps obtained above to get the diagram (3.17).

□

3.4 Proof of Proposition 3.1.1

The proof of Proposition 3.1.1 is essentially diagram chasing in (3.17). We begin with analyzing three of the four vertical maps.

Proposition 3.4.1. *The maps π_1, π_1^K and π_2^K in (3.17) are isomorphisms.*

Proof. We already know this for the maps π_1^K and π_2^K, which are Equations (3.8) and (3.9), respectively. Remember that π_2^K is the map taking the irreducible K-module $\mathcal{L}_i^K \bar{Z}$ of highest weight $w(\mu+\rho_c)-\rho_c$ into X.

Let us now examine the map π_1. Theorem 2.3.1 tells us that $H^i(\mathfrak{u}, X)$ is a direct sum of distinct discrete series for L. Proposition 3.2.1 shows that these have no higher Ext groups between them. This means that in the spectral sequence (3.3), we have

$$E_2^{j,i} = \mathrm{Ext}^j_{(\mathfrak{l},L\cap K)}\left(Z, H^i(\mathfrak{u}, X)\right) \tag{3.30}$$

for all $j > 0$. The spectral sequence collapses, and we get

$$E_\infty^i = E_2^{0,i} = \mathrm{Hom}_{(\mathfrak{l},L\cap K)}\left(Z, H^i(\mathfrak{u}, X)\right) \tag{3.31}$$

proving that π_1 is an isomorphism. □

We are now ready for the final step in proving Proposition 3.1.1. Let Z be our discrete series that occurs in $H^N(\mathfrak{u}, X)$, defining a nonzero element

$$\xi \in \mathrm{Hom}_{(\mathfrak{l},L\cap K)}\left(Z, H^N(\mathfrak{u}, X)\right) \tag{3.32}$$

We are assuming that Z has a nonzero restriction to $H^N(\mathfrak{u}\cap\mathfrak{k}, X)$, giving us a nonzero map

$$\rho(\xi) = \bar{\xi} \in \mathrm{Hom}_{(\mathfrak{l}\cap\mathfrak{k},L\cap K)}\left(\bar{Z}, H^N(\mathfrak{u}\cap\mathfrak{k}, X)\right) \tag{3.33}$$

Since the vertical maps π_2^K and π_1^K are isomorphisms, we get an element

$$\bar{\psi} = \pi_2^K \circ (\pi_1^K)^{-1}(\bar{\xi}) \in \mathrm{Hom}_{(\mathfrak{k},K)}\left(\mathcal{L}_N^K \bar{Z}, X\right) \tag{3.34}$$

which is our map taking the irreducible K-module $\mathcal{L}_N^K \bar{Z}$ of highest weight $w(\nu+\rho_c)-\rho_c$ into X.

Since π_1 has an inverse, we have an element

$$\psi = \pi_2 \circ \pi_1^{-1}(\xi) \in \mathrm{Hom}_{(\mathfrak{g},K)}\left(\mathcal{L}_N Z, X\right) \tag{3.35}$$

By commutativity of the diagram, $\rho(\psi) = \bar{\psi}$, proving that ψ is nonzero on the K-type $w(\nu+\rho_c)-\rho_c$ as asserted.

Chapter 4

Edge

4.1 Edge - a simplified case

We shall now analyze the K-types that occur in $\mathcal{L}_i Z$. To simplify the notation, we let

$$\begin{aligned} \mathrm{ind}_{\mathfrak{q}\cap\mathfrak{k}}^{\mathfrak{k}} &= P_{\mathfrak{q},L\cap K}^{\mathfrak{g},L\cap K} & \qquad \mathrm{ind}_{\mathfrak{q}}^{\mathfrak{g}} &= P_{\mathfrak{q}\cap\mathfrak{k},L\cap K}^{\mathfrak{k},L\cap K} \\ \Pi &= P_{\mathfrak{g},L\cap K}^{\mathfrak{g},K} & \qquad \Pi^K &= P_{\mathfrak{k},L\cap K}^{\mathfrak{k},K} \end{aligned}$$

The most immediate tool for analyzing $\mathcal{L}_i Z$ is a spectral sequence converging to $\mathcal{L}_i Z|_K$, taken from Chapter V of [8]. We shall sketch the idea behind it here, and show how to incorporate multiplication by $\mathfrak{g}$ into it.

We define the module

$$M = \mathrm{ind}_{\mathfrak{q}\cap\mathfrak{k}}^{\mathfrak{k}} Z = U(\mathfrak{g}) \otimes_{U(\mathfrak{q})} Z \tag{4.1}$$

Recall that $\mathcal{L}Z = \Pi M$. The functor $P_{\mathfrak{q},L\cap K}^{\mathfrak{g},L\cap K}$ is exact, so $\mathcal{L}_i Z = \Pi_i M$. As a K-module, this is the same as $\Pi_i^K M$, which is what our spectral sequence will calculate.

We define a filtration of M by submodules

$$M_j = \mathrm{span}\,\{X_1 \cdots X_{j'} u \otimes z \mid z \in Z,\ u \in U(\mathfrak{k}),\ X_k \in \mathfrak{g},\ j' \leq j\} \tag{4.2}$$

The pieces of the associated graded are

$$\bar{M}_j = U(\mathfrak{k}) \otimes_{U(\mathfrak{q}\cap\mathfrak{k})} \left[S^j(\mathfrak{g}/\mathfrak{k}+\mathfrak{q}) \otimes Z_{\mathfrak{q}}\right] = \mathrm{ind}_{\mathfrak{q}}^{\mathfrak{g}} \left[S^j(\mathfrak{g}/\mathfrak{k}+\mathfrak{q}) \otimes Z_{\mathfrak{q}}\right] \tag{4.3}$$

The spectral sequence is a double complex that calculates $\Pi_i^K[M]$ using this filtration.

Theorem 4.1.1. *There is a spectral sequence converging to $\mathcal{L}_{a+b}Z|_K$, with $b \geq 0$, $a \geq -b$, differential of bidegree $(-r, r-1)$ and E_1 terms*

$$E_1^{a,b} = \Pi_{a+b}^K \left(\mathrm{ind}_{\mathfrak{q}}^{\mathfrak{g}} \left(S^b(\mathfrak{g}/\mathfrak{k}+\mathfrak{q}) \otimes_{\mathbb{C}} Z_{\mathfrak{q}}\right)\right) \tag{4.4}$$

If $\mathfrak{u}\cap\mathfrak{k}$ acted by zero on $S^b(\mathfrak{g}/\mathfrak{k}+\mathfrak{q})$, we could write the E_1 term as

$$E_1^{a,b} = \mathcal{L}_{a+b}^K \left(S^b(\bar{\mathfrak{u}}\cap\mathfrak{p}) \otimes_{\mathbb{C}} Z_{\mathfrak{q}}\right) \tag{4.5}$$

In this special case, section 2.2.3 tells us exactly what K-representations can occur on the right-hand side. Let ν be a weight in $S^b(\bar{\mathfrak{u}}\cap\mathfrak{p}) \otimes_{\mathbb{C}} Z_{\mathfrak{q}}$. If ν lies in C_w for some w of length $a+b$, then $E_1^{a,b}$ will contain the K-type $w(\nu+\rho_c)-\rho_c$. Otherwise there is no contribution to $E_1^{a,b}$ from ν.

Unfortunately, in general $S^b(\mathfrak{g}/\mathfrak{k}+\mathfrak{q})$ is not a trivial module for $\mathfrak{u}\cap\mathfrak{k}$. To deal with this case, we shall use an analog of our approach from Section 2.2. We can define a filtration of $S(\mathfrak{g}/\mathfrak{k}+\mathfrak{q})$ such that $\mathfrak{u}\cap\mathfrak{k}$ acts by zero on the associated graded. The conclusion is that the K-types in $E_1^{a,b}$ are a subset of the K-types occurring in (4.5). Not all of them survive the spectral sequence, so there is even fewer in $\mathcal{L}_{a+b}Z|_K$. Equation (4.5) provides an upper bound estimate, which will be good enough for our purposes.

To summarize, the K-types in $\mathcal{L}_N Z$ are a subset of the K-types in the "estimated" module

$$\mathcal{L}_N Z^{\mathrm{est}} = \mathcal{L}_N^K \left(S(\bar{\mathfrak{u}}\cap\mathfrak{p}) \otimes_{\mathbb{C}} Z_{\mathfrak{q}}\right) \tag{4.6}$$

Let us analyze what these K-types might be, and what that tells us about possible edges.

Suppose that $Z \subset H^N(\mathfrak{u}, X)$ restricts to an $L \cap K$-representation $\bar{Z} \subset H^N(\mathfrak{u} \cap \mathfrak{k}, X)$, with highest weight $\nu \in C_w$. This tells us that X has a K-type $\mu = w(\nu + \rho_c) - \rho_c$. As a first-step analysis of where the other K-types might lie, let us treat a very special case and assume that the only length N chamber where weights of $S(\bar{\mathfrak{u}} \cap \mathfrak{p}) \otimes_{\mathbb{C}} Z_{\mathfrak{q}}$ lie is C_w, so they are the only ones contributing to $\mathcal{L}_N Z^{\text{est}}$.

