Equivariant Coherent Sheaves on the Nilpotent Cone for Complex Reductive Lie Groups

by

Pramod Narahari Achar

S.B., Massachusetts Institute of Technology, 1997

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

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2001

Author . . . .
Department of Mathematics
May 2, 2001

Certified by . . . .
David A. Vogan, Jr.
Professor of Mathematics
Thesis Supervisor

Accepted by . . . .
Tomasz Mrowka
Chairman, Department Committee on Graduate Students

Equivariant Coherent Sheaves on the Nilpotent Cone for Complex Reductive Lie Groups

by

Pramod Narahari Achar

Submitted to the Department of Mathematics
on May 2, 2001, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Let G be a connected complex reductive Lie group. We propose a certain bijection between the set of dominant integral weights of G, and the set of pairs consisting of a nilpotent coadjoint orbit and a finite-dimensional irreducible representation of the isotropy group of the orbit. A constructive proof of this bijection is given for the groups $GL(n, \mathbb{C})$, and the bijection is established by direct calculation in a handful of particular groups. Partial progress is made on a general proof for $Sp(2n, \mathbb{C})$.

Thesis Supervisor: David A. Vogan, Jr.
Title: Professor of Mathematics

Acknowledgments

I would like to thank my elder siblings in representation theory, Dana Pascovici and Thom Pietraho, as well as my coeval cousins, Anthony Henderson and Kevin McGerty, for indulging me on numerous occasions and answering questions I was too embarrassed to ask David.

I owe much gratitude to Eric Sommers for taking an interest in my work at a time when I was still too young to appreciate his.

I would like to thank my mentor, Robert Militello, for inspiring me to go into mathematics, and for impressing upon me the value of good teaching.

Above all, I would like to thank my advisor, David Vogan, for having patiently supplied me with knowledge, guidance, and suggestions over the past four years, as well as for overlooking my assertion that G_2 has seven positive roots. Without him, the present work would not have been possible.

Contents

List of Tables

CHAPTER 1

Introduction

Let G be a connected complex reductive Lie group, and $\mathfrak{g}$ its Lie algebra, with universal enveloping algebra $U(\mathfrak{g})$. Let K be a compact real form of G, with Lie algebra $\mathfrak{k}$. Let T be a maximal (complex) torus in G, and let $\Lambda(G)$ (resp. $\Lambda_+(G)$) denote the weight lattice (resp. set of dominant weights) of T. Let $\mathcal{N}^*$ denote the nilpotent cone in $\mathfrak{g}^*$, and let $\mathcal{M}(G)$ denote the category of Harish-Chandra modules of G. Finally, let $\mathcal{C}(G)$ denote the category of finitely-generated $(S(\mathfrak{g}_{\mathbb{C}}/\mathfrak{k}_{\mathbb{C}}), K)$-modules supported on $\mathcal{N}^*$, where $S(V)$ denotes the symmetric algebra on a vector space V.

If we pick a finite-dimensional K-invariant generating subspace for a given Harish-Chandra module in $\mathcal{M}(G)$, we can form its associated graded module, an $(S(\mathfrak{g}), K)$-module. In fact, it turns out that this module is supported on $\mathcal{N}^*$, and that $\mathfrak{k}$ acts on it by 0, so we actually get an object in $\mathcal{C}(G)$. This operation is not functorial, since it depends on the choice of a finite-dimensional generating subspace. Nevertheless, it gives rise to a well-defined homomorphism

$$[\mathrm{gr}] : K\mathcal{M}(G) \to K\mathcal{C}(G)$$

between the Grothendieck groups of these two categories [13].

Studying $K\mathcal{C}(G)$ can tell us things about objects in $\mathcal{M}(G)$. The following theorem gives us one description of $K\mathcal{C}(G)$.

Theorem 1 ([13]). *Let $\mathcal{O}_1, \ldots, \mathcal{O}_m$ be the nilpotent orbits in $\mathcal{N}^*$; let $e_1, \ldots, e_m$ be representatives thereof, and let G^{e_i} be the isotropy group for each e_i. Let (τ, V_τ) be an irreducible algebraic representation of G^{e_i}. There is a module $N(e_i, \tau)$ in $\mathcal{C}(G)$, whose G-module structure is given by $\mathrm{Ind}_{G^{e_i}}^G \tau$, and which meets the following conditions.*

(a) The associated variety of $N(e_i, \tau)$ is $\overline{\mathcal{O}_i}$.

(b) The collection $\{N(e_i, \tau)\}$, as i ranges over all orbits and τ over all irreducible representations of the appropriate isotropy group, constitutes a $\mathbb{Z}$-basis for $K\mathcal{C}(G)$.

Another description of the holomorphic G-module structure of modules in $\mathcal{C}(G)$ is furnished by the following theorem.

Theorem 2 ([14]). *Let M be an object in $\mathcal{C}(G)$. For $\sigma \in \Lambda(G)$, let $\mathbb{C}_\sigma$ be the one-dimensional T-representation of weight σ. Then, in $K\mathcal{C}(G)$, there is a unique expression for M of the form*

$$M = \sum_{\sigma \in \Lambda_+(G)} m_\sigma(M) \,\mathrm{Ind}_T^G \mathbb{C}_\sigma, \tag{1.1}$$

where only finitely many of the $m_\sigma(M)$'s are nonzero.

This theorem essentially says that the $\mathrm{Ind}_T^G \mathbb{C}_\sigma$, as σ ranges over dominant weights, constitute a basis for $K\mathcal{C}(G)$. (In Chapter 2, we will see explicitly how to produce a sum of objects of $\mathcal{C}(G)$ such that their holomorphic G-module structure is given by $\mathrm{Ind}_T^G \mathbb{C}_\sigma$.)

What do these two theorems have to do with each other? Theorem 1 describes a basis with respect to which associated varieties are readily accessible—indeed, the asssociated varieties of the basis elements are essentially given as part of the notation by which the basis elements are indexed.

Theorem 2, on the other hand, seems to present to us, in its language of weights, something more akin to infinitesimal characters.

Now, unitarity results for reductive Lie groups often come about from a comparison of infinitesimal characters and associated varieties, so it would be nice to understand the relationship between the bases arising from these two theorems. In particular, an explicit description of this relationship would make feasible certain computations to establish the unitarity of various representations introduced by Arthur, leading to the verification of the Arthur conjectures [2], [3].

The object of the present work is to study this relationship. Combining the two theorems gives us equations of the form

$$N(e_i, \tau) = \sum_{\sigma \in \Lambda_+(G)} m_\sigma(e_i, \tau) \operatorname{Ind}_T^G \mathbb{C}_\sigma. \tag{1.2}$$

Let us define a map between the indexing sets for the two bases:

$$\gamma : \{(e_i, \tau)\} \to \Lambda_+(G),$$

where e_i ranges over G-conjugacy classes of nilpotent elements, and for each e_i, τ ranges over irreducible algebraic representations of G^{e_i}, is defined by

$$\gamma(e_i, \tau) = \text{the largest } \sigma \text{ such that } m_\sigma(e_i, \tau) \neq 0.$$

Of course, it is not clear that this definition even makes sense; indeed, we shall revise it in Chapter 2 before attempting to do anything rigorous with it. Nevertheless, we can say something about the intuition by which we hope to be guided. The hope is that there is an "upper-triangular" relationship between the two bases under consideration, and that, in particular, picking off the largest term on the right-hand side of (1.2) actually gives a bijection.

Related bijections between $\Lambda_+(G)$ and $\{(e_i, \tau)\}$ have been considered by Bezrukavnikov [4], [5] and by Ostrik [12]. The existence of such bijections was originally conjectured by Lusztig [10].

In Chapter 2, we develop the necessary background of facts from algebraic geometry that are needed in order to manipulate sheaves on the nilpotent cone. In the course of developing this background, we shall see what kinds of tools and tricks are needed in trying to compute expressions of the form (1.2). An understanding of these tools and tricks will enable us to give a precise definition of γ in Section 2.1. By the end of the chapter, we formulate some guesses on how to pick off the largest term of an expression of the form (1.1) for a given pair (e_i, τ), and we actually compute γ in some extreme cases.

Chapters 3–5 are all set in the restricted context of $GL(n, \mathbb{C})$. In Chapter 3, we introduce a new formalism for manipulating and studying the relationships between weights of G, representations of isotropy groups of orbits, and the representation theory of certain Levi subgroups that serves as the intermediary between these two. We establish the basic properties of the artifacts of this formalism, called *weight diagrams*, and we introduce a number of operations that can be performed on weight diagrams while leaving certain properties invariant.

In Chapter 4, we put the weight diagrams to work in two different algorithms. One of these, called "α," combines the operations on weight diagrams from Chapter 3 to transform any weight diagram whatsoever into one that meets certain very stringent requirements. It turns out, in fact, that those requirements are so stringent that a weight diagram meeting them can be reconstructed from very little information. Such a reconstruction is accomplished by the second algorithm in the chapter, called, sensibly enough, "β."

Chapter 5 is devoted to establishing the main result of the present work:

THEOREM 3. *For the group $G = GL(n, \mathbb{C})$, the map γ is a bijection. This map and its inverse may be explicitly computed by algorithms.*

Approximately speaking, α computes γ, and β computes its inverse.

In Chapter 6, we attempt to duplicate the work of Chapters 3–5 in the context of $Sp(2n, \mathbb{C})$. We do not yet obtain an analogue of Theorem 3, nor do we have any reason to believe that such a result should not hold. Indeed, we introduce a weight diagram formalism appropriate to the $Sp(2n, \mathbb{C})$

context, and we establish some of the basic properties thereof. Under some assumptions (*viz.* the guesses at the end of Chapter 2, and some restrictions on what orbits we look at), we are able to twist the algorithm for α from Chapter 4 into something that makes sense in the $Sp(2n, \mathbb{C})$ context, and we show how to use it to compute γ under certain circumstances.

Finally, the appendices show various examples worked out. Appendix A gives examples of the two algorithms of Chapter 4, and Appendix B contains tables of explicit calculations of γ for an assortment of groups.

CHAPTER 2

Equivariant Coherent Sheaves on the Nilpotent Cone

In this chapter, we develop some general ideas using the algebraic geometry of the nilpotent cone. Given an irreducible representation (τ, V_τ) of G^e, let $\mathcal{V}_\tau$ denote the vector bundle $G \times_{G^e} V_\tau$ over $G/G^e \simeq \mathcal{O}_e$. In the first section, we investigate how to relate $N(e,\tau) = \Gamma(\mathcal{O}_e, \mathcal{V}_\tau)$ to the cohomology of some G-equivariant coherent sheaf supported on $\overline{\mathcal{O}_e}$. In the second, we tackle the computation of those cohomology groups using other techniques. We obtain formulas explicitly enough that in the last two sections, we are able to carry out some calculations, as well as conjecture how those calculations might work in a more general setting.

We shall have occasion to refer to the structure sheaves of various varieties, but the standard notation for structure sheaves poses a problem, since we use the letter $\mathcal{O}$ for orbits. Instead, we shall use the calligraphic letter $\mathcal{S}$ for structure sheaves.

For the most part, vector bundles and representations will be considered to be virtual; that is, we work with equivalence classes in the Grothendieck group of the appropriate category, rather than the objects themselves. As a consequence, even when we have a representation (of, say, a parabolic subgroup) that is not completely reducible, we are at liberty to write it as a sum of irreducible representations. Occasionally, we may without comment regard a representation of a Levi factor of some larger group as a representation of the whole group by letting its unipotent radical act trivially.

2.1 Relating Orbits to their Closures

In this section, we carry out some calculations that more or less follow [11]. Let P_e be a parabolic subgroup of G containing G^e, and suppose that it has Levi decomposition $L_e U_e$. Recall that G^e also has a decomposition as a semidirect product of a reductive group and unipotent group; let G^e_{red} denote its reductive part. Finally, suppose that $\mathfrak{v}_e$ is an $\mathrm{Ad}(P_e)$-invariant subspace of $\mathfrak{g}$ with the properties that

$$\mathrm{Ad}(G)(\mathfrak{v}_e) = \overline{\mathcal{O}_e}$$

and

$$\mathcal{O}_e \cap \mathfrak{v}_e \text{ is a single } P_e\text{-orbit.}$$

In particular, the first of these assumptions implies that $\mathcal{O}_e \cap \mathfrak{v}_e$ is dense in $\mathfrak{v}_e$. Next, let $\mu : G \times_{P_e} \mathfrak{v}_e \to \overline{\mathcal{O}_e}$ be the moment map, given by $\mu(g, v) = \mathrm{Ad}(g)(v)$. It follow from the assumptions that $G^e \subseteq P^e$ and that $\mathcal{O}_e \cap \mathfrak{v}_e$ is a single P_e-orbit that μ is birational and one-to-one over the open orbit $\mathcal{O}_e$.

(The supposition of the existence of such a $\mathfrak{v}_e$ is not vacuous. For example, given e, we can always produce a Jacobson-Morozov $\mathfrak{sl}(2)$-triple $\{h, e, f\}$, and form the h-eigenspace decomposition of $\mathfrak{g}$:

$$\mathfrak{g} = \bigoplus_{i \in \mathbb{Z}} \mathfrak{g}_i.$$

Then we can take P_e to be the connected subgroup of G with Lie algebra $\bigoplus_{i \geq 0} \mathfrak{g}_i$. A result of Kostant [9] implies that $\mathfrak{v}_e = \bigoplus_{i \geq 2} \mathfrak{g}_i$ has the properties listed above.)

Our strategy is encoded in the commutative diagram below. We begin by trying to extend V_τ to a virtual representation (η, W_η) of P_e. This gives rise to a G-equivariant vector bundle $\mathcal{W}_\eta = G \times_{P_e} W_\eta$ over G/P_e. We pull this back to a vector bundle over $G \times_{P_e} \mathfrak{v}_e$, and then push forward its sheaf of sections to get a G-equivariant sheaf $\mu_* \pi^* \mathcal{W}_\eta$ over $\overline{\mathcal{O}_e}$. We establish the relationship between $\Gamma(\mathcal{O}_e, \mathcal{V}_\tau)$ and this sheaf over the orbit closure, and then we (almost) explicitly compute the space of global sections of the latter using Lie algebra cohomology.

$$\begin{tikzcd}
 & & \pi^*\mathcal{W}_\eta \arrow[d] & \mathcal{W}_\eta \arrow[d] \\
 & \mu^{-1}(\mathcal{O}_e) \subseteq \arrow[d] & G \times_{P_e} \mathfrak{v}_e \arrow[r, "\pi"] \arrow[d, "\mu"] & G/P_e \\
\mathcal{V}_\tau \arrow[r] & \mathcal{O}_e \quad \subseteq & \overline{\mathcal{O}_e} & \mu_*\pi^*\mathcal{W}_\eta \arrow[l]
\end{tikzcd}$$

Lemma 2.1.1. *If $\eta|_{G^e} \simeq \tau$, then $\mu_* \pi^* \mathcal{W}_\eta|_{\mathcal{O}_e}$ is isomorphic to the sheaf of sections of $\mathcal{V}_\tau$.*

Proof. For any open set $U \subseteq \mathcal{O}_e$, we need to check that $\Gamma(U, \mathcal{V}_\tau) \simeq \Gamma(U, \mu_* \pi^* \mathcal{W}_\eta)$. The latter space is the same as $\Gamma(\mu^{-1}(U), \pi^* \mathcal{W}_\eta)$. Since $\mathcal{W}_\eta = G \times_{P_e} W_\eta$, it is easy to check that

$$\pi^* \mathcal{W}_\eta \simeq G \times_{P_e} (\mathfrak{v}_e \times W_\eta).$$

Now, $\mu^{-1}(U)$ is some open subset of $G \times_{P_e} \mathfrak{v}_e$; it consists of equivalence classes of pairs (g, v) with $\mathrm{Ad}(g)(v) \in U$. Recall that the moment map μ is one-to-one over $\mathcal{O}_e$. So for any given $u \in U$, if we pick some $g \in G$ such that $\mathrm{Ad}(g)(e) = u$, then one set of representatives for the element $\mu^{-1}(u)$ is $\{(gh, e) \mid h \in G^e\}$. A section of $\pi^* \mathcal{W}_\eta$ over $\mu^{-1}(U)$ can be thought of as a map assigning to each (gG^e, e) an equivalence class of triples $\{(gh, h^{-1}e, h^{-1}w) \mid h \in G^e)\}$, where $w \in W_\eta$. Written this way, the middle component is superfluous. A section of $\pi^* \mathcal{W}_\eta$ can be interpreted as a map assigning to each coset $gG^e \in U$ an equivalence class of pairs $\{(gh, h^{-1}w) \mid h \in G^e\}$.

We have just described precisely what it means to have a section of $G \times_{G^e} W_\eta$ over U. We conclude that if $\eta|_{G^e} \simeq \tau$, then $\Gamma(U, \mathcal{V}_\tau) \simeq \Gamma(\mu^{-1}(U), \mu_* \pi^* \mathcal{W}_\eta)$. $\square$

The sheaf $\mu_* \pi^* \mathcal{W}_\eta$ will serve as an intermediary between $N(e, \tau)$ and the orbit closure $\overline{\mathcal{O}_e}$. We just saw that $\Gamma(\mathcal{O}_e, \mu_* \pi^* \mathcal{W}_\eta) \simeq \Gamma(\mathcal{O}_e, \mathcal{V}_\tau)$; the following proposition helps us relate these spaces to $\Gamma(\overline{\mathcal{O}_e}, \mu_* \pi^* \mathcal{W}_\eta)$.

The spaces of sections we examine in this proposition, and henceforth in the chapter, are objects in the category $\mathcal{C}(G)$, but we seek an expression of the form in (1.1), where the $S(\mathfrak{g}_\mathbb{C}/\mathfrak{k}_\mathbb{C})$-module structure is not shown. We therefore stick to considering these objects only as K-modules, or equivalently as algebraic G-modules. We always have complete reducibility for such modules, and we freely make use of this fact in what follows.

Proposition 2.1.2. *Let X be a closed union of nilpotent orbits, and let $\mathcal{O}_e \subseteq X$ be an orbit of maximal dimension. Let $Z = X \smallsetminus \mathcal{O}_e$. If $\mathcal{F}$ is a G-equivariant coherent sheaf on X, then there is an equivalence of algebraic G-modules*

$$\Gamma(X, \mathcal{F}) \simeq \Gamma(\mathcal{O}_e, \mathcal{F}) + \sum_{\mathcal{O}_{e'} \subseteq Z} (\Gamma(\mathcal{O}_{e'}, \mathcal{G}_{e'}) - \Gamma(\mathcal{O}_{e'}, \mathcal{H}_{e'})),$$

where each $\mathcal{G}_{e'}$ and $\mathcal{H}_{e'}$ is a G-equivariant coherent sheaf on $\mathcal{O}_{e'}$.

Proof. We prove this by induction on the dimension of $\mathcal{O}_e$. If $\mathcal{O}_e = \mathcal{O}_0$ is the zero orbit, the result is obvious: in this case, $\overline{\mathcal{O}_0} = \mathcal{O}_0$, so we just have

$$\Gamma(\overline{\mathcal{O}_0}, \mathcal{F}) = \Gamma(\mathcal{O}_0, \mathcal{F}).$$

For higher-dimensional orbits, let $j : \mathcal{O}_e \hookrightarrow X$ and $i : Z \hookrightarrow X$ be the inclusions. We proceed by induction on the number of orbits of maximal dimension in X. The argument is almost the same in both the base case and the general case; below, we include a remark at the one point where they differ.

Let $\Gamma_Z(X, \mathcal{F})$ denote the space of global sections of $\mathcal{F}$ whose support is contained in Z. $\Gamma_Z(X, \cdot)$ is a left-exact functor; we let $H^i_Z(X, \mathcal{F})$ denote its right-derived functors, the cohomology groups of X with supports in Z and coefficients in $\mathcal{F}$. There is a long exact sequence of cohomology groups

$$0 \to H^0_Z(X, \mathcal{H}) \to H^0(X, \mathcal{F}) \to H^0(\mathcal{O}_e, \mathcal{F}|_{\mathcal{O}_e})) \to H^1_Z(X, \mathcal{F}) \to H^1(X, \mathcal{F}) \to \cdots .$$

Since X is an affine variety, its higher cohomology groups with coefficients in a coherent sheaf vanish. Thus $H^1(X, \mathcal{F}) = 0$. The above long exact sequence yields, therefore, the four-term exact sequence

$$0 \to \Gamma_Z(X, \mathcal{F}) \to \Gamma(X, \mathcal{F}) \to \Gamma(\mathcal{O}_e, \mathcal{F}|_{\mathcal{O}_e}) \to H^1_Z(X, \mathcal{H}) \to 0.$$

Then, if we take the alternating sum of the sequence in the Grothendieck group, we obtain

$$\Gamma(X, \mathcal{F}) \simeq \Gamma(\mathcal{O}_e, \mathcal{F}|_{\mathcal{O}_e}) + \Gamma_Z(X, \mathcal{F}) - H^1_Z(X, \mathcal{F}). \tag{2.1}$$

We need to understand the last two terms of this formula. These are both finitely-generated modules for the coordinate ring $\Gamma(X, \mathcal{S}_X)$ of X (see [8], Ex. II.5.6 and Ex. III.3.3), so they correspond to coherent equivariant sheaves on X. Let us call these sheaves $\mathcal{H}^0$ and $\mathcal{H}^1$, respectively. Each $\mathcal{H}^q$ ($q = 0, 1$) is supported on Z, so perhaps we should consider the sheaf $i^*\mathcal{H}$. This is a coherent sheaf on Z; moreover, it is G-equivariant because the inclusion i is a G-equivariant map. Now, the inductive hypothesis is applicable to Z, because either X contains one orbit of maximal dimension and Z contains only orbits of strictly lower dimension, or X contains multiple orbits of maximal dimension, and Z contains fewer. Let $\mathcal{O}_{e_1}$ be an orbit of maximal dimension in Z. Then, using the inductive hypothesis, we write

$$\begin{aligned} \Gamma(Z, i^*\mathcal{H}^q) &\simeq \Gamma(\mathcal{O}_{e_1}, i^*\mathcal{H}^q|_{\mathcal{O}_{e_1}}) + \sum_{\mathcal{O}_{e''} \subseteq Z \smallsetminus \mathcal{O}_{e_1}} (\Gamma(\mathcal{O}_{e''}, \mathcal{G}'_{e''}) - \Gamma(\mathcal{O}_{e''}, \mathcal{H}'_{e''})) \\ &\simeq \sum_{\mathcal{O}_{e'} \subseteq Z} (\Gamma(\mathcal{O}_{e'}, \mathcal{G}'_{e'}) - \Gamma(\mathcal{O}_{e'}, \mathcal{H}'_{e'})), \end{aligned} \tag{2.2}$$

where we have taken $\mathcal{G}'_{e_1} = i^*\mathcal{H}^q$ and $\mathcal{H}'_{e_1} = 0$. We would like to substitute this expression into (2.1), but we cannot do so until we understand the relationship between $\Gamma(X, \mathcal{H}^q)$ and $\Gamma(Z, i^*\mathcal{H}^q)$. Now, the latter space is given by

$$\Gamma(X, \mathcal{H}^q) \otimes_{\Gamma(X, \mathcal{S}_X)} \Gamma(Z, \mathcal{S}_Z).$$

But Z and X are both affine varieties, and we have the explicit description of $\Gamma(Z, \mathcal{S}_Z)$ as the quotient of $\Gamma(X, \mathcal{S}_X)$ by the ideal of regular functions that vanish on Z. Since $\Gamma(X, \mathcal{H}^q)$ has support contained in Z, it follows that $\Gamma(Z, \mathcal{S}_Z)$ is the quotient of $\Gamma(X, \mathcal{S}_X)$ by an ideal contained in the annihilator of $\Gamma(X, \mathcal{H}^q)$. Thus, the above expression is simply formalism for regarding $\Gamma(X, \mathcal{H}^q)$ as a $\Gamma(Z, \mathcal{S}_Z)$-module; but as modules for $\Gamma(X, \mathcal{S}_X)$, or as algebraic G-modules, we have equivalences

$$\begin{aligned} \Gamma_Z(X, \mathcal{F}) &= \Gamma(X, \mathcal{H}^0) = \Gamma(Z, i^*\mathcal{H}^0) \\ H^1_Z(X, \mathcal{F}) &= \Gamma(X, \mathcal{H}^1) = \Gamma(Z, i^*\mathcal{H}^1) \end{aligned}$$

Now we can substitute the expressions (2.2) into (2.1), and the result follows. □

To convert from the preceding proposition back to the language of vector bundles, we shall need the following fact (see [8], Ex. II.5.8):

LEMMA 2.1.3. *If $\mathcal{F}$ is a G-equivariant coherent sheaf on $\mathcal{O}_e$, then $\mathcal{F}$ is the sheaf of sections of some finite-dimensional G-equivariant vector bundle on $\mathcal{O}_e$.*

Proof. For any $x \in \mathcal{O}_e$, let k_x denote the residue field of the local ring $\mathcal{S}_x$. Define a function $\phi : \mathcal{O}_e \to \mathbb{Z}$ by

$$\phi(x) = \dim_{k_x} \mathcal{F}_x \otimes_{\mathcal{S}_x} k(x).$$

This function is constant on $\mathcal{O}_e$ by virtue of G-equivariance. It follows with the help of Nakayama's Lemma that if ϕ is constant, $\mathcal{F}$ is locally free; *i.e.*, $\mathcal{F}$ is the sheaf of sections of a vector bundle. $\square$

Now, equivariant vector bundles over a single orbit just correspond to representations of the isotropy group of that orbit. Since we are at ease with virtual representations, we henceforth allow virtual equivariant vector bundles and virtual equivariant coherent sheaves on orbits. Of course, some care is required in manipulating the latter, since the global sections functor Γ is not exact.

