Combinatorics for the representation theory of real reductive groups

Fokko du Cloux

July 5, 2005

These are notes for the third meeting of the *Atlas of reductive Lie groups* project at AIM, in Palo Alto. They describe how to take the description of the representation theory of a real reductive Lie group (cf. Jeff Adams' notes from last year) to finite combinatorial terms, that can be implemented in a computer.

These ideas evolved during my stay at MIT last fall, and benefited immensely from innumerable discussions with David Vogan, as well as from an intensive session with Jeff during a visit to Maryland.

1 Real forms and strong real forms

1.1. Throughout these notes, G will denote a connected complex reductive algebraic group. Define a *pinning* of G to be the datum of a maximal torus T, a Borel B containing T, and a set $\{X_\alpha\}_{\alpha\in\Pi}$, where Π is the set of simple roots for T in G defined by B, and $X_\alpha \in \mathrm{Lie}(G)$ is a root vector for the simple root α. We will fix once and for all such a pinning $\mathcal{P}$.

To the choice of T also corresponds a *root datum* $(X, R, X^\vee, R^\vee)$. Here X is the character group of T, $X^\vee = \mathrm{Hom}(\mathbf{C}^\times, T)$ its cocharacter group, and R and $R^\vee$ are the roots and coroots for T in G, respectively. The choice of B turns our root datum into a *based* root datum $(X, \Pi, X^\vee, \Pi^\vee)$, where Π and $\Pi^\vee$ are the sets of simple roots and coroots, respectively. It is worth noticing that conversely, given two lattices in duality which we might as well take to be $\mathbf{Z}^n$, and given two finite subsets Π and $\Pi^\vee$ of $\mathbf{Z}^n$, together with a bijection $\alpha \to \alpha^\vee$ from Π to $\Pi^\vee$, we get a based root datum if and only if the matrix $(<\alpha, \beta^\vee>)_{\alpha,\beta\in\Pi}$ is a Cartan matrix (i.e., after permutation of Π, a block-diagonal matrix whose diagonal entries are either 0 or one of the familiar Cartan types A_n–G_2.)

1.2 Exercise. — Show that up to $\mathbf{GL}(n, \mathbf{Z})$-conjugation there are exactly three types of root data with rank $n \geq 2$ and semisimple rank one (*i.e.*, $|\Pi| = 1$.) In other words, up to $\mathbf{GL}(n, \mathbf{Z})$-conjugation (using the transpose inverse action on the dual side) there are only three types of pairs $(\alpha, \alpha^\vee)$ s.t. $< \alpha, \alpha^\vee >= 2$.

1.3. In principle, a *real form* of G should be an antiholomorphic involution σ of G; and also in principle, one should consider two real forms to be equivalent if they are conjugate under the full automorphism group $\mathrm{Aut}(G)$. However, it turns out that the appropriate notion of equivalence is rather equivalence under conjugation by G itself (for (quasi)simple G, this makes a difference only in type D_n with n even; for general semisimple or reductive G, however, the difference is big.)

Also, G-conjugacy classes of antiholomorphic involutions are in $(1,1)$-correspondence with G-conjugacy classes of ordinary involutions of G. (To set up such a correspondence, choose a compact real form of G, with antiholomorphic involution σ_0. Then any σ may be conjugated to commute with σ_0; similarly any involution θ may be conjugated to commute with σ_0; then the map $\sigma \to \theta = \sigma\sigma_0$ sets up a bijection between anti-involutions commuting with σ_0 and involutions commuting with σ_0.) In this correspondence, the group $K = G^\theta$ of fixed points of θ in G is the complexification of a maximal compact subgroup of G^σ (the group of real points of G for σ.) In particular, the component group of G^σ is isomorphic to the component group of K.

Henceforth, we view real forms of G as G conjugacy classes of ordinary involutions. In this picture, the identity involution corresponds to the compact real form. Note that the real reductive Lie groups we are dealing with in this way are the *full* groups of real points of complex connected reductive groups defined over $\mathbf{R}$. To get to results about open subgroups of such (e.g., their identity components), requires some minor modifications which we will not go into here.

1.4. Consider the exact sequence

$$1 \to \mathrm{Int}(G) \to \mathrm{Aut}(G) \to \mathrm{Out}(G) \to 1$$

where $\mathrm{Int}(G) = G/Z(G)$. We say that two involutions θ and θ' are in the same inner class, or are inner to each other, if they have the same image in $\mathrm{Out}(G)$. Obviously this defines a partition of the real forms of G.

It is well-known that the group $\mathrm{Int}(G)$ acts simply transitively on the set of pinnings (cf. 1.1.) Therefore $\mathrm{Out}(G)$ may be identified with the stabilizer in $\mathrm{Aut}(G)$ of our fixed pinning $\mathcal{P}$. Clearly, the permutation of the X_α induced by an element of $\mathrm{Out}(G)$ is entirely determined by the corresponding permutation of the set of simple roots Π; therefore $\mathrm{Out}(G)$ may be identified with a subgroup of the group of automorphisms of the based root datum $(X, \Pi, X^\vee, \Pi^\vee)$. Conversely, any automorphism of the root datum may be lifted to an automorphism of G *(reference??)*; hence $\mathrm{Out}(G)$ is exactly the group of automorphisms of the based root datum.

Note that when G is semisimple, any automorphism of the based root datum is entirely determined by the automorphsim of the Dynkin diagram it induces; but here we do not have surjectivity in general (think of the case where $G = \mathbf{PSL}(2) \times \mathbf{SL}(2)$, and the automorphism interchanges the two factors of the Dynkin diagram.) However, we do have surjectivity when G is simply connected, or adjoint.

The basic starting datum for the description of representations that we are using in these notes is the complex group G, together with a given inner class of real forms, i.e., a given involution γ of the based root datum. It is natural to consider the inner class up to conjugacy in $\mathrm{Aut}(G)$; it turns out that there are then in all cases finitely many possibilities for any given G. (In the case where G is a torus, this hinges on the fact that there are finitely many conjugacy classes of involutions in each group $\mathbf{GL}(r, \mathbf{Z})$, an interesting exercise to which we will come back later.)

1.5 Example. — Let $G = \mathbf{SL}(n)$, $n \geq 3$. Then there is one non-trivial automorphism of the Dynkin diagram, so there are two inner classes of real forms. The inner class corresponding to the identity automorphism may be called (for any G) the *equal rank* inner class, as it corresponds to those real forms for which $\mathrm{rk}(K) = \mathrm{rk}(G)$. We will see that in our example, this inner class is made up of the real forms $\mathbf{SU}(p, q)$, $p + q = n$, $0 \leq q \leq n/2$.

For the other inner class, again we will see that there are either one or two real forms, depending on the parity of n : if n is odd, there is just one real form, *viz.* the split form $\mathbf{SL}(n, \mathbf{R})$; if n is even, there is in addition the form $\mathbf{SL}(n/2, \mathbf{H})$.

1.6. Given our choice of pinning $\mathcal{P}$ in 1.1, and given an inner class of real forms, there is a unique representative θ_{fund} of this inner class which belongs to $\mathrm{Aut}(G, \mathcal{P})$. We say that the real form corresponding to θ_{fund} is the *fundamental form* of the inner class (for instance, for the equal rank inner class, the fundamental form is the compact one, and θ_{fund} is the identity.)

Now we consider the semidirect product

$$G^\Gamma = G \rtimes \mathbf{Z}_2 = G \coprod G.\delta$$

where the semidirect product is through the automorphism θ_{fund}, *i.e.*, the action of $\text{int}(\delta)$ on G is θ_{fund}.

A *strong involution* for G in the chosen inner class is an element $x \in G.\delta$ such that $x^2 \in Z$; the involution corresponding to x is the involution $\theta_x = \text{int}(x)$. A *strong real form* is a G-conjugacy class of strong involutions. Clearly, the map $x \to \text{int}(x)$ goes over to a surjection from strong real forms to real forms, and also clearly this map is bijective when G is adjoint.

1.7. Let W be the Weyl group of $(X, R, X^\vee, R^\vee)$. Then W acts simply transitively on the set of bases, so our group $\text{Out}(G) = \text{Aut}(X, \Pi, X^\vee, \Pi^\vee)$ is also isomorphic to $\text{Aut}(X, R, X^\vee, R^\vee)/W$. Viewing our involution $\gamma \in \text{Out}(G)$ as an involution of X, we may consider the involution $-{}^t\gamma$ of $X^\vee$; of course this will almost never fix $\Pi^\vee$, but there is a unique involution $\gamma^\vee$ congruent to it modulo W that will (in fact it should be clear that $\gamma^\vee = -{}^t\gamma.w_0$, where w_0 is the longest element in W.) In this way, we define an inner class for the dual group $G^\vee$, with based root datum $(X^\vee, \Pi^\vee, X, \Pi)$, which we call the inner class dual to γ. In the example of $\mathbf{SL}(n)$ which we began in 1.5, the dual group is $\mathbf{PSL}(n)$, and the duality interchanges the "equal rank" and "split" inner classes.