Pick an $L \cap K$-type ν' in $S(\bar{\mathfrak{u}} \cap \mathfrak{p}) \otimes_{\mathbb{C}} Z_{\mathfrak{q}}$, where $\nu' \in C_w$. It gives us a K-type $\mu' = w(\nu' + \rho_c) - \rho_c$ in $\mathcal{L}_N Z^{\text{est}}$. The weight $\nu' - \nu$ is in the span of $\Delta(\bar{\mathfrak{q}})$, so $\mu' - \mu$ is in the span of $\Delta(w\bar{\mathfrak{q}})$. In other words, all K-types of X on a $\Delta(w\mathfrak{l})$ hyperplane through μ, or in the half-space in the direction of $w\bar{\mathfrak{u}}$ away from it. This is precisely saying that μ is on a $w\mathfrak{q}$-edge.

In general, the highest weights in $S(\bar{\mathfrak{u}} \cap \mathfrak{p}) \otimes_{\mathbb{C}} Z_{\mathfrak{q}}$ will lie in multiple W_K^1-chambers of length N, which makes the situation significantly more complicated. Indeed, let $w' \in W_K$ be another element of length N, and suppose that there is an $L \cap K$-type $\nu' \in S(\bar{\mathfrak{u}} \cap \mathfrak{p}) \otimes_{\mathbb{C}} Z_{\mathfrak{q}}$ that lies in $C_{w'}$. Then ν' will contribute the K-type $w'(\nu' + \rho_c) - \rho_c$ to $\mathcal{L}_N Z^{\text{est}}$, which could lie anywhere in the dominant chamber.

Unfortunately, there is no easy way to control the position of the K-types coming from chambers other than C_w. We shall need to introduce the step algebra $\mathcal{S}(\mathfrak{g}, \mathfrak{k})$, which acts on the set of $\mathfrak{k}$-highest weight vectors in any $\mathfrak{g}$-module, and will give us information about $X^{\mathfrak{k}+}$.

We shall analyze the action of $\mathcal{S}(\mathfrak{g}, \mathfrak{k})$ on $X^{\mathfrak{k}+}$ by observing how it acts on $(\mathcal{L}_i Z)^{\mathfrak{k}+}$. Before we get to all of that, we need to understand how $\mathfrak{g}$-action on $\mathcal{L}_i Z$ fits in with the spectral sequence (4.4). On $M = U(\mathfrak{g}) \otimes_{U(\mathfrak{q})} Z_{\mathfrak{q}}$, an element of $\mathfrak{g}$ acts by multiplying the first factor. This gives the $\mathfrak{g}$ structure on $\mathcal{L}_i Z = \Pi_i[M]$:

$$U(\mathfrak{g}) \otimes \Pi_i M = \Pi_i(U(\mathfrak{g}) \otimes M) \longrightarrow \Pi_i M \tag{4.7}$$

The multiplication by an element $t \in \mathfrak{p}$ does not preserve the spectral sequence filtration. Somewhat contrary to intuition, we shall *add* an extra filtration on M, and show that there is a reasonable way to track what t does to both of them combined.

Let Z be an irreducible $L \cap K$-representation in $H^N(\mathfrak{u}, X)$, and $\bar{Z}$ be an $L \cap K$-type in Z. Recall that in Section 4.1 we filtered M by the "number of $\mathfrak{p}$ roots away from $Z_{\mathfrak{q}}$". We filter it even further, adding the "$\mathfrak{p}$-distance" from $\bar{Z}$ inside Z:

$$M_{i,j} = \operatorname{span}\{X_1 \cdots X_{i'} u \otimes z \mid z \in Z_j,\ u \in U(\mathfrak{k}),\ X_k \in \mathfrak{g},\ i' \leq i\} \tag{4.8}$$

Multiplication by an element in $\mathfrak{g}$ sends $M_{i,j}$ to $M_{i+1,j} + M_{i,j+1}$. It is therefore reasonable to define

$$M[d] = \oplus_{i+j=d} M_{i,j} \tag{4.9}$$

Note that $M[d]$ is a subset of M_d in our original filtration. Tracking the $\mathfrak{g}$-action simplifies to $\mathfrak{g} \otimes M[d] \to M[d+1]$.

We are interested specifically in how $\mathfrak{g}$ acts on $\mathcal{L}_N^K \bar{Z} \subset \mathcal{L}_N Z$. Let v be an element in $\mathcal{L}_N Z$; it has a representative

$$v_1 \in \Pi_N^K(M[0]) \subset E_1^{N,0} \tag{4.10}$$

Let $t \in \mathfrak{g}$. The E_1 representative v_1' of tv lies in $\Pi_N^K M[1]$. The K-types occurring in it must lie in $\Pi_N^K(M[0])$ or in $\Pi_N^K(M[1]/M[0])$ The former is $\mathcal{L}_N^K \bar{Z}$; the latter is estimated by $\mathcal{L}_N^K((S^1(\bar{\mathfrak{u}} \cap \mathfrak{p}) \otimes \bar{Z}) + S^1(\mathfrak{l} \cap \mathfrak{p})\bar{Z})$, which is again estimated by $\mathcal{L}_N^K(S^1(\bar{\mathfrak{q}} \cap \mathfrak{p}) \otimes \bar{Z})$.

Proposition 4.1.2. *Let Z be an irreducible $L \cap K$-representation in $H^N(\mathfrak{u}, X)$, and $\bar{Z}$ an $L \cap K$-type in Z. Define the filtration $\{M[d]\}_{d \leq 0}$ of M as above. Then*

$$\mathfrak{g} \otimes \mathcal{L}_N^K \bar{Z} \subset \Pi_N^K M[1] \tag{4.11}$$

The K-types in $\Pi_N^K M[1]$ are estimated by

$$\mathcal{L}_N^K \left(S^{\leq 1}(\bar{\mathfrak{q}} \cap \mathfrak{p}) \otimes \bar{Z}\right) \tag{4.12}$$

4.2 Step algebra

The step algebra was first defined by Mickelsson in [10]. Denote by $U(\mathfrak{g})\mathfrak{k}_+ \subset U(\mathfrak{g})$ the left ideal generated by the positive root vectors in $\mathfrak{k}$. For any $\mathfrak{g}$-module V, denote by V^+ the set of $\mathfrak{k}_+$ highest weight vectors in V. The step algebra consist of elements in $U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+$ which act on V^+ for any V. The precise definition is:

$$\mathcal{S}(\mathfrak{g}, \mathfrak{k}) = \{u \in U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+ \mid xu \equiv 0 \pmod{U(\mathfrak{g})\mathfrak{k}_+} \;\; \forall x \in \mathfrak{k}_+\} \tag{4.13}$$

Let $\beta_1, \ldots, \beta_n$ be the roots in $\mathfrak{p}$, and $t_1, \ldots, t_n$ the corresponding root vectors. We say that $\beta_i < \beta_j$ if $\beta_j - \beta_i$ is a sum of positive compact roots. This defines a partial ordering on the roots. We can refine it to a total ordering, referred to as lexicographic: Let $\{h_1, \ldots, h_r\}$ be a basis of $\mathfrak{h}$ such that $\alpha(h_i) \geq 0$ for all $\alpha \in \Delta_+^{\kappa}$ and $i \leq r$. We say that $\beta_i \prec \beta_j$ if there is an index q such that $\beta_i(h_p) = \beta_j(h_p)$ for all $p < q$ and $\beta_i(h_q) < \beta_j(h_q)$. This total order $\prec$ is compatible with the partial order $<$ on the roots of p, in the sense that $\beta_i < \beta_j$ implies $i < j$. The ordering is not unique, but the lowest noncompact root is always first, and the highest one last.

For each root $\beta_i \in \mathfrak{p}$, Mickelsson constructed an *elementary step* $s_i \in \mathcal{S}(\mathfrak{g}, \mathfrak{k})$ of weight β_i:

$$s_i = t_i p_i + \sum_{\beta_j > \beta_i} u_i^j t_j p_i^j \tag{4.14}$$

Here p_i and p_i^j are polynomials in $\mathfrak{h}$, and u_i^j are polynomials in $\mathfrak{k}_-$.

Example 1. Let $\mathfrak{g} = \mathfrak{su}(2,1)$, with simple roots α and β. The maximal compact subalgebra is $\mathfrak{k} = s(u(2) \oplus u(1))$, along the root α. In this section, we shall follow Mickelsson's numbering of the roots of $\mathfrak{p}$, which is different from our previous examples. The highest root is $\beta_4 = \alpha + \beta$, $\beta_3 = \beta$; then $\beta_2 = -\beta$ and finally $\beta_1 = \beta_2 - \alpha$ is the lowest root.

Choose a basis $\{e, f, h, h'\}$ for $\mathfrak{k}$ where e, f, h are a standard $\mathfrak{sl}_2$ triple and $\mathfrak{h} = \langle h, h \rangle$ span $\mathfrak{h} \subset \mathfrak{k}$. The noncompact part $\mathfrak{p} \subset \mathfrak{g}$ splits into two 2-dimensional representations under $\mathfrak{k}$. One is spanned by $\{t_3, t_4\}$ and the other by $\{t_1, t_2\}$ where as usual t_i has weight β_i.

Let us find the elementary steps s_1 and s_2. Let E be a representation of $\mathfrak{k}$ with highest weight μ, such that $\mu(h) = m$; i.e. E has highest weight m with respect to our $\mathfrak{su}_2 \subset \mathfrak{k}$. For every $i \in [0, m]$, pick $v_{m-i\alpha} \in E_m$ a basis vector of weight $m - i\alpha$.