Consider the sheaf $\mathcal{F} = \mu_* \pi^* \mathcal{W}_\eta$ over $\overline{\mathcal{O}_e}$. Suppose that $\eta|_{G^e} \simeq \tau$, so that $\Gamma(\mathcal{O}_e, \mathcal{F}) \simeq N(e, \tau)$. If we apply Proposition 2.1.2 to $\mathcal{F}$ and combine it with the preceding lemma, we obtain

$$\Gamma(\overline{\mathcal{O}_e}, \mu_* \pi^* \mathcal{W}_\eta) \simeq N(e, \tau) + \sum_{i=1}^{m} c_i \Gamma(\mathcal{O}_{e_i}, \mathcal{G}_{e_i}),$$

where each $c_i \in \{+1, -1\}$, and the $\mathcal{O}_{e_i}$'s are orbits contained in the boundary of $\mathcal{O}_e$. (We have merely combined the positive and negative boundary terms from Proposition 2.1.2 into a single summation.) Every term on the right-hand side of this expression can be written as a finite sum $\sum m_\sigma \operatorname{Ind}_T^G \sigma$. In the end, we will be able to obtain such a sum explicitly by analyzing the left-hand side of this equation, but we will not be able to separate that sum into terms arising from $N(e, \tau)$ and terms arising from vector bundles on smaller orbits. This fact motivates us to resolve our earlier approximate description of γ with the following definition:

Definition 2.1.4. Consider all expressions of the form

$$N(e, \tau) + \sum_{i=1}^{m} c_i \Gamma(\mathcal{O}_{e_i}, \mathcal{G}_{e_i}) = \sum_{\sigma \in \Lambda_+(G)} m_\sigma \operatorname{Ind}_T^G \sigma, \tag{2.3}$$

where each $\mathcal{O}_{e_i}$ is some orbit contained in the boundary of $\mathcal{O}_e$, $\mathcal{G}_{e_i}$ is a G-equivariant coherent sheaf on $\mathcal{O}_{e_i}$, and the c_i are integers. For each choice of $\mathcal{G}_{e_i}$'s and c_i's, there is some largest σ that appears on the right-hand side. We define

$$\gamma : \{(e, \tau)\} \to \Lambda_+(G)$$

by taking $\gamma(e, \tau)$ to be the minimum largest weight that can occur on the right-hand side, where the minimum is taken over all choices of $\mathcal{G}_{e_i}$'s and c_i's.

(There is no harm in allowing the c_i's to be any integers, rather than just $+1$ or -1, in this definition.)

Earlier it was promised that we would use Lie algebra cohomology to do computations on $\overline{\mathcal{O}_e}$. These techniques will not help us compute $\Gamma(\overline{\mathcal{O}_e}, \mu_* \pi^* \mathcal{W}_\eta)$ directly, but they will help us with the related sum

$$\sum_{q \geq 0} (-1)^q H^0(\overline{\mathcal{O}_e}, R^q \mu_* \pi^* \mathcal{W}_\eta). \tag{2.4}$$

We can analyze this module with the help of Proposition 2.1.2. We know that the restriction of μ to $\mu^{-1}(\mathcal{O}_e)$ is an isomorphism with $\mathcal{O}_e$. Therefore, $(\mu|_{\mu^{-1}(\mathcal{O}_e)})_*$ is exact, and its right-derived functors are zero. In other words, for any sheaf $\mathcal{G}$ on $G \times_{P_e} \mathfrak{v}_e$ and any $i > 0$, we have that $R^i \mu_*(\mathcal{G})|_{\mathcal{O}_e}$ is the zero sheaf.

Let $\mathcal{F}_i = R^i \mu_* \pi^* \mathcal{W}_\eta$. From the preceding discussion, we know that for $i > 0$, $\mathcal{F}_i$ must be supported on $\partial \mathcal{O}_e$. Applying Proposition 2.1.2 to $\mathcal{F}_i$, we see that $\Gamma(\overline{\mathcal{O}_e}, \mathcal{F}_i)$ is isomorphic to a sum of $\Gamma(\mathcal{O}_{e'}, \mathcal{G}_{e'})$'s over smaller orbits. So the expression in (2.4) is a G-module of the form shown on the left-hand side of Equation (2.3).

Our strategy for computing γ will be to show that for an appropriate choice of virtual P_e-representation (η, W_η) whose restriction to G^e is isomorphic to τ, the expression in (2.4) achieves the minimum possible largest σ when written in the form of the right-hand side of (2.3). We will obtain an explicit formula for (2.4), so once we understand what it means to choose η "appropriately," we will be able to compute $\gamma(e, \tau)$ explicitly as well.

We begin by tackling the computation of (2.4), to which end we employ the Leray spectral sequence:

$$H^p(\overline{\mathcal{O}_e}, R^q \mu_* \pi^* \mathcal{W}_\eta) \Longrightarrow H^{p+q}(G \times_{P_e} \mathfrak{v}_e, \pi^* \mathcal{W}_\eta).$$

We actually want to compute certain entries in the E_2 term of this spectral sequence, not the module to which the sequence converges. So we shall work backwards, obtaining the limit of the spectral sequence by other means, and then trying to deduce something about the E_2 term. Fortunately, we can conclude at the outset that most of the entries in the E_2 term vanish: because $\overline{\mathcal{O}_e}$ is an affine variety, its higher cohomology with coefficients in any coherent sheaf vanishes. And $R^q\mu_*\pi^*\mathcal{W}_\eta$ is coherent because μ is proper and $\pi^*\mathcal{W}_\eta$ is coherent. That is,

$$H^p(\overline{\mathcal{O}_e}, R^q\mu_*\pi^*\mathcal{W}_\eta) = 0 \qquad \text{for } p > 0.$$

We conclude that

$$H^0(\overline{\mathcal{O}_e}, R^q\mu_*\pi^*\mathcal{W}_\eta) = H^q(G \times_{P_e} \mathfrak{v}_e, \pi^*\mathcal{W}_\eta). \tag{2.5}$$

We still need to compute the right-hand side of the above equation. In this we are aided by the lemma below. If F is a vector space or a representation, we use $S^i(F)$ to denote the i-th symmetric power of F, and $S(F)$ to denote the entire symmetric algebra.

LEMMA 2.1.5. *We have the following relationship between cohomology groups over $G \times_{P_e} \mathfrak{v}_e$ and G/P_e:*

$$H^q(G \times_{P_e} \mathfrak{v}_e, \pi^*\mathcal{W}_\eta) = H^q(G/P_e, G \times_{P_e} (W_\eta \otimes S(\mathfrak{v}_e^*))). \tag{2.6}$$

Proof. The map π is the projection map of a vector bundle (with fiber $\mathfrak{v}_e$), so it is an affine morphism of varieties. Therefore,

$$H^q(G \times_{P_e} \mathfrak{v}_e, \pi^*\mathcal{W}_\eta) \simeq H^q(G/P_e, \pi_*\pi^*\mathcal{W}_\eta).$$

We just need to show that $\pi_*\pi^*\mathcal{W}_\eta$ is the sheaf of sections of $G \times_{P_e} (W_\eta \otimes S(\mathfrak{v}_e^*))$. But we also have that because the sheaf of sections of $\mathcal{W}_\eta$ is locally free of finite rank,

$$\pi_*\pi^*\mathcal{W}_\eta \simeq \pi_*(\mathcal{S}_{G\times_{P_e}\mathfrak{v}_e}) \otimes \mathcal{W}_\eta.$$

Now, it is easy to check that $\pi_*(\mathcal{S}_{G\times_{P_e}\mathfrak{v}_e})$ is just the sheaf of sections of the vector bundle in which the fiber over $x \in G/P_e$ is the ring of regular functions on $\pi^{-1}(x)$. But $\pi^{-1}(x) \simeq \mathfrak{v}_e$, and the regular functions on the latter space are given by the symmetric algebra of its dual. In other words, $\pi_*(\mathcal{S}_{G\times_{P_e}\mathfrak{v}_e})$ is the sheaf of sections of $G \times_{P_e} S(\mathfrak{v}_e^*)$. It follows that $\pi_*\pi^*\mathcal{W}$ is the sheaf of sections of $G \times_{P_e} (W \otimes S(\mathfrak{v}_e^*))$. $\square$

Now, a theorem of Bott [6] almost computes the right-hand side of (2.6) for us. Specifically, if U is a completely reducible representation of P_e and $\mathcal{U} = G \times_{P_e} U$ is the corresponding vector bundle over G/P_e, then

$$H^q(G/P_e, \mathcal{U}) = \bigoplus_{(\sigma, M_\sigma)\in\widehat{G}} H^q(\mathfrak{p}_e, \mathfrak{l}_e, \mathrm{Hom}(M_\sigma, U)) \otimes M_\sigma. \tag{2.7}$$

In our case, $W_\eta \otimes S(\mathfrak{v}_e^*)$ is not completely reducible (because $\mathfrak{v}_e^*$ is not). But an Euler-Poincaré-principle version of (2.7) still holds. That is, if we take the alternating sum of (2.6) with respect to q, the right-hand side is actually computed by the alternating sum of (2.7). Combining this with (2.5) and (2.4), we conclude that

$$\bigoplus_{q\geq 0}(-1)^q H^0(\overline{\mathcal{O}_e}, R^q\mu_*\pi^*\mathcal{W}_\eta) = \bigoplus_{\substack{q\geq 0\\(\sigma, M_\sigma)\in\widehat{G}}} (-1)^q H^q(\mathfrak{p}_e, \mathfrak{l}_e, \mathrm{Hom}(M_\sigma, W_\eta \otimes S(\mathfrak{v}_e^*)) \otimes M_\sigma.$$

Finally, we observe that in the above equation, we do not need to know the Lie algebra cohomology groups themselves, just their alternating sum. The Euler-Poincaré principle helps us out once more: tracing through the definitions of Lie algebra cohomology, one can show that

$$\bigoplus_{q\geq 0}(-1)^q H^q(\mathfrak{p}_e, \mathfrak{l}_e, U) = \bigoplus_{q\geq 0}(-1)^q \mathrm{Hom}_{L_e}(\textstyle\bigwedge^q \mathfrak{u}_e, U).$$

DEFINITION 2.1.6. Suppose that for each nilpotent orbit in G we have fixed a particular $\mathfrak{sl}(2)$-triple containing a representative element of the orbit, and that we have thence defined $\mathfrak{u}_e$, $\mathfrak{v}_e$, and $P_e = L_e U_e$ as above. For any (possibly virtual) representation (η, W_η) of P_e, we define a certain representation $\Theta(\eta)$ of G as follows:

$$\Theta(\eta) = \bigoplus_{\substack{q \geq 0 \\ (\sigma, M_\sigma) \in \widehat{G}}} (-1)^q \operatorname{Hom}_{L_e}(\textstyle\bigwedge^q \mathfrak{u}_e, \operatorname{Hom}(M_\sigma, W_\eta \otimes S(\mathfrak{v}_e^*))) \otimes M_\sigma. \tag{2.8}$$

For future reference, we summarize the developments in this section with the following result.

PROPOSITION 2.1.7. *If $\eta|_{G^e} \simeq \tau$, then $\Theta(\eta)$ is an approximation to $N(e, \tau)$, in the sense that*

$$\Theta(\eta) = N(e, \tau) + \sum_{\substack{\mathcal{O}_{e'} \subseteq \partial \mathcal{O}_e \\ (\tau', V_{\tau'}) \in \widehat{G}^{e'}}} c_{\tau'} N(e', \tau'),$$

where only finitely many of the $c_{\tau'}$'s are nonzero.

Proof. The preceding computations show that $\Theta(\eta)$ is isomorphic to the module in (2.4). Recall that that module has a decomposition of the sort seen in (2.3), by repeated application of Proposition 2.1.2. Lemma 2.1.3 tells us that the terms of the form $\Gamma(\mathcal{O}_{e_i}, \mathcal{G}_{e_i})$ are all spaces of sections of equivariant vector bundles on smaller orbits. Each $\Gamma(\mathcal{O}_{e_i}, \mathcal{G}_{e_i})$ can be identified as the space of sections of some G-equivariant vector bundle $\mathcal{V} = G \times_{G^{e_i}} V$. V in turn can be written as a finite sum of irreducible representations of G^e. If $V = \sum c_{\tau'} V_{\tau'}$, then

$$\Gamma(\mathcal{O}_{e_i}, \mathcal{G}_{e_i}) = \sum c_{\tau'} \Gamma(\mathcal{O}_{e_i}, \mathcal{V}_{\tau'}) = \sum c_{\tau'} N(e_i, \tau').$$

We sum up such expressions over all the $\mathcal{O}_{e'} \subseteq \partial \mathcal{O}_e$, and the proposition follows. ☐

2.2 Computations over the Orbit Closure

Let us recap the strategy we have been pursuing thus far. Proposition 2.1.7 gives us a formula enticingly similar in form to that in (2.3). In the latter, of course, the part of the sheaf supported on the boundary of the orbit is supposed to be carefully chosen so as to minimize the largest weight occurring on the right-hand side. If we could compute the expressions (2.8), then perhaps minimizing that largest term in (2.3) would just be a matter of making the "right" choice of virtual representation η.

In this, section, we show how to compute $\Theta(\eta)$ in the case that η is an *irreducible* representation; additivity then lets it be computed for any virtual representation. We start with some manipulations of the formula in (2.8).

$$\begin{aligned}
\Theta(\eta) &= \sum (-1)^q \operatorname{Hom}_{L_e}(\textstyle\bigwedge^q \mathfrak{u}_e, \operatorname{Hom}(M_\sigma, W_\eta \otimes S(\mathfrak{v}_e^*))) \otimes M_\sigma \\
&= \sum (-1)^q \operatorname{Hom}_{L_e}(\textstyle\bigwedge^q \mathfrak{u}_e, M_\sigma^* \otimes W_\eta \otimes S(\mathfrak{v}_e^*)) \otimes M_\sigma \\
&= \sum (-1)^q \operatorname{Hom}_{L_e}(\textstyle\bigwedge^q \mathfrak{u}_e \otimes S(\mathfrak{v}_e), M_\sigma^* \otimes W_\eta) \otimes M_\sigma
\end{aligned}$$

Next, we need the following general facts from commutative algebra.

LEMMA 2.2.1. *Let (π, F) be a finite-dimensional representation of a group H over a field k, not necessarily irreducible, and let $F \simeq F' \oplus F''$ be an H-equivariant decomposition of it. Then we have the following equivalence of representations:*

$$\textstyle\bigwedge^q F \simeq \displaystyle\sum_{r+s=q} \textstyle\bigwedge^r (F') \otimes \bigwedge^s F''.$$

Corollary 2.2.2. *In the context of the preceding lemma, we have the following equivalence of virtual representations:*

$$\sum_q (-1)^q \bigwedge^q F \simeq \left(\sum_r (-1)^r \bigwedge^r F'\right) \otimes \left(\sum_s (-1)^s \bigwedge^s F''\right).$$

Proof. The right-hand side can be rewritten as

$$\sum_q (-1)^q \sum_{r+s=q} \bigwedge^r F' \otimes \bigwedge^s F'',$$

and the result follows. □

Lemma 2.2.3. *Let (π, F) be a finite-dimensional representation of a group H over a field k. Then we have the following equivalence of virtual representations:*

$$\sum_{q\geq 0} (-1)^q \bigwedge^q F \otimes S^{m-q}(F) \simeq \begin{cases} k & \text{if } m = 0; \\ 0 & \text{if } m > 0. \end{cases}$$

(Here k is considered as the trivial representation of H.)

Proof. We prove this by interpreting the above expression as the alternating sum of the de Rham complex of polynomial functions on the space F^*. That is, let

$$C^q(F^*) = \bigwedge^q F \otimes S(F).$$

If F has a basis $x_1, \ldots, x_n$, then we suppose that elements of $C^q(F^*)$ can be written in the form

$$\sum (\text{polynomial in } x_1, \ldots, x_n)\, dx_{i_1} \wedge \cdots \wedge dx_{i_q},$$

where $d : C^q(F^*) \to C^{q+1}(F^*)$ has its usual meaning. What is the cohomology of this complex? It is easy to check that $H^q(F^*) = 0$ for $q \geq 1$, but in $C^0(F^*)$, the "constant functions" on F^* are in the kernel of d, but there is no image. So $H^0(F^*) = k$. Moreover, we can identify this space with the zeroth exterior power of F. Taking the alternating sum of the cohomology groups, we get

$$\sum_{q\geq 0} (-1)^q H^q(F^*) = \bigwedge^0 F = k.$$

We would now like to conclude that $\sum (-1)^q C^q(F^*) = 0$ by the Euler-Poincaré principle, but since the $C^q(F^*)$ are infinite-dimensional, some care is required when treating them in a virtual context. To that end, we impose a grading on the de Rham complex as follows:

$$C_i^q(F^*) = \bigwedge^q F \otimes S^{i-q}(F).$$

Then, d respects the grading, so we get graded cohomology groups as well. All the $H_i^q(F^*)$ are zero except when $q = i = 0$. We are now able to apply the Euler-Poincaré principle in this graded setting, and obtain

$$\sum_{q\geq 0} (-1)^q C_i^q(F^*) = \begin{cases} k & \text{if } i = 0; \\ 0 & \text{if } i > 0. \end{cases}$$

This is precisely the statement of the lemma. □

Applying the preceding lemmas to the expression for $\Theta(\eta)$ results in the following simplification.

Note that we permit ourselves to decompose $\mathfrak{u}_e$ over L_e as $\mathfrak{u}_e/\mathfrak{v}_e \oplus \mathfrak{v}_e$.

$$\begin{aligned}
\Theta(\eta) &= \sum_{(\sigma,M_\sigma)\in\widehat{G}} \operatorname{Hom}_{L_e}\left(\sum_q (-1)^q {\textstyle\bigwedge}^q \mathfrak{u}_e \otimes S(\mathfrak{v}_e), M_\sigma^* \otimes W_\eta\right) \otimes M_\sigma \\
&= \sum_{(\sigma,M_\sigma)\in\widehat{G}} \operatorname{Hom}_{L_e}\left(\sum_r (-1)^r {\textstyle\bigwedge}^r (\mathfrak{u}_e/\mathfrak{v}_e) \otimes \sum_s (-1)^s {\textstyle\bigwedge}^s \mathfrak{v}_e \otimes S(\mathfrak{v}_e), M_\sigma^* \otimes W_\eta\right) \otimes M_\sigma \\
&= \sum_{(\sigma,M_\sigma)\in\widehat{G}} \operatorname{Hom}_{L_e}\left(\sum_r (-1)^r {\textstyle\bigwedge}^r (\mathfrak{u}_e/\mathfrak{v}_e), M_\sigma^* \otimes W_\eta\right) \otimes M_\sigma \qquad (2.9)
\end{aligned}$$

This equation expresses $\Theta(\eta)$ as a sum over all finite-dimensional irreducible G-representations, with coefficients given by the dimensions of certain Hom_{L_e} groups. To mold this formula into a more tractable form, we turn our attention to those groups.

If λ is a dominant weight of L_e, let $(\eta_\lambda, W_\lambda)$ denote the irreducible L_e-representation of highest weight λ. Write $\sum(-1)^r {\textstyle\bigwedge}^r(\mathfrak{u}_e/\mathfrak{v}_e)$ as a sum of irreducible representations; say

$$\sum(-1)^r {\textstyle\bigwedge}^r(\mathfrak{u}_e/\mathfrak{v}_e) \simeq \sum_{i=1}^k m_i W_{\xi_i}, \qquad (2.10)$$

with the $m_i \in \mathbb{Z}$ (possibly negative) and the ξ_i distinct and dominant. Then, the coefficient of M_σ in (2.9) is given by

$$\sum_{i=1}^k m_i \dim \operatorname{Hom}_{L_e}(W_{\xi_i}, M_\sigma^* \otimes W_\eta). \qquad (2.11)$$

The dimension of the Hom group in this form of the equation is just the multiplicity of W_{ξ_i} in the L_e-representation $M_\sigma^* \otimes W$.

Let W_{L_e} denote the Weyl group of L_e, and let ρ_{L_e} denote half the sum of its positive roots.

Proposition 2.2.4. *Let $(\eta_\lambda, W_\lambda)$ be an irreducible representation of L_e with highest weight λ, and let $\lambda_1, \ldots, \lambda_k$ be the highest weights of the irreducible constituents of $\sum(-1)^r {\textstyle\bigwedge}^r(\mathfrak{u}_e/\mathfrak{v}_e)$, and let those constituents have multiplicities $m_1, \ldots, m_k$, respectively, as in (2.10). Then*

$$\Theta(\eta_\lambda) \simeq \sum_{i=1}^k \sum_{w\in W_{L_e}} m_i \operatorname{Ind}_T^G \mathbb{C}_{\lambda+\rho_{L_e}-w(\xi_i+\rho_{L_e})}. \qquad (2.12)$$

Proof. We need to evaluate (2.11) more explicitly; *i.e.*, we need to compute the multiplicity of W_{ξ_i} in $M_\sigma^* \otimes W_\lambda$. For this, we recall the following formula.

$$\langle W_{\xi_i}, M_\sigma^* \otimes W_\lambda\rangle = \sum_{w\in W_{L_e}} (-1)^w \dim M^*_{\sigma, w(\xi_i+\rho_{L_e})-(\lambda+\rho_{L_e})}. \qquad (2.13)$$

Here $M^*_{\sigma,\nu}$ is the ν-weight space of M^*_σ.

This reasoning lets us compute as follows:

$$\begin{aligned}
\Theta(\eta_\lambda) &= \sum_{(\sigma,M_\sigma)\in\widehat{G}} \langle W_{\xi_i}, M_\sigma^* \otimes W_\lambda\rangle M_\sigma \\
&= \sum_{(\sigma,M_\sigma)\in\widehat{G}} \sum_{w\in W_{L_e}} (-1)^w \dim M^*_{\sigma, w(\xi_i+\rho_{L_e})-(\lambda+\rho_{L_e})}.
\end{aligned}$$

Now, the dimensions of the weight spaces in a dual of a given representation are such that $\dim M^*_{\sigma,\nu} = \dim M_{\sigma,-\nu}$. In addition, the dimension of a weight space of M_σ can be regarded as the multiplicity of a certain T-representation in the restriction $M_\sigma|_T$. With these observations in mind, we rewrite as follows.

$$\Theta(\eta) = \sum_{(\sigma, M_\sigma)\in\widehat{G}} \sum_{w\in W_{L_e}} (-1)^w \langle \mathbb{C}_{\lambda+\rho_{L_e}-w(\xi_i+\rho_{L_e})}, M_\sigma|_T\rangle M_\sigma.$$

The proposition follows by Frobenius reciprocity. □

Now, we could of course replace any weight appearing on the right-hand side of (2.12) by any W_G-conjugate of itself, where W_G denotes the Weyl group of G, without changing the validity of the equation. It follows that for any virtual representation (η, W_η) of P_e, we can now write down an expression

$$\Theta(\eta) = \sum_{\sigma\in\Lambda_+(G)} m_\sigma(\eta) \operatorname{Ind}_T^G \mathbb{C}_\sigma, \tag{2.14}$$

in which every weight σ is taken to be dominant; that is, we know how to write down the expression in (1.1).

2.3 Further Simplification in Special Cases

Recall our hope from the beginning of Section 2.2 that the "right" choice of virtual representation η with $\eta|_{G^e} \simeq \tau$ would permit $\gamma(e,\tau)$ to be read off from a computation of $\Theta(\eta)$. What might constitute such a "right" choice? An elementary approach to constructing such an η would be to start with an irreducible P_e-representation η_0 such that τ occurs as a summand in $\eta_\lambda|_{G^e}$, and then subtract off "smaller" P_e-representations as necessary.

We naturally prefer easy computations to hard ones, and finding the right "smaller" P_e-representations to subtract off might well be hard. Perhaps, then, the "right" choice of virtual representation η comes down to a "right" choice of irreducible representation η_0 as a starting point. That is, we hope that if we choose η_0 properly, the largest term of $\Theta(\eta)$ is actually the largest term of $\Theta(\eta_0)$, in which case we do not to worry about the full virtual structure of η, and we only even need to compute the expression (2.12) for one irreducible.

Even for a single irreducible, however, writing out (2.12) is not trivial. In the following two subsections, we explore how that formula might be simplified under two different sorts of assumptions about $\mathfrak{u}_e$ and $\mathfrak{v}_e$. We shall also make explicit, under each of these assumptions, what it might mean to make the "right" choice of irreducible P_e-representation.

2.3.1 Richardson orbits Richardson orbits are orbits that are induced from the zero orbit in some Levi subalgebra. If $\mathcal{O}_e$ is such an orbit, and $P_e = L_eU_e$ is taken to be a parabolic such that $\mathcal{O}_e$ is induced from the zero orbit in $\mathfrak{l}_e$, we have that $\mathcal{O}_e$ meets $\mathfrak{u}_e$ in a dense open subset. We can therefore take $\mathfrak{v}_e = \mathfrak{u}_e$.

Another requirement that must be met for the work of the preceding two sections to be applicable is the containment condition $G^e \subseteq P_e$. This condition does not always hold, so what follows cannot be applied to all Richardson orbits in all reductive groups. However, the condition is satisfied in the case of one important example: that of $GL(n,\mathbb{C})$, in which every orbit is Richardson. (In the next chapter, we shall fix this as our choice of parabolic for computations in that group.)

In this setting, the only highest L_e-weight on $\mathfrak{u}_e/\mathfrak{v}_e$ is 0, so the formula in (2.12) reduces to

$$\sum_{w\in W_{L_e}} (-1)^w \operatorname{Ind}_T^G \mathbb{C}_{\lambda+\rho_{L_e}-w\rho_{L_e}}. \tag{2.15}$$

We shall find it convenient to be able to refer to the weights appearing on the right-hand side of the formula with an abbreviated notation, so we introduce such a notation now. If λ is a dominant weight of L_e, then for any $w \in W_{L_e}$, we define $E_w\lambda$ to be the weight

$$\lambda + \rho_{L_e} - w\rho_{L_e},$$

In particular, we are often interested in the largest term occurring in (2.15). The size of these terms is given by

$$\|E_w\lambda\|^2 = \|\lambda + \rho_{L_e}\|^2 + 2\langle\lambda + \rho_{L-e}, -w\rho_{L_e}\rangle + \|\rho_{L_e}\|^2.$$

(For the last term, we have $\|-w\rho_{L_e}\|^2 = \|\rho_{L_e}\|^2$.) Only the middle term depends on w, and this term is maximized when $-w\rho_{L_e}$ is dominant. Since $-\rho_{L_e}$ is always in the W_{L_e}-orbit of ρ_{L_e}, $E_w\lambda$ achieves its maximum size when it is equal to $\lambda + 2\rho_{L_e}$. We introduce the additional notation

$$E\lambda = \lambda + 2\rho_{L_e},$$

and we remark again that there always exists some $w \in W_{L_e}$ such that $E_w\lambda = E\lambda$.

Recall our hope that if (η, W_η) is chosen "appropriately" and has $\eta|_{G^e} \simeq \tau$, then $\gamma(e,\tau)$ may be computed by examining $\Theta(\eta)$. We make this hope more precise now.

We continue to assume that $\mathfrak{u}_e = \mathfrak{v}_e$, so that (2.12) takes the form (2.15). Let us suppose that there is a unique largest σ occurring in (2.14) with $m_\sigma(\eta) \neq 0$, and let $\xi(\eta)$ denote this largest weight σ. We now have an explicit formula for $\Theta(\eta)$ in such a form that it is easy to read off $\xi(\eta)$: if $\eta = \sum c_i\eta_{\lambda_i}$, where the η_{λ_i}'s are irreducible representations with highest weights λ_i, then

$$\xi(\eta) = \text{the largest } E\lambda_i.$$

Note that even if the supposition of the existence of a unique such σ fails, and there are multiple weights of maximal size that occur with nonzero coefficient, then there is nevertheless still a well-defined maximal size. By abuse of notation, we let $\|\xi(\eta)\|^2$ denote this maximal size. Thus, $\|\xi(\eta)\|^2$ is defined for any η, even when $\xi(\eta)$ is not.

