MODULI OF REPRESENTATIONS OF FINITE DIMENSIONAL ALGEBRAS

By A. D. KING

1. Introduction

In this paper, we present a framework for studying moduli spaces of finite dimensional representations of an arbitrary finite dimensional algebra A over an algebraically closed field k. (The abelian category of such representations is denoted by mod-A.) Our motivation is twofold. Firstly, such moduli spaces should play an important role in organising the representation theory of wild algebras. Secondly, such moduli spaces can be identified with moduli spaces of vector bundles on special projective varieties. This identification is somewhat hidden in earlier work ([6], [7]) but has become more explicit recently ([4], [12]). It can now be seen to arise from a 'tilting equivalence' between the derived category of mod-A and the derived category of coherent sheaves on the variety.

It is well-established that mod-A is equivalent to the category of representations of an arrow diagram, or 'quiver', Q by linear maps satisfying certain 'admissible' relations. Thus, the problem of classifying A-modules with a fixed class in the Grothendieck group $K_0(\text{mod-}A)$, represented by a 'dimension vector' α, is converted into one of classifying orbits for the action of a reductive algebraic group $GL(\alpha)$ on a subvariety $V_A(\alpha)$ of the representation space $\mathcal{R}(Q, \alpha)$ of the quiver.

Now, the moduli spaces provided by classical invariant theory ([1], [18]) are not interesting in this context. This is because the classical theory only picks out the closed $GL(\alpha)$-orbits in $V_A(\alpha)$, which correspond to semisimple A-modules, and the quiver Q is chosen so that there is only one semisimple A-module of each dimension vector. On the other hand, we can apply Mumford's geometric invariant theory, with the trivial linearisation twisted by a character χ of $GL(\alpha)$, which restricts our attention to an open subset of $V_A(\alpha)$, consisting of semistable representations. Within this open set there are more closed orbits and the corresponding algebraic quotient is then a more interesting moduli space. In fact, this approach also has a classical flavour, since it involves the relative (or semi-) invariants of the $GL(\alpha)$ action.

The main purpose of this paper is to show that the notions of semistability and stability, that arise from the geometric invariant theory, coincide with more algebraic notions, expressed in the language of mod-A. Indeed, the definition is formulated for an arbitrary abelian category as follows:

DEFINITION 1.1. Let $\mathscr{A}$ be an abelian category and $\theta\colon K_0(\mathscr{A}) \to \mathbb{R}$ an additive function on the Grothendieck group. (We shall call θ a character of $\mathscr{A}$.) An object $M \in \mathscr{A}$ is called θ-semistable if $\theta(M) = 0$ and every subobject $M' \subseteq M$ satisfies $\theta(M') \geq 0$. Such an M is called θ-stable if the only subobjects M' with $\theta(M') = 0$ are M and 0.

In the light of the relation to vector bundles, mentioned in the first paragraph, it is important to note that this definition has as another special case Mumford's definition of stability for vector bundles on a smooth projective curve [15] and its subsequent generalisation, as μ-stability, to other projective varieties [22].

The central result of the paper (Theorem 4.1 together with Propositions 4.2 and 5.2) shows that, when θ takes integer values, there is an associated character χ_θ for $GL(\alpha)$ such that the algebraic quotient of $V_A(\alpha)$ by $GL(\alpha)$ with respect to the linearisation χ_θ, is a coarse moduli space for families of θ-semistable modules up to S-equivalence. The notion of S-equivalence was originally introduced by Seshadri [21] for vector bundles on curves. Two θ-semistable modules are S-equivalent if they have the same composition factors in the full abelian subcategory of θ-semistable modules. The simple objects in this subcategory are the θ-stable modules.

We denote the moduli space thus constructed by $\mathscr{M}_A(\alpha, \theta)$ and show that it is a projective variety (Proposition 4.3), which may be a little surprising, since it is constructed as the quotient of an affine variety. General theory also implies that $\mathscr{M}_A(\alpha, \theta)$ contains an open set $\mathscr{M}_A^s(\alpha, \theta)$ which is a coarse moduli space for families of θ-stable modules up to isomorphism. We show (Proposition 4.4) that if A is hereditary, then $\mathscr{M}_A^s(\alpha, \theta)$ is nonempty for some θ if and only if α is a 'Schur root' (as defined in [9]). We also give a criterion (Proposition 5.3) for the existence of a universal family over $\mathscr{M}_A^s(\alpha, \theta)$. Finally, when $k = \mathbb{C}$, we give a more analytic construction of $\mathscr{M}_A(\alpha, \theta)$ as a 'symplectic quotient' (or Marsden–Weinstein reduction).

The paper is organised as follows. In Section 2, we present the main results from geometric invariant theory in a form tailored to the applications in the rest of the paper. In Section 3, we prove the central result for representations of quivers. This is the key case, since there is no interaction between the notions of stability and the relations which determine the A-modules. In Section 4, we explain how the results of Section 3 apply to an arbitrary algebra A and we prove some of the additional properties of the moduli space mentioned above. In Section 5, we discuss families of A-modules and show that $\mathscr{M}_A(\alpha, \theta)$ is a moduli space in the strict sense of the term. In Section 6, we describe the symplectic quotient construction of the moduli space.

General references for the material discussed in the paper are [2] and [19] for the representation theory of algebras and quivers, [16] and [17] for geometric invariant theory and moduli spaces, and [13] for symplectic quotients. In addition, [14] contains much on the application of invariant theory to representation theory.

2. A numerical criterion

In this section, we describe the version of Mumford's numerical criterion that applies when a reductive algebraic group G acts linearly on a finite dimensional vector space $\mathcal{R}$. The polarising line bundle L, required by geometric invariant theory, is necessarily trivial, but the lift of the G-action to L can be given by an arbitrary character χ of G. To the best of our knowledge, an explicit description of this case (in particular the statements of Propositions 2.5 and 2.6, which we subsequently need) does not appear in the literature and we feel that it merits a separate explanation.