Now let E_m lie inside some $\mathfrak{g}$ representation X. The steps s_1 and s_2 are supposed to move the highest weight vector v_m to highest weight vectors in K-types $\mu+\beta_1$ and $\mu+\beta_2$, respectively.

Let E_1 be the two-dimensional representation of $\mathfrak{k}$ spanned by $\{t_1, t_2\}$. The tensor product of E and E_1 splits as a direct sum $E_1 \otimes E = E_{\mu+\beta_2} \oplus E_{\mu-\beta_2}$. The extra dimension of $\mathfrak{h}$ is of no consequence to finding the steps; let us drop it from the notation and just write $E_1 \otimes E = E_{m+1} \oplus E_{m-1}$.

The highest weight vector of E_{m+1} is $t_2 \otimes v_m$. In other words, applying t_2 to any highest weight vector in X again gives a highest weight vector; naturally, for t_2 itself is a highest weight vector with respect to α. Therefore $s_2 = t_2$.

The situation with t_1 is a little more complicated. The product $t_1 \otimes v_m$ is not a highest weight vector; it has a component in both E_{m-1} and E_{m+1}:

$$(m+1)t_2 \otimes v_m = (mt_2 \otimes v_m - t_1 \otimes v_{m-2}) + (t_2 \otimes v_m - t_1 \otimes vm - 2) \tag{4.15}$$

The first summand is a highest weight vector z_{m-1} of E_{m-1}; the second is equal to $f(t_1 \otimes v_m) \subset E_{m+1}$. We have

$$z_{m-1} = (m+1)t_2 \otimes v_m + f(t_1 \otimes v_m) \tag{4.16}$$

It follows that we can let $s_2 = t_2(h+1) + ft_1$.

Note that the polynomial $p_2 = (h+1)$ in the example is nonzero on every highest weight of $\mathfrak{k}$. This is important to ensure that s_2 actually has a nonzero component along t_2. In general the polynomials p_i satisfy the following condition: if $p_i(\mu) = 0$ then there is a $y \in W_K$ such that

$$y(\mu + \beta_i + \rho_c) = \mu + \beta_j + \rho_c \tag{4.17}$$

for some $\beta_j > \beta_i$ a root of $\mathfrak{p}$.

For the present we only need know the value of p_i on highest weights. We shall need it soon for weights in all chambers C_w. Let us analyze all the cases now.

Proposition 4.2.1. *Let β_i be a root of $\mathfrak{p}$, and μ a weight in C_w for some $w \in W_K$. Suppose $\mu + \beta_i \in C_w$. Then $p_i(\mu) \neq 0$.*

Proof. Suppose that $p_i(\mu) = 0$. In condition (4.17) we must have $w \neq 1$ for $\beta_j \neq \beta_i$. This means that the weight

$$\mu + \beta_j = y(\mu + \beta_i + \rho_c) - \rho_c \tag{4.18}$$

lies in some chamber $C_{w'}$ for $w' \neq w$. This is impossible by Proposition 2.2.1. Therefore $p_i(\mu) \neq 0$

□

Note that the proof only works for simply-laced Lie algebras. The proposition remains true in the non-simply laced case for weights in the fundamental chamber. For other chambers there is a complication which corresponds to some unusual properties of discrete series for those algebras; we shall discuss it in more detail in Chapter 5.

Mickelsson proved that the products of these s_i span the step algebra. We define $\mathcal{S}_0(\mathfrak{g}, \mathfrak{k}) \subset \mathcal{S}(\mathfrak{g}, \mathfrak{k})$ to be the linear span of the products $s_1^{l_1} \cdots s_n^{l_n}$ that are in our particular order. The following result was proved by van den Hombergh in [5]:

Theorem 4.2.2. *Let $\mathcal{M}$ be $\mathfrak{k}$-finite $U(\mathfrak{g})$-module and v a nonzero weight vector in $\mathcal{M}^+$; then $U(\mathfrak{g}) \cdot v = U(\mathfrak{k})\mathcal{S}_0(\mathfrak{g}, \mathfrak{k}) \cdot v$.*

Consider our discrete series X, and $\bar{Z}$ in the image of the restriction map ρ. Let Y be the K-submodule $\mathcal{L}_N^K \bar{Z}$ of highest weight $\mu = w(\nu + \rho_c) - \rho_c$ that occurs inside X, and let v be the highest weight vector in Y. Since discrete series are irreducible, Theorem 4.2.2 proves that $U(\mathfrak{k})\mathcal{S}_0(\mathfrak{g}, \mathfrak{k}) \cdot v$ is all of X. In other words, $\mathcal{S}_0(\mathfrak{g}, \mathfrak{k})$ acts irreducibly on X^+.

Ultimately, we want to show that X has no highest weights in the $w(\mathfrak{u} \cap \mathfrak{p})$ direction from $w(\nu + \rho_c) - \rho_c$. With that in mind, we start by applying the outward direction steps to v.

Proposition 4.2.3. *With the above notation, let β be a root in $w(\mathfrak{u} \cap \mathfrak{p})$, and s the corresponding step in $\mathcal{S}(\mathfrak{g}, \mathfrak{k})$. Then we have $s \cdot v = 0$.*

Proof. Proposition 4.1.2 tells us that the K-types that $s \cdot v$ lies in are a subset of

$$\mathcal{L}_N^K \left(S^{\leq 1}(\bar{\mathfrak{q}} \cap \mathfrak{p}) \otimes \bar{Z} \right) \tag{4.19}$$

The $L \cap K$-types of $S^{\leq 1}(\bar{\mathfrak{q}} \cap \mathfrak{p}) \otimes \bar{Z}$ are of the form $\nu + \Delta(\bar{\mathfrak{q}} \cap \mathfrak{p})$. Proposition 2.2.1 shows that the only chamber of length N that these can lie in is C_w. We conclude that the K-types in (4.19) are of the form

$$w(\nu + \beta' + \rho_c) - \rho_c = w(\nu + \rho_c) - \rho_c + \beta' = \mu + w(\beta') \tag{4.20}$$

This is of course equal to the weight $\mu' = \mu + \beta$ of v', which is a contradiction because we assumed that $\beta \in \Delta(w(\mathfrak{u} \cap \mathfrak{p}))$ while $w\beta' \in w(\bar{\mathfrak{q}} \cap \mathfrak{p})$ $\square$

Let us now examine the extremely favorable case where all the roots in $w(\mathfrak{u} \cap \mathfrak{p})$ are the last in our ordering; i.e. $w(\mathfrak{u} \cap \mathfrak{p}) = \{\beta_{m+1}, \ldots \beta_n\}$ for some $m \leq n$. Then we have $s_1^{l_1} \cdots s_n^{l_n} \cdot v = 0$ whenever $l_r > 0$ for any $r > m$. By Theorem 4.2.2, every highest weight vector of X can be written as a linear combination of $s_1^{l_1} \cdots s_n^{l_m} \cdot v$. The weight of this product is $\mu + \sum_{i=1}^{m} l_i \beta_i$, which lies in the $w\bar{\mathfrak{q}}$ direction from μ and shows that μ is on a $w\mathfrak{q}$-edge.

Unfortunately, this scenario is not very common. Indeed, consider any edge through the lowest K-type μ_0. We know that all the K-types lie in the cone spanned by the positive noncompact roots starting at μ_0. The outward directions from our edge will therefore all be negative, and cannot be ordered last.

The elementary steps cannot simply be reordered to account for this. We will therefore "redirect" the edge by defining a step algebra action on $H^{l(w)}(\mathfrak{k}_+, X)$. We will show that the action preserves $H^w(\mathfrak{k}_+, X)$ and acts irreducibly on it. As stated before, the weights in $H^w(\mathfrak{k}_+, X)$ are in one-to-one correspondence to the highest weights of X. In particular, μ occurs on an "edge" in direction of $\mathfrak{l}$ with outward directions $\mathfrak{u} \cap \mathfrak{p}$. The roots in $\mathfrak{u} \cap \mathfrak{p}$ can be ordered last, putting us in a perfect position to prove Theorem 1.4.1.

4.3 A little category theory

For a $\mathfrak{g}$-module V, we can express the space of its $\mathfrak{k}$-highest weights as follows:

$$V^{\mathfrak{k}_+} = H^0(\mathfrak{k}_+, V) = \mathrm{Hom}_{\mathfrak{k}_+}(\mathbb{C}, V) = \mathrm{Hom}_{\mathfrak{g}}(U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+, V) \tag{4.21}$$

The set of natural transformations of $\mathrm{Hom}_{\mathfrak{g}}(U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+, \_\_)$ as a functor from $\mathcal{C}(\mathfrak{g}, K)$ to Vect is

$$\mathbb{S} = \mathrm{Hom}_{\mathfrak{g}}(U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+,\ U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+) \tag{4.22}$$

via action on the first factor of $\mathrm{Hom}_{\mathfrak{g}}(U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+, \_\_)$.

Proposition 4.3.1. *There is an algebra isomorphism $\Phi : \mathcal{S}(\mathfrak{g}, \mathfrak{k}) \to \mathbb{S}$, defined by $s \mapsto \phi_s = [1 \mapsto s]$.*

Proof. First note that any element $\phi \in \mathbb{S}$ is completely defined by its action on 1: if $\phi(1) = t$ then $\phi(u) = ut$ for any $u \in U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+$.

