BRANCHING TO A MAXIMAL COMPACT SUBGROUP

DAVID A. VOGAN, JR.

Abstract. Suppose $G = \mathbf{G}(\mathbb{R})$ is the group of real points of a complex connected reductive algebraic group, and K is a maximal compact subgroup of G. The classical branching problem in this setting is to determine the restriction to K of a standard representation of G. Implicit in the branching problem are the problems of parametrizing both irreducible representations of K and standard representations of G. We address these three problems, looking for answers amenable to computer implementation.

Contents

Date: June 14, 2007.
Supported in part by NSF grant DMS-0554278.

1. General introduction

Suppose G is a real reductive Lie group and K is a maximal compact subgroup. (We will make our hypotheses more precise and explicit beginning in Definition 2.3 below.) Harish-Chandra proved that any (nice) irreducible representation of G decomposes on restriction to K into irreducible components having finite multiplicity. Our goal is to describe an algorithm for computing those multiplicities: that is, for solving the "branching problem" from G to K.

The traditional wisdom is that this branching problem was solved with the proof of the Blattner Conjecture by Hecht and Schmid more than thirty years ago. What the Blattner formula actually accomplishes is a reduction of the branching problem from G to K to various branching problems from K to subgroups $S \subset K$. (The subgroups S that arise are the intersections of K with parabolic subgroups of G.) For these compact group branching problems, one can hope to use the Weyl character formula in one form or another. Often this works very well; there are many widely available software packages for doing such calculations. There are two fundamental difficulties, however. First, if K is disconnected, then applying the theory of highest weights and the Weyl character formula to K requires some care. Second, if S is disconnected, these problems are multiplied.

A good example to keep in mind is $G = GL_n(\mathbb{R})$, the group of invertible $n \times n$ matrices with real entries. The maximal compact subgroup is the orthogonal group $K = O_n$. The diagonal matrices in K form a subgroup $S = O_1^n$ isomorphic to $\{\pm 1\}^n$. A principal series representation of G is determined by (among other things) a character δ of S. If μ is an irreducible representation of K, then the multiplicity of μ in this principal series is equal to the multiplicity of δ in $\mu|_S$. That is, to solve the branching problem for GL_n (even for principal series) we must solve the branching problem from O_n to O_1^n. This latter problem is "elementary" but painful, because the (abelian) group S is not contained in a maximal torus of K. In this special case there are more or less classical solutions available, but finding generalizations sufficient to deal cleanly with general branching from G to K looks very difficult.

We will approach the problem indirectly. Recall first of all (from [8], for example) that each irreducible representation π of G has a (very small) finite set

$$A(\pi) \subset \text{irreducible representations of } K \tag{1.1}$$

of *lowest K-types*. These are the representations that actually appear in the restriction of π to K, and which have minimal norm with respect to that requirement. (The precise definition of the norm of an irreducible representation of K is a little complicated, but it is approximately the length of the highest weight.) It turns out that the cardinality of $A(\pi)$ is always a power of two; the power is bounded by the rank of K.

For our purposes, the most convenient way to parametrize irreducible representations of K is by listing certain irreducible representations of G in which they appear as lowest K-types. Here is a quick-and-dirty version of the result we need (taken from [10]); a more complete and careful statement is in Theorem 11.9 below.

Theorem 1.2. *Suppose G is a real reductive group with maximal compact subgroup K and Cartan involution θ. Then there are natural bijections among the following three sets.*

(1) *Tempered irreducible representations π of G, having real infinitesimal character.*
(2) *Irreducible representations τ of K.*
(3) *Discrete final limit parameters Φ (Definition 11.2) attached to θ-stable Cartan subgroups H of G, modulo conjugation by K.*

The bijection from (1) *to* (2) *sends π to the unique lowest K-type $\tau(\pi)$ of π. That from* (3) *to* (1) *is the Knapp-Zuckerman parametrization of irreducible tempered representations: if Φ is a discrete final limit parameter, we write $\pi(\Phi)$ for the corresponding tempered representation, and $\tau(\Phi)$ for its unique lowest K-type.*

From a theoretical point of view, what is most interesting about this theorem is the bijection between (1) and (2). For our algorithmic purposes, what is most important is the bijection $\Phi \mapsto \tau(\Phi)$ from (3) (a set which is explicitly computable) to the set of representations of K we wish to study.

To a first approximation, a discrete final limit parameter attached to the θ-stable Cartan subgroup H is simply a character of the compact group $H \cap K$. We are therefore parametrizing irreducible representations of K by characters of the compact parts of Cartan subgroups of G (up to conjugation). This is in the spirit of the Cartan-Weyl parametrization of irreducible representations of a compact connected Lie group K by characters of a maximal torus T (up to conjugation); indeed we will see in Example 1.5 that the Cartan-Weyl result may be regarded as a special case of Theorem 1.2.

Example 1.3. Suppose $G = SL_2(\mathbb{R})$. Recall that $K = SO_2$, and that the irreducible representations of K are the one-dimensional characters

$$\tau_m \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} = e^{im\theta}. \tag{1.4}$$

Here is the list of the tempered irreducible representations of G having real infinitesimal character.

(1) For each strictly positive integer $n > 0$ there is a holomorphic discrete series representation $\pi(n)$ with Harish-Chandra parameter n. Its lowest K-type is τ_{n+1}, and its full restriction to K is

$$\pi(n)|_K = \tau_{n+1} + \tau_{n+3} + \tau_{n+5} + \cdots$$

(2) For each strictly positive integer $n > 0$ there is an antiholomorphic discrete series representation $\pi(-n)$ with Harish-Chandra parameter $-n$. Its lowest K-type is τ_{-n-1}, and its full restriction to K is

$$\pi(-n)|_K = \tau_{-n-1} + \tau_{-n-3} + \tau_{-n-5} + \cdots$$

(3) There is a holomorphic limit of discrete series representation $\pi(0^+)$. Its lowest K-type is τ_1, and its full restriction to K is

$$\pi(0^+)|_K = \tau_1 + \tau_3 + \tau_5 + \cdots$$

(4) There is an antiholomorphic limit of discrete series representation $\pi(0^-)$. Its lowest K-type is τ_1, and its full restriction to K is

$$\pi(0^-)|_K = \tau_{-1} + \tau_{-3} + \tau_{-5} + \cdots$$

(5) There is the spherical principal series representation π_{sph} with continuous parameter zero. Its lowest K-type is τ_0, and its full restriction to K is

$$\pi_{sph}|_K = \tau_0 + \tau_2 + \tau_{-2} + \tau_4 + \tau_{-4} + \cdots$$

The bijection of Theorem 1.2 identifies each non-trivial character τ_m as the lowest K-type of a (limit of) discrete series representation, holomorphic if $m > 0$ and antiholomorphic if $m < 0$. The trivial character τ_0 is the lowest K-type of the spherical principal series representation.

In terms of part (3) of Theorem 1.2, the discrete series and limits of discrete series correspond to the compact Cartan subgroup SO_2, and the spherical principal series to the split Cartan subgroup of diagonal matrices.

Example 1.5. Suppose that G is connected complex reductive algebraic group. A maximal compact subgroup K of G is the same thing as a compact form of G; K can be any compact connected Lie group. If T is a maximal torus in K, then the centralizer H of T in G is a θ-stable Cartan subgroup of G, and $H \cap K = T$. The inclusion of K in G provides an identification

$$W(K,T) = N_K(T)/T \simeq N_G(H)/H = W(G,H). \tag{1.6a}$$

There is a unique θ-stable complement A for T in H, so that $H \simeq T \times A$. The group A is isomorphic via the exponential map to its Lie algebra.

Fix a Borel subgroup $B = HN$ of G. The Iwasawa decomposition is $G = KAN$. Any character ξ of H extends uniquely to B, and so gives rise to a principal series representation

$$\pi_B(\xi) = \mathrm{Ind}_B^G(\xi). \tag{1.6b}$$

As a consequence of the Iwasawa decomposition,

$$\pi_B(\xi)|_K = \mathrm{Ind}_T^K(\xi|_T) = \sum_{\tau \text{ irr of } K} (\text{mult of } \xi|_T \text{ in } \tau)\tau. \tag{1.6c}$$

If ξ is unitary, then $\pi_B(\xi)$ is tempered, irreducible, and (up to equivalence) independent of the choice of B. It follows that

$$\pi_B(w \cdot \xi) \simeq \pi_B(\xi) \qquad \xi \text{ unitary, } w \in W(G,H). \tag{1.6d}$$

In this way the tempered irreducible representations of G are identified with Weyl group orbits of unitary characters of H. We may drop the subscript B since that choice does not matter.

The tempered representation $\pi(\xi)$ has real infinitesimal character if and only if $\xi|A$ is trivial. Characters of H trivial on A are the same as characters of the compact torus T. We have therefore described a bijection

$$\{\text{temp irr of } G, \text{ real infl char}\} \leftrightarrow \{\text{chars of } T \bmod W(K,T)\}\,; \tag{1.6e}$$

this is the bijection between (1) and (3) in Theorem 1.2. It follows from (1.6c) that the lowest K-type of $\pi(\xi)$ is the smallest representation of K containing the weight $\xi|T$. This is the representation $\tau(\xi|_T)$ of extremal weight ξ_T. The bijection between (2) and (3) in Theorem 1.2 is therefore precisely the Cartan-Weyl parametrization by Weyl group orbits of extremal weights.

We can now give notation for the branching problem we wish to solve. Given an irreducible representation τ of K and a tempered irreducible representation π of G with real infinitesimal character, write

$$m(\tau,\pi) = \text{multiplicity of } \tau \text{ in } \pi|_K \tag{1.7}$$
$$= \dim \operatorname{Hom}_K(\tau,\pi), \tag{1.8}$$

a non-negative integer; this is the kind of multiplicity we wish to compute. (Originally we were interested in replacing π by an arbitrary standard representation π'. But it is easy explicitly to write $\pi'|_K$ as a (small) finite sum of $\pi_j|_K$, with each π_j tempered with real infinitesimal character. (One needs a single term π_j for each lowest K-type of π'.) The special case considered in (1.7) is therefore enough to solve the general problem of branching to K.)

If G is a complex group, then the discussion in Example 1.5 shows that $m(\tau,\pi(\xi))$ is the multiplicity of the T-weight $\xi|_T$ in the K-representation τ. The branching law we are studying here therefore generalizes the classical problem of finding weight multiplicities for representations of a compact group.

We want to compute the matrix $m(\tau,\pi)$. What is interesting now is that Theorem 1.2 identifies the index set for the rows of the matrix m with the index set for the columns: we may regard m as a square matrix, indexed by any of the three sets in Theorem 1.2. The definition of lowest K type implies that m is upper triangular (with respect to the ordering of representations of K used to define lowest). The fact that lowest K-types have multiplicity one implies that m has ones on the diagonal; and of course the entries of m are integers.

As a formal consequence of these observations, we find that the matrix m is necessarily invertible. A little more explicitly,

Proposition 1.9. *The multiplicity matrix m of 1.7 has a two-sided inverse M. That is, for every irreducible representation τ of K and every tempered* *irreducible representation π of G with real infinitesimal character, there is an integer $M(\pi,\tau)$ having the following properties:*

(1) *$M(\pi,\tau) = 0$ unless τ is less than or equal to the lowest K-type of π. If equality holds, then $M(\pi,\tau) = 1$.*
(2) *If τ and τ' are two irreducible representations of K, then*

$$\sum_{\pi \text{ temp real infl char}} m(\tau,\pi)M(\pi,\tau') = \delta_{\tau,\tau'}.$$

Here the summands can be nonzero only for tempered representations π whose lowest K-type is less than or equal to τ, a finite set.
(3) *If π and π' are two irreducible tempered representations of real infinitesimal character, then*

$$\sum_{\tau \text{ irr of } K} M(\pi,\tau)m(\tau,\pi') = \delta_{\pi,\pi'}.$$

Here the summands can be nonzero only for τ less than or equal to the lowest K-type of π', a finite set.

The matrix M expresses each irreducible representation of K as an integer combination of tempered representations of G:

$$\tau = \sum_{\pi \text{ temp real infl char}} M(\pi,\tau)\pi.$$

This sum extends over tempered representations π whose lowest K-type is greater than or equal to τ, and so may in principle be infinite.

The final assertion of the proposition is just a reformulation of (2).

Inverting an upper triangular matrix with diagonal entries equal to 1 is easy; so either of the matrices m and M is readily calculated from the other. Since both matrices are infinite, this statement requires a little care to make into an algorithm. In practice we can consider only some finite upper left corner of the matrix, corresponding to the finitely many representations of K less than or equal to some bound, and the tempered representations with lowest K-type in that finite set.

Example 1.10. We return to the example of $SL_2(\mathbb{R})$. In the preceding example we wrote down the multiplicity matrix m: for example,

$$m(\tau_n,\pi_{sph}) = \begin{cases} 1 & n \text{ even} \\ 0 & n \text{ odd.} \end{cases}$$

Using the interpretation of the inverse matrix M given at the end of Proposition 1.9, it is a simple matter to compute the inverse matrix M:

$$\begin{aligned} \tau_0 &= \pi_{sph} - \pi(1) - \pi(-1) \\ \tau_1 &= \pi(0^+) - \pi(2) \\ \tau_{-1} &= \pi(0^-) - \pi(-2) \\ \tau_n &= \pi(n-1) - \pi(n+1) \qquad (n \geq 1) \\ \tau_{-n} &= \pi(-n+1) - \pi(-n-1) \qquad (n \geq 1). \end{aligned}$$

Explicitly, this means that for example

$$M(\pi(j), \tau_n) = \begin{cases} 1 & j = n-1 \\ -1 & j = n+1 \\ 0 & \text{otherwise} \end{cases}$$

for any $n \geq 2$.

The central observation we wish to make about Examples 1.3 and 1.10 is that the formulas in the second example (for the inverse matrix M) are much simpler. What we will do in this paper (essentially in Theorem 16.6) is give a closed formula for M in the general case. The formula is due to Zuckerman in the case of the trivial representation of K; the generalization here is straightforward. In order to compute an explicit branching law, one can then invert (some finite upper left corner of) the upper triangular matrix M.

Implementing this branching law within Fokko du Cloux's `atlas` computer program is a current project of the NSF Focused Research Group "Atlas of Lie groups and representations." I have therefore included in this paper comments about how the algorithms described are related to things that the `atlas` software already computes.

Here is what the formula for M looks like in the case of a complex group.

Example 1.11. Suppose that G is connected complex reductive algebraic group; use the notation of Example 1.5 above. Fix a set $\Delta^+(\mathfrak{k}, T)$ of positive roots for T in $\mathfrak{k}$; these are non-trivial characters of T. Suppose τ is an irreducible representation of K, of highest weight ξ (a character of T). We want to write τ as an integer combination of restrictions to K of tempered representations of G with real infinitesimal character; equivalently, as an integer combination of representations of K induced from characters of T. This is accomplished by the Weyl character formula for τ. We can give the answer in two forms; the equivalence of the two is the Weyl denominator formula. First, we have

$$\tau = \sum_{w \in W(K,T)} \operatorname{sgn}(w) \pi(\xi + (\rho - w\rho)).$$

Here ρ is half the sum of the positive roots of T in $\mathfrak{k}$. This formula appears to have $|W|$ terms, but some of the weights $\xi + (\rho - w\rho)$ may be conjugate by W (so there may be fewer teEmacsrms).

The second formula is

$$\tau = \sum_{S \subset \Delta^+(\mathfrak{k},T)} (-1)^{|S|} \pi(\xi + 2\rho(S)).$$

Here $2\rho(S)$ is the sum of the positive roots in S. This formula appears to have $2^{|\Delta^+|}$ terms, but there is cancellation (even before the question of W conjugacy is considered). After cancellation is taken into account, the second formula reduces precisely to the first.

2. Technical introduction

For a Lie group H we will write H_0 for the identity component, $\mathfrak{h}_0$ for the real Lie algebra, and $\mathfrak{h} = \mathfrak{h}_0 \otimes_{\mathbb{R}} \mathbb{C}$ for its complexification. If $\mathbf{H}$ is a complex algebraic group, we sometimes write simply $\mathbf{H}$ instead of $\mathbf{H}(\mathbb{C})$. When $\mathbf{H}$ is defined over $\mathbb{R}$, the group of real points will be written

$$H = \mathbf{H}(\mathbb{R}). \tag{2.1}$$

The notation allows for various kinds of ambiguity, since for example the same real Lie group H may appear as the group of real points of several distinct algebraic groups $\mathbf{H}$.

We write a subscript e for the identity component functor on complex algebraic groups. The group of complex points of the identity component is the classical identity component of the group of complex points:

$$\mathbf{H}_e(\mathbb{C}) = [\mathbf{H}(\mathbb{C})]_0. \tag{2.2a}$$

The group of real points of a connected algebraic group may however be disconnected, so we have in general only an inclusion

$$\mathbf{H}_e(\mathbb{R}) \supset [\mathbf{H}(\mathbb{R})]_0; \tag{2.2b}$$

Briefly (but perhaps less illuminatingly)

$$H_e \supset H_0. \tag{2.2c}$$

If $\mathbf{H}(\mathbb{R})$ is compact, then equality holds:

$$\mathbf{H}_e(\mathbb{R}) = [\mathbf{H}(\mathbb{R})]_0 \qquad (\mathbf{H}(\mathbb{R}) \text{ compact}). \tag{2.2d}$$

That is,

$$H_e = H_0 \qquad (H \text{ compact}). \tag{2.2e}$$

Definition 2.3. We fix once and for all a complex connected reductive algebraic group $\mathbf{G}$, together with a *real form*; that is, a complex conjugate-linear involutive automorphism

$$\sigma\colon \mathbf{G}(\mathbb{C}) \to \mathbf{G}(\mathbb{C}).$$

(When this is all translated into the `atlas` setting with several real forms, it will be convenient to define a *strong real form* to be a particular choice of representative for σ in an appropriate extended group containing $\mathbf{G}(\mathbb{C})$ as a subgroup of index two. I am avoiding that by speaking always of a single real form.) The corresponding *group of real points* is

$$\begin{aligned}\mathbf{G}(\mathbb{R},\sigma) = \mathbf{G}(\mathbb{R}) = G(\sigma) = G&\\ = \mathbf{G}(\mathbb{C})^\sigma;&\end{aligned}$$

here the second line defines the notation introduced in the first.

Definition 2.4. A *Cartan involution* for the real group G is an algebraic involutive automorphism θ of $\mathbf{G}$ subject to

1. The automorphisms σ and θ commute: $\sigma\theta = \theta\sigma$.
2. The composite (conjugate-linear) involutive automorphism $\sigma\theta$ has as fixed points a compact real form U of $\mathbf{G}$.

The second requirement is equivalent to

(2′) The group of fixed points of θ on G is a maximal compact subgroup K of G:

$$G^\theta = K.$$

Under these conditions, the complexification of the compact Lie group K is equal to $\mathbf{G}^\theta$:

$$\mathbf{K} = \mathbf{G}^\theta.$$

The action of θ on the complex Lie algebra $\mathfrak{g}$ defines an eigenspace decomposition

$$\mathfrak{g} = \mathfrak{k} + \mathfrak{s}, \tag{2.5}$$

with $\mathfrak{s}$ the -1 eigenspace of θ. This decomposition is inherited by any θ-stable real or complex subspace of $\mathfrak{g}$. The letter $\mathfrak{s}$ is chosen to suggest "symmetric," since when G is $GL_n(\mathbb{R})$ we can take θ to be negative transpose on the Lie algebra.

Almost all the real Lie groups here are going to be groups of real points of complex algebraic groups. In general an identification of a real Lie group with such a group of real points need not be unique even if it exists. But there is one very important special case when the algebraic group *is* unique. Any compact (possibly disconnected) Lie group S is a compact real form of a canonically defined complex reductive algebraic group $\mathbf{S}$. The complex algebra of regular functions on $\mathbf{S}$ is the complex algebra of matrix coefficients of finite-dimensional representations of S (and so $\mathbf{S}$ is by definition the spectrum of this ring).

This explicit description of $\mathbf{S}$ shows that if S is a closed subgroup of another compact Lie group K, then $\mathbf{S}$ is canonically an algebraic subgroup of $\mathbf{K}$ defined over $\mathbb{R}$.

The complexification of the unitary group U_n may be naturally identified with the complex general linear group $\mathbf{GL}_n(\mathbb{C})$. (Another dangerous bend here: the corresponding real form of $\mathbf{GL}_n(\mathbb{C})$ is given by inverse conjugate transpose. This is not the standard real form, given by complex conjugation, which is (one of the reasons) why U_n is not the same as $GL_n(\mathbb{R})$.) Pursuing this a bit shows that a continuous morphism from a compact group S to U_n is the restriction of a unique algebraic morphism of $\mathbf{S}$ into $\mathbf{GL}_n$. That is, continuous representations of the compact Lie group S may be identified with algebraic representations of $\mathbf{S}$.

The involutive automorphism θ is unique up to conjugation by G. Again in the setting of the atlas software we will keep a *strong Cartan involution*, a choice of representative for θ in an extended group for $\mathbf{G}$. (This extended group is not the same as the one containing σ, since that one contains conjugate-linear automorphisms).

Here is some general notation for representations.

Definition 2.6. Suppose $\mathbf{H}$ is a complex algebraic group. The set of equivalence classes of irreducible (necessarily finite-dimensional) algebraic representations of $\mathbf{H}$ is written

$$\Pi^*_{alg}(\mathbf{H}) = \Pi^*(\mathbf{H}),$$

the *algebraic dual* of $\mathbf{H}$. If $\mathbf{H}$ is a complex torus, then these representations are all one-dimensional, providing a canonical identification

$$\Pi^*(\mathbf{H}) \simeq X^*(\mathbf{H})$$

(the term on the right being the lattice of characters of $\mathbf{H}$).

Suppose $\mathbf{H}$ is reductive and defined over $\mathbb{R}$, and that H is compact. According to the remarks after Definition 2.4, restriction to H provides a canonical identification

$$\Pi^*(\mathbf{H}) \simeq \widehat{H},$$

with $\widehat{H}$ denoting the set of equivalence classes of continuous irreducible representations of H. Because of this identification, we will use the more suggestive notation

$$\Pi^*_{adm}(H) = \Pi^*(H)$$

for $\widehat{H}$, calling it the *admissible dual* of H. Then the canonical identification looks like

$$\Pi^*_{alg}(\mathbf{H}) \simeq \Pi^*_{adm}(H) \qquad (H \text{ compact}).$$

The term “continuous dual” makes more sense than “admissible dual” for H compact, but we have chosen the latter to fit with the following generalization.

Suppose that $\mathbf{H}$ is reductive and defined over $\mathbb{R}$, and that L is a maximal compact subgroup of H. The *admissible dual* of H is

$$\Pi^*_{adm}(H) = \Pi^*(H) = \text{equivalence classes of irreducible } (\mathfrak{h}, L)\text{-modules};$$

this set is often written $\widehat{H}$.

If L' is another maximal compact subgroup of H, then we know that L' is conjugate to L by $\mathrm{Ad}(h)$, for some element $h \in H$. Twisting by h carries (irreducible) $(\mathfrak{h}, L)$-modules to (irreducible) $(\mathfrak{h}, L')$-modules, and the correspondence of irreducibles is independent of the choice of h carrying L to L'. For this reason we can omit the dependence of $\Pi^*(H)$ on L.

Our goal is to parametrize $\Pi^*(G)$ and $\Pi^*(K)$ (that is, $\widehat{G}$ and $\widehat{K}$) in compatible and computer-friendly ways. Such parametrizations involve a variety of different "ρ-shifts." The most fundamental of these is that Harish-Chandra's parameter for a discrete series representation is not the differential of a character of the compact Cartan subgroup, but rather such a differential shifted by a half-sum of positive roots. There is a convenient formalism for keeping track of these shifts, that of "ρ-double covers" of maximal tori. But we have avoided it in these notes.

3. Highest weights for K

The phrase "maximal torus" can refer to a product of circles in a compact Lie group; or to a product of copies of $\mathbb{C}^\times$ in a complex reductive algebraic group; or to the group of real points of a complex maximal torus (which is isomorphic to a product of circles, $\mathbb{R}^\times$, and $\mathbb{C}^\times$). I will use all three meanings, sometimes with an adjective "compact" or "complex" or "real" to clarify.

Definition 3.1. Choose a compact maximal torus

$$T_{f,0} \subset K_0;$$

that is, a maximal connected abelian subgroup. This torus is unique up to conjugation by K_0. Then $T_{f,0}$ is a product of circles, and its complexification may be identified canonically with a complex maximal torus

$$\mathbf{T}_{f,0} \subset \mathbf{K}_e$$

The centralizer

$$\mathbf{H}_f = Z_{\mathbf{G}}(\mathbf{T}_{f,0})$$

is a θ-stable maximal torus in $\mathbf{G}$ (this requires proof, but is not difficult) that is defined over $\mathbb{R}$. The group $H_f = \mathbf{H}_f(\mathbb{R})$ is a real maximal torus of G, called a *fundamental Cartan subgroup*; it is the centralizer in G of $T_{f,0}$, and is unique up to conjugation by K_0.

Define

$$\mathbf{T}_f = \mathbf{H}_f^\theta = \mathbf{H}_f \cap \mathbf{K} = Z_{\mathbf{K}}(\mathbf{T}_{f,0}).$$

The group of real points is

$$T_f = H_f^\theta = H_f \cap K = Z_K(T_{f,0}),$$

a compact abelian subgroup of $K(\mathbb{R})$ containing the maximal torus $T_{f,0}$. In fact it is easy to see that $T_f(\mathbb{R})$ is a maximal abelian subgroup of $K(\mathbb{R})$ (and similarly over $\mathbb{C}$); it might reasonably be called a *maximally toroidal maximal abelian subgroup*, but I will follow [9] and call it a *small Cartan subgroup*.

The definition of T_f as $Z_K(T_{f,0})$ is internal to K, and so makes sense for any compact Lie group. In that generality T_f is a compact Lie group with identity component the compact torus $T_{f,0}$; but T_f need not be abelian. If we define

$$K^\sharp = \{k \in K \mid \mathrm{Ad}(k) \text{ is an inner automorphism of } \mathfrak{k}\},$$

then

$$K^\sharp = T_f K_0, \qquad K^\sharp/T_f \simeq K_0/T_{f,0};$$

all of these equalities remain true for complex points.