We shall write the action of G on $\mathcal{R}$ on the left as $G \times \mathcal{R} \rightarrow \mathcal{R}$: $(g, x) \mapsto g \cdot x$. Given a character $\chi \colon G \rightarrow k^\times$, we lift the G-action to the trivial line bundle L as follows. On the total space of L^{-1}, written explicitly as $\mathcal{R} \times k$, G acts by $g \cdot (x, z) = (g \cdot x, \chi^{-1}(g)z)$. An invariant section of L^n is then a function $f(x)z^n \in k[\mathcal{R} \times k]$, where $f(x) \in k[\mathcal{R}]$ is a relative invariant of weight χ^n. Recall that a function $f \in k[\mathcal{R}]$ is a relative invariant of weight χ if $f(g \cdot x) = \chi(g)f(x)$. (We shall write $k[\mathcal{R}]^{G,\chi}$ for the space of such relatively invariant functions.)

In this case, Mumford's definitions of semistability and stability can be phrased as in Definition 2.1 below. We use the terms 'χ-semistable'/'χ-stable' to mean semistable/stable with respect to the linearisation determined, as above, by the character χ.

Warning. We wish to allow the representation to have a kernel Δ, so we adopt a variant of the definition of 'stable' which is a compromise between Mumford's original definition and the more commonly used one which Mumford called 'properly stable'.

DEFINITION 2.1. (i) A point $x \in \mathcal{R}$ is χ-*semistable* if there is a relative invariant $f \in k[\mathcal{R}]^{G,\chi^n}$ with $n \geq 1$, such that $f(x) \neq 0$.

(ii) A point $x \in \mathcal{R}$ is χ-*stable* if there is a relative invariant $f \in k[\mathcal{R}]^{G,\chi^n}$ with $n \geq 1$, such that $f(x) \neq 0$ and, further, $\dim G \cdot x = \dim G/\Delta$ and the G-action on $\{x \in \mathcal{R} \mid f(x) \neq 0\}$ is closed.

Hence, the corresponding algebraic quotient of $\mathcal{R}$ by G, which we shall denote by $\mathcal{R}/\!/(G, \chi)$ can be described as

$$\mathcal{R}/\!/(G, \chi) = \mathrm{Proj}\left(\bigoplus_{n \geq 0} k[\mathcal{R}]^{G,\chi^n}\right)$$

which is projective over the ordinary quotient $\mathscr{R}//G = \text{Spec}\,(k[\mathscr{R}]^G)$. In particular, if $k[\mathscr{R}]^G = k$, then $\mathscr{R}//(G, \chi)$ is a projective variety.

The central point of geometric invariant theory, is that this quotient has a more geometric description, as the quotient of the open set $\mathscr{R}_\chi^{ss}$ of χ-semistable points by the equivalence relation: $x \sim y$ if and only if the orbit closures $\overline{G \cdot x}$ and $\overline{G \cdot y}$ intersect (in $\mathscr{R}_\chi^{ss}$). We shall call two semistable points or orbits identified by this relation 'GIT equivalent'. Since each orbit closure contains a unique closed orbit, the points of the quotient are in one-one correspondence with the closed orbits in $\mathscr{R}_\chi^{ss}$. In particular, there is an open subset of the quotient corresponding to the χ-stable orbits, all of which are closed.

There is also a more geometric characterisation of when points are semistable or stable, which uses the action of G lifted to $\mathscr{R} \times k$, i.e. the total space of L^{-1}, as described above.

LEMMA 2.2. *Lift* $x \in \mathscr{R}$ *to a point* $\hat{x} = (x, z) \in \mathscr{R} \times k$ *with* $z \neq 0$. *Then*

(i) *x is χ-semistable if and only if the orbit closure* $\overline{G \cdot \hat{x}} \subseteq \mathscr{R} \times k$ *is disjoint from the zero-section* $\mathscr{R} \times \{0\}$. *In particular, it is necessary that* $\chi(\Delta) = \{1\}$.

(ii) *x is χ-stable if and only if* $G \cdot \hat{x}$ *is closed and the stabiliser of* $\hat{x}$ *contains* Δ *with finite index.*

Proof. Use the fact that G is geometrically reductive, i.e. that disjoint closed G-sets in affine space can be distinguished by G-invariant functions.

Remark 2.3. The GIT equivalence relation on $\mathscr{R}_\chi^{ss}$ can also be described by: $x \sim y$ if and only if there are lifts $\hat{x}$ and $\hat{y}$ for which $\overline{G \cdot \hat{x}}$ and $\overline{G \cdot \hat{y}}$ intersect (in $\mathscr{R} \times k$). In particular, the orbits $G \cdot x$ which are closed in $\mathscr{R}_\chi^{ss}$ are precisely those whose lifted orbits $G \cdot \hat{x}$ are closed in $\mathscr{R} \times k$.

Using Lemma 2.2, we obtain a version of Hilbert's Lemma characterising semistable and stable points by the behaviour of their lifts under the action of the one-parameter subgroups of G.

LEMMA 2.4. *Let* $\hat{x} \in \mathscr{R} \times k$ *be a lift of* $x \in \mathscr{R}$, *as above. Then*

(i) *x is χ-semistable if and only if, for all one-parameter subgroups λ of G, $\lim_{t \to 0} \lambda(t) \cdot \hat{x} \notin \mathscr{R} \times \{0\}$.*

(ii) *x is χ-stable if and only if the only one-parameter subgroups λ of G, for which $\lim_{t \to 0} \lambda(t) \cdot \hat{x}$ exists, are in Δ.*

Proof. Use the 'fundamental theorem' ([10] Theorem 1.4) that any

closed G-set that meets the closure of a G-orbit contains a point in the closure of some one-parameter subgroup orbit.

This result can be reformulated in terms of the original G action on $\mathcal{R}$ by introducing the integral pairing between one-parameter subgroups λ and characters χ, defined by $\langle \chi, \lambda \rangle = m$ when $\chi(\lambda(t)) = t^m$. This yields 'Mumford's Numerical Criterion'.

