From Morse homology to Floer homology
Informal lecture notes for MATH 621

Ely Kerman

Contents

Chapter 1

The Morse Complex

The goal of this chapter is to introduce the Morse complex and to give a (nearly) complete account of the analysis required to rigorously define it. Effort will be made to at least acknowledge those details not covered here.

1.1 The data

Let M^n be a smooth compact orientable manifold. Let $f \in C^\infty(M)$ be a smooth Morse function and let g be a smooth Riemannian metric on M.

The pair (f, g) determines a unique gradient vector field $\nabla_g f$ on M via the equation

$$g(\nabla_g f, \cdot) = df(\cdot).$$

We will be interested in the flow of the negative gradient vector field $-\nabla_g f$ which we denote by ψ_s. This flow is defined for all $s \in \mathbb{R}$ and satisfies $\psi_s \circ \psi_{s'} = \psi_{s+s'}$.

By the classical Morse Lemma the set of critical points of f, $\mathrm{Crit}(f)$, is finite. We can associate to each $p \in \mathrm{Crit}(f)$ it's Hessian. This is traditionally defined to be the bilinear map

$$\begin{aligned} \mathcal{H}_f(p)\colon T_pM \times T_pM &\rightarrow \mathbb{R} \\ (V, W) &\mapsto L_{\tilde{V}}(L_{\tilde{W}} f)(p), \end{aligned}$$

where $\tilde{V}$ and $\tilde{W}$ are local extensions of V and W.

Alternatively, given a metric g one can choose an affine connection $\widetilde{\nabla}$ on TM and define the Hessian of f at $p \in \mathrm{Crit}(f)$ to be the self-adjoint linear map

$$\begin{aligned} H_f(p)\colon T_pM &\rightarrow T_pM \\ W &\mapsto \widetilde{\nabla}_W(\nabla_g f). \end{aligned}$$

The **Morse index** of $p \in \mathrm{Crit}(f)$ is defined to be dimension of the negative eigenspace of the Hessian of f at p. This integer is denoted by $\mu(p)$.

Exercise 1.1.1. *Show that $\mu(p)$ does not depend on which definition of the Hessian we use or on the extra choices made in each definition.*

1.2 The moduli spaces

Given a pair of critical points $p, q \in \mathrm{Crit}(f)$ we consider the set of integral curves of $-\nabla_g f$ which converge to p and q in forward and backward time, respectively.

$$\mathcal{M}(p,q) = \{u \colon \mathbb{R} \to M \mid \frac{du}{ds} = -\nabla_g f(u), \lim_{s\to-\infty} u(s) = p, \lim_{s\to+\infty} u(s) = q\}. \tag{1.1}$$

This set comes with a free $\mathbb{R}$-action which is defined by

$$(\tau \cdot u)(s) = \psi_\tau \circ u(s) = u(s+\tau).$$

The quotient

$$\widehat{\mathcal{M}}(p,q) = \mathcal{M}(p,q)/\mathbb{R}. \tag{1.2}$$

is the space of trajectories of $-\nabla_g f$ from p to q.

Theorem 1.2.1 (The main analytic theorem of Morse homology). *For generic data pairs (f,g), each moduli space $\widehat{\mathcal{M}}(p,q)$ has the following properties:*

1. *it is a smooth orientable manifold of dimension $\mu(p) - \mu(q) - 1$.*
2. *it has a natural compactification as a smooth manifold with corners $\overline{\widehat{\mathcal{M}}(p,q)}$ whose stratum of codimension k is*

$$\overline{\widehat{\mathcal{M}}(p,q)}_k = \bigcup_{\substack{r_1,\ldots,r_k \in \mathrm{Crit}(f) \\ p,r_1,\ldots,r_k,q \text{ distinct}}} \widehat{\mathcal{M}}(p,r_1) \times \widehat{\mathcal{M}}(r_1,r_2) \times \cdots \times \widehat{\mathcal{M}}(r_k,q).$$

Explanation of terms

A property defined for elements of a topological space X is said to be **generic** if it is satisfied by a subset of objects in X which contains a countable intersection of open dense sets.

A **manifold with corners** is a second countable Hausdorff space such that each point has a neighborhood which is homeomorphic to $\mathbb{R}^{n-k} \times [0,\infty)^k$ for some k, and the transition maps are smooth.

Comments on the proof

It will take a lot of effort for us to prove this result and we will do this in parts over the next several sections. The first statement of the theorem (modulo issues of orientation) will be proved as the *Transversality theorem.* The *Compactness theorem* will show that the "broken trajectories" forming the boundary strata are the correct objects to add in order to compactify $\widehat{\mathcal{M}}(p,q)$. The *Gluing theorem* will prove that every possible broken trajectory must appear in the compactification and that near each broken trajectory the compactified space has the structure of a manifold with corners. These are both important points because we will be counting boundary elements (see below). Matters of orientation are also nontrivial and will be discussed last.

Immediate consequences of the theorem

1. if $\mu(p) - \mu(q) \le 0$, then $\widehat{\mathcal{M}}(p,q) = \emptyset$
2. if $\mu(p) - \mu(q) = 1$, then $\widehat{\mathcal{M}}(p,q)$ is a compact zero-dimensional manifold. Hence, we can count the elements of $\widehat{\mathcal{M}}(p,q)$. We can also associate signs to these elements by choosing certain orientations. The total sum of the elements with sign will be denoted by $\#\widehat{\mathcal{M}}(p,q)$.
3. if $\mu(p) - \mu(q) = 2$, then $\overline{\widehat{\mathcal{M}}(p,q)}$ is a compact one-dimensional manifold with boundary
$$\partial(\overline{\widehat{\mathcal{M}}(p,q)}) = \bigcup_{\substack{r\in \mathrm{Crit}(f)\\ \mu(r)=\mu(q)+1}} \widehat{\mathcal{M}}(p,r) \times \widehat{\mathcal{M}}(r,q).$$
Since this is the boundary of a compact one dimensional manifold, if we orient $\overline{\widehat{\mathcal{M}}(p,q)}$ and count the boundary elements with the appropriate signs, then we should get zero, i.e.,
$$\#\partial(\overline{\widehat{\mathcal{M}}(p,q)}) = \sum_{\substack{r\in \mathrm{Crit}(f)\\ \mu(r)=\mu(q)+1}} \#\widehat{\mathcal{M}}(p,r) \cdot \#\widehat{\mathcal{M}}(r,q) = 0.$$

1.3 The definition of the complex

Let (f, g) be a generic data pair in the sense of Theorem 1.2.1 and let $\mathrm{Crit}_k(f)$ denote the set of critical points of f with index equal to k.

The chain group in degree k, $C_k(f)$, is defined to be the $\mathbb{Z}$-vector space generated by the elements in $\mathrm{Crit}_k(f)$.

The boundary map $\partial_g \colon C_k(f) \to C_{k-1}(f)$ counts negative gradient flow lines. It is defined on the basis elements $p \in \mathrm{Crit}_k(f)$ by the formula

$$\partial_g(p) = \sum_{q \in \mathrm{Crit}_{k-1}(f)} \#\widehat{\mathcal{M}}(p, q) \cdot q.$$

Lemma 1.3.1. $\partial_g \circ \partial_g = 0$

Proof. For every $p \in \mathrm{Crit}_k(f)$

$$\begin{aligned}
\partial_g \circ \partial_g(p) &= \partial_g \left(\sum_{r \in \mathrm{Crit}_{k-1}(f)} \#\widehat{\mathcal{M}}(p, r) \cdot r \right) \\
&= \sum_{r \in \mathrm{Crit}_{k-1}(f)} \#\widehat{\mathcal{M}}(p, r) \left(\sum_{q \in \mathrm{Crit}_{k-2}(f)} \#\widehat{\mathcal{M}}(r, q) \cdot q \right) \\
&= \sum_{q \in \mathrm{Crit}_{k-2}(f)} \left(\sum_{r \in \mathrm{Crit}_{k-1}(f)} \#\widehat{\mathcal{M}}(p, r) \cdot \#\widehat{\mathcal{M}}(r, q) \right) q \\
&= 0
\end{aligned}$$

□

The homology of the complex for the pair (f, g) is called the Morse homology,

$$HM_*(f, g) = H_*(C(f), \partial_g) = \frac{\ker \partial_g \colon C_*(f) \to C_{*-1}(f)}{\mathrm{im}\, \partial_g \colon C_{*+1}(f) \to C_*(f)}.$$

It is defined for generic pairs (f, g). We will prove that the Morse homology is independent of the generic pair (f, g) and equal to the singular homology of M with coefficients in $\mathbb{Z}$.

1.4 The transversality theorem

The property of being a Morse function is generic in $C^k(M)$ for $k \geq 2$ (see, for example [4], Prop. 5.5). In fact, as described in § 1.4.10, it is a generic property in $C^\infty(M)$. Using this fact we will fix a smooth Morse function f and two critical points $p, q \in \mathrm{Crit}(f)$.

The following result implies the first part of Theorem 1.2.1.

Theorem 1.4.1. *(Transversality) For a generic C^k-smooth metric g on M $(k \geq 2)$, the moduli space $\mathcal{M}(p, q)$ is a smooth manifold of dimension $\mu(p) - \mu(q)$.*

One can describe the moduli space $\mathcal{M}(p, q)$ as a potential manifold in two ways.

1.4.1 The geometric (dynamical) picture

For each $p \in \mathrm{Crit}(f)$ we can define the **descending manifold**

$$\mathcal{D}(p) = \{x \in M \mid \lim_{s \to -\infty} \psi_s(x) = p\}$$

and the **ascending manifold**

$$\mathcal{A}(p) = \{x \in M \mid \lim_{s \to \infty} \psi_s(x) = p\}.$$

Theorem 1.4.2. *If $p \in \mathrm{Crit}(f)$ is nondegenerate, then $\mathcal{D}(p)$ and $\mathcal{A}(p)$ are smooth submanifolds of M. More precisely, $\mathcal{D}(p)$ is an embedded open disc of dimension $\mu(p)$ and $\mathcal{A}(p)$ is an embedded open disc of dimension $n - \mu(p)$*

For a proof of this see [5] §6.3. The smoothness condition is stronger than one might hope to prove using the classical Hartman-Grobman theorem (which states that the flow of a vector field near a hyperbolic (nondegenerate) fixed point is **topologically** equivalent to the linearized flow.

It is clear from the definitions that

$$\mathcal{M}(p, q) = \mathcal{D}(p) \cap \mathcal{A}(p).$$

In particular, each gradient trajectory $u \in \mathcal{M}(p, q)$ can be identified with its unique initial value $u(0)$ which must lie in $\mathcal{D}(p) \cap \mathcal{A}(p)$. The following result is equivalent to Theorem 1.4.1.