Clearly $K^\sharp$ is a normal subgroup of K, of finite index. We may therefore define the *R-group of K* as

$$R(K) = K/K^\sharp \simeq \mathbf{K}/\mathbf{K}^\sharp.$$

The R-group is a finite group, equal to the image of K in the outer automorphism group of the root datum of K_0. It will play an important part in our description of representations of K. For the special groups K appearing in our setup, we will see that the R-group is a product of copies of $\mathbb{Z}/2\mathbb{Z}$; the number of copies is bounded by the cardinality of a set of orthogonal simple roots in the Dynkin diagram of $\mathbf{G}$.

Definition 3.2. Recall from Definition 2.6 that the *group of characters of* T_f is

$$\begin{aligned}\Pi^*(T_f) &= \text{continuous characters } \xi\colon T_f \to U_1\\ &\simeq \Pi^*(\mathbf{T}_f) = \text{algebraic characters } \xi\colon \mathbf{T}_f \to \mathbb{C}^\times;\end{aligned}$$

identification of the second definition with the first is by restriction to T_f. (Here U_1 is the circle group, the compact real form of $\mathbb{C}^\times$.) The group $\Pi^*(T_f)$ is a finitely generated abelian group. More precisely, restriction to $T_{f,0}$ defines a short exact sequence of finitely generated abelian groups

$$0 \to \Pi^*(T_f/T_{f,0}) \to \Pi^*(T_f) \to \Pi^*(T_{f,0}) \to 0;$$

the image is a lattice (free abelian group) of rank equal to the dimension of the torus, and the kernel is the (finite) group of characters of the component group $T_f/T_{f,0}$.

We will occasionally refer to the *lattice of cocharacters of* T_f

$$\begin{aligned}\Pi_*(T_f) &= \text{continuous homomorphisms } \xi\colon U_1 \to T_f\\ &= \text{algebraic homomorphisms } \xi\colon \mathbb{C}^\times \to \mathbf{T}_f;\end{aligned}$$

Clearly such homomorphisms automatically have image in the identity component, so $\Pi_*(T_f) = \Pi_*(T_{f,0})$. There is a natural bilinear pairing

$$\langle , \rangle\colon \Pi^*(T_f) \times \Pi_*(T_f) \to \mathbb{Z};$$

this pairing is trivial on the torsion subgroup $\Pi^*(T_f/T_{f,0})$, and descends to the standard identification of $\Pi_*(T_{f,0})$ as the dual lattice of $\Pi^*(T_{f,0})$.

Whenever V is a complex representation of T_f, we write

$$\Delta(V, T_f) = \text{set of weights of } T_f \text{ on } V$$

regarded as a multiset of elements of $\Pi^*(T_f)$.

The most important example of such a set is the *root system of* T_f *in* K

$$\Delta(K,T_f) = \Delta\left(\mathfrak{k}/\mathfrak{t}_f\right).$$

In one-to-one correspondence with the root system is the system of *coroots of* T_f *in* K

$$\Delta^\vee(K,T_f) \subset \Pi_*(T_f);$$

we write $\alpha^\vee$ for the coroot corresponding to the root α.

The difference from the classical definition of root system is that we are keeping track of the action of the disconnected group T_f. Because the roots of $T_{f,0}$ in K have multiplicity one, the restriction map from $\Pi^*(T_f)$ to $\Pi^*(T_{f,0})$ is necessarily one-to-one on this root system. That is, each root of the connected torus $T_{f,0}$ in the connected group K_0 extends to a unique root of T_f. The point of this discussion is that $\Delta(K,T_f)$ is in canonical one-to-one correspondence with a root system in the classical sense (that of $T_{f,0}$ in K_0). We may therefore speak of classical concepts like "system of positive roots" without fear of reproach.

Definition 3.3. While we are feeling bold, let us therefore choose a system

$$\Delta^+(K,T_f) \subset \Delta(K,T_f)$$

of positive roots. Because so much of what we will do depends on this choice, it will appear constantly in the notation. Simply for notational convenience, we will therefore use the shorthand

$$\Delta_c^+ = \Delta^+(K,T_f);$$

the subscript c stands for "compact." A representation (μ, E_μ) of T_f is called *dominant* (or Δ_c^+*-dominant* if for every positive root $\alpha \in \Delta_c^+$, and every character $\xi \in \Pi^*(T_f)$ occurring in E_μ, we have $\langle \xi, \alpha^\vee\rangle \geq 0$. We write

$$\Pi^*_{\Delta_c^+ - dom}(T_f) = \Pi^*_{dom}(T_f)$$

for the set of dominant weights.

Choice of Δ_c^+ is equivalent to the choice of a Borel subalgebra

$$\mathfrak{b}_c = \mathfrak{t}_f + \mathfrak{n}_c.$$

We want to use $\mathfrak{b}_c$ to construct representations of K, and for that purpose it will be convenient to take the roots of the Borel subalgebra to be the *negative* roots:

$$\Delta(\mathfrak{n}_c, T_f) = -\Delta_c^+.$$

Define

$$\mathbf{B}_c^\sharp = \mathbf{T}_f \mathbf{N}_c \subset \mathbf{K}^\sharp,$$

which we call a *small Borel subgroup of* $\mathbf{K}$. There are several things to notice here. First, $\mathbf{B}_c^\sharp$ and its unipotent radical $\mathbf{N}_c$ are not preserved by the complex structure, so they are not defined over $\mathbb{R}$. Second, the group $\mathbf{B}_c^\sharp$ is disconnected if $\mathbf{T}_f$ is; we have

$$\mathbf{B}_{c,0} = \mathbf{T}_{f,0}\mathbf{N}_c \subset \mathbf{K}_0,$$

a Borel subgroup in the usual sense.

From time to time we will have occasion to refer to the Borel subalgebra whose roots are the positive roots; the nil radical is $\sigma\mathfrak{n}_c$ (with σ the complex conjugation). The corresponding small Borel subgroup is $\sigma(\mathbf{B}_{c,0})$.

Our first goal is to parametrize representations of K in terms of elements of $\Pi^*(T_f)$, together with a little more information. The point of this is that $\Pi^*(T_f)$ is accessible to the computer. I will recall why (in a slightly more general context) in Section 6.

Before embarking on a precise discussion of representations of K, we need a few more definitions.

Definition 3.4. In the setting of Definition 3.3, the *large Cartan subgroup of* K is

$$\begin{aligned} T_{fl} &= \text{normalizer in } K \text{ of } \mathfrak{b}_c \\ &= \text{normalizer in } K \text{ of } \mathfrak{n}_c. \end{aligned}$$

A representation μ_l of T_{fl} is called *dominant* if its restriction to T_f is dominant (Definition 3.3). If μ_l is irreducible, this is equivalent to requiring that just one weight of T_f in μ_l be dominant. We write

$$\Pi^*_{\Delta_c^+-dom}(T_{fl}) = \Pi^*_{dom}(T_{fl})$$

for the set of dominant irreducible representations of T_{fl}.

Because $\mathfrak{b}_c \cap \sigma\mathfrak{b}_c = \mathfrak{t}_f$, the group T_{fl} normalizes $T_{f,0}$. Its complexification is

$$\begin{aligned} \mathbf{T}_{fl} &= \text{normalizer of } \mathbf{T}_{f,0} \text{ and } \mathbf{B}_{c,0} \text{ in } \mathbf{K} \\ &= \text{normalizer of } \mathbf{T}_{f,0} \text{ and } \mathfrak{n}_c \text{ in } \mathbf{K}. \end{aligned}$$

Define

$$\begin{aligned} \mathbf{B}_{cl} &= \text{normalizer of } \mathbf{B}_{c,0} \text{ in } \mathbf{K} \\ &= \text{normalizer of } \mathfrak{n}_c \text{ in } \mathbf{K} \\ &= \mathbf{T}_{fl}\mathbf{N}_c. \end{aligned}$$

The equality of the three definitions is an easy exercise; the last description is a semidirect product with the unipotent factor normal.

Next we introduce the flag varieties that control the representation theory of K.

Definition 3.5. One of the most fundamental facts about connected compact Lie groups is the identification (defined by the obvious map from left to right)

$$K_0/T_{f,0} \simeq \mathbf{K}_0/\mathbf{B}_{c,0}. \tag{3.6a}$$

This homogeneous space is called the *flag variety of* K_0 (or of $\mathbf{K}_0$); it may be identified with the projective algebraic variety of Borel subalgebras of $\mathfrak{k}$. Especially when we have this interpretation in mind, we may write the space as $X(K_0)$ or $X(\mathbf{K}_0)$.

For disconnected K we will be interested in at least two versions of this space. The first is the *large flag variety for* K (or $\mathbf{K}$)

$$X_{large}(K) = K/T_f \simeq \mathbf{K}/\mathbf{B}_c^\sharp. \tag{3.6b}$$

The right side is a complex projective algebraic variety. This variety has (many) K-equivariant embeddings in the flag variety for $\mathbf{G}$, corresponding to the finitely many extensions of $\mathfrak{b}_c$ to a Borel subalgebra of $\mathfrak{g}$. But we will also need to consider the *small flag variety for* K (or $\mathbf{K}$)

$$X_{small}(K) = K/T_{fl} \simeq \mathbf{K}/\mathbf{B}_{cl}. \tag{3.6c}$$

Since $\mathbf{B}_{cl}$ is by definition the normalizer of $\mathfrak{b}_c$, the small flag variety may be identified with the variety of Borel subalgebras of $\mathfrak{k}$. It follows that the inclusion of $\mathbf{K}_0$ in $\mathbf{K}$ defines a $\mathbf{K}_0$-equivariant identification

$$K_0/T_{f,0} \simeq \mathbf{K}_0/\mathbf{B}_{c,0} \simeq \mathbf{K}/\mathbf{B}_{cl} \simeq K/T_{fl}. \tag{3.6d}$$

Because T_f is a subgroup of T_{fl}, there is a natural K-equivariant surjection

$$K/T_f = X_{large}(K) \to X_{small}(K) = K/T_{fl}. \tag{3.6e}$$

The equivariant geometry here is slightly confusing. Because $T_f \cap K_0 = T_{f,0}$, K/T_f is (as an algebraic variety) a union of $\mathrm{card}(K/K^\sharp)$ copies of $K_0/T_{f,0}$. (Here we write $\mathrm{card}\, S$ for the cardinality of a set S.) The reason this is the number of copies is that $K_0 T_f = K^\sharp$. The projection of (3.6e) sends each copy isomorphically onto $K/T_{fl} \simeq K_0/T_{f,0}$.

Proposition 3.7. *Suppose we are in the setting of Definition 3.4. Then the group* T_{fl} *meets every component of* K. *We have*

$$T_{fl} \cap K_0 = T_{f,0}, \qquad T_{fl} \cap K^\sharp = T_f.$$

Consequently the inclusion of T_{fl} *in* K *defines natural isomorphisms*

$$T_{fl}/T_{f,0} \simeq K/K_0, \qquad T_{fl}/T_f \simeq K/K^\sharp = R(K)$$

(Definition 3.1).

We can now begin to talk about representations of K.

Definition 3.8. Recall from Definition 2.6 the identification

$$\begin{aligned}\Pi^*(K) &= \text{equiv. classes of continuous irr. representations of } K\\ &\simeq \text{equiv. classes of algebraic irr. representations of } \mathbf{K} = \Pi^*(\mathbf{K}).\end{aligned}$$

The identification of the second definition with the first is by restriction to K.

Typically I will write (τ, V_τ) for a representation of K (not necessarily irreducible). The *highest weight space of* τ is

$$V_\tau^h = V_\tau/\mathfrak{n}_c V_\tau,$$

the subspace of coinvariants for the Lie algebra $\mathfrak{n}_c$ in V_τ. Recall that $\mathfrak{n}_c$ corresponds to negative roots; for that reason, V_τ^h may be naturally identified with the subspace of vectors annihilated by the positive root vectors. Explicitly,

$$V_\tau^{\sigma\mathfrak{n}_c} \simeq V_\tau^h$$

by composing the inclusion of left side in V_τ with the quotient map. Of course it is this second definition of highest weight space that is more commonly stated, but the first is going to be more natural for us. It is clear from Definition 3.4 that the large Cartan subgroup T_{fl} acts on V_τ^h; write

$$\mu_l(\tau)\colon T_{fl} \to \mathrm{End}(V_\tau^h)$$

for the corresponding representation. Of course we may also regard $\mu_l(\tau)$ as a representation of T_f, $T_{f,0}$, $\mathbf{T}_{fl}$, $\mathbf{B}_{cl}$, and so on.

Proposition 3.9. *Suppose we are in the setting of Definitions 2.4 and 3.4.*

(1) *Passage to highest weight vectors (Definition 3.8) defines a bijection*

$$\Pi^*(K) \to \Pi^*_{dom}(T_{fl}), \qquad \tau \mapsto \mu_l(\tau)$$

from irreducible representations of K to dominant irreducible representations (Definition 3.4) of the large Cartan subgroup T_{fl}.

(2) *Suppose (μ_l, E_{μ_l}) is a dominant representation of T_{fl}. Extend μ_l to an algebraic representation of $\mathbf{B}_{cl}$ (on the same space E_{μ_l}) by making $\mathbf{N}_c$ act trivially. Write $\mathcal{O}(\mu_l)$ for the corresponding $\mathbf{K}$-equivariant sheaf on the small flag variety $X_{small}(K)$. Then the space*

$$V(\mu_l) = H^0(X_{small}, \mathcal{O}(\mu_l))$$

of global sections is a finite-dimensional representation of K, irreducible if and only if μ_l is. Evaluation of sections at the base point of $X_{small}(K)$ defines a natural isomorphism

$$V(\mu_l)^h \simeq E_{\mu_l};$$

This construction therefore inverts the bijection of (1).

(3) *Suppose $\mu \in \Pi^*_{dom}(T_f)$ is a dominant character of T_f. Write $\mathcal{O}(\mu)$ for the corresponding $K(\mathbb{C})$-equivariant sheaf on the large flag variety $X_{large}(K)$, and define*

$$V(\mu) = H^0(X_{large}, \mathcal{O}(\mu)),$$

a finite-dimensional representation of K. Finally define

$$\mu_l = \mathrm{Ind}_{T_f}^{T_{fl}} \mu,$$

a representation of dimension equal to the cardinality of $T_{fl}/T_f \simeq R(K)$ (Proposition 3.7). Then there is a natural isomorphism

$$H^0(X_{large}, \mathcal{O}(\mu)) \simeq H^0(X_{small}, \mathcal{O}(\mu_l));$$

that is, $V(\mu) \simeq V(\mu_l)$.

Our goal was to relate irreducible representations of K to dominant characters of T_f. What remains is to relate dominant irreducible representations of T_{fl} to dominant characters of T_f. We will see that "most" dominant irreducible representations of T_{fl} are induced from dominant irreducible characters of T_f. Part (3) of the Proposition shows how to construct irreducible representations of K from such dominant characters of T_f. We will see what happens in the remaining cases in Theorem 4.10.

4. The R-group of K and irreducible representations of the large Cartan

In this section we elucidate the structure of $R(K)$, and find a (computable) parametrization of $\widehat{K}$ in terms of T_{fl}.

Definition 4.1. In the setting of Definition 3.1, the *Weyl group of H_f in G* is

$$W(G, H_f) = N_G(H_f)/H_f.$$

This acts on $\mathbf{H}_f$ respecting the real structure and θ; it is a subgroup of the Weyl group of the root system of $\mathbf{H}_f$ in $\mathbf{G}$. Each coset in this Weyl group has a representative in K. Furthermore any element of K normalizing $T_{f,0}$ actually normalizes all of H_f; this is immediate from the definition of H_f.

The *Weyl group of T_f in K* is

$$W(K, T_f) = N_K(T_f)/T_f.$$

According to the remarks in the preceding paragraph, inclusion defines an isomorphism

$$W(K, T_f) \simeq W(G, H_f).$$

Because T_f is a normal subgroup of T_{fl} (Definition 3.4), Proposition 3.7 provides a natural inclusion

$$T_{fl}/T_f \simeq R(K) \hookrightarrow W(K, T_f).$$

This inclusion depends on our fixed choice of Δ_c^+. If necessary to avoid confusion, we may write the image as $R(K)_{dom}$ or $R(K)_{\Delta_c^+}$. A little more explicitly, it is clear from the definition of T_{fl} that

$$R(K)_{dom} = \{w \in W(K, T_f) \mid w(\Delta_c^+) = \Delta_c^+\}.$$

Finally, since $T_f \cap K_0 = T_{f,0}$, there is a natural inclusion

$$W(K_0, T_{f,0}) \hookrightarrow W(K, T_f).$$

Proposition 4.2. *In the setting of Definition 4.1, there is a semidirect product decomposition*

$$W(K, T_f) = R(K)_{dom} \ltimes W(K_0, T_{f,0}).$$

This is an immediate consequence of the fact that $W(K_0, T_{f,0})$ acts in a simply transitive way on choices of positive roots for $T_{f,0}$ in K.

We turn next to an explicit description of $R(K)_{dom}$. Define the *modular character* $2\rho_c$ of T_{fl} to be the determinant of the adjoint action of T_{fl} on the holomorphic tangent space

$$2\rho_c(t) = \det(\mathrm{Ad}(t) \text{ acting on } \mathfrak{k}/\mathfrak{b}_c) \tag{4.3a}$$

The notation is a little misleading, since there may not be a character ρ_c with square equal to $2\rho_c$. But it is traditional. I will use the same notation $2\rho_c$ for the restriction to T_f or to $T_{f,0}$. We could avoid this abuse of notation by writing $2\rho_{cl}$, $2\rho_c$, and $2\rho_{c,0}$ for the characters of T_{fl}, T_f, and $T_{f,0}$ respectively. But the ambiguity will not lead to problems.

As a character of T_f, we have

$$2\rho_c = \sum_{\alpha\in\Delta_c^+} \alpha. \tag{4.3b}$$

But T_{fl} may act non-trivially on $\Delta(K, T_f)$, so the summands need not extend to characters of T_{fl}.

For the moment the most important property of the modular character for us is

$$\Delta_c^+ = \left\{\alpha \in \Delta(K, T_f) \mid \langle 2\rho_c, \alpha^\vee\rangle > 0\right\}. \tag{4.3c}$$

More precisely,

$$\text{simple roots of } T_f \text{ in } \Delta_c^+ = \left\{\alpha \in \Delta(K, T_f) \mid \langle 2\rho_c, \alpha^\vee\rangle = 2\right\}. \tag{4.3d}$$

Proposition 4.4. *In the setting of (4.3), regard $2\rho_c$ as an element of $\mathfrak{t}_f^*$; extend it to a θ-fixed element of $\mathfrak{h}_f^*$, as is possible uniquely. Define*

$$\Delta_c^\perp = \{\beta \in \Delta(\mathbf{G}, \mathbf{H}_f) \mid \langle 2\rho_c, \beta^\vee\rangle = 0\}$$

.

(1) *The set of roots $\Delta_c^\perp$ is a root system. If we fix any positive system for $\Delta(\mathbf{G}, \mathbf{H}_f)$ making $2\rho_c$ dominant, then $\Delta_c^\perp$ is spanned by simple roots. We may therefore define*

$$W_c^\perp = W(\Delta_c^\perp),$$

a Levi subgroup of $W(\mathbf{G}, \mathbf{H}_f)$.

(2) *The stabilizer of $2\rho_c$ in $W(\mathbf{G}, \mathbf{H}_f)$ is equal to $W_c^\perp$.*

(3) *As a subgroup of $W(K, T_f)$ (Definition 4.1), $R(K)_{dom}$ is the intersection with $W_c^\perp$:*

$$R(K)_{dom} = W(K, T_f) \cap W_c^\perp.$$

(4) *The root system $\Delta_c^\perp$ is of type A_1^r. That is,*

$$\Delta_c^\perp = \{\pm\beta_1, \dots, \pm\beta_r\},$$

with $P = \{\beta_i\}$ a collection of strongly orthogonal noncompact imaginary roots. Consequently $W_c^\perp \simeq (\mathbb{Z}/2\mathbb{Z})^r$. Elements of $W_c^\perp$ are *in one-to-one correspondence with subsets $A \subset P$: to the subset A corresponds the Weyl group element*

$$w_A = \prod_{\beta_i \in A} s_{\beta_i}$$

(5) *The group $R(K)_{dom}$ acts on the collection of weights (Definition 3.3) of T_f. Each $R(K)_{dom}$ orbit of weights has a unique element μ satisfying*

$$\sum_{\beta_i \in A} \langle \mu, \beta_i^\vee \rangle \geq 0 \qquad (\textit{all } w_A \in R(K)_{dom}).$$

(6) *Suppose $\mu \in \Pi^*_{dom}(T_f)$. Define*

$$\begin{aligned} P(\mu) &= \{\beta_i \in P \mid \langle \mu, \beta_i^\vee \rangle = 0\} \\ R(K,\mu)_{dom} &= \{w_A \mid w_A \in R(K)_{dom},\ A \subset P(\mu)\}. \end{aligned}$$

Then $R(K,\mu)_{dom}$ is the stabilizer of μ in $R(K)_{dom}$.

Proof. An appropriate analogue of (1) holds with $2\rho_c$ replaced by any weight in $\Pi^*(\mathbf{H}_f)$; and (still in that generality) (2) is Chevalley's theorem on the stabilizer of a weight in the Weyl group. Part (3) is now clear from Definition 4.1. For (4), it is clear that $\Delta_c^\perp$ is preserved by θ but has no compact roots (since compact coroots cannot vanish on $2\rho_c$). Now the assertion that $\Delta_c^\perp$ is of type A_1^r, with all roots noncompact imaginary, can be proved either by a short direct argument, or by inspection of the classification of real reductive groups. Parts (5) and (6) are elementary. ☐

The $R(K)_{dom}$ orbit representatives defined in (5) are called *P-positive*, or *$\{\beta_i\}$-positive*. Of course the notion depends on the choice of $\{\beta_i\}$; that is, on the choice of one root from each pair $\pm\beta_i$.

Not all of the roots in P contribute to $R(K)_{dom}$, and occasionally it will be helpful to single out the ones (called *essential*) that do. In the notation of the preceding Proposition, set

$$P_{ess} = \bigcup_{w_A \in R(K)_{dom}} A \quad \subset \; P = \{\beta_1, \ldots, \beta_r\}. \tag{4.5}$$

We may write this set as $P_{ess}(K) \subset P(K)$ if necessary. Similarly we define $P_{ess}(\mu)$. Clearly the notion of P-positive depends only on P_{ess}, so we may call it *P_{ess}-positive.*

Corollary 4.6. *In the setting of Definition 4.1, each orbit of $W(K,T_f)$ on $\Pi^*(T_f)$ has a unique representative which is dominant for Δ_c^+ (Definition 3.3 and P-positive (Proposition 4.4).*

Definition 4.7. Suppose $\mu \in \Pi^*_{dom}(T_f)$. In terms of the isomorphism $R(K)_{dom} \simeq T_{fl}/T_f$ of Proposition 3.7, define $T_{fl}(\mu)$ to be the inverse image of $R(K,\mu)_{dom}$ (Proposition 4.4). With respect to the natural action of T_{fl} on characters of its normal subgroup T_f, we have

$$T_{fl}(\mu) = \{x \in T_{fl} \mid x \cdot \mu = \mu\}.$$

We are interested in representations of $T_{fl}(\mu)$ extending the character μ of T_f. We will show in a moment (Lemma 4.8) that the quotient group $T_{fl}(\mu)/\ker(\mu)$ is abelian, so irreducible representations of this form are one-dimensional. We can therefore define an *extension of* μ to be a character

$$\widetilde{\mu} \in \Pi^*(T_{fl}(\mu))$$

with the property that

$$\widetilde{\mu}|_{T_f} = \mu.$$

There is a natural simply transitive action of $\Pi^*(R(K,\mu)_{dom})$ on extensions of μ, just by tensor product of characters. We call $\widetilde{\mu}$ an *extended dominant weight.*

Lemma 4.8. *The commutator subgroup of* $T_{fl}(\mu)$ *is contained in the product of the images of the coroots*

$$\beta^\vee \colon U(1) \to T_f,$$

as β *runs over the noncompact imaginary roots in* $P_{ess}(\mu)$ *(Proposition 4.4). Since* μ *is trivial on all of these coroots, the quotient* $T_{fl}(\mu)/\ker(\mu)$ *is abelian.*

Example 4.9. Suppose $\mathbf{G} = \mathbf{GL}_{2n}$ with the standard real form (complex conjugation of matrices) $G = GL_{2n}(\mathbb{R})$. We choose the Cartan involution $\theta(g) = {}^t g^{-1}$, so that $K = O_{2n}$ is the compact orthogonal group. We may choose as maximal torus

$$T_{f,0} = [SO_2]^n,$$

embedded as block-diagonal matrices. If we identify $\mathbb{R}^2$ with $\mathbb{C}$ as usual, then $SO_2 \subset GL_2(\mathbb{R})$ is identified with U_1, the group of multiplications by complex numbers of absolute value 1. It follows that

$$Z_{GL_2(\mathbb{R})}(SO_2) \simeq \mathbb{C}^\times,$$

the multiplicative group of nonzero complex numbers; and that

$$H_f = Z_{GL_{2n}(\mathbb{R})}([SO_2]^n) \simeq [\mathbb{C}^\times]^n,$$

a fundamental Cartan subgroup of G. The maximal compact subgroup of $[\mathbb{C}^\times]^n$ is $[U_1]^n$, so we find

$$T_f = T_{f,0} = [SO_2]^n.$$

Characters of SO_2 may be naturally identified with $\mathbb{Z}$, so $\Pi^*(T_f) = \mathbb{Z}^n$.

If we write $e_1, \ldots, e_n$ for the standard basis of $\mathbb{Z}^n$, the compact root system is

$$\Delta(K, T_f) = \{\pm e_i \pm e_j, \mid 1 \le i \ne j \le n\}.$$

As positive roots we may choose

$$\Delta_c^+ = \{e_i \pm e_j, \mid 1 \le i < j \le n\}.$$

Adding these characters gives

$$2\rho_c = (2n-2, 2n-4, \ldots, 2, 0) \in \mathbb{Z}^n.$$

The roots of K are precisely the restrictions to T_f of the complex roots of H_f. Evidently none of these complex roots is orthogonal to $2\rho_c$. The imaginary roots are $\{\pm 2e_j\}$, all noncompact. Only the last of these is orthogonal to $2\rho_c$, so the set P of Proposition 4.4 is $\{2e_n\}$.

