COMPLEXITY IN INVARIANT THEORY

by

Harlan Kadish

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in The University of Michigan
2011

Doctoral Committee:

Professor Harm Derksen, Chair
Professor Mel Hochster
Professor Karen E. Smith
Professor Wayne E. Stark

For my family, for all their support and encouragement,
because they understand why it's neat to just have a talking frog.

ACKNOWLEDGMENTS

Most of all I owe my thanks to my adviser, Harm Derksen, who introduced me to invariant theory and guided these first steps. His creativity, excitement, vast knowledge, and enthusiasm to solve any problem inspires me. It has been an honor to work with him.

I am grateful for my dissertation readers: for Karen Smith and Gopal Prasad, from whose classes I learned algebraic geometry, commutative algebra, and how to teach mathematics; for Mel Hochster, for joining Harm in my preliminary exam and my dissertation defense; and for Wayne Stark for providing cognate assistance. David Wehlau led me to the history of efforts to count generating invariants.

I am especially grateful for the support of the Algebraic Geometry RTG grant. I also wish to thank the friendly professors and staff at the University of Michigan Mathematics Department, my mathematical home, and to thank my undergraduate mentors Edray Goins and David Wales, who led me to its doorstep.

Lastly, I remember my friends Roy, Ricardo, Nina, Michelle, Patrick, Clara, and Rob, for five years of study, cooking, and adventure.

TABLE OF CONTENTS

CHAPTER I

Introduction

This thesis considers computational questions about representations of algebraic groups. Let an algebraic group G act rationally on a finite-dimensional vector space V, that is, by a homomorphism $G \to GL(V)$ defined by regular functions. For example, for an algebraically closed field k, $GL_n(k)$ acts on the $n \times n$ matrices over k by conjugation. Two related questions arise: (1) How can we tell if two points in V lie in the same orbit? (2) What is the geometry of the quotient of V by G, that is, of the set of orbits? Working in the algebraic category, one looks for answers in the ring $k[V]$ of polynomial functions on V. Specially, invariant theory studies the subring $k[V]^G$ of functions that are invariant under translation by G. In 1890 David Hilbert showed that $\mathbb{C}[V]^G$ is finitely generated for the classical matrix groups [19], and in 1893 he outlined a procedure to compute its generators [20]. These results ended an era of furious computation, and his Basis Theorem sent research on a more astract course for the next several decades. The rise of computer technology and new tools in commutative algebra sparked a renaissance, for example, Sturmfels modern and more detailed formulation of Hilbert's algorithm of 1893 [44, p. 177]. Now computational invariant theory considers the complexity of problems in invariant theory and develops algorithms to solve them.

Although degree bounds for the generators of various invariant rings have been known for decades, little is said about the cardinality of minimal generating sets. Estimates of such would provide lower bounds for the runtime of algorithms that compute invariants. For a semisimple group G, choose an irreducible representation of highest weight λ, and consider the irreducible representations of highest weight $n\lambda$. The results herein prove that as n goes to infinity, the cardinality of a minimal set of generating invariants grows faster than any polynomial in n. The same holds when SL_2 acts on the space V_n of binary forms of degree n, the polynomials of the form $a_n x^n + a_{n-1}x^{n-1}y + \cdots + a_0 y^n$. Combinatorial methods yield sub-exponential upper bounds for the growth of generating sets for torus invariants on the binary forms.

On the other hand, this thesis establishes an algorithm that distinguishes orbits in polynomial time, for any algebraic group G and finite-dimensional representation V. Previously known algorithms take longer to run or place restrictions on the group. This new algorithm outputs a finite set $\mathcal{C}$ of functions with the property that $x \in V$ is not in the orbit of y if and only if there exists a function $f \in \mathcal{C}$ such that that $f(x) \neq f(y)$. The functions are constructed with the polynomial operations in $k[V]$ and with a new "quasi-inverse" that computes the multiplicative inverse of a function where defined. For fixed G, the size of $\mathcal{C}$ and the number of steps the algorithm takes are bounded by polynomials in the dimension of V and the degrees appearing in the homomorphism $G \to GL(V)$. Rings of such "quasi-regular functions" are explored in detail, as well. Thus the problem of separating orbits has polynomial complexity. What is more, it follows that the quasi-inverse is a sufficient generalization of the polynomial functions to allow us to separate all orbits for any group G.

CHAPTER II

Background

2.1 Algebraic Group Actions

The objects of study are linear algebraic groups, their representations, their orbits, and the rings of invariant polynomials on representations. Throughout, k will be an algebraically closed field. A *linear algebraic group* is a group whose elements form an affine algebraic variety (possibly not irreducible), such that the group multiplication and the inverse operation correspond to algebraic morphisms: for example, GL_n, SL_n, and other classical matrix groups. What follows is less concerned with the structure of algebraic groups than with their representations.

A *representation* of a group G is a vector space V and a map $\rho : G \to GL(V)$. One refers to the space V or the function ρ as a "representation of G" when the other piece is understood. Henceforth all representations V will have finite dimension and G is a linear algebraic group acting rationally: that is, ρ is a morphism of varieties. A representation V is *irreducible* if it is nontrivial and contains no proper subrepresentation.

The study of representations subsumes the study of group actions on varieties for the following reason.

Proposition II.1. *[6, p. 239] Let G be a linear algebraic group acting rationally on*

an affine variety X. Then there exists a representation V of G and a G-equivariant, closed embedding $X \hookrightarrow V$.

The *orbit* of a point $p \in V$ is the set $G \cdot p = \{g \cdot p \mid g \in G\}$.

Proposition II.2. *[27, p. 60] If a linear algebraic group acts rationally, then its orbits are smooth, irreducible, and open in their closures. The boundary of an orbit is a union of orbits of strictly smaller dimension. Thus closed orbits exist.*

2.2 Invariant Subrings are Finitely Generated

If G acts rationally on an affine variety X over $k = \overline{k}$, then G acts on the coordinate ring $k[X]$ as follows: for $f(x) \in k[X]$ and $\sigma \in G$, $\sigma \cdot f(x) = f(\sigma^{-1} \cdot x)$. (Recall the assumption $k = \overline{k}$, though this is not always necessary). The subring in $k[X]$ of invariant functions is denoted $k[X]^G$. The invariant subring is finitely generated for the reductive groups (see below), but not for general groups. Another, more useful sufficient condition for finite generation is linear reductiveness. An algebraic group G is *linearly reductive* over k if for every rational representation V over k and every $v \in V^G - \{0\}$ there exists a linear invariant function $f \in (V^*)^G$ such that $f(v) \neq 0$.

For an equivalent definition, define a *Reynolds operator* to be a G-equivariant linear projection $\mathcal{R}\colon k[X] \to k[X]^G$. Here, a map f between sets S, T on which G acts is called *G-equivariant* if $f(g \cdot s) = g \cdot f(s)$ for all $s \in S, g \in G$. When G is finite and $\mathrm{char}(k) \nmid |G|$ or G is compact over $\mathbb{C}$ with a Haar measure, the Reynolds operator is just averaging over G. Explicit formulas for other groups exist classically. Linear reductiveness relates to the Reynold's operator as follows.

Theorem II.3. *[6, p. 46] The following are equivalent for a linear algebraic group G over a field k:*

(a) G is linearly reductive over k.

(b) For every affine variety X over k with a G action, there exists a unique Reynolds operator $\mathcal{R}\colon k[X] \to k[X]^G$.

(c) For every rational representation V and subrepresentation $W \subseteq V$, there exists a subrepresentation $W' \subseteq V$ such that $V = W \oplus W'$.

(d) Every rational representation V decomposes into a direct sum of irreducible representations $V = V_1 \oplus \cdots \oplus V_n$.

In 1890, Hilbert proved that $\mathbb{C}[V]^G$ is finitely generated for GL_n, SL_n, and other classical groups [19]. In fact the existence of a Reynolds operator is sufficient to apply the method of his proof.

Theorem II.4. *[6, p. 49] If G is a linearly reductive group and V a rational representation, then $k[V]^G$ is finitely generated.*

The result holds when a linearly reductive group acts on an affine variety X, by finding a closed embedding $i\colon X \hookrightarrow V$ into a finite-dimensional vector space V [6, p. 48]. Indeed, any finite-dimensional, G-stable subspace $W \subseteq k[X]$ is the image under i^* of a G-stable subspace $Z \subseteq k[V]$. Since G is linearly reductive, Z has a unique decomposition into subrepresentations, and $i^*(Z^G) = W^G$. Then $i^*\colon k[V]^G \xrightarrow{\sim} k[X]^G$.

Other properties of an algebraic group G can ensure that a representation V has $k[V]^G$ finitely generated. Define the *radical* $R(G)$ of a linear algebraic group to be the largest, normal, connected, solvable subgroup. When the subgroup of unipotent elements of $R(G)$ is trivial, G is called *reductive*. Examples of reductive groups include $GL_n, SL_n, O_n, SO_n, SP_n$, finite groups, tori, and any group with $R(G) = \{1\}$ [6, p. 50]. Groups in this last class are called *semisimple*. A linear algebraic group is called *geometrically reductive* if for every rational representation V and nonzero fixed point $v \in V$, there exists a nonconstant homogeneous $f \in k[V]^G$ such that

$f(v) \neq 0$. Note linear reductive implies geometrically reductive.

These three notions of reductive are related as follows. In 1963 Nagata showed that $k[X]^G$ is finitely generated when a geometrically reductive group acts on an affine variety X [40]. With Miyata he next showed that in characteristic zero, reductive and linearly reductive are equivalent, and that geometrically reductive implies reductive [41]. In 1974 Haboush responded that reductive implies geometrically reductive [15], so that the two are equivalent. Thus the three notions are equivalent in characteristic zero, but linear reductiveness is stronger than the others in positive characteristic: for example, a cyclic p-group in characteristic $p > 0$ is geometrically reductive by not linearly reductive [6, p. 51]. Regardless, any of the notions is sufficient for $k[V]^G$ to be finitely generated. On the other hand, in 1959 Nagata provided an example where $k[V]^G$ is not finitely generated, and in doing so found a counterexample to Hilbert's fourteenth problem: whether every subfield $L \subset k(x_1, \ldots, x_n)$ has $L \cap k[x_1, \ldots, x_n]$ finitely generated [39].

2.3 Highest-Weight Representations and the Ring of Covariants

The structure of an algebraic group G can help distinguish its representations. As always, work over a field $k = \overline{k}$, for geometric reasoning, but in this section assume also that $\mathrm{char}(k) = 0$, to avoid trivial cases. The maximal connected solvable subgroups of G are called *Borel subgroups*. To distinguish the representations of a reductive group G, choose a Borel subgroup B. A *torus* is an algebraic group isomorphic to $(k^*)^n$, that is, the diagonal subgroup of $GL_n(k)$ for some field k. Then B can be written $B = T \ltimes U$, where U is a unipotent group and T is a torus. The torus action on a representation V decomposes V into a direct sum of weight spaces. Now, since B is connected and solvable, the Lie-Kolchin Theorem [27] guarantees

the existence of a 1-dimensional subspace stable under the action of B. If V is irreducible under the action of G, then this subspace V^B is unique, and its weight with respect to the chosen torus T will be the highest in the ordering on the abstract weights of T. A generator for $V^B = V^U$ is called a *highest weight vector* of V. The highest weight of a representation is dominant in the root system for G. In fact, if λ is a dominant weight, then there is a unique, irreducible representation, up to G-equivariant isomorphism, whose highest weight is λ [27, p. 188].

If G is reductive, then its highest weight representations can be collected and given a ring structure; the development here follows [6, p. 156]. Let G act on an affine variety X and any finite-dimensional G-module W. Define a "covariant of X with values in W" to be a G-equivariant morphism $X \to W$. The set of covariants for such X and W is denoted $\mathrm{Mor}(X, W)^G$, the morphisms invariant under the G-action. Popov and Vinberg showed that $\mathrm{Mor}(X, W)^G$ is a finitely-generated module over $k[X]^G$.

The module structure becomes more clear with the isomorphism

$$\mathrm{Mor}(X, W)^G \cong (W \otimes k[X])^G.$$

Indeed, a covariant $\phi\colon X \to W$ yields a ring homomorphism $\phi^*\colon k[W] \to k[X]$, or what is the same, $\phi^*\colon S(W^*) \to k[X]$, as the symmetric algebra on W^* is canonically isomorphic to $k[W]$. Any such map is determined by its values on W^*, and a linear map $W^* \to k[X]$ determines a ring homomorphism $S(W^*) \to k[X]$. Thus $\mathrm{Mor}(X, W)^G \cong \mathrm{Hom}(W^*, k[X])^G$ as $k[X]^G$-modules, and the above isomorphism follows.

Recall that to discuss highest weight representations of a reductive group G, one fixes a Borel subgroup B, a maximal torus $T \subset B$, and a maximal unipotent subgroup $U \subseteq B$ such that $B = T \ltimes U$. If X is an affine variety with a G action, the ring of

U-invariant polynomials $k[X]^U$ is called the *ring of covariants* of X, for the following reason: Write V_λ for the irreducible representation of G with highest weight λ, whose λ weight space is spanned by a vector v_λ. Then $G \cdot v_\lambda$ spans V_λ. Recall a covariant $\phi \in \mathrm{Mor}(X, V_\lambda)^G \cong \mathrm{Hom}(V_\lambda^*, k[X])^G$ is determined by $\phi^* : V_\lambda^* \to k[X]$. If w spans the highest weight space of V_λ^*, then ϕ^* is determined by the choice of $\phi^*(w)$, which lies in $k[X]^U$.

Grosshans showed that if G is reductive and $\mathrm{char}(k) = 0$, then a ring of covariants $k[X]^U$ is finitely generated [14]. Now consider the ring of covariants for G on itself:

Proposition II.5. *Let G be a linearly reductive group over k, and let V_λ denote an irreducible representation of G with highest weight λ. For fixed $B = T \ltimes U \subset G$, let U act on G by right-multiplication. Then if $X(T)_+$ denotes the dominant weights,*

$$k[G]^U = k[G/U] = \bigoplus_{\lambda \in X(T)_+} V_\lambda$$

as left G-modules. If $\mathrm{char}(k) = 0$, then $k[G]^U$ is a finitely-generated ring.

Proof. If U acts on G on the right, then G/U is a quasi-affine variety, and $k[G/U] = k[G]^U$. Choose a highest weight representation V_λ. Then the module of covariants of G/U with values in V_λ is

$$\mathrm{Mor}(G/U, V_\lambda)^G \cong (k[G/U] \otimes V_\lambda)^G \cong (k[G]^U \otimes V_\lambda)^G.$$

First show that the module of covariants, on the left side above, has dimension one: any morphism $\phi : G/U \to V_\lambda$ is determined by the image of the identity coset $eU \in G/U$. This element is fixed by the usual U-action on V_λ, whence $\phi(eU)$ lies in the one-dimensional space V_λ^U. Thus $(k[G]^U \otimes V_\lambda)^G$ also has dimension one, as does $\mathrm{Hom}(V_\lambda, k[G]^U)^G$, because $\dim V_\lambda < \infty$. Therefore, each V_λ occurs exactly once in $k[G]^U$.

Now, any $f \in k[G]^U$ lies in a finite-dimensional, G-stable subspace F. Since G is linearly reductive, F is a direct sum of highest weight representations, and the G-module isomorphism $k[G]^U = \bigoplus_{\lambda \in X(T)_+} V_\lambda$ is proved. Grosshans in [14] showed that $k[G]^U$ is finitely generated for reductive G in characteristic 0. □

2.4 The Categorical Quotient and Separating Invariants

When G acts rationally on an affine variety X and $k[X]^G$ is finitely generated, write $X/\!\!/G$ to denote the variety with coordinate ring $k[X]^G$, a notation established in the book [6] and papers of Derksen and Kemper. Then the dominant morphism $\pi\colon X \to X/\!\!/G$ is in fact surjective, and $X/\!\!/G$ has the quotient topology; call $X/\!\!/G$ the *categorical quotient*. Each fiber of π contains exactly one closed orbit, which lies in the closure of all orbits in that fiber. One can then apply the theorem on fiber dimension to the closed orbits to determine the dimension of invariant rings.

Since invariant polynomials are constant on orbits, they may separate the orbits of a group action in applications. If separating orbits is the goal, though, one does not need to compute generators for the entire invariant ring. A subset $S \subseteq k[X]^G$ is called *separating* if it has the following property for every pair of points $x, y \in X$: if there exists $f \in k[X]^G$ such that $f(x) \neq f(y)$ (so f separates the orbits of x and y), then there exists a $g \in S$ such that $g(x) \neq g(y)$. That is, the polynomials in S separate as many orbits as the polynomials in $k[X]^G$. When the categorical quotient exists, if x and y have distinct closed orbits, then some invariant polynomial must separate the orbits, so a member of a separating set must separate the orbits. More generally, there is the following:

Theorem II.6. *[6, p. 58] Let X be an affine variety and G a group of automorphisms of $k[X]$. Then there exists a finite separating set $S \subset k[X]^G$.*

Proof. Let $I \subset k[X] \otimes k[X]$ be the ideal generated by the elements $f \otimes 1 - 1 \otimes f$ for all $f \in k[X]^G$. Since $k[X] \otimes k[X]$ is Noetherian, the ideal I is finitely generated. Thus there exists a finite, generating subset

$$I = (f_1 \otimes 1 - 1 \otimes f_1, \ldots, f_m \otimes 1 - 1 \otimes f_m)$$

with the $f_i \in k[X]^G$. Claim that the f_i appearing here form a separating subset for $k[X]^G$. For proof, choose $x, y \in X$ such that there exists $f \in k[X]^G$ with $f(x) \neq f(y)$. It remains to show that x and y are separated by one of the f_i. Since $f \otimes 1 - 1 \otimes f \in I$, one has, for some $g_i \in k[X] \otimes k[X]$,

$$f \otimes 1 - 1 \otimes f = \sum_{i=1}^{m} g_i (f_i \otimes 1 - 1 \otimes f_i).$$

For the fixed $x, y \in X$, there is an evaluation homomorphism $\varphi \colon k[X] \otimes k[X] \to k$ defined by $\varphi \colon g \otimes h \mapsto g(x)h(y)$. Applying φ to the above equation,

$$\sum_{i=1}^{m} \varphi(g_i) \left(f_i(x) \otimes 1 - 1 \otimes f_i(y)\right) = \varphi(f \otimes 1 - 1 \otimes f) = f(x) - f(y) \neq 0.$$

So for some i, it follows $f_i(x) \neq f_i(y)$, completing the proof. □

Assume G is reductive and V is a rational representation. Then $\{0\} \subset V$ is always a closed orbit and an interesting point for the geometry of group actions. The *nullcone* $\mathcal{N}_V$ is defined as

$$\mathcal{N}_V = \{v \in V \mid f(v) = 0 \text{ for all } f \in k[V]^G_+\}.$$

That is, $\mathcal{N}_V$ is the set of points on which all homogeneous invariants vanish, whence $\mathcal{N}_V$ is the fiber $\pi^{-1}(0)$ of the categorical quotient.

Lemma II.7. *[6, p. 60] A point $v \in V$ lies in the nullcone $\mathcal{N}_V$ if and only if $\overline{G \cdot v}$ contains* 0.

A useful tool for determining if orbits are closed is the Hilbert-Mumford criterion [20, 38]. Define a *1-parameter subgroup* of an algebraic group G to be a homomorphism $\lambda\colon k^* \to G$. Sometimes the phrase "1-parameter subgroup λ" denotes the image of the map. The following formulation follows from Kempf's work on instability [35]:

Theorem II.8. *Let V be a rational representation of a reductive group G, and choose $v, w \in V$. Then $w \in \overline{G \cdot v}$ if and only if there exists a 1-parameter subgroup λ such that $w \in \overline{\lambda \cdot v}$.*

2.5 The Cohen-Macaulay Property

This property of the invariant rings of linearly reductive groups elucidates their structure and aids computation. One can define the Cohen-Macaulay property R-modules using the notion of depth, when R is graded or local, but the results below focus on the graded case, taking R as an R-module. First, recall some commutative algebra. If $R = \oplus_{d=0}^{\infty} R_d$ is a graded algebra over a field $k = R_0$, then $f_1, \ldots, f_n \in R$ are a *homogeneous system of parameters* if both

- $f_1, \ldots, f_n$ are algebraically independent over k,
- R is a finitely generated module over $k[f_1, \ldots, f_n]$.

In particular, $\dim(R) = n$. Under the assumption $k = \overline{k}$ so that k is infinite, a finitely generated graded algebra over a field always has a homogeneous system of parameters. This fact follows from Noether's Normalization Lemma [6, p. 61]. Now, a sequence $f_1, \ldots, f_n \in R$ is called R-regular (or just regular) if

$$R/(f_1, \ldots, f_n) \neq 0$$

and multiplication by f_i on $R/(f_1, \ldots, f_{i-1})R$ gives an injective map, for $i = 1, \ldots, n$. One can define the *Cohen-Macaulay property* as follows:

Proposition II.9. *[32] For a Noetherian graded algebra R over a field $k = R_0$, the following are equivalent:*

(a) R is Cohen-Macaulay.

(b) Every homogeneous system of parameters is an R-regular sequence.

(c) If $f_1, \ldots, f_n$ is a homogeneous system of parameters, then R is a free module over $k[f_1, \ldots, f_n]$.

(d) There exists a homogeneous system of paramaters $f_1, \ldots, f_n$ such that R is a free module over $k[f_1, \ldots, f_n]$.

It is immediate from this definition that a polynomial ring is Cohen-Macaulay. The property applies to invariant rings under nice group actions, thanks to the following theorems:

Theorem II.10 (Hochster, Eagon)**.** *[23] If G is a finite, linearly reductive group over k and V a finite dimensional, rational representation, then $k[V]^G$ is Cohen-Macaulay.*

A finite group is linearly reductive if and only if its order is relatively prime to the characteristic of k, with a Reynold's operator that averages functions over translation by the group. More generally,

Theorem II.11 (Hochster, Roberts)**.** *[24] If G is a linearly reductive group over k and V a finite dimensional, rational representation, then $k[V]^G$ is Cohen-Macaulay.*

The next section considers a computational application of this property.

2.6 Hilbert Functions

For an invariant ring $k[V]^G$ of a linearly reductive group, let $S = k[f_1, \ldots, f_n]$ be the subalgebra generated by a homogeneous system of parameters. Then S is isomorphic to a polynomial ring, and there exist homogeneous $g_1, \ldots, g_\ell$ such that

$$k[V]^G = S \cdot g_1 \oplus \cdots \oplus S \cdot g_\ell$$

as a free S module. Let $d_i = \deg(f_i)$ and $e_j = \deg(g_j)$. Then the Hilbert series for $k[V]^G$ can be written as a rational function:

$$\frac{\sum_j t^{e_j}}{\prod_i (1 - t^{d_i})}.$$

Note that the decomposition of $k[V]^G$ and the rational expression above are not unique.

For a graded k-algebra R generated in degree one, the Hilbert polynomial describes $\dim_k R_i$ for large enough i. Because invariant rings are generated in several degrees, more analysis is required to recover the growth rate of the Hilbert function. In the notation of Campbell, et al. [3], let $R[\ell, i]$ denote the elements of R of degree congruent to i modulo ℓ. Note that each $R[\ell, i]$ is an $R[\ell, 0]$-submodule of R. Choose ℓ such that the Hilbert series for R can be written $a(t)/(1-t^\ell)^n$, where $a(t) \in \mathbb{Z}[t]$ and $n = \dim R$. Campbell, et al. show that ℓ can be chosen as the least common multiple of the degrees of a homogenoeus system of parameters for R (their Proposition 3.1). They conclude the following:

Proposition II.12. *Let R be a graded, finitely generated k-algebra with $R_0 = k$. If R is an integral domain of dimension n, then the Hilbert polynomials of the non-trivial modules $R[\ell, i]$ have the same leading coefficient and the same degree $n - 1$, for $i = 0, \ldots, \ell - 1$.*

From the hypotheses of the proposition, it follows $\dim_k R_d = O(d^{n-1})$. Hence intuition about Hilbert polynomials applies to invariant rings generated in several degrees.

