Closure Relations of K orbits on G/B

1. Introduction

Let G be a complex, connected, reductive algebraic group defined over $\mathbb{R}$ and let $G_{\mathbb{R}}$ be the real points of G. Let $K_{\mathbb{R}}$ be a maximal compact subgroup in $G_{\mathbb{R}}$ and let K be its complexification. Choose a Cartan subgroup H of G and a Borel subgroup B of G containing H. Let $\mathfrak{h} \subset \mathfrak{b} \subset \mathfrak{g}$ be the corresponding Lie algebras and let $X = G/B$ be the flag variety. Then G acts on X in the obvious way and hence K acts on X by restriction.

Theorem 1.1. *K acts on X with finitely many orbits.*

The orbits in G/B under the action of K are denoted $K\backslash G/B$. There is a partial order on $K\backslash G/B$, called the <u>closure order</u>, defined as follows:

$$Q_1 \preceq Q_2 \iff Q_1 \subset \overline{Q_2}$$

The Atlas software command *kgb* produces a subset of the closure order for a general group $G_{\mathbb{R}}$. Our goal here is to describe an algorithm for computing the remaining closure relations from the information currently provided by Atlas.

2. The Complex Case

Much of the behavior in the general case can be anticipated from the more familiar complex case. For now, suppose that $G_{\mathbb{R}} = G$ is a complex group regarded as a real group.

Proposition 2.1. *There is an order preserving bijection between $K\backslash G^{\mathbb{C}}/B$ and $B\backslash G/B$. Here $G^{\mathbb{C}}$ refers to the complexification of the group G.*

Thus, in the complex case, we can study $K\backslash G^{\mathbb{C}}/B$ by simply studying B double-cosets in G. To this end, we have the following familiar theorem.

Theorem 2.2. *(Bruhat Decomposition) There is a bijection:*

$$W(\mathfrak{h}, \mathfrak{g}) \longleftrightarrow B\backslash G/B$$

Here, $W(\mathfrak{h}, \mathfrak{g})$ refers to the Weyl group of G. Furthermore, if we denote the orbit in $B\backslash G/B$ corresponding to w by BwB, then we have $l(w) = dim(BwB)$, where $l(w)$ denotes the length of w.

There is a very convenient way of describing the closure order (called the Bruhat order) in $B\backslash G/B$.

Theorem 2.3. *(Bruhat Order) Let $w \epsilon W(\mathfrak{h}, \mathfrak{g})$ and suppose $w = s_{\alpha_1} s_{\alpha_2} \cdots s_{\alpha_k}$ is a reduced expression for w. Then $Bw'B \preceq BwB$ if and only if $w' = s'_{\alpha_1} \cdots s'_{\alpha_k}$ where $s'_{\alpha_1} \cdots s'_{\alpha_k}$ is an ordered subset of $s_{\alpha_1} s_{\alpha_2} \cdots s_{\alpha_k}$.*

Example Let G =SL(2,$\mathbb{C}$) and let B be the set of upper triangular matrices. Then we can identify G/B with one dimensional complex projective space, $\mathbb{P}^1$. The Weyl group is S_2 and thus we expect that there should be exactly 2 orbits in the Bruhat decomposition. Clearly, the orbit corresponding to the identity is $BeB = eB$ and is just a single point. If we denote the non-trivial element of S_2 by s_α, then the orbit $Bs_\alpha B$ is everything else. The closure order is then given by the following graph:

$$\begin{array}{c} Bs_\alpha B \\ \bullet \\ | \\ | \\ \bullet \\ \mathrm{eB} \end{array}$$

Alternatively, this example shows that B acts on the parabolic subgroup corresponding to the simple root α with exactly 2 orbits (modulo B).

In fact, the general Bruhat order is 'generated' by the above example. To see this, consider the following map:

$$\pi_\alpha : G/B \to G/\mathbb{P}_\alpha$$

where $\mathbb{P}_\alpha$ denotes the parabolic subgroup corresponding to the root α. The map π_α is the natural projection map that sends B cosets in G to their corresponding $\mathbb{P}_\alpha$ cosets. Now consider the map:

$$\pi^\alpha := \pi_\alpha^{-1} \circ \pi_\alpha$$

π^α converts B cosets to $\mathbb{P}_\alpha$ cosets modulo B. We then have to following geometric fact that generalizes the above example.

