Holomorphic Vector Bundles over Riemann Surfaces

Hugues Bellemare, Department of Mathematics and Statistics
McGill University, Montreal
August 15, 2021

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of

Master of Mathematics

Abstract

We study holomorphic vector bundles over Riemann surfaces. After recalling the basic concepts of the theory, we prove that every holomorphic vector bundle over a non-compact Riemann surface is trivial using methods from functional analysis. We then turn over to the case of a compact Riemann surface X, where we study an infinite-dimensional universal space parametrizing the holomorphic vector bundles over X of the same rank and degree, although with a lot of redundancies. Following the pioneering work of Atiyah and Bott, we use ideas from Morse theory to exhibit a stratification of that space that eventually gives us an inductive procedure to compute the equivariant cohomology of the minimal stratum, which consists of the "semi-stable" holomorphic vector bundles.

Abrégé

On étudie les fibrés vectoriels holomorphes sur des surfaces de Riemann. Après une révision des concepts de base de la théorie, on prouve que tout fibré vectoriel holomorphe sur une surface de Riemann non compacte est trivial à l'aide d'outils provenant de l'analyse fonctionnelle. On se tourne ensuite vers le cas d'une surface de Riemann compacte X, où l'on étudie un espace universel de dimension infinie paramétrisant les fibrés vectoriels holomorphes sur X de même rang et degré, quoiqu'avec beaucoup de redondance. Suivant les travaux d'Atiyah et Bott, on s'inspire d'idées provenant de la théorie de Morse afin d'exhiber une stratification de cet espace qui nous conduira finalement vers une méthode inductive pour calculer la cohomologie équivariante de la strate minimale, celle-ci étant composée des fibrés vectoriels holomorphes dits "semi-stables".

Acknowledgements

Of course, I thank Prof. Hurtubise for his excellent advice, his patience and all the time he took to discuss with me over the last two years, whether in-person or remotely. I also want to thank the Natural Sciences and Engineering Research Council of Canada (NSERC), the Fonds de recherche du Québec – Nature et techonologies (FRQNT) as well as the McGill Mathematics and Statistics Department for their financial support. Finally, I want to thank my family and my significant other for their ever-lasting support, even though this means exposing themselves to hear about mathematics a little too often for their taste.

Contents

Abstract i

Abrégé ii

Acknowledgements iii

1 Introduction **1**

2 Holomorphic Vector Bundles **3**

3 Non-Compact Riemann Surfaces **10**

3.1 Runge's approximation theorem 10

3.2 Two vanishing theorems 14

3.3 Triviality of holomorphic vector bundles 18

4 Compact Riemann Surfaces **23**

4.1 Review of Morse theory 23

4.1.1 Elementary Morse theory 23

4.1.2 Extension of Morse theory 25

4.1.3 Equivariant cohomology 26

4.1.4 Stratifications 27

4.2 Defining the stratification 28
4.2.1 The space of holomorphic structures 28
4.2.2 The Harder-Narasimhan filtration 30
4.2.3 The emergence of the stratification 33
4.2.4 Notions of convexity 34
4.3 The Yang-Mills functional 36
4.4 Completion of the Morse picture 38
4.5 Equivariant cohomology of $\mathscr{C}^{ss}$ 45

Chapter 1

Introduction

Holomorphic vector bundles have been around for a long time. They have been studied with tools coming from many different disciplines, ranging from differential geometry to algebraic geometry and not forgetting symplectic geometry. In this thesis, we undertake the study of these objects in the simplest case where the base manifold is a Riemann surface. When this Riemann surface is non-compact, there is not much to say: the non-compactness gives enough space to allow the existence of global holomorphic sections, so much that the holomorphic vector bundle has no choice but to be trivial. However, the story gets more complicated when the Riemann surface is compact.

One of the reasons for the interest towards holomorphic vector bundles over compact Riemann surfaces is that they are one of the first examples where the geometric invariant theory of Mumford was successful, i.e. to find an appropriate subclass of these objects (the stables ones) for which the classifying space is a nice geometric one. The first tool that was key to the future developments of the theory was the Narasimhan–Seshadri criterion (cf. [8]), which established a link between stable holomorphic (unitary) bundles and representations of a central extension of the fundamental group of the Riemann surface. Donaldson then linked the Narasimhan–Seshadri criterion to the existence of a connection whose curvature is subject to specific restrictions. Atiyah and Bott, on their side, recast the study of such bundles (of fixed rank and degree) into differential-geometric terms by considering an infinite-dimensional universal space parametrizing these objects (although with a lot of redundancies) and exhibiting a stratification of this space reminiscent to the ones encountered in Morse theory. In fact, they found that the stratification behaved as if it was the Morse stratification of the Yang–Mills functional, and their intuition was proven correct by Daskalopoulos in [2]. Using this analogy, they were able to give an inductive procedure for computing the equivariant cohomology of the minimal stratum, which consists of the semi-stable holomorphic vector bundles. They were finally able to use this information to compute the ordinary cohomology of the moduli space associated to the semi-stable holomorphic bundles in the cases where the rank is either equal to 2 or coprime to the degree. These results were finally extended to higher ranks by Zagier.

The layout of this thesis is as follows. First, since holomorphic vector bundles are going to be the main characters of the story, we collect various basic results concerning holomorphic vector bundles in chapter 2. In chapter 3, we deal with the case of a non-compact Riemann surface, where we show that the triviality of holomorphic vector bundles over them follows from their particular topological and functional-analytic (cohomological) properties. Finally, in chapter 4, we work out in more detail the case of a compact Riemann surface that was highlighted above. Throughout the thesis, no originality is claimed, and due references are given where appropriate.

Chapter 2

Holomorphic Vector Bundles

Definition 2.1. Let E and X be two topological spaces together with a continuous surjection $p : E \to X$ between them. The map $p : E \to X$, or simply E, is said to be a complex vector bundle of rank n if for every point $x \in X$:

i) the fiber $E_x := p^{-1}(x)$ has the structure of a complex vector space of dimension n;

ii) there exists a neighborhood U of x and a fiber-preserving homeomorphism h of $E_U := p^{-1}(U)$ onto $U \times \mathbb{C}^n$ such that the restriction $h|_{E_x}$ is a vector space isomorphism of the fiber E_x onto $\{x\} \times \mathbb{C}^n \cong \mathbb{C}^n$. Recall that in this case, that h preserves the fibers means that the diagram

$$\begin{tikzcd} E_U \arrow[rr, "h_U"] \arrow[dr, "p"'] & & U \times \mathbb{C}^n \arrow[dl, "\mathrm{pr}_U"] \\ & U & \end{tikzcd}$$

commutes, where pr_U is the canonical projection on U.

The map $h : E_U \to U \times \mathbb{C}^n$ is called a local trivialization of E over U. If $\mathfrak{U} = \{U_i\}_{i\in I}$ is an open cover of X and we have local trivializations $h_i : E|_{U_i} \to U_i \times \mathbb{C}^n$ for all $i \in I$, then the family $\{h_i\}_{i\in I}$ is called an atlas of E.

Definition 2.2. Let $p : E \to X$ be a complex vector bundle of rank n, where E and X are complex manifolds. An atlas $\mathfrak{A} = \{h_i : E_{U_i} \to U_i \times \mathbb{C}^n\}$ of E is said to be holomorphic if the local trivializations h_i are biholomorphic. Two atlases $\mathfrak{A} = \{h_i : E_{U_i} \to U_i \times \mathbb{C}^n\}_{i\in I}$ and $\mathfrak{A}' = \{h'_j : E_{U'_j} \to U'_j \times \mathbb{C}^n\}_{j\in J}$ are said to be compatible if the maps

$$h_i \circ (h'_j)^{-1} : (U_i \cap U'_j) \times \mathbb{C}^n \to (U_i \cap U'_j) \times \mathbb{C}^n$$

are biholomorphic for all $i \in I$, $j \in J$.

Notice that compatibility of atlases is an equivalence relation. Moreover, given two compatible atlases $\mathfrak{A}$ and $\mathfrak{A}'$, one is holomorphic if and only if the other is. Thus, either all the atlases inside the same equivalence class are holomorphic or none of them are. This is the motivation behind the following definition.

Definition 2.3. A holomorphic vector bundle is a complex vector bundle $p : E \to X$, where E and X are complex manifolds and p is holomorphic, together with an equivalence class of holomorphic atlases.

In the future, when speaking of an atlas of a holomorphic vector bundle E, we will always assume that it belongs to the underlying equivalence class of holomorphic atlases. In particular, it will always be holomorphic.

Theorem 2.4. *Let $E \to X$ be a holomorphic vector bundle and let $\{h_i : E_{U_i} \to U_i \times \mathbb{C}^n\}$ be an atlas of E. On the intersections $U_i \cap U_j$, consider the maps*

$$\varphi_{ij} \coloneqq h_i \circ h_j^{-1} : (U_i \cap U_j) \times \mathbb{C}^n \to (U_i \cap U_j) \times \mathbb{C}^n.$$

There exist holomorphic maps $g_{ij} : U_i \cap U_j \to$ GL(n,$\mathbb{C}$) (as usual, GL(n,$\mathbb{C}$) *is equipped with the topology it inherits as a subspace of $\mathbb{C}^{n^2}$) such that $\varphi_{ij}(x,t) = (x, g_{ij}(x)t)$ for all $(x,t) \in (U_i \cap U_j) \times \mathbb{C}^n$. Moreover, on triple intersections $U_i \cap U_j \cap U_k$, these maps satisfy the cocycle condition $g_{ij}g_{jk} = g_{ik}$.*

Proof. Since h_i and h_j are fiber-preserving biholomorphisms, then so is φ_{ij}. Moreover, the restriction of φ_{ij} to a fiber gives an isomorphism of vector spaces $\varphi_{ij}|_{\{x\}\times\mathbb{C}^n} : \mathbb{C}^n \to \mathbb{C}^n$. Thus, for $x \in U_i \cap U_j$, the assignment $g_{ij}(x) \coloneqq \varphi_{ij}|_{\{x\}\times\mathbb{C}^n}$ defines an element of GL(n,$\mathbb{C}$). In other words, g_{ij} is defined exactly so that the equality $\varphi_{ij}(x,t) = (x, g_{ij}(x)t)$ holds for all $(x,t) \in (U_i \cap U_j) \times \mathbb{C}^n$. Considering g_{ij} as a function into $\mathbb{C}^{n^2}$, it is holomorphic if and only if the functions

$$\begin{aligned} g_{ij,k} : U_i \cap U_j &\longrightarrow \mathbb{C}^n \\ x \quad &\longmapsto g_{ij}(x)(e_k) \end{aligned}$$

are holomorphic for $k = 1,\dots,n$. But the functions $g_{ij,k}$ are holomorphic, each being the composition of the holomorphic maps

$$U_i \cap U_j \times \{e_k\} \xrightarrow{\varphi_{ij}} U_i \cap U_j \times \mathbb{C}^n \xrightarrow{\mathrm{pr}_{\mathbb{C}^n}} \mathbb{C}^n,$$

where $\mathrm{pr}_{\mathbb{C}^n}$ denotes the projection onto $\mathbb{C}^n$. Finally, the cocycle condition $g_{ij}g_{jk} = g_{ik}$ follows from the corresponding relation for the maps φ_{ij}. $\square$

Definition 2.5. Let $E \to X$ be a holomorphic vector bundle and $\{h_i : E_{U_i} \to U_i \times \mathbb{C}^n\}$ be an atlas of E. The maps g_{ij} constructed in the last theorem are called the transition functions, and the family (g_{ij}) is called the cocycle associated to the atlas $\{h_i\}$.

Of course, there is a reason behind the use of the word "cocycle" for the family of transition functions (g_{ij}). For an open set U of a complex manifold X, let $\mathrm{GL}(n,\mathscr{O}(U))$ denote the set of $n \times n$ matrices with coefficients in $\mathscr{O}(U)$, the holomorphic functions on U. If $V \subset U$ is an open subset, we have a natural restriction map from $\mathrm{GL}(n,\mathscr{O}(U))$ to $\mathrm{GL}(n,\mathscr{O}(V))$ that restricts each entry to the subset V. This data defines a sheaf, which we denote by $\mathrm{GL}(n,\mathscr{O})$.

Given the sheaf $\mathrm{GL}(n,\mathscr{O})$ and an atlas $\mathfrak{U} = \{U_i\}_{i\in I}$ of a complex manifold X, one can talk of (Čech) cocycles with values in $\mathrm{GL}(n,\mathscr{O})$ with respect to $\mathfrak{U}$: these are merely elements of the set $Z^1(\mathfrak{U}, \mathrm{GL}(n,\mathscr{O}))$, which consists of families

$$(f_{ij}) \in \prod_{(i,j)\in I^2} \mathrm{GL}(n,\mathscr{O}(U_i \cap U_j))$$

satisfying the cocycle condition $f_{ik} = f_{ij}f_{jk}$ on triple intersections for all $i,j,k \in I$. As Theorem 2.4 shows, to any holomorphic vector bundle $E \to X$ with atlas $\mathfrak{A}$ we can associate its cocycle of transition functions. Conversely, given a complex vector bundle of complex manifolds $E \to X$ with an atlas $\mathfrak{A} = \{h_i\}$ such that the associated transition functions g_{ij} are holomorphic, then the atlas $\mathfrak{A}$ is holomorphic and E becomes a holomorphic vector bundle. In fact, the proof of this last statement generalizes to a construction from which a lot of examples of holomorphic vector bundles naturally arise.

Theorem 2.6. *Let $p : E \to X$ be a surjective map between a set E and a complex manifold X and let $\mathfrak{U} = \{z_i : U_i \to \mathbb{C}^d\}$ be an atlas of X. Suppose that each fiber $E_x = p^{-1}(x)$ has the structure of a complex vector space of dimension n. Further suppose that there are fiber-preserving maps*

$$\begin{tikzcd} h_i : E_{U_i} = \coprod_{x\in U_i} E_x \arrow[rr] \arrow[dr, "p"'] & & U_i \times \mathbb{C}^n \arrow[dl, "pr_{U_i}"] \\ & U_i & \end{tikzcd}$$

that restricts to linear isomorphisms on fibers and whose family of associated transition functions (g_{ij}) is a cocycle in $Z^1(\mathfrak{U}, \mathrm{GL}(n,\mathscr{O}))$. Then $p : E \to X$ can be made into a holomorphic vector bundle with atlas $\mathfrak{A} = \{h_i\}$.

Proof. We begin by putting a topology on E. Over U_i, we can use the bijection h_i to pull back the topology on $U_i \times \mathbb{C}^n$: we define a subset $V \subset E_{U_i}$ to be open if and only if its image $h_i(V)$ is open in $U_i \times \mathbb{C}^n$. We declare a subset $V \subset E$ to be open if and only if its intersection with every E_{U_i} is open. With this topology, the sets E_{U_i} are open, the functions h_i are automatically homeomorphisms and the surjection p becomes continuous. This turns E into a complex vector bundle of rank n.

The homeomorphisms h_i together with the existing charts on X now induce a complex structure on E. Indeed, $\{E_{U_i}\}$ is an open cover of E and we have complex charts

$$\varphi_i = (z_i \times \mathrm{id}) \circ h_i : E_{U_i} \longrightarrow \mathbb{C}^d \times \mathbb{C}^n = \mathbb{C}^{d+n}.$$

The only thing we need to verify is that these charts are pairwise compatible. But since the transition functions g_{ij} are assumed to be holomorphic, the map

$$\varphi_i \circ \varphi_j^{-1} = (z_i \times \mathrm{id}) \circ (h_i \circ h_j^{-1}) \circ (z_j^{-1} \times \mathrm{id}) : (x,t) \longmapsto \big((z_i \circ z_j^{-1})(x), g_{ij}(x)t\big)$$

also is. The space E is thus equipped with the structure of a complex manifold for which the maps h_i are biholomorphic and the projection p is holomorphic. In other words, E is a holomorphic vector bundle over X and $\{h_i\}$ is a holomorphic atlas of E. □

Remark 2.7. There is an alternative way to define holomorphic vector bundles, this time in terms of transition functions instead of local trivializations. Namely, we could define an atlas of a complex vector bundle E over a complex manifold X to be holomorphic if the associated transition functions are holomorphic. Two holomorphic atlases would then be declared compatible if their union is again holomorphic. Finally, we would define a holomorphic vector bundle to be a complex vector bundle over a complex manifold together with an equivalence class of holomorphic atlases. Theorem 2.6 shows how to reconcile this definition with ours.

Example 2.8 (Holomorphic tangent bundle)**.** The method highlighted in Theorem 2.6 is typically used to show that the family of tangent spaces of a complex manifold X forms a holomorphic vector bundle. Indeed, let $E := \coprod_{p\in X} T_pX$ be the disjoint union of all the tangent spaces of X together with the natural projection $p : E \to X$ that maps a tangent vector to its basepoint. If $z : U \to \mathbb{C}^n$ is a complex chart on X, then for every point $p \in U$ there is a natural basis for T_pX given by

$$\partial_{z_1}|_p := \frac{\partial}{\partial z_1}\bigg|_p, \ldots, \partial_{z_n}|_p := \frac{\partial}{\partial z_n}\bigg|_p.$$

Therefore, the map $h_p : T_pX \to \mathbb{C}^n$ which sends a tangent vector to its coefficients with respect to the basis $\{\partial z_1|_p, \ldots, \partial z_n|_p\}$ is a linear isomorphism. Put together, the maps h_p give rise to a map $h_U : E_U \longrightarrow U \times \mathbb{C}^n$. After doing so for every complex chart of an atlas of X, the only thing left to verify is that the associated transition functions are holomorphic. But if $z : U \to \mathbb{C}^n$ and $w : V \to \mathbb{C}^n$ are two charts in the atlas of X, then the associated transition function is given by the Jacobian of the holomorphic function $z \circ w^{-1}$, whose entries are all holomorphic. Therefore, we can apply Theorem 2.6 to turn E into a holomorphic vector bundle over X.