The map ϕ_s is well defined: let $x \in \mathfrak{k}_+$, and $u \in U(\mathfrak{g})$. By definition, $\phi_s(ux) = uxs$. Since $s \in \mathcal{S}(\mathfrak{g}, \mathfrak{k})$, we know that $xs \equiv 0 \bmod U(\mathfrak{g})\mathfrak{k}_+$, so $\phi_s(ux) = 0$.

For $s, t \in \mathcal{S}(\mathfrak{g}, \mathfrak{k})$, we have $\phi_{st} = [1 \mapsto st]$. The map ϕ_s sends t to $\phi_s(t) = st]$, which means that $[1 \mapsto st] = [1 \mapsto s] \circ [1 \mapsto t]$. This translates to $\phi_{st} = \phi_s \phi_t$, proving that Φ is an algebra homomorphism.

Let us first prove that Φ is injective. Suppose that $\phi_s = \phi_{s'}$ for some $s' \in \mathcal{S}(\mathfrak{g}, \mathfrak{k})$. This means that $\phi_s(1) = \phi_{x'}(1)$, or $s - s' \in U(\mathfrak{g})\mathfrak{k}_+$. The step algebra is defined as a subset of $U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+$, so $s - s' = 0 \in \mathcal{S}(\mathfrak{g}, \mathfrak{k})$.

To prove surjectivity of our correspondence, suppose we have a map $\phi \in \mathbb{S}$, and let $s = \phi(1)$. By definition of $\mathbb{S}$, every element $x \in \mathfrak{k}_+$ satisfies $\phi(x) = 0$, which means that $xs \equiv 0 \mod U(\mathfrak{g})\mathfrak{k}_+$, and s belongs to the step algebra. We know that $\phi(1) = \phi_s(1)$, so $\phi = \phi_s$. □

It follows from Proposition 4.3.1 that the step algebra acts on the derived functors of $\mathrm{Hom}_{\mathfrak{g}}(U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+, \_\_)$.

Theorem 4.3.2. *Let X be a $(\mathfrak{g}, K)$ module. The the step algebra $\mathcal{S}(\mathfrak{g}, \mathfrak{k})$ acts on $H^i(\mathfrak{k}_+, X)$ for all i.*

Mickelsson's definition of the step algebra $\mathcal{S}(\mathfrak{g}, \mathfrak{k})$ only assumes that $\mathfrak{k}$ is a reductive subalgebra of $\mathfrak{g}$. Theorem 4.3.2 is true in that generality. The results that follow, concerning the action of $\mathcal{S}(\mathfrak{g}, \mathfrak{k})$ on $H^i(\mathfrak{k}_+, X)$, rely heavily on the fact that $\mathfrak{k}$ is compact.

Let f be an element of $H^w(\mathfrak{k}_+, X)$ for w of length N, and s_j one of the elementary steps. The element $(s_j f)$ lies in $H^N(\mathfrak{k}+, X)$. Proposition 2.2.1 shows that its weight cannot lie in any chamber of length N other than C_w, so $s_j \cdot f \in H^w(\mathfrak{k}_+, X)$. The step algebra is spanned by products of the s_i, so it acts on each $H^w(\mathfrak{k}_+, X)$ separately; and, in fact, this action is irreducible whenever X is an irreducible $\mathfrak{g}$-module.

Theorem 4.3.3. *Let X be a $(\mathfrak{g}, K)$-module which is irreducible under the action of $U(\mathfrak{g})$, and let f be an element in $H^w(\mathfrak{k}_+, X)$. Then $\mathcal{S}_0(\mathfrak{g}, \mathfrak{k}) \cdot f = H^w(\mathfrak{k}_+, X)$.*

Proof. We already know this for $w = 1$, and shall use it to prove all the other cases. The first ingredient we need is a result whose rather technical proof is deferred to the appendix.

Proposition 4.3.4. *Let v be an element in $H^0(\mathfrak{k}_+, X)$ corresponding to $f \in H^w(\mathfrak{k}_+, X)$. Suppose that $v' = s_i' v$ is a nonzero element of $H^0(\mathfrak{k}_+, X)$. Let $\beta_j = w^{-1}\beta_i$. Then $(s_j f)$ is a nonzero element of $H^w(\mathfrak{k}_+, X)$ corresponding to v'.*

Theorem 4.2.2 tells us that the products $s_1^{l_1} \cdots s_n^{l_n}$ span all of $H^0(\mathfrak{k}_+, X)$. If we let $\beta_{j_i} = w^{-1}\beta_i$, then it follows that $s_{j_1}^{l_{j_1}} \cdots s_{j_n}^{l_{j_n}}$ span all of $H^w(\mathfrak{k}_+, X)$.

Mickelsson proved in [10] that for simply laced Lie algebras, any product of the steps can be written as a sum of products $s_1^{i_1} \cdots s_n^{i_n}$ with polynomials in $\mathfrak{h}$ as coefficients. This means that $s_{j_1}^{k_1} \cdots s_{j_n}^{k_n}$ lies in $\mathcal{S}_0(\mathfrak{g}, \mathfrak{k})$, and therefore $\mathcal{S}_0(\mathfrak{g}, \mathfrak{k}) f = H^w(\mathfrak{k}_+, X)$ as claimed.

□

Let us show here how this finishes the proof of Theorem 1.4.1. With setup as in the theorem, let $Y = \mathcal{L}_i^K \bar{Z}$ be the K-representation of highest weight $w(\nu + \rho_c) - \rho_c$

that comes from $\bar{Z} \subset H^N(\mathfrak{k}_+, X)$. Let v be the highest weight vector in Y and f the corresponding element in $H^w(\mathfrak{k}_+, X)$. We first prove a result analogous to Proposition 4.2.3.

Proposition 4.3.5. *With the above notation, let β be a root in $\mathfrak{u}$, and $s \in \mathcal{S}(\mathfrak{g}, \mathfrak{k})$ the corresponding elementary step. Then we have $sf = 0$.*

Proof. For details on how the step algebra acts on $H^i(\mathfrak{k}_+, X)$ see the appendix. All we need to note here is that the element f comes from $H^w(\mathfrak{k}_+, Y) \subset H^w(\mathfrak{k}_+, X)$. By an argument exactly analogous to the proof of Proposition 4.2.3, we know that $(sf) \in H^w(\mathfrak{k}_+, \oplus Y_i)$, where the highest weight of Y_i is $w(\nu + \beta_i + \rho_c) - \rho_c$ for some $\beta_i \in \Delta(\bar{\mathfrak{q}} \cap \mathfrak{p})$.

The weight of $H^w(\mathfrak{k}_+, Y_i')$ is $\nu + \beta_i$. The weight of (sf) is $\nu + \beta$, so $\beta_i = \beta$ for all i, which contradicts $\beta \in \Delta(\mathfrak{u})$.

□

We shall now define a specific lexicographic ordering $\prec$ (see Section 4.2). Pick $h_1 \in \mathfrak{h}$ such that $\beta(h_1) > 0$ for all $\beta \in \mathfrak{u}$, and $\beta'(h_1) = 0$ for all $\beta' \in \mathfrak{l}$. This is certainly possible, and we can extend it by choosing arbitrary $h_2, \ldots, h_r$ that complete the basis for $\mathfrak{h}$. This ordering puts the roots in $\mathfrak{u}$ last, say $\beta_{m+1}, \ldots \beta_n$.

We are now in a position very similar to our "extremely favorable case" from Section 4.2, and can repeat the argument exactly: A product $s_1^{l_1} \cdots s_n^{l_n} \cdot f$ which is nonzero must have $l_{m+1} = \ldots = l_n = 0$. This means that there are no weights in $H^w(\mathfrak{k}_+, X)$ in the $\mathfrak{u} \cap \mathfrak{p}$ directions from μ, and therefore no highest weights in $w(\mathfrak{u} \cap \mathfrak{p})$ directions from $w(\mu + \rho_c) - \rho_c$. It follows that the K-type $w(\mu + \rho_c) - \rho_c$ is on $w\mathfrak{q}$-edge.

Chapter 5

The algorithm

5.1 Examples

5.1.1 $SO(2)$

Let G be the 1-dimensional compact torus $G = SO(2)$. In this case $G = K = H$, so there are no roots and $W_K = \{1\}$. The "discrete series" X_λ is simply a weight space of weight λ. Indeed, there is only one summand in the Blattner formula, with $\rho_n = \rho_c = 0$ and $\Delta_K = 1$, leaving the character as e^λ.

5.1.2 $SU(1,1)$

To illustrate the procedure of finding edges in nontrivial cases, we will first look at the rank one group $G = SU(1,1)$. Remember from Chapter 1 that in this case $\mathfrak{k} = \mathfrak{h}$, which means that K-types are weights, and the Weyl group $W_K = \{1\}$, with a single chamber $C_1 = i\mathfrak{h}_0^*$. The noncompact roots are $\pm\beta$, with the corresponding fundamental weight $\Lambda = \frac{1}{2}\beta$. Let X be the discrete series with parameter $\lambda = 3\Lambda$, shown in Figure 1-3. In Section 1.2 we computed the K-types as $\{4\Lambda, 6\Lambda, 8\Lambda, \ldots\}$.

