Geometric Complexity Theory II: Towards explicit obstructions for embeddings among class varieties

Dedicated to Sri Ramakrishna

Ketan D. Mulmuley *
The University of Chicago and I.I.T., Mumbai †
Milind Sohoni
I.I.T., Mumbai

(To appear in SIAM J. Comp.)

December 25, 2006

1 Abstract

In [26], henceforth referred to as Part I, we suggested an approach to the P vs. NP and related lower bound problems in complexity theory through geometric invariant theory. In particular, it reduces the arithmetic (characteristic zero) version of the $NP \not\subseteq P$ conjecture to the problem of showing that a variety associated with the complexity class NP cannot be embedded in a variety associated with the complexity class P. We shall call these *class varieties* associated with the complexity classes P and NP.

This paper develops this approach further, reducing these lower bound problems–which are all nonexistence problems–to some existence problems: specifically to proving existence of *obstructions* to such embeddings among class varieties. It gives two results towards explicit construction of such obstructions.

The first result is a generalization of the Borel-Weil theorem to a class of orbit

*The work of the first author was supported by NSF grant CCR 9800042 and, in part, by Guggenheim Fellowship.

†Visiting faculty member

closures, which include class varieties. The second result is a weaker form of a conjectured analogue of the second fundamental theorem of invariant theory for the class variety associated with the complexity class NC. These results indicate that the fundamental lower bound problems in complexity theory are, in turn, intimately linked with explicit construction problems in algebraic geometry and representation theory.

The results here were announced in [25]. We are grateful to Madhav Nori for his guidence and encouragement during the course of this work, and to C. S. Seshadri and Burt Totaro for helpful discussions.

2 Main results

We shall now state the results precisely. For the sake of completeness, we recall in Section 3 the main results from part I. The rest of this paper is self contained. All groups in this paper are algebraic and the base field is $\mathbb{C}$.

Let G be a connected, reductive group, V its (finite dimensional) linear representation, and $P(V)$ the corresponding projective space. Let $\Delta_V[v]$ denote the projective closure of the G-orbit of v in $P(V)$. It is an almost-homogeneous space in the terminology of [1]. Let $R_V[v]$ be its homogeneous coordinate ring, $I_V[v]$ its ideal, and $R_V[v]_d$, the degree d component of $R_V[v]$.

In Part I, we reduced arithmetic (characteristic zero) implications of the lower bound problems in complexity theory, such as P vs. NP, and NC vs. $P^{\#P}$, to instances of the following (Section 3):

Problem 2.1 *(The orbit closure problem)*

Given explicit points $f, g \in P(V)$, does $f \in \Delta_V[g]$? Equivalently, is $\Delta_V[f] \subseteq \Delta_V[g]$?

The goal is to show that this is not the case in the problems under consideration.

The f's and g's here depend on the complexity classes in the lower bound problem under consideration. In the context of the P vs. NP problem, the point g will correspond to a judiciously chosen P-complete problem, and f to a judiciously chosen NP-complete problem. We call $\Delta_V[g]$ and $\Delta_V[f]$ the *class varieties* associated with the complexity classes P and NP (this terminology was not used in part I). The orbit closure problem in this context is to show that the class variety associated with NP cannot be embedded in a class variety associated with P. We have oversimplified the story here. There is not just one class variety associated with a given complexity class, but a sequence of class varieties depending on the parameters of the lower bound problem under consideration. In the context of the P vs. NP problem, the goal is to show that a class variety for NP associated with a given set of parameters cannot be embedded in the class variety for P associated with the same set of parameters. This would imply that $P \neq NP$ in characteristic zero.

Class variety for the complexity class NC

We give an example of a class variety, associated with the complexity class NC, the class of problems with efficient parallel algorithms. This occurs in the context of NC vs. $P^{\#P}$ problem (Section 3.1). Here we let g be the determinant function, which is a complete function for this class. Specifically, let Y be an $m \times m$ variable matrix, which can also be thought of as a variable l-vector, $l = m^2$. Let $V = \mathrm{Sym}^m(Y)$ be the space of homogeneous forms of degree m in the l variable entries of Y, with the natural action of $G = SL(Y) = SL_l(\mathbb{C})$. Let $g = \det(Y) \in P(V)$ be the determinant form, considered as a point in the projective space. Then $\Delta_V[g]$, the orbit closure of the determinant function, is the class variety associated with NC. This is a basic example of a class variety, which the reader may wish to keep in mind throughout this paper.

For arbitrary f and g, Problem 2.1 is hopeless. But f and g in the preceding lower bound problems can be chosen judiciously, like the determinant function, to have some special properties (cf. Section 3 and part I). To state these properties, we need a few definitions.

Given a point $v \in P(V)$, let $\hat{v} \in V$ denote a nonzero point on the line representing v; the exact choice of $\hat{v}$ will not matter. Let $G_v, G_{\hat{v}} \subseteq G$ denote the stabilizers of v and $\hat{v}$, respectively. We say that v is *characterized* by its stabilizer, if $V^{G_{\hat{v}}}$, the set of points in V stabilized by $G_{\hat{v}}$, is equal to $\mathbb{C}v$, the line in V corresponding to v.

Following Mumford and Kempf [27, 10], we say that v is *stable* if the orbit $G\hat{v} \subseteq V$ is closed, and *semistable* if the closure of this orbit does not contain zero [27, 10]. We say v belongs to the *null cone* if all homogeneous G-invariants of positive degree vanish at $\hat{v}$. We also define a more general notion of *partial stability* which also applies to points in the null cone. A stable point is also partially stable by definition. Now suppose v is not stable. Let S be any closed G-invariant subset of V not containing $\hat{v}$ and meeting the boundary of the orbit $G\hat{v}$. Kempf [10] associates with v and S a canonical parabolic subgroup $P = P[S, v] \subseteq G$, call its canonical destabilizing flag. Let L be its semisimple Levi subgroup. We say that v is *partially stable* with defect zero, or more specifically, (L, P)-stable, if (1) the unipotent radical U of P is contained in G_v, and (2) v is stable with respect to the restricted action of L on V. A more general notion of partial stability allowing nonzero defect is given later (Definition 8.1).

We say that v is *excellent* if

1. it is stable or partially stable with defect zero, and
2. it is characterized by its stabilizer.

If V is an irreducible representation $V_\lambda(G)$ of G, corresponding to a dominant weight λ, then the point in $P(V)$ corresponding to the highest weight vector of G is excellent. This is the simplest example of an excellent point. In this case, the stabilizer G_v is a parabolic subgroup $P = P_\lambda$ of G, and the orbit $Gv \cong G/P$ is closed. Hence $\Delta_V[v] \cong G/P$. The algebraic geometry of G/P has been intensively studied in the literature and is well understood by now; cf. [5, 13] for surveys.

For the lower bound problems under consideration, the points f and g can be chosen so that they are either excellent or almost excellent; the meaning of almost excellent is stated in Section 3. For example, the determinant function above is excellent. In this paper, we shall develop an approach to the orbit closure problem specifically for such f and g. The goal is to understand the orbit closure problem by systematically extending the results for G/P to the (almost) excellent points that arise in this approach.

A natural approach to the orbit closure problem is the following. If f lies in $\Delta_V[g]$, then the embedding $\Delta_V[f] \hookrightarrow \Delta_V[g]$ is G-equivariant. This gives a degree preserving G-equivariant surjection from $R_V[g]$ to $R_V[f]$. Hence, if S is any irreducible representation of G, its multiplicity in $R_V[g]_d$ exceeds that in $R_V[f]_d$, for all d.

Definition 2.2 *We say that S is an* obstruction *for the pair (f, g) if, for some d,*

1. *it occurs in (a complete G-decomposition of) $R_V[f]_d$,*
2. *but not in $R_V[g]_d$.*

Existence of such an S implies that f cannot lie in $\Delta_V[g]$. In a lower bound problem, this S can be considered to be a "witness" to the computational hardness of f.

If S occurs in $R_V[g]_d$, then it is easy to show (Proposition 5.2) that its dual S^* contains a G_g-module isomorphic to $(\mathbb{C}g)^d$, the d-th tensor power of $\mathbb{C}g$. Hence:

Definition 2.3 *We say that S is a* strong obstruction *if, for some d,*

1. *it occurs in $R_V[f]_d$,*
2. *but its dual S^* does not contain a G_g-module isomorphic to $(\mathbb{C}g)^d$.*

A strong obstruction is also an obstruction.

For the (f, g)'s in the lower bound problems under consideration, strong obstructions are conjectured to exist in plenty (Section 4). But to prove their existence it is necessary to construct them more or less explicitly. Otherwise, the proof technique can not cross the natural proof barrier formulated in [28] that any technique for proving the $P \neq NP$ conjecture must cross. Explicit constructions have been used in the theory of computing earlier in different contexts. For example, explicit expanders, needed for efficient pseudo-random generation, have been constructed by Margulis [20], and Lubotzky, Phillips and Sarnak [16]. The essential difference from the situation here is that proving existence of expanders is easy, whereas proving existence of obstructions is itself the main problem.

Hence, we are lead to:

Problem 2.4 *(Explicit Construction of obstructions)*

Given f and g as in Problem 2.1, explicitly construct a (strong) obstruction for the embedding $\Delta_V[f] \hookrightarrow \Delta_V[g]$.

In the orbit closure problems under consideration, $H = G_g$ turns out to be a reductive subgroup of G. Hence, to solve Problem 2.4, we have to solve the following problems first.

Problem 2.5 *(Subgroup restriction problem)*

Let H be a reductive subgroup of a connected, reductive group G. Find an explicit decomposition a given irreducible G-representation S as an H-module.

This arises in the context of the second condition in Definition 2.3.

Problem 2.5, with H equal to the the stabilizer of the determinant function considered earlier, turns out to be equivalent to the Kronecker problem of finding an explicit decomposition of the tensor product of two irreducible representations of the symmetric group; cf. Section 3. This is an outstanding problem in the representation theory of the symmetric group [19, 5]. Other specific instances of Problem 2.5 that arise in the lower bound problems under consideration (cf. Section 3) include the well known plethysm problem [19, 5], which is an outstanding problem in the representation theory of $GL_n(\mathbb{C})$.

Problem 2.6 *(Problem in geometric invariant theory)*

Let $v \in P(V)$ be an (almost) excellent point.

Find an explicit decomposition of $R_V[v]_d$, for a given d, as a G-module.

This is needed in the context of both conditions in Definition 2.2. For this, it is desirable to solve the following problem first:

Problem 2.7 *(SFT problem)*

Let $v \in P(V)$ be an (almost) excellent point. Find an explicit set of generators for the ideal $I_V[v]$ of $\Delta_V[v]$ with good representation theoretic properties.

Problems 2.6 and 2.7 are intractable for general v's. Hence, specialization to almost excellent v's is necessary. Some additional reasonable restrictions may be necessary in these problems.

When $V = V_\lambda(G)$, v the point corresponding to the highest weight vector of $V_\lambda(G)$, and $\Delta_V[v] \cong G/P$, the second fundamental theorem (SFT) of invariant theory for G/P [13], answers Problem 2.7. By the Borel-Weil theorem for G/P [5], $R_V[v]_d = V_{d\lambda}(G)^*$. This answers Problem 2.6.

What is desired is a generalization of these results for G/P to the class varieties $\Delta_V[v]$, for the (almost) excellent v's under consideration. Before we go any further, let us point out the main difference between G/P and the class varieties:

1. Luna and vust [18] have assigned a complexity to orbit closures, which measures the complexity of their algebraic geometry. All orbit closures whose algebraic geometry has been well understood have low Luna-Vust complexity–close to zero. For example, the Luna-Vust complexity of G/P is zero. In contrast, the Luna-Vust complexity of a class variety can be polynomial in the parameters in the lower bound problem under consideration.

2. The analogue of the subgroup restriction problem (Problem 2.5), with H being the parabolic stabilizer P of the highest weight vector in $V_\lambda(G)$, is trivial. In contrast, the instances of Problem 2.5 in the context of the class varieties include the nontrivial Kronecker and plethysm problems.

This indicates that the algebraic geometry of class varieties is substantially more complex than that of G/P. For this reason, we cannot expect a full solution to Problems 2.6 and 2.7 until the outstanding Problem 2.5 in representation theory is resolved. Rather, our goal is to connect Problems 2.6 and 2.7 with the "easier" Problem 2.5 for the almost excellent v's under consideration. We prove two results in this direction.

Let us begin by considering a weaker form of Problem 2.6; i.e., we only ask which G-modules can occur in $R_V[v]$, without worrying about $R_V[v]_d$ for a specific d. This is addressed by the following result.

We call a G-module $V_\lambda(G)$ $G_{\hat{v}}$-admissible if it contains a $G_{\hat{v}}$-invariant (cf. Definition 5.1).

Theorem 2.8 *(Borel-Weil for orbit closures of partially stable points)*

Let V be a (finite dimensional) linear representation of a connected, reductive G.

(a) If $v \in P(V)$ is stable, an irreducible G-module $V_\lambda(G)$ with weight λ can occur in $R_V[v]$ iff $V_\lambda(G)$ is $G_{\hat{v}}$-admissible.

(b) Suppose v is partially stable with defect zero, specifically (L,P)-stable, as defined above. Let $S_V[v]$ be the homogeneous coordinate ring of the projective closure in $P(V)$ of the L orbit of v. Then the G-module structure of $R_V[v]$ is completely determined by the L-module structure of $S_V[v]$. A weaker statement holds a partially stable point of nonzero defect as defined in Section 8.

A precise statement of (b) is given in Section 9. We actually prove a stronger result (Theorem 9.2) that specializes to the Borel-Weil theorem [11] when v corresponds to the highest weight vector of an irreducible representation $V = V_\lambda(G)$ of a semisimple G.

When the defect is nonzero, Theorem2.8 (b) does not tell precisely which irreducible G-modules occur in $R_V[v]$ if we only knew which irreducible L-modules occur in $S_V[v]$ as a whole. But it gives a good information on this and also on which irreducible G-modules occur in $R_V[v]_d$, for a given d, provided we know precisely which irreducible L-modules occur in every degree d-component $S_V[v]_d$; this is Problem 2.6 for a stable v, with L playing the role of G.

Now we turn to the actual Problem 2.6. For this, we have to understand Problem 2.7 first. We turn to this problem next.

Let v be an excellent point. We associate with it a representation-theoretic data $\Pi_v = \cup_d \Pi_v(d)$ (cf. Definitions 7.1 and 11.1). If v is stable, $\Pi_v(d)$ is just the set of all irreducible G-submodules of $\mathbb{C}[V]$ whose duals do not contain a G_v-submodule isomorphic $(\mathbb{C}v)^d$.

Then the ideal $I_V[v]$ contains all modules in Π_v (Proposition 5.2). Let $X(\Pi_v)$ be the variety (scheme) defined by the ideal generated by the modules in Π_v. It follows that $\Delta_V[v] \subseteq X(\Pi_v)$.

Now we ask:

Question 2.9 *Suppose v is excellent. Is $X(\Pi_v) = \Delta_V[v]$ as a variety, or more strongly, as a scheme?*

The scheme theoretic equality means that the ideal $I_V[v]$ of $\Delta_V[v]$ is generated by the modules in Π_v.

If v is stable, then G_v is reductive [2, 24]. Hence, the G-modules contained in Π_v are precisely determined once we know answer to Problem 2.5, with $H = G_v$. This turns out to be so even for the partially stable v's that arise in the lower bound problems, by letting H be the reductive part of G_v. Hence, if the answer to Question 2.9 is yes, the algebraic geometry of $\Delta_V[v]$ is completely determined by the representation theory of the pair (G_v, G), and hence, Problems 2.6 and 2.7 are intimately related to Problem 2.5. Clearly, this can happen only for very special v's. The answer need not be yes even for a general excellent v.

When v corresponds to the highest weight vector of $V_\lambda(G)$, so that $\Delta_V[v] = G/P$, answer to Question 2.9 is yes. This follows from the second fundamental theorem (SFT) for G/P [13] (cf. Section 11.1).

We conjecture that this is also the case for the class variety associated with the complexity class NC described above.

Conjecture 2.10 *(Second fundamental theorem (SFT) for the orbit closure of the determinant)*

Let $\Delta_V[v]$ be the class variety associated with the complexity class NC–the orbit closure of the determinant function.

Then $X(\Pi_v) = \Delta_V[v]$ as a variety, and more strongly, as a scheme.

This is expected because of the very special nature of the determinant function. We have already remarked that it is excellent. Furthermore, its stabilizer has an additional conjectural property called G-separability (Definition 7.3). For analogous conjectures for other almost excellent class varieties, one has to address complications caused by almost excellence instead of full excellence. This is possible, and will be done elsewhere.

The following general result implies a weaker form of Conjecture 2.10 when v is the determinant function.

Theorem 2.11 *(Second Fundamental Theorem (SFT) for the orbit of an excellent point)*

Suppose V is a linear representation of a connected, reductive group G, and $v \in P(V)$ an excellent point.

(a) Suppose v is stable. Furthermore, assume that the stabilizer $G_{\hat{v}}$ is G-separable (cf. Definition 7.3). Then the orbit $Gv \subseteq P(V)$ is determined by the representation-theoretic data Π_v within some G-invariant neighbourhood U: i.e.,

$$Gv = \Delta_V[v] \cap U = X(\Pi_v) \cap U,$$

as schemes.

