Geometric complexity theory, P vs. NP, and explicit obstructions

Dedicated to Sri Ramakrishna

Ketan Mulmuley and Milind Sohoni

Abstract. Theory of computing has given rise to some fundamental mathematical problems, notably the $P \neq NP$ conjecture, and the related lower bound problems concerning formula or circuit size. We develop an approach to these problems through geometric invariant theory. The goal of this approach is to reduce the *hard nonexistence* problems under consideration to *tractable existence* problems.

Accordingly, we reduce the arithmetic (characteristic 0) version of the $P \neq NP$ conjecture, and other related lower bound problems to proving existence of *obstructions*. These are representations in the homogeneous coordinate rings of orbit-closures in geometric invariant theory [**MFK**], of a class of points which are *partially stable* and whose stabilizers have special representation-theoretic properties. However, the Luna-Vust complexity [**LV**] of these orbit closures is quite high, in contrast with the well-understood homogeneous or almost-homogeneous-spaces, such as G/P [**LLM**], toric varieties [**F3**], and spherical embeddings [**BLV**], whose Luna-Vust complexity is zero.

We take a step towards *explicit* construction of obstructions by proving two results regarding these orbit closures. The first is a generalization of the Borel-Weil theorem for G/P to these orbit-closures. Second, we conjecture a nice representation-theoretic set of generators for their ideals, and prove a weaker version of the conjecture. Such a set of generators had earlier been given for the ideal of G/P by Lakshmibai, Seshadri, Littelmann [**LS, Li3, LLM**] and Kostant (cf. [**PK**]).

Finally, using these results, we reduce, in essence, the arithmetic non-existence problems under consideration to fundmental existence and construction problems in representation theory and algebraic geometry that are conjectured to be in the complexity class P.
Keywords: Computational Complexity, Geometric Invariant Theory, Representation Theory, Classical Groups, Plethysm, Homogeneous Spaces.
AMS Classification: 68Q15, 14L24, 14L35, 20G05, 14M17.

1. Introduction

This is an expository article on geometric complexity theory [**GCT1, GCT2**], which is an approach to the P vs. NP [**Co, Le, K**] and related

The work of the first author was supported by NSF grant CCR 9800042. The paper appears in "Advances in algebra and geometry", Edited by C. Musili, Proceedings of the International Conference on Algebra and Geometry, Hyderabad, 2001.

problems via algebraic geometry and representation theory. For proofs and details, see [**GCT1, GCT2**], henceforth referred to as Part I and Part II, respectively. We shall confine ourselves to the arithmetic (characteristic zero) versions of the lower bound problems in this article–so the underlying field is always $\mathbb{C}$ in what follows. In the characteristic zero versions, issues dealing with positive characteristic, finite fields, and bit representations of numbers are ignored. Analogue of the $P \neq NC$ conjecture in this spirit, but actually in the stronger PRAM model without bit operations, was proved in [**PRAM**]. Here, we develop an approach for the harder arithmetic versions of other fundamental conjectures in complexity theory [**Co, Le, K, V**]: namely, $P \neq NP$, $P \neq P^{\#P}$, and $NC \neq P^{\#P}$ conjectures.

The goal of the approach here is to reduce these *hard non-existence* problems in complexity theory to *tractable existence* problems in algebraic geometry and representation theory. The transition from non-existence to existence in the reduction is crucial. If one were to show that some problems in algebraic geometry and representation theory are (NP)-complete, this will only reduce one set of non-existence problems to another set of similar problems, and will not be very interesting. Rather, we reduce fundamental lower bound problems in complexity theory to proving existence of *obstructions*. These are algebro-geometric and representation–theoretic objects, which serve as "witnesses" or "proofs" of non-existence of small circuits or formulae. To prove their existence one has to more or less construct them *explicitly*; otherwise, the proof technique would have the same natural limitation [**RR**] that the earlier combinatorial attempts had [**BS**]. For earlier uses of explicit constructions in a different context of pseudorandom generation, see Margulis [**Mr**], and Lubotzky et al [**LPS**].

In the first step (cf. Section 2 and Part I), using the the work of Hilbert[**Hi**], Weyl[**W**], Mumford[**MFK**], Kempf[**Ke**], and others in geometric invariant theory, we reduce the preceding fundamental lower bound problems of complexity theory in characteristic zero to nice instances–in a sense to be made precise soon–of the following orbit-closure problem.

Let G be a connected, reductive group. Let V be a finite dimensional linear representation of G, and $P(V)$ the corresponding projective space. For a $v \in V$, let $\Delta_V[v]$ denote the projective closure of the G-orbit of v in $P(V)$.

PROBLEM 1.1. *Given explicit points $f, g \in P(V)$, is $f \in \Delta_V[g]$? Equivalently, is the projective G-variety $\Delta_V[f]$ contained in the projective G-variety $\Delta_V[g]$?*

The goal is to show that $f \notin \Delta_V[g]$, i.e., $\Delta_V[f] \not\subseteq \Delta_V[g]$. This leads us to the study of algebraic geometry of the orbit closures $\Delta_V[f]$ and $\Delta_V[g]$. These are almost-homogeneous-spaces in the terminology of [**A**]. Their algebraic geometry, for general f and g, is not well understood. Luna and Vust [**LV**] define the complexity of an almost-homogeneous space $\Delta_V[g]$ to be the difference between the dimension of $\Delta_V[g]$ and the maximum dimension of a B-orbit in it, where $B \subseteq G$ is a Borel subgroup. They call an almost-homogeneous space a spherical embedding if this complexity is zero, i.e., if some B-orbit is dense in $\Delta_V[g]$. For example, a homogeneous variety of the form G/P [**F1**], where P is a parabolic subgorup of G, or a toric variety [**F3**] is a spherical embedding. The algebraic geometry of spherical embeddings is well understood [**BLV**]. For G/P, in particular, Borel-Weil theorem [**F1**], standard monomial theory [**LLM**], and the work of Kostant [**PK**] provide a complete understanding of the G-module structure of its homogeneous coordinate ring (for the standard embedding), its ideal, and the generators of this ideal.

We need similar results for the almost-homogeneous-spaces that arise in the approach here. However, their Luna-Vust complexity is quite high: it is $\Omega(n^a)$, for some constant a, where $n \to \infty$ parametrizes the size of the input and $\dim(G) = \mathrm{poly}(n)$. This is the main source of difficulty here in comparison with the analogous problems for G/P and spherical embeddings.

What gives us a handle on the problem is that the almost-homogeneous spaces that arise in the approach here are special. Specifically, it is proved in Part I that (1) the point f is partially stable (cf. Sections 6 and 9), where partial stability is an extension of the notion of stability in geometric invariant theory due to Hilbert and Mumford [**MFK**], and (2) the stabilizers of both f and g have nice representation-theoretic properties (cf. Section 4). Here f will always be in the null cone–that will be another source of difficulties.

We shall take the following approach to the orbit closure problem (Problem 1.1). Let $R_V[f]$ and $R_V[g]$ denote the homogeneous coordinate rings of the projective varieties $\Delta_V[f]$ and $\Delta_V[g]$, respectively. Let $R_V[g]_d$ denote the degree-d component of $R_V[g]$, i.e., the subspace of forms of degree d; $R_V[f]_d$ is defined similarly. The rings $R_V[g]$, $R_V[f]$, as well as their degree d components $R_V[g]_d$ and $R_V[f]_d$ are representations of G. Suppose to the contrary that $\Delta_V[f] \subseteq \Delta_V[g]$. Then we have 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 any d.

DEFINITION 1.2. *We say that S is a representation-theoretic* obstruction *for the pair (f,g) if it occurs in $R_V[f]_d$ but not in $R_V[g]_d$, for some d.*

(Here, we say that an irreducible G-module S occurs in a G-module M if the complete decomposition of M into irreducible G-modules, as guaranteed by Weyl's decomposibility result ([**F1**]), contains S; we also say that M contains S.)

