Weight Varieties

by

Allen Ivar Knutson

B.S., Caltech, 1991

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

Doctor of Philosophy in Mathematics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 1996

Author ....
Department of Mathematics
May 3, 1996

Certified by ....
Victor Guillemin
Professor of Mathematics
Thesis Supervisor

Accepted by ....
David Vogan
Chairman, Departmental Committee on Graduate Students

Weight Varieties

by

Allen Ivar Knutson

Submitted to the Department of Mathematics
on May 3, 1996, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Mathematics

Abstract

I define "weight varieties", the geometric analogues of weight spaces of irreducible representations of a Lie group G. The zero weight variety carries a natural action of the Weyl group; I give a description of the fixed-point set of each element. I reinterpret the Gel'fand-MacPherson correspondence to show that certain $GL_n(\mathbf{C})$ weight varieties can be identified with moduli spaces of polygons in $\mathbf{R}^3$, and that on them the residual Gel'fand-Cetlin system has a concrete geometrical interpretation. Lastly, I compare those with the moduli space of flat $SU(2)$ connections on an n-holed Riemann sphere.

Thesis Supervisor: Victor Guillemin
Title: Professor of Mathematics

Acknowledgments

Naturally the greatest influence on this thesis is that of my advisor, Victor Guillemin, who suggested the problem and made sure I saw it to its present state of completion. I can't thank him enough for bringing into existence such a stimulating environment for symplectic geometry as MIT has been, and I warmly acknowledge the influence of the many other symplectic geometers here, especially Ana Canas da Silva, Bertram Kostant, Shaun Martin, Michael Thaddeus, and Chris Woodward. I owe a special debt also to my Princeton advisor Lisa Jeffrey for spending a lot of time bringing me into symplectic geometry. Reaching back to Caltech, I must thank Yaser Abu-Mostafa for helping me learn to do mathematical research.

My greatest scientific influence has of course been my parents. "Without them, I wouldn't have..." is a bit tautological.

Dramatis Personae

Standard notation

$\mathbb{N} \leq \mathbb{Z} \leq \mathbb{R} \leq \mathbb{C}$	the naturals (including 0), integers, reals, and complex numbers
$G \geq B = NT$	a reductive connected complex Lie group, a Borel subgroup, its unipotent radical, and a maximal torus
$W = N_G(T)/T$	the Weyl group of G
$C_H(K)$	the centralizer of K in H (K, H subgroups of G)
$Z(G) = C_G(G)$	the center of G
$GL_n(\mathbb{C})$	the group of invertible $n \times n$ complex matrices
$U(n)$	its compact real form, the unitary $n \times n$ complex matrices
T^n	the diagonal matrices in either context; in each case a maximal torus
t^*_+	the usual (closed) positive Weyl chamber in t^*
$\mu \in t^*$ $(\lambda \in t^*_+)$	a weight (resp. a high weight) of T
V_λ	the irreducible representation of G with high weight λ
Λ_k	the high weight of the representation $\mathrm{Alt}^k\mathbb{C}^n$ of $GL_n(\mathbb{C})$, i.e. $V_{\Lambda_k} = \mathrm{Alt}^k\mathbb{C}^n$

Notation special to this thesis

$L_\pi = C_G(C_T(\pi)^0)$	the Levi subgroup associated to π
$\mathcal{O}_\lambda := B\backslash\backslash_{\lambda} G$	the flag manifold $B\backslash G$ carrying the λ line bundle
$\mathbb{C}_\mu$	the irrep of T with character μ
$(V_\lambda)^\mu := \mathrm{Hom}_T(\mathbb{C}_\mu, V_\lambda)$	the subspace of V_λ upon which T acts with character μ: the μ weight space of V_λ
$(\mathcal{O}_\lambda)^\mu := \mathcal{O}_\lambda //_{\mu} T$	the geometric invariant theory quotient of $\mathcal{O}_\lambda$ by T at μ: the μ weight variety of $\mathcal{O}_\lambda$.

Throughout, $X//G$ refers to the geometric invariant theory quotient [24] of X by G with respect to some implicit G-polarization. The basics of this are presented in Appendix 2.

$X//_{\lambda}G$ includes the specification of the G-polarization, or the twist relative to an implicit one (in which case the quotient is really with respect to a Borel B, as explained in Appendix 3).

By picking a maximal compact subgroup $K \leq G$, and a Kähler metric invariant under K, these quotients can be identified with symplectic reductions (as explained in Appendix 2), which is necessary for further identifications made in part II.

Familiarity is assumed with the theory of reductive groups, in particular the finite-dimensional representation theory, Borel subgroups, and the dual role of Levi subgroups as Levi factors of parabolics or centralizers of tori [20]. Some exposure to Hamiltonian actions on symplectic manifolds is also desirable though less necessary [3][13].

One fact used repeatedly is that for each Weyl group element $\pi \in W$, any lift $\tilde{\pi} \in G$ is semisimple. (Proof: $\tilde{\pi}^{|\pi|} \in T$.) Another is that lifts $\tilde{\chi}$ of Coxeter elements (products of all the simple reflections) are regular [19].

Contents

0. Introduction and summary of results

The geometric view of representation theory guiding this thesis begins with the Borel-Weil theorem [20], which is concerned with the ample line bundles over the space $B\backslash G$, where G (for us) is a reductive complex algebraic group and B a Borel subgroup.

Any character $\lambda : T \to \mathbb{C}^\times$ of the maximal torus gives us a line bundle over $B\backslash G$ in the following way: compose λ with the natural map $B \twoheadrightarrow T$ (since B is the semidirect product of T and $[B, B]$), and use this character of B to associate a line bundle to the principal B-bundle over $B\backslash G$ with total space G. Such a line bundle $\mathcal{L}_\lambda := G \times_B \mathbb{C}_\lambda$ comes with a G-action such that the projection to $B\backslash G$ is G-equivariant. In particular, G acts on the finite-dimensional space of holomorphic sections $H^0(B\backslash G, \mathcal{L}_\lambda)$.

Theorem (Borel-Weil). *If λ is in the closed positive Weyl chamber of* G *determined by* B, *then the space $H^0(B\backslash G, \mathcal{L}_\lambda)$ is V_λ, the unique irrep of* G *with high weight λ, and the higher sheaf cohomology vanishes.*

Corollary. *Let λ be in the interior of the positive Weyl chamber. Then $B\backslash G \cong \mathrm{Proj} \bigoplus_{n\in\mathbb{N}} V_{n\lambda}$ as varieties, where the ring is given the unique* G*-invariant multiplication (the "Cartan product").*

(For λ on a wall of the Weyl chamber, the associated projective variety is $P\backslash G$, where P is a parabolic containing B corresponding to the wall. These are the bundles that are not ample but merely "nef".)

One direct benefit of realizing this vector space as the sheaf cohomology of a line bundle is the chance to apply the Atiyah-Segal-Singer formula [4]. Given an automorphism γ of a complex-manifold-with-line-bundle $(M, \mathcal{L})$ and thus an induced action on the Dolbeault cohomology, this formula calculates the "supertrace" $\sum_i (-1)^i \operatorname{Tr} \gamma|_{H^i(\mathcal{L})}$ as an integral over the fixed point set M^γ on the base. In the original Atiyah-Bott paper this formula is used to rederive the Weyl character formula by way of Borel-Weil; a regular element $g \in G$ acts on $B\backslash G$ with $|W|$ isolated fixed points, where W is the Weyl group of G.

The goal of this thesis is to extend this circle of ideas to the computation of **weight spaces** $(V_\lambda)^\mu := \mathrm{Hom}_T(\mathbb{C}_\mu, V_\lambda)$ of an irrep V_λ – the isotypic components that the rep breaks into under a chosen maximal torus $T \leq G$.

A priori, these spaces have only two invariants: their dimension, and a possible Weyl group action (if μ is stable under some of the Weyl group) [9][16][1][2]. In this thesis we make use of the natural ring structure on $\bigoplus_{n\in\mathbb{N}} (V_{n\lambda})^{n\mu}$, and define the **$\mu$-weight variety**

$$(\mathcal{O}_\lambda)^\mu = \mathrm{Proj} \bigoplus_{n\in\mathbb{N}} (V_{n\lambda})^{n\mu} = B \underset{\lambda}{\backslash\backslash} G \underset{\mu}{//} T.$$

In particular the zero weight variety carries an action of W.

In contrast to the study of entire irreps, where there is a single space $B\backslash G$ (and perhaps certain quotients $P\backslash G$), the complex variety $B \underset{\lambda}{\backslash\backslash} G \underset{\mu}{//} T$ depends on λ and μ. In I.1 I explain the chamber structure in $\mathfrak{t}_+^* \times \mathfrak{t}^*$ determining the set of weight varieties, following [12] – the only new result here is the condition for a wall to be external.