PROPOSITION 2.5. *A point $x \in \mathcal{R}$ is χ-semistable if and only if $\chi(\Delta) = \{1\}$ and every one parameter subgroup λ of G, for which $\lim_{t \to 0} \lambda(t) \cdot x$ exists, satisfies $\langle \chi, \lambda \rangle \geqslant 0$. Such a point is χ-stable if and only if the only one-parameter subgroups λ of G, for which $\lim_{t \to 0} \lambda(t) \cdot x$ exists and $\langle \chi, \lambda \rangle = 0$, are in Δ.*

It is this criterion that we will use to identify stable and semistable representations of quivers. Note the similarity in form with Definition 1.1. It will also be useful to characterise the GIT equivalence relation on semistable points in terms of the original G-action on $\mathcal{R}$ and its one parameter subgroups.

PROPOSITION 2.6. (i) *An orbit $G \cdot x$ is closed in $\mathcal{R}_\chi^{ss}$ if and only if, for every one parameter subgroup λ with $\langle \chi, \lambda \rangle = 0$, when the limit $\lim_{t \to 0} \lambda(t) \cdot x$ exists, it is in $G \cdot x$.*

(ii) *If $x, y \in \mathcal{R}_\chi^{ss}$, then $x \sim y$ if and only if there are one parameter subgroups λ_1, λ_2 such that $\langle \chi, \lambda_1 \rangle = \langle \chi, \lambda_2 \rangle = 0$ and $\lim_{t \to 0} \lambda_1(t) \cdot x$ and $\lim_{t \to 0} \lambda_2(t) \cdot y$ are in the same closed G-orbit.*

Proof. Use Remark 2.3 and the fundamental theorem used in Lemma 2.4.

3. Representations of quivers

We now use the results of the previous section to reinterpret χ-semistability, χ-stability and GIT equivalence for representations of quivers in the language of the abelian category of such representations. We start by recalling some basics to fix notation.

A quiver Q is a diagram of arrows, specified combinatorially by two finite sets Q_0 (of 'vertices') and Q_1 (of 'arrows') with two maps $h, t \colon Q_1 \to Q_0$ which indicate the vertices at the head and tail of each arrow. A representation of Q consists of a collection of k-vector spaces W_v, for each $v \in Q_0$, together with k-linear maps $\phi_a \colon W_{ta} \to W_{ha}$, for each $a \in Q_1$. The dimension vector $\alpha \in \mathbb{Z}^{Q_0}$ of such a representation is given by $\alpha_v = \dim_k W_v$. A map between representations (W_v, ϕ_a) and (U_v, ψ_a) is given by linear maps $f_v \colon W_v \to U_v$, for each $v \in Q_0$, such that $f_{ha} \phi_a = \psi_a f_{ta}$, for each $a \in Q_1$. Such a map is an isomorphism if and only

if each f_v is. The abelian category of representations of Q will be denoted by mod-kQ, since it is the same as the category of finite dimensional representations of the path algebra kQ.

Having chosen vector spaces W_v of dimension α_v, the isomorphism classes of representations of Q with dimension vector α are in natural one-one correspondence with the orbits in the representation space

$$\mathcal{R}(Q, \alpha) = \bigoplus_{a \in Q_1} \text{Hom}\,(W_{ta}, W_{ha}).$$

of the symmetry group

$$GL(\alpha) = \prod_{v \in Q_0} GL(W_v)$$

acting by $(g \cdot \phi)_a = g_{ha}\phi_a g_{ta}^{-1}$. Note that this contains the diagonal one-parameter subgroup $\Delta = \{(t1, \ldots, t1): t \in k^\times\}$ acting trivially.

We now wish to apply Propositions 2.5 and 2.6 to the action of $GL(\alpha)$ on $\mathcal{R}(Q, \alpha)$. We first note that the characters of $GL(\alpha)$ are given by

$$\chi_\theta(g) = \prod_{v \in Q_0} \det\,(g_v)^{\theta_v}$$

for $\theta \in \mathbb{Z}^{Q_0}$. Such a function θ can also be interpreted as a homomorphism $K_0(\text{mod-}kQ) \to \mathbb{Z}$ as follows. Let $M = (U_v, \psi_a)$ be a representation of Q, set $\theta(M) = \sum_v \theta_v \dim U_v$, and observe that this is additive on short exact sequences. The condition $\chi_\theta(\Delta) = \{1\}$ becomes $\sum_v \theta_v \alpha_v = 0$, i.e. $\theta(M) = 0$ if M has dimension vector α.

We next observe that one-parameter subgroups of $GL(\alpha)$ correspond to filtrations. This requires a small modification of the argument in [14] Proposition 4.3. Let $\lambda: k^\times \to GL(\alpha)$ be a one-parameter subgroup and, for each $v \in Q_0$, make the decomposition

$$W_v = \bigoplus_{n \in \mathbb{Z}} W_v^{(n)},$$

where $\lambda(t)$ acts on the weight space $W_v^{(n)}$ as multiplication by t^n, and define the filtrations

$$W_v^{(\geqslant n)} = \bigoplus_{m \geqslant n} W_v^{(m)}.$$

Under the action of λ the components of the arrow maps $\phi_a^{(mn)}: W_{ta}^{(n)} \to W_{ha}^{(m)}$ are multiplied by t^{m-n}. Hence, $\lim_{t \to 0} \lambda(t)\phi_a$ exists if and only if $\phi_a^{(mn)} = 0$, for all $m < n$. This in turn happens if and only if ϕ_a gives a map $W_{ta}^{(\geqslant n)} \to W_{ha}^{(\geqslant n)}$, for all n, i.e. if and only if the subspaces $W_v^{(\geqslant n)}$

determine subrepresentations M_n of M, for all n. Thus a one-parameter subgroup λ, for which $\lim_{t \to 0} \lambda(t)\phi_a$ exists, determines a filtration of M,