Example 2.9. Similarly, one can apply Theorem 2.6 to show that the cotangent bundle of a complex manifold, its kth exterior power or even its exterior algebra are all holomorphic vector bundles.

Definition 2.10. Let $p : E \to X$ be a complex vector bundle and $U \subset X$ be an open subset. A section of E over U is a continuous function $f : U \to E$ such that $p \circ f = \mathrm{id}_U$. If $E \to X$ is a holomorphic vector bundle, we say that a section of E over U is holomorphic if it is holomorphic as a map between complex manifolds.

Let $E \to X$ be a complex vector bundle of rank n, and let f be a section of E over U. The condition $p \circ f = \mathrm{id}_U$ states that for all $x \in U$, $f(x)$ lives in E_x, the fiber over x. If

$h_i : E_{U_i} \to U_i \times \mathbb{C}^n$ is a local trivialization, then over $U_i \cap U$ the section f can be seen as a family of continuous functions $f_i : U_i \cap U \to \mathbb{C}^n$ in the following way: we simply define $f_i(x)$ by the relation $h_i(f(x)) = (x, f_i(x))$. In this case, we call the function f_i a representation of f with respect to the local trivialization h_i. It turns out that these representations are also useful to describe holomorphic sections.

Theorem 2.11. *Let $E \to X$ be a holomorphic vector bundle of rank n and $\{h_i : E_{U_i} \to U_i \times \mathbb{C}^n\}_{i \in I}$ be an atlas of E. A section f over an open set U is holomorphic if and only if all its representations $f_i : U_i \cap U \to \mathbb{C}^n$ are holomorphic.*

Proof. $\Rightarrow$) Suppose that the section f is holomorphic. Recall that the local trivializations h_i of a holomorphic vector bundle are holomorphic. Then, for any $i \in I$, the representation f_i is a composition of holomorphic functions

$$f_i : U_i \cap U \xrightarrow{f} E_{U_i \cap U} \xrightarrow{h_i} (U_i \cap U) \times \mathbb{C}^n \xrightarrow{\mathrm{pr}_{\mathbb{C}^n}} \mathbb{C}^n,$$

hence it is holomorphic as well.

$\Leftarrow$) Suppose that the representations f_i are all holomorphic. Without loss of generality, we can assume that $\{z_i : U_i \cap U \to \mathbb{C}^d\}$ is an atlas of U. Then the maps $\varphi_i := (z_i \times \mathrm{id}) \circ h_i$ form an atlas of E_U, so to show that f is holomorphic, we only need to show that $\varphi_i \circ f$ is. Writing down $\varphi_i \circ f$ explicitly, we get

$$(\varphi_i \circ f)(x) = (z_i \times \mathrm{id})(x, f_i(x)) = \big(z_i(x), f_i(x)\big),$$

which is seen to be holomorphic. $\square$

The space $\mathscr{F}(U)$ of holomorphic sections over U has the natural structure of a vector space. Using the natural restriction maps, this gives a sheaf $\mathscr{F}$ of vector spaces over X, called the sheaf of holomorphic sections of E. Elements in $\mathscr{F}(X)$ are often called global sections.

Suppose that $(g_{ij}) \in Z^1(\mathfrak{U}, \mathrm{GL}(n, \mathscr{O}))$ is the cocycle of transition functions associated to the atlas $\{h_i : E_{U_i} \to U_i \times \mathbb{C}^n\}$ of E. Looking at two representations f_i, f_j of a section f over an open set U, we have that

$$(x, f_i(x)) = h_i(f(x)) = (h_i \circ h_j^{-1} \circ h_j)(f(x)) = \varphi_{ij}(x, f_j(x)) = \big(x, g_{ij}(f_j(x))\big) \tag{2.12}$$

for all $x \in U_i \cap U_j \cap U$, which implies that f_i and f_j satisfy the relation $f_i = g_{ij} f_j$ on $U_i \cap U_j \cap U$. In particular, this implies that over a complex chart U_i, the space of sections $\mathscr{F}(U_i)$ is isomorphic to $\mathscr{O}(U_i)^n$.

A consequence of Theorem 2.11 is that we could have defined holomorphic sections strictly in terms of their representations. This is the approach we take to define meromorphic sections, which will come in handy in the next chapter.

Definition 2.13. Let E be a holomorphic vector bundle of rank n over a Riemann surface X and let $h : E_U \to U \times \mathbb{C}^n$ be a local trivialization around a point $p \in U$. A holomorphic section $f \in \mathscr{F}(U \setminus \{p\})$ can be seen as n-tuple of holomorphic functions $(f_1, \ldots, f_n) \in \mathscr{O}(U \setminus \{p\})^n$. The point p is called a pole of order m of f if

(i) all the functions f_k have a pole of order $\leq m$ at p or a removable singularity at p;

(ii) at least one of the f_k has a pole of order m at p.

A meromorphic section of E over an open set U is a holomorphic section $f \in \mathscr{F}(U \setminus A)$, where A is a discrete subset of U and every point $p \in A$ is a pole of f.

We conclude this chapter by giving equivalent conditions under which a holomorphic vector bundle is trivial.

Theorem 2.14. *Let $E \to X$ be a holomorphic vector bundle of rank n, $\{h_i : E_{U_i} \to U_i \times \mathbb{C}^n\}_{i \in I}$ an atlas of E and (g_{ij}) the associated cocycle of transition functions. The following statements are equivalent:*

(i) E is trivial, i.e. there exists a global trivialization $E \to X \times \mathbb{C}^n$;

(ii) There exist n global holomorphic sections $f_1, \ldots, f_n \in \mathscr{F}(X)$ that are linearly independent on each fiber, i.e. the vectors $f_1(x), \ldots, f_n(x) \in E_x = \mathbb{C}^n$ are linearly independent for every point $x \in X$;

(iii) The cocycle of transition functions (g_{ij}) is a coboundary, i.e. there exists a family of functions $(g_i) \in \prod_{i \in I} \mathrm{GL}(n, \mathscr{O}(U_i))$ such that $g_{ij} = g_i g_j^{-1}$ on $U_i \cap U_j$ for all $i,j \in I$.

Proof. *(i)* $\Rightarrow$ *(ii)*: Let $h : E \to X \times \mathbb{C}^n$ be a global trivialization of E. On $X \times \mathbb{C}^n$, there already exist n global holomorphic sections, namely the sections

$$\begin{aligned} \bar{e}_k : X &\longrightarrow X \times \mathbb{C}^n, \\ x &\longmapsto (x, e_k) \end{aligned}$$

where $\{e_1, \ldots, e_n\}$ denotes the canonical basis of $\mathbb{C}^n$. Using the biholomorphism $h : E \to X \times \mathbb{C}^n$, we can pull back these sections to get global holomorphic sections $f_k := h^{-1} \circ \bar{e}_k : X \to E$. By construction, the sections $f_1, \ldots, f_n$ are linearly independent on each fiber.

(ii) $\Rightarrow$ *(iii)*: Recall that every global holomorphic section f_k can be represented by an n-tuple of holomorphic funtions $(f^i_{lk})_{1 \leq l \leq n}$ with respect to the local trivialization h_i. We can combine these representations into the matrix $g_i := (f^i_{lk})_{1 \leq l,k \leq n}$, which really sits inside $\mathrm{GL}(n, \mathscr{O}(U_i))$ since the sections $f_1, \ldots, f_n$ are assumed to be linearly independent on each fiber. It follows from Equation (2.12) that

$$\forall k,\ f^i_k = g_{ij} f^j_k \implies g_i = g_{ij} g_j \implies g_{ij} = g_i g_j^{-1} \text{ on } U_i \cap U_j.$$

$(iii) \Rightarrow (i)$: We use the cochain (g_i) to build a global trivialization $h : E \to X \times \mathbb{C}^n$ from the local trivializations h_i. For $v \in E_{U_i}$, define $h(v) \coloneqq (x, g_i^{-1} t)$, where $h_i(v) = (x, t)$. This assignment is well-defined because if v also lies in E_{U_j} and $h_j(v) = (x, t')$, then

$$t \overset{(2.12)}{=} g_{ij} t' = g_i g_j^{-1} t' \implies g_i^{-1} t = g_j^{-1} t'.$$

By construction, the map h is biholomorphic and is compatible with the local trivializations h_i, thus providing the holomorphic vector bundle E with a global trivialization. $\square$

Chapter 3

Non-Compact Riemann Surfaces

The goal of this chapter is to show that holomorphic vector bundles over non-compact Riemann surfaces are always trivial. We will first exhibit some topological properties enjoyed by non-compact Riemann surfaces, eventually leading to Runge's approximation theorem. We will then spend some time studying the functional analysis of such surfaces, culminating with various theorems on the existence of some specific sections. Finally, we will be able to prove the triviality of line bundles, and an induction on the rank of the bundle will extend the result to general holomorphic vector bundles. All the results in this chapter are taken from [4], although they have been reordered and edited in the hope of conveying even more clarity.

3.1 Runge's approximation theorem

One of the theorems that we will use time and again in this chapter is the Runge approximation theorem. Its classical version for the complex plane asserts that on a simply connected domain $Y \subset \mathbb{C}$, every holomorphic function on Y can be approximated uniformly on compact sets by entire functions. Its statement for non-compact Riemann surfaces is almost identical:

Theorem 3.1 (Runge's approximation theorem, cf. [4, Theorem 25.5])**.** *Let X be a non-compact Riemann surface and $Y \subset X$ be an open subset whose complement contains no compact connected components. Then every holomorphic function on Y can be approximated uniformly on compact subsets of Y by holomorphic functions on X.*

In functional analysis, we often encounter the following strategy to construct global functions with desired properties: starting with a sequence of functions (themselves having certain key properties) defined on increasingly bigger domains, we modify them using approximation theorems to get functions that are closer together (and that retain the key properties) and then take the limit of these new functions. In order to use this kind of argument with Runge's approximation theorem, we need to know that non-compact Riemann surfaces X can be exhausted by open subsets Y satisfying the hypothesis of Theorem 3.1. The remainder of this section will be devoted to prove the existence of such an exhaustion. At this point, it is convenient to introduce some terminology.

Definition 3.2. Let X be a Riemann surface. For any subset $Y \subset X$, let $h(Y)$ denote the union of Y with all connected components of $X \setminus Y$ that are relatively compact in X. An open subset $Y \subset X$ is called Runge if $Y = h(Y)$, i.e. $X \setminus Y$ has no compact connected components.

Remark 3.3. We can already see from the definition that the hull operator h satisfies the two following properties:

(i) $h(h(Y)) = Y$ for every subset $Y \subset X$;

(ii) $Y_1 \subset Y_2 \implies h(Y_1) \subset h(Y_2)$ for every pair of subsets $Y_1, Y_2 \subset X$.

Before moving on to a closer study of h, we point out a fact that will be used without mention in the subsequent theorems.

Lemma 3.4. *Let Y be a relatively compact subset of a manifold X. Then there exists a relatively compact open neighbourhood U of $\overline{Y}$. Moreover, U can be taken to be connected.*

Proof. By assumption, $\overline{Y}$ is compact. Since X is locally Euclidean, every point of $\overline{Y}$ has a neighbourhood which is a relatively compact domain. By compactness, finitely many of these relatively compact domains, say $U_1, \ldots, U_k$, cover $\overline{Y}$. It follows that $U := U_1 \cup \cdots \cup U_k$ is a relatively compact open neighbourhood of $\overline{Y}$. Finally, we can further join the $\overline{U_k}$ with paths to get a compact connected set K and then repeat the construction (with K instead of $\overline{Y}$) to obtain a relatively compact open neighbourhood U' of $\overline{Y}$ that is connected. $\square$

Theorem 3.5. *Let Y be a subset of a Riemann surface X. Then the following hold:*

(i) If Y is closed, then $h(Y)$ is closed;

(ii) If Y is compact, then $h(Y)$ is compact.

Proof. (i) Let C_i, $i \in I$ denote the connected components of $X \setminus Y$. Being a manifold, X is locally connected. Since Y is closed, $X \setminus Y$ is an open subset of X, so it inherits the local connectedness of X. It follows that the connected components C_i are open. Let $I_0 = \{i \in I : C_i \text{ is relatively compact in } X\}$. Then

$$X \setminus h(Y) = \bigcup_{i \in I \setminus I_0} C_i$$

is also an open set, i.e. $h(Y)$ is closed.

(ii) Clearly, the result holds if $Y = \emptyset$. Assume from now on that Y is non-empty. Let C_i, $i \in I$ denote the connected components of $X \setminus Y$ (recall from (i) that the C_i are open in X) and let U be a relatively compact open neighbourhood of Y. We first claim that all the connected components C_i meet U. Otherwise, there would be some component C_i such that $C_i \subset X \setminus U$, hence we would have $\overline{C_i} \subset X \setminus U \subset X \setminus Y$ (here $\overline{C_i}$ denotes the closure of C_i in X). But the maximality of C_i as a connected component of $X \setminus Y$ implies that $C_i = \overline{C_i}$ and therefore that C_i would be both open and closed in X, which is impossible since X is connected.

Next, we claim that only finitely many of the components C_i meet ∂U. Indeed, since U is relatively compact, then ∂U is compact. Since $\partial U \subset X \setminus Y = \bigcup C_i$, it follows that finitely many of the C_i cover ∂U. The C_i being disjoint, the claim follows. Now, let $I_0 = \{i \in I : C_i \text{ is relatively compact in } X\}$. Then

$$h(Y) = Y \cup \bigcup_{i \in I_0} C_i.$$

Let $C_{i_1}, \dots, C_{i_k}$ denote the relatively compact connected components who intersect ∂U. Since all the other such components meet U but not ∂U, they must be contained in U. Therefore, $h(Y) \subset U \cup C_{i_1} \cup \dots \cup C_{i_k}$ is relatively compact. By (i), $h(Y)$ is closed, therefore it is compact. □

Theorem 3.6. *Let X be a non-compact Riemann surface. Then there exists a sequence of compact subsets K_i, $i \in \mathbb{N}$ with the following properties:*

- *(i)* $K_i = h(K_i)$ *for all* $i \in \mathbb{N}$*;*
- *(ii)* $K_i \subset \mathring{K}_{i+1}$ *for all* $i \in \mathbb{N}$*, where* $\mathring{K}_{i+1}$ *denotes the interior of* K_{i+1}*;*
- *(iii)* $\bigcup_{i=0}^{\infty} K_i = X$.

Proof. Since X is Hausdorff, locally compact and second countable, it admits a countable basis of relatively compact sets, say $\{U_i\}_{i \in \mathbb{N}}$. Letting $K_i' = \bigcup_{j=0}^{i} \overline{U_j}$, we get a sequence of compact sets $K_0' \subset K_1' \subset \dots$ that cover X. Let $K_0 = h(K_0')$. By Theorem 3.5, K_0 is compact. Moreover, K_0 satisfies property (i) since $h(K_0) = h(h(K_0')) = h(K_0') = K_0$. We now construct the other compact sets K_i by induction. Suppose that we have already constructed compact sets $K_0, \dots, K_n$ satisfying properties (i) and (ii) and such that $K_i' \subset K_i$ for all $i = 0, \dots, n$. Now, let M be a relatively compact open neighbourhood of the compact set $K_n \cup K_{n+1}'$ and define $K_{n+1} = h(\overline{M})$. By Theorem 3.5 and Remark 3.3, K_{n+1} is compact and $h(K_n) = K_n$. Moreover, we have $K_n, K_{n+1}' \subset M \subset \mathring{K}_{n+1}$, which completes the inductive step. Property (iii) now follows since $K_i' \subset K_i$ and the sets K_i' already cover X. □

Theorem 3.7. *Let K_1, K_2 be compact subsets of a Riemann surface X such that $K_1 \subset \mathring{K}_2$ and $h(K_2) = K_2$. Then there exists an open subset Y of X that is Runge and satisfies $K_1 \subset Y \subset K_2$.*

Proof. Since $K_1 \subset \mathring{K}_2$, every point $x \in \partial K_2$ has a coordinate neighbourhood U_x not intersecting K_1. For every point $x \in \partial K_2$, we pick a compact disk $D_x \subset U_x$ containing x in its interior. By compactness, finitely many of these disks cover ∂K_2, say $D_1, \dots, D_k$. Let

$$Y := K_2 \setminus (D_1 \cup \dots \cup D_k) = \mathring{K}_2 \setminus (D_1 \cup \dots \cup D_k),$$

which is open and satisfies $K_1 \subset Y \subset K_2$. Let C_i, $i \in I$, denote the connected components of $X \setminus K_2$. Since $h(K_2) = K_2$, none of the C_i is relatively compact. But the disks D_j are connected and intersect some C_i, so every connected component of $X \setminus Y$ contains at least one of the C_i and hence cannot be relatively compact. In other words, Y is Runge. $\square$

Theorem 3.8. *Let Y be a Runge open subset of a Riemann surface X. Then every connected component of Y is Runge.*

Proof. Let Y_i, $i \in I$, denote the connected components of Y. Since Y is open and locally connected, all the components Y_i are open. Let $A = X \setminus Y$ and let A_k, $k \in K$, be the connected components of A. The open set Y being Runge, it follows that the components A_k are closed (in X) but not compact. Fix a connected component Y_ℓ of Y and let C be a connected component of $X \setminus Y_\ell$. Suppose for a moment that C intersects some A_k. Since C is maximal upon the connected subsets of $X \setminus Y_\ell$, we must have $A_k \subset C$. But A_k is not compact, so the same goes for C. Therefore, the only thing left to prove is that $C \cap A_k \neq \emptyset$ for some $k \in K$.

