

THE SYMPLECTIC AND ALGEBRAIC GEOMETRY OF HORN'S PROBLEM

ALLEN KNUTSON

ABSTRACT. One version of Horn's problem asks for which λ, μ, ν does $H_\lambda + H_\mu + H_\nu = 0$ have solutions, where $H_{\lambda, \mu, \nu}$ are Hermitian matrices with spectra λ, μ, ν . This turns out to be a *moment map* condition in Hamiltonian geometry. Many of the results around Horn's problem proven with great effort "by hand" are in fact simple consequences of the modern machinery of symplectic geometry, and the subtler ones provable via the connection to geometric invariant theory. We give an overview of this theory (which was not available to Horn), including all definitions, and how it can be used in linear algebra.

1. INTRODUCTION

This is an expository paper on the symplectic and algebraic geometry implicit in *Horn's problem*, which asks the possible spectra of a sum of two Hermitian matrices each with known spectrum.

The connection with symplectic geometry is very straightforward: the map $(H_\lambda, H_\mu) \mapsto H_\lambda + H_\mu$ that takes a pair of Hermitian matrices with known spectra to their sum is a *moment map* for the diagonal conjugation action of $U(n)$ on a certain *symplectic manifold* (definitions to follow). This is a very restrictive property of maps, and many things can be proved about them. The proofs, at heart, are not really any different than the techniques Horn himself used to study this map. Nonetheless the framework is worth understanding in order to recognize what other linear algebra problems are likely to have answers as nice as the ones to Horn's problem. In particular the Schur-Horn theorem follows very easily from the general theorems in this area (and was a primary inspiration for them).

Some of the more esoteric connections – to algebraic geometry and representation theory – can also be seen in this context, via the Kirwan/Ness theorem (which we will also state). Again, the basic techniques used are the same, but in the Kirwan/Ness theorem one sees these techniques pushed to prove the statements in what appears to be their proper generality.

Along the way we explain the relation between Hermitian matrices, flag manifolds, and the Borel-Weil-Bott-Kostant theorem.

2. THE SCHUR-HORN THEOREM, HORN'S PROBLEM, AND HAMILTONIAN MANIFOLDS

Let $\lambda = (\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n)$ be a weakly decreasing list of real numbers, which we'll use to encode the eigenvalue spectrum of a Hermitian matrix. The Schur-Horn theorem states the following:

Date: February 1, 2008.

Theorem. Let \mathcal{O}_λ be the space of Hermitian matrices with spectrum λ . Let $\Phi : \mathcal{O}_\lambda \rightarrow \mathbb{R}^n$ take a matrix to its diagonal entries. Then the image of Φ is a convex polytope, whose vertices are the $n!$ permutations of λ .

This theorem was interpreted by Kostant in 1970 as the $U(n)$ case of a theorem for arbitrary compact Lie groups, leading the way to a much wider generalization found in 1982 by Atiyah and independently by Guillemin and Sternberg:

Theorem. Let M be a compact connected symplectic manifold, with an action of a torus T . Let $\Phi : M \rightarrow \mathfrak{t}^*$ be a moment map for this action. Then the image of Φ is a convex polytope, the convex hull of the images of the T -fixed points on M .

Of course, to see how to cast the Schur-Horn theorem in this formulation, we'll need to define "symplectic manifold" and "moment map". We will make no attempt to be encyclopædic in our references and instead direct the reader to

- [GLS] – for all matters symplectic or convex
- [F] – for the algebraic geometry of flag manifolds
- [MFK] – for geometric invariant theory
- chapter 8 of [MFK] – for the Kirwan/Ness theorem

and references therein.

2.1. Symplectic manifolds. Let M be a manifold, and ω an *antisymmetric* inner product on the tangent spaces to M , sort of a skew Riemannian metric. Since the inner product of any vector with itself is zero, we can't talk about positive definiteness, and so we instead ask for nondegeneracy – that for any tangent vector \vec{v}_1 , there exists another vector \vec{v}_2 with $\omega(\vec{v}_1, \vec{v}_2)$ nonzero. Surprisingly, this forces M to be even dimensional.

There is a standard example: \mathbb{R}^{2d} with basis $\{\vec{x}_1, \dots, \vec{x}_d, \vec{y}_1, \dots, \vec{y}_d\}$, where $\omega(\vec{x}_i, \vec{x}_j) = \omega(\vec{y}_i, \vec{y}_j) = 0$, but $\omega(\vec{x}_i, \vec{y}_j) = -\omega(\vec{y}_j, \vec{x}_i) = \delta_{ij}$ (Kronecker delta).

If (M, ω) is locally isomorphic to \mathbb{R}^{2d} with its standard ω , we say that M is a **symplectic manifold**, and call ω its **symplectic form**.¹ This is roughly analogous to studying Riemannian manifolds that are locally Euclidean.

As with a Riemannian metric, we can talk about the **symplectic gradient** X_f of a function f , also called the Hamiltonian vector field X_f associated to a Hamiltonian f . It is defined uniquely by the equation

$$D_{\vec{v}(m)}f = \omega(\vec{v}(m), X_f(m))$$

where $\vec{v}(m), X_f(m)$ are tangent vectors to the point $m \in M$, and $D_{\vec{v}(m)}f$ is the directional derivative of f in the direction $\vec{v}(m)$. That this defines X_f uniquely follows from ω 's nondegeneracy.

It is easy to show that symplectic gradients have an immense advantage over Riemannian gradients: the derivative of ω along X_f vanishes. In integral form this says that the time t flow map from M to M , given by following X_f for time t , takes ω to itself.

Example 1. Let $M = S^2$, and ω be the area form, taking a pair of vectors to the oriented area of the parallelogram they define. Let f be the height function on M , normalized to take the North pole to 1 and the South pole to -1 . Then the Riemannian gradient points

¹The usual definition is in terms of de Rham forms, and has this formulation as Darboux's theorem.

down everywhere along longitude lines, as in figure 1, whereas the symplectic gradient points sideways everywhere along latitude lines – and generates rotations, which are of course area-preserving.




FIGURE 1. The Riemannian gradient vs. the symplectic gradient of the height function on S^2 .

This gives us a handy way to describe (certain) flows on a symplectic manifold – in terms of functions, which opens up the world of e.g. Morse theory. One example of this connection: a critical point of a function f is one where all the directional derivatives are zero, which (by ω 's nondegeneracy) is equivalent to the vanishing there of X_f . So the critical points of f are exactly the fixed points of the flow generated by its symplectic gradient X_f .

2.2. Moment maps. Now let M be a symplectic manifold with symplectic form ω , and an action of a connected Lie group K ; for us K will always be the circle group S^1 , a product T of circle groups, or the unitary group $U(n)$.

If we assume K acts smoothly on M , the elements of K nearby the identity give us diffeomorphisms of M very close to the identity; differentiating this picture, we find that each tangent vector to the identity of K gives a vector field on M . Denote the tangent space to the identity, also known as the Lie algebra of K , by \mathfrak{k} , and its dual by \mathfrak{k}^* . (This funny letter is a Fraktur k .)

Since K acts on itself by conjugation, fixing the identity, it acts on \mathfrak{k} ; this is called the **adjoint representation**, and the induced action on \mathfrak{k}^* the **coadjoint representation**.

We say that a map $\Phi : M \rightarrow \mathfrak{k}^*$ is a **moment map** for the action of K on M if

1. Φ is equivariant, i.e. $\forall k \in K, m \in M, \Phi(k \cdot m) = k\Phi(m)k^{-1}$
2. for each $\vec{k} \in \mathfrak{k}$, the vector field induced on M by \vec{k} equals the symplectic gradient of $\langle \vec{k}, \Phi \rangle$. (We are using here the natural pairing between \mathfrak{k} and \mathfrak{k}^* .)