1.8. The fundamental real form may be characterized by the fact that there exists a Borel preserved by θ_{fund}. It turns out that there is also a unique real form satisfying the dual condition, namely that there is a Borel B such that $\theta(B) = \overline{B}$ is the opposite Borel. This is called the *quasisplit* real form in the class. We will often use notation like θ_{qs}, x_{qs}, for data pertaining to this real form.

2 The one-sided parameter space

2.1. We are now going to describe the main combinatorial construction. The set thus obtained plays a fundamental technical role for the main classification problems that we have to deal with.

Recall the notation of 1.6. The set we are interested in is the set $\mathcal{X}$ of triples (x, T', B') up to G-conjugacy, where

(a) $x \in G.\delta$, and $x^2 \in Z(G)$ (i.e., x is a strong real form representative as defined in 1.6);

(b) T' is a maximal torus of G, and B' is a Borel containing T';

(c) the involution $\theta_x = \text{int}(x)$ normalizes T'.

Of course, all pairs $T' \subset B'$ are conjugate in G. So we may as well assume that $T' = T$ and $B' = B$, where (T, B) is the pair chosen in 1.1. By construction T is stable under $\mathrm{int}(\delta)$, so $\mathrm{int}(x)$ normalizes T if and only if x belongs to $N.\delta$, where N is the normalizer of T in G. With these remarks, our set $\mathcal{X}$ becomes the set of elements $x \in N.\delta$ such that $x^2 \in Z(G)$, up to T-conjugacy (because the stabilizer of the pair (T, B) in G is exactly T.)

2.2. It is in fact not hard to see that the set $\mathcal{X}$ is in canonical bijection with the disjoint union over all strong real forms of G of $B \times K$-orbits in G (or equivalently, K-orbits in G/B, or B-orbits in G/K), as described by Richardson and Springer [3].

Indeed, fix a strong involution x, which we may assume to lie in $N.\delta$. Then we may write for $g \in G$:

$$g.x.g^{-1} = g.\theta_x(g^{-1}).x$$

which shows that the G-conjugacy class of x identifies with the image of the map $g \to \tau_x(x) := g.\theta_x(g^{-1})$. Now of course $\tau_x(G)$ is isomorphic to G/K, and it is shown in [3] (or rather, in [4]) that the B-conjugation orbits in $\tau_x(G)$ (which correspond to the left B-orbits in G/K_x) are in $(1,1)$ correspondence with the T-conjugation orbits in $\tau_x(G) \cap N$. After multiplication by x, this amounts to taking the elements in $\mathcal{X}$ which correspond to the strong real form defined by x.

2.3. We have denoted W the Weyl group of (G, T). We may form the semidirect product

$$W^\Gamma = W \coprod W.\delta$$

just as in 1.6. When $w.\delta$ is an involution in W^Γ, we say that $w \in W$ is a *twisted involution* for the involution of W induced by δ (which can be read off from the Dynkin diagram involution corresponding to δ). The natural map $N \to W$ restricts to a map $x \to \tau_x$ (not the same τ_x as in 2.2!) from $\mathcal{X}$ to $\mathcal{I}$, where $\mathcal{I}$ is the set of elements $w.\delta$, with w a twisted involution.

One way to interpret this map is to look at the action of $W.\delta$ on the torus, and to note that τ_x acts on T as the restriction of the involution θ_x to the torus. We will say that $\mathcal{I}$ is the set of *root datum involutions* for the given G and inner class.

It turns out that the map $x \to \tau_x$ is surjective (5.6). For each $\tau \in \mathcal{I}$, we denote $\mathcal{X}_\tau$ the fiber of $\mathcal{X}$ over τ.

2.4 Proposition. — *Fix $\tau \in \mathcal{I}$. Denote T_τ the subgroup of elements $t \in T$ such that $t.\tau(t)$ is central in G, and $T^{-\tau}$ the subgroup of elements t such that $\tau(t) = t^{-1}$. Then T_τ acts simply transitively on the set of strong involutions in $N.\delta$ lying over τ, and $T_\tau/T_\circ^{-\tau}$ acts simply transitively on the fiber $\mathcal{X}_\tau$.*

Proof. — Fix a strong involution x in $N.\delta$ lying over τ. Then any element of $N.\delta$ lying over τ can be uniquely written in the form $t.x$ with $t \in T$. We have $(t.x)^2 = t.\mathrm{int}(x)(t).x^2$, and $\mathrm{int}(x)(t) = \tau(t)$, so $(t.x)^2$ is central if and only if $t.\tau(t)$ is central, whence our first claim. Clearly T_τ contains $T^{-\tau}$.

Let $s \in T$. Then the conjugation action of s on $t.x$ may be written as

$$s.t.x.s^{-1} = s.t.\mathrm{int}(x)(s^{-1})x = s.t.\tau(s^{-1}).x = s.\tau(s^{-1}).t.x$$

because T is commutative, and therefore the action is just left multiplication by $s.\tau(s^{-1})$. Clearly for each $s \in T$ the element $s.\tau(s^{-1})$ is in $T^{-\tau}$, and even in $T_\circ^{-\tau}$ because T is connected. On the other hand, if $s \in T^{-\tau}$, we have $s.\tau(s^{-1}) = s^2$, so the image of the map $s \to s.\tau(s^{-1})$ is exactly the identity component of $T^{-\tau}$, and we are done.

2.5 Corollary. — *When G is semisimple, the set $\mathcal{X}$ is finite, and each fiber $\mathcal{X}_\tau$ carries a simply transitive action of a finite abelian group, canonically defined by τ. For general G, the same conclusion holds for each set $\mathcal{X}(z) := \{x \in \mathcal{X} \mid x^2 = z\}$, with $z \in Z(G)$ fixed (note that x^2 depends only on the conjugacy class of x, and may therefore be defined at the level of $\mathcal{X}$); $\mathcal{X}_\tau(z)$ carries a simply transitive action of the elementary abelian two-group $T^{-\tau}/T_\circ^{-\tau}$.*

Proof. — From the proof of the proposition, we see that $(t.x)^2 = (t'x)^2$ if and only if $t.\tau(t) = t'.\tau(t')$, which may be rewritten as $t^{-1}.t' \in T^{-\tau}$. So the group $T^{-\tau}$ acts simply transitively on each set of strong involutions in $N.\delta$ with fixed square; it follows that $T^{-\tau}/T_\circ^{-\tau}$ acts simply transitively on each $\mathcal{X}_\tau(z)$.

When G is semisimple, $Z(G)$ is finite, so the whole group $T_\tau/T_\circ^{-\tau}$ is finite.

2.6 Corollary. — *Each $\mathcal{X}_\tau(z)$ has the structure of an affine space over the two-element field $\mathbf{F}_2$.*

3 Classification of Cartan subgroups and determination of real Weyl groups

3.1. It is clear from the definition that in the natural conjugation action of N on $\mathcal{X}$, the torus T acts trivially (precisely because $\mathcal{X}$ has been defined as a set of T-orbits), and so gives rise to an action of the Weyl group W. Our objective in this section is to study the orbits of this action, and to show their relation to the classification of Cartan subgroups for the various strong real forms the given inner class. This is also done in [3], where additional results are given (see in particular sect. 9 of that paper.)

3.2. The main observation is the following. Let $\mathcal{T}$ be the set of pairs (x, H), where x is a strong involution in G, and H is a maximal torus in G normalized by $\mathrm{int}(x)$, up to G-conjugation. We look at $\mathcal{T}$ in two different ways.

First, we may always conjugate by an element of G so that $H = T$. Then we see that $\mathcal{T}$ identifies with the set of strong involutions $x \in N.\delta$, up to N-conjugation; but clearly this is also the set of W-orbits in $\mathcal{X}$.

Second, we may choose a set $\{x_i\}_{i\in I}$ of representatives of G-conjugacy classes of elements x (in other words, a set of representatives of strong real forms of G in our given inner class.) Then every (x, H) is G-conjugate to an element of the form (x_i, H), and H is now determined up to conjugacy by the stabilizer of x_i in G. Since for any involution θ of G there are θ-stable tori, all x_i, $i \in I$, will occur. We note that $\mathrm{int}(g)(x_i) = x_i$ is equivalent to $\mathrm{int}(x_i)(g) = g$, so the stabilizer of x_i in G is just the fixed point group K_i of the corresponding involution $\theta_i = \mathrm{int}(x_i)$. So from this picture we see that the set $\mathcal{T}$ also identifies with the disjoint union of the sets of K_i-conjugacy classes of θ_i-stable maximal tori in G. (In the language of groups of real points that we have been avoiding, this is also the set of $(G, \sigma_i)(\mathbf{R})$-conjugacy classes of real maximal tori in $(G, \sigma_i)(\mathbf{R})$, where σ_i is an antiholomorphic involution corresponding to θ_i.)

