I

Outlines

The three parts of this chapter provide outline accounts of three different topics. While all three are central to the subject of this book, the outlines serve three different purposes. In Section 1, we give a brief account of the Seiberg–Witten invariants, or monopole invariants, of smooth, closed 4-manifolds. These invariants, discovered by Seiberg and Witten and originally described in Witten's paper [125], are now the subject of several expository papers, published lecture notes and books. Our purpose here is to review the definition and main properties of these invariants, while establishing our notation and conventions.

Section 2 covers Morse theory, and specifically the manner in which one can recover the ordinary homology of a manifold *with boundary* from a "Morse complex", constructed from the data provided by the critical points and gradient-flow lines of a suitable Morse function. There are no proofs in this section. In the main part of this book, the Floer homology of a 3-manifold will be constructed by taking these constructions of Morse theory and repeating them in an infinite-dimensional setting. Proofs of the main propositions are presented only in the more difficult context of Floer homology; the finite-dimensional constructions are presented here for motivation, to provide a framework that explains the origin of many arguments. Although some notation is introduced, no essential use is made of this material in the later chapters.

Finally in this chapter, Section 3 provides an outline of the main results of this book. We describe the principal features and properties of the monopole Floer homology groups of 3-manifolds; we explain how their construction is related to the Morse theory of Section 2, and we explain the role of Floer homology in computing the monopole invariants of closed 4-manifolds.

1 Monopole invariants of four-manifolds

1.1 Spin^c structures

Spin^c structures can be considered on manifolds of any dimension, but we will focus here on dimensions 3 and 4, the two cases we will need. We begin with 3-manifolds.

Let Y be a closed, oriented, Riemannian 3-manifold. A spin^c structure on Y consists of a unitary rank-2 vector bundle $S \to Y$ with a Clifford multiplication

$$\rho : TY \to \mathrm{Hom}(S, S).$$

Clifford multiplication is a bundle map that identifies TY isometrically with the subbundle $\mathfrak{su}(S)$ of traceless, skew-adjoint endomorphisms equipped with the inner product $\frac{1}{2}\operatorname{tr}(a^*b)$. It also respects orientation, which by convention means that

$$\rho(e_1)\rho(e_2)\rho(e_3) = 1$$

when the e_i are an oriented frame. Given any oriented frame at a point y in Y, these conditions mean that we can choose a basis for the fiber S_y such that the matrices of the linear transformations $\rho(e_i)$ are the three Pauli matrices σ_i:

$$\sigma_1 = \begin{bmatrix} i & 0 \\ 0 & -i \end{bmatrix}, \quad \sigma_2 = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \quad \sigma_3 = \begin{bmatrix} 0 & i \\ i & 0 \end{bmatrix}. \tag{1.1}$$

The action of ρ is extended to cotangent vectors using the metric, and then to forms using the rule

$$\rho(\alpha \wedge \beta) = \frac{1}{2}\big(\rho(\alpha)\rho(\beta) + (-1)^{\deg(\alpha)\deg(\beta)}\rho(\beta)\rho(\alpha)\big).$$

We also extend ρ to complex forms, so that it gives, for example, an isomorphism

$$\rho : T^*Y \otimes \mathbb{C} \to \mathfrak{sl}(S).$$

Our orientation convention means that $\rho(*\alpha) = -\rho(\alpha)$ for 1-forms α.

Because the tangent bundle of an oriented 3-manifold is always trivial, a spin^c structure always exists: we can simply take S to be the product bundle $\mathbb{C}^2 \times Y$ and then define Clifford multiplication globally by the matrices (1.1), using any trivialization of TY. To understand the classification of spin^c structures in general, the important observation is that if we are given one spin^c structure, say (S_0, ρ_0) on Y, then we can construct a new spin^c structure (S, ρ) as follows. Choose any hermitian line bundle $L \to Y$, and define

$$\begin{aligned} S &= S_0 \otimes L \\ \rho(e) &= \rho_0(e) \otimes 1_L. \end{aligned} \tag{1.2}$$

The following proposition tells us that any (S, ρ) can be obtained from (S_0, ρ_0) in this way, for a uniquely determined L, up to isomorphism.