Forgetting we already know the K-types, let us look for the edges using cohomology techniques. The group G has no compact factors so the dimension of the convex set containing the K-types of X is $D = \mathrm{rk}(\mathfrak{g}) = 1$. We have to pick a θ-stable parabolic subalgebra $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ that satisfies $\mathrm{rk}(\mathfrak{l}) = D - 1 = 0$ and $\mathfrak{u} \cap \mathfrak{k} \subset \mathfrak{k}_+$. There are no compact roots so the latter condition is vacuous. The former implies that that $\mathfrak{l} = \mathfrak{h} = \mathfrak{k}$ and $\mathfrak{u}$ runs along one of the noncompact roots. The Weyl group W_K is trivial, so Theorem 2.3.1 says that $H^\bullet(\mathfrak{u}, X)$ has exactly one summand X_{λ_1} with parameter $\lambda_1 = \lambda - \rho_u$. The algebra $\mathfrak{k}$ is a torus, so X_1 is a one-dimensional weight space for $\mathfrak{k}$ of weight $\mu_1 = \lambda_1$.

Let us first examine the case where $\Delta(\mathfrak{u}) = \{\beta\}$. The parameter is

$$\lambda_1 = \lambda - \rho_u = 3\Lambda - \Lambda = 2\Lambda \tag{5.1}$$

The noncompact root $\beta \in \mathfrak{u}$ is positive on λ_1, so according to Theorem 2.3.1 again, we know that X_1 will occur in degree 1. The subalgebra $\mathfrak{k}$ is equal to $\mathfrak{h}$, so $H^i(\mathfrak{u}\cap\mathfrak{k}, X) = 0$ for all $i > 0$. This means that the restriction map in cohomology is zero on X_1 and we do not get any edges from $\mathfrak{q}$.

Next we take $\mathfrak{q}$ with $\Delta(\mathfrak{u}) = \{-\beta\}$. We get

$$\lambda_1 = \lambda - \rho_u = 3\Lambda + \Lambda = 4\Lambda \tag{5.2}$$

In this case, the noncompact root $-\beta \in \mathfrak{u}$ is negative on λ_1, so X_1 occurs in degree zero. Its single $L \cap K$-type $\mu_1 = \lambda_1 = \mu$ lives in $C_1 = i\mathfrak{h}_0^*$. The condition of being safe is automatically satisfied for degree zero cohomology. By Proposition 2.3.3, μ_1 is in the image of ρ, giving us an edge through μ with outward direction $-\beta$.

5.1.3 $SU(2,1)$

We shall look at a discrete series for the group $SU(2,1)$, using notation from Section 1.2. We will take the same discrete series X with Harish-Chandra parameter $\lambda = 3\Lambda_\alpha + 2\Lambda_\beta$ and minimal K-type $\mu_0 = 2\Lambda_\alpha + 4\Lambda_\beta$, which is shown on Figure 1-4. Again G has no compact factors so $D = \mathrm{rk}(\mathfrak{g}) = 2$. Our parabolic subalgebras $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ should have $\mathrm{rk}(\mathfrak{l}) = 1$. The Lie algebra $\mathfrak{su}(2,1)$ has compact roots, so we also need to satisfy $\mathfrak{u} \cap \mathfrak{k} \subset \mathfrak{k}_+ = \{\alpha\}$.

Let us first consider the parabolic $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ where $\mathfrak{l}$ is in the β direction, and

$\Delta(\mathfrak{u}) = \{\alpha, \beta_2\}$. By Theorem 2.3.1, we know that there will be two discrete series Z_1 and Z_r of $\mathfrak{l}$ occurring in $H^\bullet(\mathfrak{u}, X)$, corresponding to the two elements in $W_K = \{1, r\}$. The Lie algebra $\mathfrak{l}$ is a copy of $\mathfrak{su}(1,1)$, so we knows its discrete series by what we just computed in Section 5.1.2.

The discrete series for Z_1 has Harish-Chandra parameter λ_1 and lowest $L \cap K$-type μ_1 which are equal to

$$\lambda_1 = \lambda - \rho_u = 1.5\Lambda_\alpha + 2\Lambda_\beta \quad (5.3)$$

$$\mu_1 = \lambda_1 + \frac{1}{2}\beta = \Lambda_\alpha + 3\Lambda_\beta \quad (5.4)$$

The compact root α of $\mathfrak{u}$ is positive on λ_1, and so is the noncompact root β_2. By Equation (2.16), Z_1 occurs in $H^1(\mathfrak{u}, X)$. To find its restriction to $H^1(\mathfrak{u} \cap \mathfrak{k}, X)$, we need to know which of its $L \cap K$-types $\lambda_1 + (n + \frac{1}{2})\beta$ lie in C_r, meaning that they have product at most -1 with α. A quick calculation reveals that this happens precisely for $n \geq 4$. All of those $L \cap K$-types are guaranteed to be safe by Proposition 2.2.1. By Proposition 2.3.3, they restrict to nonzero elements of $H^1(\mathfrak{u} \cap \mathfrak{k}, X)$.

By Theorem 1.4.1, we get an edge of K-types through

$$\mu = r(\lambda_1 + 4.5\beta + \rho_c) - \rho_c = r(\lambda_1 + 4.5\beta + 0.5\alpha) - 0.5\alpha = \Lambda_\alpha + 9\Lambda_\beta \quad (5.5)$$

in direction of $r(\beta) = \beta_2$ and with outward direction $r(\beta_2) = \beta$. The K-types guaranteed to occur on this edge will be precisely $\Lambda_\alpha + 9\Lambda_\beta + m\beta_2$ for m a nonnegative integer. This edge is labeled E_1 in Figure 5-1.

Let us now examine the discrete series Z_r, whose parameters are

$$\lambda_r = -4.5\Lambda_\alpha + 5\Lambda_\beta \quad (5.6)$$

$$\mu_r = -5\Lambda_\alpha + 6\Lambda_\beta \quad (5.7)$$

The compact root α of $\mathfrak{u}$ is negative on λ_r, and the noncompact root β_2 is positive on it. It follows that Z_r occurs in $H^2(\mathfrak{u}, X)$. The highest nonzero cohomology for $\mathfrak{u} \cap \mathfrak{k}$ is in degree 1, so Z_r restricts to zero, and gives no boundaries.

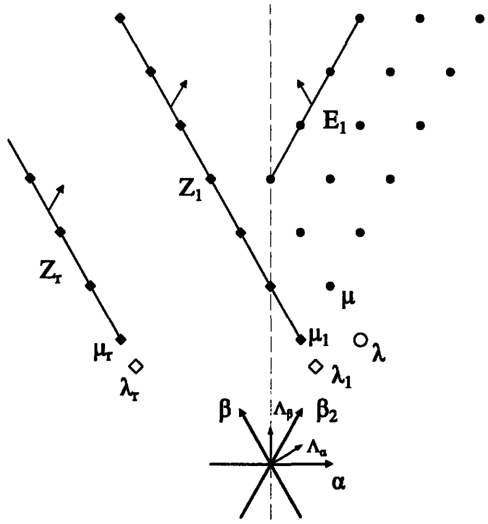

Figure 5-1: Walls, $\Delta(\mathfrak{l}) = \{\beta\}$

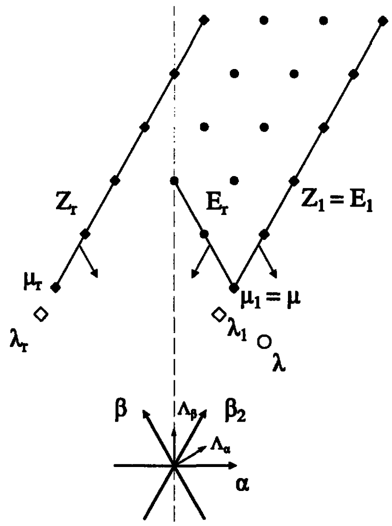

Figure 5-2: Walls, $\Delta(\mathfrak{l}) = \{\beta_2\}$

The other parabolic with $\mathrm{rk}(\mathfrak{l}) = 1$ and $\mathfrak{u} \cap \mathfrak{k}$ positive has $\mathfrak{l}$-direction given by β_2, and $\Delta(\mathfrak{u}) = \{\alpha, -\beta\}$. Again we get two discrete series Z_1 and Z_r in $H^\bullet(\mathfrak{u}, X)$, both shown in Figure 5-2. Let us start with Z_1:

$$\lambda_1 = 1.5\Lambda_\alpha + 3.5\Lambda_\beta \tag{5.8}$$

$$\mu_1 = 2\Lambda_\alpha + 4\Lambda_\beta \tag{5.9}$$

The parameter λ_1 has positive product with α and negative with $-\beta$, so Z_1 occurs in $H^0(\mathfrak{u}, X)$. Its $L \cap K$-types are $\mu_1 + n\beta_2$ for $n \in \mathbb{Z}_{\geq 0}$. They all lie in C_1, and are automatically safe. We get an edge E_1 through μ_1, in β_2-direction, with outward direction given by $-\beta$. Note that μ_1 is the lowest K-type occurring in X; the other K-types on the edge are precisely $\mu_1 + m\beta_2$ for m a positive integer.

The data for Z_r are

$$\lambda_r = -4.5\Lambda_\alpha + 6.5\Lambda_\beta \tag{5.10}$$

$$\mu_r = -4\Lambda_\alpha + 7\Lambda_\beta \tag{5.11}$$

The parameter λ_1 has negative product with both α and $-\beta$, so Z_1 occurs in $H^1(\mathfrak{u}, X)$. Its $L \cap K$-types are $\mu_r + n\beta_2$, for $n \in \mathbb{Z}_{\geq 0}$. They are in C_r for $n \in \{0, 1\}$, and both are automatically safe. They give us an edge E_r containing the K-type

$$r(\mu_r + \rho_c) - \rho_c = 2\Lambda_\alpha + 4\Lambda_\beta = \mu_0 \tag{5.12}$$

The edge runs in direction of $r(\beta_2) = \beta$, with outward direction given by $r(-\beta) = \beta_2$. The other K-type guaranteed to occur on this edge is $\mu_0 + \beta$. Both edges are shown in Figure 5-2, and complete the set of boundaries of X.