Note that for the action to have a moment map at all, the action of K on M must preserve the symplectic structure. This is almost sufficient (but not quite); in any case all the actions considered in this paper will have moment maps. If the action has a moment map it is said to be **Hamiltonian** and M is called a **Hamiltonian K -manifold**.

Since K is connected, we can recover its action from the action of the Lie algebra, which in turn we can get from the moment map; in particular the moment map uniquely determines the action. (The action does not determine the moment map uniquely even when one exists, but only up to certain translations.) And while any map $\Phi : M \rightarrow \mathfrak{k}^*$ will give us associated vector fields $X_{\langle \vec{k}, \Phi \rangle}$, very few maps will give actions of K ; moment maps are very special.

We need four particularly important facts about moment maps. (All follow easily from the definition.)

1. Let K act on M with $\Phi_K : M \rightarrow \mathfrak{k}^*$ a moment map for the action on M , and $\rho : H \rightarrow K$ be a Lie group homomorphism, making H act on M too. There is a corresponding map of Lie algebras $d\rho : \mathfrak{h} \rightarrow \mathfrak{k}$, and a dual map $(d\rho)^* : \mathfrak{k}^* \rightarrow \mathfrak{h}^*$. Then the action of H is also Hamiltonian, with moment map given by composing Φ_K and $(d\rho)^*$. The case of greatest interest is when H is a subgroup of K , and ρ is inclusion.
2. Let M be a **coadjoint orbit**, i.e. an orbit of K 's action on \mathfrak{k}^* . Then there exists uniquely a symplectic structure on M such that the inclusion map $M \hookrightarrow \mathfrak{k}^*$ is a moment map. (This was called the Kirillov-Kostant-Souriau symplectic structure until Alan Weinstein found it in 19th-century notebooks of Lie.)
3. Let M be a Hamiltonian K -manifold, and N a Hamiltonian H -manifold. Then the natural action of $K \times H$ on $M \times N$ is Hamiltonian, with moment map the direct sum of the two individual moment maps.
4. Let M be a Hamiltonian K -manifold, and N a submanifold of M invariant under K , such that the restriction of the symplectic form on M is nondegenerate on N (making N a symplectic manifold in a natural way). Then the action of K on N is also Hamiltonian, with moment map the composition of inclusion with M 's own moment map.

Since the image of a moment map is K -invariant, it must be a union of coadjoint orbits; in this way one can regard the individual coadjoint orbits as sort of "minimal" Hamiltonian K -manifolds. (This is only interesting for nonabelian groups K , insofar as the coadjoint action is trivial for abelian groups, with the coadjoint orbits just points.)

Example 2. Let $M = \mathbb{R}^2$ with the standard symplectic form, and $K = S^1$ acting on M by rotation. Then we can identify K 's tangent space with \mathbb{R} , such that the moment map is $\Phi(\vec{v}) = |\vec{v}|^2$. The image of Φ is the positive half-line.

Example 3 (using property 3 above of moment maps). Let $M = \mathbb{R}^{2d}$ with the standard symplectic form, and $K = (S^1)^d$, each circle acting on a pair of the coordinates – really the n th power of the previous example. Then we can identify K 's tangent space with \mathbb{R}^d . The image of Φ is the positive "orthant" (the d -dimensional generalization of the first quadrant in \mathbb{R}^2 , octant in \mathbb{R}^3 , etc.).

This second example gives a hint of the source of convex polytopes in the Atiyah/Guillemin-Sternberg theorem stated above; since we've defined a symplectic manifold as one that looks locally like \mathbb{R}^{2d} with the standard form, we just need to prove a "T-equivariant version" of that local normal form, in order to know that at least locally the image of the moment map is a polytope. The actual proof of A/G-S requires some more ingredients (not explored here) to show the global statement.

Example 4. Let $M = \mathbb{C}\mathbb{P}^n = \{[z_0, z_1, \dots, z_n] : z_i \text{ not all } 0\}/\mathbb{C}^\times$, complex projective n -space. This has a natural symplectic form (the Fubini-Study form) we do not pause to write down, and a Hamiltonian action of T^{n+1} , whose Lie algebra we can identify with \mathbb{R}^{n+1} . One moment map is then

$$\Phi([z_0, z_1, \dots, z_n]) = \left(\frac{|z_0|^2}{\sum_i |z_i|^2}, \frac{|z_1|^2}{\sum_i |z_i|^2}, \dots, \frac{|z_n|^2}{\sum_i |z_i|^2} \right).$$

The individual coordinates can only vary between 0 and 1, and their sum is automatically 1; it is easy to check that the image is the whole simplex. In fact this is really the $\lambda = (1, 0, 0, \dots, 0)$ case of the Schur-Horn result.

2.3. The relation with the Schur-Horn theorem. Let K be the unitary group $U(n)$, and T the subgroup of diagonal matrices, a “maximal torus” of $U(n)$. The Lie algebra $\mathfrak{u}(n)$ is the space of skew-Hermitian matrices.² We can identify the Hermitian matrices with $\mathfrak{u}(n)^*$ using the trace form, $H \mapsto \text{Tr}(iH \cdot)$. This identification is $U(n)$ -equivariant, intertwining the conjugation action with the coadjoint action; in particular it takes orbits of Hermitian matrices to coadjoint orbits.

The upshot is that we can think of coadjoint orbits of $U(n)$ as orbits of $U(n)$ acting on the space of Hermitian matrices by conjugation, which of course are the isospectral sets \mathcal{O}_λ required for the Schur-Horn theorem.

From this point, we just have to apply properties 1 and 2 of moment maps; \mathcal{O}_λ is a symplectic manifold, and the action of the diagonal subgroup T is Hamiltonian with moment map $\mathcal{O}_\lambda \hookrightarrow \mathfrak{u}(n)^* \rightarrow \mathfrak{t}^*$. Using our trace form identification, this composite is the map taking a Hermitian matrix with spectrum λ to its diagonal entries. Then the A/G-S convexity theorem tells us that the image is the convex hull of images of T -fixed points. A Hermitian matrix is fixed under conjugation by all diagonal matrices if and only if it itself is diagonal, which means its entries are a permutation of λ . This proves Schur-Horn.

(Of course this presentation is historically unfair, as indicated at the beginning – the Schur-Horn theorem was a primary inspiration for A/G-S convexity.)

2.4. Stabilizer groups. We mention one of the many easy properties of the definition of moment map. The **stabilizer algebra** of a point m for the action of a Lie algebra \mathfrak{k} is defined as the Lie subalgebra giving vector fields that vanish at m . (Correspondingly, the flows they generate leave m fixed, so the stabilizer algebra is the Lie algebra of the stabilizer group of m .)

Proposition. *Let M be a Hamiltonian K -manifold, with moment map Φ , and $m \in M$. Then the stabilizer algebra of m is the perp of the image of the differential $d\Phi$ mapping from the tangent space at m to \mathfrak{k}^* .*

(Again we’re using the natural pairing between \mathfrak{k} and \mathfrak{k}^* ; the perpendicular of a subspace of \mathfrak{k}^* is a subspace of \mathfrak{k} .)

Proof. A vector \vec{k} is in the stabilizer algebra of m iff the vector field induced on M vanishes at m , iff the symplectic gradient of $\langle \Phi, \vec{k} \rangle$ vanishes at m , iff the differential $d\langle \Phi, \vec{k} \rangle$ vanishes at m , iff $\langle d\Phi, \vec{k} \rangle$ vanishes at m , which means \vec{k} is in the perp of the image of the differential $d\Phi$ at m . \square

There are two interesting extremes of this. One is that m ’s stabilizer group in K is discrete if and only if m ’s stabilizer algebra in \mathfrak{k} is trivial if and only if the differential of Φ at m is onto. On the other hand, m is a K -fixed point if and only if all of the vector fields induced by \mathfrak{k} vanish at m if and only if the differential of Φ at m is zero.