Theorem 1.4.3. *(Transversality Version A) For a generic C^k-smooth metric g on M, the submanifolds $\mathcal{D}(p)$ and $\mathcal{A}(p)$ intersect transversally, i.e.,*

$$\mathcal{D}(p) \pitchfork \mathcal{A}(p).$$

Remark 1.4.4. *For this description of $\mathcal{M}(p,q)$ we require the flow of the negative gradient vector field to be globally defined. As we will see later, in the infinite dimensional setting of Floer theory the negative gradient flow is not even locally defined.*

1.4.2 The functional analytic picture

Right idea, wrong setting

First we consider a naive picture. Assume that g is a smooth metric and let $\mathcal{B}$ be the set of smooth maps $v\colon \mathbb{R} \to M$ such that

$$\lim_{s\to-\infty} v(s) = p \quad \text{and} \quad \lim_{s\to-\infty} v(s) = q.$$

Let $\mathcal{E} \to \mathcal{B}$ be the bundle over $\mathcal{B}$ whose fibre over $v \in \mathcal{B}$ is $\mathcal{E}_v = \Gamma^\infty(v^*TM)$, the space of smooth sections of v^*TM.

Now consider the section of the bundle $F_g\colon \mathcal{B} \to \mathcal{E}$ defined by

$$F_g(v) = (v, \partial_s v - \nabla_g f(v(s))).$$

If we let $S_{\mathcal{E}}$ denote the zero-section of $\mathcal{E}$ then

$$\mathcal{M}(p,q) = F_g^{-1}(S_{\mathcal{E}}).$$

At this point one would like to prove a statement like "for generic metrics g we have $F_g \pitchfork S_{\mathcal{E}}$." Unfortunately, the spaces $\mathcal{B}$ and $\mathcal{E}$ are infinite dimensional Fréchet manifolds. For these manifolds the inverse function theorem is, at best, extremely complicated. Moreover, there is no version of Sard's theorem with which to establish the genericity property. To overcome these difficulties we must make some better choices.

The right setting

We begin by choosing a background metric $\tilde{g}$ on M which we will use to make our measurements. Now, we let $\mathcal{B}$ be the subset of $v \in L^2_{1,loc}(\mathbb{R}, M)$ such that

1. $\lim_{s\to-\infty} = p$, and for any $R << 0$ which is sufficiently negative for $v(-\infty, R]$ to be contained in a coordinate chart around p, we have $v|_{(-\infty,R]} \in L^2_1$.

2. $\lim_{s\to\infty} = q$, and for any $R >> 0$ which is sufficiently positive for $v[R,\infty]$ to be contained in a coordinate chart around q , we have $v|_{[R,\infty)} \in L^2_1$.

By the Sobolev embedding theorem $L^2_{1,loc}(\mathbb{R}, M) \subset C^0(\mathbb{R}, M)$, so these conditions make sense. The space $\mathcal{B}$ is a smooth Banach manifold modeled on $L^2_1(\mathbb{R}, \mathbb{R}^n)$.

Let $\mathcal{E} \to \mathcal{B}$ be the Banach space bundle whose fibre over $v \in \mathcal{B}$ is $\mathcal{E}_v = L^2(v^*TM)$, the space of L^2-sections of $v^*(TM)$. We can then define the section $F_g : \mathcal{B} \to \mathcal{E}$ by

$$F_g(v) = (v, v' - \nabla_g f(v(s))),$$

where v' is the weak derivative of v.

Lemma 1.4.5. $F_g^{-1}(S_\mathcal{E}) = \mathcal{M}(p, q)$.

Proof. Since the integral curves of the negative gradient flow are C^k and converge exponentially to their end points, p and q, it follows that $\mathcal{M}(p, q) \subset F_g^{-1}(S_\mathcal{E})$.

Now suppose that $v \in F_g^{-1}(S_\mathcal{E})$, i.e., v is in $\mathcal{B}$ and

$$v' = \nabla_g f(v(s)).$$

As a map from M to TM, $\nabla_g f$ is as smooth as g, i.e., C^k-smooth. The right-hand side of the equation above is then continuous since v is continuous. Hence, v' is continuous and equal to the usual derivative $\partial_s v$. This means that v is actually C^1 and hence an integral curve of the negative gradient flow running from p to q.

(In fact, one can continue this line of argument to show that v is C^k. This is a simple example of elliptic regularity theory, a.k.a. *boot-strapping.)* □

Theorem 1.4.6. *(Transversality Version B) For a generic C^k-smooth metric g on M, the section $F_g : \mathcal{B} \to \mathcal{E}$ is transversal to the zero-section $S_\mathcal{E}$ and the inverse image $F_g^{-1}(S_\mathcal{E})$ is smooth submanifold of dimension $\mu(p) - \mu(q)$.*

We will prove this version of the transversality theorem and show that it is equivalent to Version A.

1.4.3 Outline of the proof

Let $\mathcal{A}^k$ be the Banach manifold of C^k-smooth metrics on M. Let $\mathcal{E}^* \to \mathcal{A}^k \times \mathcal{B}$ be the pullback of the bundle $\mathcal{E}$ to $\mathcal{A}^k \times \mathcal{B}$ via projection to $\mathcal{B}$. In particular, the fibre $\mathcal{E}^*_{(g,v)}$ is just $\mathcal{E}_v = L^2(v^*TM)$.

We extend F_g to a section $F : \mathcal{A}^k \times \mathcal{B} \to \mathcal{E}^*$ as follows

$$F(g, v) = (g, v, v' - \nabla_g f(v(s))).$$

Proposition 1.4.7. $F \pitchfork S_{\mathcal{E}^*}$.

By the standard inverse function theorem for Banach spaces this implies that $F^{-1}(S_{\mathcal{E}^*})$ is a smooth submanifold of $\mathcal{A}^k \times \mathcal{B}$. This is called the **universal moduli space**, and it consists of the trajectories from p to q for all the negative gradient flows of f with respect to the metrics in $\mathcal{A}^k$.

Proposition 1.4.8. *The map* $\pi \colon F^{-1}(S_{\mathcal{E}^*}) \to \mathcal{A}^k$ *is a Fredholm map of index* $\mu(p) - \mu(q)$.

Here, π is the restriction of the projection $\pi \colon \mathcal{A}^k \times \mathcal{B} \to \mathcal{A}^k$. Since the relevant Banach manifolds are separable and π is Fredholm by the proposition, we may invoke the Sard-Smale theorem. This implies that a generic $g \in \mathcal{A}^k$ is a regular value for π. The proposition also implies that the inverse image, $\pi^{-1}(g)$, of each regular metric $g \in \mathcal{A}^k$, is a smooth submanifold of dimension $\mu(p) - \mu(q)$. Since $\pi^{-1}(g) = \mathcal{M}(p,q)$, this yields Theorem 1.4.6.

1.4.4 The vertical differential

Before we embark on our planned path, let us first study the transversality issue for F_g and $S_{\mathcal{E}}$. In order for $F_g \pitchfork S_{\mathcal{E}}$ to hold, we must have

$$(dF_g)_v(T_v\mathcal{B}) + T_{(v,0)}S_{\mathcal{E}} = T_{(v,0)}\mathcal{E}$$

for all $v \in (F_g)^{-1}(S_{\mathcal{E}})$. This motivates us to consider the image of $(dF_g)_v$.

Note that for $v \in (F_g)^{-1}(S_{\mathcal{E}})$ we have the splitting

$$T_{(v,0)}\mathcal{E} = T_{(v,0)}S_{\mathcal{E}} \oplus \mathcal{E}_v.$$

Let $\pi_v \colon T_{(v,0)}\mathcal{E} \to \mathcal{E}_v$ be projection. Since F_g is a section, the relevant part of $(dF_g)_v$ is determined by the map

$$D_g^v \equiv \pi_v \circ (dF_g)_v \colon T_v\mathcal{B} \to \mathcal{E}_v,$$

which is called the **vertical differential** of F_g. In particular, if the vertical differential of F_g is onto $\mathcal{E}_v$, then the transversality condition is satisfied. (This is a completely trivial, but somewhat confusing statement with which you should make peace.)

Let us now derive an expression for the vertical differential D_g^v by finding a formula for F_g in a local trivialization of the bundle $\mathcal{E}$. Given $v \in \mathcal{B}$ we can identify a neighborhood of v with a neighborhood of $0 \in T_v\mathcal{B}$ by using the exponential map

$\exp\colon \mathcal{B}\times T_v\mathcal{B}\to\mathcal{B}$ for the background metric $\tilde{g}$. In particular, any path near v can be written uniquely in the form $s\mapsto \exp(v(s),\xi(s))$ for some $\xi\in T_v\mathcal{B}=\Gamma_1^2(v^*TM)$. In other words, the $\xi\in T_v\mathcal{B}$ near 0 are local coordinates for a neighborhood of v.

To trivialize the bundle $\mathcal{E}$ over this neighborhood we use parallel translation with respect to the Levi-Civita connection $\widetilde{\nabla}$ of $\tilde{g}$. More precisely, let $\phi_v(\xi)\colon v^*TM\to \exp(v,\xi)^*TM$ be the map defined by parallel translating each $T_{v(s)}M$ to $T_{\exp(v(s),\xi(s))}M$ along the geodesic path $t\to\exp(v(s),t\xi(s))$. This allows us to identify the fibre over ξ, $\Gamma_1^2(\exp(v,\xi)^*TM)$, with the fibre over 0, $\Gamma_1^2(v^*TM)$.

In this trivialization the map F_g has the following form

$$\begin{array}{rcl} F_g\colon \mathcal{B} & \to & \mathcal{E} \\ \xi & \mapsto & \big(\xi,(\phi_v(\xi))^{-1}(\exp(v,\xi)'+\nabla_g f(\exp(v,\xi))\big)\equiv(\xi,F_g^v(\xi)). \end{array}$$

In these coordinates

$$\begin{aligned} D_g^v(\xi) &= (dF_g^v)_0(\xi) \\ &= \frac{d}{dt}\left(F_g^v(t\xi)\right)|_{t=0} \\ &= \frac{d}{dt}\left((\phi_v(t\xi))^{-1}(\exp(v,t\xi)'+\nabla_g f(\exp(v,t\xi))\right)|_{t=0} \\ &= \widetilde{\nabla}_\xi(v'+\nabla_g f(v)) \\ &= \widetilde{\nabla}_{v'}\xi+\widetilde{\nabla}_\xi(\nabla_g f(v)). \end{aligned}$$

Note: the second last line follows from the definition of the Levi-Civita connection (see Do Carmo, chpt 2, exercise 2). The last line follows from the fact that the Levi-Civita connection is torsion free.

The proof of the Transversality theorem will depend heavily on the following result which we will prove later.

Theorem 1.4.9. *For every $v\in F_g^{-1}(S_\mathcal{E})$ the vertical differential*

$$D_g^v\colon T_v\mathcal{B}\to\mathcal{E}_v$$

is a Fredholm map with index equal to $\mu(p)-\mu(q)$.