(b) A generalized result also holds for the G-orbit of a partially stable, excellent point with defect zero.

This follows from a stronger result proved in Section 7 (stable case) and Section 12 (partially stable case).

When v corresponds to the highest weight vector in $V_\lambda(G)$, Theorem 2.11 (b), after some strengthening (cf. Section 11.1), becomes the second fundamental theorem for G/P [13]–hence the terminology.

The rest of this paper is organized as follows. In section 3, we describe how the orbit closure problem arises in complexity theory, and summarize the relevant results from part I. In Section 5 we prove some basic propositions based on the notion of admissibility. The stable case of Theorem 2.8 is proved in Section 6. The stable case of Theorem 2.11 is proved in Section 7. The stable cases illustrate the main ideas in this paper. The notion of partial stability is introduced in Section 8. The partially stable case of Theorem 2.8 is proved in Section 9. Its specialization in the context of complexity theory is given in Section 10. The partially stable case of Theorem 2.11 is proved in Section 12. Conjectural G-separability of the stabilizer of the determinant is proved in Section 13 for a special case.

Notation

We let G denote a connected reductive group. An irreducible G-representation with highest weight λ will be denoted by $V_\lambda(G)$. We say that $V_\lambda(G)$ occurs in a G-module M, or that M contains $V_\lambda(G)$, if a complete decomposition of M into G-irreducibles contains a copy of $V_\lambda(G)$. We denote the dual of M by M^*. We always denote a Levi decomposition of a parabolic subgroup $P \subseteq G$ in the form $P = TLU = KU$, where T is a torus, L is a semisimple Levi subgroup, $K = TL$ is a reductive Levi subgroup, and U is the unipotent radical. The root system of K is a subsystem of that of G. Hence a dominant weight of G can be assumed to be a dominant weight of K by restriction.

3 The orbit closure problem

In this section we describe the orbit closure problem that arises in complexity theory, and the related results; cf. Part I for details and proofs.

Let $Y = [y_0, \cdots, y_{l-1}]$ denote a variable l-vector. For $k < l$, let $X = [y_1, \cdots, y_k]$, and $\bar{X} = [y_0, \cdots, y_k]$ be its subvectors of size k and $k+1$. Let $V = \mathrm{Sym}^m(Y) = \mathrm{Sym}^m((\mathbb{C}^l)^*)$ be the space of homogeneous forms of degree m in the l variable-entries of Y, with the natural action of $G = SL(Y) = SL_l(\mathbb{C})$, and $\hat{G} = GL(Y) = GL_l(\mathbb{C})$.

Let $W = \mathrm{Sym}^n(X)$, $n < m$, be the representation of $GL(X) = GL_k(\mathbb{C})$. We have a natural embedding $\phi : W \to V$, which maps any $w \in W$ to $y^{m-n}w$, where $y = y_0$ is used as the homogenizing variable. The image $\phi(W)$ is contained in $\bar{W} = \mathrm{Sym}^m(\bar{X})$, a representation of $GL(\bar{X}) = GL_{k+1}(\mathbb{C})$.

Definition 3.1 *We say that $f = \phi(h)$ is partially stable with respect to the action of G if $h \in P(W)$ is stable with respect to the action of $SL_k(\mathbb{C})$.*

These are the only kinds of partially stable points that arise in the context of complexity theory. If the reader wishes, he may confine himself to only these kinds. When we introduce a more general definition of partial stability (Section 8), it will turn out that f is partially stable with defect one. In contrast, the (L, P)-stable points in the introduction will turn out to be partially stable points with defect zero. Note that f in Definition 3.1 belongs to the null cone of the G-action–this follows easily from the Hilbert-Mumford criterion [27].

The orbit closure problems (Problem 2.1) that arise in complexity theory (cf. Part I) have the following form:

Problem 3.2 *Given fixed forms $g \in P(V)$, and $h \in P(W) \overset{\phi}{\hookrightarrow} P(V)$, does $f = \phi(h)$ belong to $\Delta_V[g]$? That is, is $\Delta_V[f] \subseteq \Delta_V[g]$?*

The goal is to show that the specific f does not belong to $\Delta_V[g]$. The specific f and g depend on the lower bound problem under consideration, and will be either excellent (cf. Section 1), or almost excellent–the latter means: (1) the defect of partial stability may not be zero, but will be small, and (2) the point may not be fully characterized by the stabilizer, but almost (as explained in part I).

The following are two instances of the orbit closure problem that arise in complexity theory.

3.1 Arithmetic version of the NC vs. $P^{\#P}$ conjecture

In concrete terms, this says that the permanent of an $n \times n$ matrix cannot be computed by an integral circuit of depth $\log^c n$, for any constant $c > 0$ [31].

The class varieties in this context are as follows. Let Y be an $m \times m$ variable matrix, which can also be thought of as a variable l-vector, $l = m^2$. Let X be its, say, principal bottom-right $n \times n$ submatrix, $n < m$, which can be thought of as a variable k-vector, $k = n^2$. We use any entry y of Y not in X as the homogenizing variable for embedding $W = \mathrm{Sym}^n(X)$ in $V = \mathrm{Sym}^m(Y)$. Let $g = \det(Y) \in P(V)$ be the determinant form (which will also be considered as a point in the projective space), and $f = \phi(h)$, where $h = \mathrm{perm}(X) \in P(W)$. Then $\Delta_V[g]$ is the class variety associated with NC and $\Delta_V[f]$ the class variety associated with $P^{\#P}$. These depend on the lower bound parameters n and m. If we wish to make these implicit, we should write $\Delta_V[f, n, m]$ and $\Delta_V[g, m]$ instead of $\Delta_V[f]$ and $\Delta_V[g]$.

It is conjectured in part I that, if $m = 2^{O(\mathrm{polylog} n)}$ and $n \to \infty$, then $f \notin$

$\Delta_V[g]$; i.e., the class variety $\Delta_V[f, n, m]$ cannot be embedded in the class variety $\Delta_V[g, m]$. This implies the arithmetic form of the $NC \neq P^{\#P}$ conjecture.

The following result provides the connection with geometric invariant theory.

Theorem 3.3 *(cf. Part I) The point $h = perm(X) \in P(W)$ is stable with respect to the action of $SL(X) = SL_k(\mathbb{C})$ on $P(W)$ (thinking of X as a k-vector). Hence the point $f = \phi(h) \in P(V)$ is partially stable (definition 3.1) with respect to the action of $G = SL(Y) = SL_l(\mathbb{C})$, as well as $\hat{G} = GL(Y) = GL_l(\mathbb{C})$.*

Similarly, $g = \det(Y) \in P(V)$ is stable with respect to the action of G on $P(V)$, thinking of Y as an l-vector on which $SL_l(\mathbb{C})$ acts in the usual way.

Moreover, both $\mathrm{perm}(X) \in P(W)$ and $\det(Y) \in P(V)$ are characterized by their stabilizers. Hence, both h and g are excellent. But, in contrast, $f = \phi(h)$ is only almost excellent–because its defect of partial stability is one.

The stabilizer of $\det(Y)$ in $G = SL_{m^2}(\mathbb{C})$ consists of linear transformations of the form $Y \to AY^*B^{-1}$, thinking of Y as an $m \times m$ matrix, where Y^* is either Y or Y^T, $A, B \in GL_m(\mathbb{C})$. The stabilizer of $\mathrm{perm}(X)$ in $SL_{n^2}(\mathbb{C})$ is generated [22] by linear transformations of the form $X \to \lambda X \mu^{-1}$, thinking of X as an $n \times n$ matrix, where λ and μ are either diagonal or permutation matrices.

Let $H \subseteq G = SL_{m^2}(\mathbb{C})$ be the stabilizer of $\det(Y)$. Since $SL_m(\mathbb{C}) \times SL_m(\mathbb{C})$ is a subgroup of H, the subgroup restriction problem (Problem 2.5) in this context becomes:

Problem 3.4 (Kronecker problem)

Given a partition λ of height m^2, find an explicit decomposition of $V_\lambda(G)$ as an $SL_m(\mathbb{C}) \times SL_m(\mathbb{C})$-module:

$$V_\lambda(G) = \oplus_{\alpha,\beta} k^\lambda_{\alpha,\beta} V_\alpha(SL_m(\mathbb{C})) \otimes V_\beta(SL_m(\mathbb{C})),$$

where α, β range over partitions of height at most m.

The coefficients $k^\lambda_{\alpha,\beta}$'s here are the same as the Kronecker coefficients that arise in the internal product of Schur functions. The problem of decomposing the tensor product of two irreducible representations of the symmetric group S_m can be reduced to this problem [5]. This is one of the outstanding problems in the representation theory of symmetric groups.

3.2 Arithmetic (nonuniform) version of the $P \neq NP$ conjecture

This is a version of the usual $P \neq NP$ conjecture (the nonuniform version), which does not involve problems of positive characteristic, and hence, is addressed first.

Now h, g in the orbit closure problem (Problem 3.2) correspond to some integral functions that are NP-complete and P-complete, respectively. These functions have to be chosen judiciously, because most functions that arise in complexity theory, e.g. the one associated with the travelling salesman problem, do not have a nice stabilizer, as required in our approach. For a detailed definition of h and g, see part I. We shall call $\Delta_V[f]$, $f = \phi(h)$, and $\Delta[g]$ for the specific h and g here the *class varieties* associated with the complexity classes NP and P. The conjecture that $NP \not\subseteq P$ in characteristic zero is then reduced to the problem of showing that the class variety $\Delta_V[f]$ associated with NP cannot be embedded in the class variety $\Delta_V[g]$ associated with P, for the parameters of the lower bound problem under consideration.

The following is an analogue of Theorem 3.3 in this context.

Theorem 3.5 *The point $h \in P(W)$, for a suitable W, which corresponds to an NP-complete function as in [26], is stable with respect to the action of $SL(W)$ on $P(W)$. Hence, the point $f = \phi(h)$ is partially stable.*

The h here is not completely characterized by its stabilizer, but almost so; cf. part I. Hence it is almost excellent. The subgroup restriction problem Problem 2.5 that arises for the stabilizer of h is essentially the well known plethysm problem [5] in the theory of symmetric functions.

4 Why should obstructions exist?

Before we go any further, we have to argue why obstructions should exist for the pairs (f, g) that arise in the lower bound problems under consideration.

Let us begin with an observation that for an orbit closure problem that arises in complexity theory, an obstruction for the pair (f, g) cannot exist if l is sufficiently larger than k. For example, let $(f, g) = (\phi(h), g))$, where $h = \mathrm{perm}(X)$ and $g = \det(Y)$, as in Section 3.1. Then there cannot be any obstruction for $m > n!$, or for that matter, $m > 2^{cn}$ for a large enough constant c. This is because $\mathrm{perm}(X)$ has a formula of size 2^{cn} for a large enough $c > 0$ [22] (the usual formula is of size $n!$) and hence $f \in \Delta_V[g]$, for $m > 2^{cn}$ (cf. Part I).

At the other extreme, when $l = k$, so that f is a stable point of V, it follows from the étale slice theorem [27, 17] that, if $f \in \Delta_V[g]$, then some conjugate of the stabilizer of f must be contained in the stabilizer of g (cf. Part I). This will not happen for our judiciously chosen f and g. For example, when f and g are the permanent and the determinant and $m = n$–in fact, in this case, there are infinitely many obstructions to this containment (cf. part I).

The goal is to understand the transition between these two extremes.

First, let us consider the arithmetic implication of the $P^{\#P} \neq NC$ conjecture. Let $g = \det(Y)$, $f = \phi(h)$, and $h = \text{perm}(X)$ as in Section 3.1.

Proposition 4.1 *Suppose $h = perm(X)$ cannot be approximated infinitesimally closely by a circuit of depth $O(\log^c n)$, where $c > 0$ is a constant, and $n \to \infty$. Suppose $X(\Pi_g) = \Delta_V[g]$ as varieties (cf. Conjecture 2.10). Then there exists a strong obstruction for the pair (f, g), for $m \leq 2^{\log^{c/2} n}$*

Proof: It is proved in Part I that the hypothesis implies that $f \notin \Delta_V[g]$ if $m \leq 2^{\log^{c/2} n}$. Assuming $X(\Pi_g) = \Delta_V[g]$, this means $f \notin X(\Pi_g)$. Hence there exists a G-module $S \in \Pi_g$ which does not vanish on f, and hence on its orbit. So S occurs in $R_v[f]$. By the definition of Π_g, the dual S^* does not contain a G_g-module isomorphic to $(\mathbb{C}g)^d$. Hence S is a strong obstruction for the pair (f, g). Q.E.D.

Since $\text{perm}(X)$ is $\#P$-complete [31], it is not expected to have infinitesimally close approximations by circuits of $O(log^c(n))$ depth, for any constant $c > 0$. Hence, Proposition 4.1 leads to:

Conjecture 4.2 *There exist (infinitely many) strong obstructions for $(f, g) = (\phi(h), g)$, $g = \det(Y)$, $h = perm(X)$, if $m = 2^{\log^c n}$, c a constant, and $n \to \infty$.*

In turn, this conjecture implies $f \notin \Delta_V[g]$, and hence, the arithmetic implication of the $P^{\#P} \neq NC$ conjecture (Section 3.1).

In the same vein, we also make:

Conjecture 4.3 *There exist (infinitely many) obstructions for $(f, g) = (\phi(h), g)$, that occur in the context of the P vs. NP problem, if $m = poly(n)$, and $n \to \infty$, where n denotes the input size parameter, and m denotes the circuit size parameter in the nonuniform version of the P vs. NP problem.*

This would imply $f \notin \Delta_V[g]$, and hence, the arithmetic implication of the $P \neq NP$ conjecture in Section 3.2. This conjecture is motivated by similar considerations as in Proposition 4.1. The g that occurs in the context of the P vs. NP problem is not fully characterized by its stabilizer. But it is still determined by its stabilizer to a large extent. Hence, similar considerations apply.

5 Admissibility

In this section, we introduce a basic notion of admissibility and study how it influences which G-modules may appear in the homogeneous coordinate ring

$R_V[v]$ of the projective-orbit closure of a point $v \in V$. The basic propositions proved here will be useful in the proofs of the main results.

Definition 5.1 *Given a reductive subgroup $H \subseteq G$ and an H-module W, we say that a G-module M is (H, W)-admissible, if some irreducible H-submodule of M occurs in W.*

We say that M is H-admissible if it is $(H, 1_H)$-admissible, where 1_H is the trivial H-module; i.e., if it contains a (nonzero) H-invariant.

For general H, not necessarily reductive, we say that M is H-admissible if M^* contains an H-invariant.

If H is reductive, M contains an H-invariant, iff M^* does–this follows from Weyl's result on complete reducibility of a reductive group representation–and hence, the second and third statement are then equivalent.

Given a G-module S, and a subgroup $H \subseteq G$, not necessarily reductive, we shall say that S has an H-coinvariant if S is H-admissible, i.e., the dual module S^* has an H-invariant (cf. Definition 5.1).

Let $h \in P(V)$ be any point, not necessarily stable. Let $\mathbb{C}h$ be the corresponding line in V. It is one-dimensional, i.e., a character, as a G_h-module , and trivial as a $G_{\hat{h}}$-module. Let $\check{\Delta}[h] = \check{\Delta}_V[h] \subseteq V$ denote the affine cone of the projective-orbit-closure $\Delta[h]$. Its coordinate ring $\mathbb{C}[\check{\Delta}[h]]$ coincides with the homogeneous coordinate ring $R[h] = R_V[h]$ of $\Delta[h]$. Since the G-action is degree preserving, each homogeneous component $R[h]_d$ is a finite dimensional G-module.

Proposition 5.2 *Let V be a linear representation of a reductive group G. Let $h \in P(V)$ be any point, not necessarily stable, with stabilizer $G_h \subseteq G$. Let S be any irreducible G-module occurring in $R[h]$–that is, in $R[h]_d$ for some d. Then the dual module S^* must contain a G_h-submodule isomorphic to $(\mathbb{C}h)^d$, and hence both S and S^* are $G_{\hat{h}}$-admissible. In particular, a G-module $S \subseteq \mathbb{C}[V]_d$ not satisfying this constraint belongs to the ideal $I_V[h]$.*

Similarly, given an algebraic subgroup $H \subseteq G$, and an H-module M, let $B = G \times_H M$ be the induced bundle [27] with base space G/H and fibre M. Let N be any irreducible G-submodule of $\Gamma(G/H, B)$, the space of global sections of B. Then the G-module $Hom(N, M)$ must contain a nonzero H-invariant.

Proof: Not all functions in S can vanish at $\hat{h}$: Otherwise, they will vanish identically on the G-orbit of $\hat{h}$ in V, and so also on its cone, since the functions in S are homogeneous. But the cone of the affine G-orbit of h is dense in $\check{\Delta}[h]$. Hence, it would follow that the functions in S vanish on $\check{\Delta}[h]$ identically; a contradiction.

