

COMPLEX ALGEBRAIC COMPACTIFICATIONS OF THE MODULI SPACE OF HERMITIAN-YANG-MILLS CONNECTIONS ON A PROJECTIVE MANIFOLD

DANIEL GREB, BENJAMIN SIBLEY, MATEI TOMA, AND RICHARD WENTWORTH

ABSTRACT. In this paper we study the relationship between three compactifications of the moduli space of gauge equivalence classes of Hermitian-Yang-Mills connections on a fixed Hermitian vector bundle over a projective algebraic manifold of arbitrary dimension. Via the Donaldson-Uhlenbeck-Yau theorem, this space is analytically isomorphic to the moduli space of stable holomorphic vector bundles, and as such it admits an algebraic compactification by Gieseker-Maruyama semistable torsion free sheaves. A recent construction due to the first and third authors gives another compactification as a moduli space of slope semistable sheaves. In the present article, following fundamental work of Tian generalising the analysis of Uhlenbeck and Donaldson in complex dimension two, we define a gauge theoretic compactification by adding certain gauge equivalence classes of ideal connections at the boundary. Extending work of Jun Li in the case of bundles on algebraic surfaces, we exhibit comparison maps from the sheaf theoretic compactifications and prove their continuity. The continuity, together with a delicate analysis of the fibres of the map from the moduli space of slope semistable sheaves allows us to endow the gauge theoretic compactification with the structure of a complex analytic space.

1. INTRODUCTION

The study of the Hermitian-Yang-Mills (HYM) equations on a Kähler manifold joins two rather different areas of research in differential geometry. On the one hand, the Donaldson-Uhlenbeck-Yau theorem [15, 17, 62], which relates irreducible solutions of these equations to Mumford-Takemoto stable holomorphic vector bundles, is an important example of a more general relationship between nonlinear geometric PDEs in complex geometry and algebraic geometric notions of stability coming from Geometric Invariant Theory (GIT). On the

Date: 1st September 2020.

2010 Mathematics Subject Classification. 14D20, 14J60, 32G13, 53C07.

Key words and phrases. Hermitian-Yang-Mills connections, stability, moduli of coherent sheaves, Kobayashi-Hitchin correspondence, Donaldson-Uhlenbeck compactification.

other hand, the HYM equations are an archetypal instance of equations appearing in mathematical gauge theory. Indeed, in the case of Kähler surfaces they correspond to the anti-self-duality equations that were so profitably exploited by Donaldson and others in the 1980s and 1990s. The purpose of this paper is to develop further one aspect of the link between these two points of view in the case of higher dimensional projective algebraic manifolds: namely, the precise relationship between degenerations of locally free sheaves and singular limits of HYM connections.

Concretely, we are concerned with the moduli space $M_{\text{HYM}}^*(E, h, a_J)$ of gauge equivalence classes of smooth irreducible h -unitary connections that are solutions to the HYM equations on a fixed hermitian vector bundle (E, h) of rank r on a projective algebraic manifold $X \subset \mathbb{P}^N$ of dimension n that induce a fixed connection a_J on $J = \det E$. If $M_{\text{HYM}}^*(E, h, a_J)$ is nonempty it will generally fail to be compact. Both in differential and algebraic geometry, when one studies such noncompact moduli spaces it is often useful to construct a compactification by adding singular objects of one kind or another at the boundary.

1.1. Algebro-geometric modular compactifications. The first point of view mentioned at the beginning of the introduction suggests adjoining "singular" holomorphic vector bundles, or torsion free coherent sheaves, in order to compactify the moduli space $M^s(E, \mathcal{J})$ of slope stable holomorphic structures on E with fixed determinant \mathcal{J} , corresponding to a_J . In fact, it is possible to implement this idea in two different ways, corresponding to two different natural restrictions imposed on the sheaves allowed. The most well-studied solution, going back to Gieseker and Maruyama [26, 41], uses Gieseker-Maruyama (GM) semistable sheaves \mathcal{E} with the same topological invariants as E . This is a natural choice because all such sheaves are quotients $\mathcal{H} \rightarrow \mathcal{E} \rightarrow 0$ of a certain fixed vector bundle $\mathcal{H} \rightarrow X$ (see (2.14)), and GM-semistability turns out to be the same as GIT stability for a particular linearised line bundle on the Quot scheme $\text{Quot}(\mathcal{H}, c(E))$ parametrising such quotients. A projective scheme $M^{\text{Gss}}(E, \mathcal{J})$, containing $M^s(E, \mathcal{J})$ as a Zariski open subset may then be constructed as a GIT quotient, and taking the closure of $M^s(E, \mathcal{J})$ provides a compactification $\overline{M}^{\text{GM}}(E, \mathcal{J})$.

It is also reasonable to ask if there is a moduli space which allows degeneration to torsion free sheaves which are merely slope semistable. This has been analysed in the case of surfaces in [37, 38] (see also [33]), and very recently in higher dimensions by the first and third authors in [30]. All relevant slope semistable sheaves are again quotients of \mathcal{H} , and there is furthermore a certain natural linearised line bundle over the Quot scheme whose restriction to the

open subset $R^{\mu_{\text{ss}}}$ of slope semistable sheaves is equivariantly semiample for the case where X is a surface. In dimensions greater than two, semiamplicity of the analogous line bundle \mathcal{L}_{n-1} can be established on the weak normalisation $(R^{\mu_{\text{ss}}})^{wn}$.¹ We refer the reader to the introduction of [30] for an explanation of why the passage to the weak normalisation is necessary for technical, and very likely also philosophical, reasons. Given this, one then defines the moduli space $M^{\mu_{\text{ss}}}(E, \mathcal{J})$ of slope semistable sheaves with fixed determinant \mathcal{J} , in imitation of the GIT construction, to be the projective spectrum of the ring of invariant sections of powers of \mathcal{L}_{n-1} . An application of Langton's Theorem implies that this defines a complete, hence projective, variety. Considering the separation properties of sections, one then shows that $M^s(E, \mathcal{J})^{wn}$ embeds as a Zariski open subset of $M^{\mu_{\text{ss}}}(E, \mathcal{J})$. By taking its closure inside $M^{\mu_{\text{ss}}}(E, \mathcal{J})$ we thus obtain a second compactification $\overline{M}^{\mu}(E, \mathcal{J})$.

Points of $\overline{M}^{\text{GM}}(E, \mathcal{J})$ and $\overline{M}^{\mu}(E, \mathcal{J})$ may be represented by torsion free GM- or slope semistable sheaves, respectively, but it is important to understand their isomorphism classes; that is, when do two sheaves represent the same point in the moduli space? For $\overline{M}^{\text{GM}}(E, \mathcal{J})$, this is given by the notion of s -equivalence. For $\overline{M}^{\mu}(E, \mathcal{J})$, in the case of algebraic surfaces the correct identifications were found by Jun Li in [38]. If $\dim X \geq 3$, the issue is more subtle and was analysed in [30]. The characterisation of points in these spaces, together with the fact that GM-semistability implies slope semistability, leads to the definition of a birational morphism $\Xi : \overline{M}^{\text{GM}}(E, \mathcal{J})^{wn} \rightarrow \overline{M}^{\mu}(E, \mathcal{J})$ that extends the identity on $M^s(E, \mathcal{J})^{wn}$, but usually collapses points at the boundary. A more detailed summary of the theory described in the preceding paragraphs is given in Section 2.8 below.

1.2. Gauge theoretic compactification in higher dimensions. From the gauge theoretic point view, a natural candidate for a compactification would be to add "singular" HYM connections. These may be understood from the foundational results of Uhlenbeck [60, 61, 59] which provide weak convergence along a subsequence assuming appropriate bounds on curvature (see also [47, 50]). Essential to the present paper is the work of Tian [57], who proves that a sequence of smooth HYM connections $\{A_i\}$ converges (subsequentially) to a limiting HYM connection A_{∞} smoothly away from a holomorphic subvariety $S \subset X$ of codimension at least two. Moreover, using the weak limit of the Yang-Mills energy density, one can assign a positive integer multiplicity to the

¹Recall that the *weak normalisation* of a complex space Z is a reduced complex space Z^{wn} on which the first Riemann Extension Theorem holds, together with a homeomorphic, holomorphic map $Z^{wn} \rightarrow Z$. Weak normalisation is a functor on the category of complex spaces, which by definition factors through the reduction functor.

irreducible components of S of pure codimension two. In this way the limit produces an effective holomorphic cycle \mathcal{C} . The complement in S of the support of \mathcal{C} consists of nonremovable singularities of the limiting connection A_∞ . Hence, there is a decomposition $S = S_b \cup S(A_\infty)$, where S_b is a union of pure codimension two subvarieties, and A_∞ extends as a smooth connection locally on the complement of $S(A_\infty)$ ([57, 56]; Theorems 3.5 and 3.7 below). Furthermore, by results of Bando and Siu ([3]; Theorem 3.6 below), the holomorphic bundle defined by A_∞ away from S actually extends to all of X as a slope polystable reflexive sheaf \mathcal{E}_∞ , and $S(A_\infty) = \text{sing}(\mathcal{E}_\infty)$. Note that we may have $S_b \cap S(A_\infty) \neq \emptyset$.

This suggests that ideal boundary points should be represented by triples of the form $(A, \mathcal{C}, S(A))$, or equivalently pairs $(\mathcal{E}, \mathcal{C})$, where A is a finite energy HYM connection defined on the complement of $|\mathcal{C}| \cup S(A)$, \mathcal{C} is an effective codimension two holomorphic cycle, $S(A)$ is a holomorphic subvariety of codimension at least 3 which is exactly the set of nonremovable singularities of A , and \mathcal{E} is a slope polystable reflexive sheaf with $S(A) = \text{sing}(\mathcal{E})$, subject to the cohomological condition $\text{ch}_2(E) = \text{ch}_2(\mathcal{E}) + [\mathcal{C}]$.

This description of the "ideal connections" appearing at infinity should be compared to the definition of the Donaldson-Uhlenbeck compactification (see [16], [20], and the summary provided in [38, Sect. 4]) used in the theory of instantons on smooth four dimensional real manifolds. In this latter situation, S is a finite set of points, and the connection A extends smoothly to the entire manifold ($S(A) = \emptyset$), albeit on a bundle with modified second Chern class. Originally, the theory is developed for either $\text{SU}(2)$ or $\text{SO}(3)$ bundles, and the second Chern number is then essentially the Yang-Mills energy; in particular, it is positive. The cohomological relation stated above then gives a uniform bound on the total length of the cycles (the sum of the number of points with multiplicity) and this is enough to obtain convergence of the singular sets to a limiting cycle. One may then apply Uhlenbeck compactness on the complement of (the support of) this limiting cycle to obtain a limiting ideal connection, and the set of all ideal connections is thereby (sequentially) compact. Taking the closure inside this compact space yields a gauge theoretic compactification.

In order to obtain an analogue of the Donaldson-Uhlenbeck compactification in higher dimensions, we mimic the procedure described in the previous paragraph using the approach of [57]. The appearance of the nonremovable singularity sets $S(A)$ is a key new feature in higher dimensions. For a sequence $(A_i, \mathcal{C}_i, S(A_i))$, the Bogomolov inequality and the relations $\text{ch}_2(E) = \text{ch}_2(A_i) + [\mathcal{C}_i]$, imply a uniform bound on the degrees of the \mathcal{C}_i , and

therefore also subsequential convergence of the \mathcal{C}_i as currents to some codimension two cycle \mathcal{C}'_∞ . Given a general sequence of ideal connections, however, it would seem there is no a priori control on the singular sets $S(A_i)$ beyond the fact that they converge in the Hausdorff topology to some compact set, the Hausdorff dimension of which is in principle only bounded in real codimension four (cf. [58, Prop. 3.3]). Although Uhlenbeck compactness applies on the complement of this closed set to obtain a limiting connection A_∞ , this will not result in an ideal connection in our sense unless we can pin down a more precise structure for the singular set. One problem is the extension of the codimension two subvarieties of the bubbling locus to all of X . A second crucial issue is the behaviour of the sets $S(A_i)$. In particular, since the lower dimensional strata $S(A_\infty)$ associated to an ideal connection should be precisely the codimension at least three subvariety given by the singularities of a reflexive sheaf, we wish first of all to rule out the possibility that the $S(A_i)$ accumulate along a set of strictly larger Hausdorff dimension.

Subsequential convergence of the $S(A_i)$ in the cycle space is implied by a uniform bound on the degrees (or volumes) of the holomorphic subvarieties given by the irreducible components of the $S(A_i)$ in each dimension. The problem that appears at this point is that it is by no means obvious that the set $\widehat{M}_{\text{HYM}}(E, h, a_J)$ of ideal connections, or indeed even the set of limits of smooth, irreducible, HYM connections satisfies this property. This is where we are forced to use the projectivity assumption on X . Thinking of points of $\widehat{M}_{\text{HYM}}(E, h, a_J)$ as a set of pairs $(\mathcal{E}, \mathcal{C})$ the above property for $S(A) = \text{sing}(\mathcal{E})$ would follow from the statement that the family of all polystable reflexive sheaves \mathcal{E} arising as the sheaf component of an ideal connection is bounded. Here we may rely on fundamental results of [31] controlling the Hilbert polynomials of such a family of sheaves; namely, only finitely many Hilbert polynomials appear as we range over a bounded family. This fact, together with the sheaf theoretic description of the singular sets, is enough to give a uniform degree bound for sequences in $\widehat{M}_{\text{HYM}}(E, h, a_J)$. Therefore the question is reduced to whether or not the sheaves \mathcal{E} form a bounded family. This is answered affirmatively, using [42, Main Theorem] (see Lemma 3.16).

Applying the compactness results of Uhlenbeck and Tian on the complement of $|\mathcal{C}'_\infty| \cup S'_\infty$, we obtain sequential compactness for the space $\widehat{M}_{\text{HYM}}(E, h, a_J)$ (see Theorem 3.17). A crucial aspect of this result, also essential for the algebraic arguments of Section 5, is the fact that the limiting set S'_∞ is actually contained in $|\mathcal{C}_\infty| \cup S(A_\infty)$, where \mathcal{C}_∞ is a cycle obtained from \mathcal{C}'_∞ and the cycle produced by bubbling of the connections. Interestingly, the proof of this

fact relies on $\bar{\partial}$ -methods and relatively recent estimates of Chakrabarti-Shaw [9].

With these results we may proceed to define the compactification $\overline{M}_{\text{HYM}}(E, h, a_J)$ to be the set of Uhlenbeck limits in $\widehat{M}_{\text{HYM}}(E, h, a_J)$ of smooth, irreducible HYM connections. A diagonalisation argument shows that this space is in fact closed, or in other words any limit of a sequence of ideal connections in $\overline{M}_{\text{HYM}}(E, h, a_J)$ again arises as an Uhlenbeck limit of connections in $M_{\text{HYM}}^*(E, h, a_J)$. From the definition of convergence emerges a natural topology on $\widehat{M}_{\text{HYM}}(E, h, a_J)$, and hence also for $\overline{M}_{\text{HYM}}(E, h, a_J)$, resulting in the following theorem.

Theorem 1.1 (GAUGE THEORETIC COMPACTIFICATION). *$\overline{M}_{\text{HYM}}(E, h, a_J)$ is a sequentially compact Hausdorff space containing $M_{\text{HYM}}^*(E, h, a_J)$ as a dense open subspace.*

We should point out here that after the proof of Theorem 1.2 below we will conclude that $\overline{M}_{\text{HYM}}(E, h, a_J)$ is actually compact, and not just sequentially so.

1.3. A complex structure on the HYM compactification. It is natural to ask how much structure the topological space $\overline{M}_{\text{HYM}}(E, h, a_J)$ inherits from the projective manifold X . As mentioned in the first paragraph, the Donaldson-Uhlenbeck-Yau Theorem gives a bijection

$$\Phi : M^s(E, \mathcal{J}) \longrightarrow M_{\text{HYM}}^*(E, h, a_J) .$$

In fact, much more is true: the space $M^s(E, \mathcal{J})$ has the structure of a possibly nonreduced complex analytic space, obtained for example by analytification from the scheme structure of the Gieseker-Maruyama moduli space,² while $M_{\text{HYM}}^*(E, h, a_J)$ has a natural real-analytic space structure. With respect to these, the map Φ is real-analytic, and may hence be used to transfer the complex structure from $M^s(E, \mathcal{J})$ to $M_{\text{HYM}}^*(E, h, a_J)$ (see [24], [39], [44], and especially [40] for this point of view). The question then arises as to what relationship the algebraic compactifications $\overline{M}^{\text{GM}}(E, \mathcal{J})$ and $\overline{M}^\mu(E, \mathcal{J})$ might have to $\overline{M}_{\text{HYM}}(E, h, a_J)$, and whether or not the gauge theoretic compactification $\overline{M}_{\text{HYM}}(E, h, a_J)$ also has the structure of a complex or even (projective) algebraic space.

In the case of projective surfaces, these two questions have been answered by Jun Li in [38] (see also [19]). In this case, a complex space closely connected

²This complex structure coincides with the one induced by the moduli space of simple holomorphic bundles and also with the one given by the moduli space of simple holomorphic structures ([34, Chapter VII] and [44]).

with $\overline{M}^\mu(E, \mathcal{J})$ turns out to be homeomorphic to the Uhlenbeck-Donaldson compactification (see Remark 4.16 for the details), and Li realizes this homeomorphism via an algebraic map $\overline{M}^{\text{GM}}(E, \mathcal{J}) \rightarrow \overline{M}^\mu(E, \mathcal{J})$. We also note that in this context Morgan [45] constructed a continuous map from $M^{\text{Gss}}(E, \mathcal{J})$ to $\overline{M}_{\text{HYM}}(E, h, a_J)$ (see also [7] for a generalisation to non-Kählerian surfaces).

This brings us to the second main result of the paper, which provides a generalisation of Jun Li's theorem to higher dimensions.

Theorem 1.2 (COMPLEX STRUCTURE). *The space $\overline{M}_{\text{HYM}}(E, h, a_J)$ can be endowed with the structure of a reduced complex analytic space such that the natural map $M_{\text{HYM}}^*(E, h, a_J)^{\text{wn}} \hookrightarrow \overline{M}_{\text{HYM}}(E, h, a_J)$ is holomorphic and embeds $M_{\text{HYM}}^*(E, h, a_J)^{\text{wn}}$ as a Zariski open and dense subset. There are moreover natural surjective holomorphic maps*

$$\overline{M}^{\text{GM}}(E, \mathcal{J})^{\text{wn}} \longrightarrow \overline{M}^\mu(E, \mathcal{J}) \xrightarrow{\overline{\Phi}} \overline{M}_{\text{HYM}}(E, h, a_J) ,$$

where the first map is birational, and the second map is finite and extends the map

$$\Phi^{\text{wn}} : M^s(E, \mathcal{J})^{\text{wn}} \xrightarrow{\cong} M_{\text{HYM}}^*(E, h, a_J)^{\text{wn}}$$

to the respective compactifications.

Once the complex structure on $\overline{M}_{\text{HYM}}(E, h, a_J)$ has been constructed, the fact that $\overline{M}^\mu(E, \mathcal{J})$ is projective, and hence algebraic, implies by standard arguments that $\overline{M}_{\text{HYM}}(E, h, a_J)$ is actually likewise algebraic (see the end of Section 5).

Corollary 1.3 (ALGEBRAICITY OF COMPLEX STRUCTURE). *The complex space $\overline{M}_{\text{HYM}}(E, h, a_J)$ constructed in Theorem 1.2 is the analytification of an algebraic space, which we also denote by $\overline{M}_{\text{HYM}}(E, h, a_J)$, such that*

$$\overline{M}^\mu(E, \mathcal{J}) \xrightarrow{\overline{\Phi}} \overline{M}_{\text{HYM}}(E, h, a_J)$$

is a morphism of algebraic spaces which embeds $M^s(E, \mathcal{J})^{\text{wn}} \subset \overline{M}^\mu(E, \mathcal{J})$ as a Zariski open subset into $\overline{M}_{\text{HYM}}(E, h, a_J)$.

1.4. The sheaf-theory-to-gauge-theory comparison map. Let us give a rough idea of the proof of Theorem 1.2, which can be found in Section 5. Broadly speaking, we follow the strategy of Li in the case of surfaces. However, as we shall see, the relationship of $\overline{M}_{\text{HYM}}(E, h, a_J)$ to $\overline{M}^\mu(E, \mathcal{J})$ is less straightforward if $\dim X \geq 3$. We begin with a comparison map

$$\overline{\Phi} : \overline{M}^\mu(E, \mathcal{J}) \longrightarrow \widehat{M}_{\text{HYM}}(E, h, a_J) ,$$

defined by sending $[\mathcal{E}] \mapsto ((\text{Gr } \mathcal{E})^{\vee\vee}, \mathcal{C}_{\mathcal{E}})$, where \mathcal{E} is a slope semistable sheaf, $\text{Gr } \mathcal{E}$ is a torsion free sheaf associated to a Jordan-Hölder filtration of \mathcal{E} , $(\text{Gr } \mathcal{E})^{\vee\vee}$ is its sheaf theoretic double dual, and $\mathcal{C}_{\mathcal{E}}$ is the support cycle of the torsion sheaf $(\text{Gr } \mathcal{E})^{\vee\vee} / \text{Gr } \mathcal{E}$. That this map is well defined is a consequence of the geometry of $M^{\mu\text{ss}}(E, \mathcal{J})$ (see Section 2.8). In fact, as a direct result of Propositions 4.5 and 4.7 in Section 4, we will see that the image of this map actually lies in $\overline{M}_{\text{HYM}}(E, h, a_J)$ (see Lemma 4.12).

The first point is that $\overline{\Phi}$ is continuous. We prove this by showing that it respects limits in the respective spaces. The existence of subsequential limits in $\overline{M}_{\text{HYM}}(E, h, a_J)$ is guaranteed by Theorem 1.1. We first prove the continuity for degenerations of the form $\mathcal{E}_i \rightarrow \mathcal{E}_{\infty}$, where the \mathcal{E}_i are slope stable and locally free and \mathcal{E}_{∞} is slope semistable. The result follows from continuity on sequences of this type, together with a diagonalisation argument. Since a point in $\overline{M}_{\text{HYM}}(E, h, a_J)$ is a pair $(\mathcal{E}, \mathcal{C})$, the proof reduces to two natural continuity-type statements; one for the limiting reflexive sheaf, and one for the holomorphic cycle.

In the sheaf component, the key point is that in the projective setting any Uhlenbeck limit \mathcal{E}_{∞} of a sequence of smooth, irreducible, HYM connections can be identified with the double dual of a polystable quotient $\mathcal{H} \rightarrow \widehat{\mathcal{E}}_{\infty} \rightarrow 0$ in $\text{Quot}(\mathcal{H}, c(E))$ (see Section 4.2). The continuity essentially follows by applying this result to the HYM connections in the complex gauge orbit of the \mathcal{E}_i . Here we follow the argument of [38], with some important modifications. The idea, which is reminiscent of methods used in [15, 45, 13, 53, 7], is to find a subsequential limit of an L^2 -orthonormal basis of sections of \mathcal{E}_i twisted by a high power of the polarisation. The associated maps $\mathbf{e}_i^{(k)} : \mathcal{H}(k) \rightarrow \mathcal{E}_i(k)$ give rise to a subsheaf $\widetilde{\mathcal{E}}_{\infty} \subset \mathcal{E}_{\infty}$ of full rank. On the other hand, one can take a limit of the sequence in $\text{Quot}(\mathcal{H}, c(E))$, and we show that $\widetilde{\mathcal{E}}_{\infty}$ coincides with this limit.

For continuity in the cycle component, our proof is rather different to that of [38], and we adapt instead an argument of [54]. The main part of the proof is to show the equality of the multiplicity of an irreducible component of the support cycle of the quotient $(\text{Gr } \mathcal{E}_{\infty})^{\vee\vee} / \text{Gr } \mathcal{E}_{\infty}$ with the multiplicity defined by the limit of the Yang-Mills energy densities for A_i (see Proposition 4.9). We accomplish this by means of a slicing argument from [54] together with a slight modification of the Bott-Chern formula proven in that reference. The delicate point here is that in order to isolate the irreducible component in question, we must work on a modification of X by a sequence of blowups that achieves certain useful properties. The argument in [54] is applicable to the support cycle of the quotient of the associated graded of a holomorphic vector bundle

and its double dual (in that case this double dual arises as an Uhlenbeck limit). The main trick is then to notice that by virtue of the preceding paragraph, the Uhlenbeck limit \mathcal{E}_∞ is the double dual of a quotient of \mathcal{H} , and so the cycle in question is the one associated to a kind of graded object to which the singular Bott-Chern argument applies. With this in hand, the proof proceeds in a manner similar to that of [54] (see Section 4.3).

1.5. The induced equivalence relation. The geometric description of the points of the space $M^{\mu_{\text{SS}}}(E)$ given in [30] (see also Section 2.8) implies that the map $\bar{\Phi}$ is in fact finite-to-one. A major difference to the case of surfaces is that in general the fibres may potentially consist of more than one point. We define a finite equivalence relation on $\bar{M}^\mu(E, \mathcal{J})$ by considering two points to be equivalent precisely when they lie in the same fibre. Once the continuity of $\bar{\Phi}$ has been established, the quotient of $\bar{M}^\mu(E, \mathcal{J})$ by this equivalence relation is then homeomorphic to $\bar{M}_{\text{HYM}}(E, h, a_J)$. In order to endow the quotient with a complex structure, we need to know that our equivalence relation has sufficiently nice properties so that the complex structure on $\bar{M}^\mu(E, \mathcal{J})$ descends to the quotient.

Fortunately, there is a criterion due to H. Cartan for when the quotient of a complex space is itself complex; namely, that points can be separated by local holomorphic functions that collapse the relation. We show that this indeed holds for the relation on $\bar{M}^\mu(E, \mathcal{J})$ defined above by a more detailed study of sections of the line bundle \mathcal{L}_{n-1} used to define $M^{\mu_{\text{SS}}}(E)$. The sheaf theoretic moduli space was defined using all equivariant sections of \mathcal{L}_{n-1} , which give a globally defined map to projective space, and whose image is by definition $M^{\mu_{\text{SS}}}(E, \mathcal{J})$. It is possible that we thereby separate too many sheaves and so obtain a space that is larger than $\bar{M}_{\text{HYM}}(E, h, a_J)$. We rectify this by considering instead only those sections which are "lifted" (in a certain precise sense) from certain well chosen complete intersection curves inside of a fixed complete intersection surface. The associated linear systems W give rise to locally defined morphisms $\nu : U_S \rightarrow \mathbb{P}(W^*)$ with $U_S \subset \bar{M}^\mu(E, \mathcal{J})$ open. Taking products of these maps for finitely many choices of surfaces, we obtain maps defined on certain open neighbourhoods of a given point of $\bar{M}^\mu(E, \mathcal{J})$ that identify precisely the fibres of $\bar{\Phi}$.

The chief difficulty in implementing this strategy is verifying first of all that the morphisms so obtained are actually constant on the fibres of $\bar{\Phi}$ and secondly that by restricting to sections arising from these fixed linear subsystems, we continue to separate sheaves not lying in the same fibre of $\bar{\Phi}$. The first point is established by showing that the image of a sheaf under the map ν is determined by the image under the corresponding (global) map on the corresponding

moduli space on the surface S of the restriction of its associated ideal connection to S (see Proposition 5.25). Here we appeal to Jun Li's result mentioned above, that on a surface the algebraic moduli space is the same as the space of ideal connections. The collapsing then follows, since if two sheaves give the same ideal connection in $\overline{M}_{\text{HYM}}(E, h, a_J)$, their restrictions to the surface are still equal in the moduli space on S . The second point is verified by showing that any ideal connection is completely determined by its restriction to finitely many such complete intersection surfaces S (see Lemma 5.27).

1.6. Open questions and further directions. The proof of Theorem 1.2 outlined above leaves unanswered several important questions. First, one would like to know whether or not the strong form of Li's theorem holds in higher dimensions; that is, whether or not $\overline{M}_{\text{HYM}}(E, h, a_J)$ is a scheme or even projective. In contrast to the surface case this does not follow immediately from the method of proof, since $\overline{\Phi}$ is not bijective. There is no good general theory for deciding whether the quotient of a scheme by a finite equivalence relation is again a scheme, even when the quotient is an algebraic space (see [35] for a thorough discussion).

Second, it would be interesting to find examples where the finite-to-one property described here actually occurs. A related question is whether the gauge theoretic compactification can be realized as the completion of $M_{\text{HYM}}^*(E, h, a_J)$ for an appropriate metric. In [18], Donaldson proves this for the moduli space of anti-self-dual equations on a compact four manifold, but generalising the result to higher dimensions requires more information about the structure of a neighbourhood of an ideal HYM connection. Note that there has been recent progress in this direction (cf. [64, 46, 10, 11, 12]).

Finally, we remark that, via the use of boundedness results, projectivity of X is essential to the construction of the compactification $\overline{M}_{\text{HYM}}(E, h, a_J)$ and its complex structure that is presented here. In principle, one would like to define a gauge theoretic compactification in various other scenarios along the lines of the programme in [21], including HYM connections on nonprojective Kähler manifolds, where the tools used in this paper are not available. This seems to require at the very least some kind of analytic control on the higher codimension singular sets $S(A_i)$.

1.7. Organisation of the paper and conventions. In Section 2 we give basic background about the moduli spaces, coherent sheaves, the Quot scheme, and the compactifications using algebraic geometry. The whole of Section 3 is devoted to the proof of Theorem 1.1; namely, the construction of

$\overline{M}_{\text{HYM}}(E, h, a, J)$ and the proof of compactness. In Section 4 we prove continuity of $\overline{\Phi}$. In Section 5 we analyse the fibres of $\overline{\Phi}$ in detail and prove Theorem 1.2. For more details, see the table of contents below.

If not explicitly stated otherwise, X will always denote a smooth projective variety over the complex numbers \mathbb{C} . For the purposes of a clear notation, C^∞ bundles will be denoted with straight characters E , etc., whereas coherent sheaves will be denoted in script, e.g. $\mathcal{E} \rightarrow X$. We shall confuse the notation of a holomorphic vector bundle with its sheaf of germs of holomorphic sections. When the distinction becomes necessary, we will generally use \mathcal{F} for elements of Quot schemes, and \mathcal{E} for a holomorphic bundle associated to a $\bar{\partial}$ -operator on E giving an HYM connection, or to an Uhlenbeck limit of these, or more generally to the reflexive sheaf arising from an ideal connection (see Sections 2 and 3). We use mathscript letters, e.g. \mathcal{F} , to denote families of sheaves over a parameter space. If $\mathcal{L} \rightarrow X$ is a holomorphic line bundle, we will denote by $|\mathcal{L}|$ the associated linear system. If $\mathcal{O}_X(1)$ is the very ample line bundle on $X \subset \mathbb{P}^N$, and S is a subvariety of X , we will denote the restriction of $\mathcal{O}_X(1)$ to S by $\mathcal{O}_S(1)$. Analytic subsets are also called holomorphic subvarieties; in particular, the latter are not assumed to be irreducible. Dimensions and codimensions \dim and codim are always over the complex numbers \mathbb{C} unless otherwise indicated.

Acknowledgments. D.G.'s research has been supported by the DFG-Collaborative Research Center SFB/TRR 45 "Periods, moduli spaces and arithmetic of algebraic varieties" (subproject M08-10 "Moduli of vector bundles on higher-dimensional varieties"). B.S. and M.T. are grateful for the hospitality of the Max Planck Institut für Mathematik in Bonn and the Department of Mathematics at the University of Maryland, where a portion of this work was completed. B.S. also received support from the Simons Center for Geometry and Physics when he was a postdoc there, and from the PAI grant "DyGest" when he was a postdoc at the Université Libre de Bruxelles. R.W.'s research was supported by grants from the National Science Foundation.

Finally, the authors warmly thank Jun Li, Mei-Chi Shaw, Carlos Simpson, and Katrin Wehrheim for helpful comments and suggestions concerning this research. We are grateful to Xuemiao Chen and Song Sun for pointing out an error in the first version of this paper. We also thank the referee for detailed comments, which have helped improve the exposition.

CONTENTS

1. Introduction	1
1.1. Algebro-geometric modular compactifications	2
1.2. Gauge theoretic compactification in higher dimensions	3
1.3. A complex structure on the HYM compactification	6
1.4. The sheaf-theory-to-gauge-theory comparison map	7
1.5. The induced equivalence relation	9
1.6. Open questions and further directions	10
1.7. Organisation of the paper and conventions	10
2. HYM connections, reflexive sheaves, and their moduli	13
2.1. Moduli spaces	13
2.2. The Donaldson-Uhlenbeck-Yau theorem	16
2.3. The Hilbert polynomial and GM-semistability	16
2.4. Jordan-Hölder filtrations	17
2.5. Cycle spaces, singular sets, and support cycles	18
2.6. The Quot scheme, natural subschemes, and convergence	25
2.7. Flat families and line bundles over Quot schemes	27
2.8. Compactifications by sheaves	28
3. The gauge theoretic compactification	32
3.1. Uhlenbeck limits and admissible connections	32
3.2. Analytic multiplicities	35
3.3. Ideal HYM connections	38
3.4. Diagonalisation	45
3.5. The compactification	48
4. Comparison of analytic and algebraic moduli spaces	49
4.1. Definition of the comparison between moduli spaces	50
4.2. Identification of the limiting sheaf	52
4.3. Identification of the limiting cycle	57
4.4. Continuity of $\bar{\Phi}$	63
5. A complex structure on \bar{M}_{HYM}	64
5.1. The fibres of $\bar{\Phi}$ and the equivalence relation on \bar{M}^μ	65
5.2. Restriction theorems	66
5.3. Admissible flags and neighbourhoods in \bar{M}_{HYM}	68
5.4. Extension of sections and flat restriction	73
5.5. Lifting sections from curves I: the method of [30]	76
5.6. Lifting sections from curves II: the method of [38]	78
5.7. The 2-dimensional case	81
5.8. The higher dimensional case	82
References	92

2. HYM CONNECTIONS, REFLEXIVE SHEAVES, AND THEIR MODULI

In this preliminary section we fix basic notation and provide background material. We give precise definitions of the moduli spaces discussed in the Introduction. We collect the necessary lemmas on compact analytic cycles on (projective) Kähler manifolds. Finally, we formulate a general construction used to produce line bundles on base spaces of flat families, and we go on to define the two algebraic compactifications referred to previously.

2.1. Moduli spaces. Let $X \subset \mathbb{P}^N$ be as in the Introduction. Throughout this section and for the rest of the paper, $(E, h) \rightarrow X$ will be a C^∞ complex vector bundle of rank r with hermitian metric h . Recall that a $\bar{\partial}$ -operator $\bar{\partial}_E : \Omega^0(X, E) \rightarrow \Omega^{0,1}(X, E)$ is *integrable* if $\bar{\partial}_E \wedge \bar{\partial}_E = 0$. In this case, the Newlander-Nirenberg theorem guarantees the existence of local frames annihilated by $\bar{\partial}_E$, and $\bar{\partial}_E$ therefore defines a holomorphic structure on E . Conversely, every holomorphic structure on E defines a unique $\bar{\partial}$ -operator via the Leibniz formula: $\bar{\partial}_E(f \cdot s) = \bar{\partial}f \otimes s$, for holomorphic sections s and functions f . We sometimes specify this relationship by $\mathcal{E} = (E, \bar{\partial}_E)$.

In the presence of the metric h , a $\bar{\partial}$ -operator defines a unique *Chern connection* $A = (\bar{\partial}_E, h)$ that is compatible with $\bar{\partial}_E$ in the sense that $\bar{\partial}_A := (d_A)^{0,1} = \bar{\partial}_E$, and A is *unitary*, that is,

$$d\langle s_1, s_2 \rangle_h = \langle d_A s_1, s_2 \rangle_h + \langle s_1, d_A s_2 \rangle_h ,$$

for any local smooth sections s_1, s_2 of E . We shall say that a unitary connection A is integrable if $\bar{\partial}_A$ is an integrable $\bar{\partial}$ -operator.

Let $\det E = J$ and $\det h = h_J$ for a hermitian line bundle $J \rightarrow X$. We suppose J carries an integrable $\bar{\partial}$ -operator $\bar{\partial}_J$ with Chern connection $a_J = (\bar{\partial}_J, h_J)$, and we denote the associated holomorphic bundle by $\mathcal{J} := (J, \bar{\partial}_J)$. As stated in the Introduction, we will always work with the space $\mathcal{A}_{hol}(E, \bar{\partial}_J)$ of integrable $\bar{\partial}$ -operators on E which induce the fixed operator $\bar{\partial}_J$ on J .

For the purposes of this paper, it will be convenient to renormalise the restriction of the Fubini-Study Kähler form on \mathbb{P}^N to give a Kähler metric ω on X with $\text{vol}(X, \omega) = 2\pi$. Let $\mathcal{O}_X(1)$ be the restriction of the hyperplane line on \mathbb{P}^N to X . We will denote the underlying C^∞ line bundle of $\mathcal{O}_X(1)$ by L , its $\bar{\partial}$ -operator by $\bar{\partial}_L$, and by h_L the hermitian metric on L whose Chern connection $a_L = (\bar{\partial}_L, h_L)$ has curvature equal to $-2\pi i \lambda \cdot \omega$. Here, $c_1(\mathcal{O}_X(1)) = \lambda \cdot \omega$, for some constant $\lambda > 0$. We will mostly omit ω from the notation, but we note here that all of the constructions depend on the Kähler class in an interesting and important manner.

Let $\mathcal{A}(E, h, a_J)$ denote the infinite dimensional space of h -unitary connections on E inducing the connection a_J on J . This is an affine space modelled

on $\Omega^1(X, \mathfrak{g}_E)$, where \mathfrak{g}_E denotes the bundle of traceless skew-hermitian endomorphisms of E . For $A \in \mathcal{A}(E, h, a_J)$, the curvature $F_A = d_A \wedge d_A \in \Omega^2(X, \mathfrak{g}_E)$. The space of integrable connections will be denoted $\mathcal{A}^{1,1}(E, h, a_J) \subset \mathcal{A}(E, h, a_J)$. Then $A \in \mathcal{A}^{1,1}(E, h, a_J)$ if and only if F_A is of type $(1, 1)$, and the Chern connection construction yields the identification $\mathcal{A}^{1,1}(E, h, a_J) \simeq \mathcal{A}_{hol}(E, \bar{\partial}_J)$.

Choose local holomorphic coordinates $\{z^i\}$ on X and write

$$\omega = \sqrt{-1} \sum g_{i\bar{j}} dz^i \wedge d\bar{z}^j .$$

Given $A \in \mathcal{A}^{1,1}(E, h)$, we have an expression $F_A = \sum_{i,j} F_{i\bar{j}} dz^i \wedge d\bar{z}^j$. If Λ is the adjoint of the Lefschetz operator defined by wedging with ω , then the hermitian endomorphism

$$\sqrt{-1} \Lambda F_A = \sum_{i,j} g^{i\bar{j}} F_{i\bar{j}} \in \Omega^0(X, \sqrt{-1} \mathfrak{g}_E) ,$$

is called the *Hermitian-Einstein tensor* (abbreviated *HE tensor*). Note that Λ is characterized by the property that

$$(\Lambda \Omega) \omega^n = n(\Omega \wedge \omega^{n-1}) . \quad (2.1)$$

for any $(1, 1)$ form Ω .

The integrable connections important for this paper are those for which the HE tensor is a constant multiple of the identity.

Definition 2.1. A *Hermitian-Yang-Mills* (or *HYM*) *connection* is an integrable, unitary connection $A \in \mathcal{A}^{1,1}(E, h, a_J)$ satisfying

$$\sqrt{-1} \Lambda F_A = \mu \cdot \mathbf{I}_E , \quad (2.2)$$

for some $\mu \in \mathbb{R}$. The subspace of HYM-connections will be denoted:

$$\mathcal{A}_{\text{HYM}}(E, h, a_J) \subset \mathcal{A}^{1,1}(E, h, a_J) .$$

The group of unitary gauge transformations is defined by

$$\mathcal{G}(E, h) = \{g \in \Omega^0(X, \text{End } E) \mid gg^* = \mathbf{I}_E , (\det g)(x) = 1 , \forall x \in X\} .$$

Then $\mathcal{G}(E, h)$ acts on $\mathcal{A}^{1,1}(E, h, a_J)$ (on the right) by conjugation:

$$d_{g(A)} = g^{-1} \circ d_A \circ g ,$$

and this induces an action on the curvature given by $F_{g(A)} = g^{-1} \circ F_A \circ g$. In particular, the subspaces

$$\mathcal{A}_{\text{HYM}}^*(E, h, a_J) \subset \mathcal{A}_{\text{HYM}}(E, h, a_J) \subset \mathcal{A}^{1,1}(E, h, a_J)$$

of (irreducible) Hermitian-Yang-Mills connections are preserved by $\mathcal{G}(E, h)$.

Definition 2.2. The *moduli spaces* of HYM connections and irreducible HYM connections on E are the sets

$$\begin{aligned} M_{\text{HYM}}(E, h, a_J) &= \mathcal{A}_{\text{HYM}}(E, h, a_J) / \mathcal{G}(E, h) , \\ M_{\text{HYM}}^*(E, h, a_J) &= \mathcal{A}_{\text{HYM}}^*(E, h, a_J) / \mathcal{G}(E, h) , \end{aligned} \text{ respectively.}$$

Note that by taking the trace of both sides of (2.2), integrating over X , and using (2.1) and the normalisation of the volume, one sees that the constant μ must be given explicitly by

$$\mu = \frac{2\pi}{\text{vol}(X)} \mu(E) = \mu(E) ,$$

where the *degree*, $\text{deg}(E)$, and *slope*, $\mu(E)$, are defined by

$$\text{deg}(E) := \int_X c_1(E) \wedge \frac{\omega^{n-1}}{(n-1)!} , \quad \mu(E) := \frac{\text{deg}(E)}{\text{rank } E} , \quad (2.3)$$

respectively. For a torsion free coherent sheaf $\mathcal{E} \rightarrow X$ of rank r , we have a line bundle $\det \mathcal{E} = (\Lambda^r \mathcal{E})^{\vee\vee}$. Setting $c_1(\mathcal{E}) = c_1(\det \mathcal{E})$, the definitions of the degree and slope thus make sense for any torsion free sheaf.

Definition 2.3. A torsion free coherent sheaf $\mathcal{E} \rightarrow X$ is ω -*slope stable* (resp. *semistable*) if any coherent subsheaf $\mathcal{S} \subset \mathcal{E}$ with $0 < \text{rank } \mathcal{S} < \text{rank } \mathcal{E}$ satisfies $\mu(\mathcal{S}) < \mu(\mathcal{E})$ (resp. \leq). A torsion free sheaf is called *polystable* if it is a direct sum of stable sheaves of the same slope.

As indicated, the notion of ω -slope stability in Definition 2.3 depends on the cohomology class $[\omega]$ of the polarization as soon as $\dim X \geq 2$. In this paper, however, the class $[\omega]$ will be fixed throughout, and we shall therefore refer to ω -(semi)slope stability simply as slope (semi)stability, or more often, as μ -(semi)stability.

The group $\mathcal{G}(E, h)$ has a natural complexification:

$$\mathcal{G}^{\mathbb{C}}(E) = \{g \in \text{End } E \mid (\det g)(x) = 1 , \forall x \in X\} .$$

Then $\mathcal{G}^{\mathbb{C}}(E)$ acts on the space of holomorphic structures $\mathcal{A}_{\text{hol}}(E, \bar{\partial}_J)$ on E by $g(\bar{\partial}_E) = g^{-1} \circ \bar{\partial}_E \circ g$. Two holomorphic structures give rise to isomorphic holomorphic vector bundles if and only if they are related by a complex gauge transformation. Therefore, the subspaces

$$\mathcal{A}_{\text{hol}}^s(E, \bar{\partial}_J) \subset \mathcal{A}_{\text{hol}}^{ps}(E, \bar{\partial}_J) \subset \mathcal{A}_{\text{hol}}^{ss}(E, \bar{\partial}_J) \subset \mathcal{A}_{\text{hol}}(E, \bar{\partial}_J)$$

of stable, polystable, and semistable holomorphic structures on E are preserved by the action of $\mathcal{G}^{\mathbb{C}}(E)$.