1.4.5 Proof of Proposition 1.4.7

We want to prove that the section $F\colon \mathcal{A}^k\times\mathcal{B}\to\mathcal{E}^*$ is transversal to the zero-section $S_{\mathcal{E}^*}$. Using the local trivialization for $\mathcal{E}$ described in the previous section we can write F near $\mathcal{A}^k\times\{v\}$ as

$$F(g,\xi)=(g,\xi,F_g^v(\xi)).$$

Let us relabel the fibre component as $F^v(g,\xi)$ to emphasize the fact that it is now considered to be a function of ξ and g.

The vertical differential of F is the map

$$\begin{aligned}(dF^v)_{(g,0)}\colon T_{(g,v)}(\mathcal{A}^k \times \mathcal{B}) &\to \mathcal{E}^*_{(g,v)}\\ (\eta,\xi) &\mapsto D^v_g(\xi) + \partial_\eta(\nabla_g f(v)),\end{aligned}$$

where ∂_η refers to the derivative in the metric variable in the η-direction.

To prove the transversality statement of Proposition 1.4.7 it suffices to show that the vertical differential of F is onto for every $(g,v) \in F^{-1}(S_{\mathcal{E}^*})$. Here is an outline of how one establishes this.

1. Assume that $(dF^v)_{(g,0)}$ is not onto $\mathcal{E}^*_{(g,v)} = L^2(v^*TM)$.

2. Show that the image of $(dF^v)_{(g,0)}$ is closed. This follows from the (presently unproven) fact that D^v_g is Fredholm.

3. If the image of $(dF^v)_{(g,0)}$ is closed and not onto, then it follows from the Hahn-Banach theorem (and the Riesz representation theorem) that there is a nonzero $W \in L^2(v^*TM)$ which **annihilates** the image of $(dF^v)_{(g,0)}$ in the following sense

   $$\int_{\mathbb{R}} \langle (dF^v)_{(g,0)}(\eta,\xi), W\rangle \, ds = 0 \quad \text{for all} \quad (\eta,\xi) \in T_{(g,v)}(\mathcal{A}^k \times \mathcal{B}) \tag{1.3}$$

   We will obtain our contradiction by showing that, in fact, any $W \in L^2(v^*TM)$ which satisfies this equation must be $0 \in L^2(v^*TM)$.

4. Letting $\eta = 0$ in (1.3), we get

   $$\int_{\mathbb{R}} \langle (D^v_g)(\xi), W\rangle \, ds = 0 \quad \text{for all} \quad \xi \in T_v\mathcal{B}.$$

   This, by definition, means that W is a weak solution of $(D^v_g)^*W = 0$ where $(D^v_g)^*$ is the formal adjoint of the operator (D^v_g). We will see later that $(D^v_g)^*$ is an operator of the same type as (D^v_g) whose weak solutions are actually strong solutions. Hence, W is C^k-smooth.

5. Letting $\xi = 0$ in (1.3), we get

   $$\int_{\mathbb{R}} \langle \partial_\eta(\nabla_g f(v)), W\rangle \, ds = 0 \quad \text{for all} \quad \eta \in T_g\mathcal{A}^k. \tag{1.4}$$

Exercise 1.4.10. *Fix* $s_0 \in \mathbb{R}$ *and show that for any vector* V *in the fibre of* v^*TM *over* s_0*, there is an* $\eta \in T_g\mathcal{A}^k$ *such that* $\partial_\eta(\nabla_g f(v))(s_0) = V$*.*

By our assumption, W is nonzero so we can fix an s_0 such that $W(s_0) \neq 0$.

The previous exercise implies that we can choose an η such that $\partial_\eta(\nabla_g f(v))(s_0) = W(s_0)$. The integrand in (1.4) evaluated at s_0 is then positive.

Choosing a smooth bump function $\beta\colon \mathbb{R} \to \mathbb{R}$ supported near s_0, set $\tilde{\eta}(s) = \beta(s)\eta(s)$. Using the C^k-smoothness of W, it is easy to check that

$$\int_{\mathbb{R}} \langle \partial_{\tilde{\eta}}(\nabla_g f(v)), W \rangle \, ds > 0.$$

This contradicts (1.4), so $(dF^v)_{(g,0)}$ must be onto.

1.4.6 Proof of Proposition 1.4.8

We need to prove that the maps

$$d\pi_{(g,v)} \colon T_{(g,v)}F^{-1}(S_{\mathcal{E}^*}) \to T_g\mathcal{A}^k$$

are Fredholm of index $\mu(p) - \mu(q)$ for all $(g,v) \in F^{-1}(S_{\mathcal{E}^*})$.

We show here that this is a consequence of Theorem 1.4.9.

Claim 1 The kernel of $d\pi_{(g,v)}$ is isomorphic to the kernel of D_g^v, i.e.,

$$\ker(d\pi_{(g,v)} \colon T_{(g,v)}F^{-1}(S_{\mathcal{E}^*}) \to T_g\mathcal{A}^k) \simeq \ker(D_g^v \colon T_v\mathcal{B} \to \mathcal{E}_v)$$

Proof. We know that

$$T_{(g,v)}F^{-1}(S_{\mathcal{E}^*}) = (dF_{(g,v)})^{-1}(T_{(g,v)}S_{\mathcal{E}^*}).$$

In our local trivialization, this looks like

$$T_{(g,v)}F^{-1}(S_{\mathcal{E}^*}) = \{(\eta,\xi) \in T_{(g,v)}(\mathcal{A}^k \times \mathcal{B}) \mid D_g^v(\xi) + \partial_\eta(\nabla_g f(v)) = 0\}.$$

Since $d\pi_{(g,v)}(\eta,\xi) = \eta$, we have

$$\begin{aligned} \ker(d\pi_{(g,v)}) &= \{(0,\xi) \in T_g\mathcal{A}^k \times T_v\mathcal{B} \mid D_g^v(\xi) = 0\} \\ &\simeq \{\xi \in T_v\mathcal{B} \mid D_g^v(\xi) = 0\} \\ &= \ker(D_g^v). \end{aligned}$$

□

Claim 2 The cokernel of $d\pi_{(g,v)}$ is isomorphic to the cokernel of D_g^v,

$$\operatorname{coker}(d\pi \colon T_{(g,v)}F^{-1}(S_{\mathcal{E}^*}) \to T_g\mathcal{A}^k) \simeq \operatorname{coker}(D_g^v \colon T_v\mathcal{B} \to \mathcal{E}_v)$$

Proof. Consider the following inclusions (and identification)

$$\operatorname{coker}(d\pi) \subset T_g\mathcal{A}^k \subset T_{(g,v)}(\mathcal{A}^k \times \mathcal{B})$$

and

$$\operatorname{coker}(D_g^v) \subset \mathcal{E}_v \simeq \mathcal{E}^*_{(g,v)}.$$

Exercise 1.4.11. *Show that the vertical differential $(dF^v)_{(g,0)} \colon T_{(g,v)}(\mathcal{A}^k \times \mathcal{B}) \to \mathcal{E}^*_{(g,v)}$ induces an isomorphism of the cokernels.*

□

Claim 3

It remains to prove that if the image of D^v_g is closed, then so is the image of $d\pi_{(g,v)}$.

Proof.

$$\begin{aligned} d\pi_{(g,v)}(T_{(g,v)}F^{-1}(S_{\mathcal{E}^*})) &= \{\eta \in T_g\mathcal{A}^k \mid D^v_g(\xi) + \partial_\eta(\nabla_g f(v)) = 0\} \\ &= \{\eta \in T_g\mathcal{A}^k \mid \partial_\eta(\nabla_g f(v)) \in \mathrm{im}(D^v_g)\}. \end{aligned}$$

Since $\mathrm{im}(D^v_g)$ is closed, the result follows from the fact that $\eta \to \partial_\eta(\nabla_g f(v))$ is continuous.

□

1.4.7 Proof of Theorem 1.4.9

It just remains to show that the vertical differentials

$$\begin{aligned} D^v_g\colon \Gamma^2_1(v^*TM) &\to L^2(v^*TM) \\ \xi &\mapsto \widetilde{\nabla}_{v'}\xi + \widetilde{\nabla}_\xi(\nabla_g f(v)) \end{aligned}$$

are Fredholm with index $\mu(p) - \mu(q)$ for every $v \in F_g^{-1}(S_\mathcal{E})$. Here, we follow closely the discussion in section 5.3 of [4] which, in turn, is based on the paper [6].

To simplify matters we trivialize the bundle $v * TM$ using parallel translation with respect to $\widetilde{\nabla}$. For this choice we have

$$\begin{aligned} D^v_g\colon L^2_1(\mathbb{R}, \mathbb{R}^n) &\to L^2(\mathbb{R}, \mathbb{R}^n) \\ \xi &\mapsto (\partial_s - A_s)\xi \end{aligned}$$

where $A_s\colon \mathbb{R}^n \to \mathbb{R}^n$ is defined in this trivialization by $A_s(W) = -\widetilde{\nabla}_W(\nabla_g f(v(s)))$.

The family of linear maps A_s depends C^k-smoothly on the parameter s and converges exponentially as $s \to \pm\infty$ to limits $A^\pm$. Note that

$$A^- = -H_f(p) \quad \text{and} \quad A^+ = -H_f(q).$$

It follows that the dimension of the positive eigenspace of A^-, $\dim(E^+(A^-))$, is equal to $\mu(p)$. Similarly, $\dim(E^+(A^+)) = \mu(q)$.

Theorem 1.4.12 (see [6]). *The operator $\partial_s - A_s$ is Fredholm and*

$$\mathrm{ind}(\partial_s - A_s) = -SF\{A_s\}.$$

The term $SF\{A_s\}$ is called the **spectral flow** of the family of operators A_s. Intuitively, it counts the number of eigenvalues of A_s which pass from negative to positive as s goes from $-\infty$ to ∞. Since we know the number of positive eigenvalues of the limits $A^{\pm}$ we can write

$$SF\{A_s\} = \dim(E^+(A^+)) - \dim(E^+(A^-)) = \mu(q) - \mu(p)$$

From this it follows that D_g^v is Fredholm and has index $\mu(p) - \mu(q)$.

Remark 1.4.13. *In the Floer setting the analogues of the Morse indices are infinite, but the spectral flow (the "difference" between the indices) is finite. This is why the spectral flow is often referred to as a relative index.*

1.4.8 Proof of Theorem 1.4.12

First let us acknowledge the actual details involved in proving this result and then discuss the more easily accessible aspects of the proof. The fact that the operator $\partial_s - A_s$ has a closed range and a a finite dimensional kernel can be established using the following result:

Lemma 1.4.14 (Abstract closed range lemma). *Suppose that X, Y and Z are Banach spaces, $D\colon X \to Y$ is a bounded linear operator, and $K\colon X \to Z$ is a compact linear operator. If*

$$\|x\|_X \leq c(\|Dx\|_Y + \|Kx\|_Z),$$

then D has a closed range and a a finite dimensional kernel.