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" for the assertion that $f \notin \Delta_V[g]$. The representation-theoretic obstructions are expected to be in plenty for the (f,g) that occur in the lower bound problems under consideration (cf. Part II). However, for complexity-theoretic reasons [**RR**], to show their existence one has to more or less construct them explicitly

Explicit construction of obstructions requires a good understanding of the G-module structures of the homogeneous coordinate rings $R_V[g]$ and $R_V[f]$.

The first result in Part II gives a fairly precise characterization of the G-modules that can occur in these coordinate rings. Specifically, let v be any stable [**MFK, Ke**] or a partially stable point (cf. Definitions 6.1 and 9.1) in the projective space $P(V)$. The result (Theorem 10.1, and its stronger version in Part II) tells which G-modules can occur in the homogeneous coordinate ring of the projective closure of the orbit of v in $P(V)$. When v corresponds to the highest weight vector of V, this result specializes to the Borel-Weil theorem [**Kn**] for G/P.

The preceding result does not, however, tell which G-modules can occur in the specific degree-d component $R_V[v]_d$ of $R_V[v]$. This finer information is needed in the explicit construction of obstructions, and this, in turn leads to the problem of finding a set of generators with good representation-theoretic properties for the ideal $I_V[v]$ of $\Delta_V[v]$ (cf. Section 5).

When v is the highest weight vector in V, so that $\Delta_V[v] \cong G/P$, such a set of generators has been given earlier by Lakshmibai, Seshadri, and Littelmann [**LS, Li3, LLM**] and Kostant (cf. Lemma 3 in [**PK**]). For the nonspherical $\Delta_V[v]$'s that occur in the lower bound problems under consideration, we conjecture such a set of generators (cf. Conjecture 7.1). The second result in Part II (Theorems 7.3 and 10.3) proves a weaker version of this conjecture. It shows that the G-orbit of a stable [**MFK**] or a partially stable point $v \in P(V)$ with defect zero (Definition 9.1), with a characterizing G-separable stabilizer (Definitions 4.3 and 4.1), is scheme-theoretically defined in a G-invariant neighbourhood by a set of relations defined only on the basis of the representation-theoretic properties of the stabilizer.

Using these results, in conjunction with the stability result in Part I (Theorem 3.1), we reduce (cf. Section 8 and Part II) the problem of explicit construction of obstructions, in essence, to fundmental existence and associated explicit construction problems in representation theory and algebraic geometry (Problems 8.1 and 8.2) that are conjectured to be in the complexity class P (the class of polynomial-time decidable problems). This transition from hard non-existence problems to tractable existence problems, along with the associated explicit construction problems, is then, in a nutshell, the guiding principle of the approach here.

The fundamental problems and conjectures in representation theory, alluded to above, are instances of the following *subgroup restriction problem.* Let G be a connected reductive group and H its explicitly described connected, reductive subgroup. Given an irreducible representation $V = V_\lambda(G)$ of G with highest weight λ, find an explicit decomposition of V, as an H-module, into irreducible H-submodules. The associated existence problem is: given highest weights α and β of H and G respectively, does $V_\alpha(H)$ occur within $V_\beta(G)$? The goal is to find a polynomial time algorithm for this decision problem, i.e., show that it belongs to the complexity class P. (Here by polynomial, we mean polynomial in the dimension of G and the bitlengths of the weights.)

One special case of the subgroup restriction problem that is well understood is when $G = H \times H$, H semisimple, and H is embedded in G diagonally. Then the problem is equivalent to finding an explicit decomposition of the tensor product of two irreducible H-modules into irreducible H-submodules. For $H = SL_n(\mathbb{C})$, an explicit decomposition rule for the tensor product was given by Littlewood and Richardson [**Ma**]. For an arbitrary semisimple group an explicit decomposition rule was given much later by Kashiwara [**Ks**] and Littelmann [**Li2**]. A semi-explicit basis for the H-submodules of the tensor product was given by the work of Kashiwara

[**Ks**] and Lusztig [**Ls**] on canonical (global crystal) bases for the representations of the quantum groups. The Littlewood-Richardson, or more generally, Kashiwara-Littelmann explicit decomposition rule also implies (cf. Part II) a polynomial time algorithm for deciding, given highest weights α, β and λ of a connected reductive group G, if $V_\lambda(G)$ occurs within $V_\alpha(G) \otimes V_\beta(G)$.

Two instances of the subgroup restriction problem that arise in the approach here are the following.

The first is the plethysm problem in the representation theory of $SL_n(\mathbb{C})$. Here $H = SL_n(\mathbb{C})$, and $G = GL(V_\lambda) = GL_m(\mathbb{C})$, where V_λ is an m-dimensional, irreducible $SL_n(\mathbb{C})$-Weyl-module [**F1**] corresponding to a Young diagram λ, and the embedding of H into G is given by the homomorphism map of the representation. Thus, given an irreducible $GL_m(\mathbb{C})$-Weyl-module V_μ, the problem is decompose V_μ into irreducible $SL_n(\mathbb{C})$-submodules.

The second instance is equivalent to the tensor-product decomposition problem in the representation theory of the symmetric group S_n. The problem is to decompose the tensor product of two irreducible S_n-modules (Specht modules) [**F1**] into irreducible S_n-submodules.

Both these instances of the subgroup restriction problem are still open despite an intensive research on them [**Ma**]. In both cases, we conjecture that the problem of deciding if $V_\alpha(H)$ occurs within $V_\beta(G)$ belongs to the complexity class P. This is what is needed in the explicit construction of obstructions (Section 8).

The rest of this article is organized as follows.

In Section 2, we show how the fundamental lower bound problems in complexity theory can be reduced to the orbit closure problems in geometric invariant theory. In Section 3, we describe the central notion of stability in geometric invariant theory, due to Hilbert and Mumford [**MFK**], and state the main results in Part I regarding stability of various forms that occur in these instances of the orbit-closure problems (cf. Problem 1.1). The special properties of their stabilizers are described in Section 4. In Section 5, we describe our approach to the orbit closure problem via explicit construction of obstructions. In Sections 6 and 7, we state the main results in part II, which constitute a step towards explicit construction of obstructions. Using these results, we connect in Section 8 the problem of explicit construction of obstructions to existence and associated construction problems in representation theory and algebraic geometry, such as the subgroup-restriction problem described above, which are conjectured to be polynomial-time-decidable. Two main results (Theorems 6.2 and 7.3), required in the complexity theoretic-setting, actually hold in a more general setting of arbitrary connected reductive groups. These may have wider applications in geometric invariant theory, since the notion of partial stability (Section 9) gives an extension, to the points in the null cone, of Mumford's notion of stability [**MFK**]. In Section 10 we briefly state these results in a general setting. Finally, in Section 11 we summarize the main problems that loom ahead in the future course of this approach.

In Sections 2–8, G will usually denote the special linear group, which is all that we need for the complexity theory applications. In Sections 9 and 10 it denotes any connected, reductive group.

2. Orbit closure problems in complexity theory

We shall now describe the instances of the general orbit closure problem (Problem 1.1) that arise in the approach here. See [**St**] for an earlier work on invariant theory and orbit closure problems that arise in the context of bilinear complexity of matrix multiplication.

Let G be $SL_l(\mathbb{C})$, where $l \to \infty$ is a parameter depending on the input size. Let Y be a variable l-vector. Let $V = \mathrm{Sym}^m(Y) = \mathrm{Sym}^m((\mathbb{C}^l)^*)$ be the representation of $G = SL_l(\mathbb{C})$ consisting of homogeneous forms of degree m in the l variable entries of Y. For $k < l$, let $X = [y_1, \cdots, y_k]$ be the subvector of $Y = [y_0, \ldots, y_{l-1}]$ of size k. Let $W = \mathrm{Sym}^n(X)$, $n < m$ be the space of forms of degree n in the variable entries of X. It is a representation of $SL(X) = SL_k(\mathbb{C})$. We have an embedding $\phi : W \to V$, which maps any form $w \in W$ to $y^{m-n}w$, where $y = y_0$ is used as the homogenizing variable.

By abuse of notation, given a nonzero homogeneous form f in V, the corresponding point in $P(V)$ will also be denoted by f. If there is a possibility of confusion, we shall state if $f \in V$ or $f \in P(V)$.