2.7 Computation and Degree Bounds

2.7.1 Invariants of Reductive Groups

Several algorithms exist to compute generators for an invariant ring. These algorithms require the acting group to be finite, linearly reductive, or reductive. Sturmfels in 1993 filled in the details of Hilbert's procedure for $G = GL_n$ of a century earlier [20, 44]. The following denotes the vanishing sets, in a variety X, of polynomials or an ideal I:

$$\mathbb{V}(f_1, \ldots, f_n) = \{x \in X \mid f_i(x) = 0 \ \forall \ i\}$$

$$\mathbb{V}(I) = \{x \in X \mid f(x) = 0 \ \forall \ f \in I\}$$

Hilbert observed this property of the nullcone:

Theorem II.13. *[20] For a finite-dimensional representation V of $G = GL_n$ over $\mathbb{C}$, let $f_1, \ldots, f_n$ be homogeneous invariant polynomials such that $\mathbb{V}(f_1, \ldots, f_n) = \mathcal{N}_V$. Then $\mathbb{C}[V]^G$ is a finitely generated $\mathbb{C}[f_1, \ldots, f_n]$-module.*

The same statement holds for any reductive G [8, p. 226], and in either case, the integral closure of $k[f_1, \ldots, f_n]$ is $k[V]^G$. Now, for a maximal torus $T \subset G$, the Hilbert-Mumford criterion yields that $\mathcal{N}_{V,G} = G \cdot \mathcal{N}_{V,T}$. Using this relation, Hilbert computes generators for $k[V]^T$ and from these the f_i defining the null cone. Sturmfels then computes the integral closure of $k[f_1, \ldots, f_n]$ with multiple Gröbner basis calculations.

Derksen's 1999 algorithm computes generators of $k[V]^G$ when G is linearly reductive [4], using only one Gröbner basis. Let $\psi : G \times V \to V \times V$ be an explicit map

to the graph of the action, given by $(g, v) \mapsto (v, gv)$. The algorithm computes a Gröbner for the ideal $\mathfrak{b}$ vanishing on the closure of the image of ψ. Let $I \subset k[V]$ be the ideal generated by a homogeneous generating set for $k[V]^G$. Then the projection of $\mathfrak{b}$ to $k[V]$ gives generators $f_1, \ldots, f_n$ for I. Applying the Reynolds operator to each of the f_i gives generators for $k[V]^G$. Note this algorithm requires a subroutine to compute the Reynolds operator. Recall that for finite groups with $|G| \nmid \mathrm{char}(k)$ if $\mathrm{char}(k) > 0$, the Reynolds operator is just averaging. Classical approaches for GL_n and SL_n employ differential operators on $k[G]$; for general semisimple groups, that is, when the radical $R(G)$ is trivial, one employs an operator in the dual to the Lie algebra of G [6, 4.5].

Derksen's algorithm, though elegant, cannot produce generators when G is reductive in positive characteristic. Kemper's 2003 algorithm [33] works for any reductive group. He and Derksen combined and optimized their approaches in 2009 [7], providing algorithms that compute $k[X]^G$ in the case where reductive G acts on any affine variety X and the case where G is unipotent (hence non-reductive) and connected and X is irreducible. As an example, for Kemper's original procedure in positive characteristic, let $A \subseteq k[V]$ be a subalgebra, and define

$$\widehat{A} = \{f \in k[V] \mid f^{p^r} \in A \text{ for some } r \in \mathbb{Z}^+\}$$

to be the *purely inseparable closure* of A in $k[V]$. When k has characteristic zero, define $\widehat{A} = A$. Recall that for the action of any linear algebraic group G on a affine variety X, there exists a finitely generated, separating subalgebra $A \in k[X]^G$. Separating and generating invariants enjoy the following relationship:

Theorem II.14. *Let V be a rational representation of a reductive group G, let $A \subseteq k[V]^G$ be a finitely generated, separating subalgebra, and let $\widehat{\widehat{A}}$ be the purely*

inseparable closure of its normalization. Then

$$\widehat{\widetilde{A}} = k[V]^G.$$

A technical lemma in the proof requires G to be reductive, and this lemma fails if the only assumption is that $k[V]^G$ is finitely generated.

With this fact, Kemper's algorithm takes as input an embedding of the reductive group G and an explicit polynomial morphism $G \to GL(V)$. With two Gröbner basis calculations, it computes generators for the separating subalgebra A. Kemper then provides a known algorithm for computing the normalization $\widetilde{A}$ and a new algorithm for the inseparable closure $\widehat{\widetilde{A}}$. These procedures require additional Gröbner bases calculations. They output generators for $k[V]^G$ in any characteristic. On the other hand, Kemper provides an example where $k[V]^G$ is finitely generated, but G is not reductive, so $k[V]^G$ is not integral over a separating subalgebra. In such cases, his 2003 algorithm fails.

2.7.2 Degree Bounds

The following notion of *degree bound* assists in computation of generators for a graded ring $R = \oplus_{d \geq 0} R_d$:

$$\beta(R) := \min\{D \mid R \text{ is generated by } \oplus_{d=0}^{D} R_d\}.$$

If one knows $\beta(k[V]^G)$ or an upper bound for it, then one knows in which degrees an algorithm should search for generating invariants. Starting classically, let V_d be the k-space of degree d polynomials, that is $k[x,y]_d$ or $S(V_1)^d$. This V_d is called the space of *binary forms* of degree d. Let $SL_2(k)$ act on V_d, for any d, as follows:

$$\begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \cdot f(x,y) = f(\alpha x + \gamma y, \beta x + \delta y).$$

In fact, every irreducible representation of SL_2 is isomorphic to some V_d. Indeed, let $T \cong \{t \mid t \in k^*\}$ be the diagonal torus of the Borel subgroup of upper-triangular matrices in SL_2, and let W be a representation of highest weight $\lambda(t) = t^d$, d an positive integer. Then V_d has the same highest weight as W, namely, the weight of $x^d \in V_d$. So V_d and W must be isomorphic representations. Camille Jordan proved the following [28, 29]:

Theorem II.15. *Fix $d \geq 2$. Then the ring of covariants of the representation $W = \oplus_{d_i \leq d} V_{d_i}$, where each $d_i \leq d$, is generated in degree less than d^6.*

In particular, $\beta(k[V_d]^G) \leq d^6$, a smaller degree bound than for any other semisimple group. Kraft and Weyman provide a modern formulation of Jordan's method in [36].

In 1916, Noether produced the following degree bound for finite groups such that the characteristic of k is 0 or larger than $|G|$ [42]. Fleischmann and Fogarty improved her result to this modern formulation,

Theorem II.16. *[11, 12, 42] Let V be any representation of a finite group G. If the characteristic of k does not divide $|G|$, then $\beta(k[V]^G) \leq |G|$.*

The characteristic of k causes problems when the Reynolds operator averages polynomials over their G-translations.

Recall that Hilbert's work lead him to consider defining equations $f_1, \ldots, f_n$ for the null cone $\mathcal{N}_V$, the set of points whose orbit closure contains 0 in a representation V. If the f_i are algebraically independent (e.g. upon application of Noether's Normalization Lemma), then they form a homogeneous system of parameters for $k[V]^G$. Let $\sigma(V)$ denote the largest degree necessary for defining equations for N_V. Vladimir Popov employed the Hochster-Roberts theorem to show that if G is semisimple,

$$\beta(k[V]^G) \leq \dim V \cdot \operatorname{lcm}\{1, 2, \ldots, \sigma(V)\},$$

and he provided an intricate bound for $\sigma(V)$ [8, p. 229]. For example, for the binary forms his result obtains $\sigma(k[V_d]^G) = \frac{2^7}{3^2}d(d+1)^6$. For other degree bounds, see Hiss [21] and Kazarnovskii [31].

In 2001, Derksen [5] improved Popov's degree bound:

Theorem II.17. *Let G be a linearly reductive group and V a rational representation of dimension s. Then*

$$\beta(k[V]^G) \leq \max\{2, \frac{3}{8}s(\sigma(V))^2\}.$$

Derksen also produced a more explicit upper bound for $\sigma(V)$ by considering the degree of the generic orbit as a variety. This method is emulated by the work on separating orbits, below. Choose an embedding of the group G, so that its coordinate ring is $k[z_1, \ldots, z_\ell]/(h_1, \ldots, h_r)$ for polynomials h_i. Let $m = \dim(G)$, let $M = \max\{\deg(h_i) \mid \forall i\}$, and let N denote the maximum degree of the polynomials appearing in the representation $\rho\colon G \to GL(V)$.

Proposition II.18. *In the notation above, if the kernel of ρ is finite, then*

$$\sigma(V) \leq M^{\ell-m}N^m.$$

Note that this bound is polynomial in M and N, and one would choose ℓ as small as possible.

2.7.3 Invariants of Tori

The new work below considers torus invariants in particular. David Wehlau in 1993 produced degree bounds for the case that $G = T$ is a torus of rank r [46]. Let V be a representation of T, with dimension n and weights $w_1, \ldots, w_m$. The *character group* $X(T)$ of T is the set of algebraic homomorphism $T \to k^*$, and it is isomorphic

to $\mathbb{Z}^r$. Furthermore, $X(T) \otimes_{\mathbb{Z}} \mathbb{R} \cong \mathbb{R}^r$. As the weights hence correspond to points in $\mathbb{R}^r$, it makes sense to talk about the volume of the convex hull $\mathcal{C}_V$ of the weights w_i.

For a monomial $x_1^{a_1} \cdots x_n^{a_n} \in k[x_1, \ldots, x_n]$ the vector $(a_1, \ldots, a_n)$ gives a point in $\mathbb{N}^n$. Hence a monomial is invariant if $\sum_i a_i w_i = 0$. Wehlau finds a degree bound [8, p. 231]:

Theorem II.19. *In the above notation,*

$$\beta(k[V]^T) \leq (n-r)(r!)\mathrm{vol}(\mathcal{C}_V).$$

Sturmfels in [44, p. 19] explains an algorithm to compute $k[V]^T$ using a Gröbner basis, but Derksen and Kemper give an algorithm to compute generators for $k[V]^T$ without computing any Gröbner bases [6, p. 159]. If T has rank r, the algorithm considers points in a sufficiently large, finite set $\mathcal{C} \subset \mathbb{Z}^r$ containing the weights of the representation. For each $w \in \mathcal{C}$, let I_w be the ideal generated by monomials of weight w. The algorithm begins with the coordinate functions on V and, for each $w \in \mathcal{C}$, multiplies and tests monomials to produce a minimal set of (monomial) generators for each I_w. In particular, the generators of I_0 will generate $k[V]^T$.

When the rank of T is 1, the algorithm suggests a better degree bound for the generators. In this case, the set $\mathcal{C}$ is the convex hull of the weights of the representation V. Let $m = x_{i_1} \cdots x_{i_d}$ be a generating invariant, and let $m_j = x_{i_1} \cdots x_{i_j}$. The algorithm constructs m by building on $m_1, m_2, \ldots$ in such a way that the m_i all have distinct weights in $\mathcal{C}$. Therefore, the maximal degree of a generator is $1 + 2 \cdot \max\{|w| : w \text{ a weight of } V\}$.

2.8 Counting Invariants

2.8.1 SL_2 Invariants

Victor Kac employs the "Luna Slice Method" to reduce questions of one representation to a "better" representation [30]; in the article he works over $k = \mathbb{C}$. Let V be a representation of a reductive algebra group G. For a point $p \in V$, let G_p denote its stabilizer in G, and let T_p denote the tangent space at p to the orbit $G \cdot p$. Consider T_p as a linear, G_p-stable subspace of V, and find a G-stable complement S_P with $V = S_p \oplus T_p$. If $G \cdot p$ is closed, then G/G_p is affine, by the Matsushima criterion [37], whence G_p is reductive. The action of G_p on S_p is a *slice representation*, and there exists a categorical quotient $S_p \to S_p /\!\!/ G_p$. The categorical quotient $\pi \colon V \to V /\!\!/ G$ restricts to $S_p \to V /\!\!/ G$. Since the fibers of this map are G_p-stable, one has a morphism $\pi_p \colon S_p /\!\!/ G_p \to V /\!\!/ G$. Kac proves the following:

Proposition II.20. *In the notation above, the size of a minimal generating set for $k[V]^G$ is at least as large as the size of a minimal generating set for $k[S_p]^{G_p}$.*

To bound from below the size of a minimal generating set for $k[V]^G$, one can study a simpler slice representation S_p by choosing p wisely.

Kac considers the action of $G = SL_2(\mathbb{C})$ on the binary forms, chooses a slice representation such that G_p is a finite, cyclic group, then gives a combinatorial lower bound for the size of minimal generating sets for invariants. The combinatorial characterization plays a central role in the new upper bounds for torus invariants below.

To sketch Kac's argument, first assume d is odd. Recall V_d is spanned by the

degree-d monomials in two variables. Choose $p = x^{d-1}y + xy^{d-1}$. Then

$$G_p = \left\{ \begin{pmatrix} \zeta^k & 0 \\ 0 & \zeta^{-k} \end{pmatrix} \mid \zeta = e^{2\pi/(d-2)},\ k = 1, \ldots, d-2 \right\}.$$

Kac shows $G \cdot p$ is closed; a similar argument appears in Chapter 3 below. So Kac considers the slice representation S_p, and he shows that a generator σ of the cyclic group G_p acts on a basis of S_p with eigenvalues $1, \zeta, \zeta^2, \ldots, \zeta^{d-3}$.

Let $k[S_p] = k[x_0, x_1, \ldots, x_{d-3}]$, so that $\sigma \cdot x_i = \zeta_i x_i$. Then a minimal generating set of $k[S_p]^{G_p}$ consists of monomials $m = x_0^{a_0} \cdots x_{d-3}^{a_{d-3}}$ such that

$$\sum_{i=1}^{d-3} a_i \cdot i \equiv 0 \mod d-2$$

and such that m is not divisible by another different invariant. Kac notes first that every partition of n provides the degrees a_i for such a monomial. Then he sees that x_i^{d-2} for every i relatively prime to $d-2$ is also invariant. So writing $p(k)$ =(number of partitions of k) and $\phi(k)$ =(number of numbers $1, \ldots, k-1$ relatively prime to k), Kac counts at least $p(d-2) + \phi(d-2) - 1$ generators for $k[S_p]^{G_p}$. The -1 appears because otherwise x_1^{d-2} is counted twice. Therefore, when d is odd, this number also provides a lower bound for the size of a minimal generating set for $k[V_d]^{SL_2}$. The analysis proceeds similarly for even d. Now, Hardy and Ramanujan [16] (and independently, Uspensky [45]) found the asymptotic growth rate

$$p(n) \sim \frac{1}{4\sqrt{3}n} e^{\pi\sqrt{2n/3}},$$

so these generating sets exhibit non-polynomial but sub-exponential growth.

In 1988, Roger Howe produced a more explicit estimates of the size of "fundamental generating sets" for $R_d = k[V_d]^{SL_2}$ [25]. Let $\mathfrak{m}$ denote the maximal homogeneous ideal of R_d. By the Graded Nakayama Lemma, a set $S \subset R_d$ generates R_d if and

only if the image of S in $R_d/\mathfrak{m}$ generates $R_d/\mathfrak{m}$. If S has minimal size, Howe calls S a fundamental generating set. He proves the following:

Theorem II.21. *Let Γ_d denote the number of fundamental invariants of R_d, and let $R_d(k)$ denote the degree-k piece of R_d. Then for fixed degree k,*

(a) $\Gamma_d/(\dim R_d(k)) \to 1$ *as* $d \to \infty$.

(b) For constants c_k, the number of fundamental generators in degree k for $k \geq 4$ is asymptotically

$$\begin{cases} \frac{1}{2}(k!)^{-1}c_k n^{k-3}/(k-3) & nk \text{ even,} \\ 0 & nk \text{ odd.} \end{cases}$$

The paper [25] includes formulas for the c_k, in terms of binomial coefficients. By "asymptotically" Howe means "the difference between the two expressions is small in comparison with either," when n is large enough. He concludes "that almost all invariants of a fixed degree are eventually fundamental."

2.8.2 Torus and Cyclic Group Invariants

More generally, the positive integer vector solutions to an equation

$$\sum_{i=1}^{n-1} a_i \cdot i \equiv 0 \mod n$$

relate to torus invariants in the following way:

Proposition II.22. *Let $T = k^*$ act on x_i with weight t^i. Identify $\mathbb{Z}_n$ with the nth roots of unity in T. The evaluation homomorphism*

$$\text{ev}\colon f(x_1, \ldots, x_n, x_{-n}) \mapsto f(x_1, \ldots, x_n, 1)$$

provides a $\mathbb{Z}_n$-equivariant isomorphism

$$k[x_1, \ldots, x_n, x_{-n}]^T \to k[x_1, \ldots, x_n]^{\mathbb{Z}_n}.$$

Proof. Surjectivity is clear. For injectivity, it suffices to show that the ideal

$$(x_{-n} - 1) \subset k[x_1, \ldots, x_n, x_{n-1}].$$

contains no T-invariants. Recall every T-invariant is a sum of invariant monomials. If $f \in (x_{-n} - 1)$ is a T-invariant, then half of the monomials of f have nonzero weight, which is absurd. □

John C. Harris and David Wehlau [17] consider the general problem of producing all solutions $A = (a_1, \ldots, a_r) \in \mathbb{N}^r$ to an equation

$$w_1 x_1 + w_2 x_2 + \cdots + w_r x_r = 0 \mod n,$$

where the w_i are integers. They note that finding solutions to this equation is equivalent to finding solutions to Kac's equation,

$$x_1 + 2x_2 + \cdots + (n-1)x_{n-1} = 0 \mod n,$$

and they point out that the set of solutions forms a monoid. To state their result, they call a solution *decomposable* if it can be written as a sum of two non-trivial solutions, and indecomposable otherwise. There are only finitely many indecomposable solutions: if, say, $a_i \geq n$, then one may subtract off the extremal solution $(0, \ldots, n, \ldots, 0)$ that is non-zero in the ith place.

The degree of a solution A is $\deg(A) = \sum a_i$. The indecomposable solutions $A = (a_1, \ldots, a_{n-1})$ (and the variable x_n) correspond to generators $x_1^{a_1} \cdots x_{n-1}^{n-1}$ for $k[x_1, \ldots, x_n]^{\mathbb{Z}_n}$ in the proposition above. Lastly, define the *multiplicity* of a solution A to be

$$m(A) = \frac{a_1 + 2a_2 + \cdots + (n-1)a_{n-1}}{n}.$$

The solutions with multiplicity one correspond to the partitions of n, and these solutions are all indecomposable. Given such a solution, one can produce other indecomposable solutions, of possibly higher multiplicity, with the following permutation action. Let $H_n = \mathbb{Z}_n^*$ be the group of units in the ring $\mathbb{Z}_n$. Then $A = (a_1, \ldots, a_{n-1})$ is a solution if and only if $hA = (ha_1, \ldots, ha_{n-1})$ is a solution. Note that A and hA will have the same degree. In fact, H_n is the full group of automorphisms of the monoid of solutions, but not every solution is in the orbit of a solution with multiplicity one. So define the *level* of a solution A to be

$$\ell(A) = \min\{m(hA) \mid h \in H_n\}.$$

Harris and Wehlau first prove the following.

Proposition II.23. *Let A be a solution of multiplicity one and degree $k \geq \lceil n/2 \rceil + 1$. Then,*

(a) The H_n-orbit of A contains no other solution of multiplicity one.

(b) H_n acts faithfully on the orbit of A, whence the orbit has size $\phi(n)$.

In particular, they conclude that if $k \geq \lceil n/2 \rceil + 1$, then there are exactly $p(n-k)\phi(n)$ solutions in degree k. Note that $p(n-k)$ is the number of partitions of n into k parts. This count provides a lower bound for the number of indecomposable solutions to $\sum_i a_i x_i = 0 \mod n$. What is more, computing the H_n action on partitions of n provides an efficient algorithm for computing solutions in high degree. Wehlau and Harris further characterize these solutions as below:

Theorem II.24. *The following conjectures, due to A. Elashvili, are equivalent:*

(a) If A has degree $\geq \lfloor n/2 \rfloor + 2$, then $\ell(A) = 1$.

(b) If $k \geq \lfloor n/2 \rfloor + 2$, then there are exactly $p(n-k)\phi(n)$ solutions in degree k.

These conjectures differ subtly from the proposition above. The proposition describes the orbits of solutions with level one in high degree. The first conjecture claims that in fact all solutions in high degree have level one, in which case the number of indecomposable solutions in high degree would be known. Note that if n is odd, then the degree requirements are the same throughout. If n is even, then the conjectures require degree one higher than the proposition.

2.9 Algebraic Complexity

2.9.1 Complexity of Algorithms and Problems

The goal of computational invariant theory is to write algorithms to solve problems in invariant theory, such as distinguishing orbits or computing generators of invariant rings. To describe the *complexity* of a computation is to describe the number of steps or amount of computer memory space necessary to complete the computation. One can gain information about complexity indirectly, for example, by determining the number of cases an algorithm must consider or determining the minimum size of an output. When implementing an algorithm on a computer, complexity considerations have implications for the amount of memory the algorithm uses or the time it takes to run.

Each of these parameters (number of steps, number of cases to consider, size of output) depends on the size of the input to the algorithm. For example, an algorithm $\Gamma(G, V)$ to compute generating invariants may accept as input any reductive algebraic group G and any of its representations V. The number of steps Γ requires to run could depend on the dimensions of G and V, among other parameters. Indeed, the word "algorithm" is often shorthand for "family of algorithms" that accept inputs of different sizes and properties.

Thus to describe the complexity of an algorithm, one specifies which inputs con-

tribute fixed costs of run time and memory space, and which inputs are allowed to vary in the family. One says, for example, that Gaussian elimination can compute the reduced row echelon form of an $n \times n$ matrix over $\mathbb{Q}$ with $O(n^3)$ algebraic operations like $+, -$, and $\times$. Note that such a complexity estimate assumes all computations in $\mathbb{Q}$ require the same amount of processing time. Here, the "big O" notation $O(f(x))$ describes the order of growth for the function $f(x)$. One writes $g(x) = O(f(x))$ if there exists $x_0 > 0$ and a constant $c \geq 0$ such that $g(x) < cf(x)$ for all $x \geq x_0$, that is, "for sufficiently large x." For example, $5e^n + 4n^{99} + 3\log n = O(e^n)$.

One can determine the complexity of a problem with a two-part process. First, one describes the size of the output or the number of times some particular calculation must be made by any algorithm. This analytic work produces a lower bound, say $O(f)$. Then, one writes an algorithm that solves the problem, aiming for complexity similar to $O(f)$. The existence of such an algorithm provides an upper bound for the complexity of the problem. The lower and upper bounds then suggest the complexity of the problem itself.

2.9.2 Straight Line Programs

One framework that defines complexity more formally is that of straight line programs [2]. With notation inspired by applications to algebraic geometry, let V be a set, F a field, and let R be an F-subalgebra of the F-valued functions on V. Let $A = (a_{-m}, \ldots, a_{-1}) \in R^m$ be a finite, ordered subset of R. Consider a tape of cells with $a_i \in A$ in position i. A *straight line program* Γ is a finite, ordered list of instructions $\Gamma = (\Gamma_0, \ldots, \Gamma_{\ell-1})$. Each instruction Γ_i is of the form $(\star; j, k)$ or $(\star; j)$, where $\star$ is an operation and j, k are positive integers referring to tape entries in positions $i - j$ and $i - k$, that is, j and k cells before i, respectively. The *length* $\ell = |\Gamma|$ measures the complexity of the computation.

To execute Γ on input A, for $i = 0, \ldots, \ell - 1$ write a_i in tape position i as follows:

$$a_i = \begin{cases} a_{i-j} + a_{i-k} & \text{if } \Gamma_i = (+; j, k) \\ a_{i-j} - a_{i-k} & \text{if } \Gamma_i = (-; j, k) \\ a_{i-j} \cdot a_{i-k} & \text{if } \Gamma_i = (\times; j, k) \qquad \text{where } j, k < i. \\ c & \text{if } \Gamma_i = (\text{const}; c) \text{ for } c \in F \\ a_{i-j} & \text{if } \Gamma_i = (\text{recall}; j) \end{cases}$$

The "recall" instruction of position j serves to collect relevant computations at the end of the tape. Define the *order-d output* of Γ by $\text{Out}_d(\Gamma, A) = (a_{\ell-d}, \ldots, a_{\ell-1}) \in R^d$, where $\ell = |\Gamma|$. We omit the d where convenient. A straight line program hence defines a function $R^m \to R^d$.