One applies this lemma to the Banach spaces $X = L_1^2(\mathbb{R}, \mathbb{R}^n)$, $Y = L^2(\mathbb{R}, \mathbb{R}^n)$, $Z = L^2([-T, T], \mathbb{R}^n)$, and the maps $D = \partial_s - A_s$ and $K =$ inclusion. The verification of the required estimate is described in §2 of [6].

In fact, it is easy to identify the kernel of $\partial_s - A_s$ and hence to determine its dimension. Assuming that the range of $\partial_s - A_s$ is closed, it is also easy to do the same for the cokernel of $\partial_s - A_s$.

- **The kernel of $\partial_s - A_s$**

  Assume that $u \in C^0(\mathbb{R}, \mathbb{R}^n)$ satisfies

  $$(\partial_s - A_s)u = 0.$$

  Then u is a solution of the ODE, $\partial_s u = A_s u$, which is linear and has C^k-smooth coefficients. It follows that u is C^k-smooth and is uniquely determined by its value at $s = 0$. Let u_h be the unique solution with initial value $u_h(0) = h$. It is defined for all $s \in R$. More generally, there is a fundamental solution matrix $\Phi_s \in \mathbb{R}^{n \times n}$ such that $u_h(s) = \Phi_s(h)$ for all $h \in \mathbb{R}^n$ and for all $s \in R$.

Exercise 1.4.15. *Show that if* $A_s = A$ *then the fundamental solution matrix* Φ_s *can be written in the form* $\exp(Qs)$, *where* Q *is a diagonal matrix whose entries are the eigenvalues of* A. *Describe the possible asymptotic behaviors of the trajectories.*

In order to determine the kernel of $\partial_s - A_s$, we must specify which of the solutions u_h is in $L^2_1(\mathbb{R}, \mathbb{R}^n)$. At the very least, we require $\lim_{s\to\pm\infty} u_h(s) = 0$. Accordingly, we define

$$\mathcal{H}^+ = \{h \in \mathbb{R}^n \mid \lim_{s\to+\infty} u_h(s) = 0\}$$

and

$$\mathcal{H}^- = \{h \in \mathbb{R}^n \mid \lim_{s\to-\infty} u_h(s) = 0\}.$$

Lemma 1.4.16. *The map*

$$\begin{array}{rcl} \mathcal{H}^- \cap \mathcal{H}^+ & \to & \ker(\partial_s - A_s) \\ h & \mapsto & u_h \end{array}$$

is an isomorphism.

Proof. The map is clearly injective.

It is onto, since $u \in \ker(\partial_s - A_s) \subset L^2_1$ implies that u goes to zero as $s \to \pm\infty$. Hence, $u = u_h$ for $h = u(0)$.

To see that the map is into, we note that when $|s|$ is large, $u_h(s)$ behaves like a solution of the linear flow determined by the self adjoint operators $A^\pm$. Hence, if $|u_h|$ tends to zero it must do so exponentially (together with its derivatives), see Exercise 1.4.15. Consequently, u_h belongs to $L^2_1(\mathbb{R}, \mathbb{R}^n)$. □

Lemma 1.4.17.

$$\dim(\mathcal{H}^+) = \dim(E^-(A^+)) \ \ \textit{and} \ \ \dim(\mathcal{H}^-) = \dim(E^+(A^-)).$$

Proof. The map

$$\begin{array}{rcl} \mathcal{H}^+ & \to & E^-(A^+) \\ h & \mapsto & \lim_{s\to\infty} |h| \dfrac{u_h(s)}{|u_h(s)|} \end{array}$$

is an isomorphism.

□

- **The cokernel of $\partial_s - A_s$**

  Given the fact that the image of $\partial_s - A_s$ is closed, we have the following isomorphism:

$$\operatorname{coker}(\partial_s - A_s) \simeq \operatorname{im}(\partial_s - A_s)^{\perp}.$$

If the formal adjoint of $\partial_s - A_s$ exists, then the we would have

$$\operatorname{im}(\partial_s - A_s)^{\perp} = \ker((\partial_s - A_s)^*).$$

Exercise 1.4.18. *Show that $(\partial_s - A_s)^*$ exists and is equal to $-\partial_s - A_s$.*

By the exercise, it suffices to calculate the dimension of $\ker(\partial_s + A_s)$.

Lemma 1.4.19. *The map*

$$\begin{array}{rcl} \ker(\partial_s + A_s) & \to & (\mathcal{H}^+)^{\perp} \cap (\mathcal{H}^-)^{\perp} \\ \tilde{f} & \mapsto & \tilde{f}(0) \end{array}$$

is an isomorphism.

Proof. First we show that the map is well-defined. For every $h \in \mathcal{H}^{\pm}$ we have

$$\begin{array}{rcl} \partial_s \langle \tilde{u}, u_h \rangle & = & \langle \partial_s \tilde{u}, f_u \rangle - \langle \tilde{u}, \partial_s u_h \rangle \\ & = & \langle -A_s \tilde{u}, u_h \rangle - \langle \tilde{u}, A_s u_h \rangle \\ & = & 0. \end{array}$$

Since $\tilde{u}$ is in $\ker(\partial_s + A_s)$ we know that $\lim_{s \to \pm\infty} \tilde{u}(s) = 0$. Also, since $h \in \mathcal{H}^{\pm}$ we have $\lim_{s \to \pm\infty} u_h(s) = 0$. It follows that

$$\langle \tilde{u}, u_h \rangle = 0 \text{ for all } s \in \mathbb{R}.$$

Setting $s = 0$, we get

$$\langle \tilde{u}(0), h \rangle = 0.$$

Thus, $\tilde{u}(0) \in (\mathcal{H}^+)^{\perp} \cap (\mathcal{H}^-)^{\perp}$.

Note that $\partial_s + A_s^*$ is an operator of the same general form as $\partial_s - A_s$. In particular, its kernel consists of certain solutions of an ODE which can be specified by their initial values. Hence, the map $\tilde{u} \to \tilde{u}(0)$ is injective.

The fact that the map is surjective can be proven in the same way as Lemma 1.4.17.

□

- By the previous discussion it follows that $\partial_s - A_s$ is Fredholm and

$$\begin{aligned}
\operatorname{ind}(\partial_s - A_s) &= \dim(\mathcal{H}^+ \cap \mathcal{H}^-) - \dim((\mathcal{H}^+)^\perp \cap (\mathcal{H}^-)^\perp) \\
&= \dim(\mathcal{H}^+ \cap \mathcal{H}^-) - (n - \dim(\mathcal{H}^+ \cup \mathcal{H}^-)) \\
&= \dim(\mathcal{H}^+) + \dim(\mathcal{H}^-) - n \\
&= \dim(E^-(A^+)) + \dim(E^+(A^-)) - n \\
&= \dim(E^+(A^-)) - \dim(E^+(A^+)).
\end{aligned}$$

1.4.9 Aside 1: Transversality Version B $\Rightarrow$ Transversality Version A

Let x be any point in $\mathcal{D}(p) \cap \mathcal{A}(q)$. Then there exists a unique gradient trajectory $v \in \mathcal{M}(p,q)$ such that $v(0) = x$.

For the trivialization of v^*TM above, one can check that

$$\mathcal{H}^+ = T_x\mathcal{D}(p) \quad \text{and} \quad \mathcal{H}^- = T_x\mathcal{A}(p).$$

Transversality theorem 1.4.6 implies that for generic $g \in \mathcal{A}^k$ the vertical differential $D_g^v = \partial_s - A_s$ is onto.

Hence,

$$\operatorname{coker}(D_g^v) = \operatorname{coker}(\partial_s - A_s) = 0.$$

This implies, by the isomorphism of Lemma 1.4.19, that

$$\begin{aligned}
& (\mathcal{H}^+)^\perp \cap (\mathcal{H}^-)^\perp = 0 \\
\Rightarrow\ & \mathcal{H}^+ + \mathcal{H}^- = \mathbb{R}^n \\
\Rightarrow\ & T_x\mathcal{D}(p) + T_x\mathcal{A}(p) = T_xM \\
\Rightarrow\ & \mathcal{D}(p) \pitchfork \mathcal{A}(q).
\end{aligned}$$

1.4.10 Aside 2: genericity in the smooth category

1.5 Compactness

By the Transversality theorem we know that for a generic $g \in \mathcal{A}^k$ the moduli space $\mathcal{M}(p,q)$ is a smooth manifold of dimension $\mu(p) - \mu(q)$. In fact, we know that $\mathcal{M}(p,q)$ is a smooth submanifold of the Banach manifold $\mathcal{B} \subset L_1^2(\mathbb{R}, M)$ from which it inherits its topology.

Each moduli space $\mathcal{M}(p,q)$ is inherently noncompact due to the natural free R-action it admits. For example, given $u \in \mathcal{M}(p,q)$ the sequence (u_k) where $u_k(s) = u(s+k)$ has no convergent subsequence. Accordingly, in this section we consider the quotient moduli spaces $\widehat{\mathcal{M}}(p,q) = \mathcal{M}(p,q)/\mathbb{R}$ and construct natural compactifications for them.

Definition 1.5.1. *A sequence $(\hat{u}_n) \subset \widehat{\mathcal{M}}(p,q)$* **converges** *to $\hat{u} \in \widehat{\mathcal{M}}(p,q)$ if for any lifts u_n of the $\hat{u}_n$ and u of $\hat{u}$, there are shifts τ_n such that*

$$\tau \cdot u_n \to u(s)$$

in $\mathcal{M}(p,q)$. (Recall that $\tau \cdot u(s) = u(s+\tau)$.)

The process of compactification

To compactify a topological space X one must first characterize (classify) the ways in which sequences in X can fail to converge. Then one adds these missing limits, ∂X, to the original space and defines a new notion of convergence for $\bar{X} = X \coprod \partial X$ which extends the notion of convergence in X.

The standard example of this is the case $X = \mathbb{R}$ where sequences can diverge to either $\pm\infty$ and one defines convergence in the compactification $\bar{\mathbb{R}} = \mathbb{R} \coprod \{\pm\infty\}$ by defining open the neighborhoods of $\pm\infty$.

1.5.1 Basic properties of negative gradient trajectories

Let $u\colon \mathbb{R} \to M$ be a solution of

$$\partial_s u = -\nabla_g f(u). \tag{1.5}$$

We describe here some useful properties of negative gradient trajectories which will help us build the compactification.