Definition 2.4. The moduli space $M^s(E, \mathcal{J})$ of stable holomorphic structures on E is defined as:

$$M^s(E, \mathcal{J}) = \mathcal{A}_{hol}^s(E, \bar{\partial}_J) / \mathcal{G}^{\mathbb{C}}(E) .$$

Remark 2.5. In order to lighten the notation if no confusion is likely to arise, for the rest of the paper we shall most often drop (E, h, a_J) , $\bar{\partial}_J$ and \mathcal{J} from the notation. So for example:

$$M_{\text{HYM}} := M_{\text{HYM}}(E, h, a_J) , \quad M_{\text{HYM}}^* := M_{\text{HYM}}^*(E, h, a_J) , \quad M^s := M^s(E, \mathcal{J}) .$$

2.2. The Donaldson-Uhlenbeck-Yau theorem. Viewed as an equation for the hermitian metric h on a fixed holomorphic bundle \mathcal{E} , eq. (2.2) is a quasi-linear second order elliptic PDE. It is therefore nontrivial to determine when solutions exist. Famously, the obstruction is related to stability.

Theorem 2.6 (DONALDSON [17], UHLENBECK-YAU [62]). *Fix $A \in \mathcal{A}^{1,1}$. There exists a HYM connection (resp. irreducible HYM connection) in the $\mathcal{G}^{\mathbb{C}}$ orbit of A if and only if $\mathcal{E} = (E, \bar{\partial}_A) \in \mathcal{A}_{hol}$ is polystable (resp. stable).*

This key result implies a corresponding statement at the level of moduli spaces. Namely, there exists a natural bijection

$$\Phi : M^s \xrightarrow{\cong} M_{\text{HYM}}^* . \tag{2.4}$$

Furthermore, M^s is a Hausdorff complex analytic space (possibly nonreduced), and M_{HYM}^* has the structure of a real analytic space, such that this map restricts to a real analytic isomorphism ([34, Chapter VII], [44, Prop. 4.2], [24], and [40, Thm. 4.1.1]).

2.3. The Hilbert polynomial and GM-semistability. Recall that $L \rightarrow X$ denotes the underlying complex line bundle of $\mathcal{O}_X(1)$ and that $c_1(L) = \lambda \cdot \omega$. For a complex vector bundle $E \rightarrow X$ we set

$$\tau_E(m) := \int_X \text{ch}(E \otimes L^m) \text{td}(X) . \tag{2.5}$$

Then $\tau_E(m)$ is a polynomial of degree n in m . The first two terms will be important:

$$\begin{aligned} \tau_E(m) &= m^n (2\pi\lambda^n \text{rank}(E)) + m^{n-1} \lambda^{n-1} \left(\deg E + \frac{\text{rank } E}{2} \deg TX \right) \\ &\quad + O(m^{n-2}) \end{aligned} \tag{2.6}$$

The definition (2.5) and expansion (2.6) extend to coherent sheaves \mathcal{E} , and by the Hirzebruch-Riemann-Roch Theorem and Kodaira Vanishing, for all sufficiently large natural numbers m the following holds:

$$h^0(X, \mathcal{E}(m)) = \chi_{\mathcal{E}}(m) := \chi(\mathcal{E}(m)) = \int_X \text{ch}(\mathcal{E}(m)) \text{td}(X) .$$

We refer to $\chi_{\mathcal{E}}(m)$ as the *Hilbert polynomial* of \mathcal{E} . The Hilbert polynomial is clearly topological: for any holomorphic structure $\mathcal{E} = (E, \bar{\partial}_E)$ on E we have $\chi_{\mathcal{E}}(m) = \tau_E(m)$.

An alternative notion of stability that will play an important role in the subsequent discussion is due to Gieseker and Maruyama. Let the *reduced Hilbert polynomial* be defined by

$$p_{\mathcal{E}}(m) := \frac{\chi(\mathcal{E}(m))}{\text{rank } \mathcal{E}} .$$

Then we say \mathcal{E} is *GM-stable* (resp. *GM-semistable*) if for any subsheaf $\mathcal{S} \subset \mathcal{E}$ with $0 < \text{rank } \mathcal{S} < \text{rank } \mathcal{E}$ we have $p_{\mathcal{S}}(m) < p_{\mathcal{E}}(m)$ (resp. \leq) for $m \gg 0$.

The following relationship between GM- and slope stability follows immediately from (2.6):

$$\mu\text{-stable} \implies \text{GM-stable} \implies \text{GM-semistable} \implies \mu\text{-semistable} .$$

2.4. Jordan-Hölder filtrations. If \mathcal{E} is torsion free and μ -semistable, then it has a *Seshadri filtration* (also called a *Jordan-Hölder filtration*). This is a filtration by coherent subsheaves

$$0 = \mathcal{E}_0 \subset \mathcal{E}_1 \subset \cdots \subset \mathcal{E}_{\ell-1} \subset \mathcal{E}_{\ell} = \mathcal{E} , \quad (2.7)$$

so that the successive quotients $Q_i = \mathcal{E}_i / \mathcal{E}_{i-1}$ are torsion free and slope stable. Moreover, $\mu(Q_i) = \mu(\mathcal{E})$, $i = 1, \dots, \ell$. We will write $\text{Gr } \mathcal{E} = \bigoplus_i Q_i$ for the graded object associated to the above filtration.

Such a filtration (and even its associated graded $\text{Gr } \mathcal{E}$) is not uniquely determined by \mathcal{E} and ω (see [7] for examples). On the other hand, we will see below that one can extract from it certain natural algebraic-geometric data that is unique.

By analogy with the discussion for μ -semistability, a GM-semistable sheaf \mathcal{E} has a Jordan-Hölder filtration by subsheaves whose successive quotients are torsion free GM-stable with reduced Hilbert polynomial equal to that of \mathcal{E} . Such a filtration is not unique, but the associated graded $\text{gr } \mathcal{E}$ is unique. Given two GM-semistable sheaves \mathcal{E}_1 and \mathcal{E}_2 , we say they are *s-equivalent* if $\text{gr } \mathcal{E}_1 = \text{gr } \mathcal{E}_2$.

Notice that since a GM-semistable sheaf is also μ -semistable, such a sheaf admits two types of Jordan-Hölder filtrations and graded objects, which in general will differ.

2.5. Cycle spaces, singular sets, and support cycles.

2.5.1. *The cycle space.* Write $\mathcal{C}_p(X)$ for the set of all analytic p -cycles $\mathcal{C} = \sum_i n_i Z_i$ on X , where the irreducible subvarieties $Z_i \subset X$ all have dimension p and $n_i \in \mathbb{N}$. Let $\mathcal{C}(X) = \cup_{p=0}^n \mathcal{C}_p(X)$. This is a complex space called the *cycle space*. We define the *degree* of a cycle $\mathcal{C} \in \mathcal{C}_p(X)$ by

$$\deg(\mathcal{C}) := \sum_i n_i \int_{Z_i} \frac{\omega^p}{p!},$$

where the integral is performed over the nonsingular locus of Z_i . In other words, the degree is the weighted sum of the volumes of the Z_i with respect to the Kähler metric ω . We will write $|\mathcal{C}| = \cup_i Z_i$ for the support of \mathcal{C} , and $[\mathcal{C}] = \sum_i n_i [Z_i]$ for the cohomology class in $H^{2(n-p)}(X, \mathbb{Q})$ defined by \mathcal{C} . When we occasionally allow negative integers n_i in the definition such objects will be called *generalised cycles*. The following fact will be important later.

Theorem 2.7 (cf. [6, 4]). *A subset $\mathcal{S} \subset \mathcal{C}(X)$ is relatively compact if and only if there is K such that $\deg(\mathcal{C}) \leq K$ for all $\mathcal{C} \in \mathcal{S}$.*

2.5.2. *Singular sets.* For any coherent analytic sheaf $\mathcal{E} \rightarrow X$, define the set of singular points to be

$$\text{sing}(\mathcal{E}) = \{x \in X \mid \mathcal{E} \text{ is not locally free at } x\}.$$

In general, $\text{sing}(\mathcal{E})$ is an analytic subvariety of codimension at least 1. If \mathcal{E} is torsion free it is of codimension at least 2, and if \mathcal{E} is reflexive it is of codimension at least 3. Another description of this set is given by

$$\text{sing}(\mathcal{E}) = \bigcup_{i>0} \text{supp}(\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X)),$$

(see [48, Ch. 2, Sect. 1]). There is a unique *torsion filtration* (see [33, Def. 1.1.4]):

$$\mathcal{T}_0(\mathcal{E}) \subset \mathcal{T}_1(\mathcal{E}) \subset \cdots \subset \mathcal{T}_d(\mathcal{E}) = \mathcal{E},$$

where $d = \dim(\text{supp}(\mathcal{E}))$, and $\mathcal{T}_i(\mathcal{E})$ is the maximal subsheaf of \mathcal{E} fulfilling $\dim \text{supp}(\mathcal{T}_i(\mathcal{E})) \leq i$. By construction, the support of $\mathcal{Q}_i(\mathcal{F}) = \mathcal{T}_i(\mathcal{E})/\mathcal{T}_{i-1}(\mathcal{E})$ is a pure codimension i subvariety if it is nonzero. Notice that if \mathcal{E} is not a torsion sheaf, then $d = n$, and $\mathcal{T}_{n-1}(\mathcal{E}) \subset \mathcal{E}$ is the torsion subsheaf.

Now suppose \mathcal{E} is torsion free, so that we have an injection $\mathcal{E} \hookrightarrow \mathcal{E}^{\vee\vee}$. Since $\mathcal{E}^{\vee\vee}$ is reflexive, it has homological dimension strictly less than $n - 1$, and

therefore $\mathcal{E}xt^i(\mathcal{E}^{\vee\vee}, \mathcal{O}_X) = 0$ for $i \geq n - 1$ (see [34, Prop. V.4.14 (b)] and [48, Ch. 2, Lemma 1.1.1]). Hence,

$$\text{sing}(\mathcal{E}^{\vee\vee}) = \bigcup_{i=1}^{n-2} \text{supp}(\mathcal{E}xt^i(\mathcal{E}^{\vee\vee}, \mathcal{O}_X)) . \quad (2.8)$$

We define the *codimension k singular set of \mathcal{E}* to be

$$\text{sing}_{n-k}(\mathcal{E}) := \text{supp}(\mathcal{Q}_{n-k}(\mathcal{E}^{\vee\vee}/\mathcal{E})) \cup \bigcup_{i=1}^{n-2} \text{supp}(\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}^{\vee\vee}, \mathcal{O}_X))) .$$

This set is exactly the union of all the irreducible components of $\text{sing}(\mathcal{E})$ of codimension k . Notice also that

$$\text{sing}_{n-k}(\mathcal{E}^{\vee\vee}) = \bigcup_{i=1}^{n-2} \text{supp}(\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}^{\vee\vee}, \mathcal{O}_X))) . \quad (2.9)$$

From (2.8) and (2.9) we therefore obtain the following third description of the set $\text{sing}(\mathcal{E})$ when \mathcal{E} is torsion free:

$$\text{sing}(\mathcal{E}) = \bigcup_{k=2}^n \bigcup_{i=1}^{n-2} \text{supp}(\mathcal{Q}_{n-k}(\mathcal{E}^{\vee\vee}/\mathcal{E})) \cup \text{supp}(\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}^{\vee\vee}, \mathcal{O}_X)) ,$$

which by rearranging the terms can be seen to be equal to $\text{supp}(\mathcal{E}^{\vee\vee}/\mathcal{E}) \cup \text{sing}(\mathcal{E}^{\vee\vee})$.

2.5.3. Support cycles. Consider a general torsion sheaf $\mathcal{T} \rightarrow X$ such that $\text{supp}(\mathcal{T})$ has codimension p , and write Z_j for the irreducible components of codimension p . We will write \mathcal{I} for the ideal sheaf of $Z = \text{supp}(\mathcal{T})$ with the induced reduced structure. Then, there is some power \mathcal{I}^N so that $\mathcal{I}^N \mathcal{T} = 0$; this leads to the following filtration of \mathcal{T} :

$$0 = \mathcal{I}^N \mathcal{T} \subset \dots \subset \mathcal{I}^{k+1} \mathcal{T} \subset \mathcal{I}^k \mathcal{T} \subset \dots \subset \mathcal{I} \mathcal{T} \subset \mathcal{T} .$$

We will write

$$\text{gr}_{\mathcal{I}}(\mathcal{T}) := \bigoplus_{k=0}^{N-1} \mathcal{I}^k \mathcal{T} / \mathcal{I}^{k+1} \mathcal{T}$$

for the associated graded of this filtration. Notice that \mathcal{I} annihilates $\text{gr}_{\mathcal{I}}(\mathcal{T})$, so the latter is an \mathcal{O}_Z -module, whose first summand $\mathcal{T}/\mathcal{I}\mathcal{T}$ is precisely the restriction $\mathcal{T}|_Z$. Regarding the restriction $\text{gr}_{\mathcal{I}}(\mathcal{T})|_{Z_j}$ as a sheaf of \mathcal{O}_{Z_j} -modules, we note that the fibre dimension

$$\dim(\text{gr}_{\mathcal{I}}(\mathcal{T})|_{Z_j}(z)) = \dim((\text{gr}_{\mathcal{I}}(\mathcal{T})|_{Z_j})_z / \mathfrak{m}_z((\text{gr}_{\mathcal{I}}(\mathcal{T})|_{Z_j})_z)) \quad (2.10)$$

is constant on a dense open subset of Z_j . We call this natural number the *rank*. We may now associate to each Z_j the multiplicity

$$m_j^{\mathcal{T}} = \text{rank}(\text{gr}_{\mathcal{I}}(\mathcal{T})|_{Z_j}) .$$

Thus, associated to the sheaf \mathcal{T} is the effective analytic (algebraic) cycle

$$\mathcal{C}_{\mathcal{T}} = \sum_j m_j^{\mathcal{T}} \cdot Z_j \in \mathcal{C}_{n-p}(X) , \quad (2.11)$$

which we call the *support cycle* of \mathcal{T} .

To a semistable sheaf \mathcal{E} , we may associate a canonical element $\mathcal{C}_{\mathcal{E}} \in \mathcal{C}_{n-2}(X)$ as follows. Consider the torsion sheaf

$$\mathcal{T}_{\mathcal{E}} := (\text{Gr } \mathcal{E})^{\vee\vee} / \text{Gr } \mathcal{E} .$$

Since $\text{Gr } \mathcal{E}$ is torsion free, $\text{sing}(\text{Gr } (\mathcal{E}))$ has codimension at least 2, and admits the decomposition

$$\text{sing}(\text{Gr } \mathcal{E}) = \text{supp}((\text{Gr } \mathcal{E})^{\vee\vee} / \text{Gr } \mathcal{E}) \cup \text{sing}((\text{Gr } \mathcal{E})^{\vee\vee}) .$$

In particular, $\text{supp}((\text{Gr } \mathcal{E})^{\vee\vee} / \text{Gr } \mathcal{E})$ has codimension at least 2 as well. We may therefore define the *support cycle of \mathcal{E}* by

$$\mathcal{C}_{\mathcal{E}} := \mathcal{C}_{\mathcal{T}_{\mathcal{E}}} \in \mathcal{C}_{n-2}(X) . \quad (2.12)$$

Remark 2.8. The pair $((\text{Gr } \mathcal{E})^{\vee\vee}, \mathcal{C}_{\mathcal{E}})$ is uniquely determined by \mathcal{E} and the polarisation defining semistability (see [7, Appendix]).

As in [54] (see also the latest preprint version), we define the *multiplicities* of \mathcal{E} to be the positive integers $m_j^{\mathcal{E}} = \text{rank}(\text{gr}_{\mathcal{I}}(\mathcal{T}_{\mathcal{E}})|_{Z_j^{\mathcal{E}}})$ that appear as the coefficients of $\mathcal{C}_{\mathcal{E}}$. In [30, Rem. 5.3] the multiplicities are defined differently. Namely, let S be a general complete intersection surface intersecting $Z_j^{\mathcal{E}}$ transversally in a finite number of smooth points $\{z_1, \dots, z_k\}$, and set

$$n_j^{\mathcal{E}}(S, z_i) = \ell_{\mathcal{O}_{S, z_i}}(\mathcal{T}_{\mathcal{E}}|_S)_{z_i} ,$$

where ℓ denotes the length of a module. Since we will rely on results from both [54] and [30] that were proven using these two definitions, we point out the following simple lemma.

Lemma 2.9. *For general S as above, $m_j^{\mathcal{E}} = n_j^{\mathcal{E}}(S, z_i)$. In particular, $n_j^{\mathcal{E}}(S, z_i)$ is independent of z_i and S .*

Proof. We begin with the observation that, since the ideal \mathcal{I} that cuts out the support of $\mathcal{T}_{\mathcal{E}}$ annihilates $\text{gr}_{\mathcal{I}}(\mathcal{T}_{\mathcal{E}})$, then if ι is the natural inclusion of the support in X , we have $\text{gr}_{\mathcal{I}}(\mathcal{T}_{\mathcal{E}}) = \iota_*(\iota^* \text{gr}_{\mathcal{I}}(\mathcal{T}_{\mathcal{E}}))$. In other words, for a sufficiently small open set U containing $z_i \in Z_j$, $\text{gr}_{\mathcal{I}}(\mathcal{T})$ may be expressed

as an \mathcal{O}_X -module as $T \otimes_{\mathbb{C}} \mathcal{O}_{Z_j}$, for a \mathbb{C} -vector space T of dimension $m_j^{\mathcal{E}}$. Furthermore, $T = \bigoplus T_k$, where T_k is the vector space corresponding to each summand of $\mathrm{gr}_{\mathcal{I}}(\mathcal{T}_{\mathcal{E}})$. We consider the exact sequence

$$0 \longrightarrow T \otimes_{\mathbb{C}} \mathcal{I}_{U \cap Z_j} \longrightarrow T \otimes_{\mathbb{C}} \mathcal{O}_U \longrightarrow \mathrm{gr}_{\mathcal{I}}(\mathcal{T}_{\mathcal{E}})|_U \longrightarrow 0 .$$

By our assumptions, tensoring with \mathcal{O}_S we get an exact sequence of \mathcal{O}_S -modules:

$$0 \longrightarrow T \otimes_{\mathbb{C}} \mathcal{I}_{z_i} \longrightarrow T \otimes_{\mathbb{C}} \mathcal{O}_{U \cap S} \longrightarrow \mathrm{gr}_{\mathcal{I}}(\mathcal{T}_{\mathcal{E}}) \otimes_{\mathcal{O}_U} \mathcal{O}_{U \cap S} \longrightarrow 0 .$$

The length of the torsion quotient on the right hand side (the dimension of its fibre over z_i) is by definition $\ell_{\mathcal{O}_{S, z_i}}(\mathrm{gr}_{\mathcal{I}}(\mathcal{T}_{\mathcal{E}})|_S)_{z_i}$. The relevant fibre over z_i is clearly isomorphic so T , so this number is equal to $\dim T$.

On the other hand, considering the exact sequences

$$0 \rightarrow \mathcal{I}^{k+1} \mathcal{T}|_U \rightarrow \mathcal{I}^k \mathcal{T}|_U \rightarrow T_k \otimes \mathcal{O}_{Z_j \cap U} \rightarrow 0 ,$$

and again using the assumptions on Z_j and S , we have $\mathrm{Tor}_1^{\mathcal{O}_X}(\mathcal{O}_{Z_j}, \mathcal{O}_S) = 0$, and so tensoring by \mathcal{O}_S we obtain

$$0 \rightarrow \mathcal{I}^{k+1} \mathcal{T}|_{U \cap S} \rightarrow \mathcal{I}^k \mathcal{T}|_{U \cap S} \rightarrow T_k \otimes \mathcal{O}_{z_i} \rightarrow 0 .$$

Using additivity of the length in exact sequences, we obtain

$$n_j(S, z_i) = \ell_{\mathcal{O}_{S, z_i}}(\mathcal{T}_{\mathcal{E}}|_S)_{z_i} = \sum_{k=0}^{N-1} \dim T_k = \dim T = m_j^{\mathcal{E}} . \quad \square$$

We will repeatedly use the following result from [54] (see also the newest preprint version) which is certainly also well-known in the algebraic geometry literature. It relates the support cycle defined above to the first nonzero part of the Chern character of \mathcal{T} .

Lemma 2.10. *Let \mathcal{T} be a torsion sheaf with $\mathrm{codim}(\mathrm{supp}(\mathcal{T})) = p$. Then $ch_k(\mathcal{T}) = 0$ for $k < p$, and $ch_p(\mathcal{T})$ coincides with the Poincaré dual in $H^{2p}(X, \mathbb{Q})$ of the homology class of the cycle $\mathcal{C}_{\mathcal{T}}$ defined in (2.11).*

2.5.4. *Generalised cycles.* For inductive arguments, the following more general notion of associated cycle is often useful.

Definition 2.11. Let $\mathcal{E} \rightarrow X$ be coherent such that its torsion part $\mathrm{Tors}(\mathcal{E})$ has support in codimension ≥ 2 . We define the (*generalised*) *codimension 2 cycle of \mathcal{E}* as

$$\mathcal{C}(\mathcal{E}) := \mathcal{C}_{\ker(\mathcal{E} \rightarrow \mathcal{E}^{\vee\vee})} - \mathcal{C}_{\mathrm{coker}(\mathcal{E} \rightarrow \mathcal{E}^{\vee\vee})} .$$

We further set, $\gamma(\mathcal{E}) := (\mathcal{E}^{\vee\vee}, -\mathcal{C}(\mathcal{E}))$. If \mathcal{E} torsion free, we set $\mathcal{Q}_{\mathcal{E}} := \mathcal{E}^{\vee\vee}/\mathcal{E}$ and we put $\widehat{\mathcal{Q}}_{\mathcal{E}} := \mathcal{Q}_{\mathcal{E}}/\mathcal{T}(\mathcal{Q}_{\mathcal{E}})$, where $\mathcal{T}(\mathcal{Q}_{\mathcal{E}})$, where $\mathcal{T}(\mathcal{Q}_{\mathcal{F}}) = \mathcal{T}_{n-3}(\mathcal{Q}_{\mathcal{F}})$ is the

maximal subsheaf of $\mathcal{Q}_{\mathcal{F}}$ of dimension less than or equal to $n - 3$; cf. [33, Def. 1.1.4]. Notice that since $\mathcal{Q}_{\mathcal{E}}$ has support in codimension at least 2, $\widehat{\mathcal{Q}}_{\mathcal{E}}$ is pure of codimension 2 or vanishes. In this case $\mathcal{C}(\mathcal{E}) = -\mathcal{C}_{\mathcal{Q}_{\mathcal{E}}} = -\mathcal{C}_{\widehat{\mathcal{Q}}_{\mathcal{E}}}$. Moreover, we will write $[\gamma(\mathcal{E})]$ for the class of $\gamma(\mathcal{E})$ under the natural equivalence relation taking isomorphism classes in the first component and equality of cycles in the second component.

Remark 2.12. Note that if \mathcal{E} is polystable, then $\mathcal{E} = \text{Gr } \mathcal{E}$, $\mathcal{Q}_{\mathcal{E}}$ is equal to the torsion sheaf $\mathcal{T}_{\mathcal{E}}$ previously defined, and $\mathcal{C}(\mathcal{E}) = -\mathcal{C}_{\mathcal{E}}$ so that $\gamma(\mathcal{E}) = (\mathcal{E}^{\vee\vee}, \mathcal{C}_{\widehat{\mathcal{Q}}_{\mathcal{E}}}) = (\mathcal{E}^{\vee\vee}, \mathcal{C}_{\mathcal{Q}_{\mathcal{E}}}) = (\mathcal{E}^{\vee\vee}, \mathcal{C}_{\mathcal{E}})$.

The following three results gather properties of codimension 2 cycles.

Lemma 2.13. *If $0 \rightarrow \mathcal{E}_1 \rightarrow \cdots \rightarrow \mathcal{E}_m \rightarrow 0$ is an exact sequence of coherent sheaves on X with supports of codimension ≥ 2 , then $\sum_j (-1)^j \mathcal{C}(\mathcal{E}_j) = 0$.*

Proof. This is easily checked by cutting with general complete intersection surfaces $S \subset X$ and using the fact that the multiplicities of $\mathcal{C}(\mathcal{E}_j)$ are equal to the lengths of the skyscraper sheaves $\mathcal{E}_j|_S$. \square

Lemma 2.14. *Let $\alpha : \mathcal{E}_1 \rightarrow \mathcal{E}_2$ be a morphism between coherent sheaves on X inducing an isomorphism $\alpha^{\vee\vee} : \mathcal{E}_1^{\vee\vee} \xrightarrow{\simeq} \mathcal{E}_2^{\vee\vee}$ between the double duals. Then*

$$\mathcal{C}(\mathcal{E}_1) = \mathcal{C}(\mathcal{E}_2) + \mathcal{C}(\ker \alpha) - \mathcal{C}(\text{coker } \alpha) .$$

Proof. The idea is to reduce the situation to an application of Lemma 2.13 on several exact sequences of coherent sheaves of codimension at least 2. By decomposing the sequence

$$0 \rightarrow \ker \alpha \rightarrow \mathcal{E}_1 \rightarrow \mathcal{E}_2 \rightarrow \text{coker } \alpha \rightarrow 0$$

into short exact sequences one first reduces the question to the cases where $\ker \alpha = 0$ or $\text{coker } \alpha = 0$. One then compares these short exact sequences with the morphisms $\mathcal{E}_1 \rightarrow \mathcal{E}_1^{\vee\vee}$ and $\mathcal{E}_2 \rightarrow \mathcal{E}_2^{\vee\vee}$. \square

The next basic result will be needed in Sections 4.2 and 4.3; see [7, Appendix] and [12, Cor. 2.23] for very similar results.

Proposition 2.15. *Let \mathcal{E} be a torsion free sheaf such that $\mathcal{E}^{\vee\vee}$ is polystable, and let $\mathcal{T} = \mathcal{E}^{\vee\vee}/\mathcal{E}$. Then \mathcal{E} is μ -semistable. Moreover, if $\text{Gr } \mathcal{E}$ is the associated graded to any Jordan-Hölder filtration of \mathcal{E} , then $(\text{Gr } \mathcal{E})^{\vee\vee} \cong \mathcal{E}^{\vee\vee}$ and $\mathcal{C}_{\mathcal{E}} = \mathcal{C}_{\mathcal{T}}$.*

Proof. Since $\mathcal{E}^{\vee\vee}$ is polystable it is in particular μ -semistable. Suppose $\mathcal{F} \subset \mathcal{E}$ is a proper subsheaf with $\mu(\mathcal{F}) > \mu(\mathcal{E})$. Then \mathcal{F} is also a subsheaf of $\mathcal{E}^{\vee\vee}$ with $\mu(\mathcal{F}) > \mu(\mathcal{E}^{\vee\vee})$, since \mathcal{E} and $\mathcal{E}^{\vee\vee}$ coincide in codimension one and hence have

the same slope. This violates semistability of $\mathcal{E}^{\vee\vee}$, and so \mathcal{E} is μ -semistable, proving the first claim.

Recall the definition of the Jordan-Hölder filtration (2.7). We claim that taking the saturations $\widehat{\mathcal{E}}_i$ of the \mathcal{E}_i inside $\mathcal{E}^{\vee\vee}$ gives a Jordan-Hölder filtration for $\mathcal{E}^{\vee\vee}$. First of all, $\widehat{Q}_i = \widehat{\mathcal{E}}_i/\widehat{\mathcal{E}}_{i-1}$ is torsion free by construction. Moreover, there is an isomorphism

$$(\widehat{Q}_i)^{\vee\vee} \cong (Q_i)^{\vee\vee} = (\mathcal{E}_i/\mathcal{E}_{i-1})^{\vee\vee}, \quad (2.13)$$

since these sheaves are isomorphic away from $\text{supp}(\mathcal{T})$, which has codimension greater than or equal to two. In fact, notice that since a saturated subsheaf of a reflexive sheaf is reflexive, we actually have $\widehat{\mathcal{E}}_i \cong \mathcal{E}_i^{\vee\vee}$, by normality. The stability of \widehat{Q}_i then follows since Q_i is stable by definition, and a sheaf is stable if and only if its double dual is, see [48, Chap. II, Lemma 1.2.4]. Therefore $\widehat{\mathcal{E}}_\bullet$ defines a Jordan-Hölder filtration.

It follows from (2.13) that $(\text{Gr } \mathcal{E})^{\vee\vee} \cong \bigoplus_i (\widehat{Q}_i)^{\vee\vee}$. The right hand side of this isomorphism is the double dual of the associated graded object for the Jordan-Hölder filtration for $\mathcal{E}^{\vee\vee}$ by $\widehat{\mathcal{E}}_i \cong \mathcal{E}_i^{\vee\vee}$. Since this sheaf is uniquely associated to $\mathcal{E}^{\vee\vee}$ by Remark 2.8, and $\mathcal{E}^{\vee\vee}$ is polystable with possibly different Jordan-Hölder filtration whose successive quotients are the direct summands in the decomposition of $\mathcal{E}^{\vee\vee}$, we obtain the isomorphism $(\text{Gr } \mathcal{E})^{\vee\vee} \cong \mathcal{E}^{\vee\vee}$.

In order to prove the claim regarding cycles, we start by observing that $(\mathcal{E}/\mathcal{E}_1)^{\vee\vee} \cong \mathcal{E}^{\vee\vee}/\mathcal{E}_1^{\vee\vee}$, as $\mathcal{E}^{\vee\vee}/\mathcal{E}_1^{\vee\vee}$ is a direct summand of $\text{Gr}(\mathcal{E})^{\vee\vee}$ and hence reflexive. Therefore, we obtain the following diagram with exact rows:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \mathcal{E}_1 & \longrightarrow & \mathcal{E} & \longrightarrow & \mathcal{E}/\mathcal{E}_1 & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow \alpha & & \\ 0 & \longrightarrow & \mathcal{E}_1^{\vee\vee} & \longrightarrow & \mathcal{E}^{\vee\vee} & \longrightarrow & (\mathcal{E}/\mathcal{E}_1)^{\vee\vee} & \longrightarrow & 0 \end{array}$$

Notice that α is injective, as $\mathcal{E}/\mathcal{E}_1$ is torsion free. An application of the snake lemma yields an exact sequence

$$0 \rightarrow \mathcal{E}_1^{\vee\vee}/\mathcal{E}_1 \rightarrow \mathcal{E}^{\vee\vee}/\mathcal{E} \rightarrow (\mathcal{E}/\mathcal{E}_1)^{\vee\vee}/(\mathcal{E}/\mathcal{E}_1) \rightarrow 0.$$

Lemma 2.13 therefore implies that $\mathcal{C}(\mathcal{E}) = \mathcal{C}(\mathcal{E}_1) + \mathcal{C}(\mathcal{E}/\mathcal{E}_1)$. Observe that $\mathcal{E}/\mathcal{E}_1$ is again torsion free with polystable double dual, so that Remark 2.12 allows us to conclude by induction. \square

2.5.5. *Some boundedness results.* The following will be used later on.

Proposition 2.16. *Let \mathcal{S} be a set of semistable sheaves \mathcal{E} that is bounded in the sense of [31]. Then the set of all possible Seshadri graduations $\text{Gr } \mathcal{E}$ for $\mathcal{E} \in \mathcal{S}$ is also bounded.*

Proof. Let r be the maximal rank of the sheaves appearing in \mathcal{S} . Let \mathcal{E} any sheaf in \mathcal{S} . Then the length ℓ of any Seshadri filtration (2.7) of \mathcal{E} is bounded by r . We have $\text{Gr } \mathcal{E} = \bigoplus_{i=1}^{\ell} \mathcal{E}_i/\mathcal{E}_{i-1}$, and it is enough to show that the isomorphism classes of all quotients $\mathcal{E}_i/\mathcal{E}_{i-1}$ arising from such filtrations form a bounded set. If \mathcal{E} is such that $\ell = 1$, we have $\text{Gr } \mathcal{E} = \mathcal{E}_1/\mathcal{E}_0 = \mathcal{E}$, and the statement is clear.

Consider now the subset of \mathcal{S} consisting of those \mathcal{E} with $\ell > 1$. Then for such sheaves \mathcal{E} , the isomorphism classes of quotients $\mathcal{E}_\ell/\mathcal{E}_{\ell-1} = \mathcal{E}/\mathcal{E}_{\ell-1}$ belong to a bounded set by Grothendieck's lemma, [33, Lemma 1.7.9], since their slope equals to $\mu(\mathcal{E})$ and is therefore bounded. The sheaves $\mathcal{E}_{\ell-1}$ may now be viewed as kernels of the projections $\mathcal{E}_\ell \rightarrow \mathcal{E}_\ell/\mathcal{E}_{\ell-1}$ so their isomorphism classes form a bounded set \mathcal{S}' . We continue by applying the above procedure to the set \mathcal{S}' and conclude by descending induction on r . \square

In the next section we will also require a boundedness result for the codimension k singular sets of a bounded set of polystable reflexive sheaves. We first need the following elementary proposition.

Proposition 2.17. *Let \mathcal{S} be a set of polystable reflexive sheaves \mathcal{E} that is bounded in the sense of [31]. Then for each k and each i , the set*

$$\{\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X))\}_{\mathcal{E} \in \mathcal{S}}$$

is also bounded.

Proof. Since \mathcal{S} is bounded, by [31, p. 251], we know that the sets of sheaves $\{\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X)\}_{\mathcal{E} \in \mathcal{S}}$, are bounded. This means that for any i, j, k , the associated sets $\{\mathcal{T}_j(\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X))\}_{\mathcal{E} \in \mathcal{S}}$, and therefore also $\{\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X))\}_{\mathcal{E} \in \mathcal{S}}$, are bounded. \square

Corollary 2.18. *In the setup of Proposition 2.17, the set $\{\text{deg sing}_{n-k}(\mathcal{E})\}_{\mathcal{E} \in \mathcal{S}}$, is finite for each k , where we regard $\text{sing}_{n-k}(\mathcal{E})$ as an element of $\mathcal{C}_{n-k}(X)$ by assigning the weight 1 to each of its irreducible components. In particular, the number of irreducible components of the sets $\text{sing}_{n-k}(\mathcal{E})$ is bounded. As a result, the set $\{\text{sing}_{n-k}(\mathcal{E})\}_{\mathcal{E} \in \mathcal{S}}$ is relatively compact in $\mathcal{C}_{n-k}(X)$.*

Proof. By (2.9), we have

$$\text{sing}_{n-k}(\mathcal{E}) = \bigcup_{i=1}^{n-2} \text{supp}(\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X))),$$

and by the preceding proposition the family of sheaves appearing on the right hand side is bounded. By [31, Thm. 2.1], this implies in particular that for each i the set of Hilbert polynomials $\{\chi_{\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X))}(m)\}_{\mathcal{E} \in \mathcal{S}}$, is finite. The

$(n - k)$ -th coefficients of these polynomials are precisely

$$\int_X \text{ch}_k(\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X))) \wedge \frac{\omega^{n-k}}{(n-k)!},$$

and these are therefore finite in number. On the other hand,

$$\text{ch}_k(\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X))) = \text{PD}[\mathcal{C}_{\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X))}]$$

are the (Poincaré duals of) the support cycles of these torsion sheaves (see Lemma 2.10). Therefore, the coefficients are given by the degrees of the cycles $\mathcal{C}_{\mathcal{Q}_{n-k}(\mathcal{E}xt^i(\mathcal{E}, \mathcal{O}_X))}$, and so there are only finitely many such degrees. Since the sum of these cycles is exactly the cycle associated to $\text{sing}_{n-k}(\mathcal{E})$, except with possibly larger multiplicities, the set of degrees $\text{deg}(\text{sing}_{n-k}(\mathcal{E}))$ is also finite. The second statement follows directly from Theorem 2.7. \square

2.6. The Quot scheme, natural subschemes, and convergence. The starting point for forming moduli spaces of sheaves is that the set of slope (or GM-) semistable sheaves with fixed Chern classes $c_i(E)$ is a bounded family (see for example [42]). In practice this means that there is a natural number $m_0 \gg 0$, such that for any $m > m_0$ and any such sheaf \mathcal{E} , $H^i(X, \mathcal{E}(m)) = 0$ for $i > 0$, and $\mathcal{E}(m)$ is globally generated.

Write $V = \mathbb{C}^{\tau_E(m)}$ and set

$$\mathcal{H} = V \otimes \mathcal{O}_X(-m). \quad (2.14)$$

If we choose an isomorphism $H^0(X, \mathcal{E}(m)) \cong V$ with m chosen as above, then we have a surjection $\mathcal{H} \rightarrow \mathcal{E} \rightarrow 0$. Notice that the map we obtain depends on our choice of isomorphism. Hence, we see that to any slope or GM-semistable sheaf with Chern classes $c_i(E)$ we can associate a point in the *Quot scheme* $\text{Quot}(\mathcal{H}, \tau_E)$, which is defined as the set of equivalence classes of quotients $q_{\mathcal{E}} : \mathcal{H} \rightarrow \mathcal{E} \rightarrow 0$, where $\mathcal{E} \rightarrow X$ is a coherent sheaf with Hilbert polynomial equal to τ_E . Quotients q_1 and q_2 are equivalent if $\ker q_1 = \ker q_2$. This is equivalent to the existence of a commutative diagram

$$\begin{array}{ccccc} \mathcal{H} & \xrightarrow{q_1} & \mathcal{E} & \longrightarrow & 0 \\ \parallel & & \downarrow \varphi & & \\ \mathcal{H} & \xrightarrow{q_2} & \mathcal{E} & \longrightarrow & 0 \end{array}$$

where φ is an isomorphism.

Observe that there is an action $\text{GL}(V) \curvearrowright \text{Quot}(\mathcal{H}, \tau)$, namely, the map $g \cdot q_{\mathcal{E}} : \mathcal{H} \rightarrow \mathcal{E} \rightarrow 0$ is given by composing $q_{\mathcal{E}}$ with g . This action amounts to the fact that there is an ambiguity due to the choice of isomorphism $H^0(X, \mathcal{E}(m)) \cong V$ (i.e. a choice of basis for $H^0(X, \mathcal{E}(m))$). Since \mathbb{C}^* acts trivially (rescaling the

basis vectors by the same constant results in the same kernel), it will suffice to consider the action of $\mathrm{SL}(V)$.

By the construction in [31], $\mathrm{Quot}(\mathcal{H}, \tau_E)$ is a projective scheme. It is furthermore a fine moduli space, so in particular there is a universal quotient sheaf $q_{\mathcal{U}} : \mathcal{H} \rightarrow \mathcal{U} \rightarrow 0$, where \mathcal{H} is the pullback of \mathcal{H} via the projection map, and $\mathcal{U} \rightarrow X \times \mathrm{Quot}(\mathcal{H}, \tau_E)$ is a flat family so that $\mathcal{U}|_{X \times \{q_E\}} = \mathcal{E}$, and $q_{\mathcal{U}}|_{X \times \{q_E\}} = q_E$. There are subspaces $\mathrm{Quot}(\mathcal{H}, (c_1, \dots, c_{\min(r,n)})) \subset \mathrm{Quot}(\mathcal{H}, \tau_E)$ consisting of those quotients with fixed Chern classes $(c_1, \dots, c_{\min(r,n)})$. Since the Chern classes of a flat family (in particular those of \mathcal{U}) are locally constant, we have a decomposition:

$$\mathrm{Quot}(\mathcal{H}, \tau_E) = \coprod \mathrm{Quot}(\mathcal{H}, (c_1, \dots, c_{\min(r,n)})) ,$$

where the union is over all tuples of Chern classes whose associated Hilbert polynomial is τ_E . We will write

$$\mathrm{Quot}(\mathcal{H}, c(E)) := \mathrm{Quot}(\mathcal{H}, (c_1(E), \dots, c_{\min(r,n)}(E))) .$$

Let $R^{\mu\mathrm{ss}}$ (resp. R^{Gss}) $\subset \mathrm{Quot}(\mathcal{H}, c(E))$ denote the subscheme of quotients $q : \mathcal{H} \rightarrow \mathcal{E}$ satisfying:

- (1) \mathcal{E} is torsion free;
- (2) $\det \mathcal{E} \simeq \mathcal{J}$;
- (3) \mathcal{E} is μ -semistable (resp. GM-semistable);
- (4) q induces an isomorphism $V \xrightarrow{\simeq} H^0(X, \mathcal{E}(m_0))$.

The spaces $R^{\mu\mathrm{ss}}$ and R^{Gss} are preserved by the action of $\mathrm{SL}(V)$, and there is an inclusion $R^{\mathrm{Gss}} \hookrightarrow R^{\mu\mathrm{ss}}$.

In Section 3 we will need a result concerning the meaning of convergence of a sequence in the space $\mathrm{Quot}(\mathcal{H}, \tau_E)$ in the analytic topology. For the following, fix a hermitian structure on \mathcal{H} .

Lemma 2.19. *Let $q_i : \mathcal{H} \rightarrow \mathcal{E}_i \rightarrow 0$, $q : \mathcal{H} \rightarrow \mathcal{E} \rightarrow 0$ be points in $\mathrm{Quot}(\mathcal{H}, \tau)$. suppose each \mathcal{E}_i is locally free with underlying C^∞ -bundle smoothly isomorphic to E . Let π_i denote the orthogonal projections to $\ker q_i$, and π the orthogonal projection to $\ker q$ on the open set $X \setminus \mathrm{sing} \mathcal{E}$ where \mathcal{E} is locally free. If $q_i \rightarrow q$ in the analytic topology of $\mathrm{Quot}(\mathcal{H}, \tau)$, then*

- (1) $\pi_i \rightarrow \pi$ smoothly on $X \setminus \mathrm{sing} \mathcal{E}$, and
- (2) on $X \setminus \mathrm{sing} \mathcal{E}$, the underlying C^∞ -bundle of \mathcal{E} is smoothly isomorphic to E .

Proof. Recall from [31] that $\mathrm{Quot}(\mathcal{H}, \tau_E)$ admits an embedding into a Grassmannian as follows. By choosing $k \gg 0$ we may assume

$$H^1(X, \ker q \otimes \mathcal{O}_X(k)) = \{0\}$$

for all $q \in \text{Quot}(\mathcal{H}, \tau)$. Hence, q defines a point $P_q \in \text{Gr}(r, N)$, where $N = \dim W$, where $W := H^0(X, \mathcal{H}(k))$, and $r = N - \tau(k)$. By choosing hermitian structures, we may regard P_q as the element of $\text{End } W$ given by orthogonal projection to $K_q := H^0(X, \ker q \otimes \mathcal{O}_X(k))$. Applying this to the situation in the statement of the lemma, convergence $q_i \rightarrow q$ in the analytic topology implies smooth convergence $P_{q_i} \rightarrow P_q$ in $\text{End } W$. This gives a sequence of smooth bundle morphisms $W \otimes \mathcal{O}_X \rightarrow \mathcal{H}(k)$ gotten from the composition

$$W \otimes \mathcal{O}_X \xrightarrow{P_{q_i}} W \otimes \mathcal{O}_X \longrightarrow \mathcal{H}(k),$$

whose images are precisely $\ker q_i \otimes \mathcal{O}_X(k)$. Since the bundle morphisms converge smoothly, items (1) and (2) clearly follow. \square

2.7. Flat families and line bundles over Quot schemes. The line bundles in question come from the following general construction (see [33, Section 8.1] for more details). Let S be a scheme over \mathbb{C} . Given an S -flat family of coherent sheaves $\mathcal{E} \rightarrow X \times S$, the determinant of cohomology $\lambda_{\mathcal{E}} : K(X) \rightarrow \text{Pic}(S)$ is defined by

$$\lambda_{\mathcal{E}}(u) := \det((p_S)_!(p_X^*(u) \otimes \mathcal{E})). \quad (2.15)$$

Here, $K(X) = K_0(X) = K^0(X)$ is the Grothendieck group of holomorphic vector bundles (and of coherent sheaves) on X , $p_X : X \times S \rightarrow X$ and $p_S : X \times S \rightarrow S$ are the projections, and $(p_S)_! : K^0(X \times S) \rightarrow K^0(S)$ is defined on a class $[F]$ represented by a sheaf \mathcal{F} as

$$(p_S)_!([F]) = \sum_{i=0}^{\dim X} (-1)^i [R^i(p_S)_* \mathcal{F}] \in K^0(S),$$

and is then extended to $K^0(X)$ by linearity. Moreover, in our situation, it can be shown that $(p_S)_!(p_X^*(u) \otimes \mathcal{E}) \in K_0(S)$, so that applying the determinant in (2.15) makes sense; see the discussion preceding [33, Cor. 2.1.11] for details. Therefore, we may construct line bundles on the parameter space S by specifying classes in $K(X)$. One natural way to do this is to consider classes arising from complete intersections as follows; cf. [30, Section 3.2].