For generic λ and μ these spaces (when nonempty) are dense compact Hausdorff subsets of the very non-Hausdorff double coset space $B\backslash G/T$. In section I.2 I give a description of the Weyl group action on this space, determining the fixed points of each given element π; it remains to be seen which lie in a given zero weight variety. Nevertheless, there are consequences for the character of the Weyl group on the zero weight space (via the orbifold Atiyah-Segal-Singer formula), spelled out in section I.3.

The remainder of the thesis is a study of certain very specific weight varieties $B \underset{p\Lambda_2}{\backslash\backslash} GL_n(\mathbb{C}) \underset{\mu}{//} T$ (where Λ_2 is the second fundamental weight, so $\mathcal{O}_{\Lambda_2}$ is a 2-Grassmannian); this is an excerpt of a paper with

Jean-Claude Hausmann [18]. In section II.1 I use the Gel'fand-MacPherson correspondence [11] to give an alternate interpretation of $(\mathcal{O}_{\wedge_2})^{\mu}$ as a moduli space of polygons in $\mathbb{R}^3$. I show that the Gel'fand-Cetlin system on the Grassmannian descends to the geometrically intuitive "bending flows" on the polygon space.

In section II.3 I relate these polygon spaces to the now famous moduli space of flat SU(2) connections on a multiply punctured 2-sphere.

The Appendices contain the basics of geometric invariant theory for affine and projective varieties, and some of their relation to symplectic quotients.

I. Cultural facts about weight varieties

There is no vast literature on weight varieties, so for the convenience of the reader I give here a to-my-knowledge complete summary of what is known. This section can be safely skipped; only the Gel'fand-Cetlin system is used later, where it is reintroduced.

We mention up front the alternate, symplectic, origin of flag manifolds $B\backslash G$ – they arise as generic orbits of G's maximal compact subgroup K acting on the dual of its Lie algebra. (The nongeneric orbits are K-diffeomorphic to $P\backslash G$ for the various parabolics P.) Even better, the symplectic structure that $B\backslash\backslash_{\lambda} G$ pulls back via its projective embedding is the same as the Lie-Kirillov-Kostant-Souriau symplectic structure on the coadjoint orbit through λ. What we will use this identification $\mathcal{O}_\lambda \cong$ [orbit through λ] for is to give the natural map $\mathcal{O}_\lambda \to t^*$, composing inclusion into k^* with the transpose of the inclusion $t \to k$.

The first, albeit indirect, result on weight varieties is Kostant's linear convexity theorem [23], which says which weight varieties $(\mathcal{O}_\lambda)^\mu$ are nonempty; the weight μ must be in the convex hull of the orbit of λ under the Weyl group. (For GL_n this result is the much older Schur-Horn theorem.) Today this is just seen as an easy corollary of the more general Atiyah/Guillemin-Sternberg convexity theorem for Hamiltonian torus actions on symplectic manifolds [3][13].

For generic μ, the weight variety will be an orbifold (a general fact about symplectic reductions of manifolds). What is really a very surprising fact is that for GL_n they are actually *manifolds*. This is true only for GL_n, a factoid that is relegated to a footnote in the seminal paper [7].

Another phenomenon in the $GL_n(\mathbb{C})$ case is the Gel'fand-Cetlin system, which gives a complete system of Poisson commuting functions on a flag variety, extending those generating the action of the maximal torus [14]. In particular the new ones descend to a complete system on the weight varieties.

$GL_n(\mathbb{C})$ has one last trick up its sleeve: a partial flag manifold with a dense orbit of the torus – I refer of course to CP^{n-1}. The manifold $\mathcal{O}_\lambda$ of full flags is a bundle over this; the bundle map takes a flag, a list of subspaces, to its first entry, the 1-d subspace. A fiber of this $GL_n(\mathbb{C})$-equivariant symplectic fibration is the space of flags starting with that line, or equivalently, the full flags in the quotient of $\mathbb{C}^n$ by that line.

The minimal coupling theorem [12] says roughly that "if the fibers of an equivariant symplectic fibration are small enough, the symplectic quotients also have a bundle map *with the same fiber*". "Small enough" in this case means that λ should be in a certain open neighborhood in t^*_+ of the edge containing the first fundamental weight. (The last fundamental weight also has such a neighborhood, if one uses a different map to projective space – instead of picking the line from a flag, pick out the hyperplane.)

In the symplectic fibration at hand, the fiber is a flag manifold one dimension down, the reduced total space is the weight variety, and the reduced base is a point – hence an identification of the weight variety with the one-lower flag manifold! This actually holds for an open set of μ's, also, producing a region in the center of the Weyl polytope where all the weight varieties are symplectomorphic. This phenomenon is explored in detail in the book [12].

I.1. The chamber structure on $\{(\lambda, \mu)\}$

To begin with, fix a regular high weight $\lambda \in t^*_+$, and consider the moment map $\mathcal{O}_\lambda \to t^*$ of the torus action. (The regular λ are those such that $B\backslash\backslash_{\lambda} G \cong B\backslash G$; everything that follows except for the uniqueness statements is true also for non-regular λ.) The image of $\mathcal{O}_\lambda$ has more structure than merely that of a convex polytope; it is divided into chambers of μ's, within each of which the quotient $(\mathcal{O}_\lambda)^\mu$ does not depend (up to complex isomorphism) on the μ chosen [7][15]. The purpose of this section is to find the walls separating the chambers (the images of the submanifolds with positive-dimensional stabilizer).

Lemma. *The fixed points of a subtorus* $T' \leq T$ *acting on the right on* $B\backslash G$ *are*

$$\bigcup_{w \in W} B w C_G(T').$$

Proof. Pick $t \in T'$ generating it topologically; the t-fixed points are then the same as the T'-fixed points, and $C_G(t) = C_G(T')$.

The fixed point set of t are those

$$\{Bg \in B\backslash G : Bgt = Bg\} = \{Bg \in B\backslash G : Bgtg^{-1} = B\} = \{Bg \in B\backslash G : gtg^{-1} \in B\}$$

(the last does not depend on the choice of g in Bg).

Any torus in B can be B-conjugated to lie in T. This corresponds to left-multiplying g by an element of B, which doesn't change the element Bg. So it is enough to consider those g such that $gtg^{-1} \in T$.

The G-conjugates of t in T are all already conjugate under $N(T)$, so

$$\{g : gtg^{-1} \in T\} = N(T)\left\{h : hth^{-1} = t\right\}$$

but this last is just $C_G(t)$. □

(This line of argument is repeated in the next section, toward a slightly different purpose.)

Corollary. *The only* T*-stabilizers which occur are the centers of Levi subgroups, and they all do actually occur as the identity components of stabilizers.*

Proof. The fixed points $\bigcup_{w \in W} BwC_G(T')$ are invariant under $Z(C_G(T'))$, since the latter is a subgroup of T.

It remains to find a point on $BC_G(T')$ whose stabilizer is no larger than this. The rest of this proof is a little more convenient in the compact world; until its conclusion take T, L to refer to what used to be their maximal compact subgroups.

For each Levi subgroup L, pick a lift $\tilde{\chi}$ of a Coxeter element (a product of all the simple reflections) of the Weyl group of L. Since $\tilde{\chi}$ is semisimple we can use some $l \in L$ to conjugate it into the maximal torus, $l\tilde{\chi}l^{-1} \in T$.

Then consider the torus stabilizer of the point $T\,l$:

$$\{t \in T : T\,lt = T\,l\} = \{t \in T : ltl^{-1} \in T\} = T \cap l^{-1}T\,l$$

However, since $\tilde{\chi}$ is regular [19], this latter torus $l^{-1}Tl$ containing $\tilde{\chi}$ is just $C_L(\tilde{\chi})$, so its intersection with T is $C_T(\tilde{\chi})$, whose identity component is $Z(L)$. □

Luckily, all we need in order to find the walls separating the chambers are the fixed points with minimal stabilizer; these correspond to the maximal Levi subgroups.

Example: $GL_n(\mathbb{C})$. The choice of T corresponds to a decomposition of $\mathbb{C}^n$ as a sum of lines; T is the subgroup stabilizing each line. Each equivalence relation on $\{1, \ldots, n\}$ corresponds to a way to sum up lines into subspaces, thus a coarser decomposition of $\mathbb{C}^n$, and a subgroup preserving each subspace; these are the Levi subgroups. The maximal ones correspond to partitions into exactly two cells, of which there are $(2^n - 2)/2$.

Very few of these are actually Levi factors of parabolics containing B – only those corresponding to equivalence relations whose cells are intervals in $\{1, \ldots, n\}$. There are 2^{n-1} of these, of which $n - 1$ are maximal.

N.B. The reader may be tempted to generalize from this example that Levi subgroups are exactly those reductive subgroups with rank equal to that of the whole group. Unfortunately this is not true; $SL_3(\mathbb{C}) \leq G_2$ is not a Levi subgroup.