So we have proved the following

3.3 Theorem. — *The set of W-orbits in $\mathcal{X}$ is in natural $(1,1)$-correspondence with the disjoint union over all strong real forms of G of the set of K_x-conjugacy classes of θ_x-stable Cartan subgroups in G, where x is a representative of the strong real form, $\theta_x = \mathrm{int}(x)$, and $K_x = G^{\theta_x}$.*

3.4. In practice, the W-orbits in $\mathcal{X}$ will be computed by picking a set of representatives for the set of W-orbits in $\mathcal{I}$ (an elementary Weyl group computation), and then for each such representative τ, computing the W^τ-orbits in the fiber $\mathcal{X}_\tau$, where W^τ denotes the stabilizer of τ in W, also the set of $w \in W$ such that $\tau(w) = w$. The delicate issue here is the choice of basepoint in the fiber; we will come back to that in 6.10.

3.5. The construction in 3.2 also allows us to compute real Weyl groups. Given G, an involution θ and a θ-stable Cartan subgroup H, we denote $W(K, H)$ the group $N_K(H)/Z_K(H)$, with $K = G^\theta$. (The natural definition in terms of groups of real points gives rise to the same group.) Now let x be a strong involution such that $\mathrm{int}(x) = \theta$. Then $N_K(H)$ is the stabilizer in G

of the pair (x, H); and of course $Z_K(H)$ is just $H \cap K$. If we go over to the first picture, where $H = T$ and $x \in N.\delta$, we see that $W(K, H)$ is the image in W of the centralizer of x in N, which is also $Z_N(x)T/T$. But $Z_N(x)T$ is also the stabilizer in N of the T-orbit of x; so $W(K, H)$ may be identified with the stabilizer in W of the image of x in $\mathcal{X}$.

In other words, we have proved the following:

3.6 Theorem. — *In the description of Theorem 3.3, the real Weyl group corresponding to a given θ_x-stable Cartan H is isomorphic to the stabilizer in W of any element of the corresponding W-orbit in $\mathcal{X}$.*

3.7 Proposition. — *For any given strong real form with representative x, the map $H \to \tau$ from K_x-conjugacy classes of θ_x-stable Cartan subgroups to W-conjugacy classes of root datum involutions is* injective.

Proof. — This result is essentially Proposition 2.5 in [3]. Let us recall the main idea of the proof. The statement is equivalent to saying that if x and $x' = g.x.g^{-1}$ are two strong involutions lying over the same root datum involution τ, then they are N-conjugate. Now the hypothesis is that there exists $t \in T$ such that $g.x.g^{-1} = g.\theta_x(g^{-1}).x = t.x$. Now we apply Proposition 2.3 from [3], which says that if $t = g.\theta_x(g^{-1})$ for some $g \in G$, then there is also an $n \in N$ such that $t = n.\theta_x(n^{-1})$; this will translate to $x' = n.x.n^{-1}$, whence our result.

3.8 Corollary. — *For any given real form of G, the set of K-conjugacy classes of θ-stable Cartans may be canonically identified with a subset of $\mathcal{I}/W$.*

3.9. One can endow the set $\mathcal{I}/W$ with a poset structure, as follows. We say that $[\tau] \leq [\tau']$ if and only if we may choose the representatives τ and τ' such that the fixed point space of τ in $\mathfrak{t} = \mathrm{Lie}(T)$ contains that of τ'. Since this condition implies in particular that $\dim(\mathfrak{t}^{\tau'}) \leq \dim(\mathfrak{t}^{\tau})$, and $\mathfrak{t}^{\tau} = \mathfrak{t}^{\tau'}$ implies $\tau = \tau'$, this is indeed an order relation. The poset thus obtained has a unique minimal element (the orbit of δ, corresponding to the fundamental Cartan in each real form), and a unique maximal element (this is reached if and only if the real form is quasisplit, and is then the unique most split Cartan for this real form.) We will see that for any real form of G in our inner class, the image in $\mathcal{I}/W$ of the set of conjugacy classes of Cartan subgroups is an interval of the form $[[\delta], [\tau_{\max}]]$, where $[\tau_{\max}]$ corresponds to the most split Cartan for the given real form.

4 Classification of real forms and strong real forms

4.1. Let us now show how Theorem 3.3 yields a classification of real forms and strong real forms in terms of W-orbits.

It is known *(reference??)* that for each involution θ in our chosen inner class, there is exactly one K-conjugacy class of fundamental Cartan subgroups, *i.e.*, θ-stable Cartan subgroups that are contained in a θ-stable Borel. This implies that the set of G-conjugacy classes of strong involutions is in bijection with the set of G-conjugacy classes of pairs (x, H), with H $\mathrm{int}(x)$-stable and fundamental. In the picture of Theorem 3.3, these Cartan subgroups are the ones that map to the orbit of δ in $\mathcal{I}$. So, in the notation of 2.3, the set of strong real forms in our inner class is in bijection with the set of W^δ-orbits in $\mathcal{X}_\delta$ (cf. 3.4).

The real forms of G in our given inner class are classified by a similar computation in the adjoint group.

4.2 Example. — Consider the case where $G = \mathbf{SL}(2)$. Here there is only one inner class of real forms, which is therefore the equal rank one. The fiber $\mathcal{X}_\delta$ is then just the group of elements in T with square ± 1, i.e., it is the subgroup $T(4) \simeq \mathbf{Z}_4$ of elements of T with order dividing four.

The action of the non-trivial element of the Weyl group is by $t \to t^{-1}$. Hence there are three orbits : $\{1\}$, $\{-1\}$ and $\{i, -i\}$ (in the obvious identification of T with $\mathbf{C}^\times$). So there are three strong real forms : two corresponding to the compact real form, and one corresponding to the split one.

4.3 Example. — (example 1.5, continued) Consider again the case of $\mathbf{SL}(n)$, $n \geq 3$. To compute the classification of real forms, we go over to the adjoint group $\mathbf{PSL}(n)$.

For the equal rank inner class, we have $\delta = \mathrm{Id}$, so that $W^\delta = W$, and $\mathcal{X}_\delta = \{t \in T \mid t^2 = 1\}$. So the set of real forms is in $(1,1)$ correspondence with the set of $\mathfrak{S}_n$-orbits in $\mathbf{Z}_2^n/\Delta$, where Δ denotes the diagonal. It is now an easy exercise to check that the number of real forms is as stated in 1.5.

For the other inner class, we hace seen in Corollary 2.6 that $\mathcal{X}_\delta$ carries a simply transitive action of the group $T^{-\delta}/T_\circ^{-\delta}$. This is also the component group of the group of real points of the real form defined by the involution $-\delta$ of T. But in this case it is very easy to determine the structure of $T(\mathbf{R})$, because the character lattice of T (which is just the root lattice of the root system) has a basis that is permuted by δ. We find that $T(\mathbf{R})$ is a complex torus when n is even, hence X^δ is a singleton, and $T(\mathbf{R})$ is the direct product of a one-dimensional split real torus with a complex torus when n is odd; hence $|X^\delta| = 2$ in this case. In the first case, it is already clear that there can only be one real form for $\mathbf{SL}(n, \mathbf{R})$ in this inner class, which must necessarily be the split form. In the second case, there could be one or two; we will see in a moment that in fact there must be two.

To compute the strong real forms, in the equal rank case, we have to compute the W-orbits in an abelian group of order $n.2^{n-1}$, with $n-2$ cyclic factors of order 2 and one cyclic factor of order $2n$. When n is odd, it is in fact clear that this group is the direct product of $Z(G)$ and $T(2)$; so in this case we just get n copies of the orbit picture in the adjoint group, and therefore n isomorphic strong real forms for each real form.

When n is even, the situation gets more interesting. For instance, when $n = 4$, one has three real forms, with orbits of cardinalities 1 for $\mathbf{SU}(2)$, 4 for $\mathbf{SU}(3,1)$ and 3 for $\mathbf{SU}(2,2)$ (the correspondence between orbits and the usual nomenclature of real forms will be explained in 4.6 below.) The strong real forms are determined by looking at the $\mathfrak{S}_4$-orbits in the group of diagonal matrices of the form $(\varepsilon_1 e^{ik\pi/4}, \varepsilon_2 e^{ik\pi/4}, \varepsilon_3 e^{ik\pi/4}, \varepsilon_4 e^{ik\pi/4})$, where $\varepsilon_j = \pm 1$ for all j, $0 \leq k \leq 3$, and the product of the signs is 1 for k even, -1 for k odd. Then for k even, there are three orbits, two of cardinality one, corresponding to strong real forms isomorphic to $\mathbf{SU}(4)$, and one of cardinality six, corresponding to a strong real form isomorphic to $\mathbf{SU}(2,2)$; for k odd, there are two orbits of cardinality four, corresponding to strong real forms isomorphic to $\mathbf{SU}(3,1)$. In particular, the possible values for $z = x^2$ for the quasisplit forms are 1 and -1. In general, there are always n strong real forms for each real form, except for the quasisplit form when there are $n/2$; in fact, one may show that if we partition the strong real forms according to the values of x^2, then there are just two types of orbit pictures, according to whether $z^{n/2} = 1$ or -1. (this example and a number of others are also given in [1].)