5.2 Algorithm

Let us now return to the case where G is any group satisfying the hypotheses of Theorem 1.4.1. Recall that D is the dimension of the minimal convex polyhedron containing the K-types of X (see Section 1.4). With notation and assumptions as in the Theorem, the algorithm for finding the edges of our discrete series X is

1. Let $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ run through all θ-stable parabolic subalgebras of $\mathfrak{g}$ that satisfy $\mathrm{rk}(\mathfrak{l}) = D - 1$ and $\mathfrak{u} \cap \mathfrak{k} \subset \mathfrak{k}_+$.

2. Use Theorem 2.3.1 to find all the discrete series of $\mathfrak{l}$ that occur in $H^\bullet(\mathfrak{u}, X)$. Let Z run through the ones in degrees $N \leq l(w_0)$, where w_0 is the longest element in W_K.

3. Since $\mathrm{rk}(\mathfrak{l}) < \mathrm{rk}(\mathfrak{g})$, we can find the $L \cap K$-types of Z by induction. Let w go through the elements of W_K of length N such that Z has a safe $L \cap K$-type $\nu \in C_w$.

4. Proposition 2.3.3 guarantees that the $L \cap K$-type ν appears in the image of Z under the restriction map in cohomology. By Kostant's Theorem, X has a K-type $\mu = w(\nu + \rho_c) - \rho_c$. Theorem 1.4.1 tells us that μ lies on a $w\mathfrak{q}$-edge.

Now that we have seen some examples as well as the algorithm in general, let us make the following observations:

1. Theorem 1.4.1 is true for any parabolic $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ satisfying its hypotheses; that is, we can use it to look for edges of arbitrary dimension. However, only the maximal-dimension edges are needed to completely describe our convex set.

2. In the last step, we looked only at a single safe $L \cap K$-type μ of Z in the chamber C_w. This is in fact sufficient. Suppose that μ' is another $L \cap K$-type in C_w that is hit by the restriction map for the same $\mathfrak{q}$. It gives an edge E' parallel to E with the same outward directions. Since both edges have a K-type of X on them, if follows that in fact $E = E'$.

5.3 Future work

Given a reductive Lie group G with simply-laced Lie algebra, and X a discrete series for G, we have an algorithm for finding the edges of the minimal convex shape containing the K-types of X.

1. What goes wrong for non simply-laced Lie algebras is the result of Theorem 4.3.3. Already the proof of Proposition 4.3.4 does not work if $\mathfrak{g}$ is not simply-laced. In fact, the polynomial u_i multiplying t_i in some of the s_i is zero on weights lying very close to the chamber boundary. The problem is not just with techniques of proof, but with what is true about the representations: For low rank Lie algebras that I examined, there actually are "holes" in the discrete series on these edges. That is, $s_i \cdot f$ has to be zero because $u_i = 0$; but there are, in fact, no elements of weight $wt(f) + \beta_i$ in $H^w(\mathfrak{k}_+, X)$. Understanding this phenomenon might allow me to extend Theorem 1.4.1 to all Lie algebras.

2. The "holes" are rather fascinating: All K-types occurring in X lie in the convex set described by Theorem 1.4.1. The Blattner formula shows that they are congruent to μ_0 modulo the lattice of noncompact roots. As mentioned above, it is not true that all K-types satisfying both conditions will have nonzero multiplicity. I conjecture that only the lattice points closest to a wall of the

fundamental Weyl chamber can "disappear", and that it only happens for non-simply laced Lie algebras.

3. The boundaries close to walls of the Weyl chamber are interesting in their own right. In Section 2.3 we discussed briefly some walls which have no safe K-types yet they are not eliminated by Sepanski's definition of an edge so we *were* supposed to find them. My guess is this never happens for "generic" Harish-Chandra parameters; but might if λ is very close to the origin, for example, or any wall of W_K. I would like to find out exactly what the definition of "generic" is that eliminates this kind of issue.

   In cases where λ is not "generic", Theorem 1.4.1 still holds, but is much less convenient to work with, for we do not understand the restriction map in cohomology well enough. I would like to study it some more, and figure out exactly what it does to the "uncertain" $L \cap K$-types: those that lie in a chamber C_w of the right length so they *might* restrict, but they are not safe so it is not guaranteed.

4. Theorem 4.3.2 fits awkwardly with the proof of Theorem 1.4.1. The representation Z arose naturally in $H^i(\mathfrak{u} \cap \mathfrak{k}, X)$, yet the step algebra acts on cohomology with respect to all of $\mathfrak{k}_+$. We can define a generalized step algebra:

$$\mathcal{S}_\mathfrak{u}(\mathfrak{g}, \mathfrak{k}) = \mathrm{Hom}_{U(\mathfrak{g})}(U(\mathfrak{g})/U(\mathfrak{g})(\mathfrak{u} \cap \mathfrak{k}),\ U(\mathfrak{g})/U(\mathfrak{g})(\mathfrak{u} \cap \mathfrak{k})) \tag{5.13}$$

   I plan to describe this algebra, and look for a generalization of Theorem 4.3.2 describing its action on $H^i(\mathfrak{u} \cap \mathfrak{k}, V)$ for a $\mathfrak{g}$ module V. I expect that this action can be tied to the restriction map ρ.

5. The proof of Proposition 3.1.1 relied heavily on the collapse of the spectral sequence (3.3). The E_2-terms of this sequence were $E_2^{b,a} = \mathrm{Ext}^b_{\mathfrak{l}, L\cap K}(Z, H^a(\mathfrak{u}, X))$. Theorem 2.3.1 showed gives $H^a(\mathfrak{u}, X)$ as a sum of discrete series for L. These have no higher Ext groups between them, so the sequence only lives in the 0-th column.

I would like to generalize X to any irreducible representation of $\mathfrak{g}$ with a regular infinitesimal character. The Kazhdan-Lusztig conjectures (appearing in [6] and [14], and proved in [9] and [15]) give a precise description of $H^a(\mathfrak{u}, X)$. These certainly may have extensions between them, so the spectral sequence (3.3) is more complicated. I shall try adapting the techniques in my thesis to examine this scenario.

Appendix A

Step algebra and higher cohomology

A.1 $\mathbb{S}$ action on $H^w(\mathfrak{k}_+, X)$

Let us first write down the complexes that compute $H^i(\mathfrak{u} \cap \mathfrak{k}, X)$. We shall use them to analyze the action of $\mathcal{S}(\mathfrak{g}, \mathfrak{k})$ on $H^w(\mathfrak{k}_+, X)$. Most of what follows is taken from Chapter 4 of [7].

Given a Lie algebra $\mathfrak{g}$, the *Koszul resolution* is the complex $V_\bullet \to \mathbb{C}$ where $V_n = U(\mathfrak{g}) \otimes_{\mathbb{C}} \wedge^n \mathfrak{g}$. The differential is

$$\partial(u \otimes v_1 \wedge \ldots \wedge v_n) = \sum_{i=1}^{n} (-1)^{i+1} (uv_i \otimes v_1 \wedge \ldots \wedge \hat{v}_i \wedge \ldots \wedge v_n) + \tag{A.1}$$
$$+ \sum_{k<l} (-1)^{k+l} (u \otimes [v_k, v_l] \wedge v_1 \wedge \ldots \wedge \hat{v}_k \wedge \ldots \wedge \hat{v}_l \wedge \ldots \wedge v_n)$$

Given a $\mathfrak{g}$-module X, we can make a projective resolution $V_\bullet \otimes X \to X$ and an injective resolution $X \to \mathrm{Hom}_{\mathbb{C}}(V_\bullet, X)$.

Let us for a second take the compact Lie algebra $\mathfrak{k}$ and $X = E_\mu$. We shall use the above injective resolution for $\mathfrak{k}_+$: $X \to I_\bullet^K$ where $I_i^K = \mathrm{Hom}_{\mathbb{C}}(U(\mathfrak{k}_+) \otimes \wedge^i \mathfrak{k}_+, X)$.

Applying the functor $(\ )^{\mathfrak{k}_+}$ gives

$$(I_i^K)^{\mathfrak{k}_+} = \mathrm{Hom}_{U(\mathfrak{k}_+)}\left(U(\mathfrak{k}_+)\otimes \wedge^i\mathfrak{k}_+, X\right) = \mathrm{Hom}_{\mathbb{C}}(\wedge^i\mathfrak{k}_+, X) \tag{A.2}$$

. Let v be the highest weight vector of X, and $w \in W_K$ an element of length N.

