Lectures on Toric Varieties

Nicholas Proudfoot

Department of Mathematics, University of Oregon, Eugene, OR 97403

Contents

1 Four constructions

Let V be a real vector space of dimension d, and let $V_{\mathbb{Z}} \subset V$ be a lattice. Let $a_1, \ldots, a_n \in V_{\mathbb{Z}}^* \smallsetminus \{0\}$ be a collection of nonzero linear functions on V that take integer values on $V_{\mathbb{Z}}$. Let $\lambda_1, \ldots, \lambda_n \in \mathbb{Z}$ be any integers, and let

$$P := \{v \in V \mid a_i(v) + \lambda_i \geq 0 \text{ for all } i\}.$$

This is called a **rational polyhedron**; if it is bounded, it is called a **rational polytope**. Since all of our polyhedra and polytopes will be rational, we'll drop the adjective.

For all i, let

$$F_i := \{v \in P \mid a_i(v) + \lambda_i = 0\}$$

be the i^{th} **facet** of P. A **face** of P is an intersection of facets. A face of dimension 1 is called an **edge**, and a face of dimension 0 is called a **vertex**. We will always assume that P has at least one vertex.

We intuitively expect that P has dimension d, each facet F_i has dimension $d-1$, each a_i is primitive (not a multiple of another element of $V_{\mathbb{Z}}^*$, and the set of all facets is distinct. However, none of these statements is implied by the definitions, and we will not assume them. For example, if $a_1 = -a_2$ and $\lambda_1 = -\lambda_2$, then P will be contained in a hyperplane and therefore have dimension smaller than d. If $a_1 = a_2$ and $\lambda_1 = \lambda_2$, then F_1 will equal F_2. If $a_1 = a_2$ and $\lambda_1 = \lambda_2 + 1$, then F_2 will be empty. If $a_3 = a_1 + a_2$ and $\lambda_3 = \lambda_1 + \lambda_2$, then F_3 will be contained in $F_1 \cap F_2$, and therefore have smaller than expected dimension.

If all of our expectations *do* hold, then the data of $a_1, \ldots, a_n$ and $\lambda_1, \ldots, \lambda_n$ are completely determined by P. If not, then whenever we say we have a rational polyhedron, what we really mean is that we have a collection $a_1, \ldots, a_n$ and $\lambda_1, \ldots, \lambda_n$ such that the associated polyhedron P has at least one vertex. It is a fact that the varieties constructed from the four procedures described below depend only on P and not on the additional data, and this will be obvious in the cases of the first and fourth constructions. It is not at all obvious in the cases of the second or third construction, and in fact the examples that we get by allowing repeated or empty facets are very useful for illustrating what happens.

1.1 First construction

Let $T := V^*/V_{\mathbb{Z}}^*$. Since $V^* \cong \mathbb{R}^d$ and $V_{\mathbb{Z}}^* \cong \mathbb{Z}^d$, we have $T \cong (S^1)^d$. This isomorphism is not canonical, however, since we have no canonical choice of basis for $V_{\mathbb{Z}}$. What *is* true canonically? The Lie algebra of T (the tangent space to T at the identity) is naturally isomorphic to V^*. We also have natural group isomorphisms

$$\operatorname{Hom}(T, S^1) \cong \{f \in \operatorname{Hom}(V^*, \mathbb{R}) \mid f(V_{\mathbb{Z}}^*) \subset \mathbb{Z}\} \cong V_{\mathbb{Z}}$$

and

$$\operatorname{Hom}(S^1, T) \cong \{f \in \operatorname{Hom}(\mathbb{R}, V^*) \mid f(\mathbb{Z}) \subset V_{\mathbb{Z}}^*\} \cong V_{\mathbb{Z}}^*.$$

In particular, for each i, the element $a_i \in V^*_{\mathbb{Z}}$ determines a 1-dimensional subtorus $S^1 \cong T_i \subset T$. (One way to think of this subtorus is as the image of $\mathbb{R}a_i \subset V^*$ in $T = V^*/V^*_{\mathbb{Z}}$. The fact that it closes up follows from the fact that $a_i \in V^*_{\mathbb{Z}}$.)

For each face $F \subset P$, let

$$I_F := \{i \mid F \subset F_i\}$$

be the biggest set of facets with intersection F, and let

$$T_F := \prod_{i \in I_F} T_i \subset T$$

be the subtorus generated by T_i for all $i \in I_F$. The dimension of T_F is equal to the codimension of F in P. In particular, if $F = F_i$ is a facet, then $T_F = T_i$. If F is a vertex, then $T_F = T$.

Definition 1.1. Let $X_1(P) := P \times T/\sim$, where

$$(p_1, t_1) \sim (p_2, t_2) \qquad \text{if} \qquad p_1 = p_2 \in F \text{ and } t_1^{-1}t_2 \in T_F \text{ for some face } F \subset P.$$

More colloquially, we start with $P \times T$, and over every point $p \in P$, we divide by the subtorus $T_F \subset T$, where F is the largest face in which p lies.

Some properties are obvious from this definition: T acts on $X_1(P)$ with orbit space homeomorphic to P, $X_1(P)$ is compact if and only if P is a polytope. If a face F of P is regarded as a polyhedron in its own right, then we have a natural equivariant inclusion of $X_1(F)$ into $X_1(P)$. Many other properties are not obvious. For example, it's not immediately clear that $X_1(P)$ should have a complex algebraic structure or an action of the complexification of T. All of that will come later, via the other constructions. We will, however, give a sufficient criterion for $X_1(P)$ to be a smooth manifold.

Definition 1.2. We say that P is **simple** if, for every vertex $v \in P$, $|I_v| = d$. We say that P is **Delzant** if it is simple and, for every vertex $v \in P$, the set $\{a_i \mid i \in I_v\}$ spans not only the vector space V^* (which is automatic) but also the lattice V^*_Z. Thus simplicity is a purely combinatorial condition, but Delzantness is not.

Example 1.3. Give examples in which these conditions fail.

Theorem 1.4. *If P is Delzant, then $X_1(P)$ is a smooth manifold of dimension $2d$.*

Proof. A neighborhood of each vertex looks like $\mathbb{C}^d \cong \mathbb{R}^{2d}$. Thus one only needs to check that the transition functions are smooth, which they are. □

Remark 1.5. Though we won't prove it now, if P is simple, then $X_1(P)$ is an orbifold.

Example 1.6. Chopping off a corner of a Delzant polytope corresponds to taking a topological blow-up of the toric variety.

1.2 Second construction

Let $T^n = \mathrm{U}(1)^n$ be the n-dimensional coordinate torus. As we remarked in Section 1.1, each a_i defines a homomorphism from $\mathrm{U}(1)$ to T; taken all together, they define a homomorphism from T^n to T. The fact that P has at least one vertex implies that this homomorphism is surjective. Let $K \subset T^n$ be the kernel, so that we have a short exact sequence of tori

$$1 \to K \to T^n \to T \to 1. \tag{1}$$

Since K is an $(n-d)$-dimensional subgroup of T^n, it must be isomorphic to $\mathrm{U}(1)^{n-d}$ times some (possibly trivial) finite abelian group.

Just to start to get a feel for K, we'll prove the following two lemmas. Let $\mathfrak{k} = \mathrm{Lie}(K)$ and $\mathfrak{t}^n = \mathrm{Lie}(T^n) \cong \mathbb{R}^n$. The inclusion of K into T^n induces an inclusion of $\mathfrak{k}$ into $\mathfrak{t}^n$.

Lemma 1.7. *An n-tuple $(k_1, \ldots, k_n)$ lies in $\mathfrak{k}$ if and only if $\sum k_i a_i = 0$.*

Proof. The map $(k_1, \ldots, k_n) \mapsto \sum k_i a_i$ is the derivative of the map from T^n to T. □

Lemma 1.8. *The group K is connected if and only if $\{a_1, \ldots, a_n\}$ spans $V_{\mathbb{Z}}^*$ over the integers.*

Proof. This is just a big diagram chase around the following commutative diagram:

$$\begin{tikzcd}
 & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & \\
0 \arrow[r] & \mathfrak{k}_{\mathbb{Z}} \arrow[r] \arrow[d] & \mathfrak{t}^n_{\mathbb{Z}} \arrow[r] \arrow[d] & \mathfrak{t}_{\mathbb{Z}} = V^*_{\mathbb{Z}} \arrow[d] & \\
0 \arrow[r] & \mathfrak{k} \arrow[r] \arrow[d] & \mathfrak{t}^n \arrow[r] \arrow[d] & \mathfrak{t} = V^* \arrow[r] \arrow[d] & 0 \\
1 \arrow[r] & K \arrow[r] & T^n \arrow[r] \arrow[d] & T \arrow[r] \arrow[d] & 1 \\
 & & 1 & 1 &
\end{tikzcd}$$

The vertical sequence on the left is short-exact if and only if the horizontal sequence on top is short-exact. □

We now define a map

$$\mu : \mathbb{C}^n \to \mathfrak{k}^*$$

as follows. Let $i^* : (\mathfrak{t}^n)^* \to \mathfrak{k}^*$ the projection dual to the inclusion of $\mathfrak{k}$ into $\mathfrak{t}^n$. Since T^n is endowed with a set of coordinates, we may identify $(\mathfrak{t}^n)^*$ with $\mathbb{R}^n$. We then put

$$\mu(z_1, \ldots, z_n) := i^* \left(\frac{1}{2}|z_1|^2 - \lambda_1, \ldots, \frac{1}{2}|z_n|^2 - \lambda_n \right). \tag{2}$$

Note that T^n acts on $\mathbb{C}^n$ in the obvious way, and the map μ is invariant under the action of T^n. In particular, T^n (as well as its subtorus K) acts on each fiber of μ.

Definition 1.9. Let $X_2(P) := \mu^{-1}(0)/K$.

Since T^n acts on $\mu^{-1}(0)$, $T \cong T^n/K$ acts on $X_2(P) = \mu^{-1}(0)/K$. It is easy to check that the orbit space can be naturally identified with P, but we'll save that for Section 2. It is not immediately obvious when $X_2(P)$ is smooth. Indeed, two things could go wrong: $0 \in \mathfrak{k}^*$ could fail to be a regular value of μ, or the action of K on $\mu^{-1}(0)$ could fail to be free. In Section 2 we will show that neither of these things happen when P is Delzant. In this case, we'll show that $X_2(P)$ is a symplectic manifold, and that the action of T on $X_2(P)$ is the moment map.

1.3 Third construction

Let

$$1 \to K_{\mathbb{C}} \to T^n_{\mathbb{C}} \to T_{\mathbb{C}} \to 1.$$

be the complexification of Equation (1). More precisely, let $T_{\mathbb{C}} := V^*_{\mathbb{C}}/V^*_{\mathbb{Z}}$, so that $T_{\mathbb{C}} \cong (\mathbb{C}^\times)^d$ (non-canonnically), and we have natural isomorphisms

$$\mathrm{Hom}(T_{\mathbb{C}}, \mathbb{C}^\times) \cong \{f \in \mathrm{Hom}(V^*_{\mathbb{C}}, \mathbb{C}) \mid f(V^*_{\mathbb{Z}}) \subset \mathbb{Z}\} \cong V_{\mathbb{Z}}$$

and

$$\mathrm{Hom}(\mathbb{C}^\times, T_{\mathbb{C}}) \cong \{g \in \mathrm{Hom}(\mathbb{C}, V^*_{\mathbb{C}}) \mid g(\mathbb{Z}) \subset V^*_{\mathbb{Z}}\} \cong V^*_{\mathbb{Z}}.$$

The second isomorphism tells us how to use the elements $a_1, \ldots, a_n$ to define a surjective homomorphism $T^n_{\mathbb{C}} \to T_{\mathbb{C}}$, and we define $K_{\mathbb{C}}$ to be the kernel. The group $T^n_{\mathbb{C}}$, and therefore also the subgroup $K_{\mathbb{C}}$, acts on $\mathbb{C}^n$ in the obvious way, extending the actions of T^n and K. If $\mathfrak{k}_{\mathbb{C}} = \mathrm{Lie}(K_{\mathbb{C}})$, we again have $\mathfrak{k}_{\mathbb{C}} \subset \mathfrak{t}^n_{\mathbb{C}} \cong \mathbb{C}^n$, and Lemma 1.7 holds with $\mathfrak{k}$ replaced by $\mathfrak{k}_{\mathbb{C}}$.

Remark 1.10. Another way to express the relationship between $T_{\mathbb{C}}$ and $V_{\mathbb{Z}}$ is to observe that $T_{\mathbb{C}} \cong \mathrm{Spec}\,\mathbb{C}[V_{\mathbb{Z}}]$, where $\mathbb{C}[V_{\mathbb{Z}}]$ denotes the group algebra. (If you don't know what Spec means, then you can come back to this remark after reading Section 3.1.) Indeed, we have already seen that every element of $V_{\mathbb{Z}}$ defines a function on $T_{\mathbb{C}}$ with values in $\mathbb{C}^\times \subset \mathbb{C}$, and that the sum of two elements gives rise to the product of the two functions. By choosing an isomorphism between $T_{\mathbb{C}}$ and $(\mathbb{C}^\times)^d$, we can see that these functions form an additive basis for the ring of all functions on $T_{\mathbb{C}}$. The somewhat confusing fact is that the surjective map from $V^*_{\mathbb{C}}$ to $T_{\mathbb{C}}$ is not algebraic! After choosing a basis to identify $V^*_{\mathbb{C}}$ with $\mathbb{C}^d$, the map would have the form

$$(v_1, \ldots, v_d) \mapsto (e^{2\pi i v_1}, \ldots, e^{2\pi i v_d}),$$

which is not algebraic because the exponential map is not polynomial. What this is telling us is that, if we want to think of $T_{\mathbb{C}}$ as an algebraic variety, then $T_{\mathbb{C}} := \mathrm{Spec}\,\mathbb{C}[V_{\mathbb{Z}}]$ is the "correct" definition while $V^*_{\mathbb{C}}/V_{\mathbb{Z}}$ doesn't really make sense. If we want to think of $T_{\mathbb{C}}$ as a complex analytic space, the two definitions are equivalent.

For any $z = (z_1, \dots, z_n) \in \mathbb{C}_n$, let $F_z := \bigcap_{z_i=0} F_i$. This is a face of P. For most z, none of the coordinates is zero, so $F_z = P$ is the entire polyhedron. The more coordinates of z that vanish, the smaller F_z will be. In particular, F_z is sometimes equal to the empty face. Let

$$U_\lambda := \{z \in \mathbb{C}^n \mid F_z \neq \emptyset\}.$$

This is clearly a Zariski-open subset of $\mathbb{C}^n$; it is the complement of some coordinate subspaces, therefore it has an action of the torus $T^n_{\mathbb{C}}$. Note that $K_{\mathbb{C}}$ (like K) depends only on the parameters $a_1, \dots, a_n$, while U_λ (like $\mu^{-1}(0)$) depends only on the parameters $\lambda_1, \dots, \lambda_n$.

Definition 1.11. Let $X_3(P) := U_\lambda/\sim$, where

$$z \sim z' \qquad \text{if} \quad \overline{K_{\mathbb{C}} \cdot z} \cap \overline{K_{\mathbb{C}} \cdot z'} \neq \emptyset.$$

Example 1.12. Explain the difference between this quotient and a geometric quotient. Note that when the action of $K_{\mathbb{C}}$ is locally free (which is often the case), then it really is a geometric quotient.

It is not immediately obvious that $\sim$ is an equivalence relation or that $X_3(P)$ is Hausdorff; both of these will follow from Proposition 4.10. It is clear that the complex torus $T_{\mathbb{C}} \cong T^n_{\mathbb{C}}/K_{\mathbb{C}}$ acts on $X_3(P) = U_\lambda/K_{\mathbb{C}}$. The following proposition gives a beautiful characterization of the orbits of this action. For any $z \in U_\lambda$, let $[z] \in X_3(P)$ be the equivalence class of z.

Proposition 1.13. *The assignment $[z] \mapsto F_z$ gives a bijection between $T_{\mathbb{C}}$-orbits in $X_3(P)$ and nonempty faces of P.*

In order to prove Proposition 1.13, and again to prove Theorems 5.3 and 5.4 we will need the following fact, whose proof eludes me at the moment.

Fact 1.14. Let G be an algebraic group acting linearly on a vector space W. Let z be an element of W, and z' another element of W that lies in the closure of the orbit $G \cdot z$. Then there exists an element $g \in G$ and a homomorphism $\rho \in \mathrm{Hom}(\mathbb{C}^\times, G)$ such that $z' = \lim_{t\to\infty} \rho(t)g \cdot z$.

Proof of Proposition 1.13. To see that the assignment $[z] \mapsto F_z$ is well-defined, we must show that if $z \sim z'$, then $F_z = F_{z'}$. It is enough to prove this when $z' \in \overline{K_{\mathbb{C}} \cdot z}$. Write

$$z = (z_1, \dots, z_n) \qquad \text{and} \qquad z' = (z'_1, \dots, z'_n).$$

If $z' \in \overline{K_{\mathbb{C}} \cdot z}$, then Fact 1.14 tells us that there exists an element $k = (k_1, \dots, k_n) \in \mathfrak{k}_{\mathbb{C}} \subset \mathfrak{t}^n \cong \mathbb{C}^n$ such that

$$z' = \lim_{t\to\infty} e^{tk} \cdot z = \lim_{t\to\infty} (e^{tk_1} z_1, \dots, e^{tk_n} z_n).$$

This means that, for every i, at exactly one of the following conditions must hold:

(a) $z_i = z_i' = 0$

(b) $z_i = z_i' \neq 0$ and $k_i = 0$

(c) $z_i \neq z_i' = 0$ and $k_i < 0$.

Write $\{1, \ldots, n\} = A \sqcup B \sqcup C$, where A is the set of all i satisfying condition (a), B is the set of all i satisfying condition (b), and C is the set of all i satisfying condition (c). To show that $F_z = F_{z'}$, we need to show that for all $i \in C$ and $v \in F_z$, $a_i(v) + \lambda_i = 0$.