Finally, for the non-equal rank inner class, we note that θ acts on the center by inversion (because this is clear for the split form, and the action on the center is the same for all involutions in a given inner class.) So the square of a strong involution in this class can only take the values ± 1. When n is odd, the only possible value is of course 1. When n is even, *(well, what do you know! when n is even, the square should be* 1 *for* $\mathbf{SL}(n/2, \mathbf{H})$*, and* -1 *for* $\mathbf{SL}(n, \mathbf{R})$*, right, but how to prove it?)*

4.4. Let $\tau \in \mathcal{I}$ be a root datum involution, and let θ be an involution of G in our chosen inner class, inducing τ.

The datum of τ yields the classification of roots into real, imaginary and complex, in the usual way. Now it is easy to see that once the root datum involution is fixed, the involution θ is entirely determined by its action on the root vectors X_α of the pinning $\mathcal{P}$ (just because those root vectors and the $X_{-\alpha}$ generate the Lie(G)). Of course, $\theta(X_\alpha)$ should be a root vector for the root $\theta(\alpha)$. Then one sees that if $\theta(\alpha) \neq \alpha$, *i.e.*, if α is real or complex, all the possibilities for θ are T-conjugate *(reference?? is this even right?)*. So the only real choices are for the imaginary roots α that belong to our chosen basis. For each such α, we have $\theta(X_\alpha) = \pm X_\alpha$.

In general, a *grading* of a root system Φ is a map $\mathrm{gr} : \Phi \to \mathbf{Z}_2$ which satisfies $\mathrm{gr}(\alpha) = \mathrm{gr}(-\alpha)$ and $\mathrm{gr}(\alpha+\beta) = \mathrm{gr}(\alpha) + \mathrm{gr}(\beta)$ for all α, β in Φ such that $\alpha + \beta \in \Phi$. Once a basis of Φ is chosen, it is clear that such a grading is entirely determined by the degrees of the simple roots, and that conversely, the degrees of the simple roots may be chosen arbitrarily. Hence there are 2^r possible choices, where r is the rank of the root system Φ.

Clearly θ induces a grading of the imaginary root system Φ_i, by setting $\mathrm{gr}(\alpha) = 0$ if $\theta(X_\alpha) = X_\alpha$ (the compact roots), and $\mathrm{gr}(\alpha) = 1$ when $\theta(X_\alpha) = -X_\alpha$ (the non-compact roots.) It may be shown that this grading, together with the datum of τ, defines θ up to T-conjugacy. Hence we get an injection from the set of real forms for which the Cartan of type τ is defined, to the set of W^τ-conjugacy classes of gradings of the imaginary root system. A delicate aspect of this is that this injection is usually far from being a bijection. One of our objectives is going to be to determine its image.

4.5. For the fundamental torus, and G adjoint, say, the correspondence can be made very precise. Of course, for the equal rank case, all roots are imaginary, so we just have to deal with gradings of the root system $\Phi = \Phi(G, T)$. The fundamental involution θ_{fund} induces the trivial grading where all roots are compact. If we denote $T(2)$ the subgroup of elements of order two in T, we get $T^{-\delta} = T(2)$ in this case, from which it follows easily that all gradings are allowed, and that in fact real forms are in $(1,1)$-correspondence with W-conjugacy classes of gradings.

For the general case where δ is arbitrary, the grading of the imaginary roots is entirely determined by its restriction to those imaginary roots that are in Π. The reasoning to prove this is as follows: let β be a positive imaginary root, and α simple such that $< \beta, \alpha^\vee >$ is positive. If α is imaginary, of course $\beta' = \beta - \alpha$ is again positive imaginary, and choosing root vectors $X_{\beta'}$ and $X_\beta = [X_\alpha, X_{\beta'}]$, we see that $\mathrm{gr}(\beta) = \mathrm{gr}(\beta') + \mathrm{gr}(\alpha)$ is determined inductively. Now assume α is complex. Then we also have $< \beta, \delta(\alpha)^\vee >$ positive. Now there are two cases. If α and $\delta(\alpha)$ are not adjacent, then we see that $< \beta - \alpha, \delta(\alpha)^\vee >$ is again positive, and we may write $\beta = \beta' + \alpha + \delta(\alpha)$ for a positive imaginary β'. Now we note that X_α and $X_{\delta(\alpha)}$ commute, so we may pick root vectors $X_{\beta'}$ and $X_\beta = [X_\alpha, [X_{\delta(\alpha)}, X_{\beta'}]]$. Then

$$\theta(X_\beta) = [X_{\delta(\alpha)}, [X_\alpha, (-1)^{\mathrm{gr}(\beta')} X_{\beta'}]] = (-1)^{\mathrm{gr}(\beta')} X_\beta$$

from the commutativity of X_α and $X_{\delta(\alpha)}$, and we are done. A simple case-by-case analysis reveals in fact that α and α' can be adjacent only in type A_n with n even, and then only when $\beta = \alpha + \delta(\alpha)$, i.e., essentially in type A_2. Then we may take $X_\beta = [X_\alpha, X_{\delta(\alpha)}]$ and

$$\theta(X_\beta) = \theta([X_\alpha, X_{\delta(\alpha)}]) = [X_{\delta(\alpha)}, X_\alpha] = -X_\beta$$

so the root β must be non-compact.

So here, there will be in general many gradings that are not allowed. The computation of $T^{-\delta}/T_0^{-\delta}$ that we did for $\mathbf{PSL}(n)$ in 4.3 trivially generalizes, and we get a group of order 2^r, where r is the number of δ-fixed elements in Π, where we then have to describe the action of the group W^δ.

4.6 Example. — (example 4.3, continued) Now we are in a position to identify the real forms of $\mathbf{SL}(n)$ corresponding to the orbits in 4.3. The $\mathfrak{S}_n$-orbits in $\mathbf{Z}_2^n/\Delta$ are classified by their cardinality, up to complement, i.e., by an integer q, with $0 \le q \le n/2$. Taking the representative with the ones at the end, we see that we may represent the corresponding involution by conjugation with the diagonal matrix that has $p = n - q$ ones followed by q minus ones. The action of this matrix on the X_α for α simple (which have a single 1 just above the diagonal) is trivial except for the single case where the non-zero entry is at position $(p, p+1)$ (and $q > 0$ of course.) It is easy to see that this is exactly the grading for $\mathbf{SU}(p, q)$.

For the non-equal rank case with n even, from the procedure described in 4.5, which really becomes very simple in type A, we see that the two gradings that are defined are the one for which all imaginary roots are compact, and the one for which they are all non-compact (the imaginary root system is of type A_1^m, $m = \lfloor n/2 \rfloor$.) So there are definitely going to be two distinct real forms in this case. After we have explained how to generate all Cartans for a given group using Cayley transforms (cf. 5.6), it will be apparent that the compact grading gives rise to a group with a single conjugacy class of Cartan subgroups, which must therefore be $\mathbf{SL}(n/2, \mathbf{H})$, and the other one corresponds to the split form $\mathbf{SL}(n, \mathbf{R})$.

5 Cayley transforms

5.1. We now come to the essential operation of Cayley tranform. This will enable us to move from one conjugacy class of Cartan subgroups to another, and in this manner bootstrap things from the fundamental Cartan. We will approach this operation in a purely combinatorial manner—no attempt shall be made to relate this to the usual definition of Cayley transform for which we refer to Vogan [6] (although this will of course be essential if we want to link the combinatorial picture to actual representations.)

5.2. Let x be a strong involution, and $\tau \in \mathcal{I}$ the corresponding root datum involution. Recall from 4.4 the grading of the imaginary root system Φ_i induced by x. Let α be an imaginary non-compact root, and X_α a corresponding root vector. Denote φ_α the homomorphism $\mathbf{SL}(2) \to G$ taking $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ to $\alpha^\vee$ and $\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ to X_α, and let $\sigma_\alpha = \varphi_\alpha \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$. Then σ_α belongs to N, and is a representative of the root reflection s_α. We will see later that it is possible to normalize the choice of σ_α in terms of our chosen pinning, but in any case it is clear that the various choices of σ_α are conjugate under the one-parameter subgroup T_α of T corresponding to $\alpha^\vee$. Denote $m_\alpha = \sigma_\alpha^2$, an element of order two in T.

5.3 Definition. — Let the notation be as in 5.2. The *Cayley transform* of x through α is the element $c^\alpha(x) = \sigma_\alpha.x \in N$. This is well-defined and independent of the choice of σ_α at the level of the one-sided parameter space $\mathcal{X}$, on the set $\mathcal{X}_\tau^\alpha$ of elements in $\mathcal{X}_\tau$ for which α is noncompact.