Consider the G_h-equivariant map $\phi : S \to ((\mathbb{C}h)^*)^d = (\mathbb{C}h^d)^*$ that maps every function in S to its restriction on the line $\mathbb{C}h$. It follows that this evaluation map is nonzero. Hence the dual map ϕ injects the G_h-module $\mathbb{C}h^d$ into S^*.

The argument extends to the vector bundle B by considering instead the evaluation map $\phi : N \to M$ at the base point $e \in G/H$, which must be nonzero and H-equivariant; i.e., $\phi \in \mathrm{Hom}(N, M)^H$. Q.E.D.

6 Admissibility and stability

In this section we shall prove the first statement of Theorem 2.8 concerning stable points.

Proposition 6.1 *Let $h \in P(V)$ be a point such that the stabilizer $H = G_{\hat{h}}$ of $\hat{h} \in V$ is reductive. Then every irreducible G-module occurring in $R[h]$ must be H-admissible; i.e., must contain a nonzero H-invariant.*

If $G_{\hat{h}}$ is not reductive, this still holds if H is any reductive subgroup of $G_{\hat{h}}$.

Proof: If H is reductive, then Weyl's theorem on complete decomposibility of H-modules into irreducibles implies that the existence of an H-invariant is equivalent to existence of an H-coinvariant. Hence this follows from Proposition 5.2. Q.E.D.

Conversely,

Proposition 6.2 *Suppose $h \in P(V)$ is stable. Then every H-admissible, irreducible G-module occurs in $R[h]$.*

Proof: Since h is stable, the stabilizer $H = G_{\hat{h}}$ is reductive [3, 24], and the orbit $G\hat{h} \subseteq V$ is affine and isomorphic to G/H [27]. Moreover, an explicit G-module decomposition of the coordinate ring $\mathbb{C}[G\hat{h}] = \mathbb{C}[G/H]$ can be computed as follows. First, we recall (cf. Page 48, [30]) the algebraic version of the

Peter-Weyl Theorem:

$$\mathbb{C}[G] = \oplus_S S \otimes S^*, \tag{1}$$

where S ranges over all irreducible G-modules, and S^* is the dual module. From this it follows that

$$\mathbb{C}[G/H] = \oplus_S S \otimes (S^*)^H, \tag{2}$$

where $(S^*)^H$ denotes the subspace of H-invariants in S^*. Since h is stable, the affine orbit $G\hat{h}$ is closed in V. So it is a closed G-subvariety of the cone $\check{\Delta}[h] \subseteq V$, which is also a G-variety. It follows that there is a G-equivariant surjection from

$R[h]$ to $\mathbb{C}[G\hat{h}] = \mathbb{C}[G/H]$. Both $R[h]$ and $\mathbb{C}[G/H]$ have direct sum decompositions into finite dimensional G-modules. It follows that every irreducible G-module that occurs in $\mathbb{C}[G/H]$ must occur in $R[h]$. But by the Peter-Weyl theorem, i.e., eq(2), the irreducible G-modules that appear within $\mathbb{C}[G/H]$ are precisely the H-admissible ones. Q.E.D.

Proof of Theorem 2.8 (a): Since $v \in P(V)$ is stable, $G_{\hat{v}}$ is reductive [3, 24]. Hence this follows from Propositions 6.1 and 6.2. Q.E.D.

7 SFT for the orbit of a stable, excellent point

In this section we shall now prove Theorem 2.11 for stable points. To give its precise statement, we need a few definitions.

We associate with a stable point v representation-theoretic data Π_v and $\Sigma_v \subseteq \Pi_v$ as follows.

Definition 7.1 *Suppose $v \in P(V)$ is stable.*

Let Σ_v be the set of all non-$G_{\hat{v}}$-admissible G-submodules of $\mathbb{C}[V]$– here $G_{\hat{v}}$ is necessarily reductive [3, 24].

Let $\Pi_v = \cup_d \Pi_v(d)$, where $\Pi_v(d)$ is the set of all irreducible G-submodules of $\mathbb{C}[V]$ whose duals do not contain a G_v-submodule isomorphic $(\mathbb{C}v)^d$–the dth tensor power of $\mathbb{C}v$.

Clearly $\Sigma_v \subseteq \Pi_v$. Basis elements (suitably chosen) of the G-submodules of Σ_v will be called *nonadmissible basis elements.*

Proposition 7.2 *If v is stable, the G-modules in the representation-theoretic data Π_v, and hence Σ_v, associated with v are contained in $I_V[v]$.*

This follows from Proposition 5.2.

Definition 7.3 *Given a reductive $H \subseteq G$, we say that a nontrivial, irreducible H-module L, which occurs in some G-module, is G-separable (from the trivial H-module) if there exists an irreducible non-H-admissible G-module M that contains L; we say it is strongly G-separable if there exist infinitely many such G-modules. We shall say that a subgroup $H \subseteq G$ is G-separable (strongly G-separable), if every nontrivial irreducible H-module, which occurs in some G-module, is G-separable (resp. strongly G-separable).*

For example, $SL_k(\mathbb{C}) \subseteq SL_n(\mathbb{C})$, $k > n/2$, and a semisimple $H \subseteq H \times H$ (diagonal embedding) are separable (Proposition 13.1). We conjecture that

$SL_n(\mathbb{C}) \times SL_n(\mathbb{C}) \subseteq SL(\mathbb{C}^n \otimes \mathbb{C}^n) = SL_{n^2}(\mathbb{C})$ is separable, and prove this for $n = 2$ (Proposition 13.6). We also conjecture that the stabilizers of the permanent, the determinant and other functions that arise in our lower bound applications are G-separable; the stabilizer of the determinant is very similar to the subgroup $SL_n(\mathbb{C}) \times SL_n(\mathbb{C}) \subseteq SL_{n^2}(\mathbb{C})$ above (cf. Section 3.1).

A precise statement of Theorem 2.11 now reads as follows:

Theorem 7.4 *Suppose V is a linear representation of a connected, reductive group G. Let $v \in P(V)$ be a stable point such that stabilizer $G_{\hat{v}}$ is G-separable (cf. Definition 7.3) and characterizes v.*

Then there exists a homogeneous G-invariant $\beta \in \mathbb{C}[V]$ not vanishing at v such that the ideal of Gv as a closed subvariety of the open neighbourhood $U = V_\beta = V \setminus \{\beta = 0\}$ is generated by the nonadmissible basis elements in Σ_v–in fact, it is generated by the basis of less than codim$(Gv, P(V))$) irreducible, non-$G_{\hat{v}}$-admissible G-submodules of $\mathbb{C}[V]$.

Remark: Since $\Sigma_v \subseteq \Pi_v$, this statement is slightly stronger than Theorem 2.11.

Theorem 7.4, in turn, follows from the following stronger result.

Let X be a nonsingular, affine G-variety, G a connected reductive group. Given a point $x \in X$, we shall denote by $[x] \subseteq X$ the subvariety consisting of all points in X whose stabilizers contain $H = G_x$, the stabilizer of x. Assume that x is a nonsingular point of $G \cdot [x]$; when the orbit $Gx \subseteq X$ is closed, this is automatically so, because of the étale slice theorem [17] (cf. the proof of Lemma 7.7 below.) We shall denote by N_x (resp. $N_{[x]}$) the H-module that is an H-complement of the tangent space of $G \cdot x$ (resp. $G \cdot [x]$) at x in the total tangent space to X at x; it can be thought of as the "normal" space to $G \cdot x$. (resp. $G \cdot [x]$) at x. $N_{[x]}$ is the H-submodule of N_x consisting of all nontrivial H-submodules of N_x.

Given a G-invariant $\beta \in \mathbb{C}[X]$, we shall denote by $X(\beta)$ the G-variety obtained from X by removing the divisor $\{\beta = 0\}$.

We shall denote the codimension of a subvariety $Y \subseteq X$ by codim(Y, X). We say that an open subset $U \subseteq X$ is saturated if its of the form $\psi^{-1}(U')$, where ψ is the projection from X to X/G and U' is an open subset of X/G.

Theorem 7.4 for stable points in $P(V)$ follows from the following result by letting $X = V$ and $x = \hat{v}$. When $v \in P(V)$ is characterized by the stabilizer $G_{\hat{v}}$, $[x] = \mathbb{C}v$ Passage from V to $P(V)$ is possible because the nonadmissible basis elements are homogeneous.

Theorem 7.5 *Assume that G is a connected, reductive group, and X an affine, nonsingular, irreducible G-variety X. Let $x \in X$ be a point, with stabilizer $H =$*

G_x, whose orbit Gx is closed. Suppose every H-module L that appears in $N^*_{[x]}$ is G-separable (Definition 7.3). Then, for some G-invariant $\beta \in \mathbb{C}[X]$ not vanishing at x, and non-H-admissible, irreducible G-submodules $P_i \subseteq \mathbb{C}[X]$, $1 \leq i \leq r$, with $r < codim(G \cdot [x], X)$, $Spec(\mathbb{C}[X]/J) \cap X(\beta) = G \cdot [x] \cap X(\beta)$, where J denotes the ideal generated by the P_is.

(Here we are identifying a variety with the corresponding reduced scheme supported by it.)

Proof: By Proposition 5.2, or rather its proof, the functions in every non-H-admissible P within $\mathbb{C}[X]$ must vanish on $G \cdot [x]$. We need to show show that, for some G-invariant β not vanishing at x, the zero set of J within $X(\beta)$ equals $G \cdot [x] \cap X(\beta)$ scheme-theoretically.

Étale slice theorem (page 198 in [27], [17]): Let x be a point of an affine, smooth, irreducible G-variety X, whose orbit $Gx \subseteq X$ is closed. Then there exists a smooth, affine H-variety $Y \subseteq X$ passing through x and a strongly étale map ψ from $G \times_H Y$ to a G-invariant neighbourhood of $G \cdot x$ in X of the form $X(\alpha)$, for some G-invariant $\alpha \in \mathbb{C}[X]$.

Here $Z = G \times_H Y$ denotes the induced G-equivariant fibre bundle, with base G/H and fibre isomorphic to Y [27]. Strong étale-ness of ψ means that the map ψ/G from the quotient Z/G to X/G is étale and that the induced natural map from Z to $X \times_{X/G} Z/G$, the G-variety obtained from X by base extension, is a G-isomorphism.

The slice theorem suggests that we prove our theorem in two steps. First, consider the case when X is a fibre bundle of the form $G \times_H Y$, where Y is a smooth affine variety and then make a transition to the general case. Note that $H = G_x$ is reductive since $Gx \subseteq X$ is closed and hence affine [24].

We shall need the following:

Proposition 7.6 *Let V be a finite-dimensional irreducible G-module, G connected and reductive, with basis co-ordinate functions $V_1, \ldots, V_s$, Let $g \in V$ be a point with closed, affine orbit $Gg \subseteq V$. Further, let $I(g)$ be the ideal of Gg. Let J be an ideal of $\mathbb{C}[V]$ such that*
(i) The variety of J is precisely the orbit $O = Gg$.
(ii) The ideal J is itself G-invariant.
(iii) There are elements $w_1, \ldots, w_k \in J$ such that the tangent space TO_g of the orbit O at g consists of precisely the tangent vectors in TV_g annihilated by the differential forms dw_i's.
Then $J = I(g)$; i.e., $Spec(\mathbb{C}[V]/J) = O$.

Suppose (i) is replaced by the weaker condition:

(i)' The variety of J contains the orbit $O = G \cdot g$.

Then there exists a G-invariant neighbourhood U_g of the orbit Gg such that the zero set of J restricted to U_g coincides with Gg scheme-theoretically.

Proof: The G-invariance of J and the connectedness of G implies that all associated primes of J must themselves be G-invariant. Since there are no proper G-invariant subsets of O, we conclude that there are no associated primes of J other than $I(g)$. Now (iii) may be used to apply the 'Jacobian Criterion' (Matsumura[21], Theorem 30.4) locally. The G-invariance of J shows that (iii) holds at every point $y \in O$. The global assertion then follows. Q.E.D.

Given an H-module M, we denote by $\mathbb{C}[M]$ the H-module $\sum_{i\geq 0} \mathrm{Sym}^i(M^*)$, i.e., the space of polynomial functions on M. Let N denote the tangent space to Y at x; it is an H-module. Now we prove the theorem for the variety $G \times_H N$.

Lemma 7.7 *The theorem holds when $X = G \times_H N$ and $x = (1_G, 0_N)$ is the base point on its null section G/H, with stabilizer $G_x = H$.*

Proof: In this case N can be identified with the normal space N_x at x to the orbit $G \cdot x = G/H$. Let $N = \sum_R R$ be an H-module decomposition of N into irreducibles. Then we can write $N_x = N_{[x]} + M_x$, where $\bar{N} = N_{[x]}$ is the sum of all nontrivial H-submodules R in this decomposition, and M_x is the sum of all trivial H-submodules. The subvariety $G \cdot [x] = G \cdot M_x$, and the codimension of $G \cdot [x]$ is just the dimension of $N_{[x]}$.

For any H-submodule L of N_x, consider the induced bundle $F(L^*) : G \times_H L^* \to G/H$. Let $\mathcal{O}_{F(L^*)}$ be the sheaf of germs of sections of this bundle. Let $H^0(G/H, \mathcal{O}_{F(L^*)})$ be the G-module of its global sections. These global sections are regular functions on $G \times_H L$ that are linear on each fibre. Clearly $H^0(G/H, \mathcal{O}_{F(N^*_{[x]})})$ is a G-submodule of $H^0(G/H, \mathcal{O}_{F(N^*_x)})$, whose elements are regular functions on X linear on each fibre. Since G is connected, we can apply Jacobi's criterion (Proposition 7.6) and the transitivity of G-action. Hence it suffices to show that the sections in the non-H-admissible G-submodules of $H^0(G/H, \mathcal{O}_{F(N^*_{[x]})})$, when restricted to the fibre $N^*_{[x]}$ at x, span $N^*_{[x]}$; clearly the number r of such submodules is less than $\dim(N_{[x]}) = \mathrm{codim}(G \cdot [x], X)$.

Let $N_{[x]} = \oplus_R R$ be an H-module direct sum decomposition of $N_{[x]}$, where each R is an irreducible, nontrivial H-submodule. Then $N^*_{[x]} = \oplus_R R^*$, as an H-module, so we get a natural G-module decomposition

$$H^0(G/H, \mathcal{O}_{F(N^*_{[x]})}) = \oplus_R H^0(G/H, \mathcal{O}_{F(R^*)}).$$

Hence, it suffices to show that for each R in this decomposition, there exists a non-H-admissible G-submodule of $H^0(G/H, \mathcal{O}_{F(R^*)})$, whose sections, when restricted to the fibre R^* at x, span R^*.

So, let R be any such nontrivial, irreducible H-submodule in this decomposition and $L = R^*$ its dual. By the Peter-Weyl theorem (eq.(1))

$$H^0(G/H, \mathcal{O}_{F(L)}) = \oplus_Q Q \otimes Hom(Q, L)^H, \qquad (3)$$

where Q ranges over all finite dimensional irreducible G-modules, and $Hom(Q, L)^H$ denotes the vector space of H-equivariant linear maps from Q to L. Thus the G-modules Q that appear in $H^0(G/H, \mathcal{O}_F)$ are precisely the ones that contain L. By our G-separability assumption, there exists a nonadmissible, irreducible G-module Q_L containing L. By eq.(3), $H^0(G/H, \mathcal{O}_{F(L)})$ contains a copy of Q_L. Fix one such copy; we denote it by Q_L again. The restriction of Q_L to the fibre L of $F(L)$ at x is precisely L. Hence the basis elements of Q_L when restricted to L span L. Q.E.D.

For every R that appears in the H-module decomposition of $N_{[x]}$, let $Q_L \subseteq H^0(G/H, \mathcal{O}_{F(N^*_{[x]})})$, $L = R^*$, be a fixed copy as in the proof above. Let Φ be the set of such finitely many Q_Ls, each a non-H-admissible, irreducible G-module of regular functions on $G \times_H N$. The number r of Q_L's in Φ is less than codim$(G \cdot [x], X)$. Since many R's in the H-module decomposition may be isomorphic, many Q_Ls in Φ may be isomorphic as G-modules. The proof above shows that:

Lemma 7.8 *The differentials of the basis elements of the non-H-admissible G-modules Q_L in Φ, when restricted to $N_{[x]}$, span the whole of $N_{[x]}$, and the zero set of the ideal generated by them coincides with $G \cdot [x] = G \cdot M_x$ scheme-theoretically. The number r of modules in Φ is less than codim$(G \cdot [x], G \times_H N)$.*

Now we turn to the general case. By the étale slice theorem, there exists an affine H-variety $Y \subseteq X$ passing through x and a strongly étale map ψ from $G \times_H Y$ to a G-invariant neighbourhood of $G \cdot x$ in X of the form $X(\alpha)$, for some G-invariant $\alpha \in \mathbb{C}[X]$ not vanishing at x. Since $\mathbb{C}[X(\alpha)] = \mathbb{C}[X]_\alpha$, the ideal generated by non-H-admissible, irreducible G-submodules of $\mathbb{C}[X]$ within $\mathbb{C}[X(\alpha)]$ coincides with the one generated by non-H-admissible, irreducible G-submodules of $\mathbb{C}[X(\alpha)]$. Hence, in the statement of the theorem, we can replace X by $X(\alpha)$. Strong étale-ness of ψ implies [27] that there is an analytic neighbourhood $Y_{an} \subseteq N_x$ of x in N_x–called analytic slice through x–such that $G \times_H Y_{an}$ is G-isomorphic to an analytic G-invariant neighbourhood U of the orbit of x. However, there may not be an algebraic slice with this property, and this forces us in the analytic category in what follows. Since $U \simeq G \times_H Y_{an} \subseteq G \times_H N$, each Q_L corresponds to, and can be identified with, a G-module $Q_L(U)$ of analytic functions on U. By Lemma 7.8, the zero set of the $Q_L(U)$'s in Φ within U coincides, as a complex space [8], with $G \cdot [x] \cap U$. Our goal is to show that each $Q_L(U)$ can be approximated very closely within U by an isomorphic G-submodule of $\mathbb{C}[X]$. For this, we shall need the following results from complex function theory.