²A unitary matrix has $UU^* = 1$. A unitary matrix near the identity, $1 + \epsilon S$, therefore has $(1 + \epsilon S)(1 + \epsilon S)^* = 1 + \epsilon(S + S^*) + O(\epsilon^2) = 1$. Differentiating, we get $S = -S^*$.

Corollary. Let $p \in \mathfrak{k}^*$ be a boundary point of the image of a moment map $\Phi : M \rightarrow \mathfrak{k}^*$. Then each point in $\Phi^{-1}(p)$ is stabilized by at least a circle subgroup of K .

Proof. If p is an extremal point, the image of the differential can't be onto – some directions from p lead outside the image of Φ . So the perp to the differential is positive-dimensional, and generates a positive-dimensional subgroup of K , the connected component of the stabilizer group. And any positive-dimensional compact group contains a circle group. \square

We include these proofs to show off the simple connection between the action of the Lie group and the properties of a moment map. While A/G-S convexity is based on Morse theory applied to the moment map, the results above are much more pedestrian and use only basic differential geometry.

This doesn't all come for free, of course – the important fact is that we happen to be studying isospectral sets of Hermitians (which turn out to be symplectic manifolds), and certain maps from them like “take diagonals” (which turn out to be moment maps). Most other equations one might like to study in linear algebra are not stable as the vanishing of a moment map. I personally take this framework as a guide to some of those linear algebra problems which are likely to have nice solutions.

2.5. When K is not a torus. There is a different, rather less pleasant, convexity result for noncommutative groups K . The first example is $K = SO(3)$ the group of rotations of \mathbb{R}^3 , and M a coadjoint orbit. We can identify $\mathfrak{so}(3)^*$ with \mathbb{R}^3 and so M is a sphere centered at the origin. Since this is not convex, we know we'll have to look for a slightly subtler statement than A/G-S.

In the conjugation action of $U(n)$ on the space of Hermitian matrices (not naturally a symplectic manifold), we know two nice things already; every Hermitian matrix is diagonalizable, and if we insist that the (real) diagonal entries then be in decreasing order, the diagonal matrix is unique. Inside the space of diagonal Hermitian matrices, which we can think of as the dual \mathfrak{t}^* to the Lie algebra of the maximal torus of $U(n)$, this picks out a certain cone \mathfrak{t}_+^* called a **positive Weyl chamber** for $U(n)$. Since $U(n)$ is the only example we will need we won't give the general definition of positive Weyl chamber, but merely state that for every connected compact Lie group K , there is an analogous group T and polyhedral cone \mathfrak{t}_+^* , such that each orbit in \mathfrak{k}^* intersects \mathfrak{t}_+^* in a unique element.

Theorem. Let M be a compact connected Hamiltonian K -manifold, with moment map Φ . Then the intersection of the image of Φ with the positive Weyl chamber \mathfrak{t}_+^* is a convex polytope.

This theorem is due to Kirwan, and is a fair bit harder than the case K commutative (the A/G-S theorem above). The new difficulties primarily come from the points of the moment polytope lying on the boundary of the positive Weyl chamber (in the Hermitian case: the Hermitian matrices with repeated eigenvalues), which we call the **Weyl walls**.

We can ignore these to some extent: define the **symplectic slice** of M as the preimage under Φ of the interior of \mathfrak{t}_+^* . It is then a theorem that the symplectic slice is symplectic, and a Hamiltonian T -manifold, whose image is the K -moment polytope minus the parts hitting the Weyl walls. (Since it is usually noncompact we can't use it to trivially reduce Kirwan convexity to A/G-S convexity.)

There is a slightly different statement of Kirwan convexity: instead of intersecting the image of Φ with \mathfrak{t}_+^* , we can compose Φ with the K -invariant map $\mathfrak{k}^* \rightarrow \mathfrak{t}_+^*$ that takes an element to the unique point in \mathfrak{t}_+^* in its K -orbit. (For $K = U(n)$ this takes a Hermitian matrix M to the diagonal matrix with M 's eigenvalues decreasing down the diagonal.) Then Kirwan convexity says that the image of this composite map is a convex polytope (the same as constructed before by intersecting).

With all this machinery we can give pleasant proofs of Horn's theorems (but not Horn's conjecture!) on the sum of two Hermitian matrices.

Theorem. *Let $\mathcal{O}_\lambda, \mathcal{O}_\mu$ be the spaces of Hermitian matrices with spectrum λ, μ . Let $e : \mathcal{O}_\lambda \times \mathcal{O}_\mu \rightarrow \mathbb{R}^n$ take a pair of matrices to the spectrum of their sum, listed in decreasing order. Then the image of e is a convex polytope. Also, if $e(H_\lambda, H_\mu)$ is an extremal point of the image of e and is also a strictly decreasing list, then H_λ, H_μ are simultaneously block diagonalizable.*

Proof. We have already explained how to identify \mathcal{O}_λ and \mathcal{O}_μ with coadjoint orbits of $U(n)$. Therefore by property 3 of moment maps listed above, their product has a Hamiltonian action of $U(n) \times U(n)$. (In fact it is a coadjoint orbit for this big group.)

Consider the action of the diagonal $U(n)$, i.e. conjugating both Hermitian matrices by the same unitary matrix. Then by property 1 this action is Hamiltonian, and we can compute its moment map as the transpose of the inclusion $\mathfrak{u}(n) \hookrightarrow \mathfrak{u}(n) \oplus \mathfrak{u}(n)$, composed with the $U(n) \times U(n)$ moment map, which was just inclusion of the coadjoint orbit.

The transpose of diagonal inclusion $V \rightarrow V \oplus V$ is summation $V^* \oplus V^* \rightarrow V^*$; so the moment map for $U(n)$'s action on $\mathcal{O}_\lambda \times \mathcal{O}_\mu$ just takes a pair of Hermitian matrices to their sum.

But now we have convexity, as our map e is just the map used in the alternate description of Kirwan's convexity theorem.

If $e(H_\lambda, H_\mu)$ is a strictly decreasing list of real numbers, that means it's in the interior of the positive Weyl chamber, therefore in the image of the symplectic slice, which for $\mathcal{O}_\lambda \times \mathcal{O}_\mu$ is the set of pairs (H_λ, H_μ) whose sum is diagonal with decreasing entries (already a familiar set to people studying Horn's problem). Then we can use some element of $U(n)$ to conjugate (H_λ, H_μ) into the symplectic slice, on which e is the moment map for the action of T .

Since $e(H_\lambda, H_\mu)$ is on the boundary of the image of e , we can apply the corollary from subsection 2.4, and determine that (H_λ, H_μ) is invariant under some circle in T (and not just the scalar matrices, which fix all pairs (H_λ, H_μ)). Being invariant under conjugation by a nonscalar diagonal matrix forces each of H_λ and H_μ to be block diagonal. \square

(The very careful reader will wonder why we bothered with the symplectic slice in the above, since the proposition about stabilizers didn't require that the group be a torus. But just because we knew that $e(H_\lambda, H_\mu)$ was on the boundary of the image of e , we didn't know $H_\lambda + H_\mu$ to be on the boundary of the image of $U(n)$'s moment map – and in fact it potentially wasn't, if $e(H_\lambda, H_\mu)$ was on a wall of the Weyl chamber.)

There is now an industry generalizing these convexity results to larger contexts (such as "Poisson" actions), some of which are definitely relevant to linear algebra; we don't pause to discuss these, as we won't need them for Horn's problem.