Let $H \subset X \subset \mathbb{P}^N$ be a hyperplane section and consider the class $[\mathcal{O}_H] \in K(X)$. We fix a class $c \in K(X)_{\text{num}} := K(X)/\sim$, where $u \sim v$ iff the difference $u - v$ lies in the kernel of the quadratic form induced by the Euler characteristic χ , i.e., $u - v \in \ker((a, b) \mapsto \chi(a \cdot b))$. By the Riemann-Roch Theorem, the numerical behaviour of such a class is hence determined by its associated rank r and its Chern classes $c_i \in H^{2i}(X, \mathbb{Z})$. As we consider sheaves with fixed determinant, we furthermore fix a line bundle $\mathcal{J} \in \text{Pic}(X)$ such that $c_1(\mathcal{J}) = c_1$. Now, for any integer $1 \leq i \leq n - 1$, let $X_i \in |H|$ and write

$X^{(0)} = X$ and $X^{(l)} = \cap_{i=1}^l X_i$ with $1 \leq l \leq n-1$. We will assume that $X^{(l)}$ is smooth for each l . Finally, fix a basepoint $x \in X^{(n-1)}$, whose structure sheaf \mathcal{O}_x defines a class $[\mathcal{O}_x]$ in each $K(X^{(l)})$.

Then, for each $1 \leq l \leq n-1$ we define a class in $K(X^{(l)})$ by

$$u_{n-1-l}(c|_{X^{(l)}}, [\mathcal{O}_H]_{X^{(l)}}) := -r([\mathcal{O}_H]|_{X^{(l)}})^{n-1-l} + \chi(c|_{X^{(l)}} \cdot ([\mathcal{O}_H]|_{X^{(l)}})^{n-1-l})[\mathcal{O}_x].$$

Here, multiplication "." and powers " \bullet^{n-1-l} " are taken with respect to the ring structure of $K(X^{(l)})$, which is derived from the operation of taking tensor product of locally free sheaves.

Supposing furthermore that the family $\mathcal{E}|_{X^{(l)} \times S} \rightarrow X^{(l)} \times S$ is flat over S for each l , by using the previous construction, we can produce a sequence of line bundles

$$\mathcal{L}_{\mathcal{E}, n-1-l} := \lambda_{\mathcal{E}}(u_{n-1-l}(c|_{X^{(l)}}, [\mathcal{O}_H]_{X^{(l)}})).$$

The main case of interest in this article is $l = 0$, $r = \text{rank}(E)$, $c_i = c_i(E)$, and $\mathcal{E} = \mathcal{U}$, the universal sheaf on $\text{Quot}(\mathcal{H}, \tau_E)$. By the previous discussion we obtain a line bundle $\mathcal{L}_{\mathcal{U}, n-1} \rightarrow \text{Quot}(\mathcal{H}, \tau_E)$ which we abbreviate by \mathcal{L}_{n-1} . We will discuss properties of this line bundle in the following subsection.

2.8. Compactifications by sheaves. By Theorem 2.6, the moduli space M_{HYM}^* is identified with the complex analytic space M^s . By a result of Miyajima [44], the latter is isomorphic as a complex analytic space to both M_{an}^s and M_{alg}^s , which are moduli spaces corepresenting the appropriate moduli functors (namely those associating to a complex space or \mathbb{C} -scheme the set of isomorphism classes of flat families of slope stable holomorphic or algebraic bundles which are smoothly isomorphic to E over this space).³ We will sometimes identify all three spaces and use the notation M^s .

Since M_{alg}^s corepresents a subfunctor of the moduli functor for the Gieseker moduli space M^{Gss} of GM-semistable sheaves, and forms a Zariski open subspace, M^{Gss} compactifies M^s , and therefore also M_{HYM}^* . Below we briefly recall the construction in [30] of another compactification of M_{HYM}^* based instead on slope semistable sheaves. Both of these compactifications arise as the image of a scheme in projective space under sections of line bundles produced by the construction of the previous section.

Denote by $\mathcal{L}_{\text{Gss}} \in \text{Pic}(\text{Quot}(\mathcal{H}, \tau_E))$ the bundle $\lambda_{\mathcal{U}}(\mathcal{O}_X(l))$ associated to the universal sheaf $\mathcal{U} \rightarrow X \times \text{Quot}(\mathcal{H}, c(E))$. This turns out to be ample for l sufficiently large, and it possesses an $\text{SL}(V)$ linearisation. The reason for the subscript *Gss* is that the subspace $R^{\text{Gss}} \subset \text{Quot}(\mathcal{H}, c(E))$, is exactly the set of

³In the case of M_{alg}^s we really refer to its analytification.

GIT semistable points of \mathcal{L}_{Gss} on the Zariski closure $\overline{R^{\text{Gss}}}$ with respect to this linearisation. The *Gieseker moduli space* is the GIT quotient

$$M^{\text{Gss}} := \overline{R^{\text{Gss}}} //_{\mathcal{L}_{\text{Gss}}} \text{SL}(V).$$

By definition this is the Proj of the invariant section ring of \mathcal{L}_{Gss} :

$$M^{\text{Gss}} = \text{Proj} \left(\bigoplus_k H^0 \left(\overline{R^{\text{Gss}}}, \mathcal{L}_{\text{Gss}}^k \right)^{\text{SL}(V)} \right).$$

A point in M^{Gss} is represented by a GM-semistable quotient $q_{\mathcal{E}} : \mathcal{H} \rightarrow \mathcal{E}$ with $\det \mathcal{E} \simeq \mathcal{J}$. Moreover GIT gives a way to understand the geometry of M^{Gss} . More precisely, two quotients $q_{\mathcal{E}_1}, q_{\mathcal{E}_2} \in R^{\text{Gss}}$ represent the same point of M^{Gss} if and only if they are s -equivalent. In particular, the space $M_{\text{alg}}^s \cong M^s$ embeds as a Zariski open set.

Definition 2.20. The *Gieseker compactification* \overline{M}^{GM} of M^s is defined as the Zariski closure $\overline{M}^s \subset M^{\text{Gss}}$.

There is no linearised ample line bundle for which the GIT semistable points coincide with the slope semistable sheaves. Nevertheless, several authors have considered the line bundle

$$\mathcal{L}_{n-1} = \lambda_{\mathcal{U}}(u_{n-1}(c, [\mathcal{O}_H])) \in \text{Pic}(\text{Quot}(\mathcal{H}, c(E))).$$

When $\dim X = 2$, the restriction to $R^{\mu\text{ss}}$ was studied by Le Potier and Li ([37, 38], and see also [33]) and it was found to be equivariantly semiample.

For technical reasons explained in [30, Sect. 1.2], the situation is more subtle in higher dimensions. One is more or less forced to consider the weak normalisation $(R^{\mu\text{ss}})^{\text{wn}}$. A *weakly normal* complex space \mathcal{Z} is one for which every locally defined continuous function on \mathcal{Z} which is holomorphic on restriction to the smooth points \mathcal{Z}_{reg} of \mathcal{Z} is in fact holomorphic. Every complex space \mathcal{Z} has a weak normalisation \mathcal{Z}^{wn} which is a reduced weakly normal complex space homeomorphic to \mathcal{Z} (see [30, Sect. 2.3] for a summary of the relevant theory and for references). Henceforth, we write $\mathcal{Z} = (R^{\mu\text{ss}})^{\text{wn}}$.

We consider pull back $\widetilde{\mathcal{U}} \rightarrow X \times \mathcal{Z}$ of the universal sheaf $\mathcal{U} \rightarrow X \times R^{\mu\text{ss}}$ via the map $X \times \mathcal{Z} \rightarrow X \times R^{\mu\text{ss}}$, and the line bundle $\lambda_{\widetilde{\mathcal{U}}}(u_{n-1}(c, [\mathcal{O}_H])) \in \text{Pic}(\mathcal{Z})$, which we continue to denote by \mathcal{L}_{n-1} . In [30, Thm. 3.6] it is shown that this line bundle is equivariantly semiample with respect to the $\text{SL}(V)$ -action on \mathcal{Z} induced by the natural $\text{SL}(V)$ -action on $\text{Quot}(\mathcal{H}, c(E))$. Furthermore, the equivariant section ring of \mathcal{L}_{n-1} is finitely generated in degree 1 (perhaps after passing to a sufficiently large power, see [30, Prop. 4.3]). In other words, a sufficiently large power of \mathcal{L}_{n-1} gives a map from \mathcal{Z} to a projective space, and furthermore the images of these maps stabilise. Formally imitating the

GIT construction we may define the space $M^{\mu ss}$ to be the image, that is:

$$M^{\mu ss} = \text{Proj} \left(\bigoplus_k H^0(\mathcal{Z}, \mathcal{L}_{n-1}^k)^{\text{SL}(V)} \right). \quad (2.16)$$

We will write $\pi : \mathcal{Z} \rightarrow M^{\mu ss}$ for the natural surjective map. We note that by a recent result [7, Appendix A.4], the map π is in fact a quotient map.

The space $M^{\mu ss}$ is by construction a projective scheme. Furthermore, it comes with a distinguished ample line bundle $\mathcal{O}_{M^{\mu ss}}(1)$ and enjoys a subtle universal property (for details, see [30, Sect. 4]). In particular, any flat family $\mathcal{E} \rightarrow X \times S$ of slope semistable sheaves over a weakly normal complex space S yields a classifying morphism $\psi_{\mathcal{E}} : S \rightarrow M^{\mu ss}$, so that $\psi_{\mathcal{E}}^*(\mathcal{O}_{M^{\mu ss}}(1)) = \lambda_{\mathcal{E}}(u_{n-1}(c, [\mathcal{O}_H]))^{\otimes N}$, for some power N . Again, when we refer to $M^{\mu ss}$ in the sequel we will always mean its analytification, which is a weakly normal complex space by [30, Thm. 4.7].

A point in $M^{\mu ss}$ is represented by a quotient sheaf $q_{\mathcal{E}} : \mathcal{H} \rightarrow \mathcal{E} \rightarrow 0$ with $\det \mathcal{E} \simeq \mathcal{J}$. Just as for the Gieseker moduli space, one seeks a characterisation of the points of this space in terms of properties of \mathcal{E} . Clearly, quotients corresponding to isomorphic sheaves \mathcal{E}_1 and \mathcal{E}_2 give rise to the same point in $M^{\mu ss}$ since this means that $q_{\mathcal{E}_1}$ and $q_{\mathcal{E}_2}$ belong to the same $\text{SL}(V)$ orbit, and therefore they cannot be separated by any $\text{SL}(V)$ invariant section of the line bundle \mathcal{L}_{n-1} . Hence, given a μ -semistable sheaf \mathcal{E} with the appropriate Chern classes we may speak unambiguously of the point $\pi(q_{\mathcal{E}}) = [\mathcal{E}] \in M^{\mu ss}$ corresponding to \mathcal{E} . Conversely, every point of $M^{\mu ss}$ is represented by such a sheaf \mathcal{E} . In fact, [30, Thm. 5.10] shows that the weak normalisation of $M_{an}^s \cong M_{alg}^s \cong M^s$ embeds into $M^{\mu ss}$ as a Zariski open set, similar to the case of M^{Gss} . This motivates the following definition.

Definition 2.21. Let \overline{M}^{μ} denote the Zariski closure of $(M^s)^{wn} \subset M^{\mu ss}$. We will call this the *slope compactification*.

It is important to clarify when two sheaves correspond to the same point in $M^{\mu ss}$. We have the following result from [38] ($\dim X = 2$) and [30] (higher dimensions).

Proposition 2.22. *If $\{q_i : \mathcal{H} \rightarrow \mathcal{E}_i\} \in \mathcal{Z}$, $i = 1, 2$ satisfy $[\mathcal{E}_1] = [\mathcal{E}_2] \in M^{\mu ss}$, then $\text{Gr}(\mathcal{E}_1)^{\vee\vee} \cong \text{Gr}(\mathcal{E}_2)^{\vee\vee}$ and $\mathcal{C}_{\mathcal{E}_1} = \mathcal{C}_{\mathcal{E}_2}$, where $q_i : \mathcal{H} \rightarrow \mathcal{E}_i$, $i = 1, 2$. If $\dim X = 2$, the converse also holds.*

The proposition means that for $\dim X = 2$, the points of $M^{\mu ss}$ are in bijection with isomorphism classes of pairs $(\mathcal{E}, \mathcal{C})$ consisting of a polystable reflexive sheaf and a 0-cycle satisfying $\text{ch}_2(\mathcal{E}) = \text{ch}_2(E) + [\mathcal{C}]$. When $\dim X \geq 3$, one needs more information to characterise the geometry of $M^{\mu ss}$. Suppose \mathcal{E}_1

and \mathcal{E}_2 with $q_{\mathcal{E}_1}, q_{\mathcal{E}_2} \in \text{Quot}(\mathcal{H}, c(E))$ satisfy $\text{Gr}(\mathcal{E}_1)^{\vee\vee} \cong \text{Gr}(\mathcal{E}_2)^{\vee\vee} := \mathcal{E}$ and $\mathcal{C}_{\mathcal{E}_1} = \mathcal{C}_{\mathcal{E}_2} = \mathcal{C}$. Consider the Quot scheme of torsion quotients $\text{Quot}(\mathcal{E}, \tau_{\mathcal{E}} - \tau_E)$. Then we have the following *Quot-to-Chow map*:

$$\chi : \text{Quot}(\mathcal{E}, \tau_{\mathcal{E}} - \tau_E) \longrightarrow \mathcal{C}_{n-2}(X) ,$$

given by taking a quotient \mathcal{T} to its cycle $\mathcal{C}_{\mathcal{T}}$. The conditions imposed on \mathcal{E}_1 and \mathcal{E}_2 imply that the sheaves $\mathcal{T}_{\mathcal{E}_1}$ and $\mathcal{T}_{\mathcal{E}_2}$ are in $\text{Quot}(\mathcal{E}, \tau_{\mathcal{E}} - \tau_E)$ and lie in the same fibre of χ .

Proposition 2.23 ([30, Prop. 5.8]). *If $\mathcal{T}_{\mathcal{E}_1}$ and $\mathcal{T}_{\mathcal{E}_2}$ lie in the same connected component of $\chi^{-1}(\mathcal{C})$, then $[\mathcal{E}_1] = [\mathcal{E}_2]$. In particular, for any $q_{\mathcal{E}} \in \text{Quot}(\mathcal{H}, c(E))$ giving a point $[\mathcal{E}]$ of $M^{\mu\text{ss}}$, there are at most finitely many different points $\{[\mathcal{E}_i]\}_{i=1}^k$ of $M^{\mu\text{ss}}$ with $(\text{Gr } \mathcal{E})^{\vee\vee} = (\text{Gr } \mathcal{E}_i)^{\vee\vee}$ and $\mathcal{C}_{\mathcal{E}} = \mathcal{C}_{\mathcal{E}_i}$.*

Remark 2.24. The phenomenon of finite to oneness described above is a genuine one. There are indeed examples where $M^{\mu\text{ss}}$ is a finite set of cardinality larger than one, and yet all representatives have graded objects with the same double dual and cycle (see [29, Example 3.3]). In principle, it is also possible for two sheaves to represent the same point even when their associated torsion sheaves lie in different components of $\chi^{-1}(\mathcal{C})$. The reason that this issue doesn't arise in the case $\dim X = 2$ is that in this case the fibres $\chi^{-1}(\mathcal{C})$ are connected [38], and even irreducible [22].

We finish this section by summarising the relationship between M^{Gss} and $M^{\mu\text{ss}}$. Writing \mathcal{Z}^{Gss} for the space $(R^{\text{Gss}})^{wn}$, the pullback of $\widehat{\mathcal{U}}$ to $X \times \mathcal{Z}^{\text{Gss}}$ under the natural map $\mathcal{Z}^{\text{Gss}} \rightarrow \mathcal{Z}$ induces a classifying map $\mathcal{Z}^{\text{Gss}} \rightarrow M^{\mu\text{ss}}$ which factors through a map $(M^{\text{Gss}})^{wn} \rightarrow M^{\mu\text{ss}}$, as $(M^{\text{Gss}})^{wn}$ is a good quotient of $(\mathcal{Z}^{\text{Gss}})^{wn}$ (see [30, Sect. 5.3]). Composing with the natural map $(\overline{M}^{\text{GM}})^{wn} \rightarrow (M^{\text{Gss}})^{wn}$ induced by the inclusion, we obtain a map

$$\Xi : (\overline{M}^{\text{GM}})^{wn} \longrightarrow \overline{M}^{\mu} . \quad (2.17)$$

On the level of points, this map can be expressed more explicitly in terms of sheaf theory as follows. The space $(M^{\text{Gss}})^{wn}$ is in bijection with the set $\{\text{gr } \mathcal{E} \mid q_{\mathcal{E}} \in \mathcal{Z}^{\text{Gss}}\}$. Being GM-polystable, the sheaf $\text{gr } \mathcal{E}$ is μ -semistable (with the same Chern classes as \mathcal{E}). Therefore, writing the quotient map as $\pi_G : \mathcal{Z}^{\text{Gss}} \rightarrow (M^{\text{Gss}})^{wn}$, we have $\Xi(\pi_G(q_{\mathcal{E}})) = [\text{gr } \mathcal{E}]$.

Together with [30, Thm. 5.10], this shows that Ξ is an isomorphism when restricted to $(M^s)^{wn} \subset (\overline{M}^{\text{GM}})^{wn}$. Since Ξ is in particular continuous, this means it is also surjective, and hence in fact birational, when restricted to the respective compactifications.

3. THE GAUGE THEORETIC COMPACTIFICATION

As we have seen, the moduli spaces M^s and M_{HYM}^* are not compact in general, even after adjoining strictly semistable bundles and reducible connections, respectively. Two compactifications of M^s , \overline{M}^{GM} and \overline{M}^μ , are obtained algebro-geometrically by adding torsion free sheaves, and this was described briefly in Section 2.8. In this section, we describe a third, gauge theoretic compactification associated to M_{HYM}^* . In the case $\dim X = 2$, this construction is due to Donaldson [16, 20] based on work of Uhlenbeck [60] and also Sedlacek [50]. A version of limiting Yang-Mills connections on higher dimensional manifolds appeared in Nakajima [47]. Tian [57] and Tian-Yang [58] generalise the method and obtain key results on the structure of the singular sets. Since the construction will play a central role in this paper, and since some of the details presented here differ from those in the references above, we will provide a complete description of the compactification and some of its important features over the next few subsections.

3.1. Uhlenbeck limits and admissible connections. We begin with a key definition (cf. [3] and [57, Sect. 2.3]).

Definition 3.1. Let (E, h) be a hermitian vector bundle on a Kähler manifold X , not necessarily compact, with $n = \dim X$. Then by an *admissible connection* we mean a pair (A, S) where

- (1) $S \subset X$ is a closed subset of locally finite Hausdorff $(2n - 4)$ -measure;
- (2) A is a smooth integrable unitary connection on $E|_{X \setminus S}$;
- (3) $\int_{X \setminus S} |F_A|^2 d\text{vol}_X < +\infty$;
- (4) $\sup_{X \setminus S} |\Lambda F_A| < +\infty$.

An admissible connection is called *admissible HYM* if there is a constant μ such that $\sqrt{-1}\Lambda F_A = \mu \cdot \mathbf{I}$ on $X \setminus S$.

We will sometimes abuse terminology by saying "A is an admissible connection on (E, h) ", when S is understood.

Remark 3.2. The preceding definition is closely related to, but should not be confused with, the notion of an *admissible Hermitian-Yang-Mills metric* in a given polystable reflexive sheaf \mathcal{E} on X as in [3]. We will elaborate on the relationship between the two notions in Sections 3.3 and 4.1 below.

The fundamental weak compactness result is the following.

Theorem 3.3 (Uhlenbeck, [59]). *In the setup of Definition 3.1, let (A_i, S_i) be a sequence of admissible HYM connections with a uniform bound on the L^2 -norm of curvature. Assume there is a closed set S_∞^l of finite Hausdorff*

$(2n - 4)$ -measure such that $S_i \rightarrow S'_\infty$ on compact sets in the Hausdorff sense. Then there is

- (1) a subsequence (still denoted A_i),
- (2) a closed subset $S_\infty \subset X$ of locally finite $(2n - 4)$ -Hausdorff measure containing S'_∞ ,
- (3) a HYM connection A_∞ on a hermitian bundle $E_\infty \rightarrow X \setminus S_\infty$, and
- (4) local isometries $E_\infty \simeq E$ on compact subsets of $X \setminus S_\infty$

such that with respect to the local isometries, and modulo unitary gauge equivalence, $A_i \rightarrow A_\infty$ in $C_{loc}^\infty(X \setminus S_\infty)$.

We call the limiting connection A_∞ an *Uhlenbeck limit*. The set S_∞ , which we call the *(analytic) singular set*, is the union of S'_∞ with the set

$$\bigcap_{\sigma_0 \geq \sigma > 0} \left\{ x \in X \setminus S'_\infty \mid \liminf_{i \rightarrow \infty} \sigma^{4-2n} \int_{B_\sigma(x)} |F_{A_i}|^2 \frac{\omega^n}{n!} \geq \varepsilon_0 \right\},$$

where σ_0 and ε_0 are universal constants depending only on the geometry of X .

For the definition of a gauge theoretic compactification, it will be important that the Uhlenbeck limits of smooth HYM connections be admissible HYM connections in the sense of Definition 3.1. This will be true if E_∞ is isometric to E on the complement of the singular set. We formulate the precise statement as follows.

Proposition 3.4. *Let $(E, h) \rightarrow X$ be a hermitian vector bundle on a compact Kähler manifold. Then any Uhlenbeck limit of a sequence of smooth HYM connections on (E, h) is an admissible HYM connection. Moreover, the corresponding singular set is a holomorphic subvariety of codimension at least 2.*

The remainder of this section is devoted to the description of the proof of Proposition 3.4, which follows from work of Tian, Bando-Siu, and Tao-Tian. In the following, let

$$S(A_\infty) := \left\{ x \in X \mid \lim_{\sigma \downarrow 0} \sigma^{4-2n} \int_{B_\sigma(x)} |F_{A_\infty}|^2 \frac{\omega^n}{n!} \neq 0 \right\}. \quad (3.1)$$

This is a closed set with zero $(2n - 4)$ -dimensional Hausdorff measure.

Theorem 3.5 (Tian, [57, Thm. 4.3.3]). *Let A_i be a sequence of smooth integrable HYM connections (i.e. $S_i = \emptyset$) converging in the sense of Theorem 3.3 to an Uhlenbeck limit (A_∞, S_∞) . Then S_∞ admits a decomposition into closed sets $S_\infty = S_b \cup \tilde{S}_\infty$, where S_b is a pure codimension 2 holomorphic subvariety, and \tilde{S}_∞ has zero $(2n - 4)$ -dimensional Hausdorff measure.*

The set S_b appearing in the above theorem is called the *blow-up locus* of the sequence A_i . We will also need two results on "removable" singularities for admissible connections.

Theorem 3.6 (Bando–Siu, [3]). *Let \mathcal{E} be a hermitian holomorphic vector bundle on $X \setminus S$, where S has finite Hausdorff $(2n - 4)$ -measure, and suppose the Chern connection of \mathcal{E} is admissible. Then \mathcal{E} extends uniquely as a reflexive sheaf $\widehat{\mathcal{E}}$ on X . If the Chern connection of \mathcal{E} is HYM, the hermitian structure extends smoothly to $X \setminus \text{sing}(\widehat{\mathcal{E}})$, the Chern connection is HYM there, and \mathcal{E} is polystable.*

Theorem 3.7 (Tao–Tian, [56]). *Let (A, S) be an admissible connection on the trivial bundle over a ball $B_{\sigma_0}(x) \subset X$. Suppose $x \notin S(A)$, where $S(A)$ is defined in (3.1). Then for $0 < \sigma < \sigma_0$ sufficiently small there is a unitary gauge transformation g on $B_\sigma(x) \setminus S$, such that $g(A)$ extends to a smooth HYM connection on $B_\sigma(x)$.*

We note the following important fact, which is implicit in [57, Thm. 4.3.3].

Lemma 3.8. *In Theorem 3.5, we may take $\widetilde{S}_\infty = S(A_\infty)$; i.e., the singular set decomposes into closed sets as follows: $S_\infty = S_b \cup S(A_\infty)$.*

Proof. Let $\mu_i = |F_{A_i}|^2 \omega^n / n!$ be the Yang-Mills energy densities. Then we may assume without loss of generality that we have a convergence of finite Radon measures:

$$\mu_i \rightarrow \mu = |F_{A_\infty}|^2 \frac{\omega^n}{n!} + \nu, \quad (3.2)$$

where the measure ν is absolutely continuous with respect to $H_{S_\infty}^{2n-4}$, where $H_{S_\infty}^{2n-4}(A) := H^{2n-4}(S_\infty \cap A)$ is the $(2n - 4)$ -dimensional Hausdorff measure on S_∞ . We will write $\nu(x) = \Theta(x) \cdot \mathcal{H}_{S_\infty}^{2n-4}$, the density function Θ being

$$\Theta(x) := \lim_{\sigma \downarrow 0} \sigma^{4-2n} \mu(B_\sigma(x)),$$

where the limit exists due to the monotonicity formula (see [57, Lemma 3.1.4 (a)]). Then $x \in S_\infty$ if and only if $\Theta(x) \neq 0$ (see [57, p. 222]). On the other hand, if $x \notin S_b$, then for $\sigma > 0$ sufficiently small, $\mathcal{H}^{2n-4}(S_\infty \cap B_\sigma(x)) = 0$, so $\nu(S_\infty \cap B_\sigma(x)) = 0$. If $x \notin S(A_\infty)$, then by (3.1) and (3.2) we must then also have $\Theta(x) = 0$. The result follows. \square

Let \mathcal{E}_∞ be the reflexive sheaf on X extending the holomorphic bundle $(E_\infty, \bar{\partial}_{A_\infty})$, which comes from Theorem 3.6. The following result is a restatement of [58, Thm. 1.4]. For completeness and for the convenience of the reader, we provide a condensed proof of this result.

Lemma 3.9. *The equality $S(A_\infty) = \text{sing } \mathcal{E}_\infty$ holds. In particular, S_∞ is a holomorphic subvariety.*

Proof. Fix $x \in S_\infty$. If $x \notin S(A_\infty)$, then there is $r_0 > 0$ such that $S_\infty \cap B_{r_0}(x) \subset S_b$. As S_b is analytic by Theorem 3.5, $S_\infty \cap B_{r_0}(x)$ has a neighbourhood that admits a deformation retraction to $S_\infty \cap B_{r_0}(x)$. It follows (cf. [63, Ch. 7]) that $E_\infty|_{B_{r_0}(x)}$ is isometric to $E|_{B_{r_0}(x)}$. In particular, A_∞ is identified with an admissible connection on a trivial bundle on $B_{r_0}(x)$. We may therefore apply Theorem 3.7 to conclude that the holomorphic bundle $(E, \bar{\partial}_{A_\infty})$ extends to a holomorphic bundle on $B_r(x)$ for some $r > 0$. Since the reflexive extension \mathcal{E}_∞ is unique, we conclude that $x \notin \text{sing}(\mathcal{E}_\infty)$. Conversely, if $x \notin \text{sing}(\mathcal{E}_\infty)$, then the Chern connection is smooth at x , and it follows that $x \notin S(A_\infty)$. \square

We can now give the

Proof of Proposition 3.4. By Lemma 3.9, S_∞ is a union of irreducible subvarieties, and so there is an exhaustion of $X \setminus S_\infty$ by compact subsets which are deformation retracts of $X \setminus S_\infty$. It follows that in the patching argument for the construction of an Uhlenbeck limit one can find global gauges, and so the bundle E_∞ obtained in Theorem 3.3 is isometric to (E, h) on $X \setminus S_\infty$. We refer to [63] for further details. \square

3.2. Analytic multiplicities. In this section we discuss the multiplicities that are associated to the irreducible components of the blow-up locus. For an admissible connection (A, S) on $(E, h) \rightarrow X$, the result in [57, Prop. 2.3.1] states that integration against the form

$$\text{ch}_2(A) = -\frac{1}{8\pi^2} \text{tr}(F_A \wedge F_A) \quad (3.3)$$

on $X \setminus S$ defines a closed $(2, 2)$ -current on X .

Theorem 3.10 (Tian, [57, Thm. 4.3.3]). *Suppose A_i is a sequence of smooth HYM connections on E , and A_i has Uhlenbeck limit A with blow-up locus S_b . Then to each irreducible codimension 2 component $Z_j^{an} \subset S_b$ there is a positive integer m_j^{an} such that*

$$\text{ch}_2(A_i) \longrightarrow \text{ch}_2(A_\infty) - \sum_j m_j^{an} Z_j^{an} \quad (3.4)$$

in the sense of currents.

We define the $(n - 2)$ -cycle associated to the sequence $\{A_i\}$ by $\mathcal{C}^{an} = \sum_j m_j^{an} Z_j^{an}$, so that $|\mathcal{C}| = S_b$. From (3.4) we see that $\text{ch}_2(A_\infty)$ represents $\text{ch}_2(E) + [\mathcal{C}^{an}]$ in $H^4(X, \mathbb{Q})$. In this context we will also refer to the triple $(A_\infty, \mathcal{C}^{an}, S(A_\infty))$ as an *Uhlenbeck limit* of A_i .

To elaborate on the origin of the integer multiplicities above, we recall two slicing lemmas from [54] which will be needed later on. We begin with the following definition.

Definition 3.11. Let z be a smooth point of a codimension 2 subvariety $Z \subset X$. We say that Σ is a *transverse slice* to Z at z if $\Sigma \cap Z = \{z\}$ and Σ is the restriction of a linear subspace $\mathbb{C}^2 \hookrightarrow \mathbb{C}^n$ to some coordinate ball centred at z that is transverse to $T_z Z$ at the origin.

Lemma 3.12 ([54, eq. (4.1)]). *Let T be a smooth, closed $(2, 2)$ form satisfying the equation $T = m_Z Z + dd^c \Psi$, where Ψ is a $(1, 1)$ -current, smooth away from Z , $m_Z Z$ is the current of integration over the nonsingular points of Z with multiplicity m_Z , and the equation is taken in the sense of distributions. Then for a transverse slice,*

$$m_Z = \int_{\Sigma} T - \int_{\partial \Sigma} d^c \Psi .$$

The next result shows that the analytic multiplicities may also be calculated by restricting to transverse slices.

Lemma 3.13 ([54, Lemma 4.1]). *Let A_i be a sequence of Hermitian-Yang-Mills connections on a fixed hermitian vector bundle $E \rightarrow X$, with Uhlenbeck limit $(A_{\infty}, \mathcal{C}^{an}, S(A_{\infty}))$ and corresponding blow-up locus S_b , and let Z be an irreducible codimension 2 subvariety of X . For a transverse slice Σ at a generic smooth point $z \in Z$, we have:*

$$\lim_{i \rightarrow \infty} \frac{1}{8\pi^2} \int_{\Sigma} \{\mathrm{tr}(F_{A_i} \wedge F_{A_i}) - \mathrm{tr}(F_{A_{\infty}} \wedge F_{A_{\infty}})\} = m_j^{an}$$

if $Z = Z_j^{an} \subset S_b^{an}$, and

$$\lim_{i \rightarrow \infty} \frac{1}{8\pi^2} \int_{\Sigma} \{\mathrm{tr}(F_{A_i} \wedge F_{A_i}) - \mathrm{tr}(F_{A_{\infty}} \wedge F_{A_{\infty}})\} = 0$$

otherwise.

It will be useful to have a more explicit description of the multiplicity. First, we will need the next result, which is an elementary computation that we omit.

Lemma 3.14. *Let $E \rightarrow \Sigma$ be a bundle over a smooth 4-manifold Σ with boundary $\partial \Sigma$. For connections A, B on $E \rightarrow U$, $U \subset \Sigma$ open, write $\mathrm{CS}(A, B)$ for the Chern-Simons 3-form on U satisfying*

$$d\mathrm{CS}(A, B) = \frac{1}{8\pi^2} \mathrm{tr}\{(F_A \wedge F_A) - (F_B \wedge F_B)\} .$$

Then we have the following:

- (1) For A_0, A_1, A_2 be smooth connections on $E \rightarrow U$, there is a 2-form Ω so that on U ,

$$\text{CS}(A_2, A_0) = \text{CS}(A_1, A_0) + \text{CS}(A_2, A_1) + d\Omega.$$

- (2) If A is a connection on $E \rightarrow \Sigma$ and g is a smooth gauge transformation defined in a neighbourhood U of $\partial\Sigma$, then

$$\frac{1}{8\pi^2} \int_{\partial\Sigma} \text{CS}(A, g(A)) = \frac{1}{24\pi^2} \int_{\partial\Sigma} \text{tr}((g^{-1}dg)^3) =: \text{deg}(g) \in \mathbb{Z},$$

where the right hand side is the evaluation on $\partial\Sigma$ of the pullback by g of the Cartan 3-form generating $H^3(\text{SU}(r), \mathbb{Z})$.

- (3) If $E_\infty \rightarrow \Sigma$ is a smooth vector bundle and $\phi_\infty : E|_{\partial\Sigma} \rightarrow E_\infty|_{\partial\Sigma}$ is a smooth isomorphism, and A_∞, B_∞ are connections on E_∞ , then

$$\int_{\partial\Sigma} \text{CS}(\phi_\infty^* A_\infty, \phi_\infty^* B_\infty) = \int_{\partial\Sigma} \text{CS}(A_\infty, B_\infty).$$

In particular, for g as in (2)

$$\int_{\partial\Sigma} \text{CS}(g(A), g(B)) = \int_{\partial\Sigma} \text{CS}(A, B).$$

- (4) If g_1, g_2 are gauge transformations as in (2) then

$$\text{deg}(g_1 g_2) = \text{deg } g_1 + \text{deg } g_2. \quad (3.5)$$

Returning to the situation in Lemma 3.13, let $z \in Z$. Without loss of generality assume $z \notin S(A_\infty)$. Then by Proposition 3.4 and Theorem 3.7, A_∞ locally extends to a connection on a bundle E_∞ that is isometric to E away from Z . Along the slice Σ we choose local unitary frames \mathbf{e} and \mathbf{e}_∞ of E and E_∞ , respectively. Let D_E and D_{E_∞} denote the connections on E and E_∞ that make the frames \mathbf{e} and \mathbf{e}_∞ parallel. We have:

$$\begin{aligned} \frac{1}{8\pi^2} \int_{\Sigma} \text{tr}(F_{A_i} \wedge F_{A_i}) &= \frac{1}{8\pi^2} \int_{\partial\Sigma} \text{CS}(A_i, D_E) \\ \frac{1}{8\pi^2} \int_{\Sigma} \text{tr}(F_{A_\infty} \wedge F_{A_\infty}) &= \frac{1}{8\pi^2} \int_{\partial\Sigma} \text{CS}(A_\infty, D_{E_\infty}) \end{aligned} \quad (3.6)$$

Away from $z \in \Sigma$ there is a isometry $\phi_\infty : E \rightarrow E_\infty$. By Theorem 3.3 there are gauge transformations g_i defined away from $z \in \Sigma$ such that $g_i(A_i) \rightarrow \phi_\infty^* A_\infty$. Furthermore, since $\phi_\infty^* D_\infty$ and D_E are flat connections on E over a simply connected manifold $\Sigma \setminus \{z\}$, there is a gauge transformation h on $\Sigma \setminus \{z\}$ such that $h(\phi_\infty^* D_{E_\infty}) = D_E$. Set $h_i = hg_i$. Now, using (3.6) and Lemma 3.14,

$$\frac{1}{8\pi^2} \int_{\Sigma} \left\{ \text{tr}(F_{A_i} \wedge F_{A_i}) - \text{tr}(F_{A_\infty} \wedge F_{A_\infty}) \right\} = \int_{\partial\Sigma} \text{CS}(A_i, D_E) - \text{CS}(A_\infty, D_{E_\infty})$$

$$\begin{aligned}
&= \int_{\partial\Sigma} \text{CS}(g_i(A_i), g_i(D_E)) - \text{CS}(\phi_\infty^* A_\infty, \phi_\infty^* D_{E_\infty}) \\
&= \int_{\partial\Sigma} \text{CS}(g_i(A_i), \phi_\infty^* A_\infty) + \text{CS}(\phi_\infty^* A_\infty, g_i(D_E)) - \text{CS}(\phi_\infty^* A_\infty, \phi_\infty^* D_{E_\infty}) \\
&= \int_{\partial\Sigma} \text{CS}(g_i(A_i), \phi_\infty^* A_\infty) + \text{CS}(\phi_\infty^* D_{E_\infty}, g_i(D_E)) \\
&= \int_{\partial\Sigma} \text{CS}(g_i(A_i), \phi_\infty^* A_\infty) + \text{CS}(D_E, h_i(D_E)) .
\end{aligned}$$

The first term on the right hand side above vanishes as $i \rightarrow \infty$, whereas by part (2) of Lemma 3.14, $\deg(h_i)$ is an integer which must stabilise to give the multiplicity for i sufficiently large.

3.3. Ideal HYM connections. Observe from Lemma 3.8 that limits $A_i \rightarrow A_\infty$ of smooth HYM connections satisfy the following property: any removable singularity of A_∞ lies in the codimension 2 cycle associated to $\{A_i\}$. This result motivates the next definition, which is slightly more restrictive than the one used in [57], but well-adapted to our purposes.

Definition 3.15. An *ideal HYM connection* is a triple $(A, \mathcal{C}, S(A))$ satisfying the following conditions:

- (1) $\mathcal{C} \in \mathcal{C}_{n-2}(X)$;
- (2) the pair $(A, |\mathcal{C}| \cup S(A))$ is an admissible HYM connection on the hermitian vector bundle $(E, h) \rightarrow X$, where $S(A)$ is given by eq. (3.1);
- (3) $[\text{ch}_2(A)] = \text{ch}_2(E) + [\mathcal{C}]$, in $H^4(X, \mathbb{Q})$;
- (4) A induces the connection a_J on $J = \det E$.

Moreover, we say that ideal connections $(A_1, \mathcal{C}_1, S(A_1))$ and $(A_2, \mathcal{C}_2, S(A_2))$ are *gauge equivalent* if $\mathcal{C}_1 = \mathcal{C}_2$ as cycles (so in particular $|\mathcal{C}_1| = |\mathcal{C}_2| =: Z$), and if there is a smooth unitary gauge transformation g on $X \setminus (Z \cup S(A_1) \cup S(A_2))$ such that $g(A_1) = A_2$.

Given an ideal HYM connection $(A, \mathcal{C}, S(A))$ on (E, h) , by Theorem 3.6 and Lemma 3.9 there is a polystable reflexive extension $\mathcal{E} \rightarrow X$ of the holomorphic bundle $(E, \bar{\partial}_E)$ defined in the complement of $Z \cup S(A)$, with $S(A) = \text{sing}(\mathcal{E})$. Conversely, let $(\mathcal{E}, \mathcal{C})$ be a pair consisting of a polystable reflexive sheaf and a codimension 2 holomorphic cycle so that \mathcal{E} is smoothly isomorphic to E on $X \setminus (Z \cup \text{sing}(\mathcal{E}))$, where we put $Z = |\mathcal{C}|$. By [3, Thm. 3], \mathcal{E} admits an admissible Hermitian-Einstein metric $h_{\mathcal{E}}$ that is unique up to a constant. Let g be a complex gauge transformation on $X \setminus (Z \cup \text{sing}(\mathcal{E}))$ such that $g^*(h_{\mathcal{E}}) = h$. Then a simple calculation shows (see the discussion in Section 4.1 below), that if we write $\bar{\partial}_E$ for the holomorphic structure on E associated to \mathcal{E} , the

Chern connection $A = (\bar{\partial}_{g^*(E)}, h)$ gives an admissible HYM connection $(A, Z \cup \text{sing}(\mathcal{E}))$ on (E, h) in the sense of Definition 3.1. By construction, the sheaf \mathcal{E} is the reflexive extension associated to this admissible connection, and in particular $S(A) = \text{sing}(\mathcal{E})$. The current $\text{ch}_2(A)$ defined in eq. (3.3) is closed and represents $\text{ch}_2(\mathcal{E})$, see the proof of [54, Prop. 3.3]. Hence, if we assume furthermore that $\text{ch}_2(\mathcal{E}) = \text{ch}_2(E) + [\mathcal{C}]$, the triple $(A, \mathcal{C}, S(A))$ is an ideal connection.

By construction, if $(A_1, \mathcal{C}_1, S(A_1))$ and $(A_2, \mathcal{C}_2, S(A_2))$ are gauge equivalent, their associated holomorphic bundles are isomorphic away from the analytic set $Z \cup S(A_1) \cup S(A_2)$, which has codimension ≥ 2 in X . Hence, the respective reflexive extensions \mathcal{E}_1 and \mathcal{E}_2 coming from Theorem 3.6 are isomorphic, and in particular we conclude with the help of Lemma 3.9 that $S(A_1) = S(A_2)$. Conversely, if \mathcal{E}_1 and \mathcal{E}_2 admit admissible HYM metrics in the sense of Bando–Siu, and if $\mathcal{E}_1 \cong \mathcal{E}_2$ and $\mathcal{C}_1 = \mathcal{C}_2$, then $(A_1, \mathcal{C}_1, S(A_1))$ and $(A_2, \mathcal{C}_2, S(A_2))$ are gauge equivalent. We may therefore regard an isomorphism class of ideal connections as equivalent to an isomorphism class of pairs $(\mathcal{E}, \mathcal{C})$, where \mathcal{E} is a polystable reflexive sheaf whose underlying C^∞ vector bundle on the complement of $|\mathcal{C}| \cup \text{sing} \mathcal{E}$ is isomorphic to E . This is compatible with the notation introduced in Definition 2.11. We will use this description in Section 5.

Gauge equivalence defines an equivalence relation \sim on the space of ideal HYM connections on (E, h) . We define the *moduli set of ideal HYM connections* on (E, h) to be

$$\widehat{M}_{\text{HYM}} = \widehat{M}_{\text{HYM}}(E, h, a_J) := \{\text{ideal HYM connections on } (E, h)\} / \sim . \quad (3.7)$$

Notice that there is a natural inclusion $M_{\text{HYM}}^* \subset \widehat{M}_{\text{HYM}}$.

In order to obtain a compactification, we will apply Theorem 3.3 to sequences of ideal HYM connections as well. Given any $[(\mathcal{E}, \mathcal{C})] \in \widehat{M}_{\text{HYM}}$, observe that owing to polystability of \mathcal{E} the Bogomolov inequality, see for example [3, Cor. 3], applies to give:

$$\int_X \left(c_2(\mathcal{E}) - \frac{r-1}{2r} c_1^2(\mathcal{E}) \right) \wedge \omega^{n-2} \geq 0 . \quad (3.8)$$

Using the relations $\text{ch}_2(\mathcal{E}) = \text{ch}_2(E) + [\mathcal{C}]$ and $c_1(\mathcal{E}) = c_1(E) = c_1(J)$, we obtain a bound

$$\text{deg}(\mathcal{C}) = \sum_i n_i \int_{Z_i} \frac{\omega^{n-2}}{(n-2)!} \leq C$$

that is independent of $(\mathcal{E}, \mathcal{C})$. This means that if $[(A_i, \mathcal{C}_i, S(A_i))]$ is sequence in \widehat{M}_{HYM} , there is a uniform bound $\text{deg} \mathcal{C}_i \leq C$, and so \mathcal{C}_i converges subsequentially as cycles by Theorem 2.7.

Note that since X is compact the sets $S(A_i)$ always converge subsequentially in the Hausdorff sense to some compact subset S'_∞ of X . On the other hand, to maintain the structure of an ideal connection we want the $S(A_i)$ to converge in the union of cycle spaces $\cup_{k=0}^{n-3} \mathcal{C}_k(X)$. This is guaranteed by the next lemma, which follows⁴ from [42, Main Theorem] together with Corollary 2.18.

