THE DENSE ORBIT OF N ON THE FLAG MANIFOLD, AND APPLICATIONS NOTES FOR MATH 261, FALL 2001

ALLEN KNUTSON

Let B (the "Borel subgroup" of $GL_n(\mathbb{C})$) denote the group of upper triangular matrices, and N its (commutator) subgroup consisting of those with all 1s along the diagonal.

Theorem (The Bruhat decomposition). *In each orbit NgB of N acting by left multiplication on $GL_n(\mathbb{C})/B$, there is a permutation matrix wB.*

(It's unique, but we won't prove this. It was a homework question though.)

Proof. It's easier to think in terms of $N \times B$ orbits on $GL_n(\mathbb{C})$, which are doing rightward column and upward row operations (also rescaling rows). With these it's easy to reduce to a permutation matrix, almost always the antidiagonal one (called the **long element** w_0 of the group S_n). □

In particular N has only finitely many orbits, so at least one of them must be top-dimensional (that's $\binom{n}{2}$-dimensional). Since they're all complex, the lower-dimensional ones can't disconnect the flag manifold (a disconnector would have to be real codimension 1). So there's just one top-dimensional one.

Theorem. *The open orbit of N (the one through w_0) on $GL_n(\mathbb{C})/B$ is free.*

Proof. First, a direct proof. Let $nw_0b = n'w_0b'$, so $w_0bb'^{-1}w_0 = n^{-1}n'$. The LHS is lower triangular, and the RHS is upper triangular with 1s on the diagonal, so both sides are the identity. In particular $n = n'$.

A second proof shows off the power of knowing that our action is algebraic. The dimension of the orbit is right (i.e. $\dim GL_n(\mathbb{C}) = \dim B + \dim N$), so the N-stabilizer of a point of $GL_n(\mathbb{C})/B$ must be discrete. But every nonidentity element of N has infinite order, so the stabilizer must be discrete and infinite (unless it's just the identity), which can't happen for an algebraic subgroup. (Essentially, a polynomial can't vanish infinitely often unless it's zero.) □

1. N-INVARIANT SECTIONS

Theorem (Borel). *Every irrep of B is 1-dimensional. The only irrep of N is the trivial rep.*

Proof. Let V be an irrep of B. Since T is a subgroup of B, we can consider V's weight diagram. Let λ be a weight of V with maximal $(\lambda_1, \lambda_1+\lambda_2, \lambda_1+\lambda_2+\lambda_3, \ldots)$ in lexicographic order. (There may be several such.)

Then each raising operator $e_{ij} \in \mathrm{Lie}(N)$, $i < j$ must act as 0 on the λ weight space, for otherwise it would take it to a weight space with weight larger in that order. Therefore N acts as the identity on the λ weight space.

So any vector in the λ weight space defines a 1-d B-subrep of V. So if V is irreducible, that's all of V.

The statement for N was a homework problem. □

Corollary. *Every rep of* $GL_n(\mathbb{C})$ *has at least one* 1*-d* B*-subrep. (Maybe more than one.) A sum of* k *irreps has at least a* k*-dim space of* N*-invariant vectors.*

We'll make use of the following description of the line bundle $\mathcal{O}(1)$ over $\mathbb{P}(V^*)$. We can identify $\mathbb{P}(V^*)$ with the set of hyperplanes in V (since they are the perps of lines in V^*). The tautological line bundle over $\mathbb{P}(V^*)$ has as line fiber over a hyperplane H the space $H^\perp \cong (V/H)^*$. So the fiber over H in the *anti*tautological bundle can be identified with the line V/H.

In particular, each element $\vec{v} \in V$ gives a section of this line bundle, taking a point H to the vector $\vec{v} + H$ in the fiber V/H. This section vanishes exactly on those H containing $\vec{v}$.