Let $\{\alpha_i\}$ be the roots of $\mathfrak{k}_+$ and $\{e_i\}$ the corresponding root vectors. For convenience let $\Delta^K(w) = \{\alpha_1, \ldots, \alpha_N\}$. The proof of Kostant's Theorem found for example in Chapter 6 of [7] shows that the cohomology $H^w(\mathfrak{k}_+, X)$ can be represented by a function f defined as

$$f(e_{i_1}\wedge\ldots\wedge e_{i_N}) = \begin{cases} w^{-1}v & \text{for } [i_1,\ldots,i_N] = [1,\ldots,N] \\ 0 & \text{otherwise} \end{cases} \tag{A.3}$$

The element $w^{-1}v$ is the extremal weight vector of X, of weight $w(\mu)$, and the roots $\{\alpha_1, \ldots, \alpha_N\}$ are precisely those negative on $w(\mu)$. The weight of f is $w(\mu) - \sum_{i=1}^N \alpha_i = w(\mu+\rho_c)-\rho_c$, as we already know. The function f is unique up to scalars; no other function in $\mathrm{Hom}_{\mathbb{C}}(\wedge^i\mathfrak{k}_+, X)$ can possibly have the same weight.

Let X be our discrete series. To compute $H^w(\mathfrak{k}_+, X)$ we shall use the resolution $I_i^K = \mathrm{Hom}_{\mathbb{C}}(U(\mathfrak{k}_+)\otimes\wedge^i\mathfrak{k}_+, X)$. To track down the step algebra action, we will also need a $\mathfrak{g}$-resolution of X: $X \to I_\bullet$ where $I_i = \mathrm{Hom}_{\mathbb{C}}(U(\mathfrak{g})\otimes\wedge^i\mathfrak{g}, X)$. To compute the cohomology $H^N(\mathfrak{k}_+, X)$ we apply the functor $(\ )^{\mathfrak{k}_+} = \mathrm{Hom}_{U(\mathfrak{g})}(U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+, \_)$ to get

$$\begin{aligned} I_i^{\mathfrak{k}_+} &= \mathrm{Hom}_{U(\mathfrak{g})}\left(U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+, \mathrm{Hom}_{\mathbb{C}}(U(\mathfrak{g})\otimes\wedge^i\mathfrak{g}, X)\right) \\ &= \mathrm{Hom}_{U(\mathfrak{k}_+)}\left(U(\mathfrak{g})\otimes\wedge^i\mathfrak{g}, X\right) \\ &= \mathrm{Hom}_{\mathbb{C}}\left(\mathfrak{k}_+U(\mathfrak{g})\backslash U(\mathfrak{g})\otimes\wedge^i\mathfrak{g}, X\right) \end{aligned} \tag{A.4}$$

We have an inclusion $\iota : U(\mathfrak{k}_+)\otimes\wedge^N\mathfrak{k}_+ \to U(\mathfrak{g})\otimes\wedge^N\mathfrak{g}$ which gives a restriction

$$\mathrm{Hom}_{U(\mathfrak{k}_+)}(U(\mathfrak{g})\otimes\wedge^i\mathfrak{g}, X) \longrightarrow \mathrm{Hom}_{U(\mathfrak{k}_+)}(U(\mathfrak{k}_+)\otimes\wedge^i\mathfrak{k}_+, X) \tag{A.5}$$

The two complexes compute the same cohomology so the cohomology classes exactly correspond. Given $\bar{f} \in H^w(\mathfrak{k}_+, X)$ defined on $U(\mathfrak{k}_+) \otimes \wedge^N \mathfrak{k}_+$, we can extend it to a function f on $U(\mathfrak{g}) \otimes \wedge^N \mathfrak{g}$ arbitrarily, provided that we satisfy $f(\omega) = \bar{f}(\omega_K)$ whenever $d\omega = \iota(d\omega_K)$.

Recall that $\mathbb{S} = \mathrm{Hom}_{U(\mathfrak{g})}(U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+, U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+)$. An element $s \in \mathbb{S}$ acts on the first factor in $I_i^{\mathfrak{k}_+} = \mathrm{Hom}_{U(\mathfrak{g})}(U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+, \mathrm{Hom}_{\mathbb{C}}(U(\mathfrak{g}) \otimes \wedge^i \mathfrak{g}, X))$. In that notation, our f comes from the function

$$[1 \mapsto f] \in \mathrm{Hom}_{U(\mathfrak{g})}(U(\mathfrak{g})/U(\mathfrak{g})\mathfrak{k}_+, \mathrm{Hom}_{\mathbb{C}}(U(\mathfrak{g}) \otimes \wedge^i \mathfrak{g}, X)) \tag{A.6}$$

The action by $\mathbb{S}$ is $s[1 \mapsto f] = [s \mapsto f] = [1 \mapsto (sf)]$. We need to write (sf) in the form of Equation A.3.

Let $f \in \mathrm{Hom}_{\mathbb{C}}(U(\mathfrak{g}) \otimes \wedge^i \mathfrak{g}, X)$ be a representative for a class in $H^w(\mathfrak{k}_+, X)$; we have some freedom in defining it. An element $X \in \mathfrak{g}$ acts on $f \in \mathrm{Hom}_{\mathbb{C}}(U(\mathfrak{g}) \otimes \wedge^i \mathfrak{g}, X)$ by $(Xf)(\omega) = Xf(\omega) - f(X\omega)$. Let s be an elementary step. We know s preserves $H^w(\mathfrak{k}_+, X)$, so to know the class of $f \in H^N(\mathfrak{k}_+, X)$, it is enough to know $f(e_1 \wedge \ldots \wedge e_N)$, in the notation of Equation (A.3). Each summand of s is a product of elements in $\mathfrak{k}_-$, $\mathfrak{h}$ and $\mathfrak{p}$. Let us first see how those act on f.

Proposition A.1.1. *With the above setup, let $X \in \mathfrak{p}$ or $X \in \mathfrak{k}_-$. Then we can choose $f(X \otimes e_1 \wedge \ldots \wedge e_N) = 0$*

Proof. Let $\omega = X \otimes e_1 \wedge \ldots \wedge e_N$. If $X \in \mathfrak{p}$, then $d\omega$ is a sum of elements of the form $u \otimes e'_1 \wedge \ldots e'_{N-1}$ where $e'_i \in \mathfrak{k}_+$ and u is a product of elements in $\mathfrak{g}$ that includes an element of $\mathfrak{p}$. In other words $d\omega$ is not in the image of ι so in our choice of f as an extension of $\bar{f}$ we can set $f(\omega) = 0$.

Let $X \in \mathfrak{k}_-$, of weight $-\alpha$. The weight of $d\omega$ is $\sum_{i=1}^N \alpha_i - \alpha$. Recall that $\bar{f}$ is nonzero only on elements of type $\omega_K = u_K \otimes e_1 \wedge \ldots \wedge e_N$ where $u_k \in U(\mathfrak{k}_+)$. The weight of ω_K is $\sum_{i=1}^N \alpha_i + \delta$ where δ is in the span of $\Delta(\mathfrak{k}_+)$. If $\iota(d\omega_K) = d\omega$ then $-\alpha = \delta$ which is impossible. Again we are free to define $f(\omega) = 0$ ☐

The action of $\mathfrak{h}$ on f is the unique one we already know, which simply returns the weight ν of f. Let ρ_w be half-sum of roots in $\mathfrak{k}_+$ that are negative on C_w:

$$\rho_w = \frac{1}{2} \sum_{\alpha \in \Delta^K(w)} \alpha == \alpha_1 + \ldots \alpha_N \tag{A.7}$$

An element $h \in \mathfrak{h}$ acts on f by $(hf)(e_1 \wedge \ldots \wedge e_N) = (h - 2\rho_w(h))f((e_1 \wedge \ldots \wedge e_N)$. This is the last ingredient we needed to describe the action of s on f completely.

Proposition A.1.2. *Let $f \in H^w(\mathfrak{k}_+, X)$ given by $[e_1 \wedge \ldots \wedge e_N \mapsto wv]$ in the notation of Equation (A.3). Think of $s \in \mathcal{S}(\mathfrak{g}, \mathfrak{k})$ as a polynomial $s(h_1, \ldots, h_r)$ in $\mathfrak{h}$. The class of $(sf) \in H^w(\mathfrak{k}_+, X)$ is given by*

$$e_1 \wedge \ldots \wedge e_N \mapsto s\left(h_1 - 2\rho_w(h_1), \ldots, h_r - 2\rho_w(h_r)\right)(wv) \tag{A.8}$$

A.2 Proof of Proposition 4.3.4

For calculations in this section, we shall use a different description of the step algebra. In [16], Zhelobenko extended $\mathcal{S}(\mathfrak{g}, \mathfrak{k})$ to allow denominators in $\mathfrak{h}$, showed that the resulting algebra $\mathcal{S}(\mathfrak{g}, \mathfrak{k})'$ is an image of a certain projection operator, and proved that the denominators involved are nonzero. We shall give a rough sketch his approach here.

Let $R(\mathfrak{h})$ be the field of fractions of $U(\mathfrak{h})$, and define $U'(\mathfrak{k}) = U(\mathfrak{k}) \otimes_{U(\mathfrak{h})} R(\mathfrak{h})$. We then let $F_\mu(\mathfrak{k})$ be the space of formal power series in $U'(\mathfrak{k})$ of weight μ, and define the algebra $F(\mathfrak{k}) = \oplus F_\mu(\mathfrak{k})$. We now define an element $p_\alpha \in F(\mathfrak{k})$ as

$$p_\alpha = \sum_{n=0}^{\infty} \frac{(-1)^n}{n!} f_{\alpha,n}^{-1} e_{-\alpha}^n e_\alpha^n \tag{A.9}$$

where

$$f_{\alpha,n} = \prod_{i=1}^{n-1} (h_\alpha + \rho_c(h_\alpha) + i) \tag{A.10}$$

We put a normal ordering $\{\alpha_1, \ldots, \alpha_m\}$ on the positive roots in $\mathfrak{k}$; this means that $i < j < k$ whenever $\alpha_j = \alpha_i + \alpha_k$. Let $p \in F(\mathfrak{k})$ be the product $p = p_{\alpha_1} \cdots p_{\alpha_n}$. Define $U'(\mathfrak{g}) = U(\mathfrak{g}) \otimes_{U(\mathfrak{h})} R(\mathfrak{h})$. The operator p projects $U'(\mathfrak{g})/U'(\mathfrak{g})\mathfrak{k}_+$ onto the extended step algebra $\mathcal{S}(\mathfrak{g},\mathfrak{k})' = \mathcal{S}(\mathfrak{g},\mathfrak{k}) \otimes_{U(\mathfrak{h})} R(\mathfrak{h})$. Note that for $v \in U'(\mathfrak{g})/U'(\mathfrak{g})\mathfrak{k}_+$ and $\alpha \in \Delta_+^K$, $e_\alpha v = [e_\alpha, v]$.