Lemma 3.16. *There is a constant K with the following significance. For any ideal HYM connection $(A, \mathcal{C}, S(A))$ on (E, h) and each $k = 0, 1, \dots, n-3$, the k -dimensional stratum of $S(A)$ has degree less than or equal to K , when considered as a cycle in X .*

Proof. A consequence of the statement in [42] is that a set of isomorphism classes of slope semistable reflexive sheaves of fixed rank, c_1 , and $c_2 \cup [\omega]^{n-2}$, is bounded. For sheaves \mathcal{E} associated to points in \widehat{M}_{HYM} as above, the rank and c_1 are fixed, whereas $c_2(\mathcal{E}) \cup [\omega]^{n-2}$ is bounded: from below by (3.8), and from above by $c_2(E) \cup [\omega]^{n-2}$. The aforementioned result therefore applies to our family of reflexive sheaves. We may then use Corollary 2.18 applied to the sets $\text{sing}(\mathcal{E}) = \cup_{k=3}^n \text{sing}_{n-k}(\mathcal{E})$. Since these algebraic singularity sets coincide with the analytic singularity sets $S(A)$ by Lemma 3.9, the claim follows. \square

The following is the main result of this subsection.

Theorem 3.17. *Let $(A_i, \mathcal{C}_i, S(A_i)) \in \widehat{M}_{\text{HYM}}$. Then there is a subsequence (also denoted by $\{i\}$), and an ideal HYM connection $(A_\infty, \mathcal{C}_\infty, S(A_\infty))$ such that \mathcal{C}_i converges to a subcycle of \mathcal{C}_∞ , and (up to gauge transformations) $A_i \rightarrow A_\infty$ in C_{loc}^∞ on $X \setminus (Z_\infty \cup S(A_\infty))$ where $Z_\infty := |\mathcal{C}_\infty|$. Moreover,*

$$\text{ch}_2(A_i) - \mathcal{C}_i \longrightarrow \text{ch}_2(A_\infty) - \mathcal{C}_\infty \quad (3.9)$$

in the mass norm; in particular, also in the sense of currents.

The rest of this section is devoted to the proof of Theorem 3.17. We proceed in several steps.

Step 1. By Lemma 3.16 and Theorem 2.7, we may first extract a subsequence (also denoted $\{i\}$), such that

- (1) the Yang-Mills densities converge weakly $|F_{A_i}|^2 d\text{vol}_\omega \rightarrow \mu_\infty$ to a Radon measure μ_∞ ;
- (2) there is a cycle \mathcal{C}'_∞ with $|\mathcal{C}'_\infty| = Z'_\infty$ such that $\mathcal{C}_i \rightarrow \mathcal{C}'_\infty$ as cycles;
- (3) we have Hausdorff convergence $S(A_i) \rightarrow S'_\infty$, where S'_∞ is a subvariety of codimension at least 3.

⁴We are grateful to Carlos Simpson for suggesting that Lemma 3.16 should be a consequence of Maruyama's result.

Step 2. We may apply Theorem 3.3 to see that the following holds: There is a closed subset $\tilde{S}_\infty \subset X \setminus (Z'_\infty \cup S'_\infty)$ of locally finite Hausdorff $(2n-4)$ -measure, and a smooth HYM connection A_∞ on $X \setminus (Z'_\infty \cup S'_\infty \cup \tilde{S}_\infty)$, such that $A_i \rightarrow A_\infty$ up to gauge, smoothly in the local C^∞ topology. Moreover, by Theorem 3.5, we may write $\tilde{S}_\infty = \tilde{Z}_\infty \cup S(A_\infty)$, where $S(A_\infty)$ is defined as in (3.1), and \tilde{Z}_∞ is a closed pure $(n-2)$ -dimensional analytic subvariety of $X \setminus (Z'_\infty \cup S'_\infty)$.

Step 3. We claim that $\tilde{Z}_\infty \subset X \setminus (Z'_\infty \cup S'_\infty)$ extends to X . We shall use the Bishop-Stoll removable singularities theorem. Choose a point $p \in Z'_\infty$, $\sigma > 0$, such that $B_{2\sigma}(p) \subset X \setminus (S'_\infty \cup S(A_\infty))$. Then for i sufficiently large, the connections A_i are smooth on $B_{2\sigma}(p) \setminus Z_i$ and extend to smooth connections \hat{A}_i on $B_\sigma(p)$ by Theorem 3.7. Hence, applying Theorem 3.3 we conclude that after passing to a subsequence and up to gauge \hat{A}_i converges to a limit \hat{A}_∞ in the local C^∞ topology on $B_\sigma(p)$ away from a singular set $\hat{Z}_\infty \cup S(\hat{A}_\infty)$ that is an analytic subvariety of $B_\sigma(p)$. Since \hat{A}_∞ and A_∞ agree up to gauge off a codimension 2 set, they agree up to gauge on their common domain of definition. By our choice of p it follows that $S(\hat{A}_\infty) \cap B_\sigma(p) = \emptyset$, and $\hat{Z}_\infty \cap B_\sigma(p)$ and $\tilde{Z}_\infty \cap B_\sigma(p)$ agree on the complement of $Z'_\infty \cap B_\sigma(p)$. In particular, the intersection $\text{cl}(\tilde{Z}_\infty) \cap Z'_\infty \cap B_\sigma(p)$ has codimension at least 3. Since $S'_\infty \cup S(A_\infty)$ also has codimension at least 3, the same is true for $\text{cl}(\tilde{Z}_\infty) \cap (Z'_\infty \cup S'_\infty)$. It now follows from [6, Lemma 9], that \tilde{Z}_∞ extends as a holomorphic $(n-2)$ -dimensional subvariety Z''_∞ on X .

Step 4. Set $Z_\infty = Z'_\infty \cup Z''_\infty$ as pure $(n-2)$ -dimensional subvarieties. We have local C^∞ convergence $A_i \rightarrow A_\infty$ on $X \setminus (Z_\infty \cup S(A_\infty) \cup S'_\infty)$. Hence, by the same argument as in the proof of Proposition 3.4, A_∞ is an admissible HYM connection. To prove that we have an ideal HYM connection, we need to show that $S'_\infty \subset Z_\infty \cup S(A_\infty)$, and we have to assign multiplicities to the components of Z_∞ . The latter part will be discussed in Step 5 below. The former statement is a consequence of the next lemma, which is also absolutely crucial for the argument in Section 5.

Lemma 3.18. *In the situation above, $S'_\infty \subset S(A_\infty) \cup Z_\infty$.*

Proof. Consider the reflexive sheaves $\mathcal{E}_\infty, \mathcal{E}_i$ obtained by extending the holomorphic bundles $(E, \bar{\partial}_{A_\infty})$ and $(E, \bar{\partial}_{A_i})$, respectively. By Lemma 3.9, the loci where the \mathcal{E}_i are not locally free are precisely the codimension at least 3 singular sets $S(A_i)$.

Suppose that $p \in S'_\infty$, $p \notin S(A_\infty) \cup Z_\infty$. Then we can find $x_i \in S(A_i)$ converging to p . We may also find a coordinate ball $B_{2\sigma}(p)$ whose closure lies in the complement of $S(A_\infty) \cup Z_\infty$, and such that $(\mathcal{E}_\infty)|_{B_{2\sigma}(p)}$ is a holomorphically

trivial vector bundle. Notice that by definition of Z_∞ , $Z_i \cap B_{2\sigma} = \emptyset$ for sufficiently large i . Let $\Delta_\delta^{(k)}$ denote the polydisk of radius (δ, \dots, δ) in \mathbb{C}^k and let $B_\sigma^{(k)}(0)$ the ball of radius σ in \mathbb{C}^k . We may find an annular region $U := U(\sigma, \varepsilon; \delta) \subset B_{2\sigma}(p)$ centred about p given in coordinates by

$$(B_\sigma^{(3)}(0) \setminus \overline{B_\varepsilon^{(3)}(0)}) \times \Delta_\delta^{(n-3)},$$

such that $S'_\infty \cap U = \emptyset$, and hence also $S(A_i) \cap U = \emptyset$, for sufficiently large i . For the purposes of this proof, we choose the euclidean product metric on U , and the standard hermitian structure on \mathcal{E}_∞ with respect to its trivialisation over U , scaled by an appropriate weight as explained in [51], whose results we will apply later on in the proof. The notation $\|\cdot\|$ will refer to the L^2 -norm of bundle valued forms on U with respect to these metrics, and $*$ will denote the associated hermitian conjugate.

Change notation slightly, and let $\bar{\partial}_{A_i}$ denote the $\bar{\partial}$ -operator on $\text{End } E$ inducing the holomorphic structure $\mathcal{E}_\infty^\vee \otimes \mathcal{E}_i$. Writing $\bar{\partial}_{A_i} = \bar{\partial}_{A_\infty} + a_i$ for some $\text{End } E$ -valued $(0, 1)$ -form a_i , we have $a_i \rightarrow 0$ smoothly on U . We claim that for i sufficiently large there are smooth endomorphisms u_i of E and a constant C such that

- (1) $\bar{\partial}_{A_i} u_i = -a_i$;
- (2) $\|u_i\| \leq C \|a_i\|$.

In order to prove this we verify the basic estimate; namely, that there exists a constant C such that

$$\|\phi\| \leq C (\|\bar{\partial}_{A_i} \phi\| + \|\bar{\partial}_{A_i}^* \phi\|) \quad (3.10)$$

for all ϕ , a smooth $(0, 1)$ -form with values in $\text{End } E$ satisfying the $\bar{\partial}$ -Neumann boundary conditions (cf. [32, 23]). We first note that (3.10) is valid if A_i is replaced by A_∞ :

$$\|\phi\| \leq C (\|\bar{\partial}_{A_\infty} \phi\| + \|\bar{\partial}_{A_\infty}^* \phi\|) . \quad (3.11)$$

Indeed, since the bundle and metrics are trivialized, the problem reduces to the scalar $\bar{\partial}$ -estimate, and hence (3.11) follows using standard results such as [9, Lemma 2.1] from [52, Thm. 2.2] in the case $\dim X = 3$ and [9, Thm. 1.1, Cor. 6.3] in higher dimensions⁵. On the other hand, there is some numerical constant c_0 independent of A_i such that

$$\begin{aligned} \|\bar{\partial}_{A_i} \phi\| &\geq \|\bar{\partial}_{A_\infty} \phi\| - c_0 \sup |a_i| \|\phi\| , \\ \|\bar{\partial}_{A_i}^* \phi\| &\geq \|\bar{\partial}_{A_\infty}^* \phi\| - c_0 \sup |a_i| \|\phi\| . \end{aligned}$$

Since $\sup |a_i|$ is arbitrarily small on U for sufficiently large i , the estimate (3.11) for A_∞ can be parlayed into one for A_i . This proves (3.10).

⁵The authors thank Mei-Chi Shaw for pointing us towards these references.

Given (3.10), it follows as in [51, Lemma 3.2] that we can find u_i satisfying (1) and (2). Let $\varphi_i = \text{Id}_E + u_i$. Then by (1),

$$\bar{\partial}_{A_i}\varphi_i = \bar{\partial}_{A_i}\text{Id}_E + \bar{\partial}_{A_i}u_i = (\bar{\partial}_{A_\infty} + a_i)\text{Id}_E - a_i = 0 .$$

We thus have produced a holomorphic map $\varphi_i : \mathcal{E}_\infty \rightarrow \mathcal{E}_i$ on U . Using interior elliptic estimates for $\bar{\partial}$, along with (2), it follows that for sufficiently large i , $\sup |u_i|$ is arbitrarily small on a subannular region, so that in particular $\det \varphi_i \neq 0$ on $U' := U(\sigma/2, \varepsilon'; \delta/2)$, say, for all sufficiently large i . Since \mathcal{E}_∞ is holomorphically trivial and since \mathcal{E}_i is reflexive, we may use a realisation of \mathcal{E}_i on the interior $\widehat{U} := B_\sigma^3(0) \times \Delta_\delta^{n-3}$ of U of the form stated in [34, Chap. V, Prop. 4.13] as well as a classical Hartogs-type theorem for holomorphic functions on annular regions to see that φ_i extends as a map of sheaves to \widehat{U} .

Let D_i be the zero locus of $\det \varphi_i$ in \widehat{U} . Notice that $x_i \in D_i$, for if not, \mathcal{E}_∞ and \mathcal{E}_i would be isomorphic in a neighbourhood of x_i by reflexivity, and hence \mathcal{E}_i would be locally free at x_i , thus contradicting the assumption that $x_i \in S(A_i)$. Then, all the divisors D_i have to intersect the annular region U' , in particular those with index i sufficiently large. This contradicts our choice of the set U' and completes the proof. \square

Step 5. Finally, we explain how to assign multiplicities to each irreducible component of Z_∞ to obtain the cycle \mathcal{C}_∞ . Since Z''_∞ is the extension of \widetilde{Z} , the irreducible components of Z''_∞ have assigned multiplicities from Theorem 3.10, giving a cycle \mathcal{C}''_∞ . Since we have chosen $\mathcal{C}_j \rightarrow \mathcal{C}'_\infty$ as cycles, each component $Z \subset Z'_\infty$ carries a multiplicity m'_Z . It may occur that there is additional bubbling along Z . Let $z \in Z$ be a point such that $B_{2\sigma}(z) \subset X \setminus S(A_\infty)$ intersects Z_∞ only in the smooth locus of Z , and let \widehat{A}_j be the extended connections from Step 3. Let Ω be smooth $(n-4)$ -form compactly supported on $B_\sigma(z)$. Then there is an integer m''_Z such that

$$\lim_{j \rightarrow \infty} \int_{B_\sigma(z)} \left(\text{ch}_2(A_\infty) - \text{ch}_2(\widehat{A}_j) \right) \wedge \Omega = m''_Z \cdot \int_{Z \cap B_\sigma(z)} \Omega . \quad (3.12)$$

It follows that we should assign the multiplicity of the component Z to be $m_Z = m'_Z + m''_Z$, and we write \mathcal{C}'_∞ for the cycle whose summands are the irreducible components of Z'_∞ with multiplicities defined in this way. Finally, we are in a position to prove the following

Lemma 3.19. *Define \mathcal{C}_∞ to be the cycle $\mathcal{C}'_\infty + \mathcal{C}''_\infty$. Then with this definition eq. (3.9) holds.*

Proof. Choose a smooth $(2n - 4)$ -form Ω , and fix $\varepsilon > 0$. First, we choose $r_0 > 0$ such that

$$0 \leq \int_{\mathcal{N}_{r_0}(S'_\infty \cup S(A_\infty))} \mathcal{C}_\infty \wedge \frac{\omega^{m-2}}{(m-2)!} \leq \varepsilon/2 .$$

Since the \mathcal{C}_j converge to a subcycle of \mathcal{C}_∞ , for j sufficiently large it follows that

$$0 \leq \int_{\mathcal{N}_{r_0}(S'_\infty \cup S(A_\infty))} \mathcal{C}_j \wedge \frac{\omega^{m-2}}{(m-2)!} \leq \varepsilon . \quad (3.13)$$

Next, since $S'_\infty \cup S(A_\infty)$ is a subvariety of codimension at least 3, in particular it has zero Hausdorff $(2n - 4)$ -measure. Hence, we may find finitely many $x_i \in S'_\infty \cup S(A_\infty)$, $i = 1, \dots, M$, and $0 < r_i \leq r_0/2$, such that

$$S'_\infty \cup S(A_\infty) \subset \bigcup_{i=1}^M B_{r_i}(x_i) =: U_1 ; \quad (3.14)$$

$$\sum_{i=1}^M r_i^{2n-4} \leq \varepsilon \cdot 2^{4-2n} . \quad (3.15)$$

Set $U_2 = \cup_{i=1}^M B_{2r_i}(x_i)$, and $r = \min\{r_1, \dots, r_M\}$. Note that if $y \notin U_2$, then $B_r(y) \cap U_1 = \emptyset$. Find finitely many $y_j \in \overline{U_2}^c \cap Z'_\infty$, $j = 1, \dots, N$, and $0 < s_j \leq r$, such that

$$U_2^c \cap Z'_\infty \subset \bigcup_{j=1}^N B_{s_j}(y_j) =: V .$$

Taking a partition of unity subordinate to the cover

$$\{B_{2r_i}(x_i), B_{s_j}(y_j), \mathcal{N}_r((\overline{U_2} \cup \overline{V})^c)\} ,$$

it suffices to consider the cases where Ω is compactly supported in each of the elements of the cover. From the monotonicity formula referred to previously, there is a constant Λ independent of i , such that

$$\int_{B_\sigma(x)} |F_{A_i}|^2 \frac{\omega^m}{m!} \leq \Lambda \cdot \sigma^{2n-4} ,$$

for $0 < \sigma \leq \sigma_0$. It follows that there is a constant Λ_1 such that

$$\left| \int_{B_\sigma(x)} \text{ch}_2(A_i) \wedge \Omega \right| \leq \Lambda_1 \sup |\Omega| \cdot \sigma^{2n-4} . \quad (3.16)$$

The constant σ_0 only depends on the geometry of X , and so we may assume without loss of generality that $r_0 \leq \sigma_0$. By Fatou's lemma, (3.16) also holds

for A_∞ . It follows from (3.16) that

$$\begin{aligned} \left| \int_{U_2} \text{ch}_2(A_i) \wedge \Omega \right| &\leq \sum_{i=1}^M \left| \int_{B_{2r_i}(x_i)} \text{ch}_2(A_i) \wedge \Omega \right| \leq 2^{2n-4} \Lambda_1 \sup |\Omega| \sum_{i=1}^M r_i^{2n-4} \\ &\leq \Lambda_1 \sup |\Omega| \cdot \varepsilon . \end{aligned}$$

The same holds for A_∞ in place of A_i . Since $U_2 \subset \mathcal{N}_{r_0}(S'_\infty \cup S(A_\infty))$, and using (3.13), we have for sufficiently large i ,

$$\left| \int_{U_2} (\text{ch}_2(A_i) - \mathcal{C}_i) \wedge \Omega - \int_{U_2} (\text{ch}_2(A_\infty) - \mathcal{C}_\infty) \wedge \Omega \right| \leq 2\varepsilon \cdot \sup |\Omega| (\Lambda_1 + 1) . \quad (3.17)$$

Next, suppose Ω is compactly supported in $B_{s_j}(y_j)$. For i sufficiently large we have $S(A_i) \subset U_1$, and hence $S(A_i) \cap B_{s_j}(y_j) = \emptyset$. As in Step 3, we have connections \widehat{A}_i on $B_{s_j}(y_j)$ that converge (up to gauge) to A_∞ smoothly on compact subsets in the complement of $Z_\infty \cap B_{s_j}(y_j)$. By definition of the excess multiplicity in (3.12), we have

$$\lim_{i \rightarrow \infty} \int_{B_{s_j}(y_j)} (\text{ch}_2(A_\infty) - \text{ch}_2(\widehat{A}_i) + \mathcal{C}_i) \wedge \Omega = \int_{B_{s_j}(y_j)} \mathcal{C}'_\infty \wedge \Omega .$$

Finally, on $\mathcal{N}_r(\overline{U}_2 \cap V)^c$, the A_i are smooth HYM connections and Theorem 3.10 applies directly to show that

$$\lim_{i \rightarrow \infty} \int_{\mathcal{N}_r(\overline{U}_2 \cap V)^c} (\text{ch}_2(A_\infty) - \text{ch}_2(A_i)) \wedge \Omega = \int_{\mathcal{N}_r(\overline{U}_2 \cap V)^c} \mathcal{C}''_\infty \wedge \Omega .$$

Since ε was arbitrary, the lemma follows. \square

3.4. Diagonalisation. The goal of this section is to prove the crucial result.

Proposition 3.20. *Suppose that we have a bounded sequence of ideal HYM connections that converges $(A_i, \mathcal{C}_i, S(A_i)) \rightarrow (A_\infty, \mathcal{C}_\infty, S(A_\infty))$ as in the previous section. If each $(A_i, \mathcal{C}_i, S(A_i))$ is an Uhlenbeck limit of smooth HYM connections on $(E, h) \rightarrow X$, then so is $(A_\infty, \mathcal{C}_\infty, S(A_\infty))$.*

Proof. We have sequence of gauge transformations s_i such that $s_i(A_i) \rightarrow A_\infty$ in C_{loc}^∞ away from $Z_\infty \cup S(A_\infty)$. For each i , let $\{A_{i,k}\}$ be a sequence of smooth HYM connections and gauge transformations $g_{i,k}$ on $X \setminus Z_i \cup S(A_i)$ and Uhlenbeck limits $g_{i,k}(A_{i,k}) \rightarrow A_i$ as $k \rightarrow \infty$. By a diagonalisation argument, and using the fact that $Z_\infty \cup S(A_\infty)$ is a holomorphic subvariety, we can find a subsequence $k_i \rightarrow \infty$ such that $s_i g_i(A_{i,k_i}) \rightarrow A_\infty$ in C_{loc}^∞ on $X \setminus Z_\infty \cup S(A_\infty)$, where $g_i = g_{i,k_i}$. Applying Theorem 3.3, after passing to a further subsequence which we continue to denote by i , we may assume there is an admissible connection B_∞ on $(E, h) \rightarrow X$ and gauge transformations \tilde{g}_i such that $\tilde{g}_i(A_{i,k_i}) \rightarrow B_\infty$

up to gauge in C_{loc}^∞ away from a singular set $\tilde{Z}_\infty \cup S(B_\infty)$. By Theorem 3.6, the holomorphic bundles $(E, \bar{\partial}_{A_\infty})$ and $(E, \bar{\partial}_{B_\infty})$ extend as reflexive sheaves \mathcal{E}_{A_∞} and \mathcal{E}_{B_∞} on X . Since $(Z_\infty, S(A_\infty))$ and $(Z_\infty, S(B_\infty))$ have codimension at least 2, and $\mathcal{E}_{A_\infty} \cong \mathcal{E}_{B_\infty}$ on the complement, then in fact $\mathcal{E} \cong \mathcal{E}_{A_\infty} \cong \mathcal{E}_{B_\infty}$ everywhere on X . Furthermore, by uniqueness of the Bando-Siu HYM connection, $A_\infty = B_\infty$ up to gauge, and so $S(A_\infty) = S(B_\infty)$. By further modifying \tilde{g}_i , we may assume $A_\infty = B_\infty$.

It remains to consider the blow-up loci Z_∞ and \tilde{Z}_∞ and multiplicities. Since up to gauge, $A_{i,k_i} \rightarrow A_\infty$ on $X \setminus Z_\infty \cup S(A_\infty)$, it follows the blow-locus for A_{i,k_i} is contained in Z_∞ . This means that if \tilde{Z}_j is an irreducible component of \tilde{Z}_∞ , then \tilde{Z}_j must be equal to some component $Z_j \subset Z_\infty$. For simplicity, call this component Z .

Then to prove equality of sets we only need to know that Z_∞ and \tilde{Z}_∞ have the same number of irreducible components. This will follow immediately from the fact that the cohomology classes of $[\mathcal{C}_\infty]$ and $[\tilde{\mathcal{C}}_\infty]$ agree (and are equal to $[\text{ch}_2(A_\infty)] - \text{ch}_2(E) = [\text{ch}_2(B_\infty)] - \text{ch}_2(E)$), if we show that the multiplicity associated to Z in the cycle \mathcal{C}_∞ agrees with its the multiplicity in the cycle $\tilde{\mathcal{C}}_\infty$.

In other words, to prove that \mathcal{C}_∞ is equal to the cycle $\tilde{\mathcal{C}}_\infty$ associated to \tilde{Z}_∞ , it suffices to prove that $\tilde{m}_Z = m_Z$. Let Σ be a generic slice to Z . Then by Lemma 3.13,

$$\tilde{m}_Z = \lim_{i \rightarrow \infty} \frac{1}{8\pi^2} \int_{\Sigma} \left(\text{tr}(F_{A_{i,k_i}} \wedge F_{A_{i,k_i}}) - \text{tr}(F_{A_\infty} \wedge F_{A_\infty}) \right).$$

By the discussion following Lemma 3.14, we have $\tilde{m}_Z = \deg \tilde{h}_i$, where $\tilde{h}_i = h \circ \tilde{g}_i$, and h is the gauge transformation on $\Sigma \setminus \{z\}$ defined there. Write $a_i = A_\infty - \tilde{g}_i A_{i,k_i}$ and $b_i = A_\infty - s_i g_i A_{i,k_i}$ so that $s_i g_i A_{i,k_i} - \tilde{g}_i A_{i,k_i} = a_i - b_i$. Then by Lemma 3.14 we have

$$\begin{aligned} \deg \tilde{g}_i - \deg(s_i g_i) &= \int_{\partial\Sigma} \text{CS}(s_i g_i(A_{i,k_i}), A_{i,k_i}) - \text{CS}(\tilde{g}_i(A_{i,k_i}), A_{i,k_i}) \\ &= \int_{\partial\Sigma} \text{CS}(s_i g_i(A_{i,k_i}), \tilde{g}_i(A_{i,k_i})) \\ &= \frac{1}{8\pi^2} \int_{\partial\Sigma} \text{tr} \left\{ (a_i - b_i) \wedge D_{\tilde{g}_i(A_{i,k_i})}(a_i - b_i) \right. \\ &\quad \left. + \frac{2}{3}(a_i - b_i)^3 + 2(a_i - b_i) \wedge F_{\tilde{g}_i(A_{i,k_i})} \right\}. \end{aligned}$$

Now $s_i g_i(A_{i,k_i})$ and $\tilde{g}_i(A_{i,k_i})$ converge smoothly on $\partial\Sigma$ to the same connection A_∞ , so we have $a_i, b_i \rightarrow 0$, and therefore the right hand side above converges to zero. Hence, for sufficiently large i , $|\deg s_i g_i - \deg \tilde{g}_i| < 1$. Since

the degree is an integer this implies that for all such i we have $\deg(s_i g_i) = \deg \tilde{g}_i$ and therefore $\deg \tilde{h}_i = \deg h s_i g_i$.

We claim that $\deg(h s_i g_i) = m_Z$. Recall that \mathcal{C}_∞ breaks up into two sub-cycles \mathcal{C}'_∞ and \mathcal{C}''_∞ , and so there are two cases according to whether Z is an irreducible component of $Z'_\infty = |\mathcal{C}'_\infty|$ or $Z''_\infty = |\mathcal{C}''_\infty|$. In the second case the connections A_i bubble at the point $z \in Z$ through which the slice Σ is taken. Since Σ may be chosen to be away from Z'_∞ ; and the supports $Z_i = |\mathcal{C}_i|$ are contained in a neighbourhood of Z'_∞ for large enough i , then since Z_i is the bubbling set for the sequence of connection $A_{i,k}$ converge without bubbling to A_i on Σ as $k \rightarrow \infty$. This furthermore implies that if we write E_i for the topological bundle underlying the reflexive extension \mathcal{E}_i , there are isomorphisms $E|_\Sigma \simeq E_i|_\Sigma$. By the argument following Lemma 3.14, we see that $\deg g_i$ calculates the bubbling multiplicity for $A_{i,k}$ along Z for sufficiently large i and k , and this number is zero by construction, so $\deg g_i = 0$ and $\deg(h s_i g_i) = \deg(h s_i)$. Again by the discussion after Lemma 3.14, $\deg(h s_i) = m_Z$.

Now consider the case that Z is an irreducible component of Z'_∞ . In this case $m_Z = m'_Z + m''_Z$ where m'_Z is the multiplicity assigned Z by the convergence of the cycles \mathcal{C}_i , and m''_Z is the excess multiplicity defined in Step 5 eq. (3.12). Notice that by Lemma 3.14 we have $\deg(h s_i g_i) = \deg(h g_i) + \deg s_i$. We claim that $\deg(h g_i) = m'_Z$ and $\deg s_i = m''_Z$. The second equality follows again from the argument after Lemma 3.14 (here we choose a single frame for the trivial bundle on Σ , so no additional gauge transformation h is necessary). To prove the second equality, we note that by assumption, for each i sufficiently large Z_i has an irreducible component $Z_{i,l}$ converging to Z in the Hausdorff topology as $i \rightarrow \infty$. Assume we have chosen Σ so that it is also a transverse slice to all $Z_{i,l}$ through points z_i for i sufficiently large. Since $A_{i,k}$ bubbles along $Z_{i,l}$ with Uhlenbeck limit A_i , there are gauge transformations h_i of E defined on $\Sigma \setminus \{z_i\}$ and isomorphisms $\phi_i : E|_{\Sigma \setminus \{z_i\}} \rightarrow E_i|_{\Sigma \setminus \{z_i\}}$ so that $h_i \phi_i^* D_{E_i} = D_E$ (using the same notation as in the argument following Lemma 3.14), and such that $\deg h_i$ calculates the associated multiplicity. Then for sufficiently large i , $\deg h_i g_i = m'_Z$. We claim that $\deg h_i = \deg h$ for i sufficiently large. We have

$$\begin{aligned} \deg h_i - \deg h &= \int_{\partial \Sigma} \text{CS}(h_i \phi_i^* D_{E_i}, \phi_\infty^* D_{E_\infty}) - \text{CS}(h \phi_\infty^* D_{E_\infty}, \phi_\infty^* D_{E_\infty}) \\ &= \int_{\partial \Sigma} \text{CS}(D_E, \phi_i^* D_{E_i}) - \text{CS}(D_E, \phi_\infty^* D_{E_\infty}) \\ &\quad + \int_{\partial \Sigma} \text{CS}(\phi_i^* D_{E_i}, \phi_\infty^* D_{E_\infty}). \end{aligned}$$

Since $A_i \rightarrow A_\infty$ on $\partial \Sigma$, $\phi_i^* D_{E_i} \rightarrow \phi_\infty^* D_{E_\infty}$ on $\partial \Sigma$ and therefore for large i this difference is zero. This implies $\deg(h g_i) = m'_Z$. \square

3.5. The compactification. We are now in the position to define the desired compactification of the moduli set M_{HYM}^* .

Definition 3.21. We say that a sequence $[(A_i, \mathcal{C}_i, S(A_i))] \in \widehat{M}_{\text{HYM}}$ converges to $[(A_\infty, \mathcal{C}_\infty, S(A_\infty))]$ in \widehat{M}_{HYM} if the conclusion of Theorem 3.17 holds.

We define a topology on \widehat{M}_{HYM} by giving a basis $\{\mathcal{U}_{\vec{\varepsilon}}([A, \mathcal{C}, S(A)])\}$ of open neighbourhoods of an ideal HYM connection depending on a 4-tuple $\vec{\varepsilon} = (\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4)$, $\varepsilon_i > 0$. The set $\mathcal{U}_{\vec{\varepsilon}}([A, \mathcal{C}, S(A)])$ consists of gauge equivalence classes of all ideal HYM connections $(A_1, \mathcal{C}_1, S(A_1))$ satisfying:

- (1) There exists a subcycle $\widehat{\mathcal{C}}_1$ of \mathcal{C} such that $\widehat{\mathcal{C}}_1$ and \mathcal{C}_1 are within ε_1 with respect to the mass norm.
- (2) $S(A_1)$ is contained in the ε_2 -neighborhood of $S(A) \cup |\mathcal{C}| \subset X$;
- (3) On the complement of a $2\varepsilon_2$ -neighbourhood of $S(A) \cup |\mathcal{C}| \subset X$, there is a gauge transformation g such that $g(A_1)$ and A are within ε_3 with respect to the C^∞ -topology;
- (4) the currents $\text{ch}_2(A_1) - \mathcal{C}_1$ and $\text{ch}_2(A) - \mathcal{C}$ are within ε_4 with respect to the mass norm.

Remark 3.22. By definition, a sequence $[(A_i, \mathcal{C}_i, S(A_i))]$ converges to an ideal connection $[(A_\infty, \mathcal{C}_\infty, S(A_\infty))]$ in this topology if and only if it converges in the sense of Definition 3.21. Namely a sequence converges the sense of Definition 3.21 to an ideal connection $[(A_\infty, \mathcal{C}_\infty, S(A_\infty))]$ if and only if for every basic open set $\mathcal{U}_{\vec{\varepsilon}}([A_\infty, \mathcal{C}_\infty, S(A_\infty)])$, there exists $I \gg 0$ such that $[(A_i, \mathcal{C}_i, S(A_i))] \in \mathcal{U}_{\vec{\varepsilon}}([A_\infty, \mathcal{C}_\infty, S(A_\infty)])$ for all $i > I$.

First of all, we have the following.

Theorem 3.23. *The set \widehat{M}_{HYM} defined in (3.7) and endowed with the topology described above is a first countable Hausdorff space and moreover sequentially compact.*

Proof. The first statement follows directly from the definition. Compactness follows from Theorem 3.17: Indeed, any sequence of ideal connections has a subsequence converging to an ideal HYM connection $(A_\infty, \mathcal{C}_\infty, S(A_\infty))$ with respect to the topology defined above (regarding point (2) above, note that by Lemma 3.18, the singular sets $S(A_i)$ converge in the Hausdorff sense into $S(A_\infty) \cup |\mathcal{C}_\infty|$). \square

With this understood we make the following definition.

Definition 3.24. The *gauge theoretic compactification* $\overline{M}_{\text{HYM}}$ is the (sequential) closure $\overline{M}_{\text{HYM}}^* \subset \widehat{M}_{\text{HYM}}$, that is, the set of all Uhlenbeck limits of sequences in M_{HYM}^* .

Remark 3.25. In a first countable space the notions of closure and sequential closure coincide. Hence, in the preceding definition it does not make a difference which closure we take.

Remark 3.26. If $\dim X = 2$, $\overline{M}_{\text{HYM}}$ coincides with the Donaldson-Uhlenbeck compactification (cf. [16, III(iii)] as well as the summary of the construction provided in [38, p. 444]).

Remark 3.27. By Proposition 3.20, $\overline{M}_{\text{HYM}}$ is sequentially closed; notice that this fact requires an argument and is not automatic. This means precisely that any limit $[(A_\infty, \mathcal{C}_\infty, S(A_\infty))]$ of elements in $\overline{M}_{\text{HYM}}$ is itself the limit of connections in M_{HYM}^* .

The main result of this section is the following.

Theorem 3.28. *The space $\overline{M}_{\text{HYM}}$ is Hausdorff and sequentially compact.*

Proof. The Hausdorffness is clear, since $\overline{M}_{\text{HYM}}$ is a subspace of a Hausdorff space. The sequential compactness follows from the fact $\overline{M}_{\text{HYM}}$ is a sequentially closed subspace of a sequentially compact space by Remark 3.27 and Theorem 3.23 above. \square

Remark 3.29. The results of the next section will give a slightly different proof of sequential compactness for $\overline{M}_{\text{HYM}}$. Namely, we will see that the reflexive sheaves associated to Uhlenbeck limits of sequences in M_{HYM}^* are the double duals of sheaves in a fixed Quot scheme, and therefore the boundedness in Lemma 3.16 actually follows for this particular subset of \widehat{M}_{HYM} without appealing to the results of [42].

4. COMPARISON OF ANALYTIC AND ALGEBRAIC MODULI SPACES

In this section we define the principal object of study in this paper: namely, a map $\overline{\Phi} : \overline{M}^\mu \rightarrow \widehat{M}_{\text{HYM}}$ from the closure of $(M^s)^{wn}$ inside the slope compactification $M^{\mu_{\text{ss}}}$ to \widehat{M}_{HYM} , extending (2.4). We will give the definition in Section 4.1. The main result of this section is that $\overline{\Phi}$ maps onto $\overline{M}_{\text{HYM}}$ and is continuous. More precisely, recall from Section 2.8 that \overline{M}^μ is constructed using the ring of invariant sections of some determinant line bundle on \mathcal{Z} , the weak normalisation of a locally closed subscheme of some Quot scheme parametrising slope semistable sheaves. We will study a map $\overline{\Psi} : \overline{\mathcal{Z}}_\circ \rightarrow \widehat{M}_{\text{HYM}}$,

defined on a certain closed subvariety $\overline{\mathcal{Z}}_\circ$ of \mathcal{Z} , that descends to $\overline{\Phi}$. Both $\overline{\mathcal{Z}}_\circ$ and \overline{M}^μ are (quasi-)projective and hence admit metrics. In particular, in order to prove continuity it suffices to prove sequential continuity.

Recall from Definition 3.15 that a point of \widehat{M}_{HYM} consists of two pieces of data: an admissible HYM connection and a holomorphic cycle. The proof of continuity consists of showing that the respective limits coincide in both these components. We first check on sequences in the interior of $\overline{\mathcal{Z}}_\circ$ (that is, for sequences of locally free stable sheaves) that the image of $\overline{\Psi}$ and therefore of $\overline{\Phi}$ lies in $\overline{M}_{\text{HYM}}$. Using Proposition 3.20 we will then reduce to this case. Equality in the sheaf component will follow from the fact that any Uhlenbeck limit of smooth HYM connections may be identified with the double dual of a sheaf appearing as a quotient of \mathcal{H} . This will be proven in Section 4.2. In Section 4.3 we give a small extension of the singular Bott-Chern formula from [54]. Using this, equality of the cycle components is proven, and the proof of continuity is then completed in Section 4.4 using the diagonalisation argument of Section 3.4.

4.1. Definition of the comparison between moduli spaces. Suppose first that we have a torsion free sheaf $\mathcal{F} \rightarrow X$ satisfying the conditions:

- (1) \mathcal{F} is μ -semistable and $\det \mathcal{F} \simeq \mathcal{J}$;
- (2) $\text{ch}_2(\mathcal{F}) = \text{ch}_2(E)$ in rational cohomology;
- (3) there is a smooth bundle isomorphism $\mathcal{F}|_{X \setminus \text{sing} \mathcal{F}} \simeq E|_{X \setminus \text{sing} \mathcal{F}}$, and so a C^∞ isomorphism

$$(\text{Gr } \mathcal{F})^{\vee\vee}|_{X \setminus \text{sing}(\text{Gr } \mathcal{F})} = (\text{Gr } \mathcal{F})|_{X \setminus \text{sing}(\text{Gr } \mathcal{F})} \simeq E|_{X \setminus \text{sing}(\text{Gr } \mathcal{F})} .$$

We will associate to this data a gauge equivalence class of ideal HYM connections $[(A, \mathcal{C}, S(A))]$ as follows.

First, note that by [3, Thm. 3] there exists an admissible HE metric $h_{\mathcal{F}}$ on $(\text{Gr } \mathcal{F})^{\vee\vee}$. By (3) we may write $(\text{Gr } \mathcal{F})^{\vee\vee} = \text{Gr } \mathcal{F} = (E, \bar{\partial}_E)$ on $X \setminus \text{sing}(\text{Gr } \mathcal{F})$. Let g be a complex gauge transformation defined on $X \setminus \text{sing}(\text{Gr } \mathcal{F})$ such that $g(h_{\mathcal{F}}) = h$ in the sense that for any $e_1, e_2 \in E$,

$$\langle e_1, e_2 \rangle_{h_{\mathcal{F}}} = \langle g^{-1}e_1, g^{-1}e_2 \rangle_h .$$

Then a straightforward calculation regarding the curvatures of the Chern connections $A = (g \cdot \bar{\partial}_E, h)$ and $(\bar{\partial}_E, h_{\mathcal{F}})$ yields that

$$F_{(g \cdot \bar{\partial}_E, h)} = g^{-1} \circ F_{(\bar{\partial}_E, h_{\mathcal{F}})} \circ g .$$

In particular, since $h_{\mathcal{F}}$ is HE and admissible, the pair $(A, \text{sing}(\text{Gr } \mathcal{F}))$ defines an admissible HYM connection on (E, h) in the sense of Definition 3.1.

Second, consider the support cycle $\mathcal{C} \in \mathcal{C}_{n-2}(X)$ of the torsion sheaf $\mathcal{T}_{\mathcal{F}} = (\text{Gr } \mathcal{F})^{\vee\vee} / \text{Gr } \mathcal{F}$ defined in Section 2.5.3. By [57, Prop. 2.3.1], $\text{ch}_2(A)$ is a closed current which is easily seen to represent $\text{ch}_2((\text{Gr } \mathcal{F})^{\vee\vee})$ in rational cohomology (cf. [54, proof of Prop. 3.3]). From the exact sequence

$$0 \longrightarrow \text{Gr } \mathcal{F} \longrightarrow (\text{Gr } \mathcal{F})^{\vee\vee} \longrightarrow \mathcal{T}_{\mathcal{F}} \longrightarrow 0$$

and the fact that $\text{ch}_2(\mathcal{T}_{\mathcal{F}}) = [\mathcal{C}]$ (see [54, Prop. 3.1]), we conclude using assumption (2) above that

$$[\text{ch}_2(A)] = \text{ch}_2(E) + [\mathcal{C}] . \quad (4.1)$$

Hence, if we define $S(A)$ to be equal to $\text{sing}((\text{Gr } \mathcal{F})^{\vee\vee})$, the conditions in Definition 3.15 are satisfied.

In summary, from a sheaf \mathcal{F} satisfying (1)–(3) above we obtained an ideal HYM connection $(A, \mathcal{C}, S(A))$ on (E, h) . Different choices of g give gauge equivalent ideal connections, and isomorphic sheaves \mathcal{F} give rise to the same class $[(A, \mathcal{C}, S(A))] \in \widehat{M}_{\text{HYM}}$.

Let $\overline{R}_o \subset R^{\mu\text{ss}}$ denote the Zariski closure of the Zariski open set R_o consisting of locally free μ -stable quotients $q : \mathcal{H} \rightarrow \mathcal{F}$ with $\det \mathcal{F} \simeq \mathcal{J}$ such that the underlying smooth bundle of \mathcal{F} is C^∞ -isomorphic to E . By an application of Ehresmann's Theorem to the associated family of unitary bundles, this is a finite union of connected components of the Zariski open subset of locally free μ -stable quotients $q : \mathcal{H} \rightarrow \mathcal{F}$ with $\det \mathcal{F} \simeq \mathcal{J}$. Let $\overline{\mathcal{Z}}_o \subset \mathcal{Z}$ its preimage under the weak normalisation map $\mathcal{Z} \rightarrow R_{\text{red}}^{\mu\text{ss}}$, which recall is a homeomorphism between the underlying topological spaces. Note that the Zariski closure of $(M^s)^{wn}$ defined in Definition 2.21 is equal to the image of $\overline{\mathcal{Z}}_o$ under the classifying map $\pi : \mathcal{Z} \rightarrow M^{\mu\text{ss}}$. By Lemma 2.19, if $q : \mathcal{H} \rightarrow \mathcal{F}$ is in \overline{R}_o , then \mathcal{F} satisfies conditions (1)–(3) above. Hence, from the previous discussion we have a well-defined map

$$\overline{\Psi} : \overline{\mathcal{Z}}_o \rightarrow \widehat{M}_{\text{HYM}}, \quad q_{\mathcal{F}} \mapsto [(A, \mathcal{C}, S(A))] .$$

Notice that the gauge equivalence class $[(A, \mathcal{C}, S(A))]$ depends only on $(\text{Gr } \mathcal{F})^{\vee\vee}$ and $\mathcal{C}_{\mathcal{F}}$. Consequently, by Proposition 2.22, the map $\overline{\Psi}$ descends to a map $\overline{\Phi} : \overline{M}^{\mu} \rightarrow \widehat{M}_{\text{HYM}}$ satisfying

$$\overline{\Psi} = \overline{\Phi} \circ \pi \Big|_{\overline{\mathcal{Z}}_o} \quad (4.2)$$

and extending (2.4) to the respective compactifications.

Remark 4.1. By construction, if \mathcal{F} is polystable, then Remark 2.12 implies that with the notations introduced in Definition 2.11 we have $\overline{\Phi}([\mathcal{F}]) = [\gamma(\mathcal{F})]$.

4.2. Identification of the limiting sheaf. The main goal of this section is to prove the following result, whose proof identifies an Uhlenbeck limit of smooth HYM connections with the double dual of a certain point in $\text{Quot}(\mathcal{H}, \tau)$ (in fact in $\text{Quot}(\mathcal{H}, c(E))$).