In the rest of this section, we will call the image of $BL \subseteq B\backslash G$ under the moment map "L's wall" or "the wall coming from L". (These are only the walls through λ, but the others are just Weyl translates of these.)

Corollary. *Each Levi subgroup has a distinct wall, and the external walls incident on λ come exactly from those Levi subgroups that are actually Levi factors of parabolics containing* B. *More generally, the external faces incident on a corner* Bw *are the images of* BwL, *where* L *is a Levi factor of a parabolic containing* $w^{-1}Bw$.

Proof. The "slope" of L's wall, as part of an affine subspace of $\mathfrak{t}^*$, is the perp to the Lie algebra of the stabilizer, which by the previous Lemma is $Z(L)$. So different walls can be told apart.

The external edges of the Weyl polytope incident on the high weight λ connect λ to its reflections through the r simple roots (here r is the rank of G). The Levi factor generated by T, a k-element subset of the simple roots, and their negatives, will hit those k edges and thus have an image at least k-dimensional; conversely, it is the centralizer of an $(r-k)$-dimensional torus, and thus the image is at most k-dimensional. So we've found the Levi subgroup producing that external wall, and it's a Levi factor.

The last statement comes by left-multiplying by the element w^{-1} to reduce the problem to the one already solved. □

Theorem. *Each wall separating chambers is a subset of*

$$\{(\lambda,\mu) : \mu \in \mathfrak{z}^- + \lambda^w\}$$

for a pair $(\mathfrak{z}, w)$, *where* $\mathfrak{z}$ *is the center of a maximal Levi subalgebra, and* w *is a Weyl group element. Every such* $(\mathfrak{z}, w)$ *arises, and each wall comes from a unique* $\mathfrak{z}$ *and a* w *well-defined in* W *up to left multiplication by the stabilizer of* $\mathfrak{z}$.

Proof. This is just a summary of the Lemmata, except for the uniqueness statement about the w.

The $\mathfrak{z}$-wall incident on λ^w is the image of BwL, where $\mathfrak{z}$ is the Lie algebra of the center of L. The other vertices that this hits are Bww', where $w' \in W_L$. But the Weyl group of the centralizer of $\mathfrak{z}$ is the stablizer of $\mathfrak{z}$ in the whole Weyl group. □

It would be nice to say that this subset is simply the intersection of the Weyl polytope with the affine hyperplane given, as a nice generalization of the fact that the polytope has no internal vertices. Unfortunately this is not true, making it difficult to enumerate the chambers: the smallest example comes from the maximal Levi subgroup $\begin{pmatrix} * & & & * \\ & * & * & \\ & * & * & \\ * & & & * \end{pmatrix}$, whose wall is a rectangle, two of whose edges lie wholly within the Weyl polytope.

I.2. The fixed points of Weyl group elements on $B\backslash G/T$

By its definition, $(\mathcal{O}_\lambda)^0$ comes into existence carrying an action of the Weyl group $W = N_G(T)/T$. In this section we describe one aspect of W's action, the fixed point set of each element $\pi \in W$ on $B\backslash G/T$, a non-Hausdorff space in which $(\mathcal{O}_\lambda)^0$ sits densely (for those λ for which the GIT quotient has no properly semistable points – equivalently, zero a regular value of the map $\mathcal{O}_\lambda \to \mathfrak{t}^*$).

A point BmT in $B\backslash G/T$ is invariant under π if for some lift $\tilde{\pi} \in N_G(T)$, $b \in B$, $t \in T$,

$$m\tilde{\pi} = bmt.$$

Lemma. *Fix a lift* $\tilde{\pi} \in N_G(T)$ *of the Weyl group element* π. *Let* $C_T(\pi)^0 \leq T$ *be the connected component of the identity in the centralizer of* $\tilde{\pi}$ *in* T. *Under the natural surjection* $T\backslash G/C_T(\pi)^0 \twoheadrightarrow B\backslash G/T$, *any* π*-fixed point in* $B\backslash G/T$ *lies under at least one* $\tilde{\pi}$*-fixed point in* $T\backslash G/C_T(\pi)^0$.

(Technically I must speak of $C_T(\tilde{\pi})$ rather than $C_T(\pi)$, but $C_T(\tilde{\pi})$ doesn't depend on the choice of lift. In what follows I omit the tilde where possible.) Another way of saying this is that since we only care about BmT, not the actual m, we can massage m into a normal form by multiplication by elements of B and T.

Proof. **From T to $C_T(\pi)^0$ on the right.** To adjust the t, we'll replace m by mu, $u \in T$, for a judicious choice of u.

$$\begin{aligned}(mu)\tilde{\pi} &= m\tilde{\pi}u^{\pi} = bmtu^{\pi} \\ &= b(mu)tu^{-1}u^{\pi}\end{aligned}$$

π acts diagonalizably on $\mathfrak{t}$ (being of finite order); decompose $\mathfrak{t} = \mathfrak{t}_1 \oplus \mathfrak{t}_{rest}$ into the 1-eigenspace and the sum of the other eigenspaces. Then every element t can be written (not necessarily uniquely) as a product $t_1 t_{rest}$ of elements from $\exp(\mathfrak{t}_1)$ and $\exp(\mathfrak{t}_{rest})$.

The map $u \mapsto u^{-1}u^{\pi}$ is, on the Lie algebra level, $U \mapsto (\pi - 1)U$, so maps $\mathfrak{t}$ onto $\mathfrak{t}_{rest}$. Therefore the same is true on the group level, so there exists a $u \in T$ such that $u^{-1}u^{\pi} = t_{rest}^{-1}$. Then $tu^{-1}u^{\pi} \in \exp(\mathfrak{t}_1) = C_T(\pi)^0$, as desired.

From B to T on the left. Any choice of lift $\tilde{\pi}$ is semisimple. And so is $\tilde{\pi}t^{-1}$, since they commute. Thus $b = m(\tilde{\pi}t^{-1})m^{-1}$ is too, which means it's B-conjugate to something in T; $c^{-1}bc \in T \leq B$.

$$(c^{-1}m)\pi = c^{-1}bmt = c^{-1}bc(c^{-1}m)t$$

Replace m by $c^{-1}m$, and we shift b into T. □

Corollary. *In the case that $C_T(\pi)^0 = Z(G)$, the fixed points on $B\backslash G/T$ lie under fixed points on $T\backslash G$.*

Proof. In this case we can pull the t past m and combine it with b, eliminating it. □

Weyl group elements with this property are reasonably rare. Coxeter elements (a product of all the simple reflections) have this property [19]; in the A_n series they (the n-cycles) are the only ones such. In G_2's Weyl group, by contrast, all three conjugacy classes of rotations have this property.

An equivalent description of this property (when G is semisimple, so $Z(G)^0 = 1$) is that 0 is the only weight fixed by π, which implies that the trace of π on the zero weight space of a representation equals the trace of $\tilde{\pi}$ acting on the whole representation. The Corollary is a geometric analogue of this, saying that the action of such an element on $B\backslash G/T$ is already understandable from its action of $B\backslash G$ (or even $T\backslash G$).

The following Lemma is more easily stated if G is semisimple, which doesn't affect anything else (since the center of G already lies in B), so assume it.

Lemma. *Fix $l \in L_\pi := C_G(C_T(\pi)^0)$ such that $l\tilde{\pi}l^{-1} \in T$ (which exists since $\pi \in L_\pi$ and L_π is reductive). The G-elements lying above $\tilde{\pi}$ fixed points on $T\backslash G/C_T(\pi)^0$ are*

$$N(T)\left(\bigcup_{t \in C_T(\pi)^0} Z(l\tilde{\pi}l^{-1}t)\right) l.$$

(This seemingly infinite union can actually be taken over a finite set of t's.)

Proof. We are looking for $g \in G$ such that

$$\exists b \in T, \exists t \in C_T(\pi)^0 \text{ such that } g\tilde{\pi} = bgt^{-1},$$

or equivalently,

$$\exists t \in C_T(\pi)^0 \text{ such that } g\tilde{\pi}tg^{-1} \in T.$$

Bringing in the union over t is just a restatement:

$$\bigcup_{t \in C_T(\pi)^0} \{g : g(\tilde{\pi}t)g^{-1} \in T\}$$

There are only a finite number of conjugates of $\tilde{\pi}t$ in T, conjugate to one another under $N(T)$. One of them is $l(\tilde{\pi}t)l^{-1} = l\tilde{\pi}l^{-1}\,t$ (here l commutes with t since $l \in L_\pi$).

$$\exists n \in N(T) \quad n(l\tilde{\pi}tl^{-1})n^{-1} = g\tilde{\pi}tg^{-1} = gl^{-1}l\tilde{\pi}tl^{-1}lg^{-1}$$

or equivalently, $n^{-1}gl^{-1} \in Z(l\tilde{\pi}l^{-1}t)$, and finally,

$$g \in N(T)\,Z(l\tilde{\pi}l^{-1}t)\,l.$$

The potential finiteness of the union is due to there only being finitely many Levi subgroups (such as the $Z(\,l\tilde{\pi}l^{-1}\,t))$ containing T. □

Examples.