$$\cdots \supseteq M_n \supseteq M_{n+1} \supseteq \cdots$$

indexed by $\mathbb{Z}$ and such that $M_n = M$ for $n \ll 0$ and $M_n = 0$ for $n \gg 0$. Furthermore, it is clear that every such '$\mathbb{Z}$-filtration' is associated to some (not necessarily unique) one-parameter subgroup λ for which $\lim_{t \to 0} \lambda(t)\phi_a$ exists. This limit is then equal to

$$\bigoplus_{n \in \mathbb{Z}} (W_v^{(n)}, \phi_a^{(nn)}) = \bigoplus_{n \in \mathbb{Z}} M_n/M_{n+1}$$

i.e. it is the graded representation associated to the filtration. Observe also that the filtration determined by λ is proper, i.e. some M_n is neither M nor 0, unless λ is in Δ.

Finally, provided $\sum_v \theta_v \alpha_v = 0$, the pairing $\langle \chi_\theta, \lambda \rangle$ has a simple expression in terms of the filtration $(M_n)_{n \in \mathbb{Z}}$, namely

$$\langle \chi_\theta, \lambda \rangle = \sum_{v \in Q_0} \theta_v \sum_{n \in \mathbb{Z}} n \dim W_v^{(n)}$$

$$= \sum_{n \in \mathbb{Z}} n\theta(M_n/M_{n+1})$$

$$= \sum_{n \in \mathbb{Z}} \theta(M_n).$$

Note that the last equality requires the fact that $\theta(M_n) = 0$, for all but finitely many n.

PROPOSITION 3.1. *A point in $\mathcal{R}(Q, \alpha)$ corresponding to a representation $M \in \text{mod-}kQ$ is χ_θ-semistable (resp. χ_θ-stable) if and only if M is θ-semistable (resp. θ-stable).*

Proof. The 'if' part is now immediate from Proposition 2.5. For the 'only if' part, observe that a subrepresentation $M' \subseteq M$ induces a $\mathbb{Z}$-filtration with $M_i = M'$, for (any) one value of i, and $M_n = M$ or 0, otherwise. The filtration is proper if and only if M' is proper. For a corresponding one-parameter subgroup λ, we have $\langle \chi_\theta, \lambda \rangle = \theta(M')$.

Now for any map between θ-semistable representations, the kernel, image and cokernel of the map are all θ-semistable and hence the θ-semistable representations from an abelian subcategory of mod-kQ. furthermore, the simple objects in this subcategory are precisely the θ-stable répresentations. Since this category is noetherian and artinian,

the Jordan-Hölder theorem holds ([21] Theorem 2.1) and so we call two θ-semistable representations 'S-equivalent' if they have the same composition factors in the category of θ-semistable representations. We can thus use Proposition 2.6 to show that two θ-semistable representations are S-equivalent if and only if the corresponding points in $\mathcal{R}(Q, \alpha)$ are GIT equivalent.

PROPOSITION 3.2. (i) *A θ-semistable representation M corresponds to a GL(α)-orbit which is closed in $\mathcal{R}_\chi^{ss}(Q, \alpha)$, if and only if M is a direct sum of θ-stable modules.*

(ii) *Two θ-semistable representations correspond to GIT equivalent GL(α)-orbits if and only if they have filtrations with θ-stable quotients and the same associated graded representation.*

Proof. First observe that, if M is θ-semistable, then a $\mathbb{Z}$-filtration $\{M_i\}$ corresponds to a one-parameter subgroup λ with $\langle \chi_\theta, \lambda \rangle = 0$ if and only if each M_i is θ-semistable. Now recall the earlier observation that taking the limit under a one-parameter subgroup corresponds to taking the associated graded representation of a filtration. Thus Proposition 2.6(i) becomes: M is isomorphic to the associated graded of any filtration, in particular a maximal filtration in which all the quotients are θ-stable. Proposition 2.6(ii) immediately gives part (ii) above.

4. Representations of algebras

In this section, we adapt the results just proved for quivers to prove our main theorem for an arbitrary finite dimensional algebra A. We first explain the various correspondences mentioned in the statement of the theorem.

Given any finite dimensional algebra A, the subcategory $\mathcal{P}_A \subseteq$ mod-A of projective A-modules is generated by a quiver Q, subject to certain linear relations amongst its 'paths', i.e. formal composites of arrows. The vertices of Q are in one-one correspondence with the isomorphism classes of indecomposable projective modules. Given any $M \in$ mod-A, the functor $\mathrm{Hom}_A(-, M)$ restricted to $\mathcal{P}_A$ determines and is determined by a representation of Q. The module M can be recovered from this restricted functor, because $A \in \mathcal{P}_A$, while $\mathrm{Hom}_A(A, M) = M$ with the natural action of $\mathrm{End}_A(A) = A$. In this way, mod-A can be realised as an abelian subcategory of mod-kQ. The vertices of Q are also in one-one correspondence with the isomorphism classes of simple A-modules and thus, since the Jordan–Hölder theorem holds in mod-A, the Grothendieck group $K_0(\text{mod-}A)$ is naturally the free abelian group generated by Q_0. Hence an integer-valued character θ of mod-A is just an element of $\mathbb{Z}^{Q_0}$, which determines a character of $GL(\alpha)$, as described in Section 3.

THEOREM 4.1. *Let A be a finite dimensional algebra over an algebraically closed field k and let Q be the associated quiver. Let M be a finite dimensional A-module, α its dimension vector and $x \in \mathcal{R}(Q, \alpha)$ a point of the representation space that corresponds to M. Let $\theta\colon K_0(\text{mod-}A) \to \mathbb{Z}$ be a character of $\text{mod-}A$ and $\chi_\theta\colon GL(\alpha) \to k^\times$ the corresponding character of $GL(\alpha)$.*

Then x is a χ_θ-semistable (resp. χ_θ-stable) point for the action of $GL(\alpha)$ on $\mathcal{R}(Q, \alpha)$, if and only if M is a θ-semistable (resp. θ-stable) module.