The following is a specialization of Problem 1.1 that will be used in the complexity theory setting.

Problem 2.1. *Given fixed forms $g \in P(V)$, and $h \in P(W) \overset{\phi}{\hookrightarrow} P(V)$, does $f = \phi(h)$ belong to the (projective) Zariski-closure $\Delta_V[g]$?*

The specific h and g depend on the lower bound problem under consideration, and the goal is to show that $f = \phi(h)$ does not belong to $\Delta_V[g]$.

2.1. Arithmetic version of the $P^{\#P}$ vs. NC conjecture. The nonuniform version of the $P^{\#P}$ vs. NC conjecture says that the permanent of an $n \times n$ matrix cannot be computed by a boolean polynomial-size circuit of depth $\log^c n$, for any constant c. Its arithmetic implication is that the permanent of an $n \times n$ matrix cannot be computed by an integral polynomial-size circuit of depth $\log^c n$, for any constant c–otherwise such a circuit can be reduced modulo 2 to yield a small boolean circuit. In the approach here, we treat this arithmetic version first, since it does not involve problems of positive characteristic and finite fields. The arithmetic version, in turn, can be reduced [**V**] to showing that the permanent of $n \times n$ matrix X cannot be computed by evaluating the determinant $m \times m$ matrix, whose entries are nonhomogeneous linear functions of the entries of X, with $m = 2^{\log^c n}$ (the only lower bound on m known so far is linear [**Ga**]).

The specific instance of the orbit closure problem that arises here is the following. 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, $k = n^2$. Let $V = \mathrm{Sym}^m(Y)$ be the space of homogeneous forms of degree m in the variable entries of Y, and $W = \mathrm{Sym}^n(X)$, the space of homogeneous forms of degree n in the variable entries of X. We have a natural action of $G = SL(Y) = SL_l(\mathbb{C})$ on V and the projective space $P(V)$: namely, $\sigma \in G$ maps a form $q(Y) \in P(V)$ to $q(\sigma^{-1}Y)$, where we think of Y as a variable l-vector. Similarly, we have an action of $SL(X) = SL_k(\mathbb{C})$ on $P(W)$.

Using any entry y of Y not in X as a homogenizing variable we get an embedding $\phi : W \to V$, which maps any $w \in W$ to $y^{m-n}w \in V$.

Let $g = \det(Y) \in P(V)$ be the determinant form, and $f = \phi(h)$, where $h = \mathrm{perm}(X) \in P(W)$ is the permanent form. It is shown in Part I that if $\mathrm{perm}(X)$ can be computed by a circuit of depth less than $\log^c n$, then f lies in $\Delta_V[g]$ for $m = 2^{\log^c n}$. Conversely, if f lies in $\Delta_V[g]$ then $\mathrm{perm}[X]$ can be approximated infinitesimally closely[1] by a circuit of depth $O(\log^{2c} m)$. Since, the permanent is $\#P$-complete [**V**], this is not expected to happen if $m = 2^{O(\mathrm{polylog}(n))}$. This led us to conjecture that $f \notin \Delta_V[g]$, if $m = 2^{O(\mathrm{polylog}(n))}$, as $n \to \infty$. This is the specific orbit closure problem that arises in this setting–the reader may wish to keep this example in mind all the time.

2.2. Arithmetic 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 in the approach here. The orbit closure problem that arises in the context of the usual P vs. NP problem (cf. Part I) is a harder variant of the one that arises here.

The role of the permanent function in the previous example is now played by a certain integral form $E(X)$, which considered as a form over $\mathbb{F}_2$, belongs to co-NP but is not expected to be in P. The role of the determinant function is played by a certain integral form $H(Y)$, which considered as a form over $\mathbb{F}_2$, is P-complete.

The form $h = E(X)$ is the following. Let X be a variable $n \times 3n$-matrix. Divide its columns into n groups of size 3. Accordingly, let the columns of X be the $3n$ variable vectors $\{X_i^j | 1 \leq j \leq 3,\ 1 \leq i \leq n\}$. For any function $\sigma : \{1, \ldots, n\} \to \{1, 2, 3\}$, let $\det_\sigma(X)$ denote the determinant of the matrix whose i-th column is $X_i^{\sigma(i)}$. Let $E(X) = \prod_\sigma \det_\sigma(X)$, where σ ranges over all such functions. The function $E(X)$, reduced modulo $\mathbb{F}_2$, belongs to co-NP, but is not expected to be in P. We do not know if $E(X)$ is co-NP-complete, but it is closely related to a certain co-NP-complete function; cf. Part I. To show $P \neq NP$, it suffices to show that $E(X)$ cannot be computed by a boolean circuit of size asymptotically polynomial in n. The latter has an arithmetic implication that $E(X)$ cannot have an integral computational circuit of size polynomial in n–otherwise this circuit can be reduced modulo $\mathbb{F}_2$ to yield a polynomial-size boolean circuit for computing $E(X)$ over $\mathbb{F}_2$. In the approach to the P vs. NP problem here we address this arithmetic version of the $P \neq NP$ conjecture first.

The g in the orbit closure problem is now an integral form $H(Y)$ computed by a generic universal circuit with enough symmetries, so that its stabilizer is sufficiently large and has the properties (cf. Section 4) that we require in the approach here. The reader is referred to Part I for its precise description.

[1] This means, for every $\epsilon > 0$, there exists a form $\tilde{f} \in V$, which has a circuit of depth $O(\log^{2c} m)$, such that $||f - \tilde{f}|| < \epsilon$, in the usual norm on V.

3. Stability in complexity theory

In this section we shall describe the stability properties of the specific forms that arise in the orbit closure problems (cf. Problem 2.1 and Section 2).

We begin by recalling the central notion of stability in geometric invariant theory [**MFK, P**]. Let V be a representation of a reductive group G. Then G acts on the projective space $P(V)$ of lines in V through the origin. A polynomial function on V is said to be a G-invariant if it is constant on each G-orbit in V. A point $x \in V$ is said to be *unstable* or *nilpotent* if all homogeneous invariants of positive degree vanish at x, i.e., if x belongs to the zero set of all homogeneous invariants of positive degree. This turns out to be equivalent to saying that the closure of the orbit $G \cdot x$ in V contains zero. The set of nilpotent elements in V is called the *null cone* of the G-action. This important notion, crucial to us, was introduced by Hilbert [**Hi**].

Let y be a point in $P(V)$ and $x \in V$ any nonzero point on y. Following Mumford [**MFK**] and Kempf [**Ke**], we say that y is *G-stable*–or simply *stable* if G is clear from the context–iff the orbit $G.x \subseteq V$ is (Zariski-) closed. By abuse of notation, we shall sometimes say that x is stable. We say that y is *semistable* iff the closure of the orbit $G.x$ does not contain zero, i.e., if x is not nilpotent. These definitions clearly do not depend on the choice of x on y.

Let X be the variable vector, with entries $x_1, \cdots, x_k$. Let $W = \mathrm{Sym}^n(X)$ be the space of forms of degree n in the variable entries of X. It is a representation of $G = SL_k(\mathbb{C})$. It is known [**MFK**] that most forms in $P(W)$ are stable, but, their complete classification is not known. The task of proving stability of an explicitly given form is, in general, nontrivial. For this purpose, Mumford [**MFK**] extended Hilbert's criterion for semistability, and used it to prove stability of several forms which arise in his work. This work was later extended significantly by Kempf [**Ke**]. The resulting Hilbert-Mumford-Kempf criterion for stability is the main tool for proving stability of the forms that arise in the orbit closure problems in complexity theory (Section 2). Using it, we proved in Part I, the following basic result which connects complexity theory with geometric invariant theory.

THEOREM 3.1 (cf. part I). *The point $perm(X) \in P(W)$ (cf. Section 2.1) is stable with respect to the action of $SL(X) = SL_k(\mathbb{C})$ on $P(W)$.*

The point $g = \det(Y) \in P(V)$ (cf. Section 2.1) is stable with respect to the action of $G = SL(Y) = SL_l(\mathbb{C})$ on $P(V)$.