1. π a *Coxeter element*. In this case the only $t \in C_T(\pi)^0$ is $t = 1$, and since π is regular, $Z(l\tilde{\pi}l^{-1}) = T$. So the set described in the Lemma is just $N(T)l$, which has $|W|$ elements in $T\backslash G$, and fewer in $B\backslash G/T$. (In $GL_n(\mathbb{C})$ a judicious choice $l_{ab} = \exp(2\pi i ab/n)$ lets one show that two elements Bw, $Bw' \in B\backslash GL_n(\mathbb{C})$ fall together in $B\backslash GL_n(\mathbb{C})/T$ iff they are in the same coset of $C_W(\pi)$, leading to $(n-1)!$ fixed points.)

2. *More generally, if* $C_T(\pi)^0 = \{1\}$. Now $N(T) \cap Z(l\tilde{\pi}l^{-1}) = N_{Z(l\tilde{\pi}l^{-1})}(T)$, so the set described in the Lemma is $|W/W_{Z(l\tilde{\pi}l^{-1})}|$ copies of $Z(l\tilde{\pi}l^{-1})$.

3. $G = GL_n(\mathbb{C})$, $W = S_n$. Let π be a product of k cycles of lengths $\{a_1, \dots, a_k\}$. Pick the lift $\tilde{\pi}$ to be the corresponding permutation matrix. Its eigenvalues are the union-with-multiplicity of the a_ith roots of unity, $i = 1 \dots k$. But when we pick a different lift $\tilde{\pi}t$, where $t = \mathrm{diag}(a_1\; t_1\text{'s}, a_2\; t_2\text{'s}, \dots, a_k\; t_k\text{'s}) \in C_T(\pi)^0$, the eigenvalues are $\bigcup\{$the a_ith roots of unity times $t_i\}$. There is a pretty geometric picture of this: take k concentric circles, with a_i equally spaced dots on the ith circle, one pointing north. The t_i say how to rotate the circles from true north. $Z(\tilde{\pi}t)$ increases as a set when dots on different circles line up. Rotating all the circles together corresponds to multiplying $\tilde{\pi}$ by a central element of $GL_n(\mathbb{C})$, which doesn't change its centralizer, nor the points it stabilizes in $B\backslash G$.

For example, let $\pi = \binom{1234}{1243}$. Then the two types of ts are $t = 1$, with $Z(\tilde{\pi}t) \cong GL_3 \times GL_1$, and $t = \mathrm{diag}(1,-1,1,1)$, with $Z(\tilde{\pi}t) \cong GL_2 \times GL_2$. These eventually lead to components of the fixed points set in $B\backslash GL_4/T$ of different dimensions (1 and 0).

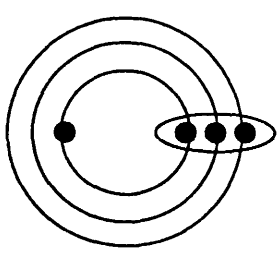

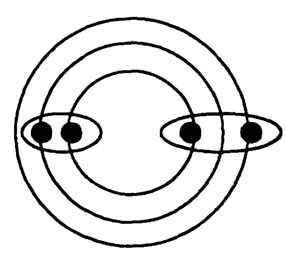

(It is not hard to give the condition that says when we avoid this phenomenon of multiple maximal centralizers: π must be a product of cycles $\{a_i\}$ such that the least common multiple of any pair equals the least common multiple of the whole set – a product of two cycles, for instance, or $\{6, 10, 15\}$.)

Corollary. *The fixed point set of* π *on* $B\backslash G/T$ *is a union over maximal centralizers* $C = Z(\,l\tilde{\pi}l^{-1}\,t)$ *of* W/W_C *copies of* $B_C\backslash C/T_C$ *(where* W_C, B_C, T_C *are the Weyl group, Borel, and torus of* C*).*

(Note that no promise is made, or can be made, as to the disjointness of this union. It seems though that the points of overlap have continuous stabilizers and are thus never "stable" from the GIT quotient point of view.)

Unfortunately it seems difficult to say which of these fixed points lie in a given $(\mathcal{O}_\lambda)^0$ inside $B\backslash G/T$. (Technically speaking $(\mathcal{O}_\lambda)^0$ is only a subset of $B\backslash G/T$ when the reduction has no properly semistable points, which happens for generic λ.) For example in $G = GL_3(\mathbb{C})$, the condition above is satisfied by the transposition $(12) \in S_3$, the Weyl group of GL_3, and there are three fixed points on $B\backslash GL_3/T$. However, for any regular λ other than multiples of the high root ρ (this is the "small fibers" condition alluded to in I.0), the reduction $(\mathcal{O}_\lambda)^0$ is $\mathbb{CP}^1$ and one of the fixed points is unstable. Which one depends on on which side of the ρ line λ lies.

I.3. Consequences for the character of W on the zero weight space

Let γ be an automorphism of a complex manifold M with a specified lift to a vector bundle $\mathcal{W}$ over the manifold. Atiyah, Segal, and Singer have given a pleasant formula [4] for the supertrace of the action on the sheaf cohomology of the bundle as an integral over the fixed point set.

For the problem at hand, M is $(\mathcal{O}_\lambda)^0$ (and so only an orbifold for G not $GL_n(\mathbb{C})$), $\mathcal{W}$ is the line bundle induced from $\mathcal{L}_\lambda$ on $\mathcal{O}_\lambda$, and γ is a Weyl group element. The orbifold Atiyah-Segal-Singer formula [6] then gives the trace of γ's action on the total cohomology.

Unfortunately I cannot prove that the higher cohomology always vanishes. Each of the walls found in I.1 has an associated distance, such that only for λ far enough away from the wall does the associated $H^{\bullet>0}((\mathcal{O}_\lambda)^0; \mathcal{L}_\lambda)$ vanish. When the canonical line bundle of $(\mathcal{O}_\lambda)^\mu$ is ample, this distance is negative – hopefully this is always the case.

For λ not on these walls, the higher cohomology will vanish for all $k\lambda$ with k large enough, leaving only the sections – the zero weight space $(V_\lambda)^0$.

In particular the character of π on the zero weight space of $V_{k\lambda}$ is a quasipolynomial in k of degree at most the dimension of the fixed point set (and actually equal to the dimension, barring cancellation between the components of top dimension).

Theorem. *For each $\pi \in S_n$ with at most two orbits, there is a constant c such that for any irreducible representation V of $GL_n(\mathbb{C})$, $\left|\mathrm{Tr}\,\pi|_{V^0}\right| < c$.*

Proof. These π are exactly the elements with isolated fixed points, as calculated in the last section. (The permutation matrix lift $\tilde{\pi}$ of a permutation with three orbits will contribute positive-dimensional components to the fixed point set.) □

Conjecture. *The permutation matrix lift $\tilde{\pi}$ is is the unique way to maximize the dimension of the components of the fixed point set. In particular the degree of polynomial growth of the character is bounded by (and still more conjecturally, equal to) the dimension of the zero weight variety of the centralizer of $\tilde{\pi}$.*

I have verified the weaker part of the conjecture for small n, but large enough to be convinced. I thank Terry Tao for some assistance on this strange problem, who proved it in the case that neither 2 nor 3 divide any of the cycle lengths.

II. 2-Grassmannians and polygon spaces

This half of the thesis is condensed from the paper [18] written jointly with Jean-Claude Hausmann. The subject matter dictates a symplectic, rather than GIT, approach to reduction, and all the groups referred to are compact.

II.1 The identification

In the case that λ is a multiple p of the second fundamental weight Λ_2, the space $\mathcal{O}_{\Lambda_2}$ is a 2-Grassmannian, and can be alternately constructed as

$$\mathcal{O}_{\Lambda_2} = U(2)\underset{p\mathbf{1}}{\backslash\backslash} M_{2\times n}(\mathbb{C})$$

or

$$\mathcal{O}_{\Lambda_2} = \{U^{-1}\operatorname{diag}(p,p,0,0,0,\dots,0)\,U : U \in U(n)\}.$$

It is not difficult to see concretely that this space is the 2-Grassmannian claimed. Compose the $U(2)$ moment map $M_{2\times n}(\mathbb{C}) \to u(2)^*$ with the trace form identification $u(2)^* \cong iu(2)$ to get

$$\Phi_2 : M_{2\times n}(\mathbb{C}) \to iu(2),$$
$$M \mapsto MM^*.$$

The symplectic reduction $U(2)\underset{p\mathbf{1}}{\backslash\backslash} M_{2\times n}(\mathbb{C})$ is defined as $U(2)\backslash S$, where

$$S := \Phi_2^{-1}(p\mathbf{1})$$

for "Stiefel manifold"; it is (p times) the set of orthonormal pairs in $\mathbb{C}^n$. Quotienting it by $U(2)$ forgets the actual pairs of vectors, and remembers only their 2-dimensional span – hence the identification with the 2-Grassmannian.

