Geometric Complexity Theory: Introduction

Dedicated to Sri Ramakrishna

Ketan D. Mulmuley [1]
The University of Chicago

Milind Sohoni
I.I.T., Mumbai

Technical Report TR-2007-16
Computer Science Department
The University of Chicago
September 2007

October 16, 2007

[1] Part of the work on GCT was done while the first author was visiting I.I.T. Mumbai to which he is grateful for its hospitality

Foreword

These are lectures notes for the introductory graduate courses on geometric complexity theory (GCT) in the computer science department, the university of Chicago. Part I consists of the lecture notes for the course given by the first author in the spring quarter, 2007. It gives introduction to the basic structure of GCT. Part II consists of the lecture notes for the course given by the second author in the spring quarter, 2003. It gives introduction to invariant theory with a view towards GCT. No background in algebraic geometry or representation theory is assumed. These lecture notes in conjunction with the article [GCTflip1], which describes in detail the basic plan of GCT based on the principle called the *flip*, should provide a high level picture of GCT assuming familiarity with only basic notions of algebra, such as groups, rings, fields etc. Many of the theorems in these lecture notes are stated without proofs, but after giving enough motivation so that they can be taken on faith. For the readers interested in further study, Figure 1 shows logical dependence among the various papers of GCT. The papers need not be read as per the logical dependence. Specifically, it is suggested that the reader without a prior background in algebraic geometry or representation theory read these notes first before reading the basic plan of GCT in [GCTflip1], because these focus on a part of GCT that is easiest to understand without such familiarity but which still gives an idea of the whole, and then go through it in a leisurely fashion.

The first author is grateful to Paolo Codenotti, Joshua Grochow, Sourav Chakraborty and Hari Narayanan for taking notes for his lectures.

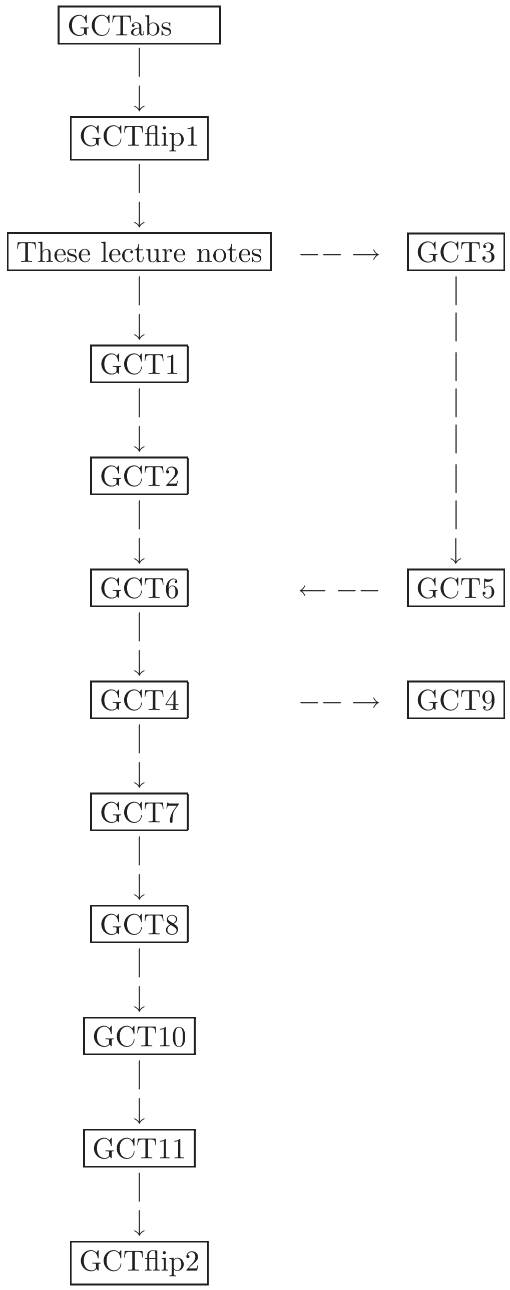

Figure 1: Logical dependence among the GCT papers

Contents

Part I

The basic structure of GCT

By Ketan D. Mulmuley

Chapter 1

Overview

Scribe: Joshua A. Grochow

Goal: An overview of GCT.

The purpose of this course is to give an introduction to Geometric Complexity Theory (GCT), which is an approach to proving **P** $\neq$ **NP** via algebraic geometry and representation theory. A basic plan of this approach is described in [GCTflip1, GCTflip2]. It is partially implemented in a series of articles [GCT1]-[GCT11]. The paper [GCTconf] is a conference announcement of GCT. The paper [Ml] gives an unconditional lower bound in a PRAM model without bit operations based on elementary algebraic geometry, and was a starting point for the GCT investigation via algebraic geometry.

The only mathematical prerequisites for this course are a basic knowledge of abstract algebra (groups, ring, fields, etc.) and a knowledge of computational complexity. In the first month we plan to cover the representation theory of finite groups, the symmetric group S_n, and $GL_n(\mathbb{C})$, and enough algebraic geometry so that in the remaining lectures we can cover basic GCT. Most of the background results will only be sketched or omitted.

This lecture uses slightly more algebraic geometry and representation theory than the reader is assumed to know in order to give a more complete picture of GCT. As the course continues, we will cover this material.

1.1 Outline

Here is an outline of the GCT approach. Consider the **P** vs. **NP** question in characteristic 0; i.e., over integers. So bit operations are not allowed, and basic operations on integers are considered to take constant time. For a similar approach in nonzero characteristic (characteristic 2 being the classical case from a computational complexity point of view), see GCT 11.

The basic principle of GCT is the called the *flip* [GCTflip1]. It "reduces" (in essence, not formally) the lower bound problems such as **P** vs. **NP** in characteristic 0 to upper bound problems: showing that certain decision problems in algebraic geometry and representation theory belong to P. Each of these decision problems is of the form: is a given (nonnegative) structural constant associated to some algebro-geometric or representation theoretic object nonzero? This is akin to the decision problem: given a matrix, is its permanent nonzero? (We know how to solve this particular problem in polynomial time via reduction to the perfect matching problem.)

Next, the preceding upper bound problems are reduced to purely mathematical positivity hypotheses [GCT6]. The goal is to show that these and other auxilliary structural constants have positive formulae. By a positive formula we mean a formula that does not involve any alternating signs like the usual positive formula for the permanent; in contrast the usual formula for the determinant involves alternating signs.

Finally, these positivity hypotheses are "reduced" to conjectures in the theory of quantum groups [GCT6, GCT7, GCT8, GCT10] intimately related to the Riemann hypothesis over finite fields proved in [Dl2], and the related works [BBD, KL2, Lu2]. A pictorial summary of the GCT approach is shown in Figure 1.1, where the arrows represent reductions, rather than implications.

To recap: we move from a negative hypothesis in complexity theory (that there does not exist a polynomial time algorithm for an **NP**-complete problem) to a positive hypotheses in complexity theory (that there exist polynomial-time algorithms for certain decision problems) to positive hypotheses in mathematics (that certain structural constants have positive formulae) to conjectures on quantum groups related to the Riemann hypothesis over finite fields, the related works and their possible extensions. The first reduction here is the *flip*: we reduce a question about lower bounds, which are notoriously difficult, to the one about upper bounds, which we have a much better handle on. This flip from negative to positive is already present in Gödel's work: to show something is impossible it suffices to show that something else is possible. This was one of the motivations for the GCT approach. The Gödelian flip would not work for the **P** vs. **NP** problem because it relativizes. We can think of GCT as a form of nonrelativizable (and non-naturalizable, if reader knows what that means) diagonalization.

In summary, this approach very roughly "reduces" the lower bound prob-

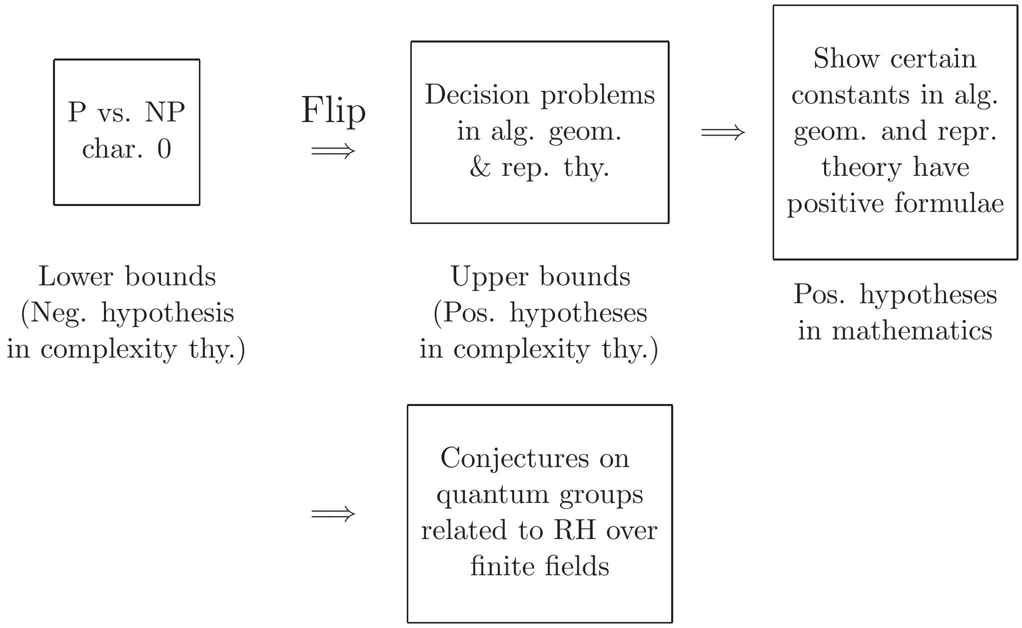

Figure 1.1: The basic approach of GCT

lems such as **P** vs. **NP** in characteristic zero to as-yet-unproved quantum-group-conjectures related to the Riemann Hypothesis over finite fields. As with the classical RH, there is experimental evidence to suggest these conjectures hold – which indirectly suggests that certain generalizations of the Riemann hypothesis over finite fields also hold – and there are hints on how the problem might be attacked. See [GCTflip1, GCT6, GCT7, GCT8] for a more detailed exposition.

1.2 The Gödelian Flip

We now re-visit Gödel's original flip in modern language to get the flavor of the GCT flip.

Gödel set out to answer the question:

Q: Is truth provable?

But what "truth" and "provable" means here is not so obvious *a priori*. We start by setting the stage: in any mathematical theory, we have the syntax (i.e. the language used) and the semantics (the domain of discussion). In this case, we have:

Syntax (language)	Semantics (domain)
First order logic	
$(\forall, \exists, \neg, \vee, \wedge, \ldots)$	
Constants	0,1
Variables	$x, y, z, \ldots$
Basic Predicates	$>, <, =$
Functions	$+,-,\times$,exponentiation
Axioms	Axioms of the natural numbers $\mathbb{N}$
	Universe: $\mathbb{N}$

A sentence is a valid formula with all variables quantified, and by a *truth* we mean a sentence that is true in the domain. By a proof we mean a valid deduction based on standard rules of inference and the axioms of the domain, whose final result is the desired statement.

Hilbert's program asked for an algorithm that, given a sentence in number theory, decides whether it is true or false. A special case of this is Hilbert's 10th problem, which asked for an algorithm to decide whether a Diophantine equation (equation with only integer coefficients) has a nonzero integer solution. Gödel showed that Hilbert's general program was not achievable. The tenth problem remained unresolved until 1970, at which point Matiyasevich showed its impossibility as well.

Here is the main idea of Gödel's proof, re-cast in modern language. For a Turing Machine M, whether the empty string ε is in the language $L(M)$ recognized by M is undecidable. The idea is to reduce a question of the form $\varepsilon \in L(M)$ to a question in number theory. If there were an algorithm for deciding the truth of number-theoretic statements, it would give an algorithm for the above Turing machine problem, which we know does not exist.

The basic idea of the reduction is similar to the one in Cook's proof that SAT is **NP**-complete. Namely, $\varepsilon \in L(M)$ iff there is a valid computation of M which accepts ε. Using Cook's idea, we can use this to get a Boolean formula:

$$\begin{array}{l}\exists m \exists \text{ a valid computation of } M \text{ with configurations of size } \leq m \\ \qquad \text{s.t. the computation accepts } \varepsilon.\end{array}$$

Then we use Gödel numbering – which assigns a unique number to each sentence in number theory – to translate this formula to a sentence in number theory. The details of this should be familiar.

The key point here is: to show that truth is undecidable in number theory (a negative statement), we show that there exists a computable reduction from $\varepsilon \stackrel{?}{\in} L(M)$ to number theory (a positive statement). This is the essence of the Gödelian flip, which is analogous to – and in fact was the original motivation for – the GCT flip.

1.3 More details of the GCT approach

To begin with, GCT associates to each complexity class such as **P** and **NP** a projective algebraic variety χ_P, χ_{NP}, etc. [GCT1]. In fact, it associates a family of varieties $\chi_{NP}(n, m)$: one for each input length n and circuit size m, but for simplicity we suppress this here. The languages L in the associated complexity class will be points on these varieties, and the set of such points is dense in the variety. These varieties are thus called *class varieties*. To show that $\mathbf{NP} \nsubseteq \mathbf{P}$ in characteristic zero, it suffices to show that χ_{NP} cannot be imbedded in χ_P.

These class varieties are in fact G-varieties. That is, they have an action of the group $G = GL_n(\mathbb{C})$ on them. This action induces an action on the homogeneous coordinate ring of the variety, given by $(\sigma f)(\mathbf{x}) = f(\sigma^{-1}\mathbf{x})$ for all $\sigma \in G$. Thus the coordinate rings R_P and R_{NP} of χ_P and χ_{NP} are

G-algebras, i.e., algebras with G-action. Their degree d-components $R_P(d)$ and $R_{NP}(d)$ are thus finite dimensional G-representations.

For the sake of contradiction, suppose $\mathbf{NP} \subseteq \mathbf{P}$ in characteristic 0. Then there must be an embedding of χ_{NP} into χ_P as a G-subvariety, which in turn gives rise (by standard algebraic geometry arguments) to a surjection $R_P \twoheadrightarrow R_{NP}$ of the coordinate rings. This implies (by standard representation-theoretic arguments) that $R_{NP}(d)$ can be embedded as a G-sub-representation of $R_P(d)$. The following diagram summarizes the implications.

complexity classes	class varieties	coordinate rings	representations of $GL_n(\mathbb{C})$

$$
\begin{tikzcd}
NP \arrow[r, rightsquigarrow] \arrow[d, hook] & \chi_{NP} \arrow[r, rightsquigarrow] \arrow[d, hook] & R_{NP} \arrow[r, rightsquigarrow] & R_{NP}(d) \arrow[d, hook] \\
P \arrow[r, rightsquigarrow] & \chi_P \arrow[r, rightsquigarrow] & R_P \arrow[r, rightsquigarrow] \arrow[u, two heads] & R_P(d)
\end{tikzcd}
$$

Weyl's theorem–that all finite-dimensional representations of $G = GL_n(\mathbb{C})$ are completely reducible, i.e. can be written as a direct sum of irreducible representations–implies that both $R_{NP}(d)$ and $R_P(d)$ can be written as direct sums of irreducible G-representations. An *obstruction* [GCT2] of degree d is defined to be an irreducible G-representation occuring (as a subrepresentation) in $R_{NP}(d)$ but not in $R_P(d)$. Its existence implies that $R_{NP}(d)$ cannot be embedded as a subrepresentation of $R_P(d)$, and hence, χ_{NP} cannot be embedded in χ_P as a G-subvariety; a contradiction.

We actually have a *family* of varieties $\chi_{NP}(n, m)$: one for each input length n and circuit size m. Thus if an obstruction of some degree exists for all $n \to \infty$, assuming $m = n^{\log n}$ (say), then $\mathbf{NP} \neq \mathbf{P}$ in characteristic zero.

Conjecture 1.1. *[GCTflip1] There is a polynomial-time, i.e., poly(n, m) time algorithm for deciding if there exists an obstruction for given n and m, and for constructing one if it exists–in particular, this is so if $m =$ poly(n) or $O(n^{\log n})$, say.*

This is the GCT flip: to show that no polynomial-time algorithm exists for an **NP**-complete problem, we hope to show that there is a polynomial time algorithm for finding obstructions. This task then is further reduced to finding polynomial time algorithms for other decision problems in algebraic geometry and representation theory.

Mere existence of an obstruction for all n would actually suffice here. For this, it suffices to show that there is an algorithm which, given n, outputs an obstruction showing that $\chi_{NP}(n,m)$ cannot be imbedded in $\chi_P(n,m)$, when $m = n^{\log n}$. But the conjecture is not just that there is an algorithm, but that there is a polynomial-time algorithm.

The basic principle here is that the complexity of the proof of existence of an object (in this case, an obstruction) is very closed tied to the computational complexity of finding that object, and hence, techniques underneath an easy (i.e. polynomial time) time algorithm for deciding existence may yield an easy (i.e. feasible) proof of existence. This is supported by much anecdotal evidence:

- An obstruction to planar embedding (a forbidden Kurotowski minor) can be found in polynomial, in fact, linear time by variants of the usual planarity testing algorithms, and the underlying techniques, in retrospect, yield an algorithmic proof of Kurotowski's theorem that every nonplanar graph contains a forbidden minor.

- Hall's marriage theorem, which characterizes the existence of perfect matchings, in retrospect, follows from the techniques underlying polynomial-time algorithms for finding perfect matchings.

- The proof that a graph is Eulerian iff all vertices have even degree is, essentially, a polynomial-time algorithm for finding an Eulerian circuit.

- In contrast, we know of no Hall-type theorem for Hamiltonians paths, essentially, because finding such a path is computationally difficult (**NP**-complete).

Analogously the goal is to find a polynomial time algorithm for deciding if there exists an obstruction for given n and m, and then use the underlying techniques to show that an obstruction always exists for every large enough n if $m = n^{\log n}$. The main mathematical work in GCT takes steps towards this goal.

Chapter 2

Representation theory of reductive groups

Scribe: Paolo Codenotti

Goal: Basic notions in representation theory.
References: [FH, F]

In this lecture we review the basic representation theory of reductive groups as needed in this course. Most of the proofs will be omitted, or just sketched. For complete proofs, see the books by Fulton and Harris, and Fulton [FH, F]. The underlying field throughout this course is $\mathbb{C}$.

2.1 Basics of Representation Theory

2.1.1 Definitions

Definition 2.1. *A* representation *of a group G, also called a G-module, is a vector space V with an associated homomorphism $\rho : G \to GL(V)$. We will refer to a representation by V.*

The map ρ induces a natural action of G on V, defined by $g \cdot v = (\rho(g))(v)$.

Definition 2.2. *A map $\varphi : V \to W$ is G-*equivariant *if the following diagram commutes:*

$$\begin{array}{ccc} V & \xrightarrow{\varphi} & W \\ \downarrow g & & \downarrow g \\ V & \xrightarrow{\varphi} & W \end{array}$$

That is, if $\varphi(g\cdot v) = g\cdot\varphi(v)$. A G-equivariant map is also called G-invariant or a G-homomorphism.

Definition 2.3. *A subspace $W \subseteq V$ is said to be a* subrepresentation*, or a G-submodule of a representation V over a group G if W is G*-equivariant*, that is if $g \cdot w \in W$ for all $w \in W$.*

Definition 2.4. *A representation V of a group G is said to be* irreducible *if it has no proper non-zero G-subrepresentations.*

Definition 2.5. *A group G is called* reductive *if every finite dimensional representation V of G is a direct sum of irreducible representation.*

Here are some examples of reductive groups:

- finite groups;
- the n-dimensional torus $(\mathbb{C}*)^n$;
- linear groups:
  - the general linear group $GL_n(\mathbb{C})$,
  - the special linear group $SL_n(\mathbb{C})$,
  - the orthogonal group $O_n(\mathbb{C})$ (linear transformations that preserve a symmetric form),
  - and the symplectic group $Sp_n(\mathbb{C})$ (linear transformations that preserve a skew symmetric form);
- Exceptional Lie Groups

Their reductivity is a nontrivial fact. It will be proved later in this lecture for finite groups, and the general and special linear groups. In some sense, the list above is complete: all reductive groups can be constructed by basic operations from the components which are either in this list or are related to them in a simple way.

2.1.2 New representations from old

Given representations V and W of a group G, we can construct new representations in several ways, some of which are described below.

- Tensor product: $V \otimes W$. $g \cdot (v \otimes w) = (g \cdot v) \otimes (g \cdot w)$.
- Direct sum: $V \oplus W$.

- Symmetric tensor representation: The subspace $Sym^n(V) \subset V \otimes \cdots \otimes V$ spanned by elements of the form

$$\sum_\sigma (v_1 \otimes \cdots \otimes v_n) \cdot \sigma = \sum_\sigma v_{\sigma(1)} \otimes \cdots v_{\sigma(n)},$$

where σ ranges over all permutations in the symmetric group S_n.

- Exterior tensor representation: The subspace $\Lambda^n(V) \subset V \otimes \cdots \otimes V$ spanned by elements of the form

$$\sum_\sigma sgn(\sigma)(v_1 \otimes \cdots \otimes v_n) \cdot \sigma = \sum_\sigma sgn(\sigma) v_{\sigma(1)} \otimes \cdots v_{\sigma(n)}.$$

- Let V and W be representations, then $\mathrm{Hom}(V, W)$ is also a representation, where $g \cdot \varphi$ is defined so that the following diagram commutes:

$$\begin{array}{ccc} V & \xrightarrow{\varphi} & W \\ \Big\downarrow g & & \Big\downarrow g \\ V & \xrightarrow{g\cdot\varphi} & W \end{array}$$

More precisely,

$$(g \cdot \varphi)(v) = g \cdot (\varphi(g^{-1} \cdot v)).$$

- In particular, $V^* : V \to \mathbb{C}$ is a representation, and is called the *dual representation.*

- Let G be a finite group. Let S be a finite G-set (that is, a finite set with an associated action of G on its elements). We construct a vector space over any field K (we will be mostly concerned with the case $K = \mathbb{C}$), with a basis vector associated to each element in S. More specifically, consider the set $K[S]$ of formal sums $\sum_{s \in S} \alpha_s e_s$, where $\alpha_s \in K$, and e_s is a vector associated with $S \in s$. Note that this set has a vector space structure over K, and there is a natural induced action of G on $K[S]$, defined by:

$$g \cdot \sum_{s \in S} \alpha_s e_s = \sum_{s \in S} \alpha_s e_{g \cdot s}.$$

This action gives rise to a representation of G.

- In particular, G is a G-set under the action of left multiplication. The representation we obtain in the manner described above from this G-set is called the *regular representation.*

2.2 Reductivity of finite groups

Proposition 2.1. *Let G be a finite group. If W is a subrepresentation of a representation V, then there exists a representation $W^\perp$ s.t. $V = W \oplus W^\perp$.*

Proof. Choose any Hermitian form H_0 of V, and construct a new Hermitian form H defined as:

$$H(v, w) = \sum_{g \in G} H_o(g \cdot v, g \cdot w).$$

Averaging is a useful trick that is used very often in representation theory, because it ensures G-invariance. In fact, H is G-invariant, that is,

$$H(v, w) = \sum_{g \in G} H_o(g \cdot v, g \cdot w) = H(h \cdot v, h \cdot w)$$

Let $W^\perp$ be the perpendicular complement to W with respect to the Hermitian form H. Then $W^\perp$ is also G-invariant, and therefore it is a G-submodule. □

Corollary 2.1. *Every representation of a finite group is a direct sum of irreducible representations.*

Lemma 2.1. *(Schur) If V and W are irreducible representations over $\mathbb{C}$, and $\varphi : V \to W$ is a homomorphism (i.e. a G-invariant map), then:*

1. *Either φ is an isomorphism or $\varphi = 0$.*
2. *If $V = W$, $\varphi = \lambda I$ for some $\lambda \in \mathbb{C}$.*

Proof.
1. Since $\mathrm{Ker}(\varphi)$, and $\mathrm{Im}\varphi$ are G-submodules, either $\mathrm{Im}(\varphi) = V$ or $\mathrm{Im}(\varphi) = 0$.
2. Let $\varphi : V \to V$. Since $\mathbb{C}$ algebraically closed, there exists an eigenvalue λ of φ. Look at the map $\varphi - \lambda I : V \to V$. By (1), $\varphi - \lambda I = 0$ (it can't be an isomorphism because something maps to 0). So $\varphi = \lambda I$.

□

Corollary 2.2. *Every representation is a* unique *direct sum of irreducible representations. More precisely, given two decompositions into irreducible representations,*

$$V = \bigoplus V_i^{a_i} = \bigoplus W_j^{b_j},$$

there is a one to one correspondence between the V_i's and W_j's, and the multiplicities correspond.

Proof. exercise (follows from Schur's lemma). □

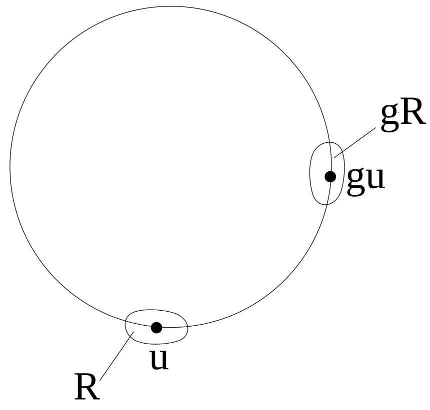

Figure 2.1: Example of a left Haar measure for the circle ($U_1(\mathbb{C})$). Left action by a group element g on a small region R around u does not change the area.

2.3 Compact Groups and $GL_n(\mathbb{C})$ are reductive

Now we prove reductivity of compact groups.

2.3.1 Compact groups

Examples of compact groups:

- $U_n(\mathbb{C}) \subseteq GL_n(\mathbb{C})$, the unitary groups (all rows are normal and orthogonal).
- $SU_n(\mathbb{C}) \subseteq SL_n(\mathbb{C})$, the special unitary group.

Given a compact group, a *left-invariant Haar measure* is a measure that is invariant under the left action of the group. In other words, multiplication by a group element does not change the area of a small region (i.e., the group action is an isometry, see figure 2.1).

Theorem 2.1. *Compact groups are reductive*

Proof. We use the averaging trick again. In fact the proof is the same as in the case of finite groups, using integration instead of summation for the averaging trick. Let H_0 be any Hermitian form on V. Then define H as:

$$H(v, w) = \int_G H(gv, gw)dG$$

where dG is a left-invariant Haar measure. Note that H is G-invariant. Let $W^{\perp}$ be the perpendicular complement to W. Then $W^{\perp}$ is G-invariant. Hence it is a G-submodule. $\square$

The same proof as before then gives us Schur's lemma for compact groups, from which follows:

Theorem 2.2. *If G is compact, then every finite dimensional representation of G is a unique direct sum of irreducible representations.*

2.3.2 Weyl's unitary trick and $GL_n(\mathbb{C})$

Theorem 2.3. *(Weyl) $GL_n(\mathbb{C})$ is reductive*

Proof. (general idea)
Let V be a representation of $GL_n(\mathbb{C})$. Then $GL_n(\mathbb{C})$ acts on V:

$$GL_n(\mathbb{C}) \hookrightarrow V.$$

But $U_n(\mathbb{C})$ is a subgroup of $GL_n(\mathbb{C})$. Therefore we have an induced action of $U_n(\mathbb{C})$ on V, and we can look at V as a representation of $U_n(\mathbb{C})$. As a representation of $U_n(\mathbb{C})$, V breaks into irreducible representations of $U_n(\mathbb{C})$ by the theorem above. To summarize, we have:

$$U_n(\mathbb{C}) \subseteq GL_n(\mathbb{C}) \hookrightarrow V = \oplus_i V_i,$$

where the V_i's are irreducible representations of $U_n(\mathbb{C})$. Weyl's unitary trick uses Lie algebra to show that every finite dimensional representation of $U_n(\mathbb{C})$ is also a representation of $GL_n(\mathbb{C})$, and irreducible representations of $U_n(\mathbb{C})$ correspond to irreducible representations of $GL_n(\mathbb{C})$. Hence each V_i above is an irreducible representation of $GL_n(\mathbb{C})$. $\square$

Once we know these groups are reductive, the goal is to construct and classify their irreducible finite dimensional representations. This will be done in the next lectures: Specht modules for S_n, and Weyl modules for $GL_n(\mathbb{C})$.

Chapter 3

Representation theory of reductive groups (cont)

Scribe: Paolo Codenotti

Goal: Basic representation theory, continued from the last lecture.

In this lecture we continue our introduction to representation theory. Again we refer the reader to the book by Fulton and Harris for full details [FH]. Let G be a finite group, and V a finite-dimensional G-representation given by a homomorphism $\rho : G \to GL(V)$. We define the *character* of the representation V (denoted χ_V) by $\chi_V(g) = Tr(\rho(g))$.

Since $Tr(A^{-1}BA) = Tr(B)$, $\chi_V(hgh^{-1}) = \chi_V(g)$. This means characters are constant on conjugacy classes (sets of the form $\{hgh^{-1} | h \in G\}$, for any $g \in G$). We call such functions *class functions*.

Our goal for this lecture is to prove the following two facts:

Goal 1 A finite dimensional representation is completely determined by its character.

Goal 2 The space of class functions is spanned by the characters of irreducible representations. In fact, these characters form an orthonormal basis of this space.

First, we prove some useful lemmas about characters.

Lemma 3.1. $\chi_{V \oplus W} = \chi_V + \chi_W$

Proof. Let $g \in G$, and let ρ, σ be homomorphisms from G into V and W, respectively. Let $\lambda_1, \ldots, \lambda_r$ be the eigenvalues of $\rho(g)$, and $\mu_1, \ldots, \mu_s$ the eigenvalues of $\sigma(g)$. Then $(\rho \oplus \sigma)(g) = (\rho(g), \sigma(g))$, so the eigenvalues of $(\rho \oplus \sigma)(g)$ are just the eigenvalues of $\rho(g)$ together with the eigenvalues of $\sigma(g)$.

Then $\chi_V(g) = \sum_i \lambda_i$, $\chi_W(g) = \sum_i \mu_i$, and $\chi_{V\oplus W} = \sum_i \lambda_i + \sum_i \mu_i$. $\square$

Lemma 3.2. $\chi_{V\otimes W} = \chi_V \chi_W$

Proof. Let $g \in G$, and let ρ, σ be homomorphisms into V and W, respectively. Let $\lambda_1, \dots, \lambda_r$ be the eigenvalues of $\rho(g)$, and $\mu_1, \dots, \mu_s$ the eigenvalues of $\sigma(g)$. Then $(\rho \otimes \sigma)(g)$ is the Kronecker product of the matrices $\rho(g)$ and $\sigma(g)$. So its eigenvalues are all $\lambda_i \mu_j$ where $1 \leq i \leq r$, $1 \leq j \leq s$.

Then, $\mathrm{Tr}((\rho \otimes \sigma)(g)) = \sum_{i,j} \lambda_i \mu_j = (\sum_i \lambda_i) \left(\sum_j \mu_j\right)$, which is equal to $\mathrm{Tr}(\rho(g))\mathrm{Tr}(\sigma(g))$. $\square$

3.1 Projection Formula

In this section, we derive a projection formula needed for Goal 1 that allows us to determine the multiplicity of an irreducible representation in another representation. Given a G-module V, let $V^G = \{v | \forall g \in G, g \cdot v = v\}$. We will call these elements G-invariant. Let

$$\phi = \frac{1}{|G|} \sum_{g \in G} g \in \mathrm{End}(V), \tag{3.1}$$

where each g, via ρ is considered an element of $\mathrm{End}(V)$.

Lemma 3.3. *The map $\phi : V \to V$ is a G-homomorphism; i.e., $\phi \in Hom_G(V, V) = (Hom(V, V))^G$.*

Proof. The set $\mathrm{End}(V)$ is a G-module, as we saw in last class, via the following commutative diagram: for any $\pi \in \mathrm{End}(V)$, and $h \in G$:

$$\begin{array}{ccc} V & \xrightarrow{\pi} & V \\ \big\downarrow h & & \big\downarrow h \\ V & \xrightarrow{h \cdot \pi} & V. \end{array}$$

Therefore $\pi \in Hom_G(V, V)$ (i.e., π is a G-equivariant morphism) iff $h \cdot \pi = \pi$ for all $h \in G$.

When ϕ is defined as in equation (3.1) above,

$$h \cdot \phi = \frac{1}{|G|} \sum_g hgh^{-1} = \frac{1}{|G|} \sum_g g = \phi.$$

Thus

$$h \cdot \phi = \phi, \ \forall h \in G,$$

and $\phi : V \to V$ is a G-equivariant morphism, i.e. $\phi \in Hom_G(V, V)$. $\square$

Lemma 3.4. *The map ϕ is a G-equivariant projection of V onto V^G*

Proof. For every $w \in W$, let

$$v = \phi(w) = \frac{1}{|G|} \sum_{g \in G} g \cdot w.$$

Then

$$h \cdot v = h \cdot \phi(w) = \frac{1}{|G|} \sum_{g \in G} hg \cdot w = v, \text{ for any } h \in G.$$

So $v \in V^G$. That is, $\mathrm{Im}(\phi) \subseteq V^G$. But if $v \in V^G$, then

$$\phi(v) = \frac{1}{|G|} \sum_{g \in G} g \cdot v = \frac{1}{|G|} |G| v = v.$$

So $V^G \subseteq \mathrm{Im}(\phi)$, and ϕ is the identity on V^G. This means that ϕ is the projection onto V^G. $\square$

Lemma 3.5.

$$\dim(V^G) = \frac{1}{|G|} \sum_{g \in G} \chi_V(g).$$

Proof. We have: $dim(V^G) = \mathrm{Tr}(\phi)$, because ϕ is a projection ($\phi = \phi|_{V^G} \oplus \phi|_{Ker(\phi)}$). Also,

$$\mathrm{Tr}(\phi) = \frac{1}{|G|} \sum_{g \in G} \mathrm{Tr}_V(g) = \frac{1}{|G|} \sum_{g \in G} \chi_V(g).$$

$\square$

This gives us a formula for the multiplicity of the trivial representation (i.e., $dim(V^G)$) inside V.

Lemma 3.6. *Let V, W be G-representations. If V is irreducible, $dim(Hom_G(V, W))$ is the multiplicity of V inside W. If W is irreducible, $dim(Hom_G(V, W))$ is the multiplicity of W inside V.*

Proof. By Schur's Lemma. $\square$

Let $C_{class}(G)$ be the space of class functions on (G), and let $(\alpha, \beta) = \frac{1}{|G|}\sum_g \overline{\alpha}(g)\beta(g)$ be the Hermitian form on C_{class}

Lemma 3.7. *If V and W are irreducible G representations, then*

$$(\chi_V, \chi_W) = \frac{1}{|G|}\sum_{g\in G} \overline{\chi_V}(g)\chi_W(g) = \begin{cases} 1 & \textit{if } V \cong W \\ 0 & \textit{if } V \ncong W. \end{cases} \tag{3.2}$$

Proof. Since $\mathrm{Hom}(V,W) \cong V^* \otimes W$, $\chi_{\mathrm{Hom}(V,W)} = \chi_{V^*}\chi_W = \overline{\chi_V}\chi_W$. Now the result follows from Lemmas 3.5 and 3.6. □

Lemma 3.8. *The characters of the irreducible representations form an orthonormal set.*

Proof. Follows from Lemma 3.7. □

If V,W are irreducible, then $\langle \chi_V, \chi_W\rangle$ is 0 if $V \neq W$ and 1 otherwise.
This implies that:

Theorem 3.1 (Goal 1)**.** *A representation is determined completely by its character.*

Proof. Let $V = \bigoplus_i V_i^{\oplus a_i}$. So $\chi_V = \sum_i a_i\chi_{V_i}$, and $a_i = (\chi_V, \chi_{V_i})$. This gives us a formula for the multiplicity of an irreducible representation in another representation, solely in terms of their characters. Therefore, a representation is completely determined by its character. □

3.2 The characters of irreducible representations form a basis

In this section, we address Goal 2.

Let R be the regular representation of G, V an irreducible representation of G.

Lemma 3.9.

$$R = \bigoplus_V End(V, V),$$

where V ranges over all irreducible representations of G.

Proof. $\chi_R(g)$ is 0 if g is not the identity and $|G|$ otherwise.

$$(\chi_R, \chi_V) = \frac{1}{|G|}\sum_{g\in G}\overline{\chi_R}(g)\chi_V(g) = \frac{1}{|G|}|G|\chi_V(e) = \chi_V(e) = dim(V)$$

□

Let $\alpha : G \to \mathbb{C}$. For any G-module V, let $\phi_{\alpha,V} = \sum_g \alpha(g)g : V \to V$

Exercise 3.1. *$\phi_{\alpha,V}$ is G equivariant (i.e. a G-homomorphism) iff α is a class function.*

Proposition 3.1. *Suppose $\alpha : G \to \mathbb{C}$ is a class function, and $(\alpha, \chi_V) = 0$ for all irreducible representations V. Then α is identically 0.*

Proof. If V is irreducible, then, by Schur's lemma, since $\phi_{\alpha,V}$ is a G-homomorphism, and V is irreducible, $\phi_{\alpha,V} = \lambda\mathrm{Id}$, where $\lambda = \frac{1}{n}\mathrm{Tr}(\phi_{\alpha,V})$, $n = dim(V)$. We have:

$$\lambda = \frac{1}{n}\sum_g \alpha(g)\chi_V(g) = \frac{1}{n}|G|(\alpha, \chi_{V^*}).$$

Now V is irreducible iff V^* is irreducible. So $\lambda = \frac{1}{n}|G|0 = 0$. Therefore, $\phi_{\alpha,V} = 0$ for any irreducible representation, and hence for any representation.

Now let V be the regular representation. Since g as endomorphisms of V are linearly independent, $\phi_{\alpha,V} = 0$ implies that $\alpha(g) = 0$. □

Theorem 3.2. *Characters form an orthonormal basis for the space of class functions.*

Proof. Follows from Proposition 3.1, and Lemma 3.8 □

If $V = \bigoplus_i V_i^{\oplus a_i}$, and $\pi_i : V \to V_i^{\oplus a_i}$ is the projection operator. We have a formula $\pi = \frac{1}{|G|}\sum_g g$ for the trivial representation. Analogously:

Exercise 3.2. $\pi_i = \frac{dimV_i}{|G|}\sum_g \overline{\chi_{V_i}}(g)g$.

3.3 Extending to Infinite Compact Groups

In this section, we extend the preceding results to infinite compact groups. We must take some facts as given, since these theorems are much more complicated than those for finite groups.

Consider compact G, specifically $U_n(\mathbb{C})$, the unitary subgroup of $G_(\mathbb{C})$. $U_1(\mathbb{C})$ is the circle group. Since $U_1(\mathbb{C})$ is abelian, all its representations are one-dimensional.