Choose an arbitrary $v \in F_z$ as well as an arbitrary $v' \in F_{z'}$; the fact that such a v' exists follows from the fact that $z' \in U_\lambda$. We have

$$\begin{aligned}
\sum_{i \in C} k_i(a_i(v) + \lambda_i) &= \sum_{i \in C} k_i a_i(v) + \sum_{i \in C} k_i \lambda_i \\
&= -\sum_{i \in A} k_i a_i(v) + \sum_{i \in C} k_i \lambda_i && \text{because } \sum k_i a_i = 0 \text{ and } k_i = 0 \text{ for all } i \in B \\
&= \sum_{i \in A} k_i \lambda_i + \sum_{i \in C} k_i \lambda_i && \text{because } v \in F_z \text{ and } z_i = 0 \text{ for all } i \in A \\
&= \sum_{i \in A} k_i a_i(v') + \sum_{i \in C} k_i a_i(v') && \text{because } v' \in F_{z'} \text{ and } z_i' = 0 \text{ for all } i \in A \cup C \\
&= \sum_{i=1}^{n} k_i a_i(v') && \text{because } k_i = 0 \text{ for all } i \in B \\
&= 0 && \text{because } \sum k_i a_i = 0.
\end{aligned}$$

Since we know that $a_i(v) + \lambda_i \geq 0$ for all i and $k_i < 0$ for all $i \in C$, the only way this is possible is if $a_i(v) + \lambda_i = 0$ for all $i \in C$. Thus $[z] \mapsto F_z$ is well-defined.

Surjectivity of this assignment is obvious; we now prove injectivity. Let z be in arbitrary element of U_λ, and let $I_z := \{i \mid z_i = 0\} \subset I_{F_z}$. It follows from the definitions of F_z and I_{F_z} that the two sets $\{a_i \mid i \in I_z\} \subset \{a_i \mid i \in I_{F_z}\}$ have the same linear span. In particular, this means that we can find an element $k = (k_1, \ldots, k_n) \in \mathfrak{k}_\mathbb{C}$ such that $k_i = -1$ for all $k \in I_{F_z} \smallsetminus I_z$ and $k_i = 0$ for all $i \notin I_{F_z}$. Let

$$z' := \lim_{t \to \infty} e^{tk} \cdot z.$$

Then $z_i' = 0$ if $i \in I_{F_z}$ and $z_i' = z_i \neq 0$ otherwise. It is clear that $F_z = F_{z'}$, hence $z' \in U_\lambda$ and $z \sim z'$. Now hit z' with an element of $T_\mathbb{C}^n$ to obtain a new element $z'' \in U_\lambda$ such that $z_i'' = 0$ if $i \in I_{F_z}$ and $z_i'' = 1$ otherwise. It is clear from the definition that $[z'']$ lies in the same $T_\mathbb{C}$-orbit as $[z'] = [z]$, and that $[z'']$ is completely determined by F_z. hence the $T_\mathbb{C}$-orbit of $[z]$ is completely determined by F_z. $\square$

Remark 1.15. We can also show that the bijection of Proposition 1.13 is compatible with dimension and closure relations. That is, a $T_\mathbb{C}$-orbit of complex dimension r is sent to a face of real dimension r, and one orbit lies in the closure of another if and only if its associated face lies in the closure of the other face. Though it is probably possible to prove these facts directly, it will follow easily from the equivalence of the first three constructions.

1.4 Fourth construction

Let $\Sigma := \{(v, r) \in V \times \mathbb{R}_{\geq 0} | \mid a_i(v) + \lambda_i r \geq 0 \text{for all } i\}$. The intersection $\Sigma_0 := \Sigma \cap V \times \{0\}$ is sometimes called the **cone of unbounded directions** in P. The complement $\Sigma \smallsetminus \Sigma_0$ is the union of all of the rays through $P \subset V \times \{1\} \subset V \times \mathbb{R}$, and Σ is the closure of this locus. Consider the semigroup $S_P := \Sigma \cap (V_{\mathbb{Z}} \times \mathbb{N})$ along with its semigroup ring $\mathbb{C}[S_P]$, which is graded by $\mathbb{N}$.

Example 1.16. A standard interval, a standard triangle, a ray, a square.

Definition 1.17. Let $X_4(P) := \operatorname{Proj} \mathbb{C}[S_P]$.

Proj will be defined in Section 3, but for the time being it is not so difficult to give an idea of what's going on, at least in certain special (but common) cases. The simplest case is when $\mathbb{C}[S_P]$ is a polynomial ring generated in degree 1; this happens when S_P is equal to the standard simplex in $\mathbb{R}^d$, in which case we have $d+1$ generators. In this case

$$\operatorname{Proj} \mathbb{C}[S_P] = \operatorname{Proj} \mathbb{C}[x_0, \ldots, x_n] = \mathbb{P}^n$$

(for now, take this as a definition of Proj). Now suppose that $\mathbb{C}[S_P]$ is generated in degree 1, which means that it is isomorphic to $\mathbb{C}[x_0, \ldots, x_n]/I_P$ for some homogeneous ideal I_P. Then $\operatorname{Proj} \mathbb{C}[S_P]$ is the variety cut out of $\mathbb{P}^n$ by the polynomials in I_P.

Example 1.18. The square.

Next suppose that $\mathbb{C}[S_P]$ is a polynomial ring generated in degrees 0 and 1, that is,

$$\mathbb{C}[S_P] = \mathbb{C}[x_0, \ldots, x_n, y_1, \ldots, y_m],$$

with $\deg x_i = 1$ and $\deg y_i = 0$. This happens, for example, in the case of the ray. Then

$$\operatorname{Proj} \mathbb{C}[S_P] = \mathbb{P}^n \times \mathbb{C}^m;$$

again, this can be taken as a definition for now. More generally, if $\mathbb{C}[S_P]$ is a quotient of the ring $\mathbb{C}[x_0, \ldots, x_n, y_1, \ldots, y_m]$ by an ideal I_P that is homogeneous in the y variables, then $\operatorname{Proj} \mathbb{C}[S_P]$ is the subvariety cut out of $\mathbb{P}^n \times \mathbb{C}^m$ by this ideal.

Example 1.19. The ray. More generally, any cone.

Example 1.20. The blow-up of $\mathbb{C}^2$ at a point. Go back and treat this in the context of the three other constructions, as well.

From this construction we see that $X_4(P)$ is an algebraic variety (and that it is projective over the affine variety $X_4(\Sigma_0)$). In particular, $X_4(P)$ is compact if and only if $\Sigma_0 = \{0\}$, which is equivalent to the condition that P is a polytope. Furthermore, the $V_{\mathbb{Z}}$-grading of $\mathbb{C}[S_P]$ induces an action of $T_{\mathbb{C}}$ on $\mathbb{C}[S_P]$ and therefore also on $X_4(P)$. On the other hand, the questions of smoothness and classification of the $T_{\mathbb{C}}$-orbits are somewhat opaque from this perspective.

2 Symplectic reduction: the first and second constructions

In this section we'll introduce the concept of symplectic reduction, and use it to better understand the second construction in the previous section. We'll show that if P is Delzant, then $X_2(P)$ is a symplectic manifold, and the action of T on $X_2(P)$ is hamiltonian. In this case $X_1(P)$ is diffeomorphic to $X_2(P)$, and this diffeomorphism identifies the projection from $X_1(P)$ to P with the moment map for the T-action.

2.1 The homeomorphism

The proof that $X_1(P)$ and $X_2(P)$ are T-equivariantly homeomorphic is quite elementary; we don't need to assume that P is Delzant, nor do we need to use the language of symplectic geometry, even though a seasoned symplectic geometer will immediately spot the moment maps lurking within the proof. Just to make the point that this is not hard, we'll get it out of the way right away.

The projection from T^n to T induces a projection $\pi : \mathfrak{t}^n \to \mathfrak{t}$ and an inclusion $\pi^* : \mathfrak{t}^* \to (\mathfrak{t}^n)^*$. Recall that we have identified $\mathfrak{t}^*$ with V and $(\mathfrak{t}^n)^*$ with $\mathbb{R}^n$. With these identifications, we have

$$\pi^*(v) = \big(a_1(v), \ldots, a_n(v)\big) \qquad \text{for all } v \in V.$$

Consider the shifted map $\pi^*_\lambda : V \to \mathbb{R}^n$ defined by the formula

$$\pi^*_\lambda(v) := \big(a_1(v) + \lambda_1, \ldots, a_n(v) + \lambda_n\big) \qquad \text{for all } v \in V.$$

It is immediate from the definition that P is the preimage of $\mathbb{R}^n_{\geq 0}$ along the map π^*_λ. Thus we will use π^*_λ to identify P with a subset of $\mathbb{R}^n_{\geq 0}$. We then have

$$\begin{aligned} P &= \big\{(r_1, \ldots, r_n) \in \mathbb{R}^n_{\geq 0} \mid (r_1, \ldots, r_n) \in \operatorname{im} \pi^*_\lambda\big\} \\ &= \big\{(r_1, \ldots, r_n) \in \mathbb{R}^n_{\geq 0} \mid (r_1 - \lambda_1, \ldots, r_n - \lambda_n) \in \operatorname{im} \pi^*\big\} \\ &= \big\{(r_1, \ldots, r_n) \in \mathbb{R}^n_{\geq 0} \mid (r_1 - \lambda_1, \ldots, r_n - \lambda_n) \in \ker i^*\big\} \\ &= \big\{(r_1, \ldots, r_n) \in \mathbb{R}^n_{\geq 0} \mid i^*(r_1 - \lambda_1, \ldots, r_n - \lambda_n) = 0\big\}. \end{aligned}$$

Consider the map $f : P \times T^n \to \mathbb{C}^n$ defined by the formula

$$f\Big((r_1, \ldots, r_n), (t_1, \ldots, t_n)\Big) := \big(t_1\sqrt{2r_1}, \ldots, t_n\sqrt{2r_n}\big).$$

It is clear from Equation (2) and the description of P above that f is a surjection onto $\mu^{-1}(0)$, and that it descends to a surjection

$$\bar{f} : P \times T \twoheadrightarrow \mu^{-1}(0)/K = X_2(P).$$

This map is not quite injective, but the failure of injectivity is easy to measure. If $r_i = 0$, then it doesn't matter what t_i is. In terms of P, the condition $r_i = 0$ is equivalent to $(r_1, \ldots, r_n)$ lying on the facet F_i. In terms of T, the statement that "it doesn't matter what t_i is" means that we identify two elements of T whose ratio lies in the subtorus $T_i \subset T$. In other words, two elements of $P \times T$ are sent to the same element of $X_2(P)$ if and only if they are equivalent in the sense of Definition 1.1. Thus $\bar{f}$ descends to a continuous bijection from $X_1(P)$ to $X_2(P)$.

If P is a polytope, and therefore $X_1(P)$ is compact, this proves that $X_1(P)$ is homeomorphic to $X_2(P)$. In the general case where P is allowed to be a polyhedron, we need to prove that the inverse map is continuous. Since $\mu^{-1}(0)$ is proper over $X_2(P)$, it suffices to prove that the composition

$$\mu^{-1}(0) \to X_2(P) \to X_1(P)$$

is continuous. This composition lifts to a map

$$\mu^{-1}(0) \to P \times T^n$$

given by the formula

$$(z_1, \ldots, z_n) \mapsto \left((\pi_\lambda^*)^{-1} \left(\frac{1}{2}|z_1|^2, \ldots, \frac{1}{2}|z_n|^2 \right), \left(e^{i \arg(z_1)}, \ldots, e^{i \arg(z_n)} \right) \right),$$

which is obviously continuous. Thus we have proven the following theorem.

Theorem 2.1. *The spaces $X_1(P)$ and $X_2(P)$ are homeomorphic. In particular, $X_2(P)$ depends only on the polyhedron P, and not on the additional data.*

Remark 2.2. We have used the word "homeomorphism" everywhere because both $X_1(P)$ and $X_2(P)$ can be singular. We know that $X_1(P)$ is smooth when P is Delzant (Theorem 1.4), and we'll prove the same statement about $X_2(P)$. It is clear that when both spaces are smooth, our homeomorphism is in fact a diffeomorphism.

2.2 Symplectic manifolds and hamiltonian actions

Definition 2.3. Let X be a real manifold. A **symplectic form** on X is a closed, non-degenerate 2-form ω on X. The pair (X, ω) is called a **symplectic manifold**. Often we will abuse notation and refer to X itself as a symplectic manifold.

What exactly is a closed, non-degenerate 2-form? "2-form" means that at every point $x \in X$, we have a skew-symmetric bilinear form $\omega_x : T_xX \times T_xX \to \mathbb{R}$ that varies smoothly with x. "Non-degenerate" means that for every $x \in X$ and $p \in T_xX$, there exists a vector $q \in T_xX$ such that $\omega_x(p,q) \neq 0$. In particular, ω_x provides an isomorphism from T_xX to T_x^*X by sending p to $\omega_x(p,-)$. If we consider every $x \in X$ at once, ω provides an isomorphism from the space of vector fields (sections of the tangent bundle) to the space of 1-forms (sections of the cotangent bundle). "Closed" means that the 3-form $d\omega$ is equal to zero; since this is not a course on differential topology and this condition will not be particularly relevant to our discussion, I won't bother to remind you what it means to take the exterior derivative of a differential form (though this is certainly something that everybody should know).

Example 2.4. Consider the vector space $\mathbb{C}^n$. At every point $z \in \mathbb{C}^n$, we have a canonical identification $T_z\mathbb{C}^n \cong \mathbb{C}^n$, so we only need to give one pairing. We put

$$\omega(z,w) := \operatorname{Im}(z_1)\operatorname{Re}(w_1) - \operatorname{Re}(z_1)\operatorname{Im}(w_1) + \ldots + \operatorname{Im}(z_n)\operatorname{Re}(w_n) - \operatorname{Re}(z_n)\operatorname{Im}(w_n).$$

If you're comfortable with the language of differential forms and you identify $\mathbb{C}^n$ with $\mathbb{R}^{2n}$, then we have just described the 2-form $dy_1 \wedge dx_1 + \ldots + dy_n \wedge dx_n$.

Example 2.5. More abstractly, if W is a complex vector space with a hermitian inner product, the imaginary part of that inner product is a symplectic form. The standard inner product on $\mathbb{C}^n$ is

$$(z,w) := z_1\bar{w}_1 + \ldots + z_n\bar{w}_n,$$

and the imaginary part is exactly the form described in Example 2.4.

Example 2.6. Consider the 2-dimensional sphere S^2. We define the symplectic form $\omega_x(y,z) := x \cdot (y \times z)$ whenever $x \in S^2 \subset \mathbb{R}^3$ is a unit vector and $y, z \in T_xS^2 = \{w \in \mathbb{R}^3 \mid x \cdot w = 0\}$. It is a fun exercise in differential forms to check that if h is the height (latitude) function and θ is the angle (longitude) function, then $\omega = d\theta \wedge dh$. (Note that θ isn't *really* a function because it is multi-valued, but $d\theta$ is a perfectly good 1-form.)

What's so good about symplectic manifolds? Let (X,ω) be a symplectic manifold, and let $f : X \to \mathbb{R}$ be a smooth function. Then df is a 1-form, and we can use the symplectic form to convert df to a vector field $\hat{f}$ such that, for all $x \in X$ and $q \in T_xX$, $\omega_x(\hat{f}_x, q) = df_x(q)$. Flowing along the vector field $\hat{f}$ preserves both the function f and the symplectic form ω. Indeed, the directional derivative of f at x along the vector $\hat{f}_x$ is $df_x(\hat{f}_x) = \omega_x(\hat{f}_x, \hat{f}_x) = 0$. The fact that flowing along $\hat{f}$ preserves ω follows from the fact that ω is closed by applying Cartan's magic formula to the Lie derivative of ω along $\hat{f}$. We'll suppress this computation in order to minimize the differential topology, but it is useful to think about the example of the height function on S^2. The function f is called a **hamiltonian** for the vector field $\hat{f}$. Note that a vector field v admits a hamiltonian if and only if the 1-form $\omega(v,-)$ is exact, in which case the hamiltonian is unique up to a locally constant function.

Let K be a real Lie group acting on X. For every point $x \in X$, we get a map $\rho_x : K \to X$ by letting each element of K act on $x \in X$. Differentiating this map gives us a linear map $d\rho_x : \mathfrak{k} \to T_xX$. This means that for every $k \in \mathfrak{k}$, we get a vector field $\hat{k}$ defined by putting

$\hat{k}_x := d\rho_x(k)$. One way to think about this is to say that an element of K gives an automorphism of X, so an infinitesimal element of K (an element of $\mathfrak{k}$) gives an infinitesimal automorphism of X (a vector field).

Definition 2.7. A **moment map** for the action of K on X is a smooth map $\mu : X \to \mathfrak{k}^*$ such that for all $k \in \mathfrak{k}$, the smooth function $\mu_k := \mu(-)(k) : X \to \mathbb{R}$ is a hamiltonian for $\hat{k}$. We also require that μ is equivariant with respect to the coadjoint action on $\mathfrak{k}^*$. (Note that when K is abelian, this means that μ is K-invariant.)

Example 2.8. The circle group $T^1 \cong \mathrm{U}(1)$ acts on $\mathbb{C}$ with moment map $\Phi : \mathbb{C}^n \to (\mathfrak{t}^n)^* \cong \mathbb{R}$ given by

$$\Phi(z) = \frac{1}{2}|z|^2 - \lambda$$

for any $\lambda \in \mathbb{R}$. More generally, T^n acts on $\mathbb{C}$ with moment map $\Phi : \mathbb{C}^n \to (\mathfrak{t}^n)^* \cong \mathbb{R}^n$ given by

$$\Phi(z_1, \dots, z_n) = \left(\frac{1}{2}|z_1|^2 - \lambda_1, \dots, \frac{1}{2}|z_n|^2 - \lambda_n\right).$$

Example 2.9. If H acts on X with moment map $\Phi : X \to \mathfrak{h}^*$ and $K \subset H$ is a subgroup with inclusion $i : \mathfrak{k} \to \mathfrak{h}$ and projection $i^* : \mathfrak{h}^* \to \mathfrak{k}^*$, then

$$\mu := i^* \circ \Phi : X \to \mathfrak{k}^*$$

is a moment map for the action of K on X. In particular, Equation (2) is a moment map for the action of K on $\mathbb{C}^n$. This holds more generally for group homomorphisms, not just inclusions.