Claim 2.3.1. *Suppose that $\mathfrak{u}_e = \mathfrak{v}_e$, and let $(\eta_\lambda, W_\lambda)$ be an irreducible P_e-representation of highest weight λ, with the following properties:*

(a) The irreducible G^e-representation (τ, V_τ) occurs in the restriction $W_\lambda|_{G^e}$.

(b) Over all irreducible P_e-representations whose restrictions to G^e contain τ as a constituent, the function $\|\xi(\cdot)\|^2$ takes its minimal value at η_λ.

Then, there is a virtual representation (η, W_η) of P_e consisting of $(\eta_\lambda, W_\lambda)$ plus other irreducible representations with smaller highest weights, and with the property that $\eta|_{G^e} \simeq \tau$. In particular,

$$\gamma(e,\tau) = \xi(\eta) = \xi(\eta_\lambda) = E\lambda.$$

This claim will serve as intuition for much of our work in $GL(n,\mathbb{C})$; ultimately, in Chapter 5, we will establish this claim in the course of proving Theorem 3.

2.3.2 Using the Jacobson-Morozov parabolic If we use the Jacobson-Morozov parabolic to define $\mathfrak{u}_e$ and $\mathfrak{v}_e$, then $\mathfrak{u}_e \neq \mathfrak{v}_e$ unless the niloptent e is even. Nevertheless, this is an easy way to produce the setup of Section 2.1 for any orbit in any group, so it is worth seeing what we can make of Proposition 2.2.4 in this setting.

Recall that the construction entails construing the nilpotent element e as an element of an $\mathfrak{sl}(2)$-triple $\{h, e, f\}$; then, $\mathfrak{u}_e$ and $\mathfrak{v}_e$ are defined as certain sums of h-eigenspaces:

$$\mathfrak{u}_e = \bigoplus_{i\geq 1}\mathfrak{g}_i \qquad\qquad \mathfrak{v}_e = \bigoplus_{i\geq 2}\mathfrak{g}_i,$$

where $\mathfrak{g}_i$ is the h-eigenspace of eigenvalue i. We can thence identify $\mathfrak{u}_e/\mathfrak{v}_e = \mathfrak{g}_1$.

Since it is possible to read off the weights that occur in $\mathfrak{g}_1$ from the weighted Dynkin diagram of an orbit, we can in principle evaluate expressions of the form (2.12) now. Nevertheless, identifying the extremal L_e-weights in the alternating sum of exterior powers of $\mathfrak{g}_1$ can be daunting computationally. The list of *all* weights in this alternating exterior algebra is easier to specify: there are $2^{\dim\mathfrak{g}_1}$ of them, and they are given as sums of subsets of the weights of $\mathfrak{g}_1$. And since the extremal weights of an irreducible representation are larger in size than any of the other weights, it is reasonable to guess that the largest term in (2.12) can be found simply by trying all possible subsets of weights in $\mathfrak{g}_1$.

Let $\mathcal{P}$ denote the set of weights occurring in $\mathfrak{g}_1$, and let $w_0 \in W_{L_e}$ be the element which exchanges the dominant and antidominant Weyl chambers.

$$\mathcal{F}(\lambda) = \left\{ \lambda + \sum_{\nu \in P} w_0(-\nu) \,\middle|\, P \subseteq \mathcal{P} \right\} \tag{2.16}$$

With the goal of stating an analogue of Claim 2.3.1, we introduce the following notation.

$$F\lambda = \text{the element of } \mathcal{F}(\lambda) \text{ of maximal size,} \tag{2.17}$$

provided of course that a unique such element exists. Even if not, we, as before, abusively employ the following notation:

$$\|F\lambda\|^2 = \max\{\|\nu\|^2 \mid \nu \in \mathcal{F}(\lambda)\}.$$

Claim 2.3.2. *Suppose that $\mathfrak{u}_e$ is the nilradical of the Jacobson-Morozov parabolic for $\mathcal{O}_e$, and $\mathfrak{v}_e$ is the sum of the eigenspaces of eigenvalue at least 2 for the semisimple element of the Jacobson-Morozov triple. Let $(\eta_\lambda, W_\lambda)$ be an irreducible P_e-representation of highest weight λ, with the following properties*

(a) The irreducible G^e-representation (τ, V_τ) occurs in the restriction $W_\lambda|_{G^e}$.

(b) Over all highest weights of all irreducible P_e-representations whose restrictions to G^e contain τ as a constituent, the function $\|F(E\cdot)\|^2$ takes its minimal value at λ.

Then, there is a virtual representation (η, W_η) of P_e consisting of $(\eta_\lambda, W_\lambda)$ plus other irreducible representations with smaller highest weights, and with the property that $\eta|_{G^e} \simeq \tau$. In particular,

$$\gamma(e, \tau) = F(E\lambda).$$

This claim guides some of the work in Chapter 6, as well as some of the computations in Appendix B.

2.4 The Zero and Principal Orbits

We conclude this chapter by giving an explicit computation of γ on the zero and principal orbits in any complex reductive group. We first recall the definition of minuscule weights, and we introduce the idea of majuscule weights by analogy.

Definition 2.4.1. Let Δ be the root system for G, and let Π be a choice of simple roots. For $\alpha \in \Delta$, let $\alpha^\vee$ denote the corresponding coroot. An integral weight λ is called *minuscule* if

$$|\langle \alpha^\vee, \lambda \rangle| \leq 1 \qquad \text{for all } \alpha \in \Delta.$$

Conversely, λ is said to be *majuscule* if

$$|\langle \alpha^\vee, \lambda \rangle| \geq 2 \qquad \text{for all } \alpha \in \Pi.$$

Let $\mathcal{O}_0$ denote the zero orbit, and $\mathcal{O}_{\text{princ}}$ the principal orbit. The isotropy group of $\mathcal{O}_0$ is of course all of G.

Theorem 2.4.2. *Let G be connected complex reductive Lie group. Let V_λ be the irreducible representation of highest weight λ. When restricted to the zero orbit, the image of γ is precisely the set of majuscule weights, and γ is given by the formula*

$$\gamma(\mathcal{O}_0, V_\lambda) = \lambda + 2E_G.$$

When restricted to the principal orbit, the image of γ is precisely the set of minuscule weights. In this case, γ gives a bijection between representations of the center of G and the set of minuscule weights.

Proof. Computing γ on the zero orbit is quite easy. The space of sections of the vector bundle over the orbit arising from the irreducible representation V_λ is again just V_λ: we are not overwhelmed by the mathematical complexity of vector bundles over a single point. Now, we just need to write V_λ itself in the form of (1.1) and read off its largest weight. That task is easy enough: one form of the Weyl character formula gives us

$$V_\lambda = \sum_{w \in W} (-1)^w \operatorname{Ind}_T^G \mathbb{C}_{\lambda+\rho_G - w\rho_G}.$$

The largest term on the right-hand side of this expression is $\lambda + 2\rho_G$. Now, ρ_G itself has inner product 1 with every simple coroot, so it follows that weights of the form $\lambda + 2\rho_G$, where λ is dominant, are precisely the majuscule weights.

For the principal orbit, as noted above, the reductive part of the isotropy group is just the center of the group $Z(G)$, and the Jacobson-Morozov Levi is the torus T. There is a well-known bijection between representations of the center of G and minuscule weights, which are in turn a set of coset representatives for

$$\Lambda(G)/\mathbb{Z}\Delta(G),$$

where $\mathbb{Z}\Delta(G)$ is the root lattice. Indeed, for any given representation of T, *i.e.*, any weight, its restriction to $Z(G)$ depends only on the coset to which it belongs in the above quotient. Now, let λ be some nonminuscule weight, and let λ' be the minuscule weight in the same coset of $\Lambda(G)/\mathbb{Z}\Delta(G)$ as λ. The T-representations λ and λ' are isomorphic upon restriction to $Z(G)$, so $\mathbb{C}_{\lambda'} - \mathbb{C}'_\lambda$ is a virtual T-representation whose restriction to $Z(G)$ is the zero representation. Applying (2.12) in this context, with $L_e = T$, we find that

$$\operatorname{Ind}_T^G \mathbb{C}_{\lambda'} - \operatorname{Ind}_T^G \mathbb{C}_\lambda \tag{2.18}$$

is an expression supported on the boundary of the principal orbit, so we can add it to any expression of the form in (2.3). In particular, if we have obtained such an expression in an attempt to compute some $\gamma(e_{\text{princ}}, \tau)$ in which the largest term is not minuscule, we can always add an expression of the form (2.18) to eliminate that largest term. It follows that over all possible expressions of the form (2.3), the one with the minimal largest weight must include only minuscule weights. Finally, since the minuscule weights are in one-to-one correspondence with the irreducible representations of $Z(G)$, we observe that if the virtual T-representation under consideration is to have the right restriction to $Z(G)$, we must have exactly one minuscule weight appearing in (2.3), *viz.*, the one corresponding to the $Z(G)$-representation under consideration. Thus, on the principal orbit, the values of γ are precisely the minuscule weights. □

CHAPTER 3

Computational Machinery for the General Linear Group

In this chapter, we introduce the machinery that will enable us to prove Theorem 3. Recall that in $GL(n, \mathbb{C})$, nilpotent orbits are indexed by partitions of n. We shall use boldface letters like $\mathbf{d}$ to denote partitions, and "absolute-value bars" as in $|\mathbf{d}|$ to denote the sum of a partition. An example of the notation we use to write down a partition explicitly, with exponents giving multiplicities, is as follows: $\mathbf{d} = [6, 3^2, 2, 1^4]$. (Thus, for this partition, $|\mathbf{d}| = 18$.)

For the next few chapters, we shall be working exclusively with $GL(n, \mathbb{C})$, so we modify the notation of Chapter 2 to one that is more convenient in this setting. We label orbits and centralizers by partitions rather than by nilpotent elements: thus, $\mathcal{O}_{[n]}$ is the principal orbit, and $G^{[1^n]}$ is the centralizer of the zero orbit (*i.e.*, the entire group).

In the context of $GL(n, \mathbb{C})$, $P_{\mathbf{d}} = L_{\mathbf{d}} U_{\mathbf{d}}$ will always denote a certain parabolic subgroup with respect to which the orbit $\mathcal{O}_{\mathbf{d}}$ is Richardson. (The group $L_{\mathbf{d}}$ is identified more explicitly below.) $G^{\mathbf{d}}$ also has a decomposition as a semidirect product of a reductive group $G^{\mathbf{d}}_{\mathrm{red}}$ and a unipotent group [7]. In both these cases, the unipotent part must act trivially in irreducible algebraic representations, and henceforth, we liberally confuse irreducible algebraic representations of the whole group with those of the Levi factor.

In particular, we adopt the following abuse of notation: if (τ, V_τ) is an irreducible representation of $G^{\mathbf{d}}$, we also write "τ" for the highest weight of V_τ regarded as a $G^{\mathbf{d}}_{\mathrm{red}}$-module. This ought not be overly confusing, since we will never speak of $G^{\mathbf{d}}$-modules that are not irreducible. With $P_{\mathbf{d}}$ and $L_{\mathbf{d}}$, on the other hand, we will be dealing with reducible representations, so we shall be somewhat less abusive in notation: if λ is a dominant weight of $L_{\mathbf{d}}$, we let $(\eta_\lambda, W_\lambda)$ be the irreducible $P_{\mathbf{d}}$-representation on which $L_{\mathbf{d}}$ acts with highest weight λ and $U_{\mathbf{d}}$ acts trivially.

Given a partition $\mathbf{d} = [k_1^{a_1}, \ldots, k_l^{a_l}]$, define $GL(\mathbf{d})$ to be the subgroup $GL(k_1)^{a_1} \times \cdots \times GL(k_l)^{a_l}$ of $GL(n)$. Let $\mathbf{d}^*$ denote the dual partition of $\mathbf{d}$. It is easy to check that $G^{\mathbf{d}}_{\mathrm{red}}$ is isomorphic to $GL(a_1) \times \cdots \times GL(a_l)$, and that $L_{\mathbf{d}}$ is isomorphic to $GL(\mathbf{d}^*)$ [7].

Let $D(\mathbf{d})$ denote the weight lattice of $GL(\mathbf{d}^*)$, and let

$$D_n = \coprod_{|\mathbf{d}|=n} D(\mathbf{d}).$$

(We write $D(\mathbf{d})$ instead of $\Lambda(GL(\mathbf{d}^*))$ for purposes of analogy with another notation that we introduce shortly.) If $\mathbf{d}^* = [m_1^{b_1}, \ldots, m_p^{b_p}]$, then a weight of $GL(\mathbf{d}^*)$ consists of b_1 m_1-tuples of integers, b_2 m_2-tuples, *etc.* Since $|\mathbf{d}^*| = n$, a weight of $GL(\mathbf{d}^*)$ is really just an n-tuple of integers, which one might think as being subdivided into smaller tuples according to the parts of $\mathbf{d}^*$.

3.1 Weight Diagrams

We shall devise a particular way of writing down elements of D_n with diagrams, but there may be multiple possible diagrams for a given element of D_n. We write $\tilde{D}_n = \coprod \tilde{D}(\mathbf{d})$ for this set of diagrams, and we shall have a collection of surjective maps $\pi : \tilde{D}(\mathbf{d}) \to D(\mathbf{d})$. For each orbit, we also get a map

$$\kappa : \tilde{D}_n \to \coprod_{|\mathbf{d}|=n} \Lambda_+(G^{\mathbf{d}}_{\mathrm{red}}),$$

that computes the restriction of an $L_{\mathbf{d}}$-weight to $G^{\mathbf{d}}_{\mathrm{red}}$. (*Caveat lector*, however: see Remark 3.1.2.) Recall also the map E which adds $2\rho_{L_{\mathbf{d}}}$ from the end of Chapter 2. We will have a lifting a this map to one at the diagram level: $E : \tilde{D}(\mathbf{d}) \to \tilde{D}(\mathbf{d})$.

In Chapter 4, we will define a certain subset $\tilde{D}^{\circ}_n$ of $\tilde{D}_n$ which will be in bijection with the set of dominant weights of $GL(n)$, and which will serve as an intermediary between the set of pairs $\{(\mathbf{d}, \tau)\}$ and $\Lambda_+(G)$ in our quest to compute γ. The proposed weight-diagram method for computing γ is encapsulated in the following diagram.

$$\coprod_{|\mathbf{d}|=n} \Lambda_+(G^{\mathbf{d}}_{\mathrm{red}}) \xleftarrow{\ \kappa\ } \tilde{D}_n \xrightarrow{\ \alpha\ } \tilde{D}^{\circ}_n \underset{\simeq}{\xleftarrow{E^{-1}\circ\beta}} \widehat{G}$$

Once we have this machinery in place, we will be well on our way to proving Claim 2.3.1 in the context of $GL(n, \mathbb{C})$, and to proving that γ is a bijection.

Definition 3.1.1. A *weight diagram* is a diagram of left-justified rows of boxes with integer entries. The set of all weight diagrams with a total of n boxes is denoted $\tilde{D}_n$. If $\mathbf{d}$ is a partition of n, then a weight diagram $X \in \tilde{D}_n$ is said to be *of shape-class* $\mathbf{d}$ if the row lengths of X are given by the parts of $\mathbf{d}$. The set of all weight diagrams of shape-class $\mathbf{d}$ is denoted $\tilde{D}(\mathbf{d})$.

We do not say "of shape $\mathbf{d}$," as one might when discussing Young tableaux, because in our case the partition does not determine the shape of the diagram. (The row lengths are not required to be nonincreasing as one goes down the diagram.)

How do diagrams in $\tilde{D}(\mathbf{d})$ relate to weights in $D(\mathbf{d})$ or $\Lambda_+(G^{\mathbf{d}}_{\mathrm{red}})$? Recall that $L_{\mathbf{d}} \simeq GL(\mathbf{d}^*)$. If $X \in \tilde{D}(\mathbf{d})$, then the column lengths of X are the parts of $\mathbf{d}^*$. We define $\pi : \tilde{D}_n \to D_n$ by regarding each column of a diagram X as the weight of the corresponding factor of $L_{\mathbf{d}} \simeq GL(\mathbf{d}^*)$.

Recall that if $\mathbf{d} = [k_1^{a_1}, \ldots, k_l^{a_l}]$, then $G^{\mathbf{d}}_{\mathrm{red}} = GL(a_1) \times \cdots GL(a_l)$. Given $X \in \tilde{D}(\mathbf{d})$, then, we can obtain a weight of $G^{\mathbf{d}}_{\mathrm{red}}$ as follows: the coordinates of the $GL(a_i)$-component of the weight will be the sums of the a_i rows of length k_i in X. For example,

$$X = \begin{array}{|c|c|c|c|} \hline 5 & 6 \\ \cline{1-2} 2 & 3 & 1 & 0 \\ \hline 8 & 2 \\ \cline{1-3} 1 & 1 & 1 \\ \cline{1-3} \end{array}$$

gives the weight $(6, 3, (11, 10))$ of $G^{[4,3,2^2]}_{\mathrm{red}} = GL(1) \times GL(1) \times GL(2)$. This row-summing procedure associates to any diagram in $\tilde{D}(\mathbf{d})$ a weight of $G^{\mathbf{d}}_{\mathrm{red}}$, and hence defines a map $\tilde{\kappa} : \tilde{D}(\mathbf{d}) \to \Lambda(G^{\mathbf{d}}_{\mathrm{red}})$. We define $\kappa : \tilde{D}(\mathbf{d}) \to \Lambda_+(G^{\mathbf{d}}_{\mathrm{red}})$ by

$$\kappa(X) = \text{the unique dominant weight in the Weyl group orbit of } \tilde{\kappa}(X).$$

For τ an irreducible algebraic representation of $G^{\mathbf{d}}$ or, equivalently, a dominant weight of $G^{\mathbf{d}}_{\mathrm{red}}$, we define

$$\tilde{D}(\mathbf{d}, \tau) = \{X \in \tilde{D}(\mathbf{d}) \mid \kappa(X) = \tau\}.$$

Remark 3.1.2. It seems that one ought to be able to restrict weights of $L_{\mathbf{d}}$ to $G^{\mathbf{d}}_{\mathrm{red}}$ without going through the artifice of weight diagrams, but κ does not descend to a map $D(\mathbf{d}) \to \Lambda(G^{\mathbf{d}}_{\mathrm{red}})$. That is, one can readily construct diagrams X, X' of the same shape-class, and whose columns are the same, but whose row-sums are not. The reason for this is that there is not a unique imbedding $G^{\mathbf{d}}_{\mathrm{red}} \hookrightarrow L_{\mathbf{d}}$. Diagrams of a given shape-class determine conjugate imbeddings, but only if one fixes a particular shape of shape-class $\mathbf{d}$ (thus choosing a specific imbedding) and considers only diagrams of that shape does κ induce a restriction map $D(\mathbf{d}) \to \Lambda_+(G^{\mathbf{d}}_{\mathrm{red}})$.

We introduce the notation X_{ir} to refer to the entry in row i, column r of the diagram X.

Definition 3.1.3. X_{ir} and X_{jr} are said to be *column-consecutive* if $j > i$ and all the positions $X_{i+1,r}, \ldots, X_{j-1,r}$ are empty. X_{jr} is then the *column-successor* of X_{ir}, and X_{ir} is the *column-*

predecessor of X_{jr}. An entry is *column-last* if it has no column-successor, and it is *column-first* if it has no column-predecessor.

The next definition we need is that of $E : D_n \to D_n$. EX is defined to be a a diagram of the same shape as X, whose entries are given by

$$\begin{aligned}(EX)_{ir} = X_{ir} + \#\{X_{jr} \mid X_{jr} < X_{ir}, \text{ or } X_{jr} = X_{ir} \text{ with } j > i\} \\ - \#\{X_{jr} \mid X_{jr} > X_{ir}, \text{ or } X_{jr} = X_{ir} \text{ with } j < i\}.\end{aligned}$$

For the example above, we compute

$$E(X) = \begin{array}{|c|c|c|c|} \hline 6 & 9 \\ \cline{1-4} 1 & 4 & 2 & 0 \\ \hline 11 & 1 \\ \cline{1-3} -2 & -2 & 0 \\ \cline{1-3} \end{array}$$

It is easy to verify that this E agrees with the E we already have for weights of the $L_{\mathbf{d}}$'s, in the following sense: if $\pi(X)$ is dominant, then so is $\pi(EX)$, and in fact

$$\pi(EX) = \pi(X) + 2\rho_{L_{\mathbf{d}}}.$$

(In general, $\pi(EX)$ is equal to $\pi(X)$ plus the sum of positive roots, with respect to a choice of positive roots that makes $\pi(X)$ dominant.)

3.2 Manipulating Elements of $\tilde{D}_n$

In this section, we lay the groundwork for the algorithm α to be defined in the next chapter. Specifically, we write down a list of three "moves" that one can perform on a diagram in $\tilde{D}_n$ to obtain a new diagram. Each move will have associated to it certain permissibility criteria. The algorithm will be essentially, "At each stage, do any move whose permissibility criteria are satisfied; stop when no move's permissibility criteria are satisfied." First, a few definitions:

DEFINITION 3.2.1. X_{ir} is *lowerable* if either it is column-last, or it has column-successor $X_{i'r}$ with $X_{ir} > X_{i'r}$. It is *raisable* if either it is column-first, or it has column-predecessor $X_{i'r}$ with $X_{ir} < X_{i'r}$.

To *lower* (resp. *raise*) X_{ir}, or X at position ir, is to construct a new diagram of the same shape and entries as X, except that its entry in position ir is obtained by subtracting 1 from (resp. adding 1 to) X_{ir}.

Let $X \in \tilde{D}_n$. We state four properties that X may have:

$\mathbf{p_1}(r)$ ($r > 1$) For any s, $1 \le s < r$, such that $s \equiv r \pmod 2$, we have

(a) if r and s are odd and $X_{is} < X_{ir}$, then X_{is} is not raisable.

(b) if r and s are even and $X_{is} > X_{ir}$, then X_{is} is not lowerable.

$\mathbf{p_2}(r)$ ($r > 1$) For any s, $1 \le s < r$, we have

(a) if $X_{is} \le X_{ir} - 2$, then X_{is} is not raisable.

(b) if $X_{is} \ge X_{ir} + 2$, then X_{is} is not lowerable.

$\mathbf{p_3}(r)$ ($r > 1$) If r is odd, every difference $EX_{ir} - EX_{i,r-1}$ is either 0 or $+1$. If r is even, every such difference is either 0 or -1.

$\mathbf{p_4}(r)$ Column r of X has entries in non-increasing order.

$$
\begin{aligned}
q_1(X) &= \|EX\|^2 \\
q_2(X) &= \sum_{r \text{ even}} \sum_i EX_{ir} \\
q_3(X) &= \sum_{r \text{ odd}} \sum_i r\, EX_{ir} - \sum_{r \text{ even}} \sum_i r\, EX_{ir} \\
q_4(X) &= -\max\{r \mid \mathbf{P_3}(r) \text{ and } \mathbf{P_4}(r) \text{ both hold}\} \\
q_5(X) &= \sum_{\{j \mid X_{jr} \text{ nonempty}\}} \tilde{q}_{5;jr}(X) \qquad \text{where } r = q_4(X)+1 \\
q_6(X) &= \sum_{\{j \mid X_{jr} \text{ nonempty}\}} \tilde{q}_{6;jr}(X) \qquad \text{where } r = q_4(X)+1,
\end{aligned}
$$

where

$$
\begin{aligned}
\tilde{q}_{5;ir}(X) &= \begin{cases} \max\{EX_{ir} - (EX_{i,r-1}+1), EX_{i,r-1} - EX_{ir}\} & \text{if } r \text{ is odd} \\ \max\{EX_{ir} - EX_{i,r-1}, (EX_{i,r-1}-1) - EX_{ir}\} & \text{if } r \text{ is even} \end{cases} \\
\tilde{q}_{6;ir}(X) &= \sum_{\substack{i'<i \\ X_{i'r} \text{ nonempty}}} \max\{0, X_{ir} - X_{i'r}\} + \sum_{\substack{i'>i \\ X_{i'r} \text{ nonempty}}} \max\{0, X_{i'r} - X_{ir}\}
\end{aligned}
$$

TABLE 3.1. Some integer-valued functions on $\tilde{D}_n$.

Every X is said to have properties $\mathbf{p_1}(1)$ and $\mathbf{p_3}(1)$, for convenience.

We also define capital-letter versions of these properties: X is said to have $\mathbf{P_1}(r)$ if it has all of $\mathbf{p_1}(1), \ldots, \mathbf{p_1}(r)$; the properties $\mathbf{P_3}(r)$ and $\mathbf{P_4}(r)$ are defined similarly. We suppose that the properties $\mathbf{P_1}(0)$, $\mathbf{P_3}(0)$, and $\mathbf{P_4}(0)$ always hold, as do $\mathbf{P_1}(1)$ and $\mathbf{P_3}(1)$.

Now, the main idea of α will be to modify a diagram X so as to minimize the size of $\|EX\|^2$, but we do not understand that goal well enough yet to directly give an algorithm for it. Instead, we will describe an algorithm that works by trying to bring X closer to having all the properties $\mathbf{P_1}(r)$, $\mathbf{P_2}(r)$, $\mathbf{P_3}(r)$, and $\mathbf{P_4}(r)$ for all r; we shall subsequently see that this has the desired consequence for $\|EX\|^2$ as well.

In Table 3.1, we define some integer-valued functions that will be used to measure the progress of the algorithm: the smaller the values of these functions, the more progress we have made. The use of function q_1 is evident; minimizing it is the purpose of the entire algorithm. Now, one implication of $\mathbf{p_1}(r)$ and $\mathbf{p_2}(r)$ is, roughly, that even-numbered columns ought to have smaller entries than odd-numbered columns; q_2 is used to make sure that the entries in even-numbered columns do not become too large. Another implication of these first two properties is that the largest entries ought to appear in the leftmost odd-numbered columns, and the smallest entries in the left-most even-numbered columns. The function q_3 tells us if we have too many excessively large or small entries too far right in the diagram.

Next, q_4 just tells us how many columns on the left-hand side of the diagram satisfy $\mathbf{p_3}(r)$ and $\mathbf{p_4}(r)$. It is negative just because we want an *increase* in the number of such columns as we make progress. For the leftmost column that does not satisfy $\mathbf{p_3}(r)$ and $\mathbf{p_4}(r)$, q_5 measures quite directly how far it is from having $\mathbf{p_3}(r)$, and q_6 does the same for $\mathbf{p_4}(r)$.

REMARK 3.2.2. X has $\mathbf{p_3}(r)$ for $r = -q_4(X)+1$ if and only if $q_5(X) = 0$, and it has $\mathbf{p_4}(r)$ for $r = -q_4(X)+1$ if and only if $q_6(X) = 0$.