For example, the function $f(x, y) = x^2 + 2xy + y^2$ in $R = \mathbb{Q}[x, y]$ can be computed with the following naive straight line program. The input is $(a_{-2}, a_{-1}) = (x, y)$. Here are the instructions:

- $\Gamma_0 = (\times, 2, 2)$
- $\Gamma_1 = (\times, 2, 2)$
- $\Gamma_2 = (\times, 4, 3)$
- $\Gamma_3 = (\text{const}, 2)$
- $\Gamma_4 = (\times, 2, 1)$
- $\Gamma_5 = (+, 5, 4)$
- $\Gamma_6 = (+, 2, 1)$

Note that the numbers in each instruction describe locations relative to the current location on the tape, so some instructions are identical. This algorithm has length

7, and its entire output is

$$(x^2, y^2, xy, 2xy, x^2 + y^2, x^2 + y^2 + 2xy).$$

Of course, the complexity of the computation of f is 2: on the same input, let $\Gamma_0 = (+, x, y)$ and $\Gamma_1 = (*, 1, 1)$.

Write $\Gamma^{(2)} \circ \Gamma^{(1)}$ for the composition of two straight line programs, in which the input of $\Gamma^{(2)}$ is $\mathrm{Out}_d(\Gamma^{(1)}, A)$ for some d depending on $\Gamma^{(2)}$. Then $\Gamma^{(2)} \circ \Gamma^{(1)}$ has input A, and we execute $\Gamma^{(2)} \circ \Gamma^{(1)}$ by concatenating the instruction lists.

Since the multiplication and division of numbers requires more memory and processor time than addition, subtraction, and the calling of constants, one can choose only to consider multiplications when determining lower bounds for the length of an algorithm. On the other hand, the convention here of counting all operations yields stronger complexity results and upper bounds. Now, programs cease to be "straight line" when they involve "branching" from IF-THEN clauses. For these programs, different inputs may require different run times and memory uses, because the algorithm performs different steps. For these algorithms, one may define the "branching complexity" as the total length of all branches of the tree of computations.

2.9.3 Examples

One hopes that the length or memory use of an algorithm is a polynomial function of the size of the input, so that the algorithm remains practical for larger and larger instances of the problem. If the length of an algorithm is polynomial in some relevant parameters of the input, one says the algorithm is *polynomial time*. Of course, a statement "algorithm Γ has complexity $O(f(n))$" ignores constants and constant coefficients in the true function $g(n)$ for the length of Γ. In an implementation of the algorithm on a computer, these constants could lead to prohibitive memory

requirements or run times for even small instances of the problem. So in fact the statement "algorithm Γ has complexity $O(f)$" describes how the complexity of the algorithm grows over larger inputs.

Several computations in algebra have polynomial or better complexity. For example, if $f, g \in k[x]$ have $\deg(fg) = d$ and k has sufficiently many roots of unity, then Fast Fourier Transform algorithms can compute $f \cdot g$ with total complexity bounded by $O(d \log d)$ [2, p. 33]. The Gaussian elimination algorithm to compute the reduced row echelon form of an $n \times n$ matrix has complexity $O(n^3)$, including operations like exchanging rows. In fact, computing the inverse, row echelon form, or determinant of an $n \times n$ matrix can be reduced to a sequence of matrix multiplications. The complexity of matrix multiplication then provides a total complexity bound for all of these computations, namely, as of 1987, $O(n^{2.38})$ [2, p. 420].

The complexity of Gröbner basis calculations, that is, the number of steps performed to compute a Gröbner basis, is unpredictable but believed to be quite large [10]. The essential process in computing a Gröbner basis is the normal form algorithm. Dubé et al. count the number of "reductions" required to write a polynomial f in a normal form with respect to some fixed basis $\mathcal{G}$ of polynomials: if L is the number of monomials in f, then the number of reductions is bounded above by $L \cdot O(1)^{\deg f}$. Furthermore, they prove the existence of $\mathcal{G}$ and f with $d > L$ such that the number of reductions is at least exponential in d. On the other hand, they note that many ideals are "highly structured." As result, the Buchberger algorithm is practical in many examples, especially in two variables, even though its complexity is theoretically exponential.

It must be said that the above discussion of computational complexity simplifies some aspects of problems while overstating others. For one, an algorithm whose

length is asymptotically polynomial may still in practice take a long time to terminate. Avner Ash, who tests and develops conjectures in algebra by computing a large number of examples with technology, puts it this way [1]: "When you say an algorithm is polynomial time, I want to know the constants." That is, he warns that an algorithm with length $10^9 n^6 + 10^6 = O(n^6)$ requires 4 billion steps when $n = 4$. Roger Howe points out that polynomial complexity of degree 6, for example, becomes time-consuming even for $n < 100$ [26]. On the other hand, Howe notes that many large objects to compute, even sets of generating invariants, in fact have simple descriptions. After all, the minimal generating sets for torus invariants have size at least $O(e^{\sqrt{d}})$, but one can describe them with linear integer equations.

CHAPTER III

Counting Generating Invariants of Semisimple Groups

The first chapter of new results considers the growth of minimal generating sets for invariant rings. For an algebraically closed field k, parameterize with integers $n \geq 0$ the family of representations $V_{n\lambda}$ with highest weight $n\lambda$. Let $S^d(V_{n\lambda})^G$ denote the degree-d invariant polynomials on $V_{n\lambda}$. We fix d and apply a ring structure to the collection of spaces $S^d(V_{n\lambda})$ for $n \geq 0$, graded now by n. It turns out that $\dim S^d(V_{n\lambda})^G$ grows like a polynomial in n whose degree is a linear function of d. Choosing high enough d, we show that the minimal cardinality of a generating set for $k[V_{n\lambda}]^G$ grows faster than any polynomial in n.

The same trick works when SL_2 acts on the space V_n of binary forms of degree n. Again, $\dim S^d(V_n)^{SL_2}$ grows as a polynomial in n with degree d as large as we want, and the minimal cardinality of a generating for $k[V_n]^{SL_2}$, the invariants on the binary forms, grows faster than any polynomial in n.

Counting generating invariants of $T \subset SL_2$ reduces to the problem of counting, for each $n \geq 1$, the $S \subseteq \{-n, -(n-1), \ldots, n\}$ such that $\sum_{a \in S} a = 0$ and no subset of S has this property (the "subset sum problem"). Olson [43] proves that the size of such S is no more than $3\sqrt{n}$. In the context of monomials, this result provides a degree bound, and one can conclude that the size of a generating set for $k[V_n]^T$ is

$O(e^{\sqrt{n}\log n})$.

For motivation and intuition, we begin with the invariants of SL_2. The Background chapter discusses results of Kac [30] and Howe [25] on the size of generating sets for $k[V_n]^{SL_2}$. The below proof that these sets must grow faster than any polynomial mirrors the computations of Howe, but the equivalent result for an arbitrary semisimple group appears to be new.

3.1 The Orbits of SL_2 Acting on Binary Forms

Let k be an algebraically closed field, and assume for Sections 3.2 and 3.3 that $\mathrm{char}(k) = 0$. Consider the classical action of SL_2 on the binary forms V_d of degree d.

Lemma III.1. *Let $X = SL_2 \cdot f$ be the orbit of a form $f \in V_d$.*

1. *If f has a factor of multiplicity $\geq d/2$ and X is closed, then f has at most two distinct roots.*
2. *A form f has root factor of multiplicity $> d/2$ if and only if f lies in the null cone.*

Proof. Only forms of even degree have roots of multiplicity $d/2$. Assume without loss that $x^{d/2} \mid f$. Then

$$f = a_d x^d + a_{d-1} x^{d-1} y + \cdots + a_{d/2} x^{d/2} y^{d/2}.$$

Consider the orbit of f under the action of the diagonal torus. Then,

$$\lim_{t \to 0} t \cdot f = a_{d/2} x^{d/2} y^{d/2}.$$

Hence if X is closed, f has only two distinct roots.

A form f has a root of multiplicity $> d/2$ if and only if X contains, say,

$$f_0 = \begin{cases} a_d x^d + \cdots + a_{d/2+1} x^{d/2+1} y^{d/2-1} & d \text{ even} \\ a_d x^d + \cdots + a_{(d+1)/2} x^{(d+1)/2} y^{(d-1)/2} & d \text{ odd} \end{cases}.$$

Such forms comprise the null cone of T, hence lie in the null cone of SL_2, by the Hilbert-Mumford Criterion. Conversely, if f lies in the null cone of SL_2, then X contains an element in the null cone of T, hence of the above form. In particular, X is not closed. □

Lemma III.2. *If the degree $d \geq 3$, then the generic orbit is closed, of dimension 3.*

Proof. The non-vanishing of the discriminant gives a dense open set of forms with no double roots. Claim the orbit of such a form is closed. First consider the diagonal torus T in SL_2. The T-weight spaces of V_d are spanned by the monomials $x^i y^{d-i}$. Since f has only single roots and degree at least 3, it involves monomials of both positive and negative weight. Hence if $\gamma\colon k^* \to T$ is a 1-parameter subgroup, then

$$\lim_{t\to 0} \gamma(t) \cdot f \text{ does not exist.}$$

Now let $\gamma\colon k^* \to SL_2$ be any 1-parameter subgroup. Find $\sigma \in SL_2$ such that $\sigma\gamma\sigma^{-1}$ lies in the diagonal torus T in SL_2. Noting that $\sigma \cdot f$ also has all single roots,

$$\begin{aligned} \lim_{t\to 0} \gamma(t) \cdot f &= \lim_{t\to 0} \gamma(t)\sigma^{-1} \cdot \sigma f \\ &= \sigma^{-1} \cdot \lim_{t\to 0} \sigma\gamma(t)\sigma^{-1} \cdot \sigma f \end{aligned}$$

which also does not exist. By the Hilbert-Mumford Criterion, the orbit of f is closed.

For $d \geq 3$, we may consider any three factors of f as a triple of points in $\mathbb{P}^1$. From the analysis of the complex plane, an element $\sigma \in SL_2$ is uniquely determined by its action on three distinct points, which it sends to a triple of distinct points. Thus the stabilizer of f is finite, and $\dim \overline{SL_2 \cdot f} = 3$. □

Lemma III.3. *For $d \geq 3$, the categorical quotient has $\dim V_d /\!\!/ SL_2 = d - 2$.*

Proof. Let $\pi\colon V_d \to V_d /\!\!/ SL_2$ be the categorical quotient, a surjection of irreducible varieties. Since the generic orbit is closed of dimension 3,

$$3 = \dim V_d - \dim V_d /\!\!/ SL_2 = d + 1 - \dim V_d /\!\!/ SL_2.$$

□

3.2 Bounding Generating Invariants for the Binary Forms

Let $V = V_1 = \{ax + by \mid a, b \in k\}$ be the binary forms of degree 1 over an algebraically closed field k. Then the space of binary forms of degree d is isomorphic to $S^d(V)$, and $S^e(S^d(V))$ is isomorphic to the space of degree-e regular functions on V_d. That is, $S^e(S^d(V)) = k[V_d]_e$.

Proposition III.4. *For $V = V_1$ and natural numbers d, e, $S^e(S^d(V)) \cong S^d(S^e(V))$.*

Proof. The linear factorization of $f \in V_d$ yields a surjective, SL_2-equivariant morphism of varieties $\pi\colon V^d \twoheadrightarrow V_d$. Let S_d act on V_d by permuting the factors, and let the torus $(k^*)^{d-1}$ act as follows:

$$(t_1, \ldots, t_{d-1}) \cdot (f_1, \ldots, f_d) = (t_1 f_1, t_1^{-1} t_2 f_2, \ldots, t_{d-2}^{-1} t_{d-1} f_{d-1}, t_{d-1}^{-1} f_d).$$

Then $\pi^{-1}(f)$ is a $T \rtimes S_d$-orbit, and we have an isomorphism

$$\pi^*\colon k[V_d] \xrightarrow{\sim} k[V^d]^{T \rtimes S_d} = \left[\overbrace{S(V) \otimes \cdots \otimes S(V)}^{d}\right]^{T \rtimes S_d}.$$

where $S(V)$ is the symmetric algebra. For details, please see [6, pg. 164].

For $(f_1, \ldots, f_d) \in V^d$, write $f_i = (a_i x + b_i y)$. If $c_0 x^d + c_1 x^d y + \cdots + c_d y^d \in V_d$, then for every j, $\pi^*(c_j) \in k[V^d]$ is homogeneous of degree d in the a_i, b_i. Since $V \cong V^*$,

the map π^* gives

$$\begin{aligned} S^e(S^d(V)) &\xrightarrow{\sim} \left[\bigoplus_{\sum e_i = de} S^{e_1}(V) \otimes \cdots \otimes S(V)^{e_d}\right]^{T \rtimes S_d} \\ &= \left[\overbrace{S^e(V) \oplus \cdots \oplus S^e(V)}^{d}\right]^{S_d} \\ &= S^d(S^e(V)). \end{aligned}$$

□

Lemma III.5. *Let R be a graded, Cohen-Macaulay domain of dimension n. If R has Hilbert-Poincaré series $\sum a(d)t^d$, then there is a constant $c > 0$ such that*

$$\limsup_d \left\{\frac{a(d)}{d^{n-1}}\right\} = c > 0,$$

and $a(d)/d^{n-1} \geq c$ for a sequence of integers d with constant difference.

Proof. Let ℓ be the least common multiple of the degrees of a set of generators for R. Let $R[\ell; i] = \oplus_m R_{m\ell+i}$, the ring of elements of degree congruent to i modulo ℓ. Then from Section 4 of [3], if R is Cohen-Macaulay, then each nontrivial $R[\ell; i]$ has Hilbert polynomial $H_i(m)$ of degree $n-1$. What is more, if R is a domain, then the leading coefficient c of each nontrivial $H_i(m)$ is equal to that of $H_0(m)$; the constant c is the degree of the $R[\ell; i]$. Thus there exist infinitely many d, with period at most ℓ, such that $a(d)/d^{n-1} = c + O(d^{-1})$, and the result follows. □

Since SL_2 acts linearly on V_n, one can find generating sets for $k[V_n]^{SL_2}$ such that each polynomial is homogenous. Call a subset Γ of a k-algebra R *minimal* of it has minimal cardinality among all generating sets. By the graded Nakayama lemma, every minimal, homogenous generating set has the same cardinality.

Proposition III.6. *As $n \to \infty$, the size of a minimal set of generators for $k[V_n]^{SL_2}$ grows faster than any polynomial in n.*

Proof. Suppose Γ_n is a generating set for $R_n := k[V_n]^{SL_2}$ of minimal cardinality. Let $\Gamma_n(d)$ be the number of generators in Γ_n of degree d. If $k[V_n]^{SL_2}$ has Poincaré series $\sum a_n(d)t^d$, then for large d,

$$\Gamma_n(d) \geq a_n(d) - \sum_{i=1}^{\lfloor d/2 \rfloor} a_n(i) \cdot a_n(d-i) = a_d(n) - \sum_{i=1}^{\lfloor d/2 \rfloor} a_i(n) \cdot a_{d-i}(n).$$

That is, we then subtract from $\dim R_d$ the number of products of elements from $R_{<d}$, assuming no relations among them. The result is a lower bound for the number of generators in degree d. We next recall $S^i(V_n) = S^n(V_i)$ to substitute $a_n(i) = a_i(n)$.

Consider the sum on the far right above. Now, $a_i(1) = a_1(i) = 0$ for all n. When $i = 2$, $a_i(2) = a_2(i)$ is 1 or 0, as $k[V_2]^{SL_2}$ is generated by the discriminant. Recall $\dim R_n = n - 2$. Thus by Lemma III.5, $\limsup_n \{a_2(n)a_{d-2}(n)/n^{d-5}\}$ is a constant (albeit a function of d). Thus $a_2(n)a_{d-2}(n) = O(n^{d-5})$. Similarly, for $3 \leq i \leq \lfloor d/2 \rfloor$, $a_i(n) \cdot a_{d-i}(n) = O(n^{i-3} \cdot n^{d-i-3}) = O(n^{d-6})$. Thus in the relation

$$\Gamma_n(d) \geq a_d(n) - \sum_{i=1}^{\lfloor d/2 \rfloor} a_i(n) \cdot a_{d-i}(n),$$

if $d \geq 6$, then the right-most term grows as $O(n^{d-5})$ By Lemma III.5, there exists $c > 0$ such that $a_d(n)/n^{d-3} \geq c$ for a sequence of integers n with constant difference. For n in this sequence, $a_d(n) = O(n^{d-3})$ for large n. Choosing d arbitrarily large forces $\Gamma_n(d)$ to grow faster than any polynomial in n. □

3.3 Counting Torus Invariants

3.3.1 Applying the Grosshans Principle

Counting torus invariants on binary forms may be easier than counting SL_2-invariants. To that end, consider

Proposition III.7 (Grosshans Principle)**.** *[14] Let an algebraic group G act rationally on V, and let $H \subseteq G$ be a subgroup. Let H and G act on $k[G]$ by right and left*

translation, respectively. Then

$$k[V]^H \cong \left(k[G]^H \otimes k[V]\right)^G.$$

Proof. Consider the morphism $G \times V \to G \times V$ by $(g, v) \mapsto (g, gv)$. Let $G \times H$ act on the source by $(g, h) \cdot (u, v) = (guh^{-1}, hv)$ and on the target by $(g, h) \cdot (u, v) = (guh^{-1}, gv)$. Indeed, the actions of G and H commute, and the map is a $(G \times H)$-equivariant isomorphism. On the left, G acts only on G, so

$$\left((k[G] \otimes k[V])^G\right)^H = k[V]^H.$$

On the right, H acts only on G, so

$$\left((k[G] \otimes k[V])^H\right)^G = \left(k[G]^H \otimes k[V]\right)^G.$$

By the equivariance of the map, we obtain two expressions for the invariants of $G \times H$. □

Let $T \subseteq SL_2$ be the diagonal torus, and

$$k[SL_2] \cong k[z_{11}, z_{12}, z_{21}, z_{22}]/(z_{11}z_{22} - z_{12}z_{21} - 1).$$

Then under the right-action of T on SL_2,

$$k[SL_2]^T \cong k[z_{11}z_{12}, z_{11}z_{22}, z_{12}z_{21}, z_{21}z_{22}]/(z_{11}z_{22} - z_{12}z_{21} - 1).$$

Lemma III.8. *There exists a SL_2-equivariant surjection*

$$k[V_2] \otimes k[V_d] \twoheadrightarrow k[SL_2]^T \otimes k[V_d].$$

Proof. Write $a_0x^2 + a_1xy + a_2y^2$. For the homomorphism $k[V_2] = k[a_0, a_1, a_2] \twoheadrightarrow k[SL_2]^T$, send

$$\begin{aligned} a_0 &\mapsto z_{11}z_{12} \\ a_1 &\mapsto z_{11}z_{22} + z_{12}z_{21} \equiv 2z_{11}z_{22} - 1 \equiv 1 + 2z_{12}z_{21} \ (z_{11}z_{22} - z_{12}z_{21} - 1) \\ a_2 &\mapsto z_{21}z_{22} \end{aligned}$$

Explicitly, if

$$\sigma^{-1} = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \in SL_2,$$

then

$$\begin{aligned} \sigma \cdot a_0 &= \alpha^2 a_0 + \alpha\beta a_1 + \beta^2 a_2 \\ \sigma \cdot a_1 &= 2\alpha\gamma a_0 + (\alpha\delta + \beta\gamma) a_1 + 2\beta\delta a_2 \\ \sigma \cdot a_2 &= \gamma^2 a_0 + \gamma\delta a_1 + \delta^2 a_2 \end{aligned}$$

and

$$\begin{aligned} \sigma \cdot (z_{11}z_{12}) &= \alpha^2 z_{11}z_{12} + \alpha\beta(z_{11}z_{22} + z_{12}z_{21}) + \beta^2 z_{21}z_{22} \\ \sigma \cdot (z_{11}z_{22}) &= \alpha\gamma z_{11}z_{12} + \alpha\delta z_{11}z_{22} + \beta\gamma z_{12}z_{21} + \beta\delta z_{21}z_{22} \\ \sigma \cdot (z_{12}z_{21}) &= \alpha\gamma z_{11}z_{12} + \beta\gamma z_{11}z_{22} + \alpha\delta z_{12}z_{21} + \beta\delta z_{21}z_{22} \\ \sigma \cdot (z_{21}z_{22}) &= \gamma^2 z_{11}z_{12} + \gamma\delta(z_{11}z_{22} + z_{12}z_{21}) + \delta^2 z_{21}z_{22}. \end{aligned}$$

Comparing the actions of SL_2 on the generators above, it follows that the defintion of $k[a_0, a_1, a_2] \twoheadrightarrow k[SL_2]^T$ is indeed surjective and SL_2-equivariant. □

Applying the Grosshans Principle,

$$k[V_2 \oplus V_d]^{SL_2} \twoheadrightarrow \left(k[SL_2]^T \otimes k[V_d]\right)^{SL_2} \cong k[V_d]^T.$$

Kac showed that number of generating invariants for $k[V_d]^{SL_2}$ is bounded below by the size of a minimal generating set for $k[V_d]^T$ [30]. The Grosshans Principle may help relate torus invariants and G invariants for representations of other groups G.

3.3.2 An Upper Bound for Torus invariants

Because the following invariant subrings are generated by monomials, we say an invariant monomial is *indecomposable* if it is not the product of non-constant

invariants. Consider the action of $T \cong k^*$ on a polynomial ring $k[x_1, \ldots, x_n, x_{-n}]$ defined by

$$t \cdot x_i = t^i x_i.$$

Then the invariants are monomials of weight zero, whose positive-weight part is a multiple of n.

Proposition III.9. *The cardinality of a minimal generating set for*

$$k[x_1, \ldots, x_n, x_{-n}]^T \quad \textit{is} \quad O\left(e^{6\sqrt{n}\log 2n}\right).$$

Proof. A minimal generating set contains only monomials of the form $x_n x_{-n}$ and $m x_{-n}^k$, where $m \in (x_1, \ldots, x_{n-1})$ is properly divisible by no monomial of weight congruent to 0 modulo n. Now, Olson shows that if $S \subset \mathbb{Z}_n$ has order at least $3\sqrt{n}$, then a subset of S has trivial sum [43]. Now, a minimal generating set of $k[x_1, \ldots, x_n, x_{-n}]^T$ can be chosen such that each (monomial) generator properly includes no invariant. So such a set can be chosen such that each generating monomial includes no more than $3\sqrt{n}$ distinct variables.