Since the group G is infinite, we can no longer sum over it. The idea is to replace the sum $\frac{1}{|G|}\sum_g f(g)$ in the previous setting with $\int_G f(g)d\mu$, where μ is a left-invariant Haar measure on G. In this fashion, we can derive analogues of the preceding results for compact groups. We need to normalize, so we set $\int_G d\mu = 1$.

Let $\rho : G \to GL(V)$, where V is a finite dimensional G-representation. Let $\chi_V(g) = \mathrm{Tr}(\rho(g))$. Let $V = \bigoplus_i V_i^{a_i}$ be the complete decomposition of V into irreducible representations.

We can again create a projection operator $\pi : V \to V^G$, by letting $\pi = \int_G \rho(g)d\mu$.

Lemma 3.10. *We have:*

$$dim(V^G) = \int_G \chi_V(g)d\mu.$$

Proof. This result is analogous to Lemma 3.5 for finite groups. □

For class functions α, β, define an inner product

$$(\alpha, \beta) = \int_G \overline{\alpha}(g)\beta(g)d\mu.$$

Lemma 3.10 applied to $\mathrm{Hom}_G(V, W)$ gives

$$(\chi_V, \chi_W) = \int_G \overline{\chi_V}\chi_W d\mu = dim(\mathrm{Hom}_G(V, W)).$$

Lemma 3.11. *If V, W are irreducible, $(\chi_V, \chi_W) = 1$ if V and W are isomorphic, and $(\chi_V, \chi_W) = 0$ otherwise.*

Proof. This result is analogous to Lemma 3.7 for finite groups. □

Lemma 3.12. *The irreducible representations are orthonormal, just as in Lemma 3.8 in the case of finite groups.*

If V is reducible, $V = \bigoplus_i V_i^{\oplus a_i}$, then

$$a_i = (\chi_V, \chi_{V_i}) = \int_G \overline{\chi_V} \chi_{V_i} d\mu.$$

Hence

Theorem 3.3. *A finite dimensional representation is completely determined by its character.*

This achieves Goal 1 for compact groups. Goal 2 is much harder:

Theorem 3.4 (Peter-Weyl Theorem)**.** *(1) The characters of the irreducible representations of G span a dense subset of the space of continuous class functions.*

(2) The coordinate functions of all irreducible matrix representations of G span a dense subset of all continuous functions on G.

By a coordinate function of a representation $\rho : G \to GL(V)$, we mean the function on G corresponding to a fixed entry of the matrix form of $\rho(g)$. For $G = U_1(\mathbb{C})$, (2) gives the Fourier series expansion on the circle. Hence, the Peter-Weyl theorem constitutes a far reaching generalization of the harmonic analyis from the circle to general $U_n(\mathbb{C})$.

Chapter 4

Representations of the symmetric group

Scribe: Sourav Chakraborty

Goal: To determine the irreducible representations of the Symmetric group S_n and their characters.

Reference: [FH, F]

Recall

Let G be a reductive group. Then

1. Every finite dimensional representation of G is completely reducible, that is, can be written as a direct sum of irreducible representations.
2. Every irreducible representation is determined by its character.

Examples of reductive groups:

- Continuous: algebraic torus $(\mathbb{C}^*)^m$, general linear group $GL_n(\mathbb{C})$, special linear group $Sl_n(\mathbb{C})$, symplectic group $Sp_n(\mathbb{C})$, orthogonal group $O_n(\mathbb{C})$.
- Finite: alternating group A_n, symmetric group S_n, $Gl_n(\mathbb{F}_p)$, simple lie groups of finite type.

4.1 Representations and characters of S_n

The number of irreducible representations of S_n is the same as the the number of conjugacy classes in S_n since the irreducible characters form a basis of the space of class functions. Each permutation can be written uniquely as a product of disjoint cycles. The collection of lengths of the cycles in a permutation is called the cycle type of the permutation. So a cycle type of a permutation on n elements is a partition of n. And in S_n each conjugacy class is determined by the cycle type, which, in turn, is determined by the partition of n. So the number of conjugacy class is same as the number of partitions of n. Hence:

$$\text{Number of irreducible representations of } S_n = \text{Number of partitions of } n \tag{4.1}$$

Let $\lambda = \{\lambda_1 \geq \lambda_2 \geq \dots\}$ be a partition of n; i.e., the size $|\lambda| = \sum \lambda_i$ is n. The Young diagram corresponding to λ is a table shown in Figure 1. It is like an inverted staircase. The top row has λ_1 boxes, the second row has λ_2 boxes and so on. There are exactly n boxes.

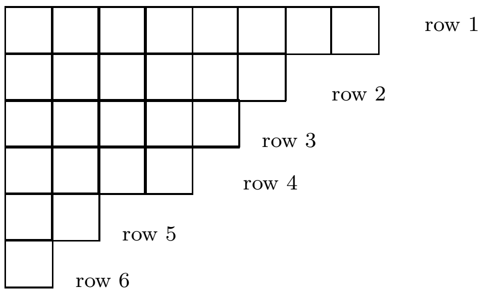

Figure 4.1: Row i has λ_i number of boxes

For a given partition λ, we want to construct an irreducible representation S_λ, called the Specht-module of S_n for the partition λ, and calculate the character of S_λ. We shall give three constructions of S_λ.

4.1.1 First Construction

A numbering T of a Young diagram is a filling of the boxes in its table with distinct numbers from $1, \dots, n$. A numbering of a Young diagram is also called a tableau. It is called a standard tableaux if the numbers are strictly increasing in each row and column. By T_{ij} we mean the value in the tableaux at i-th row and j-th column. We associate with each tableaux T a polynomial in $\mathbb{C}[X_1, X_2, \ldots, X_n]$:

$$f_T = \Pi_j \Pi_{i<i'} (X_{T_{ij}} - X_{T_{i'j}}).$$

Let S_λ be the subspace of $\mathbb{C}[X_1, X_2, \ldots, X_n]$ spanned by f_T's, where T ranges over all tableaux of shape λ. It is a representation of S_n. Here S_n acts on $\mathbb{C}[X_1, X_2, \ldots, X_n]$ as:

$$(\sigma.f)(X_1, X_2, \ldots, X_n) = f(X_{\sigma(1)}, X_{\sigma(2)}, \ldots, X_{\sigma(n)})$$

Theorem 4.1. *1. S_λ is irreducible.*

2. $S_\lambda \not\approx S_{\lambda'}$ if $\lambda \neq \lambda'$

3. The set $\{f_T\}$, where T ranges over standard tableau of shape λ, is a basis of S_λ.

4.1.2 Second Construction

Let T be a numbering of a Young diagram with distinct numbers from $\{1, \ldots, n\}$. An element σ in S_n acts on T in the usual way by permuting the numbers. Let $R(T), C(T) \subset S_n$ be the sets of permutations that fix the rows and columns of T, respectively. We have $R(\sigma T) = \sigma R(T) \sigma^{-1}$ and $C(\sigma T) = \sigma R(T) \sigma^{-1}$. We say $T \equiv T'$ if the rows of T and T' are the same up to ordering. The equivalence class of T, called the tabloid, is denoted by $\{T\}$. Its orbit is isomorphic to $S_n/R(T)$.

Let $\mathbb{C}[S_n]$ be the group algebra of S_n. Representations of S_n are the same as the representations of $\mathbb{C}[S_n]$. The element $a_T = \sum_{p \in R(T)} p$ in $\mathbb{C}[S_n]$ is called the row symmetrizer, $b_T = \sum_{q \in C(T)} sign(q) q$ the column symmetrizer, and $c_T = a_T b_T$ the Young symmetrizer.

Let

$$V_T = b_T.\{T\} = \sum_{q \in C(T)} sign(q) \{qT\}.$$

Then $\sigma.V_T = V_{\sigma T}$. Let S_λ be the span of all V_T's, where T ranges over all numberings of shape λ.

Theorem 4.2. *1. S_λ is irreducible*

2. $S_\lambda \not\approx S_{\lambda'}$ *if* $\lambda \neq \lambda'$.

3. *The set* $\{V_T | T$ *standard*$\}$ *forms a basis for* S_λ.

4.1.3 Third Construction

Let T be a canonical numbering of shape $\lambda = (\lambda_1, \ldots, \lambda_k)$. By this, we mean the first row is numbered by $1, \ldots, \lambda_1$, the second row by $\lambda_1 + 1, \ldots \lambda_1 + \lambda_2$, and so on, and the rows are increasing. Let $a_\lambda = a_T$, $b_\lambda = b_T$, and $c_\lambda = c_T = a_T.b_T$.

Then $S_\lambda = \mathbb{C}[S_n].c_\lambda$ is a representation of S_n from the left.

Theorem 4.3. 1. *S_λ is irreducible*

2. $S_\lambda \not\cong S_{\lambda'}$ *if* $\lambda \neq \lambda'$.

3. *The basis: an exercise.*

4.1.4 Character of S_λ [Frobenius character formula]

Let $i = (i_1, i_2, \ldots, i_k)$ be such that $\sum_j j i_j = n$. Let C_i be the conjugacy class consisting of permutations with i_j cycles of length j. Let χ_λ be the character of S_λ. The goal is to find $\chi_\lambda(C_i)$.

Let $\lambda : \lambda_1 \geq \cdots \geq \lambda_k$ be a partition of length k. Given k variables $X_1, X_2, \ldots, X_k$, let

$$P_j(X) = \sum X_i^j,$$

be the power sum, and

$$\Delta(X) = \Pi_{i<j}(X_j - X_i)$$

the discriminant. Let $f(X)$ be a formal power series on X_i's. Let $[f(X)]_{\ell_1, \ell_2, \ldots, \ell_k}$ denote the coefficient of $X_1^{\ell_1} X_2^{\ell_2} \ldots X_k^{\ell_k}$ in $f(X)$. Let $\ell_i = \lambda_i + k - i$.

Theorem 4.4 (Frobenius Character Formula)**.**

$$\chi_\lambda(C_i) = \left[\Delta(X) \cdot \Pi_j P_j(X)^{i_j}\right]_{\ell_1, \ell_2, \ldots, \ell_k}$$

4.2 The first decision problem in GCT

Now we can state the first hard decision problem in representation theory that arises in the context of the flip. Let S_α and S_β be two Specht modules of S_n. Since S_n is reductive, $S_\alpha \otimes S_\beta$ decomposes as

$$S_\alpha \otimes S_\beta = \bigoplus k_{\alpha\beta}^{\lambda} S_\lambda$$

Here $k_{\alpha\beta}^{\lambda}$ is called the Kronecker coefficient.

Problem 4.1. (Kronecker problem) *Given λ, α and β decide if $k_{\alpha\beta}^{\lambda} > 0$.*

Conjecture 4.1 (GCT6)**.** *This can be done in polynomial time; i.e. in time polynomial in the bit lengths of the inputs λ, α and β.*

55

Chapter 5

Representations of $GL_n(\mathbb{C})$

Scribe: Joshua A. Grochow

Goal: To determine the irreducible representations of $GL_n(\mathbb{C})$ and their characters.
References: [FH, F]

The goal of today's lecture is to classify all irreducible representations of $GL_n(\mathbb{C})$ and compute their characters. We will go over two approaches, the first due to Deruyts and the second due to Weyl.

A *polynomial representation* of $GL_n(\mathbb{C})$ is a representation $\rho : GL_n(\mathbb{C}) \to GL(V)$ such that each entry in the matrix $\rho(g)$ is a polynomial in the entries of the matrix $g \in GL_n(\mathbb{C})$.

The main result is that the polynomial irreducible representations of $GL_n(\mathbb{C})$ are in bijective correspondence with Young diagrams λ of height at most n, i.e. $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n \geq 0$. Because of the importance of Weyl's construction (similar constructions can be used on many other Lie groups besides $GL_n(\mathbb{C})$), the irreducible representation corresponding to λ is known as the *Weyl module* V_λ.

5.1 First Approach [Deruyts]

Let $X = (x_{ij})$ be a generic $n \times n$ matrix with variable entries x_{ij}. Consider the polynomial ring $\mathbb{C}[X] = \mathbb{C}[x_{11}, x_{12}, \ldots, x_{nn}]$. Then $GL_n(\mathbb{C})$ acts on $\mathbb{C}[X]$ by $(A \circ f)(X) = f(A^T X)$ (it is easily checked that this is in fact a left action).

Let T be a tableau of shape λ. To each column C of T of length r, we associate an $r \times r$ minor of X as follows: if C has the entries $i_1, \ldots, i_r$, then take from the first r columns of X the rows $i_1, \ldots, i_r$. Visually:

$$C = \begin{pmatrix} i_1 \\ \vdots \\ i_r \end{pmatrix} \longrightarrow e_C = \begin{array}{c} \\ \\ i_1 \rightarrow \\ i_2 \rightarrow \\ \vdots \\ i_r \rightarrow \end{array} \begin{array}{c} \begin{array}{ccccc} 1 & \cdots & r & & \\ \downarrow & & \downarrow & & \end{array} \\ \begin{pmatrix} \mathbf{x_{i_1,1}} & \cdots & \mathbf{x_{i_1,r}} & \cdots & x_{i_1,n} \\ \mathbf{x_{i_2,1}} & \cdots & \mathbf{x_{i_2,r}} & \cdots & x_{i_2,n} \\ \vdots & & \vdots & & \vdots \\ \mathbf{x_{i_r,1}} & \cdots & \mathbf{x_{i_r,r}} & \cdots & x_{i_r,n} \end{pmatrix} \end{array}$$

(Thus if there is a repeated number in the column C, $e_C = 0$, since the same row will get chosen twice.) Using these monomials e_C for each column C of the tableau T, we associate a monomial to the entire tableau, $e_T = \prod_C e_C$. (Thus, if in any column of T there is a repeated number, $e_T = 0$. Furthermore, the numbers must all come from $\{1, \ldots, n\}$ if they are to specify rows of an $n \times n$ matrix. So we restrict our attention to numberings of T from $\{1, \ldots, n\}$ in which the numbers in any given column are all distinct.)

Let V_λ be the vector space generated by the set $\{e_T\}$, where T ranges over all such numberings of shape λ. Then $GL_n(\mathbb{C})$ acts on V_λ: for $g \in GL_n(\mathbb{C})$, each row of gX is a linear combination of the rows of X, and since e_C is a minor of X, $g \cdot e_C$ is a linear combination of minors of X of the same size, i.e. $g(e_C) = \sum_D a^g_{C,D} e_D$ (this follows from standard linear algebra). Then

$$\begin{aligned} g(e_T) &= g(e_{C_1} e_{C_2} \cdots e_{C_k}) \\ &= \left(\sum_D a^g_{C_1,D} e_D\right) \cdots \left(\sum_D a^g_{C_k,D} e_D\right) \end{aligned}$$

If we expand this product out, we find that each term is in fact $e_{T'}$ for some T' of the appropriate shape. We then have the following theorem:

Theorem 5.1. *1. V_λ is an irreducible representation of $GL_n(\mathbb{C})$.*

2. *The set $\{e_T | T$ is a semistandard tableau of shape $\lambda\}$ is a basis for V_λ. (Recall that a semistandard tableau is one whose numbering is weakly increasing across each row and strictly increasing down each column.)*

3. *Every polynomial irreducible representation of $GL_n(\mathbb{C})$ of degree d is isomorphic to V_λ for some partition λ of d of height at most n.*

4. *Every rational irreducible representation of $GL_n(\mathbb{C})$ (each entry of $\rho(g)$ is a rational function in the entries of $g \in GL_n(\mathbb{C})$) is isomorphic to $V_\lambda \otimes \det^k$ for some partition λ of height at most n and for some integer k (where* det *is the determinant representation).*

5. *(Weyl's character formula) Define the character χ_λ of V_λ by $\chi_\lambda(g) = Tr(\rho(g))$, where $\rho : GL_n(\mathbb{C}) \to GL(V_\lambda)$ is the representation map. Then, for $g \in GL_n(\mathbb{C})$ with eigenvalues $x_1, \ldots, x_n$,*

$$\chi_\lambda(g) = S_\lambda(x_1, \ldots, x_n) := \frac{\left|x_j^{\lambda_i+n-i}\right|}{\left|x_j^{n-i}\right|}$$

(where $|y_j^i|$ is the determinant of the $n \times n$ matrix whose entries are $y_{ij} = y_j^i$, so, e.g., the determinant in the denominator is the usual van der Monde determinant, which is equal to $\prod_{i<j}(x_i - x_j)$). Here S_λ is a polynomial, called the Schur polynomial.

NB: It turns out that all holomorphic representations of $GL_n(\mathbb{C})$ are rational, and, by part (4) of the theorem, the Weyl modules classify all such representations up to scalar multiplication by powers of the determinant.

We'll give here a very brief introduction to the Schur polynomial introduced in the above theorem, and explain why the Schur polynomial S_λ associated to λ gives the character of V_λ.

Let λ be a partition, and T a semistandard tableau of shape λ. Define $x(T) = \prod_i x_i^{\mu_i(T)} \in \mathbb{C}[x_1, \ldots, x_n]$, where $\mu_i(T)$ is the number of times i appears in T. Then it can be shown [F] that

$$S_\lambda(x_1, \ldots, x_n) = \sum_T x(T),$$

where the sum is taken over all semistandard tableau of shape λ.

Proposition 5.1. *$S_\lambda(x_1, \ldots, x_n)$ is the character of V_λ, where $x_1, \ldots, x_n$ denote the eigenvalues of an element of $GL_n(\mathbb{C})$.*

Proof. It suffices to show this diagonalizable $g \in GL_n(\mathbb{C})$, since the diagonalizable matrices are dense in $GL_n(\mathbb{C})$.

So let $g \in GL_n(\mathbb{C})$ be diagonalizable with eigenvalues $x_1, \ldots, x_n$. We can assume that g is diagonal. If not, let A be a matrix that diagonalizes g. So AgA^{-1} is diagonal with $x_1, \ldots, x_n$ as its diagonal entries. If $\rho : GL_n(\mathbb{C}) \to GL(V_\lambda)$ is the representation corresponding to the module V_λ, then conjugate ρ by A to get $\rho' : GL_n(\mathbb{C}) \to GL(V_\lambda)$ defined by $\rho'(h) = A\rho(h)A^{-1}$. In particular, since trace is invariant under conjugation, ρ and ρ' have the same character. The module corresponding to ρ' is simply $A \cdot V_\lambda$, which is clearly isomorphic to V_λ since A is invertible. Thus to compute the character $\chi_\lambda(g)$, it suffices to compute the character of g under ρ', i.e., when g is diagonal, as we shall assume now.

We will show that e_T is an eigenvector of g with eigenvalue $x(T)$, i.e. $g(e_T) = x(T)e_T$. Then since $\{e_T | \text{T is a semistandard tableau of shape } \lambda\}$ is a basis for V_λ, the trace of g on V_λ will just be $\sum_T x(T)$, where the sum is over semistandard T of shape λ; this is exactly $S_\lambda(x_1, \ldots, x_n)$.

We reduce to the case where T is a single column. Suppose the claim is true for all columns C. Then since e_T is a product of e_C where C is a column, the corresponding eigenvalue of e_T will be $\prod_C x(C)$ (where the product is taken over the columns C of T), which is exactly $x(T)$.

So assume T is a single column, say with entries $i_1, \ldots, i_r$. Then e_T is simply the above-mentioned $r \times r$ minor of the generic $n \times n$ matrix $X = (x_{ij})$ (do not confuse the double-indexed entries of the matrix X with the single-indexed eigenvalues of g). Since g is diagonal, $g^t = g$. So $(g \circ e_T)(X) = e_T(g^t X) = e_T(gX)$. Thus g multiplies the i_j-th column by x_{i_j}. Thus its effect on e_T is simply to multiply it by $\prod_{j=1}^r x_{i_j}$, which is exactly $x(T)$. $\square$

5.1.1 Highest weight vectors

The subgroup $B \subset GL_n(\mathbb{C})$ of lower triangular invertible matrices, called the *Borel subgroup*, is solvable. So every irreducible representation of B is one-dimensional. A *weight vector* for $GL_n(\mathbb{C})$ is a vector v which is an eigenvector for every matrix $b \in B$. In other words, there is a function $\lambda : B \to \mathbb{C}$ such that $b \cdot v = \lambda(b)v$ for all $b \in B$. The restriction of λ to the subgroup of diagonal matrices in B is known as the *weight* of v.

As we showed in the proof of the above proposition,

$$\begin{pmatrix} x_1 & & \\ & \ddots & \\ & & x_n \end{pmatrix} e_T = x(T)e_T = x_1^{\lambda_1} \cdots x_n^{\lambda_n} e_T.$$

So e_T is a weight vector with weight $x(T)$. Thus Theorem 5.1 (2) gives a basis consisting entirely of weight vectors. We abbreviate the weight $x(T)$ by the sequence of exponents $(\lambda_1, \ldots, \lambda_n)$. We say e_T is a *highest weight vector* if its weight is the highest in the lexicographic ordering (using the above sequence notation for the weight).

Each V_λ has a unique (up to scalars) B-invariant vector, which turns out to be the highest weight vector: namely e_T, where T is canonical. For example, for $\lambda = (5, 3, 2, 2, 1)$, the canonical T is:

$$T = \begin{array}{|c|c|c|c|c|}\hline 1 & 1 & 1 & 1 & 1 \\ \hline 2 & 2 & 2 \\ \cline{1-3} 3 & 3 \\ \cline{1-2} 4 & 4 \\ \cline{1-2} 5 \\ \cline{1-1}\end{array}$$

Note that the weight of such e_T is $(\lambda_1, \ldots, \lambda_n)$, so that the highest weight vector uniquely determines λ, and thus the entire representation V_λ. (This is a general feature of highest weight vectors in the representation theory of Lie algebras and Lie groups.) Thus the irreducible representations $GL_n(\mathbb{C})$ are in bijective correspondence wih the highest weights of $GL_n(\mathbb{C})$, i.e. the sequences of exponents of the eigenvalues of the B-invariant eigenvectors.

5.2 Second Approach [Weyl]

Let $V = \mathbb{C}^n$ and consider the d-th tensor power $V^{\otimes d}$. The group $GL_n(\mathbb{C})$ acts on $V^{\otimes d}$ on the left by the diagonal action

$$g(v_1 \otimes \cdots \otimes v_d) = gv_1 \otimes \cdots \otimes gv_d \ (g \in GL(V))$$

while the symmetric group S_d acts on the right by

$$(v_1 \otimes \cdots \otimes v_d)\tau = v_{1\tau} \otimes \cdots \otimes v_{d\tau} \ (\tau \in S_d) \ .$$

These two actions commute, so $V^{\otimes d}$ is a representation of $H = GL_n(\mathbb{C}) \times S_d$. Every irreducible representation of H is of the form $U \otimes W$ for some irreducible representation U of $GL_n(\mathbb{C})$ and some irreducible representation W of S_d. Since both $GL_n(\mathbb{C})$ and S_d are reductive (every finite-dimensional representation is a direct sum of irreducible representations), their product H is reductive as well. So there are some partitions α and β and integers $m_{\alpha\beta}$ such that

$$V^{\otimes d} = \bigoplus (V_\alpha \otimes S_\beta)^{m_{\alpha\beta}},$$

where V_α are Weyl modules and S_β are Specht modules (irreducible representations of the symmetric group S_d).

Theorem 5.2. $V^{\otimes d} = \bigoplus_\lambda V_\lambda \otimes S_\lambda$, *where the sum is taken over partitions* λ *of* d *of height at most* n*. (Note that each summand appears with multiplicity one, so this is a "multiplicity-free" decomposition.)*

Now, let T be any standard tableau of shape λ, and recall the Young symmetrizer c_T from our discussion of the irreducible representations of S_d. Then $V^{\otimes d}c_T$ is a representation of $GL_n(\mathbb{C})$ from the left (since $c_T \in \mathbb{C}[S_d]$ acts on the right, and the left action of $GL_n(\mathbb{C})$ and the right action of S_d commute.)

Theorem 5.3. $V^{\otimes d}c_T \cong V_\lambda$

Thus

$$V^{\otimes d} = \bigoplus_{\lambda:|\lambda|=d} \bigoplus_{\substack{\text{std. tableau} \\ T \text{ of shape } \lambda}} V^{\otimes d}c_T,$$

where $|\lambda| = \sum \lambda_i$ denotes the size of λ. In particular, V_λ occurs in $V^{\otimes d}$ with multiplicity $\dim(S_\lambda)$.

Finally, we construct a basis for $V^{\otimes d}c_T$. A *bitableau* of shape λ is a pair (U, T) where U is a semistandard tableau of shape λ and T is a standard tableau of shape λ. (Recall that the the semistandard tableau of shape λ are in natural bijective correspondence with a basis for the Weyl module V_λ, while the standard tableau of shape λ are in natural bijective correspondence with a basis for the Specht module S_λ.)

To each bitableau we associate a vector $e_{(U,T)} = e_{i_1} \otimes \cdots \otimes e_{i_d}$ where i_j is defined as follows. Each number $1, \ldots, n$ appears in T exactly once. The number i_j is the entry of U in the same location as the number j in T; pictorially:

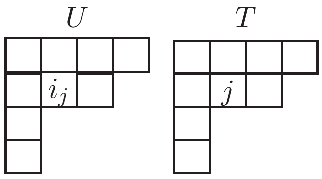

Then:

Theorem 5.4. *The set* $\{e_{(U,T)}c_T\}$ *is a basis for* $V^{\otimes d}c_T$.

Chapter 6

Deciding nonvanishing of Littlewood-Richardson coefficients

Scribe: Hariharan Narayanan

Goal: To show that nonvanishing of Littlewood-Richardson coefficients can be decided in polynomial time.

References: [DM2, GCT3, KT2]

6.1 Littlewood-Richardson coefficients

First we define Littlewood-Richardson coefficients, which are basic quantities encountered in representation theory. Recall that the irreducible representations V_λ of $GL_n(\mathbb{C})$, the Weyl modules, are indexed by partitions λ, and:

Theorem 6.1 (Weyl)**.** *Every finite dimensional representation of $GL_n(\mathbb{C})$ is completely reducible.*

Let $G = GL_n(\mathbb{C})$. Consider the diagonal embedding of $G \hookrightarrow G \times G$. This is a group homomorphism. Any $G \times G$ module, in particular, $V_\alpha \otimes V_\beta$ can also be viewed as a G module via this homomorphism. It then splits into irreducible G-submodules:

$$V_\alpha \otimes V_\beta = \oplus_\gamma c_{\alpha\beta}^{\gamma} V_\lambda. \tag{6.1}$$

Here $c_{\alpha\beta}^{\gamma}$ is the multiplicity of V_γ in $V_\alpha \otimes V_\beta$ and is known as the *Littlewood-Richardson* coefficient.

The character of V_λ is the Schur polynomial S_λ. Hence, it follows from (6.1) that the Schur polynomials satisfy the following relation:

$$S_\alpha S_\beta = \oplus_\gamma c_{\alpha\beta}^{\gamma} S_\lambda. \tag{6.2}$$

Theorem 6.2. *$c_{\alpha\beta}^{\gamma}$ is in PSPACE.*

Proof: This easily follows from eq.(6.2) and the definition of Schur polynomials. □

As a matter of fact, a stronger result holds:

Theorem 6.3. *$c_{\alpha\beta}^{\gamma}$ is in #P.*

Recall that $\#P \subsetneqq$ PSPACE.

Proof: This is an immediate consequence of the following Littlewood-Richardson rule (formula) for $c_{\alpha\beta}^{\gamma}$. To state it, we need a few definitions.

Given partitions γ and α, a skew Young diagram of shape γ/α is the difference between the Young diagrams for γ and α, with their top-left corners aligned; cf. Figure 6.2. A skew tableau of shape γ/α is a numbering of the boxes in this diagram. It is called semi-standard (SST) if the entries in each column are strictly increasing top to bottom and the entries in each row are weakly increasing left to right; see Figures 6.1 and 6.2. The row word $\mathrm{row}(T)$ of a skew-tableau T is the sequence of numbers obtained by reading T left to right, bottom to top; e.g. $\mathrm{row}(T)$ for Figure 6.2 is 13312211. It is called a reverse lattice word, if when read right to left, for each i, the number of i's encountered at any point is at least the number of $i+1$'s encountered till that point; thus the row word for Figure 6.2 is a reverse lattice word. We say that T is an LR tableau for given α, β, γ of shape γ/α and content β if

1. T is an SST,
2. $\mathrm{row}(T)$ is a reverse lattice word,
3. T has shape γ/α, and
4. the content of T is β, i. e. the number of i's in T is β_i.

1	1	2	5
2	2	3	
4			

Figure 6.1: semi-standard Young tableau

:	:	:	:	:	:	1	1
:	:	:	1	2	2		
:	:	3					
1	3						

Figure 6.2: Littlewood–Richardson skew tableau

For example, Figure 6.2 shows an LR tableau with $\alpha = (6, 3, 2)$, $\beta = (4, 2, 2)$ and $\gamma = (8, 6, 3, 2)$.

The Littlewood-Richardson rule [F, FH]: $c^{\gamma}_{\alpha\beta}$ is equal to the number of LR skew tableaux of shape γ/α and content β.

Remark: It may be noticed that the Littlewood-Richardson rule depends only on the partitions α, β and γ and not on n, the rank of $GL_n(\mathbb{C})$ (as long as it is greater than or equal to the maximum height of α, β or γ). For this reason, we can assume without loss of generalitity that n is the maximum of the heights of α, β and γ, as we shall henceforth.

Now we express $c^{\gamma}_{\alpha\beta}$ as the number of integer points in some polytope $P^{\gamma}_{\alpha\beta}$ using the Littlewood-Richardson rule:

Lemma 6.1. *There exists a polytope $P = P^{\gamma}_{\alpha\beta}$ of dimension polynomial in n such that the number of integer points in it is $c^{\gamma}_{\alpha\beta}$.*

Proof: Let $r^i_j(T)$, $i \leq n$, $j \leq n$, denote the number of j's in the i-th row of T. If T is an LR-tableau of shape γ/α with content β then these integers satisfy the following constraints:

1. Nonnegativity: $r^i_j \geq 0$.
2. Shape constraints: For $i \leq n$,

$$\alpha_i + \sum_j r^i_j = \gamma_i.$$

3. Content constraints: For $j \leq n$:
$$\sum_i r^i_j = \beta_j.$$
4. Tableau constraints:
$$\alpha_{i+1} + \sum_{k \leq j} r^{i+1}_k \leq \alpha_i + \sum_{k' < j} r^i_{k'}.$$
5. Reverse lattice word constraints: $r^i_j = 0$ for $i < j$, and for $i \leq n$, $1 < j \leq n$:
$$\sum_{i' \leq i} r^{i'}_j \leq \sum_{i' < i} r^{i'}_{j-1}.$$

Let $P^\gamma_{\alpha\beta}$ be the polytope defined by these constraints. Then $c^\gamma_{\alpha\beta}$ is the number of integer points in this polytope. This proves the lemma. $\square$

The membership function of the polytope $P^\gamma_{\alpha\beta}$ is clearly computable in time that is polynomial in the bitlengths of α, β and γ. Hence $c^\gamma_{\alpha\beta}$ belongs to $\#P$. This proves the theorem. $\square$

The complexity-theoretic content of the Littlewood-Richardson rule is that it puts a quantity, which is a priori only in PSPACE, in $\#P$. We also have:

Theorem 6.4 ([H]). *$c^\gamma_{\alpha\beta}$ is $\#P$-complete.*

Finally, the main complexity-theoretic result that we are interested in:

Theorem 6.5 (GCT3, Knutson-Tao, De Loera-McAllister). *The problem of deciding nonvanishing of $c^\gamma_{\alpha\beta}$ is in P, i. e., it can be solved in time that is polynomial in the bitlengths of α, β and γ. In fact, it can solved in strongly polynomial time [GCT3].*

Here, by a strongly polynomial time algorithm, we mean that the number of arithmetic steps $+, -, *, \leq, \ldots$ in the algorithm is polynomial in the number of parts of α, β and γ regardless of their bitlengths, and the bitlength of each intermediate operand is polynomial in the bitlengths of α, β and γ.

Proof: Let $P = P^\gamma_{\alpha\beta}$ be the polytope as in Lemma 6.1. All vertices of P have rational coefficients. Hence, for some positive integer q, the scaled polytope qP has an integer point. It follows that, for this q, $c^{q\gamma}_{q\alpha,q\beta}$ is positive. The saturation Theorem [KT] says that, in this case, $c^\gamma_{\alpha,\beta}$ is positive. Hence, P contains an integer point. This implies:

Lemma 6.2. *If $P \neq \emptyset$ then $c_{\alpha\beta}^{\gamma} > 0$.*

By this lemma, to decide if $c_{\alpha\beta}^{\gamma} > 0$, it suffices to test if P is nonempty. The polytope P is given by $Ax \leq b$ where the entries of A are 0 or 1– such linear programs are called combinatorial. Hence, this can be done in strongly polynomial time using Tardos' algorithm [GLS] for combinatorial linear programming. This proves the theorem. $\square$

The integer programming problem is NP-complete, in general. However, linear programming works for the specific integer programming problem here because of the saturation property [KT].

Problem: Find a genuinely combinatorial poly-time algorithm for deciding non-vanishing of $c_{\alpha\beta}^{\gamma}$.

Chapter 7

Littlewood-Richardson coefficients (cont)

Scribe: Paolo Codenotti

Goal: We continue our study of Littlewood-Richardson coefficients and define Littlewood-Richardson coefficients for the orthogonal group $O_n(\mathbb{C})$.
References: [FH, F]

Recall

Let us first recall some definitions and results from the last class. Let $c_{\alpha,\beta}^{\gamma}$ denote the Littlewood-Richardson coefficient for $GL_n(\mathbb{C})$.

Theorem 7.1 (last class). *Non-vanishing of $c_{\alpha,\beta}^{\gamma}$ can be decided in poly$(\langle\alpha\rangle, \langle\beta\rangle, \langle\gamma\rangle)$ time, where $\langle\ \rangle$ denotes the bit length.*

The positivity hypotheses which hold here are:

- $c_{\alpha,\beta}^{\gamma} \in \#P$, and more strongly,
- **Positivity Hypothesis 1 (PH1):** There exists a polytope $P_{\alpha,\beta}^{\gamma}$ of dimension polynomial in the heights of α, β and γ such that $c_{\alpha,\beta}^{\gamma} = \varphi(P_{\alpha,\beta}^{\gamma})$, where φ indicates the number of integer points.
- **Saturation Hypothesis (SH):** If $c_{k\alpha,k\beta}^{k\gamma} \neq 0$ for some $k \geq 1$, then $c_{\alpha,\beta}^{\gamma} \neq 0$ [Saturation Theorem].

Proof. (of theorem)
PH1 + SH + Linear programming. $\square$

This is the general form of algorithms in GCT. The main principle is that linear programming works for integer programming when PH1 and SH hold.

7.1 The stretching function

We define $\widetilde{c}_{\alpha,\beta}^{\gamma}(k) = c_{k\alpha,k\beta}^{k\gamma}$.

Theorem 7.2 (Kirillov, Derkesen Weyman [Der, Ki]). *$\widetilde{c}_{\alpha,\beta}^{\gamma}(k)$ is a polynomial in k.*

Here we prove a weaker result. For its statement, we will quickly review the theory of Ehrhart quasipolynomials (cf. Stanley [S]).

Definition 7.1. *(**Quasipolynomial**) A function $f(k)$ is called a* quasipolynomial *if there exist polynomials f_i, $1 \leq i \leq \ell$, for some ℓ such that*

$$f(k) = f_i(k) \text{ if } k \equiv i \text{ mod } \ell.$$

We denote such a quasipolynomial f by $f = (f_i)$. Here ℓ is called the period of $f(k)$ (we can assume it is the smallest such period). The degree of a quasipolynomial f is the max of the degrees of the f_i's.

Now let $P \subseteq \mathbb{R}^m$ be a polytope given by $Ax \leq b$. Let $\varphi(P)$ be the number of integer points inside P. We define the stretching function $f_P(k) = \varphi(kP)$, where kP is the dilated polytope defined by $Ax \leq kb$.

Theorem 7.3. *(Ehrhart) The stretching function $f_P(k)$ is a quasipolynomial. Furthermore, $f_P(k)$ is a polynomial if P is an integral polytope (i.e. all vertices of P are integral).*

In view of this result, $f_P(k)$ is called the Ehrhart quasi-polynomial of P. Now $\widetilde{c}_{\alpha,\beta}^{\gamma}(k)$ is just the Ehrhart quasipolynomial of $P_{\alpha,\beta}^{\gamma}$, and $c_{\alpha,\beta}^{\gamma} = \varphi(P_{\alpha,\beta}^{\gamma})$, the number of integer points in $P_{\alpha,\beta}^{\gamma}$. Moreover $P_{\alpha,\beta}^{\gamma}$ is defined by the inequality $Ax \leq b$, where A is constant, and b is a homogeneous linear form in the coefficients of α, β, and γ.

However, $P_{\alpha,\beta}^{\gamma}$ need not be integral. Therefore Theorem (7.2) does not follow from Ehrhart's result. Its proof needs representation theory.

Definition 7.2. *A quasipolynomial $f(k)$ is said to be* positive *if all the coefficients of $f_i(k)$ are nonnegative. In particular, if $f(k)$ is a polynomial, then it's positive if all its coefficients are nonnegative.*

The Ehrhart quasipolynomial of a polytope is positive only in exceptional cases. In this context:
PH2 (positivity hypothesis 2) [KTT]: The polynomial $\widetilde{c}^{\gamma}_{\alpha,\beta}(k)$ is positive.

There is considerable computer evidence for this.

Proposition 7.1. *PH2 implies SH.*

Proof. Look at:
$$c(k) = \widetilde{c}^{\gamma}_{\alpha,\beta}(k) = \sum a_i k^i.$$
If all the coefficients a_i are nonnegative (by PH2), and $c(k) \neq 0$, then $c(1) \neq 0$. □

SH has a proof involving algebraic geometry [B]. Therefore we suspect that the stronger PH2 is a deep phenomenon related to algebraic geometry.

7.2 $O_n(\mathbb{C})$

So far we have talked about $GL_n(\mathbb{C})$. Now we move on to the orthogonal group $O_n(\mathbb{C})$. Fix Q, a symmetric bilinear form on C^n; for example, $Q(V, W) = V^T W$.