Lemma. *Let* X *be a* G*-space, and* $\mathcal{L}$ *a* G*-equivariant line bundle over it. Let* σ *be a* G*-invariant section. Then* σ*'s value at a point* $x \in X$ *determines it across the whole orbit* $G \cdot x$*.*

This is essentially tautological.

Proposition. *Let* $\mathcal{L}$ *be a* $GL_n(\mathbb{C})$*-equivariant line bundle on* $GL_n(\mathbb{C})/B$*. Then either* $\mathcal{L}$ *has no holomorphic sections (other than zero), or it has a* 1*-d space of* N*-invariant sections.*

Proof. Say σ is an N-invariant section. Pick a point p in the dense N-orbit. The value of σ there determines it on the rest of the orbit (by the lemma), and therefore everywhere by extending it continuously, if that is possible.

If σ is nonzero at p, then any other N-invariant section is just a scalar multiple of σ, hence the 1-dimensionality. □

Proposition. *If* V *is an irrep of* $GL_n(\mathbb{C})$*, with high weight* λ*, then*

1. *There is an equivariant map* $\phi : GL_n(\mathbb{C})/B \to \mathbb{P}(V^*)$*.*
2. *Every irrep has a* 1*-d space of* N*-invariant sections (not more).*
3. *If we define the line bundle* $\mathcal{L}$ *as the pullback of* $\mathcal{O}(1)$ *by* ϕ*, then the space of sections* $\Gamma(GL_n(\mathbb{C})/B; \mathcal{L})$ *of* $\mathcal{L}$ *is* V*.*
4. *The* T*-action on the fiber* $\mathcal{L}_{w_0 B}$ *over* $w_0 B$ *is the weight* λ*.*

Proof. Let $H_{high} \in \mathbb{P}(V^*)$ be the hyperplane in V made by summing all of the weight spaces other than the *lowest* weight space. (Thinking in V^* terms directly, it corresponds to the line of high weight vectors in V^*.) This is obviously B-invariant, so the map $\phi : gB \mapsto g \cdot H_{high}$ is well-defined and manifestly equivariant (for left multiplication by $GL_n(\mathbb{C})$).

Pulling back the line bundle by ϕ, we get a line bundle $\mathcal{L}$ on $GL_n(\mathbb{C})/B$. Pulling back sections, we get an equivariant map

$$\phi^* : \Gamma(\mathbb{P}(V^*); \mathcal{O}(1)) \to \Gamma(GL_n(\mathbb{C})/B; \mathcal{L}).$$

This map is nonzero, because for each point $\phi(gB) \in \mathbb{P}(V^*)$, there is a section of $\mathcal{O}(1)$ not vanishing there. So not all sections pull back to the zero section. Since V is irreducible, the map ϕ^* is injective.

Therefore $\Gamma(GL_n(\mathbb{C})/B; \mathcal{L})$ is a rep containing V plus a bunch of other irreps, say k in all. Therefore its space of N-invariant vectors is $\geq$ k-dimensional. But we know from the last proposition that it's actually 1-dimensional. So $k = 1$, and $\Gamma(GL_n(\mathbb{C})/B; \mathcal{L})$ is irreducible, i.e. it's V.

In particular, V only has a 1-d space of N-invariant vectors. (We know how to find it, too – it's the high weight space.)

Finally, the point $w_0 B$ is fixed by B_- (the lower triangular matrices):

$$b_- w_0 B = w_0 (w_0 b_- w_0) B = w_0 B$$

since $w_0 b_- w_0$ is upper triangular. So it maps to a B_--fixed point in $\mathbb{P}(V^*)$. Thought of as a line in V^*, it must be the lowest weight space. The corresponding hyperplane in V is the sum of all the weight spaces except the high weight space.

Therefore, a nonzero high weight vector $\vec{v}_\lambda$ gives a section of $\mathcal{O}(1)$ which does not vanish at $w_0 B$ (see the discussion from above about how elements of V gives sections of $\mathcal{O}(1)$, and where they do and don't vanish).