An algorithm of Derksen and Kemper to construct torus invariants implies a degree bound of $2n-1$ for a generating set of $k[x_1, \ldots, x_n, x_{-n}]^T$, by computing within the convex hull of the variables' weights [6, p. 159] . This linear bound may only hold when the torus has rank 1; Wehlau provides a more general bound in [46]. An upper bound for the number of generators in degree d is

$$\binom{n}{\lfloor 3\sqrt{n}\rfloor} \cdot \binom{\lfloor 3\sqrt{n}\rfloor + d - 1}{\lfloor 3\sqrt{n}\rfloor - 1} \leq n^{\lfloor 3\sqrt{n}\rfloor} \cdot (\lfloor 3\sqrt{n}\rfloor + d - 1)^{\lfloor 3\sqrt{n}\rfloor - 1}.$$

The first term on the left counts ways of choosing $3\sqrt{n}$ variables; the second term counts monomials of degree d with $3\sqrt{n}$ variables. Summing the upper bound over degrees d up to $2n-1$ yields

$$O\left((n^{3\sqrt{n}} \cdot (3\sqrt{n} + 2n - 2)^{3\sqrt{n}}\right) \leq O((2n)^{6\sqrt{n}}) = O(e^{6\sqrt{n}\log 2n}).$$

□

The same bound holds for invariants of the cyclic group $\mathbb{Z}_n \subset T$:

Corollary III.10. *Identify $\mathbb{Z}_n$ with the nth roots of unity in T. The evaluation homomorphism*

$$\text{ev}\colon f(x_1, \ldots, x_n, x_{-n}) \mapsto f(x_1, \ldots, x_n, 1)$$

provides a $\mathbb{Z}_n$-equivariant isomorphism

$$k[x_1, \ldots, x_n, x_{-n}]^T \to k[x_1, \ldots, x_n]^{\mathbb{Z}_n}.$$

In particular, the cardinality of a minimal generating set for $k[x_1, \ldots, x_n, x_{-n}]^{\mathbb{Z}_n}$ is $O(e^{6\sqrt{n}\log 2n})$

Proof. Surjectivity is clear. For injectivity, it suffices to show that the ideal

$$(x_{-n} - 1) \subset k[x_1, \ldots, x_n, x_{n-1}].$$

contains no T-invariants. Recall every T-invariant is a sum of invariant monomials. If $f \in (x_{-n} - 1)$ is a T-invariant, then half of the monomials of f have nonzero weight, which is absurd. □

Recall our initial interest in $k[V_d]^T$. If $d = 2n$ is even, then this ring is isomorphic to

$$B_n := k[x_{-n}, x_{-n+1}, \ldots, x_0, \ldots, x_n]^T.$$

Proposition III.11. *The cardinality of a minimal generating set for B_n is*

$$O\left(ne^{12\sqrt{n}\log 2n}\right).$$

Proof. For $r \geq 1$, let mx_{-r}^k be an invariant monomial such that

$$m \in k[x_{r-1}, \ldots, x_1, x_{-1}, \ldots, x_{-(r-1)}]$$

and m is properly divisible by no monomial of weight congruent to 0 modulo r. Since x_i and x_{-r+i} have the same weight modulo r, Olson's theorem yields that m involves no more than $2 \cdot 3\sqrt{r}$ distinct variables, which may occur with multiplicity. Otherwise, if m' divides m and m' involves $6\sqrt{r}$ distinct variables, then m' is divisible by a monomial of weight congruent to 0 modulo r.

Recall that monomials of the form mx_{-r}^k in a minimal generating set have degree at most $2n-1$, for any r. Since an upper bound for the number of generating mx_{-r}^k in degree d is

$$\binom{2r}{6\sqrt{r}} \cdot \binom{6\sqrt{r}+d-1}{6\sqrt{r}-1} \leq (2r)^{6\sqrt{r}} \cdot (6\sqrt{r}+d-1)^{6\sqrt{r}-1},$$

summing these upper bounds up to degree $2n-1$ yields

$$O\left((2r)^{6\sqrt{r}} \cdot (6\sqrt{r}+2n-2)^{6\sqrt{r}}\right) \leq O((2n)^{12\sqrt{r}}) = O(e^{12\sqrt{r}\log 2n}).$$

Repeat the argument for invariants of the form $x_r^k m$ with

$$m \in k[x_{r-1}, \ldots, x_1, x_{-1}, \ldots, x_{-(r-1)}].$$

In either case, the weight of m determines the exponent k. Note that the number of invariants $x_r x_{-r}$ grows linearly, and these monomials together generate the invariants. The result follows by choosing the largest $r = n$. □

In the invariant mx_r^k considered in the above proof, the monomial m may be divisible by invariants not involving x_r. Nevertheless, the upper bound holds for monomials involving r.

Corollary III.12. *The cardinality of a minimal generating set for $k[V_d]^T$ is*

(a) $O(de^{6\sqrt{d}\log 2d})$ for odd d,

(b) $O(de^{12\sqrt{d/2}\log d})$ for even d.

Proof. For odd d there is a T-equivariant isomorphism

$$k[V_d]^T \cong k[x_{-d}, x_{-d+2}, \ldots, x_{-1}, x_1, \ldots, x_d]^T,$$

where all variables have odd weight. Following the argument above, the invariant mx_r^k, say, involves no more than $3\sqrt{r}$ variables among the $r+1$ variables $x_{r-2}, \ldots, x_{-(r-2)}$ of distinct weight modulo r. Making these adjustments to the above calculations, but retaining the degree bound $2d-1$ (from the convex hull of the variables' weights), yields an upper bound for the size of a minimal generating set:

$$(r+1)^{3\sqrt{r}} \cdot (3\sqrt{r} + 2d - 2)^{3\sqrt{r}} \leq O((2d)^{6\sqrt{r}}) = O(e^{6\sqrt{r}\log 2d}).$$

The weight r varies from $1, 3, 5 \ldots, d$, and the result follows.

For even $d = 2n$, the isomorphism

$$k[V_d]^T \cong k[x_{-n}, x_{-n+2}, \ldots, x_0, \ldots, x_n]^T$$

makes way for the previous proposition. □

3.4 Generator Counts of Representations Parameterized by Weight

3.4.1 The Ring of Covariants

Let G be a reductive algebraic group over a field k of characteristic 0. Fix a Borel subgroup $B = T \ltimes U$, where T is a maximal torus and U is the maximal unipotent subgroup in B. Let V_λ be the representation of G of highest weight λ with respect to T, which is unique up to isomorphism. We will show that whne G is semisimple, the cardinality of a minimal generating set of $k[V_{n\lambda}]^G$, as a function of n, grows faster than any polynomial.

Recall that when U acts on reductive G on the right, then $k[G]^U = \oplus_{\lambda \geq 0} V_\lambda$ as graded rings, where the latter is the direct sum of the irreducible representations V_λ

whose highest weight λ is positive. Choose a positive weight λ, and consider the subring $R_\lambda := \oplus_{n\geq 0} V_{n\lambda}$.

Lemma III.13. *The ring R_λ is finitely generated, namely, if v_{λ^*} is the lowest weight vector of $(V_\lambda)^*$, then $R_\lambda \cong k[\overline{G \cdot v_{\lambda^*}}]$.*

Proof. Let $v_{\lambda^*} \in (V_\lambda)^*$ be a *lowest* weight vector, of weight $-\lambda$, of the dual space to V_λ. Claim $R_\lambda = k[\overline{G \cdot v_{\lambda^*}}]$. Let f be the image of $(v_{\lambda^*})^*$ in $k[\overline{G \cdot v_{\lambda^*}}]$. For $n \geq 0$, the function f^n has weight $n\lambda$ under T and generates a G-module isomorphic to $V_{n\lambda}$ in $k[\overline{G \cdot v_{\lambda^*}}]$, whence $R_\lambda \hookrightarrow k[\overline{G \cdot v_{\lambda^*}}]$.

To obtain the reverse inclusion, consider the orbit map $G \to \overline{G \cdot v_{\lambda^*}}$ defined by $g \mapsto g \cdot v_{\lambda^*}$. Because this map is dominant, it gives rise to a G-equivariant injection $k[\overline{G \cdot v_{\lambda^*}}] \hookrightarrow k[G]$. Consider the stabilizer in T of v_{λ^*},

$$T_{\lambda^*} = \{t \in T \mid t \cdot v_{\lambda^*} = v_{\lambda^*}\}.$$

If $h_\mu \in k[\overline{G \cdot v_{\lambda^*}}]$ is a highest weight vector of weight μ, claim $\mu(T_{\lambda^*}) = \{1\}$. First note that if $h_\mu(v_{\lambda^*}) = 0$, then $h_\mu(T \cdot v_{\lambda^*}) = T \cdot h_\mu(v_{\lambda^*}) = \{0\}$, because h_μ is a weight vector. Let U^- be the opposite unipotent subgroup to U with respect to T; then v_{λ^*} is U^- invariant, because U^- lowers the weights of T. It follows

$$\{0\} = h_\mu(v_{\lambda^*}) = h_\mu(T \cdot v_{\lambda^*}) = h_\mu(UTU^- \cdot v_{\lambda^*})$$

because h_μ is U-invariant as a highest weight vector. Since UTU^- is dense in G, h_μ would be identically zero on $G \cdot v_{\lambda^*}$, which is absurd. Thus $h_\mu(v_{\lambda^*}) \neq 0$, and for $t \in T_{\lambda^*}$,

$$\mu(t)h_\mu(v_{\lambda^*}) = t \cdot h_\mu(v_{\lambda^*}) = h_\mu(t^{-1} \cdot v_{\lambda^*}) = h_\mu(v_{\lambda^*}).$$

Therefore $\mu(T_{\lambda^*}) = \{1\}$ and $\mu = n\lambda$, so every irreducible G-submodule of $k[\overline{G \cdot v_{\lambda^*}}]$ is one of the $V_{n\lambda}$. □

For comparison to the size of a generating set, note that the dimension of the $V_{n\lambda}$ grows as a polynomial in n:

Lemma III.14. *Let r be the number of positive roots of a reductive group G. Then* $\dim V_{n\lambda} = O(n^r)$.

Proof. Let Φ be the set of roots of G, $\delta = \frac{1}{2}\sum_{\alpha \succ 0} \alpha$ a sum over the positive roots, and $(\cdot, \cdot)$ an inner product on the space spanned by Φ, preserved by the Weyl group of reflections. Then by Weyl's formula in [27, p. 139],

$$\dim V_{n\lambda} = \frac{\prod_{\alpha \succ 0}(n\lambda + \delta, \alpha)}{\prod_{\alpha \succ 0}(\delta, \alpha)}.$$

The number r of positive roots of G satisfies $2r + \dim T = \dim G$. □

3.4.2 Generic Closed Orbits in Cartesian Products

For a finite-dimensional vector space V over an algebraically closed field k, let $\rho\colon G \to GL(V)$ be a non-trivial, rational representation of the semisimple algebraic group G.

Lemma III.15. *If G is semisimple and $\rho : G \to GL(V)$ is a representation, then the image of ρ lies in $SL(V)$. If ρ is not trivial, then* $\dim \rho(G) \geq 2$.

Proof. The image $\rho(G)$ of G in $GL(V)$ is also semisimple, and $\rho(G) = [\rho(G), \rho(G)] \subseteq [GL(V), GL(V)] = SL(V)$. If ρ is non-trivial, then $\dim \rho(G) \geq 1$, but there are no connected, semisimple algebraic groups of dimension 1 [27, p. 131]. □

Lemma III.16. *For an n-dimensional vector space V, let $X = \mathbb{P}(V)$. Let $d \geq n+1$ and let $SL(V) = SL_n$ act diagonally on X^d. If $\widetilde{X^d}$ is the affine cone over X^d, then the generic orbit of SL_n acting on $\widetilde{X^d}$ is closed, of dimension* $\dim SL_n$.

Proof. Fixing a basis for V, let $f\colon V^d \to k$ be the product of the $(n \times n)$-minors of an $n \times d$ matrix. This f defines a function on X^d and also on the affine cone $\widetilde{X^d}$.

Choose $p \in \widetilde{X^d}$ with $f(p) \neq 0$, and let $[p]$ be its image in X^d. Then $[p]$ defines d points in $\mathbb{P}(V)$, no n of which lie in the same hyperplane. Thus the stabilizer of $[p]$ in SL_n is finite (namely, the scalar matrices of SL_n such that the product of the entires is 1). It follows that the stabilizer of p is finite, whence $\dim(SL_n \cdot p) = n^2 - 1$. This dimension holds for the orbit of generic p with $f(p) \neq 0$.

Note that f is an invariant function on $\widetilde{X^d}$, because the $SL(V)$ fixes determinants. So if $q \in \widetilde{X^d}$ lies in the orbit closure of p, then $f(q) = f(p)$. Thus $\dim(SL_n \cdot q) = n^2 - 1$ as well. Since orbits in the boundary of $SL_n \cdot p$ must have strictly smaller dimension, the point q must lie in the orbit of p. Therefore, the generic orbit is closed. $\square$

Recall that V_λ is a highest-weight representation of a semisimple group G, and $R = \oplus_{n \geq 0} V_{n\lambda}$. Let $Z = \operatorname{Proj} R$, and consider the sum of tensors over k,

$$C = \bigoplus_{n \geq 0} \overbrace{V_{n\lambda} \otimes \cdots \otimes V_{n\lambda}}^{d}.$$

The dth Cartesian product of Z is

$$Z^d = \operatorname{Proj} C = \overbrace{Z \times_k \cdots \times_k Z}^{d}.$$

As $G \subseteq \operatorname{Aut}(R)$, G acts rationally on Z, hence diagonally on Z^d, hence on the affine cone $\widetilde{Z^d}$.

Lemma III.17. *In the above notation, if $\rho : G \to GL(V_\lambda)$ is an irreducible representation of highest weight λ and $d > \dim V_\lambda$, then the generic orbit of G acting on $\widetilde{Z^d}$ is closed, of dimension $\dim \rho(G)$.*

Proof. Suppose $\dim V_\lambda = n$. Note Z is a subvariety of $X = \mathbb{P}(V_\lambda^*)$: indeed, if $k[V_\lambda] = k[x_1, \ldots, x_n]$, then there is a surjection $k[x_1, \ldots, x_n] \twoheadrightarrow \oplus_{m \geq 0} V_{m\lambda}$ by sending the x_i onto an (n-dimensional) basis for V_λ. Thus $\widetilde{Z^d}$ is a closed subvariety of $\widetilde{X^d}$, and $\widetilde{Z}$ spans V_λ^* because $k[\widetilde{Z}]$ contains V_λ.

As in the proof of Lemma III.16, let $f\colon V_\lambda^d \to k$ be the product of the $(n \times n)$-minors of an $n \times d$ matrix. Since $\widetilde{Z}$ is irreducible and spans V_λ, the generic $p \in \widetilde{Z^d} \subseteq \widetilde{X^d}$ has $f(p) \neq 0$. Then for generic $p \in \widetilde{Z^d}$, the orbit $SL(V_\lambda)\cdot p$ is closed in $\widetilde{X^d}$. Recall that the stabilizer $SL(V_\lambda)_p$ of p is a finite set of scalar matrices, whence normal. It follows that the orbit $SL(V_\lambda)\cdot p$ is isomorphic to $SL(V_\lambda)/SL(V_\lambda)_p$ as a variety, so it is an algebraic group. By Lemma III.15, $\rho(G)$ is a closed subgroup of $SL(V_\lambda)$ (see [27, p. 54]), and the $\rho(G)$ action on V_λ factors through the $SL(V_\lambda)$ action. Thus for generic $p \in \widetilde{Z^d}$, $G\cdot p$ is isomorphic to $\rho(G)/\rho(G)_p$, the homomorphic image of an algebraic group. Thus $G\cdot p$ is closed in $SL(V_\lambda)\cdot p$, with dimension $\dim \rho(G)$. Since $G\cdot p \subseteq \widetilde{Z^d}$, the result follows. □

3.4.3 Counting Generators

To count generating invariants for large n, we again need to understand the degree d component of $k[V_{n\lambda}]^G$.

Proposition III.18. *Let $\rho : G \to GL(V_\lambda)$ be a non-trivial, rational representation of highest weight λ. Write $S^d(V_{n\lambda})^G \cong k[V_{n\lambda}]^G_d$, the degree-$d$ homogeneous piece of $k[V_{n\lambda}]^G$. Let $m = \dim \rho(G)$. Then there is an integer c with $1 \leq c < m$ such that for large n,*

$$\dim S^d(V_{n\lambda})^G \leq O(n^{cd-1}),$$

and when both n and d are large,

$$\dim S^d(V_{n\lambda})^G = O(n^{cd-m}),$$

Proof. Letting the symmetric group S_d permute the d factors of each n-graded piece of

$$\bigoplus_{n\geq 0} \overbrace{V_{n\lambda} \otimes \cdots \otimes V_{n\lambda}}^{d},$$

take the categorical quotient,

$$Z^d /\!\!/ S_d = \operatorname{Proj}\left(\bigoplus_{n\geq 0} S^d(V_{n\lambda})\right).$$

Next take the quotient by the G action on each copy of X:

$$\left(Z^d /\!\!/ S_d\right) /\!\!/ G = \operatorname{Proj}\left(\bigoplus_{n\geq 0} S^d(V_{n\lambda})^G\right).$$

These actions of G and S_d commute. From Lemma III.17, if $d \geq \dim V_\lambda + 1$, then the generic orbit of G acting on the cone $\widetilde{Z^d}$ is closed and of dimension $\rho(G) = m$. Therefore, for large enough d,

$$\dim(Z^d /\!\!/ S_d) /\!\!/ G = d \cdot \dim Z - m.$$

Now, Lemma III.13 yields that

$$\dim Z + 1 = \dim R = \dim k[\overline{G \cdot v_{\lambda^*}}] \leq \dim \rho(G) = m.$$

Let $c = \dim Z$. Note $c \geq 1$, because for large d, $\widetilde{Z^d}$ contains an orbit of dimension $m > 1$. Thus the Hilbert polynomial for $\oplus_n S^d(V_{n\lambda})^G$ has degree $cd - m$ for large d, and degree bounded by $cd - 1$ otherwise. □

As above, let $k[V_{n\lambda}]^G$ have Hilbert-Poincaré series $\sum_{d=0}^{\infty} a_n(d)t^d$.

Theorem III.19. *Let $\rho : G \to GL(V_\lambda)$ be a non-trivial, rational representation of highest weight λ. The minimal cardinality of a generating set for $k[V_{n\lambda}]^G$ grows faster than any polynomial in n, and hence faster than any polynomial in* $\dim V_{n\lambda}$.

Proof. Let Γ_n denote the minimal cardinality of a generating set of $k[V_{n\lambda}]^G$, and let $N = \dim V_\lambda$. From the proof above, if $d > N$, then $\dim S^d(V_{n\lambda})^G = O(n^{cd-m})$ for

large n and a constant c with $1 \leq c < m$. Then for large n,

$$\begin{aligned}
\Gamma_n &\geq a_n(d) - \sum_{i=1}^{\lfloor d/2 \rfloor} a_n(i)a_n(d-i) \\
&= a_n(d) - \sum_{i=1}^{N} a_n(i)a_n(d-i) - \sum_{i=N+1}^{\lfloor d/2 \rfloor} a_n(i)a_n(d-i) \\
&\approx a_n(d) - \sum_{i=1}^{N} a_n(i)a_n(d-i) - \sum_{i=N+1}^{\lfloor d/2 \rfloor} n^{ci-m}n^{c(d-i)-m}
\end{aligned}$$

where the approximation symbol indicates an asymptotic estimate for sufficiently large n and d. When $1 \leq i \leq N$ and n is large, we bound $a_n(i) = O(n^{ci-1})$. Assume $d - N > N, m$, so that for $i \leq N$ we may bound $a_n(d-i) = O(n^{c(d-i)-m})$. Then for such $i \leq N$, we have $a_n(i)a_n(d-i) = O(n^{cd-m-1})$, and we obtain that for large n and d,

$$\begin{aligned}
O(\Gamma_n) &\geq a_n(d) - N \cdot n^{cd-m-1} - dn^{cd-2m} \\
&\approx n^{cd-m} - N \cdot n^{cd-m-1} - dn^{cd-2m} \\
&\approx n^{cd-m}
\end{aligned}$$

Fixing d arbitrarily large, it follows that the size of a minimal generating set for $k[V_{n\lambda}]^G$ grows faster than any polynomial in n. The final assertion of the theorem follows because, by Lemma III.14, $\dim V_{n\lambda}$ grows like a polynomial in n. $\square$

CHAPTER IV

Quasi-Regular Functions

Let V be a representation of an algebraic group G over a field $k = \overline{k}$. Then for $p \in V$, the orbit $G \cdot p$ is open in its closure, and its boundary is a union of orbits. If $G \cdot q$ is an orbit in the boundary of $G \cdot p$, then every $f \in k[V]^G$ has $f(p) = f(q)$. In general, polynomial invariants cannot separate p and q if $\overline{G \cdot p} \cap \overline{G \cdot q} \neq \varnothing$. Thus if a set of more general invariant functions $V \to k$ is to separate all the orbits of a group action, then the functions in S must be able to distinguish locally closed sets. This chapter develops a set of such "quasi-regular" functions on an affine scheme and explores their properties.

4.1 The Patch Topology

For any commutative ring R with 1, consider the spectrum $\operatorname{Spec} R$ as a set, and endow it with the *patch topology* whose basis is generated by "patches" of the form $\mathbb{V}(f), \mathbb{V}(g)^c$ for $f, g \in R$. The terminology comes from Mel Hochter's 1969 paper [22], in which he identifies properties that characterize spaces in the image of the Spec functor. Note that this topology has more open sets than the Zariski topology. Let $\operatorname{QSpec} R$ denote the space $\operatorname{Spec} R$ with the patch topology. Writing $X = \operatorname{QSpec} R$, for $f \in R$ let $X_f = \mathbb{V}(f)^c$. Finite intersections of basis elements take the form

$$(f_1, \ldots, f_r) \cap X_{g_1 \cdots g_s}.$$

Thus a general element of the topology, constructed with infinite unions and finite intersections, can be written

$$\bigcup_{i\in S}\left(\mathbb{V}(f_{i,1},\ldots,f_{i,r_i})\cap X_{g_{i,1}\cdots g_{i,t_i}}\right)$$

for some index set S One can further write $g_{i,1}\cdots g_{i,t_i}=g_i$.

Proposition IV.1. *The topological space $X=\operatorname{QSpec} R$ is Hausdorff.*

Proof. Choose distinct primes $\mathfrak{p},\mathfrak{q}$ of R, and find $f\in\mathfrak{p}\backslash\mathfrak{q}$. Then $p\in\mathbb{V}(f)$, $q\in\mathbb{X}_f$, and both of these sets are open. □

The partition of $\operatorname{QSpec} R$ arising in the above proof suggests the name "'patch topology." In fact, a similar proof shows that $\operatorname{QSpec} R$ is totally disconnected.

If R is an integral domain, denote by $Q(R)$ its field of fractions.

Proposition IV.2. *For any commutative ring R with 1, $\operatorname{QSpec} R=X$ is compact under the patch topology.*

Proof. This proof imitates the classical algebraic geometry proof that $\operatorname{Spec} R$ is quasi-compact. Consider an open cover of X in the patch topology, with some index set S:

$$X=\bigcup_{i\in S}(\mathbb{V}(f_{i,1},\ldots,f_{i,r_i})\cap X_{g_i})$$

Let $R[\overline{x}]=R[x_{ij}\colon i\in S,\ 1\le j\le r_i]$, and let $I\subseteq R[\overline{x}]$ be the ideal generated by all $f_{ij}x_{ij}-g_i$. Claim $T=R[\overline{x}]/I$ is the zero ring. Assuming the claim is false, let $\mathfrak{p}\subset T$ be a nonzero prime ideal. Then $\mathfrak{p}$ pulls back to a prime $\mathfrak{q}\subset R[X]$ containing every $f_{ij}x_{ij}-g_i$. If $\mathfrak{q}$ pulls back to a prime $\mathfrak{q}'=\mathfrak{q}\cap R$, then for some $i\in S$, it follows $\mathfrak{q}'\in\mathbb{V}(f_{i,1},\ldots,f_{i,r_i})\cap X_{g_i}$, because $\mathfrak{q}'\in X$ lies in some open set in the open cover of X.

Since $\mathfrak{q}' = R \cap \mathfrak{q}$ and $\mathfrak{q}'R[\overline{x}] \subseteq \mathfrak{q}$, it follows that $\mathfrak{q}$ contains both $f_{ij}x_{ij}$ and $f_{ij}x_{ij} - g_i$, $1 \leq j \leq r_i$, and hence contains g_i. Thus $\mathfrak{q}' = R \cap \mathfrak{q}$ contains g_i; a contradiction. Therefore $T = R[\overline{x}]/I$ contains no (proper) prime ideals, proving the claim. From $I = R[\overline{x}]$ it follows that there exists a finite subset $S' \subset S$ and sets $S_i' = \{j \mid 1 \leq j \leq r_i\}$ such that

$$1 = \sum_{i \in S'} \sum_{j \in S_i'} b_{ij}(x)(f_{ij}x_{ij} - g_i) \quad \text{for some} \quad b_{ij}(x) \in R[x_{ij} \colon i \in S', j \in S_i'].$$

Also note that if $b_{ij}(0)$ is the image of $b_{ij}(x)$ under the surjection $R[x] \hookrightarrow R$ by all $x_{ij} \mapsto 0$, then $1 = \sum_{i \in S'} \sum_{j \in S_i'} b_{ij}(0) \cdot g_i$. Thus the ideal $(g_i \colon i \in S')$ is all of R.