First, we claim that $\overline{Y_i}$ intersects A for every $i \in I$ (here $\overline{Y_i}$ denotes the closure of Y_i in X). Otherwise we would have the inclusion $\overline{Y_i} \subset Y$. But then by the maximality of Y_i we would have that $\overline{Y_i} = Y_i$ and hence that X is disconnected (unless $Y_i = X$, but that case is trivially true), a contradiction. Now, suppose that $C \cap A = \emptyset$. Then $C \subset Y$ and hence C intersects Y_i for some $i \neq \ell$. Thus $\overline{Y_i} \subset C$ since Y_i is connected and C is closed. But $\overline{Y_i}$ intersects A, and therefore so does C. $\square$

Theorem 3.9. *Let X be a non-compact Riemann surface. Then there exists a nested sequence of relatively compact Runge domains $Y_0 \Subset Y_1 \Subset Y_2 \Subset \dots$ such that $\bigcup_{i=0}^{\infty} Y_i = X$.*

Proof. We claim that the result follows if we can prove that for every compact set $K \subset X$, there exists a Runge domain Y such that $K \subset Y \Subset X$. Indeed, consider an exhaustion of X by compact subsets $K_0 \subset K_1 \subset \dots$ given by Theorem 3.6. We construct the sequence $Y_0 \Subset Y_1 \Subset \dots$ by induction. Start with a Runge domain Y_0 such that $K_0 \subset Y_0 \Subset X$. Now, suppose that relatively compact Runge domains $Y_0, \dots, Y_n$ have been constructed so that $Y_0 \Subset Y_1 \Subset \dots \Subset Y_n$ and $K_i \subset Y_i$ for all $i = 1, \dots, n$. Since $K_{n+1} \cup Y_n$ is relatively compact, there exists a relatively compact open neighbourhood Y_{n+1} of $K_{n+1} \cup \overline{Y_n}$. Clearly, $Y_n \Subset Y_{n+1}$ and $K_{n+1} \subset Y_{n+1}$. This completes the induction. To show that the domains Y_i cover X, it suffices to notice that the compact sets K_i themselves cover X and that $K_i \subset Y_i$ for all $i \in \mathbb{N}$.

Therefore, the only thing left to show is that every compact set K is contained in a relatively compact Runge domain Y. Let K be a compact set, and take a compact connected set K_1

that contains K. Now let K_2 be a compact set such that $K_1 \subset \mathring{K}_2$. Then $h(K_2)$ is compact by Theorem 3.5. By Theorem 3.7, there exists a Runge open Y' such that $K_1 \subset Y' \subset h(K_2)$. We finally let Y be the connected component of Y' that contains K_1. By Theorem 3.8, Y is still Runge, and it is obviously relatively compact. $\square$

3.2 Two vanishing theorems

It is a well-known fact that the cohomology group $H^1(X,\mathscr{O})$ of a compact Riemann surface X is finite-dimensional. As done in [4], the machinery (of functional analysis) used to prove this result can be adapted to say something about relatively compact open subsets of general Riemann surfaces.

Theorem 3.10 (cf. [4, Theorem 14.9])**.** *Let X be a Riemann surface and let Y be a relatively compact open subset of X. Then the natural restriction homomorphism*

$$H^1(X,\mathscr{O}) \to H^1(Y,\mathscr{O})$$

has finite-dimensional image.

There is also a quite similar theorem for the sheaf $\mathscr{F}$ of holomorphic sections of a holomorphic vector bundle:

Theorem 3.11 (cf. [4, Theorem 29.13])**.** *Let X be a Riemann surface, Y be a relatively compact open subset of X and E be a holomorphic vector bundle over X. Then $H^1(Y,\mathscr{F})$ is finite-dimensional.*

But as far as Theorem 3.10 is concerned, this is not the end of the story for non-compact Riemann surfaces. In fact, this section is devoted to prove a much stronger statement, namely that the cohomology group $H^1(X,\mathscr{O})$ vanishes when X is a non-compact Riemann surface.

The first step towards proving this result is to reduce the problem to one of functional analysis. To do this, first recall that if we have an exact sequence of sheaves

$$0 \longrightarrow \mathscr{K} \xrightarrow{\alpha} \mathscr{G} \xrightarrow{\beta} \mathscr{H} \longrightarrow 0,$$

where $H^1(X,\mathscr{G}) = 0$, then the long exact sequence in cohomology shows that

$$H^1(X,\mathscr{K}) = \mathscr{H}(X)/\beta\mathscr{G}(X).$$

The sheaves on X that will be of importance for us are the following:

(1) the sheaf $\mathscr{O}$ of holomorphic functions;

(2) the sheaf $\mathscr{E}$ of real-differentiable functions;

(3) the sheaf $\mathscr{E}^{0,1}$ of differentiable forms locally of the form $\frac{\partial}{\partial \bar{z}} d\bar{z}$;

(4) the sheaf $\mathscr{M}$ of meromorphic functions;

(5) the sheaf $\mathscr{F}$ of holomorphic sections of a holomorphic vector bundle.

In our case, we have the exact sequence of sheaves

$$0 \longrightarrow \mathscr{O} \longrightarrow \mathscr{E} \xrightarrow{\bar{\partial}} \mathscr{E}^{0,1} \longrightarrow 0,$$

where $\bar{\partial}$ is the usual Cauchy-Riemann operator. Moreover, a classical computation in Čech cohomology yields $H^1(X,\mathscr{E}) = 0$ (cf. [4, Theorem 12.6]). Therefore,

$$H^1(X,\mathscr{O}) = \mathscr{E}^{0,1}(X)/\,\bar{\partial}\,\mathscr{E}(X),$$

so now the problem of showing that $H^1(X,\mathscr{O}) = 0$ is reduced to the problem of showing that $\mathscr{E}^{0,1}(X) = \bar{\partial}\,\mathscr{E}(X)$. In other words, we want to show that for every (0,1)-form $\omega \in \mathscr{E}^{0,1}(X)$, there exists a smooth function $f \in \mathscr{E}(X)$ such that $\bar{\partial} f = \omega$. Note that this problem always has a solution locally: this is the celebrated Dolbeault lemma.

Lemma 3.12 (Dolbeault's lemma, cf. [4, Theorem 13.2])**.** *Let $U = \{z \in \mathbb{C} : |z| < R\}$, $0 < R \leq \infty$ be an open disk in the complex plane and let $g \in \mathscr{E}(U)$. Then there exists a function $f \in \mathscr{E}(U)$ such that*

$$\frac{\partial f}{\partial \bar{z}} = g.$$

We now spend some time to show that the problem has a solution on relatively compact open sets, provided that we restrict ourselves to a suitable subset of (0,1)-forms (see the statement of Theorem 3.17 for details).

Theorem 3.13. *Let X be a Riemann surface and Y be a relatively compact open set of X. Then, for all $a \in Y$, there exists a meromorphic function $f \in \mathscr{M}(Y)$ that has a pole at a and is holomorphic on $Y \setminus \{a\}$.*

Proof. By Theorem 3.10, we know that

$$k = \dim \operatorname{Im}(H^1(X,\mathscr{O}) \to H^1(Y,\mathscr{O})) < \infty.$$

Let $a \in Y$ and consider a coordinate neighbourhood (U,z) centered at a. Let $V = X \setminus \{a\}$, so that we have an open covering $\mathcal{U} = \{U,V\}$ of X. On the intersection $U \cap V$, the functions z^{-i} are holomorphic for all $i \geq 1$, hence they define cocycles $\zeta_i \in Z^1(\mathcal{U},\mathscr{O})$. Since $\dim \operatorname{Im}(H^1(\mathcal{U},\mathscr{O}) \to H^1(\mathcal{U} \cap Y,\mathscr{O})) < k+1$, it follows that the restricted cocycles $\zeta_1,\ldots,\zeta_{k+1} \in Z^1(\mathcal{U} \cap Y,\mathscr{O})$ are linearly dependent modulo coboundaries, i.e. there exist constants $c_i \in \mathbb{C}$ not all zero and a cochain $\eta = (f,g) \in C^0(\mathcal{U}\cap Y,\mathscr{O})$ such that $\sum_{i=1}^{k+1} c_i\zeta_i = \delta\eta$, which implies that

$$\sum_{i=1}^{k+1} c_i z^{-i} = g - f \text{ on } U \cap V \cap Y.$$

Therefore, the holomorphic functions $f + \sum_{i=1}^{k+1} c_i z^{-i}$, defined on $(U\cap Y)\setminus\{a\}$, and g, defined on $V \cap Y = Y \setminus \{a\}$, glue together to give a meromorphic function h over Y that has a pole at a and is holomorphic everywhere else. □

Theorem 3.14. *Let X be a non-compact Riemann surface and Y be a relatively compact open set of X. Then there exists a holomorphic function $f \in \mathscr{O}(Y)$ that is not constant on any connected component of Y.*

Proof. Let Y' be a relatively compact domain that contains $\overline{Y}$. Since X is non-compact and connected, there exists a point $a \in Y' \setminus Y$. Applying Theorem 3.13 to Y', we get a meromorphic function $f \in \mathscr{M}(Y')$ with a single pole at a. Since $a \notin Y$, the function f restricts to a holomorphic function on Y. Finally, it cannot be constant on a connected component of Y, or otherwise it would be constant on all of Y' by the identity theorem. □

Theorem 3.15. *Let X be a non-compact Riemann surface and Y be a relatively compact open subset of X. Then $Im(H^1(X,\mathscr{O}) \to H^1(Y,\mathscr{O})) = 0$.*

Proof. Let Y' be a relatively compact open neighbourhood of $\overline{Y}$. Applying Theorem 3.10 to Y and Y', we get that

$$\dim \operatorname{Im}(H^1(Y',\mathscr{O}) \to H^1(Y,\mathscr{O})) < \infty.$$

Choose cohomology classes $\xi_1, \dots ,\xi_n \in H^1(Y',\mathscr{O})$ whose restriction to Y span the vector space $\operatorname{Im}(H^1(Y',\mathscr{O}) \to H^1(Y,\mathscr{O}))$. By Theorem 3.14, there exists a holomorphic function $f \in H^1(Y',\mathscr{O})$ that is not constant on any connected component of Y'. Since $H^1(Y',\mathscr{O})$ is an $\mathscr{O}(Y')$-module, it follows from the choice of the classes $\xi_1, \dots ,\xi_n$ that there exist constants $c_{ij} \in \mathbb{C}$ such that

$$f\xi_i = \sum_{j=1}^{n} c_{ij}\xi_j \text{ on } Y. \tag{3.16}$$

Consider the function $F = \det(fI - C)$, where I denotes the identity matrix and $C = (c_{ij})$. Apart from being holomorphic, F is also non-zero on any connected component of Y'. Indeed, if F were identically zero on some connected component of Y', then f, being a root of the characteristic polynomial of C, would be locally constant, hence constant on said connected component. Moreover, Equation (3.16) can be rewritten as $(fI-C)\xi|_Y = 0$, where ξ denotes the column vector $(\xi_1, \dots ,\xi_n)^T$. Multiplying by the adjugate matrix of $fI-C$ on both sides, it follows that $F\xi_i|_Y = 0$ for all $i = 1, \dots ,n$.

Now, let $\zeta \in H^1(Y',\mathscr{O})$ be an arbitrary cohomology class. Then ζ is represented by some cocycle $(f_{ij}) \in Z^1(\mathcal{U},\mathscr{O})$. Since F is not identically zero, we can further assume (possibly after passing to a refinement of $\mathcal{U}$) that each zero of F is contained in at most one member of the open covering $\mathcal{U} = (U_i)$. It follows that $F|_{U_i \cap U_j} \in \mathscr{O}^*(U_i \cap U_j)$. Thus there exists a cocycle $(g_{ij}) \in Z^1(\mathcal{U},\mathscr{O})$ such that $(f_{ij}) = F(g_{ij})$. Letting η denote the cohomology class of (g_{ij}), we get $\zeta = F\eta$, which further implies that $\zeta|_Y = F\eta|_Y = 0$. The result now follows since the morphism $H^1(X,\mathscr{O}) \to H^1(Y,\mathscr{O})$ is the composition of the two morphisms

$$H^1(X,\mathscr{O}) \to H^1(Y',\mathscr{O}) \to H^1(Y,\mathscr{O}).$$

□

Theorem 3.17. *Let X be a non-compact Riemann surface and Y be a relatively compact open subset of X. Then for every differential form $\omega \in \mathscr{E}^{0,1}(X)$, there exists a smooth function $f \in \mathscr{E}(Y)$ such that $\bar{\partial} f = \omega|_Y$.*

Proof. Let ω be a (0,1)-form on X. Recall from the Dolbeault lemma (Lemma 3.12) that the problem has a solution locally. Therefore, there exists an open covering $\mathcal{U} = (U_i)$ of X and functions $f_i \in \mathscr{E}(U_i)$ such that $\bar{\partial} f_i = \omega|_{U_i}$. The functions $f_j - f_i$ are holomorphic because $\bar{\partial}(f_j - f_i) = \omega|_{U_i} - \omega|_{U_i} = 0$, thus they define a cocycle $(f_j - f_i) \in Z^1(\mathcal{U},\mathscr{O})$. By Theorem 3.15, the restriction of that cocycle to Y is zero, i.e. there exists a cochain $(g_i) \in C^0(\mathcal{U},\mathscr{O})$ such that $f_j - f_i = g_j - g_i$ on $U_i \cap U_j \cap Y$. But since $\mathscr{E}$ is a sheaf, there exists a function $f \in \mathscr{E}(Y)$ such that $f|_{U_i \cap Y} = f_i - g_i$. This smooth function f is the desired solution because on $U_i \cap Y$, we have

$$\bar{\partial} f = \bar{\partial}(f_i - g_i) = \bar{\partial} f_i = \omega|_Y.$$

□

We are finally ready to attack the global problem.

Theorem 3.18. *Let X be a non-compact Riemann surface. Then $\mathscr{E}^{0,1}(X) = \bar{\partial}\mathscr{E}(X)$, i.e. for every differential form $\omega \in \mathscr{E}^{0,1}(X)$, there exists a smooth function $f \in \mathscr{E}(X)$ such that $\bar{\partial} f = \omega$.*

Proof. Let $\omega \in \mathscr{E}^{0,1}(X)$ be a global differential (0,1)-form. We know from Theorem 3.17 that for every relatively compact open subset Y of X, there exists a smooth function $g \in \mathscr{E}(Y)$ such that $\bar{\partial} g = \omega|_Y$. We now use an exhaustion process to build a global solution to the equation $\bar{\partial} f = \omega$. Let $Y_0 \Subset Y_1 \Subset Y_2 \Subset \dots$ be an exhaustion of X by relatively compact Runge domains. We know that for every $n \in \mathbb{N}$, there exists a smooth function $g_n \in \mathscr{E}(Y_n)$ such that $\bar{\partial} g_n = \omega|_{Y_n}$. Using Runge's approximation theorem, we modify the smooth functions g_n to get a sequence of smooth functions f_n such that

(i) $\bar{\partial} f_n = \omega|_{Y_n}$;

(ii) $\|f_{n+1} - f_n\|_{Y_{n-1}} \leq 2^{-n}$.

We proceed by induction. First, let $f_0 \coloneqq g_0$. Suppose now that the functions $f_0, \ldots, f_n$ have been constructed. Since f_k and g_{k+1} are both solutions to $\bar{\partial} f = \omega$ on Y_k, it follows that $\bar{\partial}(f_n - g_{n+1}) = 0$ on Y_n, i.e. $f_n - g_{n+1}$ is holomorphic on Y_n. Thus by Runge's approximation theorem, there exists a holomorphic function $h \in \mathscr{O}(X)$ such that $\|(f_n - g_{n+1}) - h\| \leq 2^{-n}$ on the compact subset $\overline{Y}_{n-1}$. Letting $f_{n+1} \coloneqq g_{n+1} + h$, we see that $\bar{\partial} f_{n+1} = \bar{\partial} g_{n+1} = \omega|_{Y_{n+1}}$ and $\|f_{n+1} - f_n\|_{Y_{n-1}} \leq 2^{-n}$, which completes the induction.

Since the Runge domains Y_n exhaust X, every point $x \in X$ is contained in almost all Y_n, so it makes sense to define $f(x) \coloneqq \lim_{n\to\infty} f_n(x)$. On Y_n, we can write

$$f = f_{n+1} + \sum_{k=n+1}^{\infty} (f_{k+1} - f_k).$$

The functions $f_{k+1} - f_k$ are holomorphic on Y_n for all $k \geq n + 1$ and the series $S = \sum_{k=n+1}^{\infty}(f_{k+1} - f_k)$ converges uniformly on Y_n, hence the series S is holomorphic on Y_n. Therefore, $f = f_{n+1} + S$ is a smooth function on Y_n and

$$\bar{\partial} f = \bar{\partial}(f_{n+1} + S) = \bar{\partial} f_{n+1} = \omega \text{ on } Y_n.$$

But this is true for all $n \in \mathbb{N}$, so it follows that f is a smooth function on all of X and $\bar{\partial} f = \omega$ on X as well. □

Corollary 3.19. *Let X be a non-compact Riemann surface. Then*

$$H^1(X,\mathscr{O}) = 0.$$

3.3 Triviality of holomorphic vector bundles

With the results of the last sections in hand, we are now ready to tackle holomorphic vector bundles over non-compact Riemann surfaces. The first item on our agenda is to generalize Theorem 3.13 to this setting, namely to show that holomorphic vector bundles (over non-compact Riemann surfaces) admit non-vanishing meromorphic sections over relatively compact open subsets.