Lemma 1.5.2 (f decreases along u). *If u is nonconstant and $s_1 > s_2$, then $f(u(s_1)) > f(u(s_2))$.*

Proof.

$$\begin{aligned} f(u(s_1)) - f(u(s_2)) &= -\int_{s_1}^{s_2} \frac{d}{ds} f(u(s))\, ds \\ &= -\int_{s_1}^{s_2} df_{u(s)}(\partial_s u)\, ds \\ &= -\int_{s_1}^{s_2} g(\nabla_g f(u), \partial_s u)\, ds \\ &= \int_{s_1}^{s_2} \|\partial_s u\|^2\, ds. \end{aligned}$$

□

Lemma 1.5.3 (Convergence to critical points). *If* $\lim_{s\to\infty} u(s) = p$*, then* $p \in \mathrm{Crit}(f)$*.*

Proof. By the proof of the previous lemma we have

$$\lim_{s\to\infty} \int_0^s \|\partial_s u\|^2 \, ds = f(u(0)) - f(u(p)) < \infty.$$

It follows that

$$\lim_{s\to\infty} \|\partial_s u(s)\| = 0.$$

Hence,

$$\|\nabla_g f(p)\| = \lim_{s\to\infty} \|\nabla_g f(u(s))\| = \lim_{s\to\infty} \|\partial_s u(s)\| = 0.$$

□

Lemma 1.5.4 (Exponential convergence to critical points). *There is an* $\epsilon > 0$ *such that if* $\lim_{s\to\infty} u(s) = p$*, then there are constants* $C > 0$ *and* $S \in \mathbb{R}$ *(depending on* u*) such that*

$$d(u(s), p) \leq Ce^{-\epsilon s} \text{ for all } s \geq S.$$

Proof. The idea of the proof is very simple. For sufficiently large s the points $u(s)$ lie in a coordinate neighborhood of p where the vector field $\nabla_g f$ is close to its linearization at p. The linearized vector field is a representation of minus the Hessian of f at p. Hence, any trajectory of the linearized vector field which converges to p as $s \to \infty$ must do so along the positive eigenspace of the Hessian at an exponential rate which is bounded from below by the smallest positive eigenvalue of the Hessian. It remains to show that similar estimates of convergence hold for the nearby vector field. See [7] Lemma 2.10, for the details.

□

Lemma 1.5.5 (Convergence on compact subsets). *Let* (u_n) *be a sequence of maps from* $\mathbb{R} \to M$ *satisfying* (1.5)*. For every* $R > 0$ *there is a subsequence* (u_{n_j}) *such that*

$$u_{n_j}|_{[-R,R]} \overset{C^k}{\to} v|_{[-R,R]}$$

where v *also satisfies* (1.5)*.*

Proof. By the compactness of M it follows that the sequence of maps u_n are uniformly bounded as are their kth-order derivatives. The lemma now follows directly from the Arzela-Ascoli theorem. □

Letting $R \to \infty$ we can rephrase this last lemma by saying that any sequence (u_n) of solutions of (1.5) has a subsequence which converges to a solution v of (1.5) in the C^k_{loc}-topology, i.e.,

$$u_{n_j} \overset{C^k_{loc}}{\to} v.$$

It is important to note that the limit v is defined on all of $\mathbb{R}$ and hence belongs to some $\mathcal{M}(p,q)$ by Lemma 1.5.3.

Lemma 1.5.6 (C^k_{loc}-convergence in $\mathcal{M}(p,q)$ implies convergence). *If a sequence $(u_n) \subset \mathcal{M}(p,q)$ converges to $v \subset \mathcal{M}(p,q)$ in the C^k_{loc}-sense, then $u_n \to v$ with respect to the usual topology of $\mathcal{M}(p,q)$*

Proof.

- First we prove that u_n converge uniformly to the points p and q. That is, for any neighborhoods U_p and U_q of p and q, respectively, there is an $S > 0$ such that for all $n \in \mathbb{N}$ we have

$$u_n(s) \in U_p \quad \text{for} \quad s < -S$$

and

$$u_n(s) \in U_q \quad \text{for} \quad s > S.$$

It suffices to show that there is an S for which the second statement holds. Assume that no such S exists, i.e., there is a sequence $s_n \to \infty$ such that $u_n(s_n) \notin U_q$. This implies that there is an $\epsilon > 0$ such that

$$\|\nabla_g f(u_n(s_n))\| \geq \epsilon. \tag{1.6}$$

Now, for any $\delta > 0$ we have

$$f(u_n(s_n)) - f(q) \geq f(u_n(s_n)) - f(u_n(s_n + \delta)).$$

We will prove later that for sufficiently small $\delta > 0$

$$f(u_n(s_n)) - f(u_n(s_n + \delta)) \geq \delta\frac{\epsilon^2}{4}. \tag{1.7}$$

Since $v \in \mathcal{M}(p,q)$ there is an s_0 such that

$$f(v(s_0)) - f(q) = \delta\frac{\epsilon^2}{8}.$$

For sufficiently large n we have $s_n > s_0$ and

$$f(u_n(s_0)) - f(q) \geq f(u_n(s_n)) - f(q) \geq \delta\frac{\epsilon^2}{4}.$$

But this means that $u_n(s_0)$ cannot converge to $v(s_0)$ which is contrary to our assumption that (u_n) converges to v in the C^k_{loc}-topology.

It remains to prove (1.7). First we note that the vector field $\nabla_g f$ is globally Lipschitz continuous. That is, there is a constant $M > 0$ such that

$$| \, \|\nabla_g f(x)\| - \|\nabla_g f(y)\| \, | \leq M \cdot d(x, y).$$

Also we have for any $u \in \mathcal{M}(p, q)$

$$\begin{aligned} d(u(s_1), u(s_2)) &\leq \int_{s_1}^{s_2} \|\partial_s u\| \, ds \\ &\leq (s_2 - s_1)^{1/2} \left(\int_{s_1}^{s_2} \|\partial_s u\|^2 \, ds \right)^{1/2} \quad \text{(by Hólder's inequality)} \\ &= (s_2 - s_1)^{1/2} (f(u(s_1)) - f(u(s_2)))^{1/2} \quad \text{(see proof of Lemma 1.5.2)} \\ &\leq (s_2 - s_1)^{1/2} (f(p) - f(q)). \end{aligned}$$

Putting these together, there is a constant $c > 0$ such that for all n we have

$$| \, \|\nabla_g f(u_n(s))\| - \|\nabla_g f(u_n(s_n))\| \, | \leq c \cdot |s - s_n|^{1/2}. \tag{1.8}$$

Now

$$f(u_n(s_n)) - f(u_n(s_n + \delta)) = \int_{s_n}^{s_n + \delta} \|\nabla_g f(u_n(s)\|^2 \, ds.$$

By (1.6) and (1.8), for $s \in [s_n, s_n + \delta]$ we have

$$\|\nabla_g f(u_n(s)\| \geq \epsilon - \delta^{1/2} \geq \epsilon/2$$

for sufficiently small δ.

Hence,

$$f(u_n(s_n)) - f(u_n(s_n + \delta)) \geq \delta \frac{\epsilon^2}{4}$$

as required.

- The previous result allows us to obtain bounds for exponential convergence of the u_n to their endpoints which are independent of n. In particular, there are constants $\epsilon, C, S > 0$ such that for all $n \in N$

$$d(u(s), p) \leq Ce^{\epsilon s} \text{ for all } s \leq -S.$$

and

$$d(u(s), q) \leq Ce^{-\epsilon s} \text{ for all } s \geq S.$$

- At this point, we can consider the convergence of u_n to v on the intervals $(-\infty, -S]$, $[-S, S]$, and $[S, \infty)$. On the intervals $(-\infty, -S]$ and $[S, \infty)$ we have exponential convergence to v which implies L^2_1-convergence. On the compact interval $[-S, S]$ we have C^k-convergence which also implies L^2_1-convergence.

□

1.5.2 Failure to converge

Now let's see what can happen to a sequence $(\hat{u}_n)$ in $\widehat{\mathcal{M}}(p, q)$. Consider a lifting of this sequence $(u_n) \subset \mathcal{M}(p, q)$. By Lemma 1.5.5, we can pass to a subsequence such that

$$u_n \overset{C^k_{loc}}{\to} v$$

where $v \colon \mathbb{R} \to M$ belongs to $\mathcal{M}(p', q')$ for some $p', q' \in \mathrm{Crit}(f)$. Since $v(s) = \lim_{n\to\infty} u_n(s)$ for every $s \in \mathbb{R}$ we have

$$f(q) \le f(q') \le f(p') \le f(p).$$

If $p' = p$ and $q' = q$, then u_n converges to $v \in \mathcal{M}(p, q)$ by Lemma 1.5.6. In this case, we would say that $\hat{u}_n \to \hat{v}$. For example, this happens automatically if there are no critical points $r \in \mathrm{Crit}(f)$ with $f(r) \in [f(q), f(p)]$.

On the other hand, it may very well happen that $p' \neq p$ and/or $q' \neq q$. In this case, the sequence $(\hat{u}_n) \subset \widehat{\mathcal{M}}(p, q)$ is divergent and its "asymptotic limit" seems, at first glance, to be the element $\hat{v}$ in the different moduli space $\widehat{\mathcal{M}}(p', q')$. However, we will show that a different choice of lifting of the $\hat{u}_n$ leads to a limit in a different moduli space. In fact, we will show that there are only finitely many possible limits for lifts of the sequence $(\hat{u}_n)$. Moreover, these limits can be assembled to form a unique **broken gradient trajectory** which will be the right choice for the "asymptotic limit" of the divergent sequence $(\hat{u}_n)$.

To show that different lifts can have different limits let's consider the case $p' \neq p$. Without loss of generality, we may assume that this implies $f(p') < f(p)$. [1] We can then choose $f(p') < a < f(p)$ and a sequence of shifts τ_n such that

$$f(u_n(\tau_n)) = a.$$

The sequence $(\tau_n \cdot u_n)$ is also a lifting of $(\hat{u}_n)$. Again by Lemma 1.5.3, it converges in the C^k_{loc} topology to some $w \in \mathcal{M}(p'', q'')$ for critical points $p'', q'' \in \mathrm{Crit}(f)$ such that

$$f(q) \le f(q'') \le f(p'') \le f(p).$$

[1] We assume that the critical points of f have distinct critical values. This is a generic property.

In fact, we will show that the f-values of w lie inside the gap between $f(p)$ and $f(p')$.

Lemma 1.5.7. *The choice of the shifts* τ_n *implies that*

$$f(p') \leq f(q'') \leq f(p'') \leq f(p).$$

Proof. It suffices to prove that $f(q'') \geq f(p')$. We first show that the sequence of shifts τ_n goes to $-\infty$ as $n \to \infty$. If this is not true then there is an $K \in \mathbb{R}$ such $\tau_n > K$ for all $n \in \mathbb{N}$. By Lemma 1.5.2, this would imply that

$$f(\tau_n \cdot u_n(0)) = f(u_n(\tau_n)) < f(u_n(K)).$$

But $f(u_n(\tau_n)) = a$, and $\lim_{n\to\infty} f(u_n(K)) = f(v(K)) < f(p')$. This contradicts the fact that $a > f(p')$, hence the τ_n are not bounded from below.