Example 2.10. In this example we'll describe the moment map for the action of $\mathrm{U}(n)$ on $\mathbb{C}^n$, which gives us (via Example 2.9) the moment map for any linear action of a compact group. First of all, we should remind ourselves what $\mathfrak{u}(n)^*$ looks like. The group $\mathrm{U}(n)$ is the set of matrices that are inverse to their adjoints, which means that the Lie algebra $\mathfrak{u}(n)$ is the set of skew-adjoint matrices. We identify $\mathfrak{u}(n)^*$ with $\mathfrak{u}(n)$ via the Killing form, which is just given by the trace of the product.

We claim that the map

$$\mu : z \mapsto -\frac{i}{2} z z^*$$

is a moment map for the action of $\mathrm{U}(n)$ on $\mathbb{C}^n$. Equivariance is clear. For every $z \in \mathbb{C}^n$, the map $d\rho_z : \mathfrak{u}(n) \to \mathbb{C}^n$ is given by $A \mapsto Az$. This means that we need to show that the function

$$\mu_A : z \mapsto -\frac{i}{2} \operatorname{tr}(Azz^*)$$

is a hamiltonian function for thr field whose value at z is Az. The derivative is

$$\begin{aligned} d\mu_A : w \mapsto -\frac{i}{2}\operatorname{tr}\left(A(zw^* + wz^*)\right) &= -\frac{i}{2}\operatorname{tr}\left(Azw^*\right) - \frac{i}{2}\operatorname{tr}\left(Awz^*\right) \\ &= -\frac{i}{2}\operatorname{tr}\left(Azw^*\right) + \frac{i}{2}\overline{\operatorname{tr}\left(zw^*A\right)} \\ &= -\frac{i}{2}\operatorname{tr}\left(w^*Az\right) + \frac{i}{2}\overline{\operatorname{tr}\left(w^*Az\right)} \\ &= \operatorname{Im}\circ\operatorname{tr}\left(w^*Az\right). \end{aligned}$$

On the other hand, we have

$$\omega_z(A_z, w) = \operatorname{Im}\left(Azw^*\right) = \operatorname{Im}\circ\operatorname{tr}\left(w^*Az\right),$$

so it works!

Let's use this to rederive the formula in Example 2.8. The map from $\mathbb{R}^n$ to $\mathfrak{u}(n)$ induced by the inclusion of T^n into $\mathrm{U}(n)$ is given by $(t_1, \ldots, t_n) \mapsto \operatorname{diag}(it_1, \ldots, it_n)$, and its dual is given by multiplying by i and projecting onto the diagonal. Yup, it works.

Example 2.11. Consider the Lie group $\mathrm{SO}(3)$ of orientation preserving isometries of $\mathbb{R}^3$. The Lie algebra of $\mathrm{SO}(3)$ is 3-dimensional, and one can check that the coadjoint action is isomorphic to the standard action of $\mathrm{SO}(3)$ on $\mathbb{R}^3$. For any positive real number r, let S^2_r be the manifold S^2 with symplectic form equal to r times the form defined in Example 2.6. It turns out that the inclusion of S^2_r into $\mathbb{R}^3$ as the sphere of radius r is the (unique) moment map for this action.

More generally, a coadjoint orbit always admits the structure of a symplectic manifold, and the inclusion of that coadjoing orbit into the dual of the Lie algebra is a moment map for the coadjoint action.

Example 2.12. In the classical example from physics, X is the manifold that parameterizes all possible states in which the universe can be at any given instant (which for some reason is supposed to be symplectic), $K = \mathbb{R}$, the action of K on X is "the march of time", and the moment map $\mu : X \to \mathfrak{k} \cong \mathbb{R}$ is "energy".

Note that the moment map, when it exists, is unique up to a K-invariant locally constant function. However, a moment map need not exist.

Example 2.13. The action of $\mathrm{U}(1)$ on T^2.

2.3 Symplectic reduction

Suppose that K acts on a symplectic manifold X with moment map $\mu : X \to \mathfrak{k}^*$. Since μ is K-equivariant, K acts on the pre-image of the K-fixed point $0 \in \mathfrak{k}^*$. We define the **symplectic reduction** of X by K to be the quotient $X/\!\!/K := \mu^{-1}(0)/K$. Note that this definition depends not just on X and K, but also on the choice of moment map μ. We have already seen an example of a symplectic reduction: the space $X_2(P)$ defined in Section 1.2 is the symplectic reduction of $\mathbb{C}^n$ by K using the moment map defined in Equation (2) and Example 2.9.

As we pointed out in Section 1.2, there are two things that could prevent $X/\!\!/K$ from being a manifold. First, 0 could fail to be a regular value of μ, in which case $\mu^{-1}(0)$ might not be a manifold. Second, even if $\mu^{-1}(0)$ is a manifold, K might fail to act freely on it, in which case the quotient could be singular. In fact,we'll see that these two potential problems are related.

We say that the action of K at a point $x \in X$ is **locally free** if the stabilizer of x is discrete. This is equivalent to saying that the Lie algebra of the stabilizer is trivial, or that the map $d\rho_x : \mathfrak{k} \to T_xX$ is injective.

Proposition 2.14. *The point $0 \in \mathfrak{k}^*$ is a regular value of μ if and only if K acts locally freely on $\mu^{-1}(0)$. In particular, if K acts freely on on $\mu^{-1}(0)$, then $X/\!\!/K$ is a manifold of dimension $\dim X - 2\dim K$.*

Proof. The statement that 0 is a regular value is the statement that the differential $d\mu_x$ is surjective for every $x \in \mu^{-1}(0)$, or equivalently that $\ker d\mu_x$ has dimension $\dim X - \dim K$. We have

$$\begin{aligned} \ker d\mu_x &= \{p \mid d\mu_x(p) = 0\} \\ &= \{p \in T_xX \mid d\mu_x(p)(k) = 0\} \qquad \text{for all } k \in \mathfrak{k} \\ &= \{p \in T_xX \mid \omega_x(\hat{k}_x, p) = 0\} \qquad \text{for all } k \in \mathfrak{k} \\ &= d\rho_x(\mathfrak{k})^\perp, \end{aligned}$$

which has dimension equal to $\dim X - \dim d\rho_x(\mathfrak{k})$. This is equal to $\dim X - \dim K$ if and only if $d\rho_x$ is injective, that is, if and only if K acts locally freely at x. □

Example 2.15. Using Examples 2.9 and 2.11, we can construct a moment map for the diagonal action of SO(3) on $\prod_{i=1}^n S^2_{r_i}$. Taking the quotient, we obtain the moduli space of spatial polygons with edge lengths $(r_1, \ldots, r_n)$. It is smooth if and only if there is no subset of edges with total length exactly equal to $\frac{1}{2}\sum r_i$.

The computation in the proof of Proposition 2.14 can be used to show that the symplectic reduction $X/\!\!/K$ has a natural symplectic form.

Theorem 2.16 (Marsden-Weinstein)**.** *If K acts freely on $\mu^{-1}(0)$, then $X/\!\!/K$ has a unique symplectic form ω_{red} with the property that the pullback of ω_{red} to $\mu^{-1}(0)$ coincides with the restriction of ω.*

Proof. Given a point $x \in \mu^{-1}(0)$, let $[x]$ denote its image in $X/\!\!/K$. Then we have a natural identification $T_{[x]}X/\!\!/K \cong \ker d\mu_x/d\rho_x(\mathfrak{k})$. Since $d\rho_x(\mathfrak{k})$ is the perpendicular space to $\ker d\mu_x$, ω_x descends to a non-degenerate skew-symmetric pairing on $T_{[x]}X/\!\!/K$. Both closedness and uniqueness follow from the fact that pulling back forms from $X/\!\!/K$ to $\mu^{-1}(0)$ is an injective operation. □

Suppose, as in Example 2.9, that H acts on X with moment map $\Phi : X \to \mathfrak{h}$ and $K \subset H$ is a subgroup, which we now assume to be normal. Let $i : \mathfrak{k} \to \mathfrak{h}$ be the inclusion, and consider the moment map $\mu := i^* \circ \Phi$ for the action of K on X. Then the quotient group K/H acts on $X/\!\!/K$, and Φ descends to a map

$$\bar{\Phi} : X/\!\!/K \to \ker(i^*) \cong (\mathfrak{h}/\mathfrak{k})^* \cong \mathrm{Lie}(H/K)^*.$$

The following proposition is straightforward to check, and is a good exercise to make sure that you understand the definitions.

Proposition 2.17. *The map $\bar{\Phi}$ is a moment map for the action of H/K on $X/\!\!/K$.*

2.4 The toric case

Let P be a rational polyhedron, and consider the action of K on $\mathbb{C}^n$ with moment map $\mu : \mathbb{C}^n \to \mathfrak{k}^*$ as in Section 1.2 and Example 2.9.

Proposition 2.18. *K acts locally freely on $\mu^{-1}(0)$ if and only if P is simple. K acts freely on $\mu^{-1}(0)$ if and only if P is Delzant.*

Proof. In Section 2.1, we implicitly showed that $\mu^{-1}(0)$ is K-equivariantly homeomorphic to the quotient space $P \times T^n / \sim$, where $\sim$ identifies (p, t) with (p, t') if $p \in F$ and t/t' lies in the coordinate subtorus $T^{I_F} \subset T^n$. (What we actually proved was the same statement *after* dividing by K, but the "upstairs" proof is the same.) Thus we only need to show that for every face F, $K \cap T^{I_F}$ is finite if and only if P is simple and trivial if and only if P is Delzant. It is clearly sufficient to prove this for vertices. The subtorus $K \subset T^n$ is by definition the kernel of the map from T^n to T. The hypothesis that P is Delzant says exactly that the restriction of this map to T^{I_v} is an isomorphism for every vertex v of P. The slightly weaker hypothesis that P is simple says that this restriction is a surjective homomorphism between two tori of dimension d, and therefore has a finite kernel. □

Corollary 2.19. *If P is Delzant, then the toric variety $X_2(P)$ is symplectic, and the composition $X_2(P) \cong X_1(P) \to P$ is a moment map for the action of T.*

Proof. The fact that $X_2(P)$ is a symplectic manifold follows from Theorem 2.16 and Proposition 2.18. To see that $X_2(P) \cong X_1(P) \to P$ is a moment map, Proposition 2.17 tells us that we only need to check that $\mu^{-1}(0) \to X_2(P) \cong X_1(P) \to P \subset V \cong \mathfrak{t}^* \subset (\mathfrak{t}^n)^* \cong \mathbb{R}^n$ is given by the formula

$$(z_1, \ldots, z_n) \mapsto \left(\frac{1}{2}|z_1|^2 - \lambda_1, \ldots, \frac{1}{2}|z_1|^2 - \lambda_n \right).$$

But this formula was established in the course of proving Theorem 2.1. □

Remark 2.20. If P is only simple, then $X_2(P)$ is a **symplectic orbifold**.

3 A quick tour of projective algebraic geometry

In this section I'll try to go over just enough algebraic geometry to prepare you for geometric invariant theory. The main focus will be on the operation Proj, which eats a graded ring an produces an algebraic variety.

3.1 Spec and Proj

Let R be a finitely generated integral domain over $\mathbb{C}$. This means that for some positive integer n, there exists a surjection $\mathbb{C}[x_1, \ldots, x_n] \twoheadrightarrow R$ whose kernel $I \subset \mathbb{C}[x_1, \ldots, x_n]$ is a prime ideal. We define $\operatorname{Spec} R$ to be the subset of $\mathbb{C}^n$ on which all of the polynomials in I vanish. You should think of $\operatorname{Spec} R$ as "the simplest space whose ring of algebraic functions is equal to R." Of course $\operatorname{Spec} R$ is not just a set; it has two important topologies, namely the Zariski topology and the analytic topology, along with sheaves of rings to go with these topologies. In these notes I'm going to ignore these technicalities and ask you to rely on your intuition.

The definition that we gave isn't very good because it involved choosing a generating set for R. The better definition is that $\operatorname{Spec} R$ is the set of maximal ideals in R. Of course, a maximal ideal of R is nothing more than a ring homomorphism from R to $\mathbb{C}$, and such homomorphisms are in bijection with n-tuples of complex numbers (the images of $x_1, \ldots, x_n$ on which the polynomials in I vanish. As an example, let $V_{\mathbb{Z}}$ be a rank d lattice, and let $R = \mathbb{C}[V_{\mathbb{Z}}^*]$. If $\{v_1, \ldots, v_d\}$ is a basis for $V_{\mathbb{Z}}^*$, then $\{\pm v_1, \ldots, \pm v_d\}$ is a generating set for R, and we obtain an embedding of $\operatorname{Spec} R$ into $\mathbb{C}^{2d}$. However, as we explained in Remark 1.10, it is most natural to think of R as the ring of algebraic functions on $T_{\mathbb{C}}$. That said, in most of our examples R will have an obvious set of generators.

Remark 3.1. One advantage of this definition is that it is clear that Spec is a contravariant functor: if $f : R \to S$ is a ring homomorphism, then we can pull back maximal ideals from S to R, so we get an induced map $f^* : \operatorname{Spec} S \to \operatorname{Spec} R$. On the other hand, an algebraic map of varieties from X to Y induces a pull back from functions on X to functions on Y. Indeed, the functor Spec and its adjoint Fun are inverse equivalences between the category of finitely generated integral domains over $\mathbb{C}$ and the category of irreducible affine algebraic varieties.

Now suppose that R is equipped with a grading

$$R = \bigoplus_{m=-\infty}^{\infty} R_m.$$

This allows us to define an action of the group $\mathbb{C}^\times$ on the ring R by putting $t \cdot f := t^m f$ for all $t \in \mathbb{C}^\times$ and $f \in R_m$. This in turn induces an action of $\mathbb{C}^\times$ on $\operatorname{Spec} R$. We refer to an element of R_m as a function of weight m on $\operatorname{Spec} R$. The functions of weight 0 are exactly the functions that are constant along orbits.

Example 3.2. $\mathbb{C}[x]$, $\mathbb{C}[x_0, \ldots, x_n]$, $\mathbb{C}[x, x^{-1}]$, $\mathbb{C}[x, y, z, w]/\langle xw - yz\rangle$, all with various gradings.

What is the fixed point set of the action of $\mathbb{C}^\times$ on $\operatorname{Spec} R$? A point $x \in \operatorname{Spec} R$ is fixed if and only if $t \cdot x = x$ for all $t \in \mathbb{C}^\times$, which is the same as saying that

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

for every $f \in R$. If $f \in R_0$, this equation is trivially satisfied. On the other hand, if $f \in R_m$ for some $m \neq 0$, then this equation is equivalent to the equation $f(x) = 0$. Thus the fixed point set is the vanishing set of the ideal generated by all functions of nonzero weight.

Example 3.3. Revisit all of the examples above. Also $S[y]$ for any S.

Now suppose that $R_m = 0$ for all $m < 0$, and let

$$R_+ = \bigoplus_{m=1}^{\infty} R_m.$$

Then $R_+ \subset R$ is an ideal, and it is exactly the vanishing ideal of the fixed point set for the action of $\mathbb{C}^\times$ on $\operatorname{Spec} R$. Thus the fixed point set is equal to $\operatorname{Spec}(R/R_+) \cong \operatorname{Spec} R_0$. Note that this is false when R has elements of negative weight!

Definition 3.4. If $R = \bigoplus_{m=0}^{\infty} R_m$ is a graded ring, then

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

Example 3.5. $\mathbb{C}[x_0, \ldots, x_n]$ where each x_i has weight 1 gives $\mathbb{P}^n$.

Suppose now that $R_0 = \mathbb{C}$ and R is generated by R_1. This means that we can find an isomorphism $R \cong \mathbb{C}[x_0, \ldots, x_n]/I$, where x_i has weight 1 for all i and I is a homogeneous ideal. Thus we have a $\mathbb{C}^\times$-equivariant embedding $\operatorname{Spec} R \hookrightarrow \mathbb{C}^{n+1}$, where the $\mathbb{C}^\times$ action on $\mathbb{C}^{n+1}$ is given by $t \cdot z = t^{-1} z$ for every $t \in \mathbb{C}^\times$ and $z \in \mathbb{C}^{n+1}$. The fixed point set $\operatorname{Spec} R_0$ maps to the origin $0 \in \mathbb{C}^{n+1}$, thus we get an embedding

$$X := \operatorname{Proj} R = (\operatorname{Spec} R \smallsetminus \{0\}) / \mathbb{C}^\times \hookrightarrow \left(\mathbb{C}^{n+1} \smallsetminus \{0\}\right) / \mathbb{C}^\times = \mathbb{P}^n.$$

We of course did something that we should not have done; that is, we chose a basis $\{x_0, \ldots, x_n\}$ for R_1. More naturally, if S is a polynomial ring generated in degree 1, then

$$\operatorname{Proj} S \cong \mathbb{P}(S_1^*) := (S_1^* \smallsetminus \{0\}) / \mathbb{C}^\times.$$

In our case, R is a quotient of the polynomial ring $\operatorname{Sym} R_1$ by a homogeneous ideal, so $\operatorname{Proj} R$ embeds into $\mathbb{P}(R_1^*)$. If we choose a basis, we can identify $\mathbb{P}(R_1^*)$ with $\mathbb{P}^n$.

More generally, suppose that R is generated by R_0 and R_1. In this case, R is a quotient of $\operatorname{Sym}_{R_0} R_1$, $\operatorname{Spec} R$ is a subspace of $\operatorname{Spec} R_0 \times R_1^*$, and $\operatorname{Proj} R$ is a subspace of $\operatorname{Spec} R_0 \times \mathbb{P}(R_1^*)$. In particular, this means that we get a map from $\operatorname{Proj} R$ to $\operatorname{Spec} R_0$ whose fibers are subvarieties of

$\mathbb{P}(R_1^*)$. Indeed, an element of $\operatorname{Spec} R_0$ is represented by a surjective homomorphism from R_0 to $\mathbb{C}$, and the fiber of $\operatorname{Proj} R$ over that element is canonically isomorphic to $\operatorname{Proj} R \otimes_{R_0} \mathbb{C}$. For this reason, the variety $\operatorname{Proj} R$ is sometimes referred to as being **projective over affine**.

Remark 3.6. If R is not generated by R_0 and R_1, then any choice of homogenous generators of R over R_0 will lead to an embedding of $\operatorname{Proj} R$ into the product of $\operatorname{Spec} R_0$ with a weighted projective space, where the weights are exactly the positive degrees of the generators. In particular, $\operatorname{Proj} R$ is always Hausdorff (which is not immediately obvious from the definition).