The three moves that can be performed in the algorithm for α are shown in Table 3.2. The distinct rows shown in the diagram for each move in this table need not actually be consecutive; however, any intermediary rows *must be shorter than* $r-1$ boxes (shorter than $s-1$ boxes for move **A**). The rows that are shown in the diagrams may or may not be longer than r boxes.

The algorithm for α will also make use of the inverse moves $\mathbf{A}^{-1}$ and $\mathbf{B}^{-1}$. (Of course, $\mathbf{C}$ and $\mathbf{C}^{-1}$ are the same move.) In practice, $\mathbf{B}$ and $\mathbf{C}$ look like the same move—they both involve exchanging two rows—but for the sake of proving various facts about the behavior of these moves, it will be convenient to have separate names for them.

Before giving the algorithm for computing α in full, explicit detail, we collect some useful lemmas about what happens to a diagram in $\tilde{D}_n$ when a certain move is performed on it under certain conditions. For $X \in \tilde{D}_n$, we write $\mathbf{A}X$, $\mathbf{B}X$, and $\mathbf{C}X$ to indicate the diagram resulting from performing moves $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ respectively. (The values of the parameters s, r, m, and $i_1, \ldots, i_m$ will be clear from context.) Finally, the symbol $\mathbf{M}$ will be used to stand for any of the moves $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, or their inverses.

Definition 3.2.3. $\mathbf{M}X$ is said to be *well-behaved of order* $\geq k$, where $1 \leq k \leq 6$, if there is some k', $k \leq k' \leq 6$, such that

$$q_l(\mathbf{M}X) = q_l(X) \qquad \text{for } l = 0, \ldots, k'-1$$

and

$$q_{k'}(\mathbf{M}X) < q_{k'}(X).$$

Lemma 3.2.4. *If $X \in \tilde{D}(\mathbf{d}, \tau)$, then $\mathbf{M}X \in \tilde{D}(\mathbf{d}, \tau)$ for any $\mathbf{M}$.*

Proof. Moves $\mathbf{A}$ and $\mathbf{A}^{-1}$ preserve the shape of the diagram, and on some rows they add $+1$ to one entry and -1 to another. Hence the row-sums of X are preserved. In the case of $\mathbf{B}$, $\mathbf{B}^{-1}$, and $\mathbf{C}$, no entries are changed; rather, we merely exchange two rows. In all cases, we see that $\kappa(\mathbf{M}X) = \kappa(X)$. □

Lemma 3.2.5. *Suppose that X has $\mathbf{p_4}(r)$ and that X_{ir} is raisable. If Y is a diagram obtained from X by raising X_{ir}, then EY is equal to EX raised at position ir.*

If X has $\mathbf{p_4}(r)$ and X_{ir} is lowerable, and if Y is X lowered at ir, then EY is EX lowered at ir.

Proof. We prove only the first part of the claim; the second part is proved similarly. E acts on each column of a diagram individually, so EX and EY will certainly agree in every column other than r. Now, in general, the amount E adds to the entries of a column depend only on the order of those entries. In particular, if a column is nonincreasing and of height h, then E adds $h-1$ to the first entry, $h-3$ to the second, etc., down to $-h+1$ to the last entry, irrespective of what those entries actually are.

Because X has $\mathbf{p_4}(r)$, column r is in fact nonincreasing. Moreover, to say X_{ir} is raisable is to say that it is strictly smaller than its column-predecessor, so it follows that column r of Y is still nonincreasing. E acts on both X and Y by adding the same numbers to corresponding entries, so EX and EY differ only where X and Y differ, with $EY_{ir} = EX_{ir} + 1$. □

Proposition 3.2.6. *Suppose X contains a sequence of rows $i_1, \ldots, i_m$ on which $\mathbf{A}$ might be performed, say at columns s and r, with $s < r$. Furthermore, suppose that $q_4(X) \leq -r$, $X_{i_m s}$ is lowerable, $X_{i_1 r}$ is raisable, and $EX_{i_k s} - EX_{i_k r} \geq 2$ for $k = 1, \ldots, m$. Then $\|\mathbf{A}X\|^2 < \|X\|^2$. Stated differently, $\mathbf{A}$ is well-behaved on X of order 1.*

Proof. We begin by noting that in any column t of X such that $\mathbf{p_4}(t)$ holds, if X_{jt} is raisable and has column-successor $X_{j't}$, then after raising X_{jt}, position $j't$ will be raisable in the new diagram: if Y is the diagram obtained by raising X_{jt}, then $X_{jt} \geq X_{j't}$ implies $Y_{jt} > Y_{j't}$.

Now, the performance of move $\mathbf{A}$ can be broken down into steps as follows:

1. Raise $X_{i_1 r}$; lower $X_{i_m s}$.

2. Raise $X_{i_2 r}$; lower $X_{i_{m-1} s}$.

$\vdots$

A	$\begin{array}{\|c\|c\|c\|c\|c\|c} \hline a_1 & \cdots & a_{s-1} & a_s & \cdots & a_r \\ \hline \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \hline a'_1 & \cdots & a'_{s-1} & a'_s & \cdots & a'_r \\ \hline \end{array} \mapsto \begin{array}{\|c\|c\|c\|c\|c\|c} \hline a_1 & \cdots & a_{s-1} & a_s - 1 & \cdots & a_r + 1 \\ \hline \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \hline a'_1 & \cdots & a'_{s-1} & a'_s - 1 & \cdots & a'_r + 1 \\ \hline \end{array}$
B	$\begin{array}{\|c\|c\|c\|c} \hline a_1 & \cdots & a_{r-1} & a_r \\ \hline a_1 & \cdots & a_{r-1} \\ \cline{1-3} \end{array} \mapsto \begin{array}{\|c\|c\|c\|c} \cline{1-3} a_1 & \cdots & a_{r-1} \\ \hline a_1 & \cdots & a_{r-1} & a_r \\ \hline \end{array}$
C	$\begin{array}{\|c\|c\|c\|c} \hline a_1 & \cdots & a_{r-1} & a_r \\ \hline a_1 & \cdots & a_{r-1} & a'_r \\ \hline \end{array} \mapsto \begin{array}{\|c\|c\|c\|c} \hline a_1 & \cdots & a_{r-1} & a'_r \\ \hline a_1 & \cdots & a_{r-1} & a_r \\ \hline \end{array}$

TABLE 3.2. Moves used in the algorithm for computing α

m. Raise $X_{i_m r}$; lower $X_{i_1 s}$.

At each step, the entries being raised and lowered are raisable and lowerable respectively (this is true in Step 1 by assumption, and in each successive step by what we noted above). It follows by Lemma 3.2.5 that $E\mathbf{A}X_{i_k r} = EX_{i_k r} + 1$ and $E\mathbf{A}X_{i_k s} = EX_{i_k s} - 1$ for each k. Therefore,

$$\begin{aligned} \|E\mathbf{A}X\|^2 &= \sum_{j,t} (E\mathbf{A}X_{jt})^2 \\ &= \sum_{j,t} (EX_{jt})^2 + \sum_{k=1}^{m} \big((E\mathbf{A}X_{i_k s})^2 + (E\mathbf{A}X_{i_k r})^2 - (EX_{i_k s})^2 - (EX_{i_k r})^2\big) \\ &= \|EX\|^2 + \sum_{k=1}^{m} ((-2EX_{i_k s} + 1) + (2EX_{i_k r} + 1)) \\ &= \|EX\|^2 - 2\sum_{k=1}^{m} (EX_{i_k s} - EX_{i_k r} - 1) \end{aligned} \tag{3.1}$$

Since $EX_{i_k s} - EX_{i_k r} \geq 2$, the summation in the last line is strictly positive. Thus $\|E\mathbf{A}X\|^2 < \|EX\|^2$. $\square$

PROPOSITION 3.2.7. *Assume the conditions of Proposition 3.2.6, but weaken the requirement on entries in* EX *to* $EX_{i_k s} - EX_{i_k r} \geq 1$. *If, in addition,* r *and* s *are both even, then* $\mathbf{A}$ *is well-behaved on* X *of order* ≥ 1.

Proof. We perform the same computation as in the proof of Proposition 3.2.6, and we arrive at equation (3.1). With the weakened inequality in the present assertion, some of the terms in the summation in (3.1) may be 0, but provided that at least one of them is positive, we still get $\|E\mathbf{A}X\|^2 < \|EX\|^2$.

If, however, $EX_{i_k s} - EX_{i_k r} = 1$ for every k, then that summation will be 0, so we shall have $q_1(\mathbf{A}X) = q_1(X)$. We know by Lemma 3.2.5 (as argued in the proof of Proposition 3.2.6) that after move $\mathbf{A}$, m entries in column s of EX are changed by -1, and m entries in column r are changed by $+1$. Thus, the total sum of elements in even-numbered columns is unchanged: $q_2(\mathbf{A}X) = q_2(X)$. Furthermore, we easily compute that q_3 changes by $s - r$, so $q_3(\mathbf{A}X) < q_3(X)$. Hence $\mathbf{A}$ is well-behaved on X, as claimed. $\square$

PROPOSITION 3.2.8. *Suppose that move* $\mathbf{A}$ *might be performed on the single row* i. *Suppose in addition that* $q_4(X) = -(r-1)$ *and that* $q_5(X) \neq 0$*; in particular, suppose that* $\tilde{q}_{5;ir}(X) > 0$.

Furthermore, suppose that $EX_{i,r-1} > EX_{ir}$, and that $X_{jr} = X_{ir}$ implies $j \geq i$. If $s < r$ and EX_{is} is lowerable, and

(a) *if $EX_{is} - EX_{ir} \geq 2$, then $q_1(\mathbf{A}X) < q_1(X)$.*

(b) *if r is odd, s is even, and $EX_{is} - EX_{ir} = 1$, then $q_1(\mathbf{A}X) = q_1(X)$, but $q_2(\mathbf{A}X) < q_2(X)$.*

(c) *if r and s are both even, and $EX_{is} - EX_{ir} = 1$, then $q_1(\mathbf{A}X) = q_1(X)$ and $q_2(\mathbf{A}X) = q_2(\mathbf{A}X)$, but $q_3(\mathbf{A}X) < q_3(X)$.*

Proof. By the assumption that $q_4(X) = -(r-1)$, we know that $\mathbf{p_4}(s)$ holds, so Lemma 3.2.5 applies to column s. We do not know enough about column r to say how E acts on it precisely, but the complicated assumptions in the proposition are specifically designed to enable us to compute the difference in size of $E\mathbf{A}X$ and EX regardless.

Let $\|EX_{*r}\|^2$ denote the size of the r-th column of EX. This quantity depends on what the entries in the r-th column are, but it is independent of their order within the column. If the column contains h entries, of which l are strictly greater than X_{ir}, it is easy to check that replacing X_{ir} by $X_{ir} + 1$ results $\|EX_{*r}\|^2$ changing by $(X_{ir} + 1 + (h-1-2l))^2 - (X_{ir} + (h-1-2l))^2$, or

$$2\,(X_{ir} + (h-1-2l)) + 1. \tag{3.2}$$

By Lemma 3.2.5, lowering X_{is} changes $\|EX_{*s}\|^2$ by $(EX_{is} - 1)^2 - (EX_{is})^2$, or

$$-2EX_{is} + 1. \tag{3.3}$$

The net change in $\|EX\|^2$ brought about by performing move $\mathbf{A}$ on X would be the sum of (3.2) and (3.3). Now, recalling the definition of E, we can compute EX_{ir} explicitly, thanks to the assumption that $X_{jr} = X_{ir}$ implies $j \geq i$:

$$\begin{aligned} EX_{ir} &= X_{ir} + \#\{X_{jr} \mid X_{jr} < X_{ir}, \text{ or } X_{jr} = X_{ir} \text{ with } j > i\} \\ &\qquad - \#\{X_{jr} \mid X_{jr} > X_{ir}, \text{ or } X_{jr} = X_{ir} \text{ with } j < i\} \\ &= X_{ir} + (h - (l+1)) - l. \end{aligned}$$

Thus, expression (3.2) is equal to $2EX_{ir} + 1$. We obtain that the sum of (3.2) and (3.3) is

$$q_1(\mathbf{A}X) - q_1(X) = 2 - 2(EX_{is} - EX_{ir}). \tag{3.4}$$

Part (a) of the proposition follows because if $EX_{is} - EX_{ir} \geq 2$, the quantity in (3.4) is strictly negative. On the other hand, if $EX_{is} - EX_{ir} = 1$, the quantity in (3.4) is 0. If r is odd and s even, lowering X_{is} obviously decreases the sum q_2 of even-column entries, so part (b) is true as well. Finally, if r and s are both odd, q_2 remains unchanged, but q_3 changes by $m(s-r)$: this establishes part (c). $\square$

Proposition 3.2.9. *Suppose rows i and i' of X are such that $\mathbf{B}$ or $\mathbf{C}$ might be performed on them: in particular, they agree in their first $r-1$ entries, and intervening rows have length less than $r-1$. Suppose furthermore that $q_4(X) = -(r-1)$, and that row i has length at least r.*

(a) *Suppose X has no entry at position $i'r$. If $\tilde{q}_{5;ir} > 0$ and $EX_{ir} < EX_{i,r-1}$, $\mathbf{B}X$ is well-behaved of order ≥ 4.*

(b) *Suppose that X does have an entry at position $i'r$. If $X_{ir} < X_{i'r}$, then $\mathbf{C}X$ is well-behaved of order ≥ 4.*

Proof. Both of these moves change the shape of the diagram without changing any entries, so q_1 is preserved. Moreover, this shape change is brought about by exchanging rows; the entries in any given column remain the same, albeit possibly rearranged. Hence q_2 and q_3 are preserved as well.

Before proceeding, we define two convenient functions:

$$w_{5;r}(X) = \sum_j \tilde{q}_{5;jr}(X) \qquad \text{and} \qquad w_{6;r}(X) = \sum_j \tilde{q}_{6;jr}(X)$$

Provided that $q_4(X) = -(r-1)$, we will of course have $w_{5;r} = q_5$ and $w_{6;r} = q_6$. Let us also note that $EX_{i',r-1} = EX_{i,r-1} - 2$.

For part (a) of the proposition, $EX_{i,r-1} - EX_{ir}$ is greater than or equal to 1 (if r is odd) or 2 (if r is even). Move $\mathbf{B}$ does not change the relative position of entry X_{ir} in column r, so $E\mathbf{B}X_{i'r} = EX_{ir}$. We have

$$E\mathbf{B}X_{i',r-1} - E\mathbf{B}X_{i'r} = EX_{i,r-1} - EX_{ir} - 2;$$

we can conclude that

$$\tilde{q}_{5;i'r}(\mathbf{B}X) = \max\{0, \tilde{q}_{5;ir}(X) - 2\}$$

and

$$w_{5;r}(\mathbf{B}X) = \max\{0, w_{5;r}(X) - 2\}.$$

In particular, $w_{5;r}(\mathbf{B}X) < w_{5;r}(X)$. This may mean that $q_4(\mathbf{B}X) = q_4(X)$ and $q_5(\mathbf{B}X) < q_5(X)$, or, if $w_{5;r}(\mathbf{B}X) = 0$ (meaning that $\mathbf{B}X$ has $\mathbf{p_3}(r)$, which X did not), it may be that $q_4(\mathbf{B}X) = -r$ and that $q_5(\mathbf{B}X)$ is unpredictable. In either case, we observe that $\mathbf{B}X$ is well-behaved of order ≥ 4.

For part (b), we break down the argument into three cases:

Case 1. $\tilde{q}_{5;ir}(X) > 0$, and $EX_{ir} > EX_{i,r-1}$. In this case, we must have $EX_{i'r} > EX_{ir} > EX_{i,r-1}$; it is easy to check that

$$\begin{aligned} \tilde{q}_{5;ir}(\mathbf{C}X) &= \tilde{q}_{5;i'r}(X) - 2 \\ \tilde{q}_{5;i'r}(\mathbf{C}X) &= \tilde{q}_{5;ir}(X) + 2. \end{aligned}$$

This equations imply that the following holds (of course, we actually have equality, but we write an inequality to accomodate the two cases considered below):

$$\tilde{q}_{5;ir}(\mathbf{C}X) + \tilde{q}_{5;i'r}(\mathbf{C}X) \leq \tilde{q}_{5;ir}(X) + \tilde{q}_{5;i'r}(X). \tag{3.5}$$

Case 2. $\tilde{q}_{5;ir}(X) = 0$. This time we have

$$\begin{aligned} \tilde{q}_{5;ir}(\mathbf{C}X) &= \tilde{q}_{5;i'r}(X) - 2 \\ \tilde{q}_{5;i'r}(\mathbf{C}X) &= 1 \text{ or } 2. \end{aligned}$$

These facts imply that (3.5) holds here as well.

Case 3. $\tilde{q}_{5;ir}(X) > 0$, and $EX_{ir} < EX_{i,r-1}$. This is the most complicated of the three cases; part of the computation has to be broken down into three sub-cases. We obtain:

$$\begin{aligned} \tilde{q}_{5;ir}(\mathbf{C}X) &= \begin{cases} \max\{0, \tilde{q}_{5;i'r}(X) - 2\} & \text{if } \tilde{q}_{5;i'r} > 0 \text{ and } EX_{i'r} > EX_{i',r-1} \\ 1 \text{ or } 2 & \text{if } \tilde{q}_{5;i'r} = 0 \\ \tilde{q}_{5;i'r} + 2 & \text{if } \tilde{q}_{5;i'r} > 0 \text{ and } EX_{i'r} < EX_{i',r-1} \end{cases} \\ \tilde{q}_{5;i'r}(\mathbf{C}X) &= \max\{0, \tilde{q}_{5;ir}(X) - 2\} \end{aligned}$$

Once again, (3.5) holds.

We finish up the argument almost as we did for part (a), but this time we have only the weaker inequality $w_{5;r}(\mathbf{C}X) \leq w_{5;r}(X)$ following from (3.5). This gives rise to the possibility that $q_4(\mathbf{C}X) = q_4(X)$ and $q_5(\mathbf{C}X) = q_5(X)$, compelling us to examine the behavior of q_6. But it is clear that move $\mathbf{C}$ brings column r closer to being nonincreasing, in the sense that $w_{6;r}(\mathbf{C}X) < w_{6;r}(X)$. If q_4 and q_5 do not change under move $\mathbf{C}$, then $q_6(\mathbf{C}X) < q_6(X)$. Hence $\mathbf{C}X$ is well-behaved of order ≥ 4. $\square$

Conditions	*Move*	*Well-Behavedness*
$q_4(X) \leq -r$.		Order ≥ 1.
$X_{i_m s}$ is lowerable, $X_{i_1 r}$ is raisable, and $EX_{i_k s} - EX_{i_k r} \geq 1$ for $k = 1, \ldots, m$.	$\mathbf{A}$	
$X_{i_1 s}$ is raisable, $X_{i_m r}$ is lowerable, and $EX_{i_k s} - EX_{i_k r} \leq -1$ for $k = 1, \ldots, m$.	$\mathbf{A}^{-1}$	
$m = 1$; $q_4(X) = -(r-1)$. $X_{i_1 s}$ is lowerable. $\tilde{q}_{3;i_1 r}(X) \neq 0$. If $X_{jr} = X_{ir}$, then $j \geq i$ (for $\mathbf{A}$) or $j \leq i$ (for $\mathbf{A}^{-1}$).		Order $= 1$.
$EX_{i_1 s} - EX_{i_1 r} \geq 2$.	$\mathbf{A}$	
$EX_{i_1 s} - EX_{i_1 r} \leq -2$.	$\mathbf{A}^{-1}$	
r odd, s even, and $EX_{i_1 s} - EX_{i_1 r} = 1$.	$\mathbf{A}$	
r even, s odd, and $EX_{i_1 s} - EX_{i_1 r} = -1$.	$\mathbf{A}^{-1}$	
$q_4(X) = -(r-1)$; $\tilde{q}_{5;i_1 r}(X) > 0$		Order ≥ 4.
$EX_{i_1, r-1} > EX_{i_1 r}$.	$\mathbf{B}$	
$EX_{i_1, r-1} < EX_{i_1 r}$.	$\mathbf{B}^{-1}$	
$X_{ir} < X_{i'r}$.	$\mathbf{C}$	

TABLE 3.3. Well-behavedness of moves under various hypotheses

The facts in Propositions 3.2.6, 3.2.8, and 3.2.9 are collected and summarized in Table 3.3.

CHAPTER 4

Construction of the Algorithms

At this stage, we are ready to put weight diagrams to work for us. Recall, from Claim 2.3.1, what it is that we are trying to do with these diagrams: we are trying to find a weight λ of $L_{\mathbf{d}}$ with a given restriction to $G_{\mathrm{red}}^{\mathbf{d}}$, such that $\|E\lambda\|^2$ is minimized. In the first section, we describe the algorithm α that does this, and we do most of the work of proving that $\|E\lambda\|^2$ is minimized, although we will not actually finish the proof until Chapter 5. We define a particular, highly constrained set of diagrams $\tilde{D}^\circ(\mathbf{d}, \tau)$; and we show that α always produces diagrams contained in this set.

In the second part, we go in the reverse direction: given a weight σ of G, we need to produce a partition $\mathbf{d}$ and a weight λ of $L_{\mathbf{d}}$ such that $E\lambda = \sigma$, and such that λ has a minimal value of $\|E\lambda\|^2$ among $L_{\mathbf{d}}$-weights having the same restriction to $G_{\mathrm{red}}^{\mathbf{d}}$. For this we give another algorithm, β, which constructs a diagram X whose entries are the coordinates of σ and which has $E^{-1}X \in \tilde{D}^\circ(\mathbf{d}, \tau)$ for some $\mathbf{d}$, τ. Again, the proof that β has the required properties will not be completed until the next chapter.

Examples of the use of both of these algorithms are given in Appendix A.

4.1 The Algorithm for α

We define $\alpha : \tilde{D}_n \to \tilde{D}_n$ to be the map computed by the following algorithm: Starting with $X \in \tilde{D}_n$, look down the leftmost column in Table 3.3 and find a hypothesis satisfied by X (for some choice of s, r, and $i_1, \ldots, i_m$), and then perform the corresponding move given in the middle column to obtain a new diagram X'. Repeat to obtain X'', etc., until some $X^{(p)}$ does not satisfy any of the conditions in the left column. Then $\alpha(X) = X^{(p)}$.

There is some ambiguity if X satisfies more than one of the hypotheses, or if it satisfies some hypothesis for more than one choice of s, r, and $i_1, \ldots, i_m$. Let us make it definite by deciding always to choose the *first* satisfied hypothesis (in the order in which they are listed in the table), together with the *lowest* choice of vector $(s, r, m, i_1, \ldots, i_m)$ (where these vectors are ordered lexicographically).

For this definition to make sense, we need the following fact.

PROPOSITION 4.1.1. *The algorithm for computing α terminates after a finite number of steps.*

Proof. The assertion is a consequence of the fact that only well-behaved moves are performed while computing α. Given a diagram X, let c be the number of columns in X, and consider the map

$$q : \tilde{D}_n \to \mathbb{N}^6$$

defined by

$$q(X) = (q_1(X), q_2(X), q_3(X), c + q_4(X), q_5(X), q_6(X)).$$

(Here, the c in the fourth coordinate appears only to make that coordinate have a nonnegative value.) Give $\mathbb{N}^6$ the lexicographical ordering; then, to say that $\mathbf{M}X$ is well-behaved is to say that $q(\mathbf{M}X) < q(X)$. Furthermore, we evidently have

$$q(X) \geq (0, 0, 0, 0, 0, 0).$$

Each move performed while computing α decreases q, and q is bounded below. Since $\mathbb{N}^6$ is well-ordered, it follows that the algorithm must stop after a finite number of steps. $\square$

Lemma 4.1.2. *Suppose X has $\mathbf{P_1}(r)$, $\mathbf{P_3}(r)$, and $\mathbf{P_4}(r)$. If a range of entries $X_{is}, \ldots, X_{ir}$ in a single row is such that none of them is lowerable (resp. raisable), then all their column-successors (resp. column-predecessors) lie in a single row. Moreover, if $s > 1$, then the column-successor (resp. column-predecessor) of $X_{i,s-1}$ lies in that same row as well.*

Proof. We prove the statement in the case that none of the entries in the range is lowerable; the other case is proved similarly. If $s = r$, the statement is trivial. Assume $s < r$, and pick any t such that $s \leq t < t+1 \leq r$. Suppose the column-successor of X_{it} is on row j_1, and that of $X_{i,t+1}$ on row j_2. There are three cases to consider:

Case 1. $EX_{it} = EX_{i,t+1}$. Hence $EX_{j_1 t} = EX_{j_2,t+1}$. But if $j_1 \neq j_2$, then $EX_{j_2,t+1}$ differs from $EX_{j_2 t}$ by at most 1, and $EX_{j_2 t}$ in turn differs from $EX_{j_1 t}$ by at least 2, so, $EX_{j_1 t}$ and $EX_{j_2,t+1}$ must differ by at least 1. Hence $j_1 = j_2$.

Case 2. The column index t is odd, and $EX_{it} = EX_{i,t+1} + 1$. Hence $EX_{j_1 t} = EX_{j_2,t+1} + 1$. Now, $EX_{j_2 t}$ must equal either $EX_{j_2,t+1}$ or $EX_{j_2,t+1} + 1$, but neither of these values differs from $EX_{j_1 t}$ by 2 or more, so it must be that $j_1 = j_2$.

Case 3. The column index t is even, and $EX_{it} = EX_{i,t+1} - 1$. This case is similar to the preceding one.