Definition 7.3. *The orthogonal group $O_n(\mathbb{C}) \subseteq GL_n(\mathbb{C})$ is the group consisting of all $A \in GL_n(\mathbb{C})$ s.t. $Q(AV, AW) = Q(V, W)$ for all V and $W \in \mathbb{C}^n$. The subgroup $SO_n(\mathbb{C}) \subseteq SL_n(\mathbb{C})$, where $SL_n(\mathbb{C})$ is the set of matrices with determinant 1, is defined similarly.*

Theorem 7.4 (Weyl)**.** *The group $O_n(\mathbb{C})$ is reductive*

Proof. The proof is similar to the reductivity of $GL_n(\mathbb{C})$, based on Weyl's unitary trick. □

The next step is to classify all irreducible polynomial representations of $O_n(\mathbb{C})$. Fix a partition $\lambda = (\lambda_1 \geq \lambda_2 \geq \dots)$ of length at most n. Let $|\lambda| = d = \sum \lambda_i$ be its size. Let $V = \mathbb{C}^n$, $V^{\otimes d} = V \otimes \cdots \otimes V$ d times, and embed the Weyl module V_λ of $GL_n(\mathbb{C})$ in $V^{\otimes d}$ as per Theorem 5.3. Define a contraction map
$$\varphi_{p,q} : V^{\otimes d} \to V^{\otimes (d-2)}$$
for $1 \leq p \leq q \leq d$ by:
$$\varphi_{p,q}(v_{i_1} \otimes \cdots \otimes v_{i_d}) = Q(v_{i_p}, v_{i_q})(v_{i_1} \otimes \cdots \otimes \widehat{v_{i_p}} \otimes \cdots \otimes \widehat{v_{i_q}} \otimes \cdots \otimes v_{i_d}),$$

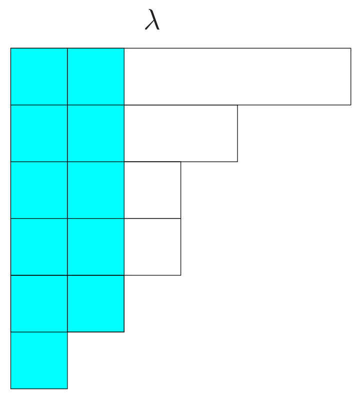

Figure 7.1: The first two columns of the partition λ are highlighted.

where $\widehat{v_{i_p}}$ means omit v_{i_p}.

It is $O_n(\mathbb{C})$-equivariant, i.e. the following diagram commutes:

$$\begin{array}{ccc} V^{\otimes d} & \xrightarrow{\varphi_{p,q}} & V^{\otimes d-2} \\ \Big\downarrow{\scriptstyle \sigma\in O_n(\mathbb{C})} & & \Big\downarrow{\scriptstyle \sigma\in O_n(\mathbb{C})} \\ V^{\otimes d} & \xrightarrow{\varphi_{p,q}} & V^{\otimes d-2} \end{array}$$

Let

$$V^{[d]} = \bigcap_{pq} ker(\varphi_{p,q}).$$

Because the maps are equivariant, each kernel is an $O_n(\mathbb{C})$-module, and $V^{[d]}$ is an $O_n(\mathbb{C})$-module. Let $V_{[\lambda]} = V^{[d]} \bigcap V_\lambda$, where $V_\lambda \subseteq V^{\otimes d}$ is the embedded Weyl module as above. Then $V_{[\lambda]}$ is an $O_n(\mathbb{C})$-module.

Theorem 7.5 (Weyl)**.** *$V_{[\lambda]}$ is an irreducible representation of $O_n(\mathbb{C})$. Moreover, the following two conditions hold:*

1. *If n is odd, then $V_{[\lambda]}$ is non-zero if and only if the sum of the lengths of the first two columns of λ is $\leq n$ (see figure 7.1).*

2. *If n is odd, then each polynomial irreducible representation is isomorphic to $V_{[\lambda]}$ for some λ.*

Let

$$V_{[\lambda]} \otimes V_{[\mu]} = \otimes_\gamma d^\gamma_{\lambda,\mu} V_{[\gamma]}$$

be the decomposition of $V_{[\lambda]} \otimes V_{[\mu]}$ into irreducibles. Here $d^{\gamma}_{\lambda,\mu}$ is called the Littlewood-Richardson coefficient of type B. The types of various connected reductive groups are defined as follows:

- $GL_n(\mathbb{C})$: type A
- $O_n(\mathbb{C})$, n odd: type B
- $Sp_n(\mathbb{C})$: type C
- $O_n(\mathbb{C})$, n even: type D

The Littlewood-Richardson coefficient can be defined for any type in a similar fashion.

Theorem 7.6 (Generalized Littlewood-Richardson rule)**.** *The Littlewood-Richardson coefficient $d^{\gamma}_{\lambda,\mu} \in \#P$. This also holds for any type.*

Proof. The most transparent proof of this theorem comes through the theory of quantum groups [K]; cf. Chapter 20. □

As in type A this leads to:

Hypothesis 7.1 (PH1)**.** *There exists a polytope $P^{\gamma}_{\lambda,\mu}$ of dimension polynomial in the heights of λ, μ and γ such that:*

1. *$d^{\gamma}_{\lambda,\mu} = \varphi(P^{\gamma}_{\lambda,\mu})$, the number of integer points in $P^{\gamma}_{\lambda,\mu}$, and*
2. *$\widetilde{d}^{\gamma}_{\lambda,\mu}(k) = d^{k\gamma}_{k\lambda,k\mu}$ is the Ehrhart quasipolynomial of $P^{\gamma}_{\lambda,\mu}$.*

There are several choices for such polytopes; e.g. the BZ-polytope [BZ].

Theorem 7.7 (De Loera, McAllister [DM2])**.** *The stretching function $\widetilde{d}^{\gamma}_{\lambda,\mu}(k)$ is a quasipolynomial of degree at most 2; so also for types C and D.*

A verbatim translation of the saturation property fails here [Z]): there exist λ, μ and γ such that $d^{2\gamma}_{2\lambda,2\mu} \neq 0$ but $d^{\gamma}_{\lambda,\mu} = 0$. Therefore we change the definition of saturation:

Definition 7.4. *Given a quasipolynomial $f(k) = (f_i)$, $index(f)$ is the smallest i such that $f_i(k)$ is not an identically zero polynomial. If $f(k)$ is identically zero, $index(f) = 0$.*

Definition 7.5. *A quasipolynomial $f(k)$ is* saturated *if $f(index(f)) \neq 0$. In particular, if $index(f) = 1$, then $f(k)$ is saturated if $f(1) \neq 0$.*

A positive quasi-polynomial is clearly saturated.

Positivity Hypothesis 2 (PH2) [DM2]: The stretching quasipolyomial $\widetilde{d}^{\gamma}_{\lambda,\mu}(k)$ is positive.

There is considerable evidence for this.

Saturation Hypothesis (SH): The stretching quasipolynomial $\widetilde{d}^{\gamma}_{\lambda,\mu}(k)$ is saturated.

PH2 implies SH.

Theorem 7.8. *[GCT5] Assuming SH (or PH2), positivity of the Littlewood-Richardson coefficient $d^{\gamma}_{\lambda,\mu}$ of type B can be decided in $poly(\langle\lambda\rangle, \langle\mu\rangle, \langle\gamma\rangle)$ time.*

This is also true for all types.

Proof. next class. □

Chapter 8

Deciding nonvanishing of Littlewood-Richardson coefficients for $O_n(\mathbb{C})$

Scribe: Hariharan Narayanan

Goal: A polynomial time algorithm for deciding nonvanishing of Littlewood-Richardson coefficients for the orthogonal group assuming SH.
Reference: [GCT5]

Let $d_{\lambda,\mu}^{\nu}$ denote the Littlewood-Richardson coefficient of type B (i.e. for the orthogonal group $O_n(\mathbb{C})$, n odd) as defined in the earlier lecture. In this lecture we describe a polynomial time algorithm for deciding nonvanishing of $d_{\lambda,\mu}^{\nu}$ assuming the following positivity hypothesis PH2. Similar result also holds for all types, though we shall only concentrate on type B in this lecture.

Let $\tilde{d}_{\lambda,\mu}^{\nu}(k) = d_{k\lambda,k\mu}^{k\nu}$ denote the associated stretching function. It is known to be a quasi-polynomial of period at most two [DM2]. This means there are polynomials $f_1(k)$ and $f_2(k)$ such that

$$d_{k\lambda,k\mu}^{k\nu} = \begin{cases} f_1(k), & \text{if k is odd;} \\ f_2(k), & \text{if k is even.} \end{cases}$$

Positivity Hypothesis (PH2) [DM2]: The stretching quasi-polynomial $\tilde{d}_{\lambda,\mu}^{\nu}(k)$ is positive. This means the coefficients of f_1 and f_2 are all non-negative.

The main result in this lecture is:

Theorem 8.1. *[GCT5] If PH2 holds, then the problem of deciding the positivity (nonvanishing) of $d^{\nu}_{\lambda\mu}$ belongs to P. That is, this problem can be solved in time polynomial in the bitlengths of λ, μ and ν.*

We need a few lemmas for the proof.

Lemma 8.1. *If PH2 holds, the following are equivalent:*

(1) $d^{\nu}_{\lambda\mu} \geq 1$.

(2) There exists an odd integer k such that $d^{k\nu}_{k\lambda\, k\mu} \geq 1$.

Proof: Clearly (1) implies (2). By PH2, there exists a polynomial f_1 with non-negative coefficients such that

$$\forall \text{ odd } k,\ f_1(k) = d^{k\nu}_{k\lambda\, k\mu}.$$

Suppose that for some odd k, $d^{k\nu}_{k\lambda\, k\mu} \geq 1$. Then $f_1(k) \geq 1$. Therefore f_1 has at least one non-zero coefficient. Since all coefficients of f_1 are nonnegative, $d^{\nu}_{\lambda\mu} = f_1(1) > 0$. Since $d^{\nu}_{\lambda\mu}$ is an integer, (1) follows. □

Definition 8.1. *Let $\mathbb{Z}_{<2>}$ be the subring of $\mathbb{Q}$ obtained by localizing $\mathbb{Z}$ at 2:*

$$\mathbb{Z}_{<2>} := \left\{ \frac{p}{q} \mid p, \frac{q-1}{2} \in \mathbb{Z} \right\}.$$

This ring consists of all fractions whose denominators are odd.

Lemma 8.2. *Let $P \in \mathbb{R}^d$ be a convex polytope specified by $Ax \leq B$, $x_i \geq 0$ for all i, where A and B are integral. Let $\mathrm{Aff}(P)$ denote its affine span. The following are equivalent:*

(1) P contains a point in $\mathbb{Z}^d_{<2>}$.

(2) $\mathrm{Aff}(P)$ contains a point in $\mathbb{Z}^d_{<2>}$.

Proof: Since $P \subseteq \mathrm{Aff}(P)$, (1) implies (2). Now suppose (2) holds. We have to show (1). Let $z \in \mathbb{Z}^d_{<2>} \cap \mathrm{Aff}(P)$.

First, consider the case when $\mathrm{Aff}(P)$ is one dimensional. In this case, P is the line segment joining two points x and y in $\mathbb{Q}^d$. The point z can be expressed as an affine linear combination, $z = ax + (1-a)y$ for some $a \in \mathbb{Q}$. There exists $q \in \mathbb{Z}$ such that $qx \in \mathbb{Z}^d_{<2>}$ and $qy \in \mathbb{Z}^d_{<2>}$. Note that

$$\{z + \lambda(qx - qy) \mid \lambda \in \mathbb{Z}_{<2>}\} \subseteq \mathrm{Aff}(P) \cap \mathbb{Z}^d_{<2>}.$$

Since $\mathbb{Z}_{<2>}$ is a dense subset of $\mathbb{Q}$, the l.h.s. and hence the r.h.s. is a dense subset of $\mathrm{Aff}(P)$. Consequently, $P \cap \mathbb{Z}^d_{<2>} \neq \emptyset$.

Now consider the general case. Let u be any point in the interior of P with rational coordinates, and L the line through u and z. By restricting to L, the lemma reduces to the preceding one dimensional case. □

Lemma 8.3. *Let*

$$P = \{x \mid Ax \leq B, (\forall i) x_i \geq 0\} \subseteq \mathbb{R}^d$$

be a convex polytope where A and B are integral. Then, it is possible to determine in polynomial time whether or not $\mathrm{Aff}(P) \cap \mathbb{Z}^d_{<2>} = \emptyset$.

Proof: Using Linear Programming [Kha79, Kar84], a presentation of the form $Cx = D$ can be obtained for $\mathrm{Aff}(P)$ in polynomial time, where C is an integer matrix and D is a vector with integer coordinates. We may assume that C is square since this can be achieved by padding it with 0's if necessary, and extending D. The Smith Normal Form over $\mathbb{Z}$ of C is a matrix S such that $C = USV$ where U and V are unimodular and S has the form

$$\begin{pmatrix} s_{11} & 0 & \dots & 0 \\ 0 & s_{22} & \dots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \dots & s_{dd} \end{pmatrix}$$

where for $1 \leq i \leq d-1$, s_{ii} divides $s_{i+1\,i+1}$. It can be computed in polynomial time [KB79]. The question now reduces to whether $USVx = D$ has a solution $x \in \mathbb{Z}^d_{<2>}$. Since V is unimodular, its inverse has integer entries too, and $y := Vx \in \mathbb{Z}^d_{<2>} \Leftrightarrow x \in \mathbb{Z}^d_{<2>}$. This is equivalent to whether $Sy = U^{-1}D$ has a solution $y \in \mathbb{Z}^d_{<2>}$. Since S is diagonal, this can be answered in polynomial time simply by checking each coordinate. □

Proof of Theorem 8.1: By [BZ], there exists a polytope $P = P^{\nu}_{\lambda,\mu}$ such that the Littlewood-Richardson coefficient $d^{\nu}_{\lambda\mu}$ is equal to the number of integer points in P. This polytope is such that the number of integer points in the dilated polytope kP is $d^{k\nu}_{k\lambda\, k\mu}$. Assuming PH2, we know from Lemma 8.1 that

$$P \cap \mathbb{Z}^d \neq \emptyset \Leftrightarrow (\exists \text{ odd } k), kP \cap \mathbb{Z}^d \neq \emptyset.$$

The latter is equivalent to

$$P \cap \mathbb{Z}^d_{<2>} \neq \emptyset.$$

The theorem now follows from Lemma 8.2 and Lemma 8.3. □

In combinatorial optimization, LP works if the polytope is integral. In our setting, this is not necessarily the case [DM1]: the denominators of the coordinates of the vertices of P can be $\Omega(l)$, where l is the total height of λ, μ and ν. LP works here nevertheless because of PH2; it can be checked that SH is also sufficient.

Chapter 9

The plethysm problem

Scribe: Joshua A. Grochow

Goal: In this lecture we describe the general *plethysm problem*, state analogous positivity and saturation hypotheses for it, and state the results from GCT 6 which imply a polynomial time algorithm for deciding positivity of a plethysm constant assuming these hypotheses.
Reference: [GCT6]

Recall

Recall that a function $f(k)$ is *quasipolynomial* if there are functions $f_i(k)$ for $i = 1, \ldots, \ell$ such that $f(k) = f_i(k)$ whenever $k \equiv i \bmod \ell$. The number ℓ is then the *period* of f. The *index* of f is the least i such that $f_i(k)$ is not identically zero. If f is identically zero, then the index of f is zero by convention. We say f is *positive* if all the coefficients of each $f_i(k)$ are nonnegative. We say f is *saturated* if $f(\mathrm{index}(f)) \neq 0$. If f is positive, then it is saturated.

Given any function $f(k)$, we associate to it the rational series $F(t) = \sum_{k \geq 0} f(k)t^k$.

Proposition 9.1. *[S] The following are equivalent:*

1. *$f(k)$ is a quasipolynomial of period ℓ.*
2. *$F(t)$ is a rational function of the form $\frac{A(t)}{B(t)}$ where $\deg A < \deg B$ and every root of $B(t)$ is an ℓ-th root of unity.*

9.1 Littlewood-Richardson Problem [GCT 3,5]

Let $G = GL_n(\mathbb{C})$ and $c^{\gamma}_{\alpha,\beta}$ the Littlewood-Richardson coefficient – i.e. the multiplicity of the Weyl module V_γ in $V_\alpha \otimes V_\beta$. We saw that the positivity of $c^{\gamma}_{\alpha,\beta}$ can be decided in $poly(\langle\alpha\rangle, \langle\beta\rangle, \langle\gamma\rangle)$ time, where $\langle\cdot\rangle$ denotes the bit-length. Furthermore, we saw that the stretching function $\widetilde{c}^{\gamma}_{\alpha,\beta}(k) = c^{k\gamma}_{k\alpha,k\beta}$ is a polynomial, and the analogous stretching function for type B is a quasipolynomial of period at most 2.

9.2 Kronecker Problem [GCT 4,6]

Now we study the analogous problem for the representations of the symmetric group (the Specht modules), called the Kronecker problem.

Let S_α be the Specht module of the symmetric group S_m associated to the partition α. Define the Kronecker coefficient $\kappa^{\pi}_{\lambda,\mu}$ to be the multiplicity of S_π in $S_\lambda \otimes S_\mu$ (considered as an S_m-module via the diagonal action). In other words, write $S_\lambda \otimes S_\mu = \bigoplus_\pi \kappa^{\pi}_{\lambda,\mu} S_\pi$. We have $\kappa^{\pi}_{\lambda,\mu} = (\chi_\lambda \chi_\mu, \chi_\pi)$, where χ_λ denotes the character of S_λ. By the Frobenius character formula, this can be computed in PSPACE. More strongly, analogous to the Littlewood-Richardson problem:

Conjecture 9.1. *[GCT4, GCT6] The Kronecker coefficient $\kappa^{\pi}_{\lambda,\mu} \in \#\boldsymbol{P}$. In other words, there is a positive #P-formula for $\kappa^{\pi}_{\lambda,\mu}$.*

This is a fundamental problem in representation theory. More concretely, it can be phrased as asking for a set of combinatorial objects I and a characteristic function $\chi : \{I\} \to \{0,1\}$ such that $\chi \in \mathbf{FP}$ and $\kappa^{\pi}_{\lambda,\mu} = \sum_I \chi(I)$. Continuing our analogy:

Conjecture 9.2. *[GCT6] The problem of deciding positivity of $\kappa^{\pi}_{\lambda,\mu}$ belongs to $\boldsymbol{P}$.*

Theorem 9.1. *[GCT6] The stretching function $\widetilde{\kappa}^{\pi}_{\lambda,\mu}(k) = \kappa^{k\pi}_{k\lambda,k\mu}$ is a quasipolynomial.*

Note that $\kappa^{k\pi}_{k\lambda,k\mu}$ is a Kronecker coefficient for S_{km}.

There is also a dual definition of the Kronecker coefficients. Namely, consider the embedding

$$H = GL_n(\mathbb{C}) \times GL_n(\mathbb{C}) \hookrightarrow G = GL(\mathbb{C}^n \otimes \mathbb{C}^n),$$

where $(g,h)(v \otimes w) = (gv \otimes hw)$. Then

Proposition 9.2. *[FH] The Kronecker coefficient $\kappa^{\pi}_{\lambda,\mu}$ is the multiplicity of the tensor product of Weyl modules $V_\lambda(GL_n(\mathbb{C})) \otimes V_\mu(GL_n(\mathbb{C}))$ (this is an irreducible H-module) in the Weyl module $V_\pi(G)$ considered as an H-module via the embedding above.*

9.3 Plethysm Problem [GCT 6,7]

Next we consider the more general plethysm problem.

Let $H = GL_n(\mathbb{C})$, $V = V_\mu(H)$ the Weyl module of H corresponding to a partition μ, and $\rho : H \to G = GL(V)$ the corresponding representation map. Then the Weyl module $V_\lambda(G)$ of G for a given partition λ can be considered an H-module via the map ρ. By complete reducibility, we may decompose this H-representation as

$$V_\lambda(G) = \bigoplus_\pi a^\pi_{\lambda,\mu} V_\pi(H).$$

The coefficients $a^\pi_{\lambda,\mu}$ are known as *plethsym constants* (this definition can easily be generalized to any reductive group H). The Kronecker coefficient is a special case of the plethsym constant [Ki].

Theorem 9.2 (GCT 6)**.** *The plethysm constant $a^\pi_{\lambda,\mu} \in PSPACE$.*

This is based on a parallel algorithm to compute the plethysm constant using Weyl's character formula. Continuing in our previous trend:

Conjecture 9.3. *[GCT6] $a^\pi_{\lambda,\mu} \in \#$ **P** and the problem of deciding positivity of $a^\pi_{\lambda,\mu}$ belongs to **P**.*

For the stretching function, we need to be a bit careful. Define $\widetilde{a}^\pi_{\lambda,\mu} = a^{k\pi}_{k\lambda,\mu}$. Here the subscript μ is not stretched, since that would change G, while stretching λ and π only alters the representations of G.

As in the beginning of the lecture, we can associate a function $A^\pi_{\lambda,\mu}(t) = \sum_{k \geq 0} \widetilde{a}^\pi_{\lambda,\mu}(k) t^k$ to the plethysm constant. Kirillov conjectured that $A^\pi_{\lambda,\mu}(t)$ is rational. In view of Proposition 9.1, this follows from the following stronger result:

Theorem 9.3 (GCT 6)**.** *The stretching function $\widetilde{a}^\pi_{\lambda,\mu}(k)$ is a quasipolynomial.*

This is the main result of GCT 6, which in some sense allows GCT to go forward. Without it, there would be little hope for proving that the positivity of plethysm constants can be decided in polynomial time. Its proof is essentially algebro-geometric. The basic idea is to show that the stretching function is the Hilbert function of some algebraic variety with nice (i.e. "rational") singularities. Similar results are shown for the stretching functions in the algebro-geometric problems arising in GCT.

The main complexity-theoretic result in [GCT6] shows that, under the following positivity and saturation hypotheses (for which there is much experimental evidence), the positivity of the plethysm constants can indeed be decided in polynomial time (cf. Conjecture 9.3).

The first positivity hypothesis is suggested by Theorem 9.3: since the stretching function is a quasipolynomial, we may suspect that it is captured by some polytope:

Positivity Hypothesis 1 (PH1). There exists a polytope $P = P^{\pi}_{\lambda,\mu}$ such that:

1. $a^{\pi}_{\lambda,\mu} = \varphi(P)$, where φ denotes the number of integer points inside the polytope,
2. The stretching quasipolynomial (cf. Thm. 9.3) $\widetilde{a}^{\pi}_{\lambda,\mu}(k)$ is equal to the Ehrhart quasipolynomial $f_P(k)$ of P,
3. The dimension of P is polynomial in $\langle\lambda\rangle, \langle\mu\rangle$, and $\langle\pi\rangle$,
4. the membership in $P^{\pi}_{\lambda,\mu}$ can be decided in poly$(\langle\lambda\rangle, \langle\mu\rangle, \langle\pi\rangle)$ time, and there is a polynomial time separation oracle [GLS] for P.

Here (4) does not imply that the polytope P has only polynomially many constraints. In fact, in the plethysm problem there may be a super-polynomial number of constraints.

Positivity Hypothesis 2 (PH2). The stretching quasipolynomial $\widetilde{a}^{\pi}_{\lambda,\mu}(k)$ is positive.

This implies:

Saturation Hypothesis (SH). The stretching quasipolynomial is saturated.

Theorem 9.3 is essential to state these hypotheses, since positivity and saturation are properties that only apply to quasipolynomials. Evidence for PH1, PH2, and SH can be found in GCT 6.

Theorem 9.4. *[GCT6] Assuming PH1 and SH (or PH2), positivity of the plethysm constant $a^{\pi}_{\lambda,\mu}$ can be decided in poly$(\langle\lambda\rangle, \langle\mu\rangle, \langle\pi\rangle)$ time.*

This follows from the polynomial time algorithm for saturated integer programming described in the next class. As with Theorem 9.3, this also holds for more general problems in algebraic geometry.

Chapter 10

Saturated and positive integer programming

Scribe: Sourav Chakraborty

Goal : A polynomial time algorithm for saturated integer programming and its application to the plethysm problem.
Reference: [GCT6]

Notation : In this class we denote by $\langle a \rangle$ the bit-length of the a.

10.1 Saturated, positive integer programming

Let $Ax \leq b$ be a set of inequalities. The number of constraints can be exponential. Let $P \subset \mathbb{R}^n$ be the polytope defined by these inequalities. The bit length of P is defined to be $\langle P \rangle = n + \psi$, where ψ is the maximum bit-length of a constraint in the set of inequalities. We assume that P is given by a separating oracle. This means membership in P can be decided in poly($\langle P \rangle$) time, and if $x \notin P$ then a separating hyperplane is given as a proof as in [GLS].

Let $f_P(k)$ be the Ehrhart quasi-polynomial of P. Quasi-polynomiality means there exist polynomials $f_i(k)$, $1 \leq i \leq l$, l the period, so that $f_P(k) = f_i(k)$ if $k = i$ modulo l. Then

$$\text{Index}(f_P) = \min\{i | f_i(k) \text{not identically 0 as a polynomial}\}$$

The integer programming problem is called *positive* if $f_P(k)$ is positive whenever P is non-empty, and *saturated* if $f_P(k)$ is saturated whenever P is non-empty.

Theorem 10.1 (GCT6). *1. $Index(f_P)$ can be computed in time polynomial in the bit length $\langle P \rangle$ of P assuming that the separation oracle works in poly-$\langle P \rangle$-time.*

2. Saturated and hence positive integer programming problem can be solved in poly-$\langle P \rangle$-time.

The second statement follow from the first.

Proof. Let $Aff(P)$ denote the affine span of P. By [GLS] we can compute the specifications $Cx = d$, C and d integral, of $Aff(P)$ in poly($\langle P \rangle$) time. Without loss of generality, by padding, we can assume that C is square. By [KB79] we find the Smith-normal form of C in polynomial time. Let it be $\bar{C}$. So,

$$\bar{C} = ACB$$

where A and B are unimodular, and $\bar{C}$ is a diagonal matrix, where the diagonal entries $c_1, c_2, \dots$ are such that with $c_i | c_{i+1}$.

Clearly $Cx = d$ iff $\bar{C}z = \bar{d}$ where $z = B^{-1}x$ and $\bar{d} = Ad$.

So all equations here are of form

$$\bar{c}_i z_i = d_i \tag{10.1}$$

Without loss of generality we can assume that c_i and d_i are relatively prime. Let $\tilde{c} = lcm(c_i)$.

Claim 10.1. $Index(f_P) = \tilde{c}$.

From this claim the theorem clearly follows.

Proof of the claim. Let $f_P(t) = \sum_{k \geq 0} f_P(k) t^k$ be the Ehrhart Series of P.

Now kP will not have an integer point unless $\tilde{c}$ divides k because of (10.1).

Hence $f_P(t) = f_{\bar{P}}(t^{\tilde{c}})$ where $\bar{P}$ is the stretched polytope $\tilde{c}P$ and $f_{\bar{P}}(s)$ is the Ehrhart series of $\bar{P}$. From this it follows that

$$Index(f_P) = \tilde{c} Index(f_{\bar{P}})$$

Now we show that $Index(f_{\bar{P}}) = 1$.

The equations of $\bar{P}$ are of the form

$$z_i = \frac{\tilde{c}}{c_i} d_i$$

where each $\frac{\tilde{c}}{c_i}$ is an integer. Therefore without loss of generality we can ignore these equations and assume the $\bar{P}$ is full dimensional.

Then it suffices to show that $\bar{P}$ contains a rational point whose denominators are all 1 modulo $\ell(\bar{P})$, the period of the quasi-polynomial $f_{\bar{P}}(s)$.

This follows from a simple density argument that we saw earlier (cf. the proof of Lemma 8.2).

From this the claim follows. □

□

10.2 Application to the plethysm problem

Now we can prove the result stated in the last class:

Theorem 10.2. *Assuming PH1 and SH, positivity of the plethysm constant $a^{\pi}_{\lambda,\mu}$ can be decided in time polynomial in $\langle\lambda\rangle, \langle\mu\rangle$ and $\langle\pi\rangle$.*

Proof. Let $P = P^{\pi}_{\lambda,\mu}$ be the polytope as in PH1 such that $a^{\pi}_{\lambda,\mu}$ is the number of integer points in P. The goal is to decide if P contains an integer point. This integer programming problem is saturated because of SH. Hence the result follows from Theorem 10.1. □

Chapter 11

Basic algebraic geometry

Scribe: Paolo Codenotti

Goal: So far we have focussed on purely representation-theoretic aspects of GCT. Now we have to bring in algebraic geometry. In this lecture we review the basic definitions and results in algebraic geometry that will be needed for this purpose. The proofs will be omitted or only sketched. For details, see the books by Mumford [Mm] and Fulton [F].

11.1 Algebraic geometry definitions

Let $V = \mathbb{C}^n$, and $v_1, \ldots, v_n$ the coordinates of V.

Definition 11.1.

- *Y is an* affine algebraic set *in V if Y is the set of simultaneous zeros of a set of polynomials in v_i's.*
- *An algebraic set that cannot be written as the union of two proper algebraic sets Y_1 and Y_2 is called* irreducible.
- *An irreducible affine algebraic set is called an* affine variety.
- *The ideal of an affine algebraic set Y is $I(Y)$, the set of all polynomials that vanish on Y.*

For example, $Y = (v_1 - v_2^2 + v_3, v_3^2 - v_2 + 4v_1)$ is an irreducible affine algebraic set (and therefore an affine variety).

Theorem 11.1 (Hilbert)**.** *$I(Y)$ is finitely generated, i.e. there exist polynomials $g_1, \ldots, g_k$ that generate $I(Y)$. This means every $f \in I(Y)$ can be written as $f = \sum f_i g_i$ for some polynomials f_i.*

Let $\mathbb{C}[V]$, the coordinate ring of V, be the ring of polynomials over the variables $v_1, \ldots, v_n$. The coordinate ring of Y is defined to be $\mathbb{C}[Y] = \mathbb{C}[V]/I(Y)$. It is the set of polynomial functions over Y.

Definition 11.2.

- *$P(V)$ is the* projective space *associated with V, i.e. the set of lines through the origin in V.*
- *V is called the* cone *of $P(V)$.*
- *$\mathbb{C}[V]$ is called the* homogeneous coordinate ring *of $P(V)$.*
- *$Y \subseteq P(V)$ is a* projective algebraic set *if it is the set of simultaneous zeros of a set of homogeneous forms (polynomials) in the variables $v_1, \ldots, v_n$. It is necessary that the polynomials be homogeneous because a point in $P(V)$ is a line in V.*
- *A projective algebraic set Y is* irreducible *if it can not be expressed as the union of two proper algebraic sets in $P(V)$.*
- *An irreducible projective algebraic set is called a* projective variety.

Let $Y \subseteq P(V)$ be a projective variety, and define $I(Y)$, the ideal of Y to be the set of all homogeneous forms that vanish on Y. Hilbert's result implies that $I(Y)$ is finitely generated.

Definition 11.3. *The* cone *$C(Y) \subseteq V$ of a projective variety $Y \subseteq P(V)$ is defined to be the set of all points on the lines in Y.*

Definition 11.4. *We define the* homogeneous coordinate ring *of Y as $R(Y) = \mathbb{C}[V]/I(Y)$, the set of homogeneous polynomial forms on the cone of Y.*

Definition 11.5. *A* Zariski open subset *of Y is the complement of a projective algebraic subset of Y. It is called a* quasi-projective variety.

Let $G = GL_n(\mathbb{C})$, and V a finite dimensional representation of G. Then $\mathbb{C}[V]$ is a G-module, with the action of $\sigma \in G$ defined by:

$$(\sigma \cdot f)(v) = f(\sigma^{-1} v), \; v \in V.$$

Definition 11.6. *Let $Y \subseteq P(V)$ be a projective variety with ideal $I(Y)$. We say that Y is a G-*variety *if $I(Y)$ is a G-module, i.e., $I(Y)$ is a G-submodule of $\mathbb{C}[V]$.*

If Y is a projective variety, then $R(Y) = \mathbb{C}[V]/I(Y)$ is also a G-module. Therefore Y is G-invariant, i.e.

$$y \in Y \Rightarrow \sigma y \in Y, \ \forall \sigma \in G.$$

The algebraic geometry of G-varieties is called geometric invariant theory (GIT).

11.2 Orbit closures

We now define special classes of G-varieties called *orbit closures*. Let $v \in P(V)$ be a point, and Gv the orbit of v:

$$Gv = \{gv | g \in G\}.$$

Let the stabilizer of v be

$$H = G_v = \{g \in G | gv = v\}.$$

The orbit Gv is isomorphic to the space G/H of cosets, called the homogeneous space. This is a very special kind of algebraic variety.

Definition 11.7. *The* orbit closure *of v is defined by:*

$$\Delta_V[v] = \overline{Gv} \subseteq P(V).$$

Here $\overline{Gv}$ is the closure of the orbit Gv in the complex topology on $P(V)$ (see figure 11.1).

A basic fact of algebraic geometry:

Theorem 11.2. *The orbit closure $\Delta_V[v]$ is a projective G-variety*

It is also called an almost homogeneous space.

Let $I_V[v]$ be the ideal of $\Delta_V[v]$, and $R_V[v]$ the homogeneous coordinate ring of $\Delta_V[v]$. The algebraic geometry of general orbit closures is hopeless, since the closures can be horrendous (see figure 11.1). Fortunately we shall only be interested in very special kinds of orbit closures with good algebraic geometry.

We now define the simplest kind of orbit closure, which is obtained when the orbit itself is closed. Let V_λ be an irreducible Weyl module of $GL_n(\mathbb{C})$, where $\lambda = (\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n \geq 0)$ is a partition. Let v_λ be the highest weight point in $P(V_\lambda)$, i.e., the point corresponding to the highest weight

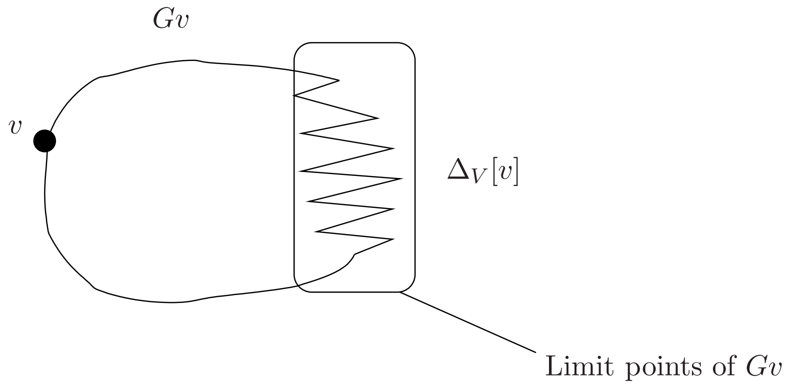

Figure 11.1: The limit points of Gv in $\Delta_V[v]$ can be horrendous.

vector in V_λ. This means $bv_\lambda = v_\lambda$ for all $b \in B$, where $B \subseteq GL_{\mathbf{n}}(C)$ is the Borel subgroup of lower triangular matrices. Recall that the highest weight vector is unique.

Consider the orbit Gv_λ of v_λ. Basic fact:

Proposition 11.1. *The orbit Gv_λ is already closed in $P(V)$.*

It can be shown that the stabilizer $P_\lambda = G_{v_\lambda}$ is a group of block lower triangular matrices, where the block lengths only depend on λ (see figure 11.2). Such subgroups of $GL_n(\mathbb{C})$ are called *parabolic subgroups*, and will be denoted by P. Clearly $Gv_\lambda \cong G/P_\lambda = G/P$.

11.3 Grassmanians

The simplest examples of G/P are *Grassmanians.*

Definition 11.8. *Let $G = Gl_n(\mathbb{C})$, and $V = \mathbb{C}^n$. The Grassmanian Gr_d^n is the space of d-dimensional subspaces (containing the origin) of V.*

Examples:

1. Gr_1^2 is the set of lines in $\mathbb{C}^2$ (see figure 11.3).
2. More generally, $P(V) = Gr_1^n$.

Proposition 11.2. *The Grassmanian Gr_d^n is a projective variety (just like $P(V) = Gr_1^n$).*

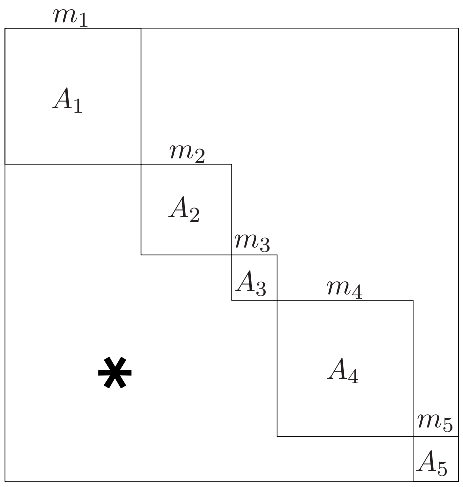

Figure 11.2: The parabolic subgroup of block lower triangular matrices. The sizes m_i only depend on λ.

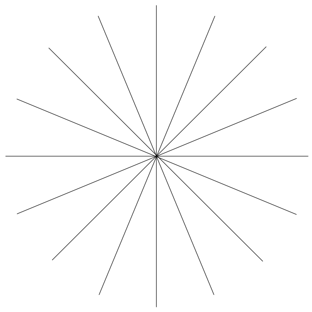

Figure 11.3: Gr_1^2 is the set of lines in $\mathbb{C}^2$.

It is easy to see that Gr^n_d is closed (since the limit of a sequence of d-dimensional subspaces of V is a d-dimensional subspace). Hence this follows from:

Proposition 11.3. *Let $\lambda = (1, \ldots, 1)$ be the partition of d, whose all parts are 1. Then $Gr^n_d \cong Gv_\lambda \subseteq P(V_\lambda)$.*

Proof. For the given λ, V_λ can be identified with the d^{th} wedge product

$$\Lambda^d(V) = \text{span}\{(v_{i_1} \wedge \cdots \wedge v_{i_d}) | i_1, \ldots, i_d \text{ are distinct}\} \subseteq V \otimes \cdots \otimes V \text{ (d times)},$$

where

$$(v_{i_1} \wedge \cdots \wedge v_{i_d}) = \frac{1}{d!} \sum_{\sigma \in S_d} sgn(\sigma)(v_{\sigma(i_1)} \otimes \cdots \otimes v_{\sigma(i_d)}).$$

Let Z be a variable $d \times n$ matrix. Then $\mathbb{C}[Z]$ is a G-module: given $f \in \mathbb{C}[Z]$ and $\sigma \in GL_n(\mathbb{C})$, we define the action of σ by

$$(\sigma \cdot f)(Z) = f(Z\sigma).$$

Now $\Lambda^d(V)$, as a G-module, is isomorphic to the span in $\mathbb{C}[Z]$ of all $d \times d$ minors of Z.

Let $A \in Gr^n_d$ be a d-dimensional subspace of V. Take any basis $\{v_1, \ldots, v_d\}$ of A. The point $v_1 \wedge \cdots \wedge v_d \in \Lambda^d(V)$ depends only on the subspace A, and not on the basis, since the change of basis does not change the wedge product. Let Z_A be the $d \times n$ complex matrix whose rows are the basis vectors $v_1, \ldots, v_d$ of A. The *Plucker map* associates with A the tuple of all $d \times d$ minors $A_{j_1,\ldots,j_d}$ of Z_A, where $A_{j_1,\ldots,j_d}$ denotes the minor of Z_A formed by the columns $j_1, \ldots, j_d$. This depends only on A, and not on the choice of basis for A.