Similarly, the form $E(X) \in Sym^d(X)$, $d = n3^n$, which arises in the P vs. NP problem (cf. Section 2.2), is stable with respect to the action of $SL(X)$ on $P(W)$.

The form $H(Y)$, which plays the role of g in the P vs. NP problem, is not stable. But the stability of g is not required in the approach here to the construction of representation-theoretic obstructions for (f, g) in Problem 2.1–only that of h is required. What is crucial, however, is that $H(Y)$ is (weakly) characterized by its stabilizer (Section 4).

4. Stabilizers and their properties

In this section we discuss the stabilizers of h and g, and their properties that we need in the approach here to the orbit closure problem (Problem 2.1). The first property we need is that h and g should be more or less characterized by their stabilizers. This will allow us to use the representation theory of the stabilizers in a powerful way.

DEFINITION 4.1. *We say that $g \in V$ is characterized by its stabilizer $H \subseteq G$, if, up to a constant multiple, it is the only form in V stabilized (left invariant) by H.*

Let us illustrate this for the forms $h = \text{perm}(X)$ and $g = \det(Y)$ that occur in the arithmetic version of the $P^{\#P} \neq NC$ conjecture (Section 2.1).

The stabilizer of $\det(Y) \in V$ (Section 2.1) in $G = SL(Y) = SL_{m^2}(\mathbb{C})$ is known to be a reductive subgroup which consists of linear transformations in G of the form (thinking of Y as an $m \times m$ matrix):

$$Y \to AY^*B, \tag{4.1}$$

where Y^* is either Y or Y^T, $A, B \in GL_m(\mathbb{C})$. That determinant is characterized by its stabilizer follows from classical invariant theory.

The stabilizer of $\text{perm}(X) \in W$ (Section 2.1) in $SL(X) = SL_{n^2}(\mathbb{C})$ is a reductive group generated by linear transformations of the form (thinking of X as an $n \times n$ matrix):

$$X \to \lambda X^* \mu, \tag{4.2}$$

where X^* is either X or X^T, λ and μ are either diagonal or permutation matrices, and $n \geq 3$. It is easy to show that the permanent is also characterized by its stabilizer.

Stabilizers of the forms $E(X)$ and $H(Y)$ that occur in the P vs. NP problem, and the way they characterize these forms, in a weaker sense, is described in Part I.

4.1. G-separability.

Another property of stabilizers that will be important for us is their G-separability. Let $G = SL(X)$. G-separability of $H \subseteq G$ is a representation-theoretic notion that is intimately linked with the subgroup restriction problem for $H \subseteq G$.

DEFINITION 4.2. *Let G be a reductive group, $H \subseteq G$ its reductive subgroup. We say that a G-module M is H-admissible if it contains a trivial H-submodule, i.e., an H-invariant.*

DEFINITION 4.3. *Given $H \subseteq G$, H and G reductive, 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 shall say that a subgroup $H \subseteq G$ is G-separable, if every nontrivial irreducible H-module, which occurs in some G-module, is 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 (cf. part II). 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 $n = 2$ (cf. part II). We also conjecture that the stabilizers of the permanent, the determinant and other functions that arise in the lower bound applications are G-separable. Partial results, which provide evidence for this conjecture, are proven in Part II. Full results may need understanding the instances the subgroup restriction problem for $H \subseteq G$ that arise in the complexity-theory setting (cf. Section 7). 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.

5. The problem of explicit construction of obstructions

Now we address the problem of explicit construction of obstructions in the context of the orbit closure problem (Problem 2.1) that arises in complexity theory. Let G, V, W, f, g and h be as defined there. Let $H \subseteq G_g$ be any reductive subgroup of the stabilizer $G_g \subseteq G$ of $g \in V$. It is easy to show (cf. Part II) that that every irreducible G-submodule of $R_V[g]$ is H-admissible. This implies the following criterion for obstructions that is basic to the approach here.

Proposition 5.1. *[A criterion for obstructions] Any G-submodule of $R_V[f]$ that is not H-admissible is an obstruction.*

To use this obstruction criterion, one needs to solve two problems:

Problem 5.2. *Find an explicit criterion for determining if a given irreducible G-module is H-admissible.*

Problem 5.3. *Understand the G-module structure of $R_V[f]$, where $f = \phi(h)$ and $h \in W = Sym^n(X)$.*

We shall see (Theorem 6.2) that Problem 5.3 can, in turn, be reduced to understanding the $SL(X)$-module structure of $R_W[h]$. To understand the $SL(X)$-module structure of $R_W[h]$, it is desirable to have a sufficiently explicit set of generators for the ideal $I_W[h]$ of $\Delta_W[h]$ with good representation-theoretic properties. The idea is that such generators may then be used for straightening–as in the standard monomial theory for G/P [**LLM**], or the Hodge algebra theory for determinantal varieties [**CEP**]–so as to obtain a standard monomial basis for $R_W[h]$, and a detailed understanding of the $SL(X)$-module structure of each $R_W[h]_d$. Thus we are led to:

Problem 5.4. *Construct a sufficiently explicit, representation-theoretic set of generators for the ideal $I_W[h]$ of the projective variety $\Delta_W[h] \subseteq P(W)$.*

We shall now turn to the study of these problems.

6. G-module structure of $R_V[f]$

We begin with the study of the G-module structure of $R_V[f]$, $f = \phi(h)$, which is required in explicit construction of obstructions (Problem 5.3). In this study, the stability of the form h will play a crucial role.

The forms h that arise in the lower bound problems under consideration are indeed stable (Theorem 3.1) with respect to the action of $SL(X)$ on $P(W)$. In contrast, the form $f = \phi(h) \in P(V)$ (cf. Problem 2.1), for the embedding $\phi : P(W) \to P(V)$, is not stable with respect to the action of

$G = SL(Y)$ on $P(V)$–in fact, it belongs to the null cone of the G-action (cf. Part I).

However, when $h = \text{perm}(X)$ or $E(X)$, the form $f = \phi(h)$ has the following properties:

1. It is stable with respect to the action of the subgroup $SL(X) \subseteq SL(Y)$ on $P(V)$. This is because h is stable with respect to the action of $SL(X)$ on $P(W)$ (Theorem 3.1).
2. Let G_f be the stabilizer of f, P the parabolic group which maps the variables occurring in f, i.e., the ones in $X \cup \{y\}$, to their linear combinations, and $U \subseteq P$ its unipotent radical. Then $U \subseteq G_f \subseteq P$.

DEFINITION 6.1. *The form $f = \phi(h)$ is called partially stable (with defect one) if it has the preceding two properties.*

The main significance of partial stability lies in the following result, which reduces the study of the G-module structure of the homogeneous coordinate ring $R_V[f]$, $f = \phi(h)$, to the study of the $SL(X)$-module structure of the homogeneous coordinate ring $R_W[h]$.

Let $\hat{G} = GL_l(\mathbb{C})$. Note that $R_V[f]$ is also a $\hat{G}$-module. It suffices to study the $\hat{G}$-module structure of $R_V[f]$. We shall denote an irreducible $\hat{G}$-module corresponding to a weight λ by $V_\lambda(\hat{G})$. If λ corresponds to a Young diagram of height less than l, then $V_\lambda(\hat{G})$, by restriction, becomes an irreducible G-module, which we shall denote by $V_\lambda(G)$. Let i^l denote the involution on the weights of $GL_l(\mathbb{C})$ so that $V_\lambda(GL_l(\mathbb{C}))^* = V_{i^l\lambda}(GL_l(\mathbb{C}))$ [**F1**]. Let $\Delta_W[h] \subseteq P(W)$ denote the projective closure of the $SL(X)$-orbit of h, and $R_W[h]$ its homogeneous coordinate ring.