It remains to show that the desired finite subcover is

$$X = \bigcup_{i \in S'} \left(\mathbb{V}(f_{i,1}, \ldots, f_{i,r_i}) \cap X_{g_i} \right).$$

Let $\mathfrak{p} \subset R$ be a prime ideal. For $n = |S'|$, write $S' = \{g_1, \ldots, g_n\}$. Then upon renumbering, we may assume $g_1, \ldots g_m \notin \mathfrak{p}$ for some $m \geq 1$ and assume $g_i \in \mathfrak{p}$ for $i > m$. We must show that for some $i \leq m$, all $f_{i,j} \in \mathfrak{p}$. Assume by way of contradiction that for every $i \leq m$ there exists an associated $f_{i,j} \notin \mathfrak{p}$. Consider the ideal

$$J = (\overline{f}_{ij}x_{ij} - \overline{g}_i \colon f_{ij} \notin \mathfrak{p}) + (x_{ij} \colon m < i \leq n, j \in S_i') \subseteq (R/\mathfrak{p})[x_{ij} \colon i \in S', j \in S_i'],$$

where overlines denote coset representatives modulo $\mathfrak{p}$, and consider the sequence

$$R/\mathfrak{p} \hookrightarrow \frac{(R/\mathfrak{p})[x_{ij} \colon i \in S', j \in S_i']}{J} \hookrightarrow Q(R/\mathfrak{p}).$$

Indeed, since $R/\mathfrak{p}$ is a domain, the middle term is isomorphic to the subring of $Q(R/\mathfrak{p})$ generated by R and $\overline{g_i}/\overline{f_{ij}}$ for $i \in S', j \in S_i'$. Since the remaining f_{ij} and g_i lie in $\mathfrak{p}$, the ideal J contains $\sum_{i \in S'} \sum_{j \in S_i'} \overline{b_{ij}(x)}(\overline{f}_{ij}x_i - \overline{g}_i) = 1$. It follows that $R/\mathfrak{p} \hookrightarrow Q(R/\mathfrak{p})$ factors through the zero ring, a contradiction. Therefore, for some

$i \leq m$, the prime $\mathfrak{p}$ does not contain g_i but does contain all f_{ij} for $j \in S_i'$. In other words, there exists $i \in S'$ such that $\mathfrak{p} \in \mathbb{V}(f_{i1}, \ldots, f_{i,r_i}) \cap X_{g_i}$ for an arbitrary choice $\mathfrak{p} \in X$, and a finite subcover exists. □

4.2 Quasi-regular Functions

Let $\phi : R \to S$ be a homomorphism. Then $\phi^* : \operatorname{QSpec} S \to \operatorname{QSpec} R$ remains a continuous map in the patch topology by the usual pull-back of prime ideals, and the localization $\phi_{\mathfrak{p}}$ induces an inclusion of fraction fields $Q(R/\phi^*(\mathfrak{p})) \to Q(S/\mathfrak{p})$. Such a homomorphism is further called a *quasi-isomorphism* if both

1. ϕ^* is a homeomorphism in the patch topology,

2. For every $\mathfrak{p} \in \operatorname{QSpec} S$, $\phi_{\mathfrak{p}} : Q(R/\phi^*(\mathfrak{p})) \to Q(S/\mathfrak{p})$ is an isomorphism.

Thus when constructing functions on $\operatorname{Spec} R$, it suffices to do so on a quasi-isomorphic spectrum. Some examples follow:

Proposition IV.3. *For $f \in R$, the homomorphism $\phi : R \to R_f \times R/(f)$ by $r \mapsto (r/1, r + (f))$ is a quasi-isomorphism.*

Proof. Recall the classical spectrum of $R_f \times R/(f)$ is homeomorphic under ϕ^* to the disjoint union of the spectra of its factors. The same holds for the QSpecs: for one thing, ϕ^* restricted to each component is homeomorphic onto its respective image. For another, the images of these components form a disjoint open cover of $\operatorname{QSpec} R$. Because $\phi_{\mathfrak{p}}$ gives the isomorphism $Q(R_f \times R/(f)/\phi^*(\mathfrak{p})) \cong Q(R/\mathfrak{p})$ for every $\mathfrak{p} \in \operatorname{Spec} R$, the result follows. □

Note that the previous proposition illustrates that the QSpecs of two rings can be quasi-isomorphic when the respective spectra are not.

Proposition IV.4. *Let R_{red} be the reduced ring of R. Then the natural map $\phi : R \to R_{red}$ is a quasi-isomorphism.*

Proof. As every prime of R contains the nilradical and ϕ is surjective, it follows that ϕ^* is a classical homeomorphism, hence a patch homeomorphism. Again $Q(R/\mathfrak{p}) \cong Q(R_{\text{red}}/\phi^*(\mathfrak{p}))$ under $\phi_{\mathfrak{p}}$, and ϕ is a quasi-isomorphism. □

In general, isomorphic spectra are quasi-isomorphic.

Proposition IV.5. *A composition of quasi-isomorphisms is a quasi-isomorphism.*

For the remainder of this section, recall that a continuous bijection of compact Hausdorff spaces is closed, whence open.

Proposition IV.6. *If $\psi : R \to S$ is a quasi-isomorphism and $\phi : R \to T$ is a homomorphism, then $\varphi : T \to S \otimes_R T$ is a quasi-isomorphism.*

Proof.

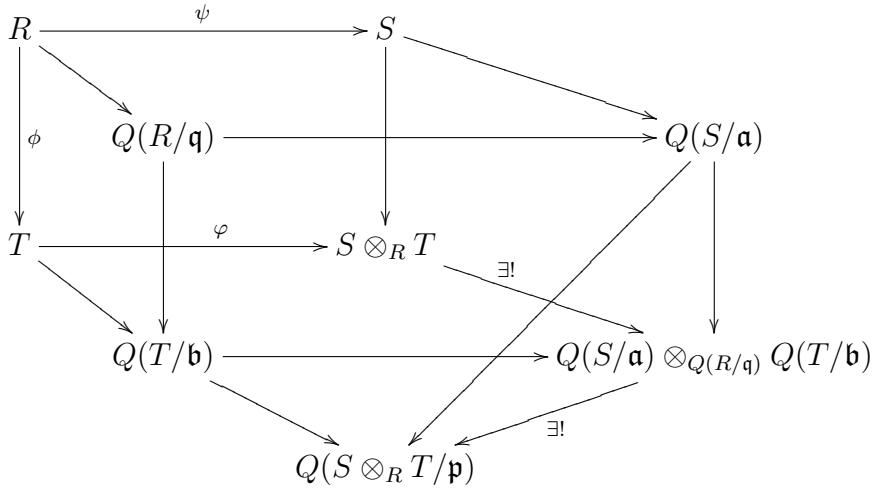

In this proof, tensors are over R unless noted otherwise. First establish a bijection between primes $\mathfrak{p} \in \operatorname{Spec} S \otimes T$ and triples $(\mathfrak{q}, \mathfrak{a}, \mathfrak{b}) \in \operatorname{Spec} R \times \operatorname{Spec} S \times \operatorname{Spec} T$ such that $\mathfrak{a}, \mathfrak{b}$ pull back to $\mathfrak{q}$. Clearly any $\mathfrak{p} \in \operatorname{Spec} S \otimes T$ produces such a triple under

the respective pull-backs. Conversely, such a triple determines the ring

$$F := Q(S/\mathfrak{a}) \otimes_{Q(R/\mathfrak{q})} Q(T/\mathfrak{b}),$$

which is isomorphic as a $Q(R/\mathfrak{q})$-algebra to $Q(T/\mathfrak{b})$, because $\psi^{-1}(\mathfrak{a}) = \mathfrak{q}$ and R, S are quasi-isomorphic. Now, S, T map to F by factoring through their respective residue fields. Hence there exists a unique map $S \otimes T \to F$. Denote by $\mathfrak{p}$ the pull-back to $S \otimes T$ of the zero ideal in F.

In summary, $\mathfrak{p} \in \operatorname{Spec} S \otimes T$ pulls back to a triple of the above form, and such a triple determines a prime of $S \otimes T$. On the one hand, if a triple determines a prime of $S \otimes T$ by the pull back of $(0) \subset F$, then that prime must pull back to the entries of the triple, by the commutativity of the diagram. On the other hand, assume a prime $\mathfrak{p} \in \operatorname{Spec} S \otimes T$ pulls back to the triple $(\mathfrak{q}, \mathfrak{a}, \mathfrak{b})$. Then as $Q(S/\mathfrak{a})$, $Q(T/\mathfrak{b})$ map to $Q(S \otimes T/\mathfrak{p})$, there exist unique maps $S \otimes T \to Q(S \otimes T/\mathfrak{p})$, $F \hookrightarrow Q(S \otimes T/\mathfrak{p})$, and $S \otimes T \to Q(S \otimes T/\mathfrak{p})$ factors uniquely as

$$S \otimes T \to Q(S/\mathfrak{a}) \otimes_{Q(R/\mathfrak{q})} Q(T/\mathfrak{b}) \hookrightarrow Q(S \otimes T/\mathfrak{p}).$$

Therefore the zero ideal of F pulls back to $\mathfrak{p}$, proving the bijection. Now, a triple $(\mathfrak{q}, \mathfrak{a}, \mathfrak{b})$ provides a prime $\mathfrak{b} \in \operatorname{Spec} T$, and as $F \cong Q(T/\mathfrak{b})$, such a prime $\mathfrak{b}$ pulls back to a unique prime $\mathfrak{p} \in \operatorname{Spec} S \otimes T$ that determines a triple containing $\mathfrak{b}$. Thus φ^* is a bijection.

What is more, the diagram shows that the inclusion $S \otimes T/\mathfrak{p} \hookrightarrow Q(S \otimes T/\mathfrak{p})$ factors through $Q(T/\mathfrak{b})$, since $Q(S/\mathfrak{a}) \otimes_{Q(R/\mathfrak{q})} Q(T/\mathfrak{b}) \cong Q(T/\mathfrak{b})$ as noted above. As $Q(S \otimes T/\mathfrak{p})$ is the smallest field containing $S \otimes T/\mathfrak{p}$, it follows that $\varphi_{\mathfrak{p}}$ is an isomorphism of residue fields. Finally, φ^* is continuous in the patch topology and is open as a bijection of compact Hausdorff spaces. $\square$

The goal of quasi-isomorphisms is that one can often define functions and homomorphisms on a quasi-isomorphic QSpec that do not exist on the original space. To this end, define a *quasi-homomorphism* $\phi : R \rightsquigarrow S$ to be a continuous map $\phi^* : \operatorname{QSpec} S \to \operatorname{QSpec} R$ together with field inclusions $\phi_{\mathfrak{p}} : Q(R/\phi^*(\mathfrak{p})) \hookrightarrow Q(S/\mathfrak{p})$ such that there exists a ring T, ring homomorphism $\psi : R \to T$, and a quasi-isomorphism $\gamma : S \to T$ with the following properties:

1. $\phi^* = \psi^* \circ (\gamma^*)^{-1}$,
2. For $\mathfrak{p} \in \operatorname{QSpec} S$ and $\mathfrak{q} = (\gamma^*)^{-1}(\mathfrak{p})$, one has $\phi_{\mathfrak{p}} = (\gamma_{\mathfrak{q}})^{-1} \circ \psi_{\mathfrak{q}}$.

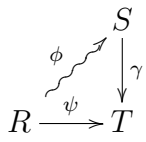

Note that T may not be unique. Indeed, the flexibility in choosing perhaps even a chain of quasi-isomorphic QSpecs will aid in the following generalization of the regular functions on $\operatorname{Spec} R$.

Consider a function f on $\operatorname{QSpec} R$ such that for every $\mathfrak{p} \in \operatorname{QSpec} R$, $f(\mathfrak{p})$ lies in $Q(R/\mathfrak{p})$. Call f a *quasi-regular function* if there exists an associated quasi-homomorphism $\phi : \mathbb{Z}[x] \rightsquigarrow R$ such that $\phi_{\mathfrak{p}}(x) = f(\mathfrak{p})$ for every $\mathfrak{p} \in \operatorname{QSpec} R$. That is, f is considered quasi-regular if there exists a quasi-isomorphic ring T such that under the quasi-isomorphism γ, f is regular on T.

Proposition IV.7. *The quasi-isomorphism ϕ above is uniquely determined by f.*

Proof. Let ϕ, φ be quasi-homomorphisms associated to f. Then both define the same map of topological spaces. If $\mathfrak{q}$ is the pull-back of $\mathfrak{p} \in \operatorname{QSpec} R$, then $\phi_{\mathfrak{p}}, \varphi_{\mathfrak{p}} : Q(\mathbb{Z}[x]/\mathfrak{q}) \to Q(R/\mathfrak{p})$ have $\phi_{\mathfrak{p}}(1) = 1 = \varphi_{\mathfrak{p}}(1)$ and $\phi_{\mathfrak{p}}(x) = f(\mathfrak{p}) = \varphi_{\mathfrak{p}}(x)$. These assignments determine $\phi_{\mathfrak{p}} = \varphi_{\mathfrak{p}}$ uniquely. Upon recalling that the choice of the ring T quasi-isomorphic to R need not be unique, the proof is complete. □

Proposition IV.8. *For any ring R, the constant functions 0, 1 are quasi-regular on* $\operatorname{QSpec} R$. *If f, g, are quasi-regular functions, then $f+g, f-g, fg$ are also quasi-regular.*

Proof. For the constant functions 0,1, define the homomorphisms $\mathbb{Z}[x] \to R$ such that $x \mapsto 0, 1$, respectively. In both cases, R itself plays the role of the ring T above.

Next assume f, g are quasi-regular, with associated quasi-homomorphisms $\mathbb{Z}[x] \rightsquigarrow R$ defined by

$$\psi_i : \mathbb{Z}[x] \to T_i, \quad x \mapsto x_i, \quad i \in \{1, 2\},$$

respectively. It follows from Proposition IV.6 that R is quasi-isomorphic to $T_1 \otimes_R T_2$. Hence define

$$\psi : \mathbb{Z}[x] \to T_1 \otimes_R T_2 \quad \text{by} \quad x \mapsto x_1 \otimes x_2.$$

Now, the rings $R,\ T_1,\ T_2$, and $T_1 \otimes T_2$ are pairwise quasi-isomorphic, and since the following diagram commutes,

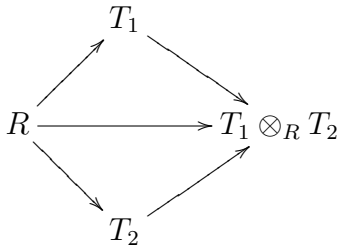

the product of x_1, x_2 in $T_1 \otimes T_2$ is precisely $x_1 \otimes x_2$, which is mapped to $f(\mathfrak{p})g(\mathfrak{p})$ upon taking the residue fields at any $\mathfrak{p} \in \operatorname{Spec} R$. The proof for $f+g$ or $f+(-g)$ proceeds similarly, by sending $\psi : x \mapsto x_1 \otimes 1 + 1 \otimes x_2$. □

One useful construction is the quasi-inverse of a quasi-regular function:

$$f^\star(\mathfrak{p}) = \begin{cases} 1/f(\mathfrak{p}) & f(\mathfrak{p}) \neq 0 \\ 0 & f(\mathfrak{p}) = 0 \end{cases}.$$

Proposition IV.9. *If f is a quasi-regular function on* $\operatorname{QSpec} R$*, then so is $f^\star$.*

Proof. Associated to f is a quasi-isomorphism $R \to T$ and a homomorphism $\mathbb{Z}[x] \to T$, with $x \mapsto g \in T$, say. Define a new quasi-homomorphism $\phi : \mathbb{Z}[x] \rightsquigarrow R$ by the quasi-isomorphism $\gamma : R \to T \to T_g \times T/(g) = S$ and the homomorphism $\psi : \mathbb{Z}[x] \to S$ sending $x \mapsto (1/g, 1)$. Then $\phi_{\mathfrak{p}}(x)$ agrees with $f^\star$ at every point of $\operatorname{QSpec} R$. □

Note further that $(f^\star)^2 f = f^\star$, $f^2 f^\star = f$, and $(f^\star)^\star = f$, where the latter requires the fact that, in the notation of the above proof, $S_{(1/g,1)} \times S/((1/g, 1)) \cong S_{(1/g,1)} \cong T_g \times T/(g)$.

4.3 QI-rings

To axiomatize a ring in which quasi-inverses exist as above, define a *QI-ring* to be a commutative ring R with 1 on which an operation $f \mapsto f^\star$ is defined and satisfies the following for every $f \in R$:

$$(f^\star)^2 f = f^\star \quad \text{and} \quad f^2 f^\star = f.$$

It follows immediately that every element f of a QI-ring is a zero-divisor: $f(1-ff^\star) = 0$. Furthermore, if $\mathfrak{q} \in \operatorname{Spec} R$ and $f \in R$, then $\mathfrak{q} \ni f$ or $\mathfrak{q} \ni (1 - ff^\star)$, but not both.

For an arbitrary commutative ring R with 1, define $\hat{R}$ to be the ring of quasi-regular functions on $\operatorname{QSpec} R$.

Proposition IV.10. *For every commutative ring R with 1, $\hat{R}$ is a QI-ring.*

Proof. Proposition IV.8 above shows that the quasi-regular functions form a ring. So it remains to verify the axioms for $f^\star$ in a QI-ring, by considering the values of $f^\star(\mathfrak{p})^2 f(\mathfrak{p})$, $f(\mathfrak{p})^2 f^\star(\mathfrak{p})$ in each $Q(R/\mathfrak{p})$. □

Ideals in QI-rings are closed under the operation $f \mapsto f^\star$:

Proposition IV.11. *For every f in a QI-ring R,*

(a) $f^\star \in (f)$

(b) $f \in (f^n)$ *for every* $n \geq 1$

Proof. Property (a) follows from the first axiom. For (b), induct on n: The case $n = 1$ is clear. If $f \in (f^{n-1})$, then $f = gf^{n-1}$ for some $g \in R$, and $gf^n = f \cdot gf^{n-1} = f^2$. Thus $(f^n) \ni f^2 f^\star = f$. □

Corollary IV.12. *Every ideal I of a QI-ring R is radical, and R is reduced.*

Proof. If $f^n \in I$, then part (b) above yields that $f \in I$. Now consider $I = (0)$. □

Corollary IV.13. *Every prime ideal $\mathfrak{p}$ of a QI-ring is maximal.*

Proof. For any f not in $\mathfrak{p}$, consider $I = (f) + \mathfrak{p}$. As remarked above, one must have $(1 - ff^\star) \in \mathfrak{p}$ and $ff^\star \in I$. Then $1 \in I$. □

In an arbitrary commutative ring with 1, every element is a quasi-regular function on $\operatorname{QSpec} R$, possibly the zero function:

Proposition IV.14. *Let $\psi : R \to \hat{R}$ be the homomorphism assigning every element of R to its associated quasi-regular function. Then the kernel of ψ is the nilradical $\sqrt{(0)}$.*

Proof. For every $f \in R$, $f(\mathfrak{p})$ is the image of f in the residue field $Q(R/\mathfrak{p})$. Thus $f(\mathfrak{p}) = 0$ for every $\mathfrak{p} \in \operatorname{QSpec} R$ if and only if f lies in every prime ideal of R. □

Lemma IV.15. *Let f be a quasi-regular function on* $\operatorname{QSpec} R$. *Then the set of points where $f = 0$ is open and closed in the patch topology.*

Proof. Consider the quasi-homomorphism associated to f:

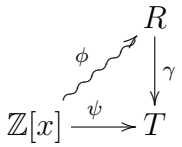

Let $\psi(x) = \xi \in T$. For $\mathfrak{p} \in \operatorname{QSpec} R$, find $\mathfrak{q} \in \operatorname{QSpec} T$ such that $\gamma^*(\mathfrak{q}) = \mathfrak{p}$. Then $f(\mathfrak{p}) = 0 \in Q(R/\mathfrak{p})$ if and only if $\xi \mapsto 0 \in Q(T/\mathfrak{q})$ if and only if $\xi \in \mathfrak{q}$ if and only if $\mathfrak{q} \in \mathbb{V}(\xi)$, an open and closed set in the patch topology. As γ^* is a homeomorphism, f vanishes on the open and closed set $\gamma^*(\mathbb{V}(\xi))$. □

Proposition IV.16. *The homomorphism $\psi : R \to \hat{R}$ induces a homeomorphism $\psi^*\colon \operatorname{Spec} \hat{R} \to \operatorname{QSpec} R$.*

Proof. Choose a prime ideal $\mathfrak{p}$ of R, and define $\mathfrak{m}_\mathfrak{p} \subset \hat{R}$ to be the ideal of quasi-regular functions vanishing at $\mathfrak{p}$. Then $\mathfrak{m}_\mathfrak{p} \not\ni 1$ is proper, contains $\psi(\mathfrak{p})$, and is prime because the product of two functions fg vanishes at $\mathfrak{p}$ if and only if $f(\mathfrak{p}) = 0$ or $g(\mathfrak{p}) = 0$ (in particular, $\mathfrak{m}_\mathfrak{p}$ is maximal). Hence the contraction of $\mathfrak{m}_\mathfrak{p}$ is a prime ideal of R containing $\mathfrak{p}$.

Define the evaluation homomorphism $\operatorname{ev}_\mathfrak{p}\colon \hat{R} \to Q(R/\mathfrak{p})$ by $f \mapsto f(\mathfrak{p})$. Then the induced evaluation map $\hat{R}/\mathfrak{m}_\mathfrak{p} \to Q(R/\mathfrak{p})$ is an isomorphism, and one has

$$R/\mathfrak{p} \xrightarrow{\overline{\psi}} \hat{R}/\mathfrak{m}_\mathfrak{p} \xrightarrow{\sim} Q(R/\mathfrak{p}).$$

To show the contraction of $\mathfrak{m}_\mathfrak{p}$ is precisely $\mathfrak{p}$, it remains to show that the first map is injective. To that end, $\overline{\psi}$ sends the class of $f \in R$ to its value at $\mathfrak{p}$ as a quasi-regular function, so the class of f lies in the kernel of $\overline{\psi}$ if and only if $f \in \mathfrak{p}$. Hence $\psi^{-1}(\mathfrak{m}_\mathfrak{p}) = \mathfrak{p}$, and ψ^* is surjective.

It suffices to show that every maximal ideal $\mathfrak{m} \subset \hat{R}$ is the (unique, maximal) ideal of quasi-regular functions vanishing at some point of $\operatorname{QSpec} R$. Then since quasi-

regular functions separate points of $\operatorname{QSpec} R$, it would follow that ψ^* is a continuous bijection of compact spaces, whence open. To this end, choose $\mathfrak{m} \subset \hat{R}$, and define

$$Z(\mathfrak{m}) = \{\mathfrak{p} \in \operatorname{QSpec} R \mid f(\mathfrak{p}) = 0 \ \forall \ f \in \mathfrak{m}\}.$$

Claim $Z(\mathfrak{m})$ is not empty: otherwise there exists an open (and closed) cover of $X = \operatorname{QSpec} R$ by the images of $X_{f_\alpha f_\alpha^*} := \mathbb{V}(f_\alpha f_\alpha^*)^c$ for a collection of $f_\alpha \in \mathfrak{m}$. Take a finite subcover, say $X = \bigcup \psi^*(X_{f_i f_i^\star}))$ for $f_i \in \mathfrak{m}$. The intersections of elements of this subcover are finite intersections, of the form

$$X_{f_{i_1} f_{i_1}^\star \cdots f_{i_k} f_{i_k}^\star} \quad \text{for} \quad f_{i_j} \in \mathfrak{m}.$$

Write $f_{i_1} f_{i_1}^\star \cdots f_{i_k} f_{i_k}^\star = g_i$. Then $X_{f_j} \backslash X_{g_i} = X_{f_j(1-g_i g_i^\star)}$, with $f_j(1 - g_i g_i^\star) \in \mathfrak{m}$. Removing intersections in this way yields a disjoint cover $X = \bigsqcup \psi^*(X_{h_\ell})$ for $h_\ell \in \mathfrak{m}$. By construction,

$$h_\ell(\mathfrak{p}) = \begin{cases} 1 & \mathfrak{p} \in X_{h_\ell} \\ 0 & \text{otherwise.} \end{cases}$$

Therefore $\sum h_\ell \in \mathfrak{m}$ is the constant function 1, which is absurd. The claim is proven.