The proposition follows from:

Claim 11.1. *The Plucker map is a G-equivariant map from Gr^n_d to $Gv_\lambda \subseteq P(V_\lambda)$ and $Gr^n_d \approx Gv_\lambda \subseteq P(V_\lambda)$.*

Proof. Exercise. Hint: take the usual basis, and note that the highest weight point v_λ corresponds to $v_1 \wedge \cdots \wedge v_d$. □

Chapter 12

The class varieties

Scribe: Hariharan Narayanan

Goal: Associate class varieties with the complexity classes $\#P$ and NC and reduce the $NC \neq P^{\#P}$ conjecture over $\mathbb{C}$ to a conjecture that the class variety for $\#P$ cannot be embedded in the class variety for NC.
reference: [GCT1]

The $NC \neq P^{\#P}$ conjecture over $\mathbb{C}$ says that the permanent of an $n \times n$ complex matrix X cannot be expressed as a determinant of an $m \times m$ complex matrix Y, $m = \text{poly}(n)$, whose entries are (possibly nonhomogeneous) linear forms in the entries of X. This obviously implies the $NC \neq P^{\#P}$ conjecture over $\mathbb{Z}$, since multivariate polynomials over $\mathbb{C}^n$ are determined by the values that they take over the subset $\mathbb{Z}^n$. The conjecture over $\mathbb{Z}$ is implied by the usual $NC \neq P^{\#P}$ conjecture over a finite field F_p, $p \neq 2$, and hence, has to be proved first anyway.

For this reason, we concentrate on the $NC \neq P^{\#P}$ conjecture over $\mathbb{C}$ in this lecture. The goal is to reduce this conjecture to a statement in geometric invariant theory.

12.1 Class Varieties in GCT

Towards that end, we associate with the complexity classes $\#P$ and NC certain projective algebraic varieties, which we call *class varieties*. For this, we need a few definitions.

Let $G = GL_\ell(\mathbb{C})$, V a finite dimensional representation of G. Let $P(V)$ be the associated projective space, which inherits the group action. Given a point $v \in P(V)$, let $\Delta_V[v] = \overline{Gv} \subseteq P(V)$ be its orbit closure. Here $\overline{Gv}$ is the closure of the orbit Gv in the complex topology on $P(v)$. It is a projective G-variety; i.e., a projective variety with the action of G.

All class varieties in GCT are orbit closures (or their slight generalizations), where $v \in P(V)$ corresponds to a complete function for the class in question. The choice of the complete function is crucial, since it determines the algebraic geometry of $\Delta_V[v]$.

We now associate a class variety with NC. Let $g = det(Y)$, Y an $m \times m$ variable matrix. This is a complete function for NC. Let $V = sym^m(Y)$ be the space of homogeneous forms in the entries of Y of degree m. It is a G-module, $G = GL_{m^2}(\mathbb{C})$, with the action of $\sigma \in G$ given by:

$$\sigma : f(Y) \longmapsto f(\sigma^{-1}Y).$$

Here $\sigma^{-1}Y$ is defined thinking of Y as an m^2-vector.

Let $\Delta_V[g] = \Delta_V[g,m] = \overline{Gg}$, where we think of g as an element of $P(V)$. This is the class variety associated with NC. If g is a different function instead of $det(Y)$, the algebraic geometry of $\Delta_V[g]$ would have been unmanageable. The main point is that the algebraic geometry of $\Delta_V[g]$ is nice, because of the very special nature of the determinant function.

We next associate a class variety with $\#P$. Let $h = perm(X)$, X an $n \times n$ variable matrix. Let $W = sym^n(X)$. It is similarly an H-module, $H = GL_k(\mathbb{C})$, $k = n^2$. Think of h as an element of $P(W)$, and let $\Delta_W[h] = \overline{Hh}$ be its orbit closure. It is called the class variety associated with $\#P$.

Now assume that $m > n$, and think of X as a submatrix of Y, say the lower principal submatrix. Fix a variable entry y of Y outside of X. Define the map $\phi : W \to V$ which takes $w(x) \in W$ to $y^{m-n}w(x) \in V$. This induces a map from $P(V)$ to $P(W)$ which we call ϕ as well. Let $\phi(h) = f \in P(V)$ and $\Delta_V[f,m,n] = \overline{Gf}$ its orbit closure. It is called the extended class variety associated with $\#P$.

Proposition 12.1 (GCT 1)**.**

1. *If $h(X) \in W$ can be computed by a circuit (over $\mathbb{C}$) of depth $\leq \log^c(n)$, c a constant, then $f = \phi(h) \in \Delta_V[g,m]$, for $m = O(2^{\log^c n})$.*

2. *Conversely if $f \in \Delta_V[g,m]$ for $m = 2^{\log^c n}$, then $h(X)$ can be approximated infinitesimally closely by a circuit of depth $\log^{2c} m$. That is, $\forall \epsilon > 0$, there exists a function $\tilde{h}(X)$ that can be computed by a circuit of depth $\leq \log^{2c} m$ such that $\|\tilde{h} - h\| < \epsilon$ in the usual norm on $P(V)$.*

If the permanent $h(X)$ can be approximated infinitesimally closely by small depth circuits, then every function in $\#P$ can be approximated infinitesimally closely by small depth circuits. This is not expected. Hence:

Conjecture 12.1 (GCT 1). *Let $h(X) = perm(X)$, X an $n \times n$ variable matrix. Then $f = \phi(h) \notin \Delta_V[g; m]$ if $m = 2^{polylog(n)}$ and n is sufficiently large.*

This is equivalent to:

Conjecture 12.2 (GCT 1). *The G-variety $\Delta_V[f; m, n]$ cannot be embedded as a G-subvariety of $\Delta_V[g, m]$, symbolically*

$$\Delta_V[f; m, n] \not\hookrightarrow \Delta_V[g, m],$$

if $m = 2^{polylog(n)}$ and $n \to \infty$.

This is the statement in geometric invariant theory (GIT) that we sought.

Chapter 13

Obstructions

Scribe: Paolo Codenotti

Goal: Define an obstruction to the embedding of the $\#P$-class variety in the NC-class-variety and describe why it should exist.
References: [GCT1, GCT2]

Recall

Let us first recall some definitions and results from the last class. Let Y be a generic $m \times m$ variable matrix, and X an $n \times n$ minor of Y (see figure 13.1). Let $g = \det(Y)$, $h = \mathrm{perm}(X)$, $f = \phi(h) = y^{m-n}\mathrm{perm}(X)$, and $V = \mathrm{Sym}^m[Y]$ the set of homogeneous forms of degree m in the entries of Y. Then V is a G-module for $G = GL(Y) = GL_l(\mathbb{C})$, $l = m^2$, with the action

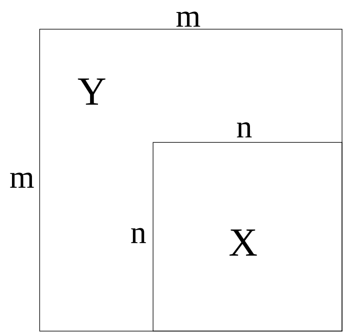

Figure 13.1: Here Y is a generic m by m matrix, and X is an n by n minor.

of $\sigma \in G$ given by

$$\sigma : f(Y) \to f(\sigma^{-1}Y),$$

where Y is thought of as an l-vector, and $P(V)$ a G-variety. Let

$$\Delta_V[f; m, n] = \overline{Gf} \subseteq P(V),$$

and

$$\Delta_V[g; m] = \overline{Gg} \subseteq P(V)$$

be the class varieties associated with $\#P$ and NC.

13.1 Obstructions

Conjecture 13.1. *[GCT1] There does not exist an embedding* $\Delta_V[f; m, n] \hookrightarrow \Delta_V[g; m]$ *with* $m = 2^{polylog(n)}$, $n \to \infty$.

This implies Valiant's conjecture that the permanent cannot be computed by circuits of polylog depth. Now we discuss how to go about proving the conjecture.

Suppose to the contrary,

$$\Delta[f; m, n] \hookrightarrow \Delta_V[g; m]. \tag{13.1}$$

We denote $\Delta_V[f; m, n]$ by $\Delta_V[f]$, and $\Delta_V[g; m]$ by $\Delta_V[g]$. Let $R_V[g]$ be the homogeneous coordinate ring of $\Delta_V[g]$. The embedding (13.1) implies existence of a surjection:

$$R_V[f] \twoheadleftarrow R_V[g] \tag{13.2}$$

This is a basic fact from algebraic geometry. The reason is that $R_V[g]$ is the set of homogeneous polynomial functions on the cone C of $\Delta_V[g]$, and any such function τ can be restricted to $\Delta_V[f]$ (see figure 13.2). Conversely, any polynomial function on $\Delta_V[f]$ can be extended to a homogeneous polynomial function on the cone C.

Let $R_V[f]_d$ and $R_V[g]_d$ be the degree d components of $R_V[f]$ and $R_V[g]$. These are G-modules since $\Delta_V[f]$ and $\Delta_V[g]$ are G-varieties. The surjection (13.2) is degree preserving. So there is a surjection

$$R_V[f]_d \twoheadleftarrow R_V[g]_d \tag{13.3}$$

for every d. Since G is reductive, both $R_V[f]_d$ and $R_V[g]_d$ are direct sums of irreducible G-modules. Hence the surjection (13.3) implies that $R_V[f]_d$ can be embedded as a G submodule of $R_V[g]_d$.

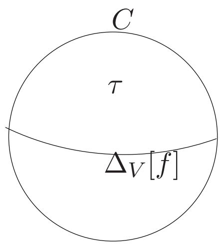

Figure 13.2: C denotes the cone of $\Delta_V[g]$.

Definition 13.1. *We say that a Weyl-module $S = V_\lambda(G)$ is an* obstruction *for the embedding (13.1) (or, equivalently, for the pair (f, g)) if $V_\lambda(G)$ occurs in $R_V[f; m, n]_d$, but not in $R_V[g; m]_d$, for some d. Here occurs means the multiplicity of $V_\lambda(G)$ in the decomposition of $R_V[f; m, n]_d$ is nonzero.*

If an obstruction exists for given m, n, then the embedding (13.1) does not exist.

Conjecture 13.2 (GCT2)**.** *An obstruction exists for the pair (f, g) for all large enough n if $m = 2^{polylog(n)}$.*

This implies Conjecture 13.1. In essence, this turns a nonexistence problem (of polylog depth circuit for the permanent) into an existence problem (of an obstruction).

If we replace the determinant here by any other complete function in NC, an obstruction need not exist. Because, as we shall see in the next lecture, the existence of an obstruction crucially depends on the exceptional nature of the class variety constructed from the determinant. The main goals of GCT in this context are:

1. understand the exceptional nature of the class varieties for NC and $\#P$, and

2. use it to prove the existence of obstructions.

13.1.1 Why are the class varieties exceptional?

We now elaborate on the exceptional nature of the class varieties. Its significance for the existence of obstructions will be discussed in the next lecture.

Let V be a G-module, $G = GL_n(\mathbb{C})$. Let $P(V)$ be a projective variety over V. Let $v \in P(V)$, and recall $\Delta_V[v] = \overline{Gv}$. Let $H = G_v$ be the stabilizer of v, that is, $G_v = \{\sigma \in G | \sigma v = v\}$.

Definition 13.2. *We say that v is* characterized by its stabilizer $H = G_v$ *if v is the only point in $P(V)$ such that $hv = v$, $\forall h \in H$.*

If v is characterized by its stabilizer, then $\Delta_V[v]$ is completely determined by the group triple $H \hookrightarrow G \hookrightarrow K = GL(V)$.

Definition 13.3. *The orbit closure $\Delta_V[v]$, when v is characterized by its stabilizer, is called a* group-theoretic variety.

Proposition 13.1. *[GCT1]*

1. *The determinant $g = \det(Y) \in P(V)$ is characterized by its stabilizer. Therefore $\Delta_V[g]$ is group theoretic.*

2. *The permanent $h = perm(X) \in P(W)$, where $W = Sym^n(X)$, is also characterized by its stabilizer. Therefore $\Delta_W[h]$ is also group theoretic.*

3. *Finally, $f = \phi(h) \in P(V)$ is also characterized by its stabilizer. Hence $\Delta_V[f]$ is also group theoretic.*

Proof. (1) It is a fact in classical representation theory that the stabilizer of $\det(Y)$ in $G = GL(Y) = GL_{m^2}(\mathbb{C})$ is the subgroup $G_{\det}$ that consists of linear transformations of the form $Y \to AY^*B$, where $Y^* = Y$ or Y^t, for any $A, B \in GL_m(\mathbb{C})$. It is clear that linear transformation of this form stabilize the determinant since:

1. $\det(AYB) = det(A)det(B)det(Y) = c\det(Y)$, where $c = \det(A)\det(B)$. Note that the constant c doesn't matter because we get the same point in the projective space.

2. $\det(Y^*) = \det(Y)$.

It is a basic fact in classical invariant theory that $\det(Y)$ is the only point in $P(V)$ stabilized by $G_{\det}$. Furthermore, the stabilizer $G_{\det}$ is reductive, since its connected part is $(G_{\det})_\circ \approx GL_m \times GL_m$ with the natural embedding

$$(G_{\det})_\circ = GL_m \times GL_m \hookrightarrow GL(\mathbb{C}^m \otimes \mathbb{C}^m) = GL_{m^2}(\mathbb{C}) = G.$$

(2) The stabilizer of perm(x) is the subgroup $G_{\text{perm}} \subseteq GL(X) = GL_{n^2}(\mathbb{C})$ generated by linear transformations of the form $X \to \lambda X^* \mu$, where $X^* = X or X^t$, and λ and μ are diagonal (which change the permanent by a constant factor) or permutation matrices (which do not change the permanent).

Finally, the discrete component of G_{perm} is isomorphic to $S_2 \rtimes S_n \times S_n$, where $\rtimes$ denotes semidirect product. The continuous part is $(\mathbb{C}^*)^n \times (\mathbb{C}^*)^n$. So G_{perm} is reductive.

(3) Similar.

$\square$

The main significance of this proposition is the following. Because $\Delta_V[g], \Delta_V[f]$, and $\Delta_W[h]$ are group theoretic, the algebraic geometric problems concerning these varieties can be "reduced" to problems in the theory of quantum groups. So the plan is:

1. Use the theory of quantum groups to understand the structure of the group triple associated with the algebraic variety.

2. Translate this understanding to the structure of the algebraic variety.

3. Use this to show the existence of obstructions.

Chapter 14

Group theoretic varieties

Scribe: Joshua A. Grochow

Goal: In this lecture we continue our discussion of group-theoretic varieties. We describe why obstructions should exist, and why the exceptional group-theoretic nature of the class varieties is crucial for this existence.

Recall

Let $G = GL_n(\mathbb{C})$, V a G-module, and $\mathbb{P}(V)$ the associated projective space. Let $v \in \mathbb{P}(V)$ be a point characterized by its stabilizer $H = G_v \subset G$. In other words, v is the only point in $\mathbb{P}(V)$ stabilized by H. Then $\Delta_V[v] = \overline{Gv}$ is called a *group-theoretic variety* because it is completely determined by the group triple

$$H \hookrightarrow G \hookrightarrow GL(V).$$

The simplest example of a group-theoretic variety is a variety of the form G/P that we described in the earlier lecture. Let $V = V_\lambda(G)$ be a Weyl module of G and $v_\lambda \in \mathbb{P}(V)$ the highest weight point (recall: the unique point stabilized by the Borel subgroup $B \subset G$ of lower triangular matrices). Then the stabilizer of v_λ consists of block-upper triangular matrices, where the block sizes are determined by λ:

$$P_\lambda := G_{v_\lambda} = \left(\begin{array}{ccc|cc|ccc} * & * & * & * & * & * & * & * \\ * & * & * & * & * & * & * & * \\ * & * & * & * & * & * & * & * \\ \hline 0 & 0 & 0 & * & * & * & * & * \\ 0 & 0 & 0 & * & * & * & * & * \\ \hline 0 & 0 & 0 & 0 & 0 & * & * & * \\ 0 & 0 & 0 & 0 & 0 & * & * & * \\ 0 & 0 & 0 & 0 & 0 & * & * & * \end{array}\right)$$

The orbit $\Delta_V[v_\lambda] = Gv_\lambda \cong G/P_\lambda$ is a group-theoretic variety determined entirely by the triple

$$P_\lambda = G_{v_\lambda} \hookrightarrow G \hookrightarrow K = GL(V).$$

The group-theoretic varieties of main interest in GCT are the class varieties associated with the various complexity classes.

14.1 Representation theoretic data

The **main principle** guiding GCT is that the algebraic geometry of a group-theoretic variety ought to be completely determined by the representation theory of the corresponding group triple. This is a natural extension of work already pursued in mathematics by Deligne and Milne on Tannakien categories [DeM], showing that an algebraic group is completely determined by its representation theory. So the goal is to associate to a group-theoretic variety some representation-theoretic data that will analogously capture the information in the variety completely. We shall now illustrate this for the class variety for NC. First a few definitions.

Let $v \in \mathbb{P}(V)$ be the point as above characterized by its stabilizer G_v. This means the line $\mathbb{C}v \subseteq V$ corresponding to v is a one-dimensional representation of G_v. Thus $(\mathbb{C}v)^{\otimes d}$ is a one-dimensional degree d representation, i.e. the representation $\rho : G \to GL(\mathbb{C}v) \cong \mathbb{C}^*$ is polynomial of degree d in the entries of the matrix of an element in G. Recall that $\mathbb{C}[V]$ is the coordinate ring of V, and $\mathbb{C}[V]_d$ is its degree d homogeneous component, so $(\mathbb{C}v)^{\otimes d} \subseteq \mathbb{C}[V]_d$.

To each $v \in \mathbb{P}(V)$ that is characterized by its stabilizer, we associate a representation-theoretic data, which is the set of G-modules

$$\Pi_v = \bigcup_d \Pi_v(d),$$

where $\Pi_v(d)$ is the set of all irreducible G-submodules S of $\mathbb{C}[V]_d$ whose duals S^* do *not* contain a G_v-submodule isomorphic to $(\mathbb{C}v)^{\otimes d *}$ (the dual of $(\mathbb{C}v)^{\otimes d}$). The following proposition elucidates the importance of this data:

Proposition 14.1. *[GCT2] $\Pi_v \subseteq I_V[v]$ (where $I_V[v]$ is the ideal of the projective variety $\Delta_V[v] \subseteq \mathbb{P}(V)$).*

Proof. Fix $S \in \Pi_v(d)$. Suppose, for the sake of contradiction, that $S \nsubseteq I_V[v]$. Since $S \subseteq \mathbb{C}[V]$, S consists of "functions" on the variety $\mathbb{P}(V)$ (actually homogeneous polynomials on V). The coordinate ring of $\Delta_V[v]$ is $\mathbb{C}[V]/I_V[v]$, and since $S \nsubseteq I_V[v]$, S must not vanish identically on $\Delta_V[v]$. Since the orbit Gv is dense in $\Delta_V[v]$, S must not vanish identically on this single orbit Gv. Since S is a G-module, if S were to vanish identically on the line $\mathbb{C}v$, then it would vanish on the entire orbit Gv, so S does not vanish identically on $\mathbb{C}v$.

Now S consists of functions of degree d. Restrict them to the line $\mathbb{C}v$. The dual of this restriction gives an injection of $(\mathbb{C}v)^{\otimes d *}$ as a G_v-submodule of S^*, contradicting the definition of $\Pi_v(d)$. $\square$

14.2 The second fundamental theorem

We now ask essentially the reverse question: when does the representation theoretic data Π_v generate the ideal $I_V[v]$? For if Π_v generates $I_V[v]$, then Π_v completely captures the coordinate ring $\mathbb{C}[V]/I_V[v]$, and hence the variety $\Delta_V[v]$.

Theorem 14.1 (Second fundamental theorem of invariant theory for G/P)**.** *The G-modules in $\Pi_{v_\lambda}(2)$ generate the ideal $I_V[v_\lambda]$ of the orbit $Gv_\lambda \cong G/P_\lambda$, when $V = V_\lambda(G)$.*

This theorem justifies the main principle for G/P, so we can hope that similar results hold for the class varieties in GCT (though not always exactly in the same form).

Now, let $\Delta_V[g]$ be the class variety for **NC** (in other words, take $g = \det(Y)$ for a matrix Y of indeterminates). Based on the main principle, we have the following conjecture, which essentially generalizes the second fundamental theorem of invariant theory for G/P to the class variety for **NC**:

Conjecture 14.1 (GCT 2)**.** $\Delta_V[g] = X(\Pi_g)$ *where $X(\Pi_g)$ is the zero-set of all forms in the G-modules contained in Π_g.*

Theorem 14.2 (GCT 2)**.** *A weaker version of the above conjecture holds. Specifically, assuming that the Kronecker coefficients satisfy a certain separation property, there exists a G-invariant (Zariski) open neighbourhood $U \subseteq P(V)$ of the orbit Gg such that $X(\Pi_g) \cap U = \Delta_V[g] \cap U$.*

There is a notion of algebro-geometric complexity called *Luna-Vust complexity* which quantifies the gap between G/P and class varieties. The Luna-Vust complexity of G/P is 0. The Luna-Vust complexity of the **NC** class variety is $\Omega(\dim(Y))$. This is analogous to the difference between circuits of constant depth and circuits of superpolynomial depth. This is why the previous conjecture and theorem turn out to be far harder than the corresponding facts for G/P.

14.3 Why should obstructions exist?

The following proposition explains why obstructions should exist to separate **NC** from $P^{\#P}$.

Proposition 14.2 (GCT 2)**.** *Let $g = \det(Y)$, $h = perm(X)$, $f = \phi(h)$, $n = \dim(X)$, $m = \dim(Y)$. If Conjecture 14.1 holds and the permanent cannot be approximated arbitrarily closely by circuits of poly-logarithmic depth (hardness assumption), then an obstruction for the pair (f, g) exists for all large enough n, when $m = 2^{\log^c n}$ for some constant c. Hence, under these conditions, $\boldsymbol{NC} \neq P^{\#P}$ over $\mathbb{C}$.*

This proposition may seem a bit circular at first, since it relies on a hardness assumption. But we do not plan to prove the existence of obstructions by proving the assumptions of this proposition. Rather, this proposition should be taken as evidence that obstructions exist (since we expect the hardness assumption therein to hold, given that the permanent is # **P**-complete), and we will develop other methods to prove their existence.

Proof. The hardness assumption implies that $f \notin \Delta_V[g]$ if $m = 2^{\log^c n}$ [GCT 1].

Conjecture 14.1 says that $X(\Pi_g) = \Delta_V[g]$. So there exists an irreducible G-module $S \in \Pi_g$ such that S does not vanish on f. So S occurs in $R_V[f]$ as a G-submodule.

On the other hand, since $S \in \Pi_g$, $S \subseteq I_V[g]$ by Proposition 14.1. So S does not occur in $R_V[g] = \mathbb{C}[V]/I_V[g]$. Thus S is not a G-submodule of $R_V[g]$, but it is a G-submodule of $R_V[f]$, i.e., S is an obstruction. ☐

Chapter 15

The flip

Scribe: Hariharan Narayanan

Goal: Describe the basic principle of GCT, called the flip, in the context of the NC vs. $P^{\#P}$ problem over $\mathbb{C}$.

references: [GCTflip1, GCT1, GCT2, GCT6]

Recall

As in the previous lectures, let $g = det(Y) \in P(V)$, Y an $m \times m$ variable matrix, $G = GL_{m^2}(\mathbb{C})$, and $\Delta_V(g) = \Delta_V[g; m] = \overline{G}g \subseteq P(V)$ the class variety for NC. Let $h = perm(X)$, X an $n \times n$ variable matrix, $f = \phi(h) = y^{m-n}h \in P(V)$, and $\Delta_V(f) = \Delta_V[f; m, n] = \overline{G}f \subseteq P(V)$ the class variety for $P^{\#P}$. Let $R_V[f; m, n]$ denote the homogeneous coordinate ring of $\Delta_V[f; m, n]$, $R_V[g; m]$ the homogeneous coordinate ring of $\Delta_V[g; m]$, and $R_V[f; m, n]_d$ and $R_V[g; m]_d$ their degree d-components. A Weyl module $S = V_\lambda(G)$ of G is an obstruction of degree d for the pair (f, g) if V_λ occurs in $R_V[f; m, n]_d$ but not $R_V[g; m]_d$.

Conjecture 15.1. *[GCT2] An obstruction (of degree polynomial in m) exists if $m = 2^{polylog(n)}$ as $n \to \infty$.*

This implies $NC \neq P^{\#P}$ over $\mathbb{C}$.

15.1 The flip

In this lecture we describe an approach to prove the existence of such obstructions. It is based on the following complexity theoretic positivity hypothesis:

PHflip [GCTflip1]:

1. Given n, m and d, whether an obstruction of degree d for m and n exists can be decided in $poly(n, m, \langle d \rangle)$ time, and if it exists, the label λ of such an obstruction can be constructed in $poly(n, m, \langle d \rangle)$ time. Here $\langle d \rangle$ denotes the bitlength of d.

2. (a) Whether V_λ occurs in $R_V[f; m, n]_d$ can be decided in $poly(n, m, \langle d \rangle, \langle \lambda \rangle)$ time.

   (b) Whether V_λ occurs in $R_V[g; m]_d$ can be decided in $poly(n, m, \langle d \rangle, \langle \lambda \rangle)$ time.

This suggests the following approach for proving Conjecture 15.1:

1. Find polynomial time algorithms sought in PHflip-2 for the basic decision problems (a) and (b) therein.

2. Using these find a polynomial time algorithm sought in PHflip-1 for deciding if an obstruction exists.

3. Transform (the techniques underlying) this "easy" (polynomial time) algorithm for deciding if an obstruction exists for given n and m into an "easy" (i.e., feasible) proof of existence of an obstruction for every $n \to \infty$ when d is large enough and $m = 2^{\text{polylog}(n)}$.

The first step here is the crux of the matter. The main results of [GCT6] say that the polynomial time algorithms for the basic decision problems as sought in PHflip-2 indeed exist assuming natural analogues of PH1 and SH (PH2) that we have seen earlier in the context of the plethysm problem. To state them, we need some definitions.

Let $S_d^\lambda[f] = S_d^\lambda[f; m, n]$ be the multiplicity of $V_\lambda = V_\lambda(G)$ in $R_V[f; m, n]$. The stretching function $\tilde{S}_d[f] = \tilde{S}_d^\lambda[f; m, n]$ is defined by

$$\tilde{S}_d^\lambda[f](k) := S_{kd}^{k\lambda}[f].$$

The stretching function for g, $\tilde{S}_d^\lambda[g] = \tilde{S}_d^\lambda[g; m]$, is defined analogously.

The main mathematical result of [GCT6] is:

Theorem 15.1. *[GCT6] The stretching functions $\tilde{S}_d^\lambda[g]$ and $\tilde{S}_d^\lambda[f]$ are quasipolynomials assuming that the singularities of $\Delta_V[f; m, n]$ and $\Delta_V[g; m]$ are rational.*

Here rational means "nice"; we shall not worry about the exact definition.

The main complexity-theoretic result is:

Theorem 15.2. *[GCT6] Assuming the following mathematical positivity hypothesis $PH1$ and the saturation hypothesis SH (or the stronger positivity hypothesis $PH2$), PHflip-2 holds.*

PH1: There exists a polytope $P = P_d^\lambda[f]$ such that

1. The Ehrhart quasi-polynomial of P, $f_P(k)$, is $\tilde{S}_d^\lambda[f](k)$.
2. $dim(P) = poly(n, m, \langle d \rangle)$.
3. Membership in P can be answered in polynomial time.
4. There is a polynomial time separation oracle [GLS] for P.

Similarly, there exists a polytope $Q = Q_d^\lambda[g]$ such that

1. The Ehrhart quasi-polynomial of Q, $f_Q(k)$, is $\tilde{S}_d^\lambda[g](k)$.
2. $dim(Q) = poly(m, \langle d \rangle)$.
3. Membership in Q can be answered in polynomial time.
4. There is a polynomial time separation oracle for Q.

PH2: The quasi-polynomials $\tilde{S}_d^\lambda[g]$ and $\tilde{S}_d^\lambda[f]$ are positive.

This implies:

SH: The quasi-polynomials $\tilde{S}_d^\lambda[g]$ and $\tilde{S}_d^\lambda[f]$ are saturated.

PH1 and SH imply that the decision problems in PHflip-2 can be transformed into saturated positive integer programming problems. Hence Theorem 15.2 follows from the polynomial time algorithm for saturated linear programming that we described in an earlier class.

The decision problems in PHflip-2 are "hyped" up versions of the plethysm problem discussed earlier. The article [GCT6] provides evidence for $PH1$

and $PH2$ for the plethysm problem. This constitutes the main evidence for PH1 and PH2 for the class varieties in view of their group-theoretic nature; cf. [GCTflip1].

The following problem is important in the context of PHflip-2:

Problem 15.1. *Understand the G-module structure of the homogeneous coordinate rings $R_V[f]_d$ and $R_V[g]_d$.*

This is an instance of the following abstract:

Problem 15.2. *Let X be a projective group-theoretic G-variety. Let $R = \bigoplus_{d=0}^{\infty} R_d$ be its homogeneous coordinate ring. Understand the G-module structure of R_d.*

The simplest group-theoretic variety is G/P. For it, a solution to this abstract problem is given by the following results:

1. The Borel-Weil theorem.

2. The Second Fundamental theorem of invariant theory [SFT].

These will be covered in the next class for the simplest case of G/P, the Grassmanian.

Chapter 16

The Grassmanian

Scribe: Hariharan Narayanan

Goal: The Borel-Weil and the second fundamental theorem of invariant theory for the Grassmanian.
Reference: [F]

Recall

Let $V = V_\lambda(G)$ be a Weyl module of $G = GL_n(\mathbb{C})$ and $v_\lambda \in P(V)$ the point corresponding to its highest weight vector. The orbit $\Delta_V[v_\lambda] := Gv_\lambda$, which is already closed, is of the form G/P, where P is the parabolic stabilizer of v_λ. When λ is a single column, it is called the *Grassmannian*.

An alternative description of the Grassmanian is as follows. Assume that λ is a single column of length d. Let Z be a $d \times n$ matrix of variables z_{ij}. Then $V = V_\lambda(G)$ can be identified with the span of $d \times d$ minors of Z with the action of $\sigma \in G$ given by:

$$\sigma : f(z) \mapsto f(z\sigma).$$

Let Gr_d^n be the space of all d-dimensional subspaces of $\mathbb{C}^n$. Let W be a d-dimensional subspace of $\mathbb{C}^n$. Let $B = B(W)$ be a basis of W. Construct the $d \times n$ matrix z_B, whose rows are vectors in B. Consider the Plücker map from Gr_d^n to $P(V)$ which maps any $W \in Gr_d^n$ to the tuple of $d \times d$ minors of Z_B. Here the choice of $B = B(W)$ does not matter, since any choice gives the same point in $P(V)$. Then the image of Gr_d^n is precisely the Grassmanian $Gv_\lambda \subseteq P(V)$.

16.1 The second fundamental theorem

Now we ask:

Question 16.1. *What is the ideal of $Gr_d^n \approx Gv_\lambda \subseteq P(V)$?*

The homogeneous coordinate ring of $P(V)$ is $\mathbb{C}[V]$. We want an explicit set of generators of this ideal in $\mathbb{C}[V]$. This is given by the second fundamental theorem of invariant theory, which we describe next.

The coordinates of $P(V)$ are in one-to-one correspondence with the $d \times d$ minors of the matrix Z. Let each minor of Z be indexed by its columns. Thus for $1 \leq i_1 < \cdots < i_d \leq n$, $Z_{i_1,\ldots,i_d}$ is a coordinate of $P(V)$ corresponding to the minor of Z formed by the columns $i_1, i_2, \ldots$. Let $\Lambda(n,d)$ be the set of ordered d-tuples of $\{1,\ldots,n\}$. The tuple $[i_1,\ldots,i_d]$ in this set will be identified with the coordinate $Z_{i_1,\ldots,i_d}$ of $P(V)$. There is a bijection between the elements of $\Lambda(n,d)$ and of $\Lambda(n,n-d)$ obtained by associating complementary sets:

$$\Lambda(n,d) \ni \lambda \leftrightsquigarrow \lambda^* \in \Lambda(n,n-d).$$

We define $sgn(\lambda,\lambda^*)$ to be the sign of the permutation that takes $[1,\ldots,n]$ to $[\lambda_1,\ldots,\lambda_d,\lambda_1^*,\ldots,\lambda_{n-d}^*]$.

Given $s \in \{1\ldots,d\}$, $\alpha \in \Lambda(n,s-1)$, $\beta \in \Lambda(n,d+1)$, and $\gamma \in \Lambda(n,d-s)$, we now define the Van der Waerden Syzygy $[[\alpha,\beta,\gamma]]$, which is an element of the degree two component $\mathbb{C}[V]_2$ of $\mathbb{C}[V]$, as follows:

$$[[\alpha,\beta,\gamma]] = \sum_{\tau\in\Lambda(d+1,s)} sgn(\tau,\tau^*)[\alpha_1,\ldots,\alpha_{s-1},\beta_{\tau_1^*},\ldots,\beta_{\tau_{d+1-s}^*}][\beta_{\tau_1},\ldots,\beta_{\tau_s},\gamma_1,\ldots,\gamma_{d-s}].$$

It is easy to show that this syzygy vanishes on the Grassmanian Gr_d^n: because it is an alternating $(d+1)$-multilinear-form, and hence has to vanish on any d-dimensional space $W \in Gr_d^n$. Thus it belongs to the ideal of the Grassmanian. Moreover:

Theorem 16.1 (Second fundamental theorem)**.** *The ideal of the Grassmanian Gr_d^n is generated by the Van-der-Waerden syzygies.*

An alternative formulation of this result is as follows. Let $P_\lambda \subseteq G$ be the stabilizer of v_λ. Let $\Pi_{v_\lambda}(2)$ be the set of irreducible G-submodules of $\mathbb{C}[V]_2$ whose duals do not contain a P_λ-submodule isomorphic to $\mathbb{C}v_\lambda^{\otimes 2*}$ (the dual of $\mathbb{C}v_\lambda^{\otimes 2}$). Here $\mathbb{C}v_\lambda$ denotes the line in $P(V)$ corresponding to v_λ, which is a one-dimensional representation of P_λ since it stabilizes $v_\lambda \in P(V)$. It can be shown that the span of the G-modules in $\Pi_{v_\lambda}(2)$ is equal to the span of the Van-der-Waerden syzygies. Hence, Theorem 16.1 is equivalent to:

Theorem 16.2 (Second Fundamental Theorem(SFT))**.** *The G-modules in $\Pi_{v_\lambda}(2)$ generate the ideal of Gr^n_d.*

This formulation of SFT for the Grassmanian looks very similar to the generalized conjectural SFT for the NC-class variety described in the earlier class. This indicates that the class varieties in GCT are "qualitatively similar" to G/P.

16.2 The Borel-Weil theorem

We now describe the G-module structure of the homogeneous coordinate ring R of the Grassmannian $Gv_\lambda \subseteq P(V)$, where λ is a single column of height d. The goal is to give an explicit basis for R. Let R_s be the degree s component of R. Corresponding to any numbering T of the shape $s\lambda$, which is a $d \times s$ rectangle, whose columns have strictly increasing elements top to bottom, we have a monomial $m_T = \prod_c Z_c \in \mathbb{C}[V]_s$, were Z_c is the coordinate of $P(V)$ indexed by the d-tuple c, and c ranges over the s columns of T. We say that m_T is (semi)-standard if the rows of T are nondecreasing, when read left to right. It is called nonstandard otherwise.

Lemma 16.1 (Straightening Lemma)**.** *Each non-standard m_T can be straightened to a normal form, as a linear combination of standard monomials, by using Van der Waerden Syzygies as straightening relations (rewriting rules).*

For any numbering T as above, express m_T in a normal form as per the lemma:

$$m_T = \sum_{\text{(Semi)-Standard Tableau } S} \alpha(S,T), m_S$$

where $\alpha(S,T) \in \mathbb{C}$.

Theorem 16.3 (Borel-Weil Theorem for Grassmannians)**.** *Standard monomials $\{m_T\}$ form a basis of R_s, where T ranges over all semi-standard tableaux of rectangular shape $s\lambda$. Hence, $R_s \cong V^*_{s\lambda}$, the dual of the Weyl module $V_{s\lambda}$.*

This gives the G-module structure of R completely. It follows that the problem of deciding if $V_\beta(G)$ occurs in R_s can be solved in polynomial time: this is so if and only if $(s\lambda)^* = \beta$, where $(s\lambda)^*$ denotes the dual partition, whose description is left as an exercise.

The second fundamental theorem as well as the Borel-Weil theorem easily follow from the straightening lemma and linear independence of the standard monomials (as functions on the Grassmanian).

Chapter 17

Quantum group: basic definitions

Scribe: Paolo Codenotti

Goal: The basic plan to implement the flip in [GCT6] is to prove PH1 and SH via the theory of quantum groups. We introduce the basic concepts in this theory in this and the next two lectures, and briefly show their relevance in the context of PH1 in the final lecture.
Reference: [KS]

17.1 Hopf Algebras

Let G be a group, and $K[G]$ the ring of functions on G with values in the field K, which will be $\mathbb{C}$ in our applications. The group G is defined by the following operations:

- multiplication: $G \times G \to G$,
- identity e: $e \to G$,
- inverse: $G \to G$.

In order for G to be a group, the following properties have to hold:

- $eg = ge = g$,
- $g_1(g_2g_3) = (g_1g_2)g_3$,
- $g^{-1}g = gg^{-1} = e$.

We now want to translate these properties to properties of $K[G]$. This should be possible since $K[G]$ contains all the information that G has. In other words, we want to translate the notion of a group in terms of $K[G]$. This translate is called a Hopf algebra. Thus if G is a group, $K[G]$ is a Hopf algebra. Let us first define the dual operations.

- Multiplication is a map:

  $$\cdot : G \times G \to G.$$

  So co-multiplication Δ will be a map as follows:

  $$K[G \times G] = K[G] \otimes K[G] \leftarrow K[G].$$

  We want Δ to be the pullback of multiplication. So for a given $f \in K[G]$ we define $\Delta(f) \in K[G] \otimes K[G]$ by:

  $$\Delta(f)(g_1, g_2) = f(g_1 g_2).$$

  Pictorially:

  $$\begin{tikzcd} G \times G \arrow[r, "\cdot"] \arrow[d, "\Delta(f)"'] & G \arrow[d, "f"] \\ k \arrow[r, equal] & k \end{tikzcd}$$

- The unit is a map:

  $$e \to G.$$

  Therefore we want the co-unit ϵ to be a map:

  $$K \xleftarrow{\epsilon} K[G],$$

  defined by: for $f \in K[G]$, $\epsilon(f) = f(e)$.