We see that any two neighboring entries in the range $X_{is}, \ldots, X_{ir}$ must have column-successors on the same row, so the lemma follows. □

Proposition 4.1.3. *If X has $\mathbf{p_3}(r)$ and $\mathbf{p_4}(r-1)$, then it also has $\mathbf{p_4}(r)$.*

Proof. Let ir, $i'r$ be column-consecutive positions in column r. Since $i < i'$ and $\mathbf{p_4}(r-1)$ holds, $X_{i,r-1} \geq X_{i',r-1}$. This in turn implies $EX_{i,r-1} - EX_{i',r-1} \geq 2$. By $\mathbf{p_3}(r)$, we know

$$|EX_{ir} - EX_{i,r-1}| \leq 1$$
$$|EX_{i'r} - EX_{i',r-1}| \leq 1.$$

It follows from these inequalities that $EX_{ir} - EX_{i'r} \geq 0$. But no two entries in a single column of EX can be closer than 2: we conclude that $EX_{ir} - EX_{i'r} \geq 2$. From this last inequality we conclude that $X_{ir} - X_{i'r} \geq 0$; hence, column r is nonincreasing, and $\mathbf{p_4}(r)$ holds. □

Corollary 4.1.4. *Suppose X has $\mathbf{P_1}(r-1)$, $\mathbf{P_2}(r-1)$, $\mathbf{P_3}(r-1)$, and $\mathbf{P_4}(r-1)$. If it does not have $\mathbf{P_4}(r)$, then it also does not have $\mathbf{P_3}(r)$.* □

Proposition 4.1.5. *Suppose X has $\mathbf{P_1}(r-1)$, $\mathbf{P_2}(r-1)$, $\mathbf{P_3}(r-1)$, and $\mathbf{P_4}(r-1)$. If it does not have $\mathbf{P_3}(r)$, then it satisfies some hypothesis in the left-hand column of Table 3.3.*

Proof. We know that $q_5(X) > 0$; find a row i such that $\tilde{q}_{5;i_0 r}(X)$ is maximal. We assume henceforth that $EX_{i_0,r-1} > EX_{i_0 r}$; the argument would proceed similarly if the inequality were reversed. Let $i \leq i_0$ be the index of the uppermost row whose entry in column r equals $X_{i_0 r}$. In other words, row i has the property that $X_{jr} = X_{ir}$ implies $j \geq i$. It is easy to check that this setup implies

$$EX_{i,r-1} - EX_{ir} \geq EX_{i_0,r-1} - EX_{i_0 r},$$

and by maximality, the preceeding line must actually be an equality. The remainder of the argument is a boring and complicated case-by-case analysis. To assist the reader in staying awake while reading it, we give here a road map of the breakdown into cases, although we do not yet define all the symbols contained herein.

1. r is odd.
   (a) $X_{i',r-1}$ is lowerable.
   (b) $X_{i',r-1}$ is not lowerable.

i. $s = 0$.

ii. $s > 0$.

A. s is even.

B. $s = r - 2$.

C. s is odd and less then $r - 2$.

2. r is even.

First, we consider the case where r is odd. Let $i' \geq i$ be the last row such that the consecutive rows $i, i+1, \ldots, i'$ all agree in the first r columns.

If $X_{i',r-1}$ is lowerable, we are done: Proposition 3.2.8 applies to rows $i, \ldots, i'$, with $s = r - 1$. (In the event that $EX_{i,r-1} - EX_{ir} = 1$, we meet the additional requirement that s be even.)

If $X_{i',r-1}$ is not lowerable, let $X_{i'',r-1}$ be its column-successor. Let s be such that rows i' and i'' of X agree in columns $s+1, s+2, \ldots, r-1$. Using Lemma 3.2.5, we see that either $s = 0$, or X_{is} ($s > 0$) is lowerable.

Suppose $s = 0$. If $X_{i''r}$ is empty, we can apply Proposition 3.2.9a and do move **B** on rows i' and i''. If $X_{i''r}$ is not empty, we must have $X_{i''r} \geq X_{ir}$ in order not to violate the maximality of $\tilde{q}_{5;ir}(X)$; but on the other hand, we cannot have $X_{i''r} = X_{ir}$, since i' is the last row to agree with row i in each of the first r columns. Hence $X_{i''r} > X_{i'r}$, and we can do move **C** by Proposition 3.2.9b.

On the other hand, if $s > 0$, we know by Lemma 3.2.5 that $X_{i''s}$ is still the column-successor of $X_{i's}$, and that $EX_{i's} - EX_{i''s} = 3$. Since $X_{i''s}$ is raisable, property $\mathbf{P_2}(r-1)$ tells us the first of the following inequalities, from which we derive the latter ones:

$$\begin{aligned} EX_{i''s} &\geq EX_{i'',r-1} - 1 \\ EX_{i''s} + 3 &\geq EX_{i'',r-1} + 2 \\ EX_{i's} &\geq EX_{i',r-1} \\ EX_{i's} &> EX_{i'r} \end{aligned} \tag{4.1}$$

Indeed, $EX_{js} > EX_{jr}$ for every $j = i, \ldots, i'$. Now we can almost apply Proposition 3.2.8 to rows $i, \ldots, i'$ with our specified s and r. If $EX_{i,r-1} - EX_{ir} \geq 2$, then $EX_{is} - EX_{ir} \geq 2$ as well, and the proposition applies. But if that difference is only 1, we need to make sure that s is even. We shall show instead that if s is odd, then $EX_{is} - EX_{ir}$ is necessarily at least 2. Observe that by property $\mathbf{P_3}(r-1)$, we must have $EX_{i's} = EX_{i',s+1} + 1$, and $EX_{i''s} = EX_{i'',s+1}$.

If $s = r - 2$, then we are done: since $EX_{i',r-1} - EX_{i'r} \geq 1$ and $EX_{i's} = EX_{i',r-1} + 1$, we conclude $EX_{i's} - EX_{i'r} \geq 2$, and hence $EX_{is} - EX_{ir} \geq 2$, as desired.

If $s < r - 2$, then because $X_{i''s}$ is raisable, property $\mathbf{P_1}(r-1)$ tells us that $EX_{i''s} \geq EX_{i'',r-2}$. In turn, property $\mathbf{P_3}(r-1)$ tells us that $EX_{i'',r-2}$ is at least as large as $EX_{i'',r-1} = EX_{i',r-1} - 2$. We compute:

$$\begin{aligned} EX_{i''s} &\geq EX_{i',r-1} - 2 \\ EX_{i''s} + 3 &\geq EX_{i',r-1} + 1 \\ EX_{i's} &> EX_{i',r-1} \end{aligned}$$

This last strict inequality implies that $EX_{i's} - EX_{i'r} \geq 2$, so once again, $EX_{is} - EX_{ir} \geq 2$, as desired.

What happens if r is even? We can repeat the above arguments until the final few steps, where we worried about the possibility that $EX_{is} - EX_{ir} = 1$. But that worry is irrelevant when r is even: now $\tilde{q}_{5;ir} > 0$ means that $EX_{i,r-1} - EX_{ir} \geq 2$, so inequality (4.1) implies directly that $EX_{is} - EX_{ir} \geq 2$ as well. □

Proposition 4.1.6. *Suppose X has $\mathbf{P_1}(r-1)$ and $\mathbf{P_2}(r-1)$, as well as $\mathbf{P_3}(r)$ and $\mathbf{P_4}(r)$. If it does not have $\mathbf{P_2}(r)$, then it satisfies some hypothesis in the left-hand column of Table 3.3.*

Proof. We assume without loss of generality that there are some positions $i's$, $i'r$ in X such that $EX_{i's} - EX_{i'r} \geq 2$, but such that $EX_{i's}$ is lowerable. (The argument is similar if the inequality is reversed and $EX_{i's}$ is raisable.) Let $i_1, i_2, \ldots, i_m = i'$ be a sequence of row indices that are column-consecutive in column r, and such that $X_{i_1 r}$ is raisable. Then Proposition 3.2.7 applies, and we can do move $\mathbf{A}$. □

PROPOSITION 4.1.7. *Suppose X has $\mathbf{P_1}(r-1)$, as well as $\mathbf{P_2}(r)$, $\mathbf{P_3}(r)$, and $\mathbf{P_4}(r)$. If it does not have $\mathbf{P_1}(r)$, then it satisfies some hypothesis in the left-hand column of Table 3.3.*

Proof. The argument is identical to that in the proof of Proposition 4.1.6, except that in the first sentence, the inequality is replaced by "$EX_{i's} - EX_{i'r} \geq 1$," and in the last sentence, we apply Proposition 3.2.6. □

Let $\mathbf{d}$ be a partition of n, and τ a representation of $G^{\mathbf{d}}_{\mathrm{red}}$. We introduce the follow two new symbols:

$$\tilde{D}_n^\circ = \{X \in \tilde{D}_n \mid X \text{ has } \mathbf{P_1}(r), \mathbf{P_2}(r), \mathbf{P_3}(r), \text{ and } \mathbf{P_4}(r) \text{ for all } r\}$$
$$\tilde{D}^\circ(\mathbf{d}, \tau) = \tilde{D}(\mathbf{d}, \tau) \cap \tilde{D}_n^\circ$$

THEOREM 4.1.8. *The image of α is contained in $\tilde{D}_n^\circ$. If $X \in \tilde{D}(\mathbf{d}, \tau)$, then $\alpha(X) \in \tilde{D}^\circ(\mathbf{d}, \tau)$. Furthermore, $\|E\alpha(X)\|^2 \leq \|EX\|^2$ for all X.*

Proof. The map α is computed by repeatedly doing moves from Table 3.3, until none is possible. Collecting the results expressed in Corollary 4.1.4 and Propositions 4.1.5, 4.1.6, and 4.1.7, we see that a move is always possible for any diagram not lying in $\tilde{D}_n^\circ$. Hence the image of α lies in $\tilde{D}_n^\circ$. Finally, recall the function $q : \tilde{D}_n \to \mathbb{N}^6$ defined in the proof of Proposition 4.1.1. The last part of the assertion is true because $q(\alpha(X)) \leq q(X)$ for any X, and this implies that $q_1(\alpha(X)) \leq q_1(X)$. □

The size-reducing property of α means that if we start with a diagram $X \in \tilde{D}(\mathbf{d}, \tau)$ such that $\|EX\|^2$ is minimal, then $\alpha(X)$ must also have minimal size. To say this symbolically, we introduce the following notation:

$$\tilde{D}(\mathbf{d}, \tau)_{\min} = \{X \in \tilde{D}(\mathbf{d}, \tau) \mid \|EX\|^2 \text{ is minimal among } \|EY\|^2 \text{ for } Y \in \tilde{D}(\mathbf{d}, \tau)\}$$
$$\tilde{D}^\circ(\mathbf{d}, \tau)_{\min} = \tilde{D}(\mathbf{d}, \tau)_{\min} \cap \tilde{D}^\circ(\mathbf{d}, \tau)$$

COROLLARY 4.1.9. *$\tilde{D}^\circ(\mathbf{d}, \tau)_{\min}$ is nonempty, and $\alpha(\tilde{D}(\mathbf{d}, \tau)_{\min}) \subseteq \tilde{D}^\circ(\mathbf{d}, \tau)_{\min}$. Furthermore, if $X \in \tilde{D}(\mathbf{d}, \tau)_{\min}$, then $E\alpha(X)$ and EX have the same integer entries.*

Proof. The first statement follows immediately from the preceding theorem and the definitions of $\tilde{D}(\mathbf{d}, \tau)_{\min}$ and $\tilde{D}^\circ(\mathbf{d}, \tau)_{\min}$. For the second statement, if $X \in \tilde{D}(\mathbf{d}, \tau)_{\min}$, then $\|E\alpha X\|^2 = \|EX\|^2$; moreover, every intermediate diagram Y obtained in the process of computing α must also have that $\|EY\|^2 = \|EX\|^2$. Let us revisit Table 3.3. What moves can be made that do not decrease the size of the diagram? Moves $\mathbf{B}$, $\mathbf{B}^{-1}$, and $\mathbf{C}$ do not decrease size; furthermore, they operate by rearranging entries in the diagram, without actually changing them. $\mathbf{A}$ and $\mathbf{A}^{-1}$ also preserve size under certain circumstances. Specifically, if these moves are applied to only one row, say row i, with $|EX_{ir} - EX_{is}| = 1$, then they are well-behaved of order 1. Although $\mathbf{A}$ and $\mathbf{A}^{-1}$ nominally alter entries in the diagram by ± 1, the fact that $|EX_{ir} - EX_{is}| = 1$ means that $\mathbf{A}$ or $\mathbf{A}^{-1}$ effectively just exchanges entries EX_{ir} and EX_{is}. Thus, any move which preserves the size of EX acts by rearranging entries in the diagram without altering them. In other words, $E\alpha X$ has the same integer entries as EX. □

4.2 The Algorithm for β

In this section we present an algorithm for inverting α. We begin with a few basic constructions dealing with finite collections of integers, in which a given integer is allowed to occur with multiplicity greater than 1. Later on, we shall sometimes use these constructions where we take the collections to be the coordinates of a weight of $GL(n)$, and sometimes where we take them to be the entries of a diagram in $\tilde{D}_n$. Henceforth, all collections of integers are assumed to be finite and to allow elements with multiplicities greater than 1.

Definition 4.2.1. The *length* of a collection of integers is the number of distinct elements it contains. This is to be distinguished from *size*, the total number of elements.

Definition 4.2.2. A collection of integers is called a *clump* if either of the following conditions holds:

i. The collection has length 1.

ii. The collection has length greater than 1, and for each member of the collection, there is another member differing from it by exactly 1.

Lemma 4.2.3. *Any collection of integers can be written uniquely as a disjoint union of clumps of maximal size.* □

The algorithm for computing $\beta : \Lambda(G) \to \tilde{D}_n$ is given below. We start with some weight $\sigma \in \Lambda(G)$.

1. Let $r = 1$, and let σ_r be the collection of integers which are coordinates of σ. Write $\sigma_r = A_{1,l_1} \amalg \cdots \amalg A_{1,l}$, where the $A_{1,i}$ are clumps of maximal size. We are about to start building the first column; below, "r" always refers to the column on which we are currently working.

2. Write down the distinct values of each clump in decreasing order. Form a set Z_r of (distinct) integers as follows: for each clump of odd length, we include in Z_r the 1st, 3rd, 5th, etc., distinct values of the clump. For each clump of even length, similarly take the odd-index distinct values if r is odd, but if r is even, take the 2nd, 4th, etc., values of each even-length clump. (Another way to think of this is that we always take alternate values from each clump; if r is odd, we must include the largest value in each clump, but if r is even, we must include the smallest.)

3. Write down the elements of Z_r vertically and in decreasing order. If $r = 1$, we are done—what we have just written down will be the first column of the diagram. Otherwise, we need to worry about what row each entry in the column belongs to. Place each element x of Z_r such that it is adjacent to a spot in column $r-1$ containing either x or $x + (-1)^r$ (this can always be done uniquely).

4. Let σ_{r+1} denote the collection obtained by removing the members of Z_r from σ_r. If σ_{r+1} is empty, we are finished drawing the diagram; otherwise, partition σ_{r+1} into disjoint maximal-size clumps $A_{r+1,1} \amalg \cdots \amalg A_{r+1,l_r}$. Advance the value of r by 1, and go to step 2.

There is something to be proved to ensure that Step 3 makes sense: namely, that each element of Z_r can be placed in a unique position next to some entry in column $r-1$ such that a certain adjacency condition is met. To that end, we make some observations.

First, note that when σ_r is divided into maximal clumps, any two entries differing by only 0 or 1 must end up in the same clump. In other words, entries in different clumps differ by at least 2. Since Z_r is constructed by taking alternate entries in every clump, it follows that any two elements of Z_r must differ by at least 2.

Let $x \in Z_r$, and suppose that we are working on Step 3. We are looking for an entry in the preceding column whose value is either x or $x+(-1)^r$. Since distinct entries in the preceding column differ from one another by at least 2, we know that at most one of x and $x+(-1)^r$ can occur in that column, and if one of them does occur, it occurs only once. In other words, if there is an appropriate entry in the column $r-1$ that meets our adjacency condition, it is unique.

Now we establish that either x or $x+(-1)^r$ actually occurs in the preceding column. Specifically, we assume that x does not occur, and show that $x+(-1)^r$ must occur. Let us rewind the algorithm to the point where we were building column $r-1$. There was some clump $A_{r-1,i}$ to which our element x belonged. Column $r-1$ includes alternate entries from clump $A_{r-1,i}$, so if it does not contain x, it must contain one or both of $x+1$ and $x-1$. If both of these occur, we are finished: one of them is $x+(-1)^r$.

But what if only one of $x \pm 1$ occurs? Suppose that only $x-1$ occurs. This can only happen if $A_{r-1,i}$ has no members equal to $x+1$; since $A_{r-1,i}$ is a clump, it must be that x is the largest value occurring in $A_{r-1,i}$. In constructing Z_{r-1}, we had not taken the largest (i.e first) value in $A_{r-1,i}$, but we had taken the second value. This is only done when the clump has even length and the column index is even. That is, $r-1$ is even, so $x+(-1)^r = x-1$. Thus $x+(-1)^r$ occurs in column $r-1$.

A similar argument shows that if $x+1$ occurs in the preceding column, but $x-1$ does not, then $r-1$ is odd, and again $x+(-1)^r = x+1$ occurs in column $r-1$.

We have established that Step 3 of the algorithm for β makes sense. In the course of establishing this, we observed that any two entries in a single column of the resulting diagram differ by at least 2. This property can be rephrased as the following statement:

PROPOSITION 4.2.4. *For any $\sigma \in \Lambda(G)$, the diagram $\beta(\sigma)$ lies in $E(\tilde{D}_n)$.* □

CHAPTER 5

Proof of the Main Theorem

In this chapter, we establish Theorem 3. In the first section, we show that β gives a bijection between dominant weights of $GL(n, \mathbb{C})$ and a certain set of weight diagrams. At that stage, we are only a stone's throw away from having a bijection between $\Lambda_+(G)$ and the set of pairs $\{(\mathbf{d}, \tau)\}$. However, it still takes some effort to show that the bijection we are computing coincides with the map γ as defined in Chapter 2. The second section is devoted to this latter goal. The map α plays an important role in this section. By the end the chapter, we shall have an explicit description of how to compute γ and its inverse using our algorithms.

5.1 Establishing the Bijection

Our goal in this section is to show that $E^{-1} \circ \beta$ gives a bijection between $\Lambda_+(G)$ and $\tilde{D}_n^\circ$.

Remark 5.1.1. Note that because adjacent entries in a given row of $\beta(\sigma)$ differ by at most 1, those entries must have belonged to the same clump when we first divided σ into clumps. In other words, clumps of σ are unions of rows in $\beta(\sigma)$. Moreover, the rows that constitute a given clump of σ must be consecutive.

Proposition 5.1.2. *For any $\sigma \in \Lambda_+(G)$, $E^{-1}\beta(\sigma) \in \tilde{D}_n^\circ$.*

Proof. Let $X = E^{-1}\beta(\sigma)$, so that $EX = \beta(\sigma)$. $\mathbf{P_3}(r)$ holds for all r because this condition about differences of adjacent entries is precisely that imposed in Step 3 of the algorithm. $\mathbf{P_4}(r)$ holds for all r because in Step 3 the columns of EX were constructed to be in decreasing order.

$\mathbf{P_1}(r)$ and $\mathbf{P_2}(r)$ are both consequences of Remark 5.1.1. Suppose, for instance, that in violation of $\mathbf{P_1}(r)$, we have $s < r$, s and r both odd, $EX_{is} < EX_{ir}$, and X_{is} raisable. If $X_{i's}$ is the column-predecessor of X_{is}, then $X_{i's} \geq X_{is} + 1$, so $EX_{i's} \geq EX_{is} + 3$. That large a difference between column-consecutive entries means that $EX_{i's}$ and EX_{is} must have come from different clumps of σ_s. We finished taking values from, say, clump $A_{s,j}$ at row i', and EX_{is} contains the first value taken from $A_{s,j+1}$. Since s is odd, EX_{is} should be the largest value in $A_{s,j+1}$, but on the other hand, EX_{ir} is a value from the same clump, but it is larger than EX_{is}, so we have a contradiction.

A similar argument establishes that $\mathbf{P_1}(r)$ holds if r is even. The same type of argument proves that $\mathbf{P_2}(r)$ holds as well, but the stronger inequality in that property means that we do not need to refer to the parity of s in the course of proving it. □

Proposition 5.1.3. *For any $X \in \tilde{D}_n^\circ$, if we regard πEX as a weight of $GL(n)$, then $\beta(\pi EX) = EX$.*

Proof. What we have to prove here is that if we forget the diagram shape of EX and regard it just as a collection of integers, then the algorithm for β recovers the original diagram. We shall show that β rebuilds the first column of EX correctly; the same argument can be iterated to show that successive columns of EX are recovered as well.

Let $\sigma = \pi EX$. Since X has $\mathbf{P_3}(r)$ for all r, adjacent entries of EX must end up in the same clump of σ. Indeed, we have the Remark-5.1.1-like statement that clumps of σ are unions of consecutive rows of EX. We need to prove that the first column of EX contains the odd-index values from every clump—that will establish that EX and $\beta(\sigma)$ have the same first column. To this end, we restrict our attention to a sequence of consecutive rows $i_0, i_0+1, \ldots, i_0+l$ whose entries constitute one clump of σ.

First, we show that $EX_{i_0+j,1} = EX_{i_0+j+1,1} + 2$ for each j. Of course, $EX_{i_0+j,1} - EX_{i_0+j+1,1}$ is at least 2, but if it were larger, that would mean $X_{i_0+j,1}$ is lowerable, and $X_{i_0+j+1,1}$ is raisable. Property $\mathbf{P_2}(r)$ implies that all successive entries in row i_0+j of EX must differ from $EX_{i_0+j+1,1}$ by at least 2, and all successive entries of row i_0+j+1 must differ from $EX_{i_0+j,1}$ by at least 2. But if these conditions held, the entries of rows i_0+j and i_0+j+1 could not belong to the same clump. Therefore $EX_{i_0+j,1} = EX_{i_0+j+1,1}+2$. It follows that the sequence $EX_{i_0,1}, EX_{i_0+1,1}, \ldots, EX_{i_0+l,1}$ is, in fact, a sequence of alternate values from the clump under examination.

Now, we show that the clump does not contain $EX_{i_0,1}+2$ as a value. Suppose this value occurs in row i_0+j. It is larger than $EX_{i_0+j,1}$, and since adjacent entries differ by at most 1, the value $EX_{i_0,1}$ must occur somewhere in row i_0+j. Then Property $\mathbf{P_2}(r)$ tells us that $EX_{i_0,1}+2$ must occur in row i_0+j-1. We iterate this argument and find that $EX_{i_0,1}+2$ occurs in each of rows $i_0+j, i_0+j-1, \ldots, i_0, i_0-1$. But that would imply that the entries of rows i_0-1 and i_0 belong to the same clump, a contradiction.

A similar argument shows that $EX_{i_0+l,1}-2$ does not occur in the clump either. Thus, the sequence of values $EX_{i_0,1}, \ldots, EX_{i_0+l,1}$ is a maximal set of alternate values from the clump. We just need to show that these are the *odd*-index values; or, that $EX_{i_0,1}$ is the largest value in the clump. We know that $EX_{i_0,1}+2$ is not a value in the clump, so we just need to check that $EX_{i_0,1}+1$ is not either. If it were, the argument used in the preceeding paragraph would show that it would occur in row i_0. Let r be the leftmost column such that $EX_{i_0,r} = EX_{i_0,1}+1$. Since adjacent entries differ by at most 1, it must be that $EX_{i_0,r} = EX_{i_0,r-1}+1$. $\mathbf{P_3}(r)$ then tells us that r must be odd. Now we apply $\mathbf{P_1}(r)$ to columns 1 and r and conclude that $EX_{i_0,1}$ is not raisable. But if $EX_{i_0-1,1} = EX_{i_0,1}+2$, and the value $EX_{i_0,1}+1$ occurs in row i_0, it would have to be that the entries of rows i_0-1 and i_0 belong to the same clump. Therefore, $EX_{i_0,1}$ is actually the largest value in its clump, and the sequence $EX_{i_0,1}, \ldots, EX_{i_0+l,1}$ is the sequence of odd-index values from this clump.

We have shown that the first column constructed by β is the same as the first column of EX. The same argument (with appropriate modifications for even-numbered columns) shows that, in general, Z_r as constructed by β in Step 2 is the set of entries in column r of EX. The one remaining detail to check is that the positioning of entries done in Step 3 is the same as the original positioning in EX. The algorithm for β does this positioning so as to satisfy $\mathbf{P_3}(r)$; and following the definition of β, we argued that this positioning can be done uniquely. Since EX satisfies $\mathbf{P_3}(r)$, the positioning found in EX must be that produced by the algorithm.

Thus, $\beta(\pi EX) = EX$. ☐

We combine the preceeding two propositions into the following result:

THEOREM 5.1.4. *$\beta : \Lambda_+(G) \to E(\tilde{D}_n^\circ)$ is a bijection, and its inverse is given by π.*

Proof. We just saw that $\beta(\pi EX) = EX$ for $X \in \tilde{D}_n^\circ$. And it is obvious that $\pi\beta(\sigma) = \sigma$, since β builds a diagram whose entries are the coordinates of a given weight, and π just forgets the diagram shape. ☐

One restatement of this result will be particularly useful to us:

COROLLARY 5.1.5. *Elements of $E(\tilde{D}_n^\circ)$ are uniquely determined by their entries. In particular, for each $\lambda \in D_n$, there is at most one $X \in \tilde{D}_n^\circ$ such that $\pi(EX) = E\lambda$.* ☐

5.2 Computing γ for $GL(n, \mathbb{C})$

Having gotten the machinery of weight diagrams and the maps α and β under our belts, let us return to the task of computing $\gamma(\mathbf{d}, \tau)$. Recall the definitions of $\Theta(\eta)$ and $\xi(\eta)$ from long ago, and the idea expressed in Claim 2.3.1. We would like to find an (η, W_η) whose restriction to $G^{\mathbf{d}}$ is isomorphic to τ and which minimizes $\|\xi(\eta)\|^2$. Recall our abuse of notation in which we write the same symbol for a given irreducible representation of $G^{\mathbf{d}}_{\text{red}}$ and its highest weight.

It was suggested near the end of Chapter 2 that we might be able to produce the desired virtual representation (η, W_η) by the following strategy. We start with an irreducible $P_{\mathbf{d}}$-representation $(\eta_\lambda, W_\lambda)$ such that $\lambda|_{G^{\mathbf{d}}_{\text{red}}} = \tau$ and such that $\|E\lambda\|^2$ is minimized: then, V_τ will be a summand of $W_\lambda|_{G^{\mathbf{d}}_{\text{red}}}$. We then subtract off "smaller" $P_{\mathbf{d}}$-representations as necessary until we are left with a virtual representation with exactly the desired restriction to $G^{\mathbf{d}}_{\text{red}}$.

We know how to find such a λ, using weight diagrams; but to make the above suggestion rigorous, we must acquire some understanding of the word "smaller."

Let us begin not with the parabolics, but with the isotropy groups $G^{\mathbf{d}}_{\text{red}}$. We define a partial order on all of $\coprod_{|\mathbf{d}|=n} \widehat{G}^{\mathbf{d}}$ as follows. If $\tau_1 \in \widehat{G}^{\mathbf{d}_1}$ and $\tau_2 \in \widehat{G}^{\mathbf{d}_2}$, then we write $\tau_1 \prec \tau_2$ if either of the following two conditions holds:

(a) $\mathbf{d}_1 \neq \mathbf{d}_2$, and $\mathcal{O}_{\mathbf{d}_1} \subseteq \overline{\mathcal{O}_{\mathbf{d}_2}}$.

(b) $\mathbf{d}_1 = \mathbf{d}_2$, and there is a sequence of irreducible representations of $G^{\mathbf{d}_1}_{\text{red}}$

$$\tau_1 = v_0, v_1, \ldots, v_m = \tau_2$$

such that for each v_i, there is some $X_i \in \tilde{D}(\mathbf{d}_1, v_i)_{\min}$ such that $V_{v_{i-1}}$ occurs with nonzero multiplicity in $W_{\pi X_i}|_{G^{\mathbf{d}_1}_{\text{red}}}$.