Proof. We first note that mod-A, as a subcategory of mod-kQ, is closed under taking arbitrary subobjects. Furthermore, a character θ of mod-A, being just an element of $\mathbb{Z}^{Q_0}$, extends to a character of mod-kQ, which we shall also denote by θ. It is then clear that an A-module M is θ-semistable (resp. θ-stable) as an element of mod-A if and only if the corresponding representation of Q is θ-semistable (resp. θ-stable) as an element of mod-kQ. But, by Proposition 3.1, this is in turn the case if and only if the corresponding point $x \in \mathcal{R}(Q, \alpha)$ is χ_θ-semistable (resp. χ_θ-stable).

We further observe that θ-semistable modules in mod-A are S-equivalent in mod-A if and only if they are S-equivalent in mod-kQ. Then Proposition 3.2 becomes

PROPOSITION 4.2. *Let x, y be χ_θ semistable points of $\mathcal{R}(Q, \alpha)$ corresponding to θ-semistable A-modules M, N. Then M and N are S-equivalent if and only if x and y are GIT equivalent.*

The representations in $\mathcal{R}(Q, \alpha)$ which are associated to A-modules form a closed $GL(\alpha)$-invariant subvariety $V_A(\alpha)$ and thus the quotient $V_A(\alpha)/\!/(GL(\alpha), \chi_\theta)$ is a quasi-projective variety whose points are in natural one-one correspondence with the S-equivalence classes of θ-semistable A-modules. We shall denote this variety by $\mathcal{M}_A(\alpha, \theta)$ and refer to it as the 'moduli space of θ-semistable A-modules of dimension α'. The justification for using the term 'moduli space' will be given in the next section.

PROPOSITION 4.3. *The space $\mathcal{M}_A(\alpha, \theta)$ is a projective variety.*

Proof. As observed near the beginning of Section 2, $V_A(\alpha)/\!/(GL(\alpha), \chi_\theta)$ will be projective over the ordinary quotient $V_A(\alpha)/\!/GL(\alpha)$. But the points of $V_A(\alpha)/\!/GL(\alpha)$ correspond to $GL(\alpha)$ orbits which are closed in $V_A(\alpha)$, which in turn correspond to semisimple A-modules. The quiver Q is chosen so that there is a unique semi-simple A-module with dimension vector α. Hence $V_A(\alpha)/\!/GL(\alpha)$ consists of a single point and thus $V_A(\alpha)/\!/(GL(\alpha), \chi_\theta)$ is projective.

In addition, standard geometric invariant theory implies that $\mathcal{M}_A(\alpha, \theta)$ contains an open set $\mathcal{M}_A^s(\alpha, \theta)$, whose points correspond to isomorphism classes of θ-stable A-modules. All other properties of $\mathcal{M}_A(\alpha, \theta)$, e.g. normality, irreducibility, even nonemptiness, will depend on the properties of the affine variety $V_A(\alpha)$ and thus on A. The most special case is when A is a hereditary algebra, i.e. when all the representations of Q correspond to A-modules. In this case, we can say that $\mathcal{M}_A(\alpha, \theta)$ is irreducible and normal and that $\mathcal{M}_A^s(\alpha, \theta)$ is smooth. We can also give a precise criterion for when $\mathcal{M}_A^s(\alpha, \theta)$ is non-empty.

PROPOSITION 4.4. *If A is a finite dimensional hereditary algebra, then there is some θ for which $\mathcal{M}_A^s(\alpha, \theta)$ is nonempty if and only if α is a Schur root.*

Proof. We need to know when $\mathcal{R}(Q, \alpha)$ has a χ_θ-stable point for some θ. For this, it is clearly necessary for some (and hence the generic) point to have a zero-dimensional stabiliser in the group $GL(\alpha)/\Delta$. (This actually means that the stabiliser is trivial because it is always connected.) Now, it turns out that, by a result of Van den Bergh ([3] Proposition 6), this condition is also sufficient. The result states that, if there is a point with zero-dimensional stabiliser, then there is an invariant affine open set in which the generic orbit is closed. This open set will be defined by the non-vanishing of a relative invariant function of some weight χ_θ and so the generic point will be χ_θ-stable.

A point has trivial stabiliser if and only if the endomorphism algebra of the corresponding module is just k. Such a module is called Schurian and a dimension vector, for which some (and hence the generic) module is Schurian, is called a Schur root [9].

Remark 4.5. The determination of the Schur roots was a problem posed by Kac and solved recently by Schofield [20]. The solution is based on giving a procedure for calculating the 'generic subvectors' of α, i.e. those dimension vectors β such that (i) every representation of dimension α has a subrepresentation of dimension β and (ii) the generic representation of dimension α has subrepresentations of no other dimension vectors. One could, in principle, apply Proposition 4.4 in reverse to determine the Schur roots. However, Schofield actually proves something rather stronger ([20] Theorem 6.1), namely that α is a Schur root if and only if the generic representation of dimension α is Θ_α-stable, i.e. $\Theta_\alpha(\beta) > 0$ for all generic subvectors β. Here Θ_α is a 'canonical character' associated to α, given by $\Theta_\alpha(\beta) = \varepsilon(\beta, \alpha) - \varepsilon(\alpha, \beta)$, where ε is the Euler inner product

$$\varepsilon(\alpha, \beta) = \sum_{v \in Q_0} \alpha_v \beta_v - \sum_{a \in Q_1} \alpha_{ta} \beta_{ha}$$

which is the natural bilinear form on $K_0(\text{mod-}A)$ induced by

$$\varepsilon(M, N) = \dim \text{Hom}\,(M, N) - \dim \text{Ext}^1\,(M, N).$$

(Note: A is hereditary so the higher Ext's vanish.)