Now we prove that $f(q'') \geq f(p')$ by contradiction. If $f(p') > f(q'')$ then there exists an $\epsilon > 0$ and an $s_1 \in \mathbb{R}$ such that

$$f(w(s_0)) = f(p') - 4\epsilon.$$

We can also choose an $s_1 \in \mathbb{R}$ such that

$$f(v(s_1)) = f(p') - \epsilon.$$

Since $f(\tau_n \cdot u_n(s_0)) \to f(w(s_0))$ and $f(u_n(s_1)) \to f(v(s_1))$, there is an $N \in \mathbb{N}$ such that for all $n > N$

$$f(\tau_n \cdot u_n(s_0)) < f(p') - 3\epsilon$$

and

$$f(u_n(s_1)) > f(p') - 2\epsilon.$$

In particular, this implies that

$$f(u_n(s_0 - \tau_n)) < f(u_n(s_1)) \text{ for all } n > N.$$

By Lemma 1.5.2, we then have

$$\begin{aligned} & s_0 & -\tau_n > s_1 \\ \Rightarrow\ & \tau_n & > s_1 - s_0 \end{aligned}$$

for all $n > N$. This contradicts the fact that the τ_n are not bounded from below. □

Note that if there are no $r \in \mathrm{Crit}(f)$ with $f(r) \in (f(p'), f(p))$, then $w \in \mathcal{M}(p, p')$. Otherwise, we can repeat this process to obtain lifts which converge to gradient trajectories that fit in the remaining gaps. Since there are only a finite number of critical points this is a finite process. Eventually we obtain a sequence of limits $\{v_j\}_{j=1,\ldots k}$ where $v_j \in \mathcal{M}(r_j, r_{j+1})$ and the critical points

$$p = r_0, r_1, \ldots, r_{k+1} = q$$

are all distinct. We then define the "asymptotic limit" of $(\hat{u}_n)$ to be the **broken trajectory**:

$$(\hat{v}_0, \hat{v}_1, \ldots, \hat{v}_k) \in \widehat{\mathcal{M}}(p, r_1) \times \widehat{\mathcal{M}}(r_1, r_2) \times \cdots \times \widehat{\mathcal{M}}(r_k, q).$$

There is one remaining question to consider: Is the "asymptotic limit" of the divergent sequence $(\hat{u}_n)$ unique? The following exercise implies that the answers to this question is "Yes."

Exercise 1.5.8. *Let (u_n) and (u'_n) be two lifts of a sequence $(\hat{u}_n) \in \widehat{\mathcal{M}}(p, q)$. We know that*

$$u_n \overset{C^k_{loc}}{\to} v \in \mathcal{M}(p', q')$$

and

$$u'_n \overset{C^k_{loc}}{\to} v' \in \mathcal{M}(p'', q'').$$

Suppose there exist numbers $s, s' \in \mathbb{R}$ such that $f(v(s)) = f(v'(s'))$. Prove that this implies $\hat{v} = \hat{v}'$.

1.5.3 The compactification

At this point we can compactify $\widehat{\mathcal{M}}(p, q)$ by adding to it the broken trajectories

$$\partial \widehat{\mathcal{M}}(p, q) = \bigcup_{\substack{r_1, \ldots, r_k \in \mathrm{Crit}(f) \\ p, r_1, \ldots, r_k, q \text{ distinct}}} \widehat{\mathcal{M}}(p, r_1) \times \widehat{\mathcal{M}}(r_1, r_2) \times \cdots \times \widehat{\mathcal{M}}(r_k, q).$$

Then on the compactification

$$\overline{\widehat{\mathcal{M}}(p, q)} = \widehat{\mathcal{M}}(p, q) \coprod \partial \widehat{\mathcal{M}}(p, q)$$

we extend the notion of convergence (for sequences in $\widehat{\mathcal{M}}(p, q)$) as follows.

Definition 1.5.9. *A sequence* $(\hat{u}_n) \subset \widehat{\mathcal{M}}(p,q)$ *converges to*

$$(\hat{v}_1, \ldots, \hat{v}_k) \in \widehat{\mathcal{M}}(p, r_1) \times \widehat{\mathcal{M}}(r_1, r_2) \times \cdots \times \widehat{\mathcal{M}}(r_k, q),$$

if for any lifts u_n *of* $\hat{u}_n$ *and* v_j *of* $\hat{v}_j$ *, there are shift sequences* $(\tau_{n,j})$ *such that*

$$\tau_{n,j} \cdot u_n \overset{C^k_{loc}}{\to} v_j$$

for $j = 1, \ldots, k$.

This definition can be easily extended to sequences with elements in $\partial\widehat{\mathcal{M}}(p,q)$.

Our previous discussion implies that $\overline{\widehat{\mathcal{M}}(p,q)}$ is indeed compact with this definition, and Lemma 1.5.6 proves that this extends the usual notion of convergence on $\widehat{\mathcal{M}}(p,q)$.

1.6 Gluing

The gluing theorem is, in essence, a converse to the compactness theorem. For simplicity (and safety), we will state the gluing theorem for broken trajectories with only one break. This suffices for the construction of the Morse complex.

Theorem 1.6.1 (Gluing). *Let* p, r, q *be critical points of* f *such that* $f(p) > f(r) > f(q)$ *and* $\mu(p) > \mu(r) > \mu(q)$. *There is a lower bound* $R_0 > 0$ *and a smooth map*

$$\#\colon \mathcal{M}(p,r) \times \mathcal{M}(r,q) \times (R_0, \infty) \to \mathcal{M}(p,q)$$

such that

1. $\#$ *induces a diffeomorphism*

   $$\widehat{\#}\colon \widehat{\mathcal{M}}(p,r) \times \widehat{\mathcal{M}}(r,q) \times (R_0, \infty) \to \widehat{\mathcal{M}}(p,q)$$

   which is onto an open set in $\widehat{\mathcal{M}}(p,q)$.

2. *As* $R \to \infty$, *the sequence* $\widehat{\#}(\hat{v}_1, \hat{v}_2, R)$ *converges to the broken trajectory* $(\hat{v}_1, \hat{v}_2)$ *in the sense of Definition 1.5.9.*

3. *If the sequence* $(\hat{u}_n) \subset \widehat{\mathcal{M}}(p,q)$ *converges to* $(\hat{v}_1, \hat{v}_2)$, *then the* $\hat{u}_n$ *are in the range of* $\widehat{\#}$ *for all sufficiently large* n.

Exercise 1.6.2. *Suppose* $\mu(p) = \mu(q) + 2$. *Show that the Gluing theorem implies that our compactification* $\overline{\widehat{\mathcal{M}}(p,q)}$ *is a compact one-dimensional manifold with (possibly empty) boundary.*

1.6.1 A brief outline of the proof

We will just describe briefly how the map # is constructed. The general idea is to first use the broken trajectories to construct an approximate solution of the negative gradient flow from p to q, and then show that there is a unique actual solution corresponding to the approximate one.

We will not provide many details or verify the various properties of the gluing map. A complete proof can be found in [7].

Step 1 : Pregluing

Given $v_1 \in \mathcal{M}(p,r)$, $v_2 \in \mathcal{M}(r,q)$ and a large constant $R > 0$, we will construct an approximate solution $v_1 \widetilde{\#}_R v_2$ of the negative gradient flow from p to q.

By Lemma 1.5.4, there is an $R > 0$ such that for each $s > R$ we can write

$$v_1(s) = \exp_r \eta(s)$$

for a unique $\eta(s) \in T_r M$. Moreover, $\eta(s)$ depends C^k-smoothly on s.

Similarly, we may write $v_2(s) = \exp_r \xi(s)$ for all $s < -R$.

Let $\beta\colon \mathbb{R} \to \mathbb{R}$, be smooth cutoff function, i.e. $\beta(s) = 0$ for $s \leq 0$ and $\beta(s) = 1$ for $s \geq 1$. We then set

$$v_1 \widetilde{\#}_R v_2(s) = \begin{cases} v_1(s+R), & s \leq -R/2 - 1 \\ \exp_r(\beta(-s - R/2)\eta(s+R), & -R/2 - 1 \leq s \leq -R/2 \\ r, & -R/2 \leq s \leq R/2 \\ \exp_r(\beta(s - R/2)\xi(s+R), & R/2 \leq s \leq R/2 + 1 \\ v_2(s-R), & s \geq R/2 + 1. \end{cases}$$

Exercise 1.6.3. *Show that $v_1 \widetilde{\#}_R v_2 \in \mathcal{B}$ and that $v_1 \widetilde{\#}_R v_2$ converges to the broken trajectory (v_1, v_2) in the C^k_{loc}-topology.*

Step 2

We now show that the approximate solutions $v_1 \widetilde{\#}_R v_2 \in \mathcal{B}$ correspond to actual solutions $u_R \in \mathcal{M}(p,q)$ for sufficiently large $R > 0$.

Here is the basic analytic machinery. Let $F\colon V \to W$ be a map of Banach spaces whose Taylor expansion around $0 \in V$ is

$$F(x) = F(0) + DF(0)x + N(x).$$

Lemma 1.6.4. *Suppose that $DF(0)$ has a finite dimensional kernel and a right inverse Q such that*

$$||QN(x) - QN(y)|| \leq C(||x|| + ||y||)||x - y|| \tag{1.9}$$

for some constant C and all x, y in some ball $B_\epsilon(0)$ with radius $\epsilon \leq 1/5C$. If $||QF(0)|| < \epsilon/2$, then there is a unique $x_0 \in B_\epsilon(0)$ satisfying $F(x_0) = 0$.

Intuitively, this result states that if $DF(0)$ is sufficiently large and $F(0)$ is sufficiently small, then there is a unique solution of $F(x) = 0$ close to 0.

Proof. To prove this lemma we consider the map

$$\begin{aligned} \psi \colon V &\rightarrow V \\ x &\mapsto -Q(F(0) + N(x)). \end{aligned}$$

Note that the fixed points of ψ correspond to zeroes of F. That is,

$$\begin{aligned} \psi(x) = x &\Rightarrow -Q(F(0) + N(x)) = x \\ &\Rightarrow -(F(0) + N(x)) = DF(0)x \\ &\Rightarrow F(0) + DF(0)x + N(x) = 0. \end{aligned}$$

It then suffices to show that ψ is a contraction mapping on some open ball $B_\epsilon(0)$.