Proposition 1.1.1. *Given a single spin^c structure (S_0, ρ_0), the construction* (1.2) *establishes a one-to-one correspondence between:*

(i) *the isomorphism classes of spin^c structures (S, ρ) on Y; and*
(ii) *the isomorphism classes of complex line bundles $L \to Y$.*

Because line bundles L are classified by their first Chern class $c_1(L) \in H^2(Y; \mathbb{Z})$, we can equivalently replace (ii) *here by:*

(iii) *the elements of $H^2(Y; \mathbb{Z})$.*

Proof. Let us show that any (S, ρ) can be obtained from (S_0, ρ_0) by tensoring with a suitable line bundle.

Given spin^c structures (S', ρ') and (S, ρ) on Y, we can define a vector bundle L on Y as the subbundle of $\mathrm{Hom}(S', S)$ consisting of homomorphisms that intertwine ρ' and ρ. This L has rank 1 (it is a line bundle): this is a manifestation of Schur's lemma and reflects the fact that only the scalar endomorphisms of S commute with the image of $\rho : TX \to \mathrm{End}(S)$. We call L the difference line bundle. If the difference line bundle is trivial, then a global section of unit length provides an isomorphism between the spin^c structures.

To apply this construction, let (S, ρ) be a spin^c structure and consider the difference line bundle L between (S_0, ρ_0) and (S, ρ). Set $S' = S_0 \otimes L$, and let ρ' be the Clifford multiplication $\rho \otimes 1_L$. Then the difference line bundle between S' and S is the trivial bundle $L^{-1} \otimes L$. So (S', ρ') and (S, ρ) are isomorphic spin^c structures. ☐

We will usually use $\mathfrak{s}$ to denote a typical spin^c structure (S, ρ). If $\mathfrak{s}_0$ is a chosen spin^c structure and L has first Chern class $l \in H^2(Y; \mathbb{Z})$, then we write

$$\mathfrak{s} = \mathfrak{s}_0 + l$$

for the spin^c structure defined by (1.2). The way we have defined it, a spin^c structure depends on a prior choice of Riemannian metric. However, if g_0 and

g_1 are two metrics on Y, and $\mathfrak{s}_0$, $\mathfrak{s}_1$ are corresponding spin^c structures, we can still compare the two: we can ask if there is a path g_t in the (contractible) space of metrics, joining g_0 to g_1, and a corresponding family (S_t, ρ_t), forming a continuous family over $[0, 1]$. We can therefore think of the set of isomorphism classes of spin^c structures as being associated to a smooth oriented manifold Y.

On an oriented 4-dimensional Riemannian manifold X, a spin^c structure again provides a hermitian vector bundle $S_X \to X$, this time of rank 4, with a Clifford multiplication

$$\rho : TX \to \mathrm{Hom}(S_X, S_X),$$

such that at each $x \in X$ we can find an oriented orthonormal frame $e_0, \dots, e_3$ with

$$\rho(e_0) = \begin{bmatrix} 0 & -I_2 \\ I_2 & 0 \end{bmatrix}, \quad \rho(e_i) = \begin{bmatrix} 0 & -\sigma_i^* \\ \sigma_i & 0 \end{bmatrix} \quad (i = 1, 2, 3) \tag{1.3}$$

in some orthonormal basis of the fiber S_x. Here I_2 is the 2-by-2 identity matrix and σ_i is as above. If we extend Clifford multiplication to (complex) forms as before, then in the same basis for S_x we have

$$\rho(\mathrm{vol}_x) = \begin{bmatrix} -I_2 & 0 \\ 0 & I_2 \end{bmatrix}$$

where $\mathrm{vol} = e_0 \wedge e_1 \wedge e_2 \wedge e_3$ is the oriented volume form. So the eigenspaces of $\rho(\mathrm{vol})$ give a decomposition of S_X into two orthogonal rank-2 bundles. We define S^+ to be the -1 eigenspace, and write