Finally, choose any $\mathfrak{p} \in Z(\mathfrak{m})$. It follows that $\mathfrak{m} \subseteq \mathfrak{m}_\mathfrak{p}$, forcing equality. □

Corollary IV.17. *Let R be a commutative ring with finitely many prime ideals. Then $\hat{R}$ is a product of (residue) fields.*

Proof. Let $\mathfrak{m}_1, \ldots, \mathfrak{m}_n$ be the maximal ideals of $\hat{R}$ corresponding to the prime ideals of R. Then these are all the prime ideals of $\hat{R}$. Because $\hat{R}$ is reduced, $\cap \mathfrak{m}_i = 0$, and of course the $\mathfrak{m}_i$ are coprime in pairs. It follows that $\hat{R} \to \prod \hat{R}/\mathfrak{m}_i$ defined by $f \mapsto (f + \mathfrak{m}_1, \ldots, f + \mathfrak{m}_n)$ is an isomorphism. □

If R has infinitely many prime ideals, then $\hat{R} \to \prod \hat{R}/\mathfrak{m}_i$ remains an injection.

Recall the homomorphism $\psi : R \to \hat{R}$ sending elements of R to quasi-regular functions on $\operatorname{QSpec} R$. Let $\tilde{R}$ be the smallest subring of $\hat{R}$ closed under $f \mapsto f^\star$ and containing $\psi(R)$. The following proposition shows that the quasi-regular functions on $\operatorname{QSpec} R$ locally look like $\tilde{R}$.

Proposition IV.18. *For every $f \in \hat{R}$ and every $\mathfrak{p} \in \operatorname{QSpec} R$, there exists a neighborhood U of $\mathfrak{p}$ and an element $g \in \tilde{R}$ such that $g(\mathfrak{q}) = f(\mathfrak{q})$ for every $\mathfrak{q} \in U$.*

Proof. Let $\alpha = f(\mathfrak{p}) \in Q(R/\mathfrak{p})$. Then α is a finite concatenation of sums, products, and fractions of cosets $\overline{f_1}, \ldots, \overline{f_n}$ in $R/\mathfrak{p}$. This concatenation defines a straight-line program Γ, taking the quasi-inverse operation as multiplicative inversion. Hence define $g \in \tilde{R}$ by $\Gamma(f_1, \ldots, f_n)$, choosing coset representatives f_i for each $\overline{f_i}$. It follows that the quasi-regular function $f - g \in \hat{R}$ has $(f-g)(\mathfrak{p}) = 0$, and the vanishing set U of $f - g$ on $\operatorname{QSpec} R$ is nonempty and open. □

An open set $\mathbb{V}(f) \cap X_g \subseteq \operatorname{QSpec} R$ in the patch topology has an indicator function in $\tilde{R}$:

$$\chi_{f,g}(\mathfrak{p}) = gg^\star(1 - ff^\star)(\mathfrak{p}) = \begin{cases} 1 & \mathfrak{p} \in U \\ 0 & \text{else.} \end{cases}$$

Of course, finite intersections of basis elements correspond to products of indicator functions. These functions allow us indeed to *patch* together the local $\tilde{R}$-representatives of quasi-regular functions.

Proposition IV.19. $\hat{R} = \tilde{R}$

Proof. By definition, $\tilde{R} \hookrightarrow \hat{R}$. Conversely, choose $f \in \hat{R}$. By the previous proposition, there exists an open cover $\operatorname{QSpec} R = \cup U_\alpha$ and $g_\alpha \in \tilde{R}$ such that $f = g_\alpha$ on U_α. As $\operatorname{QSpec} R$ is compact, choose a finite subcover. We may assume that the U_i in the subcover are basis elements of the topology, and we may further assume that

$\operatorname{QSpec} R$ is covered by a disjoint union $\sqcup_1^n W_i$, where each W_i is the finite intersection of some basis elements U_i. Each W_i has an indicator function χ_i; note $\chi_i \in \tilde{R}$ because each U_i is a basis element of the topology. Then the quasi-regular function $g = \sum_1^n \chi_i g_i$ equals f at every $\mathfrak{p} \in \operatorname{QSpec} R$, and $g \in \tilde{R}$. □

In particular, note that $\hat{\hat{R}} = \hat{R}$.

Corollary IV.20. *Suppose* $\psi^*\colon \operatorname{Spec} \hat{R} \to \operatorname{QSpec} R$ *sends* $\mathfrak{m}_\mathfrak{p} \mapsto \mathfrak{p}$. *Then* $\mathfrak{m}_\mathfrak{p}$ *is generated by* $\mathfrak{p}$ *and the* $1 - gg^\star$ *for all* $g \notin \mathfrak{p}$.

Proof. Let $I \subset \hat{R}$ be the ideal generated by $\mathfrak{p}$ and the $1 - gg^\star$ for all $g \notin \mathfrak{p}$. Then $I \subseteq \mathfrak{m}_\mathfrak{p}$ Choose $h \in \mathfrak{m}_\mathfrak{p} - I$. Then h vanishes at $\mathfrak{p}$. Consider $X_h \subset \operatorname{Spec} \hat{R}$. For every $\mathfrak{m}_\mathfrak{q} \in X_h$ (using the notation $\mathfrak{m}_\mathfrak{q} \mapsto \mathfrak{q} \in \operatorname{QSpec} R$), construct $\chi_\mathfrak{q} \in I$ such that $\chi_\mathfrak{q}(\mathfrak{q}) = 1$, as follows:

1. If $\mathfrak{q} \subsetneq \mathfrak{p}$, then find $f \in \mathfrak{p} - \mathfrak{q}$, and set $\chi_\mathfrak{q} = ff^\star$.
2. Otherwise if $\mathfrak{q} \neq \mathfrak{p}$ then find $g \in \mathfrak{q} - \mathfrak{p}$ and set $\chi_\mathfrak{q} = 1 - gg^\star$.

Then $X_h \subseteq \cup X_{\chi_\mathfrak{q}}$, for $\chi_\mathfrak{q} \in I$. Now, $X_h = \mathbb{V}(1 - hh^\star) \simeq \operatorname{Spec} \hat{R}/(1 - hh^\star)$ is compact. Thus for finitely many $\mathfrak{q}$, there exist $a_\mathfrak{q} \in \hat{R}$ such that the quasi-regular function $\chi = \sum a_\mathfrak{q} \chi_\mathfrak{q}$ equals 1 on X_h. By construction, $\chi \in I$, and

$$\chi \cdot h(x) = \begin{cases} 0 & x \in \mathbb{V}(h) \\ h(x) & x \in X_h \end{cases}.$$

Therefore $h \in I$, whence $\mathfrak{m}_\mathfrak{p} \subseteq I$. □

Lemma IV.21. *If* $\phi : R \to S$ *is a homomorphism of commutative rings, then there exists a unique* $\phi : \hat{R} \to \hat{S}$ *making the following diagram commute:*

$$\begin{array}{ccc} R & \xrightarrow{\phi} & S \\ \downarrow & & \downarrow \psi \\ \hat{R} & \xrightarrow{\hat{\phi}} & \hat{S} \end{array}$$

Proof. Recall $f \in \hat{R}$ is a straight-line program $\Gamma(f_1, \ldots, f_n)$ of QI-ring operations on finitely many $f_i \in R$, and define $\hat{\phi}(f) = \Gamma(\psi \circ \phi(f_1), \ldots, \psi \circ \phi(f_n))$. Suppose that $g = \Gamma'(g_1, \ldots, g_m)$ is another program for f. If $\mathfrak{p} \in \operatorname{QSpec} S$, then ϕ induces

$$Q(R/\phi^{-1}(\mathfrak{p})) \hookrightarrow Q(S/\mathfrak{p}).$$

As both programs for f have the same image in $Q(R/\phi^{-1}(\mathfrak{p}))$, it follows $\hat{\phi}(f)(\mathfrak{p}) = \hat{\phi}(g)(\mathfrak{p}) \in Q(S/\mathfrak{p})$, and $\hat{\phi}$ is well-defined.

This map is a homomorphism that makes the diagram commute, because addition and multiplication are possible instructions in the program Γ. If $\pi\colon \hat{R} \to \hat{S}$ also makes the diagram commute, then $\pi_{\mathfrak{m}} = \hat{\phi}_{\mathfrak{m}}$ for every prime $\mathfrak{m} \in \hat{R}$. The two maps then give the same injections of residue fields, so $\hat{\phi}$ is unique. □

Proposition IV.22. *Let R be a commutative ring, I an ideal, $\hat{I}$ its extension in $\hat{R}$. Then $\widehat{R/I} \cong \hat{R}/\hat{I}$.*

Proof. Under the inclusion $\phi\colon R/I \to \hat{R}/\hat{I}$, each prime $\mathfrak{p} \supseteq I$ of R is the image of the unique maximal vanishing ideal $\mathfrak{m}_{\mathfrak{p}} \subseteq \hat{I}$ of $\hat{R}$. As $Q(R/\mathfrak{p}) \cong Q(\hat{R}/\mathfrak{m}_p)$, it follows that $R/I \to \hat{R}/\hat{I}$ is a quasi-isomorphism. That means that every f in $\hat{R}/\hat{I}$ provides a quasi-regular function on R/I, whence there exists a homomorphism $\alpha\colon \hat{R}/\hat{I} \to \widehat{R/I}$. Now, $f \in \hat{R}/\hat{I}$ maps to the zero function if and only if f lies in $\mathfrak{m}_{\mathfrak{p}}$ for every prime $\mathfrak{p} \supseteq I$. Since all primes of $\hat{R}$ are maximal, the radical ideal $\hat{I}$ equals the intersection of all such $\mathfrak{m}_{\mathfrak{p}}$. Hence in fact $\alpha\colon R/I \to \hat{R}/\hat{I}$ is injective.

To prove α is surjective, we first claim that

$$\widehat{\hat{R}/\hat{I}} = \hat{R}/\hat{I}, \quad \text{that is,} \quad \widetilde{\hat{R}/\hat{I}} = \hat{R}/\hat{I}.$$

It suffices to show that if $(f - g) \in \hat{I}$, then $f^\star - g^\star \in \hat{I}$. To that end, note

$$f^\star - f^\star g g^\star = (f^\star)^2(1 - gg^\star)(f - g) \in \hat{I}$$

and likewise

$$-g^\star + g^\star f f^\star = (g^\star)^2(1 - ff^\star)(f - g) \in \hat{I}.$$

Thus

$$f^\star - g^\star + (g^\star f f^\star - f^\star g g^\star) \in \hat{I},$$

but also

$$(g^\star f f^\star - f^\star g g^\star) = f^\star g^\star (f - g) \in \hat{I}.$$

So indeed $f^\star - g^\star \in \hat{I}$, and the claim follows.

Hence choose $f \in \widehat{R/I}$, and suppose $f = \Gamma(f_1 + I, \ldots, f_2 + I)$ for some $f_i \in R$ and a straight-line program Γ. By the Lemma IV.21, $\phi\colon R/I \to \hat{R}/\hat{I}$ induces a unique $\hat{\phi}\colon \widehat{R/I} \to \widehat{\hat{R}/\hat{I}} = \hat{R}/\hat{I}$, and

$$\hat{\phi}(f) = \Gamma(f_1 + \hat{I}, \ldots, f_n + \hat{I}) = \Gamma(f_1, \ldots, f_n) + \hat{I}.$$

Of course, $\alpha\colon \Gamma(f_1, \ldots, f_n) + \hat{I} \mapsto f$ as a quasi-regular function on $\operatorname{QSpec} R/I$, whence α is surjective. □

In the study of equivalence relation ideals, we will consider tensor products of QI-rings.

Proposition IV.23. *There is a homomorphism $\hat{R} \otimes \hat{R} \to \widehat{R \otimes R}$ inducing $\widehat{R \otimes R} \simeq \widehat{\hat{R} \otimes \hat{R}}$ and natural bijections*

$$\operatorname{Spec} R \otimes R \leftrightarrow \operatorname{Spec} \widehat{R \otimes R} \leftrightarrow \operatorname{Spec} \hat{R} \otimes \hat{R}.$$

Proof. Consider the diagram

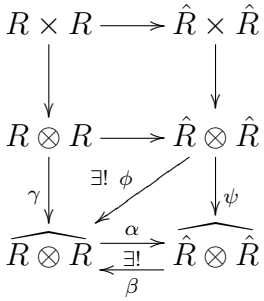

Here, the maps γ and ψ are the natural homomorphisms of a ring into its QI-ring. The map α makes the lower square commute uniquely by Lemma IV.21. That lemma also yields a commuting diagram

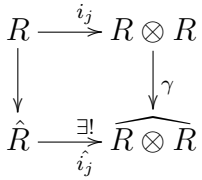

where i_j is the embedding of R into the jth factor. Hence the unique homomorphism ϕ arises from the universal property of the tensor product: the triangle above ϕ commutes. By construction of the $\hat{i}_j$,

$$\phi(\Gamma_f(f_1, \ldots, f_n) \otimes \Gamma_g(g_1, \ldots, g_m)) = \Gamma_f(f_1 \otimes 1, \ldots, f_n \otimes 1) \cdot \Gamma_g(1 \otimes g_1, \ldots, 1 \otimes g_m),$$

where $f_i, g_i \in R$ and Γ_f, Γ_g are straight-line programs of binary and q-operations. Applying Lemma IV.21 again yields the unique homomorphism β making the lower-right triangle commute.

Now, $\psi(\Gamma_f(f_1, \ldots, f_n) \otimes \Gamma_g(g_1, \ldots, g_m)) = \Gamma_f(f_1, \ldots, f_n) \otimes \Gamma_g(g_1, \ldots, g_m)$ considered as a quasi-regular function on $\operatorname{QSpec} \hat{R} \otimes \hat{R}$. To show α also makes the lower-right triangle commute, it suffices to show that in $\widehat{\hat{R} \otimes \hat{R}}$,

$$\Gamma_f(f_1, \ldots, f_n) \otimes \Gamma_g(g_1, \ldots, g_m) = \Gamma_f(f_1 \otimes 1, \ldots, f_n \otimes 1) \cdot \Gamma_g(1 \otimes g_1, \ldots, 1 \otimes g_m),$$

in fact, to show this just for $\Gamma_f(f_1, \ldots, f_n) \otimes 1$.

Choose $f_1, \ldots, f_n \in R$ and a straight-line program of q-operations Γ. Let $\hat{f} = \Gamma(f_1, \ldots, f_n) \in \hat{R}$. Consider $\hat{f} \otimes 1 \in \hat{R} \otimes \hat{R}$. When we apply ψ to $\hat{f} \otimes 1 \in \hat{R} \otimes \hat{R}$, we consider the image as a quasi-regular function on $X = \mathrm{QSpec}\, \hat{R} \otimes \hat{R}$.

If $\mathfrak{p} \in X$, then $\mathfrak{p}$ contracts to some $\mathfrak{p}_1$ and $\mathfrak{p}_2$ in the first and second factor, respectively. Let $\hat{f}(\mathfrak{p}_1) = a \in \hat{R}/\mathfrak{p}_1 \hookrightarrow Q(\hat{R} \otimes \hat{R}/\mathfrak{p})$ and evaluate

$$\Gamma(f_i \otimes 1 \mid \forall i)(\mathfrak{p}) = b \in Q(\hat{R} \otimes \hat{R}/\mathfrak{p}).$$

Recalling $\hat{f} = \Gamma(f_i \mid \forall i) \in \hat{R}$, if $a = b$ when Γ is a single q-operation on a single input $f_1 \in R$, then $a = b$ for more complicated programs as well: the binary and quasi-regular operations in Γ correspond to binary operations and inversions in $Q(\hat{R} \otimes \hat{R}/\mathfrak{p})$.

Thus without loss of generality, assume $\hat{f} = f^\star$. Then $a = 1/f$ if $f(\mathfrak{p}_1) \neq 0$, and $a = 0$ otherwise. Of course, $f(\mathfrak{p}_1) = 0$ if and only if $(f \otimes 1)(\mathfrak{p}) = 0$. If $(f \otimes 1)(\mathfrak{p}) \neq 0$, then b is the inverse of $f \otimes 1 = (f \otimes 1)(1 \otimes 1)$ in $Q(\hat{R} \otimes \hat{R}/\mathfrak{p})$, that is, b is the inverse of the image of f under $\hat{R}/\mathfrak{p}_1 \hookrightarrow Q(\hat{R} \otimes \hat{R}/\mathfrak{p})$. Hence $a = b$. Tracing through the argument, then, $\psi(\Gamma(f_i \mid \forall i) \otimes 1) \in \widehat{\hat{R} \otimes \hat{R}}$ is the same quasi-regular function as $\Gamma(f_i \otimes 1 \mid i) \in \widehat{\hat{R} \otimes \hat{R}}$. Taking sums and products of elements $\hat{f} \otimes 1, 1 \otimes \hat{g} \in \hat{R} \otimes \hat{R}$, it follows that α indeed makes the lower triangle in the main diagram commute.

To finish the proof, we show α is an isomorphism, from which follow the required bijections of spectra. Choose $\Gamma(f_i \otimes g_i \mid \forall i) \in \widehat{\hat{R} \otimes \hat{R}}$. Here, $f_i, g_i \in \hat{R}$, so, for example, $f_i = \Gamma_{f_i}(f_{ij} \mid \forall j)$ as a straight-line program on elements $f_{ij} \in R$. Since $\psi = \alpha \circ \phi$, we have

$$f_i \otimes 1 = \Gamma_{f_i}(f_{ij} \mid \forall j) \otimes 1 = \Gamma_{f_i}(f_{ij} \otimes 1 \mid \forall j),$$

and so by composing quasi-regular functions, we may assume that $f_i, g_i \in R$. Consider then $\Gamma(f_i \otimes g_i \mid \forall i) \in \widehat{R \otimes R}$. It follows immediately that $\alpha(\Gamma(f_i \otimes g_i \mid \forall i) = \Gamma(f_i \otimes g_i \mid \forall i) \in \widehat{\hat{R} \otimes \hat{R}}$, and α is surjective.

It remains to show that $\beta \circ \alpha$ is the identity on $\widehat{R \otimes R}$. We have again $\alpha(\Gamma(f_i \otimes g_i \mid \forall i) = \Gamma(f_i \otimes g_i \mid \forall i) \in \widehat{\hat{R} \otimes \hat{R}}$, for $f_i, g_i \in R$. To apply the map β, we first lift each f_i, g_i to $\hat{R}$, to obtain lifts of the $f_i \otimes g_i$ in $\hat{R} \otimes \hat{R}$.

$$\begin{array}{ccc} \hat{R} \otimes \hat{R} & \xrightarrow{\phi} & \widehat{R \otimes R} \\ \downarrow{\scriptstyle \psi} & & \| \\ \widehat{\hat{R} \otimes \hat{R}} & \xrightarrow{\beta} & \widehat{R \otimes R} \end{array}$$

Two lifts of f_i, say, differ by some $h_i \mapsto 0 \in \hat{R}$, that is, by h_i in the nilradical of R, and so h_i vanishes as a quasi-regular function on $\operatorname{QSpec} R$. Hence the lifts of the $f_i \otimes g_i$ will be well-defined as quasi-regular functions on $\operatorname{QSpec} R \otimes R$ when we apply ϕ. Lastly, β applies Γ to the $f_1 \otimes g_1, \ldots, f_n \otimes g_n$ to obtain $\Gamma(f_i \otimes g_i \mid \forall i)$. Thus $\beta \circ \alpha$ is the identity, and in particular, α is injective. The result follows. $\square$

CHAPTER V

Polynomial Bounds for Invariant Functions Separating Orbits

5.1 Introduction

5.1.1 Background

When a linear algebraic group G acts on an affine variety V over a field k, the orbit of $x \in V$ is the set

$$G \cdot x = \{g \cdot x \mid \forall g \in G\}.$$

Applications of invariant theory, such as computer vision, dynamical systems, and structural chemistry, demand constructive and more efficient techniques to distinguish the orbits of a group action. When the group acts rationally, recall that there exists a finitely generated subalgebra $S \subseteq k[V]^G$ with the following property: Let $p, q \in V$ have disjoint orbit closures, and suppose there exists $f \in k[V]^G$ such that $f(p) \neq f(q)$. Then there exists $h \in S$ such that $h(p) \neq h(q)$[6]. We say that the function h (and the algebra S) *separates* the orbit closures of p and q. Note that the functions in S, called *separating invariants*, separate as many orbits as does $k[V]^G$. Since G is a linear algebraic group, $\overline{G \cdot p} = \overline{G \cdot q}$ implies $G \cdot p = G \cdot q$, because orbits are open in their closures.

This separating subalgebra S has several weaknesses. For one, existence proofs for S may not be constructive for all algebraic groups: Kemper's algorithm to con-

struct S assumes a reductive group [33]. Even in the constructive case, although polynomial bounds exist for the degrees of generators for $k[V]^G$ under the action of a linearly reductive group [5],construction algorithms for separating invariants do not, for general G, provide good bounds on the size of a separating subset, the degrees of its elements, or the complexity of its computation. Kemper's algorithm, for example, requires two Gröbner basis calculations, a normalization algorithm, and an inseparable closure algorithm. Domokos used polarization to cut down the number of variables needed in separating invariants of reducible representations [9], while Kemper provided new bounds, when G is finite, on the required number of separating invariants [34].

As a more serious limitation, the invariant ring $k[V]^G$, and hence any subalgebra, may fail to separate orbit closures. Even when G is reductive, the polynomials in $k[V]^G$ can separate $G \cdot p$ and $G \cdot q$ if and only if $\overline{G \cdot p} \cap \overline{G \cdot q} = \emptyset$. For example, when the multiplicative group $G = k^*$ acts on $\mathbb{A}^2$ by scaling points, one finds $k[x, y]^G = k$.

5.1.2 Separating Orbits with Constructible Functions

To overcome the limitations of the invariant ring, we expand the set of regular functions on a variety to include a *quasi-inverse* $f^\star$ of a regular function f:

$$f^\star(p) = \begin{cases} 1/f(p) & f(p) \neq 0 \\ 0 & f(p) = 0 \end{cases}.$$

For $R = k[V]$, k algebraically closed, let $\widehat{R}$ denote the ring of *constructible* functions $V \to k$ obtained by defining the quasi-inverse on R. For example, if $f, g \in R$, then $(f^\star + g)^\star \in \widehat{R}$. In fact, one can show that for any $h \in \widehat{R}$, there exists finitely many locally closed sets $E_i \subseteq V$ and f_i regular on E_i such that

$$h = \sum_{i=1}^{k} f_i \cdot \chi_{E_i}$$

where χ_{E_i} is the characteristic function of a constructible subset $E_i \subseteq \operatorname{Spec} R$.

For a given group action, we seek to write down a finite set $\mathcal{C}$ of invariant, constructible functions that separate orbits. That is, if p, q lie in different orbits, then some function $f \in \mathcal{C}$ has $f(p) \neq f(q)$. Even better, we would like the construction of f at p to be reasonably simple. To measure the complexity of f, we measure its length as a straight line program over $\widehat{k[V]}$, granting unit cost to all ring operations and the quasi-inverse. Of course, the evaluation of such f at $p \in V$ requires branching, but counting the operations needed to write down f serves as an analog of classical degree bounds for invariants.

Over an algebraically closed field k, fix an embedding of an m-dimensional linear algebraic group $G \hookrightarrow \mathbb{A}^\ell$. Let $R = k[x_1, \ldots, x_n]$, let $\rho : G \hookrightarrow GL_n(k)$ be a representation, let r be the maximal dimension of an orbit, and let $N = \max\{\deg(\rho_{ij})\}$ be the degree of the representation.

Theorem V.1. *There is an algorithm to produce a finite set $\mathcal{C} \subset \widehat{R}$ of invariant, constructible functions with the following properties:*

1. *The set $\mathcal{C}$ separates orbits.*
2. *The size of $\mathcal{C}$ grows as $O(n^2 N^{(\ell+m+1)(r+1)})$.*
3. *The $f \in \mathcal{C}$ can be written as straight line programs, such that the sum of their lengths is $O(n^3 N^{3\ell(r+1)+r})$.*

Hence the problem of deciding if two points lie in the same orbit can be solved with a polynomial number of algebraic operations in the coordinates of the points.