The motivation for the second condition is that for representations of a given $G^{\mathbf{d}}$, the assertion $\tau_1 \prec \tau_2$ ought to mean that τ_1 may occur in $W_{\pi X_2}|_{G^{\mathbf{d}}_{\text{red}}}$ for $X_2 \in \tilde{D}(\mathbf{d}, \tau_2)_{\min}$, but τ_2 can never occur in any $W_{\pi X_1}|_{G^{\mathbf{d}}_{\text{red}}}$ for $X_1 \in \tilde{D}(\mathbf{d}, \tau_1)_{\min}$. The condition is stated in terms of sequences to ensure that the relation is transitive.

Of course, it is not evident that "$\prec$" is a partial order: it is not clear that it is not possible to have both $\tau_1 \prec \tau_2$ and $\tau_2 \prec \tau_1$. We establish this below, along with a few other useful facts. W_G denotes the Weyl group of G throughout.

Lemma 5.2.1. *If $X_1 \in \tilde{D}(\mathbf{d}, \tau_1)_{\min}$ and $X_2 \in \tilde{D}(\mathbf{d}, \tau_2)_{\min}$, then $\pi E X_1$ and $\pi E X_2$ (regarded as weights of the full group G) are not W_G-conjugate.*

Proof. We know that $\alpha(X_1) \in \tilde{D}^\circ(\mathbf{d}, \tau_1)_{\min}$ and $\alpha(X_2) \in \tilde{D}^\circ(\mathbf{d}, \tau_2)_{\min}$. In addition, Corollary 4.1.9 tells us that $E\alpha(X_1)$ has the same integer entries as EX_1, and likewise for X_2. If $\pi E X_1$ and $\pi E X_2$ were conjugate under W_G, then $\pi E\alpha(X_1)$ and $\pi E\alpha(X_2)$ would be as well. In other words, $E\alpha(X_1)$ and $E\alpha(X_2)$ would be distinct diagrams in $E(\tilde{D}^\circ_n)$ that have the same entries. But diagrams in $E(\tilde{D}^\circ_n)$ are determined by their entries, so it must be that $\pi E X_1$ and $\pi E X_2$ are not W_G-conjugate. ☐

Lemma 5.2.2. *If $\tau_1 \prec \tau_2$, then $\tau_2 \not\prec \tau_1$. That is, "$\prec$" defines a partial order on $\coprod_{|\mathbf{d}|=n} \widehat{G}^{\mathbf{d}}$.*

Proof. The statement is obviously true if τ_1 and τ_2 are associated with different nilpotent orbits, so we assume that are both elements of the same $\widehat{G}^{\mathbf{d}}$. Suppose that $\tau_1 \prec \tau_2$ and $\tau_2 \prec \tau_1$; let

$$\tau_1 = v_0, v_1, \ldots, v_m = \tau_2$$

be a sequence of irreducible $G^{\mathbf{d}}_{\text{red}}$-representations, and

$$X_1, \ldots, X_m$$

a sequence of diagrams such that $X_i \in D(\mathbf{d}, v_i)_{\min}$ and v_{i-1} occurs in $W_{\pi X_i}|_{G^{\mathbf{d}}_{\text{red}}}$. Conversely, let

$$\tau_2 = v'_0, v'_1, \ldots, v'_k = \tau_1$$

and

$$X'_1, \ldots, X'_k$$

be additional sequences of irreducible $G^{\mathbf{d}}_{\text{red}}$-representations and diagrams satisfying the corresponding conditions for the relation $\tau_2 \prec \tau_1$.

Now, the fact that v_{i-1} occurs in the restriction of $W_{\pi X_i}$ implies that among the weights of $W_{\pi X_i}$ is a weight μ_{i-1} whose restriction to $G^{\mathbf{d}}_{\mathrm{red}}$ (relative to the particular imbedding $G^{\mathbf{d}}_{\mathrm{red}} \hookrightarrow L_{\mathbf{d}}$ corresponding to the shape of X_i) is the weight v_{i-1}. Now, extremal weights of a representation have maximal norm among weights of the representation; moreover, this maximality is preserved by E. Thus

$$\|E\mu_{i-1}\|^2 \leq \|EX_i\|^2$$

for each i. Equality holds here if and only if μ_{i-1} is an extremal weight of $W_{\pi X_i}$. We can write each μ_i as a diagram $Y_i \in \tilde{D}(\mathbf{d}, v_i)$ of the same shape as X_{i+1}.

Since each X_i belongs to $\tilde{D}(\mathbf{d}, v_i)_{\min}$, and each Y_i to $\tilde{D}(\mathbf{d}, v_i)$, it follows that

$$\|EX_i\|^2 \leq \|E\mu_i\|^2$$

for each i. Stringing all these inequalities together, we get

$$\|E\mu_0\|^2 \leq \|EX_1\|^2 \leq \|E\mu_1\|^2 \leq \cdots \leq \|E\mu_{m-1}\|^2 \leq \|EX_m\|^2.$$

Similarly, let μ'_{i-1} be a weight of $W_{\pi X'_i}$ whose restriction to $G^{\mathbf{d}}_{\mathrm{red}}$ is v'_{i-1}; write each μ'_i as a diagram $Y'_i \in \tilde{D}(\mathbf{d}, v'_i)$ of the same shape as X'_{i+1}. We obtain another chain of inequalities:

$$\|E\mu'_0\|^2 \leq \|EX'_1\|^2 \leq \|E\mu'_1\|^2 \leq \cdots \leq \|E\mu'_{k-1}\|^2 \leq \|EX'_k\|^2.$$

But now we observe that

$$\|EX_m\|^2 \leq \|E\mu'_0\|^2 = \|EY'_0\|^2 \qquad \text{and} \qquad \|EX'_k\|^2 \leq \|E\mu_0\|^2 = \|EY_0\|^2$$

because

$$\begin{array}{ll} X'_k \in \tilde{D}(\mathbf{d}, \tau_1)_{\min} & \qquad X_m \in \tilde{D}(\mathbf{d}, \tau_2)_{\min} \\ Y_0 \in \tilde{D}(\mathbf{d}, \tau_1) & \qquad Y'_0 \in \tilde{D}(\mathbf{d}, \tau_2). \end{array}$$

This means that our two long chains of inequalities above are linked at the ends, and every inequality therein is an equality.

In particular, we find that $\|E\mu_{i-1}\|^2 = \|EX_i\|^2$. Recall that this implies that μ_{i-1} is an extremal weight of $W_{\pi X_i}$; i.e., that μ_{i-1} is conjugate to πX_i under the Weyl group of $L_{\mathbf{d}}$, and hence under W_G. It follows that πEY_{i-1} and πEX_i are also conjugate under $W_{\mathbf{d}}$ and W_G. But on the other hand, we also find that $\|EX_{i-1}\|^2 = \|E\mu_{i-1}\|^2$, so $Y_{i-1} \in \tilde{D}(\mathbf{d}, v_{i-1})_{\min}$. According to the previous lemma, this means that πEY_{i-1} cannot be W_G-conjugate to πEX_i, so we have a contradiction. Thus, it cannot be that both $\tau_1 \prec \tau_2$ and $\tau_2 \prec \tau_1$ hold. $\square$

Lemma 5.2.3. *Suppose $\tau_1, \tau_2 \in \widehat{G}^{\mathbf{d}}_{red}$. If $\tau_1 \prec \tau_2$, and if $X_1 \in \tilde{D}(\mathbf{d}, \tau_1)_{\min}$ and $X_2 \in \tilde{D}(\mathbf{d}, \tau_2)_{\min}$, then $\|EX_1\|^2 < \|EX_2\|^2$.*

Proof. The inequalities derived in the preceding proof establish that $\|EX_1\|^2 \leq \|EX_2\|^2$; we just need to show that the two cannot be equal. Let ν_1 be a weight occurring in the representation $W_{\pi X_2}$ whose restriction to $G^{\mathbf{d}}_{\mathrm{red}}$ is τ_1; write ν_1 as a diagram $Y_1 \in \tilde{D}(\mathbf{d}, \tau_1)$ of the same shape as X_2. We have

$$\|EX_1\|^2 \leq \|EY_1\|^2 \leq \|EX_2\|^2.$$

If the left-most and right-most quantities are equal, then $\|E\nu_1\|^2$ must be equal to the other two. In particular, $Y_1 \in \tilde{D}(\mathbf{d}, \tau_1)_{\min}$. But $\|EY_1\|^2 = \|EX_2\|^2$ implies that ν_1 is an extremal weight in the representation $W_{\pi X_2}$; that is, ν_1 and πX_2 are conjugate under $W_{\mathbf{d}}$ and hence under W_G. But this is a contradiction, since $Y_1 \in \tilde{D}(\mathbf{d}, \tau_1)_{\min}$ and $X_2 \in \tilde{D}(\mathbf{d}, \tau_2)_{\min}$. $\square$

Lemma 5.2.4. *If $X \in \tilde{D}^{\circ}(\mathbf{d}, \tau)_{\min}$, then $\langle V_\tau, W_{\pi X}|_{G^{\mathbf{d}}_{red}} \rangle = 1$.*

Proof. Since τ is the restriction of an extremal weight of $W_{\pi X}$, it is clear that V_τ occurs at least once in the restriction of $W_{\pi X}$ to $G^{\mathbf{d}}_{\text{red}}$. If it were to occur more than once, we would be able to find some other weight μ of $W_{\pi X}$ whose restriction to $G^{\mathbf{d}}_{\text{red}}$ is also τ. Since $X \in \tilde{D}(\mathbf{d}, \tau)_{\min}$, we know that $\|E\mu\|^2 \geq \|\pi E X\|^2$. But the extremal weights are of maximal size among the weights of a representation, so such a μ would also have to be an extremal weight. That is, μ would have be to conjugate to πX under $W_{\mathbf{d}}$.

Now, $\kappa(X) = \tau$ of course, but let us, without loss of generality, make the stronger assumption that the positive root system of $G^{\mathbf{d}}_{\text{red}}$ has been chosen such that $\tilde{\kappa}(X) = \tau$. A weight μ as described above would correspond to a diagram Y of the same shape as X (recall that the shape of X fixes an imbedding $G^{\mathbf{d}}_{\text{red}} \hookrightarrow L_{\mathbf{d}}$), which satisfies $\tilde{\kappa}(Y) = \tau$. Since μ and πX are $W_{\mathbf{d}}$-conjugate, the entries of Y are just those of X, suitably permuted by $W_{\mathbf{d}}$. Note that $\tilde{\kappa}(Y) = \tau$ is a stronger statement than $Y \in \tilde{D}(\mathbf{d}, \tau)$.

Now, $W_{\mathbf{d}}$ acts on diagrams of shape-class $\mathbf{d}$ by permuting the entries within each column but acting on each column separately. Suppose i is the first (uppermost) row in which X and X' differ. That means that some entries in row i of X have been replaced by other entries farther down in the same column to obtain row i of X'. But because $X \in \tilde{D}^\circ(\mathbf{d}, \tau)$, we know that its columns are nonincreasing. Thus $Y_{ir} \leq X_{ir}$ for each r; moreover, the inequality is strict for some r since X and X' differ in row i. Hence the sum of row i of X_1 is strictly less than the sum of row i of X, so it cannot be that $\tilde{\kappa}(Y) = \tilde{\kappa}(X)$. This contradicts the assumption that $\tilde{\kappa}(Y) = \tau$, so there can be no second weight λ' of W_λ whose restriction to $G^{\mathbf{d}}_{\text{red}}$ is τ. ∎

Now we are ready to achieve one of our goals: we show how to compute $\gamma(\mathbf{d}, \tau)$.

Proposition 5.2.5. *(a) If $X_0, X_0' \in \tilde{D}(\mathbf{d}, \tau)_{\min}$, then the weights $\pi E X_0$ and $\pi E X_0'$, regarded as weights of G, are W_G-conjugate.*

(b) $\tilde{D}^\circ(\mathbf{d}, \tau)_{\min}$ contains exactly one element for each $\mathbf{d}$ and τ. If $Z(\mathbf{d}, \tau)$ is this element, then there is a virtual representation (η, W_η) of $L_{\mathbf{d}}$ such that $\eta|_{G^{\mathbf{d}}_{red}} \simeq \tau$, and

$$\xi(\eta) = \pi(EZ(\mathbf{d}, \tau)) = \gamma(\mathbf{d}, \tau).$$

Conversely, $Z(\mathbf{d}, \tau) = E^{-1}\beta(\gamma(\mathbf{d}, \tau))$.

Proof. We prove the two parts of this proposition together, by induction on τ with respect to $\prec$. Let us begin at the zero orbit, given by $\mathbf{d} = [1^n]$. Now, $G^{[1^n]}_{\text{red}} = L_{[1^n]} = G$. Any $(\tau, V_\tau) \in \widehat{G}^{[1^n]}_{\text{red}}$ is itself an irreducible representation of $L_{[1^n]} = G$. The set $\pi E(\tilde{D}([1^n], \tau))$ is just $\{w(E\tau) \mid w \in W_G\}$, so part (a) holds. Since any element of $E(\tilde{D}^\circ([1^n], \tau))$ is determined by its entries, and since the entries of any such element must be the coordinates of $E\tau$, each $\tilde{D}^\circ([1^n], \tau)$ contains exactly one element. This element is $Z([1^n], \tau) = E^{-1}\beta(E\tau)$. Now, since there are no smaller orbits than the zero orbit, Propositions 2.1.7 and 2.2.4 tell us that

$$\Theta(\tau) \simeq N([1^n], \tau) \simeq \bigoplus_{w \in W_G} (-1)^{-w} \operatorname{Ind}_T^G \mathbb{C}_{E_w \tau}.$$

Indeed, V_τ is itself the desired virtual representation of $L_{[1^n]}$: clearly $\tau|_{G^{[1^n]}_{\text{red}}} = \tau$, and $\xi(\tau) = \pi(EZ([1^n], \tau))$. We see by inspection that the largest weight occurring on the right in this equation is $E\tau$, which is precisely πE applied to the unique element $E^{-1}\beta(E\tau)$ of $\tilde{D}^\circ([1^n], \tau)$. Because there are no smaller orbits, $\gamma([1^n], \tau)$ is just this largest weight. This establishes part (b).

Now let us move on to a larger orbit, say $\mathcal{O}_{\mathbf{d}}$. Assume that we have established the result for all smaller orbits contained in $\overline{\mathcal{O}_{\mathbf{d}}}$. We argue by induction on representations τ of $G^{\mathbf{d}}$ with respect to $\prec$. (The following argument is the general inductive step, in which it is assumed that the result is established for all representations of $G^{\mathbf{d}}$ that are smaller than τ with respect to $\prec$. In particular, the same argument also proves the result for the base case in which τ is minimal among irreducible $G^{\mathbf{d}}_{\text{red}}$-representations with respect to $\prec$.)

Let $X, X' \in \tilde{D}(\mathbf{d}, \tau)_{\min}$. Suppose that

$$\eta_{\pi X}|_{G^{\mathbf{d}}_{\mathrm{red}}} \simeq \tau \oplus \bigoplus c_i \tau_i \qquad \text{and} \qquad \eta_{\pi X'}|_{G^{\mathbf{d}}_{\mathrm{red}}} \simeq \tau \oplus \bigoplus c'_i \tau'_i.$$

All the τ_i's and τ'_i's are smaller than τ with respect to $\prec$. (The preceding lemma tells us that V_τ occurs with multiplicity 1 in both of these modules.) Now, consider the virtual representation $W_{\pi X} - W_{\pi X'}$. The restriction of this representation to $G^{\mathbf{d}}_{\mathrm{red}}$ consists entirely of representations with highest weight smaller than τ. Let us compute $\Theta(\eta_{\pi X} - \eta_{\pi X'})$ by Propositions 2.1.7 and 2.2.4.

$$\sum_{w \in W_{\mathbf{d}}} (-1)^{-w} \big(\mathrm{Ind}_T^G \mathbb{C}_{E_w(\pi X)} - \mathrm{Ind}_T^G \mathbb{C}_{E_w(\pi X')}\big) = \sum c_i N(\mathbf{d}, \tau_i) - \sum c'_i N(\mathbf{d}, \tau'_i) + \sum d_j N(\mathbf{d}''_j, \tau''_j)$$

Here the $\mathbf{d}''_j$'s are such that each $\mathcal{O}_{\mathbf{d}''_j}$ is contained in the boundary of $\mathcal{O}_{\mathbf{d}}$. Note that $N(\mathbf{d}, \tau)$ does not appear on the right-hand side. Suppose that πX and $\pi X'$ are not W_G-conjugate. Then there are two weights of maximal size on the left-hand side of this equation, $E(\pi X)$ and $E(\pi X')$. What are the maximal weights on the right-hand side? The largest weights occurring in the various $N(\mathbf{d}, \tau_i)$'s, $N(\mathbf{d}, \tau'_i)$'s, and $N(\mathbf{d}''_j, \tau''_j)$'s are all distinct, since (by the inductive hypothesis) distinct diagrams in $\tilde{D}^\circ_n$ are recovered by applying $E^{-1}\beta$ to these largest weights. So a largest weight on the right-hand side must be one of the $\gamma(\mathbf{d}, \tau_i)$'s, $\gamma(\mathbf{d}, \tau'_i)$'s or $\gamma(\mathbf{d}''_j, \tau''_j)$'s. It follows that one of these values of γ must be $E(\pi X)$ or $E(\pi X')$. But that is impossible, because $E^{-1}\beta(E(\pi X))$ and $E^{-1}\beta(E(\pi X'))$ are both diagrams in $\tilde{D}^\circ(\mathbf{d}, \tau)$, but $E^{-1}\beta$ applied to some $\gamma(\mathbf{d}, \tau_i)$, $\gamma(\mathbf{d}, \tau'_i)$, or $\gamma(\mathbf{d}''_j, \tau''_j)$ should yield a diagram in $\tilde{D}^\circ(\mathbf{d}, \tau_i)$, $\tilde{D}^\circ(\mathbf{d}, \tau'_i)$, or $\tilde{D}^\circ(\mathbf{d}''_j, \tau''_j)$ respectively. So πEX and $\pi EX'$ must be W_G-conjugate: part (a) holds.

Now, it is obvious that $E(\tilde{D}^\circ(\mathbf{d}, \tau)_{\min})$ contains exactly one element, since diagrams in $E(\tilde{D}^\circ_n)$ are determined by their entries, and all diagrams in $E(\tilde{D}(\mathbf{d}, \tau)_{\min})$ have the same entries. It then follows that $\tilde{D}^\circ(\mathbf{d}, \tau)$ contains exactly one diagram as well, which we denote by $Z(\mathbf{d}, \tau)$. How do we construct the desired virtual representation of $L_{\mathbf{d}}$? We begin by comparing again two expressions for $\Theta(\eta_{\pi X})$:

$$\sum_{w \in W_{\mathbf{d}}} (-1)^{-w} \mathrm{Ind}_T^G \mathbb{C}_{E_w(\pi X)} = N(\mathbf{d}, \tau) + \sum c_i N(\mathbf{d}, \tau_i) + \sum e_j N(\mathbf{d}''_j, \tau''_j).$$

As before, the $\mathbf{d}''_j$'s correspond to smaller orbits. By the inductive hypothesis, we have for each τ_i a virtual representation (η_i, W_i) such that $\eta_i|_{G^{\mathbf{d}}_{\mathrm{red}}} \simeq \tau_i$; moreover, by Proposition 2.2.4, we can rewrite $N(\mathbf{d}, \tau_i)$ as a sum of $\Theta(\eta_i)$ and finitely many $N(\mathbf{d}''_j, \tau''_j)$'s from smaller orbits. So we can rewrite the above equation as

$$\sum_{w \in W_{\mathbf{d}}} (-1)^{-w} \mathrm{Ind}_T^G \mathbb{C}_{E_w(\pi Z(\mathbf{d}, \tau))} - \sum c_i \Theta(\eta_i) = N(\mathbf{d}, \tau) + \sum e'_j N(\mathbf{d}''_j, \tau''_j). \tag{5.1}$$

Lemma 5.2.3 implies that $\|\xi(\eta_i)\|^2 < \|\pi(EZ(\mathbf{d}, \tau))\|^2$, so the largest weight on the left-hand side of this equation is $\pi(EZ(\mathbf{d}, \tau))$. Introduce the virtual representation $\eta_0 = \eta_{\pi(EZ(\mathbf{d}, \tau))} - \sum c_i \eta_i$; then, the left-hand side of the above equation is just $\Theta(\eta_0)$, and $\xi(\eta) = \pi(EZ(\mathbf{d}, \tau))$.

It remains to establish that $\gamma(\mathbf{d}, \tau) = \pi(EZ(\mathbf{d}, \tau))$. If this equation did not hold, we would have $\|\gamma(\mathbf{d}, \tau)\|^2 < \|\pi(EZ(\mathbf{d}, \tau)\|^2$. Moreover, there would be some expression of the form $\sum e''_j N(\mathbf{d}''_j, \tau''_j)$, supported entirely on smaller orbits, that we could add to both sides of (5.1) to achieve the minimum largest weight $\gamma(\mathbf{d}, \tau)$. Since this is smaller than $\pi(EZ(\mathbf{d}, \tau))$, it must be that the largest weight of $\sum e''_j N(\mathbf{d}'_j, \tau''_j)$ is also $\pi(EZ(\mathbf{d}, \tau))$, but with coefficient negative of that which it has in (5.1). On the other hand, we know by induction that the largest weight occurring in $\sum e''_j N(\mathbf{d}'_j, \tau''_j)$ is some $\gamma(\mathbf{d}'', \tau''_j) = \pi(EZ(\mathbf{d}'', \tau''_j))$. But no $\pi(EZ(\mathbf{d}'', \tau''_j))$ can be equal to $\pi(EZ(\mathbf{d}, \tau))$, since elements of $\tilde{D}_n$ are determined by their entries. In other words, there is no way to make the largest weight occurring in (5.1) smaller by adding or subtracting $N(\mathbf{d}''_j, \tau''_j)$'s from smaller orbits. We conclude that $\gamma(\mathbf{d}, \tau) = \pi(EZ(\mathbf{d}, \tau))$. □

THEOREM 5.2.6. *$\gamma : \{(\mathbf{d}, \tau)\} \to \Lambda_+(G)$ is a bijection.*

Proof. In the notation of the preceding proposition, we have $\gamma(\mathbf{d}, \tau) = \pi(EZ(\mathbf{d}, \tau))$. The elements $Z(\mathbf{d}, \tau) \in \tilde{D}^\circ(\mathbf{d}, \tau)_{\min}$ are all distinct, so γ is injective.

Now, let λ be a dominant weight of G. Construct the diagram $X = E^{-1}\beta(\lambda)$. Suppose $X \in \tilde{D}(\mathbf{d}, \tau)$; then, consider the $L_{\mathbf{d}}$-representation $(\eta_{\pi X}, W_{\pi X})$. If this representation is such that $\eta_{\pi X}|_{G^{\mathbf{d}}_{\mathrm{red}}} \simeq \sum c_i \tau_i$, then it follows that

$$\Theta(\eta_{\pi X}) \simeq \sum c_i \Theta(\eta_{\pi(EZ(\mathbf{d}, \tau_i))}).$$

Write both sides of this equation in the form $\sum m_\sigma \operatorname{Ind}_T^G \mathbb{C}_\sigma$; then the largest weight occurring on the left-hand side is $\pi(EX)$, while the largest on the right-hand side is some $\pi(EZ(\mathbf{d}, \tau_i))$. Thus $\pi(EX) = \pi(EZ(\mathbf{d}, \tau_i))$ for some i; and $\lambda = \gamma(\mathbf{d}, \tau_i)$. Thus, γ is surjective. $\square$

COROLLARY 5.2.7. *$\tilde{D}^\circ(\mathbf{d}, \tau) = \tilde{D}^\circ(\mathbf{d}, \tau)_{\min}$ for each $\mathbf{d}$ and τ. Moreover, for any $X \in \tilde{D}(\mathbf{d}, \tau)$, $\alpha(X)$ is the unique element of $\tilde{D}^\circ(\mathbf{d}, \tau)$.*

Proof. Again, let $Z(\mathbf{d}, \tau)$ denote the unique element of $\tilde{D}^\circ(\mathbf{d}, \tau)_{\min}$. We know that $\gamma(\mathbf{d}, \tau) = \pi EZ(\mathbf{d}, \tau)$. We just learned that γ is surjective, but we showed after constructing β that πE gives a bijection between $\tilde{D}^\circ_n$ and $\Lambda_+(G)$. So $\tilde{D}^\circ(\mathbf{d}, \tau)$ cannot contain any additional elements beyond $Z(\mathbf{d}, \tau)$. The latter half of the corollary is immediate, since for any $X \in \tilde{D}(\mathbf{d}, \tau)$, we have $\alpha(X) \in \tilde{D}^\circ(\mathbf{d}, \tau)$. $\square$

This corollary gives us a concrete way to compute γ: given $(\mathbf{d}, \tau)$, draw any weight diagram X of shape-class $\mathbf{d}$ such that $\kappa(X) = \tau$; for instance, the diagram whose first column contains the coordinates of τ and whose later columns are all zero. Then compute $\alpha(X)$ by following the algorithm. By the preceding corollary, $\alpha(X)$ will be the unique element of $\tilde{D}^\circ(\mathbf{d}, \tau)$. Then take $E(\alpha(X))$, and regard its entries as the coordinates of a dominant weight of G. The weight thus obtained is $\gamma(\mathbf{d}, \tau)$.

We can compute γ^{-1} algorithmically as well. For $\lambda \in \Lambda_+(G)$, the above results imply the following formula:

$$\gamma^{-1}(\lambda) = (\text{shape-class of } E^{-1}(\beta\lambda), \kappa(E^{-1}(\beta\lambda))).$$

Thus, we have the desired explicit description of the relationship between the bases of Theorems 1 and 2 for the groups $GL(n, \mathbb{C})$.

CHAPTER 6

Partial Results for the Symplectic Group

In this chapter, we investigate the possibility of developing an analogue of the preceding work for the group $G = Sp(2n, \mathbb{C})$. All the algebraic development in Chapter 2 is still valid in this setting, but in the intervening chapters we made a particular choice of parabolic in the case of $GL(n, \mathbb{C})$ in order to obtain the simplifying relation $\mathfrak{u}_{\mathbf{d}} = \mathfrak{v}_{\mathbf{d}}$. The parabolic in question was the one with respect to which $\mathcal{O}_{\mathbf{d}}$ was Richardson. Not every orbit in $Sp(2n, \mathbb{C})$ is Richardson, so we cannot hope to effect the same simplification in this setting. Instead, we stick to using the Jacobson-Morozov parabolic. Then, when we try to apply the formula in Proposition 2.2.4, we will not be able to ignore $\mathfrak{u}_{\mathbf{d}}/\mathfrak{v}_{\mathbf{d}}$; rather, we will give an explicit description of this space in the language of weight diagrams, and make some guesses about what part of it should be added.