3. SYMPLECTIC QUOTIENTS

Given M a Hamiltonian K -manifold with moment map $\Phi : M \rightarrow \mathfrak{k}^*$, we now have a number of theorems about the image of Φ – put differently, about which fibers of Φ are nonempty – and the action of K on the fiber. Given this setup, and a point $\mu \in \mathfrak{k}^*$ fixed by K , define the **symplectic quotient** (or **symplectic reduction**) of M at the level μ as $\Phi^{-1}(\mu)/K$. There is a reason that these quotients are nicer to study than the fibers themselves:

Theorem (Marsden-Weinstein, 1974). *The (dense, open) smooth part of $\Phi^{-1}(\mu)/K$ inherits a canonical symplectic structure.*

(The case most frequently studied is μ is a regular value of Φ , since $\Phi^{-1}(\mu)$ is then a submanifold by the inverse mapping theorem, and the action of K on it has only finite stabilizers as discussed in subsection 2.4.)

In the case $K = T$, where the conjugation action is trivial, the adjoint and coadjoint representations are also trivial. So any point $\mu \in \mathfrak{k}^*$ is fixed by T . In the case $K = U(n)$ only the scalar Hermitian matrices are fixed by conjugation.

Example 5. Let $M = \mathbb{C}^n$, with the circle group $S^1 = \{\exp(i\theta)\}$ acting by multiplication by phases. Identifying the Lie algebra of S^1 with the reals, the moment map is $\Phi(\vec{v}) = \frac{1}{2}|\vec{v}|^2$. Then if $k > 0$, the symplectic quotient $\Phi^{-1}(k)/S^1$ is $\mathbb{C}P^{n-1}$; this is the Hopf fibration, fiber-ing S^{2n-1} by circles. If $k = 0$, the symplectic quotient is a point. If $k < 0$, the symplectic quotient is empty.

Example 6. Let $M = \mathcal{O}_\lambda$ be the space of Hermitian matrices with spectrum λ , and $K = T$ the group of diagonal unitary matrices acting by conjugation, so Φ is the map taking a Hermitian matrix to its diagonal entries. Then $\Phi^{-1}(\mu)/T$ is the space of Hermitian matrices with eigenvalues λ , and diagonal entries μ , up to conjugation by diagonal unitary matrices. These spaces are studied in the doctoral theses [Kn, G], and the case $\lambda = (1, 1, 0, \dots, 0)$ is described in detail in subsection 9.1.

4. FLAG MANIFOLDS

Hopefully the previous section has convinced the reader that some of the fundamental objects of interest in this theory are the isospectral sets themselves. We will see now that these are not just real manifolds, but complex manifolds, suggesting that they may be studiable using complex algebraic geometry.

Define a **(partial) flag** in \mathbb{C}^n as an increasing list \mathcal{V} of subspaces $0 = V_0 < V_1 < \dots < V_s = \mathbb{C}^n$, with relative dimensions $\dim V_i/V_{i-1} = d_i$. If $s = n$ then \mathcal{V} is called a **full flag**. If $s = 2$ then \mathcal{V} is just given by a subspace of dimension d_1 (with the automatic V_0 and V_2).

Given a Hermitian matrix H , we can associate a partial flag \mathcal{V}_H as follows: let V_i be the sum of the eigenspaces corresponding to the i smallest eigenvalues. This flag will be full if and only if H has no repeated eigenvalues. Conversely, given a partial flag \mathcal{V} and an increasing list of s eigenvalues e_i , we can construct a Hermitian matrix $H_{\mathcal{V}}$ as the sum

$$H_{\mathcal{V}} = \sum_{i=1}^s e_i \cdot [\text{the projection onto } V_{i-1}^\perp \cap V_i].$$

So in all, the space of Hermitian matrices with spectrum λ is in 1:1 correspondence with a certain space of partial flags we denote $\text{Flags}(d_1, \dots, d_s)$ (where the $\{d_i\}$ are determined by λ 's repetition of eigenvalues), called a **flag manifold**. (In the special case $s = 2$, this is also called the **Grassmannian of d_1 -planes** and denoted $\text{Gr}_{d_1}(\mathbb{C}^n)$.)

What have we gained? The benefit of this view is that while Hermitian matrices only have an evident action of the unitary group $U(n)$, the space of flags $\text{Flags}(d_1, \dots, d_s)$ has an action of the group $GL_n(\mathbb{C})$ of all invertible $n \times n$ matrices – applying a linear transformation to a flag produces a new flag. Since $U(n)$ acts transitively on the Hermitian matrices with a given set of eigenvalues, so too will this larger group $GL_n(\mathbb{C})$ act with one orbit – given two flags, there is a linear transformation taking one to the other. (And it may even be taken to be unitary.)

The unitary matrices stabilizing a given Hermitian matrix – say, a diagonal one with strictly decreasing eigenvalues – are easy to understand; they are just the diagonal unitary matrices T . From this we can conclude that if λ has no repeated eigenvalues, we can identify the space of Hermitian matrices with spectrum λ with $U(n)/T$. For the corresponding statement for $GL_n(\mathbb{C})$ we must compute the stabilizer of a flag – say, the standard flag $(0 < \mathbb{C}^1 < \mathbb{C}^2 < \dots < \mathbb{C}^n)$ – which can be seen to be the group B of invertible upper triangular matrices. So we can identify $\text{Flags}(1, 1, \dots, 1)$ with $GL_n(\mathbb{C})/B$.

Since $GL_n(\mathbb{C})$ is not just a real Lie group but a complex Lie group, and B a complex subgroup,³ we find out that $\text{Flags}(1, 1, \dots, 1)$ is naturally a complex manifold. As we will see in the sections to come, it is actually a “complex algebraic variety”.

We invite the reader to determine the corresponding statements for the $s < n$, partial flag manifold case.

5. GEOMETRIC INVARIANT THEORY IN THE AFFINE CASE

Our interim goal is to describe the algebro-geometric analogue of symplectic quotients. We first give the basics of affine algebraic geometry, and the concept of quotient in that case; however for our linear algebra applications we'll need to work through the additional complications of projective geometry.

Before getting into algebraic geometry, let us think about embeddings of manifolds M into \mathbb{R}^n . On the one hand, if we choose n real-valued functions $\{f_i\}$ on M , such that any two points are distinguished by at least one of the functions, the map $m \mapsto (f_1(m), f_2(m), \dots, f_n(m))$ gives an injection⁴ of M into \mathbb{R}^n . Conversely, if M is a submanifold of \mathbb{R}^n , then the n coordinate functions, restricted to M , give us n functions separating points.

By taking polynomials in those n functions, we get an algebra of functions on M , some quotient of the polynomial ring $\mathbb{R}[f_1, \dots, f_n]$. For example, if M is the parabola $y = x^2$ in \mathbb{R}^2 , and $f_1 = x$ and $f_2 = y$, then the algebra of functions they will generate will be $\mathbb{R}[f_1, f_2]/(f_1^2 - f_2)$. (The function $f_1^2 - f_2$ is not zero on the whole plane, but it is the zero function on M .)

³This means that its Lie algebra is invariant under multiplication by i . By contrast, $U(n)$ is not a complex subgroup because i times a skew-Hermitian matrix is not always skew-Hermitian.

⁴Which is not good enough, if we're thinking about smooth manifolds M and want to model the smooth structure; we must also ask that the map be an immersion. This goes by the phrase “separating points”, which we're already asking, and “separating tangent vectors.”

The manifolds we are actually interested in are flag manifolds (and things we build from them), which we saw in the last section are complex. For this reason the only rings we will bother to consider henceforth will be quotients of polynomial rings with complex coefficients rather than the reals.