More explicitly, for $p \in \mathbb{A}^n$ consider the orbit map $\sigma_p \colon G \to \mathbb{A}^n$ defined by $\sigma_p \colon g \mapsto g \cdot p$. Note that $\overline{G \cdot p}$ is defined by the polynomials in the kernel of $\sigma_p^* \colon k[x_1, \ldots, x_n] \to k[G]$. These polynomials amount to algebraic relations on the

images $\sigma_p^*(x_1), \ldots, \sigma_p^*(x_n)$ in $k[G]$. One can find all such relations up to some degree d by Gaussian elimination. The coefficients of these relations vary with p, but they cannot in general be written as regular functions of p. We may nevertheless write them with constructible functions, especially utilizing the fact that $ff^\star(p) = 1$ if $f(p) \neq 0$. These constructible functions form the set $\mathcal{C}$. Essentially, the idempotent constructible functions encode if-then branching into the formulas for our relations. Now, quantifier elimination and effective Nullstellensatz algorithms can also answer the decision problem, "Given $x, y \in V$, does there exist $g \in G$ such that $g \cdot x = y$?" with slightly improved complexity bounds, but the new algorithm here offers greater algebraic and geometric intuition and applications by producing invariant, separating functions on V.

We proceed in four parts. First, given a matrix X encoding products of the $\sigma_p^*(x_i)$ and encoding $\mathbb{I}(G) = \{f \in k[z_1, \ldots, z_\ell] \mid f(G) = \{0\}\}$, up to some degree d, we produce a matrix of constructible functions that gives the entries of the reduced row echelon form of X, as functions of p. From these entries follow formulas for the kernel vectors of X and hence relations on the $\sigma_p^*(x_i)$. We next establish a degree bound for the relations sufficient to generate the ideal $\mathfrak{q}$ with $\mathbb{V}(\mathfrak{q}) = \overline{G \cdot p}$. By considering a generating set for $\mathfrak{q}$, we provide an algorithm that produces straight line programs for the functions in the set $\mathcal{C}$. We show that these functions separate orbits and have polynomial length as straight line programs in $\widehat{k[V]}$, and we establish polynomial bounds for their number in terms of n and the degree N of the representation.

5.2 Formulas for Reduced Row Echelon Form

5.2.1 Straight Line Programs

To measure the complexity of construcible functions, we adapt the framework of staight line programs over a k-algebra. For a detailed, traditional treatment, see [2].

Let V be a set, F a field, and let R be an F-subalgebra of the F-valued functions on V. Let $A = (a_{-m}, \ldots, a_{-1}) \in (\widehat{R})^m$ be a finite, ordered subset of $\widehat{R}$. Consider a tape of cells with $a_i \in A$ in position i. A *straight line program* Γ is a finite, ordered list of instructions $\Gamma = (\Gamma_0, \ldots, \Gamma_{\ell-1})$. Each instruction Γ_i is of the form $(\star; j, k)$ or $(\star; j)$, where $\star$ is an operation and j, k are positive integers referring to tape entries in positions $i - j$ and $i - k$, that is, j and k cells before i, respectively. The *length* $\ell = |\Gamma|$ measures the complexity of the computation.

To execute Γ on input A, for $i = 0, \ldots, \ell - 1$ write a_i in tape position i as follows:

$$a_i = \begin{cases} a_{i-j} + a_{i-k} & \text{if } \Gamma_i = (+; j, k) \\ a_{i-j} - a_{i-k} & \text{if } \Gamma_i = (-; j, k) \\ a_{i-j} \cdot a_{i-k} & \text{if } \Gamma_i = (\times; j, k) \\ a_{i-j}^{\star} & \text{if } \Gamma_i = (\text{qi}; j) \\ c & \text{if } \Gamma_i = (\text{const}; c) \text{ for } c \in F \\ a_{i-j} & \text{if } \Gamma_i = (\text{recall}; j) \end{cases} \qquad \text{where } j, k < i.$$

The "recall" instruction of position j serves to collect relevant computations at the end of the tape. The traditional defintion of a straight line program in a k-algebra does not include the quasi-inverse "qi" operation, but we include it here to measure the length of the construction of an $f \in \widehat{k[V]}$. Define the *order-d output* of Γ by $\text{Out}_d(\Gamma, A) = (a_{\ell-d}, \ldots, a_{\ell-1}) \in (\widehat{R})^d$, where $\ell = |\Gamma|$. We omit the d where convenient. A straight line program hence defines a constructible function $(\widehat{R})^m \to (\widehat{R})^d$.

Write $\Gamma^{(2)} \circ \Gamma^{(1)}$ for the composition of two straight line programs, in which the input of $\Gamma^{(2)}$ is $\text{Out}_d(\Gamma^{(1)}, A)$ for some d depending on $\Gamma^{(2)}$. Then $\Gamma^{(2)} \circ \Gamma^{(1)}$ has input A, and we execute $\Gamma^{(2)} \circ \Gamma^{(1)}$ by concatenating the instruction lists.

5.2.2 Outline of the Algorithm

Let $A = (a_{ij})$ be an $m \times n$ matrix over a field k. Define the *triangular reduced row echelon form* (tRREF) of A to be the $n \times n$ matrix $R_A = (r_{ij})$ whose jth row $\mathbf{r_j}$ is nonzero if and only if the reduced row echelon form (RREF) of A has a pivot in column j. In that case, $\mathbf{r_j}$ is the row of the RREF of A containing that pivot. For example,

$$\text{RREF(A)} = \begin{pmatrix} 1 & 2 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} \quad \text{corresponds to} \quad \text{tRREF}(A) = \begin{pmatrix} 1 & 2 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

This new form simplifies the identification of pivots: the (usual) RREF of A has a pivot in column j if and only if $r_{jj} = 1$ in the tRREF.

Proposition V.2. *Let (a_{ij}) be an $m \times n$ matrix with entries in any field k. Then there exists a straight line program Γ^{tR} of length $O(mn^2 + n^3)$ such that $Out_{n^2}(\Gamma^{tR}, (a_{ij}))$ are the entries of the triangular RREF of (a_{ij}). The program gives constructible functions for these entries in terms of the a_{ij}.*

The proposition does not require k to be algebraically closed, but we will need this condition for the later geometric reasoning about orbits. Note also that while the classical Gaussian elimination algorithm requires branching, the straight line program Γ^{tR} simulates branching in the computation of the quasi-inverse. The pseudo-code below proves the proposition in general terms; the subsections that follow provide specific constructions.

Algorithm V.3. Let $A = (a_{ij})$ be an $m \times n$ matrix.

1. For $i = 2, \ldots m$, if $a_{11} = 0$, exchange the first row of A with the ith row. After these steps, either $a_{11} \neq 0$, or $a_{i1} = 0$ for all i.

2. Multiply a_{11} by $a_{11}^\star$, and multiply the rest of the first row by $(1 - a_{11}a_{11}^\star + a_{11}^\star)$. This is equivalent to dividing the first row by a_{11} if $a_{11} \neq 0$.

3. For $i = 2, \ldots, m$, subtract $a_{i1} \cdot (a_{11}, \ldots, a_{1n})$ from row i. As a result, $a_{i1} = 0$ for all $i \geq 2$.

4. Let $A' = (a_{ij})_{j \geq 2}$ and $A'' = (a_{ij})_{i,j \geq 2}$, as below:

$$A = \begin{pmatrix} * & \\ 0 & A' \\ \vdots & \\ 0 & \end{pmatrix} = \begin{pmatrix} * & \cdots & * \\ 0 & & \\ \vdots & A'' & \\ 0 & & \end{pmatrix}$$

   Let A_0'' be the $m \times (n-1)$ matrix formed by appending a row of zeros to the bottom of A''; then A' and A_0'' have the same dimensions.

5. Define $B = (1 - a_{11}) \cdot A' + a_{11} \cdot A_0''$.

6. Recursively compute the tRREF of B; call it R_B, an $(n-1) \times (n-1)$ matrix.

7. Let R_A be the $n \times n$ matrix below:

$$R_A = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ 0 & R_B & \\ 0 & & \end{pmatrix}.$$

8. Let $\mathbf{r_k}$ be the kth row of $R_A = (r_{ij})$. For $k = 2, \ldots, n$, subtract $a_{1k} \cdot \mathbf{r_k}$ from the first row of R_A. This reduction produces the triangular RREF of A.

The following formulas specify straight line programs for the entries of the triangular RREF matrix R_A, and hence define Γ^{tR}.

5.2.3 Formulas for Gaussian Elimination

Recall that the first step of the algorithm exchanges the first row of (a_{ij}) with the ith row if $a_{11} = 0$, for $i = 2, \ldots, m$. Hence for an $m \times n$ input matrix X, this step requires $m - 1$ programs E_i such that $Y = \mathrm{Out}_{mn}(E_i, X)$ flips the first and ith rows if necessary. The following formulas describe the entries of $Y = (y_{ij})$:

$$\begin{aligned}
y_{11} &= x_{11} + (1 - x_{11}x_{11}^\star)x_{i1} \\
y_{1j} &= x_{1j} + (1 - x_{11}x_{11}^\star) \cdot (x_{ij} - x_{1j}) \text{ for all } j > 1 \\
y_{i1} &= x_{i1} \cdot x_{11}x_{11}^\star \\
y_{ij} &= x_{1j} + x_{11}x_{11}^\star \cdot (x_{ij} - x_{1j}) \text{ for all } j > 1 \\
y_{kj} &= x_{kj} \text{ for all } k \neq 1, i, \text{ and for all } j.
\end{aligned}$$

For example, the straight line program for y_{11} in E_i takes inputs x_{11} in position -2 and x_{i1} in position -1, and then performs the following steps:

(0) $(\mathrm{qi}; 2)$

(1) $(\times; 3, 1)$

(2) $(\mathrm{const}; 1)$

(3) $(-; 1, 2)$

(4) $(\times; 1, 5)$

(5) $(+; 7, 1)$

The formulas for the other y_{ij} have similarly obvious representations as straight line programs. If we concatenate these programs within E_i, so that all the entries of Y appear in various (known!) positions on the tape, then we can save the recall steps

for the end, and we need only compute $x_{11}^\star$, $x_{11}x_{11}^\star$, $(1-x_{11}x_{11}^\star)$, and $(x_{ij}-x_{1j})$ once. With these efficiencies, the program E_i introduces 1 quasi-inverse, 1 call to k, $3n$ additions, and $2n$ multiplications. Thus the concatenation of $E_2,\ldots,E_{m-1}$ requires $2n(m-1)$ multiplications, $3n(m-1)$ additions, $n-1$ calls to k, $n-1$ quasi-inverses, and mn recalls to collect the entries of Y in the last mn cells of the tape. Call this concatenation Γ^E; we will use it later to collect nonzero rows of a matrix.

Step (2) of the algorithm requires 1 quasi-inverse, 1 subtraction, 1 addition, n multiplications, and n recalls.

These next formulas perform step (3), on an $m\times n$ input matrix (x_{ij}):

$$\begin{aligned} y_{i1} &= 0 \text{ for all } i>1 \\ y_{ij} &= x_{ij}-x_{1j}\cdot x_{i1}\cdot x_{11}x_{11}^\star \text{ for all } i,j>1. \end{aligned}$$

These programs require $(m-1)(n-1)$ additions, $(n-1)(m-1)$ multiplications, and mn recalls. Step (5) next requires 1 subtraction, $m(n-1)$ additions, $2m(n-1)$ multiplications, and $m(n-1)$ recalls.

To perform the reductions in step (8), consider the following formula for r_{1j}, where $j\geq 2$:

$$\begin{aligned} r_{1j} := (1-r_{jj})\cdot(r_{ij}+(\quad &- \quad r_{22}\cdot r_{12}r_{2,j} \\ &- \quad r_{33}\cdot r_{13}r_{i3,j} \\ &- \quad \cdots \\ &- \quad r_{j-1,j-1}\cdot r_{1,j-1}r_{j-1,j}))\,, \end{aligned}$$

This formula sets $r_{1j}=0$ if there is a pivot in column j, that is, if $r_{jj}=1$. Otherwise, the formula subtracts from r_{1j} the effects of clearing columns $<j$. The reduction of r_{1j} requires 1 call to k, j additions/subtractions, $2(j-2)+1$ multiplications (since $j\geq 2$), and n^2 recalls, so reducing the first row has total complexity $O(n^2)$.

The above formulas specify a straight line program Γ^{tR} such that $\mathrm{Out}_{n^2}(\Gamma^{tR}, A)$ are the entries of the tRREF of A. Counting the necessary operations yields asymptotic total complexity estimates for the programs. The recursion on an $m \times t$ matrix has total complexity $O(mt + t^2)$. Summing t from 1 to n yields total complexity $O(mn^2 + n^3)$.

5.2.4 Collecting Nonzero Rows

Lastly, the main algorithm that computes orbit closures requires a program Σ that, given an indicator vector v of 0s and 1s, collects the rows i of a matrix such that the ith entry of v is 1. For example, the diagonal of R_A indicates the nonzero rows of R_A. Given R_A and its diagonal as input, the program Σ would output an $n \times n$ matrix whose first $\mathrm{rank}(A)$ rows include the traditional RREF of A. We will never need to compute the traditional RREF in practice, because the main algorithm runs more efficiently using R_A.

Recall the algorithm Γ^E that exchanges the first row of a matrix X with subsequent rows until the output has $y_{11} \neq 0$, if possible. Define Σ as follows: for an $m \times n$ input matrix X and an indicator m-vector v, form a new matrix X' by adjoining v as a column to the left side of X:

$$X' = \begin{pmatrix} v_1 & x_{11} & \cdots & x_{1n} \\ v_2 & x_{21} & \cdots & x_{2n} \\ \vdots & \vdots & \vdots & \vdots \\ v_m & x_{m1} & \cdots & x_{mn} \end{pmatrix}.$$

After applying Γ^E to X', the first row of X' with $v_i \neq 0$ becomes the first row of the output $Y = (y_{ij})$. Record $\mathbf{r_1} := (y_{12}, \ldots, y_{1,n+1})$ and apply Γ^E to the last $m - 1$ rows of this Y. Let Σ denote this series of m recursions of Γ^E. Since Γ^E applied to

an $s \times (n+1)$ matrix has total complexity $O(sn)$, the procedure Σ has complexity $O(m^2 n)$. Concatenating Σ with the straight line program for the tRREF yields the following:

Corollary V.4. *Let (a_{ij}) be an $m \times n$ matrix with entries in any field k. Then there exists a straight line program of length $O(mn^2 + m^2 n + n^3)$ for the (classical) RREF (r_{ij}) of (a_{ij}). The program gives constructible functions for r_{ij} in terms of the a_{ij}.*

5.2.5 Computing Kernels of Linear Maps

To compute the kernel up to degree d of a k-algebra homomorphism, one can write the homomorphism on elements of degree $\leq d$ as a matrix in RREF. Finding the kernel of a matrix R in RREF is equivalent to solving the system of equations $R \cdot (x_1, \ldots, x_n)^T = 0$: for every pivot r_{ij}, write an equation

$$x_j = -r_{i,j+1}x_{j+1} - r_{i,j+2}x_{j+2} - \cdots - r_{i,n}x_n.$$

Set each free variable equal to 1 in turn, set the other free variables to 0, and read off the vector of values in the pivot variables. These vectors give a basis for the kernel of R, hence of the original map. The basis is canonical because the RREF is canonical.

To compute the kernel of an $m \times n$ matrix A, we use the $n \times n$ matrix R_A containing the rows of the RREF of A: recall there is a pivot in the jth column of the RREF if and only if the row containing that pivot is jth row of $R_A = (r_{ij})$, if and only if $r_{jj} = 1$. Otherwise, $r_{jj} = 0$.

Lemma V.5. *Let R_A be the $n \times n$ tRREF of a matrix A. Then there exists a straight line program Γ^K of length $O(n^2)$ such that $Out_{n^2}(\Gamma^K, R_A)$ gives the kernel of A.*

Proof. Claim that the kernel of A, is given by the following vectors $\phi_1, \ldots, \phi_n$, in

terms of $R_A = (r_{ij})$:

$$\phi_j := (1 - r_{jj}) \cdot (-r_{1j}, -r_{2j}, \ldots, \overbrace{1}^{j\text{th place}}, \ldots, -r_{nj}).$$

Indeed, recall that the kernel of a RREF matrix has one basis vector for each non-pivot column. Namely, $\phi_j = 0$ if and only if column j of the RREF has a pivot. Otherwise, $\phi_j \neq 0$, as follows: Put the free variable $x_j := 1$. Now, $r_{kj} = 0$ unless there is a pivot in column k of the RREF. Set each pivot variable x_{kk} equal to the negation of the jth entry of the row containing that pivot.

Of course, $r_{ij} = 0$ whenever $i > j$, but such simplifications complicate the formulas without improving the asymptotic complexity. As written, each ϕ_j requires 2 calls to k, 1 addition, n scalar multiplications, and n other multiplications. Upon adding recall instructions, computing the kernel has complexity $O(n^2)$. □

5.3 Degree Bounds for Orbit Closures

We relate the degree of a variety to the degrees of polynomials that can define that variety. By bounding the degree of an orbit closure $\overline{G \cdot p}$, we can bound the degree of the defining polynomials.

Lemma V.6. *Let $V = \mathbb{V}(f_1, \ldots, f_r)$ have codimension m in $\mathbb{A}^n$. Then there exist m generic linear combinations $g_i = \sum a_{ij} f_j$ such that*

$$W := \mathbb{V}(g_1, \ldots, g_m) \supseteq V$$

and W has codimension m.

Proof. Induct on the number r of given defining equations for V. The case $r = 1$, implying $m = 1$, is clear. Assume the lemma holds for a variety defined by $r - 1$ equations, and consider $V' = \mathbb{V}(f_1, \ldots, f_{r-1})$. If V' still has codimension m, then the

result follows by the induction hypothesis. Otherwise, V' has codimension $m-1$. By the induction hypothesis, there exist $m-1$ generic linear combinations g_i of $f_1, \ldots, f_{r-1}$ such that $W' = \mathbb{V}(g_1, \ldots, g_{m-1}) \supseteq V'$ and W' has codimension $m-1$.

Since W' is defined by $m-1$ equations, every component Z_k of Y' has codimension $m-1$. It follows that on each Z_k, one of $f_1, \ldots, f_r$ is not identically zero. So for each Z_k, and for every point $p \in Z_k$, we may consider the proper hyperplane $H_{k,p} \subset \mathbb{A}^r$ defined by the vanishing of

$$x_1 f_1(p) + x_2 f_2(p) + \ldots + x_r f_r(p) \in k[x_1, \ldots, x_r].$$

Let $H_k = \cap_{p \in Z_k} H_{k,p}$. Then $\cup_k H_k$ is a closed union of finitely many subspaces of $\mathbb{A}^r$. Thus for any choice of $(a_1, \ldots, a_r)$ in the dense set $\mathbb{A}^r - \cup_k H_k$, the polynomial $g_m = \sum a_i f_i$ is not identically zero on any Z_k. Therefore $Y = \mathbb{V}(g_1, \ldots, g_{m-1}, g_m)$ contains V and has codimension m. □

Let $V \subseteq \mathbb{A}^n$ be an equidimensional affine variety of codimension m. Define the *degree* of V to be

$$\deg(V) = \#H \cap V,$$

where H is a generic linear subspace of dimension m. Heintz proves a stronger version of the following statement in [18], as well as many related results.

Proposition V.7. *Let $V \subseteq \mathbb{A}^n$ be a Zariski closed subset of degree d. Then there exists an ideal $\mathfrak{q}$, generated by polynomials of degree $\leq d$, such that $\sqrt{\mathfrak{q}} = \mathbb{I}(V)$. In particular, $\mathbb{V}(\mathfrak{q}) = V$.*

Proof. It suffices to find, for every point $p \notin V$, a polynomial f of degree $\leq d$ such that f vanishes on V but not at p. If V is a hypersurface, then $V = \mathbb{V}(f)$ with $\deg(f) = \deg(V)$, and we are done. Otherwise, assume V has codimension greater than 1. Without loss of generality, further assume that p is the origin.

To find a polynomial vanishing on V but not at the origin, we project V until an image has codimension 1. Define $\pi\colon \mathbb{A}^n \to \mathbb{P}^{n-1}$ by $\pi\colon (x_1, \ldots, x_n) \mapsto [x_1\colon \ldots : x_n]$. Since $\dim \overline{\pi(V)} \leq \dim V < n-1$, there exists a point $[L] \in \mathbb{P}^{n-1} - \overline{\pi(V)}$. Let $C(V) = \pi^{-1}(\overline{\pi(V)})$, the cone over $\overline{\pi(V)}$. Then $L = \overline{\pi^{-1}([L])}$ has $L \cap \overline{C(V)} = \{0\}$.

Assume without loss of generality that L is the x_n-axis, and consider the projection $\phi\colon \mathbb{A}^n \to \mathbb{A}^{n-1}$ along L, defined by $\phi\colon (x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_{n-1})$. Because $\overline{C(V)}$ is a cone, the restriction of ϕ to $\overline{C(V)}$ is a finite map onto $\mathbb{A}^{n-1}$. In particular, $\phi(V)$ is closed in $\mathbb{A}^{n-1}$. Since L is disjoint from V, $\phi(0) = 0$ remains outside the closed set $\phi(V)$.

Continue projecting until $\phi : \mathbb{A}^n \to \mathbb{A}^{n-m+1}$ gives $\phi(V)$ with codimension 1 (and dimension $\dim V$ after each projection). Now, $\deg(\phi(V)) \leq d$. Thus there exists a polynomial f of degree $\leq d$ such that f vanishes on $\phi(V)$ but $f(0) \neq 0$. Hence $f \circ \phi(V) = 0$ but $f \circ \phi(0) \neq 0$. As ϕ is defined by linear polynomials, the polynomial $f \circ \phi$ has degree $\leq d$, and the result follows. □

Now consider a linear algebraic group G acting on affine n-space. When we can bound the degree of an orbit closure $\overline{G \cdot x}$, then we can produce a degree bound for polynomials f_i such that $\overline{G \cdot x} = \mathbb{V}(f_1, \ldots, f_r)$. For an overview of bounds for the degrees of orbits and the (polynomial) degrees of generating invariants, see [5].

Proposition V.8. *Let G be a linear algebraic group of dimension m, embedded in $\mathbb{A}^\ell$ with ideal $\mathbb{I}(G) = (h_1, \ldots, h_s)$. Set $M = \max\{\deg(h_i)\}$.*

Suppose G acts on $\mathbb{A}^n$ with representation

$$\rho\colon G \to GL_n \quad \textit{defined by} \quad \rho\colon g \mapsto (\rho_{ij}(g)),$$

and set $N = \max\{\deg(\rho_{ij})\}$. If $\overline{G \cdot x}$ is an orbit closure with dimension r, then

$$\deg(\overline{G \cdot x}) \leq N^r M^{\ell - m}.$$

Proof. Let $d = \deg(\overline{G \cdot x})$. For a generic $(n-r)$-dimensional linear subspace $H \subseteq \mathbb{A}^n$, by definition $d = \#(\overline{G \cdot x} \cap H)$. Let $\sigma \colon g \mapsto g \cdot x$ be the orbit map. Then the degrees of the polynomials defining σ are bounded by N. Hence $\sigma^{-1}(H) = \mathbb{V}(u_1, \ldots, u_r) \subseteq G$ has $\deg(u_i) \leq N$ and has $\geq d$ irreducible components.