$$S_X = S^+ \oplus S^-.$$

Clifford multiplication by a tangent vector is an *odd* linear transformation: it interchanges the two summands, and we can write

$$\rho(e) : S^+ \to S^-.$$

If ν is a 2-form, then $\rho(\nu)$ preserves the two summands. In dimension 4, the bundle of 2-forms $\Lambda^2 X$ decomposes as a sum of the self-dual and anti-self-dual forms,

$$\Lambda^2 X = \Lambda^+ \oplus \Lambda^-,$$

the $+1$ and -1 eigenbundles of the Hodge $*$ operator. A short calculation with the matrices above shows that, if $\nu \in \Lambda^+$, then $\rho(\nu)$ restricts to zero on S^-, and vice versa. We have maps

$$\begin{aligned} \rho &: \Lambda^+ \to \mathfrak{su}(S^+) \\ \rho &: \Lambda^- \to \mathfrak{su}(S^-) \end{aligned} \tag{1.4}$$

which are bundle isometries. For $e \in T_x X$ a unit vector, the determinant of $\rho(e) : S^+_x \to S^-_x$ is a map

$$\det \rho(e) : \Lambda^2 S^+_x \to \Lambda^2 S^-_x$$

that is independent of e. So the complex line bundles $\Lambda^2 S^+$ and $\Lambda^2 S^-$ are canonically identified.

Proposition 1.1.1 continues to hold in dimension 4, but the existence of at least one spin^c structure is a slightly more subtle question than in dimension 3. The isomorphisms (1.4) mean that $w_2(\Lambda^+)$ is equal to the mod 2 reduction of $c_1(S^+)$; so the existence of a spin^c structure implies the existence of an integral lift of $w_2(\Lambda^+)$, or equivalently of $w_2(X)$ since these two are equal. This condition is also sufficient: the existence of a spin^c structure is equivalent to the existence of an integral lift of $w_2(X)$. On an orientable 4-manifold, $w_2(X)$ always has an integral lift, see [52], so spin^c structures always exist.

In any dimension, an *automorphism* of a spin^c structure (S, ρ) means a unitary bundle automorphism of S which commutes with Clifford multiplication. The group of automorphisms can be identified with the group of $\mathcal{G}$ of S^1-valued functions $u : X \to S^1$, acting by scalar multiplication. We call $\mathcal{G}$ the *gauge group* and we call its elements *gauge transformations*. The gauge group acts on sections Φ of S by

$$\Phi \mapsto u\Phi.$$

1.2 Dirac operators

Let $\mathfrak{s} = (S_X, \rho)$ be a spin^c structure on an oriented Riemannian 4-manifold X. A unitary connection A on S_X is a *spin^c connection* if ρ is parallel. This implies, in particular, that parallel transport preserves the decomposition of S_X as $S^+ \oplus S^-$. Given such a connection A, one defines the Dirac operator $D_A : \Gamma(S_X) \to \Gamma(S_X)$ as the composite

$$\Gamma(S_X) \xrightarrow{\nabla_A} \Gamma(T^*X \otimes S_X) \longrightarrow \Gamma(S_X),$$

in which the second map is constructed from the Clifford multiplication. The difference between two spin^c connections A and $\tilde{A}$, regarded as a 1-form with values in the endomorphisms of S_X, has the form

$$\tilde{A} - A = a \otimes 1_{S_X} \tag{1.5}$$

for some $a \in \Omega^1(X; i\mathbb{R})$. Conversely, if A is a spin^c connection and $a \in \Omega^1(X; i\mathbb{R})$, then $\tilde{A} = A + a \otimes 1_{S_X}$ is a spin^c connection. In this way, the spin^c connections on S_X form an affine space, with underlying vector space $\Omega^1(X; i\mathbb{R})$. If $\tilde{A}$ and A are related as above, then the corresponding Dirac operators are related by

$$D_{\tilde{A}} - D_A = \rho(a).$$

Because Clifford multiplication by 1-forms interchanges S^+ and S^-, we can write

$$D_A = D_A^+ + D_A^-,$$

where

$$\begin{aligned} D_A^+ &: \Gamma(S^+) \to \Gamma(S^-) \\ D_A^- &: \Gamma(S^-) \to \Gamma(S^+). \end{aligned}$$