The compact Weyl group $W(K, T_f)$ acts on $[SO_2]^n$ by permuting the coordinates and inverting some of them. The Weyl group $W(K_0, T_{f,0})$ is the subgroup inverting always an even number of coordinates. Consequently the large compact Cartan is

$$T_{fl} = [SO_2]^{n-1} \times O_2,$$

and the R group for K is

$$R(K)_{dom} = \{1, s_{2e_n}\}.$$

(Other choices of Δ_c^+ would replace n by some other coordinate m in this equality.) A weight $\mu = (\mu_1, \ldots, \mu_n)$ is dominant if and only if

$$\mu_1 \geq \mu_2 \geq \cdots \geq \mu_{n-1} \geq |\mu_n|.$$

It is P-positive (remark after Proposition 4.4) if and only if $\mu_n \geq 0$.

Here is a classification of irreducible representations of K.

Theorem 4.10. *Suppose G is the group of real points of a connected reductive algebraic group, K is a maximal compact subgroup (Definition 2.4), and T_f is a small Cartan subgroup of K (Definition 3.1). Fix a positive root system for T_f in K (Definition 3.3, and let $R(K)_{dom}$ be the R-group for K (acting on the character group $\Pi^*(T_f)$). Choose roots $\{\beta_i\}$ as in Proposition 4.4. Then the irreducible representations of K are in natural one-to-one correspondence with $R(K)_{dom}$ orbits of extended dominant weights $\widetilde{\mu}$ (Definition 4.7).*

Write $\tau(\widetilde{\mu})$ for the irreducible representation of K corresponding to the extended weight $\widetilde{\mu}$. Here are some properties of $\tau(\widetilde{\mu})$.

1. *The restriction of $\tau(\widetilde{\mu})$ to $K^\sharp$ (Definition 3.1) is the direct sum of the irreducible representations of highest weights $r \cdot \mu$ (for $r \in R(K)_{dom}$), each appearing with multiplicity one. These summands remain irreducible on restriction to K_0. If we write μ_0 for the restriction of μ to $T_{f,0}$, and $\tau(\mu_0)$ for the irreducible representation of K_0 of highest weight μ_0, then*

   $$\dim \tau(\widetilde{\mu}) = (\dim \tau(\mu_0))(|R(K)_{dom}/R(K,\mu)_{dom}|)$$

   The first factor is given by the Weyl dimension formula, and the second is a power of two.
2. *The highest weight of $\tau(\widetilde{\mu})$ is*

   $$(\widetilde{\mu})_l = \mathrm{Ind}_{T_{fl}(\mu)}^{T_{fl}} \widetilde{\mu}.$$

This theorem is still not quite amenable to computers: although the set of P-positive dominant characters μ of T_f can be traversed using the algebraic character lattice $\Pi^*(\mathbf{H}_f)$ and its automorphism θ (Proposition 6.3), the extensions of μ are a bit subtle. We can calculate the group $R(K,\mu)_{dom}$ (Proposition 4.4), whose character group acts in a simply transitive way on the extensions; but there is no natural base point in the set of extensions, so it is not clear how to keep track of which is which. We will resolve this problem in Theorem 11.9 below. For the moment, we will simply reformulate Theorem 4.10 in a way that does not involve a choice of positive roots for K.

Definition 4.11. Suppose $\mu \in \Pi^*(T_f)$ is any character. The Weyl group $W(K,T_f)$ of Definition 4.1 acts on $\Pi^*(T_f)$, so we can define

$$W(K,T_f)^\mu = \{w \in W(K,T_f) \mid w \cdot \mu = \mu\},$$

the stabilizer of μ. Using the isomorphism $R(K) \simeq W(K,T_f)/W(K_0,T_{f,0})$ of Proposition 4.2, this defines a subgroup

$$R(K,\mu) = W(K,T_f)^\mu / W(K_0,T_{f,0})^\mu \subset R(K).$$

It is not difficult to show that this definition agrees with the one in Proposition 4.4 when μ is dominant.

We need a notion of extended weight here, and the absence of a chosen positive root system for K complicates matters slightly. Let us write

$$\begin{aligned} N_f &= \text{normalizer of } T_f \text{ in } K^\sharp, \\ N_{fl} &= \text{normalizer of } T_f \text{ in } K. \end{aligned}$$

Then $N_f/T_f \simeq W(K^\sharp,T_f) = W(K_0,T_{f,0})$, and $N_{fl}/T_f \simeq W(K,T_f)$. The subgroups corresponding to the stabilizer of μ are

$$\begin{aligned} N_f(\mu) &= \{n \in N_f \mid n \cdot \mu = \mu\}, \\ N_{fl}(\mu) &= \{n \in N_{fl} \mid n \cdot \mu = \mu\}. \end{aligned}$$

Clearly $N_{fl}(\mu)/T_f \simeq W(K,T_f)^\mu$, and $N_f(\mu)/T_f \simeq W(K^\sharp,T_f)^\mu$.

For each root α of T_f in K there is a three-dimensional root subgroup of K locally isomorphic to $SU(2)$. The simple reflection $s_\alpha \in W(K^\sharp,T_f)$ has a representative σ_α in this root subgroup; the representative σ_α is unique up to multiplication by the coroot circle subgroup $\alpha^\vee(U_1) \subset T_f$.

A *small extension of* μ is a character $\widetilde{\mu}$ of $N_{fl}(\mu)$ subject to the following two conditions.

(1) The restriction of $\widetilde{\mu}$ to T_f is equal to μ.
(2) For each compact root α with $s_\alpha \in W(K^\sharp,T_f)^\mu$, we have $\widetilde{\mu}(\sigma_\alpha) = 1$.

The hypothesis $s_\alpha \in (K^\sharp,T_f)^\mu$ is equivalent to

$$\langle \mu, \alpha^\vee \rangle = 0;$$

that is, to μ being trivial on coroot circle subgroup $\alpha^\vee(U_1)$. Since σ_α is well-defined up to this subgroup, the requirement in (2) is independent of the choice of σ_α.

We will show in a moment (Lemma 4.12) that the quotient of $N_{fl}(\mu)$ by the kernel of μ and the various σ_α is actually abelian; so looking for one-dimensional extensions $\widetilde{\mu}$ is natural. It will turn out that small extensions of μ must exist. Once that is known, it follows immediately that $\Pi^*(R(K,\mu))$ acts in a simply transitive way on the set of small extensions of μ, just by tensor product of characters.

The Weyl group $W(K,T_f)$ acts naturally on extended weights; the stabilizer of $\widetilde{\mu}$ is equal to $W(K,T_f)^\mu$. Write

$$P_K^e(T_f) = \text{small extensions of characters of } T_f.$$

(This is the first instance of our general scheme of writing P for a set of parameters for representations.) Elements of $P_K^e(T_f)$ are also called *extended weights*.

Lemma 4.12. *In the setting of Definition 4.11, choose a positive root system $(\Delta^+)'(K,T_f)$ making μ dominant. Choose a corresponding set P' of noncompact imaginary roots as in Proposition 4.4. Define $P'(\mu) \subset P'$ as in Proposition 4.4.*

The commutator subgroup of $N_f(\mu)$ is contained in the product of the images of the coroots

$$\gamma^\vee : U(1) \to T_f,$$

as γ runs over the noncompact imaginary roots in $P(\mu)$ and the compact roots orthogonal to μ. Since μ is trivial on all of these coroots, the quotient

$$N_{fl}(\mu)/\langle \ker(\mu), \{\sigma_\alpha\}\rangle$$

is abelian.

Corollary 4.13. *Suppose G is the group of real points of a connected reductive algebraic group, K is a maximal compact subgroup (Definition 2.4), and T_f is a small Cartan subgroup of K (Definition 3.1). Suppose $\widetilde{\mu} \in P_K^e(T_f)$ is an extended weight for K (Definition 4.11). Then there is a unique irreducible representation $\tau(\widetilde{\mu}) \in \Pi^*(K)$, with the following properties.*

1. *If μ is dominant (Definition 3.3), then $\tau(\widetilde{\mu})$ is equal to the representation $\tau(\widetilde{\mu}|T_{fl}(\mu))$ defined in Theorem 4.10.*
2. *If $w \in W(K,T_f)$, then $\tau(w\cdot\widetilde{\mu}) = \tau(\widetilde{\mu})$.*

We say that $\tau(\widetilde{\mu})$ has extremal weight μ or $\widetilde{\mu}$.

Two extended weights $\widetilde{\mu}$ and $\widetilde{\gamma}$ define the same irreducible representation of K if and only if $\widetilde{\gamma} \in W(K,T_f)\cdot\widetilde{\mu}$. In this way we get a bijection

$$\Pi^*(K) \leftrightarrow P_K^e(T_f)/W(K,T_f).$$

According to Corollary 4.6, the two conditions in this Corollary specify $\tau(\widetilde{\mu})$ completely; one just has to check that the resulting representation of K is well-defined.

5. Fundamental series, limits, and continuations

The Harish-Chandra parameter for a discrete series representation is perhaps best thought of as a character of a two-fold cover (the "square root of ρ cover") of a compact Cartan subgroup. I am going to avoid discussion of coverings by shifting the parameter by ρ, the half sum of a set of imaginary roots.

Definition 5.1. In the setting of Definition 3.1, a *shifted Harish-Chandra parameter for a fundamental series representation* is a pair

$$\Phi = (\phi, \Delta^+_{im})$$

(with $\phi \in \Pi^*(H_f)$ a character and $\Delta^+_{im} \subset \Delta_{im}(G, H_f)$ a system of positive roots) subject to the following requirement:

(1-std) For every positive imaginary root $\alpha \in \Delta^+_{im}$, we have

$$\langle d\phi, \alpha^\vee \rangle > 1.$$

This requirement implies that the positive root system is entirely determined by the character ϕ. We keep it in the notation for formal consistency with notation for coherently continued representations, to be defined in a moment.

We define

$$\begin{aligned} P^s_G(H_f) &= \{\text{shifted parameters for fundamental series}\} \\ &= \{\text{pairs } \Phi = (\phi, \Delta^+_{im}) \text{ satisfying condition (1-std) above.}\} \end{aligned}$$

(The superscript s stands for "shifted"; it is a reminder that this parameter differs by a ρ shift from the Harish-Chandra parameter.)

The positivity requirement (1-std) on Φ looks a little peculiar, because ϕ differs by a ρ-shift from the most natural parameter. In terms of the infinitesimal character parameter $\zeta(\Phi)$ defined below, the positivity requirement is just

$$\langle \zeta(\Phi), \alpha^\vee \rangle > 0 \qquad (\alpha \in \Delta^+_{im}).$$

The reason we do not use $\zeta(\Phi)$ instead of ϕ is that $\zeta(\Phi)$ is only in $\mathfrak{h}^*$: it does not remember the values of ϕ off the identity component of H_f.

The imaginary roots divide into compact and noncompact as usual, according to the eigenvalue of θ on the corresponding root space:

$$\begin{aligned} \Delta^+_{im,c} &= \text{positive imaginary roots in } \mathfrak{k}, \\ \Delta^+_{im,n} &= \text{remaining positive imaginary roots.} \end{aligned}$$

Define

$$2\rho_{im} = \sum_{\alpha \in \Delta^+_{im}} \alpha, \qquad 2\rho_{im,c} = \sum_{\beta \in \Delta^+_{im,c}} \beta \qquad 2\rho_{im,n} = \sum_{\gamma \in \Delta^+_{im,n}} \gamma.$$

Each of these characters may be regarded as belonging either to the group of real characters $\Pi^*(H_f)$ or to the algebraic character lattice $\Pi^*(\mathbf{H}_f)$. The lattice of characters may be embedded in $\mathfrak{h}^*_f$ by taking differentials, so we may also regard these characters as elements of $\mathfrak{h}_f(\mathbb{C})^*$. There they may be divided by two, defining ρ_{im}, $\rho_{im,c}$, and $\rho_{im,n}$ in $\mathfrak{h}_f(\mathbb{C})^*$. We will be particularly interested in the weight

$$\zeta(\Phi) = d\phi - \rho_{im} \in \mathfrak{h}_f^*,$$

called the *infinitesimal character parameter for* Φ. Notice that it is dominant and regular for Δ_{im}^+.

More generally, a *shifted Harish-Chandra parameter for a limit of fundamental series representations* is a pair

$$\Phi = (\phi, \Delta_{im}^+)$$

(with $\phi \in \Pi^*(H_f)$ a character and $\Delta_{im}^+ \subset \Delta_{im}(G, H_f)$ a system of positive roots) subject to the following requirement:

(1-lim) For every positive imaginary root $\alpha \in \Delta_{im}^+$, we have

$$\langle d\phi, \alpha^\vee \rangle \geq 1.$$

Again this requirement implies that the positive root system is entirely determined by the character ϕ. Again we are interested in the infinitesimal character parameter

$$\zeta(\Phi) = d\phi - \rho_{im} \in \mathfrak{h}_f^*;$$

now it is dominant but possibly singular for Δ_{im}^+.

Define

$$\begin{aligned} P_G^{s,lim}(H_f) &= \{\text{shifted parameters for limits of fundamental series}\} \\ &= \{\text{pairs } \Phi = (\phi, \Delta_{im}^+) \text{ satisfying condition (1-lim) above.}\} \end{aligned}$$

We can write the positivity requirement (1-lim) in the equivalent form

$$\langle \zeta(\Phi), \alpha^\vee \rangle \geq 0 \qquad (\alpha \in \Delta_{im}^+).$$

Finally, a *shifted Harish-Chandra parameter for a continued fundamental series representation* is a pair

$$\Phi = (\phi, \Delta_{im}^+)$$

with $\phi \in \Pi^*(H_f)$, Δ_{im}^+ a system of positive imaginary roots, and no positivity hypothesis on ϕ. (This is the setting in which the positive root system needs to be made explicit.) We still use the infinitesimal character parameter

$$\zeta(\Phi) = d\phi - \rho_{im} \in \mathfrak{h}_f^*;$$

which now need not have any positivity property.

We define

$$\begin{aligned} P_G^{s,cont}(H_f) &= \{\text{shifted parameters for continued fundamental series}\} \\ &= \{\text{pairs } \Phi = (\phi, \Delta_{im}^+).\} \end{aligned}$$

For each kind of parameter there is a discrete part, which remembers only the restriction ϕ_d of the character ϕ to the compact torus T_f. We write for example

$$\begin{aligned} P^{s,lim}_{G,d}(H_f) &= \{\text{shifted discrete parameters for limits of fundamental series}\} \\ &= \{\text{pairs } \Phi_d = (\phi_d, \Delta^+_{im})\} \end{aligned}$$

with $\phi_d \in \Pi^*(T_f)$ and Δ^+_{im} a system of positive imaginary roots satisfying condition (1-lim) above.

A given discrete parameter has a distinguished extension to a full parameter, namely the one which is trivial on $A_{f,0}$ (the vector subgroup of H_f). By using this extension, we may regard $P^{s,lim}_{G,d}(H_f)$ as a subset of $P^{s,lim}_G(H_f)$.

Proposition 5.2. *In the setting of Definition 5.1, suppose* $\Phi = (\phi, \Delta^+_{im}) \in P^s_G(H_f)$ *is a shifted Harish-Chandra parameter for a fundamental series representation. Then there is a fundamental series representation* $I(\Phi)$ *for* G *attached to* Φ*. It is always non-zero, and its restriction to* K *depends only on the restriction of* Φ *to* T_f*. The infinitesimal character of* $I(\Phi)$ *corresponds (in the Harish-Chandra isomorphism) to the weight* $\zeta(\Phi) \in \mathfrak{h}^*_f$*. This fundamental series representation has a (necessarily irreducible) Langlands quotient* $J(\Phi)$*. For* Φ *and* Ψ *in* $P^s_G(H_f)$*, we have*

$$I(\Phi) \simeq I(\Psi) \iff J(\Phi) \simeq J(\Psi) \iff \Psi \in W(G, H_f) \cdot \Phi.$$

Define

$$\mu(\Phi) = [\phi - 2\rho_{im,c}]|_{T_f} \in \Pi^*(T_f).$$

Then $\mu(\Phi)$ *is dominant with respect to the compact imaginary roots* $\Delta^+_{im,c}$*. The R-group* $R(K, \mu(\Phi))$ *(Definition 4.11) is necessarily trivial, so we may identify* $\mu(\Phi)$ *with an extended weight for* K *(Definition 4.11). The corresponding representation* $\tau(\mu(\Phi))$ *of* K *(Corollary 4.13) is the unique lowest* K*-type of* $I(\Phi)$ *and of* $J(\Phi)$*.*

We will discuss the construction of $I(\Phi)$ in Section 11. For now we observe only that $I(\Phi)$ is a discrete series representation (in the strong sense that the matrix coefficients are square integrable) if and only if $H_f = T_f$ is compact. In that case $I(\Phi) = J(\Phi)$ is irreducible, and its Harish-Chandra parameter (from his classification of discrete series) is equal to $\zeta(\Phi)$.

For limits of fundamental series the situation is quite similar; the main problem is that the corresponding representation of G may vanish. Here is a statement.

Proposition 5.3. *In the setting of Definition 5.1, suppose* $\Phi = (\phi, \Delta^+_{im}) \in P^{s,lim}_G(H_f)$ *is a shifted Harish-Chandra parameter for a limit of fundamental series. Then there is a limit of fundamental series representation* $I(\Phi)$ *for* G *attached to* Φ*. Its restriction to* K *depends only on the restriction of* Φ *to* T_f*. The infinitesimal character of* $I(\Phi)$ *corresponds (in the Harish-Chandra isomorphism) to the weight* $\zeta(\Phi) \in \mathfrak{h}^*_f$*. This fundamental series* *representation has a Langlands quotient $J(\Phi)$. We have*

$$\Psi \in W(G,H_f)\cdot\Phi \Longrightarrow I(\Phi)\simeq I(\Psi) \Longleftrightarrow J(\Phi)\simeq J(\Psi).$$

Define

$$\mu(\Phi) = [\Phi - 2\rho_{im,c}]|_{T_f} \in \Pi^*(T_f).$$

Then the following three conditions are equivalent:

1. *the weight $\mu(\Phi)$ is dominant with respect to the compact imaginary roots $\Delta^+_{im,c}$;*
2. *there is no simple root for Δ^+_{im} which is both compact and orthogonal to $\phi - \rho_{im}$; and*
3. *the standard limit representation $I(\Phi)$ is non-zero.*

Assume now that these equivalent conditions are satisfied. The R-group $R(K,\mu(\Phi))$ (Definition 4.11) is necessarily trivial, so we may identify $\mu(\Phi)$ with an extended weight for K (Definition 4.11). The corresponding representation $\tau(\Phi)$ of K (Corollary 4.13) is the unique lowest K-type of $I(\Phi)$ and of $J(\Phi)$.

Here is the result for coherent continuation.

Proposition 5.4. *In the setting of Definition 5.1, suppose $\Phi = (\phi, \Delta^+_{im}) \in P^{s,cont}_G(H_f)$ is a shifted Harish-Chandra parameter for continued fundamental series. Then there is a virtual representation $I(\Phi)$ for G, with the following properties.*

1. *The restriction of $I(\Phi)$ to K depends only on the restriction of Φ to T_f.*
2. *We have*
$$I(\Phi) \simeq I(w\cdot\Phi) \qquad (w \in W(G,H_f)).$$
3. *The virtual representation $I(\Phi)$ has infinitesimal character corresponding to $\zeta(\Phi) \in \mathfrak{h}^*_f$.*
4. *If $\zeta(\Phi)$ is weakly dominant for Δ^+_{im}, then $I(\Phi)$ is equivalent (as a virtual representation) to the limit of fundamental series attached to Φ in Proposition 5.3.*
5. *Suppose V is a finite-dimensional representation of G. Recall that $\Delta(V,H_f)$ denotes the multiset of weights of H_f on V (Definition 3.2). Then*
$$I(\Phi)\otimes V \simeq \sum_{\delta\in\Delta(V,H_f)} I(\Phi+\delta).$$
Here $\Phi+\delta$ denotes the continued fundamental series parameter
$$\Phi+\delta = (\phi+\delta, \Delta^+_{im}) \in P^{s,cont}_G(H_f).$$

The result in (4) justifies the ambiguous notation $I(\Phi)$; without it, we would need to specify whether we were regarding Φ as an element of $P^{s,cont}$ or of P^s.

One of the central goals of this paper is writing down the bijection between irreducible representations of K and certain "final standard limit" representations of G. We are well on the way to this already; here is how.

Definition 5.5. A limit parameter $\Phi \in P_G^{s,lim}(H_f)$ (Definition 5.1) is called *final* if the corresponding representation $I(\Phi)$ is non-zero. According to Proposition 5.3, this is equivalent to requiring that $\mu(\Phi)$ be dominant with respect to $\Delta^+_{im,c}$. We write $P_G^{s,finlim}(H_f)$ for the set of final limit parameters for H_f.

Proposition 5.6. *The map $\Phi \mapsto \tau(\Phi)$ (Proposition 5.3) is an injection from $W(G, H_f)$ orbits of discrete final limit parameters for H_f into $\Pi^*(K)$. The image consists precisely of those irreducible representations of K for which any highest weight μ has the following two properties.*

(1) *The weight $\mu + 2\rho_c \in \Pi^*(T_f)$ is regular with respect to the system of restricted roots of T_f in G. Because of this property, we can define $\Delta^+(G, T_f)(\mu)$ to be the unique positive system making $\mu + 2\rho_c$ dominant.*
(2) *The weight $\mu + 2\rho_c - \rho(\mu) \in \mathfrak{t}_f^*$ is weakly dominant for $\Delta^+(G, T_f)(\mu)$.*

There is a similar statement just for fundamental series; the only change is that one imposes the stronger condition that $\mu + 2\rho_c - \rho$ is strictly dominant for $\Delta^+(G, T_f)$.

6. Characters of compact tori

Suppose $\mathbf{H}$ is a θ-stable maximal torus in $\mathbf{G}$ that is defined over $\mathbb{R}$. The *compact factor of H* is

$$\mathbf{T} = \mathbf{H}^\theta, \tag{6.1a}$$

the (algebraic) group of fixed points of θ on $\mathbf{H}$. Clearly this is an abelian algebraic group defined over $\mathbb{R}$. Its group of real points is

$$T = H^\theta = H \cap K, \tag{6.1b}$$

the maximal compact subgroup of the real Cartan subgroup H. Write

$$\mathbf{T}_e = (\mathbf{H}^\theta)_e, \quad T_e = \mathbf{T}_e(\mathbb{R}); \tag{6.1c}$$

this last group is a connected compact torus, equal to the identity component T_0.

So far all of this fits well with standard notation for real groups. Now we begin to deviate a little. The *split component of H* is

$$\mathbf{A} = \mathbf{H}^{-\theta}, \tag{6.1d}$$

the (algebraic) group of fixed points of $-\theta$ on $\mathbf{H}$. Again this is an abelian algebraic group defined over $\mathbb{R}$. Its group of real points is

$$A = H^{-\theta}. \tag{6.1e}$$

We have containments

$$A_0 \subset A_e \subset A, \tag{6.1f}$$

usually both proper. The complex torus $\mathbf{A}_e$ is the maximal $\mathbb{R}$-split torus in $\mathbf{H}$, so A_e is a product of copies of $\mathbb{R}^\times$. Consequently A_0 is a vector group, isomorphic to its Lie algebra by the exponential map. There is a direct product of Lie groups (the *Cartan decomposition*)

$$H = T \times A_0, \tag{6.1g}$$

but we will avoid this (first of all on the aesthetic grounds that it is not algebraic). Again in real groups one typically writes A for our A_0; this we will avoid even more assiduously.

The next goal is to describe the set of characters of T:

$$\Pi^*_{adm}(T) = \text{continuous characters } \xi\colon T \to U_1 \tag{6.2}$$

Proposition 6.3. (1) *In the setting of* (6.1)*, restriction to T defines a natural identification*

$$\Pi^*_{adm}(T) \simeq \Pi^*_{alg}(\mathbf{H})/(1-\theta)\Pi^*_{alg}(\mathbf{H})$$

given by restriction of characters of $\mathbf{H}$ to T.

(2) *The identity component $\mathbf{T}_e$ is equal to the image of the homomorphism*

$$\delta(\theta)\colon \mathbf{H} \to \mathbf{H}, \qquad \delta(\theta)(h) = h \cdot \theta(h).$$

(3) *If $\xi \in \Pi^*_{alg}(\mathbf{H})$, then*

$$\xi \circ \delta(\theta) = \xi + \theta(\xi).$$

Consequently ξ is trivial on the image of $\delta(\theta)$ if and only if $\xi+\theta\xi = 0$; that is, if and only if ξ belongs to the -1 eigenspace $\Pi^*_{alg}(\mathbf{H})^{-\theta}$ of θ on the character lattice.

(4) *Restriction to T_0 defines a natural identification*

$$\Pi^*_{adm}(T_0) \simeq \Pi^*_{alg}(\mathbf{H})/\Pi^*_{alg}(\mathbf{H})^{-\theta}.$$

The δ in (3) is meant to stand for "double."

The main point of Proposition 6.3 is the identification of characters of T with cosets in the lattice of characters of $\mathbf{H}$, a lattice to which the atlas software has access.

7. Split tori and representations of K

In Section 5 we wrote down the lowest K-type correspondence between fundamental series and certain representations of K. In this section we look at the opposite extreme case: lowest K-types of principal series representations. The general case is going to be built from these two extremes in a fairly simple way.

In this section we therefore assume that G is quasisplit, and that H_s is a θ-stable maximally split Cartan subgroup of G. What these assumptions mean is that there are no imaginary roots of H_s in G. Define

$$T_s = H_s^\theta = H_s \cap K, \qquad A_s = H_s^{-\theta} \tag{7.1a}$$

as in (6.1), the compact and split parts of H_s. Again, the quasisplit hypothesis means that every root has a non-trivial restriction to A_s, and even to the real identity component $A_{s,0}$. These restrictions form a (possibly non-reduced) root system

$$\Delta(G, A_s) = \{\alpha|_{\mathfrak{a}_s} \mid \alpha \in \Delta(G, H_s). \tag{7.1b}$$

The Weyl group of this root system is isomorphic to $W(G, H_s)$ and to $W(\mathbf{G}, \mathbf{H}_s)^\theta$:

$$W(\mathbf{G}, \mathbf{H}_s)^\theta \simeq W(G, H_s) \simeq W(G, A_s); \tag{7.1c}$$

the isomorphisms are given by restriction of the action from $\mathbf{H}$ to H_s to $A_{s,0}$.