Meanwhile, the $U(n)$ moment map (composed with minus the trace identification) is

$$\Phi_n : M_{2\times n}(\mathbb{C}) \to iu(n), \qquad M \mapsto M^*M.$$

The image of S under Φ_n is the aforementioned $\{U^{-1}\operatorname{diag}(p,p,0,0,0,\dots,0)\,U : U \in U(n)\}$. The fibers thereof are exactly the $U(2)$ orbits, producing the identification

$$U(2)\underset{p\mathbf{1}}{\backslash\backslash} M_{2\times n}(\mathbb{C}) \cong \{U^{-1}\operatorname{diag}(p,p,0,0,0,\dots,0)\,U : U \in U(n)\}.$$

So far this has been about a certain partial flag manifold; we now reintroduce weight varieties. For this we will need the moment map $\phi_n : M_{2\times n}(\mathbb{C}) \to (t^n)^*$ at whose value μ we will perform symplectic reduction; it is obtained by throwing away the off-diagonal entries of $\Phi_n(M) = M^*M$.

$$\begin{aligned}(\mathcal{O}_{\Lambda_2})^\mu &\cong \big(U(2)\underset{p\mathbf{1}}{\backslash\backslash} M_{2\times n}(\mathbb{C})\big)\underset{\mu}{//} T^n \\ &\cong U(2)\underset{p\mathbf{1}}{\backslash\backslash} M_{2\times n}(\mathbb{C})\underset{\mu}{//} T^n \\ &\cong U(2)\underset{p\mathbf{1}}{\backslash\backslash}\big(M_{2\times n}(\mathbb{C})\underset{\mu}{//} T^n\big) \\ &\cong SO(3)\backslash\backslash\big(M_{2\times n}(\mathbb{C})\underset{\mu}{//} T^n\big)\end{aligned}$$

(The actions of $U(n)$ and T^n coincide on the scalars, so to get a nonempty quotient, μ and p must be related in a certain way, spelled out below. Also, once we've reduced by T^n, there only remains an effective action of $PU(2) \cong SO(3)$.)

We can now make the further identification

$$(\mathcal{O}_{\Lambda_2})^{\mu} \cong SO(3)\backslash\backslash \prod_{i=1}^{n} (S^2)_{\mu_i}$$

where $(S^2)_{\mu_i}$ is the coadjoint orbit of $SO(3)$ with radius μ_i. (Careful: its area with the Lie-Kirillov-Kostant-Souriau symplectic form is proportional to μ_i, not to μ_i^2.)

The moment map for $SO(3)$'s action on a coadjoint orbit $(S^2)_{\mu_i}$ is simply inclusion into $so(3)^* \cong \mathbb{R}^3$. The moment map for the diagonal action on the product is the sum of the individual moment maps. We can regard $\prod_{i=1}^{n}(S^2)_{\mu_i}$ as the space of n-step polygonal paths in $\mathbb{R}^3$, starting at the origin, with the ith step being of length μ_i. In this framework the moment map takes a polygonal path to its endpoint.

The symplectic reduction of this space of paths is the zero level set (paths that end at the origin, which is to say, polygons) modulo rotation: the *moduli space of* μ*-polygons in* $\mathbb{R}^3$.

(Now we can say more about the subtlety passed over above. The μ_i are the lengths of the sides in the polygon space, whereas p is the semiperimeter. To get a nonempty quotient one must scale the μ_i so that $\sum_i \mu_i = 2p$.)

This identification – the polygonal interpretation of the GIT quotient of $\prod_{i=1}^{n}(S^2)_{\mu_i}$ by $SO(3)$ – was first noted in [22]. In [17] and [21] are constructed polytopes with chambers governing the topological behavior of the polygon spaces as a function of μ. That these polytopes should be the moment polytopes for a torus action was an idea of Sylvain Cappell. The actual torus action, which turns out to be a symplectic analogue of the Gel'fand-Cetlin system, is one of the new results in [18].

II.2. The Gel'fand-Cetlin action.

We now introduce $d : (\mathcal{O}_{\Lambda_2})^{\mu} \to \mathbb{R}^n$, where d_i is the length of the diagonal connecting the ith vertex to the origin.

(Only $n-3$ of these functions are new, as $d_1 = \mu_1$, $d_{n-1} = \mu_n$, and $d_n = 0$.)

The function d is smooth where no d_i vanishes, that is to say the polygon does not return to the origin prematurely. We call such a polygon P *prodigal*. The set of prodigal polygons is open dense in $(\mathcal{O}_{\Lambda_2})^{\mu}$ with complement of codimension 3.

There is in [21] introduced an action of a torus T^{n-3} on prodigal polygons; the ith circle acts by rotating the section of the polygon formed by the first i edges about the ith diagonal. (When that diagonal is length zero, there is no well-defined axis about which to rotate, and indeed the action cannot be extended continuously over this subset.) This action plainly preserves the level sets of the functions d, but more is true:

Theorem ([21]). *On the subspace of prodigal polygons, the function* d *is a moment map for these "bending flows".*

One important consequence of this is that the torus action also preserves the symplectic structure. It does not, seemingly, preserve the Riemannian metric nor the complex structure (the codimension of the singular set is not even).

We extend the function d to all of $\mathcal{O}_{\Lambda_2}$, and from there lift to the Stiefel manifold S. It is also convenient to regard μ as a function on S, and both functions have simple matrix-theoretic interpretations.

For $(a,b) \in \mathbb{C}^{n\times 2}$, $i = 1,\dots,n$, introduce the truncated matrices $M_i = \begin{pmatrix} a_1 & b_1 \\ \vdots & \vdots \\ a_i & b_i \end{pmatrix}$ the first i rows of (a,b). Then the 2×2 matrix

$$M_i^* M_i = \sum_{j=1}^{i} \begin{pmatrix} |a_j|^2 & \bar{a}_j b_j \\ a_j \bar{b}_j & |b_j|^2 \end{pmatrix}$$

has the eigenvalues

$$\frac{1}{2}\left(\sum_{j=1}^{i} (|a_j|^2 + |b_j|^2) \pm \sqrt{\left(\sum_{j=1}^{i} (|a_j|^2 - |b_j|^2) \right)^2 + \left| \sum_{j=1}^{i} a_j \bar{b}_j \right|^2} \right).$$

These are calculable from μ and d, since

$$\mu((a,b)) = (\dots, |a_i|^2 + |b_i|^2, \dots)$$

and

$$d((a,b)) = \left(\dots, \sqrt{\left(\sum_{j=1}^{i} (|a_j|^2 - |b_j|^2) \right)^2 + \left| \sum_{j=1}^{i} a_j \bar{b}_j \right|^2}, \dots\right)$$

So $\sum_{j=1}^{i} \mu_j$ is the sum of the two eigenvalues of $M_i^* M_i$, whereas d_i is the difference. (Note that $\mu_1 = d_1$ as promised; $M_1^* M_1$'s lesser eigenvalue is 0.)

This (2×2)-matrix $M_i^* M_i$ has the same nonzero eigenvalues as the $i \times i$ matrix $M_i M_i^*$. The latter matrix is more relevant in that it is the upper left $i \times i$ submatrix of the $U(n)$ moment MM^*.

This family of Hamiltonians – the eigenvalues of the upper left submatrices – has been studied already in [26] and is called the classical Gel'fand-Cetlin system (our main reference is [14]). The linear relations established above between them and d, μ are summed up in the following

Theorem. *The bending flows on $(\mathcal{O}_{\wedge_2})^\mu$ are the residual torus action from the Gel'fand-Cetlin system on the Grassmannian.*

The Gel'fand-Cetlin action on the flag manifold has always been rather mysterious (at least to me); it is pleasant that in this case it has a natural geometric interpretation.

The Gel'fand-Cetlin functions $\{e_{ij}\}_{j\le i}$ (the jth eigenvalue of the upper left $i \times i$ submatrix) satisfy some linear inequalities that can be established using the minimax description of eigenvalues [8].

$$e_{i,j} \le e_{i-1,j+1} \le e_{i,j+1}$$

For the polygon space functions l, d most of these say $0 \le 0$; for each $i = 0,\dots,n-1$ the nontrivial inequalities are

$$0 \le -d_i + \sum_{\iota=1}^{i} \mu_\iota \le -d_{i+1} + \sum_{\iota=1}^{i+1} \mu_\iota \le d_i + \sum_{\iota=1}^{i} \mu_\iota \le d_{i+1} + \sum_{\iota=1}^{i+1} \mu_\iota .$$

But these are transparent in our situation, as they are just the triangle inequalities!

$$\mu_{i+1} \le d_i + d_{i+1}$$
$$d_i \le \mu_{i+1} + d_{i+1}$$
$$d_{i+1} \le \mu_{i+1} + d_i$$

(The first one, $d_i \le \sum_{\iota=1}^{i} \mu_\iota$, can be proved inductively from the others starting from $d_0 = 0$.)