3.2 Line bundles on projective space

Let's look more carefully at Example 3.5. Consider the line bundle

$$\mathcal{O}_{\mathbb{P}^n}(1) := \big(\mathbb{C}^{n+1} \smallsetminus \{0\}\big) \times_{\mathbb{C}^\times} \mathbb{C} = \Big(\big(\mathbb{C}^{n+1} \smallsetminus \{0\}\big) \times \mathbb{C}\Big)\Big/\mathbb{C}^\times,$$

where $\mathbb{C}^\times$ acts by inverse scalar multiplication on both $\mathbb{C}^n$ and $\mathbb{C}$. In other words, if we write $\mathbb{C}^n = \operatorname{Spec} \mathbb{C}[x_1, \ldots, x_n]$ and $\mathbb{C} = \operatorname{Spec} \mathbb{C}[y]$, then all of the variables have weight 1 for the action of $\mathbb{C}^\times$. The way we have defined it, $\mathcal{O}_{\mathbb{P}^n}(1)$ is a space that maps to $\mathbb{P}^n$ by projection onto the first factor, and the fiber of the projection over any point is a line. More generally, consider the line bundle

$$\mathcal{O}_{\mathbb{P}^n}(r) := \big(\mathbb{C}^{n+1} \smallsetminus \{0\}\big) \times_{\mathbb{C}^\times} \mathbb{C}_r = \Big(\big(\mathbb{C}^{n+1} \smallsetminus \{0\}\big) \times \mathbb{C}_r\Big)\Big/\mathbb{C}^\times,$$

where $\mathbb{C}^\times$ acts on $\mathbb{C}_r = \operatorname{Spec} \mathbb{C}[y]$ by $t \cdot y = t^r y$ for all $t \in \mathbb{C}^\times$. The bundle $\mathcal{O}_{\mathbb{P}^n}(-1)$ is called the **tautological line bundle** because there is a natural map from $\mathcal{O}_{\mathbb{P}^n}(-1)$ to $\mathbb{C}^n$ with coordinates $x_1 y, \ldots, x_n y$ that identifies the fiber over $[z] \in \mathbb{P}^n$ with the line $\mathbb{C}[z] \subset \mathbb{C}^{n+1}$. Then bundle $\mathcal{O}_{\mathbb{P}^n}(1)$ is therefore called the **anti-tautological line bundle**.

For any two integers r and s, we can consider the tensor product of $\mathbb{C}^\times$-representations $\mathbb{C}_r \otimes \mathbb{C}_s$. The group $\mathbb{C}^\times$ acts by

$$t \cdot (z_1 \otimes z_2) = (t \cdot z_1) \otimes (t \cdot z_2) = t^r z_1 \otimes t^s z_2 = t^{r+s}(z_1 \otimes z_2),$$

thus we have an isomorphism of $\mathbb{C}^\times$-representations $\mathbb{C}_r \otimes \mathbb{C}_s \cong \mathbb{C}_{r+s}$. We also see that $\mathbb{C}_0$ is the trivial representation, and that for any r, $\mathbb{C}_{-r}$ is dual to $\mathbb{C}_r$. As a result, we obtain the analogous statement for line bundles.

Proposition 3.7. *$\mathcal{O}_{\mathbb{P}^n}(0)$ is the trivial line bundle, and we have canonical isomorphisms*

$$\mathcal{O}_{\mathbb{P}^n}(r)^* \cong \mathcal{O}_{\mathbb{P}^n}(-r) \qquad \text{and} \qquad \mathcal{O}_{\mathbb{P}^n}(r) \otimes \mathcal{O}_{\mathbb{P}^n}(s) \cong \mathcal{O}_{\mathbb{P}^n}(r+s)$$

for all $r, s \in \mathbb{Z}$.

We can also describe the space of sections of $\mathcal{O}_{\mathbb{P}^n}(r)$ for any r. A section of $\mathcal{O}_{\mathbb{P}^n}(r)$ is a map from $\mathbb{P}^n$ to $\mathcal{O}_{\mathbb{P}^n}(r)$ such that, when we compose it with the projection back down to $\mathbb{P}^n$, we get the identity endomorphism of $\mathbb{P}^n$.

Proposition 3.8. *If $n > 0$, then the space of sections of $\mathcal{O}_{\mathbb{P}^n}(r)$ is canonically isomorphic to the degree r part of $\mathbb{C}[x_0, \ldots, x_n]$. In particular, $\mathcal{O}_{\mathbb{P}^n}(r)$ has no sections when r is negative, and lots of sections when r is positive.*

Proof. A section of $\mathcal{O}_{\mathbb{P}^n}(r)$ is the same as a $\mathbb{C}^\times$-equivariant map from $\mathbb{C}^{n+1} \smallsetminus \{0\}$ to the product $\big(\mathbb{C}^{n+1} \smallsetminus \{0\}\big) \times \mathbb{C}_r$ that's equal to the identity on the first coordinate. This last condition allows us to drop the first coordinate entirely and look at $\mathbb{C}^\times$-equivariant maps from $\mathbb{C}^{n+1} \smallsetminus \{0\}$ to $\mathbb{C}_r$. This means that we are looking for functions f on $\mathbb{C}^{n+1} \smallsetminus \{0\}$ such that for all $t \in \mathbb{C}^\times$ and $z \in \mathbb{C}^{n+1} \smallsetminus \{0\}$,

$$(t \cdot f)(z) = f(t^{-1} \cdot z) = t^{-1} \cdot f(z) = t^r f(z) = (t^r f)(z).$$

In other words, we are looking for functions on $\mathbb{C}^{n+1} \smallsetminus \{0\}$ of weight r. Since $n > 0$, functions on $\mathbb{C}^{n+1} \smallsetminus \{0\}$ are the same as functions on $\mathbb{C}^{n+1}$, and we're done. $\square$

Remark 3.9. Just for fun, let's examine the two low-dimensional cases not treated by Proposition 3.8. If $n = 0$, then $\mathbb{P}^0$ is a point, and $\mathcal{O}_{\mathbb{P}^0}(r)$ is the trivial bundle for all r, which has a 1-dimensional space of sections. So in this case the statement of Proposition 3.8 fails when $r < 0$. If $n = -1$, then $\mathbb{P}^{-1} = \emptyset$, and $\mathcal{O}_{\mathbb{P}^{-1}}(r)$ has a unique section for all r. So in this case Proposition 3.8 fails when $r = 0$.

Proposition 3.10. *If $n \geq 0$, then the ring $\mathbb{C}[x_0, \ldots, x_n]$ is canonically isomorphic to the ring of functions on the tautological line bundle $\mathcal{O}_{\mathbb{P}^n}(-1)$. Elements of degree r are functions that restrict to a polynomial of degree r on each fiber.*

Proof. It is clear that the ring of functions on $\mathcal{O}_{\mathbb{P}^n}(-1)$ is graded by fiber-wise degree. A function that is constant on fibers is nothing more than a function on the base, of which there are only constants. A function that is linear on fibers is a section of the dual bundle $\mathcal{O}_{\mathbb{P}^n}(1)$. More generally, a function that has degree r on fibers is a section of the tensor product of r copies of the dual bundle. Proposition 3.7 tells us that this tensor product is isomorphic to $\mathcal{O}_{\mathbb{P}^n}(r)$, and if $n > 0$ then Proposition 3.8 tells us that the sections of $\mathcal{O}_{\mathbb{P}^n}(r)$ are the degree r elements of $\mathbb{C}[x_1, \ldots, x_n]$. The reader is invited to work out the $n = 0$ case for his or herself. $\square$

3.3 Line bundles on Proj(R)

Much of what we did in Section 3.2 can be carried out when $\mathbb{C}[x_1, \ldots, x_n]$ is replaces by a more general non-negatively graded ring R. It works best when R is generated by R_0 and R_1, so we will make that assumption for the rest of this section.

Let $X = \operatorname{Proj} R$, and for any integer r, let

$$\mathcal{O}_X(r) := (\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0) \times_{\mathbb{C}^\times} \mathbb{C}_r.$$

It is clear that, when R is a polynomial ring generated in degree 1, this agrees with our definition in Section 3.2. Let us show that $\mathcal{O}_X(r)$ is a line bundle on X. If we project onto the first factor, we get a map from $\mathcal{O}_X(r)$ to X. Let x be an element of X, and let $\tilde{x}$ be a lift of x to $\operatorname{Spec} R$. The fiber over x of the projection is isomorphic to the quotient of $\mathbb{C}_r$ by the stabilizer group of $\tilde{x} \in \operatorname{Spec} R$. Since $\tilde{x} \notin \operatorname{Spec} R_0$, we know that this stabilizer group is finite. We would like to show that this stabilizer group is trivial.

Again, since $\tilde{x} \notin \operatorname{Spec} R_0$, there must be an element of R_+ that does not vanish at $\tilde{x}$. Since R is generated by R_0 and R_1, there must be an element $f \in R_1$ that does not vanish at $\tilde{x}$. Then for every $t \in \mathbb{C}^\times$,

$$f(t \cdot \tilde{x}) = (t^{-1} \cdot f)(\tilde{x}) = (t^{-1} f)(\tilde{x}) = \mathfrak{t}^{-1} f(\tilde{x}) \neq f(\tilde{x}),$$

so the stabilizer group of $\tilde{x}$ is trivial.

Remark 3.11. If R were not generated in degree 1, we could do something similar with orbi-bundles over orbifolds.

It is almost the case that Proposition 3.10 carries over unchanged to this setting, with $\mathbb{C}[x_0, \ldots, x_n]$ replaced by R and $\mathbb{P}^n$ replaced by X. The only tricky part is finding the right analogue of the $n \geq 0$ condition. The purpose of this condition was to ensure that all functions on $\mathbb{C}^{n+1} \smallsetminus \{0\}$ of non-negative weight come from $\mathbb{C}^{n+1}$. We need to replace that with the condition that all functions on $\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$ of non-negative weight come from $\operatorname{Spec} R$. This is true when $\operatorname{Spec} R$ is smooth and $R_0 \subsetneq R$; see Example 3.14 for a singular example in which the condition fails.

Proposition 3.12. *The ring of functions on $\mathcal{O}_X(-1)$ is canonically isomorphic to the ring of functions on $\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$ with non-negative weight. For any non-negative integer r, a function on $\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$ of weight r for the action of $\mathbb{C}^\times$ corresponds to a function on $\mathcal{O}_X(-1)$ that restrict to a polynomial of degree r on each fiber.*

Corollary 3.13. *The ring R maps to the ring of functions on $\mathcal{O}_X(-1)$. If $R_0 \subsetneq R$ (that is, if $\operatorname{Spec} R_0 \neq \operatorname{Spec} R$) then this map is an inclusion. If in addition $\operatorname{Spec} R$ is smooth, then this inclusion is an isomorphism.*

Example 3.14. Let $R = \mathbb{C}[x^4, x^3y, xy^3, y^4] \subset \mathbb{C}[x, y]$, where x and y both have degree 1. Then $\operatorname{Spec} R$ is the image of the map from $\mathbb{C}^2$ to $\mathbb{C}^4$ taking (u, v) to (u^4, u^3v, uv^3, v^4), which is 2-dimensional, and $\operatorname{Spec} R_0$ is the image of the origin along the same map, which is a singular point. To get the ring of functions on $\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$, we need to take R and invert x^3y and xy^3, since these guys vanish only at the origin. In particular, we have the function $(x^3y)^{-1}(x^4)(xy^3) = x^2y^2$, which is not the restriction of an element of R.

Remark 3.15. The restriction map from R to functions on $\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$ in Example 3.14 is *almost* surjective. The function x^2y^2 is not in the image, but its square is. In general, the ring of functions of non-negative weight on $\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$ is isomorphic to the integral closure of R.

Let us rephrase some of these observations in more geometric language. Consider the action of $\mathbb{C}^\times$ on $\mathcal{O}(-1)$ is by scalar multiplication on the fibers. There is a natural $\mathbb{C}^\times$-equivariant map

$$\pi : \mathcal{O}_X(-1) \to \operatorname{Spec} R$$

taking a point of $\mathcal{O}_X(-1)$ to the maximal ideal in R consisting of elements which, when regarded as functions on $\mathcal{O}_X(-1)$, vanish at that point. Equivalently, an element of $\mathcal{O}_X(-1)$ is represented by a pair

$$(\tilde{x}, z) \in (\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0) \times \mathbb{C}_{-1}.$$

If $z \neq 0$, then we map this element to $z^{-1} \cdot \tilde{x} \in \operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$. (Let's check that this makes sense. Since $z \in \mathbb{C}_{-1}$, we have $t \cdot z = tz$ for all $t \in \mathbb{C}^\times$. Then $(\tilde{x}, z)$ should map to the same place as $(t \cdot \tilde{x}, t \cdot z)$, which it does.) If $z = 0$, we map it to the limit as t approaches zero of $t^{-1} \cdot \tilde{x}$. To see that this limit exists, we need to check that for all $f \in R$, the limit

$$\lim_{t \to 0} f(t^{-1} \cdot \tilde{x}) = \lim_{t \to 0} (t \cdot f)(\tilde{x})$$

exists. Indeed, this limit is equal to $f_0(\tilde{x})$, where f_0 is the projection of f onto R_0.

Proposition 3.16. *The preimage of* $\operatorname{Spec} R_0 \subset \operatorname{Spec} R$ *along the map* $\pi : \mathcal{O}_X(-1) \to \operatorname{Spec} R$ *is the zero section of* $\mathcal{O}_X(-1)$*, and* π *restricts to an isomorphism from the complement of the zero section to* $\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$*. If* $R_0 \subsetneq R$*, then* π *is surjective.*

Proof. All but the last statement are immediate consequences of the geometric description of π above. Note that the map from the zero section (which is isomorphic to X) to $\operatorname{Spec} R_0$ is exactly the map discussed near the end of Section 3.1. Surjectivity comes from the description that we gave of the fibers of that map, which are manifestly nonempty. $\square$

Example 3.17. Redo the example of $\mathcal{O}_{\mathbb{P}^n}(-1)$ mapping to $\mathbb{C}^n$.

The following proposition follows immediately from the definitions.

Proposition 3.18. *Suppose that* R *is generated by* R_1*, so that* $X \hookrightarrow \mathbb{P}(R_1^*)$*. For every integer* r*,* $\mathcal{O}_X(r)$ *is the restriction of* $\mathcal{O}_{\mathbb{P}(R_1^*)}(r)$ *to* X*.*

Remark 3.19. One way to express the phenomenon in Example 3.14 is to say that the restriction map from sections of $\mathcal{O}_{\mathbb{P}^3}(1)$ to sections of $\mathcal{O}_X(1)$ is not surjective. The standard terminology is that the inclusion $X \subset \mathbb{P}^3$ is not **projectively normal**.

4 Geometric invariant theory: the third and fourth constructions

Roughly speaking, geometric invariant theory (GIT) is a prescription for taking quotients in algebraic geometry. Both the third and fourth construction from Section 1 would be described by an expert as "the GIT presentation of the toric variety associated to P." This is somewhat confusing, since the two constructions are superficially very different. Our goal for this section is to define GIT quotients in general, and explain why the the third and fourth construction in Section 1 are both part of the same theory.

4.1 Reductive group actions

An algebraic group G is **reductive** if it has a compact subgroup $K \subset G$ with the property that, for any linear representation W of G, the inclusion $W^K \subset W^G$ is an equality.[1] This subgroup K is unique up to conjugation, and is known as a **maximal compact subgroup**. Examples include $\mathrm{GL}(n;\mathbb{C})$, which has maximal compact subgroup $\mathrm{U}(n)$, and $\mathrm{SL}(n;\mathbb{C})$, which has maximal compact subgroup $\mathrm{SU}(n)$. The most important example for us will be the torus $T^n_{\mathbb{C}}$ (and its subtori), which has a unique maximal torus T^n. The relevance of our definition is that we have an "averaging" function $W \to W^K = W^G$ taking w to $\int_K (k \cdot w)\, dk$ with respect to Haar measure; note that we have no such projection when G is not reductive. This feature of reductive groups is sometimes known as **Weyl's unitary trick**.

An action of G on a ring R induces an action on $\operatorname{Spec} R$. Furthermore, if R admits a grading

$$R = \bigoplus_{m=0}^{\infty} R_m$$

that is compatible with the action of G (that is, the actions of G and $\mathbb{C}^\times$ must commute), then G acts on $\operatorname{Proj} R$, as well as on all of the line bundles $\mathcal{O}(r)$ if R is generated in degrees 0 and 1. The following lemma will be useful to us in the next section.

Lemma 4.1. *Suppose that G is a reductive group acting on an affine variety $X = \operatorname{Spec} R$, and that $Y, Z \subset X$ are disjoint subvarieties preserved by the action of G. Then there exists a G-invariant function $f \in R$ that takes that value 0 on Y and 1 on Z.*

Proof. Let $I, J \subset R$ be the ideals of functions that vanish on Y and Z, respectively. The condition that $Y \cap Z = \emptyset$ is equivalent to the statement that $I + J = R$, thus there exists a pair of functions $f \in I$ and $g \in J$ such that $f + g = 1$. By averaging f and g, we obtain the functions that we want. □

4.2 GIT quotients

Let G be a reductive group acting on R; suppose that we want to take a quotient of $\operatorname{Spec} R$ by G. A naive quotient would be really bad—for example, the orbits of G on $\operatorname{Spec} R$ are usually not closed. Here's another idea; the following construction is called an **affine GIT quotient**.

Definition 4.2. $(\operatorname{Spec} R)/\!\!/G := \operatorname{Spec} R^G$.

In order to justify the idea that this construction can be regarded as some kind of a "quotient," we state and prove the following proposition.

Proposition 4.3. *There is a surjective map from $\operatorname{Spec} R$ to $(\operatorname{Spec} R)/\!\!/G$. Two points in $\operatorname{Spec} R$ lie in the same fiber of this map if and only if the closures of their G-orbits intersect nontrivially.*

[1] This is not the standard definition, but it will do for our purposes.

Before proving Proposition 4.3, let's use it to see both how well behaved and how poorly behaved this quotient can be.