Theorem 3.20. *Let $Y \Subset X$ be a relatively compact open subset of a non-compact Riemann surface X and $E \to X$ be a holomorphic vector bundle. Then, for any $a \in Y$, there exists a meromorphic section of E over Y which has a pole at a and is holomorphic on $Y \setminus \{a\}$.*

Proof. Denote by n the rank of E and by k the dimension of $H^1(Y,\mathscr{F})$ (which is finite from Theorem 3.11). Let $a \in Y$ and $h : E_{U_1} \to U_1 \times \mathbb{C}^n$ be a local trivialization over a chart $U_1 \subset Y$ centered at a. Let $U_2 \coloneqq Y \setminus \{a\}$. Then $\mathfrak{U} = (U_1,U_2)$ is an open cover of Y. Recall that a holomorphic section s over $U_1 \setminus \{a\}$ can be represented by an n-tuple of holomorphic functions $(s_1, \ldots, s_n) \in \mathscr{O}(U_1 \setminus \{a\})^n$. Using this representation, we can now repeat the argument of Theorem 3.13 (with the cocycles $\zeta_i = (z^{-i}, \ldots, z^{-i}) \in Z^1(\mathfrak{U}, \mathscr{F})$) to get a meromorphic section f over Y that has a pole at a and is holomorphic everywhere else. $\square$

Before moving on to the next theorem, we first recall two classical theorems on the complex plane.

Theorem 3.21 (Weierstrass' factorization theorem)**.** *Let $A \subset \mathbb{C}$ be a closed discrete subset. Then there exists an entire function with zeros of chosen multiplicity at every point of A, and only there.*

Theorem 3.22 (Mittag-Leffler's theorem)**.** *Let $A \subset \mathbb{C}$ be a closed discrete subset. For every point $a \in A$, let $p_a(z)$ be a polynomial in $\frac{1}{z-a}$. Then there exists a meromorphic function $f \in \mathscr{M}(\mathbb{C})$ such that for every $a \in A$, the function $f(z) - p_a(z)$ has only a removable singularity at a. In particular, the principal part of f at a is exactly $p_a(z)$.*

Simply put, these two theorems respectively assert the existence of a function with prescribed zeros (of prescribed multiplicities) or with prescribed poles (of prescribed order). As it turns out, there is a common generalization of both these theorems to the setting of non-compact Riemann surfaces. To state it in a natural way, we recall the notion of a divisor on a Riemann surface.

Definition 3.23. Let X be a Riemann surface. A divisor D on X is a function $D : X \to \mathbb{Z}$ such that for every compact subset $K \subset X$, the set $\{x \in K : D(x) \neq 0\}$ is finite.

To any nonzero meromorphic function $f \in \mathscr{M}^*(X)$ we can associate the divisor (f) which sends a point $x \in X$ to the order of f at x. Divisors arising this way are called principal divisors. In this language, Weierstrass' theorem and Mittag-Leffler's theorem asserts that on the complex plane, certain divisors are in fact principal divisors. This is no coincidence: as it turns out, every divisor on a non-compact Riemann is principal.

Theorem 3.24 (cf. [4, Theorem 26.5])**.** *Let D be a divisor on a non-compact Riemann surface X. Then there exists a meromorphic function $f \in \mathscr{M}^*(X)$ such that $D = (f)$.*

Given a meromorphic section on a holomorphic vector bundle, Theorem 3.24 enables us to construct a holomorphic section out of it.

Theorem 3.25. *Let $E \to X$ be a holomorphic vector bundle of rank n over a non-compact Riemann surface X. If E has a non-trivial global meromorphic section, then it also has a global holomorphic section with no zeros.*

Proof. Let f be a non-zero global meromorphic section of E and let A be the closed discrete set on which f has poles. Consider a point $a \in A$. Let $h_a : E_U \to U \times \mathbb{C}^n$ be a local trivialization of E over a and let $(f_1, \dots, f_n) \in \mathscr{M}(U)^n$ be the representation of f with respect to h_a. Define $D(a)$ to be the minimum of the order of the functions f_k at a. Doing this for every $a \in A$, we get a divisor D on X. By Theorem 3.24, there exists a meromorphic function $\varphi \in \mathscr{M}^*(X)$ such that $(\varphi) = -D$. It follows that φf is a global holomorphic section of E that has no zeros. □

Remark 3.26. In particular, Theorem 3.25 can be used to construct a nowhere vanishing holomorphic section from a non-trivial holomorphic section of E over X.

We are finally ready to prove that holomorphic vector bundles over non-compact Riemann surfaces are trivial. We begin with line bundles. We give two proofs: one based on Runge's approximation theorem (cf. [4, Theorem 30.3]) and a more succint one based on sheaf cohomology.

Theorem 3.27. *Every holomorphic line bundle E over a non-compact Riemann surface X is trivial.*

First proof. Recall from Theorem 2.14 that $E \to X$ is trivial if and only if we can find a global holomorphic section over X that doesn't vanish anywhere. We will thus put all our efforts towards constructing such a section. By Theorem 3.9, let $\emptyset \neq Y_0 \Subset Y_1 \Subset Y_2 \Subset \dots$ be an exhaustion of X by relatively compact Runge domains. Over any of the Runge domains Y_k there exists a (non-trivial) meromorphic section by Theorem 3.20 and thus a non-vanishing holomorphic section by Theorem 3.25. It follows from Theorem 2.14 that E is trivial over each Y_k. In particular, our discussion on sections and their representations implies that $\mathscr{F}(Y_k) \cong \mathscr{O}(Y_k)$, i.e. there is a one-to-one correspondence between holomorphic sections of E over Y_k and their representations, which are holomorphic functions on Y_k in this case.

Using this identification, Runge's approximation theorem (Theorem 3.1) tells us that every holomorphic section of E over Y_k can be approximated uniformly by a holomorphic section over Y_{k+1}. Let $f_0 \in \mathscr{F}(Y_0)$ be a holomorphic section with no zeros and fix a point $a \in Y_0$. Let $0 < \varepsilon < \frac{|f_0(a)|}{2}$ and let $K \subset Y_0$ be a compact set containing a. Using induction and Runge's approximation theorem, we can construct a sequence of holomorphic sections $(f_k)_{k \geq 1}$, where $f_k \in \mathscr{F}(Y_k)$, such that $\|f_1 - f_0\|_K < \frac{\varepsilon}{2}$ and $\|f_k - f_{k-1}\|_{\overline{Y}_{k-2}} < \frac{\varepsilon}{2^k}$ for all $k \geq 2$. Since the sequences $(f_l|_{Y_k})_{l>k}$ converge uniformly in $\mathscr{F}(Y_k)$ for every k and the Runge domains Y_k exhaust X, the limit of the sequence (f_k) gives a global holomorphic section $f \in \mathscr{F}(X)$. Moreover, this section does not vanish identically because

$$\begin{aligned}
|f_0(a)| &\leq |f_1(a) - f_0(a)| + \cdots + |f_k(a) - f_{k-1}(a)| + |f_k(a)| \\
&\leq \|f_1 - f_0\|_K + \cdots + \|f_k - f_{k-1}\|_{\overline{Y}_{k-2}} + |f_k(a)| \\
&< \frac{\varepsilon}{2} + \cdots + \frac{\varepsilon}{2^k} + |f_k(a)| \\
&< \varepsilon + |f_k(a)|
\end{aligned}$$

which implies that

$$|f(a)| = \lim_{k\to\infty} \geq |f_0(a)| - \varepsilon > \frac{|f_0(a)|}{2} > 0.$$

Therefore, one last application of Theorem 3.25 gives the existence of a global holomorphic section that doesn't vanish anywhere. $\square$

Second proof. Let $\underline{\mathbb{Z}}$ denote the constant sheaf on X associated to $\mathbb{Z}$ and let $\mathscr{O}^*$ denote the sheaf of non-vanishing holomorphic functions on X. These two are related to the sheaf $\mathscr{O}$ through the exponential sequence

$$0 \longrightarrow \underline{\mathbb{Z}} \xrightarrow{\alpha} \mathscr{O} \xrightarrow{\beta} \mathscr{O}^* \longrightarrow 0,$$

where α is the inclusion map of $\underline{\mathbb{Z}}$ into $\mathscr{O}$ and β is given by $\exp(2\pi i f)$. The long exact sequence in cohomology gives the exact sequence

$$H^1(X,\mathscr{O}) \longrightarrow H^1(X,\mathscr{O}^*) \longrightarrow H^2(X,\mathbb{Z}).$$

Since $H^1(X,\mathscr{O}) = 0$ (cf. Corollary 3.19) and $H^2(X,\mathbb{Z}) = 0$ (because X is non-compact), it follows that $H^1(X,\mathscr{O}^*) = 0$. But this cohomology group is exactly the Picard group of X. The result follows. $\square$

We can finally extend Theorem 3.27 to holomorphic vector bundles of higher rank.

Theorem 3.28. *Every holomorphic vector bundle E over a non-compact Riemann surface X is trivial.*

Proof. The theorem will be proved using induction on the rank n of the holomorphic vector bundle E. The case $n = 1$ is exactly Theorem 3.27. Let's now assume that every holomorphic vector bundle of rank $n-1$ is trivial, and let E be a holomorphic vector bundle of rank n. The proof essentially goes in two steps.

Step 1: Suppose for a moment that there exists a nowhere vanishing global holomorphic section $f_n \in \mathscr{F}(X)$. Since E is locally trivial, then by Theorem 2.14 there is an open cover $\mathfrak{U} = (U_i)_{i\in I}$ of X and holomorphic sections $f_1^i,\dots,f_{n-1}^i \in \mathscr{F}(U_i)$ such that $f_1^i(x),\dots,f_{n-1}^i(x),f_n(x)$ are linearly independent at every point $x \in U_i$. For every $i \in I$, let f^i denote the column vector with entries $f_1^i,\dots,f_{n-1}^i$. The goal now is to modify the local "frames" of $n-1$ sections f^i so that they agree on the intersections $U_i \cap U_j$. The first thing to observe is that on $U_i \cap U_j$, the local frames f^i and f^j are related by

$$\begin{pmatrix} f^i \\ f_n \end{pmatrix} = \begin{pmatrix} g^{ij} & a^{ij} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} f^j \\ f_n \end{pmatrix}, \tag{3.29}$$

where $g^{ij} \in \mathrm{GL}(n-1,\mathscr{O}(U_i \cap U_j))$ and a^{ij} is a column vector whose entries are elements of $\mathscr{O}(U_i \cap U_j)$. Since Equation (3.29) also holds on triple intersections $U_i \cap U_j \cap U_k$, a direct calculation shows that $g^{ik} = g^{ij}g^{jk}$ and $a^{ik} = g^{ij}a^{jk} + a^{ij}$. Let F be the subbundle of E locally spanned by the local frames f^i. Then F is a holomorphic vector bundle of rank $n-1$ and (g^{ij}) is its associated cocycle of transition functions. By our induction hypothesis, F is trivial. It follows from Theorem 2.14 that (g^{ij}) is a coboundary, i.e. there exists a family of matrices $(g^i) \in \prod_{i\in I} \mathrm{GL}(n-1,\mathscr{O}(U_i))$ such that $g^{ij} = g^i(g^j)^{-1}$ on $U_i \cap U_j$.

We are now ready for our first modification of the local frames f^i. Let $\tilde{f}^i := (g^i)^{-1} f^i$. Then it follows from Equation (3.29) that

$$\begin{pmatrix} \tilde{f}^i \\ f_n \end{pmatrix} = \begin{pmatrix} I & b^{ij} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \tilde{f}^j \\ f_n \end{pmatrix}, \tag{3.30}$$

where b^{ij} is the column vector $(g^i)^{-1}a^{ij} \in \mathscr{O}(U_i \cap U_j)^{n-1}$. On triple intersections $U_i \cap U_j \cap U_k$, another calculation shows that $b^{ik} = b^{ij} + b^{jk}$. In other words, each entry of b^{ij} is a cocycle in $Z^1(\mathfrak{U},\mathscr{O})$. But $H^1(X,\mathscr{O}) = 0$ by Corollary 3.19. Thus we can find holomorphic column vectors $b^i \in \mathscr{O}(U_i)^{n-1}$ such that $b^{ij} = b^i - b^j$ on $U_i \cap U_j$. We now apply a second modification to the local frames f^i: let $\hat{f}^i := \tilde{f}^i - b^i f_n$. Equation (3.30) now implies that

$$\begin{pmatrix} \hat{f}^i \\ f_n \end{pmatrix} = \begin{pmatrix} \hat{f}^j \\ f_n \end{pmatrix} \text{ on } U_i \cap U_j.$$

Therefore, the modified local frames $\hat{f}^i$ glue together to give an $(n-1)$-tuple $(f_1,\ldots,f_{n-1})$ of global holomorphic sections. By construction, the sections $f_1,\ldots,f_{n-1},f_n$ are linearly independent on each fiber. The holomorphic vector bundle E is thus trivial.

Step 2: It only remains to prove the existence of a nowhere vanishing global holomorphic section of E. By Theorem 3.20 and Theorem 3.25 there always exists such a section over a relatively compact open set $Y \Subset X$. Appealing to the argument of Step 1, we conclude that E is trivial over Y. We can then use an exhaustion of X by Runge domains combined with Runge's approximation theorem to construct a non-trivial holomorphic section of E over X, just as we did in the proof of Theorem 3.27. An application of Theorem 3.25 finally yields the desired section. □

Remark 3.31. The induction step of Theorem 3.28 can be synthesized using the language of extensions. Once we know that E has a nowhere vanishing global holomorphic section, we deduce that it has the form of an extension

$$0 \longrightarrow \mathscr{O} \longrightarrow E \longrightarrow E' \longrightarrow 0,$$

where E' is some holomorphic bundle of rank $k-1$. By induction, E' is trivial and the extension becomes

$$0 \longrightarrow \mathscr{O} \longrightarrow E \longrightarrow \mathscr{O}^{\oplus(k-1)} \longrightarrow 0. \tag{3.32}$$

But the space of extensions of E' is classified by $\mathrm{Ext}^1(E',\mathscr{O})$, and the latter is isomorphic to

$$\mathrm{Ext}^1(E',\mathscr{O}) \simeq H^1(X, \mathrm{Hom}(E',\mathscr{O})) \simeq H^1(X,\mathscr{O}^{\oplus(k-1)}) = 0.$$

In other words, the extension (3.32) splits and $E \simeq \mathscr{O}^{\oplus k}$ is trivial.

Chapter 4

Compact Riemann Surfaces

Looking back at the previous chapter, we notice that most of the tools used relied in some way or the other on the non-compactness of the Riemann surface. This tells us that completely different ideas are going to be needed for the compact case. Roughly speaking, we will study a universal space of holomorphic vector bundles over a fixed compact Riemann surface X, i.e. the space parametrizing the holomorphic vector bundles over X. Constantly using inspiration from Morse theory, we will then exhibit a stratification on this space and see that it closely resembles the ones usually encountered in Morse theory. In the end, this will enable us to get some results concerning the equivariant cohomology of the minimal stratum. All the results here are taken from [1], and the reader wishing to know more is kindly invited to read the article for a deeper dive in the subject.

4.1 Review of Morse theory

Since a lot of our intuition will rely on ideas from Morse theory, we devote this section to skim through the basics of the theory. More details can be found in [1, §1].

4.1.1 Elementary Morse theory

Let M be a compact differentiable manifold and let $f : M \to \mathbb{R}$ be a smooth function. A critical point of f is a point at which df vanishes and a critical value of f is a value of f whose preimage contains a critical point. At a critical point p, the Hessian $H_p f$ is a well-defined quadratic form on T_pM. If we have local coordinates x^i centered at p, then the matrix of the quadratic form H_pf is given by

$$H_pf = \left(\frac{\partial^2 f}{\partial x^i \partial x^j}\right).$$

A critical point p is called non-degenerate if $\det H_pf \neq 0$. The smooth function f is called a Morse function if all its critical points are non-degenerate. At a non-degenerate critical point p, the number of negative eigenvalues in a diagonalization of H_pf is called the index of p and is denoted by $\lambda_p(f)$. The Morse lemma asserts that for a non-degenerate critical point p of index λ_p, there exists local coordinates $x^1, \ldots, x^n$ centered at p such that

$$f = f(p) - (x^1)^2 - \cdots - (x^{\lambda_p})^2 + (x^{\lambda_p+1})^2 + \cdots + (x^n)^2.$$

A Morse function f can be used to understand the topology of the underlying manifold M via the fundamental structural theorem of non-degenerate Morse theory. Letting M_a ($a \in \mathbb{R}$) denote the set of points $\{m \in M : f(m) \leq a\}$, this theorem essentially says that $M_a \sim M_b$ if there are no critical values between a and b, while $M_b \sim M_a \cup e_\lambda$ if there is a single critical point of index λ in $M_b \setminus M_a$. Here, the relation $\sim$ is that of homeomorphism and $M_a \cup e_\lambda$ is the manifold obtained from M_a by attaching a cell e_λ of dimension λ to it.

To any Morse function f we associate its Morse polynomial

$$M_t(f) \coloneqq \sum_p t^{\lambda_p(f)},$$

where the sum runs over the necessarily finite number of critical points p of f. There are topological bounds for $M_t(f)$. Indeed, if we let

$$P_t(M; K) = \sum t^i \dim H^i(M; K)$$

be the Poincaré polynomial of M relative to a coefficient field K, then a Morse function f always satisfies

$$M_t(f) - P_t(M; K) = (1+t)R(t), \tag{4.1}$$

where $R(t)$ is a polynomial with non-negative coefficients. In particular, the coefficients of $M_t(f)$ dominate those of $P_t(M)$, hence the name Morse inequalities for Equation (4.1). The function f is called a K-perfect Morse functions if the Morse inequalities are in fact equalities, i.e. $M_t(f) = P_t(M; K)$, and is called perfect if the latter equality holds for all fields K.