5.1. The spectrum of a ring. Let's now think about the reverse direction: given a quotient $R = \mathbb{C}[f_1, \dots, f_n]/\mathcal{I}$ of polynomial ring by an ideal, we can define **the spectrum**⁵ $\text{Spec } R$ as the subset of \mathbb{C}^n where all the polynomials in the ideal \mathcal{I} vanish.

Of course, this is a bad definition, as the notation $\text{Spec } R$ only refers to the abstract ring R , and not the particular way of presenting it as a quotient of a polynomial ring. (There may be many of these, just as there are many embeddings of M into affine spaces, of many dimensions.) There is an alternate, equivalent, definition of $\text{Spec } R$ as the set of maximal ideals⁶ of R , which doesn't require us to choose an embedding. But just as we generally deal with a manifold by picking coordinates on it, we generally deal with **affine varieties** $\text{Spec } R$ by embedding them in affine space.

Given a map on rings $R \rightarrow S$, there is an induced map⁷ from $\text{Spec } S$ to $\text{Spec } R$, taking a maximal ideal to its preimage. The most fundamental example of this is $\mathbb{C}[x_1, \dots, x_n] \rightarrow R$, which induces the inclusion $\text{Spec } R \hookrightarrow \mathbb{C}^n$. Not every function from $\text{Spec } S$ to $\text{Spec } R$ arises in this way, only those preserving in some sense the structure of algebraic variety. While it is possible to give a definition of such functions without explicit reference to R and S we will not need this.

5.2. Group actions. Now take the situation that a complex Lie group G acts on our ring R by ring automorphisms. For example, if $R = \mathbb{C}[x_1, \dots, x_n]$ (so $\text{Spec } R$ is just \mathbb{C}^n itself), then $\text{GL}_n(\mathbb{C})$ acts on R by linear changes of the $\{x_i\}$, and the induced action on higher-order polynomials.

Since G acts by ring automorphisms, it takes maximal ideals to maximal ideals, so acts on the set $\text{Spec } R$. We will be interested in forming the quotient $(\text{Spec } R)/G$, with the natural map $\text{Spec } R \rightarrow (\text{Spec } R)/G$.

Since we're thinking in terms of algebraic geometry, this really means we're looking for a ring S with $\text{Spec } S = (\text{Spec } R)/G$; equivalently, S should be the ring of functions on $(\text{Spec } R)/G$. Each element of S therefore pulls back to a G -invariant function on $\text{Spec } R$. So S will have to map back to the G -invariant functions on $\text{Spec } R$, which we denote R^G ; this is the clear candidate for S . Since R^G is the "ring of invariants", we will call $\text{Spec } (R^G)$ the **geometric invariant theory quotient** (or GIT quotient) of $\text{Spec } R$ by G , and denote it $(\text{Spec } R)//G$.

It is not quite true, though, that $(\text{Spec } R)//G = (\text{Spec } R)/G$; rather we only have a map $(\text{Spec } R)/G \rightarrow (\text{Spec } R)//G$. It is already a tricky theorem that when G is "reductive" (like $\text{GL}_n(\mathbb{C})$, and all the other groups we will be using) that (1) R^G is finitely generated, so

⁵If T is a linear transformation $V \rightarrow V$, and $R = \mathbb{C}[T]/(\text{the characteristic polynomial of } T)$, then $\text{Spec } R$ is the spectrum in the usual sense. It is rather amusing to trace in this way the modern-day algebraic geometry terminology back to Rydberg lines of atoms!

⁶One of the great insights of algebraic geometry in this century is that for almost all purposes one should instead work with the prime ideals, a somewhat larger set than the maximal ideals. But we will be avoiding any of the contexts in which this distinction becomes important.

⁷For the rings we will consider, which are finitely generated; geometrically, this corresponds to us being able to embed our varieties in finite-dimensional vector spaces

$(\text{Spec } R)//G$ can be embedded in a finite dimensional space, and (2) the natural map from $(\text{Spec } R)/G \rightarrow (\text{Spec } R)//G$ is onto.

Example 7. Let $G = \mathbb{C}^\times$, the nonzero complex numbers under multiplication, acting on $R = \mathbb{C}[x_1, \dots, x_n]$ by rescaling each coordinate the same way. Then the induced action on G on \mathbb{C}^n is also by rescaling. The ordinary set-theoretic quotient is not Hausdorff; it's projective space union the point $\vec{0}$ in the closure of every other point. The only invariant polynomials in R^G are the constants, so $\text{Spec } (R^G) = \text{pt}$. (This will turn out to be related to the $k = 0$ case of example 5.)

Obviously this is unsatisfying; we'd rather the quotient be projective space itself (somehow including the $k > 0$ case of example 5). But projective space is not $\text{Spec } S$ for any S , by Liouville's theorem – any function on projective space is constant! We will need subtler constructions than Spec to make projective spaces, and more to the point, flag manifolds.

6. GEOMETRIC INVARIANT THEORY IN THE PROJECTIVE CASE, AND THE KIRWAN/NESS THEOREM

Let $R = \bigoplus_{k \in \mathbb{N}} R_k$ now be a *graded* ring, meaning that the product of an element of R_k with an element of R_m lands in R_{k+m} . The standard example is R a polynomial ring, with R_k the homogeneous polynomials of degree k .

We will define $\text{Proj } R$ in several equivalent ways. The simplest, least useful for visualizing examples, is that $\text{Proj } R$ the set of maximal *graded* ideals of R , where a graded ideal is one equal to the direct sum of its intersections with the R_n .

For the second, note that R has a natural action of \mathbb{C}^\times , acting on R_k by rotating it with speed k . There is a natural \mathbb{C}^\times -invariant surjection $R \twoheadrightarrow R_0$ (which would not be true if R were \mathbb{Z} -graded instead of \mathbb{N} -graded). So there is a map backwards $\text{Spec } R_0 \hookrightarrow \text{Spec } R$. Then

$$\text{Proj } R := (\text{Spec } R \setminus \text{Spec } R_0)/\mathbb{C}^\times.$$

Example 8. Let $R = \mathbb{C}[x_1, \dots, x_n]$ with R_n the homogeneous polynomials of degree n . Then as discussed in example 7, R_0 is just the constants, and the point $\text{Spec } R_0$ includes into $\text{Spec } R = \mathbb{C}^n$ as the origin $\vec{0}$. Then $\text{Proj } R = (\mathbb{C}^n \setminus \{\vec{0}\})/\mathbb{C}^\times$ is just the usual definition of $\mathbb{C}P^{n-1}$.

So Proj of a polynomial ring (with each variable given degree 1) is projective space, just like Spec of a polynomial ring is affine space. And just as writing a ring R as the quotient of a polynomial ring, $\mathbb{C}[x_1, \dots, x_n] \twoheadrightarrow R$, dually gives us an inclusion of $\text{Spec } R \hookrightarrow \mathbb{C}^n$, writing a graded ring R as the quotient of a polynomial ring by a graded ideal \mathcal{I} dually gives us an inclusion $\text{Proj } R \hookrightarrow \mathbb{C}P^{n-1}$. This gives a third description of $\text{Proj } R$, in the case R is presented as $\mathbb{C}[x_1, \dots, x_n]/\mathcal{I}$ and all the x_i are degree 1; as the points in $\mathbb{C}P^{n-1}$ where the elements of \mathcal{I} all vanish. In this case $\text{Proj } R$ is called a **projective variety**.

The description we will use most is the second, in terms of $\text{Spec } R$ and $\text{Spec } R_0$.

6.1. Another relation between Spec and Proj . Given an ungraded ring R_0 , we can define $R := R_0[l]$, where l is a new variable assigned formal degree 1. Then $\text{Spec } R = \text{Spec } R_0 \times \mathbb{C}$, so $\text{Spec } R \setminus \text{Spec } R_0 = \text{Spec } R_0 \times \mathbb{C}^\times$, and $\text{Proj } R = \text{Spec } R_0$. Upshot: anything we can make as a Spec , we can also make as a Proj .