Example 4.4. Consider two different actions of $\mathbb{C}^\times$ on $\mathbb{C}^2$, as well as the anti-diagonal action of $\mathbb{C}^\times$ on the complement of a coordinate line in $\mathbb{C}^2$.

Proof of Proposition 4.3. To see that the natural map from $\operatorname{Spec} R$ to $\operatorname{Spec} R^G$ is surjective, we need to show that for any maximal ideal $\mathfrak{m} \subset R^G$, there exists a maximal ideal in R whose intersection with R^G is $\mathfrak{m}$. Let $I = R \cdot \mathfrak{m} \subset R$. We claim that I is a proper ideal, that is, that I does not contain 1. Indeed, suppose that I did contain 1. This means that there exists $f_1, \ldots f_n \in R$ and $m_1, \ldots m_n \in \mathfrak{m}$ such that $f_1 m_1 + \ldots f_n m_n = 1$. By averaging each f_i over the maximal compact subgroup $K \subset G$, we may assume that each f_i is G-invariant, but this contradicts the fact that $\mathfrak{m}$ is a proper ideal in R^G. Any maximal ideal of R containing I has the property that its intersection with R^G is a proper ideal of R^G containing $\mathfrak{m}$, and is therefore equal to $\mathfrak{m}$.

Next we show that the fibers are what we say they are. Since any G-invariant function is constant on closures of orbits, elements of R^G can't distinguish between two points whose orbit closures intersect, therefore two such points live in the same fiber. Conversely, if x and y are elements of $\operatorname{Spec} R$ whose orbit closures do *not* intersect; then Lemma 4.1 says that we can find a G-invariant function that distinguishes x from y, thus they have different images in $X /\!\!/ G$. $\square$

It turns out that we can get much cooler quotients (even of affine varieties!) using Proj rather than Spec. Suppose that G is a reductive group acting on $R = \bigoplus_{m=0}^\infty R_m$. We can take Definition 4.2 as motivation for the following definition of a **projective GIT quotient**.

Definition 4.5. $(\operatorname{Proj} R) /\!\!/ G := \operatorname{Proj} R^G$.

Before proceeding, we warn the reader to be careful of the terrible notation in this definition! We define that quotient of $\operatorname{Proj} R$ by G in terms of the ring R, but it is not possible to recover the ring R from the space $\operatorname{Proj} R$. Under suitable hypotheses, Corollary 3.13 says that R is the ring of functions on the total space of the tautological line bundle. Thus in this case, the projective GIT quotient is determined by the action of G on the line bundle $\mathcal{O}(-1)$. A lift of a G-action from X to $\mathcal{O}_X(-1)$ is often called a **linearization** of the action.

We would now like to state and prove an analogue of Proposition 4.3, but the situation is somewhat more complicated. The problem is that there is not a natural map from $\operatorname{Proj} R$ to $\operatorname{Proj} R^G$. Indeed, the inclusion of graded rings from R^G into R induces a $\mathbb{C}^\times$ equivariant surjection from $\operatorname{Spec} R$ to $\operatorname{Spec} R^G$. However, there are some elements of $\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$ that land in $\operatorname{Spec} R_0^G \subset \operatorname{Spec} R^G$.

Example 4.6. Let $R = \mathbb{C}[x_1, \ldots, x_n, y]$, where x_i has degree 0 for all i and y has degree 1. Let $G = \mathbb{C}^\times$ act with weight 1 on each x_i and weight 0 on y.

Definition 4.7. Suppose that G is a reductive group acting on R, and let $X = \operatorname{Proj} R$. Let x be an element of X, and let $\tilde{x}$ be any lift of x to $\operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$. The element x is called **semistable**[2]

[2] Note that this condition does not depend on the choice of $\tilde{x}$.

if there exists an element $f \in R_+^G$ of such that $f(\tilde{x}) \neq 0$. Equivalently, x is semistable if its image in $\operatorname{Spec} R^G$ does not lie in $\operatorname{Spec} R_0^G$. A point which is not semistable is celled **unstable**.

We will denote the semistable set, which is clearly open, by $X^{\mathrm{ss}} \subset X$, and we denote its complement by $X^{\mathrm{us}} \subset X$. Note that the definition is exactly tailored so that there will be a well-defined map from X^{ss} to $X/\!\!/G$; we will analyze this map in greater detail shortly. Note also that, just as the definition of $X/\!\!/G$ is not determined by the action of G on X, the definition of X^{ss} is also not determined by the action of G on X. If R is generated by R_0 and R_1, then Remark 3.15 tells us that it is determined by the action of G on $\mathcal{O}_X(-1)$. Indeed, we give the geometric interpretation of the semistable locus below.

Proposition 4.8. *Suppose that R is generated by R_0 and R_1. Let x be an element of X, and let ℓ be a nonzero element of the line $\mathcal{O}_X(-1)_x$. Then x is semistable if and only if the closure of $G \cdot (x, \ell)$ in $\mathcal{O}_X(-1)$ is disjoint from the zero section.*

Proof. Let $\tilde{x} = \pi(x, \ell) \in \operatorname{Spec} R \smallsetminus \operatorname{Spec} R_0$, where $\pi : \mathcal{O}_X(-1) \to \operatorname{Spec} R$ is the map discussed in Proposition 3.16. This map is clearly G-equivariant, thus Proposition 3.16 implies that the closure of $G \cdot (x, \ell)$ in $\mathcal{O}_X(-1)$ is disjoint from the zero section if and only if the closure of $G \cdot \tilde{x}$ is disjoint from $\operatorname{Spec} R_0$. By Lemma 4.1, this is equivalent to the existence of a G-invariant function that vanishes on $\operatorname{Spec} R_0$ and takes the value 1 at $\tilde{x}$. Since the ideal of functions that vanish on $\operatorname{Spec} R_0$ is equal to R_+, we're done. □

Example 4.9. Vary the weight of y in Example 4.6.

Now let's go ahead and prove the projective analogue of Proposition 4.3.

Proposition 4.10. *There is a surjective map from X^{ss} to $X/\!\!/G$. Two points in X^{ss} lie in the same fiber of this map if and only if the closures of their G-orbits in X^{ss} (not in X) intersect nontrivially.*

Proof. The fact that X^{ss} surjects onto $X/\!\!/G$ follows immediately from Proposition 4.3 and the definition of X^{ss}. This map is clearly G-invariant; the fact that it is constant not just on orbits but also on closures of orbits follows from the fact that $X/\!\!/G = \operatorname{Proj} R^G$ is Hausdorff.

Let x and y be elements of X^{ss}, and lift them to element $\tilde{x}$ and $\tilde{y}$ of $\operatorname{Spec} R$. If the closures of the G-orbits through x and y don't intersect in X^{ss}, then the closures of the $G \times \mathbb{C}^\times$-orbits through $\tilde{x}$ and $\tilde{y}$ don't intersect in $\operatorname{Spec} R \smallsetminus Z(R_+^G)$, where $Z(R_+^G)$ is the vanishing locus of the ideal generated by all G-invariant functions of positive weight. (Note that this locus is bigger than $\operatorname{Spec} R_0 = Z(R_+)$; the difference is exactly the stuff lying over X^{us}.) We would now like to apply Lemma 4.1, but it's a little tricky because $\operatorname{Spec} R \smallsetminus Z(R_+^G)$ is not affine. The lemma tells us that we can find an open neighborhood of $\tilde{x}$ and $\tilde{y}$ and a $G \times \mathbb{C}^\times$ invariant function defined on that neighborhood that distinguishes $\tilde{x}$ from $\tilde{y}$. This is exactly what we need to show that $\tilde{x}$ and $\tilde{y}$ have different $\mathbb{C}^\times$-orbits in $\operatorname{Spec} R^G$, or equivalently that x and y have different images in $\operatorname{Proj} R^G = X/\!\!/G$. □

Consider the relation $\sim$ on X^{ss} given by putting $x \sim y$ if the closures of the G-orbits through x and y intersect. The following statement follows from Proposition 4.10.

Corollary 4.11. *The projective GIT quotient $X/\!\!/G$ is homeomorphic to $X^{\mathrm{ss}}/\!\sim$.*

Remark 4.12. Technically we have only shown that there is a continuous bijection from $X^{\mathrm{ss}}/\!\sim$ to $X/\!\!/G$. We will leave continuity of the inverse in the noncompact case as a painful exercise. In fact, we will do this quite a few more times in this section and the next.

Example 4.13. Lots of actions of $\mathbb{C}^\times$ on $\mathbb{C}^2$.

Proposition 4.10 has another nice corollary that will be very important in Section 5. We call a point $x \in X^{\mathrm{ss}}$ **polystable** if its G-orbit in X^{ss} is closed, and we denote the polystable locus by X^{ps}.

Corollary 4.14. *Every G-orbit in X^{ss} has a unique polystable orbit in its closure. Thus the natural inclusion of X^{ps} into X^{ss} induces a homeomorphism*[3]

$$X^{\mathrm{ps}}/G \to X/\!\sim\; \cong X/\!\!/G.$$

Proof. The fact that every G-orbit in X^{ss} has at least one polystable orbit in its closure follows immediately by induction on the dimension of the orbit. For uniqueness, suppose that x' and x'' are polystable points in the closure of the orbit through x. Then $x' \sim x \sim x''$, and Proposition 4.10 imples that $\sim$ is an equivalence relation. Thus $x' \sim x''$, which means that they lie in the same G-orbit. $\square$

This is a convenient time to make one more definition. We won't use it much, but it's important to know. We call a point $x \in X$ **stable** if it is polystable and its stabilizer subgroup in G is finite; we denote the locus of stable points by X^{s}. It is a fact that the stable locus is open (either empty or dense); note that this is not in general true of the polystable locus. Thus the geometri quotient X^{s}/G is an orbifold and an open (dense or empty) subset of $X/\!\!/G$.

4.3 The toric invariant ring (the fourth construction)

We are now ready to realize the construction of $X_3(P)$ in Section 1.3 as a GIT quotient. As in Section 1, let $\lambda_1, \ldots, \lambda_n$ be integers. Let $R = \mathbb{C}[x_1, \ldots, x_n, y]$ with $\deg x_i = 0$ for all i and $\deg y = 1$; then

$$\operatorname{Proj} R \cong \operatorname{Spec} \mathbb{C}[x_1, \ldots, x_n] = \mathbb{C}^n.$$

Let $T^n_{\mathbb{C}}$ act on R by putting

$$(t_1, \ldots, t_n) \cdot x_i = t_i^{-1} x_i \text{ for all } i \qquad \text{and} \qquad (t_1, \ldots, t_n) \cdot y := t_1^{\lambda_1} \ldots t_n^{\lambda_n} y.$$

We thus obtain an action of $T^n_{\mathbb{C}}$ on R such that the induced action on $\operatorname{Proj} R$ is the obvious one.

[3] See Remark 4.12.

Remark 4.15. It may be easier to think about what we just did geometrically. We have

$$\operatorname{Spec}\mathbb{C}[x_1,\ldots,x_n] = \mathbb{C}^n \cong \Big((\mathbb{C}^n \times \mathbb{C}) \smallsetminus (\mathbb{C}^n \times \{0\}) \Big) \Big/ \mathbb{C}^\times = \operatorname{Proj}\mathbb{C}[x_1,\ldots,x_n,y].$$

We know how we want $T^n_{\mathbb{C}}$ to act on $\mathbb{C}^n$, and we can achieve this by letting it act any which way on the $\mathbb{C}$ in $\mathbb{C}^n \times \mathbb{C}$. This means that we know how $T^n_{\mathbb{C}}$ has to act on $x_1,\ldots,x_n$, but we get to choose how it acts on y. That's what $\lambda_1,\ldots,\lambda_n$ are for (note that $\mathbb{Z}^n \cong \operatorname{Hom}(T^n_{\mathbb{C}},\mathbb{C}^\times)$ is exactly the set that indexes possible actions).

Now suppose that, in addition to $\lambda_1,\ldots,\lambda_n \in \mathbb{Z}$, we have a real vector space V with lattice $V_{\mathbb{Z}}$ and a collection of primitive vectors $a_1,\ldots,a_n \in V^*_{\mathbb{Z}}$. This allows us to define a polytope $P \subset V$ along with facets $F_i \subset P$ for all i. We also have a subtorus $K_{\mathbb{C}} \subset T^n_{\mathbb{C}}$, defined as the kernel of the map from $T^n_{\mathbb{C}}$ to $T_{\mathbb{C}} = V^*/V^*_{\mathbb{Z}}$ given by $a_1,\ldots,a_n$. Since $K_{\mathbb{C}}$ sits inside of $T^n_{\mathbb{C}}$, the action of $T^n_{\mathbb{C}}$ on R induces an action of $K_{\mathbb{C}}$.

Recall from Section 1.4 our definition of the cone $\Sigma \subset V \times \mathbb{R}_{\geq 0}$, the semigroup $S_P = \Sigma \cap (V_{\mathbb{Z}} \times \mathbb{N})$, and the and the semigroup ring $\mathbb{C}[S_P]$, which is graded by the last coordinate.

Proposition 4.16. *We have a natural isomorphism $\mathbb{C}[S_P] \cong R^{K_{\mathbb{C}}}$, and therefore*

$$X_4(P) \;=\; \operatorname{Proj}\mathbb{C}[S_P] \;\cong\; \operatorname{Proj} R^{K_{\mathbb{C}}} \;=\; \mathbb{C}^n /\!\!/ K_{\mathbb{C}}.$$

Proof. Recall that in Section 2.1 we identified P with an affine linear slice of $\mathbb{R}^n_{\geq 0}$, with the lattice points of P going to the intersection of the slice with $\mathbb{N}^n$. Here we will use the same trick to identify Σ with a linear slice of $\mathbb{R}^n_{\geq 0} \times \mathbb{R}_{\geq 0}$, with S_P going to the intersection of the slice with $\mathbb{N}^n \times \mathbb{N}$. The lattice points in this slice will correspond exactly to $K_{\mathbb{C}}$-invariant monomials in R, and we'll be done.

We will make the simplifying assumption that the vectors $a_1,\ldots,a_n$ span $V^*_{\mathbb{Z}}$ over the integers. Recall from Lemma 1.8 that this implies that K is connected (and the same for $K_{\mathbb{C}}$), and therefore that $\mathfrak{k}^*_{\mathbb{Z}} \cong \operatorname{Hom}(K_{\mathbb{C}},\mathbb{C}^\times)$.

Consider the linear map $\phi : V \times \mathbb{R} \to \mathbb{R}^n \times \mathbb{R}$ given by the formula

$$\phi(v,r) := (a_1(v) + \lambda_1 r, \ldots, a_n(v) + \lambda_n r, r).$$

Note that ϕ is injective; if we know $\phi(v,r)$, then we know r, which means that we know $a_i(v)$ for all i, and therefore we know v. The preimage of $\mathbb{R}^n_{\geq 0} \times \mathbb{R}_{\geq 0}$ is equal to Σ (by definition), and the preimage of $\mathbb{N}^n \times \mathbb{N}$ is equal to S_P. (Here we use our assumption to conclude that if $a_i(v)$ is an integer for every i, then $v \in V_{\mathbb{Z}}$). Thus we may identify $\mathbb{C}[S_P]$ with the monomial ring

$$\mathbb{C}[x_1^{r_1} \ldots x_n^{r_n} y^r \mid (r_1,\ldots,r_n,r) \in \operatorname{im}(\phi)].$$

We must now show that these are exactly the $K_{\mathbb{C}}$-invariant monomials.

Observe that we have an exact sequence

$$0 \to V \times \mathbb{R} \xrightarrow{\phi} \mathbb{R}^n \times \mathbb{R} \xrightarrow{\psi} \mathfrak{k}^* \to 0,$$

where

$$\psi : (r_1, \ldots, r_n, r) := i^*(\lambda_1 r - r_1, \ldots, \lambda_n r - r_n).$$

The vector $(\lambda_1 r - r_1, \ldots, \lambda_n r - r_n) \in \mathbb{Z}^n \cong (\mathfrak{t}^n)^* \cong \operatorname{Hom}(T^n_{\mathbb{C}}, \mathbb{C}^\times)$ is precisely the weight of the monomial $x_1^{r_1} \ldots x_n^{r_n} y^r$ for the action of $T^n_{\mathbb{C}}$, and $i^*(\lambda_1 r - r_1, \ldots, \lambda_n r - r_n) \in \mathfrak{k}^*_{\mathbb{Z}} \cong \operatorname{Hom}(K_{\mathbb{C}}, \mathbb{C}^\times)$ is the weight of the same monomial for the action of $K_{\mathbb{C}}$. Thus the image of ϕ is equal to the kernel of ψ, which is the indexing set for monomials that are fixed by $K_{\mathbb{C}}$. □

Remark 4.17. The proof is somewhat more difficult if we do not assume that $a_1, \ldots, a_n$ spans $V^*_{\mathbb{Z}}$. In this case $\mathbb{C}[S_P]$ only includes into the monomial ring $\mathbb{C}[x_1^{r_1} \ldots x_n^{r_n} y^r \mid (r_1, \ldots, r_n, r) \in \operatorname{im}(\phi)]$, because some of those monomials will come from non-integral elements in Σ. On the other hand, those monomials will in general only be fixed by the connected component of the identity in $K_{\mathbb{C}}$, rather than by the whole group. It turns out that the monomials fixed by all of $K_{\mathbb{C}}$ are exactly those that come from S_P. Can anyone think of a proof of this fact?

Example 4.18. Take $n = d = 1$, let a_1 be 2 times a primitive vector, and let $\lambda_1 = 0$. Then $K_{\mathbb{C}} = \{\pm 1\}$ and $R^{K_{\mathbb{C}}} = \mathbb{C}[x_1^2, y]$.

Example 4.19. Some examples with $\mathbb{C}^\times$ acting on $\mathbb{C}^n$.

4.4 The toric semistable locus (the third construction)

Recall our definition of the set $U_\lambda \subset \mathbb{C}^n$ from Section 1.3:

$$U_\lambda := \{z \in \mathbb{C}^n \mid F_z \neq \emptyset\}, \qquad \text{where} \qquad F_z := \bigcap_{z_i = 0} F_i.$$

Proposition 4.20. *$U_\lambda \subset \mathbb{C}^n = \operatorname{Proj} R$ is the semistable locus for the action of $K_{\mathbb{C}}$ on R.*

Proof. By definition, $z = (z_1, \ldots, z_n)$ if and only if there is a monomial $x_1^{r_1} \ldots x_n^{r_n} y^r$ such that

- $(r_1, \ldots, r_n, r)$ lies in the image of ϕ (the monomial has to be $K_{\mathbb{C}}$-invariant)
- $r > 0$ (the monomial has to lie in R_+)
- $r_i = 0$ whenever $z_i = 0$ (the monomial can't vanish at z).