Proposition 4.2. *Let A_i be a sequence of HYM connections on (E, h) , giving holomorphic bundles \mathcal{E}_i . Then for any Uhlenbeck limit A_∞ with reflexive extension \mathcal{E}_∞ , there is a point $\widehat{q}_\infty : \mathcal{H} \rightarrow \widehat{\mathcal{E}}_\infty \rightarrow 0$ in $\text{Quot}(\mathcal{H}, c(E))$, such that $(\widehat{\mathcal{E}}_\infty)^{\vee\vee}$ is isomorphic to \mathcal{E}_∞ .*

Let $\mathcal{E}_i = (E, \bar{\partial}_{A_i})$. We realize $\mathcal{E}_i(m)$ as a quotient by choosing a basis \mathbf{e}_i of holomorphic sections that is orthonormal with respect to the L^2 metric induced by the hermitian structure $h \otimes h_L^m$ on $\mathcal{E}_i(m)$. Denote the similar construction of a choice of a basis for a further twist by:

$$\mathbf{e}_i^{(k)} : \mathcal{O}_X^{\tau(m+k)} \longrightarrow \mathcal{E}_i(m+k). \quad (4.3)$$

We begin with a preliminary result.

Lemma 4.3. *There is a subsequence, also denoted $\{i\}$, such that for any $k \geq 0$, the basis $\mathbf{e}_i^{(k)}$ converges smoothly away from S_∞ to a nonzero holomorphic map*

$$\mathbf{e}_\infty^{(k)} : \mathcal{O}_X^{\tau(m+k)} \longrightarrow \mathcal{E}_\infty(m+k).$$

Moreover, if $\mathcal{E}_\infty^{(k)} \subset \mathcal{E}_\infty$ denotes the image of $\mathbf{e}_\infty^{(k)}$, twisted by $\mathcal{O}_X(-m-k)$, then there is a coherent subsheaf $\widetilde{\mathcal{E}}_\infty \subset \mathcal{E}_\infty$, such that

- (1) for k sufficiently large, $\mathcal{E}_\infty^{(k)} = \widetilde{\mathcal{E}}_\infty$;
- (2) $(\widetilde{\mathcal{E}}_\infty)^{\vee\vee} \cong \mathcal{E}_\infty$;
- (3) $\chi_{\widetilde{\mathcal{E}}_\infty}(\ell) \geq \tau(\ell)$, for ℓ sufficiently large.

Proof. For k large, we denote the orthonormal basis $\mathbf{e}_i^{(k)} = \{e_{i,j}^{(k)}\}_{j=1}^{\tau(m+k)}$. Let B_i denote the connection on $\mathcal{E}_i(m+k)$ induced by A_i on \mathcal{E}_i and the connection a_L on L . Notice that (see Section 2.1)

$$\sqrt{-1}\Lambda F_{B_i} = \sqrt{-1}\Lambda F_{A_i} + 2\pi(m+k)\lambda\omega \cdot \mathbf{I}.$$

Then since $e_{i,j}^{(k)}$ is B_i -holomorphic, the Bochner formula gives

$$\begin{aligned} \sqrt{-1}\bar{\partial}\partial\langle e_{i,j}^{(k)}, e_{i,j}^{(k)} \rangle &= \sqrt{-1}\bar{\partial}\langle \partial_{B_i} e_{i,j}^{(k)}, e_{i,j}^{(k)} \rangle \\ &= \sqrt{-1}(\langle \bar{\partial}_{B_i} \partial_{B_i} e_{i,j}^{(k)}, e_{i,j}^{(k)} \rangle + \langle \partial_{B_i} e_{i,j}^{(k)}, \partial_{B_i} e_{i,j}^{(k)} \rangle) \\ &= \sqrt{-1}(\langle F_{B_i} e_{i,j}^{(k)}, e_{i,j}^{(k)} \rangle - \langle \partial_{B_i} e_{i,j}^{(k)}, \partial_{B_i} e_{i,j}^{(k)} \rangle), \end{aligned}$$

Hence,

$$\Delta(|e_{i,j}^{(k)}|^2) = 2\Delta_\partial(|e_{i,j}^{(k)}|^2) = \Lambda((2\sqrt{-1}\bar{\partial}\partial(|e_{i,j}^{(k)}|^2)))$$

$$\begin{aligned}
 &= 2\sqrt{-1}\langle \Lambda F_{B_i} e_{i,j}^{(k)}, e_{i,j}^{(k)} \rangle - 2|\partial_{B_i} e_{i,j}^{(k)}|^2 \\
 &\leq 2\sqrt{-1}\langle \Lambda F_{B_i} e_{i,j}^{(k)}, e_{i,j}^{(k)} \rangle,
 \end{aligned}$$

where we have used the identity $\sqrt{-1}\Lambda\langle\beta, \beta\rangle = |\beta|^2$ for any $\beta \in \Omega^{1,0}(X, E \otimes L^{m+k})$. We therefore have

$$\Delta|e_{i,j}^{(k)}|^2 \leq C(\mu(E) + 2\pi n\lambda(m+k))|e_{i,j}^{(k)}|^2 \leq C|e_{i,j}^{(k)}|^2,$$

for a constant C independent of i . By a result of Morrey (cf. [27, Thm. 9.20]), there is a uniform bound

$$\sup_X |e_{i,j}^{(k)}| \leq C \cdot \|e_{i,j}^{(k)}\|^2 \leq C,$$

independent of i and for each j . Since the analytic singular set S has measure zero, we may find a cover $\{B_\sigma(x_\alpha)\}_{x_\alpha \in S}$ such that

$$C \sum_i \text{vol}(B_\sigma(x_\alpha)) < \frac{1}{4}.$$

Write $K_\sigma = X \setminus \cup_\alpha B_\sigma(x_\alpha)$, and note that this together with the L^∞ estimate given above we have that

$$\|e_{i,j}^{(k)}\|_{L^2(X)} = \|e_{i,j}^{(k)}\|_{L^2(\cup_\alpha B_\sigma(x_\alpha))} + \|e_{i,j}^{(k)}\|_{L^2(K_\sigma)} \leq \|e_{i,j}^{(k)}\|_{L^2(K_\sigma)} + \frac{1}{2},$$

so that $\|e_{i,j}^{(k)}\|_{L^2(K_\sigma)} \geq \frac{1}{2}$.

By Theorem 3.3, the A_i converge smoothly outside of a holomorphic subvariety S to a connection A_∞ on a bundle E_∞ which is smoothly isometric to (E, h) , and correspondingly the connections B_i converge smoothly to a connection B_∞ on the complement of this set. Write $\bar{\partial}_{B_\infty} = \bar{\partial}_{B_i} + \beta_i$ for $\beta_i \in \Omega^{0,1}(K_\sigma, \text{End } E)$. Then holomorphicity of the sections implies $\bar{\partial}_{B_\infty} e_{i,j}^{(k)} = \beta_i e_{i,j}^{(k)}$. Since $\beta_i \rightarrow 0$ smoothly, elliptic regularity gives C^∞ estimates on $e_{i,j}^{(k)}$, and we may extract a convergent subsequence to a limit $e_{\infty,j}^{(k)}$, which is holomorphic with respect to B_∞ . By the lower bound on $\|e_{i,j}^{(k)}\|_{L^2(K_\sigma)}$ it follows that $e_{\infty,j}^{(k)} \neq 0$. Repeating this argument for an exhaustion of $X \setminus S$ by sets K_σ constructed by taking the radii of the balls B_σ to zero, and using the fact that the set S is a holomorphic subvariety (and so we may take this exhaustion to be by deformation retracts of $X \setminus S$), it follows from a diagonalisation argument that we obtain nonzero holomorphic sections of $\mathcal{E}_\infty(m+k)$ on $X \setminus S$. Since \mathcal{E}_∞ is reflexive, these extend to all of X and so we obtain for each j a nontrivial holomorphic section $e_{\infty,j}^{(k)} \in H^0(\mathcal{E}_\infty(m+k))$.

For each pair (j, l) , apply the dominated convergence theorem to conclude that

$$\delta_{jl} = \lim_{i \rightarrow \infty} \int_X \langle e_{i,j}^{(k)}, e_{i,l}^{(k)} \rangle \frac{\omega^n}{n!} = \int_X \langle e_{\infty,j}^{(k)}, e_{\infty,l}^{(k)} \rangle \frac{\omega^n}{n!} ,$$

so that the $\{e_{\infty,j}^{(k)}\}_{j=1}^{\tau(m+k)}$ is an L^2 -orthonormal subset of $H^0(\mathcal{E}_\infty(m+k))$. In other words, we have constructed the promised nonzero holomorphic map

$$\mathbf{e}_\infty^{(k)} : \mathcal{O}_X^{\oplus \tau(m+k)} \rightarrow \mathcal{E}_\infty(m+k) .$$

By a further diagonalisation argument, we can arrange for the convergence along a subsequence, also denoted $\{i\}$, for all k sufficiently large. We now check the stated properties of $\mathbf{e}_\infty^{(k)}$.

Let $\mathcal{E}_\infty^{(k)} \subset \mathcal{E}_\infty$ denote the coherent subsheaf

$$\mathrm{Im}(\mathcal{O}_X^{\oplus \tau(m+k)} \xrightarrow{\mathbf{e}_\infty^{(k)}} \mathcal{E}_\infty(m+k)) \otimes \mathcal{O}_X(-m-k) .$$

We claim that for all k sufficiently large, $\mathcal{E}_\infty^{(k)} \subset \mathcal{E}_\infty^{(k+1)}$. Choose a point $p \in X$, U a neighbourhood of p , a local section s of $\mathcal{E}_\infty^{(k)}$ on U , and a global section σ of $\mathcal{O}_X(1)$ that is nonvanishing at p . By definition, we may write

$$s = \sum_{j=1}^{\tau(m+k)} f^j e_{\infty,j}^{(k)} \otimes \sigma^{-m} , \quad (4.4)$$

for some $f^j \in \mathcal{O}_U$. Since $\{e_{i,j}^{(k+1)}\}$ is an L^2 -orthonormal basis for $H^0(X, \mathcal{E}_i(m+k+1))$, for each i we may write

$$e_{i,j}^{(k)} \otimes \sigma = \sum_{q=1}^{\tau(m+k+1)} \langle e_{i,j}^{(k)} \otimes \sigma, e_{i,q}^{(k+1)} \rangle e_{i,q}^{(k+1)} , \quad (4.5)$$

where $\langle \cdot, \cdot \rangle$ is the L^2 -inner product. By the same argument as above,

$$\lim_{i \rightarrow \infty} \langle e_{i,j}^{(k)} \otimes \sigma, e_{i,q}^{(k+1)} \rangle \longrightarrow \langle e_{\infty,j}^{(k)} \otimes \sigma, e_{\infty,q}^{(k+1)} \rangle .$$

The right and left hand sides of (4.5) therefore converge smoothly away from Z^{an} , and therefore $e_{\infty,j}^{(k)} \otimes \sigma$ lies in the image of $\mathbf{e}_\infty^{(k+1)}$. This, in turn means that s in (4.4) may be written

$$s = \sum_{j=1}^{\tau(m+k+1)} g^j e_{\infty,j}^{(k+1)} \otimes \sigma^{-m-1} ,$$

for some $g^j \in \mathcal{O}_U$. Since the right hand side is in $\mathcal{E}_\infty^{(k+1)}$ by definition, the claim follows.

By the Noetherian property of coherent subsheaves of a coherent sheaf [55, Tag 01Y7], the chain $\mathcal{E}_\infty^{(k)} \subset \mathcal{E}_\infty^{(k+1)} \subset \dots \subset \mathcal{E}_\infty$, must stabilise for sufficiently

large k . We set $\tilde{\mathcal{E}}_\infty = \mathcal{E}_\infty^{(k)}$, $k \gg 0$, and this proves part (1) of the lemma. For part (3), notice that for $\ell \gg 0$,

$$\chi_{\tilde{\mathcal{E}}_\infty}(\ell) = \chi_{\mathcal{E}_\infty^{(\ell-m)}}(\ell) = h^0(\mathcal{E}_\infty^{(\ell-m)}(\ell)) \geq h^0(\mathcal{E}_i(\ell)) = \tau_E(\ell). \quad (4.6)$$

From the Hirzebruch-Riemann-Roch Theorem and (2.6), we see that $\text{rank } \tilde{\mathcal{E}}_\infty = \text{rank } E = \text{rank } \mathcal{E}_\infty$. Hence, $\mathcal{T} = \mathcal{E}_\infty / \tilde{\mathcal{E}}_\infty$ is a torsion sheaf, and $\det \mathcal{E}_\infty \cong \det \tilde{\mathcal{E}}_\infty \otimes \det \mathcal{T}$. But then (4.6) and (2.6) also imply

$$\deg \tilde{\mathcal{E}}_\infty \geq \deg E = \deg \mathcal{E}_\infty = \deg \tilde{\mathcal{E}}_\infty + \deg(\det \mathcal{T}).$$

Since the last term on the right hand side above is nonnegative, it must be that $\deg(\det \mathcal{T}) = 0$, and hence $\text{codim}(\text{supp}(\mathcal{T})) \geq 2$. Therefore, $(\tilde{\mathcal{E}}_\infty)^{\vee\vee} \cong \mathcal{E}_\infty^{\vee\vee}$. Since \mathcal{E}_∞ is reflexive, (2) holds. This completes the proof of the lemma. \square

Henceforth, we assume k has been chosen sufficiently large according to Lemma 4.3, and we drop k from the notation. Since $\text{Quot}(\mathcal{H}, c(E))$ is a projective scheme, in the analytic topology the sequence \mathbf{e}_i in (4.3) converges to a limiting quotient. By definition of the equivalence of quotients, this means there is a sequence of quotients $\hat{q}_i : \mathcal{H} \rightarrow \hat{\mathcal{E}}_i \rightarrow 0$, isomorphisms $\varphi_i : \hat{\mathcal{E}}_i \xrightarrow{\sim} \mathcal{E}_i$, and a commutative diagram

$$\begin{array}{ccc} \mathcal{H} & \xrightarrow{\hat{q}_i} & \hat{\mathcal{E}}_i \longrightarrow 0 \\ \parallel & & \downarrow \varphi_i \\ \mathcal{H} & \xrightarrow{\mathbf{e}_i} & \mathcal{E}_i \longrightarrow 0 \end{array}$$

Furthermore, there is $\hat{q}_\infty \in \text{Quot}(\mathcal{H}, c(E))$ such that $\hat{q}_i \rightarrow \hat{q}_\infty$.

Lemma 4.4. *There is a map $\varphi_\infty : \hat{\mathcal{E}}_\infty \rightarrow \tilde{\mathcal{E}}_\infty$ making the following diagram commute:*

$$\begin{array}{ccc} \mathcal{H} & \xrightarrow{\hat{q}_\infty} & \hat{\mathcal{E}}_\infty \longrightarrow 0 \\ \parallel & & \downarrow \varphi_\infty \\ \mathcal{H} & \xrightarrow{\mathbf{e}_\infty} & \tilde{\mathcal{E}}_\infty \longrightarrow 0 \end{array}$$

Moreover, φ_∞ is an isomorphism.

Proof. Since \hat{q}_∞ is surjective, to define the map φ_∞ it suffices to show that $\ker \hat{q}_\infty \subset \ker \mathbf{e}_\infty$. In fact, it is enough to prove the inclusion away from a proper subvariety. For then the image $\mathbf{e}_\infty(\ker \hat{q}_\infty)$ would be a torsion subsheaf of $\tilde{\mathcal{E}}_\infty$, and since $\tilde{\mathcal{E}}_\infty \subset \mathcal{E}_\infty$ is torsion free, such a sheaf must vanish. Now away from a set of codimension at least 2 we have smooth convergence: $\mathbf{e}_i \rightarrow \mathbf{e}_\infty$, and $\hat{\pi}_i \rightarrow \hat{\pi}_\infty$ (see Lemma 2.19). Hence, $0 = \mathbf{e}_i \circ \hat{\pi}_i \rightarrow \mathbf{e}_\infty \circ \hat{\pi}_\infty$, and the

result follows. The map φ_∞ is therefore well-defined. We claim it is in fact an isomorphism. First, since

$$\text{rank im } \widehat{q}_\infty = \text{rank } \widehat{\mathcal{E}}_\infty = \text{rank } \mathcal{E}_\infty = \text{rank } \widetilde{\mathcal{E}}_\infty = \text{rank im } \mathbf{e}_\infty ,$$

we have $\text{rank}(\ker \widehat{q}_\infty) = \text{rank}(\ker \mathbf{e}_\infty)$. By the discussion above, $\ker \widehat{q}_\infty \subset \ker \mathbf{e}_\infty$, and so

$$\ker \mathbf{e}_\infty / \ker \widehat{q}_\infty \cong \widehat{q}_\infty(\ker \mathbf{e}_\infty) \subset \mathcal{T} \subset \widehat{\mathcal{E}}_\infty , \quad (4.7)$$

where $\mathcal{T} \subset \widehat{\mathcal{E}}_\infty$ is the torsion subsheaf. Since $\widetilde{\mathcal{E}}_\infty$ is torsion free, φ_∞ extends to a map $\overline{\varphi}_\infty : \widehat{\mathcal{E}}_\infty / \mathcal{T} \rightarrow \widetilde{\mathcal{E}}_\infty$. On the other hand, from (4.7) there is a well-defined map: $\eta : \widetilde{\mathcal{E}}_\infty \rightarrow \widehat{\mathcal{E}}_\infty / \mathcal{T}$ that is a generic inverse to φ_∞ . Both maps η and $\overline{\varphi}_\infty$ are injective, since their kernels would be torsion sheaves, and $\widetilde{\mathcal{E}}_\infty, \widehat{\mathcal{E}}_\infty / \mathcal{T}$ are torsion free. Hence, there is an exact sequence

$$0 \longrightarrow \widetilde{\mathcal{E}}_\infty \xrightarrow{\eta} \widehat{\mathcal{E}}_\infty / \mathcal{T} \longrightarrow \widehat{\mathcal{T}} \longrightarrow 0 ,$$

for a torsion sheaf $\widehat{\mathcal{T}}$. It follows that,

$$\begin{aligned} \chi_{\widetilde{\mathcal{E}}_\infty}(\ell) &= \chi_{\widehat{\mathcal{E}}_\infty}(\ell) - \chi_{\mathcal{T}}(\ell) - \chi_{\widehat{\mathcal{T}}}(\ell) \\ \chi_{\widetilde{\mathcal{E}}_\infty}(\ell) - \tau(\ell) &= -\chi_{\mathcal{T}}(\ell) - \chi_{\widehat{\mathcal{T}}}(\ell) . \end{aligned}$$

By Lemma 4.3 (3), the left hand side of the last equation above is nonnegative, whereas the right hand side must be nonpositive. So $\chi_{\mathcal{T}}(\ell) = \chi_{\widehat{\mathcal{T}}}(\ell) = 0$, and therefore \mathcal{T} and $\widehat{\mathcal{T}}$ are zero. Hence, φ_∞ is an isomorphism. \square

Proof of Proposition 4.2. Immediate from Lemmas 4.4 and 4.3 (2). \square

As a consequence of the main result above, the following proves the first step in the continuity of $\overline{\Phi}$.

Proposition 4.5. *Let $q_i \in \overline{\mathcal{Z}}_\circ$ be a sequence whose quotients \mathcal{F}_i are locally free and μ -stable. Assume that $q_i \rightarrow q_\infty \in \overline{\mathcal{Z}}_\circ$ in the analytic topology, so that in particular the corresponding quotient \mathcal{F}_∞ is torsion free and μ -semistable. Let A_i be the corresponding sequence of HYM connections on E . Then for any Uhlenbeck limit A_∞ of $\{A_i\}$ extending to a polystable reflexive sheaf \mathcal{E}_∞ , we have $\text{Gr}(\mathcal{F}_\infty)^{\vee\vee} \cong \mathcal{E}_\infty$.*

Proof. As in Proposition 4.2, write $q_i : \mathcal{H} \rightarrow \mathcal{E}_i \rightarrow 0$ for the sequence in $\overline{\mathcal{Z}}_\circ$ corresponding to the connections A_i in the $\mathcal{G}^\mathbb{C}$ orbit of \mathcal{F}_i . We have $\mathcal{E}_\infty \cong (\widehat{\mathcal{E}}_\infty)^{\vee\vee}$ for some quotient $\widehat{q}_\infty : \mathcal{H} \rightarrow \widehat{\mathcal{E}}_\infty \rightarrow 0$, which is a limit of $\widehat{q}_i : \mathcal{H} \rightarrow \widehat{\mathcal{E}}_i \rightarrow 0$ with $\widehat{\mathcal{E}}_i \cong \mathcal{F}_i$.

Note that we have proven in Lemma 4.4 above that

$$\widehat{\mathcal{E}}_\infty \cong \widetilde{\mathcal{E}}_\infty \subset \mathcal{E}_\infty . \quad (4.8)$$

Since \mathcal{E}_∞ is polystable (having an admissible Hermitian-Yang-Mills connection), and hence in particular μ -semistable, this implies that $\widehat{\mathcal{E}}_\infty$ is μ -semistable, cf. Proposition 2.15. Note also that because $\widehat{\mathcal{E}}_\infty$ is isomorphic to \mathcal{E}_∞ away from a codimension two subvariety, their determinant line bundles are equal, and in particular the determinant line bundle of $\widehat{\mathcal{E}}_\infty$ is the same as that of $\widehat{\mathcal{E}}_i \cong \mathcal{F}_i$. Therefore, $\widehat{\mathcal{E}}_\infty$ defines a point in $M^{\mu ss}$, and hence by construction also in \overline{M}^μ .

By construction, we have $[\widehat{\mathcal{E}}_i] = [\mathcal{F}_i]$ in \overline{M}^μ . By continuity of the map $\pi : \mathcal{Z} \rightarrow M^{\mu ss}$, we therefore have $[\widehat{\mathcal{E}}_\infty] = [\mathcal{F}_\infty]$ in \overline{M}^μ . It hence follows from Proposition 2.22 that $\text{Gr}(\mathcal{F}_\infty)^{\vee\vee} \cong \text{Gr}(\widehat{\mathcal{E}}_\infty)^{\vee\vee}$. On the other hand, the fact that \mathcal{E}_∞ is polystable together with (4.8) allows to apply Proposition 2.15 to obtain an isomorphism $\text{Gr}(\widehat{\mathcal{E}}_\infty)^{\vee\vee} \cong \mathcal{E}_\infty$, which completes the proof. \square

4.3. Identification of the limiting cycle. Having identified the reflexive extension of an Uhlenbeck limit of a sequence of smooth HYM connections, we now wish to determine the singular set as well. We will use the following singular version of the usual Bott-Chern formula, which is a slight generalisation of [54, Thm. 1.3].

Proposition 4.6. *Let $q_{\mathcal{E}}, q_{\mathcal{F}} \in \text{Quot}(\mathcal{H}, \tau)$, with \mathcal{F} torsion free and \mathcal{E} locally free. Assume that the underlying smooth vector bundles F and E of \mathcal{F} and \mathcal{E} on $X \setminus \text{sing } \mathcal{F}$ are isomorphic. Set $\mathcal{E}_\infty = \mathcal{F}^{\vee\vee}$, and let $\mathcal{C}_{\mathcal{F}}$ be the support cycle. Then for any smooth hermitian metric h on \mathcal{E} and admissible metric h_∞ on \mathcal{E}_∞ , there is an equation of currents*

$$\text{ch}_2(\mathcal{E}_\infty, h_\infty) - \text{ch}_2(\mathcal{E}, h) = \mathcal{C}_{\mathcal{F}} + dd^c \Psi, \quad (4.9)$$

where Ψ is a $(2, 2)$ -current, smooth outside the support of $\mathcal{C}_{\mathcal{F}}$.

Proof. Choose representatives of the points in the Quot scheme:

$$\begin{aligned} q_{\mathcal{E}} : 0 &\longrightarrow \mathcal{K} \longrightarrow \mathcal{H} \longrightarrow \mathcal{E} \longrightarrow 0, \\ q_{\mathcal{F}} : 0 &\longrightarrow \mathcal{K}_\infty \longrightarrow \mathcal{H} \longrightarrow \mathcal{F} \longrightarrow 0. \end{aligned}$$

Let H, k be fixed hermitian metrics on \mathcal{H} and \mathcal{K} , respectively. By the Bott-Chern formula applied to the sequence $q_{\mathcal{E}}$, there is a smooth $(1, 1)$ -form Ψ_1 such that

$$\text{ch}_2(\mathcal{K} \oplus \mathcal{E}, k \oplus h) - \text{ch}_2(\mathcal{H}, H) = dd^c \Psi_1. \quad (4.10)$$

Let k_∞ be an admissible metric on the reflexive sheaf \mathcal{K}_∞ . Clearly, $k_\infty \oplus h_\infty$ is then an admissible metric on $\mathcal{K}_\infty \oplus \mathcal{E}_\infty$. By [54, Thm. 1.3] applied to the sequence $q_{\mathcal{F}}$, there is a $(1, 1)$ -form Ψ_2 such that

$$\text{ch}_2(\mathcal{K}_\infty \oplus \mathcal{E}_\infty, k_\infty \oplus h_\infty) - \text{ch}_2(\mathcal{H}, H) = dd^c \Psi_2 + \mathcal{C}_{\mathcal{F}}. \quad (4.11)$$

By the assumption on F and E , the Chern connection for $(\bar{\partial}_{\mathcal{E}_\infty}, h_\infty)$ on $X \setminus \text{sing } \mathcal{F}$ defines an admissible connection on the smooth bundle E . It follows from the proof of [54, Prop. 3.3] that $\text{ch}_2(\mathcal{E}_\infty, h_\infty)$ is a closed current representing $\text{ch}_2(\mathcal{E}_\infty)$ in rational cohomology. The essential point here is that although the proof contained in [54, Prop. 3.3] does not apply to an arbitrary torsion free sheaf, it does apply as long as the underlying smooth bundle extends smoothly to all of X , which is guaranteed by the hypotheses. From (4.11), it then follows that $\text{ch}_2(\mathcal{K}_\infty, k_\infty)$ is also a closed current. Appealing once more to the proof of [54, Prop. 3.3] (recall that \mathcal{K}_∞ is reflexive), $\text{ch}_2(\mathcal{K}_\infty, k_\infty)$ represents $\text{ch}_2(\mathcal{K}_\infty)$ in rational cohomology. Since \mathcal{K} and \mathcal{K}_∞ have the same Chern classes, it follows that there is a $(1, 1)$ -form Ψ_3 such that

$$\text{ch}_2(\mathcal{K}_\infty, k_\infty) - \text{ch}_2(\mathcal{K}, k) = dd^c \Psi_3 . \quad (4.12)$$

Finally, (4.9) follows by combining eqs. (4.10), (4.11), and (4.12). \square

The following result may be viewed as the analog of the main result in [54].

Proposition 4.7. *In the notation of Proposition 4.5, the cycle $\mathcal{C}_{\mathcal{F}_\infty}$ is equal to the analytic cycle \mathcal{C}^{an} associated to the Uhlenbeck limit A_∞ . Moreover, $\text{sing}(\mathcal{E}_\infty) = S(A_\infty)$.*

The rest of this section is devoted to the proof of Proposition 4.7. To simplify the notation, write: $\mathcal{C}^{alg} := \mathcal{C}_{\mathcal{F}_\infty} = \sum_j m_j^{alg} Z_j^{alg}$.

Lemma 4.8. *In rational cohomology, $[\mathcal{C}^{alg}] = [\mathcal{C}^{an}]$.*

Proof. By Proposition 4.5, the HYM connection on $\text{Gr}(\mathcal{F}_\infty)^{\vee\vee}$ is the Uhlenbeck limit A_∞ , and therefore eq. (4.1) implies that $[\text{ch}_2(A_\infty)] = \text{ch}_2(E) + [\mathcal{C}_{\mathcal{F}_\infty}]$. But recall also that (3.4) implies that $[\text{ch}_2(A_\infty)] = \text{ch}_2(E) + [\mathcal{C}^{an}]$ and the result follows. \square

Now we wish to go further and prove equality as cycles. The first observation is that by the proof of Proposition 4.5 and Lemma 4.4, in \bar{M}^μ we have $[\mathcal{F}_\infty] = [\tilde{\mathcal{E}}_\infty]$, where $\tilde{\mathcal{E}}_\infty \cong \hat{\mathcal{E}}_\infty$ is the μ -semistable sheaf that appears in Lemma 4.3. Then it follows from Proposition 2.22 that $\mathcal{C}^{alg} = \mathcal{C}_{\tilde{\mathcal{T}}_\infty}$, where $\tilde{\mathcal{T}}_\infty$ is the torsion sheaf $\text{Gr}(\tilde{\mathcal{E}}_\infty)^{\vee\vee}/\text{Gr}(\tilde{\mathcal{E}}_\infty)$. By Proposition 2.15, we obtain $\mathcal{C}^{alg} = \mathcal{C}_{\tilde{\mathcal{T}}_\infty} = \mathcal{C}_{\mathcal{T}_\infty}$ where $\mathcal{T}_\infty = \mathcal{E}_\infty/\tilde{\mathcal{E}}_\infty$. Hence, it suffices to show $\mathcal{C}^{an} = \mathcal{C}_{\mathcal{T}_\infty}$. The main technical result of this subsection is the following, which together with Lemma 4.8 will prove Proposition 4.7.

Proposition 4.9. *Let $Z \subset \text{supp}(\mathcal{T}_\infty)$ be an irreducible component of codimension 2. Write m_Z for the multiplicity of Z in the cycle $\mathcal{C}_{\mathcal{T}_\infty}$. Choose a generic*

slice Σ . Then in the notation of Proposition 4.5, we have:

$$m_Z = \lim_{i \rightarrow \infty} \frac{1}{8\pi^2} \int_{\Sigma} \{\text{tr}(F_{A_i} \wedge F_{A_i}) - \text{tr}(F_{A_\infty} \wedge F_{A_\infty})\}.$$

In particular, since m_Z is by definition positive, then by Lemma 3.13, $Z \subset S_b$. Therefore, Z is a component of the cycle \mathcal{C}^{an} with $m_Z^{an} = m_Z$.

The proof of Proposition 4.9 will be an adaptation of the proof of Proposition 4.2 of [54]. It will use the Bott-Chern formula of Section 4.3 and the slicing lemmas of Section 3.2. Additionally we will require the following lemma, whose proof is identical to that of Lemma 4.3 of [54].

Lemma 4.10. *Let $\mathcal{V} \rightarrow X$ be a holomorphic vector bundle, $\mathcal{V} \rightarrow \mathcal{F} \rightarrow 0$ a quotient, and $\mathcal{K} \subset \mathcal{V}$ the kernel. Write $\text{Gr}(\mathcal{V}) = \mathcal{K} \oplus \mathcal{F}$. Let $Z \subset \text{supp}(\text{Gr}(\mathcal{V})^{\vee\vee} / \text{Gr}(\mathcal{V}))$ be an irreducible component of codimension 2. Then there exists a sequence of blow-ups along smooth complex submanifolds $\pi : \hat{X} \rightarrow X$ having centre \mathbf{C} and exceptional divisor $\mathbf{E} = \pi^{-1}(\mathbf{C})$, and a subsheaf $\hat{\mathcal{K}} \subset \pi^*\mathcal{V}$ giving rise to an associated graded object $\text{Gr}(\pi^*\mathcal{V}) = \hat{\mathcal{K}} \oplus \hat{\mathcal{F}}$ with singular set $\text{sing}(\text{Gr}(\pi^*\mathcal{V}))$ such that the following properties are satisfied:*

- (1) $\text{Gr}(\pi^*\mathcal{V}) \cong \text{Gr}(\mathcal{V})$ on $\hat{X} - \mathbf{E} = X - \mathbf{C}$;
- (2) $\text{codim}(Z \cap \mathbf{C}) \geq 3$;
- (3) $\text{codim}(\pi(\text{sing}(\text{Gr}(\pi^*\mathcal{V})) - Z)) \geq 3$.

Proof of Proposition 4.9. Recall that \mathcal{E}_∞ was an Uhlenbeck limit of a sequence of Hermitian-Einstein vector bundles $\mathcal{E}_i = (E, \bar{\partial}_{A_i})$, which in Section 3.3 we realised as quotients $\mathcal{H} \rightarrow \mathcal{E}_i \rightarrow 0$. In Lemma 4.4 we saw that there was a resulting quotient which we have been calling $\tilde{\mathcal{E}}_\infty$, and which we showed to be isomorphic to a certain limiting quotient which we called $\hat{\mathcal{E}}_\infty$. For notational reasons, during this proof we will rename this quotient $q_{\mathcal{Q}_\infty} : \mathcal{H} \rightarrow \mathcal{Q}_\infty \rightarrow 0$, so that in particular $\mathcal{Q}_\infty^{\vee\vee} \cong \mathcal{E}_\infty$. We may therefore consider the exact sequences

$$\begin{aligned} 0 &\longrightarrow \mathcal{K}_i \longrightarrow \mathcal{H} \longrightarrow \mathcal{E}_i \longrightarrow 0, \\ 0 &\longrightarrow \mathcal{K}_\infty \longrightarrow \mathcal{H} \longrightarrow \mathcal{Q}_\infty \longrightarrow 0. \end{aligned}$$

If we write $\text{Gr}_\infty(\mathcal{H})$ for the associated graded object of this latter sequence, then notice that $\text{supp}(\mathcal{T}_\infty) = \text{supp}((\text{Gr}_\infty(\mathcal{H})^{\vee\vee}) / \text{Gr}(\mathcal{H}))$, since \mathcal{K}_∞ is saturated and therefore reflexive.

Now we are in a position to apply Lemma 4.10 and we obtain first of all an exact sequence

$$0 \longrightarrow \hat{\mathcal{K}}_\infty \longrightarrow \pi^*\mathcal{H} \longrightarrow \hat{\mathcal{Q}}_\infty \longrightarrow 0$$

having the properties stated there.

Consider $\sum_j \hat{m}_j^{alg} \hat{W}_j$, the analytic cycle associated with the support of the sheaf $\text{Gr}_\infty(\pi^*\mathcal{H})^{\vee\vee} / \text{Gr}_\infty(\pi^*\mathcal{H})$. By the first two parts of the Lemma 4.10, we know that in fact, one of these irreducible components, say \hat{W}_1 is the proper transform \hat{Z} of Z , and so $\pi(\hat{W}_1) = Z$ and so since at a generic point of \hat{Z} the sheaf $\text{Gr}_\infty(\pi^*\mathcal{H})^{\vee\vee} / \text{Gr}_\infty(\pi^*\mathcal{H})$ is isomorphic with $(\text{Gr}_\infty \mathcal{H})^{\vee\vee} / \text{Gr}_\infty(\mathcal{H})$, we also have that $\hat{m}_1^{alg} = m_Z$.

Since X is Kähler, the class $[Z]$ is homologically nontrivial, there is a class say $\alpha \in H_4(X, \mathbb{Z})$ such that the intersection product $\alpha \cdot [Z] \neq 0$. By a classical result of Thom, an integral multiple of any 4-dimensional class can be represented by an embedded 4-dimensional submanifold Σ , and by construction $[\Sigma] \cdot [Z] \neq 0$. We may furthermore choose Σ so that it intersects Z transversely in its smooth locus, and Σ is analytic in a neighbourhood of $\Sigma \cap Z$. Therefore $\Sigma \cap Z = \{z_1, \dots, z_k\}$, where this set is nonempty. By possibly moving Σ slightly, we may suppose that Σ intersects Z in each point z_i as a transverse slice. By Lemma 4.10 (3), all the other \hat{W}_j must map to a set of codimension greater than or equal to 3, and so must be contained in \mathbf{E} . Also, by the codimension statement in (3), we may assume that Σ misses the set $\pi(\text{sing}(\text{Gr}_\infty(\pi^*\mathcal{H})) - Z)$, and the proper transform $\hat{\Sigma}$ of Σ only intersects $\text{sing}(\text{Gr}_\infty(\pi^*\mathcal{H}))$ in $\pi^{-1}(Z)$. On the other hand, by part (ii) we may assume that Σ misses $Z \cap \mathbf{C}$, so $\pi^{-1}(\Sigma)$ and therefore $\hat{\Sigma}$ misses $\mathbf{E} \cap \pi^{-1}(Z)$, and so $\hat{\Sigma}$ empty intersection with all of the \hat{W}_j except for $\hat{Z} = \hat{W}_1$.

Since the quotient \mathcal{Q}_∞ is torsion free, the kernel \mathcal{K}_∞ is saturated, and therefore reflexive, and so has at most codimension 3 singularities. Therefore, we may assume the 4-dimensional submanifold Σ misses the set $\text{sing } \mathcal{K}_\infty$. In other words the restrictions $\mathcal{K}_i|_\Sigma$ and $\mathcal{K}_\infty|_\Sigma$ are smooth vector bundles on Σ . We have proven that \mathcal{Q}_∞ is isomorphic to a quotient in $\text{Quot}(\mathcal{H}, c(E))$, and in particular the total Chern class of \mathcal{Q}_∞ is the same as that of the \mathcal{E}_i , and therefore the total Chern class of \mathcal{K}_∞ is the same as that of the \mathcal{K}_i . If $\iota : \Sigma \hookrightarrow X$ is the inclusion map, then by naturality, we have $c(\mathcal{K}_i|_\Sigma) = \iota^*(c(\mathcal{K}_i)) = \iota^*(c(\mathcal{K}_\infty)) = c(\mathcal{K}_\infty|_\Sigma)$. Now the point is that the total Chern class determines the topological type of a smooth vector bundle on a 4-manifold, so $\mathcal{K}_i|_\Sigma$ and $\mathcal{K}_\infty|_\Sigma$ are smoothly isomorphic.

As in Proposition 4.6, we fix a smooth hermitian metric k on \mathcal{K}_i (smoothly these bundles are all the same), and an admissible metric on k_∞ on \mathcal{K}_∞ . We will denote by $A_{\mathcal{K}_i}$ and $A_{\mathcal{K}_\infty}$ the corresponding Chern connections. Since $A_{\mathcal{K}_i}$ and $A_{\mathcal{K}_\infty}$ restricted to Σ live on the same smooth bundle, there are Chern-Simons forms $\text{CS}(A_{\mathcal{K}_\infty}|_\Sigma, A_{\mathcal{K}_i}|_\Sigma)$ on Σ such that

$$\text{ch}_2(\mathcal{K}_\infty, k_\infty) - \text{ch}_2(\mathcal{K}_i, k) = d \text{CS}(A_{\mathcal{K}_i}|_\Sigma, A_{\mathcal{K}_\infty}|_\Sigma).$$

We will write $B_\infty = A_{\mathcal{K}_\infty} \oplus A_\infty$ which is the Chern connection associated to $(\text{Gr}_\infty \mathcal{H})^{\vee\vee} = (\mathcal{K}_\infty \oplus \mathcal{E}_\infty, k_\infty \oplus h_\infty)$, and B_i for the Chern connection on $\text{Gr}(\mathcal{H}) = (\mathcal{K}_i \oplus \mathcal{E}_i, k \oplus h)$. We will furthermore write B for the Chern connection of (\mathcal{H}, H) . As in (4.11),

$$\text{ch}_2(\mathcal{K}_i \oplus \mathcal{E}_i, k \oplus h) - \text{ch}_2(\mathcal{H}, H) = dd^c \Psi_i = d \text{CS}(B_i, B) .$$

Now for each of the $z_k \in Z \cap \Sigma$, take a ball $B_{2\varepsilon}(z_k)$, choosing ε small enough so that $B_{2\varepsilon}(z_k) \subset X - \mathbf{C}$ (which can again be achieved by part (ii)), and $B_{2\varepsilon}(z_k) \cap Z^{alg} \subset Z - \text{sing}((\text{Gr}_\infty \mathcal{H})^{\vee\vee})$. Then if we let ψ be a cut-off function which is 1 on the $B_\varepsilon(z_k)$ and supported in the $B_{2\varepsilon}(z_k)$. We define admissible metrics $\hat{k}_\infty = \psi k_\infty + (1 - \psi)\hat{k}_0$ on $\hat{\mathcal{K}}_\infty$ and $\hat{h}_\infty = \psi h_\infty + (1 - \psi)\hat{h}_0$ on $\hat{\mathcal{E}}_\infty := \hat{\mathcal{Q}}_\infty^{\vee\vee}$ where \hat{k}_0 and \hat{h}_0 arbitrary admissible metrics on $\hat{\mathcal{K}}_\infty$ and $\hat{\mathcal{E}}_\infty$ respectively. Denote by $A_{\hat{\mathcal{K}}_\infty}$ and \hat{A}_∞ the Chern connections of these metrics. By construction we have connection $A_{\hat{\mathcal{K}}_\infty} = A_{\mathcal{K}_\infty}$ and $\hat{A}_\infty = A_\infty$ on $\cup_{z_k} B_\varepsilon(z_k)$. Proposition 4.6 gives an equation of currents

$$\text{ch}_2(\hat{\mathcal{K}}_\infty \oplus \hat{\mathcal{E}}_\infty, A_{\hat{\mathcal{K}}_\infty} \oplus \hat{A}_\infty) - \text{ch}_2(\pi^* \mathcal{H}, \pi^* B) = \sum_j \hat{m}_j^{alg} \hat{W}_j + dd^c \hat{\Psi} . \quad (4.13)$$

On $\hat{\Sigma}$ and away from \hat{Z} , we also have $dd^c \hat{\Psi} = d \text{CS}(\hat{B}, \hat{B}_\infty)$, where $\hat{B}_\infty = A_{\hat{\mathcal{K}}_\infty} \oplus \hat{A}_\infty$ and $\hat{B} = \pi^* B$. Now by Lemma 3.12 and (4.13), we obtain:

$$\begin{aligned} ([\Sigma] \cdot [Z])m_Z &= \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \{ \text{tr}(F_B \wedge F_B) - \text{tr}(F_{B_\infty} \wedge F_{B_\infty}) \} \\ &\quad - \frac{1}{8\pi^2} \int_{\Sigma \cap \partial(\cup_{z_k} B_\varepsilon(z_k))} d^c \hat{\Psi} \\ &= \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \{ \text{tr}(F_B \wedge F_B) - \text{tr}(F_{B_\infty} \wedge F_{B_\infty}) \} \\ &\quad + \frac{1}{8\pi^2} \int_{\hat{\Sigma} \cap (\cup_{z_k} B_\varepsilon(z_k))^c} dd^c \hat{\Psi} \\ &= \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \{ \text{tr}(F_B \wedge F_B) - \text{tr}(F_{B_\infty} \wedge F_{B_\infty}) \} \\ &\quad + \frac{1}{8\pi^2} \int_{\hat{\Sigma} \cap (\cup_{z_k} B_\varepsilon(z_k))^c} d \text{CS}(\hat{B}, \hat{B}_\infty) \\ &= \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \{ \text{tr}(F_B \wedge F_B) - \text{tr}(F_{A_{\mathcal{K}_\infty}} \wedge F_{A_{\mathcal{K}_\infty}}) \} \\ &\quad - \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \text{tr}(F_{A_\infty} \wedge F_{A_\infty}) - \frac{1}{8\pi^2} \int_{\Sigma \cap \partial(\cup_{z_k} B_\varepsilon(z_k))} \text{CS}(B, B_\infty) . \end{aligned}$$

Similarly, we have

$$\begin{aligned}
& \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \left\{ \text{tr}(F_B \wedge F_B) - \text{tr}(F_{A_{\mathcal{K}_i}} \wedge F_{A_{\mathcal{K}_i}}) - \text{tr}(F_{A_i} \wedge F_{A_i}) \right\} \\
&= \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \left\{ \text{tr}(F_B \wedge F_B) - \text{tr}(F_{B_i} \wedge F_{B_i}) \right\} \\
&= \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} d\text{CS}(B, B_i) = \frac{1}{8\pi^2} \int_{\Sigma \cap \partial(\cup_{z_k} B_\varepsilon(z_k))} \text{CS}(B, B_i) .
\end{aligned}$$

Hence,

$$\begin{aligned}
([\Sigma] \cdot [Z])m_Z &= \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \left\{ \text{tr}(F_{A_{\mathcal{K}_i}} \wedge F_{A_{\mathcal{K}_i}}) - \text{tr}(F_{A_{\mathcal{K}_\infty}} \wedge F_{A_{\mathcal{K}_\infty}}) \right\} \\
&+ \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \left(\text{tr}(F_{A_i} \wedge F_{A_i}) - \text{tr}(F_{A_\infty} \wedge F_{A_\infty}) \right) \\
&- \frac{1}{8\pi^2} \int_{\Sigma \cap \partial(\cup_{z_k} B_\varepsilon(z_k))} \text{CS}(B, B_\infty) + \frac{1}{8\pi^2} \int_{\Sigma \cap \partial(\cup_{z_k} B_\varepsilon(z_k))} \text{CS}(B, B_i) \\
&= \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \left(\text{tr}(F_{A_i} \wedge F_{A_i}) - \text{tr}(F_{A_\infty} \wedge F_{A_\infty}) \right) \\
&+ \frac{1}{8\pi^2} \int_{\Sigma \cap \partial(\cup_{z_k} B_\varepsilon(z_k))} \left(\text{CS}(B, B_i) - \text{CS}(B, B_\infty) + \text{CS}(A_{\mathcal{K}_i}|_\Sigma, A_{\mathcal{K}_\infty}|_\Sigma) \right) \\
&= \frac{1}{8\pi^2} \int_{\Sigma \cap (\cup_{z_k} B_\varepsilon(z_k))} \left(\text{tr}(F_{A_i} \wedge F_{A_i}) - \text{tr}(F_{A_\infty} \wedge F_{A_\infty}) \right) \\
&+ \frac{1}{8\pi^2} \int_{\Sigma \cap \partial(\cup_{z_k} B_\varepsilon(z_k))} \left(\text{CS}(B_\infty, B_i) + \text{CS}(A_{\mathcal{K}_i}|_\Sigma, A_{\mathcal{K}_\infty}|_\Sigma) \right) ,
\end{aligned}$$

where we have used Lemma 3.14 (1). Now by Lemma 3.13, the first term on the right hand side above converges to $([\Sigma] \cdot [Z])m_Z^{an}$ as $i \rightarrow \infty$, whereas the expressions for $\text{CS}(A_{\mathcal{K}_\infty}|_\Sigma, A_{\mathcal{K}_i}|_\Sigma)$ and $\text{CS}(B_i, B_\infty)$ show that the second integral vanishes. Hence, $([\Sigma] \cdot [Z])m_Z = ([\Sigma] \cdot [Z])m_Z^{an}$. Since $[\Sigma] \cdot [Z] \neq 0$, this implies $m_Z = m_Z^{an}$. \square

We now have the

Proof of Proposition 4.7. It follows from Proposition 4.9 that $\mathcal{C}^{alg} = \mathcal{C}_{\mathcal{F}_\infty} = \mathcal{C}_{\mathcal{T}_\infty}$ is a subcycle of \mathcal{C}^{an} , so we may write

$$\mathcal{C}^{an} - \mathcal{C}^{alg} = \sum_j (m_j^{an} - m_j^{alg}) Z_j^{an} , \tag{4.14}$$

with $m_j^{an} \geq m_j^{alg}$. On the other hand, by Lemma 4.8, $[\mathcal{C}^{an} - \mathcal{C}^{alg}]$ vanishes in cohomology. Hence,

$$0 = [\mathcal{C}^{an} - \mathcal{C}^{alg}] \cdot [\omega^{n-2}] = \sum_j (m_j^{an} - m_j^{alg}) \int_{Z_j} \omega^{n-2} \geq \sum_j (m_j^{an} - m_j^{alg}),$$

and therefore $m_j^{alg} = m_j^{an}$ for all j . This proves the first statement. The second statement follows as in the proof of Lemma 3.9. \square

4.4. Continuity of $\bar{\Phi}$. We now put together the results of the previous several subsections to prove continuity of the map $\bar{\Phi}$. The main result of this section is the following.