Cartan-Oka Theorem: ([8]) Let A be a Stein space, and B its closed analytic subspace. Then every holomorphic function on B extends to a holomorphic function on A.

Let $\mathcal{O}_A$, $\mathcal{O}_B$ be the sheaves of germs of holomorphic functions on A and B respectively, and $\mathcal{I}_B$ the sheaf of ideals of B. Then by Oka's theorem $\mathcal{I}_B$ is coherent, and since A is Stein, its higher cohomology $H^i(A, \mathcal{I}_B)$, $i > 0$, vanishes (Cartan's theorem B). Hence, this result follows from the long exact cohomology sequence associated with the exact sequence of sheaves

$$0 \to \mathcal{I}_B \to \mathcal{O}_A \to \mathcal{O}_B \to 0,$$

where we consider $\mathcal{O}_B$ as a sheaf on A via extension by zero.

We shall denote the ring of holomorphic functions on an analytic variety W by $\mathcal{O}(W)$.

Lemma 7.9 *Let A be a Stein G-space, and B its closed analytic G-subspace, G a connected reductive group. Let M be a finite dimensional G-submodule of $\mathcal{O}(B)$. Then there exists a G-equivariant extension map $\phi : M \to \mathcal{O}(A)$.*

Here, we say that ϕ is an extension map if, for any $s \in M$, the restriction of $\phi(s)$ to B coincides with s.

Proof: Fix a basis $s_1, \ldots, s_l$ of M. By the Cartan-Oka theorem, each s_i can be extended to a holomorphic function $\hat{s}_i$ on A. Let $\rho : M \to \mathcal{O}(A)$ be a linear map defined by setting $\rho(\sum_i b_i s_i) = \sum_i b_i \hat{s}_i$. Though ρ need not be G-equivariant, it may be converted into a G-equivariant map by Weyl's unitary trick [5]. Specifically, regard $Hom(M, \mathcal{O}(A))$ as a G-module in the natural way. Fix a maximal compact subgroup $E \subseteq G$. Let de denote left-invariant Haar measure on E, and let

$$\phi = \int_E e(\rho) de.$$

Then ϕ is an E-equivariant extension. Since M is finite dimensional, it follows from the unitary trick that ϕ is G-equivariant as well. Q.E.D.

Lemma 7.10 *Let W be a linear representation of connected, reductive G, V a linear space with trivial G action, $D \subseteq V$ a ball around the origin. Let $A = W \times D$. Then:*

(1) Any holomorphic function a on A has a unique power series expansion of the form

$$a(w, v) = \sum_{i,j} a_i^j w^j v^i. \quad (4)$$

Here $i = (i_1, \cdots, i_r)$, $r = \dim(V)$, and $j = (j_1, \cdots, j_q)$, $q = \dim(W)$, are tuples of nonnegative integers; and $v^i = v_1^{i_1} \cdots v_r^{i_r}$, and $w^j = w_1^{j_1} \cdots w_q^{j_q}$, where $v_1, \ldots, v_r$ are the coordinates of V, and $w_1, \ldots, w_q$ are the coordinates of W.

(2) For any $k = (i, d)$, the map $\delta_k : \mathcal{O}(A) \to \mathbb{C}[W \times V] = \mathbb{C}[W] \otimes \mathbb{C}[V]$, which maps a to $\sum_{j:j_1+\cdots+j_q=d} a_i^j w^j v^i$ is G-equivariant.

Proof: The first statement follows because A is a proper Reinhardt domain in $W \times V$ (cf. [6], page 20).

Each $w_1^{j_1} \cdots w_q^{j_q}$ is a polynomial (regular) function on W and is contained in the finite dimensional G-submodule in $\mathbb{C}[W]$ of homogeneous forms of degree $d = j_1 + \cdots + j_q$. Since the G-action on D is trivial, it follows that each δ_k is G-equivariant. Q.E.D.

Let X be an affine, smooth G-variety, G a connected reductive group. Let ψ be the projection from X to its quotient X/G. Let x be a point in X with closed orbit $Gx \subseteq X$, and $\bar{x} = \psi(x)$ its projection. Embed the affine variety X/G in a linear space V, with $\bar{x}$ at its origin. Suppose $U_{\bar{x}}$ is a Stein neighbourhood of $\bar{x}$ in X/G, such that $U_{\bar{x}} = D \cap X/G$, where $D \subseteq V$ is a ball around $\bar{x}$. Let $U = \psi^{-1}(U_{\bar{x}})$.

Lemma 7.11 *Let Q be a finite dimensional G-submodule of $\mathcal{O}(U)$. Then there exist G-equivariant linear maps $\rho_k : Q \to \mathbb{C}[X]$ such that any $s \in Q$ admits a power series expansion $s = \sum_{k=0}^{\infty} \rho_k(s)$ that converges everywhere in U.*

Proof: We can embed X G-equivariantly as a closed affine G-subvariety of some linear representation W of G ([10], Lemma 1.1). Let $A = W \times D$, which is Stein. It has a G-action, the action on D being trivial. Let $B \subseteq A$ be the closed analytic G-subspace consisting points (x, u), with $x \in X$, $u \in U_{\bar{x}}$, and $\psi(x) = u$. It is isomorphic to U. So s corresponds to a holomorphic function on B, which we shall denote by s again. Thus we can regard $Q \subseteq \mathcal{O}(B)$.

By Lemma 7.9, there exists a G-equivariant extension map $\phi : Q \to \mathcal{O}(A)$. Let δ_k be the G-equivariant map of Lemma 7.10 applied to A. Finally, let $\alpha : \mathbb{C}[W \times V] \to \mathbb{C}[X]$ be the G-equivariant restriction map corresponding to the G-equivariant embedding $X \to W \times V$, which maps $x \in X$ to $(x, \psi(x))$. Let $\rho_k = \alpha \circ \delta_k \circ \phi$. Then $s \in Q$ has a G-equivariant power series expansion

$$s = \sum_k \rho_k(s),$$

that converges everywhere in U. Q.E.D.

Now we return to the proof of Theorem 7.5. Let ψ be the strongly étale map from $G \times_H Y_{an}$ to a G-invariant neighbourhood U of the orbit $G \cdot x$. Here

$Y_{an} \subseteq N_x$ is an analytic slice, and U is of the form $\psi^{-1}U_{\bar{x}}$, where $U_{\bar{x}}$ is an analytic neighbourhood of $\bar{x} = \psi(x)$. We can assume that $U_{\bar{x}}$ is Stein, of the form $D \cap X/G$ as in Lemma 7.11, for a small enough ball D around $\bar{x}$ in $V \supseteq X/G$. Let $Q_L \in \Phi$ be the finitely many, irreducible, non-H-admissible G-submodules of the ring of regular functions on $G \times_H N$ as in Lemma 7.8; their number r is less than $\mathrm{codim}(G \cdot [x], G \times_H N) = \mathrm{codim}(G \cdot [x], X)$. We shall denote the restriction of Q_L to $G \times_H Y_{an}$ by Q_L again. It corresponds to a G-module of analytic functions on U, which we shall denote by $Q_L(U)$; the analytic functions in $Q_L(U)$ though may not extend to the whole of X.

Now we come to the crux of the proof. The G-module $Q_L(U)$ is isomorphic to Q_L, and hence, finite dimensional. Hence we may apply Lemma 7.11. Let $\rho_k(L)$ denote the G-equivariant projection from $Q_L(U)$ to $\mathbb{C}[X]$ therein. Let $\tilde{\rho}_k(L) = \sum_{j \le k} \rho_j(L)$. When k is large enough $\tilde{\rho}_k(L)(Q_L(U))$ will be a good approximation to $Q_L(U)$. Let $Q_L^k \subseteq \mathbb{C}[X]$ be the G-module that is the image of this G-equivariant projection $\tilde{\rho}_k$. Since $Q_L(U) \simeq Q_L$ is irreducible, Q_L^k is either zero, or isomorphic to Q_L. When k is large enough, Q_L^k is isomorphic to Q_L–hence it is non-H-admissible, and vanishes on $G \cdot [x]$.

Since U is G-isomorphic to $G \times_H Y_{an} \subseteq G \times_H N$, it follows from Lemma 7.8 that the the differentials of the basis functions in all the $Q_L(U)$ in Φ, when restricted to $N_{[x]}$, span the whole of $N_{[x]}^*$. We approximate each $Q_L(U)$ by $Q_L^k \subseteq \mathbb{C}[X]$ for a large enough k. When k is large enough, the differentials of the basis functions in Q_L^k, when restricted to $N_{[x]}$, will also span the whole of $N_{[x]}^*$. But each Q_L^k is a non-H-admissible, irreducible G-submodule of $\mathbb{C}[X]$. Thus it follows that the differentials of the basis functions of the non-H-admissible, irreducible G-submodules Q_L^k of $\mathbb{C}[X]$, for k large enough, span $N_{[x]}^*$. Because of the transitivity of the G-action, the same holds for all points in the orbit of x. Since G is connected, and all Q_L's are G-modules, it now follows from the Jacobian criterion (proposition 7.6, or rather its proof) and the fact that $U \simeq G \times_H Y_{an}$, that the zero-set of the basis functions of these Q_L^k's within U coincides with $G \cdot [x] \cap U$ scheme theoretically (i.e., as a complex space [8]). Since Q_L^ks are G-submodules of $\mathbb{C}[X]$, there exists a Zariski-open G-invariant neighbourhood $U' \supseteq U$ such that the zero set of Q_L^ks within U' coincides with $G \cdot [x] \cap U'$ scheme theoretically. It remains to show that U' can be chosen to be of the form $X(\beta)$, for some G-invariant β. The projection $\psi(U')$ into X/G is a constructible [9] set that contains $U_{\bar{x}}$. Hence $\psi(U')$ contains a Zariski-open affine neighbourhood of the form $(X/G)_\alpha$ for some G-invariant α not vanishing at x. Its inverse $\psi^{-1}(X/G)_\alpha$ is of the form $X(\alpha)$ and has the required properties. Q.E.D.

Remark: Suppose every H-module that appears in $N_{[x]}^*$ is not G-separable, as assumed in Theorem 7.5. Then one can similarly prove a weaker assertion that for some G-invariant analytic neighbourhood U (as in the proof above) of Gx, $\mathrm{Spec}(\mathbb{C}[X]/J) \cap U$, as a complex space [7], is a subspace of $G \times_H \mathrm{Spec}(I)$, where $\mathrm{Spec}(I)$ is a subscheme of $N_{[x]}$ and $I \subseteq \mathbb{C}[N_{[x]}]$ is the ideal generated by the G-separable H-submodules of $\mathbb{C}[N_{[x]}]$.

8 Partial stability

Let V be a linear representation of G. Let $P = KU$ be a parabolic subgroup of G, and R a reductive subgroup of K.

Definition 8.1 *We say that $v \in P(V)$ is (R, P)-stable (*partially stable*) if (1) it is stable with respect to the restricted action of R on V, and (2) $U \subseteq G_v \subseteq P$.*

Here $U \subseteq G_v$ implies that $U \subseteq G_{\hat{v}}$. The defect $\delta(v)$ of v is defined to be the difference between the ranks of the root systems of R and K. In our applications, the defect will be small–in fact, just one–and R will always be a semisimple Levi subgroup of a parabolic subgroup of K–so that the root system of R will always be a subsystem of that of K.

A stable point of V is (G, G)-stable. A point $v \in P(V)$ is (R, P)-stable iff it is an (R, K)-stable point of $P(Y)$, where $Y = V^U$ is the K-module of U-invariants in V.

Example 1: The simplest example of a partially stable point with defect zero is the point $v = v_\lambda \in P(V)$ that corresponds to the highest weight vector of an irreducible G-representation $V = V_\lambda(G)$. The stabilizer $P = G_v$ is parabolic, and v is clearly (L, P)-stable, where L is a semisimple Levi subgroup of P.

Example 2: let $f = \phi(h)$ be as in Definition 3.1, with h stable. Then f is (R, P)-stable, with defect one, with respect to the action of G (as well as $\hat{G}$), where: P is a parabolic subgroup of G (resp. $\hat{G}$), whose elements transforms the variables in $\bar{X}$ to their linear combinations, thus preserving an appropriate flag $\mathbb{C}^{k+1} \subseteq \mathbb{C}^l$, and $R \simeq SL_k(\mathbb{C}) \times SL_{l-k-1}$ is naturally embedded in the semisimple Levi subgroup of P isomorphic to $SL_{k+1}(\mathbb{C}) \times SL_{l-k-1}(\mathbb{C})$.

Definition 8.2 *Given dominant weights α and β of R and K, we shall say that $\alpha \lhd_R^K \beta$, or $\beta \rhd_R^K \alpha$, if $V_\alpha(R)$ occurs in $V_\beta(K)$, dropping the superscript or subscript whenever possible.*

In the definition of (R, P)-stability the group R will usually be such that

$$\tilde{L} \subseteq R \subseteq \tilde{K} \subseteq K, \tag{5}$$

for some parabolic subgroup $\tilde{P} = \tilde{T}\tilde{L}\tilde{U} = \tilde{K}\tilde{U}$ of K, as in Example 2. Then, using Littelmann's restriction rule [14], one can determine how any irreducible representation $V_\beta(K)$ explicitly decomposes as a $\tilde{K}$-module (and hence as an R-module). This, in turn, gives an explicit relationship between α and β in Definition 8.2.

In Example 2 above, $K \simeq GL_{1+k}(\mathbb{C}) \times GL_{l-1-k}$ and $R \simeq SL_k(\mathbb{C}) \times SL_{l-1-k}(\mathbb{C})$. In this case, Littelmann's restriction rule reduces to a variant of the well known Pieri's branching rule [5], which gives an explicit decomposition of $V_\mu(GL_{1+k}(\mathbb{C}))$ as a $GL_k(\mathbb{C})$ module.

For a connected reductive group D, we shall denote by i_D the canonical involution of its dominant weights so that $V_\lambda(D)^* = V_{i_D\lambda}(D)$. Let $v \in P(V)$ be an (R,P)-stable point as above. Let W and Y be respectively the smallest K-submodule and R-submodule of V containing $\hat{v}$.

Definition 8.3 *We say that a dominant weight β of G lies over a weight μ of R at v and degree d if*

1. *$V_\mu(R)$ and $V_{\beta'}(K)$ occur in $R_Y[v]_d$ and $R_W[v]_d$ respectively, where $\beta' = i_K(i_G\beta)$, and*
2. *$\mu \lhd_R^K \beta'$.*

We say that a dominant weight β of G lies over a weight μ of R at v if this is so at some d.

This definition does not depend on the choice of a Levi subgroup $K \supseteq R$ of P, because $U \subseteq G_v$. When the defect is zero, and R satisfies eq.(5), 2. just says that the weight β', restricted to R, is equal to μ. The number of β lying over μ at a fixed d depends on the defect; it is small if the defect is small.

9 Borel-Weil for a partially stable point

In this section we shall prove Theorem 2.8 (b) for partially stable points. Its precise statement is as follows.

Theorem 9.1 *Suppose v is (R,P)-stable (cf. Definition 8.1). Then $V_\lambda(G)$ can occur in $R_V[v]$ only if λ lies over some $R_{\hat{v}}$-admissible dominant weight μ of R at v (cf. Definition 8.3). Conversely, for every $R_{\hat{v}}$-admissible dominant weight μ of R, $R_V[v]$ contains $V_\lambda(G)$ for some dominant weight λ of G lying over μ at v.*

This will follow from the following stronger result.

Suppose $v \in P(V)$ is partially stable, specifically (R,P)-stable, where $P = TLU = KU$, and $R \subseteq K$.

Let W and Y be respectively the smallest K-submodule, R-submodule of V containing $\hat{v}$. Let $\mathcal{O}(d)$ be the twisting sheaf on $P(V)$ and $\mathcal{O}_{\Delta_V[v]}(d)$, $\mathcal{O}_{\Delta_W[v]}(d)$ the corresponding invertible sheaves on $\Delta_V[v]$ and $\Delta_W[v]$ respectively. Let $\Gamma(\Delta_V[v], \mathcal{O}_{\Delta_V[v]}(d))$ and $\Gamma(\Delta_W[v], \mathcal{O}_{\Delta_W[v]}(d))$ be the G and K-modules of their of global sections. Clearly $R_V[v]_d \subseteq \Gamma(\Delta_V[v], \mathcal{O}_{\Delta_V[v]}(d))$ for all $d \geq 0$. We have an equality for all $d \geq 0$ iff $\Delta_V[v]$ is projectively normal. (cf. page 126, Hartshorne [9]). Similarly, $R_W[v]_d \subseteq \Gamma(\Delta_W[v], \mathcal{O}_{\Delta_W[v]}(d))$ for all $d \geq 0$, with equality if $\Delta_W[v]$ is projectively normal.