Pulling back by ϕ, this means that we are looking for $(v, r) \in S_P$ such that $r > 0$ and $a_i(v) + \lambda_i r = 0$ whenever $z_i = 0$. This last equation is equivalent to the condition that $\frac{1}{r} v \in F_z$, thus z is stable if and only if F_z contains an element of $V_{\mathbb{Q}}$. Since all of the facets have integral slope, this happens if and only if $F_z \neq \emptyset$. □

Combining Propositions 4.16 and 4.20 with Proposition 4.10, we obtain the following result.

Corollary 4.21. *$X_3(P)$ is $T_{\mathbb{C}}$-equivariantly isomorphic to $X_4(P)$.*

We'll conclude this section by showing that a Delzant polyhedron gives rise to a smooth toric variety (in the algebraic setting). The proof is very similar to the proof of Proposition 2.18.

Proposition 4.22. *If P is Delzant, then $K_{\mathbb{C}}$ acts freely on the semistable locus. In particular, we have $(\mathbb{C}^n)^{\mathrm{s}} = (\mathbb{C}^n)^{\mathrm{ps}} = (\mathbb{C}^n)^{\mathrm{ss}}$, and $X_3(P) = (\mathbb{C}^n)^{\mathrm{s}}/K_{\mathbb{C}}$ is smooth.*

Proof. We need to show that every semistable point $z \in \mathbb{C}^n$ has trivial stabilizer in $K_{\mathbb{C}}$. An element $(t_1, \ldots, t_n) \in T^n_{\mathbb{C}}$ stabilizes z if and only if $t_i = 1$ whenever $z_i \neq 0$. That means that the stabilizer of z is largest when lots of its coordinates are 0, that is, when F_z is a vertex. Suppose that F_z is a vertex, and let $I = \{i \mid z_i = 0\}$, so that the stabilizer of z is $T^I_{\mathbb{C}} \subset T^n_{\mathbb{C}}$. By Definition 1.2, $\{a_i \mid i \in I\}$ is a basis for $V^*_{\mathbb{Z}}$. This means that the map from $T^I_{\mathbb{C}}$ to $T_{\mathbb{C}}$ is an isomorphism, and therefore that $T^I_{\mathbb{C}} \cap K_{\mathbb{C}}$ is the trivial group. □

5 The Kempf-Ness theorem: the second and third constructions

The Kempf-Ness theorem says that, if G is a reductive algebraic group acting linearly on a complex vector space, the the GIT quotient by G is homeomorphic to the symplectic quotient by the maximal compact subgroup $K \subset G$. When G and K are abelian, this says exactly that the symplectic toric variety $X_1(P) \cong X_2(P)$ is homeomorphic to the algebraic toric variety $X_3(P) \cong X_4(P)$. Note that it will take a little bit of work to make this statement precise; in particular, we need to match up the freedom in the symplectic quotient (the choice of moment map) with the freedom in the GIT quotient (the choice of linearization).

5.1 The simplest case

Let's begin with the case where $G = \mathbb{C}^\times$ and $K = \mathrm{U}(1)$. Let G act on $\mathbb{C}^n$ be the forumla

$$t \cdot (z_1, \ldots, z_n) = (t^{\alpha_1} z_1, \ldots, t^{\alpha_n} z_n),$$

where $\alpha_1, \ldots, \alpha_n \in \mathbb{Z}$. Note that any linear action of G has this form.

Remark 5.1. This is a toric action if and only if the integers $\alpha_1, \ldots, \alpha_n$ are relatively prime, that is, if and only if no nontrivial element in G acts trivially.

Let $R = \mathbb{C}[x_1, \ldots, x_n, y]$ with $\deg x_i = 0$ and $\deg y = 1$, so that $\operatorname{Proj} R = \mathbb{C}^n$. We would like to lift the action of G on $\mathbb{C}^n$ to an action on R, that is, to choose a linearization of the G-action on $\mathbb{C}^n$. Such an action must have the form

$$t \cdot x_i = t^{-\alpha_i} x_i \qquad \text{and} \qquad t \cdot y = t^{\lambda} y$$

for some $\lambda \in \mathbb{Z}$. Meanwhile, the subgroup $K \subset G$ acts on $\mathbb{C}^n$ with moment map

$$\mu(z_1, \ldots, z_n) = \frac{1}{2}\alpha_1 |z_1|^2 + \ldots + \frac{1}{2}\alpha_n |z_n|^2 - \lambda$$

(see Examples 2.8 and 2.9).

Proposition 5.2 (Kempf-Ness for $\mathbb{C}^\times$)**.** *We have an inclusion $\mu^{-1}(0) \subset (\mathbb{C}^n)^{\mathrm{ps}}$, and for every polystable element $z \in \mathbb{C}^n$, $(G \cdot z) \cap \mu^{-1}(0)$ is a single K-orbit. Thus we obtain a homeomorphism*[4]

$$X_2(P) = \mathbb{C}^n /\!\!/ K = \mu^{-1}(0)/K \to (\mathbb{C}^n)^{\mathrm{ps}}/G \;\cong\; \mathbb{C}^n /\!\!/ G = X_3(P).$$

Before proving Proposition 5.2 in general, let's check that it works when $\alpha_i = 1$ for all i. If $\lambda < 0$, then $\mu^{-1}(0) = \emptyset = (\mathbb{C}^n)^{\mathrm{ss}}$. If $\lambda = 0$, then $\mu^{-1}(0) = \{0\} = (\mathbb{C}^n)^{\mathrm{ps}}$. If $\lambda > 0$, then $\mu^{-1}(0)$ is a sphere of dimension $2n-1$, and $(\mathbb{C}^n)^{\mathrm{ps}} = (\mathbb{C}^n)^{\mathrm{ss}} = \mathbb{C}^n \smallsetminus \{0\}$. We have $\mathbb{C}^\times = \mathrm{U}(1) \cdot \mathbb{R}_{>0}$. The action of $\mathbb{R}_{>0}$ on $(\mathbb{C}^n)^{\mathrm{ps}}$ can be used to take every element to a unique element of $\mu^{-1}(0)$, and then we divide by $\mathrm{U}(1)$ to get $\mathbb{P}^{n-1}$ in both cases.

The proof for arbitrary $\alpha_1, \ldots, \alpha_n$ works according to the same principle, namely that the action of $\mathbb{R}_{>0} \subset \mathbb{C}^\times$ can be used to shrink $(\mathbb{C}^n)^{\mathrm{ps}}$ down to $\mu^{-1}(0)$.

Proof of Proposition 5.2. First of all, we may as well assume that $\alpha_i \neq 0$ for all i, otherwise $\mathbb{C}^\times$ acts trivially on the i^{th} coordinate and we can just factor it out. For any $z \in \mathbb{C}^n$, define $\Psi_z : \mathbb{R} \to \mathbb{R}$ by the formula

$$\Psi_z(x) = \frac{1}{2}|e^{\alpha_1 x} z_1|^2 + \ldots + \frac{1}{2}|e^{\alpha_n x} z_n|^2 - 2\lambda x.$$

We then have

$$\Psi_z'(x) = \alpha_1 |e^{\alpha_1 x} z_1|^2 + \ldots + \alpha_n |e^{\alpha_n x} z_n|^2 - 2\lambda = 2\mu(e^x \cdot z)$$

and

$$\Psi_z''(x) = \alpha_1^2 |e^{\alpha_1 x} z_1|^2 + \ldots + \alpha_n^2 |e^{\alpha_n x} z_n|^2 \geq 0.$$

The equation for the first derivative tells us that x is a critical point of Ψ_z if and only if $\mu(e^x \cdot z) = 0$, and the equation for the second derivative tells us that $\Psi_z(x)$ is convex. If $z \neq 0$, then the function is strictly convex (the second derivative is always positive), and therefore has at most one critical point. In other words, the action of $\mathbb{R}_{>0} \subset \mathbb{C}^\times$ may or may not be able to take z to an element of $\mu^{-1}(0)$, but if so, that element is unique. This statement is of course true in the degenerate case when $z = 0$.

The next question that we need to ask is: for which z does Ψ_z have a critical point? We would like to show that Ψ_z has a critical point if and only if z is polystable, in which case we'll be done. Let's assume first that $\lambda \neq 0$. In this case $0 \in \mathbb{C}^n$ is not semistable. Since 0 is the only point that could possibly lie at the boundary of any orbit, this means that all orbits in the semistable locus are closed, in other words, $(\mathbb{C}^n)^{\mathrm{ss}} = (\mathbb{C}^n)^{\mathrm{ps}}$.

Since Ψ_z is convex and $\lambda \neq 0$, Ψ_z has a critical point if and only if

$$\lim_{x \to -\infty} \Psi_z(x) = \infty = \lim_{x \to \infty} \Psi_z(x).$$

The first equality says that, either there exists some i such that $z_i \neq 0$ and $\alpha_i < 0$, or $\lambda < 0$.

[4] See Remark 4.12.

The second equality says that, either there exists some i such that $z_i \neq 0$ and $\alpha_i > 0$, or $\lambda > 0$. Together, they say that there exists some i such that $z_i \neq 0$ and α_i has the opposite sign from λ. This is exactly the condition that there exists a G-invariant monomial of positive degree that does not vanish on z (namely $x_i^{|\lambda|}y^{|\alpha_i|}$). Thus Ψ_z has a critical point if and only if z is semistable. Since all semistable points are polystable, we are done.

Now suppose that $\lambda = 0$. In this case, every element of $\mathbb{C}^n$ is semistable. This can be seen algebraically, by noting the the monomial y is a positive degree invariant element of R that vanished nowhere on $\mathbb{C}^n$, or geometrically, using Proposition 4.8. Not every element is polystable, however; z fails to be polystable if $z \neq 0$ but z has 0 in its orbit closure, which happens if and only if α_i has the same sign for every i such that $z_i \neq 0$. The function Ψ_z has a critical point if and only if

$$\lim_{x\to-\infty} \Psi_z(x) = \infty = \lim_{x\to\infty} \Psi_z(x) \qquad \text{or} \qquad z = 0.$$

This means that either $z = 0$ or there exists i and j such that $z_i \neq 0 \neq z_j$ and α_i and α_j have opposite signs. This is equivalent to polystability, so we are done. $\square$

5.2 The toric case

In this section we will state and prove the Kempf-Ness theorem in the case that matters to us. That is, we start with the data $a_1, \ldots, a_n \in V_{\mathbb{Z}}^* \smallsetminus \{0\}$ and $\lambda_1, \ldots, \lambda_n \in \mathbb{Z}$, which determine a polyhedron $P \subset V$. We define $K \subset T^n$ and $G = K_{\mathbb{C}} \subset T^n_{\mathbb{C}}$ as in Section 1, with everything in sight acting on $\mathbb{C}^n$. We use $\lambda_1, \ldots, \lambda_n$ to define both the moment map for the action of K and the linearization of the action of $K_{\mathbb{C}}$. More precisely, we define $\mu : \mathbb{C}^n \to \mathfrak{k}^*$ by the formula

$$\mu(z_1, \ldots, z_n) = i^* \left(\frac{1}{2}|z_1|^2 - \lambda_1, \ldots, \frac{1}{2}|z_n|^2 - \lambda_n \right),$$

and we define the action of $K_{\mathbb{C}} \subset T^n_{\mathbb{C}}$ on $R = \mathbb{C}[x_1, \ldots, x_n, y]$ by putting

$$(t_1, \ldots, t_n) \cdot y = t_1^{\lambda_1} \ldots t_n^{\lambda_n} y.$$

The statement of our Theorem will be word-for-word the same as the statement of Proposition 5.2, which was a special case.

Theorem 5.3 (Toric Kempf-Ness). *We have an inclusion $\mu^{-1}(0) \subset (\mathbb{C}^n)^{\mathrm{ps}}$, and for every polystable element $z \in \mathbb{C}^n$, $(G \cdot z) \cap \mu^{-1}(0)$ is a single K-orbit. Thus we obtain a homeomorphism*[5]

$$\mathbb{C}^n /\!\!/ K = \mu^{-1}(0)/K \to (\mathbb{C}^n)^{\mathrm{ps}}/G \;\cong\; \mathbb{C}^n /\!\!/ G.$$

Proof. Our proof will model the proof of Proposition 5.2. For every $z \in \mathbb{C}^n$, consider the map

[5] See Remark 4.12.

$\Psi_z : \mathfrak{k} \to \mathbb{R}$ given by the formula

$$\Psi_z(\beta) = \frac{1}{2}|e^{\beta_1}z_1|^2 + \ldots + \frac{1}{2}|e^{\beta_n}z_n|^2 - 2(\lambda_1\beta_1 + \ldots + \lambda_n\beta_n)$$

for all $\beta \in \mathfrak{k} \subset \mathfrak{t}^n \cong \mathbb{R}^n$. Then for any $\gamma \in \mathfrak{k}$, we have

$$\partial_\gamma \Psi_z(\beta) = \frac{d}{dt}\Big|_{t=0} \Psi_z(\beta + t\gamma) = \gamma_1|e^{\beta_1}z_1|^2 + \ldots + \gamma_n|e^{\beta_n}z_n|^2 - 2(\lambda_1\gamma_1 + \ldots + \lambda_n\gamma_n) = 2\mu(e^\beta \cdot z)(\gamma)$$

and

$$\partial^2_\gamma \Psi_z(\beta) = \gamma_1^2|e^{\beta_1}z_1|^2 + \ldots + \gamma_n^2|e^{\beta_n}z_n|^2 \geq 0.$$

Thus Ψ_z is convex, and the critical points are precisely the elements $\beta \in \mathfrak{k}$ such that $\mu(e^\beta \cdot z) = 0$.

As in the $z = 0$ case from Proposition 5.2, it is possible for Ψ_z to fail to be strictly convex, in which case it will have more than one critical point. Indeed, we have $\partial^2_\gamma \Psi_z(\beta) = 0$ if and only if for every i, we have either $z_i = 0$ or $\gamma_i = 0$. This is equivalent to saying that the element $e^\gamma \in K_\mathbb{C}$ fixes z; we denote the set of such γ by $\mathfrak{k}_z$. It is nothing but the Lie algebra of the stabilizer subgroup of z in K.

Thus we find that Ψ_z has a critical point if and only if there exists an element $\beta \in \mathfrak{k}$ such that $\mu(e^\beta \cdot z) = 0$. That β will not be unique; rather, it will be unique up to translation by the linear space $\mathfrak{k}_z$. But this means that $e^\beta \cdot z$ will be unique. Thus, as in the case of Proposition 5.2, we have reduced the proof of Theorem 5.3 to showing that Ψ_z has a critical point if and only if z is polystable.

The function Ψ_z has a critical point if and only if the following two conditions hold:

(a) for every $\beta \in \mathfrak{k}_z$, $\beta_1\lambda_1 + \ldots + \beta_n\lambda_n = 0$ (otherwise $\Psi_z(\beta)$ can be made arbitrarily negative by choosing $\beta \in \mathfrak{k}_z$, and therefore cannot have a minimum)

(b) for every $\beta \in \mathfrak{k} \smallsetminus \mathfrak{k}_z$, we have $\lim\limits_{t\to\infty} \Psi_z(t\beta) = \infty$.

(Compare this to the analogous statement in the proof of Proposition 5.2, which had two cases depending on whether or not $\lambda = 0$.) Condition (b) is equivalent to the following:

(b′) for every $\beta \in \mathfrak{k} \smallsetminus \mathfrak{k}_z$, either ($\exists\, i$ such that $\beta_i > 0$ and $z_i \neq 0$) or $(\beta_1\lambda_1 + \ldots + \beta_n\lambda_n < 0)$,

which is in turn equivalent to the following:

(b″) for every $\beta \in \mathfrak{k} \smallsetminus \mathfrak{k}_z$, either $\lim\limits_{t\to\infty} e^{t\beta} \cdot z$ does not exist or $(\beta_1\lambda_1 + \ldots + \beta_n\lambda_n < 0)$.

Now let's think about what it means for z to be polystable. We know from Proposition 4.8 that z is *semistable* if and only if the $K_\mathbb{C}$-orbit through $(z, 1) \in \mathbb{C}^n \times \mathbb{C} \cong \mathcal{O}_{\mathbb{C}^n}(-1)$ does not have any points in $\mathbb{C}^n \times \{0\}$ lying in its closure. By Fact 1.14, it is enough to check this on one-parameter subgroups, and we can certainly restrict our attention to one-parameter subgroups in the "noncompact" directions. For directions $\beta \in \mathfrak{k}_z$, this means that $\beta_1\lambda_1 + \ldots + \beta_n\lambda_n \geq 0$. Since $-\beta \in \mathfrak{k}_z$ as well, we might as well say $\beta_1\lambda_1 + \ldots + \beta_n\lambda_n = 0$; that's condition (a). For $\beta \in \mathfrak{k} \smallsetminus \mathfrak{k}_z$, it means that either $\lim_{t\to\infty} e^{t\beta} \cdot z$ does not exist, or $\beta_1\lambda_1 + \ldots + \beta_n\lambda_n \geq 0$. However, if this limit exists and $\beta_1\lambda_1 + \ldots + \beta_n\lambda_n = 0$, then the orbit through z will not be closed in the semistable locus; indeed, the limit will lie in its closure. Thus, for z to be *polystable*, for all $\beta \in \mathfrak{k} \smallsetminus \mathfrak{k}_z$, either the limit does not exist or $\beta_1\lambda_1 + \ldots + \beta_n\lambda_n < 0$. That's condition (b$''$), and that completes the proof. $\square$

5.3 The general (nonabelian) case

Let G be a reductive group acting linearly on a vector space W, which means that we have a homomorphism $\rho : G \to \mathrm{GL}(W)$; for simplicity we will assume that ρ is injective. If we choose a hermitian form on W, then $K := \rho^{-1}(\mathrm{U}(W))$ is the maximal compact subgroup of G; recall from Example 2.5 that the imaginary part of the hermitian form is symplectic.