Theorem 4.11. *The map $\bar{\Phi} : \bar{M}^\mu \rightarrow \widehat{M}_{\text{HYM}}$ defined at the end of Section 4.1 is continuous and surjects onto \bar{M}_{HYM} .*

We begin with the following.

Lemma 4.12. *The images of the maps $\bar{\Psi}$ and $\bar{\Phi}$ lie in the gauge theoretic compactification \bar{M}_{HYM} .*

Proof. It follows directly from Propositions 4.5 and 4.7 that if $q_i \rightarrow q_\infty$ in $\bar{\mathcal{Z}}_\circ$, where $q_i : \mathcal{H} \rightarrow \mathcal{F}_i \rightarrow 0$ is a locally free, stable quotient for each i , then $\lim_{i \rightarrow \infty} \bar{\Psi}(q_i) = \bar{\Psi}(q_\infty)$. Since $\bar{\Psi}(q_i) \in M_{\text{HYM}}^*$ by construction, it therefore follows that $\bar{\Psi}(q_\infty) \in \bar{M}_{\text{HYM}}$. Since any element of $\bar{\mathcal{Z}}_\circ$ is a limit of this kind, this implies that $\text{im}(\bar{\Psi}) \subset \bar{M}_{\text{HYM}}$ and hence that $\text{im}(\bar{\Phi}) \subset \bar{M}_{\text{HYM}}$ by (4.2). \square

Proposition 4.13. *The map $\bar{\Psi} : \bar{\mathcal{Z}}_\circ \rightarrow \bar{M}_{\text{HYM}}$ is continuous.*

Proof. Let $q_i : \mathcal{H} \rightarrow \mathcal{F}_i \rightarrow 0$ be a sequence of quotients in $\bar{\mathcal{Z}}_\circ$ converging to $q_\infty : \mathcal{H} \rightarrow \mathcal{F}_\infty \rightarrow 0$. We must show $\bar{\Psi}(q_i) \rightarrow \bar{\Psi}(q_\infty)$. By compactness it suffices to show that any subsequential limit of $\bar{\Psi}(q_i)$ equals $\bar{\Psi}(q_\infty)$. Hence, without loss of generality we can assume $\bar{\Psi}(q_i)$ has an Uhlenbeck limit in \bar{M}_{HYM} . Now we don't assume that the \mathcal{F}_i are locally free, but for each i there is a sequence of quotients $q_{i,k} : \mathcal{H} \rightarrow \mathcal{F}_{i,k} \rightarrow 0$ with $\mathcal{F}_{i,k}$ locally free and μ -stable of topological type E , so that $q_{i,k} \rightarrow q_i$ in $\bar{\mathcal{Z}}_\circ$. As in the proof of the previous lemma, we have $\lim_{k \rightarrow \infty} \bar{\Psi}(q_{i,k}) = \bar{\Psi}(q_i)$. Choose a subsequence $\{k_i\}$ such that $q_{i,k_i} \rightarrow q_\infty$ in $\bar{\mathcal{Z}}_\circ$. As in the proof of Proposition 3.20, we may arrange that $\bar{\Psi}(q_{i,k_i})$ has the same Uhlenbeck limit as $\bar{\Psi}(q_i)$. Again using Propositions 4.5 and 4.7 we conclude that $\bar{\Psi}(q_{i,k_i}) \rightarrow \bar{\Psi}(q_\infty)$. The result follows. \square

Proof of Theorem 4.11. We claim that $\pi|_{\bar{\mathcal{Z}}_\circ} : \bar{\mathcal{Z}}_\circ \rightarrow \bar{M}^\mu$ is a quotient map in the analytic topology. Once this is established, $\bar{\Phi}$ is continuous if and only if $\bar{\Psi}$ is (cf. (4.2)), and the latter holds by Proposition 4.13. For surjectivity, we

then note that by Proposition 3.20, any point in $\overline{M}_{\text{HYM}}$ is a limit of points in the image of Φ . Since \overline{M}^μ is compact and $\overline{\Phi}$ is continuous, surjectivity follows as well.

In order to see that $\pi|_{\overline{\mathcal{Z}}_o} : \overline{\mathcal{Z}}_o \rightarrow \overline{M}^\mu$ is a quotient map in the analytic topology we remark that as in the proof of [7, Thm. A.7] it follows from Langton's theorem, from the defining properties of the Quot scheme, and from the observation that any holomorphic map from a smooth variety into a reduced scheme Y lifts to the weak normalisation⁶ Y^{wn} of Y , that the pair $(\overline{\mathcal{Z}}_o^{\text{cl}}, \pi|_{\overline{\mathcal{Z}}_o} : \overline{\mathcal{Z}}_o \rightarrow \overline{M}^\mu)$, where $\overline{\mathcal{Z}}_o^{\text{cl}}$ denotes the closure of $\overline{\mathcal{Z}}_o$ inside the weak normalisation of the Quot scheme, fulfills property (\mathfrak{L}) in the sense of [7, Def. A.8]. The claim hence follows from [7, Proposition A.9] \square

Remark 4.14. As noted in the Introduction, since $\overline{M}_{\text{HYM}}$ is the image of a compact space under a continuous map, it is in fact compact rather than just sequentially compact.

Corollary 4.15 (cf. [38, Thm. 5]). *If $\dim X = 2$, the map $\overline{\Phi} : \overline{M}^\mu \rightarrow \overline{M}_{\text{HYM}}$ is a homeomorphism.*

Proof. Theorem 4.11 implies that $\overline{\Phi}$ is continuous, proper, and onto; hence in order to prove the lemma it suffices to show that it is one-to-one. This however follows from [30, Thm. 5.5 and Prop. 5.8]. \square

Remark 4.16. As a consequence of the previous corollary, on surfaces we may identify the two moduli spaces \overline{M}^μ and $\overline{M}_{\text{HYM}}$ using the map $\overline{\Phi}$. In particular, we may think of elements of \overline{M}^μ as equivalence classes $[(\mathcal{E}, \mathcal{C})]$ of pairs. Moreover, we note that it follows from the corollary that in the surface case \overline{M}^μ is homeomorphic (but not necessarily biholomorphic) to the compactification constructed by Jun Li in [38, §3], the potential difference in complex structures stemming from the fact that we consider the determinant line bundle on the weak normalisation of the Quot scheme, not the Quot scheme itself.

5. A COMPLEX STRUCTURE ON $\overline{M}_{\text{HYM}}$

In this section we prove Theorem 1.2 by analysing the natural finite equivalence relation on the space \overline{M}^μ provided by the map $\overline{\Phi}$. We do this by studying sections of the line bundle \mathcal{L}_{n-1} in more detail.

In Section 5.1 we define a natural candidate for the structure sheaf on $\overline{M}_{\text{HYM}}$ and state a criterion for this to be the structure sheaf of an actual complex

⁶Note also that the complex space $(Y^{wn})^{an}$ associated with Y^{wn} is the weak normalisation of the complex space Y^{an} associated with Y .

space. In the following subsections we verify this criterion in our case by constructing certain saturated neighbourhoods (see Definitions 5.7 and 5.6 and Proposition 5.14 of Section 5.3) around each point of \overline{M}^μ and morphisms to projective spaces which are constant on this relation (see Section 5.8 and in particular Lemma 5.24 and Proposition 5.25). These maps are constructed by "lifting" sections from certain well chosen curves which lie inside of certain complete intersection surfaces S in X to obtain sections of \mathcal{L}_{n-1} , (see Sections 5.5 and 5.6). We then check the maps are constant on the relation by comparing them to the corresponding (global) morphism on the moduli space of μ -semistable sheaves on S (see Proposition 5.25 and Section 5.7). The lifting procedure can only be performed away from a certain proper subvariety of $\overline{\mathcal{Z}}_\circ$, due to the loss of flatness of the universal sheaf upon successive restrictions. We therefore require that we be able to construct this non-flatness locus in such a way that the sections we obtain from lifting extend over it. We give various results about restriction and flatness in Sections 5.2 and 5.4. Finally, to show the criterion holds, we need to know that (products of) the locally defined maps in question separate the points of $\overline{M}_{\text{HYM}}$. This will follow from Proposition 5.25 once we know that an element of $\overline{M}_{\text{HYM}}$ is determined by its restriction to a finite number of appropriately chosen complete intersection surfaces S_i (see Corollary 5.13 and Lemma 5.27).

5.1. The fibres of $\overline{\Phi}$ and the equivalence relation on \overline{M}^μ . We begin by introducing some standard terminology regarding equivalence relations.

Definition 5.1. An equivalence relation R on a locally compact (Hausdorff) space X is called *proper* if for every compact set $K \subset X$, its saturation $R(K)$ under the equivalence relation is compact. Equivalently, when endowed with the quotient topology, X/R is locally compact (in particular, Hausdorff), and the quotient map $\pi : X \rightarrow X/R$ is proper. A proper equivalence relation with finite equivalence classes is called *finite*.

Remark 5.2. If X and Y are locally compact spaces, and if $f : X \rightarrow Y$ is a continuous proper map, then the equivalence relation defined by f , i.e., $x \sim y$ if and only if $f(x) = f(y)$, is proper.

Returning to the specific situation at hand, by Proposition 2.23, the map $\overline{\Phi} : \overline{M}^\mu \rightarrow \overline{M}_{\text{HYM}}$ is finite to one. Explicitly, for each $[(\mathcal{E}, \mathcal{C})]$ in $\overline{M}_{\text{HYM}}$, there is an injection $\overline{\Phi}^{-1}([(\mathcal{E}, \mathcal{C})]) \hookrightarrow \pi_0(\chi^{-1}(\mathcal{C}))$ where we recall that

$$\chi : \text{Quot}(\mathcal{E}, \tau_{\mathcal{E}} - \tau_E) \longrightarrow \mathcal{C}_{n-2}(X)$$

is the natural Quot to Chow morphism defined in Section 2.8.

We define an equivalence relation $R \subset \overline{M}^\mu \times \overline{M}^\mu$ on the space \overline{M}^μ , defined by the property $([\mathcal{F}_1], [\mathcal{F}_2]) \in R$ if and only if $[\mathcal{F}_1]$ and $[\mathcal{F}_2]$ are in the same fibre of $\overline{\Phi}$. By Remark 4.14, $\overline{M}_{\text{HYM}}$ is compact, and hence locally compact. By construction there is a bijection $\overline{M}^\mu/R \simeq \overline{M}_{\text{HYM}}$, where \overline{M}^μ/R is the topological quotient of \overline{M}^μ given by identifying all points in an equivalence class of R . By Theorem 4.11 and the universal property of the quotient topology, this is a homeomorphism. We conclude that R is a finite equivalence relation.

In order to put a complex structure on $\overline{M}_{\text{HYM}}$ it therefore suffices to understand the equivalence relation R and under what conditions the quotient of a complex analytic space by a proper equivalence relation can be endowed with a sheaf of functions making it into a complex space. As a first step in this direction, we make X/R into a ringed space: given a proper equivalence relation R on a reduced complex space X and an open subset $U \subset X/R$, we set $\mathcal{O}_{X/R}(U) := \mathcal{O}_X(\pi^{-1}(U))^R$, where the latter is the algebra of R -invariant holomorphic functions on $\pi^{-1}(U)$.

In order to see that the resulting ringed space $(X/R, \mathcal{O}_{X/R})$ is in fact a complex space, we will use the following fundamental result of Henri Cartan, see [8, Main Theorem].

Theorem 5.3 (CARTAN'S CRITERION). *Consider a proper equivalence relation R on a reduced complex space X with quotient map $\pi : X \rightarrow X/R$. In order that the ringed space X/R be a complex space, it suffices that each point of X/R has an open neighbourhood V such that the R -invariant holomorphic maps $\pi^{-1}(V) \rightarrow Z$ (Z being a complex space) separate the equivalence classes in $\pi^{-1}(V)$.*

Then, in order to prove that $\overline{M}_{\text{HYM}}$ is a complex space it suffices to prove that the equivalence classes of the relation $R \subset \overline{M}^\mu \times \overline{M}^\mu$ locally over $\overline{M}_{\text{HYM}}$ can be separated by R -invariant holomorphic maps. We will in fact prove the stronger claim that locally over $\overline{M}_{\text{HYM}}$ there exists holomorphic maps to some complex spaces having as fibres exactly the given equivalence classes (see Claim 5.23). The proof of this fact will occupy the remainder of Section 5.

5.2. Restriction theorems. Throughout the remainder of this section we will need to consider the restriction of sheaves to various subvarieties of X . We begin with the following crucial semistable restriction theorem.

Theorem 5.4 (Langer, [36]). *Let X be a smooth projective variety, and let $\mathcal{O}_X(1)$ be an ample line bundle on X . Let $\{\mathcal{F}_i\}_{i \in \mathcal{S}}$ be a family of μ -(semi)stable sheaves on X with fixed rank and Chern classes. Then, there is a positive integer k_0 which is independent of i , so that for any $k \geq k_0$ and any smooth*

divisor $D \in |\mathcal{O}_X(k)|$, if $(\mathrm{Gr} \mathcal{F}_i)|_D$ is torsion free for some Seshadri graduation of \mathcal{F}_i , then $\mathcal{F}_i|_D$ remains μ -(semi)stable.

This follows immediately from [36, Thm 5.2 and Cor. 5.4] by noticing that the right hand side of the inequality provided there depends only on the rank and Chern classes of the sheaf in question.

A higher-dimensional, slightly weaker analogue of the following result will be proven in Proposition 5.12 and Corollary 5.13 below.

Proposition 5.5. *Let $\{\mathcal{F}_i\}_{i \in \mathcal{S}}$ be a set of μ -semistable sheaves on a polarised surface $(S, \mathcal{O}_S(1))$ with fixed rank and Chern classes. Then there is a positive integer $k_0 \gg 0$ which is independent of i , such that for all $k \geq k_0$ there is a finite subset $\Sigma \subset |\mathcal{O}_S(k)|$ consisting of smooth curves such that for any two sheaves $\mathcal{F}_{i_1}, \mathcal{F}_{i_2}$, there exists a curve $C \in \Sigma$ such that*

- (1) $\mathrm{sing}(\mathrm{Gr}(\mathcal{F}_{i_1})) \cap C = \mathrm{sing}(\mathrm{Gr}(\mathcal{F}_{i_2})) \cap C = \emptyset$,
- (2) $\mathcal{F}_{i_1}|_C$ and $\mathcal{F}_{i_2}|_C$ remain semistable,
- (3) $\mathrm{Gr}(\mathcal{F}_{i_1})^{\vee\vee} \cong \mathrm{Gr}(\mathcal{F}_{i_2})^{\vee\vee}$ if and only if $\mathcal{F}_{i_1}|_C$ and $\mathcal{F}_{i_2}|_C$ are s -equivalent.

Proof. Since the family $\{\mathcal{F}_i\}_{i \in \mathcal{S}}$ is in particular a bounded family, the arguments of Section 2.5.5 imply that the number of singular points of the sheaves $\{\mathrm{Gr}(\mathcal{F}_i)\}_{i \in \mathcal{S}}$ is bounded as i ranges over the set \mathcal{S} . Here we have chosen arbitrary Seshadri filtrations of the \mathcal{F}_i , but we recall that the cycles associated to $\mathrm{Gr}(\mathcal{F}_i)^{\vee\vee}/\mathrm{Gr}(\mathcal{F}_i)$, and in particular the sets $\mathrm{sing}(\mathrm{Gr}(\mathcal{F}_i))$, are independent of this choice. Fix an upper bound m for the maximum number of points of such a singular set. Fix a number k that satisfies the hypotheses of Theorem 5.4, and choose $n > 4m$ curves in $|\mathcal{O}_X(k)|$ so that no three have a common intersection point. Then for any two fixed sheaves \mathcal{F}_{i_1} and \mathcal{F}_{i_2} , there must be a curve C among the chosen curves that misses the singularities of both $\mathrm{Gr}(\mathcal{F}_{i_1})$ and $\mathrm{Gr}(\mathcal{F}_{i_2})$, and in particular $(\mathrm{Gr} \mathcal{F}_{i_1})|_C$ and $(\mathrm{Gr} \mathcal{F}_{i_2})|_C$ are locally free (and in particular torsion free) and therefore $\mathcal{F}_{i_1}|_C$ and $\mathcal{F}_{i_2}|_C$ are semistable. Moreover, the summands of each of $\mathrm{Gr}(\mathcal{F}_{i_1})|_C$ and $\mathrm{Gr}(\mathcal{F}_{i_2})|_C$ remain slope stable, all of equal slope, and are the quotients of the restricted filtrants of the respective Seshadri filtrations. This means that the restrictions of the Seshadri filtrations to C remain Seshadri filtrations for $\mathcal{F}_{i_1}|_C$ and $\mathcal{F}_{i_2}|_C$. Notice that on C , the notion of Seshadri filtration is the same for slope or GM-stability. We therefore obtain

$$\begin{aligned} \mathrm{Gr}(\mathcal{F}_{i_1})^{\vee\vee}|_C \cong \mathrm{Gr}(\mathcal{F}_{i_2})^{\vee\vee}|_C &\iff \mathrm{Gr}(\mathcal{F}_{i_1})|_C \cong \mathrm{Gr}(\mathcal{F}_{i_2})|_C \\ &\iff \mathrm{Gr}(\mathcal{F}_{i_1}|_C) \cong \mathrm{Gr}(\mathcal{F}_{i_2}|_C) \iff \mathrm{gr}(\mathcal{F}_{i_1}|_C) \cong \mathrm{gr}(\mathcal{F}_{i_2}|_C) . \end{aligned}$$

We claim that

$$\mathrm{Gr}(\mathcal{F}_{i_1})^{\vee\vee}|_C \cong \mathrm{Gr}(\mathcal{F}_{i_2})^{\vee\vee}|_C \iff \mathrm{Gr}(\mathcal{F}_{i_1})^{\vee\vee} \cong \mathrm{Gr}(\mathcal{F}_{i_2})^{\vee\vee} .$$

One direction is obvious. For the other, we consider the bounded family of locally free sheaves $\{\mathcal{E}_{i_1, i_2} := \mathcal{H}om(\mathrm{Gr}(\mathcal{F}_{i_1})^{\vee\vee}, \mathrm{Gr}(\mathcal{F}_{i_2})^{\vee\vee})\}_{i_1, i_2 \in \mathcal{S}}$. By possibly increasing the size of k , we can ensure using Serre vanishing and duality that $H^1(S, \mathcal{E}_{i_1, i_2}(-kH)) = 0$ for all choices of i_1 and i_2 . Now considering the exact sequence

$$0 \longrightarrow \mathcal{E}_{i_1, i_2}(-kH) \longrightarrow \mathcal{E}_{i_1, i_2} \longrightarrow \mathcal{E}_{i_1, i_2}|_C \longrightarrow 0 ,$$

and the induced long exact sequence in cohomology, we see that $H^0(\mathcal{E}_{i_1, i_2})$ surjects onto $H^0(\mathcal{E}_{i_1, i_2}|_C)$. Therefore, any isomorphism $\mathrm{Gr}(\mathcal{F}_{i_1})^{\vee\vee}|_C \cong \mathrm{Gr}(\mathcal{F}_{i_2})^{\vee\vee}|_C$ can be lifted to a map $\mathrm{Gr}(\mathcal{F}_{i_1})^{\vee\vee} \rightarrow \mathrm{Gr}(\mathcal{F}_{i_2})^{\vee\vee}$, which turns out to be an isomorphism for example by the arguments laid out in step 3 and 4 of the proof of Proposition 5.1 in the preprint version of [28]. \square

Proposition 5.5 is very close to being the same as [30, Lemma 5.4], but the important point that we wish to bring out here is that we may choose k to be independent of the sheaves in question as we range over a bounded family. The conclusion, however, is somewhat weaker. Namely, for each pair of sheaves in the family we only ask that there be one curve in the chosen finite set having the stated property holds, rather than this holding for a generic curve in the linear system.

5.3. Admissible flags and neighbourhoods in $\overline{M}_{\mathrm{HYM}}$. In this subsection we give some preliminary definitions required to understand the case of higher dimensions in the next subsection. In the following we will try to exploit properties of the moduli space of semistable sheaves on projective surfaces. We will reduce ourselves to the surface case by successively cutting with (smooth) hypersurfaces in X . Recall that a hypersurface X' in X is *regular* for a coherent sheaf \mathcal{F} on X , or *\mathcal{F} -regular*, if the natural morphism $\mathcal{F}(-X') \rightarrow \mathcal{F}$ is injective. This is the case if and only if X' contains none of the associated points of \mathcal{F} (cf. [33, page 8]).

In order to motivate the next definition recall that by [34, Formula V 3.20] the singularity sets of a coherent sheaf \mathcal{S} as defined in [34, V.(5.5)] are related to the supports of the local Ext sheaves as follows:

$$S_m(\mathcal{S}) = \bigcup_{d \geq n-m} \mathrm{supp}(\mathcal{E}xt_X^d(\mathcal{S}, \omega_X)), \quad 0 \leq m < n ;$$

cf. the discussion in Section 2.5.2 above.

Definition 5.6. Let $[(\mathcal{E}, \mathcal{C})]$ be a point in $\overline{M}_{\text{HYM}}$, where \mathcal{E} is a polystable reflexive sheaf and \mathcal{C} is a codimension 2 cycle on X . A smooth surface S embedded in X will be called *admissible for* $[(\mathcal{E}, \mathcal{C})]$ if $\dim(S \cap |\mathcal{C}|) \leq 0$ and $\dim(S \cap \text{sing}(\mathcal{E})) \leq 0$. An admissible surface S for $[(\mathcal{E}, \mathcal{C})]$ will be called *fully admissible for* $[(\mathcal{E}, \mathcal{C})]$ if additionally \mathcal{E} is locally free in a neighbourhood of S .

Definition 5.7. Let $[(\mathcal{E}, \mathcal{C})]$ be a point in $\overline{M}_{\text{HYM}}$, where \mathcal{E} is a polystable reflexive sheaf and \mathcal{C} is a codimension 2 cycle on X . A flag of smooth complete intersections $(X^{(l)})_{1 \leq l \leq n-2}$, $X^{(l)} := \cap_{i=1}^l X_i$, of hypersurfaces X_1, \dots, X_{n-2} in X will be called *admissible for* $[(\mathcal{E}, \mathcal{C})]$ if for each $l = 1, \dots, n-2$ both $X^{(l)} \cap |\mathcal{C}|$ and $X^{(l)} \cap \text{sing}(\mathcal{E})$ are of codimension at least two in $X^{(l)}$. An admissible flag $(X^{(l)})_{1 \leq l \leq n-2}$ for $[(\mathcal{E}, \mathcal{C})]$ will be called *fully admissible for* $[(\mathcal{E}, \mathcal{C})]$ if in addition for each $l = 1, \dots, n-2$ the complete intersection $X^{(l)}$ is $\mathcal{E}|_{X^{(l-1)}}$ -regular and $\mathcal{E}xt_{X^{(l-1)}}^q(\mathcal{E}|_{X^{(l-1)}}, \omega_{X^{(l-1)}})$ -regular for all $q \geq 0$, where we use the notation $X^{(0)} = X$, $X^{(1)} = X'$. The flag $(X^{(l)})_{1 \leq l \leq n-2}$ is said to be of *multi-degree* $(d_l)_{1 \leq l \leq n-2}$ if the hypersurfaces X_l are taken in the linear systems $|\mathcal{O}_X(d_l)|$, where $\mathcal{O}_X(1)$ is a fixed ample line bundle on X .

Definition 5.8. Let \mathcal{F} be a torsion free sheaf on X . Similar to the above, a flag of complete intersections $(X^{(l)})_{1 \leq l \leq n-2}$, $X^{(l)} := \cap_{i=1}^l X_i$ of smooth general hypersurfaces X_1, \dots, X_{n-2} in X will be called *fully admissible for* \mathcal{F} if for each $l = 1, \dots, n-2$ the complete intersection $X^{(l)}$ is $\mathcal{F}|_{X^{(l-1)}}$ -regular and $\mathcal{E}xt_{X^{(l-1)}}^q(\mathcal{F}|_{X^{(l-1)}}, \omega_{X^{(l-1)}})$ -regular for all $q \geq 0$. Every member $X^{(l)}$ of a fully admissible flag will be called *fully admissible for* \mathcal{F} .

Remark 5.9. For later usage, we note the following facts:

- (1) A flag of complete intersections $(X^{(l)})_{1 \leq l \leq n-2}$, $X^{(l)} := \cap_{i=1}^l X_i$ of smooth general hypersurfaces X_1, \dots, X_{n-2} in X is admissible for $[(\mathcal{E}, \mathcal{C})]$ if and only if its last member $X^{(n-2)}$ is admissible for $[(\mathcal{E}, \mathcal{C})]$.
- (2) If a flag $(X^{(l)})_{1 \leq l \leq n-2}$ is fully admissible for $[(\mathcal{E}, \mathcal{C})]$, then its last member $X^{(n-2)}$ is fully admissible for $[(\mathcal{E}, \mathcal{C})]$.
- (3) If a flag $(X^{(l)})_{1 \leq l \leq n-2}$ is admissible for $[(\mathcal{E}, \mathcal{C})]$, then the condition that $X^{(l)}$ contains no irreducible components of $|\mathcal{C}| \cap X^{(l-1)}$ implies that $|\mathcal{C}| \cap X^{(l)}$ is the support of a codimension 2 cycle on $X^{(l)}$, cf. [25, Sect. 2.3]: Indeed, the first hypersurface $X' \subset X$, whose associated Cartier divisor we will denote by \mathcal{D} , contains no irreducible component Z_j of $|\mathcal{C}|$. Hence, $\mathcal{C} \cdot \mathcal{D} := \sum_j n_j (Z_j \cap X')$ is a well-defined Cartier divisor on $|\mathcal{C}|$. On the other hand, $\mathcal{C} \cdot \mathcal{D}$ also defines a well-defined intersection cycle on X' , which we will also denote by $\mathcal{C}|_{X'}$. Since it is a divisor on $|\mathcal{C}|$ it must be a codimension 2 cycle \mathcal{C}' on X' . The hypothesis now similarly allows us to inductively construct codimension 2 cycles $\mathcal{C}^{(l)} :=$

$\mathcal{C}^{(l-1)} \cdot \mathcal{D}^{(l)}$ on $X^{(l)}$, where $\mathcal{D}^{(l)}$ is the divisor associated to the hypersurface $X^{(l)} \subset X^{(l-1)}$.

(4) If a flag $(X^{(l)})_{1 \leq l \leq n-2}$ is fully admissible for $[(\mathcal{E}, \mathcal{C})]$, then the restrictions $\mathcal{E}|_{X^{(l)}}$ are reflexive sheaves on their respective supports $X^{(l)}$ by [33, Cor. 1.1.14]. We will denote the induced pair $(\mathcal{E}|_{X^{(l)}}, \mathcal{C}^{(l)})$ consisting of a polystable sheaf and a codimension 2 cycle on $X^{(l)}$ by $(\mathcal{E}, \mathcal{C})|_{X^{(l)}}$.

(5) If a flag $(X^{(l)})_{1 \leq l \leq n-2}$ is fully admissible for a torsion free sheaf \mathcal{F} , then for each $l = 1, \dots, n-2$ the intersection $X^{(l)} \cap \text{sing}(\mathcal{F})$ is of codimension at least two in $X^{(l)}$.

The following result establishes a link between restriction of cycles and cycles associated with restricted sheaves via the construction given in Section 2.5.3. For simplicity we state it only for two-codimensional coherent sheaves.

Lemma 5.10. *Let \mathcal{A} be a pure coherent sheaf of codimension two on X and X' a smooth hyperplane section of X not containing any irreducible component of the support cycle $\mathcal{C}_{\mathcal{A}}$. Then, we have the following equality of cycles on X' :*

$$\mathcal{C}_{(\mathcal{A}|_{X'})} = (\mathcal{C}_{\mathcal{A}})|_{X'} . \quad (5.1)$$

Proof. Using the definition of multiplicity as in [30, Rem. 5.3], the statement is easily checked when X' is a general element in its linear system.

Suppose now that X' contains no irreducible component of the support cycle $\mathcal{C}_{\mathcal{A}}$ of \mathcal{A} . This and the purity of \mathcal{A} imply the existence of an exact sequence of the form

$$0 \rightarrow \mathcal{A}(-X') \rightarrow \mathcal{A} \rightarrow \mathcal{A}|_{X'} \rightarrow 0 .$$

Each element X'_t in a sufficiently small Zariski open neighbourhood T of X' in its complete linear system will be smooth and will not contain any irreducible component of the support cycle $\mathcal{C}_{\mathcal{A}}$. By the above and by [33, Lemma 2.1.4] it follows that the family of sheaves $\mathcal{A}|_{X'_t}$ is flat over T .

By the observation made at the beginning of the proof, for $t \in T$ general we have the equality

$$\mathcal{C}_{(\mathcal{A}|_{X'_t})} = (\mathcal{C}_{\mathcal{A}})|_{X'_t} . \quad (5.2)$$

Moreover, support cycles vary continuously in flat families of coherent sheaves, cf. [5], as do the intersection cycles appearing on the right hand side of (5.2), cf. part (3) of Remark 5.9. Thus, we get the desired equality (5.1) by passing to the limit on both sides of (5.2). \square

Next, we will introduce a terminology that is useful in formulating some of the technical results below. If $|H|$ is a basepoint free linear system on X and $Y \subset X$ is a subvariety, we will denote by $|H|_Y$ the restricted linear

system, i.e., the divisors arising from sections in the image of the restriction map $H^0(X, \mathcal{O}_X(H)) \rightarrow H^0(Y, (\mathcal{O}_X(H))|_Y)$.

Definition 5.11. Let $\mathcal{O}_X(1)$ be an ample line bundle on X with corresponding ample divisor H , and d_1, \dots, d_{n-2} positive natural numbers. A property P is said to hold for *any sufficiently general tuple* $(X_1, \dots, X_{n-2}) \in |d_1 H| \times \dots \times |d_{n-2} H|$ if there exists a nonempty open subset $U_1 \subset |d_1 H|$ such that for every $X_1 \in U_1$, there exists a nonempty open subset $U_2 \subset |d_2 H|_{X_1}$ such that for every $X_2 \in U_2$, there exists a nonempty open subset $U_3 \subset |d_3 H|_{X^{(2)}}$, ..., such that for every $X_{n-2} \in U_{n-2}$ property P holds for the tuple (X_1, \dots, X_{n-2}) .

If H is a very ample divisor on X , if d_1, \dots, d_{n-2} are given positive natural numbers, and if \mathcal{F} is a torsion free sheaf on X , then, since the restriction of $|dH|$ to any subvariety of X stays basepoint free for any positive natural number d , prime avoidance holds for these restricted linear systems, and hence any sufficiently general tuple $(X_1, \dots, X_{n-2}) \in |d_1 H| \times \dots \times |d_{n-2} H|$ is fully admissible for \mathcal{F} .

The following two results are higher-dimensional analogues of Proposition 5.5 above and are important for producing the holomorphic maps required by Cartan's criterion, Theorem 5.3; see the proof of Claim 5.23 below.

Proposition 5.12. *Let H be a very ample divisor on X , and let $k \geq 2$ be a positive integer. Let moreover \mathcal{S} be a bounded set of reflexive sheaves on X and let $[(\mathcal{E}_0, \mathcal{C}_0)]$ be a point in $\overline{M}_{\text{HYM}}$, where \mathcal{E}_0 is a polystable reflexive sheaf and \mathcal{C}_0 is a codimension 2 cycle on X . Then, there exists a positive natural number N such that for any sequence d_1, \dots, d_{n-2} of positive natural numbers and for any choice of N sufficiently general tuples $(X_1, \dots, X_{n-2}) \in |d_1 H| \times \dots \times |d_{n-2} H|$ with associated flags $((X^{(l)})_{1 \leq l \leq n-2})$ the following holds:*

- (1) *The flags are fully admissible for $[(\mathcal{E}_0, \mathcal{C}_0)]$.*
- (2) *If Σ is the chosen finite set of flags, then for any choice of $k-2$ sheaves $\mathcal{E}_3, \dots, \mathcal{E}_k$ from \mathcal{S} as well as any choice of two points $[(\mathcal{E}_1, \mathcal{C}_1)]$ and $[(\mathcal{E}_2, \mathcal{C}_2)]$ from $\overline{M}_{\text{HYM}}$ some element of Σ is in addition fully admissible for $[(\mathcal{E}_1, \mathcal{C}_1)]$ and $[(\mathcal{E}_2, \mathcal{C}_2)]$ as well as fully admissible for all the sheaves $\mathcal{E}_3, \dots, \mathcal{E}_k$.*

Proof. Let n be the dimension of X , which we may suppose to be bigger than 2. We start by choosing the first members $X' \in |d_1 H|$ of the flags in Σ . The first constraint on X' is given by the full admissibility for $[(\mathcal{E}_0, \mathcal{C}_0)]$, which is a Zariski open condition. In order to additionally ensure admissibility for $[(\mathcal{E}_1, \mathcal{C}_1)]$ and $[(\mathcal{E}_2, \mathcal{C}_2)]$ as well as for $\mathcal{E}_3, \dots, \mathcal{E}_k$, it is sufficient that X' contains neither of the associated points of the sheaves $\mathcal{E}xt_X^q(\mathcal{E}_i, \omega_X)$, $i \in \{1, \dots, k\}$, $q \geq 0$, nor any of

the irreducible components of the cycles $\mathcal{C}_1, \mathcal{C}_2$; cf. part (5) of Remark 5.9. For an element $\mathcal{E} \in \mathcal{S}$ denote by $A(\mathcal{E})$ the set of subvarieties of X arising from such associated components, and define $A(\tilde{\mathcal{E}}, \tilde{\mathcal{C}})$ analogously for $(\tilde{\mathcal{E}}, \tilde{\mathcal{C}}) \in \overline{M}_{\text{HYM}}$. Since \mathcal{E} and $(\tilde{\mathcal{E}}, \tilde{\mathcal{C}})$ run through bounded families, the arguments of Section 2.5.5 imply that the number of elements of the corresponding sets $A(\mathcal{E})$ and $A(\tilde{\mathcal{E}}, \tilde{\mathcal{C}})$ are simultaneously bounded by some $m_1 \in \mathbb{N}$.

Choose $n_1 > knm_1$ smooth hypersurfaces in $|d_1H|$ fully admissible for $(\mathcal{E}_0, \mathcal{C}_0)$ and such that no $n+1$ of them have a common intersection point, both of which are open conditions. Then, it follows that for any choice of two points $(\mathcal{E}_1, \mathcal{C}_1)$ and $(\mathcal{E}_2, \mathcal{C}_2)$ from $\overline{M}_{\text{HYM}}$ and any choice of $k-2$ sheaves $\mathcal{E}_3, \dots, \mathcal{E}_k$ from \mathcal{S} , there must be a hypersurface X' among the chosen ones containing none of the elements of $A(\mathcal{E}_1, \mathcal{C}_1) \cup A(\mathcal{E}_2, \mathcal{C}_2) \cup A(\mathcal{E}_3) \cup \dots \cup A(\mathcal{E}_k)$.

The argument for higher codimension is analogous and hence omitted. \square

The following is a consequence of the proof of Proposition 5.12.

Corollary 5.13. *Let H be a very ample divisor on X , and let $[(\mathcal{E}_0, \mathcal{C}_0)]$ be a point in $\overline{M}_{\text{HYM}}$, where \mathcal{E}_0 is a polystable reflexive sheaf and \mathcal{C}_0 is a codimension 2 cycle on X . Then, there exist a positive natural numbers d_0 and N such that for any tuple $(d_1, d_2, \dots, d_{n-2})$ of natural numbers $d_j \geq d_0$ and for any choice of N sufficiently general tuples $(X_1, \dots, X_{n-2}) \in |d_1H| \times \dots \times |d_{n-2}H|$ with associated flags $((X^{(l)})_{1 \leq l \leq n-2})$ the following holds:*

- (1) *All the flags are fully admissible for $[(\mathcal{E}_0, \mathcal{C}_0)]$.*
- (2) *If Σ is the chosen finite set of flags, then for any choice of two points $[(\mathcal{E}_1, \mathcal{C}_1)]$ and $[(\mathcal{E}_2, \mathcal{C}_2)]$ from $\overline{M}_{\text{HYM}}$ some element $(X^{(l)})$ of Σ is in addition fully admissible for $[(\mathcal{E}_1, \mathcal{C}_1)]$ and $[(\mathcal{E}_2, \mathcal{C}_2)]$.*
- (3) *If $X^{(n-2)}$ is a 2-dimensional member of any of the chosen flags, we have*

$$\mathcal{E}_1 \cong \mathcal{E}_2 \quad \text{if and only if} \quad \mathcal{E}_1|_{X^{(n-2)}} \cong \mathcal{E}_2|_{X^{(n-2)}}. \quad (5.3)$$

Proof. As the reflexive sheaves \mathcal{E} appearing as first entries for a tuple $[(\mathcal{E}, \mathcal{C})] \in \overline{M}_{\text{HYM}}$ range over a bounded set \mathcal{S} , the corresponding set of reflexive homomorphism sheaves $\mathcal{H}om(\mathcal{E}_1, \mathcal{E}_2)$ is likewise bounded. We may therefore choose a positive natural number d_0 such that the conclusion of the "general Enriques-Zariski Lemma", [43, Prop. 3.2], holds for any of the sheaves $\mathcal{H}om(\mathcal{E}_1, \mathcal{E}_2)$, where \mathcal{E}_1 and \mathcal{E}_2 are in \mathcal{S} .

Let $d_1 \geq d_0$. Then, the first step of the proof of Proposition 5.12 shows that for a general choice of n_1 smooth hypersurfaces $X'_j \in |d_1H|$ the following property holds: for any choice $[(\mathcal{E}_1, \mathcal{C}_1)], [(\mathcal{E}_2, \mathcal{C}_2)] \in \overline{M}_{\text{HYM}}$, there exists an index j such that X'_j is fully admissible for $[(\mathcal{E}_1, \mathcal{C}_1)], [(\mathcal{E}_2, \mathcal{C}_2)]$ and also for $\mathcal{H}om(\mathcal{E}_1, \mathcal{E}_2)$. It follows that $\mathcal{E}_1|_{X'_j}, \mathcal{E}_2|_{X'_j}$, hence $\mathcal{H}om(\mathcal{E}_1|_{X'_j}, \mathcal{E}_2|_{X'_j})$, and also

$\mathcal{H}om(\mathcal{E}_1, \mathcal{E}_2)|_{X'_j}$ are reflexive. Moreover, admissibility implies that X'_j does not contain any component of the singularity set of \mathcal{E}_1 or \mathcal{E}_2 , and therefore the reflexive sheaves $\mathcal{H}om(\mathcal{E}_1|_{X'_j}, \mathcal{E}_2|_{X'_j})$ and $\mathcal{H}om(\mathcal{E}_1, \mathcal{E}_2)|_{X'_j}$ agree on an open subset of X'_j whose complement has codimension at least 2 in X'_j , and hence they coincide. Therefore, if $\mathcal{E}_1|_{X'_j} \cong \mathcal{E}_2|_{X'_j}$ via an isomorphism φ , by the choice of d_1 and by [43, Prop. 3.2] we can lift φ to a homomorphism $\Phi \in \mathcal{H}om(\mathcal{E}_1, \mathcal{E}_2)$, which turns out to be an isomorphism for example by the arguments laid out in step 3 and 4 of the proof of Proposition 5.1 in the preprint version of [28]. We conclude that $\mathcal{E}_1 \cong \mathcal{E}_2$ if and only if $\mathcal{E}_1|_{X'_j} \cong \mathcal{E}_2|_{X'_j}$.

Noticing that the set of reflexive sheaves $\{\mathcal{E}|_{X'_j} \mid X'_j \text{ is fully admissible for } \mathcal{E}\}$ is again bounded, we may argue in a similar fashion to see that the claims holds for any sufficiently general tuple $(X_1, \dots, X_{n-2}) \in |d_1 H| \times \dots \times |d_{n-2} H|$, as long as $d_j \geq d_0$ for all $j = 1, \dots, n-2$. \square

The following crucial result allows us to define open neighbourhoods of a given equivalence class in \overline{M}^μ on which the equivalence relation induced by the comparison map $\overline{\Phi}$ is controlled by invariant holomorphic functions.

Proposition 5.14. *Let S be a surface embedded in X . Then the set*

$$U'_S := \{[(\mathcal{E}, \mathcal{C})] \in \overline{M}_{\text{HYM}} \mid S \text{ is admissible for } [(\mathcal{E}, \mathcal{C})]\} \subset \overline{M}_{\text{HYM}}$$

is open in $\overline{M}_{\text{HYM}}$ with respect to the topology defined in Section 3.5. In particular, the preimage $U_S := \overline{\Phi}^{-1}(U'_S)$ is open and R -saturated in \overline{M}^μ .