5.4. To see that Definition 5.3 makes sense, the first thing to check is that $\sigma_\alpha.x$ is again a strong involution. In fact, we will even show that $c^\alpha(x)$ is conjugate to x in G. Indeed, it is clear that θ_x normalizes $G_\alpha = \varphi_\alpha(\mathbf{SL}(2))$, and that its action on that subgroup is conjugation by $t_\alpha = \varphi_\alpha \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}$. Therefore we may write $x = t_\alpha.x'$ where x' commutes with G_α. Then to prove that $\sigma_\alpha.x$ and x are conjugate in G, it suffices to check that t_α and $\sigma_\alpha.t_\alpha$ are conjugate in G_α, which is an easy exercise in $\mathbf{SL}(2)$. So c^α preserves strong real forms for G.

The root datum involution induced by $\sigma_\alpha.x$ is of course $s_\alpha.\tau$. We have seen that the conjugation action of T on the strong involutions lying over τ is by multiplication with an element in $T_\circ^{-\tau}$. But clearly $T_\circ^{-\tau}$ commutes with G_α (as α is imaginary), and $T_\circ^{-\tau} \subset T_\circ^{-s_\alpha.\tau}$; so the Cayley transform goes over to $\mathcal{X}$; and since α is real for $s_\alpha.\tau$, we have $T_\alpha \subset T_\circ^{-s_\alpha.\tau}$ as well, so that the map induced at the level of $\mathcal{X}$ is indeed independent of the choice of σ_α.

5.5. Let us now show that the Cayley transform c^α is *surjective* from $\mathcal{X}_\tau^\alpha$ to $\mathcal{X}_{s_\alpha.\tau}$, and at most two-to-one. Keep the notation of 5.4. For the surjectivity, let y be a strong involution in N lying over $s_\alpha.\tau$. Multiplying by an appropiate element of $T_\alpha \subset T_\circ^{-s_\alpha.\tau}$, we may assume that $\theta_y := \operatorname{int}(y)$ takes X_α to $X_{-\alpha} = \operatorname{int}(\sigma_\alpha)(X_\alpha)$. In other words, we may write $y = \sigma_\alpha.t_\alpha.x'$ where x' commutes with G_α, and is then clear that $x = t_\alpha.x'$ is a strong involution lying over τ whose Cayley transform is y.

For the second statement, suppose that $x_1 = t_\alpha.x_1'$ and $x_2 = t_\alpha.x_2'$ are strong involutions in N lying over τ, such that $y_1 = \sigma_\alpha.x_1'$ and $y_2 = \sigma_\alpha.x_2'$ are congruent modulo $T_\circ^{-s_\alpha.\tau}$. In particular, this means that $y_1^2 = z = y_2^2$, and therefore also $z = x_1^2 = x_2^2$. So x_1 and x_2 differ by an element $t \in T^{-\tau}$. But because of our hypothesis on the y's, t must belong to $T_\circ^{-s_\alpha.\tau} = T_\circ^{-\tau}.T_\alpha$. From the condition $t\tau(t) = 1$ it follows that $t \in T_\circ^{-\tau}T_\alpha(2)$, and there are indeed at most two possibilities for t modulo $T_\circ^{-\tau}$, depending on whether $T_\alpha(2)$ is contained in $T_\circ^{-\tau}$ or not. We note that this condition depends solely on α and τ; therefore c^α is either two-to-one on all of $\mathcal{X}_\tau^\alpha$, or one-to-one on all of it.

Also, one has either $\mathcal{X}_\tau^\alpha = \mathcal{X}_\tau$, or it is "of index two" in $\mathcal{X}_\tau$ (for instance, if $\mathcal{X}_\tau$ is finite, this means that $|\mathcal{X}_\tau^\alpha|$ is either equal to $|\mathcal{X}_\tau|$, or equal to one half of it.) The first case happens when the character α is trivial on the group T_τ introduced in 2.4; the second when it is non-trivial (notice that α takes values ± 1 on T_τ).

When $\mathcal{X}_\tau$ is finite, the conclusion is that the cardinality of $\mathcal{X}_{s_\alpha.\tau}$ is either equal to that of $\mathcal{X}_\tau$, or drops by a factor of two or four; and the "typical" situation is a drop by a factor of four. Of course this is not sustainable in general as we move through a sequence of Cayley transforms towards the quasisplit root datum involution (particularly when the split form is equal rank); the rule of thumb is that there is about "half" the room required, and therefore it might be expected that the $\mathcal{X}_\tau$ tend to stabilize (and often become singletons) from the point where τ is about "half split". This turns out to be true in many examples, as may be checked using the `cartan` command in the Atlas software package.

5.6. We make some further elementary remarks about Cayley transforms. First of all, we have a simple compatibility between Cayley transforms and the conjugation action (also called cross-action) of W: this is simply $w \times c^\alpha(x) = c^{w.\alpha}(x)$ for all $\xi \in \mathcal{X}_\tau$, where we use a $\times$ to denote the action of W induced by the conjugation action of N on the set of strong involutions in

N. This shows that Cayley transforms through arbitrary imaginary roots are just cross-conjugates of Cayley transforms for imaginary roots in Pi.

The Cayley transform is defined at the level of root datum involutions by $c^\alpha(\tau) = s_\alpha.\tau$, whenever α is imaginary for τ. It is well-known (see for instance[2]) that any root datum involution can be obtained from δ through a sequence of conjugations and Cayley transforms for simple roots, or equivalently, through a sequence of Cayley transforms for arbitrary roots (the corresponding roots will then necessarily be pairwise orthogonal.) In particular, noting that for any imaginary root there will always be strong real forms that make it non-compact *(reference??)*, this proves the surjectivity of the canonical map from $\mathcal{X}$ to $\mathcal{I}$.

5.7. Recall from 4.4 the grading of the imaginary root system defined by any strong involution x. If α is a noncompact imaginary root, one would like to describe the grading defined by $c^\alpha(x)$. This is done in Vogan [6], Definition 5.2 and Lemma 10.9 :

5.8 Proposition. — *Let $x \in N.\delta$ be a strong involution, and let α be a noncompact imaginary root for x. Then the imaginary roots for $c^\alpha(x)$ are the imaginary roots for x orthogonal to α, and we have*

$$\mathrm{gr}_{\sigma_\alpha.x}(\beta) = \begin{cases} \mathrm{gr}_x(\beta) & \textit{if } \alpha+\beta \textit{ is not a root} \\ \mathrm{gr}_x(\beta)+1 \mod 2 & \textit{if } \alpha+\beta \textit{ is a root} \end{cases}$$

(i.e., the grading is preserved if $\alpha+\beta$ is not a root, and reversed otherwise.)

Proof. — Choose root vectors X_α and X_β in $\mathrm{Lie}(G)$. Note that $\alpha+\beta$ is not a root if and only if X_α and X_β commute, which is equivalent to the fact that the corresponding three-dimensional groups G_α and G_β commute. But then it is clear that σ_α acts trivially on X_β, hence the first case.

If $\alpha+\beta$ is a root, the X_β is the zero-weight space of a three-dimensional representation of G_α (because no string of roots can be longer than four.) So the action of σ_α on X_β is the same as its action on the Cartan in $\mathrm{Lie}(G_\alpha)$, *i.e.*, by -1, and the grading of β is reversed.

6 Cocycles and gradings

6.1. We have seen that it is important to understand the orbits of W acting on the one-sided parameter space $\mathcal{X}$. This immediately reduces to the understanding of the W^τ orbits on the fiber $\mathcal{X}_\tau$, when τ runs through a set of representatives of Weyl group orbits in the set $\mathcal{I}$ of root datum involutions for our given inner class.

So fix an element $\tau \in \mathcal{I}$. The stabilizer W^τ of τ in W is described in [6], as follows. Let $\Phi = \Phi(G, T)$ be the root system of G with respect to T; let Φ_r and Φ_i be the subsystems of real and imaginary roots w.r.t. τ, respectively. Let Φ_c be the set of roots that are orthogonal both to ρ_r and ρ_i, where as usual ρ denotes the half-sum of positive roots. Then Φ_c is a complex root system: it splits up into the direct sum of two root systems interchanged by τ. Now we have ([6], Proposition 3.12) :

$$W^\tau = W_c \ltimes (W_i \times W_r)$$

where W_i and W_r are the Weyl groups of Φ_i and Φ_r respectively, and W_c is the "diagonal subgroup" of the Weyl group of Φ_c: it is generated by the elements $s_\alpha.s_{\tau(\alpha)}$, where α runs through the set of simple roots of one of the two factors of Φ_c (note that these generators have a -1 eigenspace of dimension 2, so they are not reflections; hence W^τ is not a Coxeter group in general.)

6.2 Proposition. — *The actions of W_c and W_r on the fiber $\mathcal{X}_\tau$ are trivial.*

Proof. — (as explained to me by David Vogan). Recall the notation from 5.4. Let α be a real root, and G_α be the corresponding three-dimensional subgroup. Let x be a strong involution in N lying over τ. Then θ_x induces the split involution on G_α, and we may in fact assume that τ_x induces the same involution as $\mathrm{int}(\sigma_\alpha)$ (which is T_α-conjugate to the $\sigma_\alpha.t_\alpha$ we used in 5.5). Then we may write $x = \sigma_\alpha.x'$, where x' commutes with G_α, and it is clear that σ_α commutes with x, which proves that the action of s_α on $\mathcal{X}_\tau$ is indeed trivial.