There is some possible subtlety or confusion here arising from the disconnectedness of A_s. We have written $\Delta(G, A_s)$ instead of $\Delta(G, A_{s,0})$ simply because the former is shorter; but it is not entirely clear what should be meant by a root system inside the character group of a disconnected reductive abelian group like $\mathbf{A}_s$. One reassuring fact is that if α_1 and α_2 are roots in $\Delta(G, H_s)$ having the same restriction to $\mathfrak{a}_s$, then they also have the same restriction to $\mathbf{A}_s$. (The reason is that the hypothesis implies that either $\alpha_1 = \alpha_2$, or $\alpha_1 = -\theta\alpha_2$.) This means that there can be no confusion about the set $\Delta(G, A_s)$: it is the set of orbits of $\{1, -\theta\}$ on $\Delta(G, H_s)$.

Definition 7.2. In the setting of (7.1), a *(shifted) parameter for a principal series representation* is a pair

$$\Phi = (\phi, \emptyset)$$

with $\phi \in \Pi^*(H_s)$ a character and $\emptyset$ a set of positive roots for the (empty set of) imaginary roots of H_s in G. We retain the word "shifted" and the ordered pair structure for formal consistency with Definition 5.1 even though the shift (by imaginary roots) is now zero. In the same way, we may also call these same pairs *shifted parameters for limits of principal series* or *shifted parameters for continued principal series* even though the distinction among these concepts (which depends on imaginary roots) is empty. We define

$$P_G^s(H_s) = P_G^{s,lim}(H_s) = P_G^{s,cont}(H_s) = \Pi^*(H_s) \times \{\emptyset\},$$

the set of shifted parameters for principal series representations. We will need the weight

$$\zeta(\Phi) = d\phi - \rho_{im} = d\phi \in \mathfrak{h}_s^*,$$

called the *infinitesimal character parameter for* Φ. We may occasionally write $\zeta_{\mathfrak{g}}(\Phi)$ for clarity.

The *(shifted) discrete parameters for principal series representations* are the restrictions to T_s of parameters for principal series:

$$P^s_{G,d}(H_s) = \Pi^*(T_s) \times \{\emptyset\}.$$

As in the case of fundamental series, each discrete parameter has a distinguished extension to a full parameter (the one that is trivial on $A_{s,0}$), allowing us to regard

$$P^s_{G,d}(H_s) \subset P^s_G(H_s).$$

Proposition 7.3. *In the setting of Definition 7.2, suppose $\Phi = (\phi, \emptyset) \in P^s_G(H_s)$ is a (shifted) Harish-Chandra parameter for a principal series representation. Choose a system of positive roots $\Delta^+(G, A_s)$ making the real part of $d\phi|_{\mathfrak{a}_s}$ weakly dominant, and let $B_s \supset H_s$ be the corresponding Borel subgroup of G. Put*

$$I(\Phi) = \mathrm{Ind}_{B_s}^G \Phi$$

(normalized induction). Then $I(\Phi)$ is a principal series representation of G attached to Φ; its equivalence class is independent of the choice of $\Delta^+(G, A_s)$. It is always non-zero, and its restriction to K is

$$I(\Phi)|_K = \mathrm{Ind}_{T_s}^K \phi|_{T_s}$$

(which depends only on the restriction of Φ to T_s). The infinitesimal character of $I(\Phi)$ corresponds to the weight $\zeta(\Phi) \in \mathfrak{h}_s^*$. This principal series representation has a (possibly reducible) Langlands quotient $J(\Phi)$. We have

$$I(\Phi) \simeq I(\Psi) \iff J(\Phi) \simeq J(\Psi) \iff \Psi \in W(G, H_s) \cdot \Phi.$$

Now regard Φ as a shifted parameter for continued principal series, and suppose V is a finite-dimensional representation of G. Recall the multiset $\Delta(V, H_s)$ of weights of H_s on V (Definition 3.2). Then

$$I(\Phi) \otimes V \simeq \sum_{\delta \in \Delta(V, H_s)} I(\Phi + \delta)$$

as virtual representations.

In the coherent continuation formula mentioned last, both sides are actual representations of G. Nevertheless, the equality may be true only as virtual representations: composition factors may be arranged differently on the two sides. The simplest example has $I(\Phi)$ the nonspherical unitary principal series for $SL_2(\mathbb{R})$ (with $J(\Phi) = I(\Phi)$ the sum of two limits of discrete series), and V the two-dimensional representation. The trivial representation appears twice as a quotient on the right side of the formula, but not at all as a quotient on the left. Part of the difficulty is that the Borel subgroup B_s is dominant only for one of the two parameters $\Phi + \delta$ (with δ a weight of V).

In this proposition, in contrast to Definition 3.3, the positive roots are really those in B_s (and not their negatives).

To understand the K-types of the principal series $I(\Phi)$, we need to understand which representations of K can contain the character $\phi|_{T_s}$ of T_s. Since we understand representations of K in terms of their highest weights, this amounts to understanding the relationship between T_s and the Cartan subgroup T_f of K. In order to discuss this, we need the Knapp-Stein R-group for principal series.

Definition 7.4. In the setting of Definition 7.2, fix a parameter $\Phi = (\phi, \emptyset) \in P^s_G(H_s)$. Define

$$W(G, H_s)^\Phi = \{w \in W(G, H_s) \mid w \cdot \Phi = \Phi\},$$

the stabilizer of Φ (that is, of the character ϕ) in the real Weyl group. The set of *good roots* for A_s in G is

$$\Delta_\Phi(G, A_s) = \{\alpha \in \Delta(G, A_s) \mid \phi \text{ is trivial on } \alpha^\vee\}.$$

Since α is just the restriction of a root, the meaning of the condition "ϕ is trivial on $\alpha^\vee$" requires some explanation. First, we can construct (not uniquely) a homomorphism from $SL_2(\mathbb{R})$ into G using H_s and the root α. (In the special case of a real root, this homomorphism is described a little more precisely in (7.6a) below.) Restricting this homomorphism to the diagonal torus $\mathbb{R}^\times$ of $SL_2(\mathbb{R})$ gives a homomorphism

$$\alpha^\vee \colon \mathbb{R}^\times \to A_s$$

which *is* uniquely determined; and the first requirement we want to impose is that $\phi \circ \alpha^\vee$ is trivial. When α is not the restriction of a real root, then the simple real rank one subgroup of G corresponding to α is covered by $SU(2,1)$ or $SL_2(\mathbb{C})$. In each of these real groups the maximally split Cartan is $\mathbb{C}^\times$, so the real coroot $\alpha^\vee$ extends to

$$\alpha^\vee_{\mathbb{C}} \colon \mathbb{C}^\times \to H_s.$$

(The extension is well-defined up to composition with complex conjugation.) In these cases (to call α good) we impose the requirement that $\phi \circ \alpha^\vee_{\mathbb{C}}$ is trivial.

The good roots form a subroot system of the restricted roots, and the corresponding Weyl group

$$W_0(G, H_s)^\Phi = W(\Delta_\Phi(G, A_s))$$

is a normal subgroup of $W(G, H_s)^\Phi$. The *R-group of* Φ is by definition the quotient

$$R(\Phi) = W(G, H_s)^\Phi / W_0(G, H_s)^\Phi.$$

Proposition 7.5. *Suppose we are in the setting of Definition 7.4.*

1. *The irreducible constituents of the Langlands quotient representation $J(\Phi)$ all occur with multiplicity one. There is a natural simply transitive action of the character group $\Pi^*(R(\Phi))$ on these constituents.*
2. *Write Φ_d for the discrete part of Φ (the restriction of ϕ to T_s, extended to be trivial on $A_{s,0}$). Then there is a natural inclusion $W(G, H_s)^\Phi \hookrightarrow W(G, H_s)^{\Phi_d}$, which induces an inclusion $R(\Phi) \hookrightarrow R(\Phi_d)$.*

(3) *Each irreducible summand of* $I(\Phi_d) = J(\Phi_d)$ *contains a unique lowest* K*-type of* $I(\Phi_d)$. *There is a natural simply transitive action of* $\Pi^*(R(\Phi_d))$ *on this set* $A(\Phi_d)$ *of* K*-types.*
(4) *The sets* $A(\Phi_d)$ *partition a certain subset of* $\Pi^*(K)$. *We have*

$$A(\Phi_d) = A(\Psi_d) \iff \Psi_d \in W(G, H_s) \cdot \Phi_d.$$

In order to understand the lowest K-types of these principal series representations, we must therefore describe $A(\Phi_d)$ as a set of (extended) highest weights. The first step is to relate T_f (where extended highest weights are supposed to live) to T_s (where Φ_d lives). So choose a set

$$\beta_1, \ldots, \beta_r \tag{7.6a}$$

of strongly orthogonal real roots of H_s in G, of maximal cardinality. Attached to each real root α we can find an algebraic group homomorphism defined over $\mathbb{R}$

$$\psi_\alpha \colon \mathbf{SL}_2 \to \mathbf{G}, \tag{7.6b}$$

in such a way that

(1) ψ_α carries the diagonal Cartan subgroup of SL_2 into H_s, and the upper triangular subgroup into the α root subgroup; and
(2) $\psi_\alpha({}^t g^{-1}) = \theta(\psi_\alpha(g))$.

Such a homomorphism ψ_α is unique up to conjugation in $\mathbf{SL}_2$ by

$$\begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} \tag{7.6c}$$

In particular, the element

$$m_\alpha = \psi_\alpha \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix} = \alpha^\vee(-1) \in T_s \tag{7.6d}$$

is well-defined.

Since the roots $\{\beta_i\}$ are strongly orthogonal and real, the groups $\psi_{\beta_i}(SL_2)$ commute with each other and with $T_{s,0}$.

Definition 7.7. A shifted (limit) parameter $\Phi = (\phi, \emptyset) \in P^s_G(H_s)$ (Definition 5.1) is called *final* if for every real root α of H_s in G,

$$\text{either} \quad \langle d\phi, \alpha^\vee \rangle \neq 0 \quad \text{or} \quad \phi(m_\alpha) = 1.$$

(If Φ is discrete, the requirement is that $\phi(m_\alpha) = 1$ for every real root α.) We write $P_G^{s,finlim}(H_f)$ for the set of final limit parameters for H_f.

For the remainder of this section we consider the important special case

$$G \text{ split}, \quad \Phi \in P^{s,finlim}_{G,d} \text{ discrete final limit parameter.} \tag{7.8a}$$

This is the same thing as a character

$$\phi \in \Pi^*(T_s), \qquad \phi(m_\alpha) = 1 \qquad (\text{all } \alpha \in \Delta(G, H_s)). \tag{7.8b}$$

Here m_α is the element of order 2 defined in (7.6d).

Proposition 7.9. *In the setting (7.8), the principal series representation $I(\Phi)$ has a unique lowest K-type $\tau(\Phi)$. This representation of K is trivial on the identity component of $G^{der} \cap K$, with G^{der} the derived group of G. The map $\Phi \mapsto \tau(\Phi)$ is an injection from discrete final limit parameters for H_s into $\Pi^*(K)$. (The group $W(G, H_s)$ acts trivially on discrete final limit parameters for H_s, so we could also say "$W(G, H_s)$ orbits of discrete final limit parameters...") The image consists precisely of those irreducible representations of K which are trivial on $(G^{der} \cap K)_0$. These are the representations of K for which any highest weight μ has the property that $\langle \mu, \beta^\vee \rangle = 0$ for any root $\beta \in \Delta(G, T_f)$.*

8. Parametrizing extended weights

We have a computer-friendly parametrization of characters of T_f available from Proposition 6.3:

$$\Pi^*(T_f) = \Pi^*_{alg}(\mathbf{H}_f)/(1 - \theta_f)\Pi^*_{alg}(\mathbf{H}_f). \tag{8.1}$$

Here I have written θ_f for the action of θ on the fundamental Cartan $\mathbf{H}_f$; this is in some sense a compromise with the software point of view, in which the Cartan is always fixed and only the Cartan involution is changing. In order to parametrize representations of K, we need not just a character μ of T_f, but an extension $\widetilde{\mu}$ of that character to a slightly larger group. We therefore need a computer-friendly parametrization of such extensions, and this section seeks to provide such a parametrization.

We begin with the special case of Proposition 7.9: assume first of all that

$$G \text{ is split, with } \theta\text{-stable split Cartan } H_s. \tag{8.2a}$$

This means that

$$\theta_s(\alpha) = -\alpha \qquad (\alpha \in \Delta(G, H_s)). \tag{8.2b}$$

On the fundamental Cartan, it is equivalent to assume that there is an orthogonal set of imaginary roots

$$\{\gamma_1, \dots, \gamma_m\} \subset \Delta_{im,n}(G, H_f) \tag{8.2c}$$

spanning the root system $\Delta(G, T_f)$, and for which the successive Cayley transforms are defined: γ_i is noncompact if and only if it is preceded by an even number of γ_j to which it is not strongly orthogonal. The successive Cayley transforms identify $\Pi^*_{alg}(\mathbf{H}_s)$ with $\Pi^*_{alg}(\mathbf{H}_f)$ in such a way that

$$\theta_f = \prod_{j=1}^{m} s_{\gamma_j} \circ \theta_s. \tag{8.2d}$$

In particular, θ_f acts on the roots by minus the product of the reflections in the γ_j. We also want to impose a condition on the weights of T_f that we consider: we look only at

$$\{\mu \in \Pi^*(T_f) \mid \langle \mu, \alpha^\vee \rangle = 0 \quad (\alpha \in \Delta(G, T_f)\}. \tag{8.2e}$$

This restriction is equivalent (since the $\gamma_j^\vee$ span $\mathfrak{t}_f$ modulo the center of $\mathfrak{g}$) to

$$\{\mu \in \Pi^*(T_f) \mid \langle \mu, \gamma_j^\vee \rangle = 0 \quad (j = 1, \ldots, m)\}. \tag{8.2f}$$

We call these weights "G-spherical," and write the set of them as

$$P_{K,G-sph}(T_f). \tag{8.2g}$$

The corresponding set of extended weights (Definition 4.11) is written

$$P^e_{K,G-sph}(T_f). \tag{8.2h}$$

Here is an explanation of the terminology.

Lemma 8.3. *Suppose G is split, and (τ, E) is an irreducible representation of K having extremal weight $\mu \in \Pi^*(T_f)$ (Corollary 4.13). Recall from* (2.5) *the decomposition $\mathfrak{g} = \mathfrak{k} + \mathfrak{s}$ of the Lie algebra into the $+1$ and -1 eigenspaces of θ. Then E can be extended to a $(\mathfrak{g}, K)$-module on which $\mathfrak{s}$ acts by zero if and only if μ is G-spherical. That is, μ is G-spherical if and only if any irreducible representation of K of extremal weight μ is trivial on the ideal $[\mathfrak{s}, \mathfrak{s}] \subset \mathfrak{k}$.*

Recall from Proposition 7.9 the discrete final limit parameters for H_s:

$$P^{s,finlim}_{G,d}(H_s) = \{\Phi \in \Pi^*(T_s) \mid \phi(m_\alpha) = 1 \quad (\alpha \in \Delta(G, H_s))\}. \tag{8.4}$$

The element $m_\alpha \in T_s$ was defined in (7.6d).

Proposition 8.5. *Suppose G is split. The following sets of representations of K (called G-spherical) are all the same.*

1. *Representations having a G-spherical extremal weight (cf.* (8.2)*).*
2. *Representations trivial on $[\mathfrak{s}, \mathfrak{s}] \subset \mathfrak{k}$ (cf. Lemma 8.3).*
3. *Lowest K-types of principal series with discrete final limit parameters (Proposition 7.9).*

This establishes a bijection $\Phi \mapsto \tau(\Phi)$ from discrete final limit parameters for H_s to G-spherical representations of K. Taking into account the parametrization of $\Pi^*(K)$ by extended highest weights gives bijections

$$\begin{array}{ccccc} P^e_{K,G-sph}(T_f) & \longleftrightarrow & \Pi^*_{G-sph}(K) & \longleftrightarrow & P^{s,finlim}_{G,d}(H_s), \\ \widetilde{\mu}(\Phi) & \longleftrightarrow & \tau(\Phi) & \longleftrightarrow & \Phi \end{array}$$

Proposition 6.3 provides a computer-friendly parametrization of the discrete final limit parameters of H_s:

$$\begin{aligned} P^{s,finlim}_{G,d}(H_s) = \{\lambda \in \Pi^*_{alg}(\mathbf{H}_s) \mid \langle \lambda, \alpha^\vee \rangle \in 2\mathbb{Z} \ (\alpha \in \Delta(G, H_s)\} \\ /(1 - \theta_s)\Pi^*_{alg}(\mathbf{H}_s). \end{aligned} \tag{8.6}$$

The condition on λ in the numerator arises from the fact that $\lambda(m_\alpha) = (-1)^{\langle \lambda, \alpha^\vee \rangle}$. The description of characters of T_f in (8.1) immediately specializes to a description of the G-spherical characters:

$$\begin{aligned} P_{K,G-sph}(T_f) = \{\lambda_0 \in \Pi^*_{alg}(\mathbf{H}_f) \mid \langle \lambda_0, \gamma_j^\vee \rangle = 0 \ (j = 1, \ldots, m)\} \\ /(1 - \theta_f)\Pi^*_{alg}(\mathbf{H}_f). \end{aligned} \tag{8.7}$$

An equivalent description of the numerator is

(8.8)
$$P_{K,G-sph}(T_f) = \{\lambda_0 \in \Pi^*_{alg}(\mathbf{H}_f) \mid \langle \lambda_0, \alpha^\vee + \theta\alpha^\vee \rangle = 0 \ (\alpha \in \Delta(G, H_f))\} \\ /(1-\theta_f)\Pi^*_{alg}(\mathbf{H}_f).$$

What we are going to do is describe a natural surjective map

$$P^{s,finlim}_{G,d}(H_s) \twoheadrightarrow P_{K,G-sph}(T_f) \tag{8.9a}$$

Since the first set is identified by Proposition 8.5 with G-spherical extended weights, this amounts to

$$P^e_{K,G-sph}(H_s) \twoheadrightarrow P_{K,G-sph}(T_f) \tag{8.9b}$$

This will be precisely the map $\widetilde{\mu} \mapsto \mu$ restricting an extended weight to T_f.

In order to describe the map, it is helpful to introduce some auxiliary lattices attached to a maximal torus $\mathbf{H} \subset \mathbf{G}$. To every $\lambda \in \Pi^*_{alg}(\mathbf{H})$ one can attach a $\mathbb{Z}/2\mathbb{Z}$-grading of the coroot system, by

$$\epsilon(\lambda) \colon \Delta^\vee \to \{0,1\}, \qquad \epsilon(\lambda)(\alpha^\vee) = \langle \lambda, \alpha^\vee \rangle \pmod 2. \tag{8.10}$$

The map of (8.9) is closely related to these gradings. In terms of the gradings, we define

$$\begin{aligned} \Pi^*_{ev}(\mathbf{H}) &= \{\lambda \in \Pi^*_{alg}(\mathbf{H}) \mid \epsilon(\lambda) = 0\} \\ R(\mathbf{H}) &= \mathbb{Z}\Delta(\mathbf{G},\mathbf{H}) = \text{root lattice} \\ R_{ev}(\mathbf{H}) &= \{\phi \in R(\mathbf{H}) \mid \epsilon(\phi) = 0\} \\ \Pi^*_R(\mathbf{H}) &= \{\lambda \in \Pi^*_{alg}(\mathbf{H}) \mid \exists \phi \in R \text{ with } \epsilon(\lambda) = \epsilon(\phi)\} \\ &= \epsilon^{-1}(\epsilon(R)). \end{aligned} \tag{8.11}$$

We can immediately use these definitions to rewrite (8.6) (which we know is parametrizing G-spherical representations of K) as

$$P^{s,finlim}_{G,d}(H_s) = \Pi^*_{ev}(\mathbf{H}_s)/(1-\theta_s)\Pi^*_{alg}(\mathbf{H}_s). \tag{8.12}$$

To analyze (8.7) in a parallel way, we begin with a lemma.

Lemma 8.13. *Suppose $\{\gamma_1, \dots, \gamma_m\}$ is a maximal orthogonal set of roots in $\Delta(\mathbf{G},\mathbf{H})$. Suppose ϵ is a grading of the coroots with values in $\mathbb{Z}/2\mathbb{Z}$, and that*

$$\epsilon(\gamma_j^\vee) = 0, \qquad (j = 1, \dots, m).$$

Then there is a subset $A \subset \{\gamma_1, \dots, \gamma_m\}$ so that ϵ is equal to the grading ϵ_A defined by the sum of the roots in this subset:

$$\epsilon(\alpha^\vee) \equiv \sum_{\gamma \in A} \langle \gamma, \alpha^\vee \rangle \pmod 2$$

I am grateful to Jeff Adams and Becky Herb for providing a (fairly constructive) proof of this lemma, which appears in the next section.

Lemma 8.14. *In the setting of* (8.2), *consider* $P_{K,G-sph}(T_f)$ *as described by* (8.7). *According to Lemma 8.13, any element* $\lambda_0 \in \Pi^*_{alg}(\mathbf{H}_f)$ *representing an element of* $P^*_{K,G-sph}$ *(that is, orthogonal to all* γ_j*) must belong to* $\Pi^*_R(\mathbf{H}_f)$*. This defines a map*

$$P_{K,G-sph}(T_f) \to \Pi^*_R(\mathbf{H}_f)/[(1-\theta_f)\Pi^*(\mathbf{H}_f) + R].$$

This map is an isomorphism.

Proof. We first prove surjectivity. Suppose $\lambda \in \Pi^*_R(\mathbf{H}_f)$; we need to show that the class of λ is in the image of our map. Since the denominator includes the root lattice R, we may replace λ by some $\lambda_1 = \lambda + \psi$ with $\psi \in R$ and $\lambda_1 \in \Pi^*_{ev}(\mathbf{H}_f)$. Since λ_1 takes even values on all coroots, we can define

$$\lambda_0 = \lambda_1 - \sum_{j=1}^{m} (\langle \lambda_1, \gamma_j^\vee \rangle / 2)\gamma_j$$

(again modifying λ_1 by an element of R). Clearly λ_0 is orthogonal to all the γ_j, so it represents a class in $P_{K,G-sph}(T_f)$ mapping to the class of λ.

Next we prove injectivity. Suppose λ_0 represents a class in $P_{K,G-sph}(T_f)$ mapping to zero; that is, that

$$\lambda_0 = \xi - \theta_f(\xi) + \psi, \tag{8.15a}$$

with $\psi \in R$.

The term $\xi - \theta_f(\xi)$ is orthogonal to all imaginary roots, including the γ_j, so $\langle \psi, \gamma_j^\vee \rangle = 0$ for all j. By (8.2d), this is equivalent to $\theta_f(\psi) = -\psi$. Choose a θ-stable system of positive roots. List the simple imaginary roots as $\alpha_1, \ldots, \alpha_r$, and the simple complex roots as $\beta_1, \theta_f\beta_1, \ldots, \beta_s, \theta_f\beta_s$. Because these $r + 2s$ roots are a basis for the root lattice R, it is clear that the -1 eigenspace of θ_f on R has as basis the s elements $\beta_i - \theta_f\beta_i$. In particular, ψ is a sum of these elements, so

$$\psi \in (1-\theta_f)R \subset (1-\theta_f)\Pi^*(\mathbf{H}_f). \tag{8.15b}$$

Now (8.15) shows that λ_0 represents 0 in $P_{K,G-sph}(T_f)$, proving injectivity. □

Proposition 8.16. *Suppose that we use a Cayley transform to identify* $\Pi^*_{alg}(\mathbf{H}_f)$ *with* $\Pi^*_{alg}(\mathbf{H}_s)$*. This identification is unique up to the action of* $W(G, H_s)$*. This identification sends* $P_{K,G-sph}(T_f)$ *(as described in Lemma 8.14) to*

$$\Pi^*_R(\mathbf{H}_s)/[(1-\theta_s)\Pi^*(\mathbf{H}_s) + R].$$

The Weyl group $W(G, H_s)$ *acts trivially on this quotient, so the identification with* $P_{K,G-sph}(T_f)$ *is well-defined. This quotient in turn may be identified (by the natural inclusion) with*

$$\Pi^*_{ev}(\mathbf{H}_s)/[(1-\theta_s)\Pi^*(\mathbf{H}_s) + R_{ev}].$$

In this way $P_{K,G-sph}(T_f)$ *is naturally identified as a quotient of*

$$\Pi^*_{ev}(\mathbf{H}_s)/(1-\theta_s)\Pi^*(\mathbf{H}_s) \simeq P^{s,finlim}_{G,d}(H_s).$$

Taking into account the identification of Proposition 8.5, we have given a surjective map

$$P^e_{K,G-sph}(T_f) \to P_{K,G-sph}(T_f)$$

from extended (extremal) weights of G-spherical representations of K, to (extremal) weights. This map is just restriction to T_f.

Proof. The action of Weyl group elements on algebraic characters is by addition of elements of the root lattice R. Since we are dividing by R in the first formula, this shows that the Weyl group action is trivial on the quotient, as claimed. Since θ_s and θ_f differ by the product of reflections in the γ_j, it also follows that

$$(1-\theta_s)\Pi^*(\mathbf{H}_s) + R = (1-\theta_f)\Pi^*(\mathbf{H}_f) + R$$

(using the Cayley transform identifications). This provides the first formula of the proposition. The map from the second formula to the first comes from the inclusion of Π^*_{ev} in Π^*_R. Surjectivity is immediate from the definitions, and injectivity is straightforward.