- Inverse is a map:

  $$(\ )^{-1} : G \to G.$$

  We want the dual antipode S to be the map:

  $$K[G] \leftarrow K[G]$$

  defined by: for $f \in K[G]$, $S(f)(g) = f(g^{-1})$.

The following are the abstract axioms satisfied by Δ, ϵ and S.

1. Δ and ϵ are algebra homomorphisms.

$$\Delta : K[G] \to K[G] \otimes K[G]$$

$$\epsilon : K[G] \to K.$$

2. co-associativity: Associativity is defined so that the following diagram commutes:

$$\begin{tikzcd}
G \times G \times G \arrow[r, equal] \arrow[d, "\cdot \times \mathrm{id}"'] & G \times G \times G \arrow[d, "\mathrm{id} \times \cdot"] \\
G \times G \arrow[d, "\cdot"'] & G \times G \arrow[d, "\cdot"] \\
G \arrow[r, equal] & G
\end{tikzcd}$$

Similarly, we define co-associativity so that the following dual diagram commutes:

$$\begin{tikzcd}
K[G] \otimes K[G] \otimes K[G] \arrow[r, equal] & K[G] \otimes K[G] \otimes K[G] \\
K[G] \otimes K[G] \arrow[u, "\Delta \otimes \mathrm{id}"] & K[G] \otimes K[G] \arrow[u, "\mathrm{id} \otimes \Delta"'] \\
K[G] \arrow[u, "\Delta"] \arrow[r, equal] & K[G] \arrow[u, "\Delta"']
\end{tikzcd}$$

Therefore co-associativity says:

$$(\Delta \otimes \mathrm{id}) \circ \Delta = (\mathrm{id} \otimes \Delta) \circ \Delta.$$

3. The property $ge = g$ is defined so that the following diagram commutes:

$$\begin{tikzcd}
e \times G \arrow[r, equal] \arrow[d, "e \times \mathrm{id}"'] & G \arrow[d] \\
G \times G \arrow[d, "\cdot"'] & \mathrm{id} \arrow[d] \\
G \arrow[r, equal] & G
\end{tikzcd}$$

We define the co of this property so that the following diagram commutes:

$$
\begin{tikzcd}
K & \times & K[G] \arrow[rr, equal] & & K[G] \\
K[G] \arrow[u, "\epsilon"] & \times & K[G] \arrow[u, "\mathrm{id}"'] & & \mathrm{id} \arrow[u] \\
& K[G] \arrow[u, "\Delta"'] \arrow[rrr, equal] & & & K[G] \arrow[u]
\end{tikzcd}
$$

That is, $\mathrm{id} = (\epsilon \otimes \mathrm{id}) \circ \Delta$. Similarly, $ge = g$ translates to: $\mathrm{id} = (\mathrm{id} \otimes \epsilon) \circ \Delta$. Therefore we get

$$\mathrm{id} = (\epsilon \otimes \mathrm{id}) \circ \Delta = (\mathrm{id} \otimes \epsilon) \circ \Delta.$$

4. The last property is $gg^{-1} = e = g^{-1}g$. The first equality is equivalent to requiring that the following diagram commute:

$$
\begin{tikzcd}
G \arrow[r, equal] \arrow[d, "\mathrm{diag}"'] & G \arrow[dd] \\
G \times G \arrow[d, shift right=2, "()^{-1}"'] \arrow[d, shift left=2, "\mathrm{id}"] & \\
G \times G \arrow[d, "\cdot"] & e \arrow[d] \\
G \arrow[r, equal] & G
\end{tikzcd}
$$

Where $\mathrm{diag} : G \to G \times G$ is the diagonal embedding. The co of diag is $m : K[G] \leftarrow K[G] \otimes K[G]$ defined by $m(f_1, f_2)(g) = f_1(g) \cdot f_2(g)$. So the co of this property will hold when the following diagram commutes:

$$
\begin{tikzcd}
 & K[G] \arrow[rrr, equal] & & & K[G] \\
K[G] \arrow[r, phantom, "\otimes"] \arrow[ur, "m"] & {} & k[G] & & K \arrow[u, "\nu"] \\
K[G] \arrow[u, "S"] \arrow[r, phantom, "\otimes"] & {} & K[G] \arrow[u, "\mathrm{id}"'] & & \\
 & K[G] \arrow[u, "\Delta"'] \arrow[rrr, equal] & & & K[G] \arrow[uu, "\epsilon"]
\end{tikzcd}
$$

Where ν is the embedding of K into $K[G]$. Therefore the last property we want to be satisfied is:

$$m \circ (S \otimes \mathrm{id}) \circ \Delta = \nu \circ \epsilon.$$

For $e = g^{-1}g$, we similarly get:

$$m \circ (\mathrm{id} \otimes S) \circ \Delta = \nu \circ \epsilon.$$

Definition 17.1 (Hopf algebra)**.** *A K-algebra A is called a* Hopf algebra *if there exist homomorphisms $\Delta : A \otimes A \to A$, $S : A \to A$, $\epsilon : A \to K$, and $\nu : A \to K$ that satisfy* (1) − (4) *above, with A in place of $K[G]$.*

We have shown that if G is a group, the ring $K[G]$ of functions on G is a (commutative) Hopf algebra, which is non-co-commutative if G is non-commutative. Thus for every usual group, we get a commutative Hopf algebra. However, in general, Hopf algebras may be non-commutative.

Definition 17.2. *A* quantum group *is a (non-commutative and non-co-commutative) Hopf algebra.*

A nontrivial example of a quantum group will be constructed in the next lecture.

Next we want to look at what happens to group theoretic notions such as representations, actions, and homomorphisms, in the context of Hopf algebras. These will correspond to co-representations, co-actions, and co-homomorphisms.

Let us look closely at the notion of co-representation. A representation is a map $\cdot : G \times V \to V$, such that

- $(h_1 h_2) \cdot v = h_1 \cdot (h_2 \cdot v)$, and

- $e \cdot v = v$.

Therefore a (right) co-representation of A will be a linear mapping $\varphi : V \to V \otimes A$, where V is a K-vector space, and φ satisfies the following:

- The following diagram commutes:

$$\begin{tikzcd} V \otimes A \otimes A & V \otimes A \arrow[l, "\mathrm{id} \otimes \Delta"'] \\ V \otimes A \arrow[u, "\varphi \otimes \mathrm{id}"] & V \arrow[l, "\varphi"] \arrow[u, "\varphi"'] \end{tikzcd}$$

  That is, the following equality holds:

$$(\varphi \otimes \mathrm{id}) \circ \varphi = (\mathrm{id} \otimes \Delta) \circ \varphi.$$

- The following diagram commutes:

$$\begin{tikzcd} V \otimes K & V \otimes K \arrow[l, "\mathrm{id}"'] \\ V \otimes A \arrow[u, "\mathrm{id} \otimes \epsilon"] & V \arrow[l, "\varphi"] \arrow[u, equal] \end{tikzcd}$$

  That is, the following equality holds:

$$(\mathrm{id} \otimes \epsilon) \circ \varphi = \mathrm{id}$$

In fact all usual group theoretic notions can be "Hopfified" in this sense [exercise].

Let us look now at an example. Let

$$G = GL_n(\mathbb{C}) = GL(\mathbb{C}^n) = GL(V),$$

where $V = \mathbb{C}^n$. Let M_n be the matrix space of $n \times n$ $\mathbb{C}$-matrices, and $\mathcal{O}(M_n)$ the coordinate ring of M_n,

$$\mathcal{O}(M_n) = \mathbb{C}[U] = C[\{u^i_j\}],$$

where U is an $n \times n$ variable matrix with entries u^i_j. Let $\mathbb{C}[G] = \mathcal{O}(G)$ be the coordinate ring of G obtained by adjoining $det(U)^{-1}$ to $\mathcal{O}(M_n)$. That is, $\mathbb{C}[G] = \mathcal{O}(G) = \mathbb{C}[U][\det(U)^{-1}]$, which is the $\mathbb{C}$ algebra generated by u^i_j's and $\det(U)^{-1}$.

Proposition 17.1. *$\mathbb{C}[G]$ is a Hopf algebra, with Δ, ϵ, and S as follows.*

- *Recall that the axioms of a Hopf algebra require that*

$$\Delta : \mathbb{C}[G] \to \mathbb{C}[G] \otimes \mathbb{C}[G],$$

$$\Delta(f)(g_1, g_2) = f(g_1 g_2).$$

  Therefore we define

$$\Delta(u^i_j) = \sum_k u^i_k \otimes u^k_j,$$

  where U denotes the generic matrix in M_n as above.

- *Again, it is required that*

$$\epsilon(f) = f(e).$$

  Therefore we define

$$\epsilon(u^i_j) = \delta_{ij},$$

  where δ_{ij} is the Kronecker delta function.

- *Finally, the antipode is required to satisfy $S(f)(g) = f(g^{-1})$. Let $\widetilde{U}$ be the cofactor matrix of U, $U^{-1} = \frac{1}{\det(U)}\widetilde{U}$, and $\widetilde{u}^i_j$ the entries of $\widetilde{U}$. Then we define S by:*

$$S(u^i_j) = \frac{1}{\det(U)} \widetilde{u}^i_j = (U^{-1})^i_j.$$

Chapter 18

Standard quantum group

Scribe: Paolo Codenotti

Goal: In this lecture we construct the standard (Drinfeld-Jimbo) quantum group, which is a q-deformation of the general linear group $GL_{\mathbf{n}}(\mathbb{C})$ with remarkable properties.
Reference: [KS]

Let $G = GL(V) = GL(\mathbb{C}^n)$, and $V = \mathbb{C}^n$. In the earlier lecture, we constructed the commutative and non co-commutative Hopf algebra $\mathbb{C}[G]$. In this lecture we quantize $\mathbb{C}[G]$ to get a non-commutative and non-co-commutative Hopf algebra $\mathbb{C}_q[G]$, and then define the standard quantum group $G_q = GL_q(V) = GL_q(n)$ as the virtual object whose coordinate ring is $\mathbb{C}_q[G]$.

We start by defining $GL_q(2)$ and $SL_q(2)$, for $n = 2$. Then we will generalize this construction to arbitrary n. Let $\mathcal{O}(M_2)$ be the coordinate ring of M_2, the set of 2×2 complex matrices, $\mathbb{C}[V]$ the coordinate ring of V generated by the coordinates x_1 and x_2 of V which satisfy $x_1x_2 = x_2x_1$. Let

$$U = \begin{bmatrix} \mathrm{a} & \mathrm{b} \\ \mathrm{c} & \mathrm{d} \end{bmatrix}$$

be the generic (variable) matrix in M_2. It acts on $V = \mathbb{C}^2$ from the left and from the right. Let

$$x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}.$$

The left action is defined by

$$x \to x' := Ux.$$

Let

$$x' = \begin{bmatrix} x_1' \\ x_2' \end{bmatrix}.$$

Similarly, the right action is defined by

$$x^T \to (x'')^T := x^T U.$$

Let

$$x'' = \begin{bmatrix} x_1'' \\ x_2'' \end{bmatrix}.$$

The action of M_2 on V satisfies

$$x_1' x_2' = x_2' x_1', \text{ and}$$

$$x_1'' x_2'' = x_2'' x_1''.$$

Now instead of V, we take its q-deformation V_q, a quantum space, whose coordinates x_1 and x_2 satisfy

$$x_1 x_2 = q x_2 x_1, \tag{18.1}$$

where $q \in \mathbb{C}$ is a parameter. Intuitively, in quantum physics if x_1 and x_2 are position and momentum, then $q = e^{i\hbar}$ when $\hbar$ is Planck's constant. Let $\mathbb{C}_q[V]$ be the ring generated by x_1 and x_2 with the relation (18.1). That is,

$$\mathbb{C}_q[V] = \mathbb{C}[x_1, x_2] / < x_1 x_2 - q x_2 x_1 > .$$

It is the coordinate ring of the quantum space V_q. Now we want to quantize $M(2)$ to get $M_q(2)$, the space of quantum 2×2 matrices, and $GL(2)$ to $GL_q(2)$, the space of quantum 2×2 nonsingular matrices. Intuitively, $M_q(2)$ is the space of linear transformations of the quantum space V_q which preserve the equation (18.1) under the left and right actions, and similarly, $GL_q(2)$ is the space of non-singular linear transformation that preserve the equation (18.1) under the left and right actions. We now formalize this intuition.

Let $U = \begin{pmatrix} \mathrm{a} & \mathrm{b} \\ \mathrm{c} & \mathrm{d} \end{pmatrix}$ be a quantum matrix whose coordinates do not commute. The left and right actions of U must preserve 18.1.

[Left action:] Let the left action be $\varphi_L : x \to Ux$, and $Ux = x'$. Then we must have:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} ax_1 + bx_2 \\ cx_1 + dx_2 \end{pmatrix} = \begin{pmatrix} x_1' \\ x_2' \end{pmatrix}.$$

[Right action:] Let the right action be $\varphi_R : x^T \to x^T U$, and let $x'' = (x^T U)^T = U^T x$. Then we must have:

$$\begin{pmatrix} x_1 & x_2 \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} ax_1 + cx_2 \\ bx_1 + dx_2 \end{pmatrix} = \begin{pmatrix} x_1'' \\ x_2'' \end{pmatrix}.$$

The preservation of $x_1 x_2 = q x_2 x_1$ under left multiplication means

$$x_1' x_2' = q x_2' x_1'.$$

That is,

$$(ax_1 + bx_2)(cx_1 + dx_2) = q(cx_1 + dx_2)(ax_1 + bx_2). \tag{18.2}$$

The left hand side of (18.2) is

$$acx_1^2 + bcx_2x_1 + adx_1x_2 + bdx_2^2 = acx_1^2 + (bc + adq)x_2x_1 + bdx_2^2.$$

Similarly, the right hand side of (18.2) is

$$q(cax_1^2 + (da + cbq)x_2x_1 + bdx_2^2).$$

Therefore equation (18.2) implies:

$$\begin{aligned} ac &= qca \\ bd &= qdb \\ bc + adq &= da + qcb. \end{aligned}$$

That is,

$$\begin{aligned} ac &= qca \\ bd &= qdb \\ ad - da - qcb + q^{-1}bc &= 0. \end{aligned}$$

Similarly, since $x_1'' x_2'' = q x_2'' x_1''$, we get:

$$\begin{aligned} ab &= qba \\ cd &= qdc \\ ad - da - qbc + q^{-1}cb &= 0. \end{aligned}$$

The last equations from each of these sets imply $bc = cb$.

So we define $\mathcal{O}(M_q(2))$, the coordinate ring of the space of 2×2 quantum matrices $M_q(2)$, to be the $\mathbb{C}$-algebra with generators a, b, c, and d, satisfying the relations:

$$ab = qba, \quad ac = qca, \quad bd = qdb, \quad cd = qdc,$$
$$bc = cb, \quad ad - da = (q - q^{-1})bc.$$

Let

$$U = \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} u_1^1 & u_2^1 \\ u_1^2 & u_2^2 \end{pmatrix}.$$

Define the quantum determinant of U to be

$$D_q = \det(U) = ad - qbc = da - q^{-1}bc.$$

Define $\mathbb{C}_q[G] = \mathcal{O}(GL_q(2))$, the coordinate ring of the virtual quantum group $GL_q(2)$ of invertible 2×2 quantum matrices, to be

$$\mathcal{O}(GL_q(2)) = \mathcal{O}(M_q(2))[D_q^{-1}],$$

where the square brackets indicate adjoining.

Proposition 18.1. *The coordinate ring $\mathcal{O}(GL_q(2))$ is a Hopf algebra, with*

$$\Delta(u_j^i) = \sum_k u_k^i \otimes u_j^k,$$

$$S(u_j^i) = \frac{1}{D_q}\widetilde{u}_j^i = (U^{-1})_j^i,$$

$$\epsilon(u_j^i) = \delta_{ij},$$

where $\widetilde{U} = [\tilde{u}_j^i]$ is the cofactor matrix

$$\widetilde{U} = \begin{pmatrix} d & -q^{-1}b \\ -qc & a \end{pmatrix}.$$

(defined so that $U\widetilde{U} = D_q I$) and $U^{-1} = \tilde{U}/D_q$ is the inverse of U.

This is a non-commutative and non-co-commutative Hopf algebra.

Now we go to the general n. Let V_q be the n-dimensional quantum space, the q-deformation of V, with coordinates x_i's which satisfy

$$x_i x_j = q x_j x_i \quad \forall i < j. \tag{18.3}$$

Let $\mathbb{C}_q[V]$ be the coordinate ring of V_q defined by

$$\mathbb{C}_q[V] = \mathbb{C}[x_1, \ldots, x_n] / < x_i x_j - q x_j x_i > .$$

Let $M_q(n)$ be the space of quantum $n \times n$ matrices, that is the set of linear transformations on V_q which preserve (18.3) under the left as well as the right action. The left action is given by:

$$\begin{bmatrix} x_1 \\ \dots \\ x_n \end{bmatrix} = x \to Ux = x',$$

where U is the $n \times n$ generic quantum matrix. Similarly, the right action is given by:

$$x^T \to x^T U = (x'')^T.$$

Preservation of (18.3) under the left and right actions means:

$$x'_i y'_j = q x'_j x'_i, \quad \text{for} \quad i < j$$

$$x''_i y''_j = q x''_j x''_i, \quad \text{for} \quad i < j.$$

After straightforward calculations, these yield the following relations on the entries $u_{ij} = u^i_j$ of U:

$$\begin{aligned} u_{jk}u_{ik} &= q^{-1}u_{ik}u_{jk} && (i<j) \\ u_{kj}u_{ki} &= q^{-1}u_{ki}u_{kj} && (i<j) \\ u_{jk}u_{i\ell} &= u_{i\ell}u_{jk} && (i<j, k<\ell) \\ u_{jl}u_{ik} &= u_{ik}u_{j\ell} - (q-q^{-1})u_{jk}u_{i\ell} && (i<j, k<\ell). \end{aligned} \tag{18.4}$$

The quantum determinant is defined as

$$D_q = \sum_{\sigma \in S_n} (-q)^{\ell(\sigma)} u^{i_{\sigma(1)}}_{j1} \dots u^{i_{\sigma(n)}}_{jn},$$

where $\ell(\sigma)$ denotes the length of the permutation σ, that is, the number of inversions in σ. This determinant formula is the same as the usual formula substituting $(-q)$ for (-1).

We define the coordinate ring of the space $M_q(n)$ of quantum $n \times n$ matrices by

$$\mathcal{O}(M_q(n)) = \mathbb{C}[U]/ < (18.4) >, \text{ and}$$

and the coordinate ring of the virtual quantum group $GL_q(n)$ by

$$\mathbb{C}_q[G] = \mathcal{O}(GL_q(n)) = \mathcal{O}(M_q(n))[D_q^{-1}].$$

We define the quantum minors and, using these, the quantum co-factor matrix $\widetilde{U}$ and the quantum inverse matrix $U^{-1} = \widetilde{U}/D_q$ in a straightforward fashion (these constructions are left as exercises).

Theorem 18.1. *The algebra $\mathcal{O}(GL_q(n))$ is a Hopf algebra, with*

$$\Delta(u^i_j) = \sum_k u^i_k \otimes u^k_j$$

$$\epsilon(u^i_j) = \delta_{ij}$$

$$S(u^i_j) = \frac{1}{D_q}\widetilde{u}^i_j = (U^{-1})^i_j$$

$$S(D_q^{-1}) = D_q.$$

We also denote the quantum group $GL_q(n)$ by G_q, $GL_q(\mathbb{C}^n)$ or $GL_q(V)$. It has to be emphasized that this is only a virtual object. Only its coordinate ring $\mathbb{C}_q[G]$ is real. Henceforth, whenever we say representation or action of G_q, we actually mean corepresentation or coaction of $\mathbb{C}_q[G]$, and so forth.

Chapter 19

Quantum unitary group

Scribe: Joshua A. Grochow

Goal: Define the quantum unitary subgroup of the standard quantum group.
Reference: [KS]

Recall

Let $V = \mathbb{C}^n$, $G = GL_n(\mathbb{C}) = GL(V) = GL(\mathbb{C}^n)$, and $\mathcal{O}(G)$ the coordinate ring of G. The quantum group $G_q = GL_q(V)$ is the virtual object whose coordinate ring is

$$\mathcal{O}(G_q) = \mathbb{C}[U]/\langle \text{ relations } \rangle,$$

where U is the generic $n \times n$ matrix of indeterminates, and the relations are the quadratic relations on the coordinates u_i^j defined in the last class so as to preserve the non-commuting relations among the coordinates of the quantum vector space V_q on which G_q acts. This coordinate ring is a Hopf algebra.

19.1 A q-analogue of the unitary group

In this lecture we define a q-analogue of the unitary subgroup $U = U_n(\mathbb{C}) = U(V) \subseteq GL_n(\mathbb{C}) = GL(V) = G$. This is a q-deformation $U_q = U_q(V) \subseteq G_q$ of $U(V)$. Since G_q is only a virtual object, U_q will also be virtual. To define U_q, we must determine how to capture the notion of unitarity in the setting of Hopf algebras. As we shall see, it is captured by the notion of a Hopf $*$-algebra.

Definition 19.1. *A* $*$-vector space *is a vector space V with an involution $* : V \to V$ satisfying*

$$(\alpha v + \beta w)^* = \overline{\alpha} v^* + \overline{\beta} w^* \quad (v^*)^* = v$$

for all $v, w \in V$, and $\alpha, \beta \in \mathbb{C}$.

We think of $*$ as a generalization of complex conjugation; and in fact every complex vector space is a $*$-vector space, where $*$ is exactly complex conjugation.

Definition 19.2. *A* Hopf $*$-algebra *is a Hopf algebra (A, Δ, ϵ, S) with an involution $* : A \to A$ such that $(A, *)$ is a $*$-vector space, and:*

1. $(ab)^* = b^* a^*$, $1^* = 1$
2. $\Delta(a^*) = \Delta(a)^*$ *(where $*$ acts diagonally on the tensor product $A \otimes A$: $(v \otimes w)^* = (v^* \otimes w^*)$)*
3. $\epsilon(a^*) = \overline{\epsilon(a)}$

There is no explicit condition here on how $*$ interacts with the antipode S.

Let $\mathcal{O}(G) = \mathbb{C}[G]$ be the coordinate ring of G as defined earlier.

Proposition 19.1. *Then $\mathcal{O}(G)$ is a Hopf $*$-algebra.*

Proof. We think of the elements in $\mathcal{O}(G)$ as $\mathbb{C}$-valued functions on G and define $* : \mathcal{O}(G) \to \mathcal{O}(G)$ so that it satisfies the three conditions for a Hopf $*$-algebra, and

(4) For all $f \in \mathcal{O}(G)$ and $g \in U \subseteq G$, $f^*(g) = \overline{f(g)}$

Let u_i^j be the coordinate functions which, together with D^{-1}, $D = \det(U)$, generate $\mathcal{O}(G)$. Because of the first condition on a Hopf $*$-algebra (relating the involution $*$ to multiplication), specifying $(u_i^j)^*$ and D^* suffices to define $*$ completely. We define

$$(u_i^j)^* = S(u_j^i) = (U^{-1})_j^i$$

and $D^* = D^{-1}$. We can check that this satifies (1)-(4). Here we will only check (4), and leave the remaining verification as an exercise. Let g be an element of the unitary group U. Then $(u_i^j)^*(g) = S(u_j^i)(g) = (g^{-1})_j^i = (\overline{g})_i^j$, where the last equality follows from the fact that g is unitary (i.e. $g^{-1} = g^\dagger$, where $\dagger$ denotes conjugate transpose). $\square$

Thus, we have defined a map $f \mapsto f^*$ purely algebraically in such a way that the restriction of f^* to the unitary group U is the same as taking the complex conjugate $\overline{f}$ on U.

Proposition 19.2. *The coordinate ring $\mathbb{C}_q[G] = \mathcal{O}(G_q)$ of the quantum group $G_q = GL_q(V)$ is also a Hopf $*$-algebra.*

Proof. The proof is syntactically identical to the proof for $\mathcal{O}(G)$, except that the coordinate function u_i^j now lives in $\mathcal{O}(G_q)$ and the determinant D becomes the q-determinant D_q. The definition of $*$ is: $(u_i^j)^* = S(u_j^i)$ and $D_q^* = D_q^{-1}$, essentially the same as in the classical case. ☐

Intuitively, the "quantum subgroup" U_q of G_q is the virtual object such that the restriction to U_q of the involution $*$ just defined coincides with the complex conjugate.

19.2 Properties of U_q

We would like the nice properties of the classical unitary group to transfer over to the quantum unitary group, and this is indeed the case. Some of the nice properties of U are:

1. It is compact, so we can integrate over U.
2. we can do harmonic analysis on U (viz. the Peter-Weyl Theorem, which is an analogue for U of the Fourier analysis on the circle U_1).
3. Every finite dimensional representation of U has a G-invariant Hermitian form, and thus a unitary basis – we say that every finite dimensional representation of U is *unitarizable*.
4. Every finite dimensional representation X of U is completely reducible; this follows from (3), since any subrepresentation $W \subseteq X$ has a perpendicular subrepresentation $W^\perp$ under the G-invariant Hermitian form.

Compactness is in some sense the key here. The question is how to define it in the quantum setting. Following Woronowicz, we define compactness to mean that every finite dimensional representation of U_q is unitarizable. Let us see what this means formally.

Let A be a Hopf $*$-algebra, and W a corepresentation of A. Let $\rho : W \to W \otimes A$ be the corepresentation map. Let $\{b_i\}$ be a basis of W. Then, under

ρ, $b_i \mapsto \sum_j b_j \otimes m_i^j$ for some $m_i^j \in A$. We can thus define the *matrix of the (co)representation* $M = (m_i^j)$ in the basis $\{b_i\}$. We define M^* such that $(M^*)_i^j = (M_j^i)^*$. Thus, in the classical case (i.e. when $q = 1$), $M^* = M^\dagger$.

We say that the corepresentation W is *unitarizable* if it has a basis $B = \{b_i\}$ such that the corresponding matrix M_B of corepresentation satisfies the unitarity condition: $M_B M_B^* = I$. In this case, we say B is a unitary basis of the corepresentation W.

Definition 19.3. *A Hopf $*$-algebra A is* compact *if every finite dimensional corepresentation of A is unitarizable.*

Theorem 19.1 (Woronowicz)**.** *The coordinte ring $\mathbb{C}_q[G] = \mathcal{O}(G_q)$ is a compact Hopf $*$-algebra. This implies that every finite dimensional representation of G_q, by which we mean a finite dimensional coorepresentation of $\mathbb{C}_q[G]$, is completely reducible.*

Woronowicz goes further to show that we can q-integrate on U_q, and that we can do quantized harmonic analysis on U_q; i.e., a quantum analogue of the Peter-Weyl theorem holds.

Now that we know the finite dimensional representations of G_q are completely reducible, we can ask what the irreducible representations are.

19.3 Irreducible Representations of G_q

We proceed by analogy with the Weyl modules $V_\lambda(G)$ for G. Recall that every polynomial irreducible representation of $G = GL_n(\mathbb{C})$ is of this form.

Theorem 19.2.

1. *For all partitions λ of length at most n, there exists a q-Weyl module $V_{q,\lambda}(G_q)$ which is an irreducible representation of G_q such that*
$$\lim_{q \to 1} V_{q,\lambda}(G_q) = V_\lambda(G).$$
2. *The q-Weyl modules give all polynomial irreducible representations of G_q.*

19.4 Gelfand-Tsetlin basis

To understand the q-Weyl modules better, we wish to get an explicit basis for each module $V_{q,\lambda}$. We begin by defining a very useful basis – the Gel'fand-Testlin basis – in the classical case for $V_\lambda(G)$, and then describe the q-analogue of this basis.

By Pieri's rule [FH]

$$V_\lambda(GL_n(\mathbb{C})) = \bigoplus_{\lambda'} V_{\lambda'}(GL_{n-1}(\mathbb{C}))$$

where the sum is taken over all λ' obtained from λ by removing any number of boxes (in a legal way) such that no two removed boxes come from the same column. This is an orthogonal decomposition (relative to the $GL_n(\mathbb{C})$-invariant Hermitian form on V_λ) and it is also multiplicity-free, i.e., each $V_{\lambda'}$ appears only once.

Fix a G-invariant Hermitian form on V_λ. Then the *Gel'fand-Tsetlin* basis for $V_\lambda(GL_n(\mathbb{C}))$, denoted GT^n_λ, is the unique orthonormal basis for V_λ such that

$$GT^n_\lambda = \bigcup_{\lambda'} GT^{n-1}_{\lambda'},$$

where the disjoint union is over the λ' as in Pieri's rule, and $GT^{n-1}_{\lambda'}$ is defined recursively, the case $n = 1$ being trivial.

The dimension of V_λ is the number of semistandard tableau of shape λ. With any tableau T of this shape, one can also explicitly associate a basis element $GT(T) \in GT^n_\lambda$; we shall not worry about how.

We can define the Gel'fand-Tsetlin basis $GT^n_{q,\lambda}$ for $V_{q,\lambda}(G_q(\mathbb{C}^n))$ analogously. We have the q-analogue of Pieri's rule:

$$V_{q,\lambda}(G_q(\mathbb{C}^n)) = \bigoplus_{\lambda'} V_{q,\lambda'}(G_q(\mathbb{C}^{n-1}))$$

where the decomposition is orthogonal and multiplicity-free, and the sum ranges over the same λ' as above. So we can define $GT^n_{q,\lambda}$ to be the unique unitary basis of $V_{q,\lambda}$ such that

$$GT^n_{q,\lambda} = \bigcup_{\lambda'} GT^{n-1}_{q,\lambda'}.$$

With any semistandard tableau T, one can also explicitly associate a basis element $GT_q(T) \in GT^n_{q,\lambda'}$; details omitted.

Chapter 20

Towards positivity hypotheses via quantum groups

Scribe: Joshua A. Grochow

Goal: In this final brisk lecture, we indicate the role of quantum groups in the context of the positivity hypothesis PH1. Specifically, we sketch how the Littlewood-Richardson rule – the gist of PH1 in the Littlewood-Richardson problem – follows from the theory of standard quantum groups. We then briefly mention analogous (nonstandard) quantum groups for the Kronecker and plethysm problems defined in [GCT4, GCT7], and the theorems and conjectures for them that would imply PH1 for these problems.
References: [KS, K, Lu2, GCT4, GCT6, GCT7, GCT8]

Let $V = \mathbb{C}^n$, $G = GL_n(\mathbb{C}) = GL(V)$, $V_\lambda = V_\lambda(G)$ a Weyl module of G, $G_q = GL_q(V)$ the standard quantum group, V_q the q-deformation of V on which $GL_q(V)$ acts, $V_{q,\lambda} = V_{q,\lambda}(G_q)$ the q-deformation of $V_\lambda(G)$, and $GT_{q,\lambda} = GT^n_{q,\lambda}$ the Gel'fand-Tsetlin basis for $V_{q,\lambda}$.

20.1 Littlewood-Richardson rule via standard quantum groups

We now sketch how the Littlewood-Richardson rule falls out of the standard quantum group machinery, specifically the properties of the Gelfand-Tsetlin basis.

20.1.1 An embedding of the Weyl module

For this, we have to embed the q-Weyl module $V_{q,\lambda}$ in $V_q^{\otimes d}$, where $d = |\lambda| = \sum \lambda_i$ is the size of λ. We first describe how to embed the Weyl module V_λ of G in $V^{\otimes d}$ in a standard way that can be quantized.

If $d = 1$, then $V_\lambda(G) = V = V^{\otimes 1}$. Otherwise, obtain a Young diagram μ from λ by removing its top-rightmost box that can be removed to get a valid Young diagram, e.g.:

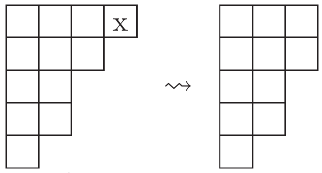

λ μ

In the following, the box must be removed from the second row, since removing from the first row would result in an illegal Young diagram:

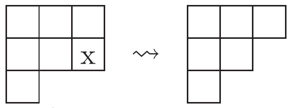

λ μ

By induction on d, we have a standard embedding $V_\mu(G) \hookrightarrow V^{\otimes d-1}$. This gives us an embedding $V_\mu(G) \otimes V \hookrightarrow V^{\otimes d}$. By Pieri's rule [FH]

$$V_\mu(G) \otimes V = \bigoplus_\beta V_\beta(G),$$

where the sum is over all β obtained from μ by adding one box in a legal way. In particular, $V_\lambda(G) \subset V_\mu(G) \otimes V$. By restricting the above embedding, we get a standard embedding $V_\lambda(G) \hookrightarrow V^{\otimes d}$.

Now Pieri's rule also holds in a quantized setting:

$$V_{q,\mu} \otimes V_q = \bigoplus_\beta V_{q,\beta}(G),$$

where β is as above. Hence, the standard embedding $V_\lambda \hookrightarrow V^{\otimes d}$ above can be quantized in a straightforward fashion to get a standard embedding $V_{q,\lambda} \hookrightarrow V_q^{\otimes d}$. We shall denote it by ρ. Here the tensor product is meant to be over $\mathbb{Q}(q)$. Actually, $\mathbb{Q}(q)$ doesn't quite work. We have to allow square roots of elements of $\mathbb{Q}(q)$, but we won't worry about this. For a semistandard tableau b of shape λ, we denote the image of a Gelfand-Tsetlin basis element $GT_{q,\lambda}(b) \in GT_{q,\lambda}$ under ρ by $GT^\rho_{q,\lambda}(b) = \rho(GT_{q,\lambda}(b)) \in V_q^{\otimes d}$.

20.1.2 Crystal operators and crystal bases

Theorem 20.1 (Crystallization). *[DJM] The Gelfand-Tsetlin basis elements crystallize at $q = 0$. This means:*

$$\lim_{q \to 0} GT^{\rho}_{q,\lambda}(b) = v_{i_1(b)} \otimes \cdots \otimes v_{i_d(b)}, \tag{20.1}$$

for some integer functions $i_1(b), \ldots, i_d(b)$, and

$$\lim_{q \to \infty} GT^{\rho}_{q,\lambda}(b) = v_{j_1(b)} \otimes \cdots \otimes v_{j_d(b)}, \tag{20.2}$$

for some integer functions $j_1(b), \ldots, j_d(b)$.

The phenomenon that these limits consists of monomials, i.e., simple tensors is known as *crystallization*. It is related to the physical phenomenon of crystallization, hence the name. The maps $b \mapsto \overline{i}(b) = (i_1(b), \ldots, i_d(b))$ and $b \mapsto \overline{j}(b) = (j_1(b), \ldots, j_d(b))$ are computable in *poly*($\langle b \rangle$) time (where $\langle b \rangle$ is the bit-length of b).

Now we want to define a special *crystal basis* of $V_{q,\lambda}$ based on this phenomenon of crystallization. Towards that end, consider the following family of $n \times n$ matrices:

$$E_i = \begin{pmatrix} 0 & 0 & & & \cdots & 0 \\ & \ddots & \ddots & & & \vdots \\ & & 0 & 1 & \cdots & 0 \\ & & & 0 & \cdots & 0 \\ & & & & \ddots & \vdots \\ & & & & & 0 \end{pmatrix},$$

where the only nonzero entry is a 1 in the i-th row and $(i+1)$-st column. Let $F_i = E_i^T$. Corresponding to E_i and F_i, Kashiwara associates certain operators $\hat{E}_i$ and $\hat{F}_i$ on $V_{q,\lambda}(G_q)$. We shall not worry about their actual construction here (for the readers familiar with Lie algebras: these are closely related to the usual operators in the Lie algebra of G associated with E_i and F_i).

If we let $\hat{E}_i$ act on $GT^{\rho}_{q,\lambda}(b)$, we get some linear combination

$$\hat{E}_i(GT^{\rho}_{q,\lambda}(b)) = \sum_{b'} a^b_{b'}(q) GT^{\rho}_{q,\lambda}(b'),$$

where $a^b_{b'}(q) \in \mathbb{Q}(q)$ (actually an algebraic extension of $\mathbb{Q}(q)$ as mentioned above). Essentially because of crystallization (Theorem 20.1), it turns out that $\lim_{q\to 0} a^b_{b'}(q)$ is always either 0 or 1, and for a given b, this limit is 1 for at most one b', if any. A similar result holds for $\hat{F}_i(GT^\rho_{q,\lambda}(b))$. This allows us to define the *crystal operators* (due to Kashiwara):

$$\widetilde{e}_i \cdot b = \begin{cases} b' & \text{if } \lim_{q\to 0} a^b_{b'}(q) = 1, \\ 0 & \text{if no such } b' \text{ exists,} \end{cases}$$

and similarly for $\widetilde{f}_i$. Although these operators are defined according to a particular embedding $V_{q,\lambda} \hookrightarrow V_q^{\otimes d}$ and a basis, they can be defined intrinsically, i.e., without reference to the embedding or the Gel'fand-Tsetlin basis.

Now, let W be a finite-dimensional representation of G_q, and R the subring of functions in $\mathbb{Q}(q)$ regular at $q = 0$ (i.e. without a pole at $q = 0$). A *lattice* within W is an R-submodule of W such that $\mathbb{Q}(q) \otimes_R L = W$. (Intuition behind this definition: $R \subset \mathbb{Q}(q)$ is analogous to $\mathbb{Z} \subset \mathbb{Q}$. A lattice in $\mathbb{R}^n$ is a $\mathbb{Z}$-submodule L of $\mathbb{R}^n$ such that $\mathbb{R} \otimes_{\mathbb{Z}} L = \mathbb{R}^n$.)

Definition 20.1. *An (upper)* crystal basis *of a representation W of G_q is a pair (L, B) such that*

- *L is a lattice in W preserved by the Kashiwara operators $\hat{E}_i$ and $\hat{F}_i$, i.e. $\hat{E}_i(L) \subseteq L$ and $\hat{F}_i(L) \subseteq L$.*
- *B is a basis of L/qL preserved by the crystal operators $\widetilde{e}_i$ and $\widetilde{f}_i$, i.e., $\widetilde{e}_i(B) \subseteq B \cup \{0\}$ and $\widetilde{f}_i(B) \subseteq B \cup \{0\}$.*
- *The crystal operators $\widetilde{e}_i$ and $\widetilde{f}_i$ are inverse to each other wherever possible, i.e., for all $b, b' \in B$, if $\widetilde{e}_i(b) = b' \neq 0$ then $\widetilde{f}_i(b') = b$, and similarly, if $\widetilde{f}_i(b) = b' \neq 0$ then $\widetilde{e}_i(b') = b$.*

It can be shown that if $W = V_{q,\lambda}(G_q)$, then there exists a unique $b \in B$ such that $\widetilde{e}_i(b) = 0$ for all i; this corresponds to the highest weight vector of $V_{q,\lambda}$ (the weight vectors in $V_{q,\lambda}$ are analogous to the weight vectors in V_λ; we do not give their exact definition here). By the work of Kashiwara and Date et al [K, DJM] above, the Gel'fand-Tsetlin basis (after appropriate rescaling) is in fact a crystal basis: just let

$$\begin{aligned} L = L_{GT} &= \text{ the } R\text{-module generated by } GT_{q,\lambda}, \textit{and} \\ B_{GT} &= \overline{GT_{q,\lambda}(b)}, \end{aligned}$$

where $\overline{GT_{q,\lambda}(b)}$ is the image under the projection $L \mapsto L/qL$ of the set of basis vectors in $GT_{q,\lambda}(b)$.