Proof. By the definition of the topology of $\overline{M}_{\text{HYM}}$, the claim is equivalent to the statement that if S is admissible for $[(\mathcal{E}, \mathcal{C})]$ and $[(\mathcal{E}_i, \mathcal{C}_i)]$ is a sequence of (isomorphism classes of) ideal connections converging to $[(\mathcal{E}, \mathcal{C})]$, then S remains admissible for $[(\mathcal{E}_i, \mathcal{C}_i)]$ for i sufficiently large. This in turn breaks into two separate conditions on the pairs $[(\mathcal{E}_i, \mathcal{C}_i)]$: first, that $\dim(S \cap |\mathcal{C}_i|) \leq 0$, and second, that $\dim(S \cap \text{sing}(\mathcal{E}_i)) \leq 0$, whenever these two properties are satisfied for \mathcal{C} and \mathcal{E} . Since the \mathcal{C}_i converge to a subcycle of \mathcal{C} in the cycle space, their supports converge to the support of this subcycle in the Hausdorff sense, from which the first item follows. On the other hand, the second item is a consequence of Lemma 3.18. \square

5.4. Extension of sections and flat restriction. Below we construct $\text{SL}(V)$ -invariant sections of the line bundle $\mathcal{L}_{n-1} \rightarrow \mathcal{Z}$ away from a subvariety $\mathcal{T} \subset \mathcal{Z}$. Since we want sections on all of \mathcal{Z} , we will require the following lemma from [30], which gives a criterion for when these sections extend over \mathcal{T} .

Lemma 5.15 ([30, Lemma 2.12]). *Let G be a connected algebraic group, \mathcal{R} a weakly normal G -variety, and $\mathcal{L} \rightarrow \mathcal{R}$ a G -linearised line bundle. Then there*

exists a finite set of closed, irreducible G -invariant subvarieties $\{\mathcal{R}_i\}_{i=1}^m$ of \mathcal{R} with the following property: if $\mathcal{T} \subset \mathcal{R}$ is a closed G -invariant subvariety with $\text{codim}(\mathcal{T} \cap \mathcal{R}_i) \geq 2$ in \mathcal{R}_i , then any section $\sigma \in H^0(\mathcal{R} \setminus \mathcal{T}, \mathcal{L})^G$ extends to a section over all of \mathcal{R} .

The following is a slight modification of [30, Cor 3.2].

Lemma 5.16. *Let G be a connected algebraic group, and $\mathcal{E} \rightarrow X \times \mathcal{R}$ be a G -linearised flat family over a parameter space \mathcal{R} , and H a very ample polarisation. Let $[(\mathcal{E}, \mathcal{C})] \in \overline{M}_{\text{HYM}}$. Suppose $\{\mathcal{R}_j\}_{j=1}^m$ is a finite set of closed, irreducible G -invariant subvarieties of \mathcal{R} so that for each i there is a point $r_i \in \mathcal{R}_i$ such that \mathcal{E}_{r_i} is torsion free. Then, for any choice of positive integers d_1, \dots, d_{n-2} and any sufficiently general tuple $(X_1, \dots, X_{n-2}) \in |d_1 H| \times \dots \times |d_{n-2} H|$ the following holds:*

- (1) *The associated flag $(X^{(l)})_{1 \leq l \leq n-2}$ is fully admissible for $[(\mathcal{E}, \mathcal{C})]$, (and so in particular the end term $S = X^{(n-2)}$ is fully admissible for $[(\mathcal{E}, \mathcal{C})]$).*
- (2) *There exists a closed subvariety $\mathcal{T} \subset \mathcal{R}$ such that for each j we have $\text{codim}_{\mathcal{R}_j}(\mathcal{T} \cap \mathcal{R}_j) \geq 2$, and for each l the family $\mathcal{E}|_{X^{(l)} \times \mathcal{R}}$ remains flat over $\mathcal{R} \setminus \mathcal{T}$.⁷*
- (3) *Each of the sheaves $\mathcal{E}_{r_i}|_{X^{(l)}}$ remains torsion free.*

Proof. We will strictly follow the strategy of [30, proof of Lemma 3.1] with the additional constraint that the flags will be required to be fully admissible for $[(\mathcal{E}, \mathcal{C})]$. In line with Definition 5.11, the conditions we have to impose on the flag in order to fulfill the claims made in the formulation of the Lemma are analysed inductively, in particular, the problem of maintaining flatness of the successive restrictions of the family \mathcal{E} on $X \times \mathcal{R}$ to the members of a flag $X^{(l)} \times \mathcal{R}$. For this purpose, we define

$$\mathcal{R}_{tf} := \{r \in \mathcal{R} \mid \mathcal{F}_r \text{ is torsion free on } X\}.$$

If $X' \subset X$ is a smooth hypersurface, then the restriction of \mathcal{E} to $X' \times \mathcal{R}_{tf}$ remains flat. In fact, flatness of the restricted family at a point $r \in \mathcal{R}$ is equivalent to the \mathcal{E}_r -regularity of X' by [33, Lemma 2.1.4] and this in turn is implied by the torsion freeness of \mathcal{E}_r .

We wish to construct a closed subvariety $\mathcal{T}' \subset \mathcal{R}$ whose intersection with any of the subvarieties \mathcal{R}_i has codimension at least 2 in \mathcal{Z}_i . This is secured by firstly choosing some "main reference" fibers \mathcal{E}_{r_i} , where the $r_i \in \mathcal{R}_i$ are as in the statement of the lemma, and then choosing our hypersurface X' such that $\mathcal{E}_{r_i}|_{X'}$ remains torsion free on X' . This is true for every element X' in a dense open set of the linear system $|H|$ by [33, Lemma 1.1.12 and Corollary 1.1.14].

⁷This subvariety may depend on the flag.

Now notice that $\mathcal{R}_i \setminus \mathcal{R}_{tf}$ is not dense in \mathcal{R}_i for any i , and that, by the previous discussion, the restriction $\mathcal{E}|_{X' \times (\mathcal{R}_{tf} \cap \mathcal{R}_i)}$ remains flat. However, it still may happen that $\text{codim}_{\mathcal{R}_i}(\mathcal{R}_i \setminus \mathcal{R}_{tf}) = 1$. To construct the set \mathcal{T}' , whenever a set $(\mathcal{R}_i) \setminus \mathcal{R}_{tf}$ has codimension one in (\mathcal{R}_i) we choose some "secondary reference" point $s'_i \in \mathcal{R}_i \setminus \mathcal{R}_{tf}$ and ask that the hypersurface X' also be $\mathcal{U}_{s'_i}$ -regular, in addition to having the property that $\mathcal{E}_{r'_i}|_{X'}$ remains torsion free on X' . This can still be achieved for X' in a dense open subset of $|H|$, again by [33, Lemma 1.1.12 and Corollary 1.1.14]. By the same reference, a dense open subset of such X' 's is in addition fully admissible for $[(\mathcal{E}, \mathcal{C})]$ in the sense of Definition 5.7. The construction of the set $\mathcal{T}' \subset \mathcal{R}$ then proceeds in the same way as in the proof of Lemma 3.1 of [30], by setting

$$\mathcal{T}' = \mathcal{R} \setminus (\mathcal{R}_{tf} \cup \bigcup_{i \in I} N(X', r'_i)) ,$$

where $N(X', r'_i)$ is an open neighbourhood of r'_i in \mathcal{R} , where the restriction to X' remains flat. This set has the desired property with respect to the \mathcal{R}_i . Set $X' = X^{(1)}$.

We will repeat this procedure to find successive elements of the flag. Let $\mathcal{R}^{(1)} \subset \mathcal{R}$ be the flatness locus of this restricted family, (i.e. $\mathcal{R}^{(1)} = \mathcal{R} \setminus \mathcal{T}'$), so that $\mathcal{E}|_{X^{(1)} \times \mathcal{R}^{(1)}}$ is flat. We again choose secondary reference points $r_i^{(2)} \in \mathcal{R}^{(1)} \cap \mathcal{R}_i$. Then, for a general hypersurface $X_2 \in |H|$ the intersection $X^{(2)} = X^{(1)} \cap X_2$ is $\mathcal{E}_{r_i^{(2)}}$ -regular for each i and also fully admissible for $[(\mathcal{E}, \mathcal{C})]$. As above, we can then construct a nonflatness locus $\mathcal{T}^{(2)} \subset \mathcal{R}^{(1)}$ fulfilling $\text{codim}_{\mathcal{R}_i}(\mathcal{T}^{(2)} \cap \mathcal{R}_i) \geq 2$ and such that, setting $\mathcal{R}^{(2)} = \mathcal{R}^{(1)} \setminus \mathcal{T}^{(2)}$, the restricted family $\mathcal{E}|_{X^{(2)} \times \mathcal{R}^{(2)}}$ remains flat. Now inductively, for each successive $\mathcal{R}^{(l-1)}$ we find secondary reference points $r_i^{(l)} \in \mathcal{R}^{(l-1)} \cap \mathcal{R}_i$. Then, for a general hypersurface $X_l \in |H|$ the intersection $X^{(l)} = X^{(l-1)} \cap X_l$ is $\mathcal{E}_{r_i^{(l)}}$ -regular and also fully admissible for $[(\mathcal{E}, \mathcal{C})]$. To each of these is associated a nonflatness locus $\mathcal{T}^{(l)} \subset \mathcal{R}^{(l-1)}$ with $\text{codim}_{\mathcal{R}_i}(\mathcal{T}^{(l)} \cap \mathcal{R}_i) \geq 2$. We then define $\mathcal{R}^{(l)} = \mathcal{R}^{(l-1)} \setminus \mathcal{T}^{(l)}$ so that $\mathcal{E}|_{X^{(l)} \times \mathcal{R}^{(l)}}$ remains flat. We repeat this procedure until we get down to $X^{(n-2)} = S$.

We conclude that a sufficiently general flag $\{X^{(l)}\}$ is fully admissible for $[(\mathcal{E}, \mathcal{C})]$ (and in particular for S), and setting $\mathcal{T} = \bigcup_{l=1}^{n-2} \mathcal{T}^{(l)}$, for the subsets $\mathcal{T}^{(l)}$ introduced in the procedure described above, that $\mathcal{E}|_{X^{(l)} \times \mathcal{R} \setminus \mathcal{T}}$ remains flat for all l , as desired. \square

In the previous proof, note that we may keep the same "main reference" sheaves \mathcal{E}_{r_i} on which we ask torsion freeness at any section, whereas we may need to choose new sets of "secondary reference" points $r_i^{(l)} \in (\mathcal{R}^{(l)} \cap \mathcal{R}_i)$ on which we ask regularity for the next section $X^{(l+1)}$, $1 \leq l \leq n-3$. In particular,

if \mathcal{E} is the universal sheaf $\mathcal{U} \rightarrow X \times \mathcal{Z}$ the restriction to X' remains flat since \mathcal{U}_s is by definition torsion free for all $s \in \mathcal{Z}$. In this case, in the notation used above we have $\mathcal{Z}' = \mathcal{Z}$, however not all the fibers of the restricted family will remain torsion free on X' , so we may not apply the same argument to get flatness for the next restriction to $X^{(2)} \times \mathcal{Z}'$. It is therefore necessary to follow the argument above for all further restrictions.

In the same vein, if \mathcal{F}_r is torsion free for all r , and $X = S$ is a surface then there is no nonflatness locus \mathcal{T} , since flatness is preserved by restricting to any curve $C \subset S$.

Remark 5.17. If the \mathcal{E}_{r_i} are not just torsion free but semistable, for any d_1, \dots, d_{n-1} and any sufficiently general tuple $(X_1, \dots, X_{n-1}) \in |d_1 H| \times \dots \times |d_{n-1} H|$ the restriction of any of the associated graded sheaves $\text{Gr}(\mathcal{E}_{r_i})$ to each term of the flag $X^{(l)}$ will remain torsion free. Hence, if the sequence d_1, \dots, d_{n-1} is sufficiently increasing, then by Langer's Theorem, Theorem 5.4, the restriction of each \mathcal{E}_{r_i} to $X^{(n-2)}$ will remain semistable. We will say in this case that the flag *satisfies Langer's conditions*.

The following lemma follows directly from the proof of Lemma 5.16, but as we shall see, it is useful to separate out the case of curves.

Lemma 5.18. *Let $\mathcal{U} \rightarrow X \times \mathcal{Z}$ be the universal sheaf, and let $[(\mathcal{E}, \mathcal{C})] \in \overline{\mathcal{M}}_{\text{HYM}}$. Let $\mathcal{Z}_i \subset \mathcal{Z}$ be the finite set of holomorphic subvarieties provided by Lemma 5.15, let $\{X^{(l)}\}_{1 \leq l \leq n-2}$ be a sufficiently general flag of complete intersections with $X^{(n-2)} = S$ together with the subvariety \mathcal{T} produced by Lemma 5.16; and let $\mathcal{Z}^{(l)} \subset \mathcal{Z}$ be the subspaces from the proof of Lemma 5.16. We also write $z_i \in \mathcal{Z}_i$ for the main reference points, so that in particular $\mathcal{U}_{z_i}|_S$ is torsion free. Suppose $C \subset S$ is a smooth curve such that there are points $z_i^{(n-1)} \in \mathcal{Z}^{(n-2)} \cap \mathcal{Z}_i$ such that C is $U_{z_i^{(n-1)}}|_S$ -regular. Then, there is a closed holomorphic subvariety $\mathcal{T}_C \subset \mathcal{Z}$ such that*

- (1) $\text{codim}_{z_i}(\mathcal{T}_C \cap \mathcal{Z}_i) \geq 2$ for each i , and
- (2) $\mathcal{U}|_{C \times \mathcal{Z} \setminus \mathcal{T}_C \cup \mathcal{T}}$ remains flat.

5.5. Lifting sections from curves I: the method of [30]. The construction of the space $\overline{\mathcal{M}}^\mu$ in [30] requires the existence of sections of $\mathcal{L}_{n-1}^{\otimes k}$, for some power of k . In particular if one considers all such sections for all powers of k then these provide a map to projective space defining $\overline{\mathcal{M}}^\mu$. In comparing this space to $\overline{\mathcal{M}}_{\text{HYM}}$, instead of considering all sections we will need to restrict our attention to a certain linear system of sections that comes from "lifting" sections from complete intersection curves. We therefore need to review the procedure by which sections of $\mathcal{L}_{n-1}^{\otimes k}$ are constructed in the first place.

The main idea is that for a smooth curve $C \subset X$, GM- and slope semistability of sheaves on C are equivalent, which means that the point in the Quot scheme parametrising quotients on C corresponding to a μ -semistable sheaf on C is GIT semistable for the appropriate determinant line bundle on C , provided that this line bundle can be constructed (see Section 2.6), and in particular this is the case for any complete intersection curve $C = X^{(n-1)}$.

What is required for the construction is that the restriction of the universal family to C remain flat. By Lemma 5.18, this can be achieved away from a certain proper subvariety $\mathcal{T} \cup \mathcal{T}_C \subset \mathcal{Z}$, provided C is well chosen. GIT then produces nonvanishing sections of this determinant line bundle (see Lemma 5.19 below), which can then be lifted to give sections of a power of \mathcal{L}_{n-1} away from $\mathcal{T} \cup \mathcal{T}_C$ (see Lemma 5.20 below). One then has to know that $\mathcal{T} \cup \mathcal{T}_C$ can be chosen in such a way that these sections extend over $\mathcal{T} \cup \mathcal{T}_C$, to give sections on all of \mathcal{Z} . This will be obtained from the stated properties of $\mathcal{T} \cup \mathcal{T}_C$ in Lemma 5.18 by applying Lemma 5.15.

We now make this more precise. Consider the universal family $\mathcal{U} \rightarrow X \times \mathcal{Z}$, and choose a flag $\{X^{(l)}\}_{1 \leq l \leq n-2}$ as in Lemma 5.16. In particular, write $X^{(n-2)} = S$ for the end-term of this flag. Now we extend this flag to a flag $\{X^{(l)}\}_{1 \leq l \leq n-1}$ by choosing a curve $X^{(n-1)} = C \subset S$.

Let $c \in K(X)_{\text{num}}$ be a class with $c_i = c_i(c) = c_i(E)$. Write $c^{(n-1)} = c|_C$, and $\tau_{E,C}$ for the Hilbert polynomial determined by $c^{(n-1)}$ with respect to $\mathcal{O}_C(1)$. Take $m^{(n-1)}$ to be a (large) positive integer, and write $V_C = \mathbb{C}^{\tau_{E,C}(m^{(n-1)})}$, and correspondingly $\mathcal{H}_C = V_C \otimes \mathcal{O}_C(-m^{(n-1)})$, and let $Q_C \subset \text{Quot}(\mathcal{H}_C, \tau_{E,C})$ be the subscheme consisting of $m^{(n-1)}$ -regular quotients $q_F : \mathcal{H}_C \rightarrow F$ with determinant $\mathcal{J}|_C$ such that q_F induces an isomorphism $V_C \xrightarrow{\cong} H^0(F(m^{(n-1)}))$. Applying the discussion of Section 2.6 to the universal quotient $\mathcal{H}_C \otimes \mathcal{O}_C \rightarrow \widehat{\mathcal{U}}_C \rightarrow 0$ we obtain a line bundle $\mathcal{L}_{\widehat{\mathcal{U}},0} := \mathcal{L}_{0,C} \rightarrow Q_C \in \text{Pic}(Q_C)$.

The point is that because the two notions of stability are equivalent on C , GIT semistability of points in Q_C suffices to produce sections of a sufficiently large tensor power of $\mathcal{L}_{0,C}$. Explicitly, we have the following result, see [33, top of p. 223], [30, Lemma 3.8], and especially [49] for detailed explanations.

Lemma 5.19. *For a sufficiently large choice of $m^{(n-1)}$, for any quotient $q_{\mathcal{F}} : \mathcal{H}_C \rightarrow \mathcal{F}$ giving a point $q_{\mathcal{F}} \in Q_C$, the sheaf \mathcal{F} is semistable with the property that the natural map $V_C \rightarrow H^0(\mathcal{F}(m^{(n-1)}))$ is an isomorphism iff $q_{\mathcal{F}}$ is GIT semistable for the $\text{SL}(V_C)$ linearisation of $\mathcal{L}_{0,C}$, i.e., iff there exists a positive integer κ , and an $\text{SL}(V_C)$ invariant section $\sigma \in H^0(\mathcal{L}_{0,C}^{\otimes \kappa})$ that does not vanish at $q_{\mathcal{F}}$. Furthermore, suppose $q_{\mathcal{F}_1}, q_{\mathcal{F}_2} \in Q_C$ are quotients satisfying any of the equivalent conditions above, so that in particular the sheaves \mathcal{F}_1 and \mathcal{F}_2 are*

semistable. Then there is an $\mathrm{SL}(V_C)$ invariant section $\sigma \in H^0(\mathcal{L}_{0,C}^{\otimes \kappa})$ for some positive integer κ such that $0 = \sigma(\mathcal{F}_1) \neq \sigma(\mathcal{F}_2)$, (that is, sections of some tensor power of $\mathcal{L}_{0,C}$ separate $q_{\mathcal{F}_1}$ and $q_{\mathcal{F}_2}$) iff \mathcal{F}_1 and \mathcal{F}_2 are not s -equivalent.

We now state the following key lemma, which produces sections of a tensor power of \mathcal{L}_{n-1} from the sections given by the previous lemma. This is implicit in [30], but we find it useful to formulate an explicit statement concerning this point here.

Lemma 5.20. *Let $C \subset S$ be a curve satisfying the hypotheses of Lemma 5.18, so that in particular C is the end-term of a flag $\{X^{(l)}\}_{1 \leq l \leq n-1}$ such that $\mathcal{U}|_{C \times \mathcal{Z} \setminus (\mathcal{T} \cup \mathcal{T}_C)}$ is flat. There exist positive integers κ_0, κ_{n-1} such that for all $k \in \mathbb{N}_{>0}$ there exist natural linear "lifting" maps*

$$\bar{L}_{\mathcal{U},C} : H^0(Q_C, \mathcal{L}_{0,C}^{\otimes k\kappa_0})^{\mathrm{SL}(V_C)} \longrightarrow H^0(\mathcal{Z}, \mathcal{L}_{n-1}^{\otimes k\kappa_{n-1}})^{\mathrm{SL}(V)}$$

such that the following holds: If $(q_{\mathcal{F}} : \mathcal{H} \rightarrow \mathcal{F}) \in \mathcal{Z} \setminus (\mathcal{T} \cup \mathcal{T}_C)$ is a quotient, $q_{\mathcal{F}_C} : \mathcal{H}_C \rightarrow \mathcal{F}_C$ is a quotient realisation of $\mathcal{F}_C = \mathcal{F}|_C$, and $\sigma \in H^0(Q_C, \mathcal{L}_{0,C}^{\otimes k\kappa_0})^{\mathrm{SL}(V_C)}$ is a section, then $\bar{L}_{\mathcal{U},C}(\sigma)(q_{\mathcal{F}}) \neq 0$ if and only if $\sigma(q_{\mathcal{F}_C}) \neq 0$. In particular, if $q_{\mathcal{F}} \in \mathcal{Z}$ is a quotient such that $\mathcal{F}|_{X^{(l)}}$ remains semistable for all members of the flag, then for k large enough there is a section of $H^0(\mathcal{Z}, \mathcal{L}_{n-1}^{\otimes k\kappa_{n-1}})^{\mathrm{SL}(V)}$ that does not vanish at $q_{\mathcal{F}}$.

Proof. The construction of the lifting map is in [30, Sect. 3.3], the numbers κ_{n-1} and κ_0 are determined by [30, Eq. (3.15)]. This gives a section in the space $H^0(\mathcal{Z} \setminus \mathcal{T} \cup \mathcal{T}_C, \mathcal{L}_{n-1}^{\otimes k\kappa_{n-1}})^{\mathrm{SL}(V)}$. By the construction of the sets \mathcal{T} and \mathcal{T}_C , and Lemma 5.15, these sections extend to all of \mathcal{Z} . The second statement follows directly from the existence of the lifting map and Lemma 5.19. \square

We will refer to elements of $H^0(\mathcal{Z}, \mathcal{L}_{n-1}^{\otimes k\kappa_{n-1}})^{\mathrm{SL}(V)}$ that are lifted from sections in $H^0(Q_C, \mathcal{L}_{0,C}^{\otimes k\kappa_0})^{\mathrm{SL}(V_C)}$ in this way as θ -functions of level $k\kappa_{n-1}$ lifted from C . Note that since the moduli space $M^{\mu_{ss}}$ is constructed using sections of the $\mathcal{L}_{n-1}^{\otimes k\kappa_{n-1}}$, any θ -function σ with $\sigma(q_{\mathcal{F}}) \neq 0$ descends to a section of some tensor power⁸ of $\mathcal{O}_{M^{\mu_{ss}}}(1)$, which does not vanish at $[\mathcal{F}]$. We will continue to refer to these induced sections as θ -functions (lifted from C).

5.6. Lifting sections from curves II: the method of [38]. We collect here some information on an alternative way of "lifting" sections (cf. [38, p. 433-4]), and its compatibility with the construction described in the previous subsection. This will later allow us to use very fine results concerning separation properties of sections in determinant line bundles for families of sheaves on surfaces obtained by Jun Li in [38] (see the proof of Proposition 5.21 below).

⁸See Diagrams (5.5) and (5.6) below for a computation of the exact tensor power.

As mentioned in Section 2.8, the space $M^{\mu ss}$ comes equipped with an ample line bundle $\mathcal{O}_{M^{\mu ss}}(1)$ and a number $N = N_{n-1}$ with the property that any flat family $\mathcal{E} \rightarrow X \times T^\circ$ of semistable sheaves (with T° weakly normal) gives rise to a unique classifying morphism $\psi_{\mathcal{E}} : T^\circ \rightarrow M^{\mu ss}$, so that

$$\psi_{\mathcal{E}}^*(\mathcal{O}_{M^{\mu ss}}(1)) = \lambda_{\mathcal{E}}(u_{n-1}(c, [\mathcal{O}_H]))^{\otimes N}. \quad (5.4)$$

Without loss of generality, as explained in [30, Sect. 4.2], we may assume that the section ring of $\mathcal{O}_{M^{\mu ss}}(1)$ is generated in degree one. This induces a "lifting" map $\psi_{\mathcal{E}}^* : H^0(M^{\mu ss}, \mathcal{O}_{M^{\mu ss}}(1)) \rightarrow H^0(T^\circ, \lambda_{\mathcal{E}}(u_{n-1}(c, [\mathcal{O}_H]))^{\otimes N})$. It is clear that if the family $\mathcal{E} \rightarrow X \times T^\circ$ is linearized with respect to some G -action on T° , then $\psi_{\mathcal{E}}$ is G -invariant and so are the sections in $\psi_{\mathcal{E}}^*(H^0(M^{\mu ss}, \mathcal{O}_{M^{\mu ss}}(1)))$. The above applies in particular to the case when $T^\circ = \mathcal{Z}$, $\mathcal{E} = \widehat{\mathcal{U}}$ and $G = \mathrm{SL}(V)$ and the induced map $\psi_{\widehat{\mathcal{U}}}^*$ is an isomorphism. It also applies to the case when $X = C$ is a curve, $T^\circ = Q_C^{ss}$, \mathcal{E} is the restriction of the universal quotient bundle $\widehat{\mathcal{U}}$ to $C \times Q_C^{ss}$ and $G = \mathrm{SL}(V_C)$ giving an isomorphism $\psi_{\widehat{\mathcal{U}}}^*$.

Let now $P \rightarrow T^\circ$ denote the frame bundle associated to the locally free sheaf $pr_{2,*}(\mathcal{E}(l))$. We have a commutative diagram of natural morphisms

$$\begin{array}{ccc} P & \xrightarrow{q} & \mathcal{Z} \\ \downarrow p & & \downarrow \\ T^\circ & \longrightarrow & M^{\mu ss} \end{array}$$

As in [33, page 224] and [30, proof of Lemma 3.8] by pulling back through q and pushing down through p for all $\kappa \in \mathbb{N}_{>0}$ we get a linear map

$$L_{\mathcal{E}} : H^0(\mathcal{Z}, \mathcal{L}_{n-1}^{\otimes \kappa N})^{\mathrm{SL}(V)} \longrightarrow H^0(T^\circ, \lambda_{\mathcal{E}}(u_{n-1}(c, [\mathcal{O}_H]))^{\otimes \kappa N})$$

which by construction sits in the following commutative diagram

$$\begin{array}{ccc} & & H^0(\mathcal{Z}, \mathcal{L}_{n-1}^{\otimes \kappa N})^{\mathrm{SL}(V)} \\ & \swarrow L_{\mathcal{E}} & \uparrow \psi_{\widehat{\mathcal{U}}}^* \\ H^0(T^\circ, \lambda_{\mathcal{E}}(u_{n-1}(c, [\mathcal{O}_H]))^{\otimes \kappa N}) & \xleftarrow{\psi_{\mathcal{E}}^*} & H^0(M^{\mu ss}, \mathcal{O}_{M^{\mu ss}}(\kappa)) \end{array}$$

showing that in this case the two ways of lifting sections are compatible.

Let now X be n -dimensional and let $\{X^{(l)}\}_{1 \leq l \leq n-1}$ be a flag of complete intersections with $X^{(n-2)} = S$ and $X^{(n-1)} = C$. Denote by $E_{X^{(i)}}$ and $\mathcal{J}_{X^{(i)}}$ the restrictions of the smooth bundle E and holomorphic bundle \mathcal{J} to $X^{(i)}$. Let T° be a weakly normal base of a flat family $\mathcal{E} \rightarrow X \times T^\circ$ of semistable sheaves, such that its restrictions $\mathcal{E}^{(l)} \rightarrow X^{(l)} \times T^\circ$ to all terms of the flag

remain flat families of semistable sheaves. By the same arguments as above and in a canonical way we obtain a map:

$$L_{\mathcal{E},C} : H^0(Q_C, \mathcal{L}_{0,C}^{\otimes \kappa_0})^{\text{SL}(V_C)} \rightarrow H^0(T^\circ, \lambda_{\mathcal{E}}(u_{n-1}(c, [\mathcal{O}_H]))^{\otimes \kappa_1}).$$

As before, if the family $\mathcal{E} \rightarrow X \times T^\circ$ is linearized with respect to some G -action on T , then the image of $L_{\mathcal{E},C}$ lies inside $H^0(T^\circ, \lambda_{\mathcal{E}}(u_{n-1}(c, [\mathcal{O}_H]))^{\otimes \kappa_1})^G$.

From the above, we also obtain a commutative diagram

$$\begin{array}{ccc}
H^0(M_C^{\mu_{ss}}, \mathcal{O}_{M_C^{\mu_{ss}}}(\kappa_0 N)) & \xrightarrow{\psi_{\mathcal{E}^*}^*} & \\
\downarrow \psi_{\mathcal{U}}^* & & \\
H^0(Q_C, \mathcal{L}_{0,C}^{\otimes \kappa_0 N N_0})^{\text{SL}(V_C)} & \xrightarrow{L_{\mathcal{E}|_C, C}} H^0(T^\circ, \lambda_{\mathcal{E}|_C}(u_0(c|_C, [\mathcal{O}_H]_C))^{\otimes \kappa_0 N N_0}) & \cong \downarrow \text{by [30, Eq. (3.5)]} \\
\downarrow \bar{L}_{\mathcal{U}, C} & \searrow L_{\mathcal{E}, C} & \\
H^0(\mathcal{Z}, \mathcal{L}_{n-1}^{\otimes \kappa_{n-1} N_0 N})^{\text{SL}(V)} & \xrightarrow{L_{\mathcal{E}}} H^0(T^\circ, \lambda_{\mathcal{E}}(u_{n-1}(c, [\mathcal{O}_H]))^{\kappa_{n-1} N_0 N}) & \\
\uparrow \psi_{\mathcal{U}}^* & & \\
H^0(M^{\mu_{ss}}, \mathcal{O}_{M^{\mu_{ss}}}(\kappa_{n-1} N_0)) & \xrightarrow{\psi_{\mathcal{E}}^*} &
\end{array} \tag{5.5}$$

where N_0 plays the role of N for the moduli space $M_C^{\mu_{ss}}$ (see eq. (5.4)). Here we have used that $H^0(Q_C^{\text{ss}}, \mathcal{L}_{0,C}^{\otimes \kappa_0})^{\text{SL}(V_C)} \cong H^0(Q_C, \mathcal{L}_{0,C}^{\otimes \kappa_0})^{\text{SL}(V_C)}$ in order to construct the map $\psi_{\mathcal{U}}^*$ (see [14, Sect. 8.2]).

Moreover these maps are compatible with those arising from the intermediate spaces $X^{(l)}$, and in particular from S , in the following sense. Indeed, using the fact that T° is also base for the restricted flat family $\mathcal{E}_{S \times T^\circ}$ of semistable sheaves on S and the isomorphism of determinant line bundles proved in [30, Sect. 3.2.2] we get for the appropriate powers of determinant sheaves a commutative diagram of the form

$$\begin{array}{ccc}
H^0(M^{\mu_{ss}}, \mathcal{O}_{M^{\mu_{ss}}}(\kappa_{n-1} N_0 N_1)) & \longrightarrow & H^0(T^\circ, \lambda_{\mathcal{E}}(u_{n-1}(c, [\mathcal{O}_H]))^{\otimes \kappa_{n-1} N_0 N_1 N}) \\
\uparrow & & \uparrow \cong \\
H^0(M_C^{\mu_{ss}}, \mathcal{O}_{M_C^{\mu_{ss}}}(\kappa_0 N N_1)) & \longrightarrow & H^0(T^\circ, \lambda_{\mathcal{E}|_C}(u_0(c|_C, [\mathcal{O}_H]_C))^{\otimes \kappa_0 N N_1 N_0}) \\
\downarrow & & \downarrow \cong \\
H^0(M_S^{\mu_{ss}}, \mathcal{O}_{M_S^{\mu_{ss}}}(\kappa_1 N_0 N)) & \longrightarrow & H^0(T^\circ, \lambda_{\mathcal{E}|_S}(u_1(c|_S, [\mathcal{O}_H]_S))^{\otimes \kappa_1 N_0 N N_1}),
\end{array} \tag{5.6}$$

where N_1 is defined similarly to N and N_0 (see again (5.4)).

In the situation of the proof of Lemma 5.20 we may set T° to be equal to $\mathcal{Z} \setminus \mathcal{T} \cup \mathcal{T}_C$ and $\mathcal{E} = \mathcal{U}$. In this case, the map $L_{\mathcal{E}}$ in (5.5) is the restriction map.

5.7. The 2-dimensional case. The next proposition deals with moduli spaces of slope-semistable sheaves on surfaces and is implicitly contained in Li's paper [38] although not explicitly stated there.

As remarked at various junctures, the construction of the moduli space $M^{\mu ss}$ proceeds by finding nonvanishing equivariant sections of tensor powers of \mathcal{L}_{n-1} . Proving the separation properties of this moduli space therefore amounts to finding such sections that separate the appropriate quotients (see Section 2.8). The point of the following proposition is that in the case of surfaces, this may in fact be achieved by only considering the θ -functions that are lifted from a finite number of fixed curves.

Proposition 5.21. *Let $(S, \mathcal{O}_S(1))$ be a polarized smooth complex projective surface, $(r, c_1, c_2) \in \mathbb{Z}_{>0} \times NS(S) \times H^4(S, \mathbb{Z})$ fixed topological invariants, let $M_S^{\mu ss}$ be the moduli space of slope semistable sheaves on S with the given topological invariants and with fixed determinant \mathcal{J}_S . Let $A \subset S$ be finite. Then there exist positive integers k and l (depending on k), and smooth curves $C^{(1)}, \dots, C^{(n)} \in |\mathcal{O}_X(k)|$ not intersecting A such that the linear system*

$$W_S \subset H^0(M_S^{\mu ss}, \mathcal{O}_{M^{\mu ss}}(l\kappa_1 N_0))$$

defined as the span of all θ -functions on $M_S^{\mu ss}$ of level $l\kappa_1 N_0 N_1$ lifted from the curves $C^{(1)}, \dots, C^{(n)}$ gives an injective morphism

$$\nu_S : M_S^{\mu ss} \hookrightarrow \mathbb{P}(W_S^*).$$

Proof. Let \mathcal{Z} be the weak normalisation of

$$R_S^{\mu ss} \subset \text{Quot}(\mathcal{H}, \tau(r, c_1, c_2)),$$

and consider the family $\{\mathcal{F}_i\}_{i \in \mathcal{Z}}$ of sheaves parametrised by \mathcal{Z} . Choose $k \gg 0$, n , and smooth curves $C^{(1)}, \dots, C^{(n)} \in |kH|$ satisfying the conclusion of Proposition 5.5. Note that the curves may be chosen such that they in addition avoid the set A . Let $\mathcal{F} = \mathcal{F}_i$ for some $i \in \mathcal{Z}$. By Proposition 5.5 there exists an index $\alpha \in \{1, \dots, n_0\}$ such that the sheaf $\mathcal{F}|_{C^{(\alpha)}}$ remains semistable and hence by Lemma 5.19 determines a section in $H^0(Q_{C^{(\alpha)}}, \mathcal{L}_{0, C^{(\alpha)}}^{lN_1 N_0 \kappa_0})^{\text{SL}(V_{C^{(\alpha)}})}$ that does not vanish at $\mathcal{F}|_{C^{(\alpha)}}$. Consequently, by Lemma 5.20 we obtain a section in $H^0(\mathcal{Z}, \mathcal{L}_{n-1}^{l\kappa_1 N_1 N_0})^{\text{SL}(V)}$ not vanishing at any quotient associated to \mathcal{F} . This implies that there exists a positive integer l such that the linear system W_S defined as the span of all θ -functions of level $l\kappa_1 N_1 N_0$ lifted from the curves $C^{(1)}, \dots, C^{(n)}$ has no base points on $M_S^{\mu ss}$ and thus gives a morphism $\nu_S : M_S^{\mu ss} \rightarrow \mathbb{P}(W_S^*)$.

Moreover, it follows from Proposition 5.5 and Lemmas 5.19 and 5.20 that, after increasing l if necessary, this morphism separates points of the form

$[(\mathcal{E}_1, \mathcal{C}_1)], [(\mathcal{E}_2, \mathcal{C}_2)]$ as soon as \mathcal{E}_1 and \mathcal{E}_2 are not isomorphic. Namely, if $\mathcal{E}_1 = (\text{Gr } \mathcal{F}_2)^{\vee\vee}$ and $\mathcal{E}_2 = (\text{Gr } \mathcal{F}_1)^{\vee\vee}$, then by Proposition 5.5 there exists an $\alpha \in \{1, \dots, n\}$ such that the restrictions of \mathcal{F}_1 and \mathcal{F}_2 to $C^{(\alpha)}$ are not s -equivalent, and therefore by Lemma 5.19 there is a section of

$$H^0(Q_{C^{(\alpha)}}, \mathcal{L}_{0, C^{(\alpha)}}^{\otimes l \kappa_0 N_1 N_0})^{\text{SL}(V_{C^{(\alpha)}})}$$

that separates them on $C^{(\alpha)}$, and hence there is an induced lifted section in $H^0(\mathcal{Z}, \mathcal{L}_{n-1}^{l \kappa_1 N_1 N_0})^{\text{SL}(V)}$ which separates them on S by Lemma 5.20.

We thus only need to show that by increasing l again if necessary and by adding θ -functions lifted from further curves $C^{(\alpha)} \in |kH|$ avoiding A to this linear system also points of the form $[(\mathcal{E}, \mathcal{C}_1)], [(\mathcal{E}, \mathcal{C}_2)]$ with $\mathcal{C}_1 \neq \mathcal{C}_2$ get separated. By a Noetherian induction argument for this it suffices to show that two *fixed* distinct points $[(\mathcal{E}, \mathcal{C}_1)], [(\mathcal{E}, \mathcal{C}_2)] \in W_S$ may be separated by adding θ -functions lifted from further curves as above. This will be done using [38, Lemma 3.6] by noticing that its proof actually shows that the added curves may be chosen to avoid A .⁹ Indeed, in order to prove the desired separation, Li uses θ -functions lifted from two *general* members D_{t_1}, D_{t_2} of a pencil of curves generated by two elements $D_0, D_1 \in |kH|$ chosen as follows (see [38, p. 441]): if $\mathcal{C}_1 = \sum_{i=1}^s m'_i P'_i$, $\mathcal{C}_2 = \sum_{i=1}^s m''_i P''_i$ and $P = P'_1 = P''_1$ is such that $m'_1 \neq m''_1$ then it is asked that $P \in D_0$ and $P \notin D_1$, but otherwise D_0, D_1 are general in $|kH|$, in particular disjoint from $A \setminus \{P\}$. Thus general elements D_{t_1}, D_{t_2} in this pencil will be disjoint from A . \square

Remark 5.22. In particular, in our global setup, for any smooth surface $S \subset X$ we obtain by restriction an injective map, still denoted ν_S , giving an embedding of $\overline{M}_S^\mu \subset M_S^{\mu, \text{ss}}$ into $\mathbb{P}(W_S^*)$. Here, S is polarised by $\mathcal{O}_S(1) = \mathcal{O}_X(1)|_S$, the invariants (r, c_1, c_2) are determined by $E|_S$, and the fixed determinant \mathcal{J}_S is given by $\mathcal{J}|_S$.

5.8. The higher dimensional case. The following claim deals with the higher dimensional case and is the main technical statement of this section. It will be proved in the sequel.

Claim 5.23. *For every point $p \in \overline{M}^\mu$ there exist an open R -saturated neighbourhood $U \subset \overline{M}^\mu$, complex vector spaces W_j , $1 \leq j \leq m$ and holomorphic*

⁹See [33, Lemma 8.3.4] for a comparison of the determinant line bundles used. Strictly speaking, Jun Li works in the case of rank two sheaves with zero first Chern class, but his computations remain valid in the general case as well, cf. [33, pp. 229/230]. The possibility of avoiding A in the general case can also be concluded from the less direct argument for separation presented in [33, proof of Prop. 8.2.13] (in the situation discussed there, let C run through all smooth curves avoiding A).

maps $\nu_j : U \rightarrow \mathbb{P}(W_j^*)$ such that the fibers of

$$\nu = \nu_1 \times \dots \times \nu_m : U \longrightarrow \mathbb{P}(W_1^*) \times \dots \times \mathbb{P}(W_m^*)$$

are precisely the equivalence classes of R inside U .

The idea will be to consider different surfaces S_1, \dots, S_m arising as last terms of flags of complete intersections in X , which are moreover fully admissible for $p = [(\mathcal{E}, \mathcal{C})]$, and their corresponding respective domains of admissibility $U'_j \subset \overline{M}_{\text{HYM}}$, $U_j \subset \overline{M}^\mu$, $j = 1, \dots, m$, and to construct morphisms $\nu_j : U_j \rightarrow \mathbb{P}(W_j)$ induced by restriction to the surfaces S_j in a way to be made precise. Then, we will simply define $\nu_1 \times \dots \times \nu_m$ on $U = \bigcap_{j=1}^m U_j$ and show that the collection of ν_j 's has the desired collapsing and separation properties.

We now describe the construction of the surfaces S_j and of the morphisms ν_j . While we start by explaining how to construct one surface, later we will see why several such surfaces may be needed. We begin by fixing our setup.

Setup. Let $p \in \overline{M}^\mu$ and $\overline{\Phi}(p) = [(\mathcal{E}, \mathcal{C})]$ be its image in $\overline{M}_{\text{HYM}}$. Let d_1, \dots, d_{n-2} be a sequence of positive natural number, and $(X^{(l)})_{1 \leq l \leq n-2}$ be a flag of complete intersections of hypersurfaces $X_i \in |\mathcal{O}_X(d_i)|$ such that

- (1) $(X^{(l)})_{1 \leq l \leq n-2}$ satisfies the conclusions of Lemma 5.16 with respect to $(\mathcal{E}, \mathcal{C})$, the universal family $\mathcal{U} \rightarrow X \times \mathcal{Z}$, the subvarieties $\mathcal{Z}^{(i)} \subset \mathcal{Z}$ provided by Lemma 5.15, and chosen main reference points z_i in $\mathcal{Z}^{(i)}$;
- (2) the flag satisfies Langer's conditions; i.e., the sequence of integers d_1, \dots, d_{n-2} is sufficiently increasing for a repeated application of Theorem 5.4, and the restrictions of the main reference sheaves to S are semistable (cf. Remark 5.17).

The following will provide us with the building blocks for the R -invariant maps we need to construct in order to prove Claim 5.23.

Lemma 5.24. *In the Setup, if we denote the last term of the flag by $X^{(n-2)} = S \subset X$ and we let $C^{(1)}, \dots, C^{(n_0)}$ be the finite set of curves and l be the natural number provided by Proposition 5.21, then the linear system W of sections of $\mathcal{L}_{n-1} \rightarrow \mathcal{Z}$ provided by lifting theta functions of level $l\kappa_0 N_0 N_1 N_{n-1}$ from the $C^{(\alpha)}$ to \mathcal{Z} gives a rational map $\eta : \mathcal{Z} \dashrightarrow \mathbb{P}(W_S^*)$, into the projective space associated with the linear system W_S , which descends to a rational map*

$$\nu : \overline{M}^\mu \dashrightarrow \mathbb{P}(W_S^*)$$

that is well-defined at the images in \overline{M}^μ of the points $z_i \in \mathcal{Z}$.

Proof. Since $\mathcal{U} \rightarrow X \times \mathcal{Z}$ has the property that \mathcal{U}_z is torsion free for each $z \in \mathcal{Z}$, choosing arbitrary main reference points $z_i \in \mathcal{Z}^{(i)}$, we see that the

assumptions of Lemma 5.16 are satisfied. We therefore obtain flatness of \mathcal{U} on restriction to $S \times \mathcal{Z} \setminus \mathcal{T}$, with $\mathcal{T} \subset \mathcal{Z}$ the closed subvariety obtained in Lemma 5.16. Moreover, we will assume that the conclusion of Remark 5.17 holds; i.e., the restrictions of the main reference sheaves to S are semistable.