THEOREM 6.2 (cf. part II). *Suppose $V_\lambda(\hat{G})$ is an irreducible $\hat{G}$-module that occurs in the complete decomposition of $R_V[f]_d$ into irreducible $\hat{G}$-modules. Then the weight λ* lies over *(cf. Definition 6.3) some weight μ of $SL(X)$ such that $V_\mu(SL(X))$ occurs in the complete decomposition of $R_W[h]_d$ into irreducible $SL(X)$-modules.*

Conversely, for every μ such that $V_\mu(SL(X))$ occurs in the complete decomposition of $R_W[h]_d$ into irreducible $SL(X)$-modules, there is a weight λ lying over μ such that $V_\lambda(\hat{G})$ occurs in the complete decomposition of $R_V[f]_d$ into irreducible $\hat{G}$-modules.

Here lying over means that the Young diagrams λ and μ are combinatorially related as follows:

DEFINITION 6.3. *Let $V_\lambda(\hat{G})$ be an irreducible $\hat{G}$-module occurring in $R_V[f]_d$, and $V_\mu(SL(X))$ an irreducible $SL(X)$-module occurring in $R_W[h]_d$. We say that λ lies over μ iff*

1. *$i^l(\lambda)$ is a Young diagram with md boxes and height at most $k+1$, where m and k are as in Problem 2.1; and*
2. *$V_{\lambda'}(GL_{k+1}(\mathbb{C}))$, with $\lambda' = i^{k+1} \circ i^l(\lambda)$, contains $V_\mu(SL_k(\mathbb{C}))$, where we consider $SL_k(\mathbb{C}) = SL(X)$ as a subgroup of $GL_{k+1}(\mathbb{C})$ in a natural way. This means μ and λ' are related by (a variant of) the Pieri's branching rule* [**F1**].

Since Pieri's rule is quite simple, so is the combinatorial relationship between such λ and μ.

Theorem 6.2 reduces the study of the G-module structure of $R_V[f]$ to the study of the $SL(X)$-module structure of $R_W[h]$. In other words, the problem of existence of representation theoretic obstructions is reduced to the problem of existence of coarse representation theoretic obstructions in the following sense.

DEFINITION 6.4. *We say that $V_\mu(SL(X))$ occurring in $R_W[h]$ is a coarse representation-theoretic obstruction for the pair (f, g), if for every λ lying over μ, $V_\lambda(G)$ is not H-admissible, where H is any reductive subgroup of the stabilizer of $g \in V$.*

By Theorem 6.2 and Proposition 5.1, existence of a coarse representation theoretic obstruction implies existence of a representation-theoretic obstruction (Definition 1.2). The idea is to go on in this fashion, reducing the problem of existence of one kind of obstructions to the problem of existence of coarser obstructions, until at some stage an obstruction of that kind can be explicitly demonstrated. This will then serve as a "witness" for the assertion that $f \notin \Delta_V[g]$.

REMARK 6.1. Does there exist a more refined version of Theorem 6.2, with a more restrictive definition of "lying over"? If so, this will correspondingly give rise to a less coarse definition of an obstruction than in Definition 6.4.

For constructing coarse obstructions, we need to understand the $SL(X)$-module structure of $R_W[h]$. By Theorem 3.1, the forms h that arise in the lower bound problems under consideration are stable. The importance of stability becomes evident in the following result, which tells precisely which irreducible $SL(X)$-modules can occur in the complete decomposition of $R_W[h] = \oplus_d R_W[h]_d$ into irreducible $SL(X)$-modules, assuming that h is stable.

The following is a consequence of the algebraic version of the Peter-Weyl theorem [**Sp**] in conjunction with a simple evaluation trick (cf. Part II).

PROPOSITION 6.5. *If $h \in P(W)$ is stable, an irreducible $SL(X)$-module $V_\lambda(SL(X))$ can occur in $R_W[h]$ iff $V_\lambda(SL(X))$ is $SL(X)_h$-admissible (cf. Definition 4.2), where $SL(X)_h \subseteq SL(X)$ denotes the stabilizer of the form h, considered as a point in W.*

7. Generators for the ideal of $\Delta_W[h]$

Proposition 6.5 does not tell which irreducible $SL(X)$-modules can occur in a $R_W[h]_d$, for a fixed d. This finer information is needed in the construction of coarse obstructions, when one tries to deduce information regarding the G-submodules of $R_V[f]_d$ lying over the $SL(X)$-submodules of $R_W[h]$ using Theorem 6.2. This is because the definition of "lying over" in Theorem 6.2 is degree-sensitive. As we saw in Section 5, the study of the $SL(X)$-module structure of $R_W[h]_d$ leads to Problem 5.4, which is studied in this section.

Let h and W be as in Problem 5.4. In this section, we shall let $G = SL(X)$, and $H \subseteq G$ the stabilizer of $h \in W$.

With a stable form $h \in P(W)$, we associate a representation-theoretic data Σ_h: It is the set of all irreducible, non-H-admissible (Definition 4.2) G-submodules of $\mathbb{C}[W] = \oplus_i \mathrm{Sym}^i(W^*)$, where W^* denotes the dual of W. Observe that this data depends only on the embedding space $P(W)$ and the stabilizer H of h. Let $\tilde{\Sigma}_h$ be the set of any appropriately chosen basis elements of all G-submodules in Σ_h. The elements of $\tilde{\Sigma}_h$ will be called *nonadmissible basis elements.*

Proposition 6.5 implies that $\tilde{\Sigma}_h \subseteq I_W[h]$. We conjecture that:

CONJECTURE 7.1. *For $h = \det(X)$, or perm(X), the zero set $X(\tilde{\Sigma}_h)$ of $\tilde{\Sigma}_h$ coincides with $\Delta_W[h]$ as a variety, and even more strongly, as a scheme. The latter means the ideal $I_W[h]$ of $\Delta_W[h] \subseteq P(W)$ is generated by $\tilde{\Sigma}_h$*

A conjectured set of generators for $I_W[h]$, when $h = E(X)$, is described in Part II.

DEFINITION 7.2. *We shall say that a form $h \in P(W)$ and the corresponding orbit closure $\Delta_W[h] \subseteq P(W)$ are* excellent *if the form h has the following properties:*

1. *The form h is stable (cf. Section 3).*
2. *The form h is characterized, i.e., uniquely determined by its stabilizer (Definition 4.1).*
3. *The stabilizer of h is G-separable (Definition 4.3).*

By Theorem 3.1 and Section 4, the forms $\det(X)$ and $\mathrm{perm}(X)$ are excellent, assuming that their stabilizers are $SL(X)$-separable. One of the main results in part II is a weaker version of Conjecture 7.1, which, actually, holds for arbitrary excellent forms:

THEOREM 7.3 (cf. part II). *If the form $h \in P(W)$ is excellent, then for some G-invariant Zariski neighbourhood U of the orbit $G \cdot h$ of h, $G \cdot h = U \cap \Delta_W[h] = U \cap X(\tilde{\Sigma}_h)$ as a scheme. In other words, the ideal of $U \cap \Delta_W[h]$, as a closed subvariety of U, is generated by $\tilde{\Sigma}_h$.*

8. Towards explicit construction of obstructions

In Section 6, we reduced the problem of existence of an obstruction to existence of a coarse obstruction. For the problems under consideration, it is expected (cf. Part II) that there are plenty of such coarse obstructions of even small degree $d = \mathrm{poly}(m, l)$, where m and l are as in Problem 2.1. Explicit construction of a coarse obstruction raises the following problems.

8.1. Subgroup restriction problem. First, we need to make the definition of $\tilde{\Sigma}_h$ sufficiently explicit (cf. Problem 5.4 and Conjecture 7.1). For this, we need to solve the following instances of the subgroup restriction problem:

1. We need an explicit decomposition of the $SL(X)$-module $\mathbb{C}[W] = \oplus_i \mathrm{Sym}^i(W^*)$ into irreducible $SL(X)$-submodules. The problem of decomposing $\mathrm{Sym}^i(W^*) \cong \mathrm{Sym}^i(\mathrm{Sym}^n(X))$ into irreducible $SL(X)$-submodules is an important special case of the plethysm problem for the group $SL(X)$. We also seek a "semi-explicit" basis–in the spirit

of the canonical bases of Kashiwara[**Ks**] and Lusztig[**Ls**]–for these G-submodules.