In [14] it is left as an exercise to show that the minimax inequalities are the *only* inequalities satisfied, equivalently, that every point in the convex polytope $\Gamma_m \subset \mathbb{R}^m \times \mathbb{R}^m$ defined by them (and $d_0 = d_m = 0$ and $\sum_i \mu_i = 2$) is realized by some Hermitian matrix. We show this directly:

Theorem. *The image of $\mathcal{O}_{\Lambda_2}$ under the map* (μ, d) *is the whole polytope* Γ_m.

Proof. We construct the polygons directly, vertex by vertex. We must place each new vertex on the intersection of two S^2's, one of radius d_{i+1} from the origin, the other of radius μ_{i+1} from the previous vertex. The inequalities $\mu_{i+1} \leq d_i + d_{i+1}$ and $d_{i+1} \leq \mu_{i+1} + d_i$ rule out one S^2 containing the other; the third inequality $d_i \leq \mu_{i+1} + d_{i+1}$ rules out their being separated balls. So they intersect in an S^1, a point or the whole S^2, anywhere on which we may place the new vertex. □

Remarks:

1. While the map μ is equivariant with respect to the usual action of S_n on $\mathbb{R}^n$, the map d can only be made equivariant under the involution $[i \leftrightarrow (n-i)]$, and the polytope Γ_n is correspondingly less symmetric than the S_n-invariant moment polytope of the 2-Grassmannian (the "2-hypersimplex" [11][10]).

2. The direct construction above actually produces planar polygons, if desired (the two S^1's intersect in an S^0, a point or the whole S^1). That the image of (μ, d) is the same when restricted to planar polygons has the flavor of a more general theorem of Duistermaat [5] on restricting moment maps to the fixed-point sets of antisymplectic involutions. In fact Duistermaat's theorem does not apply directly, because the subset where d is smooth (and thus a moment map) is noncompact.

3. The theorem guarantees that the bending torus acts simply transitively on the fiber over an interior point of Γ_m, making this fiber a torus $U(1)^{m-3}$. Over a prodigal boundary point of Γ_m, the fiber is still a product of 0- or 1-spheres, but fewer of them.

III.3. Moduli spaces of flat connections on a punctured sphere

The moduli space of μ-polygons in $\mathbb{R}^3$ has an obvious generalization to geodesic polygons in other Riemannian manifolds. (Referring to "moduli space" seems a bit disingenuous when not dealing with a homogeneous space, though.) In this section I look at geodesic μ-polygons in S^3's of varying radii R.

In the limit $R \to \infty$, the space of μ-polygons in S^3 converges in some senses (not made precise here) to the space of μ-polygons in $\mathbb{R}^3$.

Another interpretation of this space comes by observing that S^3 has a group structure, $SU(2)$. Each step in a geodesic polygon can thus be interpreted as a group element g_i, the closure property being $\prod g_i = 1$, and "mod out by rotation" being "mod out by conjugation". This is now recognizable as the moduli space of flat $SU(2)$ connections on a n-holed Riemann sphere, with the g_i being the holonomies around the punctures.

Most statements about the $\mathbb{R}^3$ polygons have evident generalizations to the moduli space of flat connections, with only a few surprises.

The functions d defined in the last section correspond to taking the holonomy of paths that go around several punctures (or rather, the arccosine of half the trace of the holonomy – the arccosine is what makes these functions only continuous and not globally smooth), as observed in [22].

Doing symplectic reduction with respect to a d_i, in the connection framework, corresponds to cutting the Riemann surface into two disjoint spheres, each with a new puncture and the same prescribed holonomy about it – consequently, the reduction is the product of two smaller moduli spaces of flat connections.

In the polygon space framework, it corresponds to fixing a diagonal length and forgetting the angle at which the two halves of the polygon are glued together along it. All that remains are the two new polygons with their new (equal) edges. So the reduction is the product of two smaller polygon spaces.

The greatest subtlety comes in generalizing the action of S_n on the zero weight varieties – these are polygon spaces where all the edges are the same length. (Note that in $\mathbb{R}^3$ it doesn't make much difference what length that edge is, except to rescale the metric by that factor.) Thinking of a polygon as a bunch of vectors adding up to zero, the action of S_n is to permute those vectors.

This doesn't work when the polygon is a bunch of noncommuting group elements multiplying to 1. Instead, one gets an action of the braid group, whose generators act as

$$(a, b) \mapsto (aba^{-1}, a)$$

preserving the product ab. In the limit that a and b are very small, and so nearly commute, the action degenerates to that of S_n.

This seemingly ad hoc action of the braid group is very natural in the flat connections on the n-punctured sphere picture; the braid group is the n-punctured sphere's mapping class group.

The fact that one must go from S_n to the braid group is not the only surprise; the other is that the action of the braid group on the moduli space is only symplectic and not complex or Riemannian.

Appendix 1. Homogenous coordinate rings

Let $E \to X$ be a complex line bundle over a complex algebraic variety – for example the total space of the tautological line bundle over projective space, $\tilde{\mathbb{C}}^n = \{(\vec{v}, l \in \mathbb{CP}^n : \vec{v} \in l\} \to \mathbb{CP}^n$. This means in particular that there is an action of $\mathbb{C}^\times$ on E, preserving and rescaling each fiber.

Usually we will also have a group G acting on E and X compatible with the projection and the linear structure – it commutes with the action of $\mathbb{C}^\times$. (It goes without saying that we want the action to preserve the complex structure.) The action of G on E is called a **lift of the action** on X: not every action of G on X has a lift (one may have to centrally extend the group), nor is it unique (one can globally "twist" by an action of the group on each line through a map $G \to GL_1(\mathbb{C}) \cong \mathbb{C}^\times$). For example, the action of $PSL_n(\mathbb{C})$ on $\mathbb{CP}^{n-1}$ does not have a lift; the group must be replaced by $SL_n(\mathbb{C})$. But since $SL_n(\mathbb{C})$ has no nontrivial homomorphisms to $\mathbb{C}^\times$, the lift, the usual action of $SL_n(\mathbb{C})$ on $\tilde{\mathbb{C}}^n$, is unique.

Let R be the ring of algebraic functions on the total space E. We now ask that these separate the points of $E - X$ (and for technical reasons, that they "separate tangent vectors" – for any two tangent vectors at the same point, some function should have different directional derivatives). This is equivalent to asking that this line bundle (or rather, its dual) be "ample". Since $G \times \mathbb{C}^\times$ acts on E, it acts on R by ring automorphisms.

Decomposing R into weight spaces R_n under $\mathbb{C}^\times$ makes R a **graded ring** (and each R_n is automatically a representation of G). The functions living in the R_n piece are those that grow like z^n along each fiber of the projection $E \to X$. Two things of note: since our fibers are actually $\mathbb{C}$'s not just $\mathbb{C}^\times$'s, the functions are not allowed to have poles at the zero section; $R_n = 0$ for $n < 0$. Second, the functions in R_0 are constant along the fibers, so determined by their values on the zero section; R_0 is the functions on X. In the case X compact, this can only be the constants, so $R_0 = \mathbb{C}$.

Notes. 1. Affine varieties are those with ample trivial bundles; even on a trivial bundle, though, we may have different choices of lift.

2. Functions on a line bundle scaling linearly along the fibers – that is, R_1 – give and are given by sections of the dual line bundle. For example, the tautological line bundle over projective space PV, being the blowup of V at the origin, has many functions linear on the fibers, all given by elements of V^*: this means that the space of sections of the dual line bundle is also V^*. This is the more common picture, with the unfortunate effect that one must construct the dual bundle to the more easily constructed tautological bundle over projective space. I will hide this problem by referring to the "usual" bundle over projective space, which to those who want functions on their line bundles will mean the tautological one, and to those who want sections will mean its dual.

3. Each weight n function on the kth tensor power of E gives a weight nk function on E itself, by composing with $\vec{v} \in C \mapsto \vec{v} \otimes \ldots \otimes \vec{v}$. In the R picture, this corresponds to throwing out each R_m for which k doesn't divide m, then contracting the grading by the factor k. This trick is called the **kth Veronese embedding**. Putting this together with note 2, we can interpret R_k as sections of the kth tensor power of the dual line bundle.