The following result shows that the G-module structure of $R_V[v]$ is ultimately related to the R-module structure of $R_Y[v]$. In turn, we already know which R-modules can occur in $R_Y[v]$ since $v \in P(Y)$ is stable with respect to the action of R (Theorem 2.8 (a)).

Theorem 9.2 *(*Borel-Weil for partially stable points*)*

Suppose $v \in P(V)$ is (R,P)-stable as above. Then

1. *The G-module structure of $R_V[v]$ is equivalent to the K-module structure of $R_W[v]$: Specifically, the multiplicity of a G-module $V_\lambda(G)$ in $R_V[v]_d^*$ is equal to the multiplicity of the K-module $V_\lambda(K)$ in $R_W[v]_d^*$, where λ is regarded as a dominant weight of K by restriction. Moreover, a K-module $V_\alpha(K)$ can occur in $R_W[v]_d^*$ only if α is also a dominant weight of G.*

2. *The multiplicity of $V_\lambda(G)$ in the module $\Gamma(\Delta_V[v], \mathcal{O}_{\Delta_V[v]}(d))^*$ of global sections of $\mathcal{O}_{\Delta_V[v]}(d)$ is less than or equal to the multiplicity of $V_\lambda(K)$ in $\Gamma(\Delta_W[v], \mathcal{O}_{\Delta_W[v]}(d))^*$. If $\Delta_W[v]$ is projectively normal then the two multiplicities are equal, for all λ and $d \geq 0$, and $\Delta_V[v]$ is also projectively normal.*

3. *A K-module $V_\beta(K)$ can occur in $R_W[v]_d$ only if for some dominant weight $\alpha \lhd_R^K \beta$ of R, $V_\alpha(R)$ occurs in $R_Y[v]_d$. Conversely, for every R-module $V_\alpha(R)$ occurring in $R_Y[v]_d$, there exists a dominant weight $\beta \rhd_R^K \alpha$ of K such that $V_\beta(K)$ occurs in $R_W[v]_d$.*

4. *Finally, an R-module $V_\mu(R)$ occurs in $R_Y[v]$–i.e. in some $R_Y[v]_d$–iff it is $R_{\hat{v}}$-admissible.*

Remark 1 In the third statement, it is desirable that we have an explicit criterion for deciding if $\alpha \lhd_R^K \beta$. When R satisfies eq.(5) in Section 8, such a criterion is given by Littlemann's rule as pointed out there.

Remark 2: When G is semisimple and simply connected, and v corresponds to the highest weight vector in $V = V_\lambda(G)$, $\Delta_V[v] = G/P$, and $\Delta_W[v]$ is just the point v. Hence $\Gamma(\Delta_W[v], \mathcal{O}_{\Delta_W[v]}(d))^* = V_{d\lambda}(K)$, for $d \geq 0$. The second statement now implies that $\Gamma(\Delta_V[v], \mathcal{O}_{\Delta_V[v]}(d))^* = \Gamma(G/P, \mathcal{O}_{G/P}(d))^* = V_{d\lambda}(G)$, for $d \geq 0$– which is the Borel-Weil theorem [11].

We will first prove two propositions. For that we need the following lemma from representation theory.

Lemma 9.3 *(cf. Theorem 5.104 [11]) Let $V_\lambda(G)$ be an irreducible representation of a connected reductive group G with highest weight λ. Let $P = KU \subseteq G$ be a parabolic subgroup. Then $V_\lambda(G)^U = V_\lambda(K)$; here $V_\lambda(G)^U$ is the subspace of U-invariants in $V_\lambda(G)$.*

Let $z \in P(V)$ be a point whose stabilizer $G_z \subseteq G$ contains U, so that the stabilizer $G_{\hat{z}} \subseteq G$ of $\hat{z} \in V$ also contains U. Let Z be the smallest K-submodule of V containing $\hat{z}$. Let i denote the embedding of Z in V. The following result shows that $R_Z[z]$ and $R_V[z]$ are closely related.

Proposition 9.4 *(a) The multiplicity of an irreducible module $V_\lambda(G)$ in $R_V[z]_d^*$ is equal to the multiplicity of $V_\lambda(K)$ in $R_Z[z]_d^*$. Moreover, $V_\alpha(K)$ can occur in $R_Z[z]_d^*$ only if α is also a dominant weight of G.*

(b) The multiplicity of $V_\lambda(G)$ in $\Gamma(\Delta_V[z], \mathcal{O}_{\Delta_V[z]}(d))^*$ is less than or equal to the multiplicity of $V_\lambda(K)$ in $\Gamma(\Delta_Z[z], \mathcal{O}_{\Delta_Z[z]}(d))^*$. If $\Delta_Z[z]$ is projectively normal then the two multiplicities are equal for all λ and $d \geq 0$, and $\Delta_V[z]$ is also projectively normal.

Proof: Since the stabilizer $G_{\hat{z}}$ contains U, and U is normalized by K, the stabilizer of every point in Z contains U; in other words, the action of U on Z is trivial. Thus Z can be considered a P-module. The embedding map $i : Z \to V$ is then P-equivariant. By restriction, we get a P-equivariant, closed embedding $i : \check{\Delta}_Z[z] \to \check{\Delta}_V[z]$, where $\check{\Delta}_V[z] \subseteq V$ and $\check{\Delta}_Z[z] \subseteq Z$ denote the affine cones of $\Delta_V[z]$ and $\Delta_Z[z]$. Hence, the corresponding surjection $i^* : R_V[z] \to R_Z[z]$ is P-equivariant. Since it is degree-preserving, by restriction, we get a P-equivariant surjection $i^* : R_V[z]_d \to R_Z[z]_d$ for every d. Since the action of U on Z is trivial, we get the dual injection $i : (R_Z[z]_d)^* \to (R_V[z]_d^*)^U$.

Let M be any irreducible G-submodule of $R_V[z]_d$. Not all functions in M can vanish at z–otherwise arguing as in the proof of Proposition 5.2, we can conclude that the functions in M vanish identically on the affine cone $\check{\Delta}_V[z]$, which is not possible. It follows that the restriction map i^* is nonzero on M. Thus $N = i^*(M)$ is a nonzero K-module, with trivial U-action. Dually, this means $i(N^*)$ is a nonzero K-submodule of $(M^*)^U$. If $M^* = V_\lambda(G)$, then $(M^*)^U = V_\lambda(K)$

(Lemma 9.3), and hence irreducible. So $(M^*)^U \simeq i(N^*)$. Thus the injection $i : (R_Z[z]_d)^* \to (R_V[z]_d^*)^U$ is an isomorphism. Hence the multiplicity of $V_\lambda(G)$ in $R_V[z]_d^*$ is equal to the multiplicity of $V_\lambda(K)$ in $R_Z[z]_d^*$, and, moreover, $V_\alpha(K)$ can occur in $R_Z[z]_d^*$ only if α is also a dominant weight of G. This proves (a)

The proof of (b) is similar. The embedding map $i : Z \to V$ induces a P-equivariant map $i^* : \Gamma(\Delta_V[z], \mathcal{O}_{\Delta_V[z]}(d))^* \to \Gamma(\Delta_Z[z], \mathcal{O}_{\Delta_Z[z]}(d))^*$, which need not be surjection in general. Let M be any irreducible G-submodule of $\Gamma(\Delta_V[z], \mathcal{O}_{\Delta_V[z]}(d))$. One shows similarly that $N = i^*(M)$ is a nonzero K-submodule of $\Gamma(\Delta_Z[z], \mathcal{O}_{\Delta_Z[z]}(d))$, with trivial U-action, and $(M^*)^U \simeq i(N^*)$. This proves the first statement of (b).

If $\Delta_Z[z]$ is projectively normal, i.e., its homogeneous coordinate ring is integrally closed, then $\Gamma(\Delta_Z[z], \mathcal{O}_{\Delta_Z[z]}(d)) = R_Z[z]_d$, for $d \geq 0$ (cf. page 126 Hartshorne [9]). Hence i^* is a surjection, for $d \geq 0$, since the restriction $i^* : R_V[z]_d \to R_Z[z]_d$ is surjective, and $R_V[z]_d \subseteq \Gamma(\Delta_V[z], \mathcal{O}_{\Delta_V[z]}(d))$. Now we prove equality of multiplicities as in (a). Since, the multiplicity of every $V_\lambda(G)$ in $\Gamma(\Delta_V[z], \mathcal{O}_{\Delta_V[z]}(d))$ or $R_V[z]_d$ is now the same, both being equal to the multiplicity of $V_\lambda(K)$ in $\Gamma(\Delta_Z[z], \mathcal{O}_{\Delta_Z[z]}(d))^*$, it now follows that $\Gamma(\Delta_V[z], \mathcal{O}_{\Delta_V[z]}(d)) = R_V[z]_d$, for all $d \geq 0$. Hence $R_V[z]$ is integrally closed and $\Delta_V[z]$ is projectively normal (cf. page 126, Hartshorne [9]). Q.E.D.

Now let W be any linear representation of a connected, reductive group K, and $R \subseteq K$ a reductive subgroup. Fix a point $y \in P(W)$. Let Y be the smallest R-submodule of W containing $\hat{y}$.

Proposition 9.5 *An irreducible K-module $V_\beta(K)$ can occur within $R_W[y]_d$ only if an R-module $V_\alpha(R)$, with $\alpha \lhd_R^K \beta$ occurs within $R_Y[y]_d$. Conversely, if an R-module $V_\alpha(R)$ occurs in $R_Y[y]_d$ then there exists a K-module $V_\beta(K)$, with $\beta_R^K \rhd \alpha$, in $R_W[y]_d$.*

Proof: The embedding $r : Y \to W$ is R-equivariant. Hence, we have an R-equivariant, closed embedding $r : \check{\Delta}_Y[y] \to \check{\Delta}_W[y]$ of the affine cone of $\Delta_Y[y]$, and the corresponding R-equivariant surjection $r^* : R_W[y] \to R_Y[y]$. Since this surjection is degree preserving, by restriction, we get an R-equivariant surjection $r^* : R_W[y]_d \to R_Y[y]_d$ for each d.

Let $V_\beta(K)$ be any irreducible K-module in $R_W[y]_d$. Arguing as in the Proof of Proposition 9.4, we can conclude that its image under r^* is nontrivial. The image can thus be identified with an R-submodule of $V_\beta(K)$. If an R-module $V_\alpha(R)$ occurs in this image, then by definition (cf. Section8), $\alpha \lhd \beta$. Conversely, for every R-module $V_\alpha(R)$ that appears in $R_Y[y]_d$, there is a K-module $V_\beta(K)$ in $R_W[y]_d$ whose image contains $V_\alpha(K)$, and hence we must have $\beta \rhd \alpha$. Q.E.D.

Proof of Theorem 9.2: The first and the second statements follow from Proposition 9.4, letting $z = v$, $Z = W$. The third statement follows from Proposition 9.5. The fourth statement follows the statement (a) of Theorem 2.8, since, by definition of partial stability, v is a stable point of Y with respect to the action of R. Q.E.D.

Proof of Theorem 9.1: Suppose $V_\lambda(G)$ occurs in $R_V[v]_d$; i.e., $V_{i_G\lambda}(G)$ occurs in $R_V[v]_d^*$. Then by the first statement of Theorem 9.2, $V_{i_G\lambda}(K)$ occurs in $R_W[v]_d^*$. That is, $V_{i_K(i_G\lambda)}(K)$ occurs in $R_W[v]$. It now follows from the third and fourth statements of Theorem 9.2 that λ lies over some $R_{\hat{v}}$-admissible weight μ of R.

Conversely, it follows from Theorem 9.2 similarly that, for every $R_{\hat{v}}$-admissible dominant weight μ of R, $R_V[v]$ contains $V_\lambda(G)$ for some dominant weight λ of G lying over μ at v. Q.E.D.

10 Application in complexity theory

We now specialize the Borel-Weil theorem for partially stable points (Section 9) to the orbit closure problem that arises in complexity theory (Section 3). We follow the notation of Section 3. Now $V = \mathrm{Sym}^m(Y)$ is a linear representation of $G = SL(Y) = SL_l(\mathbb{C})$, and $W = \mathrm{Sym}^n(X)$ is a representation of $SL(X) = SL_k(\mathbb{C})$. Let $\hat{G} = GL_l(\mathbb{C})$. Let i^l denote the involution on the weights of $GL_l(\mathbb{C})$ so that $V_\lambda(GL_l(\mathbb{C}))^* = V_{i^l\lambda}(GL_l(\mathbb{C}))$, for a weight λ. Recall that every weight λ of $GL_l(\mathbb{C})$ or its dual $i^l(\lambda)$ corresponds to a Young diagram of height at most l. Every weight of $GL_l(\mathbb{C})$ that occurs in $\mathbb{C}[V]_d^* = \mathrm{Sym}^d(V) = \mathrm{Sym}^d(\mathrm{Sym}^m(Y))$ corresponds to a Young diagram of size md–this will be used implicitly in what follows.

Theorem 10.1 *(a) Suppose $g \in P(V)$ is stable with respect to the action of G. Then a Weyl module $V_\lambda(G)$ occurs in $\Delta_V[g]$ iff it is $G_{\hat{g}}$-admissible.*

(b) Suppose $f \in P(V)$ is of the form $\phi(h)$, $h \in P(W)$. Then $V_\lambda(\hat{G})$ can occur in $R_V[f]_d$ only if (1) the weight $i^l(\lambda)$ corresponds to a Young diagram with md boxes and height at most $k+1$, and (2) $V_{\lambda'}(GL_{k+1}(\mathbb{C}))$, with $\lambda' = i^{k+1} \circ i^l(\lambda)$, contains some $SL_k(\mathbb{C})_{\hat{h}}$-admissible module $V_\mu(SL_k(\mathbb{C}))$, where we consider $SL_k(\mathbb{C})$ as a subgroup of $GL_{k+1}(\mathbb{C})$ in a natural way. This means μ and λ' are related by (a variant of) Pieri's branching rule.

Conversely, for every $SL_k(\mathbb{C})_{\hat{h}}$-admissible module $V_\mu(SL_k(\mathbb{C}))$, there exists a d and λ satisfying (1) and (2) above such that $V_\lambda(\hat{G})$ occurs in $R_V[g]_d$.

Proof: (a) follows from Theorem 2.8 (a).

(b) The point $f \in P(V)$ is partially stable with defect one with respect to the action of $\hat{G} = GL_l(\mathbb{C})$ on $P(V)$: specifically (R,P)-stable, with R and P as specificed in Section 3. Now we apply Theorem 9.2 for the action of $\hat{G}$ on $P(V)$.

We will only clarify why the height of $i^l(\lambda)$ is at most $k+1$. The reductive Levi subgroup of P under consideration is $K \simeq GL_{k+1} \times GL_{l-k-1}$, and the subgroup $1 \times GL_{l-k-1}$, where 1 denotes the identity in GL_{k+1}, is contained in the stabilizer $K_{\hat{f}}$. Suppose $V_\lambda(G)$ occurs in $R_V[f]_d$. The irreducible K-submodule of V containing f is just $\bar{W} = \mathrm{Sym}^m(\bar{X})$ defined in Section 3. Hence, by Theorem 9.2, $V_{i_K \circ i^l \lambda}(K)$ is a nonzero K-submodule of $R_{\bar{W}}[f]_d$, where i_K is the involution on the weights of K. By Proposition 6.1, $V_{i_K \circ i^l \lambda}(K)$, and hence, $V_{i^l \lambda}(K)$ must be $K_{\hat{f}}$-admissible, and hence, $1 \times GL_{l-k-1}$-admissible. For any $V_\alpha(GL_l(\mathbb{C}))$, where α is a Young diagram of height $\leq l$, the K-module $V_\alpha(K)$, with the same weight, is equal to $V_{\alpha_1}(GL_{k+1}) \otimes V_{\alpha_2}(GL_{l-k-1})$, where α_1 consists of the first $k+1$ rows of α and α_2 consists of the its remaining $l-k-1$ rows; here an empty row is treated as a row with zero length. Let $\alpha = i^l(\lambda)$. Then $V_{\alpha_2}(GL_{l-k-1})$ must be trivial since $V_\alpha(K)$ is $1 \times GL_{l-k-1}$-admissible: thus $\alpha_2 = 0$, and $\alpha_1 = \alpha$. It follows that the length of α is at most $k+1$. The number of boxes in $i^l(\lambda)$ must be md since every irreducible $\hat{G}$-representation occurring in $\mathbb{C}[V]_d^* = \mathrm{Sym}^d(\mathrm{Sym}^m(Y))$ has degree md.

The rest follows from Theorem 9.2 and Theorem 9.1; details are left to the reader. Q.E.D.