Remark 4.6. The notions of generic representation and generic subvector, described in the previous remark, are very useful. They also allow us to calculate when $\mathcal{M}_A(\alpha, \theta)$ is non-empty and show that, when it is, the generic representations form a dense open subset. Hence, the birational equivalence class of $\mathcal{M}_A(\alpha, \theta)$ is independent of θ. The birational transformations which occur as θ varies have been studied in more general contexts in [5] and [8].

5. Families of A-modules and representations of Q

To justify using the term 'moduli space' to describe $\mathcal{M}_A(\alpha, \theta)$ we must show that it has a suitable universal property with respect to families of A-modules. The appropriate notion of family is as follows.

DEFINITION. 5.1. A family of A-modules over a connected variety (or more generally scheme) B is a locally-free sheaf $\mathcal{F}$ over B together with a k-algebra homomorphism $A \to \text{End}\,(\mathcal{F})$. On the other hand, a family of representations of Q is simply a representation of Q in the category of locally free sheaves over B, i.e. a locally free sheaf $\mathcal{W}_v$, for each $v \in Q_0$, together with the appropriate sheaf maps, for each arrow.

The association between A-modules and representations of Q extends naturally to families. In particular, all the A-modules in a family, i.e. all the fibres $\mathcal{F}_x$, have the same dimension vector, given by the ranks of the locally free sheaves $\mathcal{W}_v$.

PROPOSITION 5.2. $\mathcal{M}_A(\alpha, \theta)$ *is a coarse moduli space for families of θ-semistable modules of dimension vector α, up to S-equivalence.*

Proof. If all the modules in a family $\mathcal{F}$ are θ-semistable, then there is a uniquely defined set-theoretic map $B \to \mathcal{M}_A(\alpha, \theta)$, taking a point x to the point representing the S-equivalence class of the fibre $\mathcal{F}_x$. The content of the proposition is that this map is algebraic. If we associate to $\mathcal{F}$ the corresponding family of representations of Q, then, about any point $x \in B$, there is an affine open neighbourhood U on which each locally free sheaf $\mathcal{W}_v$ is (algebraically) trivial. Choosing particular trivialisations, we thus obtain an algebraic map $U \to \mathcal{R}_\theta^{ss}(Q, \alpha)$ whose image is in $V_A(\alpha)$. Composing this map with the quotient map, gives the map $B \to \mathcal{M}_A(\alpha, \theta)$ described above. This map is thus algebraic on U and the sets U cover B.

The second question to ask is whether the moduli space itself carries a universal family, i.e. whether it is a 'fine' moduli space. One would expect obstructions to the existence of such a family to arise from semistable points and from the fact that $GL(\alpha)$ always acts with a kernel Δ. If we ignore the semistable points, then we can get around the second problem in some cases.

PROPOSITION 5.3. *When α is an indivisible vector, $\mathcal{M}_A^s(\alpha, \theta)$ is a fine moduli space for families of θ-stable A-modules.*

Proof. There is a tautological family of representations of Q over $\mathcal{R}(Q, \alpha)$ given by the trivial sheaves with fibre W_v. Each total space $\mathcal{R}(Q, \alpha) \times W_v$ carries an action of $GL(\alpha)$ in which Δ acts with weight 1 along the fibre, i.e. the W_v factor. In the open set $\mathcal{R}_\theta^s(Q, \alpha)$ of θ-stable representations, the stabiliser of each orbit is just Δ. Hence, to make the tautological family descend to the geometric quotient $\mathcal{R}_\theta^s(Q, \alpha)/GL(\alpha)$, we need to multiply the $GL(\alpha)$-action on $\mathcal{R}(Q, \alpha) \times W_v$ by a homomorphism $GL(\alpha) \to \Delta$ (note that Δ is in the centre of $GL(\alpha)$) so that Δ now acts trivially along the fibre. Choosing such a homomorphism amounts to choosing a character which has weight -1 when restricted to Δ. Such a character is determined by $\psi \in \mathbb{Z}^{Q_0}$ such that $\sum_{v \in Q_0} \psi_v \alpha_v = -1$ and for such a ψ to exist it is necessary and sufficient for α to be indivisible, i.e. not a non-trivial multiple of another integer vector. The universal representation of Q thus constructed restricts to the subvariety $\mathcal{M}_A^s(\alpha, \theta)$ of $\mathcal{R}_\theta^s(Q, \alpha)/GL(\alpha)$ and induces the required universal family of A-modules.

Remark 5.4. If α is indivisible then for a set of values of $\theta \in \mathbb{R}^{Q_0}$ which is open (and dense in the set of values for which $\mathcal{M}_A^s(\alpha, \theta)$ is nonempty) there will be no strictly semistable points. For such values of θ, we see that $\mathcal{M}_A^s(\alpha, \theta) = \mathcal{M}_A(\alpha, \theta)$ and hence is a projective variety. When A is hereditary then it is also smooth. It is then natural to ask how much of the cohomology of this smooth projective variety comes from the Chern classes of (the summands of) the universal bundle—and natural to conjecture that the cohomology ring is generated by these Chern classes.*

6. The symplectic quotient construction of the moduli space

In this section, the base field $k = \mathbb{C}$. We give an analytic criterion for determining which representations of a quiver Q are direct sums of θ-stable representations, i.e., by Proposition 3.2(i), correspond to closed

---

* This conjecture has been proved by the author and C. Walter in "On Chow Rings of Fine Moduli Spaces of Modules".

$GL(\alpha)$ orbits in $\mathscr{R}_\theta^{ss}(Q, \alpha)$. This enables us to give an alternative construction of the moduli space $\mathscr{M}_A(\alpha, \theta)$ as a 'symplectic quotient'.