4. The effect of asking that R consist of algebraic, rather than arbitrary analytic, functions on E is that R is actually equal to the direct sum $\bigoplus_{n \in N} R_n$ rather than merely containing it.

5. We can reverse this process, and obtain a space with a line bundle from some $\mathbb{N}$-graded rings R. First construct the affine variety $\operatorname{Spec} R$. Rip out the points fixed by the $\mathbb{C}^\times$ action, corresponding to the ideal $\bigoplus_{n>0} R_n$. In the good case $\mathbb{C}^\times$ will act freely on what is left; our space X is then the quotient $(\operatorname{Spec} R - \operatorname{Spec} R_0)/\mathbb{C}^\times$, born with a $\mathbb{C}^\times$ bundle over it to which to associate a line bundle; this quotient is called $\operatorname{Proj} R$. (In the bad cases $\operatorname{Proj} R$ is born with a sheaf that is not a line bundle.) If we do this

with the polynomial ring $R = \mathrm{Sym}^\bullet(V)$, we get $\mathrm{Spec}\, R = V^*$, and $(\mathrm{Spec}\, R - \{\vec{0}\})/\mathbb{C}^\times = P(V^*) \cong P^*(V)$, the hyperplanes in V.

6. In the case of affine varieties there is a perfect correspondence between maps $\mathrm{Spec}\, R \to \mathrm{Spec}\, S$ and $S \to R$. There is no such correspondence for Proj, due to the removal of the zero section from the Spec. The following example illustrates the characteristic failure. Given a linear map $V \to W$, we have a ring homomorphism $\mathrm{Sym}^\bullet(V) \to \mathrm{Sym}^\bullet(W)$, and thus a map of varieties $W^* \to V^*$. However, many nonzero elements of W^* may go to $\vec{0}$ in V^*, so the map does not descend to $W^* - \{\vec{0}\} \to V^* - \{\vec{0}\}$, nor to a map $P^*(W) \to P^*(V)$ on the quotients. The map is well-defined away from those hyperplanes in W containing the image of V, and in general, homogeneous maps $S \to R$ between graded rings will give only rational maps $\mathrm{Proj}\, R \to \mathrm{Proj}\, S$, blowing up at those points where every pulled-back nonconstant homogeneous section vanishes.

7. In the case that $R_0 = \mathbb{C}$ and R_1 generates R, we can realize $\mathrm{Proj}\, R$ as a subvariety of $P^*(R_1)$. Write $R = \mathrm{Sym}^\bullet(R_1)/(r_1, \ldots, r_m)$ where the r_m are homogeneous relations generating the ideal (which you can find in the kernels of the maps $\mathrm{mult} : \mathrm{Sym}^j(R_1) \twoheadrightarrow R_j$). $\mathrm{Proj}\, R$ is the the common vanishing locus on $P^*(R_1)$ of these polynomials r_i in the homogeneous coordinates. Conversely, every subvariety of projective space arises in this way, and thus comes with a line bundle; this is the restriction/pullback of the one on projective space itself. (This was a case where the ring map $\mathrm{Sym}^\bullet(R_1) \to R$ does induce an actual map of Proj's.)

More generally we pick generators for R that don't all live in degree 1, and we realize Proj R as a subvariety of a product of an affine space and a weighted projective space.

Appendix 2. Geometric invariant theory and symplectic reduction

In the language of Appendix 1, let $X = \mathrm{Proj}\, R$, where R is an $\mathbb{N}$-graded ring carrying an action of a group G. Inside R there is a subring R^G of G-invariants; what is its Proj? The inclusion $R^G \to R$ induces a rational epimorphism of X to $\mathrm{Proj}\, R^G =: X//G$, called the **geometric invariant theory** (or **GIT**) **quotient** of X by G. Unlike the space X/G, it depends on the choice of G-line bundle over X.

The **unstable set** X^{us} of X is defined as the place where this rational map is undefined; its complement is called the **semistable set** X^{ss}. By note 6, X^{us} is the subset over which every nonconstant homogeneous invariant function on E vanishes. In good cases $X//G = X^{ss}/G$; More generally one has a map from $X^{ss} \to X//G$ whose fibers are G-stable but not necessarily G-orbits. The elements of such fibers are called **properly semistable** (and not **stable**).

Example: let X be a vector space with basis $\{x^1, \ldots, x^d\}$, $G = \mathbb{C}^\times$. Then

$$R = \mathbb{C}[x_1^{(0)}, \ldots, x_d^{(0)}, y^{(1)}] = \bigoplus_{n \in N} y^n \mathrm{Sym}^\bullet(X^*),$$

where $\{x_i\}$ is the dual basis, and the superscript indicates the degree. The lifts of G to the line bundle Spec R are indexed by the weight w chosen for the element y – each of the $\{x_i\}$ has weight 1. The ring of invariants $R^G = \bigoplus_{n \in N} y^n \mathrm{Sym}^{-nw}(X^*)$ is the $-w$th Veronese of $\mathbb{C}[x_1 z, x_2 z, \ldots, \ldots x_n z]$, where $y = z^w$, for $w < 0$. The unstable set is $\vec{0} \in X$, and the GIT quotient is indeed $\mathrm{Proj}\, \mathbb{C}[x_1 z, x_2 z, \ldots, \ldots x_n z] = PX = X - \vec{0}/G$, carrying the $-w$th tensor power of the usual line bundle on projective space. For $w < 0$, the invariants are only $\mathbb{C}$, the whole space is unstable, and the quotient is empty ($\mathrm{Proj}\, \mathbb{C}$). For $w = 0$, the invariants are $\mathbb{C}[y]$, the whole space is semistable, and the quotient is a point ($\mathrm{Proj}\, \mathbb{C}[y]$).

N.B. While points with positive-dimensional stabilizer are never stable, the converse is not true – points with discrete stabilizer may still be unstable in some polarization. For example take the antidiagonal action of $\mathbb{C}^\times$ on $\mathbb{C}^2$. In the trivial polarization (l weight zero), the axes are properly semistable and all fall together in a point. For polarizations in which the weight on the line bundle has the same sign as the weight on the x-coordinate, the x-axis (including the origin) is unstable, but the punctured y-axis is stable. In the remaining polarizations (switch x and y), the y-axis is unstable, and the punctured x-axis is stable.

The easiest way to calculate $X//G$ is to find a section of the map $X^{ss} \to X//G$. These will only very rarely exist. Less ambitiously, we can hope to find a subset of X^{ss} preserved by some smaller group K, such that each fiber in X^{ss} contains a unique K-orbit. (If X^{ss} is a G-bundle over $X//G$, this amounts to reducing its structure group.) In the case that G is reductive and X Kähler (for instance, projective), this is possible for G's maximal compact subgroup K (not surprising when X^{ss} is a G-bundle, since every G-bundle is reducible to the maximal compact). It is picked out by a moment map $\mu : X \to k^*$ for K's action on X and E.

Definition. *Endow a symplectic manifold X with symplectic form ω with the usual Poisson bracket $\{f, g\} = \omega^{-1} \cdot df \wedge dg$ on $C^\infty(M)$, a Lie algebra structure. A map $\mu : X \to k^*$ is a* **moment map** *for K's action if the induced comoment map $\mu^T : k \to C^\infty(M)$ is a Lie algebra homomorphism.*

Three not completely obvious, but very standard, facts about moment maps [24]: they uniquely determine the action of K, can only exist if K acts preserving the symplectic structure, and do exist uniquely if K is semisimple.

Theorem ((Kirwan, Kempf/Ness, Guillemin/Sternberg) [24]). *The zero level set of μ has a unique K-orbit in each fiber of $X^{ss} \to X//G$;*

$$X//G \cong \mu^{-1}(0)/K.$$

Several things about this theorem are remarkable: while the fibers of $X^{ss} \to X//G$ are not always G-orbits, the fibers of $\mu^{-1}(0) \to X//G$ are always K-orbits. It is rather surprising from some viewpoints that the moment map is defined on all of X rather than just X^{ss}. It is also remarkable that the quotient of the noncomplex manifold $\mu^{-1}(0)$ by the noncomplex group K magically regains a complex structure. (Though built into this was the fact that K was preserving X's complex structure.)

To actually apply the theorem, we need to be able to calculate the moment maps for K acting on X, E. But the moment maps for linear groups acting on affine times projective space are well-known –

$$\begin{array}{cc} \mathbb{C}^m \to u(m)^* & \mathbb{CP}^{n-1} \to u(n)^* \\ v \mapsto v^*v & [v] \mapsto \dfrac{v^*v}{|v|^2} \end{array}$$

and we can compose $X \to \mathbb{C}^m \times \mathbb{CP}^{n-1} \to u(m)^* \times u(n)^* \to k^*$ to get the moment map for linear actions preserving closed subvarieties of $\mathbb{C}^m \times \mathbb{CP}^{n-1}$.