Claim 1. $\psi \colon B_\epsilon(0) \rightarrow B_\epsilon(0)$ for all $\epsilon < 1/5C$.

$$\begin{aligned} ||\psi(x)|| &= ||Q(F(0) + N(x))|| \\ &\leq ||QF(0)|| + ||N(x)|| \\ &< \epsilon/2 + C\epsilon^2 \\ &< \epsilon. \end{aligned}$$

Claim 2. ψ is contraction mapping on $B_\epsilon(0)$ for all $\epsilon < 1/5C$.

$$\begin{aligned} ||\psi(x) - \psi(y)|| &= ||QN(x) - QN(y)|| \\ &\leq C2\epsilon||x - y|| \\ &< 2/5||x - y||. \end{aligned}$$

∎

Let us now apply this result to construct the gluing map. Recall that elements of $\mathcal{M}(p,q)$ correspond to the inverse image $F_g^{-1}(S_{\mathcal{E}})$. Near our approximate solution $v_1 \widetilde{\#}_R v_2 \in \mathcal{B}$ we have local coordinates ξ in which $F_g(\xi) = (\xi, F_g^{v_1\widetilde{\#}_R v_2}(\xi))$. The elements of $\mathcal{M}(p,q)$ close to $v_1 \widetilde{\#}_R v_2$ then correspond to zeroes of the vertical component $F_g^{v_1\widetilde{\#}_R v_2}$. After much hard work, one can establish the estimates for the right inverse of the vertical differential necessary to apply Lemma 1.6.4. We then obtain a unique element $u_R \in \mathcal{M}(p,q)$ close to $v_1 \widetilde{\#}_R v_2$. This defines the gluing map as follows

$$\#(v_1, v_2, R) = u_R.$$

1.7 Orientations

The definition of the Morse boundary operator ∂_g requires us to count the elements of the zero-dimension moduli spaces $\widehat{\mathcal{M}}(p,q)$ with sign. These signs are determined by comparing two orientations on these spaces. In particular, if these orientations on $\hat{u} \in \widehat{\mathcal{M}}(p,q)$ agree, then $\hat{u}$ is counted with a $+1$, otherwise it is counted with a -1. One of the orientations is canonically determined by the flow. The other orientation must then be chosen in such a way that opposite signs are given to the broken trajectories which constitute the boundary of the same component of a one-dimensional moduli space. This leads one to the notion of a coherent orientation for the moduli spaces.

1.7.1 Fredholm operators and determinant line bundles

Let X and Y be finite dimensional real vector spaces. Here we adopt the following notational convenience

$$\Lambda^{\max} X \equiv \Lambda^{\dim X} X.$$

We then define the determinant of the (ordered) pair of vector spaces (X, Y) to be

$$\operatorname{Det}(X, Y) \equiv (\Lambda^{\max} X) \otimes (\Lambda^{\max} Y)^*.$$

Now, consider Banach spaces V and W and the space $\mathcal{F}(V, W)$ of linear Fredholm maps from V to W. For $F \in \mathcal{F}(V, W)$ we set

$$\operatorname{Det}(F) = \operatorname{Det}(\ker F, \operatorname{coker} F).$$

Suppose that B is a topological space and we are given a continuous map $f \colon B \to \mathcal{F}(V, W)$, i.e., a continuous family of Fredholm maps indexed by B. We can then form

$$\mathrm{Det}(f) = \bigcup_{b \in B} \{b\} \times \mathrm{Det}(f(b)).$$

Proposition 1.7.1. Det(f) *is a real line bundle over* B.

This is somewhat surprising since the dimensions of vector spaces $\ker f(b)$ and coker $f(B)$ can vary wildly with b even though the index of $f(b)$ is locally constant. For a proof of this fact see the appendix in [3].

Definition 1.7.2. *An orientation for the family of Fredholm operators given by* $f \colon B \to \mathcal{F}(V, W)$*, is a nonvanishing section of* Det(f)*.*

1.7.2 The relevant class of Fredholm operators

Let A^- and A^+ be nondegenerate self-adjoint operators on $\mathbb{R}^n$, and consider the set operators

$$\Theta(A^-, A^+) = \{\partial_s - A_s \mid A_s \in \mathrm{End}(\mathbb{R}^n), \lim_{s \to \pm\infty} A_s = A^{\pm}\}.$$

By Theorem 1.4.12, we know that each $F \in \Theta(A^-, A^+)$ is a Fredholm operator from $L^2_1(\mathbb{R}, \mathbb{R}^n)$ to $L^2(\mathbb{R}, \mathbb{R}^n)$ with

$$\mathrm{ind}(F) = -SF\{A_s\} = \dim(E^+(A^-)) - \dim(E^+(A^+)).$$

In particular, $\Theta(A^-, A^+)$ is a continuous family of operators in $\mathcal{F}(L^2_1(\mathbb{R}, \mathbb{R}^n), L^2(\mathbb{R}, \mathbb{R}^n))$ indexed by the space of smooth curves in $\mathrm{End}(\mathbb{R}^n)$ which converge to $A^{\pm}$ at $\pm\infty$.

Lemma 1.7.3. $\Theta(A^-, A^+)$ *is contractible.*

Proof. Fix $F_0 \in \Theta(A^-, A^+)$ and let F be any element of $\Theta(A^-, A^+)$. Then the straight-line path $F_\tau = (1-\tau)F_0 + \tau F$ connects F_0 to F within $\Theta(A^-, A^+)$. Consider the map

$$\begin{aligned} \kappa \colon [0,1] \times \Theta(A^-, A^+) &\to \Theta(A^-, A^+) \\ (\tau, F) &\mapsto F_\tau. \end{aligned}$$

Since $\kappa(0, \cdot) = id$ and $\kappa(1, \cdot) = F_0$, it remains to show that κ is continuous. This is proved in [7] (Lemma 2.15). $\square$

It follows from Lemma 1.7.3 that $\mathrm{Det}(\Theta(A^-, A^+))$ is trivial and so $\Theta(A^-, A^+)$ is orientable. To determine an orientation we just need to orient $Det(F)$ for some $F \in \Theta(A^-, A^+)$. A choice of orientation for $\mathrm{Det}(\Theta(A^-, A^+))$ will be denoted by $\beta(A^-, A^+)$, in general, and by $\beta(F)$ if it is induced by a choice of orientation for a given $F \in \Theta(A^-, A^+)$.

Lemma 1.7.4. *If $A^+ = B^-$, then a choice of orientations $\beta(A^-, A^+)$ and $\beta(B^-, B^+)$, determines a canonical* **glued orientation** *$\beta(A^-, A^+)\#\beta(B^-, B^+)$ on $\Theta(A^-, B^+)$. Moreover, if $A^+ = B^-$ and $B^+ = C^-$, then this canonical orientation is associative in the following sense*

$$\big(\beta(A^-, A^+)\#\beta(B^-, B^+)\big)\,\#\beta(C^-, C^+) = \beta(A^-, A^+)\#\,\big(\beta(B^-, B^+)\#\beta(C^-, C^+)\big)\,.$$

Proof. Choose $F = \partial_s - A_s \in \Theta(A^-, A^+)$ and $G \in= \partial_s - B_s \in \Theta(B^-, B^+)$ such that

$$A_s = \begin{cases} A^-, & \text{for } s \leq -1 \\ A^+, & \text{for } s \geq 1. \end{cases}$$

and

$$B_s = \begin{cases} B^-, & \text{for } s \leq -1 \\ B^+, & \text{for } s \geq 1. \end{cases}$$

For $R > 1$ we can glue these operators to form

$$F\#_R G = \partial_s - C_s \in \Theta(A^-, B^+)$$

where

$$C_s = \begin{cases} A_{s+R}, & \text{for } s \leq 0 \\ B_{s-R}, & \text{for } s \geq 0. \end{cases}$$

It then suffices to prove that there is a canonical isomorphism

$$\text{Det}(F) \otimes \text{Det}(G) \simeq \text{Det}(F\#_R G). \tag{1.10}$$

More precisely, the given orientations $\beta(A^-, A^+)$ and $\beta(B^-, B^+)$ induce orientations $\beta(F)$ and $\beta(G)$ of $\text{Det}(F)$ and $\text{Det}(G)$, respectively. By (1.10), this would give a canonical orientation on $\text{Det}(F\#_R G)$ which determines the desired glued orientation $\beta(A^-, A^+)\#\beta(B^-, B^+)$ on $\Theta(A^-, B^+)$.

We omit the full proof of (1.10), and only briefly discuss the case when F and G are surjective. Under the surjectivity assumption, (1.10) becomes

$$\Lambda^{\max}\ker(F\#_R G) \simeq \Lambda^{\max}\ker(F) \otimes \Lambda^{\max}\ker(G).$$

Using the canonical isomorphism

$$(\Lambda^{\max}\mathbb{R}^k) \otimes (\Lambda^{\max}\mathbb{R}^l) \simeq \Lambda^{\max}\mathbb{R}^{k+l}$$

it suffices to show that

$$\ker(F) \times \ker(G) \simeq \ker(F\#_R G).$$

One can easily define a monomorphism

$$\phi_R\colon \ker(F) \times \ker(G) \to \ker(F\#_R G).$$

by gluing together elements in $\ker(F) \times \ker(G)$ in the obvious way.

Since $\mathrm{ind}(F\#_R G) = \mathrm{ind}(F) + \mathrm{ind}(G)$, when F and G are surjective we have

$$\dim(\ker(F\#_R G)) \geq \dim(\ker(F)) + \dim(\ker(G)).$$

One can then prove that ϕ_R is onto by showing that when R is sufficiently large

$$\dim(\ker(F\#_R G)) \leq \dim(\ker(F)) + \dim(\ker(G)).$$

The proof of this fact can be found in [7], Proposition 2.50. The central observation is that as $R \to \infty$ the restriction of $u \in \ker(F\#_R G)$ to $(-\infty, 0)$ converges (after translation) to an element of $\ker(F)$. □

1.7.3 Coherent orientations for our moduli spaces

We first will describe how to orient each of the moduli spaces $\widehat{\mathcal{M}}(p,q)$. Then we describe the notion of a coherent orientation for the collection of moduli spaces $\bigcup_{p,q\in \mathrm{Crit}(f)} \widehat{\mathcal{M}}(p,q)$. Finally we prove that a coherent orientation exists.

Orienting a fixed moduli space $\widehat{\mathcal{M}}(p,q)$

Recall that

$$\mathcal{M}(p,q) = \widehat{\mathcal{M}}(p,q) \times \mathbb{R}.$$

Hence, it suffices to orient the moduli spaces $\mathcal{M}(p,q)$.

Each $\mathcal{M}(p,q)$ corresponds to the inverse image $F_g^{-1}(S_{\mathcal{E}})$. Near each $v \in \mathcal{M}(p,q)$ we have local coordinates ξ in which $F_g(\xi) = (\xi, F_g^v(\xi))$. The elements of $\mathcal{M}(p,q)$ close to v then correspond to zeroes of the vertical component F_g^v and we have

$$T_v\mathcal{M}(p,q) = \ker(D_g^v).$$

Since D_g^v is surjective for our choice of generic g, we also have

$$\Lambda^{\max} \ker(D_g^v) \simeq \mathrm{Det}(D_g^v),$$

and so an orientation of $\mathrm{Det}(D_g^v)$ determines an orientation of $T_v\mathcal{M}(p,q)$.