11 Representation theoretic data associated with a partially stable point

We extend the definition of the representation theoretic data (Definition 7.1) to the partially stable case, and illustrate its significance with an application to G/P.

Definition 11.1 *Suppose $v \in P(V)$ is (R,P) stable, $P = KU$, we say that a G-submodule $M \subseteq \mathbb{C}[V]_d$ is admissible, with respect to v and d, if $(M^*)^U$ is (1) $(K, Sym^d(W))$-admissible, where W is the smallest K-submodule of V containing $\hat{v}$, and (2) it is also $R_{\hat{v}}$-admissible. Let Σ_v be the set of all nonadmissible G-submodules of $\oplus_d \mathbb{C}[V]_d = \mathbb{C}[V]$.*

Let $\Sigma_v(d) \subseteq \mathbb{C}[V]_d$ be the union of nonadmissible G-submodules of $\mathbb{C}[V]_d$.

Basis elements of the G-submodules in Σ_v will be called *nonadmissible basis elements*. The following is a generalization of Proposition 7.2.

Proposition 11.2 *Suppose v is (R,P)-stable. Then the G-modules in the representation-theoretic data Σ_v associated with v are contained in $I_V[v]$.*

Proof: Let $P = KU$. Fix any irreducible G-submodule $S \subseteq R_V[v]_d$. The result will follow if we show that every such S is admissible with respect to v and d

(Definition 11.1). It follows from the first statement of Proposition 5.2 that S^* must contain a $G_{\hat{v}}$-invariant. Since $U \subseteq G_{\hat{v}}$, this implies that $(S^*)^U$ contains an $R_{\hat{v}}$-invariant.

Let W be the smallest K-submodule of V containing $\hat{v}$. It remains to show that $(S^*)^U$ is $(K, \mathrm{Sym}^d(W))$-admissible. Since v, and hence $\hat{v}$, is stabilized by U, and U is normalized by K, W is also a P-module with trivial U-action. Let $\Phi = G \cdot W \subseteq V$. Consider the induced vector bundle $G \times_P W$ ([27]) with base space G/P and fibre W. Then Φ is the image of the natural G-equivariant map $\phi : G \times_P W \to V$ that maps (g, x), $g \in G$, $x \in W$ to $gx \in V$. We also have the associated map $\tilde{\phi} : G \times_P P(W) \to P(V)$. Since $\tilde{\phi}$ is proper, its image $\tilde{\Phi}$ is closed. The G-variety Φ is just the affine cone of $\tilde{\Phi}$, and is closed. $\Delta_V[v]$ is a closed G-subvariety of $\tilde{\Phi}$, and its affine cone $\check{\Delta}_V[v]$ is a closed G-subvariety of Φ. Hence, $R_V[v]$ is a G-summand of the homogeneous coordinate ring $R[\tilde{\Phi}]$ of $\tilde{\Phi}$. So every irreducible G sub-module of $R_V[v]$ can be thought of as an irreducible G-submodule of $R[\tilde{\Phi}]$. An element of $R[\tilde{\Phi}]_d$ is a regular function on Φ of degree d. Its pull back via ϕ is a global section of the bundle $B = G \times_P (\mathrm{Sym}^d W)^*$. Hence, an irreducible G-submodule S of $R_V[v]_d$ corresponds to a nonzero irreducible G-submodule of $\Gamma(G/P, B)$. The second statement of Proposition 5.2 applied to B, in conjunction with Schur's lemma, implies that, given any such S, S^* must contain a P-submodule isomorphic to a P-submodule of $\mathrm{Sym}^d(W)$; i.e., $(S^*)^U$ must be $(K, \mathrm{Sym}^d(W))$-admissible. Q.E.D.

11.1 Example: G/P

Proposition 11.2 suggests we study to what extent the data Σ_v determines the ideal $I_V[v]$. In this section we shall show that for G/P the data Σ_v determines $I_v[v]$ completely. This observation was a starting point for Theorem 2.11 and Conjecture 2.10.

Let G be a simply connected, semisimple group G and $P \subseteq G$ its parabolic subgroup, with Levi decomposition $P = KU$. Consider any embedding of G/P in $P(V)$, where $V = V_\lambda(G)$ is an irreducible G-representation, and λ is a dominant weight lying in the interior of the face of the dominant Weyl chamber in correspondence [5] with P. Let $v \in P(V)$ correspond to its highest weight vector. Then G/P must actually be the orbit of v in $P(V)$ [5]; i.e., $\Delta_V[v] \simeq G/P$. Recall that v is (L, P)-stable, with defect zero, where L is the semisimple Levi subgroup of P (Example 1 in Section 8).

Basis elements of $\Sigma_v(2)$ are equivalent to the Grassman-Plücker syzygies in the case of Grassmanian and, more generally, the quadratic straightening relations of the standard monomial theory [13] in the ideal of G/P:

Proposition 11.3 *1.* $\mathbb{C}[V]_d = V_{d\lambda}(G)^* \oplus \Sigma_v(d)$.

2. $R_V[v]_d = V_{d\lambda}(G)^*$.

3. $I_V[v]$ *is generated by the basis elements of* $\Sigma_v(2)$*, the nonadmissibility data of degree two.*

Remark: The second statement is one part of the Borel-Weil theorem (cf. Section 9). Compare its proof here with the one based on Bruhat decomposition [11].

Proof: 1. Since $\mathbb{C}[V]_d^* \simeq \mathrm{Sym}^d(V_\lambda)$ contains a unique highest weight vector with weight $d\lambda$, its G-module decomposition is of the form

$$\mathbb{C}[V]_d^* = V_{d\lambda} + \sum_\mu V_\mu, \tag{6}$$

where each μ is some dominant weight smaller than $d\lambda$, in the usual ordering on the weights. Let $W = \mathbb{C}_\lambda$ be the one-dimensional representation (character) of P corresponding to the weight λ, so that $\mathrm{Sym}^d(W) = \mathbb{C}_{d\lambda}$. We want to show (cf. Definition 11.1) that each $V_\mu = V_\mu(G)$, $\mu \neq d\lambda$, is not admissible at v; i.e., V_μ^U is not $(K, \mathbb{C}_{d\lambda})$-admissible, or in other words, that V_μ, as a P-module, can not contain $\mathbb{C}_{d\lambda}$ as a P-submodule (with trivial U-action): otherwise let $w \in V_\mu$ be a basis vector of this one-dimensional module. Since w is invariant under the unipotent subgroup of P, it must be the highest weight vector of V_μ, and μ must belong to the interior of the face of the dominant Weyl chamber that corresponds to P [5]. Moreover, as a P-module, the line $\mathbb{C}w$ corresponding to w cannot be isomorphic to $\mathbb{C}_{d\lambda}$ unless $\mu = d\lambda$. Hence $V_\mu^* \subseteq \Sigma_v(d)$ (Definition 11.1). This proves 1.

2. By Proposition 11.2, $\Sigma_v(d) \subseteq I_V[v]$, for all d. Hence, this follows from 1. since $R_V[v]_d$ is clearly nonzero.

3. This is now a consequence of the second fundamental theorem for G/P in the standard monomial theory (cf. Theorem 7.5 in [13]), which states that the ideal $I_V[v]$ is generated by the functions in $\mathbb{C}[V]_2$ that vanish on $\Delta_V[v]$. By 1. these are contained in $\Sigma_v(2) \subseteq I_V[v]$. Q.E.D.

12 SFT for the orbit of a partially stable, excellent point

Now we shall prove Theorem 2.11 for partially stable points with defect zero, by reducing it to the stable case that we have already proved. Let V be a linear representation of G. Let $P \subseteq G$ be a parabolic subgroup with Levi decomposition $P = KU = TLU$. We shall assume that the group R in the definition of

(R,P)-stability satisfies the restriction in eq.(5), as it does in our applications (cf. Section 3).

A precise statement of Theorem 2.11 in the partially stable case is as follows.

Theorem 12.1 *Let $V = V_\lambda(G)$. Let $v \in P(V)$ be an (R,P)-stable point with defect zero. Let W be the smallest K-submodule of V containing $\hat{v}$. Assume that (1) $L \subseteq R \subseteq K$. (2) $R_{\hat{v}} \subseteq R$ is R-separable and characterizes v, considered as a point in $P(W)$. Then the orbit $Gv \subseteq P(V)$ is determined by the representation-theoretic data Σ_v (Definition 11.1) within some G-invariant neighbourhood of the orbit. Specifically, there exists a G-invariant neighbourhood $Z \subseteq P(V)$ such that Gv is a closed subvariety of Z and the zero set (scheme) in Z of the basis elements of the G-modules in Σ_v coincides with Gv.*

For example, suppose $W = \mathrm{Sym}^n(X)$ is embedded via ϕ in $V = \mathrm{Sym}^m(Y)$, as in Section 3. Suppose (1) f is a stable point in $P(W)$ with respect to the action of $R = SL(X) = SL_{n^2}(\mathbb{C})$, (2) $R_{\hat{f}}$ characterizes f and is R-separable. Then $\phi(f)$ is a partially stable point of the type above.

Let $\Phi = G \cdot W \subseteq V$ as in the proof of Proposition 11.2. As we observed there, it is the image of the natural G-equivariant map $\phi : G \times_P W \to V$ that maps (g,x), $g \in G$, $x \in W$ to $gx \in V$, and we also have the associated map $\tilde{\phi} : G \times_P P(W) \to P(V)$. Since $\tilde{\phi}$ is proper, its image $\tilde{\Phi}$ is closed. The G-variety Φ is just the affine cone of $\tilde{\Phi}$. Let $R[\tilde{\Phi}]$ be the homogeneous coordinate ring of $\tilde{\Phi}$.

Our goal is to show that the orbit Gv of v is determined scheme-theoretically by the representation theoretic data within some G-invariant neighbourhood of the orbit. Since Gv is contained in $\tilde{\Phi}$, our first goal is to understand the geometry of $\tilde{\Phi}$. Once this is done, we shall be able to reduce the present case to the stable case that has already been analyzed.

When v corresponds to the highest weight vector of $V = V_\lambda(G)$, $\tilde{\Phi} = \Delta_V[v] = G/P$. Hence we wish to generalize the results in Section 11.1.

The geometry of $\tilde{\Phi}$

We say that $V_\alpha(G)$ is (K,U,W,d)-admissible if $(V_\alpha(G)^*)^U$ contains an irreducible K-submodule that also occurs in $\mathrm{Sym}^d(W)$, and non-(K,U,W,d)-admissible otherwise.

Proposition 12.2 *Every G-submodule in $R[\tilde{\Phi}]_d$ is (K,U,W,d)-admissible. Hence, every non-(K,U,W,d)-admissible G-submodule of $\mathbb{C}[V]_d$ belongs to the homogeneous ideal of $\tilde{\Phi}$.*

The proof is an easy modification of that of Proposition 11.2.

The following is a generalization of Proposition 11.3. Recall that W is a P-module with trivial U-action (cf. proof of Proposition 11.2).

Proposition 12.3 *(1) As a G-module, $R[\tilde{\Phi}]_d$ is isomorphic to the space $\Gamma_d = \Gamma(G/P, Sym^d(W^*))$ of global sections of the vector bundle $G \times_P Sym^d(W^*)$.*

(2) $\mathbb{C}[V]_d = \Gamma_d \oplus (\oplus_\beta V_\beta(G))$, where, $V_\beta(G)^*$, for any β, cannot contain an irreducible K-submodule that also occurs in $Sym^d(W)$. In particular, each $V_\beta(G)$ is non-(K,U,W,d)-admissible, and hence belongs to the ideal of $\tilde{\Phi}$.

(3) The ideal of $\tilde{\Phi}$ is generated (actually spanned) by non-(K,U,W,d)-admissible G-submodules of $\mathbb{C}[V]_d$.

For the proof of this proposition we shall need a lemma. Let $P = KU = TLU$ be the Levi decomposition as above. We think of the root system of K as a subsystem of that of G. Let l be any linear functional l on the weight space of G with respect to which the usual ordering of the roots of G is defined; here it is assumed that l is irrational with respect to the weight lattice. Let

$$V_\lambda(G) = V_\lambda(K) \oplus \bigoplus_\mu V_\mu(K) \tag{7}$$

be a decomposition of $V_\lambda(G)$ as a K-module. Let v_β be the highest weight vector of $V_\beta(K)$ occuring in this decomposition with respect to l. Let $w_T(\beta) = w_T(v_\beta)$ denote its T-weight, i.e., the weight with respect to the central torus $T \subseteq K$.

The following is a complement to Lemma 9.3. Let ϕ be the projection of the dominant weights of G onto the largest face F of the dominant Weyl chamber that is orthogonal (in the Killing norm) to the simple roots of $\mathcal{K}$, the Lie algebra of K. Note that (1) $w_T(\alpha) = w_T(\phi(\alpha))$, for any dominant weight, since $w_T(\gamma) = 0$ for any simple root γ of K, and (2) $w_T(\phi(\alpha)) \neq w_T(\phi(\beta))$, if $\phi(\alpha) \neq \phi(\beta)$. Order the projected weights in F according to the restriction of l to F. This induces an order on the T-weights $w_T(\alpha)$s.

Lemma 12.4 *For every μ in eq.(7), $w_T(\mu) = w_T(v_\mu)$ is less than $w_T(\lambda) = w_T(v_\lambda)$ for an appropriate l.*

Proof: Let W denote the Weyl group of G. For a simple root g, let W_g be the reflection in the hyperplane perpendicular to g.

The weights of $V_\lambda(G)$ are contained in the convex hull C of the conjugates of λ under the Weyl group elements [5]. Let A be the affine space, perpendicular to F, spanned by λ and $W_g(\lambda)$'s, where g ranges over the simple roots of $\mathcal{K}$. Its intersection with C is a face of C–call it L; it is the smallest face of C containing λ and $W_g(\lambda)$, for each simple root g of $\mathcal{K}$.

Claim 12.5 *The weight vectors of $V_\lambda(G)$, whose weights are contained in L, span the irreducible K-submodule $V_\lambda(K) \subseteq V_\lambda(G)$ with weight λ.*

Proof of the claim: Let $\mathcal{G}, \mathcal{K}$ denote the Lie algebras of G and K, and $U(\mathcal{G}), U(\mathcal{K})$ the corresponding universal enveloping algebras. We know that $V_\lambda(G)$ is spanned by αv_λ, where v_λ is the highest weight vector of $V_\lambda(G)$ and $\alpha \in U(\mathcal{G})$ ranges over all monomials in the negative roots of $\mathcal{G}$. If we order the roots appropriately, the Poincare-Birkoff-Witt theorem implies that α is of the form $\alpha_1\alpha_2$ where $\alpha_2 \in U(\mathcal{K})$ is a monomial in the negative roots of $\mathcal{K}$ and α_1 is a monomial in the remaining negative roots of $\mathcal{G}$. Then αv_λ is nonzero with weight in L iff $\alpha_1 = 1$. But $\alpha_2 v_\lambda$, as α_2 ranges over all monomials in the negative roots of $\mathcal{K}$ clearly span $V_\lambda(K) \subseteq V_\lambda(G)$. This proves the claim.

It follows from the claim that no $\mu \neq \lambda$ in eq.(7) can belong to L. We shall choose an irrational l such that the weights of $V_\lambda(G)$ within L have higher l-coordinates than the remaining weights of $V_\lambda(G)$; it clearly exists.

Consider the restriction of the linear function l to F. Then $l(\phi(\alpha))$ is higher than $l(\phi(\beta))$ for any weight β of $V_\lambda(G)$ not contained in L. Since no $\mu \neq \lambda$ in eq.(7) can belong to L, the result follows. Q.E.D.

Proof of Proposition 12.3: Since w is P-stable, its stabilizer contains U. Since U is normalized by K, it follows that every point in W is also stabilized by U. By Lemma 9.3, $W = W_\lambda = V_\lambda(G)^U = V_\lambda(K)$.

The decomposition of $V = V_\lambda$ as a K-module is of the form:

$$V = V_\lambda = W_\lambda \oplus \bigoplus_\mu W_\mu,$$

where, for each μ, $w_T(v_\mu) < w_T(v_\lambda)$ (Lemma 12.4). Let $W' = \bigoplus_\mu W_\mu$. By induction, and using the formula

$$\mathbb{C}[V]_d^* = \mathrm{Sym}^d(V) = \mathrm{Sym}^d(W_\lambda \oplus W') = \sum_{i+j=d} \mathrm{Sym}^i(W_\lambda) \otimes \mathrm{Sym}^j(W'),$$

it follows that $\mathbb{C}[V]_d^*$ has a K-module decomposition of the form

$$\mathbb{C}[V]_d^* = \mathrm{Sym}^d(V) = \mathrm{Sym}^d(W) \oplus W_d, \tag{8}$$

where the T-weight of the highest-weight-vector of each K-submodule of W_d is strictly smaller than the T-weight of the highest-weight-vector of each K-submodule in $\mathrm{Sym}^d(W)$. Hence no irreducible K-module can occur in both $\mathrm{Sym}^d(W)$ and W_d, considered as abstract K-modules, i.e. $\mathrm{Hom}(\mathrm{Sym}^d(W), W_d)^K = 0$.