Consider an element

$$\lambda \in \mathrm{Hom}(G, \mathbb{C}^\times) \cong (\mathfrak{g}^*)^G \cong (\mathfrak{k}^*)^K.$$

First we use λ to define a linearization of the action of G on W. Let $R = \operatorname{Sym} W^* \otimes \mathbb{C}[y]$, so that $\operatorname{Proj} R \cong \operatorname{Spec}(\operatorname{Sym} W^*) \cong W$. We extend the action of G from $\operatorname{Sym} W^*$ to R by putting

$$g \cdot y = \lambda(g) y \qquad \text{for all } g \in G.$$

Next we use λ to define a moment map for the action of K on W. Define $\mu : W \to \mathfrak{k}^*$ by putting

$$\mu(w)(\alpha) = \omega(w, \alpha \cdot w) - \lambda(\alpha) \qquad \text{for all } w \in W \text{ and } \alpha \in \mathfrak{k}.$$

This is a moment map by Examples 2.9 and 2.10. The statement of the general Kempf-Ness theorem will now be identical to the statements of Proposition 5.2 and Theorem 5.3. The proof is also exactly the same; we reproduce it just so that you will have the benefit of seeing the argument formulated in slightly more abstract terms. (In particular, without choosing a basis for W.)

Theorem 5.4 (General Kempf-Ness)**.** *We have an inclusion $\mu^{-1}(0) \subset W^{\mathrm{ps}}$, and for every polystable element $z \in W$, $(G \cdot z) \cap \mu^{-1}(0)$ is a single K-orbit. Thus we obtain a homeomorphism*[6]

$$W/\!\!/K = \mu^{-1}(0)/K \to W^{\mathrm{ps}}/G \;\cong\; W/\!\!/G.$$

Proof. The idea of the proof is identical to that of Theorem 5.3, so we'll be a little bit sketchy. For all $z \in W$, define $\Psi_z : \mathfrak{k} \to \mathbb{R}$ by

$$\Psi_z(\beta) = \frac{1}{2}|| \exp(i\beta) \cdot w||^2 - 2\lambda(\beta).$$

By the same argument that we used above, Ψ_z is convex, and we need to show that $\Psi_z(\beta)$ has a critical point if and only if z is polystable. Again as above, Ψ_z has a minimum if and only if the following two conditions hold:

[6] See Remark 4.12.

(a) for every $\beta \in \mathfrak{k}_z$, $\lambda(\beta) = 0$

(b) for every $\beta \in \mathfrak{k} \smallsetminus \mathfrak{k}_z$, we have $\lim_{t\to\infty} \Psi_z(t\beta) = \infty$.

Once again we can rephrase condition (b) in the following manner:

(b$''$) for every $\beta \in \mathfrak{k} \smallsetminus \mathfrak{k}_z$, either $\lim_{t\to\infty} e^{t\beta} \cdot z$ does not exist or $\lambda(\beta) < 0$.

The argument that conditions (a) and (b$''$) are together equivalent to polystability of z is identical to the argument in the abelian case. □

Remark 5.5. Theorem 5.4 can be adapted to treat quotients of a projective space rather than quotients of a vector space. The most general version deals with quotients of an arbitrary closed subvariety of the product of a vector space and a projective space; that is, anything that can be written as $\operatorname{Proj} R$ for a graded ring R. Two nice expositions of the projective case can be found in lecture notes of Richard Thomas and Chris Woodward.

6 Cohomology

In this section we discuss the cohomology of a smooth toric variety. We begin with its Betti numbers which, as we will see, have a beautiful combinatorial interpretation. We'll then move on to the ring structure, which will give us a good excuse to learn the basics of equivariant cohomology.

6.1 The h-polynomial of a simplicial complex

A **simplicial complex** is a collection Δ of subsets of a fixed finite set S. The only axiom that this collection is required to satisfy is that if $I \in \Delta$ then any subset of I is also in Δ. Elements of Δ are called **faces**. If Δ is a simplicial complex, let $f_k(\Delta)$ be the number of faces of order k (that is, the number of faces containing exactly k elements of I).

Given a polyhedron P with facets $F_1, \ldots, F_n$, let $S = \{1, \ldots, n\}$, and let

$$\Delta(P) := \left\{ I \;\middle|\; \bigcap_{i\in I} F_i \neq \emptyset \right\}.$$

Though this definition makes sense for any P, it is easiest to understand when P is simple. In this case, faces of $\Delta(P)$ of order k correspond bijectively to faces of P of codimension k via the map $F \mapsto I_F$ and its inverse $I \mapsto \bigcap_{i\in I} F_i$. Thus $f_k(P) := f_k(\Delta(P))$ is equal to the number of codimension k faces of P.

Example 6.1. Do a few specific examples. Show how things are less nice in the non-simple case.

Let Δ be a simplicial complex, and let d be the order of the largest face. We define the numbers $h_k(\Delta)$, $k = 0, \ldots, d$, and the corresponding **h-polynomial**, by the following equation:

$$\sum_{k=0}^{d} h_k(\Delta) q^k := h_\Delta(q) := \sum_{k=0}^{d} f_k(\Delta) q^k (1-q)^{d-k}.$$

It is clear that each h-number is a linear combination of the f-numbers. We can also go the other way: our formula implies that

$$\sum_{k=0}^{d} f_k(\Delta)p^k = \sum_{k=0}^{d} h_k(\Delta)p^k(1+p)^{d-k}$$

(just plus in $p = \frac{q}{1-q}$), therefore each f-number is a linear combination of the h-numbers, as well. In other words, we have done nothing but repackage the same information in a slightly different collection of numbers. This repackaging may look completely unmotivated at first, but it is somewhat justified by the following construction.

Let

$$\mathbb{C}[\Delta] := \mathbb{C}[x_i \mid i \in S] \Big/ \left\langle \prod_{i \in I} x_i \;\middle|\; I \notin \Delta \right\rangle$$

be the **Stanley-Reisner ring** of Δ. This ring is graded, with $\deg(x_i) = 1$ for all i.

Proposition 6.2. *The ring $\mathbb{C}[\Delta]$ has Hilbert series $\frac{h_\Delta(q)}{(1-q)^d}$.*

Proof. We have

$$\frac{h_\Delta(q)}{(1-q)^d} = \sum_{i=0}^{d} f_i(\Delta)\left(\frac{q}{1-q}\right)^i = \sum_{I \in \Delta}\left(\frac{q}{1-q}\right)^{|I|}.$$

It's clear that the summand indexed by I is exactly the Hilbert series of the linear span of the monomials in $\mathbb{C}[\Delta]$ with support I. $\square$

Example 6.3. Do a few specific examples.

It is a general problem in combinatorics to classify the h-polynomials that can arise from various specific types of simplicial complexes. In the next section we will give a complete classification of h-polynomials associated to simple polytopes.

6.2 The Poincaré polynomial of a toric variety

In this section we prove the following result, which relates the h-polynomial of a Delzant polytope P to the topology of the toric variety $X(P)$. For any space X, let

$$\mathrm{Poin}_X(t) := \sum_{k=0}^{\dim X} \dim H^k(X;\mathbb{Q})\, t^k$$

be the Poincaré polynomial of X.

Theorem 6.4 (Danilov)**.** *For any Delzant polytope P,* $\mathrm{Poin}_{X(P)}(t) = h_{\Delta(P)}(t^2)$.

Example 6.5. Do some examples.

Proof of Theorem 6.4. Choose an element $a \in V^*$ that is non-constant on every edge of P, and for each face $F \subset P$ let $v_F \in F$ be the unique point at which a attains its minimum. For each vertex v, let

$$P_v := \bigcup_{v_F = v} \mathring{F} \subset P$$

and let $C_v \subset X(P)$ be the preimage of P_v. It is clear from the first construction that C_v is homeomorphic to an ball of dimension $2d_v := 2\dim(P_v)$, and it's not hard to see that these guys provide a cell decomposition of $X(P)$. Since all of the cells are even-dimensional, all of the boundary maps are zero, so we can compute the Poincaré polynomial simply by counting cells of various dimensions. That is, we have

$$\mathrm{Poin}_{X(P)}(t) = \sum_v t^{\dim(C_v)} = \sum_v t^{2d_v}.$$

For any v, P_v is made up of one vertex (namely v), d_v edge interiors, $\binom{d_v}{2}$ 2-face interiors, and more generally $\binom{d_v}{k}$ k-face interiors. We therefore have

$$t^{2d_v} = \big((t^2-1)+1\big)^{d_v} = \sum_{k=0}^{d_v} \binom{d_v}{k} (t^2-1)^k = \sum_{v_F=v} (t^2-1)^{\dim F}.$$

Putting it all together, we have

$$\mathrm{Poin}_{X(P)}(t) = \sum_v t^{2d_v} = \sum_v \sum_{v_F=v} (t^2-1)^{\dim F} = \sum_F (t^2-1)^{\dim F} = \sum_{k=0}^{d} f_k(\Delta(P))(t^2-1)^{d-k}.$$

Since $X(P)$ is smooth, compact, and oriented, it satisfies Poincaré duality, which means that $\mathrm{Poin}_{X(P)}(t) = t^{2d}\,\mathrm{Poin}_{X(P)}(t^{-1})$. This means that we have

$$\mathrm{Poin}_{X(P)}(t) = t^{2d} \sum_{k=0}^{d} f_k(\Delta(P))(t^{-2}-1)^{d-k} = \sum_{k=0}^{d} f_k(\Delta(P)) t^{2k} (1-t^2)^{d-k} = h_{\Delta(P)}(t^2).$$

This completes the proof. □

Next, we use Theorem 6.4 to prove the exact same statement for (unbounded) polyhedra.

Theorem 6.6. *For any Delzant polyhedron P,* $\mathrm{Poin}_{X(P)}(t) = h_{\Delta(P)}(t^2)$.

Example 6.7. Do some non-compact examples.

Proof of Theorem 6.6. Let $\bar{P}$ be a polytope attained from P by intersecting with a single half-space that contains all of the vertices of P, and let $\partial\bar{P}$ be the intersection of P with the boundary of the half space. Let $X = X(P)$, $\bar{X} = X(\bar{P})$, and $\partial X = X(\partial\bar{P})$. We then have $\bar{X} \smallsetminus \partial\bar{X} \cong X$.

Consider the long exact sequence in cohomology associated to the pair $(\bar{X}, \partial\bar{X})$. We have

$H^*(\bar{X}, \partial\bar{X}) \cong H^*_c(X)$, so this sequence takes the form

$$0 = H^{2j-1}(\partial\bar{X}) \to H^{2j}_c(X) \to H^{2j}(\bar{X}) \to H^{2j}(\partial\bar{X}) \to H^{2j+1}_c(X) \to H^{2j+1}(\bar{X}) = 0. \qquad (3)$$

Since $\partial\bar{X}$ is a CW-subcomplex of $\bar{X}$ with all cells of even dimension, the map $H^{2j}(\bar{X}) \to H^{2j}(\partial\bar{X})$ is surjective. This leaves us with $H^{2j}_c(X) \cong \ker\left(H^{2j}(\bar{X}) \to H^{2j}(\partial\bar{X})\right)$ and $H^{2j+1}_c(X) = 0$ for all j. We now use Poincaré duality for the noncompact manifold X to conclude that $H^k(X)$ is dual to $H^{2d-k}_c(X)$, and therefore

$$\mathrm{Poin}_X(t) = t^{2d}\sum_{k=0}^{2d} \dim H^k_c(X) t^{-k} = t^{2d}\,\mathrm{Poin}_{\bar{X}}(t^{-1}) - t^{2d}\,\mathrm{Poin}_{\partial\bar{X}}(t^{-1}) = \mathrm{Poin}_{\bar{X}}(t) - t^2\,\mathrm{Poin}_{\partial\bar{X}}(t).$$

By Theorem 6.4, we have $\mathrm{Poin}_{\bar{X}}(t) = h_{\Delta(\bar{P})}(t^2)$ and $\mathrm{Poin}_{\partial\bar{X}}(t) = h_{\Delta(\partial\bar{P})}(t^2)$. Putting it all together, we have

$$\begin{aligned}
\mathrm{Poin}_X(t) &= \mathrm{Poin}_{\bar{X}}(t) - t^2\,\mathrm{Poin}_{\partial\bar{X}}(t) \\
&= \sum_{F\subset\bar{P}} t^{2(d-\dim F)}(1-t^2)^{\dim F} - t^2 \sum_{F\subset\partial\bar{P}} t^{2(d-1-\dim F)}(1-t^2)^{\dim F} \\
&= \sum_{F\subset\bar{P}} t^{2(d-\dim F)}(1-t^2)^{\dim F} - \sum_{F\subset\partial\bar{P}} t^{2(d-\dim F)}(1-t^2)^{\dim F} \\
&= \sum_{F\subset P} t^{2(d-\dim F)}(1-t^2)^{\dim F} \\
&= h_{\Delta(P)}(t^2).
\end{aligned}$$

This completes the proof. □

Remark 6.8. Our proofs of Theorems 6.4 and 6.6 use the definition of the h-polynomial, which is a little bit lame. Proposition 6.2 provides a more natural way to think about the h-polynomial, and it would be nice to prove Theorems 6.4 and 6.6 from this perspective. Indeed, we will do this in Section 6.5 (see in particular Remark 6.22) using equivariant cohomology.

Theorem 6.4 has some very nice numerical consequences. A sequence $g_0, g_1, \ldots, g_m$ is called an **M-sequence** if there exists a commutative, graded ring R, generated in degree 1, such that $g_k = \dim R_k$.

Theorem 6.9 (Stanley)**.** *Let P be a Delzant polytope, and let $h_k = h_k(\Delta(P))$. Then*

- $h_k = h_{d-k}$ *for all* k
- $h_0, h_1 - h_0, h_2 - h_1, \ldots, h_{\lfloor \frac{d}{2}\rfloor} - h_{\lfloor \frac{d}{2}\rfloor - 1}$ *is an M-sequence.*

Proof. The first collection of equalities comes from Theorem 6.4 and Poincaré duality. The second comes from the Lefschetz hyperplane theorem, which says that for all $k \le \lfloor \frac{d}{2} \rfloor$, the map from $H^{2k-2}(X(P))$ to $H^{2k}(X(P))$ given by multiplication by the Euler class $e := e(\mathcal{O}_X(1))$ is injective.

Thus $h_k - h_{k-1}$ is the dimension of the degree $2k$ piece of the ring $H^*(X(P))/\langle e\rangle$. Thus, to conclude that we have an M-sequence, we only need to show that $H^*(X(P))$ is generated in degree 2. Indeed, it's easy to see that the classes $\{[X(F)] \mid F \subset P \text{ a facet}\}$ generate $H^*(X(P))$ multiplicatively (more on this in the next section), so we're done. $\square$

Remark 6.10. In fact, Theorems 6.4 and 6.9 can be extended from Delzant polytopes (smooth compact toric varieties) to simple polytopes (compact toric orbifolds). The ideas of the proofs are the same, though the details are somewhat more technical.

Remark 6.11. Theorem 6.9 was originally conjectured by McMullen and proven by Stanley (whose main contribution was to interpret Danilov's work in the appropriate combinatorial language). In fact, McMullen had conjectured the converse of Theorem 6.9, as well; that is, he conjectured that any $h_0, \ldots, h_d$ with those two properties are the h-numbers of some simple polytope. This result was proven by Billera and Lee a year or two after Stanley's proof of Theorem 6.9.

6.3 Cohomology basics

In this section we review a few properties of cohomology that everyone should know. Let X be an oriented manifold. Every (not necessarily compact) closed submanifold $Y \subset X$ of codimension k defines a class $[Y] \in H^k(X)$. You should think of this class as the class that takes a closed k-cycle to the intersection of that k-cycle with Y (though of course this only works when the k-cycle is transverse to Y). Formally, it is the pushforward to X of the class $1 \in H^*(Y)$. Here are some properties that this construction has.

(a) If Y_1 and Y_2 intersect transversely, then $[Y_1] \cdot [Y_2] = [Y_1 \cap Y_2]$.

(b) If X is a finite CW-complex and the closures of the cells are all submanifolds, then these submanifolds additively generate $H^*(X)$.

(c) If $f : Z \to X$ is transverse to $Y \subset X$, then $f^*[Y] = [f^{-1}Y] \in H^*(Z)$.

Let's think about the implications of these properties for toric varieties.

Proposition 6.12. *For any Delzant polyhedron P, the classes $\{[X(F)] \mid F \subset P \text{ a face}\}$ additively generate $H^*(X(P))$. For any two faces F and F' that intersect transversely in P,*

$$[X(F)] \cdot [X(F')] = [X(F \cap F')].$$

In particular, the classes $\{[X(F)] \mid F \subset P \text{ a facet}\}$ multiplicatively generate $H^*(X(P))$.

Proof. If P is a polytope, then the additive generation follows from property (b), and the rest follows from property (a). If P is an unbounded polyhedron, then property (b) no longer applies. However, if we can prove additive generation by some other means, then the rest of the proposition will again follow from property (a).

Consider the polytopes $\bar{P}$ and $\partial\bar{P}$ that we introduced in the proof of Theorem 6.6. As before, we let $X = X(P)$, $\bar{X} = X(\bar{P})$, and $\partial\bar{X} = X(\partial\bar{P})$, and we observe that $\bar{X} \smallsetminus \partial\bar{X}$ is diffeomorphic to X. Property (c) implies that the face classes in $H^*(\bar{X})$ restrict to the face classes in $H^*(X)$, thus it is sufficient to prove that the restriction map from $H^*(\bar{X})$ to $H^*(X)$ is surjective. By Poincaré duality, this is equivalent to the statement that $H^*_c(X)$ includes into $H^*(\bar{X})$, and we saw this from the exact sequence (3). $\square$

Remark 6.13. Proposition 6.12 holds for simplicial polyhedra, as well, provided that we work with rational coefficients. Once again the proof is more technical because it involves working with orbifolds, but the ideas are the same.