When α is in Φ_c, we may reason as follows. Pick σ_α as before, and choose $\sigma_{\tau(\alpha)} = x.\sigma_\alpha.x^{-1}$. Note that $\alpha + \tau(\alpha)$ is not a root, because it would have to be imaginary and orthogonal to ρ_i; therefore the two groups G_α and $G_{\tau(\alpha)}$ commute, and in particular σ_α and $\sigma_{\tau(\alpha)}$ commute. But clearly $x.\sigma_\alpha.\sigma_{\tau(\alpha)}.x^{-1} = \sigma_{\tau(\alpha)}.\sigma_\alpha$, so the generators of W_c also act trivially on $\mathcal{X}_\tau$, and we are done.

6.3. It follows from 6.2 that the orbits of W^τ on $\mathcal{X}_\tau$ are really just the W_i-orbits. Moreover, the action of W preserves squares, and so we wave an action of W_i on each $\mathcal{X}_\tau(z)$; recall from 2.6 that this set has the structure of an affine space over the two-element field $\mathbf{F}_2$, preserved by W. The corresponding linear action is the action of W on $T^{-\tau}/T_\circ^{-\tau}$, which may be readily computed in terms of lattices.

There is an obvious action of the center on each $\mathcal{X}_\tau$ by left multiplication. Since this obviously commutes with G-conjugation, it will preserve the action of W_i (but multiplication by z' takes $\mathcal{X}_\tau(z)$ to $\mathcal{X}_\tau(z'\delta(z')z)$). So the orbit picture in $\mathcal{X}_\tau(z)$ for two values of z that are congruent modulo the image Z_δ of Z under the homomorphism $z' \to z'\delta(z')$ will be identical. To understand all possible situations, it will be enough to have z run through a set of representatives modulo Z_δ.

6.4 Proposition. — *Each Z-conjugacy class in $\mathcal{X}_\tau$ has a representative in $\mathcal{X}_\tau(z)$ where z belongs to the center Z_1 of the derived group of G; one may even ask that the order of z be a power of two.*

Proof. — We may write $T = T_1.\mathrm{Rad}(G)$, where T_1 is the identity component of the intersection of T with the derived group, and $\mathrm{Rad}(G)$ is the identity component of Z. Recall that we denoted T_τ the group of elements $t \in T$ such that $(1+\tau)(t) \in Z$. But then it is clear that if we write $t = t_1.z$, with $t_1 \in T_1$ and $z \in \mathrm{Rad}(G)$, $(1+\tau)(t)$ is congruent to $(1+\tau)(t_1)$ modulo Z_δ, and of course $(1+\tau)(t_1) \in Z_1^\delta$. Let p be a prime divisor of the the order of Z_1, and let $Z_1(p)$ be the corresponding Sylow subgroup. If p is odd, we have that $Z_1(p)^\delta = (Z_1(p)^d)^2 \subset Z_\delta$, so we may even arrange that $(1+\tau)(t_1) \in Z_1(2)$.

6.5 Corollary. — *There are only finitely many orbit pictures for the various $\mathcal{X}_\tau(z) \subset \mathcal{X}(\tau)$.*

6.6 Example. — Consider again the example of $\mathbf{SL}(n)$. For the equal rank case, it is always true that all elements of $Z(G)$ are attained as x^2 (because this is obvious for the fundamental torus). And $z \to z^2$ is surjective from Z to Z if and only if n is odd; otherwise there are two cosets for $\mathbf{Z}_\delta$. So for the fundamental fiber, we should expect to have either one orbit picture repeated n times, or two orbit pictures repeated $n/2$ times each.

When n is odd, this remains true for all other values of τ: there will always be n identical orbit pictures lying over each τ, isomorphic to the corresponding picture for the adjoint group. When n is even, things are more subtle. When $m = n/2$ is odd, one of the two pictures contains $\mathbf{SU}(2m)$, $\mathbf{SU}(2m-2,2)$, …, $\mathbf{SU}(m+1,m-1)$, each twice, and the other $\mathbf{SU}(2m-1,1)$, …, $\mathbf{SU}(m,m)$, each twice except the last one. The first set of forms corresponds to values of x^2 that are congruent to zero modulo Z^τ, and the other ones to those that are congruent to one. Then one sees that all values in Z are reached on $\mathcal{X}_\tau$ except when τ is maximally split, where only $\mathbf{SU}(m,m)$ survives, and therefore only the values in Z that are *not* in Z_δ are reached there (this is apparent already in the elementary case of $\mathbf{SL}(2)$.)

Also, it is not true in this case that one gets the picture for the adjoint group in each $\mathcal{X}_\tau(z)$. The map from $\mathcal{X}_\tau(z)$ to the corresponding $\mathcal{X}_{\tau,\mathrm{adjoint}}$ (only one possible value of z here!) is always two-to-one except for the most split case; one of the two families maps to one-half of $\mathcal{X}_{\tau,\mathrm{adjoint}}$, the other one to the other half.

One can make a similar analysis when m is even. The main difference is that here the quasisplit forms are now in the same family as the compact ones, and therefore correspond to the class of elements with square in Z_δ, instead of the opposite.

For the non-equal rank case, things are simpler, because δ now acts on Z through inversion, so the group Z^δ is either trivial (when n is odd), or has two elements, ± 1 (when n is even.) And Z_δ is trivial in all cases. It is not hard *(I hope!)* to see that in the even case, both values in Z^δ are reached (the value -1 corresponding to $\mathbf{SL}(n, \mathbf{R})$.)

6.7 Example. — To see a case where Z is not finite, consider the case of $\mathbf{GL}(2)$. The main difference with the case of $\mathbf{SL}(2)$ is that the fundamental torus is now *complex*, which means that $T^\delta(2) = T^{-\delta}(2)$, and therefore the two sets $\mathcal{X}_\delta(1)$ and $\mathcal{X}_\delta(-1)$ are both singletons. Also, of course, $\mathcal{X}_s$ is a singleton, for the non-trivial element s of the Weyl group.

In particular, there are only two strong real forms of $\mathbf{GL}(2)$ for this inner class : the group $\mathbf{SU}(2).\mathbf{R}^\times$ generated by $\mathbf{SU}(2)$ and real dilations in $\mathbf{GL}(2, \mathbf{C})$, and $\mathbf{GL}(2, \mathbf{R})$.

6.8. It turns out that in general there is a very nice description of the action of W_i on $\mathcal{X}_\tau$ in terms of the grading of the imaginary roots associated to a strong involution $x \in N.\delta$. Denote gr_x this grading. Clearly it is unchanged under conjugation by T; therefore we may speak of gr_ξ for any $\xi \in \mathcal{X}_\tau$. Recall that for each $z \in Z$ such that $\mathcal{X}_\tau(z)$ is non-empty, there is a simply transitive action of the component group $T^{-\tau}/T_\circ^{-\tau}$ on $\mathcal{X}_\tau$, which we will denote additively.

6.9 Proposition. — *Let $\xi \in \mathcal{X}_\tau(z)$, and let α be an imaginary root for τ. Then the action of s_α on ξ is given by*

$$s_\alpha.\xi = \xi + \mathrm{gr}_\xi(\alpha) m_\alpha$$

where the action of m_α is through its image in $T^{-\tau}/T_\circ^{-\tau}$. In particular, the Weyl group of the root system $\Phi_i(0)$ consisting of the roots that are compact for ξ is contained in the stabilizer of ξ.

Proof. — Recall the notation from 5.4. Let $x \in N.\delta$ be a representative of x. If α is compact for x, then θ_x is trivial on G_α, so x and σ_α commute, and the action of s_α on ξ is indeed trivial. If α is non-compact, we have seen in 5.5 that we may write $x = t_\alpha.x'$, where x' commutes with G_α. But it is clear that $\sigma_\alpha.t_\alpha.\sigma_\alpha^{-1} = t_\alpha^{-1} = m_\alpha.t_\alpha$, so it follows that $s_\alpha.\xi = \xi + m_\alpha$, and we are done.

6.10. Proposition 6.9 makes it possible to describe each $\mathcal{X}_\tau(z)$ as a W_i-set, just from the knowledge of the grading associated to one of its elements. The map $W_i \to T^{-\tau}/T_\circ^{-\tau}$ defined by $w \to w.\xi - \xi$ satisfies an obvious cocycle condition, and can be readily computed from the actions of the generators.

Moreover it is known (cf. [6] Proposition 6.12 ??) that for every τ, the grading of the imaginary root system where all simple roots are noncompact is always allowed (and corresponds to the quasisplit real form.) If one is willing to allow translations by the full group $T_\tau/T_\circ^{-\tau}$ from Proposition 2.4, this will yield a canonical description of $\mathcal{X}_\tau$ as a W_i-space for every given τ. (If one is willing to allow only translations by $T^{-\tau}$, the situation is more delicate, as it is not always true that every $\mathcal{X}_\tau(z)$ contains a strong real form which is quasisplit; this is apparent already in the example of $\mathbf{SL}(n)$ in Example 6.6.)