Now consider a G-module decomposition

$$\mathbb{C}[V]_d^* \simeq \mathrm{Sym}^d(V_\lambda(G)) = \sum_\mu c_\mu^\lambda V_\mu(G), \tag{9}$$

where all $c_\mu^\lambda \geq 0$, and μ ranges over all dominant weights of G less than or equal to $d\lambda$. We do not know this decomposition explicitly; finding an explicit decomposition is a special case of the unsolved plethysm problem [5]. It follows from eq.(9) that

$$\mathbb{C}[V]_d^{*U} \simeq \mathrm{Sym}^d(V_\lambda(G))^U = \sum_\mu c_\mu^\lambda V_\mu(G)^U = \sum_\mu c_\mu^\lambda V_\mu(K), \tag{10}$$

where the last step follows from Lemma 9.3. Since $W = V^U$, $\mathrm{Sym}^d(W)$ is a U-submodule of $\mathrm{Sym}^d(V)$. Hence it follows from eq.(10) that each weight β of K that occurs in $\mathrm{Sym}^d(W)$ with nonzero multiplicity d^β, also occurs as a weight of G in $\mathrm{Sym}^d(V)$ with multiplicity at least d^β. On the other hand, by the Borel-Weil theorem and Lemma 9.3 (cf. also Frobenius reciprocity [4]),

$$(\Gamma_d^*)^U = \mathrm{Sym}^d(W). \tag{11}$$

It follows that as a G-module

$$\mathbb{C}[V]_d^* = \Gamma_d^* \oplus \bigoplus_\mu c_\mu V_\mu(G), \tag{12}$$

for suitable μ's. On the other hand, comparing this equation with eq.(8) it follows that no $V_\mu(G)$ here can contain an irreducible K-submodule that also occurs in $\mathrm{Sym}^d(W)$. This proves the second statement of the proposition.

It remains to show that Γ_d is a G-submodule of $R[\Phi]_d$. By eq.(11), $\Gamma_d^U = \mathrm{Sym}^d(W^*)$. Hence by the second statement, in conjunction with Lemma 9.3, this is equivalent to showing that $\mathrm{Sym}^d(W^*)$ is a K-submodule of $R[\Phi]_d^U$. This is clear, since we have the canonical U-equivariant embedding of W within Φ, the U-action on W being trivial. Q.E.D.

When $\tilde{\Phi} = G/P$, by the standard monomial theory, we know that nonadmissible basis elements of degree two generate the ideal of G/P (Section 11.1). Analogously, in the context of Proposition 12.3, one can ask for a degree bound c such that the basis elements of non-(K, U, W, d)-admissible G-submodules of $\mathbb{C}[V]_d$, $d \leq c$, generate the ideal of $\tilde{\Phi}$. This seeks an extension of the standard monomial theory to $\tilde{\Phi} = GW$.

Reduction to the stable case

Now we are ready to prove Theorem 12.1. Let $V = V_\lambda(G)$. Let v be an (R, P)-stable point with defect zero, as hypothesized, and $W = V_\lambda(K)$ be the smallest K-submodule of V containing $\hat{v}$. Since $L \subseteq R \subseteq K$, K and R are both products of the form $LT(K)$, $LT(R)$ respectively, where $T(K)$ and $T(R)$ are tori. Hence an irreducible K-module is also an irreducible R-module. In particular, W is an irreducible R-module with the action of the torus $T(R)$ being determined by a character–i.e., the action of $T(R)$ on $P(W)$ is trivial. Hence, any R-invariant subset of $P(W)$ is also K-invariant, and in particular, $Rv = Kv \subseteq P(W)$.

The orbit $Gv \subseteq P(V)$ is contained in $\tilde{\Phi}$. By Proposition 12.3, the ideal of $\tilde{\Phi}$ is generated (actually spanned) by the non-(K, U, W, d)-admissible G-submodules of $\mathbb{C}[V]_d$. These submodules are contained in the nonadmissibility data Σ_v associated with v (cf. Definition 11.1). Let $\hat{\Sigma}_v$ be the set of remaining G-submodules of $\mathbb{C}[V]$ in Σ_v. A G-submodule $M \subseteq \mathbb{C}[V]_d$ belongs to $\hat{\Sigma}_v$ iff $(M^*)^U$ is not $R_{\hat{v}}$-admissible. We shall show that there exists a G-invariant neighbourhood Z of Gv in $\tilde{\Phi}$ such that Gv is a closed subvariety of Z and Gv is determined within Z by the data $\hat{\Sigma}_v$; i.e., the zero set of the (basis elements of) the G-modules in $\hat{\Sigma}_v$, restricted to Z, coincides with Gv scheme-theoretically.

Consider the G-equivariant map $\tilde{\phi} : G \times_P P(W) \to \tilde{\Phi}$.

Claim 12.6 *$\tilde{\phi}^{-1}(v)$ is a point.*

Proof of the claim: Suppose to the contrary. Then there exists $g \notin P$ and a $w \in P(W)$ such that $\phi(g, w) = v$, i.e., $gw = v$, and hence, $w = g^{-1}(v)$. Since v is (R, P)-stable, $U \subseteq G_v \subseteq P$ (Definition 8.1). Since $w \in P(W)$, and the U-action on W is trivial,

$$U \subseteq G_w = (G_v)^{g^{-1}} \subseteq P^{g^{-1}}.$$

Thus both P and $P^{g^{-1}}$ contain U. This implies that $P = P^{g^{-1}}$ (by Lemma 5.2.5 (ii) in [30], and Corollary 11.17 (iii) in [2]). Thus g^{-1} normalizes P. Since the normalizer of P is P itself (Theorem 11.16 in [2]), it follows that $g \in P$; a contradiction.

Let us denote the point $\phi^{-1}(v)$ by $\tilde{v}$. Since $\tilde{\phi}$ is surjective, to show that Gv is scheme-theoretically determined within a G-invariant neighbourhood by the data $\hat{\Sigma}_v$, it suffices to show that $\tilde{\phi}^{-1}(Gv) = G \cdot \tilde{\phi}^{-1}(v) = G\tilde{v} \subseteq G \times_P P(W)$ is determined scheme-theoretically within some G-invariant neighbourhood by the set $\tilde{\phi}^{-1}(\hat{\Sigma}_v)$ of the pull backs of the G-modules in $\hat{\Sigma}_v$. But since $G_{\tilde{v}} = G_v \subseteq P$, the normal space to $G\tilde{v}$ can be identified with the normal space to its restriction to the slice $\tilde{\phi}^{-1}(P(W)) \simeq P(W)$, which in turn, corresponds to the normal space to the orbit $Rv = Kv \subseteq P(W)$. By the Jacobian criterion (Proposition 7.6), it now suffices to show that the set $\tilde{\phi}^{-1}(\hat{\Sigma}_v)_{P(W)}$ of the restrictions of the modules in $\tilde{\phi}^{-1}(\hat{\Sigma}_v)$ to the fixed slice $\tilde{\phi}^{-1}(P(W)) \simeq P(W)$ of the bundle $G \times_P (W)$ determines the orbit of $Rv = Kv \subseteq P(W)$ within some K-invariant neighbourhood of this orbit.

By Proposition 12.3, $R[\Phi]_d$ is isomorphic to the space $\Gamma_d = \Gamma(G/P, \mathrm{Sym}^d(W^*))$ of global sections of the bundle $G \times_P \mathrm{Sym}^d(W^*)$. By the Borel-Weil theorem and

Lemma 9.3 (see also the Frobenius reciprocity in [4]), the set of restrictions of the modules in Γ_d to the slice $P(W)$ can be identified with Γ_d^U: if $M \in R[\Phi]$ and is isomorphic to $V_\lambda(G)$, then the restriction of $\tilde{\phi}^{-1}(M)$ to $P(W)$ corresponds to M^U, which is isomorphic to $V_\lambda(K)$ (Lemma 9.3). Hence, the restrictions of the modules in $\tilde{\phi}^{-1}(\hat{\Sigma}_v)$ to the slice $P(W)$ consists of precisely the K-modules in $\mathbb{C}[W]$ that do not contain any $R_{\hat{v}}$-invariant. Since K and R are of the form $LT(K)$, $LT(R)$, an irreducible K-module is also an irreducible R-module, and the subspace of $\mathbb{C}[W]$ spanned by non-$R_{\hat{v}}$-admisible K-submodules coincides with the subspace spanned by non-$R_{\hat{v}}$-admisible R-submodules. Thus, $\tilde{\phi}^{-1}(\hat{\Sigma}_v)_{P(W)}$ consists of precisely the non-$R_{\hat{v}}$-admisible R-modules in $\mathbb{C}[W]$. Since $v \in P(W)$ is stable with respect to the action of R on $P(W)$, we can now apply Theorem 7.4 for the stable case. It implies that $Rv \subseteq P(W)$ has an R-invariant, and hence, K-invariant, neighbourhood Y such that Rv as a subvariety of Y is determined scheme-theoretically by $\tilde{\phi}^{-1}(\hat{\Sigma}_v)_{P(W)}$.

This shows that $\tilde{\phi}^{-1}(\hat{\Sigma}_v)_{P(W)}$ determines the orbit of $Rv = Kv \subseteq P(W)$ within a K-invariant neighbourhood of the orbit.

This proves Theorem 12.1. Q.E.D.

13 G-separability

We now study the notion of G-separability (Definition 7.3), which is of interest in the context of Theorem 2.11.

Proposition 13.1 *1. A semisimple group H, embedded in $G = H \times H$ diagonally, is strongly G-separable.*

2. $H = SL_k(\mathbb{C})$ is a strongly G-separable subgroup of $G = SL_n(\mathbb{C})$ if $k > (n+1)/2$.

3. $H = SL_k(\mathbb{C}) \times SL_l(\mathbb{C}) \subseteq G = SL_{k+l}(\mathbb{C})$, with natural embedding, is strongly G-separable.

Remark: The last statement can be generalized to semisimple Levi subgroups of maximal parabolic subgroups of classical simple groups, if one uses, instead of the decomposition formula in eq.(13) below, Littelmann's restriction rule [14].

Proof: (1) By Schur's lemma, a G-module $V_\alpha(H) \otimes V_\beta(H)$, where $\otimes$ denotes the external tensor product here, is H-admissible iff $V_\beta(H) \simeq V_\alpha(H)^*$, i.e. $\beta = i_H(\alpha)$, where i_H is the involution on dominant H-weights (Section 8). Any nontrivial representation $V_\lambda(H)$ occurs in the non-H-admissible G-module $V_\lambda(H) \otimes 1_H$, where 1_H denotes the trivial H-module. So H is clearly G-separable.

Strong G-separation follows from the following more general fact.

Claim 13.2 *$V_\lambda(H)$ occurs in the non-H-admissible G-module $V_\delta(H)\otimes V_\rho(H)$, $\delta=\lambda+\beta$ $\rho=i_H(\beta)$, for any dominant H-weight β.*

Proof of the claim: By Schur's lemma, this is equivalent to showing that

$$\begin{aligned}\mathrm{Hom}(V_\delta(H)\otimes V_\rho(H), V_\lambda(H)) &= V_{\lambda+\beta}(H)^*\otimes V_\rho(H)^*\otimes V_\lambda(H)\\ &= V_{\lambda+\beta}(H)^*\otimes V_\beta(H)\otimes V_\lambda(H)\end{aligned}$$

contains an H-invariant. By Schur's lemma again, this is equivalent to showing that $V_{\lambda+\beta}(H)$ occurs in $V_\beta(H)\otimes V_\lambda(H)$, which is clear.

(2) Consider a nontrivial $V_\lambda(SL_k(\mathbb{C}))$, where λ is a Young diagram of height h less than k. We shall exhibit a non-H-admissible $V_\mu(SL_n(\mathbb{C}))$ containing it. If h is greater than $n-k$, we let $\mu=\lambda$. Otherwise, let μ be a Young diagram obtained by adding $n-k-h+1$ boxes to the first column of λ. Its height is $n-k+1<k$. By Pieri's branching rule, it is easy to see that $V_\mu(SL_n(\mathbb{C}))$ contains $V_\lambda(SL_k(C))$ but not the trivial representation of $SL_k(\mathbb{C})$. More generally, if μ' is a Young diagram obtained by appropriately extending, i.e., adding boxes to the first $n-k$ rows of μ, then $V_{\mu'}(SL_n(\mathbb{C}))$ contains $V_\lambda(SL_k(C))$ but not the trivial representation of $SL_k(\mathbb{C})$. There are infinitely many such μ's. So $SL_k(\mathbb{C})$ is strongly separable.

(3) Assume that $k\geq l$, the other case being similar. Consider a nontrivial H-module $L=V_\alpha(SL_k(\mathbb{C}))\otimes V_\beta(SL_l(\mathbb{C}))$, where α and β correspond to Young diagrams of height less than k and l respectively. We shall exhibit a non-H-admissible G-module $V_\lambda(G)$ containing it. We identify α and β with the partitions: $\alpha=(\alpha_1,\alpha_2,\ldots)$, where α_i denotes the length of the ith row of the corresponding Young diagram, and $\beta=(\beta_1,\beta_2,\ldots)$. We proceed by cases.

Case 1: Either α does not correspond to a rectangular Young diagram of height l, or β is not trivial.

Let $\lambda=\alpha+\beta=(\alpha_1+\beta_1,\ldots)$. Note that the height of λ is less than k. We have [5],

$$V_\lambda(GL_{l+k}(\mathbb{C}))=\sum_{\rho,\delta} N^\lambda_{\rho,\delta}V_\rho(GL_k(\mathbb{C}))\otimes V_\delta(GL_l(\mathbb{C})), \tag{13}$$

where $N^\lambda_{\rho,\delta}$ denotes the Littlewood-Richardson coefficient. From this it easily follows that $V_\lambda(SL_{k+l}(\mathbb{C}))$ contains the representation $V_\alpha(SL_k(\mathbb{C}))\otimes V_\beta(SL_l(\mathbb{C}))$ of $SL_k(\mathbb{C})\times SL_l(\mathbb{C})$. But it cannot contain the trivial H-representation: If $\rho\neq 0$ and δ (possibly zero) correspond to rectangular Young diagrams with height k and l respectively–so that $V_\rho(SL_k(\mathbb{C}))$ and $V_\delta(SL_l(\mathbb{C}))$ are trivial–then $N^\lambda_{\rho,\delta}$ is easily seen to be zero; otherwise the height of λ will be at least k. On the other hand, if $\rho=0$, then $\lambda=\delta$. Since the height of β is less than l, the definition of λ then implies that $\alpha=\delta$ and $\beta=0$; a contradiction.

More generally, let α' be any Young diagram obtained from α by adding columns of length k. Let $\lambda' = \alpha' + \beta$. Then $V_{\lambda'}(SL_{k+l}(\mathbb{C}))$ also contains $V_\alpha(SL_k(\mathbb{C})) \otimes V_\beta(SL_l(\mathbb{C}))$ but not the trivial representation of $SL_k(\mathbb{C}) \times SL_l(\mathbb{C})$. Moreover, there are infinitely many such λ's.

Case 2: α is rectangular of height l and width w, and $\beta = 0$.

We can assume that $k > l$; otherwise $V_\alpha(SL_k(\mathbb{C}))$ too will be trivial. For any integer $r \geq 0$, let λ be the Young diagram whose first r columns are of height k, the $(r+1)$-st column is of length $l+1$, the columns numbered $r+2, \ldots, r+w$ are of height l, and the column numbered $r+w+1$ is of height $l-1$. Then it follows from eq.(13) that $V_\lambda(GL_{l+k}(\mathbb{C}))$ contains $V_\rho(GL_k(\mathbb{C})) \otimes V_\delta(GL_l(\mathbb{C}))$, where ρ is obtained from α by adding to its left r columns of length k, and δ consists of a single column of height l. Clearly $V_\rho(GL_k(\mathbb{C})) \otimes V_\delta(GL_l(\mathbb{C}))$ is isomorphic to $V_\alpha(SL_k(\mathbb{C})) \otimes V_\beta(SL_l(\mathbb{C}))$ as an $SL_k(\mathbb{C}) \times SL_l(\mathbb{C})$-module. But it does not contain the trivial $SL_k(\mathbb{C}) \times SL_l(\mathbb{C})$-module; this too follows from eq.(13). Moreover, there are infinitely many such λ.

This proves strong G-separability of H. Q.E.D.