We present the initial result in the more general context of Section 2 and as such it generalises a celebrated result of Kempf & Ness [11], which is the special case $\chi = 0$. Thus, let $\mathscr{R}$, G and χ be as in Section 2. Let K be a maximal compact subgroup of G (since G is reductive, it is isomorphic to the complexification of K), let $\mathfrak{k}$ be the Lie algebra of K and let $(\,,)$ be a Hermitian inner product on $\mathscr{R}$, preserved by K. Let $\mu\colon \mathscr{R} \to (i\mathfrak{k})^*$, where $*$ denotes the real linear dual, be the 'moment map' for the action of K, given by $\mu_x(A) = (Ax, x)$ for $A \in i\mathfrak{k}$, and let $\mathrm{d}\chi$ be the restriction to $i\mathfrak{k}$ of the derivative of χ at the identity in G (which takes real values as required).

THEOREM 6.1. *The set $\mu^{-1}(\mathrm{d}\chi)$ meets each G-orbit, which is closed in $\mathscr{R}_\chi^{ss}$, in precisely one K-orbit and meets no other G-orbit.*

COROLLARY 6.2. *The natural map $\mu^{-1}(\mathrm{d}\chi)/K \to \mathscr{R}/\!/(G, \chi)$ is a bijection.*

Proof of Theorem. We introduce the function $N\colon \mathscr{R} \times \mathbb{C} \to [0, \infty)$, defined by

$$N(x, z) = |z|\, e^{\frac{1}{2}\|x\|^2}$$

where $\|x\|$ is the norm coming from the chosen inner product. (The function N can be thought of as a norm on the line bundle L^{-1} and, as such, it is the 'Kähler potential' for the metric on $\mathscr{R}$.)

Let $\hat{x} = (x, z)$, for some $z \neq 0$, and define, for $A \in i\mathfrak{k}$,

$$m_x(A) = \frac{\mathrm{d}}{\mathrm{d}t}\bigg|_{t=0} \log N(e^{tA} \cdot \hat{x}) = (Ax, x) - \mathrm{d}\chi(A)$$

$$m_x^{(2)}(A) = \frac{\mathrm{d}^2}{\mathrm{d}t^2}\bigg|_{t=0} \log N(e^{tA} \cdot \hat{x}) = 2\, \|Ax\|^2$$

Clearly, $m_x = 0$ if and only if N restricted to $G \cdot \hat{x}$ has a critical point at $\hat{x}$. The action of K on $\mathscr{R} \times \mathbb{C}$ preserves N, so N induces a function on $G \cdot \hat{x}/K$. The formula for $m_x^{(2)}$ shows that this function is strictly convex (except in directions along the fibres $\{x\} \times \mathbb{C}$, where it is linear). Hence, any critical point must be a minimum and there can be at most one critical point. Thus, to prove the theorem, it suffices to show that N is minimised on an orbit $G \cdot \hat{x}$ if and only if that orbit is closed, which, by Remark 2.3, is equivalent to the orbit $G \cdot x$ being closed in $\mathscr{R}_\chi^{ss}$. This we do in the next two lemmas.

LEMMA 6.3. *Let V be any closed subvariety of $\mathscr{R} \times \mathbb{C}$ disjoint from $\mathscr{R} \times \{0\}$. Then the function N restricted to V is proper (and thus achieves its infimum).*

Proof. Since V and $\mathscr{R} \times \{0\}$ are closed and disjoint, there is a

polynomial which takes the value 1 on V and vanishes on $\mathcal{R} \times \{0\}$, i.e. V is contained in a hypersurface $z P_1(x) + \cdots + z^n P_n(x) = 1$, where $P_i \in k[\mathcal{R}]$. Now, N is proper if $N^{-1}([0, B])$ is compact for all B, i.e. when $N(x, z) \le B$ and $(x, z) \in V$, then $|z| \le R_1$ and $\|x\| \le R_2$, for some R_1, R_2 depending only on B and V. We can clearly take $R_1 = B$, so it remains to bound $\|x\|$. But, if $|z| < Be^{-\frac{1}{2}\|x\|^2}$, then

$$|z P_1(x) + \cdots + z^n P_n(x)| \le B|P_1(x)| e^{-\frac{1}{2}\|x\|^2} + \cdots + B^n |P_n(x)| e^{-n/2\|x\|^2}.$$

We can certainly choose R_2, depending only on B and $\{P_i\}$, so that, if $\|x\| > R_2$, then $|P_i(x)| < \frac{1}{n} B^{-i} e^{i/2\|x\|^2}$ for $i = 1, \ldots, n$. But we would then have $|z P_1(x) + \cdots + z^n P_n(x)| < 1$ and so $(x, z) \notin V$.

LEMMA 6.4. *Let O be a G-orbit in $\mathcal{R} \times \mathbb{C}$ disjoint from $\mathcal{R} \times \{0\}$. If the restriction of N to O achieves its infimum, then O is closed.*

Proof. Suppose that the infimum is achieved at a point (x, z). If the orbit is not closed, then we can find a one-parameter subgroup λ such that $\lim_{t \to 0} \lambda(t) \cdot (x, z)$ exists but is not in O, i.e. if we make the decomposition $x = \sum_{n \in \mathbb{Z}} x_n$ so that $\lambda(t) \cdot x = \sum_{n \in \mathbb{Z}} t^n x_n$, then $x_n = 0$ for all $n < 0$, $x_n \ne 0$ for some $n > 0$ and $\langle \chi, \lambda \rangle \le 0$. By conjugating λ if necessary, we may also assume that this decomposition of x is orthogonal. But then

$$N(\lambda(t) \cdot (x, z)) = |z| e^{\frac{1}{2}\|x_0\|^2} |t|^{-\langle \chi, \lambda \rangle} e^{\frac{1}{2}\sum_{n>0} |t|^n \|x_n\|^2}$$

and this will clearly decrease as $t \to 0$, contradicting the assumption that the infimum of N was achieved at (x, z). Thus the orbit O must have been closed.

This completes the proof of Theorem 6.1.