Appendix 3. The shifting trick, and G-invariants vs. B-invariants

Let $\lambda : G \to \mathbb{C}^\times$ be a 1-D rep of G (thus $G/[G,G]$, and infinitesimally given by an element of $(\mathfrak{g}/[\mathfrak{g},\mathfrak{g}])^*$). Let G act on $X = \operatorname{Proj} R$ and on R. There are several ways to think of the other lifts of G to the line bundle over X.

The first is to consider the ring of functions not *invariant* under G, but transforming under G like λ:

$$R^{G,\lambda} = \bigoplus_{n \in \mathbb{N}} (R_n)^{G, n\lambda}.$$

But we can also calculate the subspace $(R_n)^{G,n\lambda}$ as actual invariants, in $R_n \otimes \mathbb{C}_{-n\lambda}$. All together, this is

$$R^{G,\lambda} = (R \otimes \operatorname{Sym}^\bullet(\mathbb{C}_{-\lambda}))^G$$

which geometrically looks like $(X \times pt)//G$, where the innocuous-looking point carries a nontrivial action of G on its line. This is called the **shifting trick**.

It is worth noting the semidirect way that shifting behaves under Veronese embedding: the kth Veronese of the λ-shift is the $k\lambda$-shift of the kth Veronese. Since for most purposes one is willing to finesse the differences between a variety and its kth Veronese, this allows one to shift by *rational* λ: just Veronese enough until the desired shift becomes integral. In this way one can refer to generic values of λ, like "For generic values of λ, there are no properly semistable points in X^{ss}."

Honestly speaking, one should use a subscript under the // to indicate which lift is being chosen. In most cases there is an obvious choice (in particular when G is semisimple and the lift is thus unique); it is customary to leave off the subscript in this case, and otherwise label it by the difference λ.

For semisimple G shifting is not an issue, as G has no 1-d reps. However if we are willing to give up semisimplicity we can use B instead of G and gain new lifts, since B can have nontrivial characters ($B/[B,B] \cong T$).

The subspace V^N invariant under a maximal unipotent subgroup N are the same as those annihilated by a maximal nilpotent subalgebra $\mathfrak{n}$ – the high weight vectors. So the invariants under $B = TN$ are those high weight vectors of zero weight, which are the same as the G-invariants: $V^B = V^G$ and $X//B = X//G$, as long as the B-lift used is the one that extends to an action of G.

For other choices of lift, though, we will get things that are not G-quotients. It would be nice to have a version that only refers to the semisimple group G, though, since the symplectic theorems are only applicable in this case. The heuristic answer is to decide "$X/B = (X \times B\backslash G)/G$" and then make sense of $B\backslash G$ once and for all. The following analysis will make that precise.

What we need is a way of picking out the high weight vectors in V with weight λ. This can be accomplished as in the 1-d case with $(V \otimes (V_\lambda)^*)^G$. The ring version is

$$R^{G,\lambda} = (R \otimes \bigoplus_{n \in N} (V_{n\lambda})^*)^G,$$

or geometrically, $X \underset{\lambda}{//} G = (X \times \text{Proj} \bigoplus_{n \in N} (V_{n\lambda})^*)//G$. This Proj turns out to be $B\backslash G$ (by the Borel-Weil theorem).

More concretely, this ring $\bigoplus_{n \in N} (V_{n\lambda})^*$ is generated by its degree 1 part V_λ^*, so we expect to locate $B\backslash G$ in $P(V_\lambda)$ – and indeed, it's the orbit of the line of highest weight vectors. In fact this variety is sometimes only $P\backslash G$, where P is a subgroup properly containing B; P is in fact the largest sugroup to which the action on the line bundle can be extended.

Appendix 4. A commutative ring/algebraic geometry/symplectic geometry Rosetta Stone

I must say at the outset that the correspondences here are not always exact: each theory is richer than its intersection with the other two. In particular, not every symplectic manifold can be given a compatible complex structure, nor can every complex manifold be given a compatible symplectic structure. (Indeed, there exist manifolds that can be given each structure individually, but not compatibly.)

What is remarkable is the number of pairs of constructions that only use the data from one of the viewpoints and yet produce the same answer. That's what this table is about.

commutative algebra	*algebraic geometry*	*symplectic geometry*
	a variety X	a manifold X
a graded ring R with $\mathrm{Proj}\, R = X$	a choice of ample line bundle $\mathcal{L}$	an (integral) symplectic structure ω
	a holomorphic G-action	a symplectic K-action ($G = K^{\mathbb{C}}$)
a graded ring G-action	a lift to the line bundle	a moment map
R^G	the GIT quotient	the symplectic quotient
R_1	sections of the line bundle	"the quantization"
degree of $H(i) = \dim R_i$	complex dimension	half real dimension
leading coefficient of $H(i)$	"multiplicity"	symplectic volume
replacing each R_i by R_{ki}	replacing $\mathcal{L}$ by $\mathcal{L}^{\otimes k}$ ("Veronese")	replacing ω by $k\omega$
decompose into G-isotypic components		the Gel'fand-Cetlin functions $X \to \mathfrak{t}^*_+$

References

[1] S. Ariki, On the representations of the Weyl groups of type D corresponding to the zero weights of the representations of $SO(2n, \mathbb{C})$. The Arcata Conference on Representations of Finite Groups, in Proc. Sympos. Pure Math. 47, (1986), 327-342.

[2] S. Ariki, Representations of Weyl groups on zero weight spaces of $\mathfrak{g}$-modules. Algebraic and topological theories (1986) 546-568

[3] M. Atiyah, Convexity and commuting Hamiltonians. Bull. London Math. Soc. 14 (1982) 1-15

[4] M. Atiyah & G. Segal, The index of elliptic operators II. Ann. Math. 87 (1968)

[5] J.J. Duistermaat, Convexity and tightness for restrictions of Hamiltonian functions to fixed-point sets of an antisymplectic involution. Trans. AMS 275 (1983) 417–429

[6] J.J. Duistermaat, The heat kernel Lefschetz fixed point formula for the spin-c Dirac operator. Progress in Nonlinear Differential Equations and their Applications, 18. Birkhauser (1996)

[7] J.J. Duistermaat & G. Heckman, On the variation in the cohomology of the symplectic form of the reduced phase space. Invent. Math. 69 (1982) 259-268

[8] J. Franklin, Matrix theory. Prentice–Hall, Inc. (1968)

[9] D. Gay, Characters of the Weyl group of $SU(n)$ on zero weight spaces and centralizers of permutation representations. Rocky Mountain J. Math. 6 (1976) 449–455

[10] I. Gel'fand, M. Goresky, R. MacPherson & V. Serganova, Combinatorial geometries, convex polyedra and Schubert cells. Adv. Math. 63 (1987) 301–316

[11] I. Gel'fand & R. MacPherson, Geometry in Grassmannians and a generalization of the dilogarithm. Adv. Math. 44 (1982) 279–312

[12] V. Guillemin, E. Lerman, & S. Sternberg, Symplectic fibrations and multiplicity diagrams. Cambridge University Press (1996)

[13] V. Guillemin & S. Sternberg, Convexity properties of the moment mapping. II Invent. Math. 3 (1984) 533-546

[14] V. Guillemin & S. Sternberg, The Gelfand-Cetlin system and quantization of the complex flag manifolds. J-Funct-Anal 52(1) (1983) 106-128

[15] V. Guillemin & S. Sternberg, Birational equivalence in the symplectic category. Inventiones Math. 97 (1989) 485–522

[16] E. Gutkin, Representations of the Weyl group in the space of vectors of zero weight. Uspehi Mat. Nauk 28 (1973) 237-238

[17] J.-C. Hausmann, Sur la topologie des bras articulés. In "Algebraic Topology, Poznan", Springer Lectures Notes 1474 (1989) 146–159

[18] J.-C. Hausmann & A. Knutson, Polygon spaces and Grassmannians. Submitted to J. Diff. Geom

[19] J.E. Humphreys, Reflection groups and Coxeter groups. Cambridge studies in advanced mathematics (1990)

[20] J.C. Jantsen, Representations of algebraic groups. Academic Press (1987)

[21] M. Kapovich & J. Millson, On the moduli space of polygons in the Euclidean plane. J. Diff. Geom. 42 (1995) 430-464

[22] M. Kapovich & J. Millson, The symplectic geometry of polygons in Euclidean space. Preprint (Dec. 1994)

[23] B. Kostant, On convexity, the Weyl group, and the Iwasawa decomposition. Ann. Sci. Ec. Norm. Sup. 6 (1973) 413-455

[24] D. Mumford, J. Fogarty, & F. Kirwan, Geometric Invariant Theory Springer-Verlag (1994)

[25] S. Plakson, [illegible]. (1987)

[26] A. Thimm, Integrable geodesic flows on homogeneous spaces. Ergodic Theory and Dyn. Sys. 1 (1980) 495–517