If we are given a spin^c connection A, then the associated line bundles $\Lambda^2 S^+$, $\Lambda^2 S^-$ inherit connections too. The canonical isomorphism between these line bundles respects the connections. We give this connection a name:

Notation 1.2.1. If A is a spin^c connection on the spin bundle $S_X = S^+ \oplus S^-$ on X, we write A^{t} for the associated connection in the line bundle $\Lambda^2 S^+ = \Lambda^2 S^-$. So if $\tilde{A}$ and A are related by (1.5), then

$$\tilde{A}^{\mathrm{t}} = A^{\mathrm{t}} + 2a.$$

◊

In dimension 3, we define a spin^c connection B for the spin^c bundle $S \to Y$ in the same way. The spin^c connections are again an affine space, now over $\Omega^1(Y; i\mathbb{R})$, for we can write

$$\tilde{B} = B + b \otimes 1_S, \tag{1.6}$$

just as in (1.5). For each spin^c connection B, we have a Dirac operator

$$D_B : \Gamma(S) \to \Gamma(S).$$

We write B^{t} for the associated connection on the line bundle $\Lambda^2 S$. There is no decomposition of this operator as there is in dimension 4.

In any dimension, the full Dirac operator is *self-adjoint*. In dimension 4, this means that D_A^- is the adjoint of D_A^+. The Dirac operator is elliptic, so if the underlying manifold is compact, then the operator is *Fredholm*: it has finite-dimensional kernel and cokernel. In dimension 3, because it is self-adjoint, the Dirac operator has index zero. On a compact 4-manifold, the complex index of the operator D_A^+ (the difference in the complex dimensions of the kernel and cokernel) is given by the Atiyah–Singer index theorem,

$$\text{index } D_A^+ = \frac{1}{8}\big(c_1(S^+)^2[X] - \sigma(X)\big), \tag{1.7}$$

where $\sigma(X)$ is the signature of X. (We write $\alpha[X]$, typically, for the evaluation of a cohomology class α on the fundamental class.)

The gauge group $\mathcal{G}$ acts on the space of spin^c connections A on S, by pull-back. If $u : X \to S^1 \subset \mathbb{C}$ is a gauge transformation, we write the action as

$$\begin{aligned} A &\mapsto u(A) \\ &= A - u^{-1}du. \end{aligned} \tag{1.8}$$

1.3 The Seiberg–Witten equations

On an oriented Riemannian 4-manifold X with spin^c structure $\mathfrak{s}_X$, the Seiberg–Witten equations, or monopole equations, are equations for a pair (A, Φ) consisting of a spin^c connection A and a section Φ of the associated spin bundle S^+. The equations are the following:

$$\begin{aligned} \frac{1}{2}\rho(F^+_{A^{\mathrm{t}}}) - (\Phi\Phi^*)_0 &= 0 \\ D_A^+\Phi &= 0. \end{aligned} \tag{1.9}$$