We now apply Theorem 6.1 when $\mathcal{R} = \mathcal{R}(Q, \alpha)$, $G = GL(\alpha)$ and $\chi = \chi_\theta$. In order to define an inner product on $\mathcal{R}(Q, \alpha)$, we choose one on each vector space W_v and then use the standard operator inner product induced on each summand $\mathrm{Hom}(W_{ta}, W_{ha})$, i.e. $(\phi_a, \psi_a) = \mathrm{tr}(\phi_a \psi_a^*)$, where $*$ now denotes the adjoint map. The maximal compact subgroup K of $GL(\alpha)$ which preserves this inner product is $U(\alpha) = \prod_v U(W_v)$, so it is $\oplus_v \mathrm{Herm}(W_v)$, where $\mathrm{Herm}(W_v)$ is the space of hermitian endomorphisms of W_v. The action of $A \in i\mathfrak{k}$ on $\phi \in \mathcal{R}(Q, \alpha)$ is $(A\phi)_a = A_{ta}\phi_a - \phi_a A_{ha}$ and hence the moment map is given by

$$(A\phi, \phi) = \sum_{a \in Q_1} \mathrm{tr}(A_{ta}\phi_a\phi_a^* - \phi_a A_{ha}\phi_a^*)$$

$$= \sum_{v \in Q_0} \mathrm{tr}\left(A_v\left(\sum_{ha=v} \phi_a\phi_a^* - \sum_{ta=v} \phi_a^*\phi_a\right)\right)$$

On the other hand

$$d\chi_\theta(A) = \sum_{v \in Q_0} \theta_v \, \mathrm{tr} \, (A_v)$$

$$= \sum_{v \in Q_0} \mathrm{tr} \, (A_v \theta_v 1)$$

Thus, the equation $\mu(\phi) = d\chi_\theta$ becomes

$$\sum_{ha=v} \phi_a \phi_a^* - \sum_{ta=v} \phi_a^* \phi_a = \theta_v 1 \qquad \text{for all } v \in Q_0 \qquad (M)$$

Combining Theorem 6.1 and Proposition 3.2(i), we obtain.

PROPOSITION 6.5. *A representation $\phi \in \mathcal{R}(Q, \alpha)$ is a direct sum of θ-stable representations if and only if it satisfies (M) for some choice of inner products on W_v. This choice is unique up to an automorphism of the representation, i.e. up to a scale factor in each θ-stable summand.*

Corollary 6.2 translates into the fact that the moduli space $\mathcal{M}_A(\alpha, \theta)$ can be constructed by taking the solution space of (M) inside the variety $V_A(\alpha) \subseteq \mathcal{R}(Q, \alpha)$ and dividing by $U(\alpha)$. Notice that, by taking the trace of each equation in (M) and summing over all $v \in Q_0$, we recover the condition $\sum_v \theta_v \alpha_v = 0$.

REFERENCES

1. M. Artin, 'On Azumaya algebras and finite dimensional representations of rings', *J. Algebra* 11 (1969), 532–563.
2. D. J. Benson, *Representations and Cohomology I*, Camb. Stud, in Adv. Math., vol. 30, C.U.P., 1991.
3. C. Bessenrodt & L. Le Bruyn, 'Stable rationality of certain PGL$_n$-quotients', *Invent. Math* 104 (1991), 179–199.
4. A. I. Bondal, *Helices, representations of quivers and Koszul algebras*, Helices and Vector Bundles (A. N. Rudakov, ed.), L.M.S. Lecture Notes 148, C.U.P., 1990, pp. 75–95.

5. M. Brion and C. Procesi, 'Action d'un tore dans une variété projective', Birkhäuser, *Prog. in Math.* 92 (1990), 509–539.
6. J.-M. Drezet and J. Le Potier, 'Fibrés stables et fibrés exceptionnels sur $\mathbf{P}_2$', *Ann. Sci. Ec. Norm. Sup.* 18 (1985), 193–244.
7. K. Hulek, 'On the classification of stable rank-r vector bundles over the projective plane', Birkhäuser, *Prog. in Math.* 7 (1980), 113–144.
8. V. Guillemin & S. Sternberg, 'Birational equivalence in symplectic geometry', *Invent. Math.* 106 (1989), 485–522.
9. V. Kac, 'Infinite root systems, representations of graphs and invariant theory', *Invent. Math.* 56 (1980), 57–92.
10. G. Kempf, 'Instability in invariant theory', *Ann. Math.* 108 (1978), 299–316.
11. G. Kempf & L. Ness, *On the lengths of vectors in representation spaces*, Springer L.N.M. 732 (1982), 233–243.
12. A. D. King, 'Tilting bundles on some rational surfaces', preprint.
13. F. C. Kirwan, *Cohomology of Quotients in Symplectic and Algebraic Geometry*, Princeton U.P., 1984.
14. H. Kraft, *Geometric methods in representation theory*, Representations of Algebras (M. Auslander and E. Luis, eds.), L.N.M. 944, 1980, pp. 180–258.
15. D. Mumford, *Projective invariants of projective structures and applications*, Proc. I.C.M., Stockholm (1962), 526–530.
16. D. Mumford & J. Fogarty, *Geometric Invariant Theory*, 2nd ed., Springer-Verlag, 1982.
17. P. E. Newstead, *Introduction to Moduli Problems and Orbit Spaces*, T.I.F.R. Lecture Notes, Springer-Verlag, 1978.
18. C. Procesi, 'Finite dimensional representations of algebras', *Israel J. Math.* 19 (1974), 169–182.
19. C. M. Ringel, *Tame Algebras and Integral Quadratic Forms*, L.N.M. 1099, Springer-Verlag, 1984.
20. A. Schofield, 'General representations of quivers', *Proc. Lond. Math. Soc.* 65 (1992), 46–64.
21. C. S. Seshadri, 'Space of unitary vector bundles on a compact Riemann surface', Ann. Math. 85 (1967), 303–336.
22. F. Takemoto, 'Stable vector bundles on an algebraic surface', Nagoya Math. J. 47 (1972), 29–48.

Department of Pure Maths,
University of Liverpool, P.O. Box 147,
Liverpool, L69 3BX, U.K.