

Morse-Conley-Floer Homology

Contents

1	Introduction	1
1.1	Dynamics and Topology	1
1.2	Classical Morse theory, the half-space approach	4
1.3	Morse homology	7
1.4	Conley theory	11
1.5	Local Morse homology	16
1.6	Morse-Conley-Floer homology	18
1.7	Functoriality for flow maps in Morse-Conley-Floer homology .	21
1.8	A spectral sequence in Morse-Conley-Floer homology	23
1.9	Morse-Conley-Floer homology in infinite dimensional systems .	25
1.10	The Weinstein conjecture and symplectic geometry	29
1.11	Closed characteristics on non-compact hypersurfaces	32
I	Morse-Conley-Floer Homology	39
2	Morse-Conley-Floer homology	41
2.1	Introduction	41
2.2	Isolating blocks and Lyapunov functions	45
2.3	Gradient flows, Morse functions and Morse-Smale flows	49
2.4	Morse homology	55
2.5	Morse-Conley-Floer homology	65
2.6	Morse decompositions and connection matrices	68
2.7	Relative homology of blocks	74
3	Functoriality	79
3.1	Introduction	79
3.2	Chain maps in Morse homology on closed manifolds	86

CONTENTS

3.3	Homotopy induced chain homotopies	97
3.4	Composition induced chain homotopies	100
3.5	Isolation properties of maps	103
3.6	Local Morse homology	109
3.7	Morse-Conley-Floer homology	114
3.8	Transverse maps are generic	116
4	Duality	119
4.1	Introduction	119
4.2	The dual complex	120
4.3	Morse cohomology and Poincaré duality	120
4.4	Local Morse homology	122
4.5	Duality in Morse-Conley-Floer homology	122
4.6	Maps in cohomology	123
4.7	Shriek maps in Morse-Conley-Floer homology	125
4.8	Other algebraic structures	125
4.9	Some identities	127
4.10	Coincidence theory	127
5	Spectral Sequence	129
5.1	The birth of a spectral sequence	130
5.2	The Morse-Conley relations via the spectral sequence	132
5.3	Interpretation of the spectral sequence for Morse-Bott functions	134
5.4	A counter example to misstated Morse-Bott inequalities.	136
II	Periodic Orbits on Non-Compact Hypersurfaces	139
6	Periodic orbits	141
6.1	Introduction	141
6.2	Contact type conditions	144
6.3	The variational setting	150
6.4	The Palais-Smale condition	154
6.5	Minimax characterization	161
6.6	Some remarks on the geometry and topology of the loop space	164
6.7	The topology of the hypersurface and its shadow	166
6.8	The link	174
6.9	Estimates	184
6.10	Proof of the main theorem	187

6.11 Appendix: construction of the hedgehog	187
6.12 Appendix: local computations	196
7 Nederlandse samenvatting	199
7.1 Overzicht	199
7.2 Deel I: Morse-Conley-Floer homologie	199
7.3 Deel II: Periodieke banen op niet compacte hyperoppervlakken	200
Bibliography	203

Preface

This thesis consists of two parts. In Part I a homology theory for flows on finite dimensional manifolds is developed, and in Part II criteria for the existence of periodic orbits on non-compact hypersurfaces in Hamiltonian dynamics are found. The main connection between the two parts are the use of Morse theory in the study of dynamical systems. In Chapter 1 we give an overview of the setting and the history of the subjects, and discuss the results of this thesis more in depth. In Chapter 2 we define Morse-Conley-Floer homology and prove some of its basic properties. This chapter is based on the paper [RV14]. In Chapter 3 we treat functorial aspects of Morse homology, local Morse homology, and Morse-Conley-Floer homology. We apply Morse-Conley-Floer homology to the study Morse decompositions and degenerate variational systems in Chapter 5. In Part II, that consists of Chapter 6, we study the existence of periodic orbits on non-compact energy hypersurfaces. This chapter is based on the paper [VDBPRV13].

Chapter 1

Introduction

Voordat ik een fout maak
Maak ik die fout niet

Johan Cruijff

1.1 Dynamics and Topology

In this thesis two separate problems are studied. We develop a homology theory to study differential equations in general, and we find periodic solutions to a particular differential equation in Hamiltonian mechanics. The main connection between the results is the use of topological tools –in particular Morse theory– in the study of dynamical systems. An overview of the results of this thesis is given in Sections 1.6,1.8,1.7 and 1.11. Let us first comment on the symbiosis of dynamical systems theory and topology.

1.1.1 Topology

Topology is the study of that part of shape which is invariant under continuous deformations. It should not be confused with geometry, although the two are intimately related. In Figure 1.1 three different spaces are sketched. The sphere and the heart can be deformed into each other and are topologically the same. The torus cannot be deformed into either one of them. A lot of effort is devoted to understand which spaces are essentially the same in this regard. It is hard to imagine all possible deformations of a given space, and




Figure 1.1: A sphere, a heart and a torus. The first two spaces are homeomorphic (they have the same topology), but the torus has a different topology. To distinguish spaces with different topology, topological invariants are used. In this case the fundamental group distinguishes the spaces: The loop drawn on the right hand side of the torus cannot be shrunk to a point along the surface, while any loop drawn on the sphere or the heart can be contracted to a point.

to show that two spaces are distinct one defines and computes invariants. Invariants are usually algebraic objects, such as numbers or groups, that do not change under any deformation. The sphere and the torus are distinguished by the fundamental group. Any loop on the sphere can be contracted along the surface to a point, but this is impossible for some loops on the torus. The fundamental group describes exactly which loops are contractible and which are not. Because the fundamental groups of the sphere and the torus differs, it follows that these spaces are inequivalent. The Euler characteristic is another invariant which distinguishes the spaces. The Euler characteristic of a 2-dimensional space is computed as follows

$$\chi(M) = \text{number of triangles} - \text{number of edges} + \text{number of vertices.} \quad (1.1)$$

where the numbers are computed for a triangulation of the space, see Figure 1.2. For the sphere one computes, independent of the chosen triangulation, that the Euler characteristic is 2, while for the torus the Euler characteristic is 0, which also shows that the spaces differ.




Figure 1.2: Triangulations of the sphere and the torus. From a triangulation the Euler characteristic can be computed, cf. Equation (1.1). The Euler characteristic does not depend on the triangulation, but only on the surface. The sphere has Euler characteristic 2 and the torus has Euler characteristic 0.

1.1.2 Topology and dynamical systems

Many phenomena in mathematics, physics, biology, chemistry and the social sciences are described by differential equations. A differential equation relates a quantity with its various derivatives. For example the rate of growth of a population is proportional to the size of the population, and the acceleration experienced by the moon is related to the distance of the moon to the earth. Differential equations describing evolution in time lead to dynamical systems.

The study of (algebraic) topology has its origins in dynamical systems theory [Poi10]. Any non-trivial dynamical system is too complex to understand in full detail. The exact evolution of a dynamical system is extremely sensitive to initial conditions and model parameters. We do not pretend we can measure the model parameters and initial conditions exactly, so it is important to understand which part of the dynamical behavior is preserved under small perturbations of the system in question. Topology is precisely about such robust information. In Part I we define topological invariants for dynamical systems, which capture qualitative information of the gradient behavior of arbitrary flows. In Part II we study a problem in conservative dynamics. Topological tools are used to detect periodic orbits of a specific system.

In this thesis topology is used as a tool to study dynamical systems, but the synergy of topology and dynamical systems goes both ways: well chosen dynamical systems can give deep insight into pure topological problems, see

for example [Bot56, Mil63, Mil65].

1.1.3 Outline of the introduction

We first give some background material for Part I: In Section 1.2 we discuss classical Morse theory and Morse homology in Section 1.3. Conley theory is reviewed next in Section 1.4. We then proceed, using the notion of local Morse homology of Section 1.5 to Section 1.6 to the the the first result of the thesis, which is the definition of Morse-Conley-Floer homology. Functorial properties and applications of Morse-Conley-Floer homology are given in Sections and 1.7 and 1.8.

Part II is concerned with the search of periodic orbits of Hamiltonian systems on non-compact energy hypersurfaces. We briefly discuss the setting and history of Hamiltonian dynamics and the Weinstein conjecture in Section 1.10 after which we describe the contribution of this thesis to this problem in Section 1.11.

1.2 Classical Morse theory, the half-space approach

Morse realized [Mor25] that there are strong connections between the topology of a space and the number and type of critical points of functions defined on this space. Probably the best exposition of the classical part of the theory is the treatise by Milnor [Mil63], and for the more topologically inclined there is [Mat02]. A non-technical introduction are the wonderful lectures by Bott [Bot82].

1.2.1 Attaching a handle

The relation between the critical points of a function and the topology its domain is best shown by an example, so consider the landscape in Figure 1.3. The landscape is the graph of a function $f : M = [0, 1]^2 \rightarrow \mathbb{R}$, which returns to each point the height the terrain. The topology of the sublevel set $f_c = \{p \in M \mid f(p) \leq c\}$ changes when c passes a *critical value*. A value c is critical if there exists a point $x \in M$ such that $df(x) = 0$ and $f(x) = c$. Here the critical values are the heights of the peaks, passes and valleys. The changes in topology can be extremely complicated, but if one assumes that f is a Morse function the changes can be understood relatively easily.




Figure 1.3: The island in the top left corner is the graph of a function $f : [0, 1]^2 \rightarrow \mathbb{R}$. On the side sublevel sets $f^c = \{p \in [0, 1]^2 \mid f(p) \leq c\}$ are depicted, for various values of c . The topology of the sublevel changes each time the c passes a critical value. In this particular case there are 7 critical values, which are the values of the valleys, passes and peaks.

Definition 1.2.1. A function $f : M \rightarrow \mathbb{R}$ is *Morse* if all its critical points are non-degenerate. That is if $x \in M$ is a critical point, i.e. $df(x) = 0$ then the Hessian¹ is non-singular, i.e. $\det \text{Hess } f \neq 0$. The number of negative eigenvalues of the Hessian at x is the *Morse index* of x and is denoted $|x|$.

Assuming that the function f is Morse, and if $f^{-1}([c - \epsilon, c + \epsilon])$ is compact, and the only critical value in $[c - \epsilon, c + \epsilon]$ is at c and comes from a single critical point x , then the sublevel set $f_{c+\epsilon}$ can be obtained from the sublevel set $f_{c-\epsilon}$ by attaching the space $D^{|x|} \times D^{n-|x|}$, where $D^{|x|} = \{p \in \mathbb{R}^{|x|} \mid |p| \leq 1\}$.

¹The Hessian can be defined with the help of a metric and its Levi-Civita connection as $\text{Hess } f(X, Y) = \langle \nabla_X \nabla_Y f, Y \rangle$. At critical points the Hessian does not depend on the choice of metric.

Such a space is referred to as an $|x|$ -handle. The Morse index $|x|$ captures the type of the critical point i.e., if its a minimum $|x| = 0$, maximum $|x| = \dim M$ or a saddle $0 < |x| < \dim M$. The attachment of the handles is shown in Figure 1.3. The technical heart of the construction, the attachment of the $|x|$ -handle, is done with the help of dynamical systems, namely the gradient flow², generated by $x' = -\nabla_g f(x)$ and its renormalized variant generated by $x' = -\frac{\nabla_g f(x)}{|\nabla_g f(x)|^2}$, for some auxiliary metric g . These dynamical systems will play a great role for Morse homology, cf. Section 1.3.

1.2.2 Globalization

One can now build up the domain of f as follows. If we start from a very small value c –below the minimum of f – and we go up past the maximum of f , we get a nice filtration of spaces

$$\emptyset = f_{c_1} \subset f_{c_2} \subset \dots \subset f_{c_i=\max f} = M. \quad (1.2)$$

where c_i are the critical values³ of f . The filtration is also shown in Figure 1.3. We understand the topological change at each critical level c_i well: it is given by the attachment of a $|x_i|$ -handle, where x_i is the critical point of f such that $f(x_i) = c_i$. Starting from nothing, by attaching handles at each filtration step, in the end we obtain the full domain. From this procedure it now follows that topological complexity of a manifold forces Morse functions to have critical points, because the complexity must be captured by the filtration in Equation (1.2). A precise statement are the Morse inequalities, which relate the number of critical points c_k of index k , to the Betti numbers b_k , which characterize the topology of the domain, cf. [Hat02]. For an island in the sea (which has a connected coast) as in Figure 1.3, the Morse relations simply state

$$\text{number of tops} - \text{number of passes} + \text{number of valleys} = 1. \quad (1.3)$$

This number equals the Euler characteristic of a disc, cf. Equation (1.1). It is convenient to express the general Morse inequalities as a formal equality between polynomials.

²A flow is the solution operator of an autonomous differential equation satisfying some technical hypotheses. It is a map $\phi : \mathbb{R} \times M \rightarrow M$ satisfying $\phi(x, 0) = x$ and $\phi(t, \phi(s, x)) = \phi(t + s, x)$ for all $x \in M$ and $s, t \in \mathbb{R}$

³Here we ignore the issues with the boundary of the domain in Figure 1.3. With some work these can be incorporated in this description.

Theorem 1.2.2 (Morse Inequalities). *Let $f : M \rightarrow \mathbb{R}$ be a Morse function on a closed manifold M . Let b_k be Betti numbers of the domain, and c_k the number of critical points of index k . Denote by $P_t = \sum_k b_k t^k$ the Poincaré polynomial of M , and by $M_t = \sum_k c_k t^k$ the Morse polynomial. Then there exists a polynomial Q_t with non-negative coefficients such that*

$$M_t = P_t + (1 + t)Q_t.$$

Relation (1.3) follows by taking $t = -1$. We see that the topology of the domain forces critical points. This is one way to use the Morse inequalities, but the relations go both ways. Sometimes Morse theory is the only way to get a handle on a manifold, and deep topological facts have been proved using Morse theory [AB83, Bot56, Mil63, Mil65].

1.3 Morse homology

Morse homology is an orthogonal description of the changes of the topology of the sublevel sets when a function passes a critical value. We do not attempt to give a full treatment of Morse homology, since good expositions exist elsewhere. A standard reference is Schwarz' book [Sch93], but the treatments of Audin and Damian [AD14] and Banyaga and Hurtubise [BH04] are enjoyable. In this summary, and in Chapter 2, we closely follow the dynamical systems oriented approach described in Weber [Web06]. A part of the early history of the homological viewpoint of Morse theory is recalled colorfully by Bott [Bot88].

1.3.1 The gradient flow

We remarked before that the gradient flow

$$x' = -\nabla_g f(x). \tag{1.4}$$

plays a crucial role in the heart of classical Morse theory as it allows one to attach handles. Morse homology studies this equation more directly. The dynamics described by this equation is that of steepest descent of the graph of f : imagine a hiker walking around on the island drawn by the graph of f , cf. Figure 1.3, who tries to find her home in the valley (the minimum of f). The hiker can achieve her goal by walking in the steepest direction downhill.




Figure 1.4: The gradient flow of the function whose graph is shown in Figure 1.3. The orbits tend to critical points in forward and backward time. Because this domain has a boundary, orbits can also escape to the boundary, which does not happen on closed manifolds.

Measuring steepness is done with the auxiliary metric g . Depending on the starting point, the hiker will most likely walk to the valley or to the beach. It is possible, but unlikely, to arrive at the passes this way, by starting at a non-generic point. To reach a top, the hiker only needs to climb uphill in the steepest direction. In Figure 1.4 we depicted the gradient flow, i.e. the path the hiker traces, corresponding to the function of Figure 1.3. The solutions of the gradient system (1.4) have the following very nice properties.

Proposition 1.3.1. *Let M be a closed manifold, $f : M \rightarrow \mathbb{R}$ a Morse function, and g a metric. For each $p \in M$, let γ_p be the solution of (1.4) with $\gamma_p(0) = p$. Then the solution exists for all time, is unique and the limits $x = \lim_{t \rightarrow -\infty} \gamma_p(t)$, and $y = \lim_{t \rightarrow \infty} \gamma_p(t)$ exist and are critical points of f . The function f decreases along the flow*

$$\frac{d}{dt}f(\gamma_p(t)) \leq 0,$$

with equality if and only if p is a critical point of f .

Hence to locate one critical point of f , one can start anywhere on the manifold, follow the gradient flow to arrive at a critical point. Generically one arrives at a minimum (or in backwards time) at maximum of f , as mentioned before, it is possible to end up at a saddle.

1.3.2 Counting rigid solutions

In a seminal paper Witten [Wit82] showed how a large part of the topology of the domain can be described by counting only a finite number of special, so called rigid, solutions to (1.4). This is not directly possible for all systems of the form (1.4), one needs to be in a generic situation. To obtain rigid counts one extra assumption is needed: the Morse-Smale condition, which we explain now. Denote by ψ the flow of System (1.4). Given a critical point of f , consider the stable and unstable manifolds

$$W^u(x) = \{p \in M \mid \lim_{t \rightarrow -\infty} \psi(t, p) = x\},$$

$$W^s(x) = \{p \in M \mid \lim_{t \rightarrow \infty} \psi(t, p) = x\}.$$

If f is Morse, the stable and unstable manifolds are embedded open discs of dimension equal to the index $|x|$ and co-index $n - |x|$ of x respectively, see for example [BH04, Theorem 4.2]. Given two critical points x and y of f , the intersection $W(x, y) = W^u(x) \cap W^s(y)$ parametrizes all orbits connecting x and y . We want that these moduli spaces to be manifolds of the “right” dimension. This can be achieved by demanding that all stable and unstable manifolds intersect transversely, a notion introduced by Smale [Sma67].

Definition 1.3.2 (The Morse-Smale condition). We say that (f, g) is a *Morse-Smale pair* if f is a Morse function, g is a metric such that for all critical points x, y of f the unstable and stable manifolds intersect transversely, i.e. for all points $p \in W(x, y)$

$$T_p W^u(x) + T_p W^s(y) = T_p M.$$

The corresponding flow is said to be a *Morse-Smale flow*.

At first this condition seems hard to satisfy, as we need to do this for all unstable and stable manifolds at the same time. In fact it is not: it is a generic property. The space of all Morse-Smale pairs contains a countable intersection of open and dense subsets in the space of all pairs of smooth functions and

metrics, hence is dense, cf. [BH04, Theorem 6.6] or the original papers [Sma63, Kup63]. If the Morse-Smale condition is satisfied, it is not hard to see that $W(x, y)$ is a manifold of dimension $\dim W(x, y) = |x| - |y|$. The space $W(x, y)$ parameterizes the parameterized orbits and carries a proper and free action of \mathbb{R} by time reparameterizations. Hence one can quotient out this action to obtain the moduli space of unparameterized orbits $M(x, y) = W(x, y)/\mathbb{R}$, whose dimension is $\dim M(x, y) = |x| - |y| - 1$. If $|x| = |y| + 1$, the space is zero dimensional, and equals a finite number of points by the compactness of M . Orbits of this type, connecting index difference one points, do not come in families, they are isolated. Therefore these orbits are also said to be *rigid*. The unstable manifolds are embedded discs, hence orientable. A choice of orientation of all unstable manifolds gives an orientation of the spaces $W(x, y)$ and $M(x, y)$. Thus one can count with orientation sign the number of elements in $M(x, y)$ if $|x| = |y| + 1$, and we denote this count⁴ $n(x, y)$.

We finally are in a position to define Morse homology. Let \mathfrak{o} be a choice of orientation of all unstable manifolds, and denote by $C_k(f)$ the free module generated by the critical points of index k of f . Define the differential $\partial_k = \partial_k(f, g, \mathfrak{o}) : C_k(f) \rightarrow C_{k-1}(f)$ on the generators by $\partial_k(x) = \sum_{|y|=|x|-1} n(x, y)y$. The following theorem is true:

Theorem 1.3.3.

$$\partial_{k-1}\partial_k = 0.$$

The proof is a marvelous mix of dynamical systems theory and differential geometry. It shows that the one dimensional moduli spaces of unparameterized orbits $M(x, z)$ which connect critical points with index difference two ($|x| - |z| = 2$) can be compactified by adjoining broken orbits in $M(x, y)$ and $M(y, z)$ with $|y| = |x| - 1$. Figure 1.5 shows this process. Because $\partial^2 = 0$ it makes sense to define the Morse homology as $HM_k(M, f, g, \mathfrak{o}) = \ker \partial_k / \text{im } \partial_{k+1}$. The precise relation with the topology of the domain is now the following theorem.

Theorem 1.3.4. *If M is closed, then the Morse homology does not depend on f, g and \mathfrak{o} , and*

$$HM_*(M, f, g, \mathfrak{o}) \cong H_*(M, \mathbb{Z}),$$

where the right hand side is the singular homology of the domain.

⁴One can also count this number over $\mathbb{Z}/2\mathbb{Z}$ by forgetting the orientations. This is usually easier, but contains less information. The boundary operator then counts if there is an even or odd number of connecting orbits.




Figure 1.5: The reason why the boundary map squares to zero. For $|x| = |z| - 2$, the moduli space $M(x, z) = W(x, z)/\mathbb{R}$ is one dimensional. In the picture this moduli space can be identified with one of the lines orthogonal to the orbits. The non-compactness of this space comes from breaking orbits, in the picture at y and y' . Adding all those broken orbits turns $M(x, z)$ into a compact manifold with boundary. The count ∂^2 exactly counts the boundary components with sign, which must be zero.

The Morse inequalities, Theorem 1.2.2 are now an immediate corollary.

1.4 Conley theory

Morse theory strongly uses the gradient flow (1.4) to study the function f . Some of the ideas of Morse theory can be applied to study general dynamical systems⁵. Conley first realized this, and an overview of the theory is described in his monograph [Con78]. We focus on two aspects of his theory: the index and the idea that dynamical behavior splits in two distinct types of behavior, the gradient part, and the recurrent part.

⁵Here we restrict our attention to differentiable dynamical systems on a finite dimensional domain, i.e. those dynamical systems generated by an ordinary differential equation. Iterations of maps, and more general differential equations also are approachable through Conley theory, but this would take us too far afield.

1.4.1 The Conley index

Consider again the gradient flow of a Morse function. Around an equilibrium x , we understand what happens in the flow, cf. Figure 1.6. The flow is repelling in $|x|$ directions, where $|x|$ is the index of the critical point x , and attracting in $n - |x|$ directions. We can measure the index locally by the number of negative eigenvalues, counted with multiplicity, of the Hessian of f at x . However we can also measure the index of x from through a less local procedure. Take a neighborhood of the equilibrium and keep track of which points flow inward and which points flow outward at the boundary of this neighborhood. If one collapses the points which flow outward, one obtains a space –the Conley index of x – which has the pointed homotopy type of a sphere $S^{|x|}$ with a base point, cf. Figure 1.6. The dimension of this sphere determines the index of the critical point. One can now see that the index does not change if the flow is perturbed inside the neighborhood⁶. This measurement also works if the critical point is degenerate, where the linearization does not give an answer! Consider the perturbation of the monkey saddle, cf. Figure 1.7. The Conley index of the monkey saddle is the wedge of two circles. For all practical purposes, the Conley index counts the monkey saddle as *two* critical points of index one. Indeed, a small perturbation of the monkey saddle creates two equilibria, both of index 1. So far we discussed the index of an equilibrium a gradient flow. However, this can also be done in a general flow which is not necessarily of gradient type. In a more general dynamical system $x' = X(x)$, generated by any vector field X , one is not only interested in equilibria, but also for example in periodic orbits, attractors and repellers, which did not exist for the gradient flow. These are all examples of *invariant sets*, sets of initial conditions which are mapped to themselves by the flow ϕ . These invariant sets can be extremely complicated and fractal –for example, any closed set on the real line is an invariant set of a *gradient flow*– and invariant sets behave poorly under perturbations. However, the class of isolated invariant sets has a much better behavior.

Definition 1.4.1. A compact set $N \subset M$ is an *isolating neighborhood* of the flow if its *maximal invariant set*

$$\text{Inv}(N, \phi) := \{p \in M \mid \phi(t, p) \in N \text{ for all } t \in \mathbb{R}\},$$

is contained in the interior of N . An invariant set S is an *isolated invariant set* if there exists an isolating neighborhood N such that $\text{Inv}(N, \phi) = S$.

⁶The precise invariance is captured in the notion of continuation, cf. Definition 2.5.4.




Figure 1.6: On the left the gradient flow of a Morse function around an equilibrium is depicted. The Conley index is obtained from an index pair (N, L) by collapsing L to a point, which is shown to the right. The space obtained is a thickened sphere, which is homotopy equivalent to the sphere $S^{|x|}$, with $|x|$ the index of the critical point.

An *index pair* for an isolated invariant set S is roughly a pair (N, L) of spaces such that N is an isolating neighborhood, and L is an *exit set*. Any point in N that leaves N must do so through L , and any point in L that leaves N does so without passing through $N \setminus L$. Conley defined his index as the (pointed) homotopy type of an index pair, by collapsing the exit set to a point. He showed the following remarkable theorem.

Theorem 1.4.2. *Let S be an isolated invariant set of the flow ϕ .*

- *Then S admits an index pair (N, L) .*
- *The Conley index –the pointed homotopy type $(N/L, [L])$ – is independent of the choice of index pair.*
- *The Conley index is invariant under continuation of the flow ϕ . That is if ϕ_λ with $\lambda \in [0, 1]$ is a homotopy of flows such that N is an isolating neighborhood for all flows ϕ_λ simultaneously, then the Conley index of $S_0 = \text{Inv}(N, \phi_0)$ and $S_1 = \text{Inv}(N, \phi_1)$ are the same.*
- *If an isolated invariant set is the empty set, the Conley index is trivial.*




Figure 1.7: The Conley index can capture the dynamical behavior of degenerate critical points (and much more). On the left the gradient flow of a monkey saddle $f(x, y) = x^3 - 3xy^2$ is depicted. The Conley index has the homotopy type of the wedge of two circles. If one perturbs the function, as depicted on the right, the degenerate critical point x bifurcates into two critical points of index one: y and z . The Conley index of the perturbed system is the same as the Conley index of the unperturbed system.

Thus isolated invariant sets have a good behavior under perturbations as they leave the index invariant. The sets themselves can still be very complicated. However, since it is invariant under such a large class of perturbations/deformations, it is possible to compute the Conley index in many non-trivial situations, by continuing it to a systems that is much better understood. If the isolated invariant set is non-trivial, i.e. its Conley index is non-trivial, this must also be true for the original system of interest. This principle can be used to detect isolated invariant sets in non-trivial situations.

One of the most useful properties of the gradient flow is that the function f strictly decreases along non-constant solutions, cf. Theorem 1.3.1. Of course such a function cannot exist for any flow. For instance a periodic orbit is an obstruction for such a function to exist.

However, a flow does display a form of gradient-like behavior locally around isolated invariant sets. This is described by Lyapunov functions.

Theorem 1.4.3. *Let S be an isolated invariant set of a flow ϕ on a manifold M . Then*

S admits a Lyapunov function, i.e. there exist a smooth function $f_\phi : M \rightarrow \mathbb{R}$,

(i) $f_\phi|_S = \text{constant}$.

(ii) $\frac{d}{dt}f_\phi(\phi(t, p)) < 0$ for all $p \in N \setminus S$, for some isolating neighborhood N .

The theorem states that locally outside the isolated invariant set the flow is simple as it displays gradient-like behavior. Understanding in general what part of the flow is gradient-like, and what part is not is a major theme of Conley theory. The behavior of Lyapunov functions around the isolated invariant sets determines an invariant for S . This is the Morse-Conley-Floer homology, which is the content of Chapter 2, see also Section 1.6.

1.4.2 Recurrent and gradient-like dynamics

A major result due to Conley is the fundamental theorem of dynamical systems⁷. It roughly states that a dynamical system displays two types of behavior: Points that are (chain) recurrent, and all other points, which display a gradient-like behavior. The precise statement would take us too far afield, the interested reader can consult Conley's monograph [Con78], or the review papers of Mischaikow [Mis99] and Mischaikow and Mrozek [MM02]. The chain recurrent set –the set which displays the recurrent behavior– has some undesirable features. The set is poorly behaved under continuation as it is not necessarily an isolated invariant set and may have an infinite number of connected components. Much better behaved are finite approximations of the chain recurrent set: the Morse decompositions.

Definition 1.4.4. Let M be a closed manifold equipped with a flow ϕ . Let (I, \leq) be finite partially ordered set. A collection $\{S_i\}_{i \in I}$ of subsets of M is a *Morse decomposition* if

- (i) Each S_i is an isolated invariant set for the flow ϕ .
- (ii) For each $p \in M \setminus \bigcup_{i \in I} S_i$, there exist $i < j \in I$ such that $\alpha(p) \subset S_j$ and $\omega(p) \subset S_i$.⁸

The S_i 's are *Morse sets*.

⁷So baptized by D. Norton [Nor95].

⁸The α and ω limit set of p are all points that are reached by sequences $\phi(t_k, p)$ with $t_k \rightarrow -\infty$ and $t_k \rightarrow +\infty$ respectively.

Morse decompositions are modeled after gradient flows of Morse functions. If (f, g) is a Morse-Smale pair, then we can set $S_i = \text{Crit}_i f$ or all critical points separately. The orbits of the gradient flow generated by $-\nabla f(x)$ always originate and end at critical points. Moreover, both the index and the value of f strictly decreases in a Morse-Smale system, which are both admissible partial orders of the Morse decompositions. There are many algebraic structures on the set of all possible Morse decompositions [VKM13, KMV13], revealing the possibility or impossibility of the existence of certain connecting orbits.

A connection between Morse decompositions and the Conley index are the Morse-Conley relations [CZ84]. These relations are a vast generalization of the Morse relations.

Theorem 1.4.5. *For an isolated invariant set define the Poincaré polynomial by $P_t(S) = \sum_k b_k(S)t^k$, with $b_k(S)$ the k -th Betti number of the Conley index of S (as a pointed homotopy type). If $\{S_i\}_{i \in I}$ is a Morse decomposition of an isolated invariant set, then there exist a polynomial Q_t with non-negative coefficients such that*

$$\sum_{i \in I} P_t(S_i) = P_t(S) + (1 + t)Q_t.$$

In the example in Figure 1.6 we computed the Conley index of a critical point x of the gradient flow of a Morse function to be $(S^{|x|}, \text{pt})$. Now $H_k(S^{|x|}, \text{pt}) = \mathbb{Z}$ if $k = |x|$ and zero otherwise, hence $P_t(\{x\}) = t^{|x|}$. The set of all critical points is a Morse decomposition for the gradient flow –with respect to any metric–, from which we obtain the classical Morse relations, Theorem 1.2.2.

1.5 Local Morse homology

Let us return to Morse homology. If M is not closed on can still define groups generated by the critical points, but the differential might count an infinite number of connections. Even when the count is finite it need not be the case that $\partial^2 = 0$, see Figure 1.8 for an example. However suppose $N \subset M$ is an isolating neighborhood for the System (1.4), and (f, g) is a Morse-Smale pair on N . This means that all critical points of f in N are non-degenerate, and all intersections occurring in N are transverse. Then one can define the local Morse homology, cf. Section 3.1.1. One takes as generators all critical points




Figure 1.8: On non-compact domains, the relation $\partial^2 = 0$ does not need to hold. In this case ∂^2 of the red point equals the sum of the green points. The orbits seem to break at infinity. The absence of compactness can be understood as a defect, we miss critical points and connecting orbits at infinity. Compactification schemes can solve these issues, but are quite subtle. However, the relation $\partial^2 = 0$ does hold on isolated invariant sets of the flow. In the situation depicted, there is no isolated invariant set containing the red, blue and green points and the connections between them.

in N , and counts intersections of stable and unstable manifolds that occur completely in N . This count is well defined. Then again $\partial^2(f, g, \sigma, N) = 0$, and the local Morse homology is invariant under isolation preserving homotopies. The Morse homology recovers exactly the homology of the Conley index of the isolated invariant set contained in N , which follows from Theorem 2.7.1.

Local Morse homology generalizes certain facts known about the Morse Homology for manifolds with boundary. Assuming f is defined on a compact manifold with boundary, and that the gradient of f is transverse to the boundary. It follows that f does not have critical points on the boundary. Let ν be the outward pointing normal. Then M is an isolating neighborhood, and




Figure 1.9: Morse-Conley-Floer homology in action. On the left hand side, the periodic orbit is an isolated invariant set of the flow ϕ . An isolating neighborhood is depicted in gray. Let f_ϕ be the distance to the periodic orbits, which is a Lyapunov function for this isolated invariant set. The gradient flow of f_ϕ is depicted in the middle. This flow is clearly not Morse-Smale, since the Lyapunov function is not even Morse, all points on the circle are critical. On the right hand side the gradient flow of a small perturbation of the Lyapunov function is shown. The local Morse homology is easily computed. The complex has two generators: the critical points x and y with $|x| = 1$ and $|y| = 0$. The differential is zero (counted appropriately with signs). The Morse-Conley-Floer homology is therefore $HI_k(S, \phi) \cong \mathbb{Z}$ for $k = 0$ and $k = 1$ and $HI_k(S, \phi) \cong 0$ otherwise. This is the singular homology of the pair $(S^1 \cup \text{pt}, \text{pt})$, which is the Conley index of the periodic orbit.

$\partial M_- = \{p \in \partial M \mid df(p)v < 0\}$ is an exit set. It follows that

$$HM_*(f, g, \mathfrak{o}, N) \cong H_*(M, \partial M_-).$$

In Chapter 2 the local Morse homology of Lyapunov functions will be used to define invariants for isolated invariant sets of arbitrary flows.

1.6 Morse-Conley-Floer homology

The main result of this thesis is the definition of Morse-Conley-Floer homology, as well as proofs of some of its basic properties, which is the content of Chapter 2. Morse-Conley-Floer homology is an invariant for isolated invariant sets of arbitrary flows on manifolds. The invariant is defined in the spirit

of Morse homology –by counting rigid solutions– and recovers the Conley index on finite dimensional manifolds.

Let us discuss the idea, before stating the main results. The aim is to define a homology theory, similar to Morse homology, by counting rigid solutions of an arbitrary flow on a closed manifold, which relates to topological invariants of the domain. It is not clear how to do this directly. What are the generators of the chain complex? It cannot only be the equilibria, there exist many flows without any equilibria, whose domain has non-trivial topology. Another idea could be to take equilibria and periodic orbits. However, it is well known that on the three sphere S^3 there exist flows without equilibria and periodic orbits [Kup94]⁹ hence this idea does not seem too promising.

Any isolated invariant set admits a Lyapunov function f_ϕ , cf. Theorem 1.4.3. Let S be an isolated invariant set, with isolating neighborhood N on which f_ϕ is Lyapunov. Lyapunov functions are in general very degenerate as they are constant on the full isolated invariant set. A small perturbation f of the Lyapunov function f_ϕ will be Morse, and for a generic metric g , (f, g) is Morse-Smale¹⁰. We denote by $\mathcal{Q}^\alpha = (f^\alpha, g^\alpha, N^\alpha, \sigma^\alpha)$ –a Morse-Conley-Floer quadruple– a choice of such data, where σ^α denotes a choice of orientation of the unstable manifolds in N^α . We now can compute the local Morse homology of the Morse-Conley-Floer quadruple \mathcal{Q}^α , which turns out to be an invariant for (S, ϕ) .

Theorem 1.6.1. *Let \mathcal{Q}^α and \mathcal{Q}^β be Morse-Conley-Floer quadruples for S . Then there is a canonical isomorphism $\Phi_*^{\beta\alpha} : HM_*(\mathcal{Q}^\alpha) \rightarrow HM_*(\mathcal{Q}^\beta)$ induced by a continuation map, which has functorial behavior:*

i) $\Phi_*^{\alpha\alpha} = \text{id}$.

ii) If \mathcal{Q}^γ is another choice of data, then

$$\Phi_*^{\gamma\alpha} = \Phi_*^{\gamma\beta} \circ \Phi_*^{\beta\alpha}.$$

This allows us to define an invariant for (S, ϕ) .

Definition 1.6.2. Let S be an isolated invariant set of the flow ϕ . The Morse-Conley-Floer homology of (S, ϕ) is defined as

$$HI(S, \phi) = \varprojlim HM_*(f, g, N, \sigma),$$

⁹This construction of the Kuperberg plugs generalizes to any three dimensional manifold.

¹⁰Small here is in the sense that the gradient flows are isolated homotopic, cf. Definition 2.4.3.

1. INTRODUCTION

where the inverse limit runs over all admissible quadruples of data and the canonical isomorphisms.

In Chapter 2, we prove the following basic properties of this homology theory:

- (i) Morse-Conley-Floer homology $HI_k(S, \phi)$ is of finite rank for all k and $HI_k(S, \phi) = 0$, for all $k < 0$ and $k > \dim M$.
- (ii) If $S = \emptyset$ for some isolating neighborhood N , i.e. $\text{Inv}(N, \phi) = \emptyset$, then $HI_*(S, \phi) \cong 0$. Thus $HI_*(S, \phi) \neq 0$ implies that $S \neq \emptyset$, which is an important tool for finding non-trivial isolated invariant sets.
- (iii) The Morse-Conley-Floer homology satisfies a global continuation principle. If isolated invariant sets (S_0, ϕ_0) and (S_1, ϕ_1) are related by continuation then

$$HI_*(S_0, \phi_0) \cong HI_*(S_1, \phi_1). \quad (1.5)$$

This allows for the computation of the Morse-Conley-Floer homology in non-trivial examples.

- (iv) Let $\{S_i\}_{i \in I}$, indexed by a finite poset (I, \leq) , be a Morse decomposition for S . The sets S_i are Morse sets and are isolated invariant sets by definition. Then,

$$\sum_{i \in I} P_t(S_i, \phi) = P_t(S, \phi) + (1 + t)Q_t. \quad (1.6)$$

where $P_t(S, \phi)$ is the Poincaré polynomial of $HI_*(S, \phi)$, and Q_t is a polynomial with non-negative coefficients.

- (v) Let S be an isolated invariant set for ϕ and let B be an isolating block for S . Then

$$HI_*(S, \phi) \cong H_*(B, B_-; \mathbb{Z}), \quad (1.7)$$

where $B_- = \{x \in \partial B \mid X(x) \text{ is outward pointing}\}$ ¹¹ and is called the ‘exit set’.

¹¹A vector $X(x)$ is outward pointing at a point $x \in \partial B$ if $X(x)h < 0$, where the function $h : B \rightarrow [0, \infty)$ is any boundary defining function for B_- . An equivalent characterization is $g(X(x), \nu(x)) > 0$, where ν is the outward pointing g -normal vector field on B_- . These conditions do not depend on h nor g .

An isolating block is an index pair which has the structure of a manifold with piecewise smooth boundary. These always exist and we refer to Section 2.2 for the details. Point (v) states that the Morse-Conley-Floer homology is isomorphic to the homology of the classical Conley index. This might be disappointing. However, in view of the success of Morse homology in infinite dimensional settings as discussed in Section 1.9.1, we expect that these ideas generalize to true infinite dimensional settings, where classical Conley theory does not measure all dynamical behavior, cf. Section 1.9.3. Point (iv), in view of the isomorphism of Point (v), are of course the Morse-Conley relations given in Theorem 1.4.5.

1.7 Functoriality for flow maps in Morse-Conley-Floer homology

An account of the history of homology theory is given by Hilton [Hil88]. Originally homology¹² groups were not thought of as *groups*. The information contained in them was captured in numerical invariants. Originally this was the Euler characteristic, defined in Equation (1.1), but later more refined numbers such as the *Betti numbers* and *torsion coefficients* became known. It was Emmy Noether who realized that it is much better to think of the homology in terms of abelian groups. Homology groups of compact manifolds are finitely generated abelian groups, which can be classified by the rank and the invariant factors (describing the torsion subgroup). The isomorphism class of the homology group as a (graded) abelian group is then precisely captured by the Betti numbers the torsion numbers.