There are criteria for establishing perfection. First, there is the lacunary principle: if $\{\lambda_p(f)\}$ contains no consecutive integers, then f is perfect. There is also the completion principle: if all the critical points of f are completable, then f is perfect. Let's describe what we mean for a critical point p to be completable. Near p, f is of the form

$$f = f(p) - (x^1)^2 - \cdots - (x^k)^2 + (x^{k+1})^2 + \cdots + (x^n)^2,$$

where k is the index of p. The set

$$v_p^- = \{x : (x^1)^2 + \cdots + (x^k)^2 \leq \varepsilon, x^{k+1} = \cdots = x^n = 0\}$$

is a disc near p whose boundary ∂v_p^- is a $(k-1)$-sphere in the space $M_{c-\varepsilon} = \{m \in M : f(m) \leq c - \varepsilon\}$. The point p is called completable if the sphere ∂v_p^- bounds a singular chain in $M_{c-\varepsilon}$.

The Morse inequalities as well as the completion principle are both consequences of the fundamental structural theorem (more precisely the exact sequences relating the cohomology of M_b and M_a), while the lacunary principle is a consequence of the completion principle.

4.1.2 Extension of Morse theory

We now move on beyond the realm of elementary Morse theory and adapt our definitions so that they become better suited for more general purposes.

Let $N \subset M$ be a (regular) submanifold. Equip the tangent bundle of M with a Riemannian metric. This allows to speak of the normal bundle $v(N)$ of N (where v stands for vertical). Then N is said to be a non-degenerate critical manifold for f if

$$\begin{gathered} df \equiv 0 \text{ along } N \text{ and} \\ H_N f \text{ is non-degenerate on the normal bunde } v(N) \text{ of } N. \end{gathered}$$

Here again the fact that $df \equiv 0$ along N ensures that the Hessian of f is well-defined on $v(N)$. In this extended context, a function f on X is said to be a Morse function if its critical set is a union of non-degenerate critical manifolds.

Since critical sets may now contain submanifolds of different dimensions, we need to adapt our way of counting these critical submanifolds in order to form the Morse polynomial. To do so, begin by equipping M with a Riemannian metric, so that $v(N)$ also inherits such a metric. Then the Hessian $H_N f$ defines a canonical self-adjoint endomorphism

$$A_N : v(N) \to v(N) \text{ via the formula } (A_N x, y) = H_N f(x,y).$$

Since $H_N f$ is non-degenerate, the eigenvalues of A_N are all non-zero and hence decomposes $v(N)$ into an orthogonal direct sum $v(N) = v^+(N) \oplus v^-(N)$ spanned by the positive and negative eigenvalues of A_N respectively. The index λ_N of N is then simply the fibre dimension of $v^-(N)$. Choosing a coefficient field K, we count a non-degenrate critical manifold N of f with the polynomial $M_t(f,N) = \sum t^i \dim H_c^i(v^-(N))$, where H_c^i denotes the compactly supported cohomology. We can now define the Morse polynomial of a Morse function f by the natural formula

$$M_t(f) = \sum_N M_t(f,N).$$

A convincing reason to adopt this way of counting critical manifolds is that the Morse inequalities persist, enabling us to talk about K-perfect Morse functions in this extended sense as well. Another nice property is the functorial nature of this approach under pullbacks.

More precisely, if $E \to M$ is a smooth fibration, then f is non-degenerate on M iff $\pi^* f$ is non-degenerate on E, and in that case the index of N equals the index of $\pi^{-1}(N)$.

The last thing to cover is the completion principle. The picture used earlier naturally generalizes to our actual setting. Put into the form of a diagram, we have

$$\begin{tikzcd}
H_*(v_\varepsilon^-(N)) \arrow[r] & \tilde{H}_*(v_\varepsilon^-(N),\partial v_\varepsilon^-(N)) \arrow[r, "\partial"] & \tilde{H}(\partial v_\varepsilon^-(N)) \arrow[d] \\
& H_{*-\lambda_N}(N) \arrow[u, "\pi^{-1}"] \arrow[r, dashed] & \tilde{H}(M_{f(c)-\varepsilon})
\end{tikzcd}$$

where $v_\varepsilon^-(N)$ denotes the set in the exponential image of $v^-(N)$ for $f \geq c - \varepsilon$. The latter is a λ_N-disk-bundle over N if $\varepsilon > 0$ is small enough, so that π^{-1} is the so-called Thom isomorphism. We then say that N is K-completable if the dashed arrow in the diagram is zero. This definition reduces to the previous one for a non-degenerate critical point and we still have the completion principle.

Theorem 4.2. *If all the critical manifolds of f are K-completable then f is a K-perfect Morse function on M.*

Now, the interesting fact is that in this setting, a critical manifold can be self-completing. Indeed, since the top row in the diagram is exact, a class $\alpha \in H_{*-\lambda_N}(N)$ will certainly go to zero under the dashed arrow if $\pi^{-1}(\alpha)$ is in the image of $H_*(v_\varepsilon^-(N))$. This phenomenon can only occur if the bundle $v^-(N)$ is non-trivial over N. Later, we will work in an infinite-dimensional setting, and we will be so lucky that all critical manifolds will be self-completing, a feature that is impossible to realize in compact finite-dimensional settings.

4.1.3 Equivariant cohomology

When a space X is acted on by a Lie Group G, we would like to compute the cohomology of X/G, but that space often lacks nice topological properties as soon as the action is not free. A remedy to that problem is to replace X by a larger but homotopy equivalent space on which the action will be free. One way to proceed is to consider the universal bundle $EG \to BG$ for G. Since EG admits a free G-action, we can equip the space $EG \times X$ with the diagonal action $(e,x)\cdot g = (e\cdot g, g^{-1}\cdot x)$, which is also free. The product space $EG \times X$ is then homotopy equivalent to X since EG is contractible, and its orbit space $X_G \coloneqq (EG \times X)/G$ under the free G-action is then a nice topological space, called the homotopy quotient of X. This construction allows us to define to the equivariant cohomology $H_G(X)$ of X as the cohomology of its homotopy quotient $H(X_G)$.

4.1.4 Stratifications

For now we focussed on the homological aspects of Morse theory, but these actually come from a more detailed picture offered by the Morse function. Consider the vector field $\operatorname{grad} f$ which is dual to the differential df. The gradient flow of f is given by the paths of steepest descent, that is the trajectories of $-\operatorname{grad} f$. In the classical theory, if every critical point is non-degenerate, then every trajectory converges to some critical point. Taking all trajectories converging to a given critical point p, we get a cell $M^+(p)$ of M called the stable manifold of p. There is also the analogous unstable manifold $M^-(p)$ obtained using $-f$ instead of f. It is now clear (or at least plausible) from the Morse Lemma that the dimension of $M^-(p)$ (which is the codimension of $M^+(p)$) is equal to the Morse index of p. Thus, f gives a cell decomposition $M = \bigcup_p M^+(p)$ and the Morse inequalities follow at once using these cells to compute the homology of M.

More generally, in our extended Morse theory, given a non-degenerate critical manifold N we have a stable manifold $M^+(N)$ that is a cell-bundle over N, and if all critical manifolds are non-degenerate we get a corresponding stratification $M = \bigcup_N M^+(N)$, called the Morse stratification. For the equivariant case, given a G-invariant f, we can always pick a G-invariant metric and the gradient flow is then G-invariant so that the Morse stratification is also G-invariant.

There is a natural preorder $\prec$ on the critical manifolds N of a Morse function f given by $N_1 \prec N_2$ if the boundary of $M^+(N_1)$ intersects $M^+(N_2)$. In particular, if $N_1 \prec N_2$, then $f(N_1) < f(N_2)$. Taking the transitive relation $<$ associated to $\prec$, we get a partial order with the property that

$$\overline{M^+(N)} \subset \bigcup_{N' \succeq N} M^+(N') \subset \bigcup_{N' \geq N} M^+(N').$$

Using the Morse stratification, we can get valuable information about the homology of M. The way we will proceed actually doesn't care that the stratification comes from a Morse function, only that it satisfies the above property along with some other minor restrictions.

Start with an explicit finite stratification of M

$$M = \bigcup_\lambda M_\lambda,$$

where each M_λ is a locally closed submanifold of M and the index set of λ is partially ordered so that

$$\overline{M_\lambda} \subset \bigcup_{\mu \geq \lambda} M_\mu \tag{4.3}$$

holds for all λ. To compute the homology of M we can start with the open strata (given by minimal λ) and add inductively the other strata using the exact cohomology sequence for a pair $(U, U \setminus V)$ where V is a closed submanifold of U. Here are the details. We say that a subset I of indices is open if $\lambda \in I$ and $\mu \leq \lambda$ implies that $\mu \in I$ and closed if $\lambda \in I$ and

$\mu \geq \lambda$ implies that $\mu \in I$. Of course, the choice of words open and closed is no coincidence: I is closed iff its complement I^c is open. Moreover, the subspace $M_I = \bigcup_{\lambda \in I}$ is open (resp. closed) if I is open (resp. closed): this is a consequence of the stratification property. If I is open and $\lambda \in I^c$ is minimal, then $J = I \cup \lambda$ is open and the inductive step goes from M_I to M_J. Again from (4.3) we have that $M_\lambda = M_J \setminus M_I$ is a closed submanifold of M_J along with the corresponding exact sequence

$$\cdots \to H^{q-k}(M_\lambda) \to H^q(M_J) \to H^q(M_I) \to \dots$$

where we have used the Thom isomorphism $H^{q-k} \cong H^q(M_J, M_I)$ with $k = k_\lambda = \operatorname{codim} M_\lambda$. Given a field K of coefficients, we say that the stratification is perfect over K if $P_t(M) = \sum t^{k_\lambda} P_t(M_\lambda)$, which happens if the long exact sequence above breaks up into short exact sequences for all q and all λ. As for G-invariant stratifications, we simply say that it is equivariantly perfect (or G-perfect) if the corresponding equivariant cohomology sequences break up.

However, we will be interested in applying this argument to infinite-dimensional spaces, although the strata will still have finite codimension. In that context, the indexing set will be countably infinite, so to ensure that our inductive process still apply we will require the following finiteness property (called condition A): for every finite subset I there are a finite number of minimal elements of the complement I^c. Although the induction will never end, only finitely many steps will be needed to compute $H^q(M)$ if we require the additional condition (called condition B): for each integer q there are only finitely many indices $\lambda \in I$ for which $\operatorname{codim} M_\lambda < q$.

4.2 Defining the stratification

4.2.1 The space of holomorphic structures

Our goal is the following: to classify the holomorphic vector bundles over X. More precisely, we would like to construct a space parametrizing all such bundles (up to isomorphism), the so-called moduli space, and then describe its geometric properties. We first reduce the problem to the one of classifying the holomorphic vector bundles of the same topological type, i.e. of the same rank and degree (first Chern class). To do this, it is enough to study the space $\mathscr{C}$ of holomorphic structures on a fixed smooth complex vector bundle E over X of rank n and degree k. Indeed, smooth complex vector bundles over surfaces are determined (up to isomorphism) by their rank and their degree. Therefore, once we fix the smooth complex vector bundle E, all holomorphic vector bundles of rank n and degree k are smoothly isomorphic to E and their holomorphic structure can be transported to a holomorphic structure on E.

The next step is to identify holomorphic structures on E with Dolbeault operators on E. In the following, let $\bar{\partial}$ denote the usual Dolbeault operator on X.

Definition 4.4. A Dolbeault operator on a smooth complex vector bundle $E \to X$ is a $\mathbb{C}$-linear operator $\bar{\partial}_E : \Gamma(E) \to \Omega^{0,1}(X,E)$ satisfying the Leibniz condition

$$\bar{\partial}_E(fs) = \bar{\partial}(f) \otimes s + f\,\bar{\partial}_E(s)$$

for every section $s \in \Gamma(E)$ and every smooth function $f \in \mathscr{E}(X)$.

Remark 4.5. In general, for smooth complex vector bundles over complex manifolds of higher dimension, a Dolbeault operator is an operator $\bar{\partial}_E : \Omega^{p,q}(X,E) \to \Omega^{p,q+1}(X,E)$ as above that is also required to satisfy $\bar{\partial}_E^2 = 0$. In our case, that condition is already satisfied since X has complex dimension 1.

Given a holomorphic vector bundle $E \to X$, there is a natural Dolbeault operator $\bar{\partial}_E$ associated to it, which is defined as follows. For a local holomorphic frame $e_1, \ldots, e_n$ of E over U, we have

$$\bar{\partial}_E\left(\sum_i s^i e_i\right) := \sum_i \bar{\partial}(s^i) \otimes e_i.$$

That $\bar{\partial}_E$ is well-defined is a consequence of the transition functions being holomorphic. Indeed, if $f_1, \ldots, f_n$ is another local holomorphic frame of E over U, then $f_i = \sum_{j=1}^n g_{ij} e_j$, where $g \in \mathrm{GL}_n(\mathscr{O}(U))$, hence

$$\begin{aligned}
\bar{\partial}_E\left(\sum_{i=1}^n t^i f_i\right) &= \bar{\partial}_E\left(\sum_{i=1}^n t^i \left(\sum_{j=1}^n g_{ij} e_j\right)\right) = \bar{\partial}_E\left(\sum_{j=1}^n \left(\sum_{i=1}^n t^i g_{ij}\right) e_j\right) \\
&= \sum_{j=1}^n \sum_{i=1}^n \bar{\partial}(t^i) g_{ij} \otimes e_j + \sum_{j=1}^n \sum_{i=1}^n t^i\, \bar{\partial}(g_{ij}) \otimes e_j \\
&= \sum_{i=1}^n \left(\bar{\partial}(t^i) \otimes \sum_{j=1}^n g_{ij} e_j\right) = \sum_{i=1}^n \bar{\partial}(t^i) \otimes f_i
\end{aligned}$$

Conversely, the Koszul–Malgrange theorem asserts that given a Dolbeault operator $\bar{\partial}_E$ on a smooth complex vector bundle $E \to X$, there exists a unique holomorphic structure on E such that its associated Dolbeault operator is $\bar{\partial}_E$. More precisely, one can show that given a Dolbeault operator $\bar{\partial}_E$, there exist local frames of "holomorphic" sections everywhere. Here, a section s is said to be holomorphic if $\bar{\partial}_E(s) = 0$ (note that $\bar{\partial}_E$ can be restricted to any open subset $U \subset X$ since it is a local operator). This implies the existence of a smooth trivialization $\{\phi_i : \pi^{-1}(U_i) \to U_i \times \mathbb{C}^n\}$, where $\phi_i^{-1}(-,e_k)$ is a holomorphic section for all $i \in I$ and $k = 1, \ldots, n$. But this trivialization is in fact holomorphic since given the local holomorphic frames $s_k = \phi_i^{-1}(-,e_k)$ and $t_k = \phi_j^{-1}(-,e_k)$ $(k = 1, \ldots, n)$ over U_i and U_j respectively, the transition function $g_{ij} : U_i \cap U_j \to \mathrm{GL}_n(\mathbb{C})$ satisfies

$$t_\ell = \sum_{k=1}^n (g_{ij})_{k\ell} s_k \quad \forall \ell = 1, \ldots, n$$

and applying $\bar{\partial}_E$ to both sides yield

$$0 = \bar{\partial}_E(t_\ell) = \sum_{\ell=1}^{n} \bar{\partial}((g_{ij})_{k\ell}) \otimes s_k + \sum_{\ell=1}^{n} (g_{ij})_{k\ell}\, \bar{\partial}_E(s_k) = \sum_{\ell=1}^{n} \bar{\partial}((g_{ij})_{k\ell}) \otimes s_k \quad \forall \ell = 1, \dots, n$$
$$\implies \bar{\partial}((g_{ij})_{k\ell}) = 0 \quad \forall k, \ell = 1, \dots, n.$$

This establishes the one-to-one correspondence between holomorphic structures on E and Dolbeault operators on E. Now, consider two Dolbeault operators $\bar{\partial}_E$ and $\bar{\partial}'_E$ on E. It follows from the Leibniz condition that $\bar{\partial}_E - \bar{\partial}'_E$ is $C^\infty(X)$-linear:

$$(\bar{\partial}_E - \bar{\partial}'_E)(fs) = \bar{\partial}(f) \otimes s + f\, \bar{\partial}_E(s) - \Big(\bar{\partial}(f) \otimes s + f\, \bar{\partial}'_E(s)\Big) = f(\bar{\partial}_E - \bar{\partial}'_E)(s).$$

Therefore, $\bar{\partial}_E - \bar{\partial}'_E$ correspond to a bundle map $E \to \bigwedge^{0,1} T^*X \otimes E$ and can be seen as a (0,1)-form on X with values in the bundle $E^* \otimes E = \mathrm{End}(E)$ of smooth endomorphisms of E. In particular, this implies that $\mathscr{C}$ is a complex affine space whose vector space of translations is $\Omega^{0,1}(X, \mathrm{End}(E))$.

Now, the automorphism group $\mathrm{Aut}(E)$ of E acts on $\mathscr{C}(E)$ and the orbits are by definition isomorphism classes of holomorphic vector bundles of fixed rank and degree. We now want to describe the orbit structure of that action. In order to get a good moduli space (e.g. to avoid non-Hausdorff phenomena), we will need to consider (semi-)stable holomorphic structures.