2. We need an explicit criterion for deciding if a given irreducible G-module $V_\lambda(G)$ is H-admissible or not. This is an instance of the subgroup restriction problem for $H \subseteq G$.

The second problem also arises in the context of the obstruction-criterion (cf. Proposition 5.1 and Problem 5.2)–where $G = SL(Y)$ and H is the stabilizer of the form g in the orbit closure problem–and also in deciding if μ and λ in Definition 6.4 are, respectively, admissible (cf. Proposition 6.5) and non-admissible.

The specific instances of the subgroup restriction problem for $H \subseteq G$ that arise in the complexity-theory setting are the following.

First, let H be the stabilizer of $\det(Y)$ in $G = SL(Y)$ (Section 2.1). Then $SL_m(\mathbb{C}) \times SL_m(\mathbb{C}) \subseteq H \subseteq G = SL_{m^2}(\mathbb{C})$. The subgroup restriction problem for $SL_m(\mathbb{C}) \times SL_m(\mathbb{C}) \subseteq SL_{m^2}(\mathbb{C})$ is known [**F1**] to be equivalent to the problem of decomposing the tensor product of two irreducible (Specht) modules of the symmetric group (cf. Section 1).

Next, let H be the stabilizer of $\mathrm{perm}(X)$ in $G = SL(X)$ (Section 2.1). Then $S_n \times S_n \subseteq H$, and the subgroup restriction problem for $S_n \times S_n \subseteq G$ is linked to the plethysm problem for the special linear group [**Ma**]. The subgroup restriction problem that arises when H is the stabilizer of $E(X)$ (Section 2.2) is similar.

The decision problem associated with the subgroup restriction problem for $H \subseteq G$ is:

PROBLEM 8.1. *Given highest weights α and β of H and G, respectively, does $V_\alpha(H)$ occur in $V_\beta(G)$?*

We conjecture that the preceding instances of this problem that occur in the complexity-theory setting are polynomial-time-decidable, i.e., belong to the complexity class P (cf. Section 1).

8.2. Decision problem associated with $R_W[h]$. To construct coarse obstructions (Definition 6.4), one needs to understand the $SL(X)$-module structure of $R_W[h]$. Specifically:

PROBLEM 8.2. *Given an $SL(X)$-submodule P of $\mathbb{C}[W]_d$, does P occur in $R_W[h]_d$?*

Like Problem 8.1, we conjecture that this problem is also polynomial-time decidable for the specific h's that occur in the lower bound problems under consideration; by polynomial, we mean polynomial in n, k and d (as defined in Problem 2.1).

For the point $h = \mathrm{perm}(X)$ that arises in the context of the arithmetic implications of the $P^{\#P} \neq NC$ and $P \neq NP$, Conjecture 7.1 says that $I_W[h]$, where $W = \mathrm{Sym}^n(X)$, is generated by the nonadmissible basis elements. If this is so, one may be able to use the nonadmissible basis elements as straightening relations, as mentioned in Section 5, to obtain a a good understanding of the $SL(X)$-module structure of $R_W[h]_d$, thereby obtaining a polynomial-time decision procedure for Problem 8.2.

This, in conjunction with a good understanding of the subgroup restriction problem for the instances under consideration, may enable one to isolate an $R_W[h]_d$, with a small $d = \text{poly}(n,m)$, and an explicit $SL(X)_h$-admissible module $V_\mu(SL(X))$ occurring in it such that μ is a coarse obstruction; cf part II for details. Considerations for the arithmetic version of the $P \neq NP$ conjecture are similar (cf. Part II).

9. Partial stability

Theorems 6.2 and 7.3, stated earlier, actually hold in a more general setting of geometric invariant theory. To state these results in a general form, we begin by defining the central notion of partial stability, which was introduced earlier in a specialized context (Definition 6.1).

Let G be a connected reductive group. An irreducible G-representation with highest weight λ will be denoted by $V_\lambda(G)$. 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.

The notion of partial stability, unlike stability, will also apply to points in the null cone. Let V be a linear representation of G. For a point $v \in P(V)$, we shall denote by $\hat{v} \in V$ a nonzero point on the line representing v; the exact choice of $\hat{v}$ will not matter. The stabilizers of v and $\hat{v}$ in G will be denoted by G_v and $G_{\hat{v}}$, respectively.

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 [**Ke**] associates with v and S a canonical parabolic subgroup $P = P[S,v] \subseteq G$–canonical destabilizing flag–which corresponds to an essentially unique one-parameter subgroup of G that drives $\hat{v}$ most rapidly towards S. Let L be its semisimple Levi subgroup. We say that v is (L,P)-stable (partially stable with defect zero) 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.

More generally, we now define the notion of partial (R,P)-stability with a nonnegative defect. Let $P = KU$ be a parabolic subgroup of G. Let R be a reductive subgroup of K. In fact, we will assume that R is either a semisimple Levi subgroup of a parabolic subgroup of K, or more generally, such that

$$\tilde{L} \subseteq R \subseteq \tilde{K} \subseteq K, \tag{9.1}$$

for some parabolic subgroup $\tilde{P} = \tilde{T}\tilde{L}\tilde{U} = \tilde{K}\tilde{U}$ of K. Here the root system of $\tilde{K}$ will be a subsystem of that of K.

DEFINITION 9.1. *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 the complexity-theory applications, the defect will be small–in fact, just one.

If the defect is zero, there is a canonical choice for P in the preceding definition: namely, the parabolic group $P[S,v]$ defined above, where S is the unique [**MFK**] closed orbit contained in the affine orbit-closure of v in V.

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: One natural choice for R and P in the preceding definition is as follows. Suppose v in $P(V)$ is not stable. Let S be a G-invariant closed subset of V, not containing $\hat{v}$, which meets the boundary of the orbit $G\hat{v}$. Let $P[S,v]$ be the canonical parabolic subgroup, as in Kempf [**Ke**], associated with v and S. Let $L[S,v]$ be its semisimple Levi subgroup. Let $\bar{P}[S,v] \supseteq P[S,v]$ be the smallest parabolic subgroup of G whose unipotent radical is contained in the stabilizer G_v. It exists because if $P_1 \supseteq P[S,v]$ and $P_2 \supseteq P[S,v]$ are two parabolic subgroups with this property, then so is their intersection–this follows from the explicit characterization of parabolic subgroups in terms of roots (cf. Theorem 4.3.4, [**Sp**]). If v belongs to the null cone, it will be semistable with respect to the restricted action of $L[S,v]$, but, in general, v need not be stable with respect to the restricted action of $L[S,v]$. When v is stable with respect to this restricted action, it is $(L[S,v],\bar{P}[S,v])$-stable, and its defect is $\text{rank}(\bar{L}[S,v]) - \text{rank}(L[S,v])$.

Example 3: This is the example whose instances arise in the complexity-theoretic setting (cf. Definition 6.1). Let $Y = [y_0, \cdots, y_{l-1}]$ denote a variable l-vector. Let $V = \text{Sym}^m(Y) = \text{Sym}^m((\mathbb{C}^l)^*)$ be the space of homogeneous forms of degree m in l variables, with the natural action of $SL_l(\mathbb{C})$ and $GL_l(\mathbb{C})$. Let

$$f(Y) = g(y_0, \dots, y_{\delta-1})h(y_\delta, \dots, y_{\delta+k-1}),$$

where g and h are forms in the shown variables, and $h \in P(V)$ is stable with respect to the action of $SL_k(\mathbb{C})$ on $P(V)$ induced by its action on the variable k-vector $[y_\delta, \dots, y_{\delta+k-1}]$. Then $f \in P(V)$ is partially stable with defect δ with respect to the action of $SL_l(\mathbb{C})$ as well as $GL_l(\mathbb{C})$. Specifically, let $P \subseteq GL_l(\mathbb{C})$ (or $SL_l(\mathbb{C})$) be the parabolic group of transformations in $GL_l(\mathbb{C})$ (resp. $SL_l(\mathbb{C})$) that map the variables $y_0, \dots, y_{\delta+k-1}$ to their linear combinations. It preserves an appropriate flag $\mathbb{C}^{\delta+k} \subseteq \mathbb{C}^l$. The semisimple Levi subgroup, in case of the $GL_l(\mathbb{C})$-action, is isomorphic to $GL_{\delta+k}(\mathbb{C}) \times GL_{l-\delta-k}(\mathbb{C})$–in case of the $SL_l(\mathbb{C})$-action it is the subgroup of transformations in $GL_{\delta+k}(\mathbb{C}) \times GL_{l-\delta-k}(\mathbb{C})$ with determinant one. The semisimple Levi subgroup L is isomorphic to $SL_{\delta+k}(\mathbb{C}) \times SL_{l-\delta-k}(\mathbb{C})$. Let $R \subseteq L \subseteq K$ be the subgroup, isomorphic $SL_k(\mathbb{C}) \times SL_{l-\delta-k}(\mathbb{C})$, of transformations in L that map the variables $y_\delta, \dots, y_{\delta+k-1}$ (and $y_{\delta+k}, \dots, y_{l-1}$) to linear combinations among themselves. Then f is (R,P)-stable, with defect δ. Note that R is the semisimple Levi subgroup of a parabolic subgroup of K.