6.4 Equivariant cohomology basics

Let K be a topological group acting on a space X. Let EK be a contractible space on which K acts freely, and let $X_K := (X \times EK)/K$. The space X_G is often called the **Borel space** or the **homotopy quotient** of X by K. If K acts freely on X, then X_K is homotopy equivalent to X/K. At the other extreme, if K acts trivially on X, then $X_K \cong X \times BK$, where $BK = EK/K$ is the homotopy quotient of a point. More generally, X_K is a fiber bundle over BK with fiber X, and the nontriviality of the fiber bundle exactly encodes the nontriviality of the action of K on X.

Definition 6.14. We define the **equivariant cohomology ring** $H^*_K(X) := H^*(X_K)$.

Remark 6.15. If G is a reductive algebraic group and K is its maximal compact subgroup, then EG is a contractible space on which $K \subset G$ acts trivially, and any X_K is a G/K-bundle over X_G. Since G/K is contractible (think about the case of a torus), $H^*_K(X)$ is canonically isomorphic to $H^*_G(X)$. So the moral is that equivariant cohomology doesn't care whether we work with a compact group or with its complexification.

The most important case is when X is a single point and $X_K = BK$. For example, if $K = U(1)$, we can take $EG = S^{2\infty+1}$ with the Hopf action, so $BK = \mathbb{P}^\infty$. We then have $H^*_K(pt) = \mathbb{Z}[u]$, where $u \in H^2_K(pt)$ is the Euler class of the line bundle $\mathcal{O}_{\mathbb{P}^\infty}(1) = (EK \times \mathbb{C})/\mathbb{C}^\times$. If $K = T^n$, then $BK = (\mathbb{P}^\infty)^n$, and $H^*_K(pt) = \mathbb{Z}[u_1, \ldots, u_n]$. More invariantly, we claim that if T is a torus, then $H^*_T(pt)$ is canonically isomorphic to $\operatorname{Sym} \mathfrak{t}^*_{\mathbb{Z}}$. The map in degree 2 is given by taking $a \in \mathfrak{t}^*_{\mathbb{Z}} \cong \operatorname{Hom}(T, U(1))$ to the Euler class of the line bundle $(ET \times \mathbb{C}_a)/T$. This is clearly a group homomorphism, and one can see that it is an isomorphism by reducing to the case of a 1-dimensional torus, which we have already done.

The reason why this is the most important case is that for any K-space X, X_K maps to BK, and therefore $H^*_K(X) = H^*(X_K)$ is an algebra over $H^*_K(pt) = H^*(BK)$. More generally, any K-equivariant map $f : X \to Y$ induces a map $X_K \to Y_K$, and therefore a homomorphism $f^* : H^*_K(Y) \to H^*_K(X)$. Since every K-space admits a unique equivariant map to a point, f^* is compatible with the algebra structure. In other words, $H^*_K(-)$ is a contravariant functor from K-spaces to $H^*_K(pt)$-algebras. This is true even when K is the trivial group, but it's a much more interesting statement in the equivariant setting, since $H_K^*(pt)$ can be nonzero in arbitrarily high degree.

If $Y \subset X$ is a K-equivariant closed submanifold of codimension k, then $Y_K \subset X_K$ has codimension k, so we can define

$$[Y]_K := [Y_K] \in H^*(X_K) = H_K^k(X).$$

Just as in the non-equivariant case, we have $[Y_1]_K \cdot [Y_2]_K = [Y_1 \cap Y_2]_K$ if they intersect transversely, and $f^*[Y]_K = [f^{-1}Y]_K$ if $f : Z \to X$ is a K-equivariant map that's transverse to Y. The inclusion of X into X_K (as a fiber of the projection to BK) induces a **forgetful map** $H_K^*(X) \to H^*(X)$, and it is clear that the class $[Y]_K$ maps to the class $[Y]$.

Remark 6.16. An important special case of this construction is when $K = T$ is a torus and $X = \mathbb{C}_a$ is the one-dimensional representation of T defined by $a \in \mathfrak{t}_{\mathbb{Z}}^* \cong \mathrm{Hom}(T, U(1))$. Then the isomorphism

$$H_T^2(\mathbb{C}_a) \cong H_T^2(pt) \cong \mathfrak{t}^*$$

takes the class $[0]_T$ to a. This follows from the fact that the Euler class of a vector bundle is (by definition) equal to the class in the total space represented by the zero section.

Let K be a connected group, so that BK is simply-connected. (In general, the fundamental group of BK is isomorphic to the quotient of K by the connected component containing the identity.) Since X_K is a fiber bundle over BK with fiber X, there is a spectral sequence with E_2 page $H^*(BK) \otimes H^*(X)$. We say that X is **equivariantly formal** if this spectral sequence collapses. This is equivalent to saying that $H_K^*(X) \cong H_K^*(pt) \otimes H^*(X)$ as a module over $H_K^*(pt)$ (non-canonically, and not as a ring!), and that the forgetful map to $H^*(X)$ is the surjection given by setting all positive-degree elements of $H_K^*(pt)$ equal to zero.

Example 6.17. Let $K = U(1)$ acting on $X = \mathbb{P}^1$. In this case we know in advance that the spectral sequence will collapse, because both $H^*(X)$ and $H_K^*(pt)$ exist entirely in even degree! Let's compute the equivariant cohomology ring explicitly.

Let $n, s \in \mathbb{P}^1$ be the two fixed points (the north and south poles). Then we get a map $\mathbb{Z}[x,y]/\langle xy \rangle \to H_{U(1)}^*(\mathbb{P}^1)$ sending x to $[n]_{U(1)}$ and y to $[s]_{U(1)}$. We will show that this map is in fact an isomorphism.

To prove injectivity, we need to show that for all positive k, the k^{th} powers of $[n]_{U(1)}$ and $[s]_{U(1)}$ are nonzero in $H_{U(1)}^*(\mathbb{P}^1)$. Consider the restriction map

$$H_{U(1)}^*(\mathbb{P}^1) \to H_{U(1)}^*(\mathbb{P}^1 \smallsetminus s) \cong H_{\mathrm{U}(1)}^*(\mathbb{C}_1).$$

The open inclusion is transverse to n, so this restriction map takes $[n]_{\mathrm{U}(1)}$ to $[0]_{\mathrm{U}(1)}$, which is the generator of $H_{\mathrm{U}(1)}^*(\mathbb{C}_1)$ by Remark 6.16. Thus $[n]_{\mathrm{U}(1)}^k$ must be nonzero for all k. The argument for $[s]_{\mathrm{U}(1)}$ is identical.

To prove surjectivity, it is enough to show that the subring generated by $[n]_{U(1)}$ and $[s]_{U(1)}$ surjects onto $H^*(\mathbb{P}^1)$ and contains the image of $H_{\mathrm{U}(1)}^*(pt)$. The first statement is clear, since both

$[n]_{U(1)}$ and $[s]_{U(1)}$ map to the class of a point. Since $[n]_{U(1)} - [s]_{U(1)}$ lies in the kernel of this projection, it must be a multiple of the generator $u \in H^2_{\mathrm{U}(1)}(pt)$. Again, consider the restriction map

$$H^*_{U(1)}(\mathbb{P}^1) \to H^*_{U(1)}(\mathbb{P}^1 \smallsetminus s) \cong H^*_{\mathrm{U}(1)}(\mathbb{C}_1) \cong H^*_{\mathrm{U}(1)}(pt).$$

The class $[n]_{U(1)} - [s]_{U(1)}$ maps to the generator of $H^2_{\mathrm{U}(1)}(pt)$, therefore it must be equal to u.

So now we have seen that $H^*_{U(1)}(\mathbb{P}^1)$ is isomorphic to $\mathbb{Z}[x,y]/\langle xy\rangle$, and that the algebra structure over $H^*_{\mathrm{U}(1)}(pt) \cong \mathbb{Z}[u]$ is given by $u \mapsto x - y$. As a module over $\mathbb{Z}[u]$, it is a free module generated by 1 and x, and therefore isomorphic to $\mathbb{Z}[u] \otimes H^*(\mathbb{P}^1)$. But it is also a free module generated by 1 and y, or by 1 and $5x - 6y$, and so on. In other words, it is isomorphic to $\mathbb{Z}[u] \otimes H^*(\mathbb{P}^1)$, but not in a canonical way. Furthermore, it is not isomorphic as a ring; indeed, $\mathbb{Z}[u] \otimes H^*(\mathbb{P}^1)$ has a nonzero class in degree 2 whose square is trivial, but $\mathbb{Z}[x,y]/\langle xy\rangle$ does not.

We end with the following result, which will be crucial in the next section. Consider a normal subgroup $N \subset K$ with quotient $Q = K/N$. We may take $EQ \times EK$ as our contractible space on which K acts freely, and we have a map $(EQ \times EK)/K \to EQ/K = EQ/Q$, which induces a map from $H^*_Q(pt)$ to $H^*_K(pt)$. If K is a torus, then this is nothing other than the natural inclusion from $\operatorname{Sym} \mathfrak{q}^*_{\mathbb{Z}}$ to $\operatorname{Sym} \mathfrak{k}^*_{\mathbb{Z}}$.

Lemma 6.18. *Suppose that X is a K-space and that $N \subset K$ acts freely. Then $H^*_K(X) \cong H^*_Q(X/N)$ as $H^*_Q(pt)$-algebras.*

Proof. We have

$$X_K = (EK \times X)/K \cong \Big((EK \times X)/N\Big)\Big/Q,$$

thus

$$H^*_K(X) = H^*(X_K) \cong H^*\Big(\Big((EK \times X)/N\Big)\Big/Q\Big) \cong H^*_Q\Big((EK \times X)/N\Big) \cong H^*_Q(X/N),$$

where the last isomorphism comes from the fact that $(EK \times X)/N$ is a vector bundle over X/N. $\square$

6.5 Equivariant cohomology of toric varieties

In this section we will compute $H^*_T(X(P))$ as an algebra over $\operatorname{Sym} \mathfrak{t}^*_{\mathbb{Z}} \cong H^*_T(pt)$, where P is a Delzant polyhedron and T is the torus acting on the toric variety $X(P)$. We will find that the ideas that we need for this computation are no more sophisticated than those that we used in Example 6.17, which was the simplest nontrivial case.

Theorem 6.19. *Consider the map $\mathbb{Z}[u_1, \ldots, u_n] \to H^*_T(X(P))$ taking u_i to $[X(F_i)]_T$. This map is surjective with kernel*

$$I := \Big\langle \prod_{i \in S} u_i \;\Big|\; \bigcap_{i \in S} F_i = \emptyset \Big\rangle.$$

Furthermore, the $\operatorname{Sym} \mathfrak{t}^*_{\mathbb{Z}}$-algebra structure is given by the natural inclusion

$$\operatorname{Sym} \mathfrak{t}^*_{\mathbb{Z}} \subset \operatorname{Sym}(\mathfrak{t}^n)^*_{\mathbb{Z}} \cong \mathbb{Z}[u_1, \ldots, u_n].$$

Proof. Let's start with the last statement. That is, we will show that if we give $\mathbb{Z}[u_1, \ldots, u_n]$ the structure of a $\operatorname{Sym} \mathfrak{t}^*_{\mathbb{Z}}$-algebra in the indicated way, then the map from $\mathbb{Z}[u_1, \ldots, u_n]$ to $H^*_T(X(P))$ taking u_i to $[X(F_i)]_T$ is an algebra homomorphism. To see this, note that the map can be interpreted geometrically as follows:

$$\mathbb{Z}[u_1, \ldots, u_n] \cong \operatorname{Sym}(\mathfrak{t}^n)^*_{\mathbb{Z}} \cong H^*_{T^n}(pt) \cong H^*_{T^n}(\mathbb{C}^n) \to H^*_{T^n}(\mu^{-1}(0)) \cong H^*_{T^n/K}(\mu^{-1}(0)/K) = H^*_T(X_2(P)).$$

The fact that this is our map follows from the observation that if $Z_i = \{z \in \mathbb{C}^n \mid z_i = 0\}$, then u_i is identified with $[Z_i]_{T^n} \in H^*_{T^n}(\mathbb{C}^n)$, and $\big(Z_i \cap \mu^{-1}(0)\big)/K = X(F_i)$. The fact that this map is an algebra homomorphism is a consequence of Lemma 6.18.

Just as in Example 6.17, we know that $X(P)$ is equivariantly formal for degree reasons. Thus to prove surjectivity, it is enough to show that the composition $\mathbb{Z}[u_1, \ldots, u_n] \to H^*_T(X(P)) \to H^*(X(P))$ is surjective. This follows from Proposition 6.12.

It is clear that I is contained in the kernel. To prove that I is equal to the kernel, consider a polynomial $f \in \mathbb{Z}[u_1, \ldots, u_n]$ that is not contained in I. This means that at least one monomial in f is supported on a face of $\Delta(P)$. For notational convenience, let's assume that some monomial is supported on $\{1, \ldots, d\}$ and that $F_1 \cap \ldots \cap F_d$ is a vertex of P, and let x be the corresponding element of $X(P)^T$. We will show that the composition

$$\mathbb{Z}[u_1, \ldots, u_n] \to H^*_T(X(P)) \to H^*_T(x)$$

has kernel $\langle u_{d+1}, \ldots, u_n \rangle$. Since f is not contained in this ideal, it does not map to zero in $H^*_T(x)$, and therefore cannot map to zero in $H^*_T(X(P))$.

It is easy to see that u_i is in the kernel of the composition when $i > d$, since $X(F_i)$ does not contain x. To see that these is nothing else in the kernel, it is enough to see that the composition is surjective; this follows from the fact that it is a homomorphism of $H^*_T(pt)$-algebras. $\square$

Corollary 6.20. *We have a natural isomorphism*

$$H^*(X(P)) \cong \mathbb{Z}[u_1, \ldots, u_n] \Big/ I + J,$$

where

$$I = \Big\langle \prod_{i \in S} u_i \;\Big|\; \bigcap_{i \in S} F_i = \emptyset \Big\rangle$$

and J is the ideal generated by $\mathfrak{t}^*_{\mathbb{Z}} \subset (\mathfrak{t}^n)^*_{\mathbb{Z}} \cong \mathbb{Z}\{u_1, \ldots, u_n\}$.

Example 6.21. Do a planar polytope with 5 facets.

Remark 6.22. Another corollary of Theorem 6.19 is a new and more elegant proof of Theorems 6.4 and 6.6. Theorem 6.19 says that $H^*_T(X(P))$ is isomorphic to the Stanley-Reisner ring of $\Delta(P)$ with degrees doubled. We know from Proposition 6.2 that this ring has Hilbert series

$$\frac{h_{\Delta(P)}(t^2)}{(1-t^2)^{2d}}.$$

On the other hand, since $H_T^*(X(P)) \cong H^*(X(P)) \otimes \operatorname{Sym} \mathfrak{t}_{\mathbb{Z}}^*$ as a graded vector space, it has Hilbert series $\operatorname{Poin}_{X(P)}(t) \cdot (1-t^2)^{-d}$. Hence we must have $h_{\Delta(P)} = \operatorname{Poin}_{X(P)}(t)$.

6.6 What's true more generally?

There were three important ideas that we used to proof Theorem 6.19. The first was that $X(P)$ is equivariantly formal for the action of T. The second was that the natural map from the T^n-equivariant cohomology of $\mathbb{C}^n$ to the T-equivariant cohomology of $X(P)$ is surjective. The third was that nonzero classes in the T-equivariant cohomology of $X(P)$ can be detected by restricting to fixed points of the T-action. In this section I will state (without proof) the most general versions of each of these statements that I know to be true. Let's start with equivariant formality.

Theorem 6.23 (Atiyah-Bott)**.** *Let K be a compact connected group acting on a compact symplectic manifold X with a moment map. Then X is equivariantly formal with rational coefficients.*

Remark 6.24. Note that this also applies to reductive groups acting linearly on smooth projective varieties, since such varieties are all symplectic and the action of the maximal compact subgroup is always hamiltonian.

Note that Theorem 6.23 is not actually sufficient to get equivariant formality for the action of T on $X(P)$. One problem is that $X(P)$ need not be compact, and the other is that we wanted formality over $\mathbb{Z}$, not just over $\mathbb{Q}$. For this we used the following, much easier theorem.

Theorem 6.25. *Let K be a connected group acting on a topological space X with no cohomology in odd degree. Then X is equivariantly formal with integer coefficients.*

Proof. For degree reasons, the spectral sequence has no nontrivial maps! □

Next we move on to surjectivity. Let K be a compact group acting on a symplectic manifold X with moment map $\Phi : X \to \mathfrak{k}^*$. Let $N \subset K$ be a normal subgroup and let $Q = K/N$ be the quotient. Let $\mu = i^* \circ \Phi : X \to \mathfrak{n}^*$ be the moment map for the action of N.

Theorem 6.26 (Kirwan)**.** *If Φ is proper and N acts freely on $\mu^{-1}(0)$, then the restriction map*

$$H_K^*(X) \to H_K^*(\mu^{-1}(0)) \cong H_Q^*(X/\!\!/N)$$

is surjective.

Remark 6.27. Once again, the Kempf-Ness theorem can be used to infer a similar result about linear actions of reductive algebraic groups on smooth projective varieties. Such a result was also proven in a purely algebraic setting by Ellingsrud and Strømme.

Finally, let's continue to localization. I'm not sure to whom this should be attributed, but certainly some form of it goes back to Borel.

Theorem 6.28. *If K is a compact connected group and X is an equivariantly formal K-space, then the restriction map $H_K^*(X) \to H_K^*(X^K)$ is injective.*

Remark 6.29. When K is a torus, Chang and Skjelbred gave a beautiful characterization of the image of this map; this characterization was later "popularized" by Goresky, Kottwitz, and MacPherson. There is a class of spaces known as **GKM spaces** for which X^K is finite and the image of $H_K^*(X)$ in $H_K^*(X^K) \cong \bigoplus_{x \in X^T} \operatorname{Sym} \mathfrak{k}_{\mathbb{Z}}^*$ is particularly easy to describe. In the case of toric varieties, we have

$$H_T^*(X(P)) \cong \bigoplus_{x \in X(P)^T} \operatorname{Sym} V_{\mathbb{Z}} = \bigoplus_{\substack{v \in P \\ \text{a vertex}}} \operatorname{Sym} V_{\mathbb{Z}},$$

and the Chang-Skjelbred-Goresky-Kottwitz-MacPherson condition says that (f_v) lies in the image if and only if, whenever v and v' are connected by an edge, $v - v' \in V$ divides $f_v - f_{v'} \in \operatorname{Sym} V$.