Definition 4.6. Let $E \to X$ be a holomorphic vector bundle. The first Chern class of E is the 2-form $\frac{i}{2\pi}\,\mathrm{trace}(F)$ on X, where F is the curvature of any connection on E. The degree of E, denoted by $\deg(E)$, is defined to be the integral of its first Chern class on the fundamental cycle of X.

Definition 4.7. The slope of a (non-zero) holomorphic vector bundle E is the ratio

$$\mu(E) = \frac{\deg(E)}{\mathrm{rk}(E)},$$

where $\mathrm{rk}(E)$ denotes the rank of E. Moreover, E is said to be (semi-)stable if for any (non-zero) proper holomorphic subbundle F of E, we have $\mu(F) \overset{(\leq)}{<} \mu(E)$.

4.2.2 The Harder-Narasimhan filtration

Following the paper of Harder–Narasimhan (cf. [6]), we exhibit a very useful canonical filtration of holomorphic vector bundles.

Lemma 4.8. *If $0 \to E' \to E \to E'' \to 0$ is a short exact sequence of holomorphic vector bundles and E is semi-stable, then*

$$\mu(E') \leq \mu(E) \leq \mu(E'').$$

Proof. The inequality $\mu(E') \leq \mu(E)$ follows readily from the semi-stability of E. Next, we know that the degree and the rank of these bundles satisfy

$$\deg(E) = \deg(E') + \deg(E''), \quad \mathrm{rk}(E) = \mathrm{rk}(E') + \mathrm{rk}(E'').$$

Keeping in mind the first inequality, it follows that

$$\begin{aligned}
\mu(E) = \frac{\deg(E') + \deg(E'')}{\mathrm{rk}(E') + \mathrm{rk}(E'')} &\implies \frac{\mathrm{rk}(E') + \mathrm{rk}(E'')}{\mathrm{rk}(E')}\mu(E) = \mu(E') + \frac{\deg(E'')}{\mathrm{rk}(E')} \\
&\implies \frac{\mathrm{rk}(E'')}{\mathrm{rk}(E')}\mu(E) \leq \frac{\deg(E'')}{\mathrm{rk}(E')} \\
&\implies \mu(E) \leq \mu(E'').
\end{aligned}$$

□

Definition 4.9. A proper subbundle F of a holomorphic vector bundle E is said to be maximal if for every subbundle F' of E strictly containing F, we have $\mu(F) > \mu(F')$.

Lemma 4.10. *If $0 \to F \to F' \to Q \to 0$ is an exact sequence of holomorphic vector bundles and F is maximal, then*

$$\mu(F) > \mu(F') > \mu(Q).$$

In particular, if F is a maximal subbundle of a holomorphic bundle E, then for every proper subbundle Q of E/F, we have $\mu(F) > \mu(Q)$.

Proof. The first inequality readily follows from the maximality of F, while the second can be proved as follows:

$$\begin{aligned}
\mu(F') = \frac{\deg(F) + \deg(Q)}{\mathrm{rk}(F) + \mathrm{rk}(Q)} &\implies \frac{\mathrm{rk}(F) + \mathrm{rk}(Q)}{\mathrm{rk}(F)}\mu(F') = \mu(F) + \frac{\deg(Q)}{\mathrm{rk}(F)} \\
&\implies \frac{\mathrm{rk}(Q)}{\mathrm{rk}(F)}\mu(F') > \frac{\deg(Q)}{\mathrm{rk}(F)} \\
&\implies \mu(F') > \mu(Q).
\end{aligned}$$

□

Lemma 4.11. *Let F_1 and F_2 be a semi-stable and a maximal subbundle of E, respectively. If F_1 is not contained in F_2, then $\mu(F_2) > \mu(F_1)$.*

Proof. By assumption, the canonical map $p : F_1 \to E/F_2$ is non-zero. Moreover, it factorizes through an exact diagram

$$\begin{tikzcd}
F_1 \arrow[r] \arrow[d, "p"'] & F_1' \arrow[r] \arrow[d, "f"] & 0 \\
E/F_2 & F_1'' \arrow[l] & 0 \arrow[l]
\end{tikzcd}$$

where f is of maximal rank (cf. [8, §4]). In particular, $\mu(F_1'') \geq \mu(F_1')$. It follows from Lemma 4.8 and Lemma 4.10 that $\mu(F_1) \leq \mu(F_1')$ and $\mu(F_2) > \mu(F_1'')$. Combining the three inequalities yields the desired result. □

Lemma 4.12. *Let F_1, F_2 be maximal semi-stable subbundles of a holomorphic vector bundle E. Then $F_1 = F_2$.*

Proof. Suppose that F_1 is not contained in F_2. Lemma 4.11 implies that $\mu(F_2) > \mu(F_1)$. Appealing to Lemma 4.11 again, this time with the roles of F_1 and F_2 reversed, we must have $F_2 \subsetneq F_1$. But the semi-stability of F_1 forces $\mu(F_2) \leq \mu(F_1)$, a contradiction. Thus, $F_1 \subset F_2$. The other inclusion is proved similarly. □

Lemma 4.13. *Let E be a holomorphic vector bundle. Then the values of $\mu(F)$, where F is a subbundle of E, are bounded from above.*

Proof. We prove the result by induction on the rank of E. For the case of line bundles, there is nothing to prove. Now, fix a subbundle F of E. Then for any other subbundle F' we have that the bundles obtained from $F' \cap F$ and $F'/(F \cap F') \subset E/F$ both have bounded slope, and therefore so does F. □

Theorem 4.14. *If E is a holomorphic vector bundle that is not semi-stable, then there exists a unique maximal semi-stable subbundle F of E.*

Proof. The uniqueness of F is exactly the content of Lemma 4.12. For the existence, consider $m := \sup \mu(F)$, where the supremum is taken over all the subbundles F of E. Since the values of $\mu(F)$ are discrete and bounded from above (Theorem 4.13), the supremum is attained.

Among those subbundles with $\mu(F) = m$, choose one, say F_0, that is of maximal rank. Note that F_0 must be a proper subbundle of E, because otherwise E would be semi-stable. Moreover, F_0 is semi-stable by definition since for every proper subbundle $F' \subset F$ we have $\mu(F') \leq m = \mu(F_0)$. Now, if F' is a subbundle of E that strictly contains F, then $\mathrm{rk}(F) < \mathrm{rk}(F')$. Of course, $\mu(F') \leq m$, but in fact $\mu(F') < m$ because otherwise this would contradict the maximality of the rank of F_0. Therefore, F_0 is maximal as well. □

Theorem 4.15 (Harder–Narasimhan filtration)**.** *Let E be a holomorphic vector bundle. There exists a unique filtration*

$$0 = F_0 \subset F_1 \subset \cdots \subset F_r = E$$

such that each quotient $D_i := F_i/F_{i-1}$ is semi-stable and maximal in E/F_{i-1}. In particular,

$$\mu(D_1) > \mu(D_2) > \cdots > \mu(D_r).$$

Proof. We simply apply Theorem 4.14 repeatedly. If E is semi-stable, then we are done. If not, then we can find a maximal semi-stable subbundle F_1 of E. If E/F_1 is semi-stable, then we are done. If not, then we can find a maximal semi-stable subbundle F_2' of E/F_1. The preimage F_2 of F_2' under the projection $E \to E/F_1$ is again semi-stable, and the maximality of F_1 implies that $\mu(F_1) > \mu(F_2) > \mu(F_2/F_1)$. Repeating this process yields the desired filtration. Finally, the uniqueness can be proved inductively on the rank of E. Indeed, F_1 is unique from Theorem 4.14 while the quotients F_i/F_1, $i \geq 2$, form a filtration of E/F_1 and are therefore also uniquely determined by induction. $\square$

Definition 4.16. Let $0 = F_0 \subset F_1 \subset \cdots \subset F_r = E$ be the Harder-Narasimhan filtration of a holomorphic vector bundle E of rank n. If we let k_i and n_i denote respectively the degree and the rank of the quotients $D_i = F_i/F_{i-1}$, then the type of E is the vector $\mu = (\mu_1, \ldots, \mu_n)$, where the first n_1 entries are equal to $\mu(D_1) = k_1/n_1$, the next n_2 entries are equal to $\mu(D_2) = k_2/n_2$, and so on until the last n_r entries which are equal to $\mu(D_r) = k_r/n_r$.

Note that the sequence of pairs (n_i,k_i), $i = 1, \ldots, r$ can be read from the type μ of E by taking the ith different value appearing in μ and multiplying it by the number of times it appears inside μ.

Definition 4.17. Let $\mathscr{C}_\mu$ denote the subspace of all holomorphic vector bundles of type μ. In particular, if $\mu = (k/n, \ldots, k/n)$, then $\mathscr{C}_\mu$ is the space $\mathscr{C}^{ss}$ of semi-stable holomorphic bundles of degree k and rank n.

By uniqueness, two isomorphic complex structures on E must lead to the same Harder–Narasimhan filtration on E. In other words, the subspace $\mathscr{C}_\mu$ is preserved by the action of $\mathrm{Aut}(E)$ and is consequently a union of orbits.

4.2.3 The emergence of the stratification

Recall that the infinitesimal variations of a holomorphic vector bundle are classified by $H^1(M, \mathrm{End}(E))$. In our context, this gets interpreted as follows: the orbit in $\mathscr{C}$ corresponding to a given holomorphic bundle E is locally a submanifold of finite codimension and its normal can be identified with $H^1(X, \mathrm{End}(E))$. This is because an infinitesimal gauge transformation (i.e. a global endomorphism ϕ of E) alters $\bar{\partial}_E$ by the addition of $\bar{\partial}_E \phi$ and the cokernel of $\Omega^0(\mathrm{End}(E)) \xrightarrow{\bar{\partial}_E} \Omega^{0,1}(\mathrm{End}(E))$ is just $H^1(X, \mathrm{End}(E))$.

Similarly, we can identify the normal to $\mathscr{C}_\mu$, which should be a quotient of $H^1(X, \mathrm{End}(E))$ since $\mathscr{C}_\mu$ is a union of orbits. If we let $\mathrm{End}'(E)$ denote the bundle of holomorphic endomorphisms of E that preserve its canonical filtration, then we get an exact sequence of vector bundles

$$0 \to \mathrm{End}'(E) \to \mathrm{End}(E) \to \mathrm{End}''(E) \to 0.$$

But from the long exact sequence in cohomology we see that $H^1(X, \mathrm{End}''(E))$ is indeed a quotient of $H^1(X, \mathrm{End}(E))$ and the fact that $H^1(X, \mathrm{End}'(E))$ describes variation inside $\mathscr{C}_\mu$ should convince us that $H^1(X, \mathrm{End}''(E))$ is the right candidate for the normal to $\mathscr{C}_\mu$. Moreover, $\dim H^1(X, \mathrm{End}''(E))$ depends only on μ (and not on the holomorphic structure). This follows at once from Riemann-Roch

$$\dim H^0(X,\mathscr{E}) - \dim H^1(X,\mathscr{E}) = \deg(\mathscr{E}) + (g-1)\,\mathrm{rk}(\mathscr{E})$$

together with the fact that $H^0(X, \mathrm{End}''(E)) = 0$, the latter following from: if E,D are both semi-stable and $\mu(E) > \mu(D)$, then every homomorphism $E \to D$ is zero.

All of this sketching was to shed light on the Morse theory lurking in the background. Indeed, the emerging picture is that of a stratification of the space $\mathscr{C}$ by the local submanifolds $\mathscr{C}_\mu$, which are all of finite codimension, giving some sort of cell-structure attached on $\mathscr{C}^{ss}$. The striking feature of that stratification is upon looking at the relative positions of the different pieces $\mathscr{C}_\mu$. Indeed, we will see that it is possible to order the types μ such that

$$\overline{\mathscr{C}_\mu} \subset \bigcup_{\lambda \geq \mu} \mathscr{C}_\lambda, \tag{4.18}$$

just as we have for Morse stratifications.

The first thing we need in order to get a Morse-like stratification is to order our indexing set, the possible types λ of E. This ordering has a particularly nice geometric flavour. Recall that from a type λ we can read the rank and the degree (n_i,k_i) of the quotients D_i appearing in a corresponding Harder–Narasimhan filtration of E. Using this sequence of pairs, we can form the polygon P_λ with vertices (0,0), (n_1,k_1), (n_1+n_2,k_1+k_2), $\ldots$, (n,k), which will sit in the first quadrant of the plane. Note that the polygon P_λ is convex since the quotients k_i/n_i appearing in λ are arranged in decreasing order.

Definition 4.19. We say that $\lambda \geq \mu$ if the polygon P_λ is above the polygon P_μ.

We can also describe the partial order $\geq$ on types in purely numerical terms. Indeed, notice that the polygon P_λ is the graph of the function whose value at an integer i is $\sum_{j \leq i} \mu_j$ and that interpolates linearly between integers. Therefore, $\lambda \geq \mu$ if and only if $\sum_{j \leq i} \lambda_j \geq \sum_{j \leq i} \mu_j$ for all $j = 1, \ldots, n$.

We now spend some time to show that our stratification satisifies (4.18). To do this, we will identify the space $\mathscr{C}$ of holomorphic structures on E with the space $\mathscr{A}$ of unitary connections on E, which will allow to transport the stratification of $\mathscr{C}$ to a stratification of $\mathscr{A}$. We will then define a functional on $\mathscr{A}$, the so-called Yang–Mills functional, and slowly come to the conclusion that the stratification on $\mathscr{A}$ behaves in every aspect as the Morse stratification of this functional, should it exist.

4.2.4 Notions of convexity

While the definition of the partial order on types is still fresh in our memory, we make a quick detour to expand on the convexity ideas it leads to. Pretty much all the material covered here can be found in [7] and is summarized in [1, §12]. The partial order on types actually comes from a partial order defined on all n-tuples $(\mu_1, \ldots, \mu_n)$ of real numbers: we define $\mu \leq \lambda$ if, after rearranging the entries in decreasing order, we have

$$\sum_{j\leq i} \mu_j \leq \sum_{j\leq i} \lambda_j \quad \text{for} \quad i = 1,\ldots,n-1,$$
$$\sum_{j=1}^{n} \mu_j = \sum_{j=1}^{n} \lambda_j.$$

It is shown in [7] that $\mu \leq \lambda$ is equivalent to $\mu = P\lambda$, where P is a doubly stochastic matrix. But it is a theorem of Birkhoff that the doubly stochastic matrices are the convex hull of the permutation matrices. Thus $\mu \leq \lambda$ becomes equivalent to $\widehat{S_n\mu} \subset \widehat{S_n\lambda}$, where $S_n\mu$ denotes the orbit of $\mu \in \mathbb{R}^n$ under the symmetric group on n elements S_n and $\widehat{\cdot}$ denotes the convex hull operator.

Geometric notions of convexity can be transformed into statements about convex functions using the following duality: $x \in \widehat{C}$ if and only if $\phi(x) \leq \sup_C \phi$ for every convex function $\phi : \mathbb{R}^n \to \mathbb{R}$. Thus taking ϕ to be a convex symmetric function (i.e. invariant under S_n), we have

$$\phi(\mu) \leq \phi(\lambda) \text{ for all convex symmetric functions on } \mathbb{R}^n,$$

which is in fact also equivalent to $\mu \leq \lambda$.

Now, Schur showed that if μ_j, $j = 1,\ldots,n$ are the diagonal elements of a hermitian matrix whose eigenvalues are λ_j, then $\mu \leq \lambda$. Horn proved the converse, so that $\mu \leq \lambda$ is equivalent to the λ_j being eigenvalues of a hermitian matrix with diagonal elements μ_j. This is further equivalent to $\widehat{C(\mu)} \subset \widehat{C(\lambda)}$ where $C(\lambda)$ denotes the conjugacy class of hermitian matrices with eigenvalues λ_j. This last fact implies that $\Psi(B) \leq \Psi(A)$ for all convex invariant Ψ on the space of Hermitian matrices, and is in fact equivalent to it.

These results are in no way special to $U(n)$. In general, for G a compact Lie group, the role of the hermitian matrices is now played by the Lie algebra $\mathfrak{g}$ of G, the diagonal matrices are replaced by the Lie algebra $\mathfrak{t}$ of a maximal torus T, S_n becomes the Weyl group of G and we say that $y \leq x$ if $x - y$ lies in the dual cone C^* of a fixed positive Weyl chamber C. The result takes the following form:

Theorem 4.20. *The following conditions for $x,y \in \mathfrak{t}$ are equivalent:*

- *i)* $y \leq x$
- *ii)* $\widehat{Wy} \subset \widehat{Wx}$
- *iii)* $\phi(y) \leq \phi(x)$ *for all W-invariant convex functions ϕ on $\mathfrak{t}$*
- *iv)* $\widehat{Gy} \subset \widehat{Gx}$
- *v)* $\Psi(y) \leq \Psi(x)$ *for all G-invariant convex functions Ψ on $\mathfrak{g}$*

In particular, we list two applications of convexity for later use. First, in Theorem 4.23 we will need the inequality

$$\phi\begin{pmatrix} A & B \\ C & D \end{pmatrix} \geq \phi\begin{pmatrix} \alpha & 0 \\ 0 & \delta \end{pmatrix}$$

for every convex invariant function ϕ on $\mathfrak{u}(n)$, where $\begin{pmatrix} A & B \\ C & D \end{pmatrix}$ is skew-hermitian and α, δ are the diagonal components of A,D. This follows from the Schur-Horn theorem since the second matrix is in the convex hull of the S_n-orbit of the diagonal part of the first matrix.

We will also need to know that $\phi(y) = \phi(x)$ for all convex invariant function ϕ if and only if $y = x$. This statement follows from Theorem 4.20 since $\widehat{Wy} = \widehat{Wx}$ implies that extreme points of these polyhedra must coincide, so Wx and Wy intersect and in fact coincide as well.