The extended step algebra is spanned by elements $s_i' = pt_i$. The element s_i' differs by a denominator in $\mathfrak{h}$ from Mickelsson's step s_i: in the notation of Equation (4.14), $s_i' = s_i p_i^{-1}$. By Proposition 4.2.1 we are allowed to apply these to $H^w(\mathfrak{k}_+, X)$, for $p(\nu) \neq 0$ whenever ν and $\nu + \beta_i$ are both in C_w.

Let us use this description to prove Proposition 4.3.4. We restate it here for convenience:

Proposition 4.3.4. Let v_μ be an element of weight μ in $H^0(\mathfrak{k}_+, X)$, corresponding to $f \in H^w(\mathfrak{k}_+, X)$. Let $s_i \in \mathcal{S}(\mathfrak{g},\mathfrak{k})$ be one of the elementary steps and let $\beta_j = w^{-1}\beta_i$. Suppose that $v_{\mu-\beta_i} = s_i v$ is a nonzero element of $H^0(\mathfrak{k}_+, X)$. Then $(s_j f)$ is a nonzero element of $H^w(\mathfrak{k}_+, X)$ corresponding to $v_{\mu-\beta_i}$.

Proof. The proposition is stated for the steps in the Mickelsson's algebra. In this proof we shall use the extended steps s' instead; as noted before the denominators involved are nonzero so this makes no difference.

We use notation of Equation (A.3) for f: it sends $e_1 \wedge \ldots \wedge e_N$ to $w^{-1}v$, and is zero on the other monomials in $\wedge^\bullet \mathfrak{k}_+$. Let $Y_\mu = U(\mathfrak{k}) \cdot v \subset X$. By assumption, the product $\mathfrak{p}Y \in X$ contains the representation $Y_{\mu-\beta_i} = U(\mathfrak{k})v_{\mu-\beta_i}$. Therefore $t_i v = cv_{\mu-\beta_i} + \sum_{j\neq i} u_{\mu-\beta_j}$ where $c \neq 0$ and u_δ is an element in a K-type E_δ of X. The projection p sends all the u_δ to zero, leaving $pt_i v = s_i' v = v_{\mu-\beta_i}$.

Proposition A.1.2 tells us that $(t_j f)(e_1 \wedge \ldots \wedge e_N) = t_j w^{-1} v$. Again by assumption we know that $t_j w^{-1} v = C' w v_{\mu-\beta_i} + \sum_{j\neq i} z_{\mu-\beta_j}$ where $C' \neq 0$ and z_δ lies in a K-type E_δ. The element $(wv_{\mu-\beta_i})$ is an extremal vector of weight $w\mu + \beta_j = w(\mu+\beta_i)$. For all $j \neq i$, the class $[(e_1 \wedge \ldots \wedge e_N) \mapsto z_{\mu-\beta_j}]$ is zero in cohomology; it's only the coefficient C of $(wv_{\mu-\beta_i})$ that matters in where $(pt_j f)$ maps $e_1 \wedge \ldots \wedge e_N$.

Recall that $p = p_{\alpha_1} \cdots p_{\alpha_n}$ in notation of Equation (A.9). The first observation to make is that if $\alpha \notin \Delta^K(w)$, then $e_\alpha(wv_{\mu-\beta_i}) = 0$. It follows that the only term from p_α that has an effect on the C is the constant term, which is equal to 1.

Let $\alpha \in \Delta^K(w)$, and denote $x = e_\alpha$, $y = e_{-\alpha}$, and $h = h_\alpha$; these span a subalgebra

$\mathfrak{su}_2 \subset \mathfrak{k}$. We already know that x acts on t by ad_x. Now $(wv_{\mu-\beta_i})$ is an extremal vector so $y(wv_{\mu-\beta_i}) = 0$; in other words, left multiplication by y is the same as ad_y as far as C is concerned.

We shall now make use of the fact that the Lie algebra $\mathfrak{g}$ is simply laced, which means that $\mathfrak{p}$ splits into one and two-dimensional representations of our $\mathfrak{su}_2$. If $t = t_j$ is a highest weight vector with respect to $x \in \mathfrak{su}_2$ then $[x,t] = 0$ and again only the constant term of p_α matters. The other possibility is that t is the lowest-weight vector in a two-dimensional representation, of weight -1 with respect to our $\mathfrak{su}_2$. This means that $[x,[x,t]] = 0$ so we need to calculate the first two terms of p_α on t.

To calculate $\mathrm{ad}_{yx}t$, note that $[y,t] = 0$. We have the standard commutation relation $xy - yx = h$, so $\mathrm{ad}_{yx}t = \mathrm{ad}_{-h}t = t$. Recall that

$$p_\alpha = 1 - [h + \rho_c(h) + 1]^{-1} yx \tag{A.11}$$

The action of $\mathfrak{h}$ returns the weight of its target, which is $w\mu - 2\rho_w + \beta_j$. We also know that it's equal to $w(\mu + \rho_c) - \rho_c + \beta_j$, which means that $-2\rho_w = w\rho_c - \rho_c$. We also know that $\beta_j(h) = -1$, so finally we can compute that $[h + \rho_c(h) + 1]yx$ acts by

$$w\mu(h) + \beta_j(h) - 2\rho_w(h) + \rho_c(h) + 1 = [w\mu](h) + [w\rho_c](h) \tag{A.12}$$

The root α is negative on C_w so $[w\rho_c](h) \leq -1$ and $[w\mu](h) \leq 0$, which means that $[h + \rho_c(h) + 1]^{-1}$ in Equation (A.11) returns a negative number $-m_\alpha$. Altogether, we find that the coefficient of $(wv_{\mu-\beta_i})$ in $(pt_j f)(e_1 \wedge \ldots \wedge e_N)$ is

$$C = C' \prod_{\alpha \in \Delta^K(w)\ h_\alpha(t_j) = -1} (1 + m_\alpha) \neq 0 \tag{A.13}$$

□

Bibliography

[1] W. Fulton and J. Harris, *Representation theory*, Graduate Texts in Mathematics, vol. 129, Springer-Verlag, New York, 1991.

[2] Harish-Chandra, *Representations of a semisimple Lie group on a Banach space. I*, Trans. Amer. Math. Soc. **75** (1953), 185–243.

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

[4] G. Hochschild and J.-P. Serre, *Cohomology of Lie algebras*, Ann. of Math. (2) **57** (1953), 591–603.

[5] A. van den Hombergh, *A note on Mickelsson's step algebra*, Nederl. Akad. Wetensch. Proc. Ser. A **78**=Indag. Math. **37** (1975), 42–47.

[6] D. Kazhdan and G. Lusztig, *Representations of Coxeter groups and Hecke algebras*, Invent. Math. **53** (1979), 165–184.

[7] A. W. Knapp, *Lie groups beyond an introduction*, 2nd ed., Progress in Mathematics, vol. 140, Birkhäuser Boston Inc., Boston, MA, 2002.

[8] A. W. Knapp and D. A. Vogan Jr, *Cohomological induction and unitary representations*, Princeton Mathematical Series, vol. 45, Princeton University Press, Princeton, NJ, 1995.

[9] G. Lusztig and D. A. . Vogan Jr., *Singularities of closures of K-orbits on flag manifolds*, Invent. Math. **71** (1983), 365–379.

[10] J. Mickelsson, *Step algebras of semi-simple subalgebras of Lie algebras*, Rep. Mathematical Phys. **4** (1973), 307–318.

[11] W. Schmid, *Some properties of square-integrable representations of semisimple Lie groups*, Ann. of Math. (2) **102** (1975), 535–564.

[12] M. R. Sepanski, *Boundaries of K-types, restriction of cohomology, and the multiplicity free case*, Compositio Math. **106** (1997), 31–41.

[13] ———, *K-types of* $\mathrm{SU}(1,n)$ *representations and restriction of cohomology*, Pacific J. Math. **192** (2000), 385–398.

[14] D. A. Vogan Jr., *Irreducible characters of semisimple Lie groups. II. The Kazhdan-Lusztig conjectures*, Duke Math. J. **46** (1979), 805–859.

[15] ———, *Irreducible characters of semisimple Lie groups. III. Proof of Kazhdan-Lusztig conjecture in the integral case*, Invent. Math. **71** (1983), 381–417.

[16] D. P. Zhelobenko, *Extremal projectors and generalized Mickelsson algebras on reductive Lie algebras*, Izv. Akad. Nauk SSSR Ser. Mat. **52** (1988), 758–773, 895. (Russian)