For us, it is important to know if the stabilizers of the points that arise in the context of complexity theory are separable (cf. Section 3). The connected component of the stabilizer of $\det(Y)$ in $SL_{n^2}(\mathbb{C})$, where Y is an $n \times n$-matrix, contains $SL_n(\mathbb{C}) \times SL_n(\mathbb{C}) \subseteq SL(Y) = SL_{n^2}(\mathbb{C})$ (Section 3.1). Regarding this subgroup we make the following:

Conjecture 13.3 *$SL_n(\mathbb{C}) \times SL_n(C)$ is a strongly separable subgroup of $SL_{n^2}(\mathbb{C})$.*

Here the embedding corresponds to the natural embedding $SL(V) \otimes SL(V) \subseteq SL(V \otimes V)$, $V = C^n$. Specifically, letting $V_\lambda(n)$ denote $V_\lambda(SL_n(\mathbb{C}))$ in what follows, the conjecture can be reformulated as follows:

Conjecture 13.4 *For every nontrivial Weyl module $V_\lambda(n) \otimes V_\mu(n)$ of $SL_n(\mathbb{C}) \times SL_n(\mathbb{C})$, such that $|\lambda| = |\mu| \pmod n$, there exist (infinitely many) Weyl modules $V_\rho(n^2)$ of $SL_{n^2}(\mathbb{C})$ whose decomposition as an $SL_n(\mathbb{C}) \times SL_n(\mathbb{C})$-module contains $V_\lambda(n) \otimes V_\mu(n)$ but not the trivial $SL_n(\mathbb{C}) \times SL_n(\mathbb{C})$-module.*

The restriction $|\lambda| = |\mu| \pmod n$ is required to ensure (cf. Definition 7.3) that $V_\lambda(n) \otimes V_\mu(n)$ occurs in some representation of $SL_{n^2}(\mathbb{C})$; cf. eq.(14) below.

The conjecture can be reformulated in terms of the symmetric group as follows. Let $\hat{V}_\gamma(n^2)$ be a Weyl module of $GL_{n^2}(\mathbb{C})$. Embed $GL_n(\mathbb{C}) \times GL_n(\mathbb{C}) = GL(\mathbb{C}^n) \times GL(\mathbb{C}^n)$ in $GL(\mathbb{C}^n \otimes \mathbb{C}^n) = GL_{n^2}(\mathbb{C})$. The decomposition of $\hat{V}_\gamma(n^2)$ as a $GL_n(\mathbb{C}) \times GL_n(\mathbb{C})$-module is of the form:

$$\hat{V}_\gamma(n^2) = \sum_{\alpha,\beta} c_{\alpha,\beta,\gamma} \hat{V}_\alpha(n) \otimes \hat{V}_\beta(n); \tag{14}$$

here $c_{\alpha,\beta,\gamma}$ can be nonzero only if $|\alpha| = |\beta| = |\gamma|$. To get the decomposition of $\hat{V}_\gamma$ as an $SL_n(\mathbb{C}) \times SL_n(\mathbb{C})$ module, we reduce the Young diagrams occurring on the right hand side by removing columns of length n. This does not change their sizes modulo n; this explains the restriction $|\lambda| = |\mu| \pmod{n}$ in the conjecture. By Littlewood's symmetry conditions ([5]), the coefficients $c_{\alpha,\beta,\gamma}$ do not depend on the ordering of α, β and γ

Given a Young diagram δ, $|\delta| = m$, let W_δ denote the corresponding irreducible representation, the Specht module, of the symmetric group S_m. Then the coefficient $c_{\alpha,\beta,\gamma}$ occurring in the preceding decomposition is the same as the one occurring in the decomposition of the tensor product $W_\alpha \otimes W_\beta$ as an S_m-module,

$$W_\alpha \otimes W_\beta = \sum_\gamma c_{\alpha,\beta,\gamma} W_\gamma,$$

where $m = |\alpha| = |\beta|$; cf. [5].

For any λ of height less than n, and $m = |\lambda| \pmod{n}$, let $\lambda(m)$ be the unique Young diagram of size m obtained by adding to α columns of length n. Then the preceding conjecture is equivalent to saying that:

For every nontrivial pair of Young diagrams (λ, μ) of height less than n, and such that $|\lambda| = |\mu| \pmod{n}$, there exists an $m = |\lambda| = |\mu| \pmod{n}$, $m \geq n$, and a ρ of size m such that W_ρ occurs in the decomposition of $W_{\lambda(m)} \otimes W_{\mu(m)}$ as an S_m-module, but not in the decomposition of $W_\delta \otimes W_\delta$, where δ is the rectangular Young diagram of height n and size m.

If $|\lambda| = |\mu| \neq 0 \pmod{n}$, the last restriction is vacuous, because no such δ exists, and hence:

Proposition 13.5 *If $|\lambda| = |\mu| \neq 0 \pmod{n}$, Conjecture 13.4 holds.*

So, let us assume that $|\lambda| = |\mu| = 0 \pmod{n}$ in what follows.

Proposition 13.6 *Conjecture 13.4 holds for $n = 2$.*

The main difficulty in extending the proof below to $n > 2$ is that an explicit decomposition of the tensor product of two arbitrary Specht modules is not yet known.

Proof: We need to show that for every nontrivial pair of (λ, μ) of row-shaped Young diagrams, with $|\lambda|$ and $|\mu|$ even, there exists an even m and a ρ of size m such that W_ρ occurs in the decomposition of $W_{\lambda(m)} \otimes W_{\mu(m)}$ as an S_m-module, but not in the decomposition of $W_\delta \otimes W_\delta$, where δ is the rectangular Young diagram of height 2 and width $m/2$. We shall show that that there exist such a ρ for every large enough $m \geq 4(|\lambda| + |\mu|)$. Fix such an m.

Given a Young diagram γ, we shall let γ_i denote the number of boxes in its ith row from the top. We assume that the topmost row has the highest length in the diagram. We shall denote $\lambda(m)$ and $\mu(m)$ by $\bar\lambda$ and $\bar\mu$ respectively. Since λ and μ are row shaped, we shall let λ and μ denote the lengths of their row as well. Since $|\bar\lambda| = |\bar\mu| = m$, $\bar\lambda_2 - \bar\lambda_1 = \lambda$ and $\bar\mu_2 - \bar\mu_1 = \mu$, we have $\bar\lambda_2 = m/2 - \lambda/2$ and $\bar\mu_2 = m/2 - \mu/2$. Since $\bar\lambda$, $\bar\mu$ and δ have two rows, we can use the decomposition formula of Remmel and Whitehead [29].

First, we shall try to find a required ρ with two rows. Let (a, b), $a \geq b$, denote the two-rowed Young diagram with the top row of length a and the bottom row of length b. Suppose we are given Young diagrams $(k, h), (r, l), (d, c)$ of size m. Because of Littlewood's symmetry conditions we can assume that $l \leq h \leq c$. With this condition, The formula in [29] (Thm 3.3) says that

$$c_{(r,l),(k,h),(d,c)} = (1 + w - v)\chi(w \geq v), \tag{15}$$

where $w = \lfloor (l + h - c)/2 \rfloor$, $v = \max(0, \lceil (l + h + c - m)/2 \rceil)$, and the function χ is one if $w \geq v$ and zero otherwise.

By Littlewood's symmetry condition, $c_{\delta,\delta,\rho} = c_{\rho,\delta,\delta}$. Applying the preceding formula with $(r, l) = \rho$, and $(k, h) = (d, c) = \delta = (m/2, m/2)$, we conclude that this coefficient is nonzero iff $\lfloor \rho_2/2 \rfloor \geq \lceil \rho_2/2 \rceil$. That is, iff ρ_2 is even. So we need to find a ρ, with ρ_2 odd, such that $c_{\bar\lambda,\bar\mu,\rho}$ is nonzero. Because of symmetry, we can assume that $\bar\lambda_2 \leq \bar\mu_2$. We will try to find ρ such that

$$\rho_2 \leq \bar\lambda_2. \tag{16}$$

Then setting $(k, h) = \rho$, $(r, l) = \bar\lambda$, and $(d, c) = \bar\mu$ in eq.(15), we conclude that $c_{\bar\lambda,\bar\mu,\rho} = c_{\rho,\bar\lambda,\bar\mu}$ is nonzero iff

$$\lfloor (\rho_2 + \bar\lambda_2 - \bar\mu_2)/2 \rfloor \geq \max(0, \lceil (\rho_2 + \bar\lambda_2 + \bar\mu_2 - m)/2 \rceil), \tag{17}$$

i.e., iff

$$\lfloor (\rho_2 - \lambda/2 + \mu/2)/2 \rfloor \geq \max(0, \lceil (\rho_2 - \lambda/2 - \mu/2)/2 \rceil). \tag{18}$$

We now proceed by cases.

Case 1: $\mu \neq 0$,

In this case the condition in eq.(18) can be satisfied if

$$\rho_2 \geq (\lambda + \mu)/2, \tag{19}$$

and $\mu \geq 2$, which holds since μ is nonzero and even. But there are many odd ρ_2's such that eq(16) and eq(19) are satisfied if, say, $m \geq 4(\lambda + \mu)$.

Case 2: $\mu = 0$, and $\lambda/2$ is odd.

In this case, eq.(18) is satisfied if we let $\rho_2 = \lambda/2$, which is nonzero–otherwise (λ, μ) will be trivial–and odd, as required. Since m is large enough, eq.(16) is also satisfied.

It remains to consider

Case 3: $\mu = 0$ and $\lambda/2$ is even.

In this case, the required two-rowed ρ does not exist. So we shall find an appropriate $\rho = (\rho_1, \rho_2, \rho_3, \rho_4)$ with four rows such that $\rho_3 = \rho_4$.

Given Young diagrams (k, h), (m, l), (d, c, a, a) (entries in the nonincreasing order) with m boxes such that $a > 0$ and $\lceil (h+1)/2 \rceil \leq h - c$, the Remmel-Whitehead formula ([29], Theorem 3.1) says that

$$c_{(k,h),(m,l),(d,c,a,a)} = \sum_{r=h-c}^{\min(l, \lfloor \frac{l-a+h-c}{2} \rfloor)} 1 - \sum_{r=\max(a, l+h+a-m-1)}^{\min(l, \lfloor \frac{h-1}{2} \rfloor, \lfloor \frac{l+h+a+c-m-1}{2} \rfloor)} 1. \tag{20}$$

We will set $(k, h) = (m, l) = \delta = (m/2, m/2)$ and $(d, c, a, a) = \rho = (\rho_1, \rho_2, \rho_3, \rho_4)$ in this formula. For the formula to be applicable, we need to ensure that

$$\lceil \lceil (h+1)/2 \rceil = (m/2+1)/2 \rceil \leq h - c = m/2 - \rho_2. \tag{21}$$

If, in addition,

$$\rho_2 + \rho_3 < m/2, \tag{22}$$

we get that

$$\begin{aligned} c_{\delta,\delta,\rho} &= \textstyle\sum_{m/2-\rho_2}^{\lfloor (m-\rho_3-\rho_2)/2 \rfloor} 1 - \sum_{r=\rho_3}^{\lfloor \frac{\rho_3+\rho_2-1}{2} \rfloor} 1 \\ &= \lfloor (m - \rho_3 - \rho_2)/2) \rfloor - (m/2 - \rho_2) + \rho_3 - \lfloor \tfrac{\rho_3+\rho_2-1}{2} \rfloor, \\ &= \lfloor \tfrac{\rho_2-\rho_3}{2} \rfloor - \lfloor \tfrac{\rho_2-\rho_3-1}{2} \rfloor, \end{aligned}$$

which is 1 if $\rho_2 - \rho_3$ is even, and zero otherwise.

So we need to find a ρ with $\rho_2 - \rho_3$ odd, satisfying eq.(21) and (22), such that $c_{\bar{\mu},\bar{\lambda},\rho}$ is nonzero. Since $\mu = 0$, we have $\bar{\mu}_1 = \bar{\mu}_2 = m/2$. Also, recall that $\bar{\lambda}_2 = m/2 - \lambda/2$. Set $(k, h) = \bar{\mu} = (m/2, m/2)$, $(m, l) = \bar{\lambda}$ and $(d, c, a, a) = \rho$ in eq.(20). We shall choose a four rowed ρ, with nonzero ρ_3, such that $\rho_2 - \rho_3$ is odd, ρ_2 and ρ_3 are sufficiently larger than λ, and also such that the difference between $m/2$ and $\rho_2 + \rho_3$ is sufficiently larger than λ. This is possible if m is large enough compared to λ. In this case, the upper index of the first sum in eq.(20) becomes $\lfloor \frac{\bar{\lambda}_2 - \rho_3 + m/2 - \rho_2}{2} \rfloor$, and the lower index is $m/2 - \rho_2$. So the contribution of the first term is $\lfloor \frac{\rho_2 - \rho_3 - \lambda/2}{2} \rfloor$. Since λ is nonzero, the lower index of the second sum in eq.(20) is equal to $\frac{\lambda}{2} - 1 + \rho_3$. The upper index, assuming that m is large enough, and $m/2-\rho_2-\rho_3$ is sufficiently larger than λ, becomes $\lfloor\frac{-\lambda/2+\rho_2+\rho_3-1}{2}\rfloor$. Assuming that ρ_2 and ρ_3 are sufficiently larger than λ, it is larger than the lower index. Hence the second term becomes

$$\frac{\lambda}{2}-1+\rho_3-\lfloor\frac{-\lambda/2+\rho_2+\rho_3-1}{2}\rfloor=\lambda/2-1+\lfloor\frac{\lambda/2-\rho_2+\rho_3+1}{2}\rfloor.$$

Thus

$$c_{\bar{\mu},\bar{\lambda},\rho}=\lfloor\frac{\rho_2-\rho_3-\lambda/2}{2}\rfloor+\lambda/2-1+\lfloor\frac{\lambda/2-\rho_2+\rho_3+1}{2}\rfloor=\lambda/2-1.$$

This is nonzero, since $\lambda/2$, being nonzero and even, is at least two. So we can choose a ρ, with $\rho_2-\rho_3$ odd, and subject to the preceding conditions, as required. Q.E.D.

References

[1] D. Akhiezer, Homogeneous complex manifolds, Encyclopaedia of mathematical sciences, volume 10, Springer-Verlag. 1986.

[2] A. Borel, Linear algebraic groups, Springer-Verlag, 1991.

[3] A. Borel and Harish Chandra: Arithmetic subgroups of algebraic groups. Ann. of Math. 75, 485-535 (1962).

[4] R. Bott, Homogeneous vector bundles, Annals of Math. (2) 66, 1957, 203-248.

[5] W. Fulton, J. Harris: Representation theory, Springer-Verlag, 1991.

[6] H. Grauert, K. Fritzsche, Several complex variables, Springer-Verlag, 1976.

[7] H. Grauert, R. Remmert, Coherent analytic sheaves, Springer-Verlag, 1984.

[8] H. Grauert, R. Remmert, Theory of Stein spaces, Springer-Verlag, 1979.

[9] R. Hartshorne, Algebraic geometry, Springer-Verlag, 1977.

[10] G. Kempf: Instability in invariant theory, Annals of Mathematics, 108 (1978), 299-316.

[11] A. Knapp, Lie groups beyond an introduction, Birkhäuser, 1996.

[12] V. Lakshmibai, C. S. Seshadri, Geometry of G/P-V, J. Algebra 100 (1986), 462-557.

[13] V. Lakshmibai, P. Littlemann, P. Magyar, Standard monomial theory and applications, in Representation theories and algebraic geometry, (Ed. A. Broer), Kluwer Academic Publishers, (1997), 319-364.

[14] P. Littelmann, Paths and root operators in representation theory, Ann. of Math. 142 (1995), 499-525.

[15] D. Littlewood, Products and plethysms of characters with orthogonal, symplectic ad symmetric groups, Canad. J. Math. 10 (1958), 17-32.

[16] A. Lubotzky, R. Phillips, P. Sarnak: Ramanujan graphs, Combinatorica 8 (1988), 261-277.

[17] D. Luna: Slices Etales. Bull. Soc. Math. France, M emoire 33, 81 (1973).

[18] D. Luna, Th. Vust, Plongements d'espaces homogenes, Comment. Math. Helv. 58, 186 (1983).

[19] I. MacDonald, Symmetric functions and Hall polynomials, Clarendon Press, Oxford, 1995.

[20] G. Margulis: Explicit constructions of concentrators, Problemy Inf. Trans. 9 (1973), 325-332.

[21] H. Matsumara, Commutative ring theory, Cambridge University Press, 1980.

[22] H. Minc: Permanents. Addison-Wesley, 1978.

[23] M. Nagata, Polynomial Rings and Affine Spaces. CBMS Regional Conference no. 37, American Mathematical Society, 1978.

[24] Y. Matsushima: Espaces homogènes de Stein des groupes de Lie complexes. Nagoya Math. J. 16, 205-218 (1960).

[25] K. Mulmuley, M. Sohoni, Geometric complexity theory, P vs. NP and explicit obstructions, in "Advances in Algebra and Geometry", Edited by C. Musili, the proceedings of the International Conference on Algebra and Geometry, Hyderabad, 2001.

[26] K. Mulmuley, M. Sohoni, Geometric complexity theory I: an approach to the P vs. NP and related problems, SIAM J. Comput., vol 31, no 2, pp 496-526, 2001.

[27] D. Mumford, J. Fogarty, F. Kirwan: Geometric invariant theory. Springer-Verlag, 1994.

[28] A. Razborov, S. Rudich, Natural proofs, J. Comput. System Sci., 55 (1997), pp. 24-35.

[29] J. Remmel, T. Whitehead, On the Kronecker product of Schur functions of two row shapes, Bull. Belg. Math. Soc. 1 (1994), 649-683.

[30] T. Springer, Linear algebraic groups, in Encyclopaedia of mathematical sciences, vol. 55, Springer Verlag, 1991.

[31] L. Valiant: The complexity of computing the permanent. Theoretical Computer Science 8 , 189-201 (1979).