Because of the restriction (9.1) on the group R in the definition of (R,P)-stability, we can use Littelmann's restriction rule, given below, to determine how any irreducible representation $V_\beta(K)$ explicitly decomposes as a $\tilde{K}$-module (cf. (9.1), and hence as an R-module–this turns out to be crucial in applications (cf. part II), and we review it here briefly. Let B_ρ be the set of Lakshmibai-Seshadri paths [**Li2**] that index a standard basis of an irreducible representation $V_\rho(D)$ of a semisimple group D. Let $P_D = K_D U_D \subseteq D$ be a parabolic subgroup. Then:

Littelmann's restriction rule [**Li2**]: $V_\rho(D)$ decomposes as a K_D-module into the direct sum $\oplus_\eta V_{\eta(1)}(K_D)$, where η runs over all Lakshmibai-Seshadri paths in B_ρ contained in the dominant Weyl chamber of K_D. Here $\eta(1)$ denotes the endpoint of η other than the origin.

Given dominant weights α and β of R and K, we shall say that $\alpha \triangleleft_R^K \beta$ if $V_\alpha(R)$ occurs in $V_\beta(K)$. Since R satisfies eq.(9.1), an explicit relationship between such α and β can be determined by applying the preceding rule–or rather its appropriate variant for connected reductive D–to the representation $V_\beta(K)$.

For the $GL_l(\mathbb{C})$-action in Example 3 above, $K \simeq GL_{\delta+k}(\mathbb{C}) \times GL_{l-\delta-k}$ and $R \simeq SL_k(\mathbb{C}) \times SL_{l-\delta-k}(\mathbb{C})$. Hence, Littelmann's restriction rule reduces to a variant of the well known Pieri's branching rule [**F1**], which gives an explicit decomposition of $V_\mu(GL_{\delta+k}(\mathbb{C}))$ 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 $\hat{v}$. We are now in a position to generalize the earlier Definition 6.3:

Definition 9.2. *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. *$\mu \triangleleft_R^K \beta'$.*

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, 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.

10. Results in a general setting

We are now ready to state the results of part II in geometric invariant theory regarding the general orbit closure problem (Problem 1.1). Let V, f, g be as in Problem 1.1. As in Sections 5, 6 and 7, the problem of explicit construction of obstructions in this context leads one to the following problems in geometric invariant theory.

Let V be a linear representation of a connected, reductive G. Let $v \in P(V)$ be a point that is stable or partially stable with small defect, and having a sufficiently special stabilizer (cf. Definitions 4.1 and 4.3). Then what is algebraic geometry of $\Delta_V[v]$? Specifically, what is the G-module structure of the homogeneous coordinate ring $R_V[v]$? Does the ideal $I_V[v]$ have a nice set of generators?

The simplest example of a partially stable point with defect zero, which is characterized by its stabilizer, corresponds to the highest weight vector of an irreducible G-representation $V = V_\lambda(G)$ with highest weight λ: Then $\Delta_V[v]$ consists of just one orbit, isomorphic to G/P, P a parabolic subgroup of G. Its algebraic geometry is well understood. By the Borel-Weil theorem, $R_V[v]_d \simeq V_{d\lambda}(G)^*$. The standard monomial theory of Lakshmibai, Seshadri, and Littelmann [**LS, Li2, LLM**], and the theory of canonical, global crystal bases of Lusztig[**Ls**] and Kashiwara [**Ks**] give explicit bases for $R_V[v]_d$ and their quantum analogues. An explicit set of quadratic straightening relations that generate the ideal $I_V[v]$ is given by the standard monomial theory [**LLM**] and Kostant's result. What is desired is a generalization of these results, as far as possible, to $\Delta_V[v]$, where v is an arbitrary stable or partially stable point with a small defect and a nice stabilizer. We prove two results in this direction.

First, we analyze which G-modules can occur in $R_V[v]$. The following is a generalization of Proposition 6.5 and Theorem 6.2:

THEOREM 10.1 (cf. part II). *(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 (cf. Definition 4.2).*
(b) More generally, suppose v is partially stable with a small defect, specifically (R,P)-stable (cf. Definition 9.1). 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 9.2). 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.

Part (a) follows from the algebraic version of the Peter-Weyl theorem and an evaluation trick just like Proposition 6.5. The proof of Part (b) is based on the Peter-Weyl theorem, some representation theory (Theorem 5.104 in [**Kn**]) together with the evaluation trick.

For partially stable points, we actually prove a stronger result (cf. Part II) concerning multiplicity of $V_\lambda(G)$ in $R_V[v]_d$ as well as in the module $\Gamma(\Delta_V[v], \mathcal{O}_{\Delta_V[v]}(d))$ of global sections of the twisting sheaf $\mathcal{O}_{\Delta_V[v]}(d)$. When v corresponds to the highest weight vector of an irreducible representation $V = V_\lambda(G)$ of a semisimple G, so that $\Delta_V[v] = G/P$ for some parabolic P, this result specializes to the Borel-Weil theorem [**Kn**]: namely, for $d \geq 0$, $\Gamma(G/P, \mathcal{O}_{G/P}(d)) = V_{d\lambda}(G)^*$, where $\mathcal{O}_{G/P}(d)$ corresponds to the d-th power of the hyperplane bundle.

For finer understanding of the G-module structure of an individual $R_V[v]_d$, it is desirable to know a sufficiently explicit set of generators for the ideal $I_V[v]$ with good representation-theoretic properties (cf. Section 7). We turn to this problem next.

Let v be a stable or partially stable point (with defect zero) of V. If v is stable, we can associate with it the representation-theoretic data Σ_v just as in Section 7. Recall that the stabilizer $G_{\hat{v}}$ in G of $\hat{v} \in V$ is necessarily reductive [**BH, Ms**] if v is stable.

DEFINITION 10.2. *The representation theoretic data Σ_v is defined to be the set of all irreducible, non-$G_{\hat{v}}$-admissible G-submodules of $\mathbb{C}[V]$.*

Let $\tilde{\Sigma}_v$ be the set of appropriately chosen basis elements of the G-submodules in Σ_v. These basis elements will be called *nonadmissible basis elements.* For the definition of Σ_v for partially stable points, see Part II.

By Theorem 10.1 (a) (and an analogous result for partially stable points, cf. part II), the ideal $I_V[v]$ contains all modules in Σ_v. We ask to what extent Σ_v determines the algebraic geometry of $\Delta_V[v]$: specifically, when do the G-modules in Σ_v generate the ideal of $\Delta_V[v]$? This is in the spirit of the "Tannakian" question raised by Larsen and Pink [**LP**], where they showed that the representation-theoretic dimension data associated with a reductive subgroup H of $GL(V)$–which tells the dimension of the subspace of H-invariants in each irreducible $GL(V)$-module–determines H completely if V is an irreducible, faithful representation of H, but not always otherwise.