Theorem 2.4. *Suppose $l(ws_\alpha) = l(w) + 1$. Then the collection of orbits given by $\pi^\alpha(BwB)$ can be viewed as a $\mathbb{P}^1$ fiber bundle over the orbit BwB, i.e.*

$$\begin{array}{ccc} \mathbb{P}^1 & \to & \pi^\alpha(BwB) \\ & & \downarrow \\ & & BwB \end{array}$$

In the $G = \mathrm{SL}(2,\mathbb{C})$ example, the theorem is trivially true (there's a single fiber). With the above geometric picture in hand, we can now describe the closure order.

Lemma 2.5. *Suppose $l(ws_\alpha) = l(w) + 1$. Then $\pi^\alpha(BwB)$ is a union of exactly two orbits: $\pi^\alpha(BwB) = BwB \cup Bws_\alpha B$. Furthermore, $\overline{\pi^\alpha(BwB)} = \pi^\alpha(\overline{BwB})$.*

Essentially, the lemma tells us that something in the closure of the orbit $Bws_\alpha B$ lives in the π^α fiber of something in the closure of BwB. This fact (combined with the above example) leads immediately to an inductive proof of the Bruhat order.

3. The General Case

We now turn to the general case. Suppose that $Q \in K\backslash G/B$ and assume that we have $\dim(\pi^\alpha(Q)) = \dim(Q) + 1$. Then we can still view Q as the base space of a $\mathbb{P}^1$ bundle, i.e.

$$\begin{array}{ccc} \mathbb{P}^1 & \rightarrow & \pi^\alpha(Q) \\ & & \downarrow \\ & & Q \end{array}$$

Consequently, we get an analogous statement about orbit closures:

$$\overline{\pi^\alpha(Q)} = \pi^\alpha(\overline{Q})$$

So, understanding the general case comes down to understanding the possible ways that π^α(Q) can decompose as a union of K-orbits. This of course depends on the type of root that α is, according to the following:

- If α is a real or compact imaginary root, $\dim(\pi^\alpha(Q)) \neq \dim(Q) + 1$. Since it suffices to understand only those operations that increase dimension, we can ignore this case.

- Suppose now that α is complex and $\dim(\pi^\alpha(Q)) = \dim(Q) + 1$. Then the situation is the same as in the complex case and we have $\pi^\alpha(Q) = Q \cup s_\alpha Q$. Here s_αQ denotes the orbit obtained from Q by the cross-action corresponding to the root α.

- If α is noncompact imaginary, then there are two cases. The behavior of K in the first case (called type I) is illustrated by considering the example $G_{\mathbb{R}} = \mathrm{SL}(2,\mathbb{R})$. In this case, K acts on $\pi^\alpha(Q)$ with three orbits: $\pi^\alpha(Q) = Q \cup s_\alpha Q \cup c_\alpha Q$. Here c_αQ denotes the orbit obtained from Q by the Cayley transform corresponding to the root α.

- The behavior of K on the second type of noncompact imaginary root (called type II), is illustrated by the example $G_{\mathbb{R}}$ =GL(2,$\mathbb{R}$). Here the situation is a little better: $\pi^{\alpha}(Q) = Q \cup c_{\alpha}Q$. Here again $c_{\alpha}Q$ denotes the orbit obtained from Q by the Cayley transform corresponding to the root α. Note that there are no longer distinct orbits coming from the cross-action of α.

We now have everything that we need to generalize the Bruhat order from above. We start with a definition.