Now consider the T-equivariant maps

$$\mathbb{C}\vec{v}_\lambda \hookrightarrow V \cong \Gamma(GL_n(\mathbb{C})/B; \mathcal{L}) \to \mathcal{L}_{w_0 B}$$

inclusion of a weight space, the isomorphism, restriction of sections to the T-fixed point $w_0 B$. That "does not vanish at $w_0 B$" statement says that this composite is nonzero, i.e. a T-equivariant isomorphism of these 1-d spaces.

In other words, we've managed to study the action of T on the fiber $\mathcal{L}_{w_0 B}$ by extending each element of $\mathcal{L}_{w_0 B}$ to an N-invariant section, which is to say a multiple of $\vec{v}_\lambda$. Then that's weight λ, so the original fiber was weight λ. □

Actually, this last result suggests a different way to make the line bundle (other than pulling it back from projective space). Start with a weight λ, i.e. a map $T \to GL_1(\mathbb{C})$. This gives a 1-d rep of B by $B \twoheadrightarrow B/N \cong T \to GL_1(\mathbb{C})$, call this space $\mathbb{C}_\lambda$. Then the space

$$(GL_n(\mathbb{C}) \times \mathbb{C}_\lambda)/B$$

(where B acts by $b \cdot (g, z) := (gb^{-1}, b \cdot z)$) is the total space of a $GL_n(\mathbb{C})$-equivariant line bundle over $GL_n(\mathbb{C})/B$. It usually doesn't have any sections – it does only if λ was dominant.

2. Functions on affine space as a T-rep

Proposition. *Let V be an n-dimensional representation of T. Assume that there exists a vector τ in the Lie algebra of T such that for all weights $\{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ (with repeats to indicate multiplicity) of V, $\langle \alpha_i, \tau \rangle > 0$.*

Then the polynomial ring $\mathrm{Fun}(V)$ of complex-valued polynomials on V, has finite-dimensional T-weight spaces (so, a well-defined weight diagram). In particular, the multiplicity of the weight μ is the number of tuples $(m_1, \ldots, m_n) \in \mathbb{N}^n$ such that

$$\mu = -\sum_i m_i \alpha_i.$$

(Careful about that minus sign!)

Proof. Let $(\vec{v}_i)$ be a basis of weight vectors of V, with the weights (α_i). Let w_i be the dual basis of V^*, with weights $(-\alpha_i)$. Then $\mathrm{Fun}(V)$ has a basis consisting of monomials $\{\prod_i w_i^{m_i}\}$. The weight of such a monomial is $\sum_i m_i(-\alpha_i) = -\sum_i m_i \alpha_i$.

So we get the equation, because each weight space is spanned by a subset of the monomials. The only question is why it's a finite sum. Dotting both sides with τ, we get

$$\langle \mu, \tau \rangle = -\sum_i m_i \langle \alpha_i, \tau \rangle.$$

Since every term on the RHS has the same sign (and not zero), the m_i are bounded, hence the finiteness. $\square$

The V we will care about most is the nilpotent subalgebra $\mathfrak{n}$, the Lie algebra of N. Then the T-weights on $\mathfrak{n}$ are the "positive weights" $\{x_i - x_j\}$, $i < j$. Define $K(\mu)$ to be the number of ways to write μ as such a sum; this is called the **Kostant multiplicity function**. Bert Kostant was actually a professor here for a while, but nowadays he's (emeritus) at MIT.

3. Covariant maps, and line bundles on affine space

Let V, W be two reps of a torus T, and λ a weight of T. Say that a map $\phi : V \to W$ is **λ-covariant** if

$$\phi(t \cdot \vec{v}) = (t^\lambda) t \cdot \phi(\vec{v})$$

for all $\vec{v} \in V, t \in T$. So "equivariant" is "0-covariant". One can also think of it as an ordinary equivariant map from $V \otimes \mathbb{C}_\lambda \to W$, where $\mathbb{C}_\lambda$ is a 1-d rep with weight λ.