When v is the highest weight vector of V, so that $\Delta_V[v] = G/P$, P parabolic, it follows from the standard monomial theory [**LLM**] and Kostant's result [**PK**] that the quadratic relations in Σ_v generate $I_V[v]$, thereby providing a satisfactory solution to the question. For the nonspherical $\Delta_V[v]$'s that arise in the approach here, the following result is a step in that direction. It is a generalization of Theorem 7.3.

Theorem 10.3 (cf. part II). *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 4.3) and characterizes v (cf. Definition 4.1). Then the orbit $Gv \subseteq P(V)$ is determined by the representation-theoretic data Σ_v within some G-invariant neighbourhood U of Gv in the sense: $Gv = U \cap \Delta_V[v] = U \cap X(\tilde{\Sigma}_v)$ as a scheme, where $X(\tilde{\Sigma}_v)$ is the zero set of $\tilde{\Sigma}_v$.*

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

Similar result also holds for a class of partially stable points with defect zero.

The proof is based on Luna's étale slice theorem [**Lu**], Weyl's unitary trick [**W**], and Cartan-Oka theorem [**GR1**] on the vanishing of the higher cohomology of a coherent analytic ideal sheaf.

In analogy with the G/P situation, the nonadmissible basis elements in Σ_v may actually generate the ideal of $I_V[v]$ for the specific v's that arise in our lower bound applications; for example, when v corresponds to the permanent or the determinant function (cf. Conjecture 7.1).

11. Future course and open problems

In summary, the following open problems loom ahead in the future course of this approach.

(1) Solve the instances of the subgroup restriction problem that arise in the approach (Section 8.1). This will make the definition of the representation-theoretic data Σ_h (cf. Problem 5.4) sufficiently explicit.

(2) Prove that Σ_h generates the ideal $I_W[h]$ of $\Delta_W[h]$ for the hs that arise in the approach (cf. Conjecture 7.1 and part II).

(3) Choose canonical basis elements of the modules in Σ_h–canonical in the sense of crystal bases [**Ks, Ls**] or standard monomial theory [**LLM**]. Use them as straightening relations, in the spirit of standard monomial theory, to get a finer understanding of the G-module structure of $R_W[h]$.
(4) Use this understanding to solve Problem 8.2.
(5) Finally, use all this to produce an explicit coarse (or less coarse cf. Remark 6.1) obstruction as outlined in Section 8, thereby proving the arithmetic versions of the fundamental lower bound problems considered in this paper.

References

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

[BS] R. Boppana, M. Sipser, *The complexity of finite functions*, Handbook of Theoretical Computer Science, vol. A, Edited by J. van Leeuwen, North Holland, Amsterdam, 1990, pp. 757–804.

[BH] A. Borel and Harish Chandra, *Arithmetic subgroups of algebraic groups*, Ann. of Math. 75, pp. 485–535 (1962).

[BLV] M. Brion, D. Luna, T. Vust, *Espaces homogènes sphériques*, Invent. Math. 84, pp. 617-632, (1986).

[BCS] P. Bürgisser, M. Clausen, M. Shokhrollahi, *Algebraic complexity theory*, Springer-Verlag, 1997.

[CEP] C. De Concini, D. Eisenbud, C. Procesi, *Hodge algebras*, Asterique, 91, Societe mathematique de France, Paris, 1982.

[Co] S. Cook, *The complexity of theorem-proving procedures*, Proceedings of the third annual ACM Symposium on Theory of Computing. pp. 151–158. (1971).

[F1] W. Fulton, J. Harris, *Representation theory*, Springer-Verlag, 1991.

[F2] W. Fulton, *Young Tableaux*, Cambridge university press, 1997.

[F3] W. Fulton, *Introduction to toric varieties*, Princeton university press, 1993.

[Ga] J. von zur Gathen, *Permanent and determinant*, Linear algebra and its applications, 96: pp. 87–100 (1987).

[GR1] H. Grauert, R. Remmert, *Coherent analytic sheaves*, Springer-Verlag, 1984.

[Hi] D. Hilbert, *Über die vollen Invariantensysteme*, Math. Ann. 42, pp. 313–373 (1893).

[K] R. Karp, *Reducibility among combinatorial problems*, R. E. Miller and J. W. Thatcher (eds.) Complexity of computer computations, Plenum Press, New York, 1972, pp. 85–103.

[Ks] M. Kashiwara, *On crystal bases of the q-analogue of the universal enveloping algebras*, Duke Math. J. 63 (1991), pp. 465–516.

[Ke] G. Kempf, *Instability in invariant theory*, Annals of Mathematics, 108 (1978), pp. 299–316.

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

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

[LLM] V. Lakshmibai, P. Littelmann, P. Magyar, *Standard monomial theory and applications*, in Representation theories and algebraic geometry, (Ed. A. Broer), Kluwer Academic Publishers, (1997), pp. 319–364.

[LP] M. Larsen, R. Pink, *Determining representations from invariant dimensions*, Invent. math. 102, pp. 377–398 (1990).

[Le] A. Levin, *Universal sequential search problems*, Problems of information transmission (translated from Problemy Peredachi Informatsii (Russian)) 9 (1973).

[Li1] P. Littelmann, *A littlewood-Richardson rule for symmetrizable Kac-Moody algebras*, Invent. Math. 116, pp. 329–346 (1994).

[Li2] P. Littelmann, *Paths and root operators in representation theory*, Ann. of Math. 142 (1995), pp. 499–525.

[Li3] P. Littelmann, *Contracting modules and standard monomial theory for symmetrizable Kac–Moody algebras*, JAMS 11, pp. 551-567, (1998).

[GLS] M. Grotschel, L. Lovász, A. Schriver, *Geometric algorithms and combinatorial optimization*, Springer-Verlag, 1988.

[LPS] A. Lubotzky, R. Phillips, P. Sarnak, *Ramanujan graphs*, Combinatorica 8 (1988), pp. 261–277.

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

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

[Ls] G. Lusztig, *Canonical bases arising from quantized enveloping algebras*, J. Amer. Math. Soc. 3 (1990), pp. 447–498.

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

[Mr] G. Margulis, *Explicit constructions of concentrators*, Problemy Inf. Trans. 9 (1973), pp. 325–332.

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

[PRAM] K. Mulmuley, *Lower bounds in a parallel model without bit operations*, The SIAM Journal On Computing, vol. 28, no. 4, 1999.

[GCT1] K. Mulmuley, M. Sohoni, *Geometric complexity theory I: an approach to the P vs. NP and related problems*, SIAM Journal On Computing, vol. 31, no. 2, pp 496–526, 2001.

[GCT2] K. Mulmuley, M. Sohoni, *Geometric complexity theory II: explicit obstructions*, manuscript;

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

[PK] D. Peterson, V. Kac, *Infinite flag varieties and conjugacy theorems*, Proc. Natl. Acad. Sci. USA, vol. 80, pp. 1778-1782, (1983).

[P] V. Popov, E. Vinberg, *Invariant theory*, in Encyclopaedia of mathematical sciences, vol. 55, Springer Verlag, 1991.

[RR] A. Razborov, S. Rudich, *Natural proofs*, JCSS 55 (1997), pp. 24–35.

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

[St] V. Strassen, *Relative bilinear complexity and matrix multiplication*, J. Reine Angew. Math. 375/376, pp. 406-443 (1987).

[V] L. Valiant, *The complexity of computing the permanent*, Theoretical Computer Science 8, pp. 189–201 (1979).

[W] H. Weyl, *The classical groups. Their invariants and representations*, Princeton University Press. 302 (1939).

Comp. Sci. Dept., The university of Chicago, and I.I.T., Mumbai.
E-mail address: `Ramakrishnadas@cs.uchicago.edu; Ramakrishnadas@cse.iitb.ac.in`
URL: `Ramakrishnadas.cs.uchicago.edu; www.cse.iitb.ac.in/~Ramakrishnadas`

Comp. Sci. Dept., I.I.T., Mumbai
E-mail address: `sohoni@cse.iitb.ernet.in`
URL: `www.cse.iitb.ac.in/~sohoni`