Given $v \in \mathcal{M}(p,q)$ we have a trivialization $\phi_v \colon v^*TM \to \mathbb{R} \times \mathbb{R}^n$ so that in this trivialization D_g^v has the form $\partial_s - A_s$ where

$$\lim_{s \to -\infty} A_s = A_p \quad \text{and} \quad \lim_{s \to \infty} A_s = A_q.$$

In other words, we can consider D_g^v as an element in $\Theta(A_p, A_q)$. Let us then fix an orientation $\beta(A_p, A_q)$ for $\Theta(A_p, A_q)$. Together with the trivialization ϕ_v, this determines an orientation for $T_v\mathcal{M}(p,q)$.

Proposition 1.7.5 ([7], §3.2.1). *Given any $u \in \mathcal{M}(p,q)$ there is a trivialization $\phi_u \colon u^*TM \to \mathbb{R} \times \mathbb{R}^n$ such that $\phi_u(-\infty) = \phi_v(-\infty)$, $\phi_u(\infty) = \pm\phi_v(\infty)$ and the operator D_g^u transforms to an element of $\Theta(A_p, A_q)$.*

Remark 1.7.6. *If the manifold M is orientable, then these trivializations can be obtained in a canonical way and will always satisfy $\phi_u(\infty) = \phi_v(\infty)$. More precisely, we only have $\phi_u(\infty) = -\phi_v(\infty)$ if the bundle $(u \cdot v^{-1})^*TM$ is an n-dimensional Möbius band, i.e. $w_1((u \cdot v^{-1})^*TM) \neq 0$.*

The trivializations described by Proposition 1.7.5, determine orientations for each tangent space $T_u\mathcal{M}(p,q)$. The following result implies that these fit together to form an orientation of $\mathcal{M}(p,q)$.

Lemma 1.7.7 ([7], Lemma 3.8). *Let ϕ_u and ϕ'_u be trivializations as in Proposition 1.7.5. Then they determine equivalent orientations for $T_u\mathcal{M}(p,q)$.*

To summarize, if we fix a trivialization ϕ_v for some $v \in \mathcal{M}(p,q)$, we can associate to this moduli space the model space of Fredholm operators $\Theta(A_p, A_q)$. Moreover, an orientation $\beta(A_p, A_q)$ of this model space yields an orientation for $\mathcal{M}(p,q)$. For each $u \in \mathcal{M}(p,q)$ we will denote this induced orientation of $T_u\mathcal{M}(p,q)$ by $\beta(u)$.

Coherent orientations for the $\mathcal{M}(p,q)$

To the collection of moduli spaces

$$\bigcup_{p,q \in \mathrm{Crit}(f)} \mathcal{M}(p,q)$$

we can associate a collection of model spaces

$$\bigcup_{p,q \in \mathrm{Crit}(f)} \Theta(A_p, A_q).$$

Definition 1.7.8. *A collection of orientations* $\beta(A_p, A_q)$ *for the* $\Theta(A_p, A_q)$ *is said to be* **coherent** *if the glued orientation* $\beta(A_p, A_r)\#\beta(A_r, A_q)$ *agrees with* $\beta(A_p, A_q)$ *for every* $r \in \mathrm{Crit}(f)$.

Theorem 1.7.9. *Coherent orientations always exist.*

Proof. To show that a coherent orientation of the $\Theta(A_p, A_q)$ exists we construct one. First we fix a $p_o \in \mathrm{Crit}(f)$ and choose an orientation $\beta(A_{p_0}, A_{p_0})$ for $\Theta(A_{p_0}, A_{p_0})$. Then we choose orientations $\beta(A_{p_0}, A_q)$ for the spaces of the form $\Theta(A_{p_0}, A_q)$.

We claim that the coherence condition then determines the remaining choices of orientations. In particular, the orientations of the spaces $\Theta(A_q, A_{p_0})$ are determined by the condition

$$\beta(A_{p_0}, A_q)\#\beta(A_q, A_{p_0}) = \beta(A_{p_0}, A_{p_0}).$$

Finally, the the orientations of the spaces $\Theta(A_p, A_q)$ are determined by the condition

$$\beta(A_{p_0}, A_p)\#\beta(A_p, A_q)\#\beta(A_q, A_{p_0}) = \beta(A_{p_0}, A_{p_0}).$$

□

1.7.4 Canonical orientations

For $\mu(p) - \mu(q) = 1$, the moduli space $\mathcal{M}(p, q)$ comes with a canonical orientation. To see this note that for $v \in \mathcal{M}(p, q)$ we have

$$T_v\mathcal{M}(p, q) = \ker(D_g^v).$$

Since, D_g^v is surjective and one-dimensional, we also have

$$\mathrm{Det}(D_g^v) \simeq \ker(D_g^v).$$

But

$$D_g^v(\dot{v}) = \widetilde{\nabla}_{\dot{v}}\dot{v} + \widetilde{\nabla}_{\dot{v}}(\nabla_g f(v)) = \widetilde{\nabla}_{\dot{v}}\dot{v} - \widetilde{\nabla}_{\dot{v}}\dot{v} = 0.$$

Hence, $\dot{v}$ spans $\ker(D_g^v)$ and determines a canonical orientation $\alpha(v)$ for $T_v\mathcal{M}(p, q)$.

1.7.5 Geometric orientations

When the manifold M is itself orientable, one can define coherent orientations in the geometric version of the Morse complex as follows.

Choose an orientation for M and all the descending manifolds $\mathcal{D}(p)$. By the transversality condition $\mathcal{D}(p) \pitchfork \mathcal{A}(q)$ these choices determine orientations on all the ascending manifolds. This, in turn, yields noncanonical (geometric) orientations for the intersections

$$\mathcal{M}(p,q) = \mathcal{D}(p) \cap \mathcal{A}(q).$$

These noncanonical orientations satisfy the following coherence condition.

Proposition 1.7.10. *The gluing map* $\#\colon \widehat{\mathcal{M}}(p,r) \times \widehat{\mathcal{M}}(r,q) \times (R,\infty) \to \widehat{\mathcal{M}}(p,q)$ *is orientation preserving.*

This implies that the ends of one dimensional moduli spaces $\widehat{\mathcal{M}}(p,q)$ appear with opposite signs. For a proof of this see [1] (Proposition 2.7).

In Appendix B of [7], Schwarz shows that these geometric orientations can be extended to a coherent orientation in the sense of Definition 1.7.8.

1.8 The invariance of Morse homology

Given generic data (f,g) we have constructed the Morse complex $(C_*(f), \partial_g)$. In this section we prove that the corresponding Morse homology $HM_*(f,g)$ is independent of the data. More precisely, we prove

Theorem 1.8.1. *For two generic data pairs* (f_1,g_1) *and* (f_2,g_2)*, the Morse homology* $HM_*(f_1,g_1)$ *is* **canonically** *isomorphic to* $HM_*(f_2,g_2)$.

In other words, the Morse homology is an invariant of the manifold which we denote by $HM_*(M)$.

To prove Theorem 1.8.1, we will use a homotopy from (f_s,g_s) from (f_1,g_1) to (f_2,g_2) which satisfies

$$(f_s,g_s) = \begin{cases} (f_1,g_1), & s \leq -1 \\ (f_2,g_2), & s \geq 1. \end{cases}$$

One approach would be to try and keep track of the Morse complexes $(C_*(f_s), \partial_{g_s})$ as s goes from $-\infty$ to ∞. In particular, it follows from standard singularity theory that there is a homotopy f_s such that the number of elements in $\mathrm{Crit}(f_s)$ changes at only finitely many values of s at which standard "birth" or "death" bifurcations take place. Moreover, one can choose the homotopy (f_s,g_s) such that the geometric transversality condition also fails at only finitely many values of s and in a standard way. Between these critical values of s, the homology $HM_*(f_s),g_s)$ can easily be shown to be constant. One can then use the normal forms for the changes in the complex $(C_*(f_s), \partial_{g_s})$ which occur as one passes through a critical value of s to show that the homology again remains constant. Floer's original work on Floer homology for Lagrangian intersections, [2], follows this strategy which is commonly referred to as "bifurcation analysis". For more recent applications of this approach see, [8], for example.

We will use the more popular (and elegant) "Floer-Conley continuation" argument to prove Theorem 1.8.1.Here as an outline of this method.

Step 1.

First one uses the homotopy (f_s, g_s) to construct a map

$$\sigma_{21} \colon C_*(f_1) \to C_*(f_2),$$

and proves that σ_{21} is a chain map, i.e.,

$$\partial_{g_2} \circ \sigma_{21} = \sigma_{21} \circ \partial_{g_1}. \tag{1.11}$$

From this, it follows that σ_{21} induces a **homotopy homomorphism**

$$\sigma_{21} \colon HM_*(f_1, g_1) \to HM_*(f_2, g_2).$$

In the case when $(f_1, g_1) = (f_2, g_2)$ and (f_s, g_s) is constant it will follow easily that $\sigma_{11} \colon C_*(f_1) \to C_*(f_1)$ is the identity map.

Step 2.

Next, one considers two homotopies (f_s, g_s) and $(\widetilde{f}_s, \widetilde{g}_s)$, and shows that the corresponding chain maps σ_{21} and $\widetilde{\sigma}_{21}$ are chain homotopic. That is, there is a map

$$K \colon C_*(f_1) \to C_*(f_2),$$

such that

$$\sigma_{21} - \widetilde{\sigma}_{21} = \partial_{g_2} \circ K + K \circ \partial_{g_1}. \tag{1.12}$$

This implies that σ_{21} and $\widetilde{\sigma}_{21}$ induce the same homomorphism at the homology level. We denote this homomorphism by σ_{21}.

Step 3.

We then prove the following composition rule for homotopy homomorphisms

$$\sigma_{31} = \sigma_{32} \circ \sigma_{21}.$$

The maps σ_{21} and K from Steps 1 and 2, are constructed using different moduli spaces. The required algebraic relations (1.11) and (1.12) are equivalent to the fact that the signed sum of the boundary components of the one-dimensional moduli spaces is zero. Step 3, follows from a gluing theorem.

With this machinery we can now prove Theorem 1.8.1. From the composition rule we have

$$\sigma_{11} = \sigma_{12} \circ \sigma_{21}.$$

Since σ_{11} is the identity homomorphism (see last comment in Step 1), we see that σ_{21} must be injective. Similarly, the fact that

$$\sigma_{22} = \sigma_{21} \circ \sigma_{12},$$

implies that σ_{21} is surjective.

Bibliography

[1] D.M. Austin, P.J. Braam. **Morse-Bott theory an equivariant cohomology**.

[2] A. Floer. **Morse theory for Lagrangian intersections**.

[3] A. Floer, H. Hofer. **Coherent orientations for periodic orbit problems in symplectic geometry**.

[4] M. Hutchings. **Lecture notes on Morse homology (with an eye towards Floer theory and pseudoholomorphic curves)**.

[5] J. Jost. **Riemannian geometry and geometric analysis**.

[6] J. Robbin, D. Salamon. **The spectral flow and the Maslov index**.

[7] M. Schwarz. **Morse Homology**.

[8] M.G. Sullivan.**K-theoretic invariants for Floer homology**.