By the first lemma above, there exist generic linear combinations f_j of the generators of $\mathbb{I}(G)$ such that $\mathbb{V}(f_1, \ldots, f_{\ell-m})$ is a complete intersection and contains G. Thus

$$\sigma^{-1}(H) \subseteq \mathbb{V}\left(u_1, \ldots, u_r, f_1, \ldots, f_{\ell-m}\right) \subset \mathbb{A}^{\ell}.$$

Consider the vanishing of the homogenized polynomials

$$\mathbb{V}\left(\overline{u}_1, \ldots, \overline{u}_r, \overline{f}_1, \ldots, \overline{f}_{\ell-m}\right) \subset \mathbb{P}^{\ell}.$$

By a generalization of Bézout's theorem (see [13], section 12.3.1), the number of irreducible components of this variety is (generously) bounded by

$$\prod_i \deg(\mathbb{V}(\overline{u}_i)) \cdot \prod_j \deg(\mathbb{V}(\overline{f}_j)) = \prod_i \deg(\overline{u}_i) \cdot \prod_j \deg(\overline{f}_j) \leq N^r M^{\ell-m}.$$

This number then also bounds d. □

Corollary V.9. *With the hypotheses of the previous proposition, there exist polynomials $f_1, \ldots, f_t$ such that $\overline{G \cdot x} = \mathbb{V}(f_1, \ldots, f_t)$ and*

$$\deg(f_i) \leq \deg(\overline{G \cdot x}) \leq N^r M^{\ell-m}.$$

5.4 Separating Orbits

Let $\rho : G \hookrightarrow GL_n$ act on $\mathbb{A}^n$ as in Section 3. For $p \in \mathbb{A}^n$, there exists an ideal $\mathfrak{q}$ such that $\mathbb{V}(\mathfrak{q}) = \overline{G \cdot p}$ and $\mathfrak{q}$ is generated in degree $\leq N^r M^{\ell-m}$. We will establish straight line programs for the orbit-separating set $\mathcal{C}$ by considering a generating set

for $\mathfrak{q}$. We prove that these programs define invariant functions separating the orbits of G. The length of these programs will be polynomial in the dimension n and the degree N of the representation.

5.4.1 The Orbit Separating Algorithm

Input the embedding of $G \hookrightarrow \mathbb{A}^\ell$ and the orbit map $\sigma_p \colon g \mapsto g \cdot p$ as above, which varies with p. Let $k[x_1, \ldots, x_n]$ be the coordinate ring of $\mathbb{A}^n$. Then $\ker \sigma_p^* = \mathbb{I}(G \cdot p)$, but to define $\overline{G \cdot p}$ it suffices to compute a k-basis for $\ker \sigma_p^*$ up to degree $N^r M^{\ell - m}$. The elements of this k-basis generate $\mathfrak{q}$ as an ideal.

For each $i = 1, \ldots, N^r M^{\ell-m}$, the following algorithm computes a canonical k-basis for $\ker \sigma_p^*$ in degree $\leq i$, but for each polynomial in the basis the algorithm only outputs constructible functions (of p) that give the non-zero coefficients of monomials apearing in that basis, whatever the monomials may be. Hence the algorithm forgets the generating set of the ideal $\mathfrak{q}$. This forgetting allows the algorithm to have polynomial length as a straight line program, since the number of possible monomials grows exponentially with n.

In the most precise sense, given a point $p \in \mathbb{A}^n$, the following algorithm concatenates straight line programs to output a G-invariant vector $\mathcal{C}$ over k. In fact, each entry of $\mathcal{C}$ is a straight line program in terms of the coordinates of p. Thus the algorithm prescribes a vector $\mathcal{C}$ of G-invariant constructible functions that separate orbits: points in distinct orbits produce distinct vectors. The proofs for the G-invariance and orbit separation will follow.

Choose a monomial order for the monomials spanning $k[z_1, \ldots, z_\ell]$. As a preliminary calculation, compute a Gröbner basis and a k-basis for $\mathbb{I}(G)$ up to degree $N^{r+1} M^{\ell - m}$. Let $B(d)$ denote the set of elements of the k-basis up to degree d. Also, for a vector w, let $\pi_t(w)$ denote the vector of the the first t entries of w.

Lastly, since all computations occur in $k[G]$, we must predict the dimension of $k[G]_{\leq d}$.

Lemma V.10. *Let $m = \dim G$. There exists a function $H(d)$, computable from a Gröbner basis for $\mathbb{I}(G)$, such that $H(d) = \dim_k k[G]_{\leq d}$ for all $d \geq 0$, and $H(d) \leq O(d^m)$.*

Proof. Suppose $R = k[G]$ is generated as a k-algebra by $f_1, \ldots, f_r$ of degree 1. Define $S = k[f_1 t, \ldots, f_r t, t] \subseteq R[t]$, and claim $S_d = R_{\leq d} \cdot t^d$, where S is graded by t-degree. The inclusion $\supseteq$ is clear, and if $h \in S_d$ is a homogeneous polynomial in t, then the coefficients of t^d can have R-degree no greater than d (less, for example, in the term $f_1 t \cdot t^{d-1}$). Let $H(d)$ be the dth coefficient of the Hilbert series of S, which we may compute from a Gröbner basis for $\mathbb{I}(G)$. Then $H(d) = \dim_k R_{\leq d}$. Since S has dimension bounded by $m + 1$, the Hilbert polynomial for S has degree bounded by m. Thus $H(d) \leq O(d^m)$. □

Algorithm V.11.

1. For $j = 1, \ldots, n$, let v_j be the vector of coefficients of $\sigma_p^*(x_j)$ with respect to the (ordered) monomial basis of $k[z_1 \ldots, z_\ell]$.
2. $V_1 := (v_1, \ldots, v_n)$.
3. $i := 1$, $\mathcal{C}_0 = \emptyset$.
4. Put the vectors of $V_i = (v_1, \ldots, v_{k_i})$, in order, in the first k_i columns of a matrix X_i; fill subsequent columns with $B(iN)$.
5. Compute $\mathrm{Out}(\Gamma^{tR}, X_i)$, the tRREF of X_i.
6. Compute $\beta := \mathrm{Out}(\Gamma^K, \mathrm{Out}(\Gamma^{tR}, X_i))$, a basis for $\ker X_i$.

7. Let $\mathcal{C}_i := \mathcal{C}_{i-1} \cup \{\pi_{k_i}(v) \mid v \in \beta\}$.

8. IF $N^r M^{\ell-m} = i$, THEN output $\mathcal{C} = \mathcal{C}_i$, and STOP.

9. Let Y be the matrix whose rows are the vectors in V_i. Let D be the first k_i entries on the diagonal of the tRREF X_i.

10. Compute $Y' := \mathrm{Out}(\Sigma, \{Y, D\})$, the rows of Y indicated by D.

11. Let L_i be the first $H(i)$ rows of Y'.

12. IF $k_i = \#(\text{rows of Y}) < H(i)$, THEN pad L_i with zeros so that L_i has precisely $H(i)$ vectors.

13. $V_{i+1} := L_i \cup \left(\{\sigma_p^*(x_1), \ldots, \sigma_p^*(x_n)\} \cdot \{v_j \in L_i \mid j > H(i-1).\}\right)$.

14. $i := i + 1$.

15. GOTO (4).

The final steps of each iteration require some remarks. For step (10), recall that the nonzero entries of the diagonal D of the tRREF of X_i indicate which columns of X_i are linearly independent. These are the image vectors the algorithm should preserve for the next iteration, so that it can proceed with a polynomial number of multiplications. In step (13), we multiply the $\sigma_p^*(x_i)$ only by these newfound vectors.

Step (12) can be accomplished in the context of straight line programs because we can predict the iteration i at which $k_i \geq H(i)$ first occurs, independent of the choice of p. At step (13) we multiply L_i by all $\sigma_p^*(x_i)$ because, in principle, all $\sigma_p^*(x_i)$ could be linearly independent modulo $\mathbb{I}(G)$. As i increases, the vectors in each V_i describe the images of larger monomials x^I, I a multi-index, in $k[x_1, \ldots, x_n]$. The algorithm terminates when we have considered a k-basis for the polynomials of degree up to

$N^r M^{\ell-m}$ that vanish on $\overline{G \cdot p}$. By the previous section, the elements of that k-basis generate an ideal whose radical is $\mathbb{I}(\overline{G \cdot p})$.

Proposition V.12. *The constructible functions defined by the set $\mathcal{C}$*

1. *are constant on the orbit of $p \in \mathbb{A}^n$, and hence invariant under the usual action $g \cdot f(x) = f(g^{-1} \cdot x)$ for $g \in G$,*
2. *separate orbits.*

Proof. To show that the functions defined by the straight line programs in $\mathcal{C}$ are invariant, choose $p \in \mathbb{A}^n$ and $q \in G \cdot p$. Let $X_i(p)$ be the matrix produced in step (4) of the algorithm in the ith iteration. Let $X_i^V(p)$ be the first $|V_i| = k_i$ columns of $X_i(p)$, that is, those containing the vectors in $V_i(p)$. Now, $X_1^V(p)$ and $X_1^V(q)$ have the same kernel, because (a) as maps $k[x_1, \ldots, x_n]_1 \to k[G]_{\leq N}$ they have the same basis $x_1, \ldots, x_n$ for their domain, and because (b) the kernel of each matrix must span $\mathbb{I}(G \cdot p)_1$. Thus $X_1^V(g \cdot p) = A \cdot X_1^V(p)$ for some matrix A. In particular, $X_1^V(p)$ and $X_1^V(q)$ have linearly independent columns in the same places, and hence have the same RREF.

So letting $\mathcal{C}_i(x)$ denote the kernel vectors obtained on input x in the ith iteration, we have $\mathcal{C}_1(p) = \mathcal{C}_1(g \cdot p)$. As well, let $L_i(p)$ denote the set (produced in step (11) of the algorithm) containing the linearly independent columns of $X_i^V(p)$. Then we have $L_1(p) = \{\sigma_p^*(x_{j_1}), \ldots, \sigma_p^*(x_{j_r})\}$ and $L_1(g \cdot p) = \{\sigma_{g \cdot p}^*(x_{j_1}), \ldots, \sigma_{g \cdot p}^*(x_{j_r})\}$ for the same indices $j_1, \ldots, j_s$.

Proceed by induction on i: we may assume $X_i^V(p)$ and $X_i^V(q)$ have the same RREF and hence $\mathcal{C}_i(p) = \mathcal{C}_i(q)$. We may also assume the columns of $X_i^V(p)$ and $X_i^V(q)$ represent the images of the same set of monomials $\{x^{I_1}, \ldots, x^{I_s}\}$, for multi-indicies I_j. Then the lists $V_{i+1}(p)$ and $V_{i+1}(q)$ also represent the images of the same

monomials under σ_p^* and σ_q^*, respectively. Claim again that $X_{i+1}^V(p)$ and $X_{i+1}^V(q)$ have the same RREF. By the induction hypothesis, the two matrices have the same basis for their domain, and the kernel of each must span $I(G \cdot p)_{i+1}$. These facts prove the claim, as in the base case. Thus $\mathcal{C}_{i+1}(p) = \mathcal{C}_{i+1}(q)$, and the functions in $\mathcal{C}$ are invariant.

To show the functions in $\mathcal{C}$ separate orbits, choose $p, q \in \mathbb{A}^n$ such that the functions in $\mathcal{C}$ take the same values at both points. In particular, $\mathcal{C}_1(p) = \mathcal{C}_1(q)$, so $X_1(p)$ and $X_1(q)$ have the same canonical kernel. As above, it follows that $X_1(p)$ and $X_1(q)$ have the same RREF. Two facts emerge. Crucially, the kernels of σ_p^* and σ_q^* have the same canonical k-basis for their subspaces of degree-1 elements, because the matrices $X_1^V(p)$ and $X_1^V(q)$ assume the same basis for the domain space $k[x_1, \ldots, x_n]_1$, namely, $x_1, \ldots, x_n$. We wish to prove this for all degrees i.

What is more, $L_1(p) = \{\sigma_p^*(x_{j_1}), \ldots, \sigma_p^*(x_{j_s})\}$ and $L_1(q) = \{\sigma_q^*(x_{j_1}), \ldots, \sigma_q^*(x_{j_s})\}$ for the same indices $j_1, \ldots, j_s$, because $X_1^V(p)$ and $X_1^V(q)$ have linearly independent columns in the same positions. Thus $V_2(p)$ and $V_2(q)$ list the images of the same set of monomials $x_j x_k$ under σ_p^* and σ_q^*, respectively.

Proceeding by induction, if $X_i^V(p)$ and $X_i^V(q)$ have the same RREF and list the images of the same monomials, then $X_{i+1}^V(p)$ and $X_{i+1}^V(q)$ also list the images of the same monomials. By the assumption $C_{i+1}(p) = C_{i+1}(q)$, the matrices $X_{i+1}^V(p)$ and $X_{i+1}^V(q)$ also have the same RREF. Therefore the kernels of σ_p^* and σ_q^* have the same canonical k-basis for their degree-i subspaces, completing the induction. In particular, the same ideal $(f_1, \ldots, f_s)$ defines $\overline{G \cdot p}$ and $\overline{G \cdot q}$. Since G is a linear algebraic group, it follows $G \cdot p = G \cdot q$, completing the proof. $\square$

5.4.2 Complexity Bounds

The bookkeeping that follows confirms that the complexity of the orbit separating algorithm is polynomial in n and N. First, the degree bound $N^r M^{\ell-m}$ for a generating set of $\mathfrak{q}$ requires that we compute products of $N^r M^{\ell-m}$ degree-N polynomials f_i in $k[z_1, \ldots, z_\ell]$, for $i = 1, \ldots, N^r M^{\ell-m}$. To this end, compute the monomials in the z_j up to degree $N \cdot N^r M^{\ell-m}$, with total complexity $O(N^{\ell(r+1)} M^{\ell(\ell-m)})$. Then multiply $f_1 f_2 \cdots f_i$ and f_{i+1} to obtain an implicit straight-line program for the product of $i+1$ distinct degree-N polynomials in $k[z_1, \ldots, z_\ell]$, with complexity $O(2^{2\ell-2} i^{2\ell} N^{2\ell})$. For details of polynomial multiplication, see Chapter 2 of [2].

Next consider the sizes of matrices in the algorithm. Recall that for large d, $H(d) \leq O(d^m)$. Hence in iteration i, the matrix X_i has

$$k_i = O\left(((i-1)N)^m + n \cdot [((i-1)N)^m - ((i-2)N)^m]\right)$$

columns from V_i, has $|B(iN)|$ additional columns, and has $(iN)^\ell$ rows corresponding to the monomials in $k[z_1, \ldots, z_\ell]_{\leq iN}$. Of course, $|B(iN)| = O((iN)^\ell)$, so the number of rows of X_i is $O((iN)^\ell)$, and the number of columns is $O(n(iN)^m + (iN)^\ell) \leq O(n(iN)^\ell)$. Now, computing the tRREF of an $s \times t$ matrix has complexity $O(st^2 + t^3)$. Thus the computation of tRREF(X_i) has complexity bounded by

$$O\left((iN)^\ell \cdot n^2 (iN)^{2\ell} + n^3 (iN)^{3\ell}\right) = O\left(n^3 i^{3\ell} N^{3\ell}\right).$$

The above count of the columns of X_i also yields that the computation of the kernel of tRREF(X_i) has complexity $O(n^2 i^{2\ell} N^{2\ell})$

In collecting the independent elements of V_i in step (10), the input to the procedure Σ is a $k_i \times (iN)^\ell$ matrix, where

$$k_i = O\left(((i-1)N)^m + n \cdot [((i-1)N)^m - ((i-2)N)^m]\right) \leq O(n(iN)^m).$$

On an $s \times t$ matrix, Σ has complexity $O(s^2 t)$, whence step (10) has complexity $\leq O(n^2 (iN)^{2m} \cdot (iN)^{\ell})$.

Finally, the polynomial multiplications $f_1 \cdots f_i$ proceed through $i = N^r M^{\ell - m}$, with n multiplications for each i. Their total complexity is

$$O\left(2^{2\ell-2} n (N^r M^{\ell-m})^{2\ell+1} N^{2\ell}\right) = O\left(2^{\ell-1} n N^{2\ell(r+1)+r} M^{(\ell-m)(2\ell+1)}\right).$$

Of the other computations, the programs for the tRREF have the highest cost. Summing their complexity from $i = 1$ to the degree bound, $N^r M^{\ell-m}$, yields the following:

$$O\left(n^3 (N^r M^{\ell-m})^{3\ell+1} N^{3\ell}\right) = O\left(n^3 N^{3\ell(r+1)+r} M^{(\ell-m)(3\ell+1)}\right),$$

where, again, N is the maximum polynomial degree of the representation, M is a degree bound for a generating set of $\mathbb{I}(G) \subset k[z_1, \ldots, z_\ell]$, and under this embedding G has dimension m. Since the embedding $G \hookrightarrow \mathbb{A}^\ell$ is fixed, we omit the constant power of M from the asymptotic complexity.

Finally, to bound the number of relations that the algorithm computes, we sum the column count $O(n(iN)^\ell)$ of the matrices X_i over all iterations i, and obtain

$$O\left(n N^{\ell(r+1)+r} M^{(\ell-m)(\ell+1)}\right)$$

polynomials generating the ideal $\mathfrak{q}$. In iteration i, such a polynomial has $k_i \leq O(n(iN)^m)$ terms, giving a bound for the number of constructible functions that the algorithm computes:

$$O\left(n^2 N^{(\ell+m+1)(r+1)} M^{(\ell-m)(\ell+m+1)}\right).$$

By omitting the powers of M, the main theorem follows.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] Avner Ash. Personal communication, 2011.

[2] Peter Bürgisser, Michael Clausen, and M. Amin Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1997. With the collaboration of Thomas Lickteig.

[3] H. E. A. Campbell, A. V. Geramita, I. P. Hughes, G. G. Smith, and D. L. Wehlau. Some remarks on Hilbert functions of Veronese algebras. *Comm. Algebra*, 28(3):1487–1496, 2000.

[4] Harm Derksen. Computation of invariants for reductive groups. *Adv. Math.*, 141(2):366–384, 1999.

[5] Harm Derksen. Polynomial bounds for rings of invariants. *Proc. Amer. Math. Soc.*, 129(4):955–963 (electronic), 2001.

[6] Harm Derksen and Gregor Kemper. *Computational invariant theory*. Invariant Theory and Algebraic Transformation Groups, I. Springer-Verlag, Berlin, 2002. Encyclopaedia of Mathematical Sciences, 130.

[7] Harm Derksen and Gregor Kemper. Computing invariants of algebraic groups in arbitrary characteristic. *Adv. Math.*, 217(5):2089–2129, 2008.

[8] Harm Derksen and Hanspeter Kraft. Constructive invariant theory. In *Algèbre non commutative, groupes quantiques et invariants (Reims, 1995)*, volume 2 of *Sémin. Congr.*, pages 221–244. Soc. Math. France, Paris, 1997.

[9] M. Domokos. Typical separating invariants. *Transform. Groups*, 12(1):49–63, 2007.

[10] T. Dubé, B. Mishra, and C. Yap. *Complexity of Buchberger's algorithm for Grobner bases*, 1995. `www.cs.nyu.edu/mishra/PUBLICATIONS/89.dmy-1.ps`.

[11] Peter Fleischmann. The Noether bound in invariant theory of finite groups. *Adv. Math.*, 156(1):23–32, 2000.

[12] John Fogarty. On Noether's bound for polynomial invariants of a finite group. *Electron. Res. Announc. Amer. Math. Soc.*, 7:5–7 (electronic), 2001.

[13] William Fulton. *Intersection theory*, volume 2 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1984.

[14] Frank Grosshans. Observable groups and Hilbert's fourteenth problem. *Amer. J. Math.*, 95:229–253, 1973.

[15] W. J. Haboush. Reductive groups are geometrically reductive. *Ann. of Math. (2)*, 102(1):67–83, 1975.

[16] G. H. Hardy and S. Ramanujan. Asymptotic formulæ in combinatory analysis [Proc. London Math. Soc. (2) **17** (1918), 75–115]. In *Collected papers of Srinivasa Ramanujan*, pages 276–309. AMS Chelsea Publ., Providence, RI, 2000.

[17] John C. Harris and David L. Wehlau. Non-negative integer linear congruences. *Indag. Math. (N.S.)*, 17(1):37–44, 2006.

[18] Joos Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoret. Comput. Sci.*, 24(3):239–277, 1983.

[19] David Hilbert. Ueber die Theorie der algebraischen Formen. *Math. Ann.*, 36(4):473–534, 1890.

[20] David Hilbert. Ueber die vollen Invariantensysteme. *Math. Ann.*, 42(3):313–373, 1893.

[21] K. Hiss. *Constructive invariant theory for reductive algebraic groups*. PhD thesis, Brandeis University, Waltham, 1996.

[22] M. Hochster. Prime ideal structure in commutative rings. *Trans. Amer. Math. Soc.*, 142:43–60, 1969.

[23] M. Hochster and John A. Eagon. Cohen-Macaulay rings, invariant theory, and the generic perfection of determinantal loci. *Amer. J. Math.*, 93:1020–1058, 1971.

[24] Melvin Hochster and Joel L. Roberts. Rings of invariants of reductive groups acting on regular rings are Cohen-Macaulay. *Advances in Math.*, 13:115–175, 1974.

[25] Roger Howe. *The classical groups* and invariants of binary forms. In *The mathematical heritage of Hermann Weyl (Durham, NC, 1987)*, volume 48 of *Proc. Sympos. Pure Math.*, pages 133–166. Amer. Math. Soc., Providen, 1987.

[26] Roger Howe. Personal communication, 2011.

[27] James E. Humphreys. *Linear algebraic groups*. Springer-Verlag, New York, 1975. Graduate Texts in Mathematics, No. 21.

[28] Camille Jordan. Mémoire sur les covariants des formes binaires. *J. de Math.*, 3(2):177–232, 1876.

[29] Camille Jordan. Sur les covariants des formes binaires. *J. de Math.*, 3(5):345–378, 1879.

[30] Victor G. Kac. Root systems, representations of quivers and invariant theory. In *Invariant theory (Montecatini, 1982)*, volume 996 of *Lecture Notes in Math.*, pages 74–108. Springer, Berlin, 1983.

[31] B. Ya. Kazarnovskiĭ. Newton polyhedra and Bezout's formula for matrix functions of finite-dimensional representations. *Funktsional. Anal. i Prilozhen.*, 21(4):73–74, 1987.

[32] Gregor Kemper. Computational invariant theory. In *The Curves Seminar at Queen's. Vol. XII (Kingston, ON, 1998)*, volume 114 of *Queen's Papers in Pure and Appl. Math.*, pages 5–26. Queen's Univ., Kingston, ON, 1998.

[33] Gregor Kemper. Computing invariants of reductive groups in positive characteristic. *Transform. Groups*, 8(2):159–176, 2003.

[34] Gregor Kemper. Separating invariants. *J. Symbolic Comput.*, 44(9):1212–1222, 2009.

[35] George R. Kempf. Instability in invariant theory. *Ann. of Math. (2)*, 108(2):299–316, 1978.

[36] H. Kraft and Jerzy Weyman. *Degree bounds for invariants and covariants of binary forms*, 1999. `http://jones.math.unibas.ch/~kraft/Papers/KWJordan.pdf`.

[37] Yozô Matsushima. Espaces homogènes de Stein des groupes de Lie complexes. *Nagoya Math. J*, 16:205–218, 1960.

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

[39] Masayoshi Nagata. On the 14-th problem of Hilbert. *Amer. J. Math.*, 81:766–772, 1959.

[40] Masayoshi Nagata. Invariants of a group in an affine ring. *J. Math. Kyoto Univ.*, 3:369–377, 1963/1964.

[41] Masayoshi Nagata and Takehiko Miyata. Note on semi-reductive groups. *J. Math. Kyoto Univ.*, 3:379–382, 1963/1964.

[42] Emmy Noether. Der Endlichkeitssatz der Invarianten endlicher Gruppen. *Math. Ann.*, 77(1):89–92, 1915.

[43] John E. Olson. Sums of sets of group elements. *Acta Arith.*, 28(2):147–156, 1975/76.

[44] Bernd Sturmfels. *Algorithms in invariant theory*. Texts and Monographs in Symbolic Computation. SpringerWienNewYork, Vienna, first edition, 1993.

[45] J. V. Uspensky. Asymptotic formulae for numerical functions which occur in the theory of partitions. *Bull. Acad. Sci. URSS*, 14:199–218, 1920.

[46] David L. Wehlau. Constructive invariant theory for tori. *Ann. Inst. Fourier (Grenoble)*, 43(4):1055–1066, 1993.