Definition 3.1. *The Atlas graph is the graph consisting of the following:*

- *Vertices: One vertex for each element of $K\backslash G/B$. The orbit corresponding to the vertex v will be denoted by Q_v.*
- *Edges: There is a (directed) edge from vertex v_1 to v_2 if and only if dim(Q_{v_1}) = dim(Q_{v_2}) + 1 and $Q_{v_1} \subset \pi^{\alpha}(Q_{v_2})$, for some simple root α.*

Remark 3.2. *This information is all currently provided by the Atlas software through the* kgb *command. Notice that edges in the Atlas graph have an associated simple root (denoted by α in the above definition). In what follows, this association is important and is assumed to be part of the Atlas graph.*

Remark 3.3. *The fact that the Atlas graph only contains edges connecting vertices of codimension one is justified by the following theorem.*

Theorem 3.4. *The closure order in $K\backslash G/B$ is generated by related elements whose dimensions differ by one.*

The Atlas graph is a subgraph of the closure order. Our goal can thus be restated as follows. Starting with the Atlas graph, we would like an effective algorithm (along the lines of that implied by the Bruhat order) to compute the additional codimension one closure edges not present in the Atlas graph.

Algorithm: To find the closure of a vertex v in $K\backslash G/B$ i.e. all extra edges of codimension one:

(1) Construct a reduced expression set for v. This is a set (denoted S_v) of reduced expressions for v (i.e. paths to closed vertices in the Atlas graph). Elements contained in S_v are referred to as reduced expressions. As in the complex case, the important information contained in a reduced expression is not the edges themselves, but rather the simple roots associated to the edges. We construct S_v recursively as follows:
   - At first, S_v has a single (empty) path starting at v. If v is connected to another vertex v' via a cross-action or Cayley transform of type II, then add v' to the path in S_v. If there is more than one such edge, we can choose any one we like.
   - If there are only edges corresponding to Cayley transforms of type I, we do the following. Suppose that v is connected to v' and v'' by a pair of such edges. Then we want S_v to contain two paths, one connecting v to v' and one connecting v to v''. In other words, our path to a closed orbit must 'split' here and we need to keep track of both paths. Again, if there are multiple pairs of edges of this type, then we can choose any pair that we wish.
   - Repeat this process for each path in S_v until all paths end at a closed orbit.

(2) Now that we have constructed S_v, do the following. For each reduced expression in S_v, try deleting a single element and see if this gives a reduced expression of an orbit (say v') of codimension one not already connected to v in the Atlas graph. If so, add a new edge from v to v'. Note that when comparing reduced expressions, we compare them as ordered lists of simple roots.

The above algorithm is based on the following theorem.

Theorem 3.5. *Suppose $Q_v \epsilon K \backslash G/B$ and let S_v be its reduced expression set. If $w = \alpha_1\alpha_2 \cdots \alpha_k$ is a reduced expression in S_v, then $w' = {\alpha_1}' \cdots {\alpha_k}'$ is a subexpression of w if ${\alpha_1}' \cdots {\alpha_k}'$ is an ordered subset of $\alpha_1 \cdots \alpha_k$. Then an orbit $Q_{v'} \preceq Q_v$ if and only if there is a reduced expression for $Q_{v'}$ that is a subexpression of some reduced expression in S_v.*

Proof. (sketch) The proof is by induction on the dimension of Q_v. For the base case, it suffices to consider the case where α is noncompact type I (other cases are the same as the complex case). In this case, there are two reduced expressions in S_v, each with a single element. Deleting the single element from each of these expressions gives the two closed orbits in the closure of Q_v.

For the inductive case, it again suffices to consider the noncompact type I case. To that end, suppose that α is noncompact imaginary of type I and assume that $Q_v = c_\alpha Q_{v'}$ and $Q_v = c_\alpha Q_{v''}$. Since $\dim(Q_{v'}) + 1 = \dim(Q_{v''}) + 1 = \dim(Q_v)$, we know by induction that orbits in the closure of $Q_{v'}$ and $Q_{v''}$ are given by subexpressions of elements in $S_{v'}$ and $S_{v''}$. However, the action of the Cayley transform (c_α) simply appends α to each such reduced expression. Thus, subexpressions of elements of S_v are exactly those of elements of $S_{v'}$ and $S_{v''}$ with an additional α on the end. These correspond to orbits in the fiber (with respect to π^α) over orbits in the closures of $Q_{v'}$ and $Q_{v''}$. But, by the comments at the beginning of section 3, we see that this exactly describes the closure of Q_v. □