Here $F^+_{A^{\mathrm{t}}}$ is the self-dual part of the curvature 2-form $F_{A^{\mathrm{t}}}$ of the connection A^{t},

$$\begin{aligned} F_{A^{\mathrm{t}}} &= F^+_{A^{\mathrm{t}}} + F^-_{A^{\mathrm{t}}} \\ &\in \Omega^+(X; i\mathbb{R}) \oplus \Omega^-(X; i\mathbb{R}), \end{aligned}$$

and $(\Phi\Phi^*)_0$ denotes the trace-free part of the hermitian endomorphism $\Phi\Phi^*$ of the bundle S^+,

$$\begin{aligned}(\Phi\Phi^*)_0 &= \Phi\Phi^* - \frac{1}{2}\operatorname{tr}(\Phi\Phi^*)1_{S^+} \\ &= \Phi\Phi^* - \frac{1}{2}|\Phi|^2 1_{S^+}.\end{aligned}$$

Note that F_{A^t} is an imaginary-valued 2-form, so $\rho(F_{A^t})$ is hermitian: the map ρ in (1.4) carries real self-dual forms to skew-adjoint endomorphisms of S^+.

If ω is a smooth imaginary-valued 2-form and ω^+ its self-dual part, we can also consider the monopole equations *perturbed by* ω. These are the equations

$$\begin{aligned}\frac{1}{2}\rho(F^+_{A^t} - 4\omega^+) - (\Phi\Phi^*)_0 &= 0 \\ D^+_A\Phi &= 0.\end{aligned} \tag{1.10}$$

The left-hand sides of the two equations in (1.9) define a map

$$\mathfrak{F} : \mathcal{A} \times \Gamma(S^+) \to \Gamma(i\,\mathfrak{su}(S^+) \oplus S^-), \tag{1.11}$$

where $\mathcal{A}$ denotes the affine space of all spin^c connections A, and $i\,\mathfrak{su}(S^+)$ is the bundle of hermitian endomorphisms of S^+. We can then write the monopole equations as $\mathfrak{F}(A, \Phi) = 0$. We write the perturbed equations similarly, as

$$\mathfrak{F}_\omega(A, \Phi) = 0. \tag{1.12}$$

The set of solutions (A, Φ) of the perturbed equations is invariant under the action of the gauge group $\mathcal{G}$. We will write $[A, \Phi]$ to denote the gauge-equivalence class of a pair (A, Φ): the orbit of (A, Φ) under the action of $\mathcal{G}$.

Definition 1.3.1. If X is an oriented Riemannian 4-manifold with spin^c structure $s_X = (S_X, \rho)$, and ω is an imaginary-valued 2-form, we write $N(X, \mathfrak{s}_X)$ for the quotient space of the set of solutions of the equations (1.12) by the action of $\mathcal{G}$:

$$N(X, \mathfrak{s}_X) = \{\, [A, \Phi] \mid \mathfrak{F}_\omega(A, \Phi) = 0 \,\}.$$

This is the *monopole moduli space* for $(X, \mathfrak{s}_X)$ with perturbing 2-form ω. It is a subset of the *configuration space*

$$\mathcal{B}(X, \mathfrak{s}_X) = \big(\mathcal{A} \times \Gamma(S^+)\big) \,/\, \mathcal{G}.$$

◊

The configuration space $\mathcal{B}(X, \mathfrak{s}_X)$ is Hausdorff, so the moduli space is Hausdorff also. The following result reflects the very special nature of the monopole equations. (It would not be true, for example, if the sign of the second term in (1.9) were changed.)

Theorem 1.3.2. *If the 4-manifold X is compact (without boundary), then the moduli space $N(X, \mathfrak{s}_X) \subset \mathcal{B}(X, \mathfrak{s}_X)$ is compact.* □

1.4 Regularity

From this point on, we will always assume that our 4-manifold X is *connected*. Let (A, Φ) be a solution of the equations $\mathfrak{F}_\omega(A, \Phi) = 0$ on X, as above. We can take the derivative of the map

$$\mathfrak{F}_\omega : \mathcal{A} \times \Gamma(S^+) \to i\,\mathfrak{su}(S^+) \oplus \Gamma(S^-),$$

at the point (A, Φ) in the affine space $\mathcal{A} \times \Gamma(S^+)$, to obtain a linear map

$$\mathcal{D}_{(A,\Phi)}\mathfrak{F}_\omega : \Omega^1(X; i\mathbb{R}) \times \Gamma(S^+) \to \Gamma(i\,\mathfrak{su}(S^+) \oplus S^-),$$

given by

$$(a, \phi) \mapsto \left(\frac{1}{2}\rho(d^+a) - (\phi\Phi^* + \Phi\phi^*)_0, D_A^+\phi + \rho(a)\Phi\right). \tag{1.13}$$