6.2. Maps of graded rings. Given a ring homomorphism $R \rightarrow S$, we defined a natural map from $\text{Spec } S \rightarrow \text{Spec } R$. So it is natural to assume that a grading-preserving map $R \rightarrow S$ induces a map on the corresponding Proj 's. This is not the case! Rather:

Proposition 1. *Let $f : R \rightarrow S$ be a homomorphism of graded rings. Let X be the points of $\text{Spec } S \setminus \text{Spec } S_0$ such that every function in $f(R)$ vanishes at X , and define $(\text{Proj } S)_f^{\text{us}}$ as the quotient of X by the natural action of \mathbb{C}^\times . Then f does not necessarily induce a map from $\text{Proj } S \rightarrow \text{Proj } R$, but does induce one from $\text{Proj } S \setminus (\text{Proj } S)_f^{\text{us}} \rightarrow \text{Proj } R$.*

Proof. To see what the problem is, let's trace through the definition of Proj . Since R and S are graded, i.e. have \mathbb{C}^\times -actions, $\text{Spec } R$ and $\text{Spec } S$ have \mathbb{C}^\times actions. The map between the rings being grading-preserving is equivalent to it being \mathbb{C}^\times -equivariant. Consequently, the map backwards on $\text{Spec } S \rightarrow \text{Spec } R$ is \mathbb{C}^\times -equivariant.

The problem comes when we try to rip out $\text{Spec } S_0$ and $\text{Spec } R_0$. Of course there is no problem when in restricting the map to $\text{Spec } S \setminus \text{Spec } S_0$. But the image of this may not land inside $\text{Spec } R \setminus \text{Spec } R_0$; there may be points of $\text{Spec } S \setminus \text{Spec } S_0$ that hit $\text{Spec } R_0$. Once we rip them out, then there's no problem. \square

In the context to come, the set $(\text{Proj } S)_f^{\text{us}}$ will be called the **unstable set** (hence the us).

6.3. Geometric invariant theory. We now apply what we've learned about maps of Proj 's to the case we studied in the last section, $R^G \hookrightarrow R$. Given a graded ring R with an action of a group G , define the **geometric invariant theory quotient**

$$(\text{Proj } R)//G := \text{Proj } (R^G).$$

There is a natural map from $\text{Proj } R$ minus the unstable set described by the proposition above (the set of points where all G -invariant functions vanish) to $(\text{Proj } R)//G$.

While this may seem like a totally canonical definition, there are two traps for the unwary. In many problems one starts with the variety $\text{Proj } R$ and its G -action rather than the ring R , and does not want the extra choice that comes in finding an R (of which, unlike in the affine case, there will be many). And even when one has chosen an R , there may be many ways to get G to act on it inducing the same action on $\text{Proj } R$. It turns out that both these choices matter – which is to say, the notation $(\text{Proj } R)//G$ is misleading.

Example 9. Let $X = \mathbb{C}^n$, and $G = \mathbb{C}^\times$ act by rescaling. Then $X = \text{Spec } \mathbb{C}[x_1, \dots, x_n] = \text{Proj } \mathbb{C}[x_1, \dots, x_n, l]$, where the $\{x_i\}$ are all degree 0, and l is degree 1. The action of G on our homogeneous coordinate ring R is not quite determined – while the x_i are required to all be weight 1, the action on the variable l can be any weight k , i.e. $z \cdot l = z^k l$.

If $k > 0$, there are no invariant polynomials other than constants: $R^G = \mathbb{C}$, and $\text{Proj } (R^G)$ is the empty set.

If $k = 0$, the invariant ring $R^G = \mathbb{C}[l]$, and $\text{Proj } (R^G)$ is a point.

If $k = -1$, the invariant ring $R^G = \mathbb{C}[x_1 l, x_2 l, \dots, x_n l]$ with all products of degree 1, so $\text{Proj } (R^G) = \mathbb{CP}^{n-1}$. It is a little trickier to see that this is true for *all* negative k . (It is interesting to compare this to example 5. The difference in signs can be explained but we do not do so here.)

Example 10. Let $X = \mathbb{C}^{m \times n}$, the space of $m \times n$ matrices, with $m \leq n$, and let $G = \text{GL}_m(\mathbb{C})$ act on X on by left multiplication. Then as in the previous example (the $m = 1$ case of this

one) there is a parameter k describing the action of G on the extra variable l , such that for $k > 0$ the quotient $X//G$ is empty, and for $k = 0$ the quotient is a point.

Regard X as the space of m vectors in \mathbb{C}^n . For $k < 0$, the unstable set turns out to be the set of linearly dependent m -tuples, and the GIT quotient $X//G$ is thus the Grassmannian $\text{Gr}_m(\mathbb{C}^n)$. (If you take a linearly independent set of m vectors, and quotient by the action of $\text{GL}_m(\mathbb{C})$, you forget the vectors and only remember the m -subspace they span.)

6.4. The equivalence of symplectic and GIT quotients. We spell out a moment map that's a special case of things already said: complex projective space $\mathbb{C}\mathbb{P}^{n-1}$ is a Hamiltonian $U(n)$ -manifold, with moment map

$$\begin{aligned} \mathbb{C}\mathbb{P}^{n-1} &\rightarrow \mathfrak{u}(n)^* \\ [\vec{v}] &\mapsto [\text{the rank 1 projection onto } \mathbb{C}\vec{v}] \end{aligned}$$

using the identification already given between $\mathfrak{u}(n)^*$ and the space of Hermitian matrices. In fact we are observing here that $\mathbb{C}\mathbb{P}^{n-1}$ is a coadjoint orbit of $U(n)$, the orbit $\mathcal{O}_{(1,0,\dots,0)}$.

It is an easy fact,⁸ not proven here, that any complex submanifold of $\mathbb{C}\mathbb{P}^{n-1}$ inherits a symplectic structure from $\mathbb{C}\mathbb{P}^{n-1}$. If $X \subseteq \mathbb{C}\mathbb{P}^{n-1}$ is preserved by the action of a subgroup $K \leq U(n)$, we can compute the moment map using the properties we listed of moment maps:

$$X \hookrightarrow \mathbb{C}\mathbb{P}^{n-1} \rightarrow \mathfrak{u}(n)^* \rightarrow \mathfrak{k}^*.$$

Lastly, given an action of a compact group K on a complex vector space, there is a unique extension to an action of **the complexification** $K^{\mathbb{C}}$ of K . We will not stop to define this group in detail; suffice it to say that every compact group is a subgroup of a unique complex group, such that the Lie algebra of the complex group is the complexification of the Lie algebra of the compact group. The only example that will interest us is $U(n)^{\mathbb{C}} = \text{GL}_n(\mathbb{C})$ (every matrix is uniquely the sum of a skew-Hermitian and i times a skew-Hermitian).

We are ready to state the deepest theorem in this paper, proven separately by Kirwan and Ness:

Theorem 1. *Let K act on the graded ring $R = \mathbb{C}[x_1, \dots, x_n]/\mathcal{I}$, all generators x_i of degree 1. So $\text{Proj } R$ is a subvariety of $\mathbb{C}\mathbb{P}^{n-1}$, and (if smooth) a Hamiltonian K -manifold. Let Φ be the moment map calculated above for K 's action on M . Then there is a natural identification*

$$\Phi^{-1}(0)/K \cong (\text{Proj } R)//K^{\mathbb{C}}$$

between the symplectic quotient and geometric invariant theory quotient.

In fact there is a more general result, slightly more complicated to state, in which $\text{Proj } R$ is not necessarily projective. We already saw this in the case of $K = S^1$ acting on \mathbb{C}^n , in examples 5 and 9.