4.3 The Yang-Mills functional

We first recall some general definitions and elementary results concerning smooth vector bundles. In this section, $\mathscr{F}$ will denote the space $C^\infty(X)$ of smooth functions on X. For E a smooth vector bundle over X, a connection A on E is a map

$$\begin{aligned} \nabla : \mathfrak{X}(M) \times \Gamma(E) &\to \Gamma(E) \\ (X,s) &\mapsto \nabla_X s \end{aligned}$$

such that $\nabla_X s$ is $\mathscr{F}$-linear in X, $\mathbb{R}$-linear in s and satisfies the Leibniz rule

$$\nabla_X(fs) = (Xf)s + f\nabla_X s \quad \forall f \in \mathscr{F}.$$

Since $\mathscr{F}$-linear maps $\alpha : \Gamma(E) \to \Gamma(F)$ correspond bundle maps $\varphi : E \to F$, we can also see a connection as an operator $d_A : \Omega^0(X,E) \to \Omega^1(X,E)$, where $\Omega^k(X,E)$ denotes the space of E-valued k-forms on X, satisfying the Leibniz rule

$$d_A(fs) = (df)s + fd_A s.$$

Let $e_1, \ldots, e_n$ be a frame for E over an open set U. Since any section $s \in \Gamma(U,E)$ is a linear combination $s = \sum a^j e_j$, the section $d_A s$ can be computed from $d_A e_i$ by linearity and the Leibniz rule. As a section of E over U, $d_A e_i$ is a linear combination of the sections $e_1, \ldots, e_n$ with coefficients $\omega_{ij} \in \Omega^1(U)$:

$$d_A e_i = \sum_j \omega_{ij} e_j.$$

The matrix of 1-forms $\omega = [\omega_{ij}]$ is called the connection matrix of the connection d_A relative to the frame $e_1, \ldots, e_n$ on U.

Now, let E be a smooth complex vector bundle over X, which we equip with a fixed Hermitian metric. A connection A on E is said to be unitary if it is compatible with the Hermitian metric, i.e.

$$d(s_1,s_2) = (d_A s_1,s_2) + (s_1,d_A s_2)$$

Therefore, with respect to a unitary frame over an open subset U, we see that the connection matrix of d_A is skew-hermitian, so we can see it as a 1-form on U with values in the vector space $\mathfrak{g} := \mathfrak{u}(n)$.

The group of unitary automorphisms of E acts on the space of connections via

$$\begin{aligned}(u \cdot d_A)s &= u d_A(u^{-1}s) = uu^{-1}d_A s + u(du^{-1})s \\ &= d_A s + u(-u^{-1}(du)u^{-1})s = d_A s - (du)u^{-1}s\end{aligned}$$

where we extended the operator d to matrices. Moreover, if d_A is a unitary connection, then $d_A + \eta$, where $\eta \in \Omega^1(X,\mathfrak{g})$, is also a unitary connection, where η acts on $\Omega^0(X,E)$ simply by contraction.

The curvature $F(A)$ of a connection A is the operator $d_A \circ d_A : \Omega^0(X,E) \to \Omega^2(X,E)$, where we have naturally extended d_A to $\alpha \in \Omega^k(X)$, $\beta \in \Omega^\ell(X)$ using the Leibniz rule

$$d_A(\alpha \wedge \beta) = (d_A\alpha) \wedge \beta + (-1)^k \alpha \wedge (d_A\beta).$$

A simple calculation shows that $F(A)$ is linear over smooth functions:

$$d_A(d_A(fs)) = d_A(f d_A s + df \cdot s) = df \cdot d_A s + f d_A d_A s + ddf \cdot s + (-1)df \cdot d_A s = f d_A d_A s.$$

The curvature can thus be seen as a 2-form on X with values in the bundle $\mathfrak{g}_E$ of skew-adjoint endomorphisms or, equivalently, as a $\mathfrak{g}$-valued 2-form on E (note that requiring d_A to be unitary implies that $F(A)$ is $\mathfrak{g}$-valued instead of merely $\mathrm{End}(E)$-valued). Finally, the curvature transforms as a tensor under unitary automorphisms, i.e.

$$F(u \cdot A) = uF(A)u^{-1}.$$

Let $\Omega^k(X,\mathfrak{g})$ denote the space of k-forms on X with values in the vector space $\mathfrak{g}$. Given a bilinear map $\mu : \mathfrak{g} \times \mathfrak{g} \to \mathfrak{g}$, we can use the usual exterior multiplication to extend this to a pairing $\mu : \Omega^k(X,\mathfrak{g}) \otimes \Omega^\ell(X,\mathfrak{g}) \to \Omega^{k+\ell}(X,\mathfrak{g})$ defined by

$$\mu(\alpha,\beta)_p(u_1,\ldots,u_{k+\ell}) = \sum_{\sigma \in S_{k+\ell}} (\operatorname{sgn}\sigma)\mu(\alpha_p(u_{\sigma(1)},\ldots,u_{\sigma(k)}),\beta_p(u_{\sigma(k+1)},\ldots,u_{\sigma(k+\ell)})).$$

Equivalently, if α and β are given by $\sum \alpha^i A_i$ and $\sum \beta^j A_j$, where $\{A_1,\ldots,A_n\}$ is a set of vectors in $\mathfrak{g}$, then

$$\mu(\alpha,\beta) = \sum_{i,j} (\alpha^i \wedge \beta^j)\mu(A_i,A_j).$$

Recall that on $\mathfrak{g} = \mathfrak{u}(n)$, we have the bracket operation $[X,Y] = XY - YX$ and the inner product $\langle X,Y \rangle = \mathrm{trace}(Y^* X)$ which is invariant under conjugation, the adjoint action of

$G = U(n)$ on its Lie algebra $\mathfrak{g}$. Both of these operations are bilinear on $\mathfrak{g}$ and hence extend to products on $\Omega^*(X,\mathfrak{g})$. The invariance of $\langle \cdot\,,\cdot \rangle$ means that on $\mathfrak{g}$ we have $\langle [x,y],z \rangle = \langle x,[y,z] \rangle$, and this property extends to $\Omega^k(X,\mathfrak{g})$ as well.

Suppose now that a fixed Riemannian metric and a fixed orientation are chosen on X. We then have the corresponding Hodge star operator $* : \Omega^k(X) \to \Omega^{2-k}(X)$ characterized by

$$\eta \wedge *\eta = \langle \eta,\eta \rangle_X \operatorname{vol}(X),$$

where $\langle \cdot\,,\cdot \rangle_X$ denotes the Riemannian structure on $\Omega^k(X)$ and $\operatorname{vol}(X)$ is the unique form of length 1 in the orientation of X. We then extend $*$ linearly to $\Omega^k(X,\mathfrak{g})$, giving that space a natural inner product:

$$(\eta,\zeta) = \int_X \eta \wedge *\zeta.$$

The L^2-norm of a form $\eta \in \Omega^k(X,\mathfrak{g})$ is then given by $\|\eta\|^2 = (\eta,\eta)$. This allows to finally define the Yang-Mills functional as $L(A) = \|F(A)\|^2$, the L^2-norm of the curvature. Another way to interpret the Yang-Mills functional is to see that

$$L(A) = \int_X \phi(*F(A)),$$

where $\phi(X) = \operatorname{trace}(X^*X)$ is the sum of the square of the eigenvalues. But L is not the only functional of this form: taking ϕ to be any smooth function on the Lie algebra $\mathfrak{g}$ of G that is invariant under the adjoint action (i.e. conjugation) and convex yields a functional

$$\Phi(A) = \int_X \phi(*F(A)).$$

4.4 Completion of the Morse picture

This is where the bond with Morse theory becomes much stronger. We begin by showing that the Yang–Mills functional has a critical point structure that coincides with a lot of other similar functionals. Once this is done, we prove that our stratification satisfies Equation (4.3), as if it was the Morse stratification associated to the Yang–Mills functional. But nowadays this is no surprise: Daskalopoulos showed in [2] that this is really the case. Details for the material of this section are contained in [1, §8].

Let $\mathscr{A}$ be the space of unitary connections on our fixed Hermitian vector bundle E. A unitary connection A defines a holomorphic structure by taking the (0,1)-component d''_A of the covariant derivative d_A. We get a map $\mathscr{A} \to \mathscr{C}$, which is an affine linear isomorphism since there always exists a unique connection compatible with a given Hermitian and a given Dolbeault operator. A priori $\mathscr{A}$ and thus $\mathscr{A} \to \mathscr{C}$ depends on the metric, but ultimately it doesn't change anything since two metrics differ from a complex gauge transformation (i.e.

an element of $\mathrm{Aut}(E)$). Note that $\mathrm{Aut}(E)$ may be viewed as the complexification of the group of unitary gauge transformations $\mathscr{G}$ of E.

We now make a brief digression to quickly explain the link between critical connections of the Yang-Mills functional and unitary representations. It will be seen below that critical connections are solutions to the Yang-Mills equation $d_A * F(A) = 0$. Of course, flat connections will always be critical connections, and in fact correspond to the absolute-minimal of the Yang-Mills functional. Now for our fixed Riemann surface X, flat vector bundles E over X correspond to representations of $\pi_1(X,x_0)$ into $\mathrm{GL}(n,\mathbb{C})$. Indeed, if E is a vector bundle with a flat connection d_A, then there is a natural surjective homomorphism $\pi_1(X) \to \mathrm{Hol}(\nabla)/\mathrm{Hol}^0(\nabla)$ sending the homotopy class $[\gamma]$ to the coset $P_\gamma \cdot \mathrm{Hol}^0(\nabla)$, where $\mathrm{Hol}_{x_0}(\nabla) = \{P_\gamma \in \mathrm{GL}(E_{x_0}) | \gamma \text{ is a loop based at } x_0\}$ is the holonomy group of ∇ based at x_0 and $\mathrm{Hol}^0(\nabla)$ is the subgroup coming from contractible loops γ, which is in fact a surjective homomorphism onto $\mathrm{GL}(n,\mathbb{C})$ since the flatness of the connection implies that $\mathrm{Hol}^0(\nabla) = 0$. In other words, the parallel displacement along a curve γ starting at x_0 depends only the homotopy class of γ. Conversely, given a representation $\rho : \pi_1(X,x_0) \to \mathrm{GL}(n,\mathbb{C})$, we can construct a flat vector bundle E by setting $E = \widetilde{X} \times_\rho \mathbb{C}^n$, where $\widetilde{X}$ is the universal covering of X and $\widetilde{X} \times_\rho \mathbb{C}^n$ denotes the quotient of $\widetilde{X} \times \mathbb{C}^n$ by the action of π given by

$$\gamma : (x,v) \in \widetilde{X} \times \mathbb{C}^n \mapsto (\gamma(x),\rho(\gamma)v) \in \widetilde{X} \times \mathbb{C}^n.$$

Furthermore, restricting ourselves to flat unitary connections simply means that we substitute $\mathrm{GL}(n,\mathbb{C})$ for $G = U(n)$ in the above discussion.

Going one step further, we can describe non-zero solutions of the Yang–Mills equation by using a suitable central extension $\Gamma_\mathbb{R}$ of $\pi_1(X)$, although we won't dive into the details here. The interested reader may consult [1, §6] for completeness. The important thing to know is that given a unitary representation $\rho : \Gamma_\mathbb{R} \to G$, we get an induced unitary connection A_ρ that also satisfies the Yang–Mills equations, and this mapping induces a correspondence between conjugacy classes of unitary central representations ρ of $\Gamma_\mathbb{R}$ and equivalence classes of Yang–Mills connections.

Let $\mathscr{N} \subset \mathscr{A}$ be the set of connections giving the minimum for the Yangs-Mills functional, which as we just said are $\mathscr{G}$-equivalent to those given by representations $\rho : \Gamma_\mathbb{R} \to U(n)$ with $\rho(\mathbb{R})$ central. We let $\mathscr{N}_s \subset \mathscr{N}$ be those given by irreducible representations. The reason for this choice of notation is given by the Narasimhan–Seshadri criterion for stable bundles: a holomorphic vector bundle of rank n is stable if and only if it arises from an irreducible representation $\rho : \Gamma_\mathbb{R} \to U(n)$. Moreover, isomorphic bundles correspond to equivalent representations. Going back to our stratifications, this criterion translates as: $\mathscr{N}_s \subset \mathscr{C}_s$ and the induced map

$$\mathscr{N}_s/\mathscr{G} \cong \mathscr{C}_s/\mathscr{G}^c$$

is a homeomorphism. This statement is also equivalent to the following: an indecomposable holomorphic bundle E over X is stable if and only if there is a unitary connection on E having constant central curvature $*F = -2\pi i\mu(E)$, and such a connection is unique up to isomorphism. The latter was proved by Donaldson in [3].

Since direct sums of stable bundles with same slope are semi-stable we have $\mathscr{N} \subset \mathscr{C}^{ss}$. Now transport the stratification on $\mathscr{C}$ to get stratification on $\mathscr{A}$ by strata $\mathscr{A}_\mu$. Let $\mathscr{N}_\mu$ denote the Yang-Mills connections whose curvature is of type μ. Since such connections are direct sums of connections of the form $\mathscr{N}_s$ for smaller ranks, it follows that $\mathscr{N}_\mu \subset \mathscr{A}_\mu$. Our Morse picture on $\mathscr{A}$ is about to get clearer: the $\mathscr{N}_\mu$ are the critical submanifolds whose stable manifold should be $\mathscr{A}_\mu$. But before convincing ourselves of this, we show that this critical picture holds for a much larger class of functionals than Yang-Mills (maybe here one avenue would be to explain the picture for Yang-Mills and then carry on with all the others functionals).

We now begin our study of the Morse theory of the functionals Φ. Since $\mathscr{A}$ is an affine space, we will study the behavior of of our functionals along lines $A_t = A + t\eta$, where $\eta \in \Omega^1(M,\mathfrak{g})$. One important formula for the curvature is the following:

$$F(A+\eta) = (d_A+\eta)(d_A+\eta) = d_Ad_A + d_A\eta + \eta\wedge\eta = d_Ad_A + d_A\eta + \frac{1}{2}[\eta,\eta].$$

Applying this to A_t, we get

$$F(A_t) = F(A) + td_A\eta + \frac{1}{2}t^2[\eta,\eta],$$

and hence for the Yang-Mills functional L we have

$$L(A_t) = \|F_t\|^2 = \|F\|^2 + 2t(d_A\eta,F) + t^2(\|d_A\eta\|^2 + (F,[\eta,\eta])) + \text{higher terms}, \tag{4.21}$$

where we used the notation $F = F(A)$ and $F_t = F(A_t)$ to make the calculations less cluttered (and similarly for Φ). At an extremal connection, this yields

$$d_A * F = 0.$$

Indeed, if A is critical and we denote by d_A^* the adjoint of d_A with respect to our norm on $\Omega^*(X,\mathfrak{g})$, then $0 = (d_A\eta,F) = (\eta,d_A^*F)$ for all η, hence we must have $d_A^*F = 0$. But the adjoint of d_A is given by $- * d_A *$, so the result follows.

The calculation of the variations of Φ is only slightly more complicated. Taking again a line of connections $A_t = A + t\eta$, we get

$$\begin{aligned}\Phi_t &= \int_X \phi(*F + t * d_A\eta + \frac{1}{2}t^2 * [\eta,\eta]) \\ &= \int_X \phi(*F) + t\langle\phi'(*F), *d_A\eta\rangle \mod t^2 \\ &= \Phi(A) + t\int_X \langle\phi'(*F), *d_A\eta\rangle \mod t^2.\end{aligned}$$

where $\phi' : \mathfrak{g} \to \mathfrak{g}$ is the derivative of ϕ in the sense that $\phi(X+tY) = \phi(X) + t\langle\phi'(X),Y\rangle \mod t^2$. Rewriting the coefficient of t as

$$(\phi'(*F), * d_A\eta) = (\phi'(*F), - d_A^* * \eta) = (-d_A\phi'(*F), * \eta) = (- * d_A\phi'(*F),\eta),$$

we see that by definition, the gradient of Φ at A (with respect to our metric) is given by $\operatorname{grad}\Phi = - * d_A\phi'(*F)$. Moreover, writing

$$\phi(gXg^{-1} + tY) = \phi(g(X + tg^{-1}Yg)g^{-1}) = \phi(X + tg^{-1}Yg)$$

and expanding the first and the last term with respect to t, we see that ϕ' is an equivariant map. This further implies that $d_A\phi'(s) = \phi''(s) \circ d_A s$, so together with the formula for $\operatorname{grad}\Phi$ this enables us to conclude that if a connection A is critical for L, then it is critical for Φ, and the converse holds provided ϕ'' is invertible (which occurs when ϕ is strictly convex for example).