Definition 1.4.1. A solution (A, Φ) to the perturbed monopole equations $\mathfrak{F}_\omega = 0$ is *regular* if the linearization (1.13) is a *surjective* linear operator. We say that the moduli space $N(X, \mathfrak{s}_X)$ is regular if all solutions are regular. ◊

Proposition 1.4.2. *Suppose that the oriented Riemannian 4-manifold X is compact (without boundary), and let $\mathfrak{s}_X$ be a given* spin^c *structure. Then there is an open and dense subset of the space of imaginary-valued 2-forms ω for which the corresponding moduli space $N(X, \mathfrak{s}_X)$ is regular.* □

The action of the gauge group $\mathcal{G}$ on $\mathcal{A} \times \Gamma(S^+)$ is free on the set of pairs (A, Φ) with Φ non-zero. We call such a pair *irreducible*. For a *reducible* configuration $(A, 0)$, the equations (1.9) reduce to the equations

$$F_{A^t}^+ = 4\omega^+.$$

Suppose κ is a 2-form on X that is both closed and self-dual. Then if A satisfies the equation above, we have

$$\begin{aligned}\int_X \omega \wedge \kappa &= \int_X \omega^+ \wedge \kappa \\ &= \int_X F_{A^t} \wedge \kappa \\ &= (2\pi/i)\bigl(c_1(S^+) \smile [\kappa]\bigr)[X]. \end{aligned} \tag{1.14}$$

If κ is non-zero, this is a non-trivial linear constraint on ω, which must be satisfied if a reducible solution is to exist. The closed self-dual (real) 2-forms κ form a subspace $\mathcal{H}^+$ of the space $\mathcal{H}^2$ of harmonic 2-forms, and determine a metric-dependent subspace

$$\mathcal{H}^+ \subset H^2(X;\mathbb{R}).$$

This is a maximal positive-definite subspace for the quadratic form

$$\begin{aligned} Q : H^2(X;\mathbb{R}) &\to \mathbb{R} \\ Q(\alpha) &= \alpha^2[X]. \end{aligned}$$

We write b^+ for the dimension of $\mathcal{H}^+$. Defining b^- similarly, we have $b^+ + b^- = b^2$ and $b^+ - b^- = \sigma(X)$. From the calculation (1.14), we deduce:

Lemma 1.4.3. *If X is a compact manifold with $b^+ \geq 1$, then for all ω in the complement of a proper linear subspace, the corresponding moduli space $N(X, \mathfrak{s}_X)$ contains no reducible solutions.* □

When the moduli space is regular and contains no reducibles, it is a smooth manifold whose dimension can be computed as the index of a certain operator (essentially the sum of the two operators in (27.3); see Lemma 27.1.1):

Theorem 1.4.4. *Let X be a closed, connected, oriented Riemannian manifold with a spin^c structure $\mathfrak{s}_X$. Suppose ω is chosen so that the moduli space $N(X, \mathfrak{s}_X)$ is regular and contains no reducible solutions, as we can always do when $b^+(X) \geq 1$. Then the moduli space $N(X, \mathfrak{s}_X)$ is a smooth, compact manifold, whose dimension d is given by the formula*

$$\begin{aligned} d &= \bigl(b^1(X) - b^+(X) - 1\bigr) + 2\,\mathrm{index}\, D_A^+ \\ &= \frac{1}{4}\bigl(c_1(S^+)^2[X] - 2\chi(X) - 3\sigma(X)\bigr). \end{aligned} \tag{1.15}$$

□