From this point of view there is nothing gained by introducing the homology *groups*. However, there is one major advantage of the homology groups. A map $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ between spaces induces a map $h_*^{\beta\alpha} : H_*(M^\alpha) \rightarrow H_*(M^\beta)$ between the homology groups. Moreover, the induced map satisfies the following *functoriality*:

- (i) The identity $\text{id} : M \rightarrow M$ induces the identity on homology $H_*(M) \rightarrow H_*(M)$.
- (ii) If $h^{\gamma\beta} : M^\beta \rightarrow M^\gamma$ is another map then $(h^{\gamma\beta} \circ h^{\beta\alpha})_* = h_*^{\gamma\beta} h_*^{\beta\alpha}$

¹²The statements below are true for various flavors of homology, such as simplicial, singular, and cellular homology if one is precise about the classes of maps that are admissible. We discuss this for singular homology and continuous maps.

In categorical terminology homology is a covariant functor from spaces to (graded) abelian groups. We have a good understanding of the possibility and impossibility of maps between abelian groups, which allows us to put on firm footing our intuition about the existence or non-existence of maps between topological spaces. The existence or non-existence of maps can then be translated back into properties of spaces.

The question then springs to mind in what sense Morse homology, local Morse homology and Morse-Conley-Floer homology are functorial. We answer this question in Chapter 3. The functoriality of Morse homology on closed manifolds is known, see for example [AS10, AZ08, AD14, KM07, Sch93]. The proofs of the functoriality properties proceed through isomorphisms with other homology theories where the functoriality is known. In Chapter 3 we take a different approach by studying appropriate moduli spaces and their compactifications. A Morse datum is $\mathcal{Q} = (M, f, g, \mathfrak{o})$, with a choice of manifold M , Morse-Smale pair (f, g) on M and choice of orientation of the unstable manifolds \mathfrak{o} . Crucial for functoriality are the transverse maps.

Definition 1.7.1. Let $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ be a smooth map. We say that $h^{\beta\alpha}$ is *transverse* (with respect to Morse data \mathcal{Q}^α and \mathcal{Q}^β), if for all $x \in \text{Crit } f^\alpha$ and $y \in \text{Crit } f^\beta$, we have that

$$h^{\beta\alpha}|_{W^u(x)} \pitchfork W^s(y).$$

The set of transverse maps is denoted by $T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$. We denote by $W_{h^{\beta\alpha}}(x, y) = W^u(x) \cap (h^{\beta\alpha})^{-1}(W^s(y))$.

Given Morse data \mathcal{Q}^α and \mathcal{Q}^β the set of transverse maps is residual, cf. Theorem 3.8.1. For a transverse map the moduli space $W_{h^{\beta\alpha}}(x, y)$ is an oriented manifold of dimension $|x| - |y|$. For critical points with index $|x| = |y|$ this is a finite collection of points with orientation signs ± 1 . Denote by $n_{h^{\beta\alpha}}(x, y)$ the count with orientation signs of these points. Define the induced map $h_*^{\beta\alpha} : C_*(f^\alpha) \rightarrow C_*(f^\beta)$ by

$$h_*^{\beta\alpha}(x) = \sum_{|y|=|x|} n_{h^{\beta\alpha}}(x, y)y.$$

The induced map counts points in the domain M^α whose gradient flow originates in negative time from the critical point x , and if one applies the map $h^{\beta\alpha}$, the trajectory goes to the critical point y in forward time. From a compactness and gluing analysis of the moduli space $W_{h^{\beta\alpha}}(x, y)$, not unlike how one proves

that $\partial^2 = 0$, it follows that $h^{\beta\alpha}$ is a chain map:

$$h_{k-1}^{\beta\alpha} \partial_k^\alpha = \partial_k^\beta h_k^{\beta\alpha} \quad \text{for all } k.$$

Thus the induced map descends to a map in homology $h^{\beta\alpha} : HM_*(M^\alpha, f^\alpha, g^\alpha, \mathfrak{o}^\alpha) \rightarrow HM_*(M^\beta, f^\beta, g^\beta, \mathfrak{o}^\beta)$. A further analysis shows that this map is functorial on the homology level, commutes with different choices of Morse data on the manifolds M^α and M^β , and that homotopic maps induce the same map in Morse homology. We get functoriality of Morse homology on closed manifolds.

The statements above can be transferred to local Morse homology under extra hypotheses that maps involved are not only transverse, but also *isolating*. See Definition 3.5.1 and Propositions 3.6.1 and 3.6.3. Unfortunately compositions of isolating maps need not be isolating. However the following class of maps does form a category.

Definition 1.7.2. A smooth map $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ between manifolds equipped with flows ϕ^α and ϕ^β , is a *flow map* if it is proper and equivariant. Thus preimages of compact sets are compact and

$$h^{\beta\alpha}(\phi^\alpha(t, p)) = \phi^\beta(t, h^{\beta\alpha}(p)), \quad \text{for all } t \in \mathbb{R} \quad \text{and } p \in M^\alpha.$$

These maps were introduced by McCord [McC88] to study functoriality in Conley theory. Flow maps are always isolating with respect to well chosen isolating neighborhoods of the flows ϕ^α and ϕ^β and compositions of flow maps are flow maps.

Perturbations of such maps (to a transverse map) induce maps in local Morse homology. Moreover, these can be used to define an induced map in Morse-Conley-Floer homology as well. These maps are functorial and constant under isolating homotopies. This answers the question of functoriality in Morse-Conley-Floer homology.

1.8 A spectral sequence in Morse-Conley-Floer homology

Spectral sequences are approximations of homology groups. In Chapter 5 we show the existence of a spectral sequence in Morse-Conley-Floer homology. We have not pushed for the most general statement, so the results can be improved without much effort. We expect that these ideas, which are in essence those of [Jia99], can be used to great effect in the study of variational problems

in degenerate situations, which warrants further study. For the setting of the chapter, let ϕ be a flow and assume $f_\phi : M \rightarrow \mathbb{R}$ is a smooth function on a closed manifold M . Let (b_l^-, b_l^+) be a finite set of numbers such that

$$b_1^- < b_1^+ < b_2^- < b_2^+ < \dots < b_{n-1}^- < b_{n-1}^+ < b_n^- < b_n^+.$$

Set $N_l = f_\phi^{-1}([b_l^-, b_l^+])$, and assume that these are isolating neighborhoods of the flow ϕ and that f_ϕ is a Lyapunov function for these isolating neighborhoods. Assume that all critical values of f_ϕ lie in $\bigcup_{l=1}^n (b_l^-, b_l^+)$. It follows that $b_1^- < \min f_\phi$ and $b_n^+ > \max f_\phi$. By the compactness of M such a decomposition always exist, for instance for $n = 1$. The point is that a finer decomposition of the critical values will give more information in the spectral sequence. Two cases are of interest: f_ϕ is a Lyapunov function for a Morse decomposition of M , or ϕ is the gradient flow of the –possibly degenerate–function f_ϕ .

Theorem 1.8.1. *There exist a spectral sequence $(E_{k,l}^r, \partial^r)$ such that*

$$E_{k,l}^1 = HI_k(S_l, \phi),$$

and $E_{k,l}^r \Rightarrow H_*(M)$. That is, for each k

$$\bigoplus_l E_{k,l}^\infty \cong H_k(M).$$

As is often the case with spectral sequences: the sequence approximates a known quantity, the homology the base manifold, and the approximating objects are the interesting ones. The relations given by this spectral sequence can be used to give forcing relations. For example the existence of this spectral sequence gives a different proof of the Morse-Conley relations via standard homological algebra, but more information is contained in it. In the case that ϕ is the gradient flow of a Morse-Bott function f_ϕ , see Section 5.3, the first page of the spectral sequence can be computed in terms of critical point data. In this case the spectral sequence (for Morse-Bott functions) is well known, cf. [Fuk96, AB95, Jia99]. This gives the Morse-Bott relations. In the literature these relations are sometimes misstated due to subtle orientation issues. We give a correct statement, and give a counterexample to the misstatements.

1.9 Morse-Conley-Floer homology in infinite dimensional systems

We expect that Morse-Conley-Floer homology can be used to great effect in infinite dimensional problems, which will be the subject of further research. To explain this we first discuss how classical Morse theory can be used in certain infinite dimensional problems, and how homological techniques can be applied to certain situations where the classical approach fails.

1.9.1 Morse Theory in infinite dimensions

Classical Morse theory can be applied to a variety of infinite dimensional problems. We describe one now, which closely relates to Part II of this thesis, see also [Bot82, Kli95, Mil63, PS64, Pal63]. A closed geodesic on a closed Riemannian manifold M is a path $c : S^1 \rightarrow M$ satisfying the equation $\nabla_s c'(s) = 0$. Such closed geodesics are well known to be critical points of the energy functional $\mathcal{E} = \frac{1}{2} \int_0^1 |c'(s)|^2 ds$, defined on the free loop space ΛM of maps $c : S^1 \rightarrow M$ of Sobolev regularity H^1 , cf. [Kli95]. The search for closed geodesics is then equivalent to the search for critical points of this action functional. In the original approach this is done through a finite dimensional approximation, the method of broken geodesics. The free loop space is approximated by a subspace of $M^n = M \times \dots \times M$ for sufficiently large n , and the energy functional by $E(x_1, \dots, x_n) = \sum d(x_i, x_{i+1 \bmod n})^2$. Then each component of $E(x_1, \dots, x_n) < \epsilon < \text{inj } M$ — the approximation of the loop space— will contain a minimum of the energy functional, which corresponds to a closed geodesic. Unfortunately, a priori there is nothing to stop this closed geodesic from being a constant loop. If however the fundamental group of M is non-trivial we do get plenty of non-trivial minima (one for each conjugacy class of the fundamental group). These minima are attained on different connected components of the free loop space. If M is simply connected this reasoning does not work immediately. However in this case Morse theory can be applied to produce a different critical point. Because M is closed some homotopy group of M is non-trivial. The homotopy groups of M and ΛM are closely related, which is used to show that certain sublevel sets of the energy functional are not homotopic. This shows the existence of a critical point. This is the celebrated Theorem of Lyusternik and Fet [LF51].

In hindsight the approximating procedure works in here because of two

important properties of the energy functional. The energy functional has finite index critical points, and it satisfies the Palais-Smale condition, cf. Section 1.11.3. The Palais-Smale condition shows that sequences seemingly tending to critical points actually do converge to critical points, which resolves many compactness issues and shows that the finite dimensional approximation is sensible.

If the functional would have infinite indices, this would be a different problem. The finite dimensional approximations cannot see such critical points. This is closely related to the fact that an infinite dimensional ball can be continuously retracted into its boundary. Infinite index critical points are invisible to homotopies¹³. Fortunately Morse (and Floer) homology can be used to attack these problems.

1.9.2 Morse and Floer homology in infinite dimensions

Morse homological tools also work in cases with infinite indices and co-indices, cf. [AM05, AM08, AM01, Abb97], in situations where there is a natural notion of relative index and metrics that are compatible with the index. These assumptions are in a sense necessary, there is some sort of hard barrier here. It is not that current techniques are not sufficient, but rigidity does not come for free in infinite dimensional settings. Consider the following theorem of Abbondandolo and Majer [AM04].

Theorem 1.9.1. *Let M be an infinite dimensional real Hilbert manifold, g_0 a choice of metric, and $f : M \rightarrow \mathbb{R}$ a smooth Morse function, all of whose critical points have infinite index and co-index. Let $a : \text{Crit } f \rightarrow \mathbb{Z}$ be any function. Then there exists a metric g , uniformly equivalent to g_0 , such that for all $x, y \in \text{Crit } f$, the unstable and stable manifolds intersect transversely, and $W^u(x, f, g) \cap W^s(y, f, g)$ has dimension $a(x) - a(y)$.*

Actually, slightly more is proven, which shows non-triviality of the result. If $\{(x_i, y_i) \in \text{Crit } f \times \text{Crit } f \mid i = 1, \dots, n\}$ is a finite collection of pairs of critical points satisfying $a(x_i) > a(y_i)$, and there are smooth curves $\gamma : [0, 1] \rightarrow M$ connecting the critical points with $Df(u_i)(u'_i) < 0$ on $(0, 1)$, then the metric g can be chosen such that $W^u(x_i, f, g) \cap W^s(y_i, f, g) \neq \emptyset$, for all $i = 1, \dots, n$.

¹³This issue can be partially resolved shrinking the class of admissible homotopies.

This theorem shows that it cannot be true that a Morse homology of f and M exists which does not depend on extra structure. Restricting the class of metrics is in some situations natural.

Of course, historically Morse homologies in infinite dimensional settings came after Floer's seminal papers [Flo88c, Flo89a, Flo88a, Flo88b]. Floer realized that some important functionals in symplectic geometry have infinite index and co-index, but that the moduli spaces of the formal L^2 gradient flow had many properties in common with the finite dimensional case. The L^2 gradient is not natural in this situation, since the functionals he studied were continuous on $H^{\frac{1}{2}}$, but not L^2 (this difference not come up in finite dimensional settings). Choosing the L^2 gradient has a big downside. The formal gradient flow does not pose a good Cauchy initial value problem. However, this approach has a saving grace: Floer's equation is elliptic. This means that many elliptic estimates, regularity and maximum principles can be applied to the study of his equation. It follows, through hard analysis, that the solutions to his equation (perturbed J -homomorphic curves) which exist for all times and have finite energy act in many ways the same as solutions to a true gradient flow in finite dimensions. The critical points are now endowed with relative indices given by the dimensions of the moduli spaces connecting the critical points. The compactness issues of the moduli spaces are (up to bubbling phenomena) similar to the finite dimensional Morse case, which shows that the boundary operator squares to zero. The homology –Floer homology– is shown to be independent of many data that go into definition of the homology theory. This allows one in favorable situations to compute the homology, and deep results are obtained through this theory, which are currently unattainable through other methods.

1.9.3 Strongly indefinite systems and Conley theory

The problems that makes Morse theory hard to apply in infinite dimensions also are problematic for Conley theory. The main tool to study isolated invariant sets –the Conley index– is too flexible to measure isolated invariant sets in the case the exit sets behave poorly without further thought. These are known as strongly indefinite flows, and for gradient flows this means that the function has critical points with infinite indices. In a series of papers M. Izydorek [Izy01, Izy06] developed a Conley index approach for strongly indefinite flows based on Galerkin approximations and a cohomology theory based on spectra. Using the analogues of index pairs he established a coho-

mological Conley for strongly indefinite systems. See also [Man03, Sta14] for approaches to strongly indefinite variational problems via Conley theory.

1.9.4 Morse-Conley-Floer homology in infinite dimensions

This thesis develops an intrinsic approach towards the Conley index in the finite dimensional setting. We expect that these ideas can also be successfully be applied to to study strongly indefinite flows, and the Floer theoretic counterpart, strongly indefinite elliptic problems, by computing the local Morse and Floer homology of suitable Lyapunov functions. It is not clear that such functions always exist, and this will be the subject of further research.

1.10 The Weinstein conjecture and symplectic geometry

After having discussed the main results of Part I we now go to a different topic and introduce the setting of Part II of this thesis. The problem is the existence of periodic orbits in Hamiltonian dynamics on non-compact energy hypersurfaces. We show this existence under geometric and topological assumptions in Chapter 6.

It is not hard to see that the linearization of an equilibrium point in a dynamical system has strong connections to the dynamics around the equilibrium. Something similar happens for periodic orbits and the existence of a periodic orbit has strong implications to the dynamics nearby such an orbit. A natural question is therefore: “when does a dynamical system have periodic orbits?”. This problem in full generality is probably unsolvable, and for many dynamical systems there is a negative answer. However, in an important class, Hamiltonian systems, we do expect the existence of periodic orbits. This is basically due to Poincaré’s recurrence theorem, which states that this is always almost true: Almost all orbits are almost periodic. We do not aim to give a full history of the search of periodic orbits in Hamiltonian systems, the interested reader can consult [Pas12], but give a short overview.

Hamiltonian dynamics is a mathematical formulation of frictionless motions in classical mechanics. A good introduction to this subject curiously coincides with the origins of symplectic geometry, this is the book of Arnold [Arn78]. Let X denote a phase space, this is the space of all states of a system, e.g. the positions of particles and their momenta. A Hamiltonian function $H : X \rightarrow \mathbb{R}$ gives the energy of each state. A symplectic form ω (a non-degenerate two-form) on X describes (locally) which variables in X should be thought of as configurations and which variables are their corresponding momenta. The choice of a symplectic form and Hamiltonian determines the Hamiltonian vector field via the equation $i_{X_H}\omega = -dH$, which gives the differential equation

$$x'(t) = X_H(x(t)). \tag{1.8}$$

The flow preserves the volume measured with respect to ω^n . A version of Poincaré’s Recurrence Theorem follows: If all orbits are bounded, almost all (measured with respect to the volume measure) orbits are almost periodic. It is an elementary fact that the motion described by (1.8) preserves the (time independent) Hamiltonian H . Given an initial condition x_0 with $H(x_0) = H_0$, the motion is fixed on the hypersurface $\Sigma = H^{-1}(H_0)$. For regular energy




Figure 1.10: Poincaré's recurrence theorem implies that in compact situations, almost all orbits are almost periodic, cf. the left hand side. However, if the compactness assumption is dropped, this is not necessarily true, cf. the right hand side.

hypersurfaces, i.e. along which dH does not vanish, Poincaré Recurrence can then be improved. There exists a non-trivial invariant volume measure on the hypersurface, cf. [HZ11, Section 1.4]. It follows that almost all orbits on a regular hypersurface are almost periodic if the hypersurface has finite measure. The regularity assumption is crucial here, as can be seen from the example $H = \frac{1}{2}|p|^2 + \frac{1}{2}q^2 - \frac{1}{4}q^4$. That the hypersurface has a finite measure is also crucial. This can fail if Σ is non-compact, cf. Figure 1.10. Intuitively one expects that if Poincaré recurrence holds there also must be periodic orbits, since almost all orbits almost hit themselves. It would be unlikely that all of these near misses do not close up to form periodic orbits. This is not necessarily the case: it is not true that a periodic orbit always exists on a given compact hypersurface Σ , e.g. [Gin95, GG03].

It is good to remark that the precise choice of Hamiltonian is not relevant to the existence of periodic orbits on a given hypersurface. The geometry of Σ inside the symplectic manifold (X, Ω) is sufficient to determine the motion of the system up to time reparametrization. One therefore also speaks of closed characteristics of the hypersurface.

After the first pioneering existence results of Rabinowitz [Rab78] and [Rab79] and Weinstein [Wei78] for star-shaped and convex hypersurfaces respectively, Weinstein formulated his celebrated Weinstein conjecture [Wei79].

Conjecture 1.10.1 (Weinstein Conjecture). *Any compact hypersurface Σ of contact type¹⁴ in (X, ω) , with $H^1(\Sigma) = 0$, admits a periodic orbit.*

The contact-type condition is a generalization of both convex and star-shaped. After the conjecture Viterbo [Vit87] proved the existence of closed characteristics on all compact hypersurfaces of \mathbb{R}^{2n} of so called *contact type*. All these early existence results were obtained by variational methods applied to a suitable (indefinite) action functional. More recently, Floer, Hofer, Wysocki [FHW94] and Viterbo [Vit99], provided an alternative proof of the same results (and much more) using the powerful tools of symplectic homology or Floer homology for manifolds with boundary. The contact-type condition also leads to an intrinsic version of the problem. For hypersurfaces of contact-type, it is possible to throw away all information of the ambient symplectic manifold, and have an intrinsic odd-dimensional version of the conjecture, which concerns periodic orbits of the Reeb vector field on contact manifolds. More recently Taubes showed [Tau09a, Tau09b] a more general form of the conjecture: The intrinsic version of the Weinstein conjecture is true for all closed contact three manifolds, without any assumption on the first homology groups. From another recent result by Borman, Eliashberg and Murphy [BEM14] it follows that any almost contact manifold admits a contact form in the same homotopy class for which the Weinstein conjecture holds true. It is nowadays standard to assume that the homology assumption of the Weinstein Conjecture is superfluous.

One may wonder how natural the assumption that Σ is compact is. Many Hamiltonian systems occurring in nature, e.g. the n-body problem and second order Lagrangians [AVDBV07] have non-compact energy levels. Very little is known about closed characteristics on *non-compact* energy hypersurfaces, as the variational and Floer-theoretic techniques break down if the compactness assumption is dropped. It is also not difficult to construct examples of such systems without any periodic orbits. Additional geometric and topological assumptions are needed in order to make up for the lack of compactness. In [vdBPV09] Van den Berg, Pasquotto, and Vandervorst were able to formulate a set of such assumptions that led to an existence result for the case of *mechanical hypersurfaces* in \mathbb{R}^{2n} , cf. Section 1.11.1. In Chapter 6 we extend these results to mechanical hypersurfaces in cotangent bundles of Riemannian manifolds. In the case of compact mechanical hypersurfaces Bolotin [Bol78], Benci [Ben84], and Gluck and Ziller [GZ83] show the existence of a closed

¹⁴This is a generalization of both star-shaped and convex, see Definition 6.2.1.

characteristic on Σ via closed geodesics of the Jacobi metric on the configuration manifold. A more general existence result for cotangent bundles is proved by Hofer and Viterbo in [HV88] and improved in [Vit97]: Any connected compact hypersurface of contact type over a simply connected base manifold has a closed characteristic, which confirms the Weinstein Conjecture in cotangent bundles of simply connected manifolds. However, the existence of closed characteristics for non-compact mechanical hypersurfaces is not covered by the result of Hofer and Viterbo and fails without additional geometric conditions.

1.11 Closed characteristics on non-compact hypersurfaces

1.11.1 Mechanical Hamiltonians

An important class of Hamiltonian systems are the *mechanical* systems, which arise naturally. Consider a configuration space M , which can for instance describe the positions of a number of particles, along with their orientations. The phase space T^*M is the cotangent bundle of the configuration space, and describes the possible configurations along with their momenta. The phase space carries a canonical symplectic form Ω , cf. Equation (6.1) and the discussion following it. If the forces exerted on the particles are conservative, the motion is determined by a mechanical Hamiltonian of the form

$$H(q, \vartheta_q) = \frac{1}{2}g_q^*(\theta_q, \theta_q) + V(q).$$

Here g is a choice of metric on TM , containing information about the masses of the particles in the system, and (q, θ_q) with $\theta_q \in T_q^*M$ are elements of the cotangent bundle. The first term is the kinetic energy and the second term is the potential energy.

1.11.2 Main theorem

We show that regular ($dH \neq 0$ on Σ) hypersurfaces of mechanical Hamiltonians are always of contact type with respect to the canonical symplectic form Ω , cf. Theorem 6.2.2, independent of the compactness assumptions of Σ . Under asymptotic regularity assumptions on V , the hypersurface is of *uniform contact type*, cf. Proposition 6.2.3, which is important in the proof of the fol-

lowing existence result¹⁵. In this theorem ΛM denotes the space of all H^1 loops, cf. [Kli95]. The theorem is the main result of Chapter 6.

Theorem 1.11.1. *Let $H : T^*M \rightarrow \mathbb{R}$ be a mechanical Hamiltonian, and assume*

- (i) *(M, g) is an n -dimensional complete orientable Riemannian manifold with flat ends*
- (ii) *The hypersurface $\Sigma = H^{-1}(0)$ is regular, i.e. $dH \neq 0$ on Σ*
- (iii) *The potential V is asymptotically regular, i.e. there exist a compact K and constant $V_\infty > 0$ such that*

$$|\text{grad } V(q)| \geq V_\infty, \text{ for } q \in M \setminus K \text{ and } \frac{\|\text{Hess } V(q)\|}{|\text{grad } V(q)|} \rightarrow 0, \text{ as } d(q, K) \rightarrow \infty.$$

- (iv) *That there exists an integer $0 \leq k \leq n - 1$ such that*

- $H_{k+1}(\Lambda M) = 0$ and $H_{k+2}(\Lambda M) = 0$, and
- $H_{k+n}(\Sigma) \neq 0$.

*Then Σ has a periodic orbit which is contractible in T^*M .*

It should be noted that we expect that some conditions are not necessary. For example the assumptions on the homology of the space of free loops ΛM arises due to the method employed, and the assumption of flat ends is also too strong, cf. Section 1.11.5.

The proof is based on a linking argument for the Lagrangian action functional which is defined on the free loop space. There are two main ingredients of the proof, one is analytical, and another is topological. The topological part of the proof, the linking argument, produces sequences in the loop space that seem to converge to critical points of the action functional. The analytical part of the proof shows that these sequence actually have convergent subsequences. The possible non-compactness of the hypersurface Σ makes both aspects of the proof more technical. We outline the proof now, starting with the analytical aspects.

¹⁵Here and in Chapter we write $\text{grad } f$ for the gradient of f , to relieve some ambiguity of the notation with various covariant derivatives.




Figure 1.11: \mathcal{A} does not need to satisfy the Palais-Smale condition on non-compact domains. For the standard metric on $S^1 \times \mathbb{R}$, the circles c_n are closed geodesics, which are periodic solutions to a Hamiltonian system on its cotangent bundle. All circles have the same action and period τ . The sequence (c_n, τ) does not have a convergent subsequence.

1.11.3 Critical Points, Penalization and The Palais-Smale Condition

A periodic orbit $q : [0, T] \rightarrow M$ with period T on the hypersurface $\Sigma = H^{-1}(0)$ is in one to one correspondence to a critical point of the free period action functional

$$\mathcal{A}(c, \tau) = \frac{e^{-\tau}}{2} \int |c'(s)|^2 ds - e^\tau \int V(c(s)) ds, \quad (1.9)$$

via the coordinate transformation

$$(c(s), \tau) \mapsto (q(sT), \log(T)).$$

The domain of \mathcal{A} is $\Lambda M \times \mathbb{R}$, where ΛM is the space of H^1 loops in M along with a period parameter in \mathbb{R} .

Standard techniques to find critical points of action functionals of the form (1.9) are Morse theoretical ideas, such as mountain pass or linking techniques. These techniques produce Palais-Smale sequences in the loop space¹⁶.

Definition 1.11.2. A *Palais-Smale sequence* for (1.9) is a sequence $\{(c_n, \tau_n)\}_{n \in \mathbb{N}} \in \Lambda M \times \mathbb{R}$ such that

- $|\mathcal{A}(c_n, \tau_n)|$ is bounded,

¹⁶A more ambitious plan would be to develop a full Morse homology for the approximating action functionals, show that this is invariant under the perturbations, computing the homology and conclude that there exist periodic orbits.

- $|d\mathcal{A}(c_n, \tau_n)| \rightarrow 0$ as $n \rightarrow \infty$.

The hope is that such sequences converge to critical points of the action functional, and the geometrical/topological situation forces the existence of critical points. This need not be the case as is shown in Figure 1.11.

Definition 1.11.3. If a functional has the property that all Palais-Smale sequences¹⁷ have convergent subsequences, the functional satisfies the *Palais-Smale condition*.

Unfortunately if Σ is non-compact, \mathcal{A} need not satisfy the Palais-Smale condition, because sequences of periodic orbits (or almost periodic orbits) may escape to infinity. For example, consider the situation of Figure 1.11. The circles are closed geodesics which are solutions of a natural Hamiltonian system on the cotangent bundle with the same action and period τ . Hence $|\mathcal{A}(c_n, \tau_n)|$ is bounded, and $d\mathcal{A}(c_n, \tau) = 0$. The sequence does not have a convergent subsequence as the geodesics escape to infinity.

However, under the geometrical assumptions on the potential and the geometry of the manifold M , the following *penalized* functionals

$$\mathcal{A}_\epsilon(c, \tau) = \mathcal{A}(c, \tau) + \epsilon(e^{-\tau} + e^{\frac{\tau}{2}}),$$

for $\epsilon > 0$ do satisfy the Palais-Smale condition, cf. Lemma 6.4.5. Moreover there is a selection mechanism. If $(c_\epsilon, \tau_\epsilon)$ is family of critical points of \mathcal{A}_ϵ that is a priori bounded away from zero, i.e. there exists positive constants a_1, a_2 such that

$$0 < a_1 < \mathcal{A}_\epsilon(c_\epsilon, \tau_\epsilon) < a_2 < \infty,$$

then the family of critical points of the penalized functionals has a subfamily converging to a critical point of \mathcal{A} , which is the content of Proposition 6.4.8.

Therefore we can use the topological techniques to find critical points of the action functional \mathcal{A} , and hence find closed characteristics on Σ .

1.11.4 Candidate Critical Points, the Linking Theorem

We construct Palais-Smale sequences by studying the geometry of the functional on the loop space. We construct disjoint sets $A, B \subset E$, and constants $a < b$, which satisfy the bounds $\mathcal{A}|_A < a$, and $\mathcal{A}|_B \geq b > 0$. By Lemma 6.8.4

¹⁷This also depends on a metric, which we have suppressed here because we will only deal with a fixed metric.




Figure 1.12: A sketch of the link in the loop space. The base manifold is embedded in the loop space as the constant loops, around which we construct linking sets.

the sets topologically link which means that they are entangled in the loop space. If $\epsilon > 0$ is sufficiently small the bounds are also hold for the penalized functionals \mathcal{A}_ϵ , and by standard linking theorems this produces Palais-Smale sequences for all penalized functionals, and the resulting critical values satisfy uniform bounds, which finally gives a critical point of \mathcal{A} . The linking theorems are generalizations of the mountain pass lemma, which roughly states that if one wants to go from one valley to another valley in a mountainous landscape, one has to cross a mountain pass.

How does one get the sets A and B ? This comes from understanding of the geometry of the functional of the loop space. The (unpenalized) functional (1.9) has two terms, the kinetic energy term $e^{-\tau} \mathcal{E}(c) = e^{-\tau} \|\partial c\|_{L^2}$, which is always positive for non-constant loops, and the potential energy term $-e^\tau \mathcal{W} = -e^\tau \int_0^1 V(c(s)) ds$ which can carry a positive or negative sign, depending on the loop, and the choice of potential. Roughly if a loop c stays a long time in a region where the potential is positive, $-e^\tau \mathcal{W}(c)$ is negative, and vice-versa.

We assume that the hypersurface Σ has a certain non-vanishing homology group in one of the *top half degrees*. This is strongly related to the geometry of the shadow $N = \pi(\Sigma)$ of Σ , that is, its projection to the base manifold M , which is described in Section 6.7. The shadow can also be described as all points where the potential is negative

$$N = \{q \in M \mid V(q) \leq 0\}.$$

The potential is in a linking situation. The link on the base manifold, for the potential V , is then lifted to a link for the action functional \mathcal{A} in the loop space $\Lambda M \times \mathbb{R}$. Morally set the A contains almost constant loops, so the kinetic energy $e^\tau \mathcal{E}$ is small, located close to N , such that $-c^\tau \mathcal{W}$ is negative. The set B consists of fast moving loops, so that the kinetic energy is large, which only are at N for a very short time, so that the potential energy term is positive.

In the end the linking construction, along with the analysis of the Palais-Smale condition, produces a critical point of \mathcal{A} , which is a periodic orbit of the Hamiltonian flow on Σ .

1.11.5 Current Status

A preprint appeared by Zehmis and Suhr [SZ13] that improves Theorem 6.1.1. They roughly follow same method of proof, but weaken the hypothesis of flat ends to bounded geometry, and show that the conditions on the topology of the loop space are unnecessary.

The question of existence of non-contractible loops on non-compact hypersurfaces remains open. Another interesting problem is to determine the existence of periodic orbits on hypersurfaces for Hamiltonian systems which are not mechanical. We think that the current methods that we have applied (linking and minimax) have reached their maximal potential. More abstract machinery: such as symplectic capacities, Morse homological tools and Rabinowitz-Floer homology might give more satisfying answers. Currently these methods cannot be directly applied to the non-compact situation but the penalization method might resolve some compactness issues that occur in the non-compact situation.

Part I

Morse-Conley-Floer Homology

Chapter 2

Morse-Conley-Floer homology

Is there a thing of which it is said,
"See, this is new"?
It has been already
in the ages before us.

Ecclesiastes

2.1 Introduction

In this chapter we define an analogue of Morse homology for isolated invariant sets of smooth, not necessarily gradient-like, flows. We first briefly recall how the gradient flow of a Morse function gives rise to Morse homology.

2.1.1 Morse Homology

On a smooth, closed, m -dimensional manifold M the *Morse-Smale* pairs (f, g) , consisting of a smooth function $f : M \rightarrow \mathbb{R}$ and a Riemannian metric g , are defined by the property that all critical points of f are non-degenerate and all their stable and unstable manifolds with respect to the negative g -gradient flow intersect transversally. For a given Morse-Smale pair (f, g) , the *Morse-Smale-Witten chain complex* consists of free abelian groups¹ $C_k(f)$ generated by

¹We define the Morse homology with coefficients in \mathbb{Z} . Other coefficient fields are also possible, but can also be obtained from coefficients in \mathbb{Z} by the Universal Coefficient Theorem. We drop the coefficient group from our notation.

the critical points of f of index k , and boundary operators

$$\partial_k(f, g, \mathfrak{o}) : C_k(f) \rightarrow C_{k-1}(f),$$

which count the oriented intersection between stable and unstable manifolds of the negative g -gradient flow of f . The Morse-Smale-Witten complex is a chain complex and its homology is the *Morse homology* $HM_*(f, g, \mathfrak{o})$ of a triple (f, g, \mathfrak{o}) — called a *Morse-Smale-Witten triple* —, where \mathfrak{o} is a choice of orientations of the unstable manifolds at the critical points of f . Between different triples $(f^\alpha, g^\alpha, \mathfrak{o}^\alpha)$ and $(f^\beta, g^\beta, \mathfrak{o}^\beta)$ there exists canonical isomorphisms $\Phi_*^{\beta\alpha} : HM_*(f^\alpha, g^\alpha, \mathfrak{o}^\alpha) \rightarrow HM_*(f^\beta, g^\beta, \mathfrak{o}^\beta)$ induced by the continuation map. This defines an inverse system. The Morse homology of the manifold is then defined by

$$HM_*(M) := \varprojlim HM_*(f, g, \mathfrak{o}),$$

where the inverse limit is taken over all Morse-Smale-Witten triples (f, g, \mathfrak{o}) , with the canonical isomorphisms. A similar construction can be carried out for non-compact manifolds, using Morse functions that satisfy a coercivity condition, cf. [Sch93]. For closed manifolds there exists an isomorphism to singular homology:

$$HM_*(M) \cong H_*(M; \mathbb{Z}), \tag{2.1}$$

cf. [Sal90, Sch93, BH04].

The above results still apply if we consider compact manifolds with boundary, for which $df(x)v \neq 0$, for all $x \in \partial M$, where v is an outward pointing normal on the boundary.² This implies that f has no critical points on the boundary. The boundary splits as $\partial M = \partial M_- \cup \partial M_+$, where ∂M_- is the union of the components where $df(x)v < 0$ and ∂M_+ is the union of the components where $df(x)v > 0$. In this case the Morse homology can be linked to the singular homology of M as follows

$$HM_*(M) \cong H_*(M, \partial M_-; \mathbb{Z}), \tag{2.2}$$

cf. [KM07], [Sch93]. When $\partial M_- = \emptyset$, i.e. M has no boundary or $df(x)v > 0$ for all $x \in \partial M$, we have $HM_*(M) \cong H_*(M; \mathbb{Z})$. The classical Morse relations/inequalities for Morse functions are an immediate corollary. The isomorphism in (2.2) also holds in the more general setting when the boundary

²The outward pointing normal is defined as $\nu = -\frac{\nabla_g h}{|\nabla_g h|_g}$, where $h : M \rightarrow [0, \infty)$ is smooth boundary defining function with $h^{-1}(0) = \partial M$, and $dh|_{\partial M} \neq 0$.

allows points where $df(x)v = 0$, with the additional requirement that such points are 'external tangencies' for the negative gradient flow. The latter can also be generalized to piecewise smooth boundaries, cf. Section 2.7.

2.1.2 Morse-Conley-Floer Homology

For arbitrary flows an analogue of Morse homology can be defined. Let M be a, not necessarily compact, smooth m -dimensional manifold without boundary. A smooth function $\phi : \mathbb{R} \times M \rightarrow \mathbb{R}$ is called a *flow*, or \mathbb{R} -*action* on M if:

- (i) $\phi(0, x) = x$, for all $x \in M$ and
- (ii) $\phi(t, \phi(s, x)) = \phi(t + s, x)$, for all $s, t \in \mathbb{R}$ and $x \in M$.

A smooth flow satisfies the differential equation

$$\frac{d}{dt}\phi(t, x) = X(\phi(t, x)), \quad \text{with} \quad X(x) = \left. \frac{d}{dt} \right|_{t=0} \phi(t, x) \in T_x M,$$

the associated vector field of ϕ on M . A subset $S \subset M$ is called *invariant* for ϕ if $\phi(t, S) = S$, for all $t \in \mathbb{R}$. A compact neighborhood $N \subset M$ is called an *isolating neighborhood* for ϕ if $\text{Inv}(N, \phi) \subset \text{int}(N)$, where

$$\text{Inv}(N, \phi) = \{x \in N \mid \phi(t, x) \in N, \forall t \in \mathbb{R}\},$$

is called the *maximal invariant* set in N . An invariant set S for which there exists an isolating neighborhood N with $S = \text{Inv}(N, \phi)$, is called an *isolated invariant set*. Note that in general there are many isolating neighborhoods for an isolated invariant set S . Isolated invariant sets are compact. For analytical reasons, we need isolating neighborhoods with an appropriate manifold structure, in the sense that boundary of such a neighborhood is piecewise smooth and the flow ϕ is transverse to the smooth components of the boundary. Such isolating neighborhoods are called *isolating blocks*, cf. Definition 2.2.1. Every isolated invariant set $S = \text{Inv}(N, \phi)$ admits an isolating block $B \subset N$. Isolating blocks are used to prove the existence of Lyapunov functions.

A smooth *Lyapunov function* for an isolated invariant set S is a smooth function $f_\phi : M \rightarrow \mathbb{R}$, such that $\left. \frac{d}{dt} \right|_{t=0} f_\phi(\phi(t, x)) < 0$ for $x \in N \setminus S$. Denote the set of Lyapunov functions by $\text{Lyap}(S, \phi)$. This set is non-empty, cf. Proposition 2.2.6. If (f, g) is a Morse-Smale pair, with f is an arbitrary small Morse perturbation of a Lyapunov function f_ϕ , then one can define the Morse homology for

2. MORSE-CONLEY-FLOER HOMOLOGY

the quadruple (f, g, N, \mathfrak{o}) , for some choice of orientation \mathfrak{o} of unstable manifolds of the critical points of f in N . The Morse homology $HM_*(f, g, N, \mathfrak{o})$ is independent (up to canonical isomorphisms) of the isolating block, the Lyapunov function, the Morse perturbation, the metric and the chosen orientations, which leads to the definition of the *Morse-Conley-Floer homology* of (S, ϕ) as an inverse limit

$$HI_*(S, \phi) := \varprojlim HM_*(f, g, N, \mathfrak{o}).$$

This is also an invariant for the pair (N, ϕ) , for any isolating neighborhood N for S , if one takes the inverse limit over a fixed isolating neighborhood N . The important properties of the Morse-Conley-Floer homology can be summarized as follows:

- (i) Morse-Conley-Floer homology $HI_k(S, \phi)$ is of finite rank for all k and $HI_k(S, \phi) = 0$, for all $k < 0$ and $k > \dim M$.
- (ii) If $S = \emptyset$ for some isolating neighborhood N , i.e. $\text{Inv}(N, \phi) = \emptyset$, then $HI_*(S, \phi) \cong 0$. Thus $HI_*(S, \phi) \neq 0$ implies that $S \neq \emptyset$, which is an important tool for finding non-trivial isolated invariant sets.
- (iii) The Morse-Conley-Floer homology satisfies a global continuation principle. If isolated invariant sets (S_0, ϕ_0) and (S_1, ϕ_1) are related by continuation, see Definition 2.5.4, then

$$HI_*(S_0, \phi_0) \cong HI_*(S_1, \phi_1). \quad (2.3)$$

This allows for the computation of the Morse-Conley-Floer homology in non-trivial examples.