We now turn our attention towards the Morse indices of our functionals, starting with L. From Equation (4.21), we get that the quadratic form $Q(\eta,\eta)$ associated to the Hessian of L at a critical connection A is given by

$$Q(\eta,\eta) = \|d_A\eta\|^2 + (F,[\eta,\eta]).$$

But $\|d_A\eta\|^2 = (d_A^* d_A\eta,\eta)$ and

$$(F,[\eta,\eta]) = \int_X [\eta,\eta] \wedge *F = \int_X \eta \wedge [\eta, * F] = (-1)^{m+1} \int_X \eta \wedge * *^{-1} [*F,\eta],$$

so once one uses the formula for $*^{-1}$, the quadratic form reduces to

$$Q(\eta,\eta) = (d_A^* d_A\eta + *[*F,\eta],\eta).$$

Denote by $L_A = d_A^* d_A + *[*F, \ ]$ the operator appearing in the Hessian of L. The space of solutions $L_A\eta = 0$ for $\eta \in \Omega^1(X,\mathfrak{g})$ describe the tangent space to the space of solutions of L. Notice that L (as well as the functional Φ) is gauge-invariant because of the invariance of ϕ under the adjoint action:

$$\Phi(u \cdot A) = \int_X \phi(*F(u \cdot A)) = \int_X \phi(u(*F(A))u^{-1}) = \Phi(A).$$

Consequently, a better measure of the tangent space to the space of solutions would be the quotient of the solutions of $L_A\eta = 0$ by the directions along the orbits of the action of $\mathscr{G}$, which is precisely the image of $\Omega^0(X,\mathfrak{g})$ in $\Omega^1(X,\mathfrak{g})$ under d_A. Thus the corrected tangent space N_A to the space of solutions J_A fits in the exact sequence

$$\Omega^0(X,\mathfrak{g}) \xrightarrow{d_A} J_A \longrightarrow N_A \longrightarrow 0.$$

We call N_A the null space of Q_A and its dimension is the nullity of A. We now show that this nullity is always finite. With the norm on $\Omega^1(X,\mathfrak{g})$, the orthocomplement of the image of d_A is exactly the kernel of d_A^*. Therefore, N_A is the space of solutions $\eta \in \Omega^1(X,\mathfrak{g})$ satisfying $L_A\eta = 0$ and $d_A^*\eta = 0$ or, equivalently, $d_A^* d_A + d_A d_A^* + *[*F,] = 0$ and $d_A^* = 0$. Here we see the Laplacian Δ_A of d_A appearing, which is an elliptic operator. The operator on the left is thus elliptic, so its solutions are finite-dimensional, and as a byproduct the same goes for the nullity. This argument also extends to the Morse index of A, which is the dimension of a maximal subspace in the kernel of d_A^* on which the form $\hat{Q}(\eta) = (\Delta_A\eta + *[*F,\eta],\eta)$ is negative definite. Here the ellipticity of $\Delta_A + *[*F,]$ guarantees that its spectrum is discrete and bounded below so that there are only finitely many negative eigenvalues, showing that H has finite Morse index. All in all, we conclude that the index and the nullity of a critical connection A are both finite and equal to the index and nullity of the quadratic form

$$\widehat{Q}(\eta) = (\Delta_A\eta + \widehat{F}\eta,\eta), \quad \widehat{F} = *[*F,],$$

on the kernel of d_A^* in $\Omega^1(M,\mathfrak{g})$.

The story for the Hessian of Φ is very similar. Indeed, computations show that the Hessian correspond to the self-adjoint differential operator

$$Q = \phi''(*F(A))d_A^*d_A + \text{lower order}.$$

Since Φ is gauge invariant, we can restrict (just as we did for L) to the subclass of 1-forms $\eta \in \Omega^1(X,\mathfrak{g})$ for which $d_A^*\eta = 0$. Doing so, we can replace $d_A^*d_A$ with the Laplacian $\Delta_A = d_A^*d_A + d_Ad_A^*$, so that Q becomes a second-order elliptic differential operator. The strong convexity of ϕ ensures that the leading-order terms are positive definite, which is then again enough to make the spectrum discrete and bounded below so that there are only finitely many negative eigenvalues. Therefore, we have shown that H_Φ also has finite index.

The formula for $\operatorname{grad}\Phi$ and $\operatorname{grad} L$ together with the strong convexity of ϕ ensures that

$$(\operatorname{grad}\Phi, \operatorname{grad} L) \geq 0 \tag{4.22}$$

with equality if and only if A is critical. Thus we see that Φ is strictly decreasing along the paths of steepest descent for L. In a finite-dimensional setting this would already imply that the Morse indices of L and Φ coincide, but in our situation we need to make a little adjustment. We go back to Equation (4.22) and expand the inequality at a critical connection, discarding the higher-order terms along the way, to get

$$(H_\Phi\eta,H_L\eta) \geq 0$$

with inequality if and only if η is in the null-space of H_L (which coincides with that of H_Φ). After restricting η to the negative space V of H_L, the last inequality is reduced to a finite-dimensional one, which implies that H_Φ is negative definite on V. Thus the Morse index of Φ is at least equal to that of L. Upon reversing the roles of L and Φ in the last argument, we obtain that the Morse indices of Φ and L coincide.

Now that we know that the critical point structure is common to every one of our functionals, our goal shifts to showing that all these functionals lead to the same Morse strata, and that these strata correspond with the $\mathscr{C}_\mu$ defined earlier. The reason why we might expect such a goal to be attainable is as follows. One can show that $\operatorname{grad}\Phi$ is tangent to the $\mathscr{G}^c$-orbits, and hence that it preserves the $\mathscr{C}_\mu$. Since each $\mathscr{A}_\mu$ contains a unique component $\mathscr{N}_\mu$ of the critical set of Φ, we are lead to believe that $\mathscr{A}_\mu$ could be the stable manifold associated to

$\mathscr{N}_\mu$. Of course, a necessary condition for this to hold is that the functionals Φ reach their minimum only on $\mathscr{N}_\mu$. Now, for any $A \in \mathscr{N}_\mu$ the conjugacy class of $*F$ is constant and represented by the skew-hermitian diagonal matrix Λ_μ with entries $-2\pi i\mu_j$. It follows that $\Phi(A)$ takes the constant value $\phi(\Lambda_\mu)$, which we simply denote by $\phi(\mu)$.

Theorem 4.23. *For every $A \in \mathscr{A}_\mu$ and every convex invariant function ϕ on $\mathfrak{g}$, we have $\Phi(A) \geq \phi(\mu)$.*

Proof. We will assume for simplicity that the Harder–Narasimhan filtration only has two steps. The proof for the general case only involves more serious bookkeeping.

Suppose that the type μ is of the form

$$\mu_1 = \mu_2 = \cdots = \mu_r > \mu_{r+1} = \cdots = \mu_n.$$

We write $\mu^1 = \mu_1 = k_1/m_1$ and $\mu^2 = \mu_n = k_2/m_2$ for convenience. This means that for the holomorphic structure defined by $A \in \mathscr{A}_\mu$ we have an exact sequence of vector bundles

$$0 \to D_1 \to E \to D_2 \to 0,$$

where D_j has rank m_j and degree k_j. The curvature F of A can be written in the form

$$F = \begin{pmatrix} F_1 - \eta \wedge \eta^* & d\eta \\ -d\eta^* & F_2 - \eta^* \wedge \eta \end{pmatrix},$$

where F_j is the curvature of the unitary connection of D_j, $\eta \in \Omega^{0,1}(M, \mathrm{Hom}(D_2,D_1))$, η^* is its transposed conjugate and $d\eta$ is the covariant differential (cf. [5, §5]). Now, let f_j, α_j be scalar $m_j \times m_j$ matrices such that

$$\mathrm{trace}(f_j) = \mathrm{trace}(*F_j),$$
$$\mathrm{trace}(\alpha_1) = \mathrm{trace}(*(\eta \wedge \eta^*)) = -\,\mathrm{trace}(*(\eta^* \wedge \eta)) = -\,\mathrm{trace}(\alpha_2)$$

Since ϕ is a convex invariant function, recall from the section on convexity that it satisfies

$$\phi(*F(A)) \geq \phi \begin{pmatrix} f_1 - \alpha_1 & 0 \\ 0 & f_2 - \alpha_2 \end{pmatrix}.$$

The convexity of ϕ together with the fact that M has normalized volume implies that

$$\Phi(A) = \int_X \phi(*F(A)) \geq \phi \int_X \begin{pmatrix} f_1 - \alpha_1 & 0 \\ 0 & f_2 - \alpha_2 \end{pmatrix}.$$

Since the degree k_j of D_j is given by

$$k_j = \frac{i}{2\pi} \int_X \mathrm{trace}(f_j),$$

and f_j is a scalar matrix, it follows that $\int_X f_j$ is a scalar matrix whose diagonal entries are $-2\pi i k_j/m_j = -2\pi i \mu^j$. Also, since $\eta \in \Omega^{0,1}$, it follows that $-i\,\mathrm{trace}(\alpha_1)$ is non-negative so that

$$\int_X \alpha_1 = 2\pi i a_1,$$

where a_1 is a non-negative scalar $m_1 \times m_1$ matrix. Then

$$\int_X \alpha_2 = 2\pi i a_2,$$

where a_2 is the non-positive scalar $m_2 \times m_2$ matrix such that $\mathrm{trace}(a_2) = -\,\mathrm{trace}(a_1)$. Hence we have

$$\int_X \begin{pmatrix} f_1 - \alpha_1 & 0 \\ 0 & f_2 - \alpha_2 \end{pmatrix} = -2\pi i[\mu + a],$$

where $[\cdot]$ denotes the diagonal matrix whose diagonal components are the vector's entries. Thus, we obtain

$$\Phi(A) \geq \phi(\mu + a).$$

But since $a_1 \geq 0$, $a_2 \leq 0$ and $\mathrm{trace}(a_1) = -\,\mathrm{trace}(a_2)$, it follows that $\mu + a \geq \mu$ with respect to the partial order on types. Convexity of ϕ finally yields

$$\Phi(A) \geq \phi(\mu + a) \geq \phi(\mu).$$

□

We now know that $\Phi(A) \geq \phi(\mu)$ for any convex invariant function ϕ on $\mathfrak{g}$ and connection $A \in A_\mu$. We seek to strenghten this result. Given a holomorphic vector bundle E over X and one of the functionals Φ, we define

$$\Phi(E) = \inf_A \Phi(A),$$

where A runs over all unitary connections on E. Combining the Narasimhan–Seshadri criterion with Theorem 4.23, we obtain that $\Phi(E) = \phi(\mu)$ for stable bundles E.

This can now be extended to all bundles using the maximal nature of the canonical filtration, i.e. that the quotients D_j are semi-stable. First, arguments from [9] shows that any semi-stable bundle has a filtration with stable quotients all of which have the same slope. This together with Theorem 4.23 and the result for stable bundles yield the equality $\Phi(E) = \phi(\mu)$ for semi-stable bundles as well.

Now one can show that for E with filtration of arbitrary length with quotients D_j, we have $\Phi(E) \leq \Phi(\bigoplus D_j)$. For simplicity, we give the proof in the case where the filtration of E only has two steps. Consider the holomorphic exact sequence

$$0 \to D_1 \to E \to D_2 \to 0.$$

The metric on E gives rise to a connection whose curvature is

$$F(A) = \begin{pmatrix} F_1 - \eta \wedge \eta^* & d\eta \\ -d\eta^* & F_2 - \eta^* \wedge \eta \end{pmatrix},$$

where the element η defines a cohomology class in $H^1(M, \mathrm{Hom}(D_2, D_1))$ which classifies the extension. We then replace η by $t\eta$, which changes the extension class but not the isomorphism class of E. Letting $t \to 0$ we finally get that $\Phi(E) \leq \Phi(D_1 \oplus D_2)$. A final application of Theorem 4.23 to the filtration of E yields the equality for general E.

Thus we have proven:

Theorem 4.24. *For any convex invariant ϕ we have $\Phi(E) = \phi(\mu)$, where*

$$\Phi(E) = \inf_A \int_M \phi(*F(A))$$

and A runs over all unitary connections of E.

Recall from the section on convexity that if $\phi(\mu) = \phi(\nu)$ for all convex invariant ϕ, then we must have $\mu = \nu$. This together with Theorem 4.24 gives another characterization of the type: E is of type μ iff $\Phi(E) = \phi(\mu)$ for all convex invariant ϕ. Finally, it follows that if $\mathscr{C}_\lambda$ is in the closure of $\mathscr{C}_\mu$, we must have $\Phi(\mathscr{C}_\mu) \leq \Phi(\mathscr{C}_\lambda)$, which implies that $\phi(\mu) \leq \phi(\lambda)$ and hence that $\mu \leq \lambda$, as desired.

4.5 Equivariant cohomology of $\mathscr{C}^{ss}$

Going back to our stratification on $\mathscr{C}$, we see that Condition A (cf. §4.1.4) is satisfied. One way to see it is that if we are given a finite set of types I, then we can always create a big convex polygon P strictly encompassing all the convex polygons associated to I. Then the minimal elements of I^c are all contained inside P, which contains only finitely many convex polygons inside of it. Condition B is more involved, so let's settle for a sketch of the equivariant cohomology of $\mathscr{C}$. The reader interested in the technical details is invited to consult [1, §7,13]. First, we can show that the stratification of $\mathscr{C}$ is equivariantly perfect, so the exact sequence of equivariant cohomology coming from the stratification breaks up and the equivariant Poincaré series of $\mathscr{C}$ is given by

$$P_t(\mathscr{C}) = \sum_\mu t^{2d_\mu} P_t(\mathscr{C}_\mu), \tag{4.25}$$

where $d_\mu = \sum_{\mu_i > \mu_j} (\mu_i - \mu_j + g - 1)$ is the complex codimension of $\mathscr{C}_\mu$. While trying to compute the equivariant cohomology of a stratum $\mathscr{C}_\mu$, we encounter the following phenomenon. If we choose a smooth unitary decomposition of E as $D_1 \oplus \cdots \oplus D_r$, where D_i has rank n_i and degree d_i, and we let $\mathscr{C}(n_i, d_i)$ and $\mathscr{G}(n_i, d_i)$ denote respectively the space of holomorphic structures on D_i and the group of unitary automorphisms of D_i, then we find that the equivariant cohomology of $\mathscr{C}_\mu$ is isomorphic to the tensor product of the equivariant cohomology of the semi-stable strata for the quotients D_i:

$$H^*_{\mathscr{G}}(\mathscr{C}_\mu,\mathbb{Q}) \cong \bigotimes_{1\leq i\leq r} H^*_{\mathscr{G}(n_1,d_i)}(\mathscr{C}(n_i,d_i)^{ss},\mathbb{Q}). \tag{4.26}$$

Therefore, combining (4.25) and (4.26), we obtain the formula

$$P_t^{\mathscr{G}}(\mathscr{C}) = \sum_\mu t^{2d_\mu} \prod_{1\leq i\leq r} P_t^{\mathscr{G}(n_i,d_i)}(\mathscr{C}(n_i,d_i)^{ss}).$$

Since $\mathscr{C}$ is an infinite-dimensional affine space it is contractible so that its homotopy quotient is simply $B\mathscr{G}$. Therefore, we obtain an inductive formula

$$P_t^{\mathscr{G}}(\mathscr{C}^{ss}) = P_t(B\mathscr{G}) - \sum_{\mu\neq\mu_0} t^{2d_\mu} \prod_{1\leq i\leq r} P_t^{\mathscr{G}(n_i,d_i)}(\mathscr{C}(n_i,d_i)^{ss})$$

for the equivariant Betti numbers of $\mathscr{C}^{ss}$. In particular, Atiyah and Bott spelled out the details of that procedure in the case where the rank $n = 2$ and the degree $k = 1$. Here, the inductive formula becomes

$$P_t^{\mathscr{G}}(\mathscr{C}^s) + \sum_{r=0}^{\infty} t^{2(2r+g)} P_t^{\mathscr{G}}(\mathscr{C}_r) = P_t(B\mathscr{G}), \tag{4.27}$$

where $\mathscr{C}_r$ is the stratum corresponding to unstable bundles of type $(r+1,-r)$. For the stable bundles, they find that

$$P_t^{\mathscr{G}}(\mathscr{C}^s) = \frac{P_t(N(2,1))}{1-t^2} = \frac{(1+t)^{2g}P_t(N_0(2,1))}{1-t^2},$$

where $N(2,1)$ denotes the moduli space of stable bundles with rank 2 and degree 1 and $N_0(2,1)$ denotes the moduli space for the same bundles but with fixed determinant. For $\mathscr{C}_r$, they show that

$$P_t^{\mathscr{G}}(\mathscr{C}_r) = \left(\frac{(1+t)^{2g}}{1-t^2}\right)^2,$$

while for the whole space $\mathscr{C}$ we have

$$P_t(B\mathscr{G}) = \frac{((1+t)(1+t^3))^{2g}}{(1-t^2)^2(1-t^4)}.$$

We finally substitute these into (4.27) and sum the geometric series to get the formula

$$P_t(N_0(2,1)) = \frac{(1+t^3)^{2g} - t^{2g}(1+t)^{2g}}{(1-t^2)(1-t^4)}.$$

Bibliography

[1] Atiyah, M., and Bott, R. The yang-mills equations over riemann surfaces. *Philos. Trans. Roy. Soc. London*, A 308 (1982), 523–615.

[2] Daskalopoulos, G. The topology of the space of stable bundles on a compact riemann surface. *J. Diff. Geom.*, 36 (1992), 699–746.

[3] Donaldson, S. A new proof of a theorem of narasimhan and seshadri. *J. diff. Geom.*, 18 (1983), 269.

[4] Forster, O. *Lectures on Riemann Surfaces.* Springer-Verlag, New York, 1991.

[5] Griffiths, P., and Harris, J. *Principles of algebraic geometry.* Wiley, 1978.

[6] Harder, G., and Narasimhan, M. On the cohomology groups of moduli spaces of vector bundles over curves. *Mathematische Annalen*, 212 (1975), 215–248.

[7] Horn, A. Doubly stochastic matrices and the iagonal of a rotation matrix. *Am. J. Math*, 76 (1954), 620–630.

[8] Narasimhan, M., and Seshadri, C. Stable and unitary vector bundles on a compact riemann surface. *Annals of Mathematics*, 82 (1965), 540–567.

[9] Seshadri, C. Space of unitary vector bundles on a compact riemann surface. *Ann. Math.*, 85 (1967), 303–336.