For the last claim, recall that $T_f \subset G^\sharp = G_0 \cdot Z(G)$; and furthermore $G_0 \subset G_{der,0} \cdot Z(G)$. It follows that any G-spherical character of T_f is determined by its restriction to $Z(G) \cap K$. So we only need to verify that the mappings described in the proposition do not change central character. Since at every stage we are simply adding and subtracting roots, or changing Cartans by an inner automorphism of $\mathbf{G}$, this is clear. □

It is now more or less a routine matter to parametrize extensions of a general weight $\mu \in \Pi^*(T_f)$. To do this, we first choose positive roots

$$\Delta^+_c \subset \Delta(K, T_f) \tag{8.17a}$$

in such a way that μ is dominant. This defines $2\rho_c$ as in (4.3), the sum of the positive compact roots, and the large Cartan $T_{fl} \supset T_f$ of Definition 3.4. Define $T_{fl}(\mu)$ as in Definition 4.7, the stabilizer of μ in T_{fl}. We now choose a Levi subgroup $L \supset H_f$ large enough so that

$$T_{fl}(\mu) \subset L, \tag{8.17b}$$

(so that we can compute extensions of μ inside L), but small enough that μ is L-spherical:

$$\langle \mu, \gamma^\vee \rangle = 0, \qquad (\gamma \in \Delta(L, T_f)). \tag{8.17c}$$

Finally, we require that

$$L \text{ is split, with split Cartan subgroup } H \tag{8.17d}$$

(so that we can apply Proposition 8.16).

Here is one way to find such an L. Define

$$R(K,\mu) = T_{fl}(\mu)/T_f \subset W(G, H_f) \tag{8.17e}$$

(Proposition 4.4). Recall from Proposition 4.4 the set of strongly orthogonal noncompact imaginary roots

$$P(\mu) = \{\beta_1, \ldots, \beta_s\} \tag{8.17f}$$

with the property that

$$P(\mu) \cup -P(\mu) = \{\beta \in \Delta(G, T_f) \mid \langle \mu, \beta^\vee \rangle = \langle 2\rho_c, \beta^\vee \rangle = 0\}. \tag{8.17g}$$

(We have changed notation a little from Proposition 4.4, by renumbering the roots β_i.) These are the roots of a split Levi subgroup $L_{small}(\mu) \supset H_f$, locally isomorphic to a product of s copies of $SL_2(\mathbb{R})$ and an abelian group. Obviously μ is L_{small}-spherical, and Proposition 4.4 shows that $T_{fl}(\mu) \subset L_{small}$. (We could even have accomplished this using the slightly smaller set $P_{ess}(\mu)$ described in (4.5).) I will describe another (larger) natural choice for L in Section 13.

9. Proof of Lemma 8.13

I repeat that the following argument is due to Adams and Herb. The notation will be a little less burdensome if we interchange roots and coroots, proving instead

Lemma 9.1. *Suppose $S = \{\gamma_1, \ldots, \gamma_m\}$ is a maximal orthogonal set of roots in $\Delta(\mathbf{G}, \mathbf{H})$. Suppose ϵ is a grading of the roots with values in $\mathbb{Z}/2\mathbb{Z}$, and that*

$$\epsilon(\gamma_j) = 0, \qquad (j = 1, \ldots, m).$$

Then there is a subset $A \subset S$ so that ϵ is equal to the grading ϵ_A defined by the sum of the coroots in this subset:

$$\epsilon(\alpha) \equiv \sum_{\gamma \in A} \langle \alpha, \gamma^\vee \rangle \pmod 2$$

Proof. We proceed by induction on m. If $m = 0$ then Δ is empty and the lemma is trivial. So suppose $m > 0$. Define

$$\Delta_{m-1} = \{\beta \in \Delta \mid \langle \beta, \gamma_m^\vee \rangle = 0\}, \qquad S_{m-1} = \{\gamma_1, \ldots, \gamma_{m-1}\}. \tag{9.2a}$$

the roots orthogonal to γ_m. Clearly S_{m-1} is a maximal orthogonal set of roots in Δ_{m-1}, and by hypothesis ϵ is trivial on S_{m-1}. By inductive hypothesis there is a subset $A_{m-1} \subset S_{m-1}$ so that

$$\epsilon(\alpha) \equiv \sum_{\gamma \in A_{m-1}} \langle \alpha, \gamma^\vee \rangle \pmod 2 \qquad (\alpha \in \Delta_{m-1}). \tag{9.2b}$$

Define

$$A' = A_{m-1} \subset S, \qquad A'' = A_{m-1} \cup \{\gamma_m\}. \tag{9.2c}$$

It follow immediately from (9.2a) that

(9.2d) the gradings $\epsilon_{A'}$ and $\epsilon_{A''}$ agree with ϵ on Δ_{m-1} and $\pm\gamma_m$.

We are going to show that taking A equal to one of the two subsets A' and A'' satisfies the requirement of Lemma 9.1 that $\epsilon_A = \epsilon$. To do that we need a lemma.

Lemma 9.3. *In the setting of* (9.2)*, there are two mutually exclusive possibilities.*

(1) *The coroot $\gamma_m^\vee$ takes even values on every root:*
$$\langle \alpha, \gamma_m^\vee \rangle \in 2\mathbb{Z}, \qquad (\alpha \in \Delta).$$
In this case the integer span of γ_m and Δ_{m-1} contains Δ.

(2) *There is a root $\alpha_1 \in \Delta$ so that $\langle \alpha_1, \gamma_m^\vee \rangle$ is odd. In this case the integer span of γ_m, Δ_{m-1}, and α_1 contains Δ.*

In case (1), both of the gradings $\epsilon_{A'}$ and $\epsilon_{A''}$ are equal to ϵ. In case (2), exactly one of them is equal to ϵ.

The last statement of this lemma completes the proof of Lemma 9.1 □

Proof of Lemma 9.3. Since Δ is reduced, the possible values for $\langle \alpha, \gamma_m^\vee \rangle$ are 0, ± 1, ± 2, and ± 3. Obviously the pairings of roots with $\gamma_m^\vee$ are either all even or not all even. Consider the first possibility. If α is any root, then $\langle \alpha, \gamma_m^\vee \rangle$ is equal to 0, 2, or -2. In the first case, α belongs to Δ_{m-1}. In the second, either $\alpha = \gamma_m$ or $\alpha - \gamma_m$ is a root in Δ_{m-1}; so in either case α is in the integer span of Δ_{m-1} and γ_m. In the third case, either $\alpha = -\gamma_m$ or $\alpha + \gamma_m$ is a root in Δ_{m-1}, and we get the same conclusion.

Consider now the second possibility, and fix α_1 having an odd pairing with the coroot $\gamma_m^\vee$. Possibly after replacing α_1 by $\pm\alpha_1 \pm \gamma_m$, we may assume that

$$\langle \alpha_1, \gamma_m^\vee \rangle = 1. \tag{9.4a}$$

Suppose $\beta \in \Delta$; we want to write

$$\beta = p\gamma_m + q\alpha_1 + \delta \qquad (p, q \in \mathbb{Z}, \delta \in \Delta_{m-1}). \tag{9.4b}$$

If $\langle \beta, \gamma_m^\vee \rangle$ is even, then we can achieve (9.4b) as in the first possibility (not using α_1). So we may assume that $\langle \beta, \gamma_m^\vee \rangle$ is odd. Perhaps after replacing β by $-\beta$, we may assume that $\langle \beta, \gamma_m^\vee \rangle$ is 1 or 3. Perhaps after replacing β by $\beta - \gamma_m$, we may assume that

$$\langle \beta, \gamma_m^\vee \rangle = 1. \tag{9.4c}$$

If $\langle \beta, \alpha_1^\vee \rangle > 0$, then $\delta = \beta - \alpha_1$ is a root in Δ_{m-1}, and (9.4b) follows. So we may assume

$$\langle \beta, \alpha_1^\vee \rangle \le 0. \tag{9.4d}$$

Now it follows from (9.4c) that $\beta - \gamma_m$ is a root taking the value -1 on the coroot $\gamma_m^\vee$. Furthermore

$$\langle \beta - \gamma_m, \alpha_1^\vee \rangle = \langle \beta, \alpha_1^\vee \rangle - \langle \gamma_m, \alpha_1^\vee \rangle.$$

The first term is non-positive by (9.4d) and the second strictly negative by (9.4a). It follows that $\delta = \beta - \gamma_m + \alpha_1$ is a root or zero. In the first case

$\delta \in \Delta_{m-1}$ (by (9.4a) and (9.4c)), so in any case $\beta = \gamma_m - \alpha_1 + \delta$ has the form required in (9.4b).

The last assertion of the lemma follows from (9.2d) in the first case. In the second case, the hypothesis on α_1 shows that $\epsilon_{A'}(\alpha_1) = 1 - \epsilon_{A''}(\alpha_1)$. So exactly one of these two gradings agrees with ϵ on α_1, and the assertion now follows from (9.2d). □

10. Highest weights for K and θ-stable parabolic subalgebras

The description of $\Pi^*(K)$ in terms of $\Pi^*(T_f)$ in Theorem 4.10 was mediated by the (θ-stable) Borel subalgebra $\mathfrak{b}_c$ of $\mathfrak{k}$. In this section we consider more general correspondences of the same sort, between representations of K and those of certain Levi subgroups $L \cap K$. These correspondences will be used in Section 11 to describe the lowest K-types of standard limit representations.

We therefore fix a θ-stable parabolic subalgebra

(10.1a) $$\mathfrak{q} \supset \mathfrak{u}$$

of $\mathfrak{g}$; here $\mathfrak{u}$ denotes the nil radical of $\mathfrak{q}$. Recall that σ is the antiholomorphic involution of $\mathfrak{g}$ defining the real form. We assume also that $\mathfrak{q}$ is opposite to $\sigma\mathfrak{q}$; equivalently,

(10.1b) $$\mathfrak{q} \cap \sigma\mathfrak{q} = \mathfrak{l}$$

is a Levi subalgebra of $\mathfrak{q}$. Clearly $\sigma\mathfrak{l} = \mathfrak{l}$, so $\mathfrak{l}$ is defined over $\mathbb{R}$. It is also preserved by the Cartan involution θ. We write L for the corresponding real Levi subgroup of G. It is not difficult to show that

(10.1c) $$L = \{g \in G \mid \mathrm{Ad}(g)(\mathfrak{q}) = \mathfrak{q}\}.$$

Like G, L is the group of real points of a connected reductive algebraic group, and $\theta|_L$ is a Cartan involution for L; so $L \cap K$ is a maximal compact subgroup of L. Choose a θ-stable fundamental Cartan subgroup H_f for L as in Definition 3.1. The use of the same notation as for G is justified by the following result.

Proposition 10.2. *In the setting of* (10.1)*, the fundamental Cartan H_f for L is also a fundamental Cartan subgroup for G.*

Now let us fix a choice of positive roots

(10.3a) $$\Delta_c^+(L) \subset \Delta(L \cap K, T_f).$$

Define a maximal nilpotent subalgebra of $\mathfrak{l} \cap \mathfrak{k}$ by the requirement

(10.3b) $$\Delta(\mathfrak{n}_c(\mathfrak{l}), T_f) = -\Delta_c^+(L)$$

as in Definition 3.3, so that

(10.3c) $$\mathfrak{b}_c(\mathfrak{l}) = \mathfrak{t}_f + \mathfrak{n}_c(\mathfrak{l})$$

is a Borel subalgebra of $\mathfrak{l} \cap \mathfrak{k}$. We now get a Borel subalgebra of $\mathfrak{k}$ by defining

(10.3d) $$\mathfrak{b}_c = \mathfrak{b}_c(\mathfrak{l}) + \mathfrak{u} \cap \mathfrak{k}, \qquad \mathfrak{n}_c = \mathfrak{n}_c(\mathfrak{l}) + \mathfrak{u} \cap \mathfrak{k}$$

The corresponding positive root system is

$$\Delta_c^+ = \Delta_c^+(L) \cup -\Delta(\mathfrak{u} \cap \mathfrak{k}, T_f). \tag{10.3e}$$

In the setting of (10.1), we will refer to such choices of fundamental Cartan and positive root systems as *compatible.*

The identifications

$$R(K) \simeq W(K, T_f)/W(K_0, T_{f,0}),$$
$$R(L \cap K) \simeq W(L \cap K, T_f)/W((L \cap K)_0, T_{f,0})$$

and the obvious inclusion of Weyl groups define a natural map

$$R(L \cap K) \to R(K) \tag{10.3f}$$

It is not difficult to show that this map is an inclusion.

We will be using Proposition 4.4 for both L and G, so we need to choose the sets $P(K)$ and $P(L \cap K)$ defined there in a compatible way. It turns out that we need not have $P(L \cap K) \subset P(K)$. For example, if L is locally isomorphic to $SL_2(\mathbb{R})$, then $P(L \cap K)$ is always a single root; but this root will be in P only under special circumstances. Because of this fact, it is not clear what "compatible" ought to mean for P.

Proposition 10.4. *In the setting of* (10.3)*, suppose $A \subset P(L \cap K)$ is such that $w_A \in R(L \cap K)$ (cf. Proposition 4.4). Then $A \subset P(K) \cup -P(K)$, and $w_A \in R(K)$.*

It follows from this proposition that

$$P_{ess}(L \cap K) \subset P_{ess}(K) \cup -P_{ess}(K).$$

(notation as in (4.5)). It therefore makes sense to require

$$P_{ess}(L \cap K) \subset P_{ess}(K) \tag{10.5}$$

as the compatibility requirement between $P(K)$ and $P(L \cap K)$.

Theorem 10.6. *Suppose we are in the setting of* (10.1)*; choose compatible positive root systems $\Delta^+(L \cap K, T_f) \subset \Delta^+(K, T_f)$ as in* (10.3)*, and compatible sets $P(K)$ and $P(L \cap K)$ as in* (10.5)*. Suppose that $(\tau_{\mathfrak{q}}, E_{\mathfrak{q}})$ is an irreducible representation of $L \cap K$, of extended highest weight $\widetilde{\mu_{\mathfrak{q}}}$ (Definition 4.11). Define a generalized Verma module*

$$M_K(\tau_{\mathfrak{q}}) = U(\mathfrak{k}) \otimes_{\mathfrak{q} \cap \mathfrak{k}} E_{\mathfrak{q}},$$

which is a $(\mathfrak{k}, L \cap K)$-module. Write $\mathbb{L}_k$ for the kth Bernstein derived functor carrying $(\mathfrak{k}, L \cap K)$-modules to K-modules, and $S = \dim \mathfrak{u} \cap \mathfrak{k}$. Define

$$\zeta_{\mathfrak{k}} = d\mu_{\mathfrak{q}} + \rho(\mathfrak{u} \cap \mathfrak{k}) + \rho(\mathfrak{l} \cap \mathfrak{k}) \in \mathfrak{t}_f^*,$$

a weight parametrizing the infinitesimal character of the generalized Verma module $M_K(\tau_{\mathfrak{q}})$.

1. *If $\zeta_{\mathfrak{k}}$ vanishes on any coroot of K, then $\mathbb{L}_k(M_K(\tau_{\mathfrak{q}})) = 0$ for every k. Suppose henceforth that $\zeta_{\mathfrak{k}}$ is regular for K.*

(2) *If $\zeta_{\mathfrak{k}}$ is dominant for $\Delta^+(K,T_f)$, then the generalized Verma module $M_K(\tau_{\mathfrak{q}})$ is irreducible. In this case $\mathbb{L}_k(M_K(\tau_{\mathfrak{q}})) = 0$ for $k \neq S$. The K representation $\mathbb{L}_S(M_K(\tau_{\mathfrak{q}}))$ is generated by the highest weight*

$$\mu = \mu_{\mathfrak{q}} \otimes 2\rho(\mathfrak{u} \cap \mathfrak{k}).$$

The corresponding extended group $T_{fl}(\mu)$ (Proposition 4.4) contains $T_{fl}(\mu_{\mathfrak{q}})$ (via the inclusion (10.3f)), and the representation of $T_{fl}(\mu)$ on the μ weight space is

$$\operatorname{Ind}_{T_{fl}(\mu_{\mathfrak{q}})}^{T_{fl}(\mu)} \widetilde{\mu};$$

here $\widetilde{\mu}$ denotes the representation $\widetilde{\mu_{\mathfrak{q}}} \otimes 2\rho(\mathfrak{u}\cap\mathfrak{k})$ of $T_{fl}(\mu_{\mathfrak{q}})$. In particular, if the R-groups $R(L\cap K, \mu_{\mathfrak{q}})$ and $R(K,\mu)$ are equal (under the inclusion (10.3f)), then $\mathbb{L}_S(M_K(\tau_{\mathfrak{q}}))$ is the irreducible representation of K of extended highest weight $\widetilde{\mu}$ (Definition 4.11).

(3) *In general, let $w \in W(K_0, T_{f,0})$ be the unique element so that $\zeta_{\mathfrak{k}}$ is dominant for $w(\Delta^+(K,T_f))$. Define $k_0 = k(\mu_{\mathfrak{q}}) = S - l(w)$ (with $l(w)$ the length of the Weyl group element w). The number k_0 is equal to the cardinality of the set of roots*

$$B = \{\alpha \in \Delta(\mathfrak{u}, T_f) \cap -w\Delta^+(K,T_f)\}.$$

and

$$\mu = \mu_{\mathfrak{q}} + \sum_{\alpha\in B} \alpha;$$

this weight has differential $\zeta_{\mathfrak{k}} - w\rho_c$. In this case $\mathbb{L}_k(M_K(\tau_{\mathfrak{q}})) = 0$ for $k \neq k_0$. The K representation $\mathbb{L}_{k_0}(M_K(\tau_{\mathfrak{q}}))$ includes the extremal weight μ.

11. Standard representations, limits, and continuations

Throughout this section we will fix a θ-stable real Cartan subgroup H as in (6.1), with compact part $T = H^\theta$ and split part $A = H^{-\theta}$. (Recall that in our notation A is usually not connected: it is the identity component A_0 that is the vector group of traditional notation, figuring in the direct product Cartan decomposition $H = TA_0$.) We will be constructing representations of G attached to characters of H. The constructions will rely in particular on two Levi subgroups between H and G. In Harish-Chandra's work the most important is

(11.1a) $$M = Z_G(A_0).$$

This is a reductive algebraic subgroup of G, equal to the group of real points of $\mathbf{M} = Z_{\mathbf{G}}(\mathbf{A}_e)$. The roots of H in M are precisely the imaginary roots of H in G:

(11.1b) $$\Delta(M,H) = \Delta_{im}(G,H).$$

Once again our notation is a little different from the classical notation, in which M usually denotes the interesting factor in the direct product decomposition (often called "Langlands decomposition")

$$Z_G(\text{classical } A) = (\text{classical } M) \times (\text{classical } A).$$

This notation is inconvenient for us because the classical M need not be the real points of a connected reductive algebraic group.

Clearly the Cartan subgroup H is fundamental in M, so we can use the results of Section 5 to make representations of M from characters of H. The great technical benefit of using M is that M contains a maximally split Cartan subgroup H_s of G. The two Cartan decompositions

$$G = K(A_{s,0})K, \qquad M = (M \cap K)(A_{s,0})(M \cap K)$$

allow one to compare "behavior at infinity" on M and on G directly. (Such analysis is at the heart of Harish-Chandra's work, and of Langlands' proof of his classification of $\Pi^*(G)$. We will not make explicit use of it here.)

At the same time we will use

$$L = Z_G(T_0). \tag{11.1c}$$

This is a reductive algebraic subgroup of G, equal to the group of real points of $\mathbf{L} = Z_{\mathbf{G}}(\mathbf{T}_e)$. The roots of H in L are precisely the real roots of H in G:

$$\Delta(L, H) = \Delta_{re}(G, H). \tag{11.1d}$$

The Cartan subgroup H is split in L, so the results of Section 7 tell us about representations of L attached to characters of H. The great technical benefit of using L is that L contains a fundamental Cartan subgroup H_f of G. For that reason representations of $L \cap K$ are parametrized (approximately) by characters of T_f, which in turn (approximately) parametrize representations of K.

Definition 11.2. In the setting of (11.1), a *shifted Harish-Chandra parameter for a standard representation* is a pair

$$\Phi = (\phi, \Delta^+_{im})$$

(with $\phi \in \Pi^*(H)$ a character and $\Delta^+_{im} \subset \Delta_{im}(G, H_f)$ a system of positive roots) subject to the following requirement:

(1-std) For every positive imaginary root $\alpha \in \Delta^+_{im}$, we have

$$\langle d\phi, \alpha^\vee \rangle > 1.$$

We define

$$\begin{aligned} P^s_G(H) &= \text{shifted parameters for standard representations} \\ &= \text{pairs } (\phi, \Delta^+_{im}) \text{ satisfying condition (1-std) above.} \end{aligned}$$

(The superscript s stands for "shifted"; it is a reminder that this parameter differs by a ρ shift from the Harish-Chandra parameter.)

The imaginary roots divide into compact and noncompact as usual, according to the eigenvalue of θ on the corresponding root space:

$$\begin{aligned}\Delta^+_{im,c} &= \text{positive imaginary roots in } \mathfrak{k},\\ \Delta^+_{im,n} &= \text{remaining imaginary roots.}\end{aligned}$$

Define

$$2\rho_{im} = \sum_{\alpha\in\Delta^+_{im}} \alpha, \qquad 2\rho_{im,c} = \sum_{\beta\in\Delta^+_{im,c}} \beta \qquad 2\rho_{im,n} = \sum_{\gamma\in\Delta^+_{im,n}} \gamma.$$

Each of these characters may be regarded as belonging either to the group of real characters $\Pi^*_{adm}(H)$ or to the algebraic character lattice $\Pi^*_{alg}(\mathbf{H})$. The lattice of characters may be embedded in $\mathfrak{h}^*$ by taking differentials, so we may also regard these characters as elements of $\mathfrak{h}^*$. There they may be divided by two, defining ρ_{im}, $\rho_{im,c}$, and $\rho_{im,n}$ in $\mathfrak{h}^*$.

The *infinitesimal character parameter for* Φ is

$$\zeta(\Phi) = d\phi - \rho_{im} \in \mathfrak{h}^*;$$

we write

$$\zeta(\Phi) = \lambda(\Phi) + \nu(\Phi) \qquad (\lambda(\Phi)\in\mathfrak{t}^*, \nu(\Phi)\in\mathfrak{a}^*)$$

More generally, a *shifted Harish-Chandra parameter for a standard limit representation* is a pair

$$\Phi = (\phi, \Delta^+_{im})$$

(with $\phi\in\Pi^*(H)$ a character and $\Delta^+_{im}\subset\Delta_{im}(G,H_f)$ a system of positive roots) subject to the following requirement:

(1-lim) For every positive imaginary root $\alpha\in\Delta^+_{im}$, we have

$$\langle d\phi, \alpha^\vee\rangle \geq 1.$$

We define

$$\begin{aligned}P_G^{s,lim}(H) &= \text{shifted parameters for limits of standard representations}\\ &= \text{pairs } \Phi = (\phi,\Delta^+_{im}) \text{ satisfying condition (1-lim) above.}\end{aligned}$$

Finally, a *shifted Harish-Chandra parameter for a continued standard representation* is a pair

$$\Phi = (\phi, \Delta^+_{im})$$

with $\phi\in\Pi^*(H)$, Δ^+_{im} a system of positive imaginary roots, and no positivity hypothesis. We define

$$\begin{aligned}P_G^{s,cont}(H) &= \text{shifted parameters for continued standard representations}\\ &= \text{pairs } (\phi,\Delta^+_{im}).\end{aligned}$$

For limit and continued parameters we can define the infinitesimal character parameter $\zeta(\Phi)$ in the obvious way.

For each kind of parameter there is a discrete part, which remembers only the restriction ϕ_d of the character ϕ to the compact torus T. We write for example

$$\begin{aligned} P^{s,lim}_{G,d}(H) &= \text{shifted disc. params. for limits of standard reps.} \\ &= \text{pairs } \Phi_d = (\phi_d, \Delta^+_{im}) \text{ satisfying condition (1-lim) above.} \end{aligned}$$

A given discrete parameter has a distinguished extension to a full parameter, namely the one which is trivial on A_0 (the vector subgroup of H). By using this extension, we may regard $P^{s,lim}_{G,d}(H)$ as a subset of $P^{s,lim}_{G}(H)$.

Proposition 11.3. *In the setting of Definition 11.2, suppose $\Phi \in P^s_G(H)$ is a shifted Harish-Chandra parameter for a standard representation. Then there is a standard representation $I(\Phi)$ for G attached to Φ, which may be constructed as follows. Choose a real parabolic subgroup*

$$P = MN \subset G$$

in such a way that the real part of the differential $d\phi|_{\mathfrak{a}}$ is weakly dominant for all the roots of H in N. Let $I_M(\Phi)$ be the fundamental series representation of M specified by Proposition 5.2; this is in fact a (relative) discrete series representation of M. Define

$$I(\Phi) = \mathrm{Ind}^G_{MN} I_M(\Phi) \otimes 1$$

(normalized induction), a standard representation of G. It is always non-zero, and its restriction to K, which is

$$I(\Phi)|_K = \mathrm{Ind}^K_{M\cap K} I_M(\Phi)|_{M\cap K},$$

depends only on the restriction of Φ to T. The infinitesimal character of $I(\Phi)$ corresponds to the weight $\zeta(\Phi) \in \mathfrak{h}^*$ of Definition 11.2. This standard representation has a Langlands quotient $J(\Phi)$, which is always non-zero but may be reducible. We have

$$I(\Phi) \simeq I(\Psi) \iff J(\Phi) \simeq J(\Psi) \iff \Psi \in W(G,H)\cdot\Phi.$$

For limits of standard representations the situation is quite similar; the main problem is that the corresponding representation of G may vanish. Here is a statement.

Proposition 11.4. *In the setting of Definition 11.2, suppose $\Phi \in P^{s,lim}_G(H)$ is a shifted Harish-Chandra parameter for a standard limit representation. Then there is a standard limit representation $I(\Phi)$ for G attached to Φ, which may be constructed by parabolic induction exactly as in Proposition 11.3. The infinitesimal character of $I(\Phi)$ corresponds to the weight $\zeta(\Phi) \in \mathfrak{h}^*$ of Definition 11.2. Its restriction to K depends only on the restriction of Φ to T. This standard limit representation representation has a Langlands quotient $J(\Phi)$. We have*

$$\Psi \in W(G,H)\cdot\Phi \Longrightarrow I(\Phi) \simeq I(\Psi) \iff J(\Phi) \simeq J(\Psi).$$

Define

$$\mu(\Phi) = [\phi - 2\rho_{im,c}]|_T \in \Pi^*(T).$$

Then the following three conditions are equivalent:

(1) *the weight $\mu(\Phi)$ is dominant with respect to the compact imaginary roots $\Delta^+_{im,c}$;*
(2) *there is no simple root for Δ^+_{im} which is both compact and orthogonal to $\phi - \rho_{im}$; and*
(3) *the standard limit representation $I(\Phi)$ is non-zero.*

Here is the result for coherent continuation.