- (iv) Let $\{S_i\}_{i \in I}$, indexed by a finite poset (I, \leq) , be a Morse decomposition for S , see Definition 2.6.1. The sets S_i are Morse sets and are isolated invariant sets by definition. Then,

$$\sum_{i \in I} P_t(S_i, \phi) = P_t(S, \phi) + (1 + t)Q_t.$$

where $P_t(S, \phi)$ is the Poincaré polynomial of $HI_*(S, \phi)$, and Q_t is a polynomial with non-negative coefficients. These relations are called the Morse-Conley relations and generalize the classical Morse relations for gradient flows.

- (v) Let S be an isolated invariant set for ϕ and let B be an isolating block for S , see Definition 2.2.1. Then

$$HI_*(S, \phi) \cong H_*(B, B_-; \mathbb{Z}), \quad (2.4)$$

where $B_- = \{x \in \partial B \mid X(x) \text{ is outward pointing}\}$ ³ and is called the ‘exit set’.

Note that in the case that ϕ is the gradient flow of a Morse function on a compact manifold, then Property (v) recovers the results of Morse homology, by setting $S = M$. Property (v) also justifies the terminology Morse-Conley-Floer homology, since the construction uses Morse/Floer homology and recovers the classical homological Conley index. In the subsequent sections we construct the Morse-Conley-Floer homology and prove the above properties.

2.2 Isolating blocks and Lyapunov functions

In this section we discuss the existence of isolating blocks and Lyapunov functions.

2.2.1 Isolating blocks

Isolated invariant sets admit isolating neighborhoods with piecewise smooth boundaries known as isolating blocks, see also Figure 2.1.

Definition 2.2.1. An isolating neighborhood $B \subset M$ for ϕ is called a *smooth isolating block* if B is a compact m -dimensional submanifold with piecewise smooth boundary ∂B and the boundary satisfies the following requirements:

- (i) the boundary decomposes as $\partial B = B_+ \cup B_- \cup B_\pm$, with $B_+ \cap B_- = B_\pm$ and $B_- \setminus B_\pm, B_+ \setminus B_\pm$ (when non-empty) are smooth $(m-1)$ -dimensional submanifolds of M ;
- (ii) there exist open smooth $(m-1)$ -dimensional submanifolds D_-, D_+ such that $B_+ \subset D_+, B_- \subset D_-$ and $D_- \cap D_+ = B_\pm$ is a $(m-2)$ -dimensional submanifold (when non-empty);

³A vector $X(x)$ is outward pointing at a point $x \in \partial B$ if $X(x)h < 0$, where the function $h : B \rightarrow [0, \infty)$ is any boundary defining function for B_- . An equivalent characterization is $g(X(x), \nu(x)) > 0$, where ν is the outward pointing g -normal vector field on B_- . These conditions do not depend on h nor g .




Figure 2.1: An isolating block B for an isolated invariant set S . The boundary of B decomposes into $\partial B = B_+ \cup B_- \cup B_\pm$. Here $B_\pm = B_+ \cap B_-$ consists of the four corner points.

- (iii) The flow is transverse to D_\pm , i.e. $\phi \pitchfork D_\pm$, and for any $x \in D_\pm$ there exists an $\epsilon > 0$ such that $\phi(I_\epsilon^\pm, x) \cap B = \emptyset$, where $I_\epsilon^- = (0, \epsilon)$ and $I_\epsilon^+ = (-\epsilon, 0)$.

The sets $B_- \setminus B_\pm$ and $B_+ \setminus B_\pm$ are also called *egress* and *ingress* respectively and are characterized by the property that $X \cdot \nu > 0$ on $B_- \setminus B_\pm$ and $X \cdot \nu < 0$ on $B_+ \setminus B_\pm$, where ν is the outward pointing g -normal vector field on $\partial B \setminus B_\pm$.

Remark 2.2.2. In [WY73], Wilson and Yorke call this concept an isolating block *with corners*. For the sake of brevity we will refer to such isolating neighborhoods as (smooth) isolating blocks.

All isolated invariant sets admit isolating blocks.

Proposition 2.2.3 (Wilson-Yorke [WY73]). *For any isolating neighborhood $N \subset M$ of S , there exists a smooth isolating block $B \subset N$, such that $\text{Inv}(B, \phi) = \text{Inv}(N, \phi) = S$.*

2.2.2 Lyapunov functions

The existence of isolating blocks implies the existence of global Lyapunov functions with special properties with respect to isolated invariant sets.

Definition 2.2.4. A smooth *Lyapunov function* for (S, ϕ) is a smooth function $f_\phi : M \rightarrow \mathbb{R}$ satisfying the properties:

- (i) $f_\phi|_S = \text{constant}$;
- (ii) $\frac{d}{dt}\Big|_{t=0} f(\phi(t, x)) < 0$ for all $x \in N \setminus S$, for some isolating neighborhood N for S .⁴

The set of smooth Lyapunov functions for S is denoted by $\text{Lyap}(S, \phi)$.

In the subsequent arguments we will always assume that the gradient of f defines a global flow. All arguments in the subsequent sections only care about the behavior of f inside N . Without loss of generality we can take f to be constant outside a neighborhood of N . This does not affect the Lyapunov property of f , and it ensures that any gradient of f defines a global flow. If M is closed the gradient of f always defines a global flow.

The set of smooth Lyapunov functions is non-empty, cf. Proposition 2.2.6, and also convex in the following sense.

Lemma 2.2.5. Let $f_\phi^\alpha, f_\phi^\beta : M \rightarrow \mathbb{R}$ be Lyapunov functions for S . Then,

- (i) for all $\lambda \in \mathbb{R}$ and all $\mu \in \mathbb{R}^+$, it holds that $\lambda + \mu f_\phi^\alpha \in \text{Lyap}(S, \phi)$;
- (ii) for all $\lambda, \mu \in \mathbb{R}^+$, it holds that $f_\phi^\gamma = \lambda f_\phi^\alpha + \mu f_\phi^\beta \in \text{Lyap}(S, \phi)$.

Proof. The first property is immediate. For the second property we observe that $N^\gamma = N^\alpha \cap N^\beta$ is an isolating neighborhood for S and

$$\frac{d}{dt}\Big|_{t=0} f_\phi^\gamma(\phi(t, x)) = \lambda \frac{d}{dt}\Big|_{t=0} f_\phi^\alpha(\phi(t, x)) + \mu \frac{d}{dt}\Big|_{t=0} f_\phi^\beta(\phi(t, x)) < 0,$$

for all $x \in N^\gamma \setminus S$. □

⁴Property (ii) will also be referred to as the *Lyapunov property* with respect to an isolating neighborhood N for S .




Figure 2.2: The isolating block B of S is slightly enlarged using the flow ϕ to the manifold with piecewise smooth boundary B^\dagger .

The next result is an adaption of Lyapunov functions to smooth Lyapunov functions as described in [RS92] and [WY73], which as a consequence shows that the set $\text{Lyap}(S, \phi)$ is non-empty. The following proposition is a modification of a result in [RS92].

Proposition 2.2.6. *Let $S \subset M$ be an isolated invariant set for ϕ and $B \subset M$ a smooth isolating block for S . Then, there exists a smooth Lyapunov function $f_\phi : M \rightarrow \mathbb{R}$ for S , such that the Lyapunov property of f_ϕ holds with respect to B .*

Proof. Consider the following manifold B^\dagger with piecewise smooth boundary, defined by

$$B^\dagger = \phi([-3\tau, 3\tau], B) \subset M,$$

for $\tau > 0$ sufficiently small. By construction $\text{Inv}(B^\dagger, \phi) = S$, cf. Figure 2.2. Choose a smooth cut-off function⁵ $\zeta : B^\dagger \rightarrow [0, 1]$ such that

$$\zeta^{-1}(1) = B \quad \text{and} \quad \zeta^{-1}(0) = \phi([2\tau, 3\tau], B_-) \cup \phi([-3\tau, -2\tau], B_+)$$

and define the vector field $X^\dagger = \zeta X$ on B^\dagger . The flow generated by the vector field X^\dagger is denoted by ϕ^\dagger and $\phi^\dagger|_B = \phi$.

⁵See [WY73, Theorem 1.8] for a proof of the existence.

2.3. Gradient flows, Morse functions and Morse-Smale flows

For the flow $\phi^\dagger : \mathbb{R} \times B^\dagger \rightarrow B^\dagger$ we identify the following attracting neighborhoods⁶

$$U = \phi([0, 3\tau], B) \quad \text{and} \quad V = \phi([\tau, 3\tau], B_-),$$

for the flow ϕ^\dagger on B^\dagger . This yields the attractors $A_U = \omega(U)$ and $A_V = \omega(V)$ for ϕ^\dagger in B^\dagger . Following the theory of attractor-repeller pairs⁷ (cf. [VKM13, KMV13], [RS92]) the dual repellers in B^\dagger are given by $A_U^* = \alpha(W)$ and $A_V^* = \alpha(Z)$, where W and Z are given by

$$W = B^\dagger \setminus U = \phi([-3\tau, 0], B_+) \quad \text{and} \quad Z = B^\dagger \setminus V = \phi([-3\tau, \tau], B).$$

By the theory of Morse decompositions (cf. [VKM13], [RS92]) we have

$$S = A_U \cap A_V^*,$$

cf. Figure 2.3. By Proposition 1.4 in [RS92] we obtain Lyapunov functions f_{A_U} and f_{A_V} which are smooth on $B^\dagger \setminus (A_U \cup A_U^*)$ and $B^\dagger \setminus (A_V \cup A_V^*)$ respectively. They satisfy $\frac{d}{dt} f_{A_U}(\phi(t, x)) < 0$ for all $x \in B^\dagger \setminus (A_U \cup A_U^*)$, and $f_{A_U}^{-1}(0) = A_U$, and $f_{A_U}^{-1}(1) = A_U^*$, and the same for f_{A_V} .

Define

$$f_\phi = \lambda f_{A_U} + \mu f_{A_V}, \quad \lambda, \mu > 0.$$

Clearly, $f_\phi|_S = \text{constant}$ and $\frac{d}{dt} f_\phi(\phi(t, x)) < 0$ for all $x \in B \setminus S$ and all times $t > 0$. We extend f_ϕ to a smooth function on M . \square

2.3 Gradient flows, Morse functions and Morse-Smale flows

The gradient flow of a Lyapunov function f_ϕ for (S, ϕ) gives information about the set S .

2.3.1 Gradient flows

Let g be a Riemannian metric on M and consider the (negative) g -gradient flow equation $x' = -\nabla_g f_\phi(x)$, where $\nabla_g f_\phi$ is the gradient of f_ϕ with respect

⁶A compact neighborhood $N \subset M$ is attracting if $\phi(t, N) \subset \text{int}(N)$ for all $t > 0$. Similarly, a compact neighborhood is repelling if $\phi(t, N) \subset \text{int}(N)$ for all $t < 0$.

⁷A set $A \subset M$ is an *attractor* for a flow ϕ if there exists a neighborhood U of A such that $A = \omega(U)$. The associated *dual repeller* is given by $A^* = \alpha(M \setminus U)$. The pair (A, A^*) is called an *attractor-repeller pair* for ϕ . The pairs (\emptyset, M) and (M, \emptyset) are trivial attractor-repeller pairs.

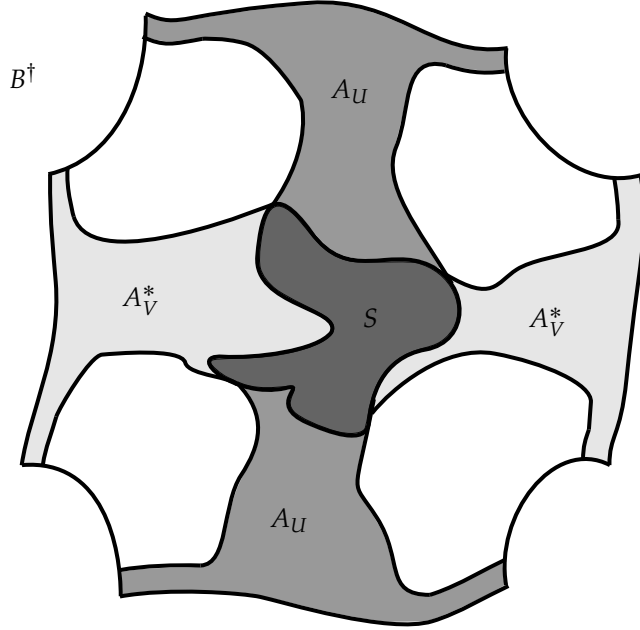


Figure 2.3: An illustration of the fact that $S = A_U \cap A_V^*$.

to the Riemannian metric g . The differential equation generates a global flow $\psi_{(f_\phi, g)} : \mathbb{R} \times M \rightarrow M$ by assumption. We say that $\psi_{(f_\phi, g)}$ is the gradient flow of the pair (f_ϕ, g) for short. The set of critical points of f_ϕ is denoted by

$$\text{Crit}(f_\phi) = \{x \in M \mid df_\phi(x) = 0\}.$$

Lemma 2.3.1. $\text{Crit}(f_\phi) \cap N \subset S$.

Proof. Because f_ϕ is a Lyapunov function for ϕ we have

$$df_\phi(x)X(x) = \left. \frac{d}{dt} \right|_{t=0} f_\phi(\phi(t, x)) < 0,$$

for all $x \in N \setminus S$. This implies that $df_\phi(x) \neq 0$ for $x \in N \setminus S$ and thus $\text{Crit}(f_\phi) \cap N \subset S$. \square

Remark 2.3.2. The reversed inclusion is not true in general. The flow ϕ on $M = \mathbb{R}$, generated by the equation $x' = \frac{x^2}{1+x^2}$, has an equilibrium at 0 and $S = \{0\}$ is an isolated invariant set. The function $x \mapsto -x$ is a Lyapunov function for this flow without any critical points.

By construction N is an isolating neighborhood for ϕ , but it is also an isolating neighborhood for the gradient flow $\psi_{(f_\phi, g)}$.

Lemma 2.3.3. *For any choice of metric g , the set $N \subset M$ is an isolating neighborhood for $\psi_{(f_\phi, g)}$ and $S_{(f_\phi, g)} = \text{Inv}(N, \psi_{(f_\phi, g)}) = \text{Crit}(f_\phi) \cap N \subset S$.*

Proof. Let $x \in \text{Inv}(N, \psi_{(f_\phi, g)})$ and let $\gamma_x = \psi_{(f_\phi, g)}(\mathbb{R}, x)$ be the orbit through x , which is bounded since $\gamma_x \subset N$, and N is compact. For the gradient flow $\psi_{(f_\phi, g)}$ the following identity holds:

$$\frac{d}{dt} f_\phi(\psi_{(f_\phi, g)}(t, x)) = -|\nabla_g f_\phi(\psi_{(f_\phi, g)}(t, x))|_g^2 \leq 0, \quad (2.5)$$

and the function $t \mapsto f_\phi(\psi_{(f_\phi, g)}(t, x))$ is decreasing and bounded. The latter implies that $f_\phi(\psi_{(f_\phi, g)}(t, x)) \rightarrow c_\pm$ as $t \rightarrow \pm\infty$. For any point $y \in \omega(x)$ there exist times $t_n \rightarrow \infty$ such that $\lim_{n \rightarrow \infty} \psi_{(f_\phi, g)}(t_n, x) = y$, and hence

$$f_\phi(y) = f_\phi\left(\lim_{n \rightarrow \infty} \psi_{(f_\phi, g)}(t_n, x)\right) = \lim_{n \rightarrow \infty} f_\phi(\psi_{(f_\phi, g)}(t_n, x)) = c_+.$$

Suppose $y \notin \text{Crit}(f_\phi) \cap N$, then for any $\tau > 0$, we have $f_\phi(\psi_{(f_\phi, g)}(\tau, y)) < c_+$, since $df_\phi(y) \neq 0$ and thus $\left.\frac{d}{dt} f_\phi(\psi_{(f_\phi, g)}(t, y))\right|_{t=0} < 0$. On the other hand, by continuity and the group property, $\psi_{(f_\phi, g)}(t_n + \tau, x) \rightarrow \psi_{(f_\phi, g)}(\tau, y)$, which implies

$$c_+ = \lim_{n \rightarrow \infty} f_\phi(\psi_{(f_\phi, g)}(t_n + \tau, x)) = f_\phi(\psi_{(f_\phi, g)}(\tau, y)) < c_+,$$

a contradiction. Thus the limit points are contained in $\text{Crit}(f_\phi) \cap N$. The same argument holds for points $y \in \alpha(x)$. This proves that $\omega(x), \alpha(x) \subset \text{Crit}(f_\phi) \cap N \subset S$ for all $x \in \text{Inv}(N, \psi_{(f_\phi, g)})$. Since f_ϕ is constant on S it follows that $c_+ = c_- = c$ and thus $f_\phi(\psi_{(f_\phi, g)}(t, x)) = c$ for all $t \in \mathbb{R}$, i.e. $f_\phi|_{\gamma_x} = c$. Suppose $x \notin \text{Crit}(f_\phi) \cap N$, then $df_\phi(x) \neq 0$. Using Equation (2.5) at $t = 0$ yields a contradiction due to the fact that f_ϕ is constant along γ_x . Therefore $\text{Inv}(N, \psi_{(f_\phi, g)}) \subset \text{Crit}(f_\phi) \cap N \subset S$. This completes the proof since $\text{Crit}(f_\phi) \cap N \subset \text{Inv}(N, \psi_{(f_\phi, g)})$. \square

2.3.2 Morse functions

A smooth function $f : M \rightarrow \mathbb{R}$ is called *Morse* on N if $\text{Crit}(f) \cap N \subset \text{int}(N)$ and $\text{Crit}(f) \cap N$ consists of only non-degenerate critical points.

Then the local stable and unstable manifolds of a critical point $x \in \text{Crit}(f) \cap N$, with respect to the gradient flow $\psi_{(f,g)}$, defined by $x' = -\nabla_g f(x)$, are given by

$$\begin{aligned} W_{\text{loc}}^s(x; N) &= \{z \in M \mid \psi_{(f,g)}(t, z) \in N, \forall t \geq 0, \lim_{t \rightarrow \infty} \psi_{(f,g)}(t, z) = x\}, \\ W_{\text{loc}}^u(x; N) &= \{z \in M \mid \psi_{(f,g)}(t, z) \in N, \forall t \leq 0, \lim_{t \rightarrow -\infty} \psi_{(f,g)}(t, z) = x\}. \end{aligned} \quad (2.6)$$

The sets $W^s(x)$ and $W^u(x)$ are the global stable and unstable manifolds respectively, which are defined without the restriction of points in N . We will write $S_{(f,g)} = \text{Inv}(N, \psi_{(f,g)})$ for the maximal invariant set of $\psi_{(f,g)}$ inside N .

Lemma 2.3.4. *Suppose f is Morse on N , then the maximal invariant set in N of the gradient flow $\psi_{(f,g)}$ is characterized by*

$$S_{(f,g)} = \bigcup_{x,y \in \text{Crit}(f) \cap N} W_{\text{loc}}^u(x; N) \cap W_{\text{loc}}^s(y; N). \quad (2.7)$$

Proof. If $z \in W_{\text{loc}}^u(x; N) \cap W_{\text{loc}}^s(y; N)$ for some $x, y \in \text{Crit}(f) \cap N$, then by definition, $\psi_{(f,g)}(t, z) \in N$ for all $t \in \mathbb{R}$, and hence $z \in S_{(f,g)}$ which shows that $W_{\text{loc}}^u(x; N) \cap W_{\text{loc}}^s(y; N) \subset S_{(f,g)}$. Conversely, if $z \in S_{(f,g)}$, then $\psi_{(f,g)}(t, z) \in S_{(f,g)} \subset N$ for all $t \in \mathbb{R}$, i.e. $\gamma_z \subset N$. By the the same arguments as in the proof of Lemma 2.3.3 the limits $\lim_{t \rightarrow \pm\infty} \psi_t(z)$ exist and are critical points of f contained in $S_{(f,g)}$ (compactness of $S_{(f,g)}$). Therefore, $z \in W_{\text{loc}}^u(x; N) \cap W_{\text{loc}}^s(y; N)$, for some $x, y \in \text{Crit}(f) \cap N$ which implies that $S_{(f,g)} \subset \bigcup_{x,y \in \text{Crit}(f) \cap N} W_{\text{loc}}^u(x; N) \cap W_{\text{loc}}^s(y; N)$. \square

2.3.3 Morse-Smale flows and isolated homotopies

Additional structure on the set of connecting orbits is achieved by the Morse-Smale property.

Definition 2.3.5. A gradient flow $\psi_{(f,g)}$ is called *Morse-Smale* on N if

- (i) $S_{(f,g)} \subset \text{int}(N)$;
- (ii) f is Morse on N , i.e. critical points in N are non-degenerate;

(iii) $W_{\text{loc}}^u(x; N) \pitchfork W_{\text{loc}}^s(y; N)$ ⁸, for all $x, y \in \text{Crit}(f) \cap N$.

In this setting the pair (f, g) is called a *Morse-Smale pair* on N .

Remark 2.3.6. In general Lyapunov functions are not Morse. We want to perturb the Lyapunov function to obtain a Morse-Smale pair, which we will use to define invariants. Not all Morse-Smale pairs are suitable for the construction of invariants as the following example shows. Let $\phi(t, x) = xe^t$ be a flow on $M = \mathbb{R}$ and let $N = [-1, 1]$ be an isolating block for ϕ , with $S = \text{Inv}(N, \phi) = \{0\}$. The function $f_\phi(x) = -\frac{1}{4}x^4$ is a Lyapunov for (S, ϕ) . Let g be the standard inner product on \mathbb{R} , then the function $f(x) = \frac{1}{2}x^2$ yields a Morse-Smale flow $\psi_{(f, g)}(t, x) = xe^{-t}$ via $x' = -f'(x) = -x$ and thus (f, g) is a Morse-Smale pair on N . The flow $\psi_{(f, g)}$ obviously displays the wrong dynamical behavior. The reason is that one cannot find a homotopy between f_ϕ and f which preserves isolation with respect to N . This motivates the following definition.

Definition 2.3.7. Let $\psi_{(h, e)}$ be the gradient flow of h with respect to the metric e , with the property that $S_{(h, e)} \subset \text{int}(N)$. A Morse-Smale pair (f, g) on N is *isolated homotopic* to (h, e) if there exists a smooth homotopy $(f_\lambda, g_\lambda)_{\lambda \in [0, 1]}$ between (f, g) and (h, e) such that $S_{(f_\lambda, g_\lambda)} \subset \text{int}(N)$ for all $\lambda \in [0, 1]$.⁹ The set of such Morse-Smale pairs is denoted by $\mathcal{S}_{MS}(h, e; N)$.

Isolation is well behaved under perturbation, cf. [Con78].

Proposition 2.3.8. *The set of flows preserving isolation of N is open in the compact-open topology.*

Proof. Let ϕ be a flow with isolating neighborhood N . Isolation implies that for all $p \in \partial N$ there exists a $t \in \mathbb{R}$ such that $\phi(t, p) \in M \setminus N$. By continuity there exist compact neighborhoods $U_p \ni p$ and $I_t \ni t$ such that $\phi(I_t, U_p) \subset M \setminus N$. The compactness of ∂N implies that a finite number of neighborhoods U_{p_i} cover ∂N . Now we have that

$$\phi \in \bigcap_i \{ \psi \in C(\mathbb{R} \times M, M) \mid \psi(I_{t_i}, U_{p_i}) \subset M \setminus N \},$$

and all flows in this open set are isolating. □

⁸The symbol \pitchfork indicates that the intersection is transverse in the following sense. For each $p \in W_{\text{loc}}^u(x; N) \cap W_{\text{loc}}^s(y; N)$, we have that $T_p W^u(x) + T_p W^s(y) = T_p M$.

⁹The flows $\psi_{(f_\lambda, g_\lambda)}$ are generated by the equations $x' = -\nabla_{g_\lambda} f_\lambda(x)$.

The following proposition shows that the set of isolated homotopies $\mathcal{S}_{MS}(h, e; N)$ is not empty.

Proposition 2.3.9. *Let $\psi_{(h,e)}$ be a gradient flow of (h, e) , with the property that $S_{(h,e)} \subset \text{int}(N)$. Then, for each f sufficiently C^2 -close to h , such that (f, e) is a Morse-Smale pair on N , then (f, e) is isolated homotopic to (h, e) . Consequently, $\mathcal{S}_{MS}(h, e; N) \cap \{(f, e) \mid \|f - h\|_{C^2} < \epsilon\} \neq \emptyset$ for every $\epsilon > 0$.*

Proof. The existence of a smooth functions f for which the Morse-Smale property holds for $-\nabla_g f$, with $g = e$, follows from the results in [AB95]. Since N is an isolating neighborhood for $\psi_{(h,e)}$, a sufficiently C^2 -small perturbation f of h implies that $\nabla_g f$ is C^1 -close to the vector field $\nabla_e h$. Now C^1 -vector fields define (local) flows, and the set of flows preserves isolation is open, by Proposition 2.3.8. Consequently, $S_{(f,g)} = \text{Inv}(N, \psi_{(f,g)}) \subset \text{int}(N)$. If the C^2 -perturbation is small enough there is a path connecting (f, g) and (h, e) preserving isolation, which proves that $(f, g) \in \mathcal{S}_{MS}(h, e; N)$. \square

Remark 2.3.10. In [AB95] it is proved that the set of smooth functions on a compact manifold is a Baire space. This can be used to strengthen the statement in Proposition 2.3.9 as follows: the set of smooth functions in a sufficiently small C^2 -neighborhood of a Morse function on N is a Baire space.

Remark 2.3.11. In [Sma61] it is proved that any gradient vector field $-\nabla_e h$ can be approximated sufficiently C^1 -close by $-\nabla_g f$, such the associated flow $\psi_{(f,g)}$ is Morse-Smale on N , cf. [Web06]. If h is Morse one can find a perturbation g of the metric e such that the flow $\psi_{(h,g)}$ is Morse-Smale, cf. [AM06].

From Lemma 2.3.3 it follows that for any Lyapunov function f_ϕ for (S, ϕ) any flow $\psi_{(f_\phi, g)}$, N is an isolating block with $\text{Inv}(N, \psi_{(f_\phi, g)}) \subset S$. The choice of metric does not play a role and provides freedom in choosing isolated homotopies:

$$\mathcal{S}_{MS}(f_\phi; N) = \bigcup_e \mathcal{S}_{MS}(f_\phi, e; N),$$

which represents the set of Morse-Smale pairs (f, g) which is isolated homotopic to (f_ϕ, e) , for some Riemannian metric e on M .

Corollary 2.3.12. *Given a Lyapunov function $f_\phi \in \text{Lyap}(S, \psi)$, then there exists a Morse-Smale pair $(f, g) \in \mathcal{S}_{MS}(f_\phi; N)$ on N with f arbitrary C^2 -close to f_ϕ .*

Proof. If $f_\phi \in \text{Lyap}(S, \phi)$ is a Lyapunov function for (S, ϕ) , then for the gradient flow $\psi_{(f_\phi, g)}$ it holds that $S_{(f_\phi, g)} = \text{Inv}(N, \psi) = \text{Crit}(f_\phi) \cap N \subset S \subset \text{int}(N)$, with g arbitrary, and thus by Proposition 2.3.9 there exist a Morse-Smale pair (f, g) , close to (f_ϕ, g) , such that $S_{(f, g)} = \text{Inv}(N, \psi_{(f, g)}) \subset \text{int}(N)$ and (f, g) is isolated homotopic to (f_ϕ, g) . □

2.4 Morse homology

From Lemma 2.3.3 it follows that N is a isolating neighborhood for the gradient flow $\psi_{(f_\phi, g)}$, for any Lyapunov function $f_\phi \in \text{Lyap}(S, \phi)$, regardless of the chosen metric g . Corollary 2.3.12 implies that there exists sufficiently C^2 -close smooth function f , such that (f, g) is a Morse-Smale pair on N and (f, g) is isolated homotopic to (f_ϕ, g) , i.e. $(f, g) \in \mathcal{S}_{MS}(f_\phi; N)$.

2.4.1 Morse homology

We follow the treatment of Morse homology given in [Web06], applied to a Morse-Smale pair $(f, g) \in \mathcal{S}_{MS}(f_\phi; N)$. Define the moduli space by

$$W_N(x, y) = W_{\text{loc}}^u(x; N) \cap W_{\text{loc}}^s(y; N). \quad (2.8)$$

By Proposition 2.3.4 and Corollary 2.3.12

$$S_{(f, g)} = \text{Inv}(N, \psi_{(f, g)}) = \bigcup_{x, y \in \text{Crit}(f) \cap N} W_N(x, y) \subset \text{int}(N),$$

and thus $W_N(x, y) \subset \text{int}(N)$ for all $x, y \in \text{Crit}(f) \cap N$. The quotient $M_N(x, y) = W_N(x, y)/\mathbb{R}$, which can be identified with $W_N(x, y) \cap f^{-1}(a)$, with $a \in (f(y), f(x))$.

The ‘full’ moduli spaces are given by $W(x, y) = W^u(x) \cap W^s(y)$. A connected component of $W(x, y)$ which intersects $W_N(x, y)$ is completely contained in $W_N(x, y)$. Since $\psi_{(f, g)}$ is Morse-Smale on N the sets $M_N(x, y)$ are smooth submanifolds (without boundary) of dimension $\dim M_N(x, y) = |x| - |y| - 1$, where $|x|$ is the Morse index of x , cf. [Web06]. The submanifolds $M_N(x, y)$ can be compactified by adjoining broken flow lines.

If $|x| = |y| + 1$, then $M_N(x, y)$ is a zero-dimensional manifold without boundary, and since $\psi_{(f, g)}$ is Morse-Smale and N is compact, the set $M_N(x, y)$




Figure 2.4: The sequence u_k converges to the broken orbit (v, w) as it approaches the end of the moduli space.

is a finite point set, cf. [Web06]. If $|x| = |y| + 2$, then $M_N(x, y)$ is a disjoint union of copies of S^1 or $(0, 1)$.

The non-compactness in the components diffeomorphic to $(0, 1)$ can be described as follows, see also Figure 2.4. Identify $M_N(x, y)$ with $W_N(x, y) \cap f^{-1}(a)$ (as before) and let $u_k \in M_N(x, y)$ be a sequence without any convergent subsequences. Then, there exist a critical $z \in \text{Crit}(f) \cap N$, with $f(y) < f(z) < f(x)$, reals a^1, a^2 , with $f(z) < a^1 < f(x)$ and $f(y) < a^2 < f(z)$, and times t_k^1 and t_k^2 , such that $f(\psi_{(f,g)}(t_k^1, u_k)) = a^1$ and $f(\psi_{(f,g)}(t_k^2, u_k)) = a^2$, such that

$$\psi_{(f,g)}(t_k^1, u_k) \rightarrow v \in W_N(x, z) \cap f^{-1}(a^1), \quad \text{and}$$

$$\psi_{(f,g)}(t_k^2, u_k) \rightarrow w \in W_N(z, y) \cap f^{-1}(a^2).$$

We say that the non-compactness is represented by concatenations of two trajectories $x \rightarrow z$ and $z \rightarrow y$ with $|z| = |x| - 1$. We write that $u_k \rightarrow (v, w)$ — geometric convergence —, and (v, w) is a broken trajectory. The following proposition adjusts Theorem 3.9 in [Web06].

Lemma 2.4.1 (Restricted transitivity). *Suppose $x, y, z \in \text{Crit}(f) \cap N$ with $\text{ind}_f(x) = \text{ind}_f(z) + 1$ and $|z| = |y| + 1$. Then there exist the (restricted) gluing embedding*

$$\#_\rho : \mathbb{R}^+ \times M_N(x, z) \times M_N(z, y) \rightarrow M_N(x, y), \quad (\rho, v, w) \mapsto v\#_\rho w,$$

such that $v\#_\rho w \rightarrow (v, w)$ in the geometric sense. Moreover, no sequences in $M_N(x, y) \setminus v\#_{\mathbb{R}^+} w$ geometrically converge to (v, w) .

Proof. In Morse homology we have the restricted gluing embedding

$$\#_\rho : \mathbb{R}^+ \times M_N(x, z) \times M_N(z, y) \rightarrow M(x, y),$$

which maps a triple (ρ, v, w) to an element $v\#_\rho w \in M(x, y)$, cf. [Sch93, Web06]. The geometric convergence $v\#_\rho w \rightarrow (v, w)$, as $\rho \rightarrow \infty$, implies that $\gamma_\rho = \psi_{(f,g)}(\mathbb{R}, v\#_\rho w)$ gets arbitrary close to the images of $\psi_{(f,g)}(\mathbb{R}, v)$ and $\psi_{(f,g)}(\mathbb{R}, w)$. Since $\psi_{(f,g)}(\mathbb{R}, v) \subset \text{int}(N)$ and $\psi_{(f,g)}(\mathbb{R}, w) \subset \text{int}(N)$ we obtain orbits $\gamma_\rho \subset \text{int}(N)$ for ρ sufficiently large. This yields the embedding into $M_N(x, z)$. \square

Following the construction of Morse homology in [Web06] (see also [Flo89b], [Sch93]) we choose orientations of the unstable manifolds $W_{\text{loc}}^u(x; N)$, for all $x \in \text{Crit}(f) \cap N$, and denote this choice by \mathfrak{o} . Define chain groups and boundary operators

$$C_k(f, N) = \bigoplus_{\substack{x \in \text{Crit}(f) \cap N \\ |x|=k}} \mathbb{Z}\langle x \rangle,$$

$$\partial_k(f, g, N, \mathfrak{o}) : C_k(f, N) \rightarrow C_{k-1}(f, N),$$

where

$$\partial_k(f, g, N, \mathfrak{o})\langle x \rangle = \sum_{\substack{y \in \text{Crit}(f) \cap N \\ |y|=k-1}} n_N(x, y)\langle y \rangle,$$

and $n_N(x, y)$ is the oriented intersection number of $W_{\text{loc}}^u(x; N) \cap f^{-1}(a)$ and $W_{\text{loc}}^s(y; N) \cap f^{-1}(a)$ with respect to \mathfrak{o} .¹⁰

The isolation of $S_{(f,g)}$ in N has strong implications for the boundary operator $\partial(f, g, N, \mathfrak{o})$.

¹⁰Over \mathbb{Z}_2 the intersection number is given by the number of elements in $M_N(x, y)$ modulo 2.

Lemma 2.4.2. $\partial^2(f, g, N, \mathfrak{o}) = 0$.

Proof. Restricted transitivity in Proposition 2.4.1 implies that all ends of the components $M_N(x, y)$ that are diffeomorphic to $(0, 1)$, are different and lie in N . The gluing map also implies that all broken trajectories in N occur as unique ends of components of $M_N(x, y)$. By the same counting procedure as in [Web06] we obtain $\partial^2(f, g, N, \mathfrak{o}) = 0$. \square

Definition 2.4.3. A quadruple $\mathcal{Q} = (f, g, N, \mathfrak{o})$ is called a *Morse-Conley-Floer quadruple* for (S, ϕ) if

- (i) N is an isolating neighborhood for S , i.e. $S = \text{Inv}(N, \phi)$;
- (ii) $(f, g) \in \mathcal{S}_{MS}(f_\phi; N)$, for a Lyapunov function $f_\phi \in \text{Lyap}(S, \phi)$;
- (iii) \mathfrak{o} is choice of orientations of the unstable manifolds of the critical points of f in N .

By Proposition 2.4.2 $(C_*(f, N), \partial_*(f, g, N, \mathfrak{o}))$ is a chain complex and the Morse homology of the quadruple $\mathcal{Q} = (f, g, N, \mathfrak{o})$ is now defined as

$$HM_*(\mathcal{Q}) := H_*(C_*(f, N), \partial_*(f, g, N, \mathfrak{o})).$$

2.4.2 Independence of Lyapunov functions and perturbations

The next step is to show that $HM_*(\mathcal{Q})$ is independent of the choice of Lyapunov function, metric and orientation, as well as the choice of perturbation f . In Section 2.4.3 we show that the homology is also independent of the isolating neighborhood N .

Theorem 2.4.4. *Let $\mathcal{Q}^\alpha = (f^\alpha, g^\alpha, N^\alpha, \mathfrak{o}^\alpha)$ and $\mathcal{Q}^\beta = (f^\beta, g^\beta, N^\beta, \mathfrak{o}^\beta)$ be Morse-Conley-Floer quadruples for (S, ϕ) , with the same isolating neighborhood $N = N^\alpha = N^\beta$. Then there exist canonical isomorphisms*

$$\Phi_*^{\beta\alpha} : HM_*(\mathcal{Q}^\alpha) \xrightarrow{\cong} HM_*(\mathcal{Q}^\beta).$$

The proof essentially follows the continuation theorem for Floer homology as given in [Flo89b]. We sketch a proof based on the same arguments in [Flo89a] and [Sch93] and the dynamical systems approach in [Web06]. The idea of the proof of this theorem is to construct higher dimensional systems which contain the complexes generated by both quadruples. The fundamental relation $\partial^2 = 0$ on the higher dimensional systems induces a map between the

Morse homologies, which is then used to construct an isomorphism between the homologies of both quadruples.

Fix a metric g on M . By assumption there are Lyapunov functions $f_\phi^\alpha, f_\phi^\beta \in \text{Lyap}(S, \phi)$ defined on the same isolating block N and Morse-Smale pairs $(f^\alpha, g^\alpha) \in \mathcal{S}_{MS}(f_\phi^\alpha; N)$ and $(f^\beta, g^\beta) \in \mathcal{S}_{MS}(f_\phi^\beta; N)$. This implies that there exists a homotopy $(f_\lambda^\alpha, g_\lambda^\alpha)$ between (f^α, g^α) and (f_ϕ^α, g) , and similarly there exists a homotopy $(f_\lambda^\beta, g_\lambda^\beta)$ between (f^β, g^β) and (f_ϕ^β, g) . By Proposition 2.2.5 the functions

$$f_{\phi, \lambda} = (1 - \lambda)f_\phi^\alpha + \lambda f_\phi^\beta,$$

are Lyapunov functions for (S, ϕ) . Since $f_{\phi, \lambda} \in \text{Lyap}(S, \phi)$, for all $\lambda \in [0, 1]$, it follows by Lemma 2.3.3 that $\text{Inv}(N, \psi_{(f_{\phi, \lambda}, g)}) = \text{Crit}(f_{\phi, \lambda}) \cap N \subset S$, where $\psi_{(f_{\phi, \lambda}, g)}$ is the gradient flow of $(f_{\phi, \lambda}, g)$, and thus N is an isolating neighborhood for all $\psi_{(f_{\phi, \lambda}, g)}$. We concatenate the homotopies and define

$$f_\lambda = \begin{cases} f_{3\lambda}^\alpha & \text{for } \lambda \in [0, \frac{1}{3}], \\ f_{\phi, 3\lambda-1} & \text{for } \lambda \in [\frac{1}{3}, \frac{2}{3}], \\ f_{3-3\lambda}^\beta & \text{for } \lambda \in [\frac{2}{3}, 1], \end{cases}$$

which is a piecewise smooth homotopy between f^α and f^β . Similarly define

$$g_\lambda = \begin{cases} g_{3\lambda}^\alpha & \text{for } \lambda \in [0, \frac{1}{3}], \\ g & \text{for } \lambda \in [\frac{1}{3}, \frac{2}{3}], \\ g_{3-3\lambda}^\beta & \text{for } \lambda \in [\frac{2}{3}, 1], \end{cases}$$

which is a piecewise smooth homotopy of metric between g^α and g^β . Both homotopies can be reparameterized to smooth homotopies via a diffeomorphism $\lambda \mapsto \alpha(\lambda)$, with $\alpha'(\lambda) \geq 0$ and $\alpha^{(k)}(\lambda) = 0$ for $\lambda = \frac{1}{3}, \frac{2}{3}$ and for all $k \geq 1$. We denote the reparameterized homotopy again by (f_λ, g_λ) . The homotopies fit into the diagram

$$\begin{array}{ccc} (f^\alpha, g^\alpha) & \xrightarrow{(f_\lambda, g_\lambda)} & (f^\beta, g^\beta) \\ \downarrow (f_\lambda^\alpha, g_\lambda^\alpha) & & \downarrow (f_\lambda^\beta, g_\lambda^\beta) \\ (f_\phi^\alpha, g) & \xrightarrow{(f_{\phi, \lambda}, g)} & (f_\phi^\beta, g) \end{array}$$

The flow of (f_λ, g_λ) is denoted by $\psi_{(f_\lambda, g_\lambda)}$. By assumption and construction $\text{Inv}(N, \psi_{(f_\lambda, g_\lambda)}) \subset \text{int}(N)$ for all $\lambda \in [0, 1]$ and therefore $(f^\alpha, g^\alpha) \in \mathcal{I}_{MS}(f^\beta, g^\beta; N)$ and vice versa. Note that f_λ is not necessarily Morse for each $\lambda \in (0, 1)$. At $\lambda = 0, 1$ the flows $\psi_{(f_\lambda, g_\lambda)}$ are Morse-Smale flows on N by assumption.