We are introducing a new kind of weight diagram in this chapter, but we will have occasion to refer to the kind used in the context of $GL(n, \mathbb{C})$. When ambiguity might result, the latter will be called *weight diagrams of type A*.

6.1 Weight Diagrams

We begin by collecting some basic facts about nilpotent orbits in the symplectic group, as found in, say, [7]. Nilpotent orbits for $Sp(2n, \mathbb{C})$ are indexed by partitions of $2n$ in which odd parts occur with even multiplicity: such a partition will be called a *C-partition*. Let $\mathbf{d} = [k_1^{a_1}, \ldots, k_l^{a_l}]$ be a C-partition, and let $\mathcal{O}_{\mathbf{d}}$ be the nilpotent orbit corresponding to it. Let $G^{\mathbf{d}}$ be the isotropy group of this orbit, and let $G^{\mathbf{d}}_{\mathrm{red}}$ denote its reductive part, as usual. We have

$$G^{\mathbf{d}}_{\mathrm{red}} \simeq \prod_{\{i|k_i \text{ odd}\}} Sp(a_i, \mathbb{C}) \times \prod_{\{i|k_i \text{ even}\}} O(a_i, \mathbb{C}).$$

Let $P_{\mathbf{d}}$ denote the Jacobson-Morozov parabolic subgroup for this orbit, and let $L_{\mathbf{d}}U_{\mathbf{d}}$ be the Levi decomposition thereof. Then $L_{\mathbf{d}}$ is given by

$$L_{\mathbf{d}} = Sp\left(\sum_{\{i|k_i \text{ odd}\}} a_i\right) \times \prod_{p\geq 2} GL\left(\sum_{\{i|k_i \,\equiv\, p (\mathrm{mod}\, 2) \text{ and } k_i \,\geq\, p\}} a_i\right).$$

We seek to mimic the $GL(n, \mathbb{C})$ development, but we face an obstacle at the very first step. Weights for $Sp(2n, \mathbb{C})$ cannot possibly be described by filling in the boxes of a diagram whose shape-class is a C-partition, since the latter has $2n$ boxes, but $Sp(2n, \mathbb{C})$ only has rank n. We instead take inspiration from the work of Garfinkle and others on "domino tableaux," and introduce the idea of a weight diagram whose constituents are dominoes rather than square boxes.

Definition 6.1.1. Consider the set of all left-justified pictures made up of 1×2 and 2×1 blocks (called *horizontal* and *vertical dominoes*, respectively), each of which contains an integer entry. The *shape-class* of such a diagram is the partition whose parts are the lengths of the rows of the diagram. (The length of a row is the number of *square* boxes the row would contain if every domino were replaced by two square boxes.) Such a diagram is called a *weight diagram of type C* if its shape-class is a C-partition.

We need a notation to refer to entries in a weight diagram of type C. We retain the type-A coordinates arising from counting square blocks, and we introduce the following conventions: A horizontal domino entry is referred to by the coordinates of its right-hand square, and a vertical domino is referred to by the coordinates of its upper square.

Now, C-partitions are such that rows of odd length come in pairs. We can manage, therefore, to use very few vertical dominoes in the picture: rows of even length can obviously be covered entirely by horizontal dominoes, and for pairs of rows of odd length, we can put one vertical domino at the left-hand edge of the diagram and fill out the rest with horizontal ones, as illustrated below.

The reason for wanting to have nearly all the dominoes be horizontal has to do with how one restricts weights from $L_{\mathbf{d}}$ to $G^{\mathbf{d}}_{\text{red}}$: we shall return to this shortly.

Definition 6.1.2. A weight diagram of type C is said to be in *standard form* if vertical dominoes only meet rows of odd length and only occur along the left-hand edge of the diagram.

Remark 6.1.3. If X is a weight diagram in standard form, there is an entry at X_{ir} if and only if r has the same parity as the length of row i.

In the case of $GL(n, \mathbb{C})$, the restriction of weights was accomplished by summing up the rows of the diagram; but in $Sp(2n, \mathbb{C})$, the rank of $G^{\mathbf{d}}_{\text{red}}$ is roughly half the number of rows. In fact, we shall need a designated grouping of (almost) all rows of a given length into ordered pairs, with each pair contributing one coordinate to the weight of $G^{\mathbf{d}}_{\text{red}}$. In a weight diagram in standard form, odd-length rows come with an *a priori* such pairing: each vertical domino determines an ordered pair of odd-length rows.

Definition 6.1.4. A weight diagram of type C is said to be *fully equipped* if it is in standard form, and if there is given a grouping of all the rows of the rows of the diagram into ordered pairs and singletons such that

(a) Both members of a given pair are rows of the same length.

(b) Each odd-length row is paired with the row with which it shares a vertical domino.

(c) There is at most one singleton row of any given even row length, and there are no singleton rows of odd length.

The set of all fully equipped weight diagrams of type C with a total of n domino entries is denoted $\tilde{D}(C_n)$. The set of all such weight diagrams of shape-class $\mathbf{d}$ is denoted $\tilde{D}(\mathbf{d}; C_n)$. In each ordered pair of rows, the first member will be called the *dorsal* member, and the second one *ventral*. Occasionally, singleton rows may be referred to as dorsal rows, although they lack ventral partners.

All weight diagrams of type C will be henceforth assumed to be fully equipped unless explicitly stated otherwise, although the pairing of rows may not always be explicitly specified. Typically, where confusion would not result, "$\tilde{D}(C_n)$" will be abbreviated to "$\tilde{D}_n$," and "$\tilde{D}(\mathbf{d}; C_n)$" to "$\tilde{D}(\mathbf{d})$."

Thus, the partition $[6, 5^2, 2, 1^2]$ may give rise to a diagram of the following shape:

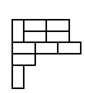

6.2 Extracting Weights from Weight Diagrams

Now, we turn our attention to getting weights of $L_{\mathbf{d}}$ and $G^{\mathbf{d}}_{\text{red}}$ from a weight diagram. As with $GL(n, \mathbb{C})$, the various factors of this Levi subgroup correspond to columns of a weight diagram. The way to obtain a weight of $L_{\mathbf{d}}$ from a weight diagram filled with integer entries is as follows: the entries in vertical dominoes constitute the coordinates of the weight of the $Sp(\cdot)$ factor, and the entries in all the horizontal dominoes whose right half lies in column r constitute the coordinates of the weight of $GL(\sum_{\{i | k_i \,\equiv\, r (\mathrm{mod}\, 2) \text{ and } k_i \,\geq\, r\}} a_i)$. For example, a weight of $L_{\mathbf{d}}$ is obtained from the following diagram as shown.

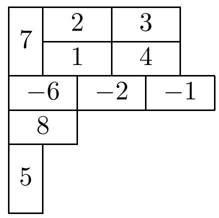

$$\mapsto \underbrace{(7,5)}_{Sp(2)}, \underbrace{(-6,8)}_{GL(2)}, \underbrace{(2,1)}_{GL(2)}, \underbrace{(-2)}_{GL(1)}, \underbrace{(3,4)}_{GL(2)}, \underbrace{(-1)}_{GL(1)}$$

By analogy with the type-A case, we write $D(\mathbf{d}; C_n)$ for the collection of weights of $L_{\mathbf{d}}$, and

$$D(C_n) = \coprod_{\substack{|\mathbf{d}|=2n \\ \mathbf{d} \text{ a } C\text{-partition}}} D(\mathbf{d}; C_n)$$

for their disjoint union. We give the name $\pi : \tilde{D}(C_n) \to D(C_n)$ to the map defined by the above procedure for reading off an $L_{\mathbf{d}}$-weight from a weight diagram.

Next we discuss how to find the restriction of an $L_{\mathbf{d}}$-weight to $G^{\mathbf{d}}_{\mathrm{red}}$, a task that is considerably more complicated than in the $GL(n, \mathbb{C})$ case. In addition to the aforementioned fact that the rank of $G^{\mathbf{d}}_{\mathrm{red}}$ is only about half the number of rows, there is the problem that some of the $G^{\mathbf{d}}_{\mathrm{red}}$ groups are disconnected (they may include orthogonal groups as factors), so their irreducible representations are not quite in bijection with the set of dominant weights.

Because of that last concern, we eschew our $GL(n, \mathbb{C})$-habit of confusing notation for irreducible representations and their highest weights. Indeed, we ought to have an established description of irreducible representations of $G^{\mathbf{d}}_{\mathrm{red}}$ before we can hope to say how to restrict weights of $L_{\mathbf{d}}$ to it, so before proceeding, let us briefly recall the representation theory of the orthogonal group. If k is odd, then $O(k)$ is just isomorphic to $\mathbb{Z}/2\mathbb{Z} \times SO(k)$, so an irreducible representation of $O(k)$ is specified by giving the highest weight of the $SO(k)$ action, along with one of "+" or "−" to indicate whether the $\mathbb{Z}/2\mathbb{Z}$ factor acts trivially or nontrivially. (We shall use "+" to denote the trivial representation of $\mathbb{Z}/2\mathbb{Z}$.) On the other hand, if k is even, let $(p_1, \ldots, p_{k/2})$ be a dominant weight of $SO(k)$. If the last coordinate $p_{k/2}$ is nonzero, then there is a unique irreducible $O(k)$-representation whose restriction to $SO(k)$ contains the irreducible representation of the given highest weight as a constituent. But if $p_{k/2} = 0$, there are two possible ways of extending the given $SO(k)$-representation to an $O(k)$-representation: we must specify one of the signs $\pm$ to indicate how the element

$$J = \begin{pmatrix} 1 & & & \\ & \ddots & & \\ & & 1 & \\ & & & -1 \end{pmatrix} \in O(k)$$

acts on the highest weight space.

Suppose, now, that our partition $\mathbf{d}$ has an even part k_i with multiplicity a_i. Then, $G^{\mathbf{d}}_{\mathrm{red}}$ contains a factor $O(a_i)$, which sits diagonally inside a product of k_i different $GL(\cdot)$'s. The size of those $GL(\cdot)$'s does not matter *per se*: we can try to understand how to restrict weights by looking at a particular example, the standard inclusion $O(a_i) \hookrightarrow GL(a_i)$. On the level of Lie algebras, the inclusion of tori looks like

$$\begin{pmatrix} h_1 & & & & \\ & -h_1 & & & \\ & & \ddots & & \\ & & & h_{a_i/2} & \\ & & & & -h_{a_i/2} \end{pmatrix} \quad \text{or} \quad \begin{pmatrix} h_1 & & & & & \\ & -h_1 & & & & \\ & & \ddots & & & \\ & & & h_{(a_i-1)/2} & & \\ & & & & -h_{(a_i-1)/2} & \\ & & & & & 0 \end{pmatrix},$$

depending on the parity of a_i. This is where the pairing of rows in the diagram comes into play: if (i_1, i_2) is an ordered pair of rows of the same length k_i, we associate them with some pair of coordinates h_i, $-h_i$ on the imbedded toral subalgebra. Note that only even-numbered columns appear when we name the horizontal dominoes that occur in these two rows. Then, the i-th coordinate of the restricted weight is given by the sum of the entries in row i_1, minus the sum of the entries in row i_2:

$$\sum_{r=1}^{k_i/2} X_{i_1,2r} - \sum_{r=1}^{k_i/2} X_{i_2,2r}.$$

What about a possible singleton row of length k_i? If there is such a singleton row, it would correspond to the "0" coordinate at the lower right-hand corner for odd-size orthogonal groups, as shown above. This means that the singleton row has no bearing on the restriction of the weight to the $SO(\cdot)$ part, but it does bear some relation to the action of the nonidentity component of $O(a_i)$.

Rather than treat this case specifically, we instead look at how to determine the action of the nonidentity component in general. If a_i is even, let (i_0', i_0) be the ordered pair of rows corresponding to the $(a_i/2)$-th (*i.e.*, the last) coordinate of an $O(a_i)$-weight; if a_i is odd, let i_0 be the singleton row. In either case, let

$$p = \sum_{r=1}^{k_i} X_{i_0 r}.$$

Now, the matrix J defined above acts by $(-1)^p$. Thus, whenever a plus or minus sign needs to be specified in addition to a highest weight to fully describe the action of $O(a_i)$, it is obtained as

$$\begin{cases} + & \text{if } p \equiv 0 \pmod 2; \\ - & \text{if } p \equiv 1 \pmod 2). \end{cases}$$

We still need to describe the restriction of weights to $Sp(\cdot)$ factors of $G^{\mathbf{d}}_{\text{red}}$. Since the symplectic group is connected, we do not face the difficulties that confronted us in the orthogonal case; we need only produce a weight. The imbedding of the torus of $Sp(\cdot)$ in a general linear group looks the same as for even orthogonal groups, and the formula for restricting a weight looks almost the same:

$$\sum_{r=1}^{(k_i+1)/2} X_{i_1,2r-1} - \sum_{r=2}^{(k_i+1)/2} X_{i_2,2r-1}.$$

This formula differs from the corresponding one for $O(\cdot)$ factors in that now k_i is odd, and the coordinates of the domino entries have odd column numbers. Moreover, the sum for row i_1 has one more term than that for row i_2: it includes the vertical domino in the first column.

An example of weight restriction may help illustrate this process:

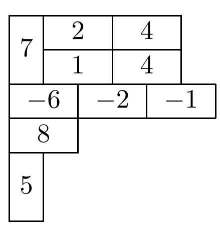

$$\mapsto \underbrace{(8)}_{Sp(2)}, \underbrace{(-)}_{O(1)}, \underbrace{(+)}_{O(1)}, \underbrace{(5)}_{Sp(2)}$$

Now that we understand restriction, we can introduce some of the notation we had in the $GL(n, \mathbb{C})$ case. We define $\kappa : \tilde{D}(\mathbf{d}) \to \widehat{G}^{\mathbf{d}}_{\text{red}}$ to be the map which assigns to each diagram the irreducible algebraic representation of $G^{\mathbf{d}}_{\text{red}}$ obtained by the above procedure. Furthermore, if τ is an irreducible algebraic representation of $G^{\mathbf{d}}_{\text{red}}$, we define

$$\tilde{D}(\mathbf{d}, \tau) = \{X \in \tilde{D}(\mathbf{d}) \mid \kappa(X) = \tau\}.$$

The final task for this section is to define the maps $E, F : \tilde{D}_n \to \tilde{D}_n$. The former is quite straightforward: recall, once again, that each column save the first corresponds to a $GL(\cdot)$ factor of

$L_{\mathbf{d}}$; for each of these, the formula is precisely the same as in type A:

$$(EX)_{ir} = X_{ir} + \#\{X_{jr} \mid X_{jr} < X_{ir}, \text{ or } X_{jr} = X_{ir} \text{ with } j > i\} \\ - \#\{X_{jr} \mid X_{jr} > X_{ir}, \text{ or } X_{jr} = X_{ir} \text{ with } j < i\},$$

provided that $r > 1$. The first column consists of vertical dominoes and corresponds to the $Sp(\cdot)$ factor of $L_{\mathbf{d}}$. In this case, the formula is

$$(EX)_{i1} = X_{i1} + 2p \cdot \#\{X_{j1} \mid |X_{j1}| < |X_{i1}|, \text{ or } |X_{j1}| = |X_{i1}| \text{ with } j > i\}$$

where

$$p = \begin{cases} +1 & \text{if } X_{i1} \geq 0, \\ -1 & \text{if } X_{i1} < 0. \end{cases}$$

The map F is significantly harder to give an analysis of. In fact, we shall not do so presently, but rather content ourselves with a description of the weights that occur in $\mathfrak{u}_{\mathbf{d}}/\mathfrak{v}_{\mathbf{d}}$. Recall that this last space is the eigenspace of eigenvalue 1 for the semisimple element of a Jacobson-Morozov triple for the orbit $\mathcal{O}_{\mathbf{d}}$. Let $\{h_{\mathbf{d}}, e_{\mathbf{d}}, f_{\mathbf{d}}\}$ be this Jacobson-Morozov $\mathfrak{sl}(2)$-triple. Now, the partition $\mathbf{d}$ is essentially a description of the decomposition of the standard representation of $\mathfrak{sp}(2n)$ on $\mathbb{C}^{2n}$ into representations of this $\mathfrak{sl}(2)$.

It is important not to confuse the two representations of this $\mathfrak{sl}(2)$, on $\mathfrak{sp}(2n)$ by the adjoint action and on $\mathbb{C}^{2n}$ by the standard representation. It is the 1-eigenspace in the former that we seek to identify, but we shall do so by examining the latter, which is described by $\mathbf{d}$. Now, the roots of $\mathfrak{sp}(2n)$ have the form $\varepsilon_i \pm \varepsilon_j$ and $\pm 2\varepsilon_i$. The values of the ε_i's on $h_{\mathbf{d}}$ are (roughly speaking) the nonnegative eigenvalues of $h_{\mathbf{d}}$ on $\mathbb{C}^{2n}$. (Of course, 0 occurs as a value of the ε_i's with only half the multiplicity it has as an eigenvalue of $h_{\mathbf{d}}$ on $\mathbb{C}^{2n}$.) So $h_{\mathbf{d}}$ has an eigenvector of the form $\varepsilon_i - \varepsilon_j$ with eigenvalue 1 in its adjoint action for each pair of nonnegative eigenvalues differing by 1 in its standard action. In addition, it has 1-eigenspaces of the form $\varepsilon_i + \varepsilon_j$ each time 1 and 0 occur as a pair of eigenvalues in the standard action.

How do we translate this into something having to do with weight diagrams? The columns of a weight diagrams correspond to factors of $L_{\mathbf{d}}$, which in turn correspond to the nonnegative eigenvalues of $h_{\mathbf{d}}$ in the standard representation. So there is an eigenvector of eigenvalue 1 for $h_{\mathbf{d}}$ in its adjoint action for each pair of entries in neighboring columns; and there is an additional such eigenvector when the pair of dominoes lie in the first and second columns.

Specifically, fix a weight diagram X. Let $Z(i, r; j, s)$ be a weight diagram of the same shape as X, all of whose entries are 0, save that it has a -1 at location (i, r), and a $+1$ at location (j, s). Let $Z'(i, r; j, s)$ denote the same diagram, but with $+1$'s at both locations. Consider the following set of diagrams.

$$\mathcal{P}(X) = \{Z(i, r; i+1, s) \mid X \text{ has entries at both } (i, r) \text{ and } (i+1, s)\} \\ \cup \{Z'(1, r; 2, s) \mid X \text{ has entries at both } (1, r) \text{ and } (2, s)\}$$

Now, diagrams of the same shape can be added entry-by-entry. The diagrams in $\mathcal{P}(X)$ are in one-to-one correspondence with the weights of $\mathfrak{u}_{\mathbf{d}}/\mathfrak{v}_{\mathbf{d}}$; according to Claim 2.3.2, we wish to add some subset of these to X such that the size of the resulting sum is maximized. In other words, the map F for type-C diagrams is defined exactly as in (2.16) and (2.17).

6.3 Even Orbits in the Symplectic Group

Because computing F is a dicey matter, we put it aside for the time being. We instead concentrate on computing γ on even nilpotent orbits, for which $\mathfrak{u}_{\mathbf{d}}/\mathfrak{v}_{\mathbf{d}} = 0$, and for which we can therefore ignore F. The even nilpotent orbits are precisely those orbits indexed by partitions $\mathbf{d}$ all of whose parts have the same parity. Moreover, even among the even orbits, the easiest to handle are those picked out by the following definition.

Definition 6.3.1. A partition is called *rectangular* if it has just one part with some multiplicity: $[k^a]$. A nilpotent orbit in a classical group is called *rectangular* if it is indexed by a rectangular partition. A rectangular partition (resp. the corresponding orbit) is called *evenly rectangular* if the multiplicity a is even.

We shall assume that Claim 2.3.1 holds, and we shall compute γ for representations of rectangular orbits in that context. Now, all the rows of the diagram come in pairs (except possibly for one singleton row), and the condition that κ have a particular value on a given diagram is a condition about the difference of the sums of the entries in paired rows. Perhaps if we rewrote the diagram in a different way, the κ condition would take the form of just requiring particular row-sums—this is the condition we had in working with weight diagrams for $GL(n, \mathbb{C})$. Then, perhaps the $GL(n, \mathbb{C})$ techniques can be used to carry out a computation in this setting.

For our first forays into computing γ, we further restrict our attention to the evenly rectangular orbits.

Definition 6.3.2. Let X be a fully equipped weight diagram of type C, corresponding to an evenly rectangular orbit. The *severing* of X is a certain weight diagram $X^{\vee}$ of type A, obtained as follows. First, we write down all the dorsal rows of X as a diagram of square boxes. Then, we juxtapose this with a diagram whose rows consist of the negated entries of ventral rows of X, such that each negated ventral row appears immediately to the right of its dorsal partner.

Conversely, a rectangular type-A diagram Y can be converted into an evenly rectangular type-C diagram. The type-C diagram X whose severing is Y is determined up to ordering of its rows; any such X with $X^{\vee} = Y$ is called a *splicing* of Y. The particular splicing which all dorsal rows first, in the order in which they occur in Y, followed by all ventral rows in the reverse order, is called the *standard splicing* of Y.

Example 6.3.3. Here is the severing of a diagram of shape-class $[5^4]$, yielding a type-A diagram of shape-class $[5^2]$. A small gap has been inserted to separate the dorsal part of each row from the negated ventral part.

<table>
<tr><td rowspan="2">5</td><td>3</td><td>2</td></tr>
<tr><td>−4</td><td>−4</td></tr>
<tr><td rowspan="2">0</td><td>2</td><td>1</td></tr>
<tr><td>−3</td><td>−1</td></tr>
</table>

$\mapsto$

<table>
<tr><td>5</td><td>3</td><td>2</td><td>4</td><td>4</td></tr>
<tr><td>0</td><td>2</td><td>1</td><td>3</td><td>1</td></tr>
</table>

The standard splicing of the latter diagram is difficult to draw, because the first row is not adjacent to its ventral partner. One might draw the following diagram, which is intended to have only four rows:

<table>
<tr><td rowspan="2">5</td><td>3</td><td>2</td></tr>
<tr><td></td><td></td></tr>
<tr><td rowspan="2">0</td><td>2</td><td>1</td></tr>
<tr><td>−3</td><td>−1</td></tr>
<tr><td></td><td>−4</td><td>−4</td></tr>
</table>

This example demonstrates the need for a better notation for drawing domino weight diagrams.

How does κ behave with respect to severing? In type A, the output of κ is just a dominant weight; but in type C, it may be a highest weight accompanied by a plus or minus sign. Let us write κ_A and κ_C to distinguish the maps κ in the type-A and type-C cases. The following lemma is easy to establish.

Lemma 6.3.4. *Let X be an evenly rectangular weight diagram of type C, with severing $X^{\vee}$. We have*

$$\kappa_A(X^{\vee}) = \textit{highest weight of } \kappa_C(X),$$

where we regard the weights merely as tuples of integers. ◻

Of course, the map E looks quite different. Or does it? As with κ, let us introduce the subscripted notations E_A and E_C to distinguish the two operations; furthermore, we have the new function $E_C^{\vee}$

on rectangular type-A diagrams defined by

$$E_C^{\vee}(X) = \big(E_C(\text{the standard splicing of } X)\big)^{\vee}.$$

(There is no particular significance to having chosen the standard splicing for the right-hand side. Any splicing would do; we have just chosen one for the sake of having $E_C^{\vee}$ be well-defined.) All the machinery developed in Chapters 3 and 4 is set up to work with E_A. If we want to compute something involving $E_C^{\vee}$, we must understand the relationship between these two maps.

Let us define yet another map on type-A weight diagrams as follows.

$$(E_0X)_{ir} = \begin{cases} (E_AX)_{ir} + (\text{number of rows of } X) & \text{if } r > 1, \\ (E_AX)_{i1} + (\text{number of rows of } X) + 1 & \text{if } r = 1. \end{cases} \tag{6.1}$$

Suppose we make the assumption that in a given column, all the entries in ventral rows are less than than any entry in any dorsal row. (This assumption is satisfied in the example above: indeed, all entries in dorsal rows are positive, and all those in ventral rows are negative.) Under this assumption, our $E_C^{\vee}$ is given exactly by the above formula for E_0. The salient property of the latter is that the amount by which it differs from E_A is constant on columns.

Proposition 6.3.5. *The propositions of Section 3.2 remain true if, in every statement, we replace E_A by a modified function E', defined as*

$$(E'X)_{ir} = E_AX_{ir} + f(r),$$

where f is some integer-valued function of the column index.

Proof. The reader may examine those proofs and determine that the actual definition of E_A is almost never used. Most of these proofs entail computations with quantities of the form $E_AX_{ir} - E_AX_{jr}$; such a difference is obviously left unchanged by columnwise addition of a quantity $f(r)$. The proof of Proposition 3.2.8 does mention the definition explicitly, but only as a stepping-stone to equation (3.4), and that equation still holds even if the definition of E_A is modified. ☐

The point of the preceding proposition is, of course, that the aforementioned E_0 is exactly in the form of E'. Now, this proposition is useless unless we can meet the assumption that all dorsal entries in a given column are smaller than all ventral entries. Instead of attacking this problem directly, we can define it out of existence, and then subsequently determine that our definition was not nonsensical. To rephrase in a less opaque manner, we shall see that it is just as well to apply the type-A machinery to the severed diagram using E_0 even when E_0 does not coincide with $E_C^{\vee}$. The following lemma tells us why.

Lemma 6.3.6. $\|E_0X\|^2 \le \|E_C^{\vee}X\|^2$.

Proof. The map $E_C^{\vee}$ just adds $2\rho_{L_{\mathfrak{d}}}$, after choosing a positive root system with respect to which πX is dominant. The map E_0 was defined, essentially, as "the map that $E_C^{\vee}$ would be if πX met certain dominance conditions." Another way of saying that is that E_0 also just adds $2\rho_{L_{\mathfrak{d}}}$, but possibly with respect to a different choice of positive root system. Now, if λ is a dominant weight, it follows that

$$\|\lambda + w \cdot 2\rho_{L_{\mathfrak{d}}}\|^2 \qquad (w \in W_{L_{\mathfrak{d}}})$$

is maximized when $w = 1$, just from the elementary theory of weights and root systems. The statement that

$$\|\lambda + w \cdot 2\rho_{L_{\mathfrak{d}}}\|^2 \le \|\lambda + 2\rho_{L_{\mathfrak{d}}}\|^2$$

is just a rephrasing of the statement of the lemma. ☐

We shall need the following lemma concerning the operation of the type-A algorithms on rectangular diagrams. We use the notation $\lfloor x \rfloor$ (resp. $\lceil x \rceil$) to denote the greatest (resp. least) integer less (resp. greater) than x.