Theorem 20.2 (Kashiwara). 1. *Every finite-dimensional G_q-module has a unique crystal basis (up to isomorphism).*

2. *Let (L_λ, B_λ) be the unique crystal basis corresponding to $V_{q,\lambda}$. Then $(L_\alpha, B_\alpha) \otimes (L_\beta, B_\beta) = (L_\alpha \otimes L_\beta, B_\alpha \otimes B_\beta)$ is the unique crystal basis of $V_{q,\alpha} \otimes V_{q,\beta}$, where $B_\alpha \otimes B_\beta$ denotes $\{b_a \otimes b_b | b_a \in B_\alpha, b_b \in B_\beta\}$.*

It can be shown that every $b \in B_\lambda$ has a weight; i.e., it is the image of a weight vector in L_λ under the projection $L_\lambda \to L_\lambda/qL_\lambda$.

Now let us see how the Littlewood-Richardson rule falls out of the properties of the crystal bases. Recall that the specialization of $V_{q,\alpha}$ at $q = 1$ is the Weyl module V_α of $G = GL_n(\mathbb{C})$, and

$$V_\alpha \otimes V_\beta = \bigoplus_\gamma c^\gamma_{\alpha,\beta} V_\gamma \tag{20.3}$$

where $c^\gamma_{\alpha,\beta}$ are the Littlewood-Richardson coefficients. The Littlewood-Richardson rule now follows from the following fact:

$$c^\gamma_{\alpha,\beta} = \#\{b \otimes b' \in B_\alpha \otimes B_\beta | \forall i,\ \widetilde{e_i}(b \otimes b') = 0 \text{ and } b \otimes b' \text{ has weight } \gamma\}.$$

Intuitively, $b \otimes b'$ here correspond to the highest weight vectors of the G-submodules of $V_\alpha \otimes V_\beta$ isomorphic to V_γ.

20.2 Explicit decomposition of the tensor product

The decomposition (20.3) is only an abstract decomposition of $V_\alpha \otimes V_\beta$ as a G-module. Next we consider the explicit decomposition problem. The goal is to find an explicit basis $\mathcal{B} = \mathcal{B}_{\alpha\otimes\beta}$ of $V_\alpha \otimes V_\beta$ that is compatible with this abstract decomposition. Specifically, we want to construct an explicit basis $\mathcal{B}$ of $V_\alpha \otimes V_\beta$ in terms of suitable explicit bases of V_α and V_β such that $\mathcal{B}$ has a filtration

$$\mathcal{B} = \mathcal{B}_0 \supseteq \mathcal{B}_1 \supseteq \cdots \supseteq \emptyset$$

where each $\langle \mathcal{B}_i \rangle / \langle \mathcal{B}_{i+1} \rangle$ is an irreducible representation of G and $\langle \mathcal{B}_i \rangle$ denotes the linear span of $\mathcal{B}_i$. Furthermore, each element $b \in \mathcal{B}$ should have a sufficiently explicit representation in terms of the basis $\mathcal{B}_\alpha \otimes \mathcal{B}_\beta$ of $V_\alpha \otimes V_\beta$. The explicit decomposition problem for the q-analogue $V_{q,\alpha} \otimes V_{q,\beta}$ is similar.

For example, we have already constructed explicit Gelfand-Tsetlin bases of Weyl modules. But it is not known how to construct an explicit basis $\mathcal{B}$

with filtration as above in terms of the Gelfand-Tsetlin bases of V_α and V_β (except when the Young diagram of either α or β is a single row).

Kashiwara and Lusztig [K, Lu2] construct certain *canonical bases* $\mathcal{B}_{q,\alpha}$ and $\mathcal{B}_{q,\beta}$ of $V_{q,\alpha}$ and $V_{q,\beta}$, and Lusztig furthermore constructs a *canonical basis* $\mathcal{B}_q = \mathcal{B}_{q,\alpha\otimes\beta}$ of $V_{q,\alpha} \otimes V_{q,\beta}$ such that:

1. $\mathcal{B}_q$ has a filtration as above,
2. Each $b \in \mathcal{B}_q$ has an expansion of the form
$$b = \sum_{b_\alpha \in \mathcal{B}_{q,\alpha}, b_\beta \in \mathcal{B}_{q,\beta}} a_b^{b_\alpha, b_\beta} b_\alpha \otimes b_\beta,$$
where each $a_b^{b_\alpha,b_\beta}$ is a polynomial in q and q^{-1} with nonnegative integral coefficients,
3. *Crystallization*: For each b, as $q \to 0$, exactly one coefficient $a_b^{b_\alpha,b_\beta} \to 1$, and the remaining all vanish.

The proof of nonnegativity of the coefficients of $a_b^{b_\alpha,b_\beta}$ is based on the Riemann hypothesis (theorem) over finite fields [Dl2], and explicit formulae for these coefficients are known in terms of perverse sheaves [BBD] (which are certain types of algebro-geometric objects).

This then provides a satisfactory solution to the explicit decomposition problem, which is far harder and deeper than the abstract decomposition provided by the Littlewood-Richardson rule. By specializing at $q = 1$, we also get a solution to the explicit decomposition problem for $V_\alpha \otimes V_\beta$. This (i.e. via quantum groups) is the only known solution to the explicit decomposition problem even at $q = 1$. This may give some idea of the power of the quantum group machinery.

20.3 Towards nonstandard quantum groups for the Kronecker and plethysm problems

Now the goal is to construct quantum groups which can be used to derive PH1 and explicit decomposition for the Kronecker and plethysm problems just as the standard quantum group can be used for the same in the Littlewood-Richardson problem.

In the Kronecker problem, we let $H = GL(\mathbb{C}^n)$ and $G = GL(\mathbb{C}^n \otimes \mathbb{C}^n)$. The Kronecker coefficient $\kappa_{\alpha,\beta}^\gamma$ is the multiplicity of $V_\alpha(H) \otimes V_\beta(H)$ in $V_\gamma(G)$:
$$V_\gamma(G) = \bigoplus_{\alpha,\beta} \kappa_{\alpha,\beta}^\gamma V_\alpha(H) \otimes V_\beta(H).$$

The goal is to get a positive # **P**-formula for $\kappa_{\alpha,\beta}^{\gamma}$; this is the gist of PH1 for the Kronecker problem.

In the plethysm problem, we let $H = GL(\mathbb{C}^n)$ and $G = GL(V_\mu(H))$. The plethysm constant $a_{\lambda,\mu}^{\pi}$ is the multiplicity of $V_\pi(H)$ in $V_\lambda(G)$:

$$V_\lambda(G) = \bigoplus_\pi a_{\lambda,\mu}^{\pi} V_\pi(H).$$

Again, the goal is to get a positive # **P**-formula for the plethysm constant; this is the gist of PH1 for the plethysm problem.

To apply the quantum group approach, we need a q-analogue of the embedding $H \hookrightarrow G$. Unfortunately, there is no such q-analogue in the theory of standard quantum groups. Because there is no nontrivial quantum group homomorphism from the standard quantum group $H_q = GL_q(\mathbb{C}^n)$ and to the standard quantum group G_q.

Theorem 20.3. *(1) [GCT4]: Let H and G be as in the Kronecker problem. Then there exists a quantum group $\hat{G}_q$ such that the homomorphism $H \to G$ can be quantized in the form $H_q \hookrightarrow \hat{G}_q$. Furthermore, $\hat{G}_q$ has a unitary quantum subgroup $\hat{U}_q$ which corresponds to the maximal unitary subgroup $U \subseteq G$, and a q-analogue of the Peter-Weyl theorem holds for $\hat{G}_q$. The latter implies that every finite dimensional representation of $\hat{G}_q$ is completely decomposible into irreducibles.*
(2) [GCT7] There is an analogous (possibly singular) quantum group $\hat{G}_q$ when H and G are as in the plethysm problem. This also holds for general connected reductive (classical) H.

Since the Kronecker problem is a special case of the (generalized) plethysm problem, the quantum group in GCT 4 is a special case of the quantum group in GCT 7. The quantum group in the plethysm problem can be singular, i.e., its determinant can vanish and hence the antipode need not exist. We still call it a quantum group because its properties are very similar to those of the standard quantum group; e.g. q-analogue of the Peter-Weyl theorem, which allows q-harmonic analysis on these groups. We call the quantum group $\hat{G}_q$ *nonstandard*, because though it is qualitatively similar to the standard (Drinfeld-Jimbo) quantum group G_q, it is also, as expected, fundamentally different.

The article [GCT8] gives a conjecturally correct algorithm to construct a canonical basis of an irreducible polynomial representation of $\hat{G}_q$ which generalizes the canonical basis for a polynomial representation of the standard quantum group as per Kashiwara and Lusztig. It also gives a conjecturally correct algorithm to construct a canonical basis of a certain q-deformation of the symmetric group algebra $\mathbb{C}[S_r]$ which generalizes the Kazhdan-Lusztig basis [KL] of the Hecke algebra (a standard q-deformation of $\mathbb{C}[S_r]$). It is shown in [GCT7, GCT8] that PH1 for the Kronecker and plethysm problems follows assuming that these canonical bases in the nonstandard setting have properties akin to the ones in the standard setting. For a discussion on SH, see [GCT6].

Part II

Invariant theory with a view towards GCT

By Milind Sohoni

Chapter 21

Finite Groups

References: [FH, N]

21.1 Generalities

Let V be a vector space over $\mathbb{C}$, and let $GL(V)$ denote the group of all isomorphisms on V. For a fixed basis of V, $GL(V)$ is isomorphic to the group $GL_n(\mathbb{C})$, the group of all $n \times n$ invertible matrices.

Let G be a group and $\rho : G \to GL(V)$ be a representation. We also denote this by the tuple (ρ, V) or say that V is a G-module. Let $Z \subseteq V$ be a subspace such that $\rho(g)(Z) \subseteq Z$ for all $g \in G$. Then, we say that Z is an **invariant subspace**. We say that (ρ, V) is **irreducible** if there is no proper subspace $W \subset V$ such that $\rho(g)(W) \subseteq W$ for all $g \in G$. We say that (ρ, V) is **indecomposable** is there is no expression $V = W_1 \oplus W_2$ such that $\rho(g)(W_i) \subseteq W_i$, for all $g \in G$.

For a point $v \in V$, the **orbit** $O(v)$, and the **stabilizer** $Stab(v)$ are defined as:

$$\begin{array}{rcl} O(v) & = & \{v' \in V | \exists g \in G \, with \, \rho(g)(v) = v'\} \\ Stab(v) & = & \{g \in G | \rho(g)(v) = v\} \end{array}$$

One may also define $v \sim v'$ if there is a $g \in G$ such that $\rho(g)(v) = v'$. It is then easy to show that $[v]_\sim = O(v)$.

Let V^* be the dual-space of V. The representation (ρ, V) induces the **dual** representation (ρ^*, V^*) defined as $\rho^*(v^*)(v) = v^*(\rho(g^{-1})(v))$. It will be convenient for ρ^* to act on the right, i.e., $((v^*)(\rho^*))(v) = v^*(\rho(g^{-1})(v))$.

When ρ is fixed, we abbrieviate $\rho(g)(v)$ as just $g \cdot v$. Along with this, there are the standard constructions of the **tensor** $T^d(V)$, the **symmetric power** $Sym^d(V)$ and the **alternating power** $\wedge^d(V)$ representations.

Of special significance is $Sym^d(V^*)$, the space of homogeneous polynomial functions on V of degree d. Let $dim(V) = n$ and let $X_1, \ldots, X_n$ be a basis of V^*. We define as follows:

$$R = \mathbb{C}[X_1, \ldots, X_n] = \oplus_{d=0}^{\infty} R_d = \oplus_{d=0}^{\infty} Sym^d(V^*)$$

Thus R is the ring of all polynomial functions on V and is isomorphic to the algebra (over $\mathbb{C}$) of n indeterminates. Since G acts on the domain V, G also acts on all functions $f : V \to \mathbb{C}$ as follows:

$$(f \cdot g)(v) = f(g^{-1} \cdot v)$$

This action of G on all functions extends the action of G on polynomial functions above. Indeed, for any $g \in G$, the map $t_g : R \to R$ given by $f \to f \cdot g$ is an algebra isomorphism. This is called the **translation map**.

For an $f \in R$, we say that f is an **invariant** if $f \cdot g = f$ for all $g \in G$. The following are equivalent:

- $f \in R$ is an invariant.
- $Stab(f) = G$.
- $f(g \cdot v) = f(v)$ for all $g \in G$ and $v \in V$.
- For all v, v' such that $v' \in Orbit(v)$, we have $f(v) = f(v')$.

If W_1 and W_2 are two modules of G and $\phi : W_1 \to W_2$ is a linear map such that $g \cdot \phi(w_1) = \phi(g \cdot w_1)$ for all $g \in G$ and $w_1 \in W_1$ then we say that ϕ is **G-equivariant** or that ϕ is a **morphism of G-modules**.

21.2 The finite group action

Let G be a finite group and (μ, W) be a representation.

Recall that a **complex inner product** on W is a map $h : W \times W \to \mathbb{C}$ such that:

- $h(\alpha w + \beta w', w'') = \overline{\alpha} h(w, w'') + \overline{\beta} h(w', w'')$ for all $\alpha, \beta \in \mathbb{C}$ and all $w, w', w'' \in W$.
- $h(w'', \alpha w + \beta w') = \alpha h(w'', w) + \beta h(w'', w')$ for all $\alpha, \beta \in \mathbb{C}$ and all $w, w', w'' \in W$.
- $h(w, w) > 0$ for all $w \neq 0$.

Also recall that if $Z \subseteq W$ is a subspace, then $Z^\perp$ is defined as:

$$Z^\perp = \{w \in W | h(w, z) = 0 \, \forall z \in Z\}$$

Also recall that $W = Z \oplus Z^\perp$.

We say that an inner product h is G**-invariant** if $h(g \cdot w, g \cdot w') = h(w, w')$ for all $w, w' \in W$ and $g \in G$.

Proposition 21.1. *Let W be as above, and Z be an invariant subspace of W. Then $Z^\perp$ is also an invariant subspace. Thus every reducible representation of G is also decomposable.*

Proof: Let $x \in Z^\perp, z \in Z$ and let us examine $(g \cdot x, z)$. Applying g^{-1} to both sides, we see that:

$$h(g \cdot x, z) = h(g^{-1} \cdot g \cdot x, g^{-1} \cdot z) = h(x, g^{-1} \cdot z) = 0$$

Thus, G preserves $Z^\perp$ as claimed. $\square$

Let h be a complex inner product on W. We define the inner product h^G as follows:

$$h^G(w, w') = \frac{1}{|G|} \sum_{g' \in G} h(g' \cdot w, g' \cdot w')$$

Lemma 21.1. *h^G is a G-invariant inner product.*

Proof: First we see that

$$h^G(w, w) = \frac{1}{|G|} \sum_{g' \in G} h(w', w')$$

where $w' = g' \cdot w$. Thus $h^G(w, w) > 0$ unless $w = 0$. Secondly, by the linearity of the action of G, we see that h^G is indeed an inner product. Finally, we see that:

$$h^G(g \cdot w, g \cdot w') = \frac{1}{|G|} \sum_{g' \in G} h(g' \cdot g \cdot w, g' \cdot g \cdot w')$$

Since as g' ranges over G, so will $g' \cdot g$ for any fixed g, we have that h^G is G-invariant. $\square$

Theorem 1.

- *Let G be a finite group and (ρ, V) be an indecomposable representation, then it is also irreducible.*

- *Every representation (ρ, V) may be decomposed into irreducible representations V_i. Thus $V = \oplus_i V_i$, where (ρ_i, V_i) is an irreducible representation.*

Proof: Suppose that $Z \subseteq V$ is an invariant subspace, then $V = Z \oplus Z^{\perp}$ is a non-trivial decomposition of V contradicting the hypothesis. The second part is proved by applying the first, recursively. $\square$

We have seen the operation of **averaging over the group** in going from the inner product h to the G-invariant inner product h^G. A similar approach may be used for constructing invariant polynomials functions. So let $p(X) \in R = \mathbb{C}[X_1, \ldots, X_n]$ be a polynomial function. We define the function $p^G : V \to \mathbb{C}$ as:

$$p^G(v) = \frac{1}{|G|} \sum_{g \in G} p(g \cdot v)$$

The transition from p to p^G is called the **Reynold's operator**.

Proposition 21.2. *Let $p \in R$ be of degree atmost d, then p^G is also a polynomial function of degree atmost d. Next, p^G is an invariant.*

Let R^G denote the set of all invariant polynomial functions on the space V. It is easy to see that $R^G \subseteq R$ is actually a *subring* of R.

Let $Z \subseteq V$ be an arbitrary subset of V. We say that Z is G-closed if $g \cdot z \in Z$ for all $g \in G$ and $z \in Z$. Thus Z is a union of orbits of points in V.

Lemma 21.2. *Let $p \in R^G$ be an invariant and let $Z = V(p)$ be the variety of p. Then Z is G-closed.*

We have already seen that $O(v)$, the orbit of v arises from the equivalence class $\sim$ on V. Since the group is finite, $|O(v)| \leq |G|$ for any v. Let O_1 and O_2 be disjoint orbits. It is essential to determine if elements of R^G can separate O_1 and O_2.

Lemma 21.3. *Let O_1 and O_2 be as above, and I_1 and I_2 be their ideals in R. Then there are $p_1 \in I_1$ and $p_2 \in I_2$ so that $p_1 + p_2 = 1$.*

Proof: This follows from the Hilbert Nullstellensatz. Since the point sets are finite, there is an explicit construction based on Lagrange interpolation. $\square$

Let G be a finite group and (ρ, V) be a representation as above. We have see that this induces an action on $\mathbb{C}[X_1, \ldots, X_n]$. Also note that this action is **homogeneous**: for a $g \in G$ and $p \in R_d$, we have that $p \cdot g \in R_d$

as well. Thus R^G, the ring of invariants, is a *homogeneous subring* of R. In other words:

$$R^G = \oplus_{d=0}^{\infty} R_d^G$$

where R_d^G are invariants which are homogeneous of degree d. The existence of the above decomposition implies that every invariant is a sum of homogeneous invariants. Now, since $R_d^G \subseteq R_d$ as a vector space over $\mathbb{C}$. Thus

$$dim_{\mathbb{C}}(R_d^G) \leq dim_{\mathbb{C}}(R_d) \leq \binom{n+d-1}{n-1}$$

We define the **hilbert function** $h(R^G)$ of R^G (or for that matter, of any homogeneous ring) as:

$$h(R^G) = \sum_{d=0}^{\infty} dim_{\mathbb{C}}(R_d^G) z^d$$

We will see now that $h(R^G)$ is actually a rational function which is easily computed. We need a lemma.

Let (ρ, W) be a representation of the finite group G. Let

$$W^G = \{w \in W | g \cdot w = w\}$$

be the set of all vector invariants in W, and this is a subspace of W. .

Lemma 21.4. *Let (ρ, W) be as above. We have:*

$$dim_{\mathbb{C}}(W^G) = \frac{1}{|G|} \sum_{g \in G} trace(\rho(g))$$

Proof: Define $P = \frac{1}{|G|} \sum_{g \in G} \rho(g)$, as the average of the representation matrices. We see that $\rho(g) \cdot P = P \cdot \rho(g)$ and that $P^2 = P$. Thus P is diagonalizable and the eigenvalues of P are in the set $\{1, 0\}$. Let W_1 and W_0 be the corresponding eigen-spaces. It is clear that $W^G \subseteq W_1$ and that W_1 is fixed by each $g \in G$. We now argue that every $w \in W_1$ is actually an invariant. For that, let $w_g = g \cdot w$. We then have that $Pw = w$ implies that

$$w = \frac{1}{|G|} \sum_{g \in G} w_g$$

Note that a change-of-basis does not affect the hypothesis nor the assertion. We may thus assume that each $\rho(g)$ is unitary, we have that $w_g = w$ for all $g \in G$. Now, the claim follows by computing $trace(P)$. $\square$

We are now ready to state **Molien's Theorem**:

Theorem 2. *Let (ρ, W) be as above. We have:*

$$h(R^G) = \frac{1}{|G|}\sum_{g\in G}\frac{1}{det(I - z\rho(g))}$$

Proof: Let $dim_{\mathbb{C}}(W) = n$ and let $\{X_1, \ldots, X_n\}$ be a basis of W^*. Since $R^G = \sum_d R_d^G$ and each $R_d^G \subseteq \mathbb{C}[X_1, \ldots, X_n]_d$. Note that each $C[X_1, \ldots, X_n]_d$ is also a representation ρ_d of G. Furthermore, it is easy to see that if $\{\lambda_1, \ldots, \lambda_n\}$ are the eigenvalues of $\rho(g)$, then the eigen-values of the matrix $\rho_d(g)$ are precisely (including multiplicity)

$$\{\prod_i \lambda_i^{d_i} | \sum_i d_i = d\}$$

Thus

$$trace(\rho_d(g)) = \sum_{\overline{d}:|\overline{d}|=d}\prod_i \lambda_i^{d_i}$$

We then have:

$$\begin{array}{rcl} h(R^G) & = & \sum_d z^d dim_{\mathbb{C}}(R_d^G) \\ & = & \sum_d z^d[\frac{1}{|G|}\sum_g trace(\rho_d(g))] \\ & = & \frac{1}{|G|}\sum_g \frac{1}{(1-\lambda_1(g)z)\ldots(1-\lambda_n(g)z)} \\ & = & \frac{1}{|G|}\sum_g \frac{1}{det(I - z\rho(g))} \end{array}$$

This proves the theorem. □

21.3 The Symmetric Group

S_n will denote the symmetric group of all bijections on the set $[n]$. The standard representation of S_n is obviously on $V = \mathbb{C}^n$ with

$$\sigma \cdot (v_1, \ldots, v_n) = (v_{\sigma(1)}, \ldots, v_{\sigma(n)})$$

Thus, regarding V as column vectors, and S_n as the group of $n \times n$-permutation matrices, we see that the action of permutation P on vector v is given by the matrix multiplication $P \cdot v$.

Let $X_1, \ldots, X_n$ be a basis of V^*. S_n acts on $R = \mathbb{C}[X_1, \ldots, X_n]$ by $X_i \cdot \sigma = X_{\sigma(i)}$. The orbit of any point $v = (v_1, \ldots, v_n)$ is the collection of all permutation of the entries of the vector v and thus the size of the orbit is bounded by $n!$.

The invariants for this action are given by the elementary symmetric polynomials $e_k(X)$, for $k = 1, \ldots, n$, where

$$e_k(X) = \sum_{i_1 < i_2 < \ldots < i_k} X_{i_1} X_{i_2} \ldots X_{i_k}$$

Given two vector v and w, if $w \notin O(v)$, then there is a k such that $e_k(v) \neq e_k(w)$. This follows from the theory of equations in one variable.

The ring R^G equals $\mathbb{C}[e_1, \ldots, e_n]$ and has no algebraic dependencies. The hilbert function of R_G may then be expressed as:

$$h(R^G) = \frac{1}{(1-z)(1-z^2)\ldots(1-z^n)}$$

It is an exercise to verify that Molien's expression agrees with the above.

A related action of S_n is the **diagonal action**: Let $X = \{X_1, \ldots, X_n\}$, $Y = \{Y_1, \ldots, Y_n\}$, and so on upto $W = \{W_1, \ldots, W_n\}$ be a family of r (disjoint) variables. Let $B = \mathbb{C}[X, Y, \ldots, W]$ be the ring of polynomials in the variables of the disjoint union.

We define the action of S_n on $X \cup Y \cup \ldots \cup W$ as $X_i \cdot \sigma = X_{\sigma(i)}$, $Y_i \cdot \sigma = Y_{\sigma(i)}$, and so on.

The matrix equivalence of this action is the action of the permutation matrices on $n \times r$ matrices A, where the action of P on A is given by $P \cdot A$.

The invariants B^G is obtained from the $r = 1$ case by a *curious* operation: Let D_{XY} denote the operator:

$$D_{XY} = Y_1 \frac{\partial}{\partial X_1} + \ldots + Y_n \frac{\partial}{\partial X_n}$$

The ring B^G is obtained from R^G by applying the operators D_{XY}, D_{XW}, D_{WX} and so on, to elements of R^G. As an example, we have

$$e_2(X) = X_1X_2 + X_1X_3 + \ldots + X_{n-1}X_n$$

We have $D_{XY}(e_2)$ as:

$$D_{XY}(e_2(X)) = \sum_{i \neq j} X_i Y_j$$

This is clearly an element of B^G.

Chapter 22

The Group SL_n

References: [FH, N]

22.1 The Canonical Representation

Let V be a vector space of dimension n, and let $x_1, \ldots, x_n$ be a basis for V. Let $X_1, \ldots, X_n$ be the dual basis of V^*.

$SL(V)$ will denote the group of all unimodular linear transformations on V. In the above basis, this group is isomorphic to that of all $n \times n$ matrices of determinant 1, or in other words $SL_n(\mathbb{C})$.

The standard representation of $SL(V)$ is obviously V itself: Given $\phi \in SL(V)$ and $v \in V$, we have $\phi \cdot v = \phi(v)$ is the action of ϕ on v.

In terms of the basis x above we may write $v = [x_1, \ldots, x_n][\alpha_1, \ldots, \alpha_n]^T$ and thus $\phi{\cdot}v$ as $[\phi{\cdot}x_1, \ldots, \phi{\cdot}x_n][\alpha_1, \ldots, \alpha_n]^T$. If $\phi{\cdot}x_i = [x_1, \ldots, x_n][a_{1i}, \ldots, a_{ni}]^T$, then we have

$$\phi \cdot v = [x_1, \ldots, x_n] \begin{bmatrix} a_{11} & \ldots & a_{1n} \\ \vdots & & \vdots \\ a_{n1} & \ldots & a_{nn} \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix}$$

We denote this matrix as A_ϕ.

We may now work with $SL_n(\mathbb{C})$ or simply SL_n. Given a vector $a = [\alpha_1, \ldots, \alpha_n]^T$, we see that the matrix multiplication $A \cdot a$ is the action of A on the column vector a.

Let us now understand the orbits of typical elements in the column space $\mathbb{C}^n$. The typical $v \in \mathbb{C}^n$ is a non-zero column vector. For any non-zero vector w, we see that there is an element $A \in SL_n$ such that $w = Av$. Furthermore, for any $B \in SL_n$, clearly $Bv \neq 0$. Thus we see that $\mathbb{C}^n$ has exactly two orbits:

$$O_0 = \{0\} \quad O_1\{v \in \mathbb{C}^n | v \neq 0\}$$

Note that O_1 is dense in $V = \mathbb{C}^n$ and its closure includes the orbit O_0 and therefore, the whole of V.

Let $R = \mathbb{C}[X_1, \ldots X_n]$ be the ring of polynomial functions on $\mathbb{C}^n$. We examine the action of SL_n on $\mathbb{C}[X]$. Recall that the action of A on X should be such that the evaluations $X_i(x_j) = \delta_{ij}$ must be conserved. Thus if the column vectors $x_i = [0, \ldots, 0, 1, 0 \ldots, 0]^T$ are the basis vectors of V and the row vectors $X_i = [0, \ldots, 0, 1, 0 \ldots, 0]$ that of V^*, then a matrix $A \in SL_n$ transforms x_i to Ax_i and X_j to $X_j A^{-1}$. Thus $X_j(x_i) = X_j/cdotx_i$ goes to $X_j A^{-1} A x_i = X_j \cdot x_i$.

Next, we examine $R = \mathbb{C}[X]$ for invariants. First note that the action of SL_n is homogeneous and thus we may assume that an invariant p is actually homogeneous of degree d. Next we see that p must be constant on O_1 and O_0. In this case, if $p(O_1) = \alpha$ then by the density of O_1 in V, we see that p is actually constant on V. Thus $R^G = R_0 = \mathbb{C}$.

22.2 The Diagonal Representation

Let us now consider the diagonal representation of the above representation. In other words, let V^r be the space of all complex $n \times r$ matrices x. The action of A on x is obvious $A \cdot x$. Let X be the $r \times n$-matrix dual to x. It transforms according to $X \to XA^{-1}$.

Let us examine the case when $r < n$ and compute the orbits in V^r. Let $x \in V^r$ be a matrix and y be a column vector such that $x \cdot y = 0$. We see that $(A \cdot x) \cdot y = 0$ as well. Thus if $ann(x) = \{y \in \mathbb{C}^r | x \cdot y = 0\}$ is the annihilator space of x, then we see that $ann(x) = ann(A \cdot x)$.

We show that the converse is also true:

Proposition 22.1. *Let $r < n$ and $x, x' \in V^r$ be such that $ann(x) = ann(x')$. Then there is an $A \in SL_n$ such that $x = A \cdot x'$.*

Proof: Use the row-echelon form construction. Make the pivots as 1 using appropriate diagonal matrices. Since $r < n$ these can be chosen from SL_n. $\square$

This decides the orbit structure of V^r for $r < n$. The largest orbit O_r is of course when $ann(x) = 0$. Thus, when $x \in V_r$ is a mtrix of rank r, we see that $ann(x) = 0$ is trivial. The *generic* element of V^r is of this form. Thus O_r is dense in V^r. For this reason, there are no non-constant invariants.

Another calculation is the computation of the closure of orbits. Let O, O' be two arbitrary orbits of V^r. We have:

Proposition 22.2. *O' lies in the closure of O if and only if $ann(O') \supseteq ann(O)$.*

Proof: One direction is clear. We prove the other direction when $ann(O) \subseteq ann(O')$ and $dim(ann(O)) = dim(ann(O')) - 1$. Thus, we may assume that $O = [x]$ and $O' = [x']$ with both x and x' such that $rowspace(x) \supseteq rowspace(x')$ with $rank(x') = rank(x) - 1$. Then, upto SL_n, we may assume that the rows $x'[1], \ldots, x'[k]$ match the first k rows of x, and that $x'[k+1] = x'[k+2] = \ldots = x'[n] = 0$. Note that $x[k+1]$ is non-zero and $x[k+2]$ exists and is zero. We then construct the matrix $A(t)$ as follows:

$$A(t) = \left[\begin{array}{c|cc|c} I_k & 0 & 0 & 0 \\ \hline 0 & t & 0 & 0 \\ 0 & 0 & t^{-1} & 0 \\ \hline 0 & 0 & 0 & I_{n-k-2} \end{array}\right]$$

We see that $A(t) \in SL_n$ for all $t \neq 0$. Next, if we let $x(t) = A(t) \cdot x$, then we see that:

$$\lim_{t \to 0} x(t) = x'$$

This shows that x' lies in the closure of O. □

We now look at the case when $r \geq n$. We have:

Proposition 22.3. *Let $r \geq n$ and $x, x' \in V^r$ be such that $ann(x) = ann(x')$. If (i) $rank(x) < n$ then there is an $A \in SL_n$ such that $x' = Ax$, (ii) if $rank(x) = n$ there is a unique $A \in SL_n$ and a $\lambda \in \mathbb{C}^*$ such that if $z = A \cdot x$, then the first $n-1$ rows of z equal those of x' and $z[n] = \lambda x'[n]$.*

The proof is easy. Let x be matrix in V^r of rank n, and C be a subset of $[r] = \{1, 2, \ldots, r\}$ such that $det(x[C]) \neq 0$.

Proposition 22.4. *Let x be as above. Then $O(x)$, the orbit of x, equals all points $x' \in V^r$ such that (i) $ann(x) = ann(x')$ and (ii) $det(X'[C]) = det(x[C])$. The set of all rank n points in V^r is dense in V^r.*

The proof is easy.

Proposition 22.5. *The orbit $O(x)$ as above is closed.*

Proof: Notice that if $A \in SL_n$ and $z = Ax$ then $det(x[C]) = det(z[C])$. We may rewrite condition (i) above as $ann(x') \supseteq ann(x)$. Condition (ii) ensures that $rank(x') = n$ and that condition (i) holds with equality. Thus is $y_1, \ldots, y_{r-n}$ are column vectors generating $ann(x)$, then the equations $x'y_s = 0$ and $det(x'[C]) = det(x[C])$ determines the orbit $O(x)$. Thus $O(x)$ is the zero-set of some algebraic equations and thus is closed. $\square$

Proposition 22.6. *Let $x \in V^r$ be such that $rank(x) < n$. Then (i) $O(x)$ equals those x' such that $ann(x) = ann(x')$, (ii) $O(x)$ is not closed and its closure $\overline{O(x)}$ equals points x' such that $ann(x') \supseteq ann(x)$.*

We now move to the computation of invariants. The space $\mathbb{C}[V^r]$ equals the space of all polynomials in the variable matrix $X = (X_{ij})$ where $i = 1, \ldots, n$ and $j = 1, \ldots, r$. It is clear that for any set C of n columns of X, we see that $det(X[C])$ is an invariant. We will denote C as $C = c_1 < c_2 < \ldots < c_n$ and $det(X[C])$ as p_C, the C-th **Plucker** coordinate. We aim to show now that these are the only invariants.

Let $C_0 = \{1 < 2 < \ldots < n\}$ and $W \subseteq V^r$ be the space of all matrices $x \in V^r$ such that $x[C_0] = diag(1, \ldots, 1, \lambda)$. Let W' be those elements of W for which $\lambda \neq 0$.

Lemma 22.1. *Let W' be as above. (i) If $x \in W'$, then $O(x) \cap W' = x$, (ii) for any $x \in V^r$ such that $det(x[C_0]) \neq 0$, there is a unique $A \in SL_n$ such that $Ax \in W'$.*

Let us call $Z' \subseteq V^r$ as those x such that $det(x[C_0]) \neq 0$. We then have the projection map

$$\pi : Z' \to W'$$

given by the above lemma. Note that Z' is SL_n-invariant: if $x \in Z'$ and $A \in SL_n$ then $A \cdot x \in Z'$ as well.

The ring $\mathbb{C}[Z']$ of regular functions on Z' is precisely $\mathbb{C}[X]_{det([X[C_0])}$, the localization of $\mathbb{C}[X]$ at $det(X[C_0])$.

We may parametrize W' as:

$$W = \left[\begin{array}{ccc|ccc} 1 & \ldots & 0 & w_{1,n+1} & \ldots & w_{1,r} \\ \vdots & & \vdots & \vdots & & \vdots \\ 0 & \ldots & w_{n,n} & w_{n,n+1} & \ldots & w_{n,r} \end{array} \right]$$

Let

$$W = \{W_{i,j} | i = 1, \ldots, n, \ j = n+1, \ldots, r\} \cup \{W_{n,n}\}$$

be a set of $(n-r) \cdot r + 1$ variables. The ring $\mathbb{C}[W']$ of regular function on W' is precisely $\mathbb{C}[W]_{W_{n,n}}$.

Let $x \in Z'$ be an arbitrary point and $A = x[C_0]$. Let $C_{i,j}$ be the set $C_0 - i + j$. The map π is given by:

$$\begin{array}{rcl} \pi(x)_{n,n} & = & det(A) \\ \pi(x)_{i,j} & = & \left\{ \begin{array}{ll} det(x[C_{i,j}])/det(A) & for\, i \neq n \\ det(x[C_{i,j}]) & otherwise \end{array} \right. \end{array}$$

The map π causes the map:

$$\pi^* : \mathbb{C}[W'] \to \mathbb{C}[Z']$$

given by:

$$\begin{array}{rcl} \pi^*(W_{n,n}) & = & det(X[C_0]) \\ \pi^*(W_{i,j}) & = & \left\{ \begin{array}{ll} det(X[C_{i,j}])/det(X[C_0]) & for\, i \neq n \\ det(X[C_{i,j}]) & otherwise \end{array} \right. \end{array}$$

Now we note that Z' is dense in V^r. Let $p \in \mathbb{C}[X]^{SL_n}$ be an invariant. Clearly p restricted to W' defines an element $p_{W'}$ of $\mathbb{C}[W']$. This then extends to Z' via π^*. Clearly $\pi^*(p_{W'})$ must match p on Z'. Thus we have that p is a polynomial in $det(X[C_{i,j}])$ possibly localized at $det(X[C_0])$.

Note that for a general C, $det(X[C])$ is already expressible in $det(X[C_0])$ and $det(X[C_{i,j}])$.

22.3 Other Representations

We discuss two other representations:

The Conjugate Action: Let $\mathcal{M}$ be the space of all $n \times n$ matrices with complex entries. we define the action of $A \in SL_n$ on $\mathcal{M}$ as follows. For an $M \in \mathcal{M}$, A acts on it by conjugation.

$$A \cdot M = AMA^{-1}$$

Note that $dim_{\mathbb{C}}(\mathcal{M}) = n^2$. Let $X = \{X_{ij} | 1 \leq i, j \leq n\}$ be the dual space to $\mathcal{M}$. The invariants for this action are $Tr(X), \ldots, Tr(X^i), \ldots, Tr(X^n)$. That these are invariants ic clear, for $Tr(AM^iA^{-1}) = Tr(M)$ for any A, M.

Also note that $Tr(X) = X_{11} + \ldots + X_{nn}$ is linear in the X's. This indicates an invariant hyperplane in $\mathcal{M}$. This is precisely the *trace zero matrices*. Thus $\mathcal{M} = \mathcal{M}_0 \oplus \mathcal{M}_1$ where $\mathcal{M}_0$ are all matrices M such that $Tr(M) = 0$. The one-dimensional complementary space $\mathcal{M}_1$ is composed of multiples of the identity matrix.