7. BACK TO FLAG MANIFOLDS

Since we've only stated the symplectic quotient equals GIT quotient theorem at the level $\Phi = 0$, we need to slightly modify the Horn problem $A + B = C$ to $A + B + C = 0$.

⁸once one proves the non-easy Darboux theorem, given the unusual way we've defined symplectic

Let $M = \mathcal{O}_\lambda \times \mathcal{O}_\mu \times \mathcal{O}_\nu$, the space of triples of Hermitian matrices with eigenvalues respectively λ, μ, ν . Then the moment map for the diagonal action of $U(n)$ on the product is the sum of the three matrices, and Horn's problem asks when this symplectic quotient is nonempty.

To jack into the symplectic vs. GIT theorem, we need to express \mathcal{O}_λ as Proj of something, or equivalently explain how to embed it in projective space such that the restriction of the symplectic structure on projective space matches that on \mathcal{O}_λ . This turns out to only be possible if λ is integral, and involves some very classical geometry.

7.1. The $U(n)$ case of the Borel-Weil-Bott-Kostant theorem. In fact we will not need the details presented in this section in what follows, but it lets us avoid some representation theory of $GL_n(\mathbb{C})$.

Recall the flag manifolds $\text{Flags}(d_1, \dots, d_s)$ defined before; we will restrict to the case $d_i \equiv 1$ and just write $\text{Flags}(\mathbb{C}^n)$, leaving the interested reader to work out the general case. Denote by $\text{Gr}_k(\mathbb{C}^n)$ the Grassmannian of k -dimensional subspaces of \mathbb{C}^n (these are very partial flag manifolds). There is a natural forgetful map from $\text{Flags}(\mathbb{C}^n)$ to each Grassmannian, taking a flag to its subspace of dimension k , forgetting all the subspaces below and above.

Since a flag is just a list of subspaces, the product map

$$\text{Flags}(\mathbb{C}^n) \rightarrow \prod_{k=1}^n \text{Gr}_k(\mathbb{C}^n)$$

is an inclusion.

Given a k -dimensional subspace A of \mathbb{C}^n , we can wedge together a basis $\{\vec{a}_i\}$ of A to get a nonzero alternating tensor in $\wedge^k \mathbb{C}^n$. If we change the basis, our element $\wedge_{i=1}^k \vec{a}_i$ only changes by a scalar factor, so gives a well-defined element of the projective space. This is called the **Plücker embedding**

$$\text{Gr}_k(\mathbb{C}^n) \rightarrow \mathbb{P}(\wedge^k \mathbb{C}^n).$$

Given a vector $\vec{v} \in V$, and a natural number a , we can tensor \vec{v} with itself a times to get a symmetric tensor $\vec{v}^{\otimes a} \in \text{Sym}^a(V)$. This descends to the projectivized spaces $\mathbb{P}V \hookrightarrow \mathbb{P}(\text{Sym}^a V)$, and is called the **a th Veronese embedding**. Choosing a family $\{a_k\}$ of naturals, we get maps

$$\mathbb{P}(\wedge^k \mathbb{C}^n) \rightarrow \mathbb{P}(\text{Sym}^{a_k}(\wedge^k \mathbb{C}^n)).$$

In the very special case that $\dim V = 1$, we don't need a to be a natural number – for $a < 0$ we define $\text{Sym}^a(V) := \text{Sym}^{-a}(V^*)$. This will be handy in the case $k = n$, where $\wedge^n \mathbb{C}^n$ is one-dimensional.

Finally, given vectors $\vec{v} \in V$ and $\vec{w} \in W$, we can tensor them together, inducing a map on projective spaces $\mathbb{P}V \times \mathbb{P}W \rightarrow \mathbb{P}(V \otimes W)$ called the **Segre embedding**. In the case at hand, this gives a map

$$\prod_{k=1}^n \mathbb{P}(\text{Sym}^{a_k}(\wedge^k \mathbb{C}^n)) \rightarrow \mathbb{P}(\otimes_{k=1}^n \text{Sym}^{a_k}(\wedge^k \mathbb{C}^n)).$$

Proposition (one small aspect of the Borel-Weil-Bott-Kostant theorem). *Let $\lambda = (\lambda_1 > \dots > \lambda_n)$ be a strictly decreasing list of integers, and $a_k = \lambda_k - \lambda_{k+1}$ for $k = 1 \dots n$ (taking λ_{n+1} to be zero). Then the composite of the above maps*

$$\text{Flags}(\mathbb{C}^n) \rightarrow \mathbb{P}(\otimes_{k=1}^n \text{Sym}^{a_k}(\wedge^k \mathbb{C}^n))$$

induces a symplectic structure on $\text{Flags}(\mathbb{C}^n)$ matching the one given by its diffeomorphism with the coadjoint orbit \mathcal{O}_λ .

(Again, we exhort the reader to think about the case of λ only weakly decreasing, the partial flag manifold case.)

Since we have exhibited \mathcal{O}_λ as a variety embedded in projective space, we can look at its homogeneous coordinate ring, which we will denote R^λ , reserving subscripts to indicate graded pieces.

7.2. The Borel-Weil theorem. It is evident that all the maps described above are equivariant with respect to the action of $\text{GL}_n(\mathbb{C})$ on \mathbb{C}^n , so $\text{GL}_n(\mathbb{C})$ therefore acts on the homogeneous coordinate rings R^λ , and in particular each graded piece $(R^\lambda)_k$ is a representation of $\text{GL}_n(\mathbb{C})$. We pause to mention three important facts about these representations.

Proposition 2 (part of Borel-Weil).

1. *Each graded piece $(R^\lambda)_k$ of each homogeneous coordinate ring for the flag manifold is an irreducible representation of $\text{GL}_n(\mathbb{C})$.*
2. *$(R^\lambda)_k$ is isomorphic, as a representation, to $(R^{k\lambda})_1$.*
3. *Every irreducible representation of $\text{GL}_n(\mathbb{C})$ arises as $(R^\lambda)_1$ for a unique λ .*

We denote $(R^\lambda)_1$ by V_λ ; for reasons we will not go into here it is called the **irreducible representation of $\text{GL}_n(\mathbb{C})$ of highest weight λ** . So we can identify R^λ with $\bigoplus_{n \in \mathbb{N}} V_{n\lambda}$.

A little bit more can be said, for those who know the classification of irreducible representations of Lie groups in terms of “highest weights”, but as representation theory is not in the title of this paper we will not explore this further.

8. EASY CONSEQUENCES

We need one more fact about the homogeneous coordinate rings R^λ for the flag manifold: they have no zero divisors. This is implied, though we will not explore how, by the (much stronger) fact that the flag manifold is connected and smooth.

We’re now ready to apply our big gun, the equivalence of symplectic and GIT quotients.

Theorem. *Let λ, μ, ν be weakly decreasing sequences of integers. Then the space*

$$\left\{ (H_\lambda, H_\mu, H_\nu) : H_\lambda + H_\mu + H_\nu = 0 \right\} / \text{U}(n)$$

can be identified with

$$\text{Proj} \bigoplus_{k \in \mathbb{N}} (V_{k\lambda} \otimes V_{k\mu} \otimes V_{k\nu})^{\text{GL}_n(\mathbb{C})}.$$

Proof. We’ve already seen that the first space is the symplectic quotient of $\mathcal{O}_\lambda \times \mathcal{O}_\mu \times \mathcal{O}_\nu$ by the diagonal action of $\text{U}(n)$, at the level 0. By the symplectic vs. GIT equivalence, this is the geometric invariant theory quotient of a certain product of flag manifolds by $\text{GL}_n(\mathbb{C})$.