LEMMA 6.3.7. *Let $\mathbf{d}$ be a rectangular partition, and let $X \in \tilde{D}^{\circ}(\mathbf{d}, \tau; A_n)$. Then, on any given row of X, two entries may differ by at most 1. Moreover, if the i-th coordinate of the highest weight of τ is p, and if X has k columns, then every entry on the i-th row of X is either $\lfloor p/k \rfloor$ or $\lceil p/k \rceil$.*

Proof. Recall the properties $\mathbf{p_1}(r), \ldots, \mathbf{p_4}(r)$ defined in Section 3.2, all of which are had by our X. Suppose, without loss of generality, that we have $X_{is} - X_{ir} \geq 2$, with $s < r$. Then, X_{is} must not be lowerable, so i is not the last row, and indeed $X_{i+1,s} = X_{is}$. Since X is a rectangular diagram, there is also an entry at $(i+1, r)$, such that $X_{i+1,r} \leq X_{ir}$. It follows that $X_{i+1,s} - X_{i+1,r} \geq 2$.

What we have shown is that if row i has two entries that differ by at least 2, then it is not the bottom row, and row $i+1$ also has two entries that differ by at least two. We could repeat this argument *ad infinitum*, but X of course only has finitely many rows, so we have a contradiction.

For the second part of the lemma, we have $\sum_r X_i r = p$. If some entry were larger than $\lceil p/k \rceil$, then the fact that no entries can differ by more than one would imply that all the entries were at least as large as $\lceil p/k \rceil$, with at least one strictly larger. This means the sum of the entries on this row would necessarily be larger than p. So no entry can be larger than $\lceil p/k \rceil$. By a similar argument, none can be smaller than $\lfloor p/k \rfloor$. □

Let $\mathbf{d}$ be an evenly rectangular partition, and let τ be an irreducible representation of $G^{\mathbf{d}}_{\mathrm{red}}$ that is wholly determined by its highest weight. (In other words, if $G^{\mathbf{d}}_{\mathrm{red}}$ is a symplectic group, we admit any τ whatsoever, but if $G^{\mathbf{d}}_{\mathrm{red}}$ is an even orthogonal group, we do not admit τ's whose highest weight has a 0 for the final coordinate.) Consider the following procedure for manipulating rectangular type-C weight diagrams.

1. Start with any diagram $X \in \tilde{D}(\mathbf{d}, \tau)$, and form its severing X^{V}.

2. Apply the type-A algorithm for α to Y, but with E_0 taking the place of E.

3. Form the standard splicing X' of $\alpha(X^{\mathrm{V}})$.

We then have the following preliminary result towards computing γ for $Sp(2n, \mathbb{C})$.

THEOREM 6.3.8. *Assume that Claim 2.3.1 holds. For $Sp(2n, \mathbb{C})$, the map γ for evenly rectangular orbits and for representations determined by their highest weights is given by*

$$\gamma(\mathbf{d}, \tau) = \pi E_C(X'),$$

where X' is the diagram produced in the above procedure.

Proof. Let $\mathbf{d'}$ be the shape-class of the severed diagram. Furthermore, let us abuse notation and let τ refer to any of the following: the given irreducible representation of $G^{\mathbf{d}}_{\mathrm{red}}$; the highest weight of said representation; or the same tuple of integers regarded as a dominant weight for the isotropy group of the $GL(n, \mathbb{C})$-orbit $\mathcal{O}_{\mathbf{d'}}$. We have already established that the moves used in the type-A algorithm still work as described in Chapter 4 when E_A is replaced by E_0, so the diagram $\alpha(X^{\mathrm{V}})$ has the following properties:

$$\alpha X^{\mathrm{V}} \in \tilde{D}(\mathbf{d'}, \tau; A_n)$$
$$\|E_0 \alpha X^{\mathrm{V}}\|^2 \leq \|E_0 Y\|^2 \qquad \text{for any } Y \in \tilde{D}(\mathbf{d'}, \tau; A_n)$$

From the latter of these and the preceding lemma, we obtain

$$\|E_0 \alpha X^{\mathrm{V}}\|^2 \leq \|E_C^{\mathrm{V}} Y\|^2 \qquad \text{for any } Y \in \tilde{D}(\mathbf{d'}, \tau; A_n).$$

Now, the standard choice of coordinates for weights in the $GL(n, \mathbb{C})$ and $Sp(2n, \mathbb{C})$ contexts is such that their norms coincide. In particular, since any diagram in $\tilde{D}(\mathbf{d'}, \tau; A_n)$ arises as the severing of some diagram in $\tilde{D}(\mathbf{d}, \tau; C_n)$, we have

$$\|E_0 \alpha X^{\mathrm{V}}\|^2 \leq \|E_C Y\|^2 \qquad \text{for any } Y \in \tilde{D}(\mathbf{d}, \tau; C_n). \tag{6.2}$$

Of course, we have in particular that $\|E_0 \alpha X^{\mathrm{V}}\|^2 \leq \|E_C X'\|^2$. Claim 2.3.1 tells us that if we actually had

$$\|E_0 \alpha X^{\mathrm{V}}\|^2 = \|E_C X'\|^2, \tag{6.3}$$

then it would follow from (6.2) that $\gamma(\mathbf{d}, \tau) = \pi E_C X'$.

Therefore, it only remains to establish (6.3). Since $\alpha X^{\mathrm{V}} = (X')^{\mathrm{V}}$, this equation would follow if we knew, for instance, that in any given column of X', every dorsal entry was smaller than every ventral entry. That last condition can be obtained by applying Lemma 6.3.7. Since τ is a dominant weight of either a symplectic or orthogonal group, all its coordinates are nonnegative. By Lemma 6.3.7, we have that every entry of $(X')^{\mathrm{V}}$ is nonnegative. Therefore, in X', all the entries in dorsal rows are nonnegative, and all those in ventral rows are nonpositive. Thus, on $(X')^{\mathrm{V}}$, E_C^{V} and E_0 coincide, so (6.3) becomes

$$\|E_C X'\|^2 \leq \|E_C Y\|^2 \qquad \text{for any } Y \in \tilde{D}(\mathbf{d}, \tau, \mathbb{C}_n).$$

The theorem follows. $\square$

APPENDIX A

Algorithms for the General Linear Group

In this appendix, we work out one example each for the two algorithms α and β for $GL(n, \mathbb{C})$. Both examples are set on the orbit labelled $\mathbf{d} = [4, 3, 2, 1^2]$ in $GL(11, \mathbb{C})$. For this orbit, we have

$$G^{\mathbf{d}}_{\text{red}} \simeq GL(1)^3 \times GL(2) \qquad \text{and} \qquad L_{\mathbf{d}} \simeq GL(5) \times GL(3) \times GL(2) \times GL(1).$$

We carry out α for the representation of $G^{\mathbf{d}}_{\text{red}}$ whose highest weight is $\tau = ((15), (14), (9), (4, 4))$; this eventually gives us that

$$\gamma([4, 3, 2, 1^2], ((15), (14), (9), (4, 4))) = (8, 7, 6, 6, 5, 4, 3, 3, 2, 2, 0).$$

Then, we start with the $GL(11)$-weight on the right-hand side of this equation and carry out β to recover the parameters on the left-hand side.

The reader should be aware of one caveat with regard to α. Recall that when α was defined, there was initially some ambiguity as to the order in which moves were to be performed; this was resolved (more or less arbitrarily) by decreeing that moves are to be performed in the order in which they are listed in Table 3.3. However, based on subsequently proved properties of α, it is clear that any sequence of well-behaved moves results in progress towards the final answer. Moreover, even moves that are not well-behaved never do irreparable damage, since all moves preserve $\tilde{D}(\mathbf{d}, \tau)$. In practice, one can sometimes decrease the total number of moves required by performing a cleverly chosen move that is not well-behaved.

In other words, in the α example, we arrive at the answer through some sequence of moves, although we do not strictly follow the prescribed sequence.

	Description	X	EX
1	Fill in $((15), (14), (9), (4, 4))$ in the first column.	$\begin{array}{llll} 15 & 0 & 0 & 0 \\ 14 & 0 & 0 \\ 9 & 0 \\ 4 \\ 4 \end{array}$	$\begin{array}{llll} 19 & 2 & 1 & 0 \\ 16 & 0 & -1 \\ 9 & -2 \\ 2 \\ 0 \end{array}$
2	Perform **A** five times with $s = 1$, $r = 2$, on rows 1–3.	$\begin{array}{llll} 10 & 5 & 0 & 0 \\ 9 & 5 & 0 \\ 4 & 5 \\ 4 \\ 4 \end{array}$	$\begin{array}{llll} 14 & 7 & 1 & 0 \\ 11 & 5 & -1 \\ 4 & 3 \\ 2 \\ 0 \end{array}$
3	Perform **A** three more times with $s = 1$, $r = 2$, on rows 1 and 2.	$\begin{array}{llll} 7 & 8 & 0 & 0 \\ 6 & 8 & 0 \\ 4 & 5 \\ 4 \\ 4 \end{array}$	$\begin{array}{llll} 11 & 10 & 1 & 0 \\ 8 & 8 & -1 \\ 4 & 3 \\ 2 \\ 0 \end{array}$

	Description	X	EX
4	Perform **A** as many times as possible with $s = 1$ or 2 and $r = 3$, first on row 2, and then on row 1.	$\begin{array}{llll} 4&5&6&0\\ 4&5&5&\\ 4&5&&\\ 4&&&\\ 4&&& \end{array}$	$\begin{array}{llll} 8&7&7&0\\ 6&5&4&\\ 4&3&&\\ 2&&&\\ 0&&& \end{array}$
5	Perform **B** on rows 2 and 3.	$\begin{array}{llll} 4&5&6&0\\ 4&5&&\\ 4&5&5&\\ 4&&&\\ 4&&& \end{array}$	$\begin{array}{llll} 8&7&7&0\\ 6&5&&\\ 4&3&4&\\ 2&&&\\ 0&&& \end{array}$
6	Perform **A** once with $s = 3$, $r = 4$ on row 1.	$\begin{array}{llll} 4&5&5&1\\ 4&5&&\\ 4&5&5&\\ 4&&&\\ 4&&& \end{array}$	$\begin{array}{llll} 8&7&6&1\\ 6&5&&\\ 4&3&4&\\ 2&&&\\ 0&&& \end{array}$
7	Perform **B** on rows 1 and 3.	$\begin{array}{llll} 4&5&5&\\ 4&5&&\\ 4&5&5&1\\ 4&&&\\ 4&&& \end{array}$	$\begin{array}{llll} 8&7&6&\\ 6&5&&\\ 4&3&4&1\\ 2&&&\\ 0&&& \end{array}$
8	Perform **A** on row 3, once with $s = 2$ and $r = 4$, and once with $s = 3$ and $r = 4$.	$\begin{array}{llll} 4&5&5&\\ 4&5&&\\ 4&4&4&3\\ 4&&&\\ 4&&& \end{array}$	$\begin{array}{llll} 8&7&6&\\ 6&5&&\\ 4&2&3&3\\ 2&&&\\ 0&&& \end{array}$
9	Perform **B** on rows 3 and 4.	$\begin{array}{llll} 4&5&5&\\ 4&5&&\\ 4&&&\\ 4&4&4&3\\ 4&&& \end{array}$	$\begin{array}{llll} 8&7&6&\\ 6&5&&\\ 4&&&\\ 2&2&3&3\\ 0&&& \end{array}$
10	Perform **B** on rows 1 and 2. The resulting diagram is now in $\tilde{D}^\circ(\mathbf{d})$, so we are finished.	$\begin{array}{llll} 4&5&&\\ 4&5&5&\\ 4&&&\\ 4&4&4&3\\ 4&&& \end{array}$	$\begin{array}{llll} 8&7&&\\ 6&5&6&\\ 4&&&\\ 2&2&3&3\\ 0&&& \end{array}$

From this last step, we read off $(8, 7, 6, 6, 5, 4, 3, 3, 2, 2, 0)$ as the desired weight of $GL(11)$. Now, we start with that weight, and we carry out β to recover a diagram from which we can find obtain the original partition $\mathbf{d}$ and weight of $G^{\mathbf{d}}_{\mathrm{red}}$.

	Description	*Clumps*	*Diagram*
1	We form two clumps.	$A_{1,1} = \{8, 7, 6, 6, 5, 4, 3, 3, 2, 2\}$ $A_{1,2} = \{0\}$	
2	Clump $A_{1,1}$ has even length, and the current column index $r = 1$ is odd, so we take the odd-index values from $A_{1,1}$. $A_{1,2}$ has odd length, so we take its odd-index values as well.	$A_{1,1} = \{\cancel{8}, 7, \cancel{6}, 6, 5, \cancel{4}, 3, 3, \cancel{2}, 2\}$ $A_{1,2} = \{\cancel{0}\}$ $Z_1 = \{8, 6, 4, 2, 0\}$	
3	We arrange the entries of Z_1 in decreasing order to form the first column. After deleting these entries from the $A_{1,\cdot}$'s, we are left with the collection of integers σ_2, which we arrange into two clumps.	$\sigma_2 = \{7, 6, 5, 3, 3, 2\}$ $A_{2,1} = \{7, 6, 5\}$ $A_{2,2} = \{3, 3, 2\}$	$\begin{array}{l} 8 \\ 6 \\ 4 \\ 2 \\ 0 \end{array}$
4	Clump $A_{2.1}$ has odd length, so we take its odd-index values. Clump A_2 has even length, and the current column, $r = 2$, is even, so we take its even-index values.	$A_{2,1} = \{\cancel{7}, 6, \cancel{5}\}$ $A_{2,2} = \{3, 3, \cancel{2}\}$ $Z_2 = \{7, 5, 2\}$	
5	We fill in the elements of Z_2 in the second column so that they are adjacent to first-column entries from which they differ by either 0 or -1. The remaining numbers form σ_3, which splits into two clumps.	$\sigma_3 = \{6, 3, 3\}$ $A_{3,1} = \{6\}$ $A_{3,2} = \{3, 3\}$	$\begin{array}{ll} 8 & 7 \\ 6 & 5 \\ 4 & \\ 2 & 2 \\ 0 & \end{array}$
6	Both clumps have odd length (indeed, length 1), so we take the odd-index values (*i.e.*, the only value) from each clump for Z_3.	$A_{3,1} = \{\cancel{6}\}$ $A_{3,2} = \{\cancel{3}, 3\}$ $Z_3 = \{6, 3\}$	
7	To build column 3, we put the entries of Z_3 next to entries in the second column from which they differ by 0 or $+1$. We also write down σ_4, which is just one clump.	$\sigma_4 = \{3\}$ $A_{4,1} = \{3\}$	$\begin{array}{lll} 8 & 7 & \\ 6 & 5 & 6 \\ 4 & & \\ 2 & 2 & 3 \\ 0 & & \end{array}$
8	Z_4 only has one element, which we position next to a third-column entry from which it differs by 0 or -1. Then σ_5 is empty, so we are finished.	$A_{4,1} = \{\cancel{3}\}$ $Z_4 = \{3\}$ $\sigma_5\{\,\}$	$\begin{array}{llll} 8 & 7 & & \\ 6 & 5 & 6 & \\ 4 & & & \\ 2 & 2 & 3 & 3 \\ 0 & & & \end{array}$

From this last diagram, we compute

$$E^{-1}X = \begin{array}{|c|c|c|c|}\cline{1-2} 4 & 5 \\ \cline{1-3} 4 & 5 & 5 \\ \cline{1-3} 4 \\ \cline{1-4} 4 & 4 & 4 & 3 \\ \cline{1-4} 4 \\ \cline{1-1}\end{array} \qquad \text{and} \qquad \begin{gathered}\mathbf{d} = [4, 3, 2, 1^2] \\ \kappa(E^{-1}X) = ((15), (14), (9), (4, 4)),\end{gathered}$$

as desired.

APPENDIX B

Computed Examples

This appendix gives tables of computed values of γ in several groups of type other than A_n. In particular, inspection of the tables for $Sp(4,\mathbb{C})$, $Sp(6,\mathbb{C})$, and G_2 shows that γ is, in fact, a bijection in each of these cases. Moreover, the values taken by γ on evenly rectangular nilpotent orbits can be seen to coincide with those tentatively claimed in Chapter 6.

A partial computation is of γ is given for F_4. In this group, γ was computed only for representations of the component group of the centralizer of each orbit. This partial computation may give some hints as to what the full γ looks like; moreover, this table lends itself to verification of certain conjectures of Sommers [1].

Table B.1. Computation of γ for $Sp(4,\mathbb{C})$

Orbit label	Reductive part of isotropy group	Jacobson-Morozov Levi subgroup L	Weight of G
$[1^4]$	$Sp(4)$	$Sp(4)$	
	(a,b)	(a,b)	$(a+2,b+2)$
$[2,1^2]$	$O(1)\times Sp(2)$	$GL(1)\times Sp(2)$	
	$\tau\otimes(a)$	$(-a,a)$	$(a+2,1)$
	$\epsilon\otimes(a)$	$(-a-1,a)$	$(a+3,0)$
$[2^2]$	$O(2)$	$GL(2)$	
	τ	$(0,0)$	$(0,1)$
	ϵ	$(0,\pm 1)$	$(2,0)$
	$(a)\oplus(-a)$	$(a,-\frac{a}{2})$	$(0,\frac{a-2}{2}+2)$
		$(a,-\frac{a\pm 1}{2})$	$(1,\frac{a-1}{2}+1)$
$[4]$	$O(1)$	$GL(1)^2$	
	τ	$(0,0)$	$(0,0)$
	ϵ	$(0,1)$	$(1,0)$

Table B.2. Computation of γ for $Sp(6,\mathbb{C})$

Orbit label	Reductive part of isotropy group	Jacobson-Morozov Levi subgroup L	Weight of G
$[1^6]$	$Sp(6)$	$Sp(6)$	
	(a,b,c)	(a,b,c)	$(a+2,b+2,c+2)$
$[2,1^4]$	$O(1)\times Sp(4)$	$GL(1)\times Sp(4)$	
	$\tau\otimes(a,b)$	$(-a-b-2,a,b)$	$(a+2,b+3,0)$
	$\epsilon\otimes(a,b)$	$(-a-b-3,a,b)$	$(a+2,b+2,1)$
$[2^2,1^2]$	$O(2)\times Sp(2)$	$GL(2)\times Sp(2)$	
	$\tau\otimes(b)$	$(0,-b-2,-b)$	$(b+2,2,0)$
	$\epsilon\otimes(b)$	$(0,-b-1,b)$	$(b+3,0,1)$
	$((a)\oplus(-a))\otimes(b)$	$(a,-b-\frac{a}{2}-1,b)$	$(b-(\frac{a}{2}-1)+2,0,(\frac{a}{2}-1)+2)$
			$(0,(\frac{a}{2}-2)-b+2,b+2)$
		$(a,-b-\frac{a-1}{2}-2,b)$	$(b-\frac{a-1}{2}+2,1,\frac{a-1}{2}+1)$
			$(1,(\frac{a-1}{2}-1)-b+1,b+2)$
$[2^3]$	$O(3)$	$GL(3)$	
	$\tau\otimes(a)$	$(\frac{a}{2},\frac{a}{2},-\frac{a}{2})$	$(0,\frac{a}{2}+2,0)$
		$(\frac{a+1}{2},\frac{a-1}{2},-\frac{a+1}{2})$	$(1,\frac{a-1}{2}+1,1)$
	$\epsilon\otimes(0)$	$(0,0,1)$	$(2,0,1)$
	$\epsilon\otimes(a)$	$(\frac{a}{2}-1,\frac{a}{2}+1,-\frac{a}{2})$	$(0,\frac{a-2}{2}+2,1)$
		$(\frac{a+1}{2},\frac{a-1}{2},-\frac{a-1}{2})$	$(1,\frac{a-1}{2}+2,0)$
$[3^2]$	$Sp(2)$	$GL(2)\times Sp(2)$	
	(a)	$(\frac{2a}{3},-\frac{2a}{3},\frac{a}{3})$	$(1,0,\frac{a}{3}+1)$
		$(\frac{2a+1}{3},-\frac{2a+1}{3},\frac{a-1}{3})$	$(0,1,\frac{a-1}{3}+1)$
		$(\frac{2a+2}{3},-\frac{2a-1}{3},\frac{a-2}{3})$	$(0,0,\frac{a-2}{3}+2)$
$[4,1^2]$	$O(1)\times Sp(2)$	$GL(1)^2\times Sp(2)$	
	$\tau\otimes(a)$	$(0,-a-1,a)$	$(a+2,1,0)$
	$\epsilon\otimes(a)$	$(\pm 1,-a-1,a)$	$(a+3,0,0)$
$[4,2]$	$O(1)^2$	$GL(1)\times GL(2)$	
	$\tau\otimes\tau$	$(0,0,0)$	$(0,1,0)$
	$\tau\otimes\epsilon$	$(0,0,1)$	$(1,1,0)$
	$\epsilon\otimes\tau$	$(1,0,0)$	$(0,0,1)$
	$\epsilon\otimes\epsilon$	$(-1,0,1)$	$(2,0,0)$
$[6]$	$O(1)$	$GL(1)^3$	
	τ	$(0,0,0)$	$(0,0,0)$
	ϵ	$(1,0,0)$	$(1,0,0)$

Table B.3. Computation of γ for G_2

Orbit label	Reductive part of isotropy group	Jacobson-Morozov Levi subgroup L	Weight of G
0	G_2	G_2	
	(a,b)	(a,b)	
A_1	A_1	$GL(2)$	
		$\mathfrak{g}_1$-weights: $(2,-3),(1,-1),(0,1),(-1,3)$	
	(a)	$(-\frac{a+2}{2},a)$	$(0,\frac{a}{2}+3)$
		$(-\frac{a+2\pm1}{2},a)$	$(1,\frac{a-1}{2}+2)$
$\widetilde{A}_1$	A_1	$GL(2)$	
		$\mathfrak{g}_1$-weights: $(-1,2),(1,-1)$	
	(a)	$(a,-\frac{3a+1}{2})$	$(\frac{a-1}{2}+2,0)$
		$(a,-\frac{3a+1\pm1}{2})$	$(\frac{a}{2}+1,1)$
$G_2(a_1)$	S_3	$GL(2)$	
	τ	$(0,0)$	$(0,1)$
	ϵ	$(1,0)$	$(0,2)$
	ζ	$(0,1)$	$(1,0)$
G_2	1	$GL(1)^2$	
	τ	$(0,0)$	$(0,0)$

TABLE B.4. Partial Computation of γ in F_4

Orbit label	Component group of isotropy group	Jacobson-Morozov Levi subgroup L	Weight of G
1	1		
	τ	0	$(2,2,2,2)$
A_1	1		
	τ	$-3\omega_1, -4\omega_1$	$(1,1,2,2)$
$\widetilde{A}_1$	S_2		
	τ	$-2\omega_4$	$(2,1,0,3)$
	ϵ	$-\omega_4$	$(2,0,2,2)$
$A_1+\widetilde{A}_1$	1		
	τ	$-\omega_2$	$(2,0,2,0)$
A_2	S_2		
	τ	0	$(0,0,2,2)$
	ϵ	$\pm\omega_1$	$(1,1,0,2)$
$\widetilde{A}_2$	1		
	τ	0	$(2,1,0,1)$
$A_2+\widetilde{A}_1$	1		
	τ	0	$(1,1,0,1)$
B_2	S_2		
	τ	$\omega_1 - 2\omega_4$	$(1,0,0,3)$
	ϵ	$-\omega_4$	$(0,1,0,2)$
$\widetilde{A}_2+A_1$	1		
	τ	$\omega_4-\omega_2$	$(0,1,1,0)$
$C_3(a_1)$	S_2		
	τ	$0, -\omega_1$	$(1,1,0,0)$
	ϵ	$\omega_3-\omega_1$	$(1,0,1,1)$
$F_4(a_3)$	S_4		
	τ	0	$(0,1,0,0)$
	ϵ	ω_2	$(0,0,2,0)$
	σ	ω_1	$(1,0,0,2)$
	η	$\omega_4, \omega_3-\omega_2$	$(1,0,1,0)$
	$\eta\otimes\epsilon$	$\omega_3, \omega_4-\omega_2$	$(0,1,0,1)$
B_3	1		
	τ	0	$(0,0,0,2)$
C_3	1		
	τ	0	$(2,0,0,0)$
$F_4(a_2)$	S_2		
	τ	0	$(0,0,1,0)$
	ϵ	$\omega_2-2\omega_4$	$(1,0,0,1)$
$F_4(a_1)$	S_2		
	τ	0	$(0,0,0,1)$
	ϵ	ω_4	$(1,0,0,0)$
F_4	1		
	τ	0	$(0,0,0,0)$

Bibliography

[1] P. Achar and E. Sommers, *Local systems and functions on orbit covers*, in preparation.

[2] J. Arthur, *Unipotent automorphic representations: conjectures*, Orbites Unipotentes et Représentations II. Groupes p-adiques et réels, Astérisque, vol. 171–172, Société Mathématique de France, 1989, pp. 13–72.

[3] D. Barbasch and D. Vogan, *Unipotent representations of complex semisimple Lie groups*, Ann. of Math. **121** (1985), 41–110.

[4] R. Bezrukavnikov, *Perverse coherent sheaves (after Deligne)*, preprint, math.AG/0005152, 2000.

[5] ——, *Quasi-exceptional sets and equivariant coherent sheaves on the nilpotent cone*, preprint, math.RT/0102039, 2001.

[6] R. Bott, *Homogeneous vector bundles*, Ann. of Math. **66** (1957), 203–248.

[7] D. H. Collingwood and W. M. McGovern, *Nilpotent Orbits in Semisimple Lie Algebras*, Van Nostrand Reinhold Mathematics Series, Van Nostrand Reinhold Co., New York, 1993.

[8] R. Hartshorne, *Algebraic Geometry*, Graduate Texts in Mathematics, no. 52, Springer-Verlag, New York–Heidelberg, 1977.

[9] B. Kostant, *The principal three-dimensional subgroup and the Betti numbers of a complex simple Lie group*, Amer. J. Math. **81** (1959), 973–1032.

[10] G. Lusztig, *Cells in affine Weyl groups. IV*, J. Fac. Sci. Univ. Tokyo Sect. 1A Math. **36** (1989), 297–328.

[11] W. McGovern, *Rings of regular functions on nilpotent orbits and their covers*, Invent. Math. **97** (1989), 209–217.

[12] V. Ostrik, *On the equivariant K-theory of the nilpotent cone*, Represent. Theory **4** (2000), 296–305.

[13] D. Vogan, *Associated varieties and unipotent representations*, Harmonic Analysis on Reductive Groups (Boston) (W. Barker and P. Sally, eds.), Birkhäuser, Boston, 1991, pp. 315–388.

[14] ——, personal communication, 1999.