The `cartan` command of the Atlas software package prints out the orbit pictures for a set of representatives τ of W-orbits in $\mathcal{I}$; for each τ, one gets a classification of the strong real forms of G in the current inner class for which this Cartan is defined (or more precisely, of those strong real forms for which x^2 belongs to the center of the derived group.) In these printouts, real form #0 is always a quasisplit one, with a corresponding element of $\mathcal{X}$ labelled as #0 as well.

7 The Tits group

7.1. We will call *Tits group*, and denote $\tilde{W}$, the subgroup of G generated by the elements σ_α introduced in 5.2. This group has been studied by Jacques Tits in [5], under the name of *extended Coxeter group*. It will play an essential role in the actual construction of the parameter set.

The following theorem contains the properties of the Tits group that we will need:

7.2 Theorem. — (Tits [5]) (a)*The kernel of the natural surjection $\tilde{W} \to W$ is the subgroup of $T(2)$ generated by the elements m_α (in particular, it is an elementary abelian 2-group.)* (b) *Let the σ_α, $\alpha \in \Pi$, be defined using the* *pinning $\mathcal{P}$ chosen in 1.1. Then the σ_α satisfy the braid relations, so that we get a canonical lifting of W as a subset of $\tilde{W}$ by taking a reduced expression for $w \in W$ and denoting $\tilde{w}$ the corresponding product of the σ_α.*

7.3. The upshot is that any element of $\tilde{W}$ may be canonically written as $n = t\sigma_{j_1} \dots \sigma_{j_r}$, where $s_{j_1} \dots s_{j_r}$ is a reduced expression of the image of n in W, and t is a product of m_α's (say for α simple, in which case the product is even unique.)

It is not difficult to describe the group generated by the m_α. This is isomorphic to one half the coroot lattice of G modulo the cocharacter lattice of T (or equivalently, the coroot lattice modulo twice the cocharacter lattice.) From this, we immediately get the conjugation action action of W on it.

It can happen that some of the m_α are trivial, but this is rather rare. It is enough to deal with the case where the root system is irreducible. Then we note that the triviality or non-triviality of m_α is constant along W-conjugacy classes of roots. So certainly it is enough to look at simple roots. Also, if there is a representation of the corresponding three-dimensional subgroup G_α where m_α acts as -1, it can of course not be trivial. In particular, if there is a root β for which $< \beta, \alpha^\vee >$ is odd, m_α is non-trivial. So the only cases where m_α can be trivial is for adjoint A_1, and for the short root in type B_2, also in the adjoint case (of course when G is simply connected, there will always be weights λ with $< \lambda, \alpha^\vee >= 1$.) Now it is not hard to check that in those two cases m_α is indeed trivial.

7.4. As we will explain in some more detail below *(reference??)*, the main ingredients for the Kazhdan-Lusztig for real reductive groups are the cross actions and the Cayley transforms (on the two-sided parameter space to be defined below.) At the level of one-sided parameters, these may in fact be defined for each strong real form (i.e., for the image of each G-conjugacy class.) A slightly larger, but much more manageable, setting is to define them on $\mathcal{X}(z)$ for a given central element z (for which $\mathcal{X}(z)$ is non-empty, of course.)

So the problem is to construct $\mathcal{X}(z)$ algorithmically. But this follows very naturally from what we have done so far. The fundamental fiber $\mathcal{X}_\delta(z)$ corresponds to a certain $T^{-\delta}/T_\circ^{-\delta}$-orbit in $\mathcal{X}_\delta$; we assume that this has been handed to us. Denote for simplicity δ_z an element in $T.\delta$ such that $\delta_z^2 = z$. Then everything we do takes place in $\tilde{W}.\delta_z$.

A rough sketch of the algorithm is as follows:

(a) maintain a first-in-first-out list of elements τ in $\mathcal{I}$, together with a strong involution x_τ in $\tilde{W}.\delta_z$ representing τ, and a subset of $T_\circ^\tau(2)$

representing a basis of $T^{-\tau}/T_\circ^{-\tau} \simeq T_\circ^\tau(2)/(T_\circ^\tau(2) \cap T_\circ^{-\tau}(2))$; initialize this list with $\tau = \delta$, $x_\delta = \delta_z$, and a basis of $T_\circ^\delta(2)/(T_\circ^\delta(2) \cap T_\circ^{-\delta}(2))$. Also keep in memory the elements τ that have been put on the list.

(b) while the list is non-empty: take the first element τ off the list. Try conjugating τ with the various σ_α, α simple. If we find a new τ, put it on the list, conjugating the data for τ by σ_α. Next, look at the Cayley transforms c^α, still for α simple (and imaginary, of course), and for which there are elements $t.x_\tau$ for which α is noncompact. Again, see if $s_\alpha.\tau$ is new. If yes, put it on the list, and take an $t.x_\tau$ for which α is non-compact (this is either all, or half of the $t.x_\tau$); take $x_{s_\alpha.\tau} = \sigma_\alpha.t.x_\tau$ in the data for the new group. Also, using the fact that $T^{-s_\alpha.\tau}/T_\circ^{-s_\alpha.\tau}$ is the quotient by the two-element subgroup generated by m_α of the kernel in $T^{-\tau}/T_\circ^{-\tau}$ of the root α, compute a basis for $T^{-\tau}/T_\circ^{-\tau}$.

(c) whenever an element τ is put on the list, put the corresponding set of parameters into a store, that will eventually contain an entry for each element of $\mathcal{X}(z)$.

In practice, the goal would be to obtain a well-defined numbering of the parameters, and in terms of this numbering, to produce tables representing the cross-actions of the simple reflections, and the direct and inverse Cayley transforms in terms of this numbering. One would probably also want to keep data such as the corresponding root datum involution, the corresponding grading, ...

8 The two-sided parameter space and the classification of representations

8.1. At long last we are now in a position to describe the parameter space for actual representations. We note that our one-sided parameter space $\mathcal{X}$ has been constructed entirely in terms of the root datum $(X, R, X^\vee, R^\vee)$ (and our chosen inner class γ.) Therefore we can similarly construct the one-sided parameter space $\mathcal{Y}$ for the dual root datum and the dual inner class $\gamma^\vee$.

Now we form the restricted product:

$$\mathcal{Z} := \mathcal{X} \times^{\mathcal{I}} \mathcal{Y}$$

as follows: $\mathcal{Z}$ is the set of pairs $(\xi, \eta) \in \mathcal{X} \times \mathcal{Y}$ such that the root datum involutions τ, τ^d induced by ξ and η satisfy $\tau^d = -\tau^\vee$. Our definition of the dual inner class ensures that this makes sense. The set $\mathcal{Z}$ is called the *two-sided parameter space* for G. Just as we did for one-sided parameter spaces, we may introduce the fiber $\mathcal{Z}_\tau$ for $\tau \in \mathcal{I}$, and $\mathcal{Z}(z, z^d)$ for $z \in Z(G)$, $z^d \in Z(G^\vee)$.

8.2. The two-sided parameter space parametrizes representations of strong real forms of G, with regular integral character, up to translation by the character lattice of T. Roughly speaking, here is how it goes (here we use the actual dual group $G^\vee$.) *(this requires more work!)*

Take a pair of strong involutions (x, y) in $N.\delta \times N^\vee.\delta^\vee$, lying over a certain root datum involution τ. We may identify $\mathfrak{t}^\vee = \mathrm{Lie}(T^\vee)$ with the dual of $\mathfrak{t} = \mathrm{Lie}(T)$. Then the center of $G^\vee$ may be canonically identified with $P/X^*(T)$, where P is the group of integral weights (which has a vector component when G is not semisimple.) Therefore, the square of y allows us to pick an integral infinitesimal character up to translation, which we may arrange to be non-singular. Now the datum of x allows us to recover a strong real form of G, together with a θ_x-stable Cartan H and a chamber in $\mathfrak{h}^*$, up to $W(K, H)$-conjugacy. Together with the chosen infinitesimal character, this gives us an element λ of $\mathfrak{h}^*$. Now the last remaining question to obtain a character of H is to extend the character defined by λ on $H_\circ^{\theta_x}$ to all of H^{θ_x}. *(here some shifting should be done that I'm not entirely clear about).* After translation, back to T, this corresponds to finding an element of the dual group of $T^\tau(2)$, whose restriction to $T^\tau(2) \cap T_\circ^{-\tau}(2)$ is given (by the square of y.) This should correspond exactly to the choice of y with the given square.

8.3. So the conclusion is as follows: the two-sided parameter space is in $(1, 1)$ correspondence with the disjoint union over all possible strong real forms of G, and all possible translation classes of regular integral infinitesimal characters, of the set of irreducible $(\mathfrak{g}, K_x)$-modules for that real form and with that infinitesimal character.