The Borel-Weil-Bott-Kostant theorem, or rather the small part of it presented here, explains how to embed $\mathcal{O}_\lambda, \mathcal{O}_\mu, \mathcal{O}_\nu$ into projective space in order to be able to apply the symplectic vs. GIT theorem. The individual coordinate rings are $\bigoplus_{n \in \mathbb{N}} V_{n\lambda}$ (likewise μ, ν); the coordinate ring of the product is, by Segre embedding, $\bigoplus_{n \in \mathbb{N}} (V_{n\lambda} \otimes V_{n\mu} \otimes V_{n\nu})^{\text{GL}_n(\mathbb{C})}$.

Then to take the GIT quotient we take Proj of the invariant subring. \square

This explains a couple of connections already noticed in the literature between the Horn's problem and the representation theory of $\text{GL}_n(\mathbb{C})$:

Corollary. *Let V_λ, V_μ, V_ν be three irreducible representations of $\text{GL}_n(\mathbb{C})$. If there is a $\text{GL}_n(\mathbb{C})$ -invariant vector in the triple tensor product, then there exist Hermitian matrices H_λ, H_μ, H_ν with the corresponding spectra and zero sum.*

Proof. If the tensor product has an invariant vector, then the coordinate ring of the GIT quotient is nontrivial in its degree 1 piece. Since this ring is a subring of a ring with no zero divisors, it itself has no zero divisors, and therefore is nontrivial in all degrees. This is enough to conclude that Proj of it, the GIT quotient, is nonempty.

By the identification in the theorem, the symplectic quotient is therefore nonempty. So there are three Hermitian matrices with the desired spectra adding to zero. \square

Corollary. *Let λ, μ, ν be weakly decreasing lists of integers. If there exist Hermitian matrices H_λ, H_μ, H_ν with the corresponding spectra and zero sum, then for some $k > 0$, the tensor product $V_{k\lambda} \otimes V_{k\mu} \otimes V_{k\nu}$ has a $\text{GL}_n(\mathbb{C})$ -invariant vector.*

Proof. The condition says the symplectic quotient is nonempty, so running the theorem in reverse, the GIT quotient is nonempty. Consequently the ring of invariants is nontrivial. Therefore in some graded piece $n > 0$ it is nontrivial, giving the desired result. \square

(Both of these corollaries appear in Klyachko's work; he essentially repeats the proof of Kirwan/Ness in this special case.)

Obviously there is a mismatch here; we would like to know that the Hermitian problem has a solution if and only if the tensor product has an invariant vector. But the big machine presented here is not powerful enough to rid us of this k . Other techniques are necessary and the first proof that k can indeed be taken to be 1 appeared in [KT].

9. CONCLUSIONS, AND AN AMUSING EXAMPLE

We saw that certain famous equations in linear algebra, such as $H_\lambda + H_\mu = H_\nu$ (or better, $H_\lambda + H_\mu + H_{-\nu} = 0$), are the conditions that a certain moment map be equal to zero. Whenever this happens we can plug into the theory of Hamiltonian actions on symplectic manifolds. The most spectacular results here establish the convexity properties of the images of these moment maps, and go some way toward determining the image.

In fact the details of the proofs are not actually very different from the hands-on techniques used e.g. by Horn himself; the benefit here is in establishing a framework that points out which problems are likely to accede to such local analysis, and in particular which are likely to lead to linear inequalities and polytopes.

In the (frequent) case that the symplectic manifolds under study are algebraic varieties, one then has a totally new viewpoint, replacing the spaces by homogeneous coordinate

rings, and the study of the moment map by invariant theory. This connects up the symplectic problem with representation theory problems, illuminating both.

Lastly, we emphasize another lesson (not really applied in this paper) from both symplectic and algebraic geometry; the questions like “for which λ, μ, ν is the following symplectic quotient nonempty?” should really be viewed as just the first step in a more detailed study of the symplectic quotient itself. What is its dimension? symplectic volume? its Betti numbers, and cohomology ring?

9.1. An amusing example. Let $X = \mathbb{C}^{2 \times n}$, the space of $2 \times n$ matrices, with its left action of $U(2)$ and right action of $U(n)$. These actions are Hamiltonian with moment maps $M \mapsto -MM^*$ and $M \mapsto M^*M$. We will only be interested in the action of the diagonal subgroup $T^n \leq U(n)$, whose moment map picks out the diagonal entries of M^*M . Note that $\text{Tr } MM^* = \text{Tr } M^*M$, so the two moment maps are not entirely unrelated.

We will symplectic quotient by both groups $U(2)$ and T^n . The symplectic quotient by $U(2)$ at level $s1$ asks that our 2 vectors in \mathbb{C}^n be orthogonal with norm-square s . So for $s > 0$, the symplectic quotient is the Grassmannian of 2-planes in \mathbb{C}^n . (Which is good, because we got this answer when we did this earlier by GIT, in example 10.) Since that is a coadjoint orbit of $U(n)$, studying its reductions by T^n is really the Schur-Horn problem we first mentioned.

If we consider this in the opposite order, we get a very different picture. In symplectic-quotienting $\mathbb{C}^{2 \times n}$ by T^n , each diagonal entry in T^n acts on its own copy of \mathbb{C}^2 . The symplectic quotient of \mathbb{C}^2 by $U(1)$ is either empty, a point, or the Riemann sphere $\mathbb{C}P^1$, depending on the level set $\vec{a} = (a_1, \dots, a_n)$ chosen in $\mathfrak{t}^n \cong \mathbb{R}^n$. Let's take each $a_i > 0$ so the symplectic quotient by T^n is a product of $\mathbb{C}P^1$'s. In fact the moment map for the residual action of $U(2)$ on each of these $\mathbb{C}P^1$ identifies it with a sphere $S_{a_i}^2$ of radius a_i .

Think of such a sphere as the set of steps in \mathbb{R}^3 of length a_i ; then the product $\prod_{i=1}^n S_{a_i}^2$ can be thought of as the set of n -step polygonal paths in \mathbb{R}^3 , with i th step length a_i . Since the moment map for the diagonal action of $U(2)$ on this product is the sum of the individual moment maps, we can think of it as the function taking a path to its endpoint.

The symplectic quotient at level zero is then the set of polygons in \mathbb{R}^3 (because the moment map condition requires that the path terminate at the origin), with edge lengths $\{a_i\}$, considered up to rotation. The connection of polygon spaces to the Schur-Horn problem was first noted in [HK].

REFERENCES

- [F] W. Fulton, Young tableaux, with applications to representation theory and geometry. London Mathematical Society Student Texts, 35. Cambridge University Press, Cambridge, 1997.
- [G] R. Goldin, The cohomology rings of weight varieties. Doctoral thesis, MIT 1999.
- [GLS] V. Guillemin, E. Lerman, S. Sternberg, Symplectic fibrations and multiplicity diagrams. Cambridge University Press, Cambridge, 1996.
- [HK] J.-C. Hausmann, A. Knutson, Polygon spaces and Grassmannians. Enseign. Math. (2) 43 (1997), no. 1-2, 173–198.
- [Kn] A. Knutson, Weight varieties. Doctoral thesis, MIT 1996.
- [KT] A. Knutson, T. Tao, The honeycomb model of $GL_n(\mathbb{C})$ tensor products I: proof of the saturation conjecture, J. Amer. Math. Soc. 12 (1999), pp. 1055-1090.

[MFK] D. Mumford, J. Fogarty, F. Kirwan. Geometric invariant theory. Third edition. Ergebnisse der Mathematik und ihrer Grenzgebiete (2) [Results in Mathematics and Related Areas (2)], 34. Springer-Verlag, Berlin, 1994.

E-mail address: `allenk@math.berkeley.edu`