The orbits are parametrized by the **Jordan canonical form** $JCF(M)$. In other words, if $JCF(M) = JCF(M')$ them $M' \in O(M)$. Furthermore, if $JCF(M)$ is diagonal then the orbit of M is **closed** in $\mathcal{M}$.
The Space of Forms: Let V be a complex vector space of dimension n andV^* its dual. Let $X_1, \ldots, X_n$ be the dual basis. The space $Sym^d(V^*)$ consists of degree d forms are formal linear combinations of the monomials $X_1^{i_1} \ldots X_n^{i_n}$ with $i_1 + \ldots + i_n = d$. The typical form may be written as:

$$f(X) = \sum_{i_1+\ldots+i_n=d} B_{i_1\ldots i_n} X_1^{i_1} \ldots X_n^{i_n}$$

Let $A \in SL_n$ be such that

$$A^{-1} = \begin{bmatrix} a_{11} & \ldots & a_{n1} \\ \vdots & & \vdots \\ a_{1n} & \ldots & a_{nn} \end{bmatrix}$$

The space $Sym^d(V^*)$ is the formal linear combination of the generaic coefficients $\{B_{\bar{i}} || \bar{i}| = d\}$. The action of A is obtained by substituting

$$X_i \to \sum_j a_{ij} X_i$$

in $f(X)$ and recoomputing the coefficients. we illustrate this for $n = 2$ and $d = 2$. Then, the generic form is given by:

$$f(X_1, X_2) = B_{20} X_1^2 + B_{11} X_1^1 X_2^1 + B_{02} X_2^2$$

Upon substitution, we get $f(a_{11}X_1 + a_{12}X_2, a_{21}X_1 + a_{22}X_2)$: Thus the new coefficients are:

$$\begin{array}{lllll} B_{20} & \to & B_{20}a_{11}^2+ & B_{11}a_{11}a_{21}+ & B_{02}a_{21}^2 \\ B_{11} & \to & B_{20}2a_{11}a_{12}+ & B_{11}(a_{11}a_{22}+a_{12}a_{21})+ & B_{02}2a_{21}a_{22} \\ B_{02} & \to & B_{20}a_{12}^2+ & B_{11}a_{12}a_{22}+ & B_{02}a_{22}^2 \end{array}$$

We thus see that the variables B move linearly, with coefficients as homogeneous polynomials of degree 2 in the coefficients of the group elements. In the above case, we know that the *discriminant* $B_{11}^2 - 4B_{20}B_{02}$ is an invariant,

These spaces have been the subject of intense analysis and their study by Gordan, Hilbert and other heralded the beginning of commutative algebra and invariant theory.

22.4 Full Reducibility

Let W be a representation of SL_n and let $Z \subseteq W$ be an invariant subspace. The reducibility question is whether there exists a complement Z' such that $W = Z \oplus Z'$ and Z' is SL_n-invariant as well.

The above result is indeed true although we will not prove it here. There are many proofs known, each with a specific objective in a specific situation, and each extremely instructive.

The simplest is possibly through the Weyl Unitary trick. In this, a suitable **compact** subgroup $U \subseteq SL_n$ is chosen. The theory of compact groups is much like that of finite groups and a complement Z' may easily be found. It is then shown that Z' is SL_n-invariant.

The second attack is through showing the fill reducibility of the module $\otimes^d V$, the d-th tensor product representation of V. This goes through the construction of the commutator of the same module regarded as a *right* S_d-module, with the symmetric group permutating the contents of the d positions. The full reducibility then follows from Maschke's theorem and that of the finite group S_d.

The oldest approach was through the construction of a symbolic reynold's operator, which is the **Cayley** Ω-process. Let $\mathbb{C}[W]$ be the ring of polynomial functions on the space W. The operator Ω is a map:

$$\Omega : \mathbb{C}[W] \to \mathbb{C}[W]^{SL_n}$$

such that if p is an invariant then $\Omega(p \cdot p') = p \cdot \Omega(p')$. The definition of Ω is nothing but **curious**.

Chapter 23

Invariant Theory

References: [FH, N]

23.1 Algebraic Groups and affine actions

An algebraic group (over $\mathbb{C}$) is an affine variety G equipped with (i) the group product, i.e., a morphism of algebraic varieties $\cdot : G \times G \to G$ which is associative, i.e. $(g_1 \cdot g_2) \cdot g_3 = g_1 \cdot (g_2 \cdot g_3)$ for all $g_1, g_2, g_3 \in G$, (ii) and algebraic inverse $i : G \to G$, i.e, $g \cdot i(g) = i(g) \cdot g = 1_G$, where 1_G is (iii) a special element $1_G \in G$, which functions as the identity, i.e., $1_G \cdot g = g \cdot 1_G = g$ for all $g \in G$. Let $\mathbb{C}[G]$ be the ring of regular functions on G. The requirements (i)-(iii) above are via morphisms of algebraic varities, and thus the group product and the inverse must be defined algebraically. Thus the product $\cdot : G \times G \to G$ results in a morphism of $\mathbb{C}$-algebras $\cdot^* : \mathbb{C}[G] \to \mathbb{C}[G] \otimes \mathbb{C}[G]$ and the inverse into another morphism $i^* : \mathbb{C}[G] \to \mathbb{C}[G]$.

The essential example is obviously SL_n. Clearly for $G = SL_n$, we have $\mathbb{C}[G] = \mathbb{C}[X]/(det(X) - 1)$, where X is the $n \times n$ indeterminate matrix. The morphism $\cdot^*$ and i^* are clearly:

$$\begin{array}{rcl} \cdot^*(X_{ij}) & = & \sum_k X_{ik} \otimes X_{kj} \\ i^*(X_{ij}) & = & det(M_{ji}) \end{array}$$

where M_{ji} is the corresponding minor of X.

Next, let Z be another affine variety with $\mathbb{C}[Z]$ as its ring of regular functions. We say that Z is a G**-variety** if there is a morphism $\mu : G \times Z \to Z$ which is a group action. Thus, not only must G act on Z, it must do so *algebraically*. This μ induces the map μ^*:

$$\mu^* : \mathbb{C}[Z] \to \mathbb{C}[G] \otimes \mathbb{C}[Z]$$

Thus every function $f \in \mathbb{C}[Z]$ goes to a finite sum:

$$\mu^*(f) = \sum_{i=1}^{k} h_i \otimes f_i$$

where $f_i \in \mathbb{C}[Z]$ and $h_i \in \mathbb{C}[G]$ for all i.

To continue with our example, consider SL_2 and $Z = Sym^d(V^*)$. $\mathbb{C}[Z] = \mathbb{C}[B_{20}, B_{11}, B_{02}]$ and $\mathbb{C}[G] = \mathbb{C}[A_{11}, A_{12}, A_{21}, A_{22}]/(A_{11}A_{22} - A_{21}A_{12} - 1)$. We have for example:

$$\mu^*(B_{20}) = A_{11}^2 \otimes B_{20} + A_{11}A_{21} \otimes B_{11} + A_{21}^2 \otimes B_{02}$$

Let $\mu : G \times Z \to Z$ and $\mu' : G \times Z' \to Z'$ be two G-varities and let $\phi : Z \to Z'$ be a morphism. We say that ϕ is **G-equivariant** if $\phi(\mu(g, z)) = \mu'(g, \phi(z))$ for all $g \in G$ and $z \in Z$. Thus ϕ commutes with the action of G.

23.2 Orbits and Invariants

Every $g \in G$ induces an algebraic bijection on Z by restricting the map $\mu : G \times Z \to Z$ to g. We denote this map by $\mu(g)$ and call it **translation by** g. The map:

$$\mu(g) : Z \to Z$$

induces the isomorphism of $\mathbb{C}$-algebras:

$$\mu(g)^* : \mathbb{C}[Z] \to \mathbb{C}[Z]$$

Given any function $f \in \mathbb{C}[Z]$, $\mu(g)^*(f) \in \mathbb{C}[Z]$ is the *translated* function and denotes the action of G on f. This makes $\mathbb{C}[Z]$ into a G-module. If $\phi : Z \to Z'$ is a G-equivariant map, then the map $\phi^* : \mathbb{C}[Z'] \to \mathbb{C}[Z]$ is also G-equivariant, for the G action on $\mathbb{C}[Z]$ and $\mathbb{C}[Z']$ as above.

We next examine the equation:

$$\mu^*(f) = \sum_{i=1}^{k} h_i \otimes f_i$$

where $f_i \in \mathbb{C}[Z]$ and $h_i \in \mathbb{C}[G]$ for all i. For a fixed g, we see that

$$\mu(g)^*(f) = \sum_{i=1}^{k} h_i(g) \otimes f_i$$

Thus we see that every translate of f lies in the k-dimensional vector space $\mathbb{C} \cdot \{f_1, \ldots, f_k\}$. Let

$$M(f) = \mathbb{C} \cdot \{\mu(g)^*(f) | g \in G\} \subseteq \mathbb{C} \cdot \{f_1, \ldots, f_k\}$$

Clearly, $M(f)$ is a G-invariant subspace of $\mathbb{C}[Z]$. This may be generalized:

Proposition 23.1. *Let $S = \{s_1, \ldots, s_m\}$ be a finite subset of $\mathbb{C}[Z]$. Then there is a finite-dimensional G-invariant subspace $M(S)$ of $\mathbb{C}[Z]$ containing $s_1, \ldots, s_m$.*

Next, let us consider $\mu : G \times Z \to Z$ and fix a $z \in Z$. We get the map $\mu_z : G \to X$. Since G is an affine variety, we see that the image $\mu_z(G)$ is a **constructible set**, whose closure is an affine variety. The image is precisely the orbit $O(z) \subseteq Z$. The closure of $O(z)$ will be denoted by $\Delta(z)$.

If $O(z) = \Delta(z)$, then we have the G-equivariant embedding $i_z : O(z) \to Z$. This gives us the map $i_z^* : \mathbb{C}[Z] \to \mathbb{C}[O(z)]$ which is a surjection. Thus there is an ideal $I(z) = ker(i_z^*)$ such that $\mathbb{C}[O(z)] \cong \mathbb{C}[Z]/I(z)$. Since the map i_z is G-equivariant, $I(z)$ is a G-submodule of $\mathbb{C}[Z]$ and is the ideal of definition of the orbit $O(z)$.

In general, if $I \subseteq \mathbb{C}[Z]$ is an ideal which is G-invariant, then the variety of I is also G-invariant and is the union of orbits.

The second construction that we make is that of the **quotient** Z/G. Since $\mathbb{C}[Z]$ is a G-module, we examine $\mathbb{C}[Z]^G$, the subring of G-invariant functions in $\mathbb{C}[Z]$. We define Z/G as the spectrum $Spec(\mathbb{C}[Z]^G)$. The inclusion $\mathbb{C}[Z]^G \to \mathbb{C}[Z]$ gives us the quotient map:

$$\pi : Z \to Z/G$$

Exercise 23.1. *Let us consider $Z = Sym^2(V^*) \cong \mathbb{C}^3$ where V is the standard representation of $G = SL_2$. As we have seen, $\mathbb{C}[Z] = \mathbb{C}[B_{20}, B_{11}, B_{02}]$. There is only one invariant $\delta = B_{11}^2 - 4B_{02}B_{20}$. Thus $\mathbb{C}[Z]^G = \mathbb{C}[\delta]$. Thus Z/G is precisely $Spec(\mathbb{C}[\delta]) = \mathbb{C}$, the complex plane. The map π is executed as follows: given a form $aX_1^2 + bX_1X_2 + cX_2^2 \equiv (a, b, c) \in Z$, we eveluate the invariant δ at the point (a, b, c). Thus $\pi : \mathbb{C}^3 \to \mathbb{C}$ is given by:*

$$\pi(a, b, c) = b^2 - 4ac$$

Clearly, if $f, f' \in Z$ such that $f = g \cdot f'$ for $g \in SL_2$ then $\delta(f) = \delta(f')$. We look at the converse: if f, f' are such that $\delta(f) = \delta(f')$ then is it that $f \in O(f')$? We begin with $f = aX_1^2 + bX_1X_2 + cX_2^2$. We assume for the moment that $a \neq 0$. If that is the case, we make the substitution $X_1 \to X_1 - \alpha X_2$

and $X_2 \to X_2$. Note that this transformation is unimodular for all α. This transforms f to:

$$a(X_1 - \alpha X_2)^2 + b(X_1 - \alpha X_2)X_2 + cX_2^2$$

The coefficient of X_1X_2 is $-2a\alpha+b$. Thus by choosing α as $b/2a$ we see that f is transformed into $a''X_1^2 - c''X_2^2$ for some a'', c''. By a similar token, even if $a = 0$ one may do a similar transform. Thus in general, if f is not the zero form, there is a point in $O(f)$ which is of the form $aX_1^2 - cX_2^2$. Thus, we may assume that both f and f' are in this form. We can simplify the form further by a diagonal element of SL_2 to put both f and f' as $X_1^2 - cX_2^2$ and $X_1^2 - c'X_2^2$. It is now clear that $\delta(f) = \delta(f')$ implies that $c = c'$. Thus the general answer is that if $f, f' \neq 0$ and $\delta(f) = \delta(f')$ then $f \in O(f')$.

Next, let us examine the form 0. we see that $\delta(0) = 0$. Thus we see that (i) for any point $d \in \mathbb{C}$, if $d \neq 0$, then $\pi^{-1}(d)$ consists of a single orbit, (ii) $\pi^{-1}(0)$ consists of two orbits, $O(0)$ and $O(X_1^2)$, the perfect square, (iii) the orbits $O(f)$ are closed when $\delta(f) \neq 0$, (iv) $O(X_1^2)$ is not closed. Its closure includes the closed orbit 0.

The above example illustrates the utility of contructing the quotient Z/G as a variety parametrizing closed orbits albeit with some deviant points. In the example above, the discrepancy was at the point 0, wherein the pre-image is closed but decomposes into two orbits.

We state the all-important theorem linking a space and its quotient in the restricted case when G is finite and Z is a finite-dimensional G-module. Thus $\mathbb{C}[Z]$ is a polynomial ring and the action of G is homogeneous.

Theorem 3. *Let G be a finite group and act on the space Z. Let $R = \mathbb{C}[Z]$ and $R^G = \mathbb{C}[Z]^G$ be the ring of invariants. Let $\pi : Z \to Z/G$ be the quotient map. Then*

(i) For any ideal $J \subseteq R^G$, we have $(J \cdot R) \cap R^G = J$.

(ii) The map π is surjective. Further, for any $x \in Z/G$, $\pi^{-1}(x)$ is a single orbit in Z.

Proof: Let $f_1, \ldots, f_k$ be elements of R^G and $f \in R^G$ be such that:

$$f = r_1 f_1 + \ldots + r_k f_k$$

where r_i are elements of R. Since G is finite, we apply the Reynolds operator

$p \to p^G$. In other words, we have:

$$\begin{aligned}
f &= \frac{1}{|G|}\sum_{g\in G} f \\
&= \frac{1}{|G|}\sum_{g\in G}\sum_i r_i f_i \\
&= \frac{1}{|G|}\sum_i f_i \sum_{g\in G} r_i \\
&= \sum_i f_i \frac{1}{|G|}\sum_{g\in G} r_i \\
&= \sum_i f_i r_i^G
\end{aligned}$$

Note that we have used the fact that if $h \in R^G$ and $p \in R$ then $(h \cdot p)^G = h \cdot p^G$. Thus we see that any element f of R^G which is expressible as an R-linear combination of elements f_i's of R^G is already expressible as an R^G-linear combination of the same elements. This proves (i) above.

Now we prove (ii). Firstly, let J be a maximal ideal of R^G. By part (i) above, $J \cdot R$ is a proper ideal of R and thus π is surjective. Let $x \in Z/G$ and $J_x \subseteq R^G$ be the maximal ideal for the point x. Let z be such that $\pi(z) = x$ and let $I_{O(x)} \subseteq R$ be the ideal of all functions in R vanishing at all points of the orbit $O(z)$ of z. We show that $rad(J_x \cdot R) = I_{O(z)}$ which proves (ii). Towards that, it is clear that (a) $J_x R \subseteq I_{O(z)}$ and (b) the variety of $J_x R$ is G-invariant. Let $O(z')$ is another orbit in the variety of $J_x \cdot R$. If $O(z') \neq O(z)$ then we already have that there is a $p \in R^G$ such that $p(O(z)) = 0$ and $p(O(z')) = 1$. Since this $p \in J_x \cdot R$, the variety of $J_x \cdot R$ excludes the orbit $O(z')$ proving our claim. $\square$

Theorem 4. *Let Z be a finite-dimesional G-module for a finite group G. Then $\mathbb{C}[Z]^G$, the ring of G-invariants, is finitely generated as a $\mathbb{C}$-algebra.*

Proof: Since the action of G is homogeneous, every invariant $f \in \mathbb{C}[Z]^G$ is the sum of homogeneous invariants. Let $I_G \subseteq \mathbb{C}[Z]$ be the ideal generated by the **positive** degree invariants. By the Hilbert Basis theorem, $I_G = (f_1, \ldots, f_k) \cdot \mathbb{C}[Z]$, where each f_i is itself an invariant. we claim that $\mathbb{C}[Z]^G = \mathbb{C}[f_1, \ldots, f_k] \subseteq \mathbb{C}[Z]$, or in other words, every invariant is a polynomial over $\mathbb{C}$ in the terms $f_1, \ldots, f_k$. To this end, let f be a homogeneous invariant. We prove this by induction over the degree d of f. Since $f \in I_G$, we have

$$f = \sum_i r_i f_i \quad \text{where for all } i,\ r_i \in \mathbb{C}[Z]$$

By applying the reynolds operator, we have:

$$f = \sum_i h_i f_i \text{ where for all } i,\ h_i \in \mathbb{C}[Z]^G$$

Now since each h_i has degree less than d, by our induction hypothesis, each is an element of $\mathbb{C}[f_1, \ldots, f_k]$. This then shows that f too is an element of $\mathbb{C}[f_1, \ldots, f_k]$. $\Box$

23.3 The Nagata Hypothesis

We now generalize the above two theorems to the case of more general groups. This generlization is possible if the group G satisfies what we call the *Nagata Hypothesis*.

Definition 23.1. *Let G be an algebraic group. We say that G satisfies the Nagata hypothesis if for every finite-dimensional module M over $\mathbb{C}$ and a G-invariant hyperplane $N \subseteq M$, there is a decomposition $M = H \oplus P$ as G-modules, where H is a hyperplane and P is an invariant (i.e., P is the trivial representation.*

The group SL_n satisfies the hypothesis, and so do the so-called reductive groups over $\mathbb{C}$.

Theorem 5. *Let G satisfy the Nagata hypothesis and let Z be an affine G-variety. The ring $\mathbb{C}[Z]^G$ is finitely generated as a $\mathbb{C}$-algebra.*

The proof will go through several steps.

Let $f \in \mathbb{C}[Z]$ be an arbitrary element. We define two modules $M(f)$ and $N(f)$.

$$\begin{aligned} M(f) &= \mathbb{C} \cdot \{s \cdot f | s \in G\} \\ N(f) &= \mathbb{C} \cdot \{s \cdot f - t \cdot f | s, t \in G\} \end{aligned}$$

Thus $M(f) \supseteq N(f)$ are finite-dimensional submodules of $\mathbb{C}[Z]$.

Lemma 23.1. *Let $f \in \mathbb{C}[Z]$ be an arbitrary element. Then there exists an $f^* \in \mathbb{C}[Z]^G \cap M(f)$ such that $f - f^* \in N(f)$.*

Proof: We prove this by induction over $dim(M(f))$. Note that $s \cdot f = f + (s \cdot f - 1 \cdot f)$ and thus $M(f) = N(f) \oplus \mathbb{C} \cdot f$ as vector spaces. Thus $N(f)$ is a G-invariant hyperplane of $M(f)$. By the Nagata hypothesis, $M(f) = f' \oplus H$, where f' is an invariant. If $f' \notin N(f)$ then $M(f) = f' \oplus N(f)$ and the lemma follows. However, if $f' \in N(f)$, then $f = f' + h$ where $h \in H$. Since H is a G-module, we see that $M(h) \subseteq H$ which is of dimension less than that of $M(f)$. Thus there is an invariant h^* such that $h - h^* \in N(h)$. Note that $f - f' = h$ with f' invariant, implies that $s \cdot h - t \cdot h = s \cdot f - t \cdot f$. Thus $N(h) \subseteq N(f)$. We take $f^* = h^*$. Examining $f - f^*$, we see that

$$f - h^* = f' + h - h^* \in N(f)$$

This proves the lemma. $\square$

Lemma 23.2. *Let $f_1, \ldots, f_r \in \mathbb{C}[Z]^G$. Then $\mathbb{C}[Z]^G \cap (\sum_i \mathbb{C}[Z] \cdot f_i) = \sum_i \mathbb{C}[Z]^G \cdot f_i$. Thus if an invariant f is a $\mathbb{C}[Z]$linear combination of invariants, then it is already a $\mathbb{C}[Z]^G$ linear combination.*

Proof: This is proved by induction on r. Say $f = \sum_{i=1}^r h_i f_i$ with $h_i \in \mathbb{C}[Z]$. Applying the above lemma to h_r, there is an $h'' \in \mathbb{C}[Z]^G$ and an $h' \in N(h_k)$ such that $h_r = h' + h''$. We tackle h' as follows. Since f is an invariant, we have for all $s, t \in G$:

$$\sum_{i=1}^{r} (s \cdot h_i - t \cdot h_i) f_i = s \cdot f - t \cdot f = 0$$

Hence:

$$(s \cdot h_r - t \cdot h_r) f_r = \sum_{i=1}^{r-1} (s \cdot h_i - t \cdot h_i) f_i$$

It follows from this that $h' f_r = \sum_{i=1}^{r-1} h'_i f_i$ for some $h'_i \in \mathbb{C}[Z]$. Substituting this in the expression

$$f = \sum_{i=1}^{r-1} h_i f_i + (h' + h'') f_r$$

we get:

$$f - h'' f_r = \sum_{i=1}^{r-1} (h_i + h'_i) f_i$$

Thus the invariant $f - h'' f_r$ is $\mathbb{C}[Z]$-linear combination of $r - 1$ invariants, and thus the induction hypothesis applies. This then results in an expression for f as a $\mathbb{C}[Z]^G$-linear combination for f. $\square$

The above lemma proves part (i) of Theorem 3 for groups G for which the Nagata hypothesis holds. The route to Theorem 4 is now a straightforward adaptation of its proof for finite groups. In the case when Z is a G-module, $\mathbb{C}[Z]^G$ would then be homogeneous and the proof of Theorem 4 holds. For general Z here is a trick which converts it to the homogeneous case:

Proposition 23.2. *Let Z be an affine G-variety. Then there is a G-module W and an equivariant embedding $\phi : Z \to W$.*

Proof: Since $\mathbb{C}[Z]$ is a finitely generated $\mathbb{C}$-algebra, $\mathbb{C}[Z] = \mathbb{C}[f_1, \ldots, f_k]$ where $f_1, \ldots, f_k$ are some specific elements of $\mathbb{C}[Z]$. By expanding the list of generators, we may assume that the vector space $\mathbb{C} \cdot \{f_1, \ldots, f_k\} \subseteq \mathbb{C}[Z]$ is a finite-dimensional G-module. Let us construct W as an isomorphic copy of this module with the dual basis $W_1, \ldots, W_k$. We construct the map $\phi^* : \mathbb{C}[W] \to \mathbb{C}[Z]$ by defining $\phi(W_i) = f_i$. Since $\mathbb{C}[W]$ is a free algebra, we indeed have the surjection:

$$\phi^* : \mathbb{C}[W] \to \mathbb{C}[Z]$$

This proves the proposition. $\square$

We are now ready to prove:

Theorem 6. *Let G satisfy the Nagata hypothesis. Let Z be an affine G-variety with coordinate ring $\mathbb{C}[Z]$. Then $\mathbb{C}[Z]^G$ is finitely generated as a $\mathbb{C}$-algebra.*

Proof: We construct the equivariant surjection $\phi^* : \mathbb{C}[W] \to \mathbb{C}[Z]$. Note that $\mathbb{C}[W]^G$ is already finitely generated over $\mathbb{C}$, say $\mathbb{C}[W]^G = \mathbb{C}[h_1, \ldots, h_k]$. We claim that $\phi^*(h_1), \ldots, \phi^*(h_k)$ generate $\mathbb{C}[Z]^G$.

Now, let $f \in \mathbb{C}[Z]^G$. By the surjectivity of ϕ^*, there is an $h \in \mathbb{C}[W]$ such that $\phi^*(h) = f$. Consider the space $N(h)$. A typical generator of $N(h)$ is $s \cdot h - t \cdot h$. Applying ϕ^* to this, we see that:

$$\begin{aligned} \phi^*(s \cdot h - t \cdot h) &= s \cdot \phi^*(h) - t \cdot \phi^*(h) \\ &= f - f = 0 \end{aligned}$$

By an earlier lemma there is an invariant h^* such that $h - h^* \in N(h)$. thus applying ϕ^* we see that $\phi^*(h^*) = \phi(h) = f$. Thus there is an invariant h^* such that $\phi^*(h^*) = f$. Now $h^* \in \mathbb{C}[h_1, \ldots, h_k]$ implies that $f = \phi^*(h^*) \in \mathbb{C}[\phi^*(h_1), \ldots, \phi^*(h_k)]$. $\square$

Chapter 24

Orbit-closures

Reference: [Ke, N]

In this chapter we will analyse the validity of Theorem 3 for general groups G with the Nagata hypothesis. The objective is to analyse the map $\pi : Z \to Z/G$.

Proposition 24.1. *Let G satisfy the Nagata hypothesis and act on the affine variety Z. Then the map $\pi : Z \to Z/G$ is surjective.*

Proof: Let $J \subseteq \mathbb{C}[Z]^G$ be a maximal ideal. By lemma 23.2 $J \cdot \mathbb{C}[Z]$ is a proper ideal of $\mathbb{C}[Z]$. This implies that π is surjective. □

Theorem 7. *Let G satisfy the Nagata hypothesis and act on the affine variety Z. Let W_1 and W_2 be G-invariant (Zariski-)closed subsets of Z such that $W_1 \cap W_2$ is empty. Then there is an invariant $f \in \mathbb{C}[Z]^G$ such that $f(W_1) = 0$ and $f(W_2) = 1$.*

Proof: Let I_1 and I_2 be the ideals of W_1 and W_2 in $\mathbb{C}[Z]$. Since their intersection is empty, by Hilbert Nullstellensatz, we have $I_1 + I_2 = 1$. Whence there are functions $f_1 \in I_1$ and $f_2 \in I_2$ such that $f_1 + f_2 = 1$. For arbitrary $s, t \in G$ we have:

$$(s \cdot f_1 - t \cdot f_1) + (s \cdot f_2 - t \cdot f_2) = 1 - 1 = 0$$

Note that $M(f_i)$ and $N(f_i)$ are submodules of the G-invariant ideal I_i. Now applying lemma 23.1 to f_1, we see that there are elements $f_i' \in N(f_i)$ such that $f_1 + f_1' \in \mathbb{C}[Z]^G$. Whence we see that:

$$(f_1 + f_1') + (f_2 + f_2') = 1$$

where $f = f_1 + f_1'$ is an invariant. Since $f_1 + f_1' \in I_1$ we see that $f(W_1) = 0$. On the other hand, $f = 1 - (f_2 + f_2') \in 1 + I_2$ and thus $f(W_2) = 1$. $\Box$

Recall that $\Delta(z)$ denotes the closure of the orbit $O(z)$ in Z.

Theorem 8. *Let G satisfy the Nagata hypothesis and act on an affine variety Z. We define the relation $\approx$ on Z as follows: $z_1 \approx z_2$ if and only if $\Delta(z_1) \cap \Delta(z_2)$ is non-empty. Then*

(i) $\approx$ *is an equivalence relation.*

(ii) $z_1 \approx z_2$ *iff* $f(z_1) = f(z_2)$ *for all* $f \in \mathbb{C}[Z]^G$.

(iii) Within each $\Delta(z)$ there is a unique closed orbit, and this is of minimum dimension among all orbits in $\Delta(z)$.

Proof: It is clear that (ii) proves (i). Towards (ii), if $\Delta(z_1) \cap \Delta(z_2)$ is empty, then by Theorem 7, there is an invariant separating the two points. Thus it remains to show that if $\Delta(z_1 0 \cap \Delta(z_2)$ is *non-empty*, and f is any invariant, then $f(z_1) = f(z_2)$. So let z be an element of the intersection. Since f is an invariant, we have $f(O(z_1))$ is a constant, say α. Since $O(z_1) \subseteq \Delta(z_1)$ is dense, and f is continuous, we have $f(z) = \alpha$. Thus $f(z_1) = f(z_2) = f(z) = \alpha$.

Now (iii) is easy. Clearly $\Delta(z)$ cannot have *two* closed distinct orbits, for otherwise they woulbe separated by an invariant. But this must take the same value on all points of $\Delta(z)$. That it is of minimum dimension follows algebraic arguments. $\Box$

Definition 24.1. *Let Z be a G-variety and $z \in Z$. We say that z is* **stable** *if the orbit $O(z) \subseteq Z$ is closed in Z.*

By the above theorem, every point x of $Spec(C[Z]^G)$ corresponds to exactly one stable point: the point whose orbit is of minimum dimension in $\pi^{-1}(x)$.

Exercise 24.1. *Consider the action of $G = SL_n$ on $\mathcal{M}$, the space of $n \times n$-matrices by conjugation. Thus, given $A \in SL_n$ and $M \in \mathcal{M}$, we have:*

$$A \cdot M = AMA^{-1}$$

Let $R = \mathbb{C}[\mathcal{M}] = \mathbb{C}[X_{11}, \ldots, X_{nn}]$ be the ring of functions on $\mathcal{M}$. The invariants R^G is generated as a $\mathbb{C}$-algebra by the forms $e_i(X) = Tr(X^i)$, for $i = 1, \ldots, n$. The forms $\{e_i | 1 \leq i \leq n\}$ are algebraically independent and *thus R^G is the polynomial ring $\mathbb{C}[e_1, \ldots, e_n]$. Clearly then $Spec(R^G) \cong \mathbb{C}^n$ and we have:*

$$\pi : \mathcal{M} \cong \mathbb{C}^{n^2} \to \mathbb{C}^n$$

Given a matrix M with eigenvalues $\{\lambda_1, \ldots, \lambda_n\}$, we have:

$$e_i(M) = \lambda_1^i + \ldots + \lambda_n^i$$

Thus, by the fundamental theorem of algebra, the image $\pi(M)$ determines the set $\{\lambda_1, \ldots, \lambda_n\}$ (with multiplicities). On the other hand, given any tuple $\mu = (\mu_1, \ldots, \mu_n)$ there is a unique set $\lambda_\mu = \{\lambda_1, \ldots, \lambda_n\}$ such that $\sum_r \lambda_r^i = \mu_i$. Clearly, for the diagonal matrix $D(\lambda_\mu) = diag(\lambda_1, \ldots, \lambda_n)$, we have that $\pi(D(\lambda)) = \mu$. This verifies that π is surjective.

For a given μ, the set $\pi^{-1}(\mu)$ are all matrices M with $Spec(M) = \lambda = \lambda_\mu$. By the Jordan canonical form (JCF), this set may be stratified by the various Jordan canonical blocks of spectrum λ. If λ has no multiplicities then $\pi^{-1}(\mu)$ consists of just one orbit: matrices M such that $JCF(M) = D(\lambda)$. For a general λ, the orbit of M with $JCF(M) = D(\lambda)$ is the unique closed orbit of minimum dimension. All other orbits contain this orbit in its closure. Thus stable points $M \in \mathcal{M}$ are the diagonalizable matrices.

As an example, consider the case when $n = 2$ and the matrix:

$$N = \begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix}$$

Consider the family $A(t) = diag(t, t^{-1}) \in SL_2$. We see that:

$$N(t) = A(t)NA(t)^{-1} = \begin{bmatrix} \lambda & t^2 \\ 0 & \lambda \end{bmatrix}$$

Thus $\lim_{t \to 0} N(t) = diag(\lambda, \lambda)$, the diagonal matrix.

Thus, we see that the invariant ring $\mathbb{C}[Z]^G$ puts a different equivalence relation $\approx$ on points in Z which is coarser than $\cong$, the orbit equivalence relation. The relation $\approx$ is more 'topological' than group-theoretic and correctly classifies orbits by their separability by invariants. The special case of Theorem 8 when Z is a representation was analysed by Hilbert in 1893. The point $0 \in Z$ is then the smallest closed orbit, and the equivalence class $[0]_\approx$ is termed as the **null-cone** of Z. We see that the null-cone consists of all points $z \in Z$ such that 0 lies in the orbit-closure $\Delta(z)$ of z. It was Hilbert who discovered that if $0 \in \Delta(z)$ then 0 lies in the orbit-closure of a 1-parameter diagonal subgroup of SL_n. To understand the intricay of Hilbert's constructions, it is essential that we understand diagonal subgroups of SL_n.

Chapter 25

Tori in SL_n

Reference: [Ke, N]

Let $\mathbb{C}^*$ denote the multiplicative group of non-zero complex numbers. A torus is the abstract group $(\mathbb{C}^*)^m$ for some m. Note that $\mathbb{C}^*$ is an abelian algebraic group with $\mathbb{C}[G] = \mathbb{C}[T, T^{-1}]$. Furthermore, $\mathbb{C}^*$ has a compact subroup $S^1 = \{z \in \mathbb{C}, |z| = 1\}$, the unit circle.

Next, let us look at representations of tori. For $\mathbb{C}^*$, the simplest representations are indexed by integers $k \in \mathbb{Z}$. So let $k \in \mathbb{Z}$. The representation $\mathbb{C}_{[k]}$ corresponds to the 1-dimensional vector space $\mathbb{C}$ with the action:

$$t \cdot z = t^k z$$

Thus a non-zero $t \in \mathbb{C}^*$ acts on z by multiplication of the k-th power. Next, for $(\mathbb{C}^*)^m$, let $\chi = (\chi[1], \ldots, \chi[m])$ be a sequence of integers. For such a χ, we define the representation $\mathbb{C}_\chi$ as follows: Let $\overline{t} = (t_1, \ldots, t_m) \in (\mathbb{C}^*)^m$ be a general element and $z \in \mathbb{C}$. The action is given by:

$$\overline{t} \cdot z = t_1^{\chi[1]} \ldots t_m^{\chi[m]} z$$

Such a χ is called a **character** of $(\mathbb{C}^*)^m$.

These 1-dimensional representations of tori are crucial in the analysis of algebraic group actions.

Let us begin by understanding the structure of algebraic homomorphisms from $\mathbb{C}^*$ to $SL_n(\mathbb{C})$. So let

$$\lambda : \mathbb{C}^* \to SL_n(\mathbb{C})$$

be such a map such that $\lambda(t) = [a_{ij}(t)]$ where $a_{ij}(T) \in \mathbb{C}[T, T^{-1}]$. An important substitution is for $t = e^{i\theta}$, and we obtain a 2π-periodic map

$$\overline{\lambda} : \mathbb{C} \to SL_n$$

We see that $\overline{\lambda}(0) = I$. Let the derivative at 0 for $\overline{\lambda}$ be X.

We have the following general lemma:

Lemma 25.1. *Let $f : \mathbb{C} \to SL_n$ be a smooth map such that $f(0) = I$ and $f'(0) = X$, where X is an $n \times n$-matrix. Then for $\theta \in \mathbb{C}$,*

$$\lim_{k \to \infty} \left[f\left(\frac{\theta}{k} \right) \right]^k = e^{\theta X}$$

The proof follows from the local diffeomorphism of the exponential map in the neighborhood of the identity matrix.

Applying this lemma to $\overline{\lambda}$ we see that $e^{\theta X}$ is in the image of $\overline{\lambda}$ for all θ. Now, since $\overline{\lambda}$ is 2π-periodic, we must have $e^{(n2\pi+\theta)X} = e^{\theta X}$. This forces (i) X to be diagonalizable, and (ii) with integer eigenvalues. This proves:

Proposition 25.1. *Let $\lambda : \mathbb{C}^* \to SL(V)$ be an algebraic homomorphism. Then the image of λ is closed and $V \cong \mathbb{C}_{[m_1]} \oplus \ldots \oplus \mathbb{C}_{[m_n]}$, for some integers $m_1, \ldots, m_n$, where $n = dim_{\mathbb{C}}(V)$.*

Based on this, we have the generalization:

Proposition 25.2. *Let $\lambda : (\mathbb{C}^*)^r \to SL(V)$ be an algebraic homomorphism. Then the image of λ is closed and $V \cong \mathbb{C}_{\chi_1} \oplus \ldots \oplus \mathbb{C}_{\chi_n}$, for some integers $\chi_1, \ldots, \chi_n$, where $n = dim_{\mathbb{C}}(V)$.*

Thus, in effect, for every homomorphism $\lambda : (\mathbb{C}^*)^r \to SL_n$, there is a fixed invertible matrix A such that for all $\overline{t} \in (\mathbb{C}^*)^r$, the conjugate $A\lambda(\overline{t})A^{-1}$ is diagonal.

A torus in SL_n is defined as an abstract subgroup H of SL_n which is isomorphic to $(\mathbb{C}^*)^r$ for some r. The **standard maximal torus** D of SL_n is the diagonal matrices $diag(t_1, \ldots, t_n)$ where $t_i \in \mathbb{C}^*$ and $t_1 t_2 \ldots t_n = 1$.

This clears the way for the important theorem:

Theorem 9. *(i) Every torus is contained in a maximal torus. All maximal tori in SL_n are isomorphic to $(\mathbb{C}^*)^{n-1}$.*

(ii) If T and T' are two maximal tori then there is an $A \in SL_n$ such that $T' = ATA^{-1}$. Thus all maximal tori are conjugate to D above.

(iii) Let $N(D)$ be the normalizer of D and $N(D)^o$ be the connected component of $N(D)$. Then $N(D)^o = D$ and $N(D)/D$ is the **Weyl group** *W, isomorphic to the symmetric group S_n.*

Definition 25.1. *Let G be an algebraic group. $\Gamma(G)$ will denote the collection of all* **1-parameter subgroups** *of G, i.e., morphisms $\lambda : \mathbb{C}^* \to G$. $X(G)$ will denote the collection of all* **characters** *of G, i.e., homomorphisms $\chi : G \to \mathbb{C}^*$.*

We consider the case when $G = (\mathbb{C}^*)^r$. Clearly, for a given $\lambda : \mathbb{C}^* \to G$, there are integers $m_1, \ldots, m_r \in \mathbb{Z}$ such that:

$$\lambda(t) = (t^{m_1}, \ldots, t^{m_r})$$

In the same vein, for the character $\chi : G \to \mathbb{C}^*$, we have integers $a_1, \ldots, a_r$ such that

$$\chi(t_1, \ldots, t_r) = \prod_i t_i^{a_i}$$

We also have the composition $\lambda \circ \chi : \mathbb{C}^* \to \mathbb{C}^*$, where by:

$$\lambda \circ \chi(t) = t^{m_1 a_1 + \ldots + m_r a_r}$$

Consolidating all this, we have:

Theorem 10. *Let $G = (\mathbb{C}^*)^r$. Then $\Gamma(G) \cong \mathbb{Z}^r$ and $X(G) \cong \mathbb{Z}^r$. Furthermore, there is the pairing*

$$<\,,\,>: \Gamma(G) \times X(G) \to \mathbb{Z}$$

which is a unimodular pairing on lattices.