If we think of V, W as T-equivariant vector bundles over a point, we can ask for a more general definition. If $\mathcal{V}, \mathcal{W}$ are T-equivariant vector bundles over a T-space X, say that $\Phi : \mathcal{V} \to \mathcal{W}$ is a **λ-covariant bundle map** if Φ takes the fiber $\mathcal{V}_x$ linearly to the fiber $\mathcal{W}_x$ (that's the "bundle map" part), and

$$\Phi(t \cdot \vec{v}) = (t^\lambda) t \cdot \Phi(\vec{v})$$

for all $\vec{v}$ in each fiber of $\mathcal{V}$.

Proposition. *Let $\mathcal{L}_1, \mathcal{L}_2$ be two T-equivariant line bundles on a space X. Let $\Phi : \mathcal{L}_1 \to \mathcal{L}_2$ be a λ-covariant isomorphism. Then the induced map $\Phi : \Gamma(X; \mathcal{L}_1) \to \Gamma(X; \mathcal{L}_2)$ on sections is a $(-\lambda)$-covariant isomorphism.*

In particular, the weight diagram of $\Gamma(X; \mathcal{L}_2)$ is just the weight diagram of $\Gamma(X; \mathcal{L}_1)$, shifted by $-\lambda$.

Recall how we define the action of T on sections $\sigma \in \Gamma(X; \mathcal{L})$:

$$(t \cdot \sigma)(x) := \sigma(t^{-1} x)$$

So σ is weight μ means exactly that

$$\sigma(t^{-1} x) = (t \cdot \sigma)(x) = t^\mu \sigma(x).$$

Proof. Let $\sigma \in \Gamma(X; \mathcal{L})$ be of weight μ. Then

$$(t \cdot \Phi(\sigma))(x) = t^{-\lambda} (\Phi(\sigma))(t^{-1} \cdot x) = t^{-\lambda} \Phi(\sigma(t^{-1} \cdot x)) = t^{-\lambda} \Phi(t^\mu \sigma(x)) = t^{\mu-\lambda} \Phi(\sigma(x)) = t^{\mu-\lambda} (\Phi(\sigma))(x)$$

so

$$t \cdot \Phi(\sigma) = t^{\mu-\lambda} \Phi(\sigma).$$

$\square$

Theorem. *The multiplicity of the weight μ in the irrep V_λ is bounded above by $K(\lambda - \mu)$.*

Proof. Let $\mathcal{L}$ denote the line bundle on $GL_n(\mathbb{C})/B$ whose sections give V_λ.

We'll compare sections of three different T-equivariant line bundles:

- the line bundle $\mathcal{L}$ on $GL_n(\mathbb{C})/B$
- the line bundle $\mathcal{L}$ on the N-orbit through $w_0 B$
- the trivial line bundle on the N-orbit through $w_0 B$.

By restriction, the sections of the first (our rep V_λ) T-equivariantly includes into the space of sections of the second. (We get an inclusion precisely because some sections on the big cell may develop poles when we try to extend them.) So what's that latter space?

Since the orbit is a *free* orbit, there is a unique element of N relating any two points. Using that, we can identify all the fibers of $\mathcal{L}$ over the orbit with the fiber over $w_0 B$. So the line bundle in question is really $N w_0 B \times \mathcal{L}_{w_0 B}$. This statement is true T-equivariantly, too, since T normalizes N.

So we have a $(-\lambda)$-covariant inclusion of V_λ into $\mathrm{Fun}(N)$. Hence the multiplicity of μ in V_λ is bounded by the multiplicity of $\mu - \lambda$ in $\mathrm{Fun}(N)$. And that's $K(\lambda - \mu)$. □

Note that this actually gives $n!$ different bounds, since the weight diagram must be S_n-invariant.