According to Lemma 5.18, for any well chosen curve C , we obtain flatness of the restriction of the universal sheaf to $C \times \mathcal{Z} \setminus (\mathcal{T} \cup \mathcal{T}_C)$, where $\mathcal{T} \cup \mathcal{T}_C \subset \mathcal{Z}$ is a nonflatness locus which has the property required by Lemma 5.15 for extension of sections. More precisely, this is the case if the curve C is chosen regular with respect to some "secondary reference" sheaves $\mathcal{U}_{z_i}^{(n-1)}|_S$. This regularity is guaranteed once the curve C contains none of the associated points of the sheaves $\mathcal{U}_{z_i}^{(n-1)}|_S$, which in turn will be the case if C avoids some fixed closed points on S , one on each associated component of the sheaves $\mathcal{U}_{z_i}^{(n-1)}|_S$. We fix A_1 to be such a (finite) subset of points of S . Moreover, let A_2 be the finite set containing the singular points of the Seshadri graduations of the restrictions $\mathcal{U}_{z_i}|_S$ and set $A := A_1 \cup A_2$.

Let k be a positive integer as guaranteed by Proposition 5.21 and such that the conclusion of Theorem 5.4 holds, and let $C^{(1)}, \dots, C^{(n_0)}$ be the curves provided by Proposition 5.21. Let

$$W \subset H^0(\mathcal{Z}, \mathcal{L}_{n-1}^{\otimes l \kappa_{n-1} N N_0 N_1})^{\text{SL}(V)}$$

be the span of all theta functions of the appropriate level obtained by lifting sections in $H^0(Q_{C^{(\alpha)}}(l), \mathcal{L}_{0, C^{(\alpha)}}^{\kappa_0 l N N_1 N_0})^{\text{SL}(V_{C^{(\alpha)}})}$ and

$$W_X \subset H^0(M^{\mu_{SS}}, \mathcal{O}_{M^{\mu_{SS}}}(\kappa_{n-1} l N_1 N_0))$$

the corresponding linear system on $M^{\mu_{SS}}$; see the discussion in Section 5.6, especially Diagrams (5.5) and (5.6), which also shows that naturally

$$\mathbb{P}(W^*) \cong \mathbb{P}(W_X^*) \cong \mathbb{P}(W_S^*) . \quad (5.7)$$

Note that these linear systems are nontrivial, since the flag $(X^{(l)})_{1 \leq l \leq n-2}$ and the curves $C^{(\alpha)}$ have been chosen so that for each i , the sheaves \mathcal{U}_{z_i} remain semistable on $C^{(\alpha)}$. By the natural identifications listed in (5.7), the linear system W gives rise to a rational map $\eta : \mathcal{Z} \dashrightarrow \mathbb{P}(W_S^*)$ defined at the z_i , which descends to the rational map $\nu : M^{\mu_{SS}} \dashrightarrow \mathbb{P}(W_S^*)$ associated with W_X , and hence by restriction yields a rational map $\nu : \overline{M}^\mu \dashrightarrow \mathbb{P}(W_S^*)$. By construction, ν is defined at the images of the z_i , as desired. \square

The following proposition gives the required collapsing property for the map ν . When combined with Corollary 5.13, it will also imply the required separation property regarding "the double-dual component" of elements of the gauge theoretic moduli space.

Proposition 5.25. *The rational map $\nu : \overline{M}^\mu \dashrightarrow \mathbb{P}(W_S^*)$ constructed in Lemma 5.24 is well-defined and hence holomorphic on the open set U_S defined in Proposition 5.14 above.*

Moreover, given any polystable sheaf \mathcal{F} with $[\mathcal{F}] \in U_S$ there is a polystable sheaf $\mathcal{F}_{\parallel S}$ on S such that the map

$$r_{U_S} : U_S \rightarrow \overline{M}_S^\mu : [\mathcal{F}] \mapsto [\mathcal{F}_{\parallel S}] ,$$

is well-defined and extends the natural rational map $U_S \dashrightarrow \overline{M}_S^\mu$ given by restricting locally free stable sheaves to S . This map makes the following into a commutative diagram, where ν_S is as defined in Proposition 5.21 and Remark 5.22:

$$\begin{array}{ccc} & \overline{M}_S^\mu & \\ & \swarrow \nu_S & \uparrow r_{U_S} \\ \mathbb{P}(W_S^*) & \xleftarrow{\nu} & U_S . \end{array} \quad (5.8)$$

The map r_{U_S} is R -invariant. In particular, $\nu : U_S \rightarrow \mathbb{P}(W_S^*)$ is R -invariant.

Proof. Let $p \in U_S$ be an arbitrary point, let \mathcal{F} be a polystable representative of p and set $\mathcal{E} := \mathcal{F}^{\vee\vee}$.

Recall from Definition 2.11 that $\gamma(\mathcal{F}) = (\mathcal{F}^{\vee\vee}, \mathcal{C}_{\mathcal{F}})$, where $\mathcal{C}_{\mathcal{F}}$ is the support cycle of the sheaf $\mathcal{Q}_{\mathcal{F}} = \mathcal{F}^{\vee\vee}/\mathcal{F}$. As also introduced there, we set $\widehat{\mathcal{Q}}_{\mathcal{F}} = \mathcal{Q}_{\mathcal{F}}/\mathcal{T}(\mathcal{Q}_{\mathcal{F}})$. Note that the associated points of $\widehat{\mathcal{Q}}_{\mathcal{F}}$ correspond to the irreducible components of $\mathcal{C}_{\mathcal{F}}$. The admissibility of the surface S for $\overline{\Phi}([\mathcal{F}])$ hence implies that the hypersurface X' is regular for $\widehat{\mathcal{Q}}_{\mathcal{F}}$. By [33, Cor. 1.1.14] the restriction $\mathcal{E}|_{X'}$ is torsion free on X' . Using this as well as the long exact sequence induced by restricting the defining sequence of $\mathcal{Q}_{\mathcal{F}}$ to X' , we obtain an inclusion $\text{supp}(\text{Tors}_{\mathcal{O}_{X'}}(\mathcal{F}|_{X'})) \subset \text{supp}(\mathcal{T}or_1^{\mathcal{O}_X}(\mathcal{Q}_{\mathcal{F}}, \mathcal{O}_{X'}))$. As X' does not contain any 2-codimensional associated point of $\mathcal{Q}_{\mathcal{F}}$, the latter subvariety is at least three codimensional in X , from which we conclude that the torsion of $\mathcal{F}|_{X'}$ on X' will be supported in codimension at least two on X' . Thus $\gamma(\mathcal{F}|_{X'})$ has a meaning according to Definition 2.11. Setting¹⁰

$$\gamma(\mathcal{F})|_{X'} := ((\mathcal{E}|_{X'})^{\vee\vee}, (\mathcal{C}_{\mathcal{F}})|_{X'} + \mathcal{C}_{(\mathcal{E}|_{X'})}) , \quad (5.9)$$

we claim that

$$[\gamma(\mathcal{F})|_{X'}] = [\gamma(\mathcal{F}|_{X'})] . \quad (5.10)$$

As \mathcal{F} and \mathcal{E} coincide outside of a subvariety of codimension two in X whose codimension two part is the support of $\mathcal{Q}_{\mathcal{F}}$, and as X' intersects this support in codimension two, equality in the sheaf component follows from reflexivity.

¹⁰The definition (5.9) extends the one given in Remark 5.9 in the fully admissible case.

In order to prove equality in the cycle component, note first that $(\mathcal{C}_{\mathcal{F}})|_{X'} = (\mathcal{C}_{\mathcal{Q}_{\mathcal{F}}})|_{X'} = (\mathcal{C}_{\widehat{\mathcal{Q}}_{\mathcal{F}}})|_{X'}$, and second that by Lemma 5.10 the last cycle is equal to $\mathcal{C}_{(\widehat{\mathcal{Q}}_{\mathcal{F}})|_{X'}}$, which is computed on X' . We thus have to show that

$$\mathcal{C}_{(\widehat{\mathcal{Q}}_{\mathcal{F}})|_{X'}} = \mathcal{C}((\widehat{\mathcal{Q}}_{\mathcal{F}})|_{X'}) = -\mathcal{C}(\mathcal{F}|_{X'}) - \mathcal{C}(\mathcal{E}|_{X'}) .$$

Applying Lemma 2.14 to the morphism $\mathcal{F}|_{X'} \rightarrow \mathcal{E}|_{X'}$ one gets $-\mathcal{C}(\mathcal{F}|_{X'}) - \mathcal{C}(\mathcal{E}|_{X'}) = \mathcal{C}((\mathcal{Q}_{\mathcal{F}})|_{X'}) - \mathcal{C}(\mathcal{T}or_1^{\mathcal{O}_X}(\mathcal{Q}_{\mathcal{F}}, \mathcal{O}_{X'}))$, so it will be enough to check that

$$\mathcal{C}((\widehat{\mathcal{Q}}_{\mathcal{F}})|_{X'}) = \mathcal{C}((\mathcal{Q}_{\mathcal{F}})|_{X'}) - \mathcal{C}(\mathcal{T}or_1^{\mathcal{O}_X}(\mathcal{Q}_{\mathcal{F}}, \mathcal{O}_{X'})) .$$

For this, we consider the following commutative diagram with exact rows and columns

$$\begin{array}{ccccccc}
& & 0 & & 0 & & \\
& & \downarrow & & \downarrow & & \\
0 & \longrightarrow & \mathcal{T}or_1^{\mathcal{O}_X}(\mathcal{T}(\mathcal{Q}_{\mathcal{F}}), \mathcal{O}_{X'}) & \longrightarrow & \mathcal{T}or_1^{\mathcal{O}_X}(\mathcal{Q}_{\mathcal{F}}, \mathcal{O}_{X'}) & \longrightarrow & 0 \\
& & \downarrow & & \downarrow & & \\
0 & \longrightarrow & \mathcal{T}(\mathcal{Q}_{\mathcal{F}})(-X') & \longrightarrow & \mathcal{Q}_{\mathcal{F}}(-X') & \longrightarrow & \widehat{\mathcal{Q}}_{\mathcal{F}}(-X') \longrightarrow 0 \\
& & \downarrow & & \downarrow & & \downarrow \\
0 & \longrightarrow & \mathcal{T}(\mathcal{Q}_{\mathcal{F}}) & \longrightarrow & \mathcal{Q}_{\mathcal{F}} & \longrightarrow & \widehat{\mathcal{Q}}_{\mathcal{F}} \longrightarrow 0 \\
& & \downarrow & & \downarrow & & \downarrow \\
0 & \longrightarrow & \mathcal{T}(\mathcal{Q}_{\mathcal{F}})|_{X'} & \longrightarrow & \mathcal{Q}_{\mathcal{F}}|_{X'} & \longrightarrow & \widehat{\mathcal{Q}}_{\mathcal{F}}|_{X'} \longrightarrow 0 \\
& & \downarrow & & \downarrow & & \downarrow \\
& & 0 & & 0 & & 0,
\end{array} \tag{5.11}$$

which incorporates the fact that $\mathcal{T}or_k^{\mathcal{O}_X}(\widehat{\mathcal{Q}}_{\mathcal{F}}, \mathcal{O}_{X'}) = 0$ for $k = 1, 2$ due to the regularity noticed at the beginning of the argument. Using again that $\mathcal{T}or_1^{\mathcal{O}_X}(\mathcal{Q}_{\mathcal{F}}, \mathcal{O}_{X'})$ has codimension at least three in X , a computation with the help of Lemma 2.13 yields the desired equality on the cycle level and finishes the proof of (5.10).

Note that the above considerations, the admissibility of S for $\gamma(\mathcal{F})$, and the fact that $\text{supp}(\mathcal{C}_{(\mathcal{E}|_{X'})}) \subset \text{sing}(\mathcal{E}) \cap X'$ together imply that S is admissible for $\gamma(\mathcal{F})|_{X'}$ as well.

We will next prove the following claim, from which all the statements of the proposition will follow relatively quickly.

Claim 5.26. *In the situation of Lemma 5.24, if \mathcal{F} is a polystable sheaf with $[\mathcal{F}] \in U_S$, then on each $X^{(l)}$ there exist a polystable sheaf $\mathcal{F}^{(l)}$ such that S is*

admissible for $\gamma(\mathcal{F}^{(l)})$ and such that for $0 \leq l \leq n-3$ we have

$$[\gamma(\mathcal{F}^{(l)})|_{X^{(l+1)}}] = [\gamma(\mathcal{F}^{(l)})|_{X^{(l+1)}}] = [\gamma(\mathcal{F}^{(l+1)})] , \quad (5.12)$$

where $\mathcal{F}^{(0)} := \mathcal{F}$, and where $\gamma(\mathcal{F}^{(l)})$ is defined as in (5.9).

Proof of Claim 5.26. We start by constructing the sheaf \mathcal{F}' on X' . If the restriction $\mathcal{F}|_{X'}$ is torsion free, we just put $\mathcal{F}' := \mathcal{F}|_{X'}$, since it is then polystable by the assumptions of Lemma 5.24 and Langer's Theorem, Theorem 5.4, as we noticed in the first paragraph of this proof. The other requirements on \mathcal{F}' are also satisfied by the discussion preceding the Claim.

If $\mathcal{F}|_{X'}$ is not torsion free, we will replace $\mathcal{F}|_{X'}$ by a torsion free sheaf \mathcal{F}' on X' such that $[\gamma(\mathcal{F}')] = [\gamma(\mathcal{F}|_{X'})]$. For this, we recall from Definition 2.21 that \overline{M}^μ is the closure of M^s inside M^{us} . It follows that $[\mathcal{F}]$ is the limit of points represented by locally free, stable sheaves and that \mathcal{F} may be realized as the central fiber of a one dimensional subfamily of \mathcal{Z} all of whose other fibers are locally free and stable (cf. the discussion at the end of Section 4.1). The restriction of this family to X' remains flat since \mathcal{F} is torsion free. By Langton's Theorem the special fiber $\mathcal{F}|_{X'}$ of the restricted subfamily may be replaced by a sheaf \mathcal{F}' which is semistable on X' . Following the proof of that theorem presented in [33, Appendix 2B] we will check that $\gamma(\mathcal{F}') = \gamma(\mathcal{F}|_{X'})$ on X' . In our case it will turn out to be in fact sufficient to follow the proof up to the point where one finds a torsion free central fiber \mathcal{F}' . For this we will proceed along the way sketched in [33, Exercise 2B2].

We will write \mathcal{F} for a flat family of sheaves on X' parameterized by the unit disk $\Delta \subset \mathbb{C}$. We suppose that all fibers over $\Delta^* = \Delta \setminus \{0\}$ of \mathcal{F} are semistable sheaves, that the fiber \mathcal{F}_0 has torsion and that $\mathcal{F}_0/\text{Tors}(\mathcal{F}_0)$ is polystable. We consider the torsion filtration $0 \subset \mathcal{T}_0(\mathcal{F}_0) \subset \dots \subset \mathcal{T}_{n-1}(\mathcal{F}_0) = \mathcal{F}_0$ of \mathcal{F}_0 as in [33, Def. 1.1.4]. The idea is to replace the central fiber \mathcal{F}_0 successively by sheaves having no zero-dimensional torsion to begin with, then no one-dimensional torsion, and so on. So supposing that $\mathcal{T}_0(\mathcal{F}_0) \neq 0$, set \mathcal{F}^1 to be the kernel of the composition of the natural projections $\mathcal{F} \rightarrow \mathcal{F}_0 \rightarrow \mathcal{F}_0/\mathcal{T}_0(\mathcal{F}_0)$. One gets two exact sequences

$$\begin{aligned} 0 &\longrightarrow \mathcal{T}_0(\mathcal{F}_0) \longrightarrow \mathcal{F}_0 \longrightarrow \mathcal{F}_0/\mathcal{T}_0(\mathcal{F}_0) \longrightarrow 0, \\ 0 &\longrightarrow \mathcal{F}_0/\mathcal{T}_0(\mathcal{F}_0) \longrightarrow \mathcal{F}_0^1 \longrightarrow \mathcal{T}_0(\mathcal{F}_0) \longrightarrow 0 \end{aligned}$$

from which one infers with the help of Lemma 2.14 that

$$\mathcal{F}_0^{\vee\vee} = (\mathcal{F}_0^1)^{\vee\vee} \quad \text{and} \quad \mathcal{C}(\mathcal{F}_0^1) = \mathcal{C}(\mathcal{F}_0) . \quad (5.13)$$

If \mathcal{F}_0^1 continues to have zero-dimensional torsion we construct the sheaf \mathcal{F}^2 starting from \mathcal{F}^1 and so on until we obtain some \mathcal{F}_0^m which has no torsion. The argument in [33, Appendix 2B] ensures that this happens for some $m \in$

\mathbb{N} . By (5.13) above we have $[\gamma(\mathcal{F}_0^{\vee\vee})] = [\gamma(\mathcal{F}_0^m)]$. We then look at one-dimensional torsion and so on. Eventually we obtain a subsheaf $\mathcal{F}' \subset \mathcal{F}$ coinciding with \mathcal{F} over Δ^* and such that its central fiber \mathcal{F}'_0 is torsion free with $[\gamma(\mathcal{F}_0)] = [\gamma(\mathcal{F}'_0)]$. Hence, if we set $\mathcal{F}' := \mathcal{F}'_0$, then from this equality together with (5.10) we infer that

$$[\gamma(\mathcal{F}')] = [\gamma(\mathcal{F}|_{X'})] = [\gamma(\mathcal{F})|_{X'}]. \quad (5.14)$$

Recall from the observation made right before Claim 5.26 that S is admissible for $\gamma(\mathcal{F}')$. Moreover, \mathcal{F}' is polystable, since $(\mathcal{F}')^{\vee\vee} \cong (\mathcal{F}^{\vee\vee}|_{X'})^{\vee\vee}$ by (5.13). Thus, the sheaf \mathcal{F}' fulfills the conditions of Claim 5.26 at level $l = 0$.

The sheaves $\mathcal{F}^{(2)}, \dots, \mathcal{F}^{(n-2)}$ are now constructed inductively by the same procedure, using the above family over the disk. \square

Now we use Claim 5.26 to prove Proposition 5.25. We again work in the setup of Lemma 5.24.

As a first observation, note that if a point p admits a locally free polystable representative \mathcal{F} , then for all α the restriction $\mathcal{F}|_{C^{(\alpha)}}$ remains semistable on $C^{(\alpha)}$. Hence, corresponding sections in \mathcal{L}_{n-1} will exist that do not vanish at p , and therefore ν is defined at such points. For such points we set $\mathcal{F}_{\parallel S} := \mathcal{F}|_S$, and the equality

$$\nu([\mathcal{F}]) = \nu_S([\mathcal{F}_{\parallel S}]) \quad (5.15)$$

holds by definition of the relevant linear systems and the compatibility of lifting maps discussed at the end of Section 5.6 (see the proof of Lemma 5.24, especially the isomorphism (5.7)). In fact, the above holds more generally for points p admitting a polystable representative \mathcal{F} such that the flag is fully admissible for \mathcal{F} , since then $\mathcal{F}|_{X^{(l)}}$ will be semistable on all the $X^{(l)}$ by the assumptions made in the formulation of Lemma 5.24 and Langer's Theorem, Theorem 5.4.

Now let $p = [\mathcal{F}] \in U_S$ be arbitrary. We define the sheaf $\mathcal{F}_{\parallel S}$ on S to be the sheaf $\mathcal{F}^{(n-2)}$ produced by Claim 5.26 in the case $l = n - 2$. Then $p = [\mathcal{F}]$ is a limit

$$p = [\mathcal{F}] = \lim_{i \rightarrow \infty} [\mathcal{F}_i] \quad (5.16)$$

of points $[\mathcal{F}_i] \in \overline{M}^\mu$ with \mathcal{F}_i locally free and μ -stable, and such that

$$\lim_{i \rightarrow \infty} [\mathcal{F}_i|_S] = [\mathcal{F}_{\parallel S}] \in \overline{M}_S^\mu. \quad (5.17)$$

As $\mathcal{F}_{\parallel S}$ is polystable, there exists an element of W_S that does not vanish at $[\mathcal{F}_{\parallel S}]$. Using eqs. (5.16) and (5.17), together with the observation made in the first paragraph of the proof, especially (5.15), and continuity of theta functions we see that the corresponding element in W_X does not vanish at p . The map

ν is therefore well-defined and hence holomorphic at p , as claimed. Moreover, again by continuity we conclude that

$$\nu(p) = \nu([\mathcal{F}]) = \nu_S([\mathcal{F}|_S]) ,$$

which proves commutativity of Diagram (5.8), once we have established that r_{U_S} is well-defined.

To see that r_{U_S} is well-defined, let \mathcal{F}_1 and \mathcal{F}_2 be polystable such that $[\mathcal{F}_1] = [\mathcal{F}_2] \in U_S$. By Proposition 2.22 and Remark 2.12 we then get $[\gamma(\mathcal{F}_1)] = [\gamma(\mathcal{F}_2)]$, which together with an recursive application of eq. (5.12) implies

$$[\gamma(\mathcal{F}_1^{(n-2)})] = [\gamma(\mathcal{F}_2^{(n-2)})].$$

We therefore obtain from the surface case of Proposition 2.22 that $[\mathcal{F}_1^{(n-2)}] = [\mathcal{F}_2^{(n-2)}]$, as needed.

To see that r_{U_S} is R -invariant, note that if $\mathcal{F}_1, \mathcal{F}_2$ are polystable such that $[\mathcal{F}_1], [\mathcal{F}_2] \in U_S$ are in the same R -equivalence class, then by definition this means that $[\gamma(\mathcal{F}_1)] = [\gamma(\mathcal{F}_2)]$ (see the remark made after eq. (4.2)). Hence, the same reasoning as in the previous paragraph implies that

$$r_{U_S}([\mathcal{F}_1]) = [\mathcal{F}_1^{(n-2)}] = [\mathcal{F}_2^{(n-2)}] = r_{U_S}([\mathcal{F}_2]).$$

The R -invariance of ν now follows by the commutativity of Diagram (5.8). This concludes the proof of Proposition 5.25. \square

To complete the proof of the main claim, we will have need of one further lemma, which says that the cycle components \mathcal{C} are determined completely by their restrictions to appropriately chosen complete intersection surfaces. Its proof, omitted here, follows easily by Noetherian induction.

Lemma 5.27. *Let X be a smooth projective variety and $\mathcal{O}_X(1)$ be an ample line bundle. Assume that $\mathcal{O}_X(k)$ is very ample. Let \mathcal{C}_0 be an element in the variety $\mathcal{C} := \mathcal{C}_{n-2,d}(X)$ of codimension 2 cycles of degree d on X and for any surface in X denote by $U_S'' \subset \mathcal{C}$ the open subset of cycles meeting S in at most finitely many points. Then for any fixed multi-degree divisible by k there exists a positive natural number m such that the following holds for every sufficiently general choice of m smooth complete intersection surfaces S_1, \dots, S_m of the given multi-degree in X :*

- (1) $\mathcal{C}_0 \in \cap_{j=1}^m U_{S_j}''$, and
- (2) for all $\mathcal{C}_1, \mathcal{C}_2 \in \cap_{j=1}^m U_{S_j}''$ we have $\mathcal{C}_1 = \mathcal{C}_2$ if and only if $\mathcal{C}_1 \cap S_j = \mathcal{C}_2 \cap S_j$ for all $j = 1, \dots, m$.

After these preparations, we are now in the position to give the proof of the main technical result of this section, Claim 5.23.

Proof of Claim 5.23. Fix a point $p \in \overline{M}^\mu$ and let U_S be one of the neighbourhoods of p provided by Proposition 5.14. We will write $p = [\mathcal{F}]$, for \mathcal{F} a polystable sheaf, and $\overline{\Phi}([\mathcal{F}]) = [(\mathcal{E}, \mathcal{C})]$.

Choose a sequence $0 \ll d_1 \ll d_2 \ll \cdots \ll d_{n-2}$ of positive integers that is sufficiently increasing such that a repeated application of Langer's Theorem, Theorem 5.4, and an application of Corollary 5.13 is possible. Let N be the natural number provided by Corollary 5.13 and assume without loss of generality that it coincides with the maximal number m provided by applying Lemma 5.27 to the cycle spaces $\mathcal{C}_{n-2,d}(X)$, where d runs through the finite set of degrees of cycles \mathcal{C} forming the second entry in a pair $[(\mathcal{E}, \mathcal{C})] \in \overline{M}_{\text{HYM}}$ (cf. Section 2.5.5). Choose m (general) flags satisfying the assumptions made in the Setup and such that furthermore the conclusion of Corollary 5.13 holds. Denote the chosen set of flags by Σ . Denote by S_1, \dots, S_m the end terms of the flags contained in Σ , and by $\nu_j : U_{S_j} \rightarrow \mathbb{P}(W_j^*) := \mathbb{P}(W_{S_j}^*)$ and $\nu_{S_j} : \overline{M}_{S_j}^\mu \hookrightarrow \mathbb{P}(W_j^*)$ the maps provided by Lemma 5.24 and Proposition 5.21, respectively; see also Proposition 5.25. Set $U := U_{S_1} \cap \cdots \cap U_{S_m}$ and let

$$\nu := \nu_1 \times \cdots \times \nu_m : U \longrightarrow \mathbb{P}(W_1^*) \times \cdots \times \mathbb{P}(W_m^*)$$

be the product map, which is R -invariant by Proposition 5.25. It therefore remains to show that ν separates different R -equivalence classes.

First, let $\mathcal{F}_1, \mathcal{F}_2$ be two polystable sheaves representing points $[\mathcal{F}_1], [\mathcal{F}_2] \in \overline{M}^\mu$ such that $\mathcal{E}_1 := \text{Gr}(\mathcal{F}_1)^{\vee\vee}$ and $\mathcal{E}_2 := \text{Gr}(\mathcal{F}_2)^{\vee\vee}$ are not isomorphic. Then, there exists a flag $(X^{(l)})$ in Σ that is fully admissible for both $[\gamma(\mathcal{F}_1)] = [(\mathcal{E}_1, \mathcal{C}_1)]$ and $[\gamma(\mathcal{F}_2)] = [(\mathcal{E}_2, \mathcal{C}_2)]$. Let $S_j \subset X$ be the corresponding surface. If we let $r_{U_{S_j}} : U_{S_j} \rightarrow \overline{M}_{S_j}^\mu$ be the map produced by Proposition 5.25, the j -th component of ν is equal to

$$\nu_j([\mathcal{F}_i]) = \nu_{S_j}(r_{U_{S_j}}([\mathcal{F}_i])) \quad \text{for } i = 1, 2. \quad (5.18)$$

Since $[\gamma(\mathcal{F}_1^{(n-2)})] = [\gamma(\mathcal{F}_1)|_S] \neq [\gamma(\mathcal{F}_2)|_S] = [\gamma(\mathcal{F}_2^{(n-2)})]$, it follows from Corollary 5.13 and Claim 5.26 that $r_{U_{S_j}}([\mathcal{F}_1]) = [\mathcal{F}_1^{(n-2)}]$ and $r_{U_{S_j}}([\mathcal{F}_2]) = [\mathcal{F}_2^{(n-2)}]$ define different points in the moduli space $\overline{M}_{S_j}^\mu$ (cf. Remark 4.16). From the defining property of ν_{S_j} (see Proposition 5.21) and from eq. (5.18) above we hence conclude that $\nu([\mathcal{F}_1]) \neq \nu([\mathcal{F}_2])$.

Second, write $[\mathcal{F}_1]$ and $[\mathcal{F}_2]$ as above for two points of \overline{M}^μ , where \mathcal{F}_1 and \mathcal{F}_2 are polystable, and let $[\gamma(\mathcal{F}_1)] = [(\mathcal{E}_1, \mathcal{C}_1)]$, $[\gamma(\mathcal{F}_2)] = [(\mathcal{E}_2, \mathcal{C}_2)]$ be such that $\mathcal{C}_1 \neq \mathcal{C}_2$. Note that by the definition of admissibility for any $[\mathcal{E}'] \in U$ with associated point $\overline{\Phi}([\mathcal{E}']) = [(\mathcal{E}, \mathcal{C})] \in \overline{M}_{\text{HYM}}$, we have $\mathcal{C} \in U_{S_j}''$ for all $j \in \{1, \dots, m\}$. By the choice of m made above there exists an index $j \in \{1, \dots, m\}$ such that $\mathcal{C}_1|_{S_j} \neq \mathcal{C}_2|_{S_j}$, and therefore $[(\mathcal{E}_1, \mathcal{C}_1)|_{S_j}] \neq [(\mathcal{E}_2, \mathcal{C}_2)|_{S_j}]$.

Applying Proposition 5.25 and using injectivity of ν_{S_j} again, similar to the first step we obtain

$$\nu_j([\mathcal{F}_1]) = \nu_{S_j}(r_{U_{S_j}}([\mathcal{F}_1])) \neq \nu_{S_j}(r_{U_{S_j}}([\mathcal{F}_2])) = \nu_j([\mathcal{F}_2]) .$$

This concludes the proof of Claim 5.23. \square

Proof of Theorem 1.2. The existence of a complex structure on $\overline{M}_{\text{HYM}}$ making the map $\overline{\Phi} : \overline{M}^\mu \rightarrow \overline{M}_{\text{HYM}}$ holomorphic follows from Cartan's criterion, Theorem 5.3, together with Claim 5.23.

To prove the other statements, recall that we endowed M_{HYM}^* with the complex structure such that $\Phi : M^s \rightarrow M_{\text{HYM}}^*$ becomes biholomorphic. Consequently, the induced map $\Phi^{wn} : (M^s)^{wn} \rightarrow (M_{\text{HYM}}^*)^{wn}$ of weak normalisations is likewise biholomorphic. From the discussion in the paragraph before Definition 2.21 we hence conclude that we have a natural map $\iota : (M_{\text{HYM}}^*)^{wn} \rightarrow \overline{M}_{\text{HYM}}$ fitting into the following commutative diagram:

$$\begin{array}{ccc} (M^s)^{wn} & \hookrightarrow & \overline{M}^\mu \\ \Phi^{wn} \downarrow \cong & & \downarrow \overline{\Phi} \\ (M_{\text{HYM}}^*)^{wn} & \xrightarrow{\iota} & \overline{M}_{\text{HYM}} . \end{array}$$

As the image of $(M^s)^{wn}$ is a $\overline{\Phi}$ -saturated subset on which the equivalence relation R built from Φ is trivial, ι is an open embedding, whose complement is the image of $\overline{M}^\mu \setminus (M^s)^{wn}$ under $\overline{\Phi}$, and hence Zariski closed.

The natural map $(\overline{M}^{\text{GM}})^{wn} \rightarrow \overline{M}^\mu$ (see (2.17)) and its birationality have already been discussed at the end of Section 2.8. \square

Proof of Corollary 1.3. Since $\overline{\Phi}$ is finite and surjective, and since \overline{M}^μ is projective and therefore in particular Moishezon, the complex space $\overline{M}_{\text{HYM}}$ is likewise Moishezon (see [1, Chap. V., Cor. 11], but note that the proof is much easier for finite maps than for the general case). It hence follows from a result of Artin (see [2, Thm. 7.3]) that $\overline{M}_{\text{HYM}}$ is the analytification of a proper algebraic space. The generalisation of Serre's GAGA to holomorphic maps between analytifications of proper algebraic spaces (see again *loc. cit.*), then implies that $\overline{\Phi}$ is (induced by) a morphism of algebraic spaces. The final statement follows from the corresponding statement in the analytic category, which is contained in Theorem 1.2, and the fact that the analytification of the algebraic weak normalisation of M^s is naturally biholomorphic to the weak normalisation of M^s as a complex space (cf. [30, Sect. 2.3]). \square

REFERENCES

1. Vincenzo Ancona and Giuseppe Tomassini, *Modifications analytiques*, Lecture Notes in Mathematics, vol. 943, Springer-Verlag, Berlin-New York, 1982. MR 673560
2. M. Artin, *Algebraization of formal moduli. II. Existence of modifications*, Ann. of Math. (2) **91** (1970), 88–135. MR 0260747
3. Shigetoshi Bando and Yum-Tong Siu, *Stable sheaves and Einstein-Hermitian metrics*, Geometry and Analysis on Complex Manifolds (River Edge, NJ) (T. Mabuchi, J. Noguchi, and T. Ochiai, eds.), World Scientific, 1994, pp. 39–50.
4. Daniel Barlet, *Convexité de l'espace des cycles*, Bull. Soc. Math. France **106** (1978), no. 4, 373–397. MR 518045
5. Daniel Barlet and Jón Magnússon, *Cycles analytiques complexes. II.*, Cours Spécialisés [Specialized Courses], Société Mathématique de France, Paris, to appear.
6. Errett Bishop, *Conditions for the analyticity of certain sets*, Michigan Math. J. **11** (1964), 289–304. MR 0168801
7. Nicholas Buchdahl, Andrei Teleman, and Matei Toma, *A continuity theorem for families of sheaves on complex surfaces*, J. Topol. **10** (2017), no. 4, 995–1028. MR 3743066
8. Henri Cartan, *Quotients of complex analytic spaces*, Contributions to function theory (Internat. Colloq. Function Theory, Bombay, 1960), Tata Institute of Fundamental Research, Bombay, 1960, pp. 1–15. MR 0139769
9. Debraj Chakrabarti and Mei-Chi Shaw, *The Cauchy-Riemann equations on product domains*, Math. Ann. **349** (2011), no. 4, 977–998. MR 2777041
10. Xuemiao Chen and Song Sun, *Analytic tangent cones of admissible Hermitian-Yang-Mills connections*, to appear in Geom. Topol., preprint <http://arxiv.org/abs/1806.11247>.
11. ———, *Algebraic tangent cones of reflexive sheaves*, IMRN, advance article (2018), <https://doi.org/10.1093/imrn/rny276>, will appear in print.
12. ———, *Reflexive sheaves, Hermitian-Yang-Mills connections, and tangent cones*, preprint (2019), <https://arxiv.org/abs/1912.09561v1>.
13. Georgios Daskalopoulos and Richard Wentworth, *Convergence properties of the Yang-Mills flow on Kähler surfaces*, J. Reine Angew. Math. **575** (2004), 69–99.
14. Igor Dolgachev, *Lectures on invariant theory*, London Mathematical Society Lecture Note Series, vol. 296, Cambridge University Press, Cambridge, 2003. MR 2004511
15. Simon Donaldson, *Anti self-dual Yang-Mills connections over complex algebraic surfaces and stable vector bundles*, Proc. London Math. Soc. **50** (1985), 1–26.
16. ———, *Connections, cohomology and the intersection forms of 4-manifolds*, J. Differential Geom. **24** (1986), no. 3, 275–341. MR 868974
17. ———, *Infinite determinants, stable bundles, and curvature*, Duke Math. J. **54** (1987), 231–247.
18. ———, *Compactification and completion of Yang-Mills moduli spaces*, Differential geometry (Peñíscola, 1988), Lecture Notes in Math., vol. 1410, Springer, Berlin, 1989, pp. 145–160. MR 1034277
19. ———, *Polynomial invariants for smooth four-manifolds*, Topology **29** (1990), no. 3, 257–315. MR 1066174
20. Simon Donaldson and Peter Kronheimer, *The geometry of four-manifolds*, Oxford Science, Clarendon Press, Oxford, 1990.

21. Simon Donaldson and Ed Segal, *Gauge theory in higher dimensions, II*, Surveys in differential geometry. Volume XVI. Geometry of special holonomy and related topics, Surv. Differ. Geom., vol. 16, Int. Press, Somerville, MA, 2011, pp. 1–41. MR 2893675
22. Geir Ellingsrud and Manfred Lehn, *Irreducibility of the punctual quotient scheme of a surface*, Ark. Mat. **37** (1999), no. 2, 245–254. MR 1714770
23. G. B. Folland and J. J. Kohn, *The Neumann problem for the Cauchy-Riemann complex*, Princeton University Press, Princeton, N.J.; University of Tokyo Press, Tokyo, 1972, Annals of Mathematics Studies, No. 75. MR 0461588
24. Akira Fujiki and Georg Schumacher, *The moduli space of Hermite-Einstein bundles on a compact Kähler manifold*, Proc. Japan Acad. Ser. A Math. Sci. **63** (1987), no. 3, 69–72. MR 894698
25. William Fulton, *Intersection theory*, second ed., Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge, vol. 2, Springer-Verlag, Berlin, 1998. MR 1644323
26. David Gieseker, *On the moduli of vector bundles on an algebraic surface*, Ann. of Math. (2) **106** (1977), no. 1, 45–60. MR 466475
27. David Gilbarg and Neil Trudinger, *Elliptic partial differential equations of second order*, second ed., Springer, New York, 1978.
28. Daniel Greb, Stefan Kebekus, and Thomas Peternell, *Étale fundamental groups of Kawamata log terminal spaces, flat sheaves, and quotients of abelian varieties*, Duke Math. J. **165** (2016), no. 10, 1965–2004. MR 3522654
29. Daniel Greb, Julius Ross, and Matei Toma, *Moduli of vector bundles on higher-dimensional base manifolds — Construction and Variation*, Internat. J. Math. **27** (2016), no. 7, 1650054 (27). MR 3605660
30. Daniel Greb and Matei Toma, *Compact moduli spaces for slope-semistable sheaves*, Algebr. Geom. **4** (2017), no. 1, 40–78. MR 3592465
31. Alexander Grothendieck, *Techniques de construction et théorèmes d’existence en géométrie algébrique. IV. Les schémas de Hilbert*, Séminaire Bourbaki, Vol. 6, Soc. Math. France, Paris, 1995, pp. Exp. No. 221, 249–276. MR 1611822
32. Lars Hörmander, *L^2 estimates and existence theorems for the $\bar{\partial}$ operator*, Acta Math. **113** (1965), 89–152. MR 0179443
33. Daniel Huybrechts and Manfred Lehn, *The geometry of moduli spaces of sheaves*, second ed., Cambridge Mathematical Library, Cambridge University Press, Cambridge, 2010. MR 2665168
34. Shoshichi Kobayashi, *Differential geometry of complex vector bundles*, Publications of the Mathematical Society of Japan, vol. 15, Iwanami Shoten and Princeton University Press, Princeton, NJ, 1987, Kanô Memorial Lectures, 5. MR 909698 (89e:53100)
35. János Kollár, *Quotients by finite equivalence relations*, Current Developments in Algebraic Geometry, MSRI Publications, vol. 59, Cambridge University Press, Cambridge, 2011, with an appendix by Claudiu Raicu, pp. 227–256.
36. Adrian Langer, *Semistable sheaves in positive characteristic*, Ann. of Math. (2) **159** (2004), no. 1, 251–276. MR 2051393
37. Joseph Le Potier, *Fibré déterminant et courbes de saut sur les surfaces algébriques*, Complex projective geometry (Trieste, 1989/Bergen, 1989), London Math. Soc. Lecture Note Ser., vol. 179, Cambridge Univ. Press, Cambridge, 1992, pp. 213–240. MR 1201385
38. Jun Li, *Algebraic geometric interpretation of Donaldson’s polynomial invariants*, J. Differential Geom. **37** (1993), no. 2, 417–466. MR 1205451

39. M. Lübke and C. Okonek, *Moduli spaces of simple bundles and Hermitian-Einstein connections*, Math. Ann. **276** (1987), no. 4, 663–674. MR 879544
40. Martin Lübke and Andrei Teleman, *The Kobayashi-Hitchin correspondence*, World Scientific Publishing Co., Inc., River Edge, NJ, 1995. MR 1370660
41. Masaki Maruyama, *Openness of a family of torsion free sheaves*, J. Math. Kyoto Univ. **16** (1976), no. 3, 627–637. MR 0429899
42. ———, *On boundedness of families of torsion free sheaves*, J. Math. Kyoto Univ. **21** (1981), no. 4, 673–701. MR 637512
43. V. B. Mehta and A. Ramanathan, *Semistable sheaves on projective varieties and their restriction to curves*, Math. Ann. **258** (1981/82), no. 3, 213–224. MR 649194
44. Kimio Miyajima, *Kuranishi family of vector bundles and algebraic description of the moduli space of Einstein-Hermitian connections*, Publ. Res. Inst. Math. Sci. **25** (1989), no. 2, 301–320. MR 1003790
45. John W. Morgan, *Comparison of the Donaldson polynomial invariants with their algebro-geometric analogues*, Topology **32** (1993), no. 3, 449–488. MR 1231956
46. Aaron Naber and Daniele Valtorta, *Energy identity for stationary Yang Mills*, Invent. Math. **216** (2019), no. 3, 847–925. MR 3955711
47. Hiraku Nakajima, *Compactness of the moduli space of Yang-Mills connections in higher dimensions*, J. Math. Soc. Japan **40** (1988), no. 3, 383–392.
48. Christian Okonek, Michael Schneider, and Heinz Spindler, *Vector bundles on complex projective spaces*, Progress in Mathematics, vol. 3, Birkhäuser, Boston, Mass., 1980. MR 561910
49. Mihai Pavel, Ph.D. thesis, IECL, Université de Lorraine, in progress.
50. Steven Sedlacek, *A direct method for minimizing the Yang-Mills functional over 4-manifolds*, Comm. Math. Phys. **86** (1982), no. 4, 515–527. MR 679200
51. Mei-Chi Shaw, *Global solvability and regularity for $\bar{\partial}$ on an annulus between two weakly pseudoconvex domains*, Trans. Amer. Math. Soc. **291** (1985), no. 1, 255–267. MR 797058
52. ———, *The closed range property for $\bar{\partial}$ on domains with pseudoconcave boundary*, Complex analysis, Trends Math., Birkhäuser/Springer Basel AG, Basel, 2010, pp. 307–320. MR 2885124
53. Benjamin Sibley, *Asymptotics of the Yang-Mills flow for holomorphic vector bundles over Kähler manifolds: the canonical structure of the limit*, J. Reine Angew. Math. **706** (2015), 123–191. MR 3393366
54. Benjamin Sibley and Richard A. Wentworth, *Analytic cycles, Bott-Chern forms, and singular sets for the Yang-Mills flow on Kähler manifolds*, Adv. Math. **279** (2015), 501–531. MR 3345190
55. The Stacks Project Authors, *Stacks Project*, <http://stacks.math.columbia.edu>, 2020.
56. Terrence Tao and Gang Tian, *A singularity removal theorem for Yang-Mills fields in higher dimensions*, J. Amer. Math. Soc. **17** (2004), no. 3, 557–593.
57. Gang Tian, *Gauge theory and calibrated geometry, I*, Ann. Math. **151** (2000), 193–268.
58. Gang Tian and Baozhong Yang, *Compactification of the moduli spaces of vortices and coupled vortices*, J. Reine. Angew. Math. **553** (2002), 17–41.
59. Karen Uhlenbeck, *A priori estimates for Yang-Mills fields*, unpublished manuscript.
60. ———, *Connections with L^p bounds on curvature*, Comm. Math. Phys. **83** (1982), no. 1, 31–42.

61. ———, *Removable singularities in Yang-Mills fields*, Comm. Math. Phys. **83** (1982), no. 1, 11–29.
 62. Karen Uhlenbeck and Shing-Tung Yau, *On the existence of Hermitian-Yang-Mills connections in stable vector bundles*, Comm. Pure Appl. Math. **39** (1986), no. 2, 257–293.
 63. Katrin Wehrheim, *Uhlenbeck compactness*, EMS Series of Lectures in Mathematics, European Mathematical Society (EMS), Zürich, 2004. MR 2030823
 64. Baozhong Yang, *The uniqueness of tangent cones for Yang-Mills connections with isolated singularities*, Adv. Math. **180** (2003), no. 2, 648–691. MR 2020554
-

DANIEL GREB, ESSENER SEMINAR FÜR ALGEBRAISCHE GEOMETRIE UND ARITHMETIK,
FAKULTÄT FÜR MATHEMATIK, UNIVERSITÄT DUISBURG–ESSEN, 45112 ESSEN, GERMANY

Email address: daniel.greb@uni-due.de

URL: <https://www.esaga.uni-due.de/daniel.greb/>

BENJAMIN SIBLEY, UNIVERSITÉ LIBRE DE BRUXELLES, CAMPUS DE LA PLAINE, BOULEVARD
DU TRIOMPHE, ACC.2, 1050 BRUXELLES, BELGIQUE

Email address: bsibley1060@gmail.com or benjamin.sibley@ulb.ac.be

MATEI TOMA, UNIVERSITÉ DE LORRAINE, CNRS, IECL, F-54000 NANCY, FRANCE

Email address: Matei.Toma@univ-lorraine.fr

URL: <http://www.iecl.univ-lorraine.fr/~Matei.Toma/>

RICHARD WENTWORTH, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MARYLAND,
COLLEGE PARK, MD 20742, USA

Email address: raw@umd.edu

URL: <http://www.math.umd.edu/~raw/>