Exercise 25.1. *Let $G = (\mathbb{C}^*)^3$ and λ and χ be as follows:*

$$\begin{aligned} \lambda(t) &= (t^3, t^{-1}, t^2) \\ \chi(t_1, t_2, t_3) &= t_1^{-1} t_2 t_3^2 \end{aligned}$$

Then, $\lambda \cong [3, -1, 2]$ and $\chi \cong [-1, 1, 2]$. We evaluate the pairing:

$$< \lambda, \chi >= 3 \cdot -1 + (-1) \cdot 1 + 2 \cdot 2 = 0$$

We now turn to the special case of $D \subseteq SL_n$, the maximal torus which is isomorphic to $(\mathbb{C}^*)^{n-1}$. By the above theorem, $\Gamma(D), X(D) \cong \mathbb{Z}^{n-1}$. However, it will more convenient to identify this space as a subset of $\mathbb{Z}^n$. So let:

$$\mathbb{Y}^n = \{[m_1, \ldots, m_n] \in \mathbb{Z}^n | m_1 + \ldots + m_n = 0\}$$

It is easy to see that $\mathbb{Y}^n \cong \mathbb{Z}^{n-1}$. In fact, we will set up a special bijection $\theta : \mathbb{Y}^n \to \mathbb{Z}^{n-1}$ defined as:

$$\theta([m_1, m_2, \ldots, m_n]) = [m_1, m_1 + m_2, \ldots, m_1 + \ldots + m_{n-1}]$$

The inverse θ^{-1} is also easily computed:

$$\theta^{-1}[a_1, \ldots, a_{n-1}] = [a_1, a_2 - a_1, a_3 - a_2, \ldots, a_{n-1} - a_{n-2}, -a_{n-1}]$$

This θ corresponds to the $\mathbb{Z}$-basis of $\mathbb{Y}^n$ consisting of the vectors $e_1 - e_2, \ldots, e_{n-1} - e_n$ where e_i is the standard basis of $\mathbb{Z}^n$. This is also equivalent to the embedding $\theta^* : (\mathbb{C}^*)^{n-1} \to D$ as follows:

$$(t_1, \ldots, t_{n-1}) \to \begin{bmatrix} t_1 & 0 & \ldots & & 0 \\ 0 & t_1^{-1}t_2 & 0 & \ldots & 0 \\ & & \vdots & & \\ 0 & \ldots & 0 & t_{n-2}^{-1}t_{n-1} & 0 \\ 0 & & \ldots & 0 & t_{n-1}^{-1} \end{bmatrix}$$

A useful computation is to consider the inclusion $D \subseteq D^*$, where $D^* \subseteq GL_n$ is subgroup of *all* diagonal matrices. Clearly $\Gamma(D) \subseteq \Gamma(D^*)$, however there is a surjection $X(D^*) \to X(D)$. It will be useful to work out this surjection explicitly via θ and θ^*. If $[m_1, \ldots .m_n] \in \mathbb{Z}^n \cong X(D^*)$, then it maps to $[m_1 - m_2, \ldots, m_{n-1} - m_n] \in \mathbb{Z}^{n-1} \cong X((\mathbb{C}^*)^{n-1})$ via θ^*. If we push this back into $\mathbb{Y}^n$ via θ^{-1}, we get:

$$[m_1, \ldots, m_n] \to [m_1 - m_2, 2m_2 - m_1 - m_3, , \ldots, 2m_{n-1} - m_{n-2} - m_n, m_n - m_{n-1}]$$

We are now ready to define the **weight spaces** of an SL_n-module W. So let W be such a module. By restricting this module to $D \subseteq G$, via Proposition 25.2, we see that W is a direct sum $W = \mathbb{C}_{\chi_1} \oplus \ldots \oplus \mathbb{C}_{\chi_N}$, where $N = dim_{\mathbb{C}}(W)$. Collecting identical characters, we see that:

$$W = \oplus_{\chi \in X(D)} \mathbb{C}_\chi^{m_\chi}$$

Thus W is a sum of m_χ copies of the module $\mathbb{C}_\chi$. Clearly $m_\chi = 0$ for all but a finite number, and is called the **multiplicity** of χ. For a given module W, computing m_χ is an intricate combinatorial exercise and is given by the celebrated **Weyl Character Formula**.

Exercise 25.2. *Let us look at SL_3 and the weight-spaces for some modules of SL_3. All modules that we discuss will also be GL_3-modules and thus D^* modules. The formula for converting D^*-modules to D-modules will be useful. This map is $\mathbb{Z}^3 \to \mathbb{Y}^3$ and is given by:*

$$[m_1, m_2, m_3] \to [m_1 - m_2, 2m_2 - m_1 - m_3, m_3 - m_2]$$

The simplest SL_n module is $\mathbb{C}^3$ with the basis $\{X_1, X_2, X_3\}$ with D^* weights $[1,0,0], [0,1,0]$ and $[0,0,1]$. This converted to D-weights give us $\{[1,-1,0], [-1,2,-1], [0,-1,1]\}$, with $\mathbb{C}_{[1,-1,0]} \cong \mathbb{C} \cdot X_1$ and so on.

The next module is $Sym^2(\mathbb{C}^3)$ with the basis X_i^2 and X_iX_j. The six D^* and D-weights with the weight-spaces are given below:

D^*-wieghts	*D-weights*	*weight-space*
$[2,0,0]$	$[2,-2,0]$	X_1^2
$[0,2,0]$	$[-2,4,-2]$	X_2^2
$[0,0,2]$	$[0,-2,2]$	X_3^2
$[0,1,1]$	$[-1,1,0]$	X_2X_3
$[1,0,1]$	$[1,-2,1]$	X_1X_3
$[1,1,0]$	$[0,1,-1]$	X_1X_2

The final example is the space of 3×3-matrices $\mathcal{M}$ acted upon by conjugation. We see at once that $\mathcal{M} = \mathcal{M}_0 \oplus \mathbb{C}\cdot I$ where $\mathcal{M}_0$ is the 8-dimensional space of trace-zero matrices, and $\mathbb{C}\cdot I$ is 1-dimensional space of multiples of the idenity matrix. Weight vectors are E_{ij}, with $1 \leq i,j \leq 3$. The D^* weights are $[1,-1,0], [1,0,-1], [0,1,-1], [-1,0,1], [0,-1,1], [-1,1,0]$ and $[0,0,0]$. The multiplicity of $[0,0,0]$ in $\mathcal{M}$ is 3 and in $\mathcal{M}_0$ is 2. Note that $E_{ii} \notin \mathcal{M}_0$. The D-weights are $[2,-3,1], [1,0,-1], [-1,3,-2]$ and its negatives, and obviously $[0,0,0]$.

The normalizer $N(D)$ gives us an action of $N(D)$ on the weight spaces. If w is a weight-vector of weight χ, $t \in D$ and $g \in N(D)$, then $g\cdot w$ is also a weight vector. Afterall $t\cdot(g\cdot w) = g\cdot t'\cdot w$ where $t' = g^{-1}tg$. Thus

$$t\cdot(g\cdot w) = \chi(t')(g\cdot w)$$

whence $g\cdot w$ must also be a weight-vector with some weight χ'. This χ' is easily computed via the action of D^*. Here the action of $N(D^*)$ is clear: if $\chi = [m_1, \ldots, m_n]$, then $\chi' = [m_{\sigma(1)}, \ldots, m_{\sigma(n)}]$ for some permutation $\sigma \in S_n$ determined by the component of $N(D^*)$ containing g. Thus the map χ to χ' for D-weights in the case of SL_3 is as follows:

$$[m_1-m_2, 2m_2-m_1-m_3, m_3-m_2] \to [m_{\sigma(1)}-m_{\sigma(2)}, 2m_{\sigma(2)}-m_{\sigma(1)}-m_{\sigma(3)}, m_{\sigma(3)}-m_{\sigma(2)}]$$

Caution: Note that though $\mathbb{Y}^3 \subseteq \mathbb{Z}^3$ is an S_3-invariant subset, the action of S_3 on $\chi \in \mathbb{Y}^3$ is **different**. Note that, e.g., in the last example above, $[2,-3,1]$ is a weight but not the 'permuted' vector $[-3,2,1]$. This is because of our peculiar embedding of $\mathbb{Z}^{n-1} \to \mathbb{Y}^n$.

Chapter 26

The Null-cone and the Destabilizing flag

Reference: [Ke, N]

The fundamental result of **Hilbert** states:

Theorem 11. *Let W be an SL_n-module, and let $w \in W$ be an element of the null-cone. Then there is a 1-parameter subgroup $\lambda : \mathbb{C}^* \to SL_n$ such that*

$$\lim_{t \to 0} \lambda(t) \cdot w = 0_W$$

In other words, if the zero-vector 0_W lies in the orbit-closure of w, then there is a 1-parameter subgroup taking it there, in the limit. We will not prove this statement here. Our objective for this chapter is to interpret the geometric content of the theorem. We will show that there is a *standard form* for an element of the null-cone. For well-known representations, this standard form is easily identified by geometric concepts.

26.1 Characters and the half-space criterion

To begin, let D be the fixed maximal torus. For any $w \in W$, we may express:

$$w = w_1 + w_2 + \ldots + w_r$$

where $w_i \in W_{\chi_i}$, the weight-space for character χ_i. Note the the above expression is unique if we insist that each w_i be non-zero. The set of characters $\{\chi_1, \ldots, \chi_r\}$ will be called the **support** of w and denoted as $supp(w)$. Let

$\lambda : \mathbb{C}^* \to SL_n$ be such that $Im(\lambda) \subseteq D$. In this case, the action of $t \in \mathbb{C}^*$ via λ is easily described:

$$t \cdot w = t^{(\lambda,\chi_1)} w_1 + \ldots + t^{(\lambda,\chi_r)} w_r$$

Thus, if $\lim_{t\to 0} t \cdot w$ exists (and is 0_W), then for all $\chi \in supp(w)$, we have $(\lambda, \chi) \geq 0$ (and further $(\lambda, \chi) > 0$).

Note that (λ, χ) is implemented as a linear functional on $\mathbb{Y}^n$. Thus, if $\lim_{t\to 0} t \cdot w$ exists (and is $)_W$) then there is a **hyperplane** in $\mathbb{Y}^n$ such that the support of w is on one side of the hyperplane (**strictly** on one side of the hyperplane). The normal to this hyperplane is given by the conversion of λ into $\mathbb{Y}^n$ notation.

On the other hand if the support $supp(w)$ enjoys the geometric/combinatorial property, then by the approximability of reals by rationals, we see that there is a λ such that $\lim_{t\to 0} t \cdot w$ exists (and is zero).

Thus for 1-parameter subgroups of D, Hilbert's theorem translates into a combinatorial statement on the lattice subset $supp(w) \subset \mathbb{Y}^n$. We call this the (**strict**) **half-space** property. In the general case, we know that given any $\lambda : \mathbb{C}^* \to SL_n$, there is a maximal torus T containing $Im(\lambda)$. By the conjugacy result on maximal tori, we know that $T = ADA^{-1}$ for some $A \in SL_n$. Thus, we may say that w is in the null-cone iff there is a translate $A \cdot w$ such that $supp(A \cdot w)$ satisfies the strict half-space property.

Exercise 26.1. *Let us consider SL_3 acting of the space of forms of degree 2. For the standard torus D, the weight-spaces are $\mathbb{C} \cdot X_i^2$ and $\mathbb{C} \cdot X_i X_j$. Consider the form $f = (X_1 + X_2 + X_3)^2$. We see that $supp(f)$ is set of all characters of $Sym^2(\mathbb{C}^3)$ and* **does not** *satisfy the combinatorial property. However, under a basis change A:*

$$\begin{array}{lcl} X_1 & \to & X_1 + X_2 + X_3 \\ X_2 & \to & X_2 \\ X_3 & \to & X_3 \end{array}$$

we see that $A \cdot f = X_1^2$. Thus $A \cdot f$ **does** *satisfy the strict half-space property. Indeed consider the λ*

$$\lambda(t) = \begin{bmatrix} t & 0 & 0 \\ 0 & t^{-1} & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

We see that

$$\lim_{t\to 0} t \cdot (A \cdot f) = t^2 X_1^2 = 0$$

Thus we see that every form in the null-cone has a standard form with a very limited sets of possible supports.

Let us look at the module $\mathcal{M}$ of 3×3-matrices under conjugation. Let us fix a λ:

$$\lambda(t) = \begin{bmatrix} t^{n_1} & 0 & 0 \\ 0 & t^{n_2} & 0 \\ 0 & 0 & t^{n_3} \end{bmatrix}$$

such that $n_1 + n_2 + n_3 = 0$. We may assume that $n_1 \geq n_2 \geq n_3$. Looking at the action of $\lambda(t)$ on a general matrix X, we see that:

$$t \cdot X = (t^{n_i - n_j} x_{ij})$$

Thus if $\lim_{t \to 0} t \cdot X$ is to be 0 then $x_{ij} = 0$ for all $i > j$. In other words, X is strictly upper-triangular. Considering the general 1-parameter group tells us that X is in the null-cone iff there is an A such that AXA^{-1} is strictly upper-triangular. In other words, X is **nilpotent**. *The 1-parameter subgroup identifies this transformation and thus the flag of nilpotency.*

26.2 The destabilizing flag

In this section we do a more refined analysis of elements of the null-cone. The basic motivation is to identify a unique set of 1-parameter subgroups which drive a null-point to zero. The simplest example is given by $X_1^2 \in Sym^2(\mathbb{C}^3)$. Let λ, λ' and λ'' be as below:

$$\lambda(t) = \begin{bmatrix} t & 0 & 0 \\ 0 & t^{-1} & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \lambda'(t) = \begin{bmatrix} t & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & t^{-1} \end{bmatrix} \quad \lambda''(t) = \begin{bmatrix} t & 0 & 0 \\ 0 & 0 & -1 \\ 0 & t^{-1} & 0 \end{bmatrix}$$

We see that all the three λ, λ' and λ'' drive X_1^2 to zero. The question is whether these are related, and to classify such 1-parameter subgroups. Alternately, one may view this to a more refined classification of points in the null-cone, such as the stratification of the nilpotent matrices by their Jordan canonical form.

There are two aspects to this analysis. Firstly, to identify a metric by which to choose the 'best' 1-parameter subgroup driving a null-point to zero. Next, to show that there is a unique equivalence class of such 'best' subgroups.

Towards the first objective, let $\lambda : \mathbb{C}^* \to SL_n$ be a 1-parameter subgroup. Without loss of generality, we may assume that $Im(\lambda) \subseteq D$. If w is a null-point then we have:

$$t \cdot w = t^{n_1} w_1 + \ldots + t^{n_k} w_k$$

where $n_i > 0$ for all i. Clearly, a measure of how fast λ drives w to zero is $m(\lambda) = min\{n_1, \ldots, n_k\}$. Verify that this really does not depend on the choice of the maximal torus at all, and thus is well-defined.

Next, we see that for a λ as above, we consider $\lambda^2 : \mathbb{C}^* \to SL_n$ such that $\lambda^2(t) = \lambda(t^2)$. It is easy to see that $m(\lambda^2) = 2 \cdot m(\lambda)$. Clearly, λ and λ^2 are intrinsically identical and we would like to have a measure invariant under such scaling. This comes about by associating a length to each λ. Let λ be as above and let $Im(\lambda) \subseteq D$. Then, there are integers $a_1, \ldots, a_n$ such that

$$\lambda(t) = \begin{bmatrix} t^{a_1} & 0 & 0 & 0 \\ 0 & t^{a_2} & 0 & 0 \\ & & \vdots & \\ 0 & 0 & 0 & t^{a_n} \end{bmatrix}$$

We define $\|\lambda\|$ as

$$\|\lambda\| = \sqrt{a_1^2 + \ldots + a_n^2}$$

We must show that this does not depend on the choice of the maximal torus D. Let $\mathcal{T}(SL_n)$ denote the collection of all maximal tori of SL_n as abstract subgroups. For every $A \in SL_n$, we may define the map $\phi_A : \mathcal{T} \to \mathcal{T}$ defined by $T \to ATA^{-1}$. The stabilizer of a torus T for this action of SL_n is clearly $N(T)$, the normalizer of T. Also recall that $N(T)/T = W$ is the (discrete) weyl group. Let $Im(\lambda) \subseteq D \cap D'$ for some two maximal tori D and D'. Since there is an A such that $AD'A^{-1} = D$, it is clear that $\|\lambda\| = \|A\lambda A^{-1}\|$. Thus, we are left to check if $\|\lambda'\| = \|\lambda\|$ when (i) $Im(\lambda), Im(\lambda') \subseteq D$, and (ii) $\lambda' = A\lambda A^{-1}$ for some $A \in SL_n$. This throws the question to invariance of $\|\lambda\|$ under $N(D)$, or in other words, symmetry under the weyl group. Since $W \cong S_n$, the symmetric group, and since $\sqrt{a_1^2 + \ldots + a_n^2}$ is a symmetric function on $a_1, \ldots, a_n$, we have that $\|\lambda\|$ is well defined.

We now define the *efficiency* of λ on a null-point w to be

$$e(\lambda) = \frac{m(\lambda)}{\|\lambda\|}$$

We immediately see that $e(\lambda) = e(\lambda^2)$.

Lemma 26.1. *Let W be a representation of SL_n and let $w \in W$ be a null-point. Let $\mathcal{N}(w, D)$ be the collection of all $\lambda : \mathbb{C}^* \to D$ such that $\lim_{t \to 0} t \cdot w = 0_W$. If $\mathcal{N}(w, D)$ is non-empty then there is a unique $\lambda' \in \mathcal{N}(w, D)$ which maximizes the efficiency, i.e., $e(\lambda') > e(\lambda)$ for all $\lambda \in \mathcal{N}(w, D)$ and $\lambda \neq (\lambda')^k$ for any $k \in \mathbb{Z}$. This 1-parameter subgroup will be denoted by $\lambda(w, D)$.*

Proof: Suppose that $\mathcal{N}(w, D)$ is non-empty. Then in the weight-space expansion of w for the maximal torus D, we see that $supp(w)$ staisfies the half-space property for some $\lambda \in \mathbb{Y}^n$. Note that the $\lambda \in \mathcal{N}(w, D)$ are parametrized by lattice points $\lambda \in \mathbb{Y}^n$ such that $(\lambda, \chi) > 0$ for all $\chi \in supp(w)$. Let $Cone(w)$ be the conical combination (over $\mathbb{R}$) of all $\chi \in supp(w)$ and $Cone(w)^\circ$ its **polar**. Thus, in other words, $\mathcal{N}(w, D)$ is precisely the collection of lattice points in the cone $Cone(w)^\circ$. Next, we see that $e(\lambda)$ is a convex function of $Cone(w)^\circ$ which is constant over rays $\mathbb{R}^+ \cdot \lambda$ for all $\lambda \in Cone(w)^\circ$. By a routine analysis, the maximum of such a function must be a unique ray with rational entries. This proves the lamma. $\square$

This covers one important part in our task of identifying the 'best' 1-parameter subgroup driving a null-point to zero. The next part is to relate D to other maximal tori.

Let $\lambda : \mathbb{C}^* \to SL_n$ and let $P(\lambda)$ be defined as follows:

$$P(\lambda) = \{A \in SL_n | \lim_{t \to 0} \lambda(t) A \lambda(t^{-1}) = I \in SL_n\}$$

Having fixed a maximal torus D containing $IM(\lambda)$, we easily identify $P(\lambda)$ as a **parabolic** subgroup, i.e., block upper-triangular. Indeed, let

$$\lambda(t) = \begin{bmatrix} t^{a_1} & 0 & 0 & 0 \\ 0 & t^{a_2} & 0 & 0 \\ & & \vdots & \\ 0 & 0 & 0 & t^{a_n} \end{bmatrix}$$

with $a_1 \geq a_2 \geq \ldots \geq a_n$ (obviously with $a_1 + \ldots + a_n = 0$). Then

$$P(\lambda) = \{(x_{ij} | x_{ij} = 0 \text{ for all } i, j \text{ such that } a_i < a_j\}$$

The **unipotent radical** $U(\lambda)$ is a normal subgroup of $P(\lambda)$ defined as:

$$U(\lambda) = (x_{ij}) \text{ where } = \begin{cases} x_{ij} = 0 & \text{if } a_i < a_j \\ x_{ij} = \delta_{ij} & \text{if } a_i = a_j \end{cases}$$

Lemma 26.2. *Let $\lambda \in \mathcal{N}(w, D)$ and let $g \in P(\lambda)$, then (i) $g\lambda g^{-1} \subseteq P(\lambda)$ and $P(g\lambda g^{-1}) = P(\lambda)$, (ii) $g\lambda g^{-1} \in \mathcal{N}(w, gDg^{-1})$, and (iii) $e(\lambda) = e(g\lambda g^{-1})$.*

This actually follows from the construction of the explicit SL_n-modules and is left to the reader. We now come to the unique object that we will define for each $w \in W$ in the null-cone. This is the parabolic subgroup $P(\lambda)$ for any 'best' λ. We have already seen above that if λ' is a $P(\lambda)$-conjugate of a best λ then λ' is 'equally best' and $P(\lambda) = P(\lambda')$.

We now relate two general equally best λ and λ'. For this we need a preliminary definition and a lemma:

Definition 26.1. *Let V be a vector space over $\mathbb{C}$. A* **flag** *$\mathcal{F}$ of V is a sequence $(V_0, \ldots, V_r)$ of nested subspaces $0 = V_0 \subset V_1 \subset \ldots \subset V_r = V$.*

Lemma 26.3. *Let $dim_{\mathbb{C}}(V) = r$ and let $\mathcal{F} = (V_0, \ldots, V_r)$ and $\mathcal{F}' = (V_0', \ldots, V_r')$ be two (complete) flags for V. Then there is a basis $b_1, \ldots, b_r$ of V and a permutation $\sigma \in S_r$ such that $V_i = \overline{\{b_1, \ldots, b_i\}}$ and $V_i' = \overline{\{b_{\sigma(1)}, \ldots, b_{\sigma(i)}\}}$ for all i.*

This is proved by induction on r.

Corollary 26.1. *Let λ and λ' be two 1-parameter subgroups and $P(\lambda)$ and $P(\lambda')$ be their corresponding parabolic subgroups. Then there is a maximal torus T of SL_n such that $T \subseteq P(\lambda) \cap P(\lambda')$.*

Proof: It is clear that there is a correspondence between parabolic subgroups of SL_n and flags. We refine the flags associated to the parabolic subgroups $P(\lambda)$ and $P(\lambda')$ to complete flags and apply the above lemma. $\square$

We are now prepared to prove Kempf'd theorem:

Theorem 12. *Let W be a representation of SL_n and $w \in W$ a null-point. Then there is a 1-parameter subgroup $\lambda \in \Gamma(SL_n)$ such that (i) for all $\lambda' \in \Gamma(SL_n)$, we have $e(\lambda) \geq e(\lambda')$, and (ii) for all λ' such that $e(\lambda) = e(\lambda')$ we have $P(\lambda) = P(\lambda')$ and that there is a $g \in P(\lambda)$ such that $\lambda' = g\lambda g^{-1}$.*

Proof: Let $\mathcal{N}(w)$ be all elements of $\Gamma(SL_n)$ which drive w to zero. Let $\Xi(W)$ be the (finite) collection of D-characters appearing in the representation W. For every $\lambda(w, T)$ such that $Im(\lambda) \subseteq D$, we may consider an $A \in SL_n$ such that $A\lambda A^{-1} \in \mathcal{N}(A \cdot w, D)$ and $e(\lambda) = e(A\lambda A^{-1})$. Since the 'best' element of $\mathcal{N}(A \cdot w, D)$ is determined by $supp(A \cdot w) \subseteq \Xi$, we see that there are only finitely many possibilities for $e(A \cdot w, A\lambda A^{-1})$ and therefore for $e(\lambda)$ for the 'best' λ driving w to zero.

Thus the length k of any sequence $\lambda(w,T_1),\ldots,\lambda(w,T_k)$ such that $e(\lambda(w,T_1)) < \ldots < e(\lambda(w,T_k))$ must be bounded by the number 2^{Ξ}. This proves (i).

Next, let $\lambda_1 = \lambda(w,T_1)$ and $\lambda_2 = \lambda(w,T_2)$ be two 'best' elements of $\mathcal{N}(w,T_1)$ and $\mathcal{N}(w,T_2)$ respectively. By corollary 26.1, we have a torus , say D, and $P(\lambda(w,T_i))$-conjugates λ_i such that (i) $e(\lambda_i) = e(\lambda(w,T_i))$ and (ii) $Im(\lambda_i) \subseteq D$. By lemma 26.1, we have $\lambda_1 = \lambda_2$ and thus $P(\lambda_1) = P(\lambda_2)$. On the other hand, $P(\lambda(w,T_i)) = P(\lambda_i)$ and this proves (ii). $\Box$

Thus 12 associates a unique parabolic subgroup $P(w)$ to every point in the null-cone. This subgroup is called the **destabilizing flag** of w. Clearly, if w is in the null-cone then so is $A \cdot w$, where $A \in SL_n$. Furthermore, it is clear that $P(A\cdot w) = AP(w)A^{-1}$.

Corollary 26.2. *Let $w \in W$ be in the null-cone and let $G_w \subseteq SL_n$ stabilize w. Then $G_w \subseteq P(w)$.*

Proof: Let $g \in G_w$. Since $g \cdot w = w$, we see that $gP(w)g^{-1} = P(w)$, and that g normalizes $P(w)$. Since the normalizer of any parabolic subgroup is itself, we see that $g \in P(w)$. $\Box$

Chapter 27

Stability

Reference: [Ke, GCT1]

Recall that $z \in W$ is stable iff its orbit $O(z)$ is closed in W. In the last chapter, we tackled the points in the null-cone, i.e., points in the set $[0_W]_\approx$, or in other words, points which close onto the stable point 0_W. A similar analysis may be done for arbitrary stable points.

Following kempf, let $S \subseteq W$ be a closed SL_n-invariant subset. Let $z \in W$ be arbitrary. If the orbit-closure $\Delta(z)$ intersects S, then we associate a unique parabolic subgroup $P_{z,S} \subseteq SL_n$ as a witness to this fact. The construction of this parabolic subgroup is in several steps.

As the first step, we construct a representation X of SL_n and a closed SL_n-invariant embedding $\phi : W \to X$ such that $\phi^{-1}(0_X) = S$, scheme-theoretically. This may be done as follows: since S is a closed sub-variety of W, there is an ideal $I_s = (f_1, \ldots, f_k)$ of definition for S. We may further assume that the vector space $\overline{\{f_1, \ldots, f_k\}}$ is itself an SL_n-module, say X. We assume that X is k-dimensional.

We now construct the map $\phi : W \to X$ as follows:

$$\phi(w) = (f_1(x), \ldots, f_k(x))$$

Note that $\phi(S) = 0_X$ and that $I_S = (f_1, \ldots, f_k)$ ensure that the requirements on our ϕ do hold.

Next, there is an adaptation of (Hilbert's) Theorem 11 which we do not prove:

Theorem 13. *Let W be an SL_n-module and let $y \in W$ be a stable point. Let $z \in [y]_\approx$ be an element which closes onto y. Then there is a 1-parameter subgroup $\lambda : \mathbb{C}^* \to SL_n$ such that*

$$\lim_{t \to 0} \lambda(t) \cdot w \in O(y)$$

Thus the limit exists and lies in the closed orbit of y.

Now suppose that $\Delta(z) \cap S$ is non-empty. Then there must be stable $y \in \Delta(z)$. We apply the theorem to $O(y)$ and obtain the λ as above. This shows that there is indeed a 1-parameter subgroup driving z into S. Next, it is easy to see that

$$\lim_{t \to 0}[\lambda(t) \cdot \phi(z)] = 0_X$$

Thus $\phi(z)$ actually lies in the null-cone of X. We may now be tempted to apply the techniques of the previous chapter to come up with the 'best' λ and its parabolic, now called $P(z, S)$. This is almost the technique to be adopted , except that this 'best' λ drives $\phi(z)$ into 0_X but $\lim_{t \to 0}[\lambda(t) \cdot z]$ (which is supposed to be in S) may not exist! This is because we are using the unproved (and untrue) converse of the assertion that 1-parameter subgroups which drive z into S drive $\phi(z)$ into 0_X.

This above argument is rectified by limiting the domain of allowed 1-parameter subgroups to (i) $Cone(supp(\phi(z))^\circ$ as before, and (ii) those λ such that $\lim_{t \to 0}[\lambda(t) \cdot z]$ exists. This second condition is also a 'convex' condition and then the 'best' λ does exist. This completes the construction of $P(z, S)$.

As before, if $G_z \subseteq SL_n$ stabilizes z then it normalizes $P(z, S)$ thus must be contained in it:

Proposition 27.1. *If G_z stabilizes z then $G_z \subseteq P(z, S)$.*

Let us now consider the **permanent** and the **determinant**. Let $\mathcal{M}$ be the n^2-dimensional space of all $n \times n$-matrices. Since *det* and *perm* are homogeneous n-forms on $\mathcal{M}$, we consider the $SL(\mathcal{M})$-module $W = Sym^n(\mathcal{M}^*)$. We recall now certain stabilizing groups of the *det* and the *perm*. We will need the definition of a certain group L'. This is defined as the group generated by the permutation and diagonal matrices in GL_n. In other words, L' is the normalizer of the complete standard torus $D^* \subseteq GL_n$. L is defined as that subgroup of L' which is contained in SL_n.

Proposition 27.2. *(A) Consider the group $K = SL_n \times SL_n$. We define the action μ_K of typical element $(A, B) \in K$ on $X \in \mathcal{M}$ as given by:*

$$X \to AXB^{-1}$$

Then (i) $\mathcal{M}$ is an irreducible representation of K and $Im(K) \subseteq SL(\mathcal{M})$, and (ii) K stabilizes the determinant.

(B) *Consider the group $H = L \times L$. We define the action μ_H of typical element $(A, B) \in H$ on $X \in \mathcal{M}$ as given by:*

$$X \to AXB^{-1}$$

Then (i) $\mathcal{M}$ is an irreducible representation of H and $Im(H) \subseteq SL(\mathcal{M})$, and (ii) H stabilizes the permanent.

We are now ready to show:

Theorem 14. *The points det and perm in the $SL(\mathcal{M}$-module $W = Sym^n(\mathcal{M}^*)$ are stable.*

Proof: Lets look at *det*, the *perm* being similar. If *det* were not stable, then there would be a closed $SL(\mathcal{M})$-invariant subset $S \subset W$ such that $det \notin S$ but closes onto S: just take S to be the unique closed orbit in $[det]_{\approx}$. Whence there is a parabolic $P(det, S)$ which, by Proposition 27.1, would contain K. This would mean that there is a K-invariant flag in $\mathcal{M}$ corresponding to $P(det, S)$. This contradicts the irreducibility of $\mathcal{M}$ as a K-module. $\square$

Bibliography

[BBD] A. Beilinson, J. Bernstein, P. Deligne, Faisceaux pervers, Astérisque 100, (1982), Soc. Math. France.

[B] P. Belkale, Geometric proofs of Horn and saturation conjectures, math.AG/0208107.

[BZ] A. Berenstein, A. Zelevinsky, Tensor product multiplicities and convex polytopes in partition space, J. Geom. Phys. 5(3): 453-472, 1988.

[DJM] M. Date, M. Jimbo, T. Miwa, Representations of $U_q(\hat{g}l(n, \mathbb{C}))$ at $q = 0$ and the Robinson-Schensted correspondence, in Physics and Mathematics of Strings, World Scientific, Singapore, 1990, pp. 185-211.

[DM1] J. De Loera, T. McAllister, Vertices of Gelfand-Tsetlin polytopes, Discrete Comput. Geom. 32 (2004), no. 4, 459470.

[DM2] J. De Loera, T. McAllister, On the computation of Clebsch-Gordon coefficients and the dilation effect, Experiment Math 15, (2006), no. 1, 7-20

[Dl2] P. Deligne, La conjecture de Weil II, Publ. Math. Inst. Haut. Étud. Sci. 52, (1980) 137-252.

[DeM] P. Deligne and J. Milne, Tannakien categories. In *Lecture Notes in Mathematics*, 900. Springer-Verlag: New York, 1982.

[Der] H. Derkesen, J Weyman, On the Littlewood-Richardson polynomials, J. Algebra 255(2002), no. 2, 247-257.

[F] W. Fulton, Young Tableaux: With Applications to Representation Theory and Geometry. *Cambridge University Press*, 1997.

[FH] W. Fulton and J. Harris, Representation Theory: A First Course. *Springer-Verlang*, 1991.

[GCTabs] K. Mulmuley, Geometric complexity theory: abstract, technical report TR-2007-12, Computer science department, The University of Chicago, September, 2007. available at http://ramakrishnadas.cs.uchicago.edu.

[GCTflip1] K. Mulmuley, On P. vs. NP, geometric complexity theory, and the flip I: a high-level view, Technical Report TR-2007-13, Computer Science Department, The University of Chicago, September 2007. Available at: http://ramakrishnadas.cs.uchicago.edu

[GCTflip2] K. Mulmuley, On P vs. NP, geometric complexity theory, and the flip II, under preparation.

[GCTconf] K. Mulmuley, M. Sohoni, Geometric complexity theory, P vs. NP and explicit obstructions, in "Advances in Algebra and Geometry", Edited by C. Musili, the proceedings of the International Conference on Algebra and Geometry, Hyderabad, 2001.

[GCT1] K. Mulmuley, M. Sohoni, Geometric complexity theory I: an approach to the P vs. NP and related problems, SIAM J. Comput., vol 31, no 2, pp 496-526, 2001.

[GCT2] K. Mulmuley, M. Sohoni, Geometric complexity theory II: towards explicit obstructions for embeddings among class varieties, to appear in SIAM J. Comput., cs. ArXiv preprint cs. CC/0612134, December 25, 2006.

[GCT3] K. Mulmuley, M. Sohoni, Geometric complexity theory III, on deciding positivity of Littlewood-Richardson coefficients, cs. ArXiv preprint cs. CC/0501076 v1 26 Jan 2005.

[GCT4] K. Mulmuley, M. Sohoni, Geometric complexity theory IV: quantum group for the Kronecker problem, cs. ArXiv preprint cs. CC/0703110, March, 2007.

[GCT5] K. Mulmuley, H. Narayanan, Geometric complexity theory V: on deciding nonvanishing of a generalized Littlewood-Richardson coefficient, Technical report TR-2007-05, Comp. Sci. Dept. The university of chicago, May, 2007.

[GCT6] K. Mulmuley, Geometric complexity theory VI: the flip via saturated and positive integer programming in representation theory and algebraic geometry, Technical report TR 2007-04, Comp.

Sci. Dept., The University of Chicago, May, 2007. Available at: http://ramakrishnadas.cs.uchicago.edu. Revised version to be available here.

[GCT7] K. Mulmuley, Geometric complexity theory VII: nonstandard quantum group for the plethysm problem (Extended Abstract), Technical report TR-2007-14, Comp. Sci. Dept., The University of Chicago, Sept. 2007. Available at: http://ramakrishnadas.cs.uchicago.edu.

[GCT8] K. Mulmuley, Geometric complexity theory VIII: On canonical bases for the nonstandard quantum groups (Extended Abstract), Technical report TR-2007-15, Comp. Sci. Dept., The University of Chicago, Sept. 2007. Available at: http://ramakrishnadas.cs.uchicago.edu.

[GCT9] B. Adsul, M. Sohoni, K. Subrahmanyam, Geometric complexity theory IX: algbraic and combinatorial aspects of the Kronecker problem, under preparation.

[GCT10] K. Mulmuley, Geometric complexity theory X: On class varieties, and the natural proof barrier, under preparation.

[GCT11] K. Mulmuley, Geometric complexity theory XI: on the flip over finite or algebraically closed fields of positive characteristic, under preparation.

[GLS] M. Grötschel, L. Lovász, A. Schrijver, Geometric algorithms and combinatorial optimzation, Springer-Verlag, 1993.

[H] H. Narayanan, On the complexity of computing Kostka numbers and Littlewood-Richardson coefficients Journal of Algebraic Combinatorics, Volume 24 , Issue 3 (November 2006) 347 - 354, 2006

[KB79] R. Kannan and A. Bachem. Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix, *SIAM J. Comput.*, 8(4), 1979.

[Kar84] N. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4(4):373–395, 1984.

[KL] D. Kazhdan, G. Lusztig, Representations of Coxeter groups and Hecke algebras, Invent. Math. 53 (1979), 165-184.

[KL2] D. Kazhdan, G. Lusztig, Schubert varieties and Poincare duality, Proc. Symp. Pure Math., AMS, 36 (1980), 185-203.

[Kha79] L. G. Khachian. A polynomial algorithm for linear programming. *Doklady Akedamii Nauk SSSR*, 244:1093–1096, 1979. In Russian.

[K] M. Kashiwara, On crystal bases of the q-analogue of universal enveloping algebras, Duke Math. J. 63 (1991), 465-516.

[Ke] G. Kempf: Instability in invariant theory, Annals of Mathematics, 108 (1978), 299-316.

[KTT] R. King, C. Tollu, F. Toumazet, Tretched Littlewood-Richardson coefficients and Kostak coefficients. In, Winternitz, P. Harnard, J. Lam, C.S. and Patera, J. (eds.) Symmetry in Physics: In Memory of Robert T. Sharp. Providence, USA, AMS OUP, 99-112, CRM Proceedings and Lecture Notes 34, 2004.

[Ki] A. Kirillov, An invitation to the generalized saturation conjecture, math. CO/0404353, 20 Apr. 2004.

[KS] A. Klimyck, and K. Schmüdgen, Quantum groups and their representations, Springer, 1997.

[KT] A. Knutson, T. Tao, The honeycomb model of $GL_n(C)$ tensor products I: proof of the saturation conjecture, J. Amer. Math. Soc. 12 (1999) 1055-1090.

[KT2] A. Knutson, T. Tao: Honeycombs and sums of Hermitian matrices, Notices Amer. Math. Soc. 48 (2001) No. 2, 175-186.

[LV] D. Luna and T. Vust, Plongements d'espaces homogenes, Comment. Math. Helv. 58, 186(1983).

[Lu2] G. Lusztig, Introduction to quantum groups, Birkhäuser, 1993.

[Ml] K. Mulmuley, Lower bounds in a parallel model without bit operations. *SIAM J. Comput.* 28, 1460–1509, 1999.

[Mm] D. Mumford, Algebraic Geometry I, Springer-Verlang, 1995.

[N] M. Nagata, Polynomial Rings and Affine Spaces. CBMS Regional Conference no. 37, American Mathematical Society, 1978.

[S] R. Stanley, Enumerative combinatorics, vol. 1, Wadsworth and Brooks/Cole, Advanced Books and Software, 1986.

[Z] A. Zelevinsky, Littlewood-Richardson semigroups, arXiv:math.CO/9704228 v1 30 Apr 1997.