Proposition 11.5. *In the setting of Definition 11.2, suppose $\Phi \in P^{s,cont}_G(H)$ is a shifted Harish-Chandra parameter for continued standard representations. Then there is a virtual representation $I(\Phi) = I_G(\Phi)$ for G with the following properties.*

(1) *If $P = MN$ is any parabolic subgroup of G with Levi factor M, then*

$$I(\Phi) = \mathrm{Ind}^G_{MN} I_M(\Phi),$$

where the inducing (virtual) representation on the right is the one described in Proposition 5.4.
(2) *The restriction of $I(\Phi)$ to K depends only on Δ^+_{im} and the restriction of ϕ to T. Explicitly,*

$$I(\Phi)|_K = \mathrm{Ind}^K_{M\cap K} I_M(\Phi)|_{M\cap K}.$$

(3) *We have*

$$I(\Phi) \simeq I(w\cdot\Phi) \qquad (w \in W(G,H)).$$

(4) *The virtual representation $I(\Phi)$ has infinitesimal character corresponding to $\zeta(\Phi, \Delta^+_{im}) \in \mathfrak{h}^*$.*
(5) *If $\zeta(\Phi)$ is weakly dominant for Δ^+_{im}, then $I(\Phi)$ is equivalent to the standard limit representation attached to Φ in Proposition 11.4.*
(6) *Suppose V is a finite-dimensional representation of G. Recall that $\Delta(V,H)$ denotes the multiset of weights of H on V (Definition 3.2). Then*

$$I(\Phi) \otimes V \simeq \sum_{\delta \in \Delta(V,H)} I(\Phi + \delta).$$

In order to describe lowest K-types of standard representations, it is convenient to give a completely different construction of them, by cohomological induction. We begin with coherent continuations of standard representations, later specializing to the standard representations themselves.

Fix a parameter

$$\Phi = (\phi, \Delta^+_{im}) \in P^{s,cont}_G(H) \tag{11.6a}$$

(Definition 11.2). Recall the Levi subgroup $L \supset H$ defined in (11.1c), corresponding to the real roots. A θ-stable parabolic subalgnebra $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ is called *weakly Φ-compatible* if

$$\Delta^+_{im} \subset -\Delta(\mathfrak{u}, H). \tag{11.6b}$$

We may also say that Φ is *weakly $\mathfrak{q}$-compatible*, or that Φ and $\mathfrak{q}$ are *weakly compatible*. Here is a way to construct (any) such $\mathfrak{q}$. Fix a generic element

$$\tau \in i\mathfrak{t}_0^* \tag{11.6c}$$

which is dominant for Δ^+_{im}. The "genericity" we require of τ is that the only roots to which τ is orthogonal are the real roots of H in G. Attached to any such generic τ there is a weakly compatible $\mathfrak{q} = \mathfrak{q}(\tau) = \mathfrak{l} + \mathfrak{u}(\tau)$ characterized by

$$\Delta(\mathfrak{u}(\tau), H) = \{\alpha \in \Delta(G, H) \mid \langle \tau, \alpha^\vee \rangle < 0\}. \tag{11.6d}$$

When Φ is actually a standard limit parameter, we will sometimes wish to require more of $\mathfrak{q}$. Recall the weight $\lambda = \lambda(\Phi) \in \mathfrak{t}^*$ attached to Φ (Definition 11.2). We say that $\mathfrak{q}$ is *strongly Φ-compatible* if it is weakly Φ-compatible, and in addition

$$\langle \lambda, \alpha \rangle \le 0, \qquad (\alpha \in \Delta(\mathfrak{u}, H)). \tag{11.6e}$$

Again we may say that Φ is *strongly $\mathfrak{q}$-compatible*, or that the pair is *strongly compatible*. Here is a way to construct (any) strongly Φ-compatible parabolic. Since λ is weakly dominant for Δ^+_{im}, the weight

$$\tau_1 = \lambda + \epsilon\tau \tag{11.6f}$$

is dominant for Δ^+_{im} and generic as long as ϵ is a small enough positive real; and in that case $\mathfrak{q}(\tau_1)$ is strongly Φ-compatible. A little more explicitly,
(11.6g)

$$\Delta(\mathfrak{u}(\tau_1), H) = \{\alpha \in \Delta \mid \langle \lambda, \alpha^\vee \rangle < 0, \text{ or } \langle \lambda, \alpha^\vee \rangle = 0 \text{ and } \langle \tau, \alpha^\vee \rangle < 0\}.$$

Notice that, just as in the definition of Borel subalgebra for $\mathfrak{k}$ in Definition 3.3, we have put the *negative* roots in the nil radical. One effect (roughly speaking) is that Verma modules constructed using $\mathfrak{q}$ and (standard limit) Φ tend (in the strongly compatible case) to be irreducible. (This statement is not precisely true, for example because the definition of strongly Φ-compatible ignores $\nu(\Phi)$ (Definition 11.2).)

The involution θ preserves $\Delta(\mathfrak{u}, H)$. The fixed points are precisely the imaginary roots Δ^+_{im}. The remaining roots therefore occur in pairs

$$\{\alpha, \theta\alpha\} \qquad (\alpha \text{ complex in } \Delta(\mathfrak{u}, H)). \tag{11.6h}$$

Evidently the two roots α and $\theta\alpha$ have the same restriction to $T = H^\theta$. In this way we get a well-defined character

$$\rho_{cplx} = \sum_{\substack{\text{pairs } \{\alpha, \theta\alpha\} \\ \text{cplx in } \Delta(\mathfrak{u}, H)}} -\alpha|_T \quad \in \Pi^*_{adm}(T). \tag{11.6i}$$

Extending this character to be trivial on A_0, we get $\rho_{cplx} \in \Pi^*_{adm}(H)$. As the notation suggests, the differential of this character is equal to half the sum of the complex roots in $\mathfrak{u}^{op}$, which we are thinking of as positive.

Now define

$$\Phi_{\mathfrak{q}} = (\phi \otimes \rho_{cplx}, \emptyset) \in P^s_L(H), \tag{11.6j}$$

a (shifted) Harish-Chandra parameter for a principal series representation of L. The corresponding infinitesimal character parameter for L is

$$\zeta_{\mathfrak{l}}(\Phi_{\mathfrak{q}}) = d\phi_{\mathfrak{q}} = d\phi + \rho_{cplx} = \zeta(\Phi) + \rho_{im} + \rho_{cplx} = \zeta_{\mathfrak{g}}(\Phi) - \rho(\mathfrak{u}); \tag{11.6k}$$

the last equality is equivalent to $\mathfrak{q}$ being weakly Φ-compatible. This is the infinitesimal character of the standard principal series representation $I_L(\Phi_{\mathfrak{q}})$.

We can now form the generalized Verma module

$$M(\Phi_{\mathfrak{q}}) = U(\mathfrak{g}) \otimes_{\mathfrak{q}} I_L(\Phi_{\mathfrak{q}}), \tag{11.6l}$$

which is a $(\mathfrak{g}, L \cap K)$-module. Its infinitesimal character is given by adding the half sum of the roots of $\mathfrak{u}$ to the infinitesimal character of $I_L(\Phi_{\mathfrak{q}})$; it is therefore equal (still in the weakly Φ-compatible case) to $\zeta(\Phi)$.

Theorem 11.7. *Suppose $\Phi \in P^{s,cont}_G(H)$ is a parameter for a continued standard representation (Definition 11.2) and $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ is a weakly Φ-compatible parabolic (see (11.6)). Write $\mathbb{L}_k$ for the kth Bernstein derived functor carrying $(\mathfrak{g}, L \cap K)$-modules to $(\mathfrak{g}, K)$-modules, and $S = \dim \mathfrak{u} \cap \mathfrak{k}$. Then the virtual representation $I(\Phi, \Delta^+_{im})$ may be constructed from the generalized Verma module $M(\Phi_{\mathfrak{q}})$ of (11.6l):*

$$I(\Phi) = \sum_{k=0}^{S} (-1)^k \mathbb{L}_{S-k}(M(\Phi_{\mathfrak{q}}))$$

Suppose from now on that Φ is a standard limit parameter (Definition 11.2), and that $\mathfrak{q}$ is strongly Φ-compatible (cf. (11.6)). Then

$$\mathbb{L}_k(M(\Phi_{\mathfrak{q}})) = \begin{cases} I(\Phi) & k = S \\ 0 & k \neq S. \end{cases}$$

Write $A(\Phi_{\mathfrak{q}})$ for the set of lowest $L \cap K$-types of the principal series representation $I(\Phi_{\mathfrak{q}})$ (Proposition 7.5), and $A(\Phi)$ for the set of lowest K-types of the standard limit representation $I(\Phi)$. If $\widetilde{\mu_{\mathfrak{q}}}$ is an extended highest weight of $\tau_{\mathfrak{q}} \in A(\Phi_{\mathfrak{q}})$ (Definition 4.11), define $\widetilde{\mu} = \widetilde{\mu_{\mathfrak{q}}} \otimes 2\rho(\mathfrak{u} \cap \mathfrak{k})$ as in Theorem 10.6. There are two possibilities.

1. *If μ is dominant for Δ^+_c, then $\widetilde{\mu}$ is an extended highest weight of the irreducible K-representation*

$$\tau = \mathbb{L}_S(M_K(\tau_{\mathfrak{q}})),$$

In this case the bottom layer map of [5] *exhibits τ as a lowest K-type of $I(\Phi)$.*

(2) *If μ is not dominant for Δ_c^+, then*

$$\mathbb{L}_S(M_K(\tau_{\mathfrak{q}})) = 0.$$

In both cases $\mathbb{L}_k(M_K(\tau_{\mathfrak{q}})) = 0$ for $k \neq S$. This construction defines an inclusion

$$A(\Phi) \hookrightarrow A(\Phi_{\mathfrak{q}}), \qquad \tau \mapsto \tau_{\mathfrak{q}}.$$

The first assertions of the theorem (those not referring to lowest K-types) are special cases of Theorem 12.8, which will be formulated and proved in section 12.

This theorem (in conjunction with Theorem 10.6) effectively computes the highest weight parameters of the lowest K-types of $I(\Phi)$ by reduction to the special case of principal series for split groups, which was treated in Section 7.

Definition 11.8. Suppose $\Phi \in P^{s,lim}(G)$ is a standard limit parameter (Definition 11.2). We say that Φ is *final* if

(1) the standard limit representation $I(\Phi)$ (Proposition 11.3) is not zero; and
(2) if we choose a strongly Φ-compatible θ-stable parabolic subalgebra $\mathfrak{q}$ as in (11.6), then $\Phi_{\mathfrak{q}}$ is a final limit parameter for L (Definition 7.7)

Conditions equivalent to (1) are given in Proposition 11.4. The second condition is equivalent to

(2′) for every real root α of H in G,

$$\text{either} \quad \langle d\phi_{\mathfrak{q}}, \alpha^\vee \rangle \neq 0 \quad \text{or} \quad \phi_{\mathfrak{q}}(m_\alpha) = 1.$$

If Φ is discrete, the requirement is that $\phi_{\mathfrak{q}}(m_\alpha) = 1$ for every real root α.)

We write $P_G^{s,finlim}(H)$ for the set of final limit parameters for H.

Theorem 11.9. *Suppose $\Phi \in P_{G,d}^{s,lim}(H)$ is a discrete standard limit parameter (Definition 11.2) and $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ is a strongly Φ-compatible parabolic (see* (11.6)*). Assume that Φ satisfies condition (2) in the definition of final (cf. Definition 11.8): that is, that $\phi_{\mathfrak{q}}(m_\alpha) = 1$ for every real root α of H in G. Then $A(\Phi_{\mathfrak{q}})$ (Proposition 7.5) consists of a single irreducible representation $\tau_{\mathfrak{q}}$ of $L \cap K$, which is trivial on $L^{der} \cap K$ and therefore one-dimensional. Write $\widetilde{\mu_{\mathfrak{q}}}$ for the (unique) extended highest weight of $\tau_{\mathfrak{q}}$ (Definition 4.11), and define $\widetilde{\mu} = \widetilde{\mu_{\mathfrak{q}}} \otimes 2\rho(\mathfrak{u} \cap \mathfrak{k})$ as in Theorem 10.6. Then Φ is final—that is, $I(\Phi) \neq 0$—if and only if μ is dominant with respect to Δ_c^+.*

Assume now that Φ is final (so that μ is dominant for K). Then the irreducible K representation $\tau(\Phi)$ of extended highest weight $\widetilde{\mu}$ is the unique lowest K-type of $I(\Phi)$.

The correspondence $\Phi \mapsto \tau(\Phi)$ is a bijection from K-conjugacy classes of discrete final limit parameters for G onto $\Pi^*(K)$.

12. More constructions of standard representations

Proposition 11.3 and Theorem 11.7 describe two constructions of standard representations of G (and, implicitly, of their restrictions to K). In order to get an algorithm for branching from G to K, we need to generalize these constructions. In fact we will make use only of the generalized construction by cohomological induction; but the generalization for real parabolic induction is easier to understand and serves as motivation, so we begin with that.

As in (6.1) and Section 11, we fix a θ-stable real Cartan subgroup

$$H \subset G, \qquad T = H^{\theta}, \qquad A = H^{-\theta}. \tag{12.1a}$$

Fix also a shifted Harish-Chandra parameter for a continued standard representation

$$\Phi = (\phi, \Delta^+_{im}) \in P_G^{s,cont}(H) \tag{12.1b}$$

(Definition 11.2). Recall that ϕ is a character of H, and Δ^+_{im} a system of positive roots for the imaginary roots of H in G. We will make extensive use of the infinitesimal character parameter

$$\zeta(\Phi) = d\phi - \rho_{im} \in \mathfrak{h}^* \tag{12.1c}$$

(Definition 11.2).

We now fix the parabolics we will use to construct continued standard representations. If P is any real parabolic subgroup of G, then $M = P \cap \theta(P)$ is a Levi subgroup. If N is the unipotent radical of P, it follows that $P = MN$ is a Levi decomposition. We want to fix such a subgroup with the property that H is contained in M:

$$P = MN \text{ real parabolic}, \qquad H \subset M. \tag{12.1d}$$

The smallest possible choice for M is the subgroup

$$M_1 = Z_G(A_0) \subset M \tag{12.1e}$$

introduced in (11.1). One way to see this containment is to notice that complex conjugation must permute the roots of H in $\mathfrak{n}$, but complex conjugation sends the roots of H in $\mathfrak{m}_1$ to their negatives. Because all the imaginary roots of H occur in M, we may regard our fixed parameter as a parameter for M:

$$\Phi = (\phi, \Delta^+_{im}) \in P_M^{s,cont}(H). \tag{12.1f}$$

In the same way, we fix a θ-stable parabolic subalgebra as in (10.1) subject to two requirements: first, that L contain H

$$\mathfrak{q} = \mathfrak{l} + \mathfrak{u} \qquad \theta\text{-stable parabolic}, \qquad H \subset L; \tag{12.1g}$$

and second, that

$$\Delta_{im}(\mathfrak{u}, \mathfrak{h}) \subset -\Delta^+_{im}. \tag{12.1h}$$

The smallest possible choice for L is the subgroup

$$L_1 = Z_G(T_0) \subset L \tag{12.1i}$$

introduced in (11.1). To see this containment, notice that the action of θ must permute the roots of H in $\mathfrak{u}$, but sends the roots of H in $\mathfrak{l}_1$ to their negatives.

A θ-stable parabolic subalgebra satisfying conditions (12.1g) and (12.1h) is called *weakly Φ-compatible*. In case $L = L_1$, this is precisely the notion introduced in (11.6); all that has changed here is that we are allowing L to be larger than L_1. Just as in the earlier setting, we may say instead that Φ is *weakly $\mathfrak{q}$-compatible*, or that Φ and $\mathfrak{q}$ are *weakly compatible*.

Just as in (11.6i), the non-imaginary roots of H in $\mathfrak{u}$ occur in complex pairs $\{\alpha, \theta\alpha\}$, and we define

$$\rho_{cplx} = \sum_{\substack{\text{pairs } \{\alpha,\theta\alpha\} \\ \text{cplx in } \Delta(\mathfrak{u},H)}} -\alpha|_T \quad \in \Pi^*_{adm}(T). \tag{12.2a}$$

Extending this character to be trivial on A_0, we get $\rho_{cplx} \in \Pi^*_{adm}(H)$. Put

$$\phi_{\mathfrak{q}} = \phi \otimes \rho_{cplx} \in \Pi^*(H), \qquad \Delta^+_{im}(L) = \Delta^+_{im} \cap \Delta(\mathfrak{l}, H). \tag{12.2b}$$

We have

$$\Phi_{\mathfrak{q}} = (\phi_{\mathfrak{q}}, \Delta^+_{im}(L)) \in P^{s,cont}_L(H), \qquad \zeta_{\mathfrak{l}}(\Phi_{\mathfrak{q}}) = \zeta_{\mathfrak{g}}(\Phi) - \rho(\mathfrak{u}). \tag{12.2c}$$

Here $\rho(\mathfrak{u}) = -d\rho_{cplx} - \rho_{im} + \rho_{im}(L)$ is half the sum of the roots of $\mathfrak{h}$ in $\mathfrak{u}$.

We will have occasion later to invert the correspondence $\Phi \to \Phi_{\mathfrak{q}}$. Suppose therefore that $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ is a θ-stable parabolic subalgebra as in (10.1), that $H \subset L$ is a θ-stable Cartan subgroup, and that

$$\Psi = (\psi, \Delta^+_{im}(L)) \in P^{s,cont}_L(H) \tag{12.3a}$$

is a Harish-Chandra parameter for a continued standard representation of L. Define

$$\Delta^+_{im} = \Delta^+_{im}(L) \cup -\Delta_{im}(\mathfrak{u}, H); \tag{12.3b}$$

this is a set of positive imaginary roots for H in G. Write

$$\psi^{\mathfrak{q}}_{\mathfrak{g}} = \psi \otimes \rho^{-1}_{cplx}. \tag{12.3c}$$

Then

$$\Psi^{\mathfrak{q}}_{\mathfrak{g}} = (\psi^{\mathfrak{q}}_{\mathfrak{g}}, \Delta^+_{im}) \in P^{s,cont}_G(H). \tag{12.3d}$$

The infinitesimal character parameters are related by

$$\zeta_{\mathfrak{g}}(\Psi^{\mathfrak{q}}_{\mathfrak{g}}) = \zeta_{\mathfrak{l}}(\Psi) + \rho(\mathfrak{u}). \tag{12.3e}$$

In order to discuss standard limit representations, we will sometimes want to impose stronger hypotheses on the relationship between Φ and $\mathfrak{q}$. Recall from Definition 11.2 the weight $\lambda(\Phi) \in \mathfrak{t}^*$. We say that $\mathfrak{q}$ is *strongly Φ-compatible* if it is weakly Φ-compatible (cf. (12.1g), (12.1h)), and in addition

$$\langle \lambda(\Phi), \alpha^\vee \rangle \le 0, \qquad \alpha \in \Delta(\mathfrak{u}, H). \tag{12.4a}$$

We will also say that Φ is *strongly* $\mathfrak{q}$*-compatible*, or that Φ and $\mathfrak{q}$ are *strongly compatible.*

Proposition 12.5. *In the setting* (12.1)*, suppose that* $\Phi = (\phi, \Delta^+_{im})$ *is a shifted Harish-Chandra parameter for a continued standard representation. Then the corresponding virtual representation* $I_G(\Phi)$ *(Proposition 11.5) may be realized as*

$$I_G(\Phi) = \operatorname{Ind}^G_{MN} I_M(\Phi).$$

In particular, its restriction to K *is*

$$I_G(\Phi)|_K = \operatorname{Ind}^K_{M\cap K} I_M(\Phi, \Delta^+_{im})|_{M\cap K}.$$

Proof. Choose a parabolic subgroup $P_1 = M_1 N_1 \subset M$ with Levi factor $M_1 = Z_G(A_0)$ as in Proposition 11.3. Then $Q_1 = M_1(N_1 N)$ is a parabolic subgroup of G. According to Proposition 11.5 (applied first to G and then to M) we have

$$I_G(\Phi) = \operatorname{Ind}^G_{M_1(N_1 N)} I_{M_1}(\Phi),$$
$$I_M(\Phi) = \operatorname{Ind}^M_{M_1 N_1} I_{M_1}(\Phi).$$

Applying induction by stages to these two formulas gives the main claim of the proposition. The second follows because

$$G = KP, \qquad P \cap K = M \cap K.$$

□

Some care is required on one point. If Φ is a standard limit parameter, then the standard limit representations $I_G(\Phi)$ and $I_M(\Phi)$ are both actual representations (not merely virtual representations). The main identity of Proposition 12.5 therefore makes sense as an identity of representations. *It need not be true.* The theorem asserts only that it is true on the level of virtual representations; that is, that both sides have the same irreducible composition factors, appearing with the same multiplicities.

To see an example of this, suppose $G = GL_4(\mathbb{R})$, and H is the split Cartan subgroup consisting of diagonal matrices. There are no imaginary roots, so we omit the empty set of positive imaginary roots from the notation. We have

$$H = \{h = (h_1, h_2, h_3, h_4) \mid h_i \in \mathbb{R}^\times\}; \tag{12.6a}$$

the h_i are the diagonal entries of h. We consider the character ϕ of H defined by

$$\phi(h) = |h_1|^{3/2}|h_2|^{-3/2}|h_3|^{1/2}|h_4|^{-1/2}. \tag{12.6b}$$

Regarded as an element of $P^s_G(H)$, this is the Harish-Chandra parameter of the trivial representation of G. The corresponding standard representation $I_G(\Phi)$ is the space of densities on the full flag variety of G (the space of complete flags in $\mathbb{R}^4$). In order to realize this, we need to use a non-standard Borel subgroup B containing H (consisting neither of upper triangular nor of lower triangular matrices).

Now define $M = GL_2 \times GL_2$. The representation $I_M(\Phi)$ has its unique irreducible quotient the finite-dimensional representation $\mathbb{C}^3 \otimes \mathbb{C}^1$ of $GL_2 \times GL_2$. The first of these is the quotient of the adjoint representation of GL_2 by the center. This quotient carries an invariant Hermitian form (related to the Killing form of the first GL_2 factor. It follows that the induced representation

$$\operatorname{Ind}_{MN}^G I_M(\Phi) \tag{12.6c}$$

has as a quotient the unitarily induced representation

$$\operatorname{Ind}_{MN}^G \mathbb{C}^3 \otimes \mathbb{C}^1. \tag{12.6d}$$

This quotient carries a non-degenerate invariant Hermitian form (induced from the one for M).

It is an easy consequence of the Langlands classification that only the Langlands quotient of a standard representation (and not some larger quotient) can carry a non-degenerate invariant Hermitian form. For this reason, the induced representation $\operatorname{Ind}_{MN}^G I_M(\Phi)$ cannot be isomorphic to the standard representation $I_G(\Phi)$.

Our next task is to realize continued standard representations by cohomological induction from the θ-stable parabolic subalgebra $\mathfrak{q}$ of (12.1) (assumed always to be weakly Φ-compatible). Just as in (11.6), we begin with the virtual representation

$$I_L(\Phi_{\mathfrak{q}}) \tag{12.7a}$$

of L. (Whatever construction we use for this virtual representation will construct it as a difference of two $(\mathfrak{l}, L \cap K)$-modules of finite length; for example the even and odd degrees respectively of some cohomologically induced representations. We will have no need to be very explicit about this.) We can therefore construct the (virtual) generalized Verma module

$$M(\Phi_{\mathfrak{q}}) = U(\mathfrak{g}) \otimes_{\mathfrak{q}} I_L(\Phi_{\mathfrak{q}}), \tag{12.7b}$$

which is a (virtual) $(\mathfrak{g}, L \cap K)$-module. Its infinitesimal character is given by adding the half sum of the roots of $\mathfrak{u}$ to the infinitesimal character of $I_L(\Phi_{\mathfrak{q}})$; it is therefore (in light of (12.2c)) equal to $\zeta(\Phi)$.

Theorem 12.8. *Suppose $\Phi = (\phi, \Delta^+_{im}) \in P_G^{s,cont}(H)$ is a parameter for a continued standard representation (Definition 11.2) and $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ is a weakly Φ-compatible parabolic (see (12.1)). Write $\mathbb{L}_k$ for the kth Bernstein derived functor carrying $(\mathfrak{g}, L \cap K)$-modules to $(\mathfrak{g}, K)$-modules, and $S = \dim \mathfrak{u} \cap \mathfrak{k}$. Then the virtual representation $I_G(\Phi)$ may be constructed from the virtual generalized Verma module $M(\Phi_{\mathfrak{q}})$ of (12.7b):*

$$I_G(\Phi) = \sum_{k=0}^{S} (-1)^k \mathbb{L}_{S-k}(M(\Phi_{\mathfrak{q}}))$$

Suppose in addition that Φ *is a standard limit parameter, and that* $\mathfrak{q}$ *is strongly* Φ*-compatible (cf.* (12.4)*). Then*

$$\mathbb{L}_k(M(\Phi_{\mathfrak{q}})) = \begin{cases} I_G(\Phi) & k = S \\ 0 & k \neq S. \end{cases}$$

Proof. We begin with the last assertion of the theorem, concerning the case when $\mathfrak{q}$ is strongly Φ-compatible. In this case the claim amounts to Theorem 11.225 of [5].

For the first assertions, we look at the family of parameters

$$\Phi + \gamma = (\phi \otimes \gamma, \Delta^+_{im}), \qquad (\gamma \in \Pi^*(\mathbf{H}))$$

The family of virtual representations $I_G(\Phi + \gamma)$ is a coherent family ([8], Definition 7.2.5); this is the how continued standard representations are defined. The family of virtual representations appearing on the right side of the first formula in the proposition is also coherent; this is proved in the same way as [8], Corollary 7.2.10. Two coherent families coincide if and only if they agree for at least one choice of γ with the property that the infinitesimal character $\zeta(\Phi + \gamma)$ is regular. If γ_0 is sufficiently negative for the roots of H in $\mathfrak{u}$—for example, if γ_0 is a large enough multiple of $-2\rho(\mathfrak{u})$—then $\mathfrak{q}$ is strongly $(\Phi + \gamma)$-compatible whenever γ is close to γ_0. In this case the equality of the two sides is a consequence of the last assertion of the theorem. Most such choices of γ will make $\zeta(\Phi + \gamma)$ regular, and so force the two coherent families to coincide. $\square$

Theorem 12.8 is phrased to construct a given (continued) standard representation of G by cohomological induction. We will also want to use it to identify the result of applying cohomological induction to a given (continued) standard representation of L, and for this a slight rephrasing is convenient.

Theorem 12.9. *Suppose* $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ *is a* θ*-stable parabolic subalgebra of* $\mathfrak{g}$*,* $H \subset L$ *a* θ*-stable Cartan subgroup, and* $\Psi = (\psi, \Delta^+_{im}(L)) \in P^{s,cont}_L(H)$ *is the Harish-Chandra parameter for a continued standard representation of* L*. Define*

$$\Psi^{\mathfrak{q}}_{\mathfrak{g}} = (\psi^{\mathfrak{q}}_{\mathfrak{g}}, \Delta^+_{im}) \in P^{s,cont}_G(H)$$

as in (12.3)*.*

1. *The parabolic* $\mathfrak{q}$ *is weakly* $\Psi^{\mathfrak{q}}_{\mathfrak{g}}$*-compatible (see* (12.1)*). We may therefore transfer* $\Psi^{\mathfrak{q}}_{\mathfrak{g}}$ *to a parameter for* L *as in* (12.2)*, and we have*
$$(\Psi^{\mathfrak{q}}_{\mathfrak{g}})_{\mathfrak{q}} = \Psi.$$
2. *Write* $I_L(\Psi)$ *for the virtual* $(\mathfrak{l}, L \cap K)$*-module attached to the parameter* Ψ*, and*
$$M(\Psi) = U(\mathfrak{g}) \otimes_{\mathfrak{q}} I_L(\Psi)$$

for the corresponding virtual $(\mathfrak{g}, L\cap K)$*-module. Write* $\mathbb{L}_k$ *for the* k*th Bernstein derived functor carrying* $(\mathfrak{g}, L\cap K)$*-modules to* $(\mathfrak{g}, K)$*-modules, and* $S = \dim \mathfrak{u}\cap\mathfrak{k}$*. Then*

$$\sum_{k=0}^{S}(-1)^k \mathbb{L}_{S-k}(M(\Psi)) = I_G(\Psi_{\mathfrak{g}}^{\mathfrak{q}}).$$

(3) *Suppose now that* Ψ *is a standard limit parameter for* L*, so that* $I_L(\Psi)$ *is an actual representation, and* $M(\Psi)$ *is an actual* $(\mathfrak{g}, L\cap K)$*-module. Write* $\lambda_L(\Psi)$ *for the weight introduced in Definition 11.2, and*

$$\lambda_G(\Psi_{\mathfrak{g}}^{\mathfrak{q}}) = \lambda_L(\Psi) + \rho(\mathfrak{u}).$$

Then $\mathfrak{q}$ *is strongly* $\Psi_{\mathfrak{g}}^{\mathfrak{q}}$*-compatible if and only if* $\lambda_G(\Psi_{\mathfrak{g}}^{\mathfrak{q}})$ *is weakly antidominant for the roots of* H *in* $\mathfrak{u}$*. If this is the case, then*

$$\mathbb{L}_k(M(\Psi)) = \begin{cases} I_G(\Psi_{\mathfrak{g}}^{\mathfrak{q}}) & k = S \\ 0 & k \neq S. \end{cases}$$

Proof. The assertions in (1) are formal and very easy. With these in hand, the rest of the theorem is just Theorem 12.8 applied to the parameter $\Psi_{\mathfrak{g}}^{\mathfrak{q}} \in P_G^{s,cont}(H)$. □

We conclude this section with some results on the restrictions to K of cohomologically induced representations. Generalizing (12.7), we begin with a virtual representation Z of L of finite length. The restriction of Z to $L\cap K$ decomposes as a (virtual) direct sum

$$Z = \sum_{\tau\in\Pi^*(L\cap K)} Z(\tau); \tag{12.10a}$$

here the finite-dimensional virtual representation $Z(\tau)$ is a (positive or negative) multiple of the irreducible representation τ. From Z we construct a virtual generalized Verma module

$$M(Z) = U(\mathfrak{g})\otimes_{\mathfrak{q}} Z. \tag{12.10b}$$

This is a virtual $(\mathfrak{g}, L\cap K)$-module. We will need to know the restriction of $M(Z)$ to $(\mathfrak{k}, L\cap K)$. This will be described in terms of the generalized Verma modules of Theorem 10.6: if E is a (virtual) $(\mathfrak{q}\cap\mathfrak{k}, L\cap K)$ module, we write

$$M_K(E) = U(\mathfrak{k})\otimes_{\mathfrak{q}\cap\mathfrak{k}} E, \tag{12.10c}$$

which is a (virtual) $(\mathfrak{k}, L\cap K)$-module. The restriction formula we need is

$$M(Z)|_{(\mathfrak{k}, L\cap K)} = \sum_{m\geq 0} M_K(Z\otimes S^m(\mathfrak{g}/(\mathfrak{k}+\mathfrak{q}))). \tag{12.10d}$$

Even if Z is a an actual representation, the restriction formula (12.10d) is true only on the level of virtual representations. The Verma module $M(Z)$

has a $(\mathfrak{k}, L\cap K)$-stable filtration for which the mth level of the associated graded module is given by the mth term in (12.10d).

It is worth noticing that the restriction formula (12.10d) depends only on $Z|_{L\cap K}$. We will deduce from this a formula for $\sum_k (-1)^k \mathbb{L}_{S-k}(M(Z))|_K$, again depending only on $Z|_{L\cap K}$. If Z is an actual representation, one might guess that each individual representation $\mathbb{L}_{S-k}(M(Z))|_K$ depends only on $Z|_{L\cap K}$. This guess is *incorrect*.

One of the main steps in our program to get branching laws for standard representations is to write irreducible representations of K explicitly as integer combinations of cohomologically induced representations. For that we will need to get rid of the infinite sum in (12.10d), and this we do using a Koszul complex. Recall from (2.5) the decomposition $\mathfrak{g} = \mathfrak{k} + \mathfrak{s}$. This is compatible with the triangular decomposition

$$\mathfrak{g} = \mathfrak{u}^{op} + \mathfrak{l} + \mathfrak{u}, \tag{12.10e}$$

so we get

$$\mathfrak{g}/(\mathfrak{k}+\mathfrak{q}) \simeq \mathfrak{u}^{op} \cap \mathfrak{s} \tag{12.10f}$$

as representations of $L\cap K$. Now consider the $(\mathfrak{q}\cap\mathfrak{k}, L\cap K)$-modules

$$E^+(\mathfrak{q}) = \sum_{p \text{ even}} \bigwedge^p \mathfrak{g}/(\mathfrak{k}+\mathfrak{q}), \qquad E^-(\mathfrak{q}) = \sum_{p \text{ odd}} \bigwedge^p \mathfrak{g}/(\mathfrak{k}+\mathfrak{q}) \tag{12.10g}$$

and the virtual $(\mathfrak{q}\cap\mathfrak{k}, L\cap K)$-module

$$E^{\pm}(\mathfrak{q}) = E^+ - E^- = \sum_p (-1)^p \bigwedge^p \mathfrak{g}/(\mathfrak{k}+\mathfrak{q}). \tag{12.10h}$$

Proposition 12.11. *Suppose we are in the setting of* (12.10)*, so that $\mathfrak{q}$ is a θ-stable parabolic subalgebra of $\mathfrak{g}$, and Z is a virtual representation of L of finite length. Then*

$$\sum_k (-1)^k \mathbb{L}_{S-k}(M(Z))|_K = \sum_{m\geq 0}\sum_k (-1)^k \mathbb{L}_{S-k}\left(M_K(Z\otimes S^m(\mathfrak{g}/(\mathfrak{k}+\mathfrak{q})))\right)$$

Suppose now that we can find a virtual $(\mathfrak{l}, L\cap K)$-module of finite length $Z^{\pm}$ with the property that

$$Z^{\pm}|_{L\cap K} = (Z|_{L\cap K})\otimes E^{\pm}$$

(see (12.10h)*). Then*

$$\sum_k (-1)^k \mathbb{L}_{S-k}(M(Z^{\pm}))|_K = \sum_k (-1)^k \mathbb{L}_{S-k}\left(M_K(Z)\right).$$

Proof. The first assertion (which is due to Zuckerman) is a version of [8], Theorem 6.3.12. For the second, the Koszul complex for the vector space $\mathfrak{g}/(\mathfrak{k}+\mathfrak{q})$ provides an equality of virtual representations of $L\cap K$

$$\sum_p (-1)^p \bigwedge^p \big(\mathfrak{g}/(\mathfrak{k}+\mathfrak{q})\big)\otimes S^{m-p}\big(\mathfrak{g}/(\mathfrak{k}+\mathfrak{q})\big) = \begin{cases} 0, & m>0 \\ \mathbb{C}, & m=0. \end{cases} \tag{12.12}$$

More precisely, the Koszul complex has the pth summand on the left in degree p, and a differential of degree -1; its cohomology is zero for $m > 0$, and equal to $\mathbb{C}$ in degree zero for $m = 0$. Tensoring with Z and applying the first assertion, we get the second assertion. $\square$

Information about how to compute $\mathbb{L}_k(M_K(E))$ (for E an irreducible representation of $L \cap K$) may be found in Theorem 10.6. We will need only the very special case of "bottom layer K-types" for standard representations, described in Theorem 11.7.

It is by no means obvious *a priori* that the virtual L representation $Z^{\pm}$ should exist. The $L \cap K$ representations E^+ and E^- need not extend to L, so we cannot proceed just by tensoring with finite-dimensional representations of L. We will be applying the proposition to continued standard representations Z, and in that case we are saved by the following result (applied to L instead of G).

Lemma 12.13. *Suppose $\Phi = (\phi, \Delta^+_{im}) \in P^{s,cont}_G(H)$ is a shifted Harish-Chandra parameter for a continued standard representation, and that (τ, E) is a finite-dimensional representation of K. Write $T = H \cap K$, and*

$$\tau|_T = \sum_{\gamma_T \in \Pi^*(T)} m_E(\gamma_T)\gamma_T;$$

here the integer multiplicities $m_E(\gamma_T)$ are finite and non-negative. For each γ_T, choose an extension $\gamma \in \Pi^*(H)$ of the character γ_T to H; for example, one can choose the unique extension trivial on A_0. Write

$$\Phi + \gamma = (\phi \otimes \gamma, \Delta^+_{im}) \in P^{s,cont}_G(H).$$

Then

$$\big(I_G(\Phi) \otimes E\big)|_K \simeq \sum_{\gamma_T \in \Pi^*(T)} m_E(\gamma_T) I_G(\Phi + \gamma)|_K.$$

What the lemma says is that we can find an expression writing

$$\text{(continued standard)}\otimes\text{(representation of } K)$$

as the restriction to K of a finite sum of continued standard representations of G. To make it explicit, we need to be able to compute the restriction of the representation of K to T. To compute this restriction for a general representation of K is beyond what standard results about compact Lie groups can provide; think of the example of $K = O(n)$ and $T = O(1)^n$ (arising from the split Cartan in $GL_n(\mathbb{R})$). In the present setting we are interested only in restricting the representations $E^{\pm}(\mathfrak{q})$ defined in (12.10) to $T \subset H \subset L$. This can be done explicitly in terms of the root decomposition of H acting $\mathfrak{g}$.

Proof. If all of the roots of H in G are imaginary, then $I_G(\Phi, \Delta^+_{im})$ is cohomologically induced from a character of H (Theorem 11.7), and the lemma follows from the first formula in Proposition 12.11. The general case may be reduced to this case by part (2) of Proposition 11.5. (Alternatively, one can begin with the case when all roots of H in G are real, and reduce to that case using Theorem 11.7 and Proposition 12.11.) □

13. From highest weights to discrete final limit parameters

Theorem 11.9 describes in a more or less algorithmic way how to pass from a discrete final limit parameter Φ to the highest weight of an irreducible representation $\tau(\Phi)$ of K. We need to be able to reverse this process: to begin with the highest weight of an irreducible representation, and extract a discrete final limit parameter. The algorithm of Theorem 11.9 uses a θ-stable parabolic subalgebra $\mathfrak{q}$ constructed (not quite uniquely) from Φ. Once we have $\mathfrak{q}$, the algorithm is easily reversible. Our task therefore is to construct $\mathfrak{q}$ directly from τ. This is very close to the basic algorithm underlying the classification of $\Pi^*(G)$ as described in [8]. In this section we will describe the small modification of that algorithm that we need.

We begin with a Δ_c^+-dominant weight

$$\mu \in \Pi^*_{dom}(T_f) \tag{13.1a}$$

(Definition 3.3), and with $2\rho_c \in \Pi^*_{dom}(T_f)$ the modular character of (4.3). We will need the set of roots

$$Q_0^{\pm}(\mu) = \{\beta \in \Delta(G, T_f) \mid \langle \mu + 2\rho_c, \beta^\vee\rangle = 0\}. \tag{13.1b}$$

These roots are (the restrictions to T_f of) noncompact imaginary roots of H_f in G; they are strongly orthogonal, and are the roots of a split Levi subgroup $L_0(\mu)$, locally isomorphic to a product of copies of $SL_2(\mathbb{R})$ and an abelian group. The roots in $Q_0(\mu)$ are going to be in the Levi subgroup $L(\mu)$ that we ultimately construct. It is convenient to choose one root from each pair $\pm\beta$ in $Q_0^{\pm}(\mu)$, and to call the resulting set $Q_0(\mu)$. Consider the Weyl group

$$W_0(\mu) = \{w_A = \prod_{\beta \in A} s_\beta \mid A \subset Q_0(\mu)\}, \tag{13.1c}$$

a product of copies of $\mathbb{Z}/2\mathbb{Z}$.

We now define

$$\Delta^+(G, T_f)(\mu) = \{\alpha \in \Delta(G, T_f) \mid \langle \mu + 2\rho_c, \alpha^\vee\rangle > 0\} \cup Q_0(\mu). \tag{13.1d}$$

This is the restriction to T_f of a θ-stable positive root system $\Delta^+(G, H_f)(\mu)$. Of course it depends on the choice of $Q_0(\mu)$; changing the choice replaces $\Delta^+(\mu)$ by $w_A\Delta^+(\mu)$ for some $A \subset Q_0(\mu)$. This positive root system defines a modular character

$$2\rho(\mu) = \sum_{\alpha \in \Delta^+(\mu)} \alpha \in \Pi^*_{alg}(\mathbf{H}_f); \tag{13.1e}$$

this character is fixed by θ, and therefore trivial on $A_{f,0}$. We write $2\rho(\mu)$ also for the restriction to T_f. Differentiating and dividing by two gives

$$\rho(\mu) \in \mathfrak{t}_f^* \subset \mathfrak{h}_f^*.$$

Construction of $\mathfrak{q}$: project the weight $\mu + 2\rho_c - \rho(\mu)$ on the $\Delta^+(G, T_f)(\mu)$ positive Weyl chamber, getting a weight $\lambda(\mu)$. Write

$$\mu + 2\rho_c - \rho(\mu) = \lambda(\mu) - \phi,$$

with ϕ a non-negative rational combination of $\Delta^+(\mu)$-simple roots. Let $\mathfrak{q}$ be the parabolic corresponding to the set of simple roots whose coefficient in ϕ is strictly positive. (Necessarily this includes the roots in $Q_0(\mu)$.) This is the one.

14. Algorithm for projecting a weight on the dominant Weyl chamber

The algorithm of Section 13 depends on projecting a certain rational weight on a positive Weyl chamber. In general (for example in the papers of Carmona and Aubert-Howe) this projection is defined in terms of orthogonal geometry. Since the language of root data avoids the orthogonal form (instead keeping roots and coroots in dual spaces) it is useful to rephrase their results in this language. At the same time I will sketch an algorithm to carry out the projection. So fix a root system in a rational or real vector space. It is convenient to label this space as V^*, and to have the coroots in V:

$$\Delta \subset V^*, \qquad \Delta^\vee \subset V. \tag{14.1a}$$

We fix also sets of positive and simple roots

$$\Delta^+ \supset \Pi. \tag{14.1b}$$

These give rise to fundamental weights and coweights

$$\begin{aligned} \chi_\alpha \in \mathbb{Q}\Delta \subset V^*, &\qquad \langle \chi_\alpha, \beta^\vee \rangle = \delta_{\alpha,\beta}, \qquad (\alpha, \beta \in \Pi) \\ \chi^\vee_{\alpha^\vee} \in \mathbb{Q}\Delta^\vee \subset V, &\qquad \langle \beta, \chi^\vee_{\alpha^\vee} \rangle = \delta_{\alpha,\beta} \qquad (\alpha, \beta \in \Pi). \end{aligned} \tag{14.1c}$$

We also need a small generalization: if $B \subset \Pi$, then

$$\Delta(B) = \mathbb{Z}B \cap \Delta \tag{14.1d}$$

is a root system in V with simple roots B. (If Δ is the root system for a reductive algebraic group $\mathbf{G}$, then $\Delta(B)$ corresponds to a Levi subgroup $\mathbf{L}(B)$.) Write $B^\vee$ for the set of simple coroots corresponding to B. We therefore get

$$\begin{aligned} \chi_\alpha(B) \in \mathbb{Q}\Delta(B) \subset V^*, &\quad \langle \chi_\alpha(B), \beta^\vee \rangle = \delta_{\alpha,\beta} \quad (\alpha, \beta \in B) \\ \chi^\vee_{\alpha^\vee}(B) \in \mathbb{Q}\Delta^\vee(B) \subset V, &\quad \langle \beta, \chi^\vee_{\alpha^\vee}(B) \rangle = \delta_{\alpha,\beta} \quad (\alpha, \beta \in B). \end{aligned} \tag{14.1e}$$

The question of constructing these elements explicitly will arise in the course of the algorithm; of course the defining equations to be solved are linear with integer coefficients.

We are interested in the geometry of cones

$$
\begin{aligned}
C &= \{\lambda \in V^* \mid \langle \lambda, \alpha^\vee \rangle \geq 0, \alpha \in \Pi\} \\
C^\vee &= \{v \in V \mid \langle \alpha, v \rangle \geq 0, \alpha \in \Pi\} \\
P &= \{\sum_{a_\alpha \geq 0} a_\alpha \alpha\} \subset V^* \\
P^\vee &= \{\sum_{b_{\alpha^\vee} \geq 0} b_{\alpha^\vee} \alpha^\vee\} \subset V
\end{aligned}
\tag{14.1f}
$$

the *positive Weyl chambers* and the *positive root (coroot) cone* respectively. Sometimes it will be convenient as well to consider the two subspaces

$$
\begin{aligned}
Z^* &= \{\lambda \in V^* \mid \langle \lambda, \alpha^\vee \rangle = 0, \alpha \in \Pi\} \\
Z &= \{v \in V \mid \langle \alpha, v \rangle = 0, \alpha \in \Pi\}.
\end{aligned}
\tag{14.1g}
$$

These are dual vector spaces under the restriction of the pairing between V and V^*.

Here is the general theoretical statement, due to Langlands. For a variety of slightly different and illuminating perspectives, see [6], Lemma 4.4; [2], Lemma IV.6.11; [3], Proposition 1.2; and [1], Proposition 1.16.

Theorem 14.2. *In the setting* (14.1)*, there is for each $\nu \in V^*$ a unique subset $A \subset \Pi$ and a unique expression*

$$\nu = \lambda - \phi \qquad (\lambda \in C, \phi \in P)$$

subject to the following requirements.

(1) *The element ϕ is a combination of roots $\alpha \in A$ with strictly positive coefficients:*

$$\phi = \sum_{\alpha \in A} a_\alpha \alpha, \qquad a_\alpha > 0.$$

(2) *The element λ vanishes on $\alpha^\vee$ for all $\alpha \in A$:*

$$\lambda = \sum_{\alpha \notin A} c_\alpha \chi_\alpha + z \qquad (c_\alpha \geq 0, z \in Z).$$

Of course there is a parallel statement for V. This formulation looks a little labored: if one is willing to fix a W-equivariant identification of V with V^* (and so to think of the pairing between them as a symmetric form on V) then the two requirements imposed on λ and ϕ may be written simply as $\langle \lambda, \phi \rangle = 0$, and there is no need to mention A. But for our purposes the most interesting output of the algorithm *is* the set A, which defines the θ-stable parabolic $\mathfrak{q} = \mathfrak{q}(\mu)$. So mentioning this set in the statement of the theorem, and constructing it directly, seem like reasonable ideas.

Proof. The uniqueness of the decomposition is a routine exercise; what is difficult is existence, and for that I will give a constructive proof. All that we will construct is the set A. This we will do inductively. At each stage of the induction, we will begin with a subset $B_m \subset A$. We will then test the simple roots outside B_m. Any root that fails the test must belong to A, and will be added to the next set B_{m+1}. If all the roots outside of B_m pass the test, then $B_m = A$, and we stop. The algorithm begins with the empty set $B_0 = \emptyset$, which is certainly a subset of A. Here is the inductive test. For each simple root $\beta \notin B_m$, construct the fundamental coweight

$$\chi^\vee = \chi^\vee_{\beta^\vee}(B_m \cup \{\beta\})$$

with respect to the simple roots consisting of B_m and β. This coweight vanishes on the roots in B_m, but takes the value 1 on the root β. (I will say a word about constructing this coweight in a moment.) The test we apply is to ask whether $\langle \nu, \chi^\vee \rangle$ is non-negative. If it is, then β passes. If β fails, then I claim that it must belong to A, and we add it to B_{m+1}. Here is why. As a fundamental coweight, $\chi^\vee$ is a nonnegative rational combination of $\beta^\vee$ and the various simple coroots in $B_m^\vee \subset A^\vee$. It therefore has a non-negative pairing with $\lambda \in C$; so if it is strictly negative on ν, then

$$0 < \langle \phi, \chi^\vee \rangle \tag{14.3}$$

$$= \sum_{\alpha \in B_m} a_\alpha \langle \alpha, \chi^\vee \rangle + \sum_{\gamma \in A - B_m} a_\gamma \langle \gamma, \chi^\vee \rangle \tag{14.4}$$

The first sum is zero by construction of $\chi^\vee$. Because $\chi^\vee$ is a non-negative combination of $\beta^\vee$ and coroots from B_m, every term in the second sum is non-positive except the term with $\gamma = \beta$. Of course that term is present only if $\beta \in A$, as we wished to show.

It remains to show that if every simple root outside B_m passes this test, then $B_m = A$. Suppose $B_m \neq A$. We may therefore find a simple root β so that

$$\langle \sum_{\gamma \in A - B_m} a_\gamma \gamma\, ,\, \beta^\vee \rangle > 0.$$

Because a simple root and a distinct simple coroot have non-positive pairing, this equation shows first of all that $\beta \in A - B_m$. We show that β fails the test. Define

$$\chi_0^\vee = \beta^\vee + \sum_{\alpha \in B_m} (-\langle \alpha, \beta^\vee \rangle) \chi^\vee_{\alpha^\vee}(B_m) \tag{14.5}$$

$$= \beta^\vee + \delta^\vee \tag{14.6}$$

This element is evidently in the $\mathbb{Q}$-span of $B_m^\vee \cup \{\beta^\vee\}$, and its pairing with the roots in B_m is zero. For its pairing with β, the coroot $\beta^\vee$ contributes 2, and the terms from the sum are all non-negative (since the elements $\chi^\vee_{\alpha^\vee}(B_m)$ are non-negative combinations of simple roots which are in B_m, and therefore distinct from β. So the pairing with β is at least two. Consequently

$$\chi_0^\vee = c\chi^\vee, \qquad \text{some } c \geq 2.$$

We are trying to prove that $\chi^\vee$ has a strictly negative pairing with ν, so it suffices to prove that $\langle \nu, \chi_0^\vee \rangle < 0$. We have already seen that $\chi_0^\vee$ is a combination of roots in A, so its pairing with λ must be zero by condition (2) of Theorem 14.2. Since $\nu = \lambda - \phi$, what we want to prove is

$$0 < \langle \phi, \chi_0^\vee \rangle. \tag{14.7}$$

Since $\chi_0^\vee$ is orthogonal to the roots in B_m, the left side is

$$\begin{aligned} &= \langle \sum_{\gamma \in A - B_m} a_\gamma \gamma, \chi_0^\vee \rangle \\ &= \langle \sum_{\gamma \in A - B_m} a_\gamma \gamma, \beta^\vee \rangle + \sum_{\gamma \in A - B_m} a_\gamma \langle \gamma, \delta^\vee \rangle. \end{aligned}$$

It follows from (14.5) that $\delta^\vee$ is a non-negative combination of roots in B_m. Consequently every term in the second sum is non-negative. The first term is strictly positive by the choice of β, proving (14.7). □

15. Making a list of representations of K

Here is a way to make a list of the irreducible representations of K that is compatible with the Langlands classification. As usual begin with a fundamental Cartan H_f, with $T_f = H_f \cap K$, and a fixed system

$$\Delta_c^+ \subset \Delta(K, T_f) \tag{15.1a}$$

of positive roots for K. As in Definition 3.3, this choice provides a Borel subalgebra

$$\mathfrak{b}_c = \mathfrak{t}_f + \mathfrak{n}_c \tag{15.1b}$$

corresponding to the negative roots.

The first serious step is to pick a representative of each T_{fl}-conjugacy class of θ-stable parabolic subalgebras $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ such that

1. the Levi subalgebra $\mathfrak{l}$ contains $\mathfrak{h}_f$;
2. the Borel subalgebra $\mathfrak{b}_c$ is contained in $\mathfrak{q}$; and
3. the Levi subgroup L has split derived group.

This is the same as picking representatives of K-conjugacy classes of θ-stable $\mathfrak{q}$ opposite to their complex conjugates, subject only to the last requirement on the (unique) θ-stable real Levi subgroup L. These parabolic subalgebras are in one-to-one correspondence with the corresponding Zuckerman unitary representations $A(\mathfrak{q})$. The larger set of T_{fl} conjugacy classes of θ-stable $\mathfrak{q}$ subject only to conditions (1), (2), and

(3′) the group L has no compact simple factor

is already picked out by the software `blocku` command. Recognizing condition (3) is equally easy.

So we fix a collection

$$\mathcal{Q} = \{\mathfrak{q}\} \tag{15.1c}$$

of representatives for the T_{fl}-conjugacy classes of parabolic subalgebras satisfying conditions (1)–(3) above. For each $\mathfrak{q}$, fix a maximally split Cartan subgroup

$$H \subset L, \tag{15.1d}$$

with $T = H \cap K$ as usual.

The goal is to list all the irreducible representations of K having a highest weight μ so that the construction of Section 13 gives the parabolic $\mathfrak{q}$. The disjoint union over $\mathfrak{q} \in \mathcal{Q}$ of these lists will be a parametrization of $\Pi^*(K)$. The list corresponding to $\mathfrak{q}$ is a certain subset of the discrete final limit parameters for H. Precisely,

$$P^{s,\mathfrak{q},finlim}_{G,d}(H) = \{\Phi \in P^{s,finlim}_{G,d}(H) \mid \Phi \text{ is strongly } \mathfrak{q}\text{-compatible}\} \tag{15.1e}$$

The notion of "strongly compatible" was introduced in (11.6). We will see that the set described by (15.1e) is essentially a set of characters of T, subject to certain parity and positivity conditions. We want to make this more explicit. First of all we invoke the notation of (11.6), writing in particular

$$\Delta^+_{im} = \Delta_{im}(G, H) \cap (-\Delta(\mathfrak{u}, H))\,. \tag{15.1f}$$

The assumption of weak compatibility from (11.6) means that

$$P^{s,\mathfrak{q},finlim}_{G,d}(H) \subset \{\Phi = (\phi, \Delta^+_{im}) \mid \phi \in \Pi^*(T)\}; \tag{15.1g}$$

that is, that we are looking only at parameters involving this particular choice of positive imaginary roots. Recall also from (11.6) the character ρ_{cplx} of T and the shift $\phi_{\mathfrak{q}} = \phi \otimes \rho_{cplx}$. The corresponding parameters for L are

$$\Phi_{\mathfrak{q}} = (\phi_{\mathfrak{q}}, \emptyset) \in P^s_{L,d}(H). \tag{15.1h}$$

With this notation, the characters ϕ of T that contribute to $P^{s,\mathfrak{q},finlim}_{G,d}(H)$ have the following characteristic properties.

(1) For each (necessarily real) root $\alpha \in \Delta(L, H)$, $\phi_{\mathfrak{q}}(m_\alpha) = 1$. Equivalently, $\phi(m_\alpha) = \rho_{cplx}(m_\alpha)$.

This is the characteristic property of $P^{s,finlim}_{L,d}(H)$. It guarantees that the principal series representation $I_L(\Phi_{\mathfrak{q}})$ has a unique lowest $L \cap K$-type $\tau_{\mathfrak{q}}(\Phi)$, which is one-dimensional and L-spherical (Proposition 8.5). Consequently $\tau_{\mathfrak{q}}(\Phi)$ has a unique weight $\mu_{\mathfrak{q}}(\Phi) \in \Pi^*(T_f)$. We write

$$\widetilde{\mu}_{\mathfrak{q}}(\Phi) \in \Pi^*(T_{fl}(\mu_{\mathfrak{q}}(\Phi)))$$

for the extended highest weight of $\tau_{\mathfrak{q}}(\Phi)$; recall (again from Proposition 8.5) that this extended highest weight may be parametrized just by $\Phi_{\mathfrak{q}}$ (and so by Φ). Define

$$\widetilde{\mu}(\Phi) = \widetilde{\mu}_{\mathfrak{q}}(\Phi) \otimes 2\rho(\mathfrak{u} \cap \mathfrak{k}) \in \Pi^*(T_{fl}(\mu_{\mathfrak{q}}(\Phi))), \qquad \mu(\Phi) = \widetilde{\mu}(\Phi)|_{T_f}$$

as in Theorem 10.6.

(2) The weight
$$\zeta(\Phi) = d\phi_{\mathfrak{q}} + \rho(\mathfrak{u})$$
takes non-positive values on every coroot of $\mathfrak{h}$ in $\mathfrak{u}$. Equivalently, Φ is strongly $\mathfrak{q}$-compatible.

This second requirement ensures (among other things) that Φ is a discrete standard limit parameter. To ensure that Φ is final, we just need to arrange $I(\Phi) \neq 0$. This is equivalent to either of the next two requirements.

(3) The weight $\mu(\Phi) \in \Pi^*(T_f)$ is dominant for $\Delta^+(K)$.
(3′) For each simple imaginary root $\alpha \in \Delta^+_{im}(G, H)$ which is compact, we have
$$\langle \zeta(\Phi), \alpha^\vee \rangle > 0.$$

Theorem 15.2. *Suppose K is a maximal compact subgroup of the reductive algebraic group G, T_f is a small Cartan subgroup of K, Δ^+_c a choice of positive roots for T_f in K, and other notation is as in* (15.1). *Fix a collection $\mathcal{Q} = \{\mathfrak{q}\}$ of representatives for the θ-stable parabolic subalgebras satisfying conditions (1)–(3) there. For each $\mathfrak{q} \in \mathcal{Q}$, fix a maximally split Cartan subgroup $H \subset L$, and define $P^{s,\mathfrak{q},finlim}_{G,d}(H)$ as in* (15.1e). *Then there is a bijection*
$$\coprod_{\mathfrak{q}\in\mathcal{Q}} P^{s,\mathfrak{q},finlim}_{G,d}(H) \leftrightarrow \Pi^*(K), \qquad \Phi \leftrightarrow \tau(\Phi).$$
Here the first term is a disjoint union of sets of discrete final limit parameters, and the bijection is that of Theorem 11.9 (described also in (15.1) *above).*

Proof. That every irreducible representation of K appears in this correspondence is clear from Theorem 11.9: to show that every θ-stable $\mathfrak{q}$ can be conjugated by K to be compatible with a fixed T_f and $\mathfrak{b}_c$ is elementary. To see that the correspondence is one-to-one, suppose that $(\mathfrak{q}_1, \Phi_1)$ and $(\mathfrak{q}_2, \Phi_2)$ are two parameters in this theorem, with $\tau(\Phi_1) = \tau(\Phi_2)$. The algorithm of Sections 13 and 14, for any $\Phi \in P^{s,\mathfrak{q},finlim}_{G,d}(H))$, reconstructs $\mathfrak{q}$ from one of the highest weights of $\tau(\Phi)$. It follows from this algorithm that $\mathfrak{q}_1$ and $\mathfrak{q}_2$ must be conjugate by T_{fl}/T_f (since this finite group acts transitively on the highest weights of an irreducible representation τ of K). By the choice of $\mathcal{Q}$, necessarily $\mathfrak{q}_1 = \mathfrak{q}_2$. □

Example 15.3. Suppose $G = Sp_4(\mathbb{R})$. In this case $K = U_2$. We can choose $T_f = H_f = U_1 \times U_1$, so that $\Pi^*(T_f) \simeq \mathbb{Z}^2$. Write $\{e_1, e_2\}$ for the standard basis elements of $\mathbb{Z}^2$. The root system is then
$$\Delta_{im,c} = \{\pm(e_1 - e_2)\}, \qquad \Delta_{im,n} = \{\pm 2e_1, \pm 2e_2, \pm(e_1 + e_2)\}.$$
We choose
$$\Delta^+_c = \{e_1 - e_2\}.$$

Since K is connected, irreducible representations are parametrized precisely by their highest weights, which are

$$\{\mu = (\mu_1, \mu_2) \in \mathbb{Z}^2 \mid \mu_1 \geq \mu_2\}.$$

There are exactly eight θ-stable parabolic subalgebras $\mathfrak{q}$ having split Levi factor and compatible with Δ_c^+. I will list them, list the dominant weights μ associated to each, and finally list the corresponding discrete final limit parameters.

The first four are Borel subalgebras containing $\mathfrak{t}_f$. For each I will list only the two simple roots in the Borel subalgebra; these are the *negatives* of the simple roots for the corresponding positive system.

$$\begin{gathered}\mathfrak{q}_1 \leftrightarrow \{-e_1 + e_2, -2e_2\}\\ \{\mu \mid \mu_1 \geq \mu_2 \geq 2\}\\ \{\phi \mid \phi_1 - 1 > \phi_2 \geq 1\}\\ \mu = \phi - (1, -1).\end{gathered}$$

$$\begin{gathered}\mathfrak{q}_2 \leftrightarrow \{-e_1 - e_2, 2e_2\}\\ \{\mu \mid \mu_1 - 1 \geq -\mu_2 \geq 0\}\\ \{\phi \mid \phi_1 > -\phi_2 > 0\}\\ \mu = \phi - (1, -1).\end{gathered}$$

$$\begin{gathered}\mathfrak{q}_3 \leftrightarrow \{e_1 + e_2, -2e_1\}\\ \{\mu \mid 0 \geq -\mu_1 \geq \mu_2 + 1\}\\ \{\phi \mid 0 > -\phi_1 > \phi_2\}\\ \mu = \phi - (1, -1).\end{gathered}$$

$$\begin{gathered}\mathfrak{q}_4 \leftrightarrow \{-e_1 + e_2, 2e_1\}\\ \{\mu \mid -2 \geq \mu_1 \geq \mu_2\}\\ \{\phi \mid -1 \geq \phi_1 > \phi_2 + 1\}\\ \mu = \phi - (1, -1).\end{gathered}$$

These four families of characters ϕ of T_f together exhaust the final standard limit characters (Definition 5.1) which are K-dominant.

The large regions corresponding to $\mathfrak{q}_1$ and $\mathfrak{q}_2$ are separated by the line $\mu_2 = 1$, and this line essentially corresponds to the next parabolic. The Levi factor is $L = U_1 \times Sp_2$, $H_{12} = U_1 \times GL_1$, so a (discrete) character ϕ of H_{12} is a pair (ϕ_1, ϵ), with $\phi_1 \in \mathbb{Z}$ and $\epsilon \in \mathbb{Z}/2\mathbb{Z}$. The corresponding parabolic and dominant weights are

$$\begin{gathered}\mathfrak{q}_{12} = \mathfrak{q}_1 + \mathfrak{q}_2, \qquad \Delta(\mathfrak{u}_1) = \{-e_1 \pm e_2, -2e_1\}\\ \{\mu \mid \mu_1 \geq \mu_2 = 1\}\\ \{\phi \mid \phi_1 \geq 1, \epsilon = 0\}\\ \mu = (\phi_1, 1).\end{gathered}$$

In exactly the same way, the regions corresponding to $\mathfrak{q}_3$ and $\mathfrak{q}_4$ are separated by the line $\mu_1 = -1$, and this line essentially corresponds to the next parabolic. The Levi factor is $L_{34} = Sp_2 \times U_1$, with split Cartan $H_{34} = GL_1 \times U_1$, so a (discrete) character ϕ of H_{34} is a pair (ϵ, ϕ_2), with $\phi_2 \in \mathbb{Z}$ and $\epsilon \in \mathbb{Z}/2\mathbb{Z}$. The corresponding parabolic and dominant weights are

$$\begin{gathered}\mathfrak{q}_{34} = \mathfrak{q}_3 + \mathfrak{q}_4, \qquad \Delta(\mathfrak{u}_1) = \{\pm e_1 - e_2, -2e_2\}\\ \{\mu \mid -1 = \mu_1 \geq \mu_2\}\\ \{\phi \mid \phi_2 \leq -1, \epsilon = 0\}\\ \mu = (-1, \phi_1).\end{gathered}$$

The two Cartans H_{12} and H_{34} are conjugate; the conjugation identifies the character (ϕ_1, ϵ) of H_{12} with (ϵ, ϕ_1) of H_{34}. We have therefore accounted for all of the even characters of H_{12} except the one with $\phi_1 = 0$. Since the imaginary root is $2e_1$, this missing character violates condition (1) of Definition 11.2 to be a standard limit character.

The regions corresponding to $\mathfrak{q}_2$ and $\mathfrak{q}_3$ are separated by the line $\mu_1 + \mu_2 = 0$, and (most of) this line corresponds to a parabolic with Levi factor $L_{23} = U_{1,1}$, with split Cartan $H_{23} = \mathbb{C}^\times$. A (discrete) character Φ of H_{23} corresponds to an integer j, best written in our root system coordinates as $(j/2, -j/2)$. The corresponding parabolic and dominant weights are

$$\begin{gathered}\mathfrak{q}_{23} = \mathfrak{q}_2 + \mathfrak{q}_3, \qquad \Delta(\mathfrak{u}_{23}) = \{-e_1 + e_2, -2e_1, 2e_2\}\\ \{\mu \mid \mu_1 = -\mu_2 > 0\}\\ \{\phi = (j/2, j/2) \mid j > 0 \text{ even}\}\\ \mu = (j/2, j/2).\end{gathered}$$

The condition on ϕ that j should be even is just $\phi_{\mathfrak{q}_{23}}(m_\alpha) = 1$ for the real root α of H_{23}. The condition $j \neq 0$ is condition (1) of Definition 11.2, required to make ϕ a standard limit character. The Weyl group of H_{23} has order 4 (generated by the involutions $z \mapsto \overline{z}$ and $z \mapsto \overline{z}^{-1}$ of $\mathbb{C}^\times$), so it includes elements sending j to $-j$. The characters attached to $\mathfrak{q}_{23}$ therefore exhaust the discrete final limit characters of H_{23} up to the conjugacy.

There remains only the trivial representation of K, corresponding to $\mu = (0,0)$. The corresponding parabolic is $\mathfrak{q}_0 = \mathfrak{g}$, with Levi factor $L_0 = Sp_4$ and split Cartan $H_0 = GL_1 \times GL_1$. The only discrete characters of H_0 are the four characters of order two; and only the trivial one is final (since the elements m_α include all the non-trivial $(\pm 1, \pm 1) \in GL_1 \times GL_1$). The trivial character ϕ corresponds to $\mu = 0$.

16. G-spherical representations as sums of standard representations

In order to complete the algorithm for branching laws from G to K, we need explicit formulas expressing each irreducible representation of K as a finite integer combination of discrete final standard limit representations. We will get such formulas by reduction to the G-spherical case (using cohomological induction as in Section 10). In the G-spherical case, we will use a formula of Zuckerman which we now recall.

We begin as in Section 8 by assuming that

$$G \text{ is split, with } \theta\text{-stable split Cartan } H_s. \tag{16.1a}$$

Fix a discrete final limit parameter

$$\Phi_0 \in P^{s,finlim}_{G,d}(H_s), \tag{16.1b}$$

with

$$\tau(\Phi_0) \in \Pi^*_{G-sph}(K) \tag{16.1c}$$

the corresponding G-spherical representation of K (Proposition 8.5). According to Lemma 8.3, there is a unique extension

$$\tau_0 = \tau_G(\Phi_0) \in \Pi^*_{adm}(G) \tag{16.1d}$$

of $\tau(\Phi_0)$ to a $(\mathfrak{g}, K)$-module (acting on the same one-dimensional space) in which $\mathfrak{s}$ acts by zero. Zuckerman's formula actually writes τ_0 as an integer combination of standard representations of G (in the Grothendieck group of virtual representations); ultimately we will be interested only in the restriction of this formula to K.

At this point we can drop the hypothesis that G be split, and work with any one-dimensional character τ_0 of G. (In fact Zuckerman's identity is for virtual representations of G, and it can be formulated with τ_0 replaced by any finite-dimensional irreducible representation of G.) We will make use of the results only for G split, but the statements and proofs are no more difficult in general.

Lemma 16.2. *Suppose τ_0 is a one-dimensional character of the reductive group G. Suppose that H is a θ-stable maximal torus in G, and that $\Delta^+ \subset \Delta(G, H)$ is a system of positive roots. Write $T = H^\theta$ for the compact part of H. Attached to this data there is a unique shifted Harish-Chandra parameter*

$$\Phi(\Delta^+, \tau_0) = (\phi(\Delta^+, \tau_0), \Delta^+_{im})$$

characterized by the following requirements.

(1) *Δ^+_{im} consists of the imaginary roots in Δ^+.*

Define

$$\rho_{im} \in \mathfrak{t}^* \subset \mathfrak{h}^*$$

as in Definition 11.2 to be half the sum of the roots in Δ^+_{im}*, and* $\rho \in \mathfrak{h}^*$ *to be half the sum of* Δ^+*. Define*

$$\rho_{cplx} = \sum_{\substack{\text{pairs } \{\alpha, \theta\alpha\} \\ \text{cplx in } \Delta^+}} \alpha|_T \quad \in \Pi^*_{adm}(T).$$

(2) *The differential of* $\phi(\Delta^+, \tau_0)$ *is equal to* $\rho + \rho_{im} + d\tau_0|_{\mathfrak{h}}$.
(3) *The restriction of* $\phi(\Delta^+, \tau_0)$ *to* T *is equal to*

$$\rho_{cplx} + 2\rho_{im} + \tau_0|_T.$$

The notation ρ_{cplx} is not inconsistent with that introduced in (11.6i), but it is a bit different. In the earlier setting, we were interested in a positive root system related to a θ-stable parabolic subalgebra, and therefore as close to θ-stable as possible. For every complex root α, either α and $\theta\alpha$ were both positive (in which case α contributed to ρ_{cplx}) or $-\alpha$ and $-\theta\alpha$ were both positive (in which case $-\alpha$ contributed to ρ_{cplx}). In the present setting there is a third possibility: that α and $\theta\alpha$ have opposite signs. In this third case there is no contribution to ρ_{cplx}. The character defined here may therefore be smaller than the one defined in (11.6i).

Proof. The first point is to observe that the two roots α and $\theta\alpha$ have the same restriction to $T = H^\theta$, so ρ_{cplx} is a well-defined character of T. It is very easy to check that

$$\rho|_{\mathfrak{t}} = d\rho_{cplx} + \rho_{im} : \tag{16.3}$$

the point is that complex pairs of positive roots $\{\beta, -\theta\beta\}$ contribute zero to $\rho|_{\mathfrak{t}}$.

Now a character χ of H is determined by its differential $d\chi \in \mathfrak{h}^*$ and by its restriction $\chi_T \in \Pi^*(T)$ to T; the only compatibility required between these two things is $d(\chi_T) = (d\chi)|_{\mathfrak{t}}$. The two requirements in the lemma specify the differential of $\phi(H, \Delta^+, \tau_0)$ and its restriction to T, and (16.3) says that these two requirements are compatible. It follows that $\phi(H, \Delta^+, \tau_0)$ is a well-defined character of H. To see that it is shifted parameter for a standard representation, use Definition 11.2 and the fact that ρ and ρ_{im} are both dominant regular for Δ^+_{im}. (Because τ_0 is a one-dimensional character of G, its differential is orthogonal to all roots.) □

In order to state Zuckerman's formula, we need one more definition: that of "length" for the shifted parameters in Lemma 16.2. All that appears in an essential way in the formulas is differences of two lengths; so there is some choice of normalization in defining them. From the point of view of the Beilinson-Bernstein picture of Harish-Chandra modules, the length is related to the dimension of some orbit of $\mathbf{K}$ on the flag variety of $\mathbf{G}$. The most natural normalizations would define the length to be *equal* to the dimension of the orbit, or equal to its codimension. In fact the original definition is equal to neither of these (but rather to the difference between the dimension of the given $\mathbf{K}$ orbit and the minimal $\mathbf{K}$ orbits). Changing the definition now would surely introduce more problems than it would solve.

Definition 16.4 ([8], Definition 8.1.4)**.** Suppose H is a θ-stable maximal torus in G (cf. (6.1)), and that $\Delta^+ \subset \Delta(G,H)$ is a system of positive roots. The *length of* Δ^+, written $\ell(\Delta^+)$, is equal to

$$\dim_{\mathbb{C}} \mathbf{B}_{c,0} - [(1/2)\#\{\alpha \in \Delta^+ \text{ complex } \mid \theta\alpha \in \Delta^+\} + \#\Delta^+_{im,c} + \dim T].$$

If $\mathfrak{b}$ is the Borel subalgebra corresponding to Δ^+, then the term in square brackets is the dimension of $\mathfrak{b} \cap \mathfrak{k}$: that is, of the isotropy group at $\mathfrak{b}$ for the action of $\mathbf{K}$ on the flag variety of $\mathbf{G}$. The first term is the maximum possible value of this dimension, attained precisely for H fundamental and Δ^+ θ-stable.

The definition in [8] is written in a form that is more difficult to explain; rewriting it in this form is a fairly easy exercise in the structure of real reductive groups. Here is Zuckerman's formula.

Theorem 16.5 (Zuckerman; cf. [8], Proposition 9.4.16)**.** *Suppose G is a reductive group and τ_0 is a one-dimensional character of G. Define*

$$r(G) = (\textit{number of positive roots for } G) - (\textit{number of positive roots for } K).$$

Then there is an identity in the Grothendieck group of virtual admissible representations of finite length

$$\tau_0 = (-1)^{r(G)} \sum_H \sum_{\Delta^+} (-1)^{\ell(\Delta^+)} I(\Phi(\Delta^+, \tau_0)).$$

The outer sum is over θ-stable Cartan subgroups of G, up to K-conjugacy; and the inner sum is over positive root systems for $\Delta(G,H)$, up to the action of the Weyl group $W(G,H)$. The shifted Harish-Chandra parameter $\Phi(\Delta^+,\tau_0)$ is defined in Lemma 16.2, and the standard representation $I(\Phi)$ in Proposition 11.3.

The number $r(G)$ is the codimension of the minimal orbits of $\mathbf{K}$ on the flag manifold of $\mathbf{G}$; so $r(G) - \ell(\Delta^+)$ is the codimension of the $\mathbf{K}$ orbit of the Borel corresponding to Δ^+. The left side of the formula is (according to the Borel-Weil theorem) the cohomology of the flag variety with coefficients in a (nearly) trivial bundle. The terms on the right side are local cohomology groups along the $\mathbf{K}$ orbits, with coefficients in this same (nearly) trivial bundle. The equality of the two sides—more precisely, the existence of a complex constructed from the terms on the right with cohomology the representation on the left—can be interpreted in terms of the "Grothendieck-Cousin complex" introduced by Kempf in [4].

Proof. From the definition of the standard representations in Proposition 11.3, it is clear that $I(\Phi + (\tau_0|_H)) \simeq I(\Phi) \otimes \tau_0$. In this way we can reduce to the case when τ_0 is trivial. That is the case treated in [8], Proposition 9.4.16. □

The terms in this sum correspond precisely to the orbits of $\mathbf{K}$ on the flag variety for $\mathbf{G}$. For the split real form of E_8, there are 320, 206 orbits. (The command `kgb` in the present version of the atlas software first counts these orbits, then provides some descriptive information about them.)

Using the branching laws in Section 12, we can immediately get a version of Zuckerman's formula for arbitrary representations of K.

Theorem 16.6. *Suppose G is a reductive group and (τ, E) is an irreducible representation of K. Let $\Phi \in P^{s,finlim}_{G,d}(H)$ be a discrete final limit parameter so that $\tau(\Phi) = \tau$ is the unique lowest K-type of $I_G(\Phi)$ (Theorem 11.9). Choose a strongly Φ-compatible parabolic subalgebra $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$ as in* (11.6)*, so that $\tau_{\mathfrak{q}}$ is a one-dimensional L-spherical character of $L \cap K$. In particular, $\tau_{\mathfrak{q}}$ is the unique lowest $L \cap K$-type of the principal series representation $I_L(\Phi_{\mathfrak{q}})$.*

$$\tau(\Phi) = (-1)^{r(L)} \sum_{H_1 \subset L} \sum_{\Delta^+} \sum_{A \subset \Delta(\mathfrak{u}^{op} \cap \mathfrak{s})} (-1)^{|A| + \ell(\Delta^+)} I_G(\Phi(\Delta^+, \tau_{\mathfrak{q}})^{\mathfrak{q}}_{\mathfrak{g}} + 2\rho(A))|_K$$

Here the outer sum is over θ-stable Cartan subgroups of L, up to conjugacy by $L \cap K$, and the next sum is over positive root systems for $\Delta(L, H_1)$, up to the action of $W(L, H_1)$. The translation from the parameter $\Phi(\Delta^+, \tau_{\mathfrak{q}})$ for L to a parameter for G is accomplished by (12.3)*. The character $2\rho(A)$ is any extension to H_1 (for example, the one trivial on $A_{1,0}$) of the character of T_1 on the sum of the weights in A.*

Since the left side in this theorem is a representation only of K and not of G, this is not an identity (as in Zuckerman's formula) of virtual representations of G. The parameters for G appearing on the right need not be standard limit parameters (although the "first term," corresponding to $H_1 = H$, is just $(\Phi_{\mathfrak{q}})^{\mathfrak{q}}_{\mathfrak{g}} = \Phi$). In order to get a formula expressing the irreducible representation $\tau(\Phi)$ of K in terms of standard limit representations of G, we therefore need to rewrite each term on the right as an integer combination of standard limit representations. The character identities of Hecht and Schmid in [7] provide algorithms for doing exactly that, which appear to be reasonable to implement on a computer.

References

[1] Anne-Marie Aubert and Roger Howe, *Géométrie des cônes aigus et application à la projection euclidienne sur la chambre de Weyl positive*, J. Algebra **149** (1992), 472–493.

[2] Armand Borel and Nolan Wallach, *Continuous Cohomology, Discrete Subgroups, and Representations of Reductive Groups*, Princeton University Press, Princeton, New Jersey, 1980.

[3] Jacques Carmona, *Sur la classification des modules admissibles irréductibles*, Noncommutative Harmonic Analysis and Lie Groups, 1983, pp. 11–34.

[4] George Kempf, *The Grothendieck-Cousin complex of an induced representation*, Adv. in Math. **29** (1978), 310–396.

[5] Anthony W. Knapp and David A. Vogan Jr., *Cohomological Induction and Unitary Representations*, Princeton University Press, Princeton, New Jersey, 1995.

[6] Robert P. Langlands, *On the classification of representations of real algebraic groups*, Representation Theory and Harmonic Analysis on Semisimple Lie Groups, 1989, pp. 101–170.

[7] Wilfried Schmid, *Two character identities for semisimple Lie groups*, Noncommutative Harmonic Analysis and Lie groups, 1977, pp. 196–225.

[8] David A. Vogan Jr., *Representations of Real Reductive Lie Groups*, Birkhäuser, Boston-Basel-Stuttgart, 1981.

[9] ———, *Unitary Representations of Reductive Lie Groups*, Annals of Mathematics Studies, Princeton University Press, Princeton, New Jersey, 1987.

[10] ———, *The method of coadjoint orbits for real reductive groups*, Representation Theory of Lie Groups, 1999.

E-mail address: `dav@math.mit.edu`

Department of Mathematics, MIT, Cambridge, MA 02139