As we will see in 9.1 below, the $G^\vee$-conjugation classes in $\mathcal{Y}$ correspond to blocks of representations.

8.4 Example. — Let us do the very simple example of $\mathbf{SL}(2)$. Here there is a single inner class, with two real forms. As we have seen in Example 4.2, the fundamental fiber has four elements, with three orbits corresponding to three strong real forms. The other fiber has a single element.

The dual group is $\mathbf{PSL}(2)$, which is also the adjoint group. So the one-sided parameter space for the dual group can be obtained by passing to the quotient: there are two elements in the fundamental fiber, corresponding to the two real forms, and one in the other fiber. Going over to $-\theta^\vee$ amounts to reversing the dual picture, and then doing the restricted product yields a space $\mathcal{Z}$ with $4.1 + 1.2 = 6$ elements. Four of these six elements correspond to the split real form $\mathbf{SL}(2, \mathbf{R})$ (the two discrete series, lying over the fundamental fiber, and the finite dimensional representation, and the non-spherical (irreducible) principal series, lying over the other fiber.) The other two correspond to the two compact strong real forms (a single representation each.) In this case, there is only one translation class of regular integral infinitesimal characters, so the picture is consistent with the interpretation in 8.3.

8.5 Example. — Consider now the example of $\mathbf{GL}(2)$, for the inner class where the radical is split (*i.e.*, the inner class containing $\mathbf{GL}(2, \mathbf{R})$.) We have described the one-sided parameter space $\mathcal{X}$ in 6.7. To describe the other side, one must take care that the dual group is still $\mathbf{GL}(2)$, but the dual inner class is now the one where the radical is *compact*, *i.e.*, the inner class containing $U(2)$. Here the fundamental torus is compact, and each of the parameter sets $\mathcal{Y}_s(1)$, $\mathcal{Y}_s(-1)$ contains four elements. Since -1 is the square of a central element in the dual group, both orbit pictures are in fact the same: each contains two compact strong real forms and one quasiplit one. Since the most split torus of the dual group is complex, $\mathcal{Y}_\delta(1)$ and $\mathcal{Y}_\delta(-1)$ are both singletons. So if we fix an infinitesimal character, we get five irreducible representations of $\mathbf{GL}(2, \mathbf{R})$, one discrete series, two finite-dimensional ones, and two irreducible principal series : essentially the combination of the pictures for the two possible infinitesimal characters for $\mathbf{PSL}(2, \mathbf{R})$.

8.6. There is a canonical identification of the Weyl groups of (G, T) and $(G^\vee, T^\vee)$. Therefore we may consider the diagonal action of W on $\mathcal{X} \times \mathcal{Y}$; practically by definition, this will preserve the restricted product, and therefore restrict to an action on $\mathcal{Z}$. This action is usually called the *cross-action* on two-sided parameters. *(here what is missing is to check that this cross-action translates to the usual cross-action on regular characters of θ-stable tori.)*

Note that the orbits of this action on each fiber $\mathcal{Z}_\tau$ will be the same as the orbits of the action of $W^\tau \times W^\tau$, because as was seen in Proposition 6.2, the W^τ-action really amounts to the W_i-action on $\mathcal{X}_\tau$, and dually to the W_r-action on $\mathcal{Y}_\tau$ (the W_r for $\mathcal{X}$ becomes the W_i for $\mathcal{Y}$.) So the orbits in $\mathcal{Z}$ are just the products of orbits in $\mathcal{X}$ by orbits in $\mathcal{Y}$.

To define the Cayley transform c^α on $\mathcal{Z}$, just set

$$c^\alpha(\xi, \eta) = (c^\alpha(\xi), c_\alpha(\eta))$$

for $\zeta = (\xi, \eta) \in \mathcal{Z}$. Here c_α is the inverse Cayley transform. This is not a function in general; it will be a one-to-two correspondence for those cases where c^α is two-to-one.

8.7 Example. — Consider the case of $\mathbf{PSL}(2)$. Here the picture is the dual picture from the one in Example 8.4. There are still six parameters, two in the fundamental fiber and four in the other, but only two strong real forms : $\mathbf{PSL}(2, \mathbf{R})$ which is the full group of real points of $\mathbf{PSL}(2)$, hence non-connected, and $\mathbf{PSU}(2)$. Now five of the six parameters correspond to $\mathbf{PSL}(2, \mathbf{R})$: one in the fundamental fiber, the unique discrete series, and four in the non-fundamental fiber, the two finite-dimensional ones, and two irreducible principal series. There are two translation classes of infinitesimal characters, one containing the discrete series and the two finite-dimensional ones, the other containing the two irreducible principal series.

Now the Cayley transform of the discrete series is going to be the pair (ζ_+, ζ_-) consisting of the two finite-dimensional representations.

8.8 Example. — Another very simple example is that of a torus. Let $G = T$ be a split torus (here the choice of inner class actually fully determines the real form.) Then δ is inversion. Here $Z^\delta = T(2)$ has 2^r elements, where r is the rank of the torus, but the image $\mathbf{Z}_\delta$ of T under $t \to t\delta(t)$ is the identity, so there is in fact only one possible value for x^2, *viz.* $x^2 = 1$, and all the corresponding x'es are T-conjugate. Hence $\mathcal{X}$ is a singleton.

On the dual side, $\delta^\vee$ is the identity, so all squares become possible in the dual torus. Moreover, each $\mathcal{Y}(z^d)$ has 2^r elements. Clearly one sees how this corresponds to choosing the derivative of a character up to translation by (the derivatives of) the algebraic characters, and then extending to the component group of T^δ.

In the case of a compact torus, the situation is reversed : there are infinitely many strong real forms, but just a single infinitesimal character up to translation : this is clear, as representations of a compact torus correspond to the elements of $X^*(T)$.

8.9. *(this needs to be expanded and made more precise!)* The Cayley transforms and cross-actions are all we need to define the action of the Hecke algebra of W on the free $\mathbf{Z}[q^{1/2}, q_{-1/2}]$-module generated by $\mathcal{Z}$. Then, we can define the length function, the order relation, and descent sets, which will give us all the necessary ingredients to set up the Kazhdan-Lusztig algorithm.

9 Duality and blocks

9.1. Both one-sided parameter spaces $\mathcal{X}$ and $\mathcal{Y}$ are naturally partitioned according to the strong real forms of G and $G^\vee$ respectively. So $\mathcal{Z}$ is partitioned as well, according to pairs of strong real forms. These classes in $\mathcal{Z}$ are called *blocks* of representations. As we have seen, cross-actions and Cayley transforms preserve blocks; hence a block is the natural (and also the minimal) setting for Kazhdan-Lusztig computations.

In other words: fix a strong real form of G; one could also fix a regular integral infinitesimal character (up to translation), *i.e.*, an element z^d of the center of $G^\vee$. Then the irreducible representations of the given real form, with the given infinitesimal character, are naturally partitioned into blocks, the partition being indexed by the strong real forms of $G^\vee$ corresponding to $\mathcal{Y}(z^d)$.

What is remarkable is that this partition is the same as the one obtained by the natural definition of block, which corresponds to the smallest equivalence relation for which two representations are equivalent when there is a non-trivial Ext between them *(reference ??).*

9.2. Even more beautiful: duality. It is clear from the definitions that interchanging G and $G^\vee$, with their corresponding inner classes, amounts to interchanging the roles of $\mathcal{X}$ and $\mathcal{Y}$ (and, in practice, reading the picture backwards, as the fundamental fiber for $\mathcal{Y}$ is at the opposite end from the fundamental fiber for $\mathcal{X}$.) This has appeared in the comparison of Examples 8.4 and 8.7.

Of course this duality preserves blocks, and can be described one block at a time. But one can not express it for the full representation theory of one real form at a time: to describe the representations of one real form of G, one needs *all* strong real forms of $G^\vee$, and conversely.

References

[1] Jeffrey D. Adams. Lifting of characters. In *Harmonic analysis on reductive groups (Brunswick, ME, 1989)*, volume 101 of *Progr. Math.*, pages 1–50, Boston, Massachussets, 1991. Birkhäuser Boston.

[2] Axel Hultman. The combinatorics of twisted involutions in Coxeter groups. *Trans. Amer. Math. Soc.*, (to appear), 2005.

[3] Roger W. Richardson and Tonny A. Springer. The Bruhat Order on Symmetric Varieties. *GeometriæDedicata*, **35**:389–436, 1990.

[4] Tonny A. Springer. Some results on algebraic groups with involutions. volume 6 of *Advanced Studies in Pure Math.*, pages 525–543. Kinokuniya/North Holland, 1985.

[5] Jacques Tits. Normalisateurs de tores I. Groupes de Coxeter étendus. *J. Algebra*, **4**:96–116, 1966.

[6] David A. Vogan, Jr. Irreducible characters of semisimple Lie groups IV. Character-multiplicity duality. *Duke Math. J.*, **49**:943–1073, 1982.