Let $r > 0$ and $0 < \delta < \frac{1}{4}$, and define the function $F : M \times \mathbb{S}^1 \rightarrow \mathbb{R}$

$$F(x, \mu) = f_{\omega(\mu)}(x) + r[1 + \cos(\pi\mu)], \quad (2.9)$$

where $\omega : \mathbb{R} \rightarrow [0, 1]$ is a smooth, even, 2-periodic function with the following properties on the interval $[-1, 1]$: $\omega(\mu) = 0$, for $-\delta < \mu < \delta$, $\omega(\mu) = 1$, for $-1 < \mu < -1 + \delta$ and $1 - \delta < \mu < 1$, and $\omega'(\mu) < 0$ on $(-1 + \delta, -\delta)$ and $\omega'(\mu) > 0$ on $(\delta, 1 - \delta)$. By identifying \mathbb{S}^1 with $\mathbb{R}/2\mathbb{Z}$, the function ω descends to a function on \mathbb{S}^1 , which is denoted by the same symbol. We consider the product metric G^\times on $M \times \mathbb{S}^1$ defined by

$$G_{(x, \mu)}^\times = (g_{\omega(\mu)})_x \oplus \frac{1}{\kappa} d\mu^2, \quad \kappa > 0.$$

For the proof of the theorem it would suffice to take $\kappa = 1$. However, in the proof of Proposition 2.5.3 we do need the parameter κ . In order to use the current proof ad verbatim, we introduce the parameter.

Lemma 2.4.5. *If*

$$r > \frac{\max_{\mu \in \mathbb{S}^1, x \in N} |\omega'(\mu) [\partial_\lambda f_\lambda(x)|_{\lambda=\omega(\mu)}]|}{\pi \sin(\pi\delta)},$$

then for each critical point $(x, \mu) \in \text{Crit}(F) \cap (N \times \mathbb{S}^1)$, either $\mu = 0$ and $|(x, 0)| = |x| + 1$, or $\mu = 1$ and $|(x, 1)| = |x|$. In particular, F is a Morse function on $N \times \mathbb{S}^1$ and

$$C_k(F, N \times \mathbb{S}^1) \cong C_{k-1}(f^\alpha, N) \oplus C_k(f^\beta, N). \quad (2.10)$$

Proof. The G -gradient vector field of F on $M \times \mathbb{S}^1$ is given by

$$\begin{aligned} \nabla_{G^\times} F(x, \mu) = \\ \nabla_{g_{\omega(\mu)}} f_{\omega(\mu)}(x) + \kappa \left(\omega'(\mu) [\partial_\lambda f_\lambda(x)|_{\lambda=\omega(\mu)}] - r\pi \sin(\pi\mu) \right) \partial_\mu. \end{aligned}$$




Figure 2.5: The continuation map $\Phi_*^{\beta\alpha} : HM_*(\mathcal{Q}^\alpha) \rightarrow HM_*(\mathcal{Q}^\beta)$ is induced by a higher dimensional gradient system with Morse function of Equation (2.9).

by the choice on r the second term is only zero on $B \times \mathbb{S}^1$ if $\mu = 0, 1$. Then the critical points at $\mu = 0$ correspond to critical points of f^α , but there is one extra unstable direction, corresponding to ∂_μ . At $\mu = 1$ the critical points come from critical points of f^β , and there is one extra stable direction, which does not contribute to the Morse index. \square

This outlines the construction of the higher dimensional system. From now on we assume r specified as above.

Proof of Theorem 2.4.4. Consider the negative gradient flow $\Psi_\kappa^\times : \mathbb{R} \times M \times \mathbb{S}^1 \rightarrow M \times \mathbb{S}^1$ is generated by the vector field $-\nabla_{G \times F}$, cf. Figure 2.5. For $\kappa = 0$ the flow Ψ_0^\times is a product flow and $N \times \mathbb{S}^1$ is an isolating neighborhood for $\{(x, \lambda) \mid x \in S_{(f_\lambda, g_\lambda)}\}$. Since isolation is preserved under small perturbations, $N \times \mathbb{S}^1$ is also an isolating neighborhood for Ψ_κ^\times , for all $\kappa > 0$ sufficiently small. Note that $C = N \times [-\frac{\delta}{2}, 1 + \frac{\delta}{2}]$ is also an isolating neighborhood of the

2. MORSE-CONLEY-FLOER HOMOLOGY

gradient flow Ψ_κ^\times . The flow Ψ^\times is not Morse-Smale on C , but since F is Morse, a small perturbation G of the product metric G^\times , makes (F, G) a Morse-Smale pair and therefore the associated negative G -gradient flow Ψ is Morse-Smale, without destroying the isolation property of C . The connections at $\mu = 0$ and $\mu = 1$ are stable, because these are Morse-Smale for $\kappa = 0$. This implies that we can still denote the boundary operators at $\mu = 0, 1$ by $\partial^\alpha = \partial(f^\alpha, g^\alpha, N, \mathfrak{o}^\alpha)$ and $\partial^\beta = \partial(f^\beta, g^\beta, N, \mathfrak{o}^\beta)$ respectively, since the perturbed metric at the end point do not change the connections and their oriented intersection numbers. We now choose the following orientation

$$\mathcal{O} = ([\partial\mu] \oplus \mathfrak{o}_0) \cup \mathfrak{o}_1,$$

for the unstable manifolds of the critical points of F in C .

By Theorem 2.4.2 the Morse homology of the Morse-Conley-Floer quadruple (F, G, C, \mathcal{O}) is well-defined. With regards to the splitting of Equation (2.10) the boundary operator $\Delta_k = \partial_k(F, G, C, \mathcal{O})$ is given by

$$\Delta_k = \begin{pmatrix} -\partial_{k-1}^\alpha & 0 \\ \Phi_{k-1}^{\beta\alpha} & \partial_k^\beta \end{pmatrix}. \quad (2.11)$$

The map $\Phi^{\beta\alpha}$ counts gradient flow lines from $(x, 0)$ to $(y, 1)$ with $|(x, 0)| = |y, 1| + 1$. Theorem 2.4.2 gives $\Delta_{k-1} \circ \Delta_k = 0$, from which we derive that $\Phi_*^{\beta\alpha}$ is a chain map

$$\Phi_{k-2}^{\beta\alpha} \circ \partial_{k-1}^\alpha = \partial_{k-1}^\beta \circ \Phi_{k-1}^{\beta\alpha},$$

hence $\Phi_*^{\beta\alpha}$ descends to a map in homology $\Phi_*^{\beta\alpha} : HM_*(\mathcal{Q}^\alpha) \rightarrow HM_*(\mathcal{Q}^\beta)$, which we denote by the same symbol. The arguments in [Web06] can be used now to show that the maps $\Psi_*^{\beta\alpha}$ only depend on the ‘end points’ \mathcal{Q}^α and \mathcal{Q}^β , and not on the particular homotopy. It is obvious that if $\mathcal{Q}^\alpha = \mathcal{Q}^\beta$ this map is the identity, both on the chain and homology level. By looking at a higher dimensional system, we establish that, if $\mathcal{Q}^\gamma = (f^\gamma, g^\gamma, N^\gamma, \mathfrak{o}^\gamma)$, with $N^\gamma = N$ is a third Morse-Conley-Floer quadruple, then the induced maps satisfy the functorial relation

$$\Phi_*^{\gamma\alpha} = \Phi_*^{\gamma\beta} \circ \Phi_*^{\beta\alpha},$$

on homology. The argument is identical to [Web06], which we therefore omit. Taking $(f^\gamma, g^\gamma, N, \mathfrak{o}^\gamma) = (f^\alpha, g^\alpha, N, \mathfrak{o}^\alpha)$, along with the observation that $\Phi_*^{\alpha\alpha} = \text{id}$, shows that the maps are isomorphisms. \square

Let $\mathcal{Q}^\alpha, \mathcal{Q}^\beta, \mathcal{Q}^\gamma$ be Morse-Conley-Floer quadruples for (S, ϕ) with $N = N^\alpha = N^\beta = N^\gamma$. This gives the above defined isomorphisms:

$$\begin{aligned}\Phi_*^{\beta\alpha} &: HM_*(\mathcal{Q}^\alpha) \rightarrow HM_*(\mathcal{Q}^\beta), \\ \Phi_*^{\gamma\beta} &: HM_*(\mathcal{Q}^\beta) \rightarrow HM_*(\mathcal{Q}^\gamma), \\ \Phi_*^{\gamma\alpha} &: HM_*(\mathcal{Q}^\alpha) \rightarrow HM_*(\mathcal{Q}^\gamma).\end{aligned}$$

The proof of the previous proposition shows that the following functorial relation holds.

Theorem 2.4.6. *It holds that $\Phi_*^{\gamma\alpha} = \Phi_*^{\gamma\beta} \circ \Phi_*^{\beta\alpha}$.*

2.4.3 Independence of isolating blocks

In section we show that Morse homology of a Morse-Conley-Floer quadruple is independent of the choice of isolating block N for the isolated invariant set S .

Theorem 2.4.7. *Let $\mathcal{Q}^\alpha = (f^\alpha, g^\alpha, N^\alpha, \mathfrak{o}^\alpha)$ and $\mathcal{Q}^\beta = (f^\beta, g^\beta, N^\beta, \mathfrak{o}^\beta)$ be Morse-Conley-Floer quadruples for (S, ϕ) . Then, there are canonical isomorphisms*

$$\Phi_*^{\beta\alpha} : HM_*(\mathcal{Q}^\alpha) \xrightarrow{\cong} HM_*(\mathcal{Q}^\beta).$$

Proof. By assumption there exist Lyapunov functions $f_\phi^\alpha, f_\phi^\beta$ on N^α, N^β and Morse-Smale pairs $(f^\alpha, g^\alpha) \in \mathcal{S}_{MS}(f_\phi^\alpha; N^\alpha)$ and $(f^\beta, g^\beta) \in \mathcal{S}_{MS}(f_\phi^\beta; N^\beta)$ be Morse-Smale pairs. We cannot directly compare the complexes as in the previous proposition, however $N = N^\alpha \cap N^\beta$ is an isolating neighborhood. The restrictions of $f_\phi^\alpha, f_\phi^\beta$ to N are Lyapunov functions. Since $(f^\alpha, g^\alpha) \in \mathcal{S}_{MS}(f_\phi^\alpha; N^\alpha)$, there exists a smooth function $f_\epsilon^\alpha : M \rightarrow \mathbb{R}$, and Riemannian metric g_ϵ^α , with $\text{Inv}(\psi_{(f_\epsilon^\alpha, g_\epsilon^\alpha)}, N^\alpha) \subset \text{int}(N)$, such that

$$(f^\alpha, g^\alpha) \in \mathcal{S}_{MS}(f_\epsilon^\alpha; N^\alpha), \quad \text{and} \quad (f_\epsilon^\alpha, g_\epsilon^\alpha) \in \mathcal{S}_{MS}(f_\phi^\alpha; N).$$

Similarly, there exists a smooth function $f_\epsilon^\beta : M \rightarrow \mathbb{R}$, and a Riemannian metric g_ϵ^β , with $\text{Inv}(\psi_{(f_\epsilon^\beta, g_\epsilon^\beta)}, N^\beta) \subset \text{int}(N)$, such that

$$(f^\beta, g^\beta) \in \mathcal{S}_{MS}(f_\epsilon^\beta; N^\beta) \quad \text{and} \quad (f_\epsilon^\beta, g_\epsilon^\beta) \in \mathcal{S}_{MS}(f_\phi^\beta; N).$$

This yields the following diagram of isomorphisms

$$\begin{array}{ccc}
 HM_*(Q^\alpha) & \xrightarrow{\Psi_*^\alpha} & HM_*(f_\epsilon^\alpha, g_\epsilon^\alpha, N^\alpha, \mathfrak{o}^\alpha) \xrightarrow{\cong} HM_*(f_\epsilon^\alpha, g_\epsilon^\alpha, N, \mathfrak{o}^\alpha) \\
 \downarrow \Phi_*^{\beta\alpha} & & \downarrow \Phi_*^{\beta\alpha, \epsilon} \\
 HM_*(Q^\beta) & \xrightarrow{\Psi_*^\beta} & HM_*(f_\epsilon^\beta, g_\epsilon^\beta, N^\beta, \mathfrak{o}^\beta) \xrightarrow{\cong} HM_*(f_\epsilon^\beta, g_\epsilon^\beta, N, \mathfrak{o}^\beta),
 \end{array}$$

where Ψ_*^α , Ψ_*^β and $\Phi_*^{\beta\alpha, \epsilon}$ are given by Theorem 2.4.4 and $\Phi_*^{\beta\alpha} = (\Psi_*^\beta)^{-1} \circ \Phi_*^{\beta\alpha, \epsilon} \circ \Psi_*^\alpha$. The equalities in the above diagram follow from the isolation properties of f_ϵ^α and f_ϵ^β with respect to N !

Suppose we use different functions $f_\epsilon^{\alpha'}$, $f_\epsilon^{\beta'}$ and metric $g_\epsilon^{\alpha'}$, $g_\epsilon^{\beta'}$ as above. Then, we obtain isomorphisms $\Psi_*^{\alpha'}$, $\Psi_*^{\beta'}$ and $\Phi_*^{\beta\alpha, \epsilon'}$, which yields the isomorphisms $\Phi_*^{\beta\alpha'} = (\Psi_*^{\beta'})^{-1} \circ \Phi_*^{\beta\alpha, \epsilon'} \circ \Psi_*^{\alpha'}$ between $HM_*(Q^\alpha)$ and $HM_*(Q^\beta)$. Define

$$\Phi_*^{\alpha'\alpha} = \Psi_*^{\alpha'} \circ (\Psi_*^\alpha)^{-1} : HM_*(f_\epsilon^\alpha, g_\epsilon^\alpha, N, \mathfrak{o}^\alpha) \rightarrow HM_*(f_\epsilon^{\alpha'}, g_\epsilon^{\alpha'}, N, \mathfrak{o}^\alpha),$$

and similarly

$$\Phi_*^{\beta'\beta} = \Psi_*^{\beta'} \circ (\Psi_*^\beta)^{-1} : HM_*(f_\epsilon^\beta, g_\epsilon^\beta, N, \mathfrak{o}^\beta) \rightarrow HM_*(f_\epsilon^{\beta'}, g_\epsilon^{\beta'}, N, \mathfrak{o}^\beta),$$

where we use the equalities $HM_*(f_\epsilon^\alpha, g_\epsilon^\alpha, N^\alpha, \mathfrak{o}^\alpha) = HM_*(f_\epsilon^\alpha, g_\epsilon^\alpha, N, \mathfrak{o}^\alpha)$ and $HM_*(f_\epsilon^\beta, g_\epsilon^\beta, N^\beta, \mathfrak{o}^\beta) = HM_*(f_\epsilon^\beta, g_\epsilon^\beta, N, \mathfrak{o}^\beta)$. From the Morse homologies of Morse-Conley-Floer quadruples with N fixed and Theorem 2.4.4 we obtain the following commutative diagram:

$$\begin{array}{ccc}
 HM_*(f_\epsilon^\alpha, g_\epsilon^\alpha, N, \mathfrak{o}^\alpha) & \xrightarrow{\Phi_*^{\beta\alpha, \epsilon}} & HM_*(f_\epsilon^\beta, g_\epsilon^\beta, N, \mathfrak{o}^\beta) \\
 \downarrow \Phi_*^{\alpha'\alpha} & & \downarrow \Phi_*^{\beta'\beta} \\
 HM_*(f_\epsilon^{\alpha'}, g_\epsilon^{\alpha'}, N, \mathfrak{o}^\alpha) & \xrightarrow{\Phi_*^{\beta\alpha, \epsilon'}} & HM_*(f_\epsilon^{\beta'}, g_\epsilon^{\beta'}, N, \mathfrak{o}^\beta).
 \end{array}$$

Combining the isomorphisms yields

$$\Phi_*^{\beta\alpha'} = (\Psi_*^\beta)^{-1} \circ (\Phi_*^{\beta'\beta})^{-1} \circ \Psi_*^{\beta'\beta} \circ \Phi_*^{\beta\alpha, \epsilon'} \circ (\Psi_*^{\alpha'\alpha})^{-1} \circ \Psi_*^{\alpha'\alpha} \circ \Psi_*^\alpha = \Phi_*^{\beta\alpha},$$

which proves the independence of isolating blocks. \square

The isomorphisms in Theorem 2.4.7 satisfy the composition law. If we consider Morse-Smale quadruples $\mathcal{Q}^\alpha, \mathcal{Q}^\beta, \mathcal{Q}^\gamma$, we find isomorphisms:

$$\begin{aligned}\Phi_*^{\beta\alpha} &: HM_*(\mathcal{Q}^\alpha) \rightarrow HM_*(\mathcal{Q}^\beta), \\ \Phi_*^{\gamma\beta} &: HM_*(\mathcal{Q}^\beta) \rightarrow HM_*(\mathcal{Q}^\gamma), \\ \Phi_*^{\gamma\alpha} &: HM_*(\mathcal{Q}^\alpha) \rightarrow HM_*(\mathcal{Q}^\gamma).\end{aligned}$$

Theorem 2.4.8. *The composition law holds: $\Phi_*^{\gamma\alpha} = \Phi_*^{\gamma\beta} \circ \Phi_*^{\beta\alpha}$.*

Proof. The intersection $N = N^\alpha \cap N^\beta \cap N^\gamma$ is an isolating neighborhood. Proceeding as in the proof of the previous theorem, and using Theorem 2.4.6 on the isolating neighborhood N , shows that the composition law holds. \square

2.5 Morse-Conley-Floer homology

Theorem 2.4.8 implies that the Morse homology $HM_*(\mathcal{Q}^\alpha)$ is an inverse system with respect to the canonical isomorphisms

$$\Phi_*^{\beta\alpha} : HM_*(\mathcal{Q}^\alpha) \xrightarrow{\cong} HM_*(\mathcal{Q}^\beta).$$

This leads to the following definition.

Definition 2.5.1. Let S be an isolated invariant set for ϕ . The *Morse-Conley-Floer homology* of (S, ϕ) is defined by

$$HI_*(S, \phi) := \varprojlim HM_*(\mathcal{Q}^\alpha) \tag{2.12}$$

with respect to Morse-Conley-Floer quadruples \mathcal{Q}^α for (S, ϕ) and the associated canonical isomorphisms.

Proposition 2.5.2. *Suppose N is an isolating neighborhood with $S = \text{Inv}(N, \phi) = \emptyset$, then $HI_*(S, \phi) = 0$.*

Proof. Let f_ϕ be a Lyapunov function for S on N . Since $\text{Crit } f_\phi \cap N \subset S$ by Lemma 2.3.1, we have that f_ϕ has no critical points in N and f_ϕ is automatically Morse-Smale on N . It follows that $HM_*(f_\phi, g, \sigma, N) = 0$, for any metric g , and therefore it follows that $HI_*(S, \phi) = 0$. \square

Local continuation gives invariance of the Morse-Conley-Floer homology with respect to nearby flows. Consider a family of flows ϕ_λ , with $\lambda \in [0, 1]$, and denote the flow at $\lambda = 0$ by ϕ^α and the flow at $\lambda = 1$ by ϕ^β .

Proposition 2.5.3. *Let $N \subset M$ be compact, and let ϕ_λ a smooth family of flows as described above such that N is an isolating neighborhood for ϕ_λ for all $\lambda \in [0, 1]$, i.e. $S_\lambda = \text{Inv}(N, \phi_\lambda) \subset \text{int}(N)$. Set $S^\alpha = \text{Inv}(N, \phi^\alpha)$ and $S^\beta = \text{Inv}(N, \phi^\beta)$. Then there exists isomorphisms*

$$HI_*(S^\alpha, \phi^\alpha) \cong HI_*(S^\beta, \phi^\beta).$$

Proof. The first step in the proof is to construct a homotopy of Lyapunov functions. As before we take $\mu \in \mathbb{S}^1 \cong \mathbb{R}/2\mathbb{Z}$ and set $\lambda = \omega(\mu)$, where ω is defined in Section 2.4.2. Define the product flow $\Phi : \mathbb{R} \times M \times \mathbb{S}^1 \rightarrow M \times \mathbb{S}^1$ by $\Phi(t, x, \mu) = (\phi_{\omega(\mu)}(t, x), \mu)$. By assumption $N \times \mathbb{S}^1$ is an isolating neighborhood for Φ , and hence contains an isolating block $B \subset N \times \mathbb{S}^1$ with $\text{Inv}(B, \Phi) = \bigcup_{\mu \in [0, 1]} S_{\omega(\mu)}$. By Proposition 2.2.6 there exists a Lyapunov function $F_\Phi : M \times \mathbb{S}^1 \rightarrow \mathbb{R}$ for the flow Φ with Lyapunov property respect to B . The fibers $B_\mu = \{x \in M \mid (x, \mu) \in B\}$ are isolating neighborhoods for the flows $\phi_{\omega(\mu)}$ and the functions $f_{\phi_{\omega(\mu)}}^H(x) = F_\Phi(x, \mu)$ are Lyapunov functions for $\phi_{\omega(\mu)}$ on B_μ . Denote the Lyapunov functions at $\mu = 0, 1$ by $f_{\phi^\alpha}^\alpha$ and $f_{\phi^\beta}^\beta$ respectively and set $B^\alpha = B_0$ and $B^\beta = B_1$. We have now established a homotopy of Lyapunov functions $f_{\phi_{\omega(\mu)}}^H$.

Choose a metric g on M . Since by Corollary 2.3.12, $\mathcal{I}_{MS}(f_{\phi^\alpha}^\alpha; B^\alpha) \neq \emptyset$ and $\mathcal{I}_{MS}(f_{\phi^\beta}^\beta; B^\beta) \neq \emptyset$, then there exist Morse-Smale pairs (f^α, g^α) and (f^β, g^β) on B^α and B^β respectively. The associated homotopies are $(f_\lambda^\alpha, g_\lambda^\alpha)$ between (f^α, g^α) and $(f_{\phi^\alpha}^\alpha, g)$, and similarly $(f_\lambda^\beta, g_\lambda^\beta)$ between (f^β, g^β) and $(f_{\phi^\beta}^\beta, g)$. Define the homotopy

$$f_\lambda = \begin{cases} f_{3\lambda}^\alpha & \text{for } \lambda \in [0, \frac{1}{3}], \\ f_{\phi_{\omega(3\lambda-1)}}^{3\lambda-1} & \text{for } \lambda \in [\frac{1}{3}, \frac{2}{3}], \\ f_{3-3\lambda}^\beta & \text{for } \lambda \in [\frac{2}{3}, 1], \end{cases}$$

which is a piecewise smooth homotopy between f^α and f^β and similarly

$$g_\lambda = \begin{cases} g_{3\lambda}^\alpha & \text{for } \lambda \in [0, \frac{1}{3}], \\ g & \text{for } \lambda \in [\frac{1}{3}, \frac{2}{3}], \\ g_{3-3\lambda}^\beta & \text{for } \lambda \in [\frac{2}{3}, 1]. \end{cases}$$

Furthermore we have the isolating neighborhood,

$$N_\lambda = \begin{cases} B^\alpha & \text{for } \lambda \in [0, \frac{1}{3}], \\ B_{3\lambda-1} & \text{for } \lambda \in [\frac{1}{3}, \frac{2}{3}], \\ B^\beta & \text{for } \lambda \in [\frac{2}{3}, 1]. \end{cases}$$

As before we tacitly assume a reparametrization of the variable λ to make the above homotopies smooth. We denote the negative gradient flows of the homotopy (f_λ, g_λ) by $\psi_{(f_\lambda, g_\lambda)}^\lambda$ and by assumption and construction (cf. Lemma 2.3.3) $\text{Inv}(N_\lambda, \psi_{(f_\lambda, g_\lambda)}^\lambda) \subset \text{int}(N_\lambda)$.

Take $\mu \in \mathbb{S}^1 \cong \mathbb{R}/2\mathbb{Z}$ and set $\lambda = \omega(\mu)$, where ω is defined in Section 2.4.2 and consider the negative gradient flow $\Psi_\kappa^\times : \mathbb{R} \times M \times \mathbb{S}^1 \rightarrow M \times \mathbb{S}^1$, given by the function $F(x, \mu) = f_{\omega(\mu)}(x) + r[1 + \cos(\pi\mu)]$ and product metric $G_{(x, \mu)}^\times = (g_{\omega(\mu)})_x \oplus \frac{1}{\kappa} d\mu^2$. For r large and κ small, and a small perturbation G of the metric G^\times we obtain a Morse-Smale pair (F, G) on $C = \{(x, \mu) : x \in N_{\omega(\mu)}\}$. We can now repeat the proof of Theorem 2.4.4 to conclude isomorphisms

$$HM_*(Q^\alpha) \cong HM_*(Q^\beta).$$

If in the above construction we choose different Morse-Conley-Floer quadruples $Q^{\alpha'}$ and $Q^{\beta'}$ we obtain the commutative square of canonical isomorphisms

$$\begin{array}{ccc} HM_*(Q^\alpha) & \xrightarrow{\cong} & HM_*(Q^{\alpha'}) \\ \downarrow \cong & & \downarrow \cong \\ HM_*(Q^\beta) & \xrightarrow{\cong} & HM_*(Q^{\beta'}) \end{array}$$

which shows that Morse-Conley-Floer homologies are isomorphic. □

The Morse-Conley-Floer homology is invariant under global deformations if the flows are related by continuation. This property allows for the computation of the index in highly non-linear situations, by continuing the flow to a flow that is better understood.

Definition 2.5.4. Isolated invariant sets S^α, S^β for the flows ϕ^α and ϕ^β respectively, are said to *related by continuation* if there exists a smooth homotopy ϕ_λ of flows, and a partition λ_i of $[0, 1]$, i.e.

$$0 = \lambda_0 < \lambda_1 < \dots < \lambda_n = 1,$$

along with compact sets $N^\alpha = N_0, \dots, N_{n-1} = N^\beta$, such that N_i are isolating neighborhoods for all flows ϕ_λ , for all $\lambda_i \leq \lambda \leq \lambda_{i+1}$, and $\text{Inv}(N_{i-1}, \phi_{\lambda_i}) = \text{Inv}(N_i, \phi_{\lambda_i})$, $\text{Inv}(N^\alpha, \phi^\alpha) = S^\alpha$, and $\text{Inv}(N^\beta, \phi^\beta) = S^\beta$.

Composing the isomorphisms yields a global continuation theorem.

Theorem 2.5.5. *Let (S^α, ϕ^α) and (S^β, ϕ^β) be related by continuation. Then there exists canonical isomorphisms*

$$HI_*(S^\alpha, \phi^\alpha) \cong HI_*(S^\beta, \phi^\beta).$$

2.6 Morse decompositions and connection matrices

If the dynamics restricted to an isolated invariant set S is not recurrent a decomposition extracting gradient dynamics exists and leads to the concept of *Morse decomposition*, which generalizes the attractor-repeller pair.

Definition 2.6.1. Let $S = \text{Inv}(N, \phi)$ be an isolated invariant set. A family $S = \{S_i\}_{i \in I}$, indexed by a finite poset (I, \leq) , consisting of non-empty, compact, pairwise disjoint, invariant subsets $S_i \subset S$, is a *Morse decomposition* for S if, for every $x \in S \setminus (\bigcup_{i \in I} S_i)$, there exists $i < j$ such that

$$\alpha(x) \subset S_j \quad \text{and} \quad \omega(x) \subset S_i.$$

The sets $S_i \subset S$ are called *Morse sets*. The set of all Morse decompositions of S under ϕ is denoted by $\text{MD}(S, \phi)$.

Let $B \subset N$ be an isolating block for S . In the proof of Proposition 2.2.6 we constructed a flow $\phi^\dagger : \mathbb{R} \times B^\dagger \rightarrow B^\dagger$ as an extension of ϕ , where B^\dagger is a compact submanifold with piecewise smooth boundary. Recall the attractor A_V and the repeller A_U^* . From the duals we have $S = A_U \cap A_V^*$, cf. Figure 2.3. Now consider an attractor A in S , and its dual repeller $R = A|_S^*$ in S . Because A, R are isolated invariant sets in S , and S is isolated invariant in B^\dagger , A, R are also isolated invariant on B^\dagger for the flow ϕ^\dagger , see [CZ84]. By a similar reasoning A , and R are also isolated invariant sets on M for the flow ϕ . The above situation is sketched in Figure 2.6.

Lemma 2.6.2 (cf. [VKM13, KMV13]). *The isolated invariant set B^\dagger has a natural Morse decomposition $A_V < S < A_U^*$. The invariant set*

$$A_V A = A_V \cup W^u(A, \phi^\dagger) \subset A_U = A_V \cup W^u(S, \phi^\dagger),$$

is an attractor in A_U and therefore an attractor in B^\dagger .




Figure 2.6: The Morse decomposition of B^\dagger , given in Lemma 2.6.2. In Figure 2.3 the sets A_U and A_V^* are depicted.

Proof. From the proof of Proposition 2.2.6 it follows that $A_V < S$ is an attractor-repeller in A_U . The set R is a repeller in S and therefore is a repeller in A_U . This provides the Morse decomposition

$$A_V < A < R,$$

for A_U . From this we derive that $A_V A < R$, with

$$A_V A = A_V \cup W^u(A, \phi^\dagger),$$

is an attractor-repeller pair in A_U and $A_V A$ is an attractor in A_U and thus $A_V A$ is an attractor in B^\dagger , which proves the lemma. \square

For a given isolated invariant sets S , with isolating block B , Lemma 2.6.2 yields the following filtration of attractors in B^\dagger of the extended flow ϕ^\dagger :

$$\emptyset \subset A_V \subset A_V A \subset A_U \subset B^\dagger.$$

The associated Morse decomposition of B^\dagger is given by

$$A_V < A < R < A_U^*,$$

indeed,

$$A_V^* \cap A_V A = A_V|_{A_V A}^* = A, \quad \text{and} \quad A_V A^* \cap A_U = A_V A|_{A_U}^* = R.$$

From the proof of Proposition 2.2.6 we have the Lyapunov functions

$$f_{A_V}, f_{A_V A}, \quad \text{and} \quad f_{A_U}.$$

Consider the positive linear combination

$$f_{\phi^\dagger}^\epsilon = \lambda f_{A_V} + \mu f_{A_U} + \epsilon f_{A_V A}$$

which is constant on the Morse sets A_V , A , R and A_U^* and decreases along orbits of ϕ^\dagger . The values at the Morse sets are: $f_{\phi^\dagger}^\epsilon|_{A_V} = 0$, $f_{\phi^\dagger}^\epsilon|_A = \lambda$, $f_{\phi^\dagger}^\epsilon|_R = \lambda + \epsilon$ and $f_{\phi^\dagger}^\epsilon|_{A_U^*} = \lambda + \mu + \epsilon$. The function $f_{\phi^\dagger}^\epsilon$ is also Lyapunov function for $(A \cup R, \phi)$, because it satisfies the Lyapunov property on a suitable chosen isolating neighborhood of $A \cup R$ for the flow ϕ , which we therefore denote by f_ϕ^ϵ . Moreover, f_ϕ^ϵ is an ϵ -perturbation of $h_\epsilon^0 = \lambda f_{A_V} + \mu f_{A_U} \in \text{Lyap}(S, \phi)$.

Lemma 2.6.3. *Let S be an isolated invariant set and let $A \subset S$ be an attractor in S . Then,*

$$P_t(A, \phi) + P_t(R, \phi) = P_t(S, \phi) + (1 + t)Q_t,$$

where $P_t(A, \phi)$, $P_t(R, \phi)$ and $P_t(S, \phi)$ are the Poincaré polynomials of the associated Morse-Conley-Floer homologies, and Q_t is a polynomial with non-negative coefficients.

Proof. Let g be a Riemannian metric on M and let f_ϕ^ϵ be the Lyapunov function for the attractor-repeller pair (A, R) as described above. By Proposition 2.3.9 there exists ϵ - C^2 close perturbation f of f_ϕ^ϵ , such that (f, g) is a Morse-Smale pair on B . Then, if $\epsilon > 0$ is sufficiently small, f is also a small perturbation of f_ϕ^0 . The next step is to consider the algebra provided by f . The latter implies that $HM_*(f, g, B, \mathfrak{o}) \cong HI_*(S, \phi)$, where the chain complex for f is given by

$$\partial_k^S : C_k(S) \rightarrow C_{k-1}(S),$$

with $\partial_k^S = \partial_k(f, g, B, \mathfrak{o})$ and $C_k(S) = C_k(f, B)$, and $H_*(C_*(S), \partial_*^S) = HM_*(f, g, B, \mathfrak{o})$.

Let B_A and B_R be a isolating blocks for A and R respectively, then the Morse-Conley-Floer homologies of A and R are defined by restricting the count of the critical points and connecting orbits of f to B_A and B_R . We obtain the chain complexes:

$$\partial_k^A : C_k(A) \rightarrow C_{k-1}(A), \quad \text{and} \quad \partial_k^R : C_k(R) \rightarrow C_{k-1}(R),$$

with $\partial_k^A = \partial_k(f, g, B_A, \mathfrak{o})$ and $C_k(A) = C_k(f, B_A)$ and similarly $\partial_k^R = \partial_k(f, g, B_R, \mathfrak{o})$ and $C_k(R) = C_k(f, B_R)$. The homologies satisfy:

$$\begin{aligned} H_*(C_*(A), \partial_*^A) &= HM_*(f, g, B_A, \mathfrak{o}) \cong HI_*(A, \phi); \\ H_*(C_*(R), \partial_*^R) &= HM_*(f, g, B_R, \mathfrak{o}) \cong HI_*(R, \phi). \end{aligned}$$

Because of the properties of f we see that $C_k(S) = C_k(A) \oplus C_k(R)$. The gradient system defined by f only allows connections from R to A and not vice versa, and therefore for the negative g -gradient flow $\psi_{(f,g)}$ defined by (f, g) , the sets

$$A_{(f,g)} = \text{Inv}(B_A, \psi_{(f,g)}) < R_{(f,g)} = \text{Inv}(B_R, \psi_{(f,g)})$$

is a Morse decomposition for $S_{(f,g)} = \text{Inv}(B, \psi_{(f,g)})$. This Morse decomposition provides additional information about the boundary operator ∂_*^S :

$$\partial_k^S = \begin{pmatrix} \partial_k^A & \delta_k \\ 0 & \partial_k^R \end{pmatrix} : C_k(A) \oplus C_k(R) \rightarrow C_{k-1}(A) \oplus C_{k-1}(R),$$

where $\delta_k : C_k(R) \rightarrow C_{k-1}(A)$ counts connections from $R_{(f,g)}$ to $A_{(f,g)}$ under ψ . For the sub-complexes $(C_*(A), \partial_*^A)$ and $(C_*(R), \partial_*^R)$ provide natural inclusions and projections

$$i_k : C_k(A) \rightarrow C_k(S) = C_k(A) \oplus C_k(R), \quad a \mapsto (a, 0),$$

and

$$j_k : C_k(S) = C_k(A) \oplus C_k(R) \rightarrow C_k(R), \quad (a, r) \mapsto r.$$

The maps i_* and j_* are chain maps. Indeed, for $a \in C_k(A)$, and $r \in C_k(R)$, we have

$$i_{k-1} \partial_k^A(a) = (0, \partial_k^A a) = \partial_k^S(0, a) = \partial_k^S i_k(a),$$

and

$$j_{k-1}\partial_k^S(a, r) = j_{k-1}(\partial_k^A a + \delta_k r, \partial_k^R r) = \partial_k^R r = \partial_k^R j_k(a, r).$$

This implies that the induced homomorphisms $i_k^* : H_k(C_*(A)) \rightarrow H_{k-1}(C_*(A))$ and $j_k^* : H_k(C_*(R)) \rightarrow H_{k-1}(C_*(R))$ are well-defined. The maps i_* and j_* define the following short exact sequence

$$0 \rightarrow C_k(A) \xrightarrow{i_k} C_k(S) \xrightarrow{j_k} C_k(R) \rightarrow 0. \quad (2.13)$$

Since i_* and j_* are chain maps, this is actually a short exact sequence of chain maps, and by the snake lemma we obtain the following long exact sequence in homology:¹¹

$$\dots \xrightarrow{\delta_{k+1}} H_k(A) \xrightarrow{i_k} H_k(S) \xrightarrow{j_k} H_k(R) \xrightarrow{\delta_k} H_{k-1}(A) \xrightarrow{i_{k-1}} \dots$$

We recall that the connecting homomorphisms $H_k(R) \xrightarrow{\delta_k} H_{k-1}(A)$ are established as follows. Let $r \in \ker \partial_k^R$, then $[(i_k^{-1} \circ \partial_k^S \circ j_k^{-1})(r)] \in H_{k-1}(A)$. Indeed, $j_k^{-1}(r) = (C_k(A), r)$ and $\partial_k^S(j_k^{-1}(r)) = (\partial_k^A C_k(A) + \delta_k r, \partial_k^R r) = (\partial_k^A C_k(A) + \delta_k r, 0)$. Furthermore, $i_k^{-1}(\partial_k^S(j_k^{-1}(r))) = \partial_k^A C_k(A) + \delta_k r = \delta_k r + \text{im } \partial_k^A \in \ker \partial_{k-1}^A$, since $\partial_{k-1}^A \delta_k r = -\delta_{k-1} \partial_k^R r = 0$. The class of $i_k^{-1}(\partial_k^S(j_k^{-1}(r)))$ is a homology class in $H_{k-1}(R)$ and we write $[\delta_k r] = \delta_k[r]$.

Since B is compact the chain complexes involved are all finite dimensional and terminate at $k = -1$ and $k = \dim M + 1$ and therefore the Poincaré polynomials are well-defined. The Poincaré polynomials satisfy

$$P_t(H_*(A)) + P_t(H_*(R)) - P_t(H_*(S)) = (1+t)Q_t,$$

where $Q_t = \sum_{k \in \mathbb{Z}} (\text{rank } \delta_k) t^k$, which completes the proof. \square

By definition ∂_*^A and ∂_*^R are zero maps on homology and therefore induced maps on homology give rise to a map

$$\Delta_k = \begin{pmatrix} 0 & \delta_k \\ 0 & 0 \end{pmatrix} : H_k(A) \oplus H_k(R) \rightarrow H_{k-1}(A) \oplus H_{k-1}(R).$$

¹¹The homologies are abbreviated as follows: $H_k(A) = H_k(C_*(A), \partial_*^A)$, $H_k(R) = H_k(C_*(R), \partial_*^R)$ and $H_k(S) = H_k(C_*(S), \partial_*^S)$.

Define $\Delta = \bigoplus_{k \in \mathbb{Z}} \Delta_k$ and

$$\Delta : \bigoplus_{k \in \mathbb{Z}} \left(H_k(A) \oplus H_k(R) \right) \rightarrow \bigoplus_{k \in \mathbb{Z}} \left(H_k(A) \oplus H_k(R) \right),$$

with the property $\Delta^2 = 0$ and is called the *connection matrix* for the attractor-repeller pair (A, R) . The above defined complex is a chain complex and the associated homology is isomorphic to $H_*(S)$, see [Mis95].

For a given Morse decomposition $S = \{S_i\}_{i \in I}$ of S the Morse sets S_i are again isolated invariant sets and therefore their Morse-Conley-Floer homology $HI_*(S_i, \phi)$ is well-defined. The associated Poincaré polynomials satisfy the Morse-Conley relations.

Theorem 2.6.4. *Let $S = \{S_i\}_{i \in I}$ be a Morse decomposition for an isolated invariant set S . Then,*

$$\sum_{i \in I} P_t(S_i, \phi) = P_t(S, \phi) + (1 + t)Q_t, \quad (2.14)$$

where $P_t(S, \phi)$ is the Poincaré polynomial of $HI_*(S, \phi)$, and Q_t is a polynomial with non-negative coefficients. These relations are called the Morse-Conley relations and generalize the classical Morse relations for gradient flows.

Proof. The Morse decomposition S is equivalent to a lattice of attractors A , see [VKM13]. We choose a chain

$$A_0 = \emptyset \subset A_1 \subset \dots \subset A_n = S$$

in A , such that $S_i = A_i \cap A_{i-1}^* = A_{i-1}|_{A_i}^*$. Each attractor $A_i \subset S$ is an isolated invariant set for ϕ in B (see [CZ84]) and for each A_i we have the attractor-repeller pair $A_{i-1} \subset A_{i-1}|_{A_i}^* = S_i$. From Lemma 2.6.3 we derive the attractor-repeller pair Morse-Conley relations for each attractor-repeller pair (A_{i-1}, S_i) in A_i :

$$P_t(A_{i-1}, \phi) + P_t(S_i, \phi) - P_t(A_i, \phi) = (1 + t)Q_t^i.$$

Summing i from $i = 1$ through to $i = n$ we obtain the Morse-Conley relations (2.14), which proves the theorem. \square

Remark 2.6.5. In Chapter 5 we give a slightly different proof of the Morse-Conley relations. The relations follow from a spectral sequence associated to the choice of a Lyapunov function compatible with the Morse decomposition of the isolated invariant set. This amounts to an order coursening to a linear order. The dependence of connection matrices on the choices of Lyapunov functions warrants further study.

2.7 Relative homology of blocks

An important property of Morse-Conley-Floer homology is that it can be computed in terms of singular relative homology of a topological pair of a defining blocks as pointed out in Property (vi) in Section 2.1.

Theorem 2.7.1. *Let S be an isolating neighborhood for ϕ and let B be an isolating block for S . Then*

$$HI_*(S, \phi) \cong H_*(B, B_-; \mathbb{Z}), \quad \forall k \in \mathbb{Z}, \quad (2.15)$$

where $B_- = \{x \in \partial B \mid X(x) \text{ is outward pointing}\}$ and is called the ‘exit set’.

Note that in the case that ϕ is the gradient flow of a Morse function, then Theorem 2.7.1 recovers the results of Morse homology. Theorem 2.7.1 also justifies the terminology Morse-Conley-Floer homology, since the construction uses Morse/Floer homology and recovers the classical homological Conley index. The following lemma states that one can choose a metric g such for any Lyapunov function f_ϕ the boundary behavior of $-\nabla_g f_\phi$ coincides with X .

Recall Definition 2.2.1 of an isolating block B . Because B is a manifold with piecewise smooth boundary, we need to be careful when we speak of the boundary behavior of B . We say $h_+ : M \rightarrow \mathbb{R}$ defines the boundary B_+ if

$$Xh_+|_{B_+} > 0, \quad dh_+|_{B_+} \neq 0, \quad \text{and} \quad h_+|_{B \setminus B_+} > 0.$$

Analogously we say that $h_- : M \rightarrow \mathbb{R}$ defines the boundary B_- if

$$Xh_-|_{B_-} < 0, \quad dh_-|_{B_-} \neq 0, \quad \text{and} \quad h_-|_{B \setminus B_-} > 0.$$

Lemma 2.7.2. *Let $f_\phi : M \rightarrow \mathbb{R}$ be a smooth Lyapunov function for (S, ϕ) with the Lyapunov property with respect to an isolating block B (cf. Proposition 2.2.6). Then there exists a metric g on M which satisfies the property: $(-\nabla_g f_\phi)h_- < 0$ for all $x \in B_-$ and $(-\nabla_g f_\phi)h_+ > 0$ for all $x \in B_+$.*

Proof. Since $X \neq 0$ on ∂B , there exists an open U containing ∂B , such that $X \neq 0$, and $Xf_\phi < 0$ on U . The span of X defines an one dimensional vector subbundle E of the tangent bundle $TM|_U$ over U . There is a complementary subbundle $E^\perp \subset TM|_U$ such that $TM|_U \cong E \oplus E^\perp$. Define the inner product e_E on E by $e_E(X, X) = 1$, and let e_{E^\perp} be any inner product on E^\perp . Declaring that E and E^\perp are orthogonal defines a metric e on $TM|_U$. Clearly $e(-\nabla_e f_\phi, X) = -Xf_\phi$, and therefore

$$-\nabla_e f_\phi = (-Xf_\phi)X + Y, \quad \text{with} \quad Y \in E^\perp.$$

Since $-Xf_\phi > 0$ on U , we can rescale the metric e to $g = (-Xf_\phi)e_E + \frac{1}{\epsilon}e_{E^\perp}$, for some $\epsilon > 0$. It follows that $-\nabla_g f_\phi = X + \epsilon Y$. For the boundary defining functions h_\pm we find

$$(-\nabla_g f_\phi)h_\pm = Xh_\pm + \epsilon Yh_\pm.$$

Because ∂B is compact there exists a uniform bound $|Y(h_\pm)(x)| \leq C$ for all $x \in \partial B$ independent of ϵ . For $\epsilon > 0$ sufficiently small the sign of $(-\nabla_g f_\phi)h_\pm$ agrees with the sign of Xh_\pm on ∂B_\pm . Via a standard partition of unity argument we extend the metric g to M , without altering it on ∂B , which gives a metric with the desired properties. \square

Proof of Theorem 2.7.1. By Lemma 2.7.2 we can choose a Riemannian metric g such that $-\nabla_g f_\phi$ has the same boundary behavior as the vector field X . Using Proposition 2.3.9 we have a small C^2 -perturbation f of f_ϕ such that $(f, g) \in \mathcal{S}_{MS}(f_\phi; B)$, via a constant homotopy in g . Since f is sufficiently close to f_ϕ the boundary behavior of $-\nabla_g f$ does not change! From the definition of the Morse-Conley-Floer homology we have that, for $\mathcal{Q} = (f, g, B, \mathfrak{o})$,

$$HI_*(S, \phi) \cong HM_*(\mathcal{Q}).$$

It remains to compute the Morse homology of the Morse-Conley-Floer quadruple \mathcal{Q} . Relating the Morse homology to the topological pair (B, B_-) is the same as in the case when $B = M$, which is described in [BH04] and [Sal90]. The arguments can be followed verbatim and therefore we only provide a sketch of the proof.

The first part of the of the proof starts with the boundary operator $\partial_*(\mathcal{Q})$. The latter can be related to the boundary operator in relative singular homology. Since the negative g -gradient flow $\psi_{(f,g)}$ is Morse-Smale all critical points $x \in \text{Crit}(f) \cap B$ are isolated invariant sets for $\psi_{(f,g)}$. Let B^x be an isolating block for $S = \{x\}$. From standard Morse theory and Wazeski's principle it follows that

$$H_k(B^x, B_-^x; \mathbb{Z}) \cong C_k(f, B^x) \cong \mathbb{Z}, \quad \text{for } k = |x|,$$

and $H_k(B^x, B_-^x; \mathbb{Z}) = 0$ for $k \neq |x|$. For critical points $y, x \in \text{Crit}(f) \cap B$, with $|x| = |y| + 1 = k$ we define the set

$$S(x, y) = W_B(x, y) \cup \{x, y\},$$

which is an isolated invariant set with isolating neighborhood N . Let c be such that $f(y) < c < f(x)$. For $T > 0$ sufficiently large, and $\epsilon > 0$ sufficiently small, define the isolating block

$$B^x = \{z \in N \mid \psi_{(f,g)}(-t, z) \in N, \quad f(\psi_{(f,g)}(-t, z)) \leq f(x) + \epsilon, \\ \forall 0 \leq t \leq T, \quad f(z) \geq c\},$$

for $\{x\}$, and the isolating block

$$B^y = \{z \in N \mid \psi_{(f,g)}(t, z) \in N, \quad f(\psi_{(f,g)}(t, z)) \geq f(y) - \epsilon, \\ \forall 0 \leq t \leq T, \quad f(z) \leq c\},$$

for $\{y\}$. The exit sets are

$$B_-^x = \{z \in B^x \mid f(z) = c\}, \\ B_-^y = \{z \in B^y \mid f(\psi_{(f,g)}(T, z)) = f(y) - \epsilon\}.$$

Define the sets $B_2^{x,y} = B^x \cup B^y$, $B_1^{x,y} = B^y \cup B_-^x$, and $B_0^{x,y} = B_-^y \cup \text{cl}(B_-^x \setminus B^x)$. The set $B_2^{x,y}$ is an isolating block for $S(x, y)$, the set $\text{cl}(B_2^{x,y} \setminus B_1^{x,y})$ is an isolating block for $\{x\}$, and the set $\text{cl}(B_1^{x,y} \setminus B_0^{x,y})$ is an isolating block for $\{y\}$. Via the triple $B_0^{x,y} \subset B_1^{x,y} \subset B_2^{x,y}$ we define the operator $\Delta_k : H_k(B_2^{x,y}, B_1^{x,y}) \rightarrow H_{k-1}(B_1^{x,y}, B_0^{x,y})$ by the commutative diagram

$$\begin{array}{ccc} H_*(B^y, B_-^y) & \xrightarrow{\Delta_k} & H_*(B^x, B_-^x) \\ \downarrow \cong & & \downarrow \cong \\ H_*(B_2^{x,y}, B_1^{x,y}) & \xrightarrow{\delta_k} & H_*(B_1^{x,y}, B_0^{x,y}). \end{array}$$

The vertical maps express the homotopy invariance of the Conley index, and the horizontal map is the connecting homomorphism in the long exact sequence of the triple. The homomorphism Δ_* can be defined on $C_*(f, B)$ directly and the analysis in [BH04] and [Sal90] shows that $\Delta_* = \partial_*(\mathcal{Q})$, which that yields that orbit counting can be expressed in terms of algebraic topology.

The next step to the apply this to the isolated invariant set $S_{(f,g)}$. Following the proofs in [BH04] and [Sal90] we construct a special Morse decomposition of $S_{(f,g)}$. Let

$$S^{k,\ell} = \bigcup \{W_B(y, x) \mid k \leq |x| \leq |y| \leq \ell\},$$

and since $\psi_{(f,g)}$ is Morse-Smale, these sets are compact isolated invariant sets contained in $S_{(f,g)}$, with $S^{k,\ell} = \emptyset$ for $k < \ell$. The sets $\{S^{k,k}\}$ form a Morse decomposition of $S_{(f,g)}$ via $S^{k,k} \leq S^{\ell,\ell}$ if and only if $k \leq \ell$. This yields a filtration of blocks B_k

$$B_- \subset B_0 \subset \dots \subset B_{m-1} \subset B_m = B,$$

such that $\text{Inv}(B_\ell \setminus B_{k-1}) = S^{\ell,k}$. We now use a modification of arguments in [BH04] and [Sal90].

As before $H_k(B_k, B_{k-1}) \cong C_k(f, B)$ and $H_\ell(B_k, B_{k-1}) = 0$ for $\ell \neq k$. Consider the triple $B_- \subset B_{k-1} \subset B_k$ and the associated homology long exact sequence

$$\rightarrow H_{\ell+1}(B_k, B_{k-1}) \rightarrow H_\ell(B_{k-1}, B_-) \rightarrow H_\ell(B_k, B_-) \rightarrow H_\ell(B_k, B_{k-1}) \rightarrow \dots$$

For $\ell \neq k-1, k$ the sequence reduces to

$$0 \rightarrow H_\ell(B_{k-1}, B_-) \rightarrow H_\ell(B_k, B_-) \rightarrow 0,$$

which shows that $H_\ell(B_{k-1}, B_-) \cong H_\ell(B_k, B_-)$ for $\ell \neq k-1, k$. By a left and right induction argument we obtain that $H_\ell(B_k, B_-) = 0$ for all $\ell > k$ and $H_\ell(B_k, B_-) \cong H_\ell(B, B_-)$ for all $\ell < k$. If we also use the homology long exact sequences for the triples $B_- \subset B_{k-2} \subset B_{k-1}$ and $B_{k-2} \subset B_{k-1} \subset B_k$ we obtain the commuting diagram

$$\begin{array}{ccccccc} & & & & 0 & & \\ & & & & \downarrow & & \\ 0 & \longrightarrow & H_k(B_k, B_-) & \longrightarrow & H_k(B_k, B_{k-1}) & \xrightarrow{\delta_k} & H_{k-1}(B_k, B_-) \\ & & & & \searrow^{\delta'_k} & & \downarrow \\ & & & & & & H_{k-1}(B_{k-1}, B_{k-2}). \end{array}$$

The following information can be deduced from the diagram. The vertical exact sequence we derive that the map $H_{k-1}(B_k, B_-) \rightarrow H_{k-1}(B_{k-1}, B_{k-2})$ is injective and thus from the commuting triangle we conclude that $\ker \delta_k = \ker \delta'_k \cong \ker \partial_k$. From the horizontal exact sequence we obtain that $H_k(B_k, B_-) \cong \ker \delta_k \cong \ker \partial_k$. From these isomorphisms we obtain the following commuting diagram of short exact sequences:

$$\begin{array}{ccccccc}
 H_{k+1}(B_{k+1}, B_k) & \xrightarrow{\delta_{k+1}} & H_k(B_k, B_-) & \xrightarrow{\cong} & H_k(B_{k+1}, B_-) & \longrightarrow & 0 \\
 \downarrow \cong & & \downarrow \cong & & \downarrow \cong & & \\
 C_{k+1}(f, B) & \xrightarrow{\partial_{k+1}} & \ker \partial_k & \xrightarrow{\cong} & H_k(B, B_-) & \longrightarrow & 0,
 \end{array}$$

where we use long exact sequence of the triple $B_- \subset B_k \subset B_{k+1}$ and the fact that $H_k(B_{k+1}, B_k) = 0$. The diagram implies that $HI_k(S, \phi) \cong \ker \partial_k / \text{im } \partial_{k+1} \cong H_k(B, B_-)$, which completes the proof. \square

Remark 2.7.3. Theorem 2.7.1 relates Morse-Conley-Floer homology to the singular homology of a pair (B, B_-) . The proof uses the fact that Morse-Conley-Floer homology is well-defined as the Morse homology of a Morse-Conley-Floer quadruple. In fact one proves that Morse-Conley-Floer homology is isomorphic to Morse homology of manifold pairs as developed in [Sch93]. To be more precise let $B \subset D$ and D is a manifold with boundary $\partial D = D_+$. Let $E \subset D$ be a manifold pair such that $B = \text{cl}(D \setminus E)$ and $B_- = B \cap E$. In order to define $HM_*(D, E)$ we consider Morse-Smale pairs (f, g) with $-\nabla_g f$ inward pointing on $\partial D = D_+$ and ‘pointing into’ E on ∂E . Then $HI_*(S, \phi) \cong HM_*(D, E) \cong H_*(D, E) \cong H_*(B, B_-)$. It is worthwhile to develop Morse homologies for blocks (B, B_-) .

Chapter 3

Functoriality in Morse and Morse-Conley-Floer homology

Isolation is [...] good for me

Lemon Tree – Fools Garden

3.1 Introduction

Morse homology is commonly used as an invariant for manifolds. Other homology theories can also be interpreted as invariants for manifolds. However the theory gains more power when one realizes that maps between manifolds induce maps between the simplicial homology groups. These induced maps satisfy an important property, they are functorial: induced maps of compositions are compositions of induced maps, and the identity induces the identity on homology.

In this chapter we address functorial properties of Morse homology, local Morse homology and Morse-Conley-Floer homology with respect to smooth maps between manifolds. Functoriality in Morse homology on closed manifolds is known [AS10, AZ08, AD14, KM07, Sch93], however no proofs are given through the analysis of moduli spaces. This analysis is done in Sections 3.2 to 3.4. These sections are of independent interest to the rest of the chapter. In Section 3.5 we discuss isolation properties of maps, which are important for the functoriality in local Morse homology and Morse-Conley-Floer

homology. This functoriality is discussed in Sections 3.6 and 3.7. In Section 3.8 we prove that the transverse maps that are important for defining the induced maps in Morse homology are generic. We start by discussing the results in more detail.

3.1.1 Morse homology and local Morse Homology

A *Morse datum* is a quadruple¹ $Q^\alpha = (M^\alpha, f^\alpha, g^\alpha, \mathfrak{o}^\alpha)$, where M^α is a choice of closed manifold and $(f^\alpha, g^\alpha, \mathfrak{o}^\alpha)$ is a Morse-Smale triple on M^α . Thus f^α is a Morse function on M^α , and g^α is a metric such that the stable and unstable manifolds of f^α intersect transversely. Finally \mathfrak{o}^α denotes a choice of orientations of the unstable manifolds. The Morse homology $HM_*(Q^\alpha)$ is defined as the homology of the chain complex of $C_*(Q^\alpha)$ which is freely generated by the critical points of f^α and graded by their index, with boundary operator $\partial_*(Q^\alpha)$ counting connecting orbits of critical points with Morse index difference of 1 appropriately with sign, cf. [BH04, Sch93, Web06] and Section 2.1.1. If Q^β is another choice of Morse datum with $M^\beta = M^\alpha := M$ there is a canonical isomorphism $\Phi_*^{\beta\alpha} : HM_*(Q^\alpha) \rightarrow HM_*(Q^\beta)$ induced by continuation. The Morse homology of the manifold M is defined by

$$HM_*(M) = \varprojlim HM_*(Q^\alpha), \quad (3.1)$$

where the inverse limit is taken over all Morse data with the canonical isomorphisms. The Morse homology $HM_*(M^\alpha)$ is isomorphic to the singular homology $H_*(M^\alpha)$.

Important in what follows is that this can also be done locally²: suppose that (f^α, g^α) is Morse-Smale on N^α , cf. Definition 2.3.5. Thus N^α is an isolating neighborhood of the gradient flow ψ^α of (f^α, g^α) such that all critical points of f^α are non-degenerate on N^α and the local stable and unstable manifolds intersect transversely. Then $\mathcal{P}^\alpha = (M^\alpha, f^\alpha, g^\alpha, N^\alpha, \mathfrak{o}^\alpha)$, with \mathfrak{o}^α a choice of orientation of the local unstable manifolds is a *local Morse datum*. The local Morse homology $HM_*(\mathcal{P}^\alpha)$ is defined by a similar counting procedure for Morse homology. However, now only critical points and connecting orbits that are contained in N^α are counted. If the gradient flows associated to two local Morse data \mathcal{P}^α and \mathcal{P}^β are isolated homotopic³ then there are canonical

¹In this chapter we need to keep track of the base manifolds, which is why we have quadruples of data.

²Here we do not need to assume M^α is closed anymore.

³We use here the fact that two gradient flows are isolated homotopic through gradient flows if and only if they are isolated homotopic through arbitrary flows, see the proof of Proposition 2.5.3.

isomorphisms $\Phi_*^{\beta\alpha} : HM_*(\mathcal{P}^\alpha) \rightarrow HM_*(\mathcal{P}^\beta)$ induced by the continuation map.

Local Morse homology is *not* a topological invariant for N^α . It measures dynamical information of the gradient flow of f^α and g^α . There is still stability under continuation however. Now, given any function f and metric g such that N is an isolating neighborhood of the gradient flow, which we do not assume to be Morse-Smale, we define the local Morse homology of such a triple via

$$HM_*(f, g, N) = \varprojlim HM_*(\mathcal{P}^\alpha).$$

The inverse limit is taken over all Morse-Smale data \mathcal{P}^α on $N^\alpha = N$ whose gradient flow is isolated homotopic to the gradient flow ψ of (f, g) on N .

If M is a compact manifold, and we take $N = M$, the local Morse homology does not depend on the function and the metric anymore, as all gradient flows are isolated homotopic to each other. In this case the local Morse homology recovers the Morse homology defined in Equation (3.1).

Results on Morse homology on compact manifolds with boundary fall in this framework. Let M be a compact manifold with boundary. Assuming that the gradient is not tangent to the boundary, we can endow the boundary components with collars to obtain a manifold \widetilde{M} without boundary. The function and metric extend to \widetilde{M} . Then $M \subset \widetilde{M}$ is an isolating neighborhood of the gradient flow. One can compute this Morse homology in terms of the singular homology of M as $HM_*(f, g, M) \cong H_*(M, \partial M_-)$, where ∂M_- is the union of the boundary components where the gradient of f points outwards. See also the discussion in Section 2.1.1.

3.1.2 Morse-Conley-Floer homology

We recall the definition of Morse-Conley-Floer homology, cf. Chapter 2, from a slightly different viewpoint. Let ϕ be a flow on M and S an isolated invariant set of the flow. Recall from Definition 2.2.4 that a Lyapunov function f_ϕ^α for (S, ϕ) is a function that is constant on S and satisfies $\frac{d}{dt} \Big|_{t=0} f_\phi^\alpha(\phi(t, p)) < 0$ on $N^\alpha \setminus S$ where N^α an isolating neighborhood for S^α . Lyapunov functions always exist for a given isolated invariant set, cf. Proposition 2.2.6. We can compute the local Morse homology of a Lyapunov function with respect to the choice of a metric e^α and N^α .

3. FUNCTORIALITY

Given another Lyapunov function f_ϕ^β which satisfies the Lyapunov property on N^β and metric e^β , there is a canonical isomorphism

$$\Phi_*^{\beta\alpha} : HM_*(f_\phi^\alpha, e^\alpha, N^\alpha) \rightarrow HM_*(f_\phi^\beta, e^\beta, N^\beta),$$

induced by continuation, cf. Theorems 2.4.7. The Morse-Conley-Floer homology of (S, ϕ) is then defined as the inverse limit over all such local Morse homologies

$$HI_*(S, \phi) = \varprojlim HM_*(f_\phi^\alpha, e^\alpha, N^\alpha).$$

The Morse-Conley-Floer homology is the local Morse homology of all the Lyapunov functions. This definition is equivalent to

$$HI_*(S, \phi) = \varprojlim HM_*(\mathcal{R}^\alpha),$$

which runs over all Morse-Conley-Floer data $\mathcal{R}^\alpha = (M, f^\alpha, g^\alpha, N^\alpha, \mathfrak{o}^\alpha)$, where $(f^\alpha, g^\alpha, \mathfrak{o}^\alpha)$ is a Morse-Smale triple on N^α whose gradient flow is isolated homotopic to the gradient flow of a Lyapunov function f_ϕ^α on N^α for some metric e^α , with respect to canonical isomorphisms $\Phi_*^{\beta\alpha} : HM_*(\mathcal{R}^\alpha) \rightarrow HM_*(\mathcal{R}^\beta)$, which was the viewpoint of Chapter 2.

3.1.3 Functoriality in Morse homology

In Sections 3.2 through 3.4 we study functoriality in Morse homology on closed manifolds. Induced maps between Morse homologies are defined by counting appropriate intersections for *transverse maps*.

Definition 3.1.1. Let $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ be a smooth map. We say that $h^{\beta\alpha}$ is *transverse* (with respect to Morse data \mathcal{Q}^α and \mathcal{Q}^β), if for all $x \in \text{Crit } f^\alpha$ and $y \in \text{Crit } f^\beta$, we have that

$$h^{\beta\alpha}|_{W^u(x)} \pitchfork W^s(y).$$

The set of transverse maps is denoted by $T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$. We write for the moduli space $W_{h^{\beta\alpha}}(x, y) = W^u(x) \cap (h^{\beta\alpha})^{-1}(W^s(y))$.

Given \mathcal{Q}^α and \mathcal{Q}^β , the set of transverse maps $T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$ is generic, cf. Theorem 3.8.1. Recall that given a critical point x we write $|x|$ for its index. The transversality assumption ensures that the moduli space $W_{h^{\beta\alpha}}(x, y)$ is an oriented manifold of dimension $|x| - |y|$, cf. Proposition 3.2.1. Hence, for $|x| = |y|$

we can compute the oriented intersection number $n_{h^{\beta\alpha}}(x, y)$ and define an induced map $h_*^{\beta\alpha} : C_*(\mathcal{Q}^\alpha) \rightarrow C_*(\mathcal{Q}^\beta)$ via

$$h_*^{\beta\alpha}(x) = \sum_{|y|=|x|} n_{h^{\beta\alpha}}(x, y)y.$$

Through compactness and gluing analysis of the moduli space $W_{h^{\beta\alpha}}(x, y)$ and related moduli spaces we prove the following properties of the induced map

- (i) The induced identity $\text{id}_*^{\alpha\alpha}$ is the identity on chain level.
- (ii) In Section 3.2 we show that $h_*^{\beta\alpha}$ is a chain map, i.e.

$$h_{k-1}^{\beta\alpha} \partial_k^\alpha = \partial_k^\beta h_k^{\beta\alpha}, \quad \text{for all } k.$$

- (iii) In Section 3.3 we study homotopy invariance. Suppose that $\mathcal{Q}^\gamma, \mathcal{Q}^\delta$ are other Morse data on $M^\alpha = M^\gamma$ and $M^\beta = M^\delta$. Suppose $h^{\delta\gamma} \in T(\mathcal{Q}^\gamma, \mathcal{Q}^\delta)$ is homotopic to $h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$, then $h_*^{\delta\gamma} \Phi_*^{\gamma\alpha}$ and $\Phi_*^{\delta\beta} h_*^{\beta\alpha}$ are chain homotopic. That is, there exists a degree +1 map $P_*^{\delta\alpha} : C_*(\mathcal{Q}^\alpha) \rightarrow C_*(\mathcal{Q}^\delta)$, such that

$$\Phi_k^{\delta\beta} h_k^{\beta\alpha} - h_k^{\delta\gamma} \Phi_k^{\gamma\alpha} = -\partial_{k+1}^\delta P_k^{\delta\alpha} - P_{k-1}^{\delta\alpha} \partial_k^\alpha, \quad \text{for all } k.$$

Here $\Phi_*^{\beta\alpha}$ denotes the isomorphism induced by continuation.

- (iv) In Section 3.4 we study compositions. If $h^{\gamma\beta} \in T(\mathcal{Q}^\beta, \mathcal{Q}^\gamma)$ is such that $h^{\gamma\beta} \circ h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\gamma)$ then $h_*^{\gamma\beta} h_*^{\beta\alpha}$ and $(h^{\gamma\beta} \circ h^{\beta\alpha})_*$ are chain homotopic, i.e. there is a degree +1 map $P_*^{\gamma\alpha} : C_*(\mathcal{Q}^\alpha) \rightarrow C_*(\mathcal{Q}^\gamma)$, such that

$$h_k^{\gamma\beta} h_k^{\beta\alpha} - \left(h^{\gamma\beta} \circ h^{\beta\alpha} \right)_k = P_{k-1}^{\gamma\alpha} \partial_k^\alpha + \partial_{k+1}^\gamma P_k^{\gamma\alpha}, \quad \text{for all } k.$$

The properties imply that the induced map $h_*^{\beta\alpha}$ descends to a map

$$h_*^{\beta\alpha} : HM_*(M^\alpha) \rightarrow HM_*(M^\beta)$$

between the Morse homologies via counting, which is functorial $h_*^{\gamma\beta} h_*^{\beta\alpha} = (h^{\gamma\beta} \circ h^{\beta\alpha})_*$ with $\text{id}_*^{\alpha\alpha} = \text{id}$ which does not depend on the homotopy class of the map $h_*^{\beta\alpha}$. The homotopy invariance and density of transverse maps allows for an extension to all smooth maps. Therefore

Theorem 3.1.2. *Morse homology is a functor $HM_* : \mathcal{Man} \rightarrow \mathcal{GrAb}$ between the category of smooth manifolds and maps to graded abelian groups. The functor sends homotopic maps to the same map between the homology groups.*

3.1.4 Isolation properties of maps

To study functoriality in local Morse homology and Morse-Conley-Floer homology isolation properties of maps are crucial. We identify the notion of *isolating map* and *isolating homotopy* in Definition 3.5.1. Isolating maps are open in the compact-open topology, cf. Proposition 3.5.2, but are not necessarily functorial: the composition of two isolating maps need not be isolating. An important class of maps that are isolating and form a category and are flow maps. Flow maps were introduced by McCord [McC88] to study functoriality in Conley index theory.

Definition 3.1.3. A smooth map $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ between manifolds equipped with flows ϕ^α and ϕ^β , is a *flow map* if it is proper and equivariant. Thus preimages of compact sets are compact and

$$h^{\beta\alpha}(\phi^\alpha(t, p)) = \phi^\beta(t, h^{\beta\alpha}(p)), \quad \text{for all } t \in \mathbb{R} \text{ and } p \in M^\alpha.$$

The following examples of flow maps are of interest.

- Let M be a closed manifold with a flow ϕ . The time- t map $\phi_t : M \rightarrow M$, $p \mapsto \phi(t, p)$ is a flow map.
- If $M^\alpha \subset M^\beta$ is a closed submanifold, invariant under the flow, then the inclusion $i^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ is a flow map.
- Let M be a closed manifold with a flow ϕ . Let G be a compact Lie group, acting freely on M . Then M/G is a smooth manifold, and if the action commutes with the flow, M/G comes equipped with a flow. The quotient map $M \rightarrow M/G$ is a flow map.

The isolation properties of flow maps are given in Proposition 3.5.4. If $h^{\beta\alpha}$ is a flow map, and N^β is an isolating neighborhood, then $N^\alpha = (h^{\beta\alpha})^{-1}(N^\beta)$ is an isolating neighborhood. Moreover $h^{\beta\alpha}$ is isolating with respect to these isolating neighborhoods. The same holds for compositions of flow maps.

3.1.5 Functoriality in Local Morse homology

In Section 3.6 we define induced maps in local Morse homology. Due to the local nature of the homology, not all maps are admissible and the notion of isolating map become crucial. The maps are computed by the same counting procedure, but now done locally, as for the functoriality in Morse homology. We sum up the functorial properties of local Morse homology from Propositions 3.6.1 and 3.6.3:

- (i) An isolating transverse map induces a chain map, hence descends to a map between the local Morse homologies.
- (ii) An isolating homotopy between transverse maps induces a chain homotopic map.
- (iii) If $h^{\gamma\beta}$ and $h^{\beta\alpha}$ and $h^{\gamma\beta} \circ h^{\beta\alpha}$ are transverse maps such that $h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}$ is an isolating map for all $R > 0$, then $(h^{\gamma\beta} \circ h^{\beta\alpha})_*$ and $h_*^{\gamma\beta} h_*^{\beta\alpha}$ are chain homotopic.

Consider the following category of isolated invariant sets of gradient flows.

Definition 3.1.4. The category of *isolated invariant sets of gradient flows* \mathcal{IGF} has as objects quadruples (M, f, g, N) of a smooth function f on M a metric g such that N is an isolating neighborhood for the gradient flow. A morphism $h^{\beta\alpha} : (M^\alpha, f^\alpha, g^\alpha, N^\alpha) \rightarrow (M^\beta, f^\beta, g^\beta, N^\beta)$, is a map that is isolated homotopic to a flow map $\tilde{h}^{\beta\alpha}$ with $N^\alpha = (\tilde{h}^{\beta\alpha})^{-1}(N^\beta)$.

These morphisms are then perturbed to transverse maps, from which the induced map can be computed, cf. Propositions 3.6.2 and 3.6.4. We obtain:

Theorem 3.1.5. *Local Morse homology is a covariant functor $HM_* : \mathcal{IGF} \rightarrow \mathcal{GrAb}$.*

Moreover the functor is constant on isolated homotopy classes of maps. The local Morse homology functor generalizes the Morse homology functor. Any map $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ between closed manifolds equipped with flows is isolated homotopic to a flow map with $M^\alpha = N^\alpha = (h^{\beta\alpha})^{-1}(N^\beta) = (\tilde{h}^{\beta\alpha})^{-1}(M^\beta)$, since the flows on both manifolds are isolated homotopic to the constant flow, and any map is equivariant with respect to constant flows.

3.1.6 Functoriality in Morse-Conley-Floer homology

Functoriality in Morse-Conley-Floer homology now follows from the results on local Morse homology, cf. Theorem 3.7.1 by establishing appropriate isolated homotopies. We have the following category of isolated invariant sets

Definition 3.1.6. The *category of isolated invariant sets* $\mathcal{I}I$ has as objects triples (M, ϕ, S) of a manifold, flow and isolated invariant set. A morphism $h^{\beta\alpha} : (M^\alpha, \phi^\alpha, S^\alpha) \rightarrow (M^\beta, \phi^\beta, S^\beta)$ is a map that is isolated homotopic for some choice of isolating neighborhoods N^α, N^β to a flow map $\tilde{h}^{\beta\alpha}$ such that $N^\alpha = (\tilde{h}^{\beta\alpha})^{-1}(N^\beta)$.

The main theorem of this chapter is.

Theorem 3.1.7. *Morse-Conley-Floer homology is a covariant functor $HI_* : \mathcal{I}I \rightarrow \mathcal{G}rAb$.*

Again the functor is constant on the isolated homotopy classes of maps and flows.

3.2 Chain maps in Morse homology on closed manifolds

In Sections 3.2 through 3.4, which discuss functoriality for Morse homology we assume that the base manifolds are closed.

3.2.1 The moduli space $W_{h^{\beta\alpha}}(x, y)$

For a transverse map, see Definition 3.1.1, the moduli spaces $W_{h^{\beta\alpha}}(x, y) = W^u(x) \cap (h^{\beta\alpha})^{-1}(W^s(y))$ are smooth oriented manifolds.

Proposition 3.2.1. *Let $h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$. For all $x \in \text{Crit } f^\alpha$ and $y \in \text{Crit } f^\beta$, the space $W_{h^{\beta\alpha}}(x, y)$ is an oriented submanifold of dimension $|x| - |y|$.*

Proof of Proposition 3.2.1. Because $h^{\beta\alpha}$ restricted to W^u is transverse to W^s Theorem 3.3. of [Hir94, Theorem 3.3, page 22] implies that $W_{h^{\beta\alpha}}(x, y) = \left(h^{\beta\alpha}|_{W^u(x)} \right)^{-1}(W^s(y))$ is an oriented submanifold of $W^u(x)$. The orientation

3.2. Chain maps in Morse homology on closed manifolds

is induced by the exact sequence of vector bundles⁴

$$0 \longrightarrow TW_{h^{\beta\alpha}}(x, y) \longrightarrow TW^u(x) \xrightarrow{dh^{\beta\alpha}} NW^s(y) \longrightarrow 0$$

where the latter is the normal bundle of $W^s(y)$. The normal bundle of $W^s(y)$ is oriented, because $W^s(y)$ is contractible, and $N_y W^s(y) \cong T_y W^u(y)$ is oriented by the choice σ^β . The codimension of $W_{h^{\beta\alpha}}(x, y)$ in $W^u(x)$ equals the codimension of $W^s(y)$ in M^β . Thus

$$|x| - \dim W_{h^{\beta\alpha}}(x, y) = \text{codim } W^s(y) = |y|,$$

from which the proposition follows. □

If $|x| = |y|$ the space $W_{h^{\beta\alpha}}(x, y)$ is zero dimensional and consist of a finite number of points carrying orientation signs ± 1 . Set $n_{h^{\beta\alpha}}(x, y)$ as the sum of these. We define the degree zero map $h_*^{\beta\alpha} : C_*(Q^\alpha) \rightarrow C_*(Q^\beta)$ by the formula

$$h_*^{\beta\alpha}(x) = \sum_{|x|=|y|} n_{h^{\beta\alpha}}(x, y)y.$$

In the following sections we prove properties (ii) . . .(iv) from Section 3.1.3 for this induced map.

3.2.2 Compactness of $W_{h^{\beta\alpha}}(x, y')$ with $|x| = |y'| + 1$

The chain map property (ii) holds by the compactness properties of the moduli space $W_{h^{\beta\alpha}}(x, y')$ with $|x| = |y'| + 1$. This space is a one dimensional manifold, but it is not necessarily compact. The non-compactness is due to breaking of orbits in the domain and in the codomain, which is the content of Lemma 3.2.2. In Figure 3.1 we depicted this breaking process in the domain. Recall that we denote by $W(x, y) = W^u(x) \cap W^s(y)$ the space of parameterized orbits, and by $M(x, y) = W(x, y)/\mathbb{R}$ the space of unparameterized orbits. The points where breaking can occur in the domain are counted by $M(x, y) \times W_{h^{\beta\alpha}}(y, y')$, with $|y| = |y'|$ and the points where breaking occurs in the codomain are counted by $W_{h^{\beta\alpha}}(x, x') \times M(x', y')$, with $|x| = |x'|$. The space $W_{h^{\beta\alpha}}(x, y')$ can be compactified by gluing in these broken orbits cf. Proposition 3.2.4. Proposition 3.2.6 then states that the resulting object is a one dimensional compact




Figure 3.1: One of the possible compactness failures of $W_{h^{\beta\alpha}}(x, y')$ with $|x| = |y'| + 1$. The sequence $p_k \in W_{h^{\beta\alpha}}(x, y')$ converges to a point $p \in W_{h^{\beta\alpha}}(y, y')$ with $|y| = |y'|$. The orbits through p_k break to (u_1, \dots, u_l) with $u_1 \in M(x, y)$.

manifold with boundary. By counting the boundary components appropriately with sign the chain map property is obtained, cf. Proposition 3.2.7.

Since M^α is assumed to be compact, and $W_{h^{\beta\alpha}}(x, y') \subset M^\alpha$, any sequence $p_k \in W_{h^{\beta\alpha}}(x, y')$ has a subsequence converging to some point $p \in M^\alpha$. In the next proposition it is stated to which points such sequences converge, see also Figure 3.1.

Proposition 3.2.2 (Compactness). *Let $h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$, $x \in \text{Crit } f^\alpha$ and $y' \in \text{Crit } f^\beta$ with $|x| = |y'| + 1$. Let $p_k \in W_{h^{\beta\alpha}}(x, y')$ be a sequence such that $p_k \rightarrow p$ in M^α . Then either one of the following is true*

- (i) $p \in W_{h^{\beta\alpha}}(x, y')$.

⁴We employ the fiber first convention in the orientation of exact sequences of vector bundles, cf. [Hir94, KM07]

3.2. Chain maps in Morse homology on closed manifolds

- (ii) *There exists $y \in \text{Crit } f^\alpha$, with $|y| = |y'|$ such that $p \in W_{h^{\beta\alpha}}(y, y')$. The orbits through p_k break to orbits (u_1, \dots, u_l) as $k \rightarrow \infty$ with $u_1 \in M(x, y)$.*
- (iii) *There exists $x' \in \text{Crit } f^\beta$, with $|x'| = |x|$ such that $p \in W_{h^{\beta\alpha}}(x, x')$. The orbits through $h^{\beta\alpha}(p_k)$ break to the orbits (u'_1, \dots, u'_l) as $k \rightarrow \infty$ with $u'_1 \in M(x', y')$.*

Proof. The spaces $W_{h^{\beta\alpha}}(x, y')$, $W_{h^{\beta\alpha}}(y, y')$ and $W_{h^{\beta\alpha}}(x, x')$ are disjoint hence the possibilities cannot occur at the same time. We choose a subsequence such that $p_k \in W(x, b)$ for some fixed $b \in \text{Crit } f^\alpha$, which is possible by the fact that there are only a finite number of such spaces by compactness. The Broken Orbit Lemma, see for example [AB95, Lemma 2.5], states that $p \in W(y, a)$ for some $y, a \in \text{Crit } f^\alpha$ with $|y| \leq |x|$, and equality if and only if $x = y$. Similarly, choosing a subsubsequence if necessary, we also assume that $h^{\beta\alpha}(p_k) \in W(a', y')$, for some fixed $a' \in \text{Crit } f^\beta$. Then $h^{\beta\alpha}(p) \in W(b', x')$, for some $b', x' \in \text{Crit } f^\beta$ with $|x'| \geq |y'|$, with equality if and only if $x' = y'$.

Now if $p \notin W^u(x)$ and $h^{\beta\alpha}(p) \notin W^s(y')$, then we have the impossibility that $p \in W_{h^{\beta\alpha}}(y, x')$, with $|y| < |x|$ and $|x'| > |y'|$. Since by transversality

$$\dim W_{h^{\beta\alpha}}(y, x') = |y| - |x'| \leq (|x| - 1) - (|y'| - 1) \leq -1.$$

Assuming that $p \notin W_{h^{\beta\alpha}}(x, y')$, only two possibilities remain. If $p \in W_{h^{\beta\alpha}}(x, x')$, with $|x'| > |y| = |x| - 1$, then from $\dim W_{h^{\beta\alpha}}(x, x') \geq 0$ it follows that $|x'| = |x|$. If $h^{\beta\alpha}(p) \in W_{h^{\beta\alpha}}(y, y')$, with $|y| < |x|$, then $\dim W_{h^{\beta\alpha}}(y, y') \geq 0$ gives that $|y'| = |y|$, see also Figure 3.1. The claim about the breaking orbits is the content of the Broken Orbit Lemma. \square

The proposition generalizes to higher index difference moduli spaces. We do not need this, and this would clutter the notation without a significant gain.

3.2.3 Gluing the ends of $W_{h^{\beta\alpha}}(x, y')$, with $|x| = |y'| + 1$

We compactify $W_{h^{\beta\alpha}}$ by gluing in the broken orbits described in Proposition 3.2.2. The following lemma is the technical heart of the standard gluing construction in Morse homology, see also Figure 3.2. We single this out because we have to construct several gluing maps, which all use this lemma. For a pair (f, g) of a function and a metric we denote by ψ its negative gradient flow.




Figure 3.2: The content of Lemma 3.2.3 is depicted. Discs transverse to the stable and unstable manifolds must intersect if they are flowed in forwards and backwards time. This intersection point is the gluing map.

Lemma 3.2.3 (Gluing). *Let $f : M \rightarrow \mathbb{R}$ be a Morse function, g a metric, and $y \in \text{Crit } f$. Write $m = \dim M$. Suppose $D^{|y|}, E^{m-|y|}$ are embedded discs of dimension $|y|$ and $m - |y|$ with $D^{|y|} \pitchfork W^s(y)$ and $E^{m-|y|} \pitchfork W^u(y)$. Assume that each intersection only consists of a single point and write $u \in D^{|y|} \cap W^s(y)$ and $v \in E^{m-|y|} \cap W^u(y)$. Then there exists an $R_0 > 0$ and an injective map $\rho : [R_0, \infty) \rightarrow M$, such that $g(\frac{d}{dR}\rho(R), -\nabla_g f) \neq 0$, and $\psi(-R, \rho(R)) \in D^{|y|}$, and $\psi(R, \rho(R)) \in E^{m-|y|}$. We have the limits*

$$\lim_{R \rightarrow \infty} \rho(R) = y, \quad \lim_{R \rightarrow \infty} \psi(-R, \rho(R)) = u, \quad \lim_{\rho \rightarrow \infty} \psi(R, \rho(R)) = v.$$

Finally there exists smaller discs $D' \subset D^{|y|}$ and $E' \subset E^{m-|y|}$ such that no orbit through $D' \setminus \bigcup_{R \in [R_0, \infty)} \psi(-R, \rho(R))$ intersects E' .

Proof. We only sketch the proof, see also Figure 3.2. More details can be found in the proof of Theorem 3.9 in [Web06]. Let $B^u \subset W^u(y)$, $B^s \subset W^s(y)$ be closed balls containing y . Write $D_R^{|y|} = \psi(R, D^{|y|})$ and $E_{-R}^{m-|y|} = \psi(-R, E^{m-|y|})$. Since the discs are transverse to the stable and unstable manifolds of y , the λ -Lemma, cf. Lemma 7.2 of Chapter 2 in [PdM82], gives a that for all t large,

3.2. Chain maps in Morse homology on closed manifolds

smaller discs $D'_R \subset D_R^{|y|}$ and $E'_{-R} \subset E_{-R}^{m-|y|}$ are $\epsilon - C^1$ close to B^u and B^s respectively. It follows through an application of the Banach fixed point theorem that there exist a $R_0 > 0$ such that D'_R and E'_{-R} intersect in a single point for each $R > R_0$ sufficiently large. Set $\rho(R) = D'_R \cap E'_{-R}$. The properties of ρ follow from the construction. \square

To prove that $\partial^2 = 0$ in standard Morse homology gluing maps are needed to compactify appropriate moduli spaces and these are constructed using the Morse-Smale condition and previous lemma as follows. Let x, y, z be critical points with $|x| = |y| - 1 = |z| - 2$ and assume $M(x, y)$ and $M(y, z)$ are non-empty. Then the Morse Smale condition gives that there exists a disc $D^{|y|}$ in $W^u(x)$ transverse to an orbit in $M(x, y)$. Because of transversality, one can also choose a disc $E^{m-|y|}$ in $W^s(z)$ transverse to an orbit in $M(y, z)$. The gluing map $\# : M(x, y) \times [R_0, \infty) \times M(y, z) \rightarrow M(x, z)$, is given by mapping R to the orbit through $\rho(R)$ as in Figure 3.2. We now use similar ideas to compactify $W_{h^{\beta\alpha}}(x, y')$ with $|x| = |y'| + 1$.

Proposition 3.2.4. *Assume $h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$. Then for critical points $x, y \in \text{Crit } f^\alpha$ and $x', y' \in \text{Crit } f^\beta$, with $|x| = |x'| = |y| + 1 = |y'| + 1$, there exists $R_0 > 0$ and gluing embeddings*

$$\begin{aligned} \#^1 : M(x, y) \times [R_0, \infty) \times W_{h^{\beta\alpha}}(y, y') &\rightarrow W_{h^{\beta\alpha}}(x, y') \\ \#^2 : W_{h^{\beta\alpha}}(x, x') \times [R_0, \infty) \times M(x', y') &\rightarrow W_{h^{\beta\alpha}}(x, y') \end{aligned}$$

Moreover, if $p_k \in W_{h^{\beta\alpha}}(x, y')$ converges to $p \in W_{h^{\beta\alpha}}(y, y')$, and the orbits through p_k break as (u_1, \dots, u_l) with $u_1 \in M(x, y)$ then p_k is in the image of $\#^1$ for k sufficiently large. Analogously, if $p_k \in W_{h^{\beta\alpha}}(x, y')$ converges to $p \in W_{h^{\beta\alpha}}(x, x')$, and the orbits through $h^{\beta\alpha}(p_k)$ break to (u_1, \dots, u_l) with $u'_1 \in M(x', y)$, then p_k is in the image of $\#^2$ for k sufficiently large.

Proof. Let $u \in W(x, y)$, and $v \in W_{h^{\beta\alpha}}(y, y')$. Since (f^α, g^α) is a Morse-Smale pair, we can choose a disc $D^{|y|}$ in $W^u(x)$ through u and transverse to $W^s(y)$. Analogously, because $h^{\beta\alpha}$ is transverse, we can find a disc $E^{m^\alpha-|y|}$ in M^α through v , and transverse to $W^u(y)$, with $h^{\beta\alpha}(E^{m^\alpha-|y|}) \subset W^s(y')$. Lemma 3.2.3 provides us with a map $\rho : [R_0, \infty) \rightarrow M^\alpha$. Denote by $\gamma_u \in M(x, y)$ the orbit through u , and set $\gamma_u \#_R^1 v = \psi^\alpha(R, \rho(R))$. Then $\gamma_u \#_R^1 v \in E^{m^\alpha-|y|}$ and $\psi^\alpha(-2R, \gamma_u \#_R^1 v) \in D^{|y|} \subset W^u(x)$, hence $\gamma_u \#_R^1 v \in W_{h^{\beta\alpha}}(x, y')$. The properties of the map ρ in Lemma 3.2.3 directly give the properties of $\#^1$.

3. FUNCTORIALITY

The construction of $\#^2$ is similar. Let $u \in W_{h^{\beta\alpha}}(x, x')$ and $v \in W(x', y')$. By transversality of $h^{\alpha\beta}$ we can choose a disc $D^{|x'|} \subset W^u(x)$ which $h^{\beta\alpha}$ maps bijectively to $h^{\beta\alpha}(D^{|x'|})$, and whose image $h^{\beta\alpha}(D)$ is transverse to $W^s(x')$. Choose a disc $E^{m^{\beta}-|x'|} \subset W^s(y')$ with $v \in E^{m^{\beta}-|x'|}$ which is transverse to $W^u(x')$. Now Lemma 3.2.3 provides us with a map $\rho : [R_0, \infty) \rightarrow M^\beta$. Set $u\#_R^2\gamma_v = (h^{\beta\alpha})^{-1}(\psi^\beta(-R, \rho(R)))$, which is well defined since $h^{\beta\alpha}|_{D^{|x'|}}$ is a diffeomorphism onto $h^{\beta\alpha}(D^{|x'|})$. The properties of $\#^2$ follow. \square

3.2.4 Orientations.

After a choice of orientations of the unstable manifolds, the moduli spaces $W_{h^{\beta\alpha}}(x, y')$ and $W(x, y)$ and $M(x, y)$ carry induced orientations cf. Proposition 3.2.1 and [Web06, Proposition 3.10]. We show that the gluing map $\#^1$ is compatible with the induced orientations, while $\#^2$ reverses the orientations.

Consider the notation of the proof of Proposition 3.2.4. Denote by $W(x, y)|_u$ the connected component of $W(x, y)$ containing $u \in W(x, y)$, with similar notation for other moduli spaces. If $|x| = |y| + 1$ the moduli space $W(x, y)$ is one dimensional and can be oriented by the negative gradient vector field. For $u \in W(x, y)$ write $[\dot{u}]$ for this induced orientation of $W(x, y)|_u$. The orientation of $W(x, y)|_u$ induced by the choice of σ^α is denoted $a[\dot{u}]$. Then $W_{h^{\beta\alpha}}(y, y')$, whose orientation induced by σ^α and σ^β by b , for $a, b \in \{\pm 1\}$. The gluing map $\#^1$ induces a map of orientations

$$\sigma^1 : \text{Or}(W(x, y)|_u) \times \text{Or}(W_{h^{\beta\alpha}}(y, y')|_v) \rightarrow \text{Or}(W_{h^{\beta\alpha}}(x, y')|_{\gamma_u\#_{R_0}^1 v}),$$

via

$$\sigma^1(a[\dot{u}], b) = ab \left[\frac{d}{dR} \Big|_{R=R_0} \gamma_u\#_R^1 v \right].$$

Similarly the gluing map $\#^2$ induces a map of orientations

$$\sigma^2 : \text{Or}(W_{h^{\beta\alpha}}(x, x')|_u) \times \text{Or}(W(x', y')|_v) \rightarrow \text{Or}(W_{h^{\beta\alpha}}(x, y')|_{u\#_{R_0}^2 \gamma_v})$$

given by

$$\sigma^2(a, b[\dot{v}]) = ab \left[\frac{d}{dR} \Big|_{R=R_0} u\#_R^2 \gamma_v \right].$$

3.2. Chain maps in Morse homology on closed manifolds

Analogously, from (3.4), (3.5) and (3.6) we get

$$\begin{aligned} T_u W^u(x) &\cong T_{\gamma_u \#_{R_0}^1 v} W^u(x) \cong T_{\gamma_u \#_{R_0}^1 v} W_{h^{\beta\alpha}}(x, y') \oplus N_{h^{\beta\alpha}(\gamma_u \#_{R_0}^1 v)} W^s(y') \\ &\cong T_{\gamma_u \#_{R_0}^1 v} W_{h^{\beta\alpha}}(x, y') \oplus T_{y'} W^u(y'). \end{aligned}$$

Combining the last two formulas we see that

$$T_{\gamma_u \#_{R_0}^1 v} W_{h^{\beta\alpha}}(x, y') \cong T_u W(x, y) \oplus T_v W_{h^{\beta\alpha}}(y, y').$$

The orientation on the right hand side is $ab[\dot{u}]$. The tangent vector to the orbit γ_u through u has a well defined limit as $t \rightarrow \infty$ which we denote by $[\dot{u}(\infty)]$, and similarly the tangent vector to γ_v has a well defined limit as $t \rightarrow -\infty$, which we denote by $[\dot{v}(-\infty)]$, cf. [Web06, Theorem 3.11] and [Sch93, Lemma B.5]. We want to compare the orientation $[\dot{u}]$ to $\left[\frac{d}{dR} \Big|_{R=R_0} \gamma_u \#_{R_0}^1 v \right]$, and are able to do this as follows. We can flow $W_{h^{\beta\alpha}}(x, y') \Big|_{\gamma_u \#_{R_0}^1 v}$, cf. Figure 3.3, which we denote by $W_{h^{\beta\alpha}}(x, y') \Big|_{\gamma_u \#_{R_0}^1 v} \times \mathbb{R}$. Note that y is in the closure of this space.

Then the orientations $[\dot{u}(\infty), \dot{v}(-\infty)]$ and $\left[\frac{d}{dR} \Big|_{R=R_0} \gamma_u \#_{R_0}^1 v, -\nabla f^\alpha \right]$ agree on this space (here we extend the manifold to its closure). Quotienting out the \mathbb{R} action, we see that the orientation $[\dot{u}] = [\dot{u}(\infty)]$ agrees with $\left[\frac{d}{dR} \rho(R) \right]$ which agrees with $\left[\frac{d}{dR} \Big|_{R=R_0} \gamma_u \#_{R_0}^1 v \right]$. The orientation map σ^1 preserves the orientation.

The proof that σ^2 is orientation reversing is analogous. Again we use the notation of Proposition 3.2.4, with $u \in W_{h^{\beta\alpha}}(x, x')$ and $v \in W(x', y')$. We have the exact sequences of oriented vector spaces

$$0 \rightarrow T_v W(x', y') \rightarrow T_v W^u(x') \rightarrow N_v W^s(y') \rightarrow 0, \quad (3.8)$$

$$0 \rightarrow T_u W_{h^{\beta\alpha}}(x, x') \rightarrow T_u W^u(x) \xrightarrow{dh^{\beta\alpha}} N_{h^{\beta\alpha}(u)} W^s(x') \rightarrow 0, \quad (3.9)$$

$$0 \rightarrow T_{u \#_{R_0}^2 \gamma_v} W_{h^{\beta\alpha}}(x, y') \rightarrow T_{u \#_{R_0}^2 \gamma_v} W^u(x) \xrightarrow{dh^{\beta\alpha}} N_{h^{\beta\alpha}(u \#_{R_0}^2 \gamma_v)} W^s(y') \rightarrow 0. \quad (3.10)$$

3. FUNCTORIALITY

By isomorphisms induced by parallel transport, analogous to the isomorphisms (3.5), (3.6), and (3.7), and from (3.9) and (3.8) we get that

$$\begin{aligned} T_u W^u(x) &\cong T_u W_{h^{\beta\alpha}}(x, x') \oplus N_{h^{\beta\alpha}(u)}(W^s(x')) \\ &\cong T_u W_{h^{\beta\alpha}}(x, x') \oplus T_v W(x', y') \oplus T_{y'} W^u(y'). \end{aligned}$$

Similarly, from (3.10) and the isomorphisms induced by parallel transport we get

$$T_u W^u(x) \cong T_{u\#2\gamma_v} W_{h^{\beta\alpha}}(x, y') \oplus T_{y'} W^u(y'),$$

which gives that

$$T_{u\#2\gamma_v} W_{h^{\beta\alpha}}(x, y') \cong T_u W_{h^{\beta\alpha}}(x, x') \oplus T_{h^{\beta\alpha}(u)} W(x', y')$$

The orientation on the right hand side is $ab[\dot{v}]$. Locally around u , $h^{\beta\alpha}$ is injective when restricted to $W_{h^{\beta\alpha}}(x, y')|_u$. We can therefore take this image and flow with ψ^β , which we denote by $h^{\beta\alpha}(W_{h^{\beta\alpha}}(x, y')|_u) \times \mathbb{R}$. The point x' lies in the closure, again the limits of tangent vectors of the orbits through $h^{\beta\alpha}(u)$ and v are well defined for $t \rightarrow \pm\infty$, and the orientations $[h^{\beta\alpha}(u)(\infty), v(-\infty)]$ and

$$\left[\frac{d}{dR} \rho(R), -\nabla f^\beta \right] = - \left[-\nabla f^\beta, \frac{d}{dR} \rho(R) \right]$$

agree. By quotienting out the flow it follows that $[\dot{v}] = [\dot{v}(-\infty)] = - \left[\frac{d}{dR} \rho(R) \right] = - \left[\frac{d}{dR} \Big|_{R=R_0} u\#2\gamma_v \right]$. The map σ^2 is orientation reversing, cf. Figure 3.4. \square

3.2.5 The induced map $h_*^{\beta\alpha}$ is a chain map.

Propositions 3.2.4 and 3.2.2, along with the considerations of the previous section directly give the following theorem.

Theorem 3.2.6. *Let $h^{\beta\alpha}$ be transverse with respect to \mathcal{Q}^α and \mathcal{Q}^β . For each $x \in \text{Crit } f^\alpha$ and $y' \in \text{Crit } f^\beta$ with $|x| = |y'| + 1$, the space*

$$\begin{aligned} \widehat{W}_{h^{\beta\alpha}}(x, y') &= W_{h^{\beta\alpha}}(x, y') \bigcup_{|y|=|y'|} M(x, y) \times W_{h^{\beta\alpha}}(y, y') \\ &\quad \bigcup_{|x'|=|x|} - (W_{h^{\beta\alpha}}(x, x') \times M(x', y')), \end{aligned}$$

3.3. Homotopy induced chain homotopies

has a natural structure as a compact oriented manifold with boundary given by the gluing maps.

Note that this proposition does not state that exactly one half of the boundary components correspond to $M(x, y) \times W_{h^{\beta\alpha}}(y, y')$ and the other half of the boundary components to $W_{h^{\beta\alpha}}(x, x') \times M(x', y')$. Still it does follow that $h_*^{\beta\alpha}$ is a chain map.

Proposition 3.2.7. *Let $h^{\beta\alpha}$ be a transverse map with respect to \mathcal{Q}^α and \mathcal{Q}^β . Then the induced map $h_*^{\beta\alpha} : C_*(\mathcal{Q}^\alpha) \rightarrow C_*(\mathcal{Q}^\beta)$ is a chain map.*

Proof. Let $x \in \text{Crit } f^\alpha$. We compute using Proposition 3.2.6

$$\begin{aligned} & (h_{k-1}^{\beta\alpha} \partial_k^\alpha - \partial_k^\beta h_k^{\beta\alpha})(x) \\ &= \sum_{|y'|=|x|-1} \left(\sum_{|y|=|x|-1} n_{h^{\beta\alpha}}(y, y') n(x, y) - \sum_{|x'|=|x|} n(x', y') n_{h^{\beta\alpha}}(x, x') \right) y' \\ &= \sum_{|y'|=|x|-1} \partial \widehat{W}_{h^{\beta\alpha}}(x, y') y' = 0. \end{aligned}$$

Because the oriented count of the boundary components of a compact oriented one dimensional manifold is zero. Hence $h_*^{\beta\alpha}$ is a chain map. \square

3.3 Homotopy induced chain homotopies

The main technical work, showing that $h_*^{\beta\alpha}$ is a chain map, is done. To show that homotopic maps induce the same maps in Morse homology, we could again define an appropriate moduli space and analyze its compactness failures. However, a simpler method is to construct a dynamical model of the homological cone. This trick is used in Morse homology to show that Morse homology does not depend on the choice of function, metric and orientation. Using the homotopy we build a higher dimensional system – the dynamical cone – where we use the fact that an induced map is a chain map to prove homotopy invariance, that is Property (iii) from Section 3.1.3.

Proposition 3.3.1 (Homotopy invariance). *Let $\mathcal{Q}^\alpha, \mathcal{Q}^\beta, \mathcal{Q}^\gamma, \mathcal{Q}^\delta$, be Morse data with $M^\alpha = M^\gamma$, and $M^\beta = M^\delta$. Let $h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$ and $h^{\delta\gamma} \in T(\mathcal{Q}^\gamma, \mathcal{Q}^\delta)$,*

3. FUNCTORIALITY

and assume the maps are homotopic. Then $h_*^{\delta\gamma}\Phi_*^{\gamma\alpha}$ and $\Phi_*^{\delta\beta}h_*^{\beta\alpha}$ are chain homotopic. That is, there exists a degree +1 map $P_*^{\delta\alpha} : C_*(\mathcal{Q}^\alpha) \rightarrow C_*(\mathcal{Q}^\delta)$, such that

$$\Phi_k^{\delta\beta}h_k^{\beta\alpha} - h_k^{\delta\gamma}\Phi_k^{\gamma\alpha} = -\partial_{k+1}^\delta P_k^{\delta\alpha} - P_{k-1}^{\delta\alpha}\partial_k^\alpha, \quad \text{for all } k.$$

Proof. Let $h_\lambda : M^\alpha \rightarrow M^\beta$ be a smooth homotopy between $h^{\beta\alpha}$ and $h^{\delta\gamma}$. Let $g_\lambda^{\gamma\alpha}$ be a smooth homotopy between g^α and g^γ , and $g_\lambda^{\delta\beta}$ a smooth homotopy between g^β and g^δ . Similarly let $f_\lambda^{\gamma\alpha}$ be a smooth homotopy between f^α and f^γ , and let $f_\lambda^{\delta\beta}$ be a smooth homotopy between f^β and f^δ . Choose $0 < \epsilon < \frac{1}{4}$ and let $\omega : \mathbb{R} \rightarrow [0, 1]$ be a smooth, even and 2-periodic function with the following properties:

$$\omega(\mu) = \begin{cases} 0 & -\epsilon < \mu < \epsilon \\ 1 & -1 \leq \mu < -1 + \epsilon \quad \text{and} \quad 1 - \epsilon < \mu \leq 1. \end{cases}$$

and $\omega'(\mu) < 0$ for $\mu \in (-1 + \epsilon, -\epsilon)$, and $\omega'(\mu) > 0$ for $\mu \in (\epsilon, 1 - \epsilon)$. We identify \mathbb{S}^1 with $\mathbb{R}/2\mathbb{Z}$. Under the identification of \mathbb{S}^1 with $\mathbb{R}/2\mathbb{Z}$, the function ω descends to a smooth function $\mathbb{S}^1 \rightarrow [0, 1]$, which is also denoted ω . Let $r > 0$. We define the functions F^α on $M^\alpha \times \mathbb{S}^1$ and F^β on $M^\beta \times \mathbb{S}^1$ via

$$\begin{aligned} F^\alpha(x, \mu) &= f_{\omega(\mu)}^{\gamma\alpha}(x) + r(1 + \cos(\pi\mu)) \\ F^\beta(x, \mu) &= f_{\omega(\mu)}^{\delta\beta}(x) + r(1 + \cos(\pi\mu)) \end{aligned}$$

For r sufficiently large the functions F^α and F^β are Morse, cf. Lemma 2.4.5, and the critical points can be identified by

$$\begin{aligned} C_k(F^\alpha) &\cong C_{k-1}(f^\alpha) \oplus C_k(f^\gamma), \quad \text{and} \\ C_k(F^\beta) &\cong C_{k-1}(f^\beta) \oplus C_k(f^\delta). \end{aligned} \tag{3.11}$$

We define $H : M^\alpha \times \mathbb{S}^1 \rightarrow M^\beta \times \mathbb{S}^1$ with $H(x, \mu) = (h_{\omega(\mu)}(x), \mu)$. The connections of the gradient flow at $\mu = 0$ and $\mu = 1$ are transverse, and the map H restricted to neighborhood at $\mu = 0$ and $\mu = 1$ also satisfies the required transversality properties. Hence we can perturb H while keeping it fixed in the neighborhoods of $\mu = 0$ and $\mu = 1$, as well as the metrics $G^\alpha = g_{\omega(\mu)}^{\gamma\alpha}(x) \oplus d\mu^2$, and $G^\beta = g_{\omega(\mu)}^{\delta\beta}(x) \oplus d\mu^2$ outside $-\epsilon < \mu < \epsilon$ and

3.3. Homotopy induced chain homotopies

$1 - \epsilon < \mu < 1 + \epsilon$ to obtain Morse-Smale flows on $M^\alpha \times \mathbb{S}^1$ and $M^\beta \times \mathbb{S}^1$, such that the map H is transverse everywhere. We orient the unstable manifolds in $M^\alpha \times \mathbb{S}^1$ by $\mathcal{O}^\alpha = (\partial\mu \oplus \mathfrak{o}^\alpha) \cup \mathfrak{o}^\gamma$, and the unstable manifolds in $M^\beta \times \mathbb{S}^1$ by $\mathcal{O}^\beta = (\partial\mu \oplus \mathfrak{o}^\beta) \cup \mathfrak{o}^\delta$.

Let $(x, 0) \in C_k(F^\alpha)$, thus $x \in C_{k-1}(f^\alpha)$. Then $W^u((x, 0)) \subset M^\alpha \times \mathbb{S}^1 \setminus \{1\}$ and $W^u((x, 0)) \cap M^\alpha \times \{0\} = W^u(x) \times \{0\}$. For $(x, 1) \in C_k(F^\alpha)$, i.e. $x \in C_k(f^\gamma)$ we have $W^u((x, 1)) = W^u(x) \times \{1\}$. Similarly for $(y, 0) \in C_k(F^\beta)$, i.e. $y \in C_{k-1}(f^\beta)$ we have $W^s((y, 0)) = W^s(y) \times \{0\}$. Finally for $(y, 1) \in C_k(F^\beta)$, i.e. $y \in C_k(f^\delta)$, we have that $W^s((y, 1)) \subset M^\beta \times (\mathbb{S}^1 \setminus \{0\})$, and $W^s((y, 1)) \cap (M^\beta \times \{1\}) = W^s(y) \times \{1\}$.

By Propositions 3.2.2 and 3.2.4 and the construction of H it is clear⁵ that we can restrict the count of the induced map H_* to $W_H(x, y) \cap M^\alpha \times [0, 1]$, see also Proposition 3.6.1. Note that $H^{-1}(M^\beta \times (\mathbb{S}^1 \setminus \{0\})) \subset M^\alpha \times (\mathbb{S}^1 \setminus \{0\})$. Let $|(x, 0)| = |(y, 0)|$, then

$$W_H((x, 0), (y, 0)) = (W^u(x) \cap (h^{\beta\alpha})^{-1}(W^s(y))) \times \{0\},$$

so $n_H((x, 0), (y, 0)) = n_{h^{\beta\alpha}}(x, y)$ as oriented intersection numbers. Similarly for $|(x, 1)| = |(y, 1)|$, we find that

$$W_H((x, 1), (y, 1)) = (W^u(x) \cap (h^{\delta\gamma})^{-1}(W^s(y))) \times \{1\},$$

which gives $n_H((x, 1), (y, 1)) = n_{h^{\delta\gamma}}(x, y)$. For $|(x, 1)| = |(y, 0)|$ we compute that $n_H((x, 1), (y, 0)) = 0$. Finally we define a map $P_k^{\delta\alpha} : C_k(f^\alpha) \rightarrow C_{k+1}(f^\delta)$ by counting the intersections of $W^u((x, 0))$ and $H^{-1}(W^s((y, 1)))$ with sign. Thus

$$P^{\delta\alpha}(x) = \sum_{|(x,0)|=|(y,1)|} n_H((x,0), (y,1))y.$$

With respect to the splitting (3.11) the induced map H_* equals

$$H_k = \begin{pmatrix} h_{k-1}^{\beta\alpha} & 0 \\ P_{k-1}^{\delta\alpha} & h_k^{\delta\gamma} \end{pmatrix}.$$

By Equation (2.11) the boundary maps Δ_k^α on $M^\alpha \times [0, 1]$ and $\Delta_k^\beta \times [0, 1]$ on M^β have the following form

$$\Delta_k^\alpha = \begin{pmatrix} -\partial_{k-1}^\alpha & 0 \\ \Phi_{k-1}^{\gamma\alpha} & \partial_k^\gamma \end{pmatrix}, \quad \Delta_k^\beta = \begin{pmatrix} -\partial_{k-1}^\beta & 0 \\ \Phi_{k-1}^{\delta\beta} & \partial_k^\delta \end{pmatrix}.$$

⁵The map is isolating cf. Definition 3.5.1

3. FUNCTORIALITY

Where the Φ 's are the maps that induce isomorphisms in Morse homology. We know that H_k is a chain map, i.e. $H_{k-1}\Delta_k^\alpha = \Delta_k^\delta H_k$. This implies

$$\begin{pmatrix} -h_{k-2}^{\beta\alpha} \partial_{k-1}^\alpha & 0 \\ -P_{k-2}^{\delta\alpha} \partial_{k-1}^\alpha + h_{k-1}^{\delta\gamma} \Phi_{k-1}^{\gamma\alpha} & h_{k-1}^{\delta\gamma} \partial_k^\gamma \end{pmatrix} = \begin{pmatrix} -\partial_{k-1}^\beta h_{k-1}^{\beta\alpha} & 0 \\ \Phi_{k-1}^{\delta\beta} h_{k-1}^{\beta\alpha} + \partial_k^\delta P_{k-1}^{\delta\alpha} & \partial_k^\delta h_k^{\delta\gamma} \end{pmatrix}.$$

The lower left corner of this matrix equation gives the desired identity. \square

3.4 Composition induced chain homotopies

To show that compositions of map induce the same map as the compositions of the induced maps in Morse homology, we take in spirit we also take a homotopy between $h^{\gamma\beta}$ and $h^{\beta\alpha}$ and $h^{\gamma\beta} \circ h^{\beta\alpha}$. This is not directly possible as the maps have different domains and codomains, but we have an approximating homotopy, which is sufficient.

For a flow $\phi : \mathbb{R} \times M \rightarrow M$ on a manifold M we write $\phi_R : M \rightarrow M$ for the time R map $\phi_R(x) = \phi(R, x)$. Let $h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$ and $h^{\gamma\beta} \in T(\mathcal{Q}^\beta, \mathcal{Q}^\gamma)$. We assume, up to possibly a perturbation of $h^{\gamma\beta}$, that the map $H : M^\alpha \times (0, \infty) \rightarrow M^\beta$ defined by

$$H(p, R) = h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}(p)$$

is transverse in the sense that

$$H|_{W^u(x) \times (0, \infty)} \pitchfork W^s(z),$$

for all $x \in \text{Crit } f^\alpha$ and $z \in \text{Crit } f^\gamma$. We then have moduli spaces

$$W_{h^{\gamma\beta}, h^{\beta\alpha}}(x, z) = \{(p, R) \in M^\alpha \times (0, \infty) \mid p \in W^u(x), \text{ and } h^{\gamma\beta} \circ \psi_R \circ h^{\beta\alpha}(p) \in W^s(z)\},$$

of dimension $|x| - |z| + 1$. The compactness issues if $R \rightarrow \infty$ are due to the breaking of orbits in M^β , which is described in the following proposition.

Proposition 3.4.1 (Compactness). *Assume the situation as above, with $|x| = |z|$. Let $(p_k, R_k) \in W_{h^{\gamma\beta}, h^{\beta\alpha}}(x, z)$ be a sequence with $R_k \rightarrow \infty$ as $k \rightarrow \infty$. Then there exists $y \in \text{Crit } f^\beta$ with $|x| = |y| = |z|$, and a subsequence (p_k, R_k) such that $p_k \rightarrow p \in W_{h^{\beta\alpha}}(x, y)$, and $\psi_{R_k}^\beta \circ h^{\beta\alpha}(p_k) \rightarrow q \in W_{h^{\gamma\beta}}(y, z)$.*




Figure 3.5: The moduli space $W_{h^{\gamma\beta}, h^{\beta\alpha}}(x, z)$ has non-compact ends if $R \rightarrow \infty$. These ends can be compactified by gluing in $W_{h^{\beta\alpha}}(x, y) \times W_{h^{\gamma\beta}}(y, z)$, for all $y \in \text{Crit } f^{\beta}$, with $|x| = |y| = |z|$ as in Proposition 3.4.2.

Proof. By compactness of M^{α} we can choose a subsequence of (p_k, R_k) such that $p_k \rightarrow p$, and, if necessary choose a subsubsequence, such that also $q_k = \psi_{R_k}^{\beta} \circ h^{\beta\alpha}(p_k) \rightarrow q$. By similar arguments as in Proposition 3.2.2, using the transversality, $p \in W_{h^{\beta\alpha}}(x', y)$, with $|x| \geq |x'| \geq |y|$ and $q \in W_{h^{\gamma\beta}}(y', z')$ with $|y'| \geq |z'| \geq |z|$ where the equality holds if and only if $x = x', z' = z$. Moreover since q_k and $h^{\beta\alpha}(p_k)$ are on the same orbit, we must have that $|y'| \geq |y|$ with equality if and only if $y' = y$. Since $|x| = |z|$, it follows that $x' = x$ and $z' = z$, and therefore also $y' = y$. \square

Proposition 3.4.2 (Gluing). *Assume the situation above. Let $x \in \text{Crit } f^{\alpha}$, and $y \in \text{Crit } f^{\beta}$ and $z \in \text{Crit } f^{\gamma}$, with $|x| = |y| = |z|$. Then there exists an $R_0 > 0$, and a gluing embedding*

$$\#^3 : W_{h^{\beta\alpha}}(x, y) \times [R_0, \infty) \times W_{h^{\gamma\beta}}(y, z) \rightarrow W_{h^{\gamma\beta}, h^{\beta\alpha}}(x, y).$$

3. FUNCTORIALITY

Moreover, if $(p_k, R_k) \in W_{h^{\gamma\beta}, h^{\beta\alpha}}(R_0, x, z)$ with $p_k \rightarrow p \in W_{h^{\beta\alpha}}(x, y)$ and $\psi^\beta(R_k, h^{\beta\alpha}(p_k)) \rightarrow q \in W_{h^{\gamma\beta}}(y, z)$, then the sequence lies in the image of the embedding for k sufficiently large.

Proof. See Figure 3.5. Let $u \in W_{h^{\beta\alpha}}(x, y)$, and $v \in W_{h^{\gamma\beta}}(y, z)$. By transversality of $h^{\beta\alpha}$, we choose a disc $D^{|y|} \subset W^u(x)$ such that $h^{\beta\alpha}|_{D^{|y|}}$ is injective and the image $h^{\beta\alpha}(D^{|y|})$ intersects $W^s(y)$ transversely in $h^{\beta\alpha}(u)$. We also choose a disc $E^{m^\beta - |y|} \subset M^\beta$, intersecting $W^u(y)$ transversely in v , whose image is contained in $W^s(z)$, cf. Figure 3.5. By Lemma 3.2.3, we get an $R'_0 > 0$ and a map $\rho : [R'_0, \infty) \rightarrow M^\beta$. Set $R_0 = R'_0/2$ and define $\#^3(u, R, v) = ((h^{\beta\alpha})^{-1}(\psi(-R, \rho(R))), 2R)$. Here we use the fact that $h^{\beta\alpha}|_{D^{|y|}}$ is bijective to $h^{\beta\alpha}(D^{|y|})$. The properties stated follow from Lemma 3.2.3. \square

Proposition 3.4.3. *Consider the situation above. Then there exist an $R > 0$ such that $h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$ and the moduli space*

$$W_{h^{\gamma\beta}, h^{\beta\alpha}}(x, z, R) := W_{h^{\gamma\beta}, h^{\beta\alpha}}(x, z) \cap M^\alpha \times (R, \infty),$$

has a compactification as a smooth oriented manifold with boundary

$$\begin{aligned} \widehat{W}_{h^{\gamma\beta}, h^{\beta\alpha}}(x, z, R) &= W_{h^{\gamma\beta}, h^{\beta\alpha}}(x, z, R) \cup -W_{h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}}(x, z) \\ &\quad \cup_{|y|=|x|} W_{h^{\beta\alpha}}(x, y) \times W_{h^{\gamma\beta}}(y, z). \end{aligned}$$

Proof. Proposition 3.4.1 shows that all sequences (p_k, R_k) in the moduli space with $R_k \rightarrow \infty$ have that the limits $p_k \rightarrow p$ and $h^{\gamma\beta} \circ \psi_{R_k}^\beta \circ h^{\beta\alpha}(p_k) \rightarrow q$ converge in the interior of $W^u(x)$ and $W^s(z)$. By compactness of the domain this implies that there is an $R > 0$ such that $W_{h^{\gamma\beta}, h^{\beta\alpha}}(x, z, R)$ has no limit points outside $W^u(x)$, and similarly that $h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}(p_k)$ cannot converge to a point outside $W^s(z)$. By parametric transversality [Hir94, Theorem 2.7 page 79] we can assume that $h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$. The space only has two non-compact ends. One is counted by $\cup_{|y|=|x|} W_{h^{\beta\alpha}}(x, y) \times W_{h^{\gamma\beta}}(y, z)$ for which we have constructed a gluing map. The other non-compact end is counted $W_{h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}}(x, z)$ where the gluing map is given by sending $p \in W_{h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}}(x, z)$ and $R' \in (1, \infty)$ to $(p, R + \frac{1}{R'}) \in W_{h^{\gamma\beta}, h^{\beta\alpha}}(x, z, R)$. \square

Proposition 3.4.4. *Suppose $h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\beta)$, $h^{\gamma\beta} \in T(\mathcal{Q}^\beta, \mathcal{Q}^\gamma)$ and $h^{\gamma\beta} \circ h^{\beta\alpha} \in T(\mathcal{Q}^\alpha, \mathcal{Q}^\gamma)$. Then $h_*^{\gamma\beta} \circ h_*^{\beta\alpha}$ and $(h^{\gamma\beta} \circ h^{\beta\alpha})_*$ are chain homotopic, i.e. there is a degree +1 map $P_*^{\gamma\alpha} : C_*(\mathcal{Q}^\alpha) \rightarrow C_*(\mathcal{Q}^\gamma)$, such that*

$$h_k^{\gamma\beta} h_k^{\beta\alpha} - \left(h^{\gamma\beta} \circ h^{\beta\alpha} \right)_k = P_{k-1}^{\gamma\alpha} \partial_k^\alpha + \partial_{k+1}^\gamma P_k^{\gamma\alpha}, \quad \text{for all } k.$$

Proof. By Proposition 3.4.3 we have that

$$\begin{aligned} & \left(\left(h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha} \right)_* - h_*^{\gamma\beta} h_*^{\beta\alpha} \right) (x) \\ &= \sum_{|z|=|x|} \left(n_{h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}}(x, z) - \sum_{|y|=|x|} n_{h^{\gamma\beta}}(y, z) n_{h^{\beta\alpha}}(x, y) \right) z \\ &= \sum_{|z|=|x|} \partial \widehat{W}(x, z, R) z = 0. \end{aligned}$$

The homotopy between $h^{\gamma\beta} \circ h^{\beta\alpha}$ and $h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}$ sending $\lambda \in [0, 1]$ to $h^{\gamma\beta} \circ \psi_{\lambda R}^\beta \circ h^{\beta\alpha}$, induces a chain homotopy by Proposition 3.3.1. That is, there a degree +1 map P_* such that

$$\left(h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha} \right)_k - \left(h^{\gamma\beta} \circ h^{\beta\alpha} \right)_k = -\partial_{k+1}^\gamma P_k^{\gamma\alpha} - P_{k-1}^{\gamma\alpha} \partial_k^\alpha \quad \text{for all } k.$$

Combining the last two equations gives the chain homotopy. \square

3.5 Isolation properties of maps

For the remainder of this chapter, we do not assume that the base manifolds are necessarily closed. We localize the discussion on functoriality in Morse homology on closed manifolds, and study functoriality for local Morse homology, as well as functoriality for Morse-Conley-Floer homology. For this we need isolation properties of maps. For a manifold equipped with a flow ϕ , denote the forwards and backwards orbit as follows

$$\mathcal{O}_+(p) := \{\phi(t, p) \mid t \geq 0\}, \quad \mathcal{O}_-(p) := \{\phi(t, p) \mid t \leq 0\}.$$

3. FUNCTORIALITY

Definition 3.5.1. Let M^α, M^β be manifolds, equipped with flows ϕ^α, ϕ^β . Let N^α, N^β be isolating neighborhoods. A map $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ is an *isolating map* (with respect to N^α, N^β) if the set

$$S_{h^{\beta\alpha}} = \{p \in N^\alpha \mid \mathcal{O}_-(p) \subset N^\alpha, \mathcal{O}_+(h^{\beta\alpha}(p)) \subset N^\beta\}, \quad (3.12)$$

satisfies the property that for all $p \in S_{h^{\beta\alpha}}$ we have that

$$\mathcal{O}_-(p) \subset \text{int } N^\alpha, \quad \mathcal{O}_+(h^{\beta\alpha}(p)) \subset \text{int } N^\beta.$$

Compositions of isolating maps need not be isolating, however the condition is open in the compact-open topology.

Proposition 3.5.2. *The condition of isolating maps is open: Let M^α, M^β be manifolds, equipped with flows ϕ^α, ϕ^β . Let N^α, N^β be isolating neighborhoods and $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ an isolating map. Then there exist an open neighborhood A of $(\phi^\alpha, h^{\beta\alpha}, \phi^\beta)$ in*

$$C = \{(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \tilde{\phi}^\beta) \in C^\infty(\mathbb{R} \times M^\alpha, M^\alpha) \times C^\infty(M^\alpha, M^\beta) \times C^\infty(\mathbb{R} \times M^\beta, M^\beta) \mid N^\alpha, N^\beta \text{ are isolating neighborhoods of flows } \tilde{\phi}^\alpha, \tilde{\phi}^\beta\}$$

equipped with the compact open topology⁶ such $\tilde{h}^{\beta\alpha}$ is isolating with respect to N^α, N^β and flows $\tilde{\phi}^\alpha$ and $\tilde{\phi}^\beta$.

Proof. Define the sets

$$S_-^\alpha = \{p \in N^\alpha \mid \mathcal{O}_-(p) \subset N^\alpha\} = \bigcap_{T \leq 0} \phi^\alpha([T, 0], N^\alpha),$$

$$S_+^\alpha = \{p \in N^\alpha \mid \mathcal{O}_+(p) \subset N^\alpha\} = \bigcap_{T \geq 0} \phi^\alpha([0, T], N^\alpha),$$

which are compact. Then $S_{h^{\beta\alpha}} = S_-^\alpha \cap (h^{\beta\alpha})^{-1}(S_+^\beta)$. Note that a map being isolating is equivalent to the following two properties of points on the boundary of the isolating neighborhoods in the domain and the codomain.

Domain: For all $p \in \partial N^\alpha$ either

- (ai) There exists a $t < 0$ such that $\phi^\alpha(t, p) \in M^\alpha \setminus N^\alpha$.

⁶The space C itself is open in the space of triples $(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \tilde{\phi}^\beta)$ with $\tilde{\phi}^\alpha$ and $\tilde{\phi}^\beta$ flows, cf. Proposition 2.3.8.

(iii) $\mathcal{O}_-(p) \subset N^\alpha$. Then

$$T = \sup\{t \in \mathbb{R}_{\geq 0} \mid \phi^\alpha([0, t], p) \subset N^\alpha\}$$

is finite⁷, and for all $q \in \phi^\alpha([0, T], p)$ there exists $s \geq 0$ such that $\phi^\beta(s, h^{\beta\alpha}(q)) \in M^\beta \setminus N^\beta$.

Codomain: For all $p \in N^\beta$ either

(bi) There exists a $t > 0$ such that $\phi^\beta(t, p) \in M^\beta \setminus N^\beta$.

(bii) $\mathcal{O}_+(p) \subset N^\beta$, and

$$T = \inf\{t \in \mathbb{R}_{\leq 0} \mid \phi^\beta([t, 0], p) \subset N^\beta\}$$

is finite and $h^{\beta\alpha}(S_-^\alpha) \cap \phi^\beta([T, 0], p) = \emptyset$.

In each of these cases we construct opens $U_p^{\alpha/\beta} \subset M^{\alpha/\beta}$ and $A_p^{\alpha/\beta} \subset C$ such that for all $q \in U_p$ and $(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \tilde{\phi}^\beta) \in A_p^{\alpha/\beta}$ such that these properties remain true for points on the boundary of the isolating neighborhoods. By compactness of N^α and N^β we can choose finite number of $U_p^{\alpha/\beta}$ that still cover $\partial N^{\alpha/\beta}$. Then

$$A = \left(\bigcap_j A_{p_j}^\alpha \right) \cap \left(\bigcap_j A_{p_j}^\beta \right)$$

is the required open set.

Let $p \in \partial N^\alpha$ and assume that (ai) holds. Then there exists a $t < 0$ such that $\phi^\alpha(t, p) \in M^\alpha \setminus N^\alpha$. By continuity there exists an open neighborhood $U_p \ni p$ such that $\phi^\alpha(t, \overline{U_p}) \subset M^\alpha \setminus N^\alpha$. Define

$$A_p^\alpha = \{(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \tilde{\phi}^\beta) \in C \mid \tilde{\phi}^\alpha(t, \overline{U_p}) \subset M^\alpha \setminus N^\alpha\}.$$

This is by definition open in the compact-open topology. Then for all $q \in U_p$ and all $(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \tilde{\phi}^\beta) \in A_p^\alpha$ there exists a $t < 0$ such that $\tilde{\phi}^\alpha(t, q) \in M^\alpha \setminus N^\alpha$.

Now let $p \in \partial N^\alpha$ and assume that (aii) holds. Choose $t \in [0, T]$ and $s > 0$ such that $\phi^\beta(s, h^{\beta\alpha}(\phi^\alpha(t, p))) \in M^\beta \setminus N^\beta$. By continuity of ϕ^β there

⁷Note that it cannot be the case that $\mathcal{O}_+(p)$ is also contained in N^α since N^α is an isolating neighborhood of the flow ϕ^α and the orbit through a boundary point must leave the isolating neighborhood at some time.

3. FUNCTORIALITY

exists a neighborhood $V_{t,p}^\beta \ni h^{\beta\alpha}(\phi^\alpha(t,p))$ such that $\phi^\beta(s, \overline{V_{t,p}^\beta}) \subset M^\beta \setminus N^\beta$. Again by continuity there exists a neighborhood $V_{t,p}^\alpha \ni \phi^\alpha(t,p)$ such that $h^{\beta\alpha}(\overline{V_{t,p}^\alpha}) \subset V_{t,p}^\beta$. Since $\phi^\alpha([0, T], p)$ is compact and covered by the $V_{t,p}^\alpha$, we can choose a finite number of t_j such that the sets $V_{t_j,p}^\alpha$ cover $\phi^\alpha([0, T], p)$. Now there exists an $\epsilon > 0$ small and a neighborhood $V_p \ni p$ such that $\phi^\alpha(T + \epsilon, \overline{V_p^\alpha}) \in (M^\alpha \setminus N^\alpha) \cap (\bigcup_j V_{t_j,p}^\alpha)$ and $\phi^\alpha([0, T + \epsilon], \overline{V_p^\alpha}) \subset \bigcup_j V_{t_j,p}^\alpha$. Set

$$U_p^\alpha = \left(\bigcap_j \phi^\alpha(-t_j, V_{t_j,p}^\alpha) \right) \cap V_p,$$

$$A_p^\alpha = \{(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \tilde{\phi}^\beta) \in C \mid \tilde{\phi}^\alpha(t + \epsilon, \overline{V_p^\alpha}) \subset (M^\alpha \setminus N^\alpha) \cap (\bigcup_j V_{t_j,p}^\alpha),$$

$$\tilde{\phi}^\alpha([0, T + \epsilon], \overline{V_p^\alpha}) \subset \bigcup_j V_{t_j,p}^\alpha, \quad \tilde{h}^{\beta\alpha}(\overline{V_{t_j,p}^\alpha}) \subset V_{t_j,p}^\beta, \quad \tilde{\phi}^\beta(s_j, \overline{V_{t_j,p}^\beta}) \subset M^\beta \setminus N^\beta\}.$$

Note that, for all $q \in U_p^\alpha$ and $(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \tilde{\phi}^\beta) \in A_p^\alpha$ and all $t \geq 0$ such that $\tilde{\phi}^\alpha([0, t], q) \in N^\alpha$, there exists an $s \geq 0$ such that $\tilde{\phi}^\beta(s, \tilde{h}^{\beta\alpha}(t, q)) \in M^\beta \setminus N^\beta$. It is not necessarily true that $\mathcal{O}_-(q) \subset N^\alpha$, but this does not matter. Now, by compactness there exists a finite number of $p_i \in \partial N^\alpha$ such that the corresponding U_{p_i} cover ∂N^α . Then $A^\alpha = \bigcap_i A_{p_i}^\alpha$ is open, and for all $p \in \partial N^\alpha$ either property (ai) or (aii) holds with respect to each $(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \tilde{\phi}^\beta) \in A^\alpha$.

Now we study the codomain. Let $p \in \partial N^\beta$ and assume that (bi) holds. Then completely analogously to (ai) there exists a $t > 0$ and an open U_p^β such that $\phi^\beta(t, \overline{U_p^\beta}) \subset M^\beta \setminus N^\beta$. Define

$$A_p^\beta = \{(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \tilde{\phi}^\beta) \in C \mid \tilde{\phi}^\beta(t, \overline{U_p^\beta}) \subset M^\beta \setminus N^\beta\}.$$

Then for all $q \in U_p^\beta$ and all $(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \tilde{\phi}^\beta) \in A_p^\beta$ there exists a $t > 0$ such that $\tilde{\phi}^\beta(t, q) \in N^\beta$.

Finally let $p \in \partial N^\beta$ and assume that (bii) holds. Thus $\mathcal{O}_+(p) \subset N^\beta$ and $T = \inf\{t \in \mathbb{R}_{\leq 0} \mid \phi^\beta([T, 0], p) \subset N^\beta\} > -\infty$ and $h^{\beta\alpha}(S_-) \cap \phi^\beta([T, 0], p) = \emptyset$. Since both sets are compact, there exists an open $V \supset \phi^\beta([T, 0], p)$ such that $h^{\beta\alpha}(S_-) \cap \overline{V} = \emptyset$. This implies that for all $q \in (h^{\beta\alpha})^{-1}(\overline{V}) \cap N^\alpha$ there exists an $s < 0$ such that $\phi^\alpha(s, q) \in M^\alpha \setminus N^\alpha$. But then there exists an open $V_q^\alpha \ni q$ such that $\phi^\alpha(s, \overline{V_q^\alpha}) \subset M^\alpha \setminus N^\alpha$. The set $(h^{\beta\alpha})^{-1}(\overline{V}) \cap N^\alpha$ is compact hence it

is covered by $V_{q_i}^\alpha$ for a finite number of q_i . Note that for all flows $\tilde{\phi}^\alpha$ with $\tilde{\phi}^\alpha(s_i, \overline{V_{q_i}^\alpha}) \subset M^\alpha \setminus N^\alpha$ we have that $\tilde{S}_-^\alpha = \{p \in N^\alpha \mid \tilde{\phi}^\alpha(t, p) \in N^\alpha, \forall t \leq 0\}$ is contained in $N^\alpha \setminus \bigcup_i \overline{V_{q_i}^\alpha}$. Let us return to the codomain. By continuity of the flow and the choice of T there exists an $\epsilon > 0$ such that $\phi^\beta(T - \epsilon, p) \subset (M^\alpha \setminus N^\alpha) \cap V$ with $\phi^\beta([T - \epsilon, 0], p) \subset V$. Fix a neighborhood $V_p^\beta \ni p$ such that $\phi^\beta(T - \epsilon, \overline{V_p^\beta}) \subset (M^\beta \setminus N^\beta) \cap V$. Now for $t \in [T, 0]$ the sets $\phi^\beta(t, V_p^\beta) \cap V$ cover $\phi^\beta([T, 0], p)$. Using compactness, choose a finite number of times t_j such that $\phi^\beta(t_j, V_p^\beta) \cap V$. For all $q \in \bigcap_j \phi(-t_j, \phi^\beta(t_j, V_p^\beta) \cap V)$ we have that $T_q = \inf\{t \in \mathbb{R}_{\leq 0} \mid \phi^\beta([t, 0], q)\} \geq T - \epsilon$, and for all $t \in T_q$, we have $\phi^\beta(t, q) \subset V$. Define

$$A_p^\beta = \{(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \phi^\beta) \mid \tilde{\phi}^\alpha(s_i, \overline{V_{q_i}^\alpha}) \subset M^\alpha \setminus N^\alpha \\ \tilde{h}^{\beta\alpha}(N^\alpha \setminus \bigcup V_{q_i}^\alpha) \subset M^\beta \setminus \overline{V}, \quad \tilde{\phi}^\beta([T - \epsilon] \times \overline{U_p^\beta}) \subset V\}.$$

Then for all $q \in U_p^\beta$ and all $(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \phi^\beta) \in A_p^\beta$

$$\tilde{h}^{\beta\alpha}(\tilde{S}_-^\alpha) \cap \tilde{\phi}^\beta([T_q, 0], q) = \emptyset.$$

The sets U_p^β cover ∂N^β . Hence we can choose a finite number of p_j such that for all $p \in \partial N^\beta$ and all $(\tilde{\phi}^\alpha, \tilde{h}^{\beta\alpha}, \phi^\beta) \in A^\beta = \bigcap_j A_{p_j}^\beta$, either (bi) or (bii) holds.

Set $A = A^\alpha \cap A^\beta$. □

Definition 3.5.3. A homotopy h_λ is *isolating* with respect to isolating homotopies of flows ϕ_λ^α and ϕ_λ^β if each h_λ is an isolating map with respect to ϕ_λ^α and ϕ_λ^β . The maps h_0 and h_1 are said to be isolated homotopic to each other.

We remark that isolating homotopies are also open in the compact open topology. An isolating homotopy h_λ can be viewed as an isolating map $H : M^\alpha \times [0, 1] \rightarrow M^\beta \times [0, 1]$, with $H(x, \lambda) = (h_\lambda(x), \lambda)$ with flows $\Phi^\alpha = (\phi_\lambda^\alpha, \lambda)$. Thus the set of isolating homotopies is open in the compact open topology by Proposition 3.5.2. Flow maps, cf. Definition 3.1.3 are isolating with respect to well chosen isolating neighborhoods.

Proposition 3.5.4. *Let $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ be a flow map. Then for each isolating neighborhood N^β of ϕ^β , $N^\alpha = (h^{\beta\alpha})^{-1}(N^\beta)$ is an isolating neighborhood, and $S^\alpha = \text{Inv}(N^\alpha, \phi^\alpha) = (h^{\beta\alpha})^{-1}(\text{Inv}(N^\beta, \phi^\beta)) = (h^{\beta\alpha})^{-1}(S^\beta)$. Moreover $h^{\beta\alpha}$ is isolating*

3. FUNCTORIALITY

w.r.t. N^α and N^β . If $h^{\gamma\beta}$ is another flow map, and N^γ is an isolating neighborhood then $h^{\gamma\beta} \circ \phi_R^\beta \circ h^{\beta\alpha}$ is isolating for all R with respect to $N^\alpha = (h^{\gamma\beta} \circ h^{\beta\alpha})^{-1}(N^\gamma)$.

Proof. We follow McCord [McC88]. Since $h^{\beta\alpha}$ is proper N^α is compact. If $p \in \text{Inv}(N^\alpha)$ then $\phi^\alpha(t, p) \in N^\alpha$ for all $t \in \mathbb{R}$. By equivariance $\phi^\beta(t, h^{\beta\alpha}(p)) \in N^\beta$ for all t and since N^β is an isolating neighborhood $h^{\beta\alpha}(p) \in \text{int}(N^\beta)$. Thus $p \in (h^{\beta\alpha})^{-1}(\text{int } N^\beta) \subset \text{int } N^\alpha$. If $p \in S^\alpha$, then for all t we have that $\phi^\alpha(t, p) \in N^\alpha$, and thus $h^{\beta\alpha}(\phi^\alpha(t, p)) \in N^\beta$. By equivariance it follows that $\phi^\beta(t, h^{\beta\alpha}(p)) \in N^\beta$ for all t . Hence $h^{\beta\alpha}(p) \in S^\beta$. Analogously, if $p \in (h^{\beta\alpha}(p))^{-1}(S^\beta)$ then $\phi^\beta(t, h^{\beta\alpha}(p)) \in S^\beta$ for all t . By equivariance it follows that $h^{\beta\alpha}(\phi^\alpha(t, p)) \in S^\beta$ and this implies that $\phi(t, p) \in (h^{\beta\alpha})^{-1}(N^\beta)$ for all t . We finally show that $h^{\beta\alpha}$ is isolating. If $p \in S_{h^{\beta\alpha}}$, then $\phi^\alpha(t, p) \in N^\alpha$ for all $t < 0$, and $\phi^\beta(t, h^{\beta\alpha}(p)) \in N^\beta$ for all $t > 0$. By equivariance $h^{\beta\alpha}(\phi^\alpha(t, p)) \in N^\beta$ for all $t > 0$, and thus $\phi^\alpha(t, p) \in N^\alpha$ for all t . Thus $p \in S^\alpha$ and $\mathcal{O}_-(p) \subset S^\alpha \subset \text{int } N^\alpha$. Similarly $h^{\beta\alpha}(p) \in S^\beta$ thus $\mathcal{O}_+(h^{\beta\alpha}(p)) \subset S^\beta \subset \text{int } N^\beta$. The proof of the latter statement follows along the same lines. \square

The previous proposition states that we can pull back isolated invariant sets and neighborhoods along flow maps. The same is true for Lyapunov functions, cf. Definition 2.2.4.

Proposition 3.5.5. *Let $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ be a flow map with respect to ϕ^α and ϕ^β . Let $S^\beta \subset M^\beta$ be an isolated invariant set, and $f^\beta : M^\beta \rightarrow M^\alpha$ a Lyapunov function satisfying the Lyapunov property with respect to an isolating neighborhood N^β . Then $f^\alpha := f^\beta \circ h^{\beta\alpha}$ is a Lyapunov function for $S^\alpha := (h^{\beta\alpha})^{-1}(S^\beta)$ satisfying the Lyapunov property on $N^\alpha = (h^{\beta\alpha})^{-1}(N^\beta)$. Moreover, for any metric e^α, e^β the map $h^{\beta\alpha}$ is isolating with respect to the gradient flows ψ^α and ψ^β of (f^α, e^α) and (f^β, e^β) .*

Proof. Since $h^{\beta\alpha}(S^\alpha) \subset S^\beta$ and $f^\beta|_{S^\beta} \equiv c$ is constant it follows that $f^\alpha|_{S^\alpha} \equiv c$ is constant. If $p \in N^\alpha \setminus S^\alpha$ then there exists a $t \in \mathbb{R}$ such that $\phi^\alpha(t, p) \subset M^\alpha \setminus N^\alpha$. Then $h^{\beta\alpha}(\phi^\alpha(t, p)) \in M^\beta \setminus N^\beta$ and by equivariance $\phi^\beta(t, h^{\beta\alpha}(p)) \in M^\beta \setminus N^\beta$. Therefore $h^{\beta\alpha}(p) \in N^\beta \setminus S^\beta$. Again by equivariance

$$\left. \frac{d}{dt} \right|_{t=0} f^\alpha(\phi^\alpha(t, p)) = \left. \frac{d}{dt} \right|_{t=0} f^\beta(h^{\beta\alpha}(\phi^\alpha(t, p))) = \left. \frac{d}{dt} \right|_{t=0} f^\beta(\phi^\beta(t, h^{\beta\alpha}(p))) < 0.$$

Thus f^α is a Lyapunov function. We now prove that $h^{\beta\alpha}$ is an isolating map with respect to the isolating neighborhoods N^α and N^β with respect to gradient flows ψ^α and ψ^β . The neighborhoods are isolating for the gradient flows

by Lemma 2.3.3. Let $P_{h^{\beta\alpha}}$ be the set of Equation (3.12) for the gradient flows. By the arguments of Lemma 2.3.3 we have that, for $p \in P_{h^{\beta\alpha}}$ that

$$\alpha(p) = \left\{ \lim_{n \rightarrow \infty} \psi^\alpha(t_n, p) \mid \lim_{n \rightarrow \infty} t_n = -\infty \right\} \subset \text{Crit } f^\alpha \cap N^\alpha,$$

and

$$\omega(h^{\beta\alpha}(p)) = \left\{ \lim_{n \rightarrow \infty} \psi^\beta(t_n, h^{\beta\alpha}(p)) \mid \lim_{n \rightarrow \infty} t_n = \infty \right\} \subset \text{Crit } f^\beta \cap N^\beta.$$

Consider $b_p : \mathbb{R} \rightarrow \mathbb{R}$ with

$$b_p(t) = \begin{cases} f^\alpha(\psi^\alpha(t, p)) & \text{for } t \leq 0 \\ f^\beta(\psi^\beta(t, h^{\beta\alpha})) & \text{for } t > 0. \end{cases}$$

The function b_p is continuous, smooth outside zero and by the Lyapunov property $\frac{d}{dt} b_p(t) \leq 0$. By Lemma 2.3.1 we know that $\text{Crit } f^\alpha \cap N^\alpha \subset S^\alpha$ and $\text{Crit } f^\beta \cap N^\beta \subset S^\beta$. Moreover $f^\alpha|_{S^\alpha} \equiv f^\beta|_{S^\beta} \equiv c$, hence $\lim_{t \rightarrow -\infty} b_p(t) = \lim_{t \rightarrow \infty} b_p(t) = c$. Hence b_p is constant and it follows that $p \in S^\alpha$. Thus the full orbit through p contained in $\text{int } N^\alpha$ and the full orbit through $h^{\beta\alpha}(p)$ is contained in $\text{int } N^\beta$. Thus $h^{\beta\alpha}$ is isolating with respect to the gradient flows. \square

3.6 Local Morse homology

We are interested in the functorial behavior of local Morse homology. Let us recall the definition of local Morse homology more in depth. Suppose N^α is an isolating neighborhood of the gradient flow of f^α an g^α . The local stable and unstable manifolds of critical points inside N^α are defined by

$$W_{\text{loc}}^u(x; N^\alpha) := \{p \in N^\alpha \mid \psi^\alpha(t, p) \in N^\alpha, \forall t < 0, \text{ and } \lim_{t \rightarrow -\infty} \psi^\alpha(t, p) = x\},$$

$$W_{\text{loc}}^s(x; N^\alpha) := \{p \in N^\alpha \mid \psi^\alpha(t, p) \in N^\alpha, \forall t > 0, \text{ and } \lim_{t \rightarrow \infty} \psi^\alpha(t, p) = x\}.$$

We write $W_{N^\alpha}(x, y) = W_{\text{loc}}^u(x, y; N^\alpha) \cap W_{\text{loc}}^s(y; N^\alpha)$, and $M_{N^\alpha}(x, y) = W_{N^\alpha}(x, y)/\mathbb{R}$. We say that the gradient flow is Morse-Smale on N^α , cf. Definition 2.3.5, if the critical points of f^α inside N^α are non-degenerate, and for each $p \in W_{N^\alpha}(x, y)$, we have that $T_p W^u(x) + T_p W^s(y) = T_p M^\alpha$. The intersection is said to be transverse and we write $W_{\text{loc}}^u(x; N^\alpha) \pitchfork W_{\text{loc}}^s(y; N^\alpha)$. Denote by $\mathcal{P}^\alpha = (M^\alpha, f^\alpha, g^\alpha, N^\alpha, \mathfrak{o}^\alpha)$ a choice of manifold (not necessarily closed), a

3. FUNCTORIALITY



Figure 3.6: All isolating maps induce chain maps in local Morse homology, but they are not necessarily functorial, as they do not capture the dynamical content. The gradient flows of $f^\alpha(x) = x^2$, $f^\beta = (x - 3)^2$, and $f^\gamma(x) = x^2$ on \mathbb{R} , with isolating neighborhood $N = [-1, 1]$ are depicted. The identity maps are isolating. We compute that $\text{id}_*^{\beta\alpha} = 0$, $\text{id}_*^{\gamma\beta} = 0$, but $\text{id}_*^{\gamma\alpha}$ is the identity. The problem is that $\text{id}^{\gamma\beta} \circ \psi_R^\beta \circ \text{id}^{\beta\alpha}$ is not isolating for all $R > 0$.

function, a metric and an isolating neighborhood of the gradient flow, such that (f^α, g^α) is Morse-Smale on N^α , and a choice of orientations of the local unstable manifolds. For such a local Morse datum \mathcal{P}^α , we define the local Morse complex

$$C_k(\mathcal{P}^\alpha) := \text{Crit}_k f^\alpha \cap N^\alpha, \quad \partial(\mathcal{P}^\alpha)(x) := \sum_{|y|=|x|-1} n_{N^\alpha} \text{loc}(x, y)y.$$

Where $n_{N^\alpha}(x, y)$ denotes the oriented intersection number in N^α . The differential satisfies $\partial^2(\mathcal{P}^\alpha) = 0$, hence we can define local Morse homology. This is not an invariant for N^α but crucially depends on the gradient flow. Compare for example the gradient flows of $f^\alpha(x) = x^2$ or $f^\beta(x) = -x^2$ on \mathbb{R} with isolating neighborhoods $N = [-1, 1]$. The homology is invariant under homotopies (f_λ, g_λ) , as long as the gradient flows preserve isolation. The canonical isomorphisms, induced by continuation are denoted by $\Phi^{\beta\alpha}$. The local Morse homology recovers the Conley index of the gradient flow.

If $h^{\beta\alpha}$ is isolating with respect to $\mathcal{P}^\alpha, \mathcal{P}^\beta$, we say it is transverse (with respect to \mathcal{P}^α and \mathcal{P}^β), if for all $p \in W_{\text{loc}}^u(x; N^\alpha) \cap (h^{\beta\alpha})^{-1}W_{\text{loc}}^s(y; N^\beta)$, we have

$$dh^{\beta\alpha} T_p W^u(x) + T_{h^{\beta\alpha}(p)} W^s(y) = T_{h(p)} M^\beta.$$

We write $W_{h^{\beta\alpha}, \text{loc}}(x, y) = W_{\text{loc}}^u(x; N^\alpha) \cap W_{\text{loc}}^s(y; N^\beta)$. The oriented intersection number is denoted by $n_{h^{\beta\alpha}, \text{loc}}(x, y)$.

Proposition 3.6.1. *Let $h^{\beta\alpha} \in T(\mathcal{P}^\alpha, \mathcal{P}^\beta)$ and suppose $h^{\beta\alpha}$ is isolating with respect to R^α and R^β .*

- Then $h_*^{\beta\alpha}(x) := \sum_{|x|=|y|} n_{h^{\beta\alpha}, \text{loc}}(x, y)y$ is a chain map.

- Suppose \mathcal{P}^γ and \mathcal{P}^δ are different local Morse data with

$$M^\alpha = M^\gamma \quad N^\alpha = N^\gamma \quad M^\beta = M^\delta, \quad \text{and} \quad N^\beta = N^\delta,$$

such that the gradient flow of \mathcal{P}^α is isolated homotopic to the gradient flow of \mathcal{P}^γ , the gradient flow of \mathcal{P}^β is isolated homotopic to the gradient flow of \mathcal{P}^δ , and $h^{\beta\alpha} \in T(\mathcal{P}^\alpha, \mathcal{P}^\beta)$ and $h^{\delta\gamma} \in T(\mathcal{P}^\gamma, \mathcal{P}^\delta)$ are isolated homotopic through these isolating homotopies. Then $\Phi_*^{\delta\beta} h_*^{\beta\alpha}$ and $h_*^{\delta\gamma} \Phi_*^{\gamma\alpha}$ are chain homotopic.

Proof. We argue that the gluing maps constructed in Proposition 3.2.4 restrict to local gluing maps

$$\begin{aligned} \#^1 : M_{N^\alpha}(x, y) \times [R_0, \infty) \times W_{h^{\beta\alpha}, \text{loc}}(y, y') &\rightarrow W_{h^{\beta\alpha}, \text{loc}}(x, y') \\ \#^2 : W_{h^{\beta\alpha}, \text{loc}}(x, x') \times [R_0, \infty) \times M_{N^\beta}(x', y') &\rightarrow W_{h^{\beta\alpha}, \text{loc}}(x, y') \end{aligned}$$

Consider $\#^1$, let $\gamma_u \in M_{N^\alpha}(x, y)$ and $v \in W_{h^{\beta\alpha}, \text{loc}}(y, y')$. Geometric convergence implies that the backwards orbit $\mathcal{O}_-(\gamma_u \#_R^1 v)$ lies arbitrary close, for R sufficiently large, to the images of $\mathcal{O}_-(u)$ and $\mathcal{O}_-(v)$. The latter are contained in $\text{int } N^\alpha$, hence $\mathcal{O}_-(\gamma_u \#_R^1 v) \subset \text{int } N^\alpha$. Similarly $\mathcal{O}_+(h^{\beta\alpha}(u \#_R^1 v)) \subset \text{int } N^\beta$ for R sufficiently large since it must converge to $\mathcal{O}_+(h^{\beta\alpha}(v))$. Thus for R_0 possibly larger than in Proposition 3.2.4 there is a well defined restriction of the gluing map $\#^1$. The situation for $\#^2$ is analogous.

For the compactness issues, observe that $S_{h^{\beta\alpha}} = S_-^\alpha \cap h^{\beta\alpha}(S_+^\beta)$ hence is compact. But it also equals the set $\bigcup_{x, y} W_{h^{\beta\alpha}, \text{loc}}(x, y)$. Therefore if $p_k \in W_{h^{\beta\alpha}, \text{loc}}(x, y)$ then it has a convergent subsequence $p_k \rightarrow p$ with $p \in W_{h^{\beta\alpha}, \text{loc}}(x', y')$. Since $\mathcal{O}_-(p_k) \subset \text{int } N^\alpha$ and $\mathcal{O}_+(h^{\beta\alpha}(p_k)) \subset \text{int } N^\beta$ it follows that the compactness issues are those described in Proposition 3.2.2, but for the local moduli spaces.

Now consider the isolating homotopy. Isolation implies that we can follow the argument of Proposition 3.3.1 to construct isolating neighborhoods on $M^\alpha \times [0, 1]$ and $M^\beta \times [0, 1]$ and a map $H^{\beta\delta}$ that is isolating with respect to these neighborhoods. Then the observation that $H^{\beta\delta}$ is a chain map implies that there exists a chain homotopy. \square

Recall that the local Morse homology of the isolating neighborhood of any pair (f, g) which is not necessarily Morse-Smale is defined as

$$HM_*(f, g, N) = \varprojlim HM_*(\mathcal{P}^\alpha)$$

3. FUNCTORIALITY

where the inverse limit runs over all local Morse data \mathcal{P}^α whose gradient flows are isolated homotopic to the gradient flows of (f, g) on N , with respect to the canonical isomorphisms.

Proposition 3.6.2. *Let (f^α, g^α) and (f^β, g^β) pairs of functions of metrics on M^α and M^β , which are not assumed to be Morse-Smale. Suppose that $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ is isolating with respect to isolating neighborhoods N^α and N^β . Then $h^{\beta\alpha}$ induces a map of local Morse homologies*

$$h_*^{\beta\alpha} : HM_*(f^\alpha, g^\alpha, N^\alpha) \rightarrow HM_*(f^\beta, g^\beta, N^\beta).$$

Suppose $h^{\delta\gamma}$ is isolated homotopic to $h^{\beta\alpha}$ through isolated homotopies of gradient flows between $(f^\alpha, g^\alpha, N^\alpha)$ and $(f^\gamma, g^\gamma, N^\gamma)$ and $(f^\beta, g^\beta, N^\beta)$ and $(f^\delta, g^\delta, N^\delta)$. Then the diagram

$$\begin{array}{ccc} HM_*(f^\alpha, g^\alpha, N^\alpha) & \xrightarrow{h_*^{\beta\alpha}} & HM_*(f^\beta, g^\beta, N^\beta) \\ \downarrow \Phi_*^{\gamma\alpha} & & \downarrow \Phi_*^{\delta\beta} \\ HM_*(f^\gamma, g^\gamma, N^\gamma) & \xrightarrow{h_*^{\delta\gamma}} & HM_*(f^\delta, g^\delta, N^\delta) \end{array}$$

commutes.

Proof. Let $\mathcal{P}^\gamma = (M^\alpha, f^\gamma, g^\gamma, N^\alpha, \sigma^\gamma)$ and $\mathcal{P}^\delta = (M^\beta, f^\delta, g^\delta, N^\beta, \sigma^\delta)$ be local Morse data such that

$$(f^\gamma, g^\gamma) \in \mathcal{I}_{MS}(f^\alpha, g^\alpha, N^\alpha), \quad (f^\delta, g^\delta) \in \mathcal{I}_{MS}(f^\beta, g^\beta, N^\beta),$$

and $h^{\beta\alpha}$ is isolating for the isolating homotopies connecting the (f^α, g^α) with (f^γ, g^γ) and (f^β, g^β) with (f^δ, g^δ) and $h^{\beta\alpha} \in \mathcal{T}(\mathcal{P}^\gamma, \mathcal{P}^\delta)$. This is possible by Proposition 3.5.2, and Corollary 2.3.12. By the density of transverse maps Theorem 3.8.1, there exist a small perturbation $h^{\delta\gamma}$ isolated homotopic to $h^{\beta\alpha}$ by this homotopy. By Proposition 3.6.1 we get a map

$$h_*^{\delta\gamma} : HM_*(\mathcal{P}^\gamma) \rightarrow HM_*(\mathcal{P}^\delta).$$

Moreover, given different local Morse data \mathcal{P}^ϵ and \mathcal{P}^ζ as above, we can construct isolating map $h^{\zeta\epsilon}$ isolating homotopic to $h^{\beta\alpha}$ and hence also $h^{\delta\beta}$ by concatenation. From 3.6.1 it follows that there is the following commutative dia-

gram

$$\begin{array}{ccc}
 HM_*(\mathcal{P}^\gamma) & \xrightarrow{h_*^{\delta\gamma}} & HM_*(\mathcal{P}^\delta) \\
 \Phi_*^{\varepsilon\gamma} \downarrow & & \downarrow \Phi_*^{\zeta\delta} \\
 HM_*(\mathcal{P}^\varepsilon) & \xrightarrow{h_*^{\zeta\varepsilon}} & HM_*(\mathcal{P}^\zeta)
 \end{array}$$

Which means that we have an induced map $h_*^{\beta\alpha} : HM_*(f^\alpha, g^\alpha, N^\alpha) \rightarrow HM_*(f^\beta, g^\beta, N^\beta)$ between the inverse limits. The arguments for the homotopy of the maps is analogous. \square

The chain maps defined above are not necessarily functorial. Consider for example Figure 3.6. The problem is that, in the proof of functoriality for Morse homology, we need the fact that $h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}$ is isolating for all $R \geq 0$ to establish functoriality. If we require this almost homotopy to be isolating the proof of functoriality follows mutatis mutandis.

Proposition 3.6.3. *Let $h^{\beta\alpha}$ and $h^{\gamma\beta}$ be transverse and isolating. Assume that $h^{\gamma\beta} \circ h^{\beta\alpha}$ is transverse and isolating, and assume that $h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}$ is an isolating homotopy for $R \in [0, R']$ for each $R' > 0$. Then $h_*^{\gamma\beta} h_*^{\beta\alpha}$ and $(h^{\gamma\beta} \circ h^{\beta\alpha})_*$ are chain homotopic.*

Proposition 3.6.4. *Let (f^α, g^α) , (f^β, g^β) , and (f^γ, g^γ) be pairs of functions and metrics on manifolds $M^\alpha, M^\beta, M^\gamma$, and suppose $h^{\gamma\beta}$ and $h^{\beta\alpha}$ are flow maps with respect to the gradient flows. Let N^γ be an isolating neighborhood of the gradient flow of (f^γ, g^γ) . Set*

$$N^\beta := (h^{\beta\alpha})^{-1}(N^\alpha), \quad N^\gamma = (h^{\gamma\beta})^{-1}(N^\beta).$$

Then these are isolating neighborhoods, the maps $h^{\gamma\beta}$, $h^{\beta\alpha}$ and $h^{\gamma\beta} \circ h^{\beta\alpha}$ are isolating and the following diagram commutes

$$\begin{array}{ccc}
 HM_*(f^\alpha, g^\alpha, N^\alpha) & \xrightarrow{h_*^{\beta\alpha}} & HM_*(f^\beta, g^\beta, N^\beta) \\
 \searrow (h^{\gamma\beta} \circ h^{\beta\alpha})_* & & \downarrow h_*^{\gamma\beta} \\
 & & HM_*(f^\gamma, g^\gamma, N^\gamma)
 \end{array}$$

3. FUNCTORIALITY

Proof. The proposition follows from combining Propositions 3.5.4, 3.6.3 and 3.6.2. \square

Theorem 3.1.5 now follows from the fact that isolated homotopic maps induce the same maps in local Morse homology.

3.7 Morse-Conley-Floer homology

We use the induced maps of local Morse homology to define induced maps for flow maps in Morse-Conley-Floer homology.

Theorem 3.7.1. *Let $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ be a flow map. Let $S^\beta \subset M^\beta$ be an isolated invariant set. Then $S^\alpha = (h^{\beta\alpha})^{-1}(S^\beta)$ is an isolated invariant set and there is an induced map*

$$h_*^{\beta\alpha} : HI_*(S^\alpha, \phi^\alpha) \rightarrow HI_*(S^\beta, \phi^\beta),$$

which is functorial: The identity is mapped to the identity and the diagram

$$\begin{array}{ccc} HI_*((h^{\gamma\beta} \circ h^{\beta\alpha})^{-1}(S^\gamma), \phi^\alpha) & \xrightarrow{h_*^{\gamma\beta}} & HI_*((h^{\beta\alpha})^{-1}(S^\gamma), \phi^\beta) \\ & \searrow_{(h^{\gamma\beta} \circ h^{\beta\alpha})_*} & \downarrow h_*^{\beta\alpha} \\ & & HI_*(S^\gamma, \phi^\gamma). \end{array}$$

commutes.

Proof. The induced map is defined as follows. Let f^β be a Lyapunov function⁸ with respect to the an isolating neighborhood N^β of (S^β, ϕ^β) . Then $N^\alpha := (h^{\beta\alpha})^{-1}(N^\beta)$ is an isolating neighborhood of $S^\alpha = (h^{\beta\alpha})^{-1}(S^\beta)$, and $f^\alpha := f^\beta \circ h^{\beta\alpha}$ is a Lyapunov function by Proposition 3.5.5. Moreover $h^{\beta\alpha}$ is isolating with respect to the gradient flows with respect to any metrics g^α and g^β . By Proposition 3.6.2 we have an induced map

$$h_*^{\beta\alpha} : HM_*(f^\alpha, g^\alpha, N^\alpha) \rightarrow HM_*(f^\beta, g^\beta, N^\beta)$$

computed by perturbing everything to a transverse situation, preserving isolation, and counting as in Proposition 3.6.1.

⁸We previously denoted this by f_{ϕ^β} but this notation is too unwieldy here

We now want to prove that the induced map passes to the inverse limit. If f^δ is another Lyapunov function with respect to an isolating neighborhood N^δ of (S^β, ϕ^β) , and $N^\gamma = (h^{\beta\alpha})^{-1}(N^\delta)$, $f^\gamma = f^\delta \circ h^{\beta\alpha}$ then we cannot directly compare the local Morse homologies. However, as in Theorem 2.4.7 the sets $N^\beta \cap N^\delta$ and $N^\alpha \cap N^\gamma$ are isolating neighborhoods, and $h^{\beta\alpha}$ is isolating with respect to these since $S_{h^{\beta\alpha}} \subset S^\alpha$ and $h^{\beta\alpha}(S_{h^{\beta\alpha}}) \subset S^\beta$.

The claim is that the following diagram in local Morse homology commutes

$$\begin{array}{ccc}
 HM_*(f^\alpha, g^\alpha, N^\alpha) & \xrightarrow{h_*^{\beta\alpha}} & HM_*(f^\beta, g^\beta, N^\beta) \\
 \downarrow & & \downarrow \\
 HM_*(f^\alpha, g^\alpha, N^\alpha \cap N^\gamma) & \xrightarrow{h_*^{\beta\alpha}} & HM_*(f^\beta, g^\beta, N^\beta \cap N^\delta) \\
 \Phi_*^{\gamma\alpha} \downarrow & & \downarrow \Phi_*^{\delta\beta} \\
 HM_*(f^\gamma, g^\gamma, N^\alpha \cap N^\gamma) & \xrightarrow{h_*^{\beta\alpha}} & HM_*(f^\delta, g^\delta, N^\beta \cap N^\delta) \\
 \downarrow & & \downarrow \\
 HM_*(f^\gamma, g^\gamma, N^\gamma) & \xrightarrow{h_*^{\beta\alpha}} & HM_*(f^\delta, g^\delta, N^\delta)
 \end{array}$$

where the vertical maps are all isomorphisms. This proves that we have a well defined map in Morse-Conley-Floer homology.

The commutativity of the square in the middle is induced by continuation as in Proposition 3.6.2. Gradient flows of different Lyapunov functions on the same isolating neighborhood are isolating homotopic by the proof of Theorem 2.4.4. Moreover since the set $S_{h^{\beta\alpha}}$ of Equation (3.12) is contained in S^α and $h^{\beta\alpha}(S^\beta) \subset S^\beta$ for any choice of Lyapunov function. The isolating homotopies of the gradient flows preserve the isolation of $h^{\beta\alpha}$.

The upper square in the diagram commutes because the set $S_{h^{\beta\alpha}}$ of Equation (3.12) is contained in S^α and $h^{\beta\alpha}(S^\alpha) \subset S^\beta$. It is possible to perturb the function f^α to a Morse-Smale function preserving the isolation in $N^\alpha \cap N^\gamma$ and similar for f^β , while also preserving isolation of $h^{\beta\alpha}$. We use the openness of isolating maps, cf. Proposition 3.5.2 and genericity of transverse maps, Theorem 3.8.1. Then the counts with respect to N^α and $N^\alpha \cap N^\gamma$ and N^β and $N^\beta \cap N^\delta$ are the same for such perturbations from which it follows that the diagram commutes. The situation for the lower square is completely analogous.

We have a well defined map in Morse-Conley-Floer homology. Functoriality might not be clear at this point, however If $h^{\gamma\beta} : M^\beta \rightarrow M^\gamma$ is another

3. FUNCTORIALITY

flow map, and f^γ is a Lyapunov function for ϕ^γ , then $f^\beta = f^\gamma \circ h^{\gamma\beta}$ and $f^\alpha = f^\gamma \circ h^{\gamma\beta} \circ h^{\beta\alpha}$ are Lyapunov functions, for the obvious isolated invariant sets and neighborhoods. Choose auxiliary metrics $g^\alpha, g^\beta, g^\gamma$, and denote the gradient flows of the Lyapunov functions by $\psi^\alpha, \psi^\beta, \psi^\gamma$. Define for $R \geq 0$ and $p \in S_{h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}}$ the map $b_{p,R} : \mathbb{R} \rightarrow \mathbb{R}$ by

$$b_{p,R}(t) = \begin{cases} f^\alpha(\psi^\alpha(t, p)) & \text{for } t < 0 \\ f^\beta(\psi^\beta(t, h^{\beta\alpha}(p))) & 0 < t < R \\ f^\gamma(\psi^\gamma(t, h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}(p))) & t > R. \end{cases}$$

The map $b_{p,R}$ is continuous and smooth outside $t = 0, R$, and outside $t = 0, R$ we see that $\frac{d}{dt}b_{p,R} \leq 0$. As in Proposition 3.5.5 we have limits $\lim_{t \rightarrow \pm\infty} b_{p,R}(t) = c$, where c is the constant with $f^\gamma|_{S^\gamma} \equiv c$. It follows that $S_{h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}} \subset S^\alpha$ and hence that the orbits through $p \in S_{h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}}$ are contained in $\text{int } N^\alpha$ and the orbits through $h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}(p)$ are contained in $\text{int } N^\gamma$ for all R . Thus $h^{\gamma\beta} \circ \psi_R^\beta \circ h^{\beta\alpha}$ is isolating for all $R \geq 0$. Perturbing everything to a transverse situation as before preserving isolation we get from Proposition 3.6.3 functoriality in Morse-Conley-Floer homology. \square

3.8 Transverse maps are generic

Theorem 3.8.1. *Let Q^α, Q^β be Morse-Smale triples. The set $\mathcal{T}(Q^\alpha, Q^\beta)$ is residual in the compact-open topology, i.e. it contains a countable intersection of open and dense sets.*

Proof. Let $x \in \text{Crit } f^\alpha$ and $y \in \text{Crit } f^\beta$. We show that

$$\mathcal{T}(x, y) = \{h \in C^\infty(M^\alpha, M^\beta) \mid h|_{W^u(x)} \pitchfork W^s(y)\}$$

is residual, from which it follows that the set of transverse maps is residual.

We first show density of $\mathcal{T}(x, y)$. The set

$$\pitchfork(W^u(x), M^\beta; W^s(y)) = \{h \in C^\infty(W^u(x), M^\beta) \mid h \pitchfork W^s(y)\},$$

is residual, cf. [Hir94, Theorem 2.1(a)], and by Baire's category theorem it is dense. We show that $h^{\beta\alpha} \in \pitchfork(M^\alpha, M^\beta; W^s(y))$ can be approximated in the

compact-open topology by maps in $\mathcal{T}(x, y)$. Because all maps in $C^\infty(M^\alpha, M^\beta)$ can be approximated by such maps $h^{\beta\alpha}$ we get the required density.

The unstable manifold $W^u(x)$ is contractible thus its normal bundle is contractible. We identify a neighborhood of $W^u(x)$ in M^α with $W^u(x) \times \mathbb{R}^{n-|x|}$. By Parametric Transversality [Hir94, Theorem 2.7 page 79], the set of $v \in \mathbb{R}^{n-|x|}$ such that $h^{\beta\alpha}|_{W^u(x) \times \{v\}} \pitchfork W^s(y)$ is dense. For a given v it is now possible to construct a flow whose time-1 map ϕ^v locally translates $W^u(x) \times \{0\}$ to $W^s(x) \times \{v\}$. Then $h^{\beta\alpha}$ can be approximated by $h^{\beta\alpha} \circ \phi^{v_k}$ with $v_k \rightarrow 0$. Thus $\mathcal{T}(x, y)$ is dense in $C^\infty(M^\alpha, M^\beta)$.

We now argue that $\mathcal{T}(x, y)$ contains a countable intersection of open sets, from which it follows that this set is residual. Consider the restriction mapping $\rho_{W^u} : C^\infty(M^\alpha, M^\beta) \rightarrow C^\infty(W^u(x), M^\beta)$, which is continuous in both the weak and strong topology. Again the transversality theorem gives that the set of transverse maps $\pitchfork(W^u(x), M^\beta; W^s(y))$ is residual, that is $\pitchfork(W^u(x), M^\beta; W^s(y)) \supset \bigcap_{k \in \mathbb{N}} U_k$ with U_k open and dense. Note that

$$\mathcal{T}(x, y) = \rho_{W^u(x)}^{-1} \left(\pitchfork(W^u(x), M^\beta; W^s(y)) \right) \supset \bigcap_{k \in \mathbb{N}} \rho_{W^u(x)}^{-1}(U_k)$$

hence contains a countable intersection of open sets. By the previous reasoning $\mathcal{T}(x, y)$ is dense, hence the open sets must be dense as well, and $\mathcal{T}(x, y)$ is residual.

Since there are only a countable number of critical points of f^α and f^β it follows that $\mathcal{T} = \bigcap_{x \in \text{Crit } f^\alpha, y \in \text{Crit } f^\beta} \mathcal{T}(x, y)$ is residual. □

Chapter 4

Duality and coincidence theory in Morse-Conley-Floer homology

Stand firm in your refusal to
remain conscious during algebra.
In real life, I assure you,
there is no such thing as algebra.

Fran Lebowitz

4.1 Introduction

We prove a Poincaré type duality theorem for Morse-Conley-Floer (co)homology, which should be compared with the results of [McC92]. The theorem expresses a duality between the Morse-Conley-Floer homology of an orientable isolated invariant set S of a flow and the Morse-Conley-Floer cohomology of S seen as an isolated invariant set of the reverse flow. Using duality and functoriality we define shriek maps¹ in Morse-Conley-Floer homology. As applications of functoriality and duality we define via standard constructions [Bre97, AS10] cup and intersection products in Morse-Conley-Floer homology. We prove projection formulas, cf. Theorem 4.9.1, and a coincidence theorem, cf. Theorem 4.10.1 for flow maps. The Lefschetz-type fixed

¹also known as *transfer* or *wrong way* maps

point theorem for flow maps due to McCord [McC88] is a corollary to the coincidence theorem.

4.2 The dual complex

We recall some definitions from homological algebra. Given a chain complex (C_*, ∂_*) the dual complex is defined via $C^k := \text{Hom}(C_k, \mathbb{Z})$ with boundary operator $\delta^k : C^{k-1} \rightarrow C^k$ the pullback of the boundary operator $\partial_k : C_k \rightarrow C_{k-1}$. Thus for $\eta \in C^{k-1}$ we have $\delta^k \eta = \eta \circ \partial_k : C_k \rightarrow \mathbb{Z}$. The homology of the dual complex C^* is the *cohomology* of C_* and is denoted by $H^k(C_*) := \ker \delta^{k+1} / \text{im } \delta^k$. There is a natural map $h : H^k(C_*) \rightarrow \text{Hom}(H_k, \mathbb{Z})$, cf. [Hat02]. If the complex C_* is free there is an exact sequence

$$0 \longrightarrow \text{Ext}(H_{n-1}(C_*), \mathbb{Z}) \longrightarrow H^n(C_*) \xrightarrow{h} \text{Hom}(H_n(C_*), \mathbb{Z}) \longrightarrow 0.$$

If we chose coefficients in $\mathbb{Z}/2\mathbb{Z}$ (or more general coefficients in a field) instead of \mathbb{Z} the first term in the exact sequence is zero and we have an isomorphism between $H^n(C_*, \mathbb{Z}/2\mathbb{Z}) \cong \text{Hom}(H_n(C_*), \mathbb{Z}/2\mathbb{Z})$.

4.3 Morse cohomology and Poincaré duality

If A, B are oriented and cooriented submanifolds of a manifold M that meet transversely, the intersection $A \cap B$ is an oriented submanifold. If the ambient manifold M is oriented, the exact sequence $0 \rightarrow TA \rightarrow TM \rightarrow NA \rightarrow 0$ coorients A , and similarly orients B . The orientation on $B \cap A$ is related to the orientation of $A \cap B$ by the formula

Check this carefully

$$\text{Or}(A \cap B) = (-1)^{(\dim M - \dim A)(\dim M - \dim B)} \text{Or}(B \cap A).$$

Now let $\mathcal{Q}^\alpha = \{M^\alpha, f^\alpha, g^\alpha, \mathfrak{o}\}$ be a Morse datum on an *oriented closed* manifold M^α . We have the exact sequence of vector spaces

$$0 \longrightarrow T_x W^u(x) \longrightarrow T_x M \longrightarrow N_x W^u(x) \longrightarrow 0.$$

Because $T_x W^u(x)$ is oriented by \mathfrak{o}^α and $T_x M$ is also oriented, this sequence orients $N_x W^u(x) \cong T_x W^s(x)$. The stable manifolds are therefore also oriented. The stable manifolds of (f^α, g^α) are precisely the *unstable manifolds* of $(-f^\alpha, g^\alpha)$ which we orient by the above exact sequence. We denote this choice by $\hat{\mathfrak{o}}^\alpha$

Definition 4.3.1. Let \mathcal{Q}^α be a Morse datum on an oriented manifold. The *dual Morse datum* $\widehat{\mathcal{Q}}^\alpha$ is defined by $\widehat{\mathcal{Q}}^\alpha = \{M^\alpha, -f^\alpha, g^\alpha, \widehat{\delta}^\alpha\}$.

Because the manifold is compact the Morse complex is finitely generated. Hence the dual complex is finitely generated. An obvious basis of $C_*(\mathcal{Q}^\alpha)$ are the critical points. An obvious basis of $C^*(\mathcal{Q}^\alpha)$ is the dual basis. That is the basis of the form

$$\eta_x(y) = \begin{cases} 1 & x = y \\ 0 & x \neq y \end{cases}$$

Note that a critical point of f^α of index $|x|$ is a critical point of $-f^\alpha$ of index $m^\alpha - |x|$.

Definition 4.3.2. Define the poincare duality map $\text{PD} : C_k(\mathcal{Q}^\alpha) \rightarrow C^{n-k}(\widehat{\mathcal{Q}}^\alpha)$ as the linear extension of the map

$$\text{PD}(x) = \eta_x.$$

Theorem 4.3.3. Let \mathcal{Q}^α be a Morse datum on an oriented closed manifold. The Poincaré duality map is a chain map hence induces a map

$$\text{PD} : H_k(\mathcal{Q}^\alpha) \rightarrow H^{m^\alpha-k}(\widehat{\mathcal{Q}}^\alpha).$$

The duality map is an isomorphism, and commutes with the canonical isomorphisms in the following way. If $\Phi_*^{\beta\alpha} : HM_*(\mathcal{Q}^\alpha) \rightarrow HM_*(\mathcal{Q}^\beta)$ and $(\widehat{\Phi}^{\beta\alpha})^* : HM^*(\widehat{\mathcal{Q}}^\beta) \rightarrow HM^*(\widehat{\mathcal{Q}}^\alpha)$ are canonical isomorphisms then

$$\text{PD}_k \Phi_k^{\beta\alpha} = (\widehat{\Phi}^{\beta\alpha})^{n-k} \text{PD}_k.$$

Proof.

□ todo

The latter proposition states that the Poincaré duality is a duality on the Morse complex of the manifold $HM_*(M^\alpha) = \varprojlim HM_*(\mathcal{Q}^\alpha)$, i.e. there is a Poincaré duality map $\text{PD}_k : HM_k(M^\alpha) \rightarrow HM^{m^\alpha-k}(M^\alpha)$.

4.4 Local Morse homology

To define Poincaré duality for local Morse homology the appropriate notion of orientation is necessary. Recall that a closed subset C of a manifold M is *orientable* if there exist a continuous section in the orientation bundle over M , cf. [Bre97, Chapter VI.7].

Definition 4.4.1. Suppose \mathcal{P}^α is a local Morse datum. The Morse datum is *orientable* if $S^\alpha = \text{Inv}(N^\alpha, \psi^\alpha)$ is orientable. A choice of section is an *orientation* of \mathcal{P}^α .

If M^α is an oriented manifold, all closed subsets are oriented, thus on an orientable manifold all local Morse data are orientable. A local Morse datum is orientable if and only if $\text{int}(N^\alpha)$ is an oriented manifold. If a local Morse datum \mathcal{P}^α is oriented we can define the dual local Morse datum $\widehat{\mathcal{P}}^\alpha = \{M^\alpha, -f^\alpha, g^\alpha, N^\alpha, \widehat{\psi}^\alpha\}$ as before. Again we have Poincaré duality isomorphisms $PD_k : HM_k(\mathcal{P}^\alpha) \rightarrow HM^{m^\alpha-k}(\widehat{\mathcal{P}}^\alpha)$. A crucial difference is now that the gradient flow of \mathcal{P}^α is in general *not* isolated homotopic to the gradient flow of $\widehat{\mathcal{P}}^\alpha$. Recall that if (f, g, N) is a triple of a function, metric and isolating neighborhood of the flow we define its local Morse homology via $HM_*(f, g, N) = \varprojlim HM_*(\mathcal{P}^\alpha)$ over all local Morse data that are isolated homotopic to the gradient flow of (f, g) in N . It is readily seen that N also isolates the gradient flow of $(-f, g)$. We have the following Poincaré duality statement

Theorem 4.4.2. Let (f, g, N) be a triple as above such that $S = \text{Inv}(N, \psi)$ is oriented. Then there is a Poincaré duality isomorphism

$$PD_k : HM_k(f, g, N) \rightarrow HM^{m^\alpha-k}(-f, g, N).$$

We need a statement about the naturality with respect to isomorphisms induced by continuation

4.5 Duality in Morse-Conley-Floer homology

If ϕ is a flow, then the *reverse flow* ϕ^{-1} defined via $\phi^{-1}(t, x) = \phi(-t, x)$ is also a flow. The following duality statement resembles an analogous theorem for the Conley index, due to McCord [McC92].

Theorem 4.5.1. *Let S be an oriented isolated invariant set of a flow ϕ . Then there are Poincaré duality isomorphisms*

$$PD_k : HI_k(S, \phi) \rightarrow HI^{m-k}(S, \phi^{-1}).$$

Proof. If f_ϕ is a Lyapunov function for (S, ϕ) then $-f_\phi$ is a Lyapunov function for (S, ϕ^{-1}) . Let N be an isolating neighborhood. Choosing a metric g , we get by Theorem 4.4.2 an isomorphism $HM_k(f_\phi, g, N) \rightarrow HM^{n-k}(-f_\phi, g, N)$. The Poincaré isomorphisms commute with the continuation isomorphisms hence combine to a Poincaré duality isomorphism as in the theorem. \square

Remark 4.5.2. Note that the orientation assumption can be simplified. An isolated invariant set on an orientable manifold is always orientable. Similarly a manifold is always orientable with $\mathbb{Z}/2\mathbb{Z}$ coefficients, hence for Morse-Conley-Floer homology with $\mathbb{Z}/2\mathbb{Z}$ coefficients the above duality statement always holds.

4.6 Maps in cohomology

The functorial constructions of Chapter 3 also hold for the cohomological counterparts. To spell this out, if $h^{\beta\alpha} \in \mathcal{T}(Q^\alpha, Q^\beta)$ this induces maps in Morse cohomology $(h^{\beta\alpha})^* : HM^*(Q^\beta) \rightarrow HM^*(Q^\alpha)$ via pullback. On generators of $C^*(Q^\beta)$ it is defined as

$$(h^{\beta\alpha})^* \eta_y := \eta_y h_{*}^{\beta\alpha} = \sum_{|x|=|y|} n_{h^{\beta\alpha}}(x, y) \eta_x.$$

Similarly if the map $h^{\beta\alpha}$ is isolating with respect to local Morse data there is an induced map in local Morse homology, and if the map is a flow map with $S^\alpha = (h^{\beta\alpha})^{-1}(S^\beta)$ there is a map in Morse-Conley-Floer cohomology

$$(h^{\beta\alpha})^* : HI^*(S^\beta, \phi^\beta) \rightarrow HI^*(S^\alpha, \phi^\alpha),$$

which is (contravariantly) functorial and independent of the isolated homotopy class.

4.6.1 Shriek maps in Morse homology

Poincaré duality allows us to define shriek maps. If M^α and M^β are orientable closed manifolds with Morse data Q^α and Q^β such that $h^{\beta\alpha} \in \mathcal{T}(\hat{Q}^\alpha, \hat{Q}^\beta)$ we

define the lower shriek maps

$$h_!^{\beta\alpha} : HM_k(\mathcal{Q}^\beta) \rightarrow HM_{m^\alpha - m^\beta + k}(\mathcal{Q}^\alpha)$$

by declaring the diagram

$$\begin{array}{ccc} HM_k(\mathcal{Q}^\beta) & \xrightarrow{h_!^{\beta\alpha}} & HM_{m^\alpha - m^\beta + k}(\mathcal{Q}^\beta) \\ PD_k \downarrow & & \uparrow PD_{m^\beta - k}^{-1} \\ HM^{m^\beta - k}(\hat{\mathcal{Q}}^\beta) & \xrightarrow{(h^{\beta\alpha})^*} & HM^{m^\beta - k}(\hat{\mathcal{Q}}^\alpha) \end{array}$$

to be commutative. Here the maps PD_k is the Poincaré duality isomorphisms of Theorem 4.3.3 on the manifold M^β and $PD_{m^\beta - k}$ is the Poincaré duality isomorphism on M^α . We decorate the notation for the stable/unstable manifold to $W^u(x; \mathcal{Q}^\alpha)$ to distinguish the unstable manifolds of \mathcal{Q}^α and $\hat{\mathcal{Q}}^\alpha$. We compute the following explicit formula for the lower shriek map on the chain level

$$\begin{aligned} h_!^{\beta\alpha} y &= PD_{m^\beta - k}^{-1} (h^{\beta\alpha})^* PD_k(y) \\ &= PD_{m^\beta - k}^{-1} \sum_{|x; \hat{\mathcal{Q}}^\alpha| = |y; \hat{\mathcal{Q}}^\beta|} \# \left(W^u(x, \hat{\mathcal{Q}}^\beta) \cap (h^{\beta\alpha})^{-1}(W^s(y, \hat{\mathcal{Q}}^\beta)) \right) \eta_x \quad (4.1) \\ &= \sum_{m^\alpha - |x; \mathcal{Q}^\alpha| = m^\beta - |y; \mathcal{Q}^\beta|} \# \left(W^s(x, \mathcal{Q}^\alpha) \cap (h^{\beta\alpha})^{-1}(W^u(y, \mathcal{Q}^\beta)) \right) x. \end{aligned}$$

Analogously we have a shriek map in Morse cohomology $(h^{\beta\alpha})^! : HM^k(\mathcal{Q}^\alpha) \rightarrow HM^{m^\beta - m^\alpha + k}(\mathcal{Q}^\beta)$ for closed orientable manifolds with $h^{\beta\alpha} \in \mathcal{T}(\hat{\mathcal{Q}}^\alpha, \hat{\mathcal{Q}}^\beta)$, by pre- and post composing the induced map in Morse homology by the appropriate Poincaré duality isomorphisms. The explicit formula is

$$(h^{\beta\alpha})^! \eta_x = \sum_{m^\beta - |y; \mathcal{Q}^\beta| = m^\alpha - |x; \mathcal{Q}^\alpha|} \# \left(W^s(x, \mathcal{Q}^\alpha) \cap (h^{\beta\alpha})^{-1}(W^u(y, \mathcal{Q}^\beta)) \right) \eta_y. \quad (4.2)$$

There are some minus signs missing depending on the degrees

Remark 4.6.1. Note that the sums run over those critical points such that $W^s(x) \cap (h^{\beta\alpha})^{-1}(W^u(y))$ is zero dimensional. The orientation assumptions on the ambient manifold are crucial to make this an oriented manifold. It is in

general a cooriented zero dimensional manifold. One cannot just define the wrong way maps by the above formula's if the manifolds are not orientable, because $W^s(x) \cap (h^{\beta\alpha})^{-1}(W^u(y))$ is *not* an oriented manifold. If the ambient manifolds are orientable this formula does work, because an orientation of $W^u(x)$ orients $W^s(x)$ and a choice of orientation of $W^u(y)$ also coorients itself.

4.7 Shriek maps in Morse-Conley-Floer homology

A flow map $h^{\beta\alpha}$ is also a flow map for the reversed flows. Therefore, just as in Morse homology, we obtain shriek maps

$$\begin{aligned} h_!^{\beta\alpha} &: HI_k(S^\beta, \phi^\beta) \rightarrow HI_{m^\alpha - m^\beta + k}(S^\alpha, \phi^\alpha), \\ (h^{\beta\alpha})^! &: HI^k(S^\alpha, \phi^\alpha) \rightarrow HI^{m^\beta - m^\alpha + k}(S^\beta, \phi^\beta). \end{aligned}$$

which on the Morse complexes of perturbed Lyapunov functions can be computed by the formulas of Equations 4.1 and 4.2, where the intersections are taken of the local stable and unstable manifolds of appropriate isolating neighborhoods.

4.8 Other algebraic structures

We define cup and intersection products. Cap products are more intricate in Morse-Conley-Floer homology. For example, the Poincaré duality isomorphism is not induced by capping with a fundamental class in Morse-Conley-Floer homology.

4.8.1 (Co)homology cross product

If \mathcal{Q}^α and \mathcal{Q}^β are Morse data, the product Morse datum $\mathcal{Q}^\gamma = \mathcal{Q}^\alpha \times \mathcal{Q}^\beta$ is the Morse datum on $M^\gamma = M^\alpha \times M^\beta$ with

$$\begin{aligned} f^\gamma(p, q) &:= f^\alpha(p) + f^\beta(q) \\ g^\gamma &:= g^\alpha \oplus g^\beta \end{aligned}$$

and the orientation $\mathfrak{o}^\gamma := \mathfrak{o}^\alpha \oplus \mathfrak{o}^\beta$ the natural one induced on $TW^u(x, y) = TW^u(x) \oplus TW^u(y)$. The homology cross product $\times : HM_k(\mathcal{Q}^\alpha) \otimes$

I do believe there is a way out here, one can define a fundamental class in $H_{m^\alpha}(N, \partial N)$ which in Morse terms would be generated by a coercive Morse function on $\text{int } N$

$HM_l(\mathcal{Q}^\beta) \rightarrow HM_{k+l}(\mathcal{Q}^\gamma)$ is defined on generators on the chain level by $x \otimes y \rightarrow (x, y)$. This product behaves well under continuation, hence works in local Morse homology as well. Noting that if S^α and S^β are isolated invariant sets of ϕ^α and ϕ^β , then $S^\alpha \times S^\beta$ is an isolated invariant set of the product flow and $f_{\phi^\alpha \times \phi^\beta}(p, q) = f_{\phi^\alpha}(p) + f_{\phi^\beta}(q)$ is a Lyapunov function. Hence the homology cross product is well defined in Morse-Conley-Floer homology. The same formulas work in the cohomology versions mutatis mutandis.

4.8.2 Cup and intersection products

Let $\Delta : M^\alpha \rightarrow M^\alpha \times M^\alpha$ be the diagonal embedding. In singular homology the cup product in Morse homology is defined as the composition

$$H^k(M^\alpha) \otimes H^l(M^\alpha) \xrightarrow{\times} H^{k+l}(M^\alpha \times M^\alpha) \xrightarrow{\Delta^*} H^{k+l}(M^\alpha)$$

Choosing Morse data $\mathcal{Q}^\alpha, \mathcal{Q}^\beta, \mathcal{Q}^\gamma$ on M^α such that $\Delta \in \mathcal{T}(\mathcal{Q}^\alpha, \mathcal{Q}^\beta \times \mathcal{Q}^\gamma)$ we have the following formula on the chain level

$$\eta_x \cup \eta_y = \sum_{|z|=|x|+|y|} \# \left(W^u(z; \mathcal{Q}^\alpha) \cap \Delta^{-1}(W^s((x, y); \mathcal{Q}^\beta \times \mathcal{Q}^\gamma)) \right) \eta_z.$$

Identifying $W^s((x, y); \mathcal{Q}^\beta \times \mathcal{Q}^\gamma) \cong W^s(x, \mathcal{Q}^\beta) \times W^s(y, \mathcal{Q}^\gamma)$, we can also write this as

$$\eta_x \cup \eta_y = \sum_{|z|=|x|+|y|} \# \left(W^u(z; \mathcal{Q}^\alpha) \cap W^s(x; \mathcal{Q}^\beta) \cap W^s(y, \mathcal{Q}^\gamma) \right) \eta_z.$$

Note that the triple oriented intersection number is well defined. By the transversality assumptions $W^u(z; \mathcal{Q}^\alpha) \pitchfork W^s(x; \mathcal{Q}^\beta)$ and $(W^u(z; \mathcal{Q}^\alpha) \cap W^s(x; \mathcal{Q}^\beta)) \pitchfork W^s(y; \mathcal{Q}^\gamma)$. Moreover $W^u(z; \mathcal{Q}^\alpha) \cap W^s(x; \mathcal{Q}^\beta)$ is oriented as it is the intersection of an oriented and a cooriented manifold. By the same argument the triple intersection is an oriented manifold as well. If M^α is oriented we also have intersection products by the composition

$$HM_k(\mathcal{Q}^\beta) \otimes HM_l(\mathcal{Q}^\gamma) \xrightarrow{\times} HM^{k+l}(\mathcal{Q}^\beta \times \mathcal{Q}^\gamma) \xrightarrow{\Delta!} HM_{k+l}(\mathcal{Q}^\alpha)$$

when $\Delta \in \mathcal{T}(\widehat{\mathcal{Q}^\alpha}, \widehat{\mathcal{Q}^\beta \times \mathcal{Q}^\gamma})$. Again on the chain level

$$x \cdot y = \sum \# \left(W^u(x; \mathcal{Q}^\beta) \cap W^u(\mathcal{Q}^\gamma) \cap W^s(z; \mathcal{Q}^\alpha) \right) z$$

Remark 4.8.1. Note that the orientation assumption on M^α is here, as we are intersecting (transversally) two oriented manifolds and a cooriented manifold which need not be oriented. Consider for example the intersection of two $\mathbb{R}P^5$'s and an $\mathbb{R}P^4$ in general position in $\mathbb{R}P^6$.

4.9 Some identities

Theorem 4.9.1 (Projection formula). *Let $h^{\beta\alpha} : M^\alpha \rightarrow M^\beta$ be a flow map, and $S^\alpha = (h^{\beta\alpha})^{-1}(S^\beta)$. Then for $x \in H^*(S^\alpha, \phi^\alpha)$ and $y \in H^*(S^\beta, \phi^\beta)$ we have the projection formula*

$$(h^{\beta\alpha})^!(x \cup (h^{\beta\alpha})^*(y)) = (h^{\beta\alpha})^!(x) \cup y.$$

If S^α and S^β are oriented, then we have the dual formula in Morse-Conley-Floer homology. That is for $x \in HI^(S^\alpha)$ and $y \in HI^*(S^\beta)$*

check this

$$(h^{\beta\alpha})_*(x \cdot (h_1^{\beta\alpha})(y)) = h_1^{\beta\alpha}(x) \cdot y.$$

4.10 Coincidence theory

Theorem 4.10.1. *Suppose that $S^\alpha = (h^{\beta\alpha})^{-1}(S^\beta)$ is an attractor.*

add

Chapter 5

A Spectral Sequence in Morse-Conley-Floer homology.

This paper, whose intent is stated in its title, gives wrong solutions to trivial problems. The basic error, however, is not new.

C. Truesdell

The existence of Lyapunov functions for flows gives rise to a spectral sequence in Morse-Conley-Floer homology. The spectral sequence is well known in the case that the flow is the gradient flow of a Morse-Bott function, cf. [Fuk96, AB95, Jia99]. We show how this spectral sequence is recovered, and also show how this spectral sequence can both be applied to general flows, as well as more degenerate variational problems. We do not aim to be overly general and the results of this chapter can be improved without much effort. In essence the ideas are those of [Jia99]¹ spelled out in the language of Morse-Conley-Floer homology. The Morse-Bott inequalities are an immediate corollary. It should be noted that some of the literature [Jia99, Hur13, BH04, BH09] contains some misstatements concerning subtle orientation issues, however earlier papers contain correct versions of the theorem [Bot54, AB95, Bis86, Nic11]. In Section 5.4 we give a counterexample to the misstatements.

¹There are some subtle orientation issues with this paper which we discuss in Section 5.4

5.1 The birth of a spectral sequence

Let ϕ be a flow and assume $f_\phi : M \rightarrow \mathbb{R}$ is a smooth function on a closed manifold M . Let (b_l^-, b_l^+) be a finite set of numbers such that

$$b_1^- < b_1^+ < b_2^- < b_2^+ < \dots < b_{n-1}^- < b_{n-1}^+ < b_n^- < b_n^+.$$

Set $N_l = f_\phi^{-1}([b_l^-, b_l^+])$, and assume that these are isolating neighborhoods of the flow ϕ and assume that f_ϕ is a Lyapunov function for these isolating neighborhoods. Assume that all critical values of f_ϕ lie in $\bigcup_{l=1}^n (b_l^-, b_l^+)$. It follows that $b_1^- < \min f_\phi$ and $b_n^+ > \max f_\phi$. By the compactness of M such a decomposition always exist, for instance for $n = 1$. The point is that a finer decomposition of the critical values will give more information in the spectral sequence that is to come. Two cases are of interest: f_ϕ is a Lyapunov function for a Morse decomposition of M , or ϕ is the gradient flow of the –possibly degenerate– function f_ϕ . Now we are able to state the main theorem of this chapter:

Theorem 5.1.1. *There exist a spectral sequence $(E_{k,l}^r, \partial^r)$ such that*

$$E_{k,l}^1 = HI_k(S_l, \phi),$$

and $E_{k,l}^r \Rightarrow H_*(M)$. That is, for each k

$$\bigoplus_l E_{k,l}^\infty \cong H_k(M).$$

Proof. Set $A = \bigcup_l N_l$. By assumption f_ϕ has no critical points in the compact set $\overline{M \setminus A}$, thus $\epsilon = \inf_{p \in \overline{M \setminus A}} |\nabla_g f(p)|_g > 0$.

It follows that each function $f^\alpha : M \rightarrow \mathbb{R}$, with $|f^\alpha - f_\phi|_{C^1} < \epsilon$ does not have critical points outside $f_\phi^{-1}(A)$. The set $\mathcal{S}_{MS}(f_\phi, N_l)$ is generic, cf. Proposition 2.3.9. Hence there exists an

$$(f^\alpha, g^\alpha) \in \bigcap_l \mathcal{S}_{MS}(f_\phi, N_l) \bigcap \mathcal{S}_{MS}(f_\phi, M)$$

with $|f^\alpha - f_\phi|_{C^1} < \epsilon$ without any critical points outside $f_\phi^{-1}(A)$. Let \mathfrak{o}^α be a choice of orientation of the unstable manifolds. We write $\mathcal{Q}^\alpha = (f^\alpha, g^\alpha, \mathfrak{o}^\alpha)$.

By assumption f^α is a Lyapunov function for each S_l . We can filter the critical points of f^α by the value of f^α

$$C_{k,l}(\mathcal{Q}^\alpha) = \mathbb{Z}\{x \in \text{Crit } f^\alpha \mid |x| \leq k, \quad f^\alpha(x) \leq b_l^+\}.$$

This is a filtration of the chain complex $C_*(f^\alpha)$ generated by the critical points of f^α . Thus $C_{k,l-1}(\mathcal{Q}^\alpha) \subset C_{k,l}(\mathcal{Q}^\alpha)$ for all l and the differential respects the filtration: $\partial C_{k,l}(\mathcal{Q}^\alpha) \subset C_{k-1,l}(\mathcal{Q}^\alpha)$. There are short exact sequences

$$0 \rightarrow C_{k,l-1}(\mathcal{Q}^\alpha) \rightarrow C_{k,l}(\mathcal{Q}^\alpha) \rightarrow C_{k,l}(\mathcal{Q}^\alpha)/C_{k,l-1}(\mathcal{Q}^\alpha) \rightarrow 0,$$

of chain groups, if we define the differential $\partial_k : C_{k,l}(\mathcal{Q}^\alpha)/C_{k,l-1}(\mathcal{Q}^\alpha) \rightarrow C_{k-1,l}/C_{k-1,l-1}$ by passing to the quotient in the natural way. By the Snake Lemma this induces long exact sequences of homology groups²

$$\dots \rightarrow H_k(C_{*,l-1}(\mathcal{Q}^\alpha)) \xrightarrow{i} H_k(C_{*,l}(\mathcal{Q}^\alpha)) \xrightarrow{j} H_k(C_{*,l}(\mathcal{Q}^\alpha), C_{*,l-1}(\mathcal{Q}^\alpha)) \xrightarrow{k} \dots$$

Summing over the filtration degree l and the grading k gives an exact couple, and hence a spectral sequence. To spell this process out, cf. [BT82, McC01], we set

$$A^1(\mathcal{Q}^\alpha) = \bigoplus_l \bigoplus_k H_k(C_{*,l-1}(\mathcal{Q}^\alpha)),$$

$$E^1(\mathcal{Q}^\alpha) = \bigoplus_l \bigoplus_k H_k(C_{*,l}(\mathcal{Q}^\alpha), C_{*,l-1}(\mathcal{Q}^\alpha)).$$

We get the exact couple

$$\begin{array}{ccc} A^1(\mathcal{Q}^\alpha) & \xrightarrow{i^1} & A^1(\mathcal{Q}^\alpha) \\ & \swarrow k^1 & \searrow j^1 \\ & E^1(\mathcal{Q}^\alpha) & \end{array}$$

This induces a differential $\partial^1 = j^1 \circ k^1$ on $E^1(\mathcal{Q}^\alpha)$. Setting $A^2(\mathcal{Q}^\alpha) = i^1(A^1(\mathcal{Q}^\alpha))$ and $E^2(\mathcal{Q}^\alpha) = H(E^1(\mathcal{Q}^\alpha), d^1)$ we get the derived exact couple

$$\begin{array}{ccc} A^2(\mathcal{Q}^\alpha) & \xrightarrow{i^2} & A^2(\mathcal{Q}^\alpha) \\ & \swarrow k^2 & \searrow j^2 \\ & E^2(\mathcal{Q}^\alpha) & \end{array}$$

²We suppress the dependence of the differentials in the notation

where the maps are defined by $i^2(i^1(a)) := i^1i^1(a)$ and $j^2(i^1(a)) = [j(a)]$. The map k^2 is defined via some diagram chasing. For $[b] \in E^2(\mathcal{Q}^\alpha)$ we have that $\partial^1 b = j^1 k^1 b = 0$, which implies that $k^1 b \in \ker j^1$, thus $k^1(b) \in \text{im}(i^1)$ which implies $k^1(b) = i^1(a)$ for some a . We set $k^2[b] = i^1(a)$. An elementary diagram chase shows that this is well-defined and again an exact couple. In this way get the spectral sequence $(E^r(\mathcal{Q}^\alpha), \partial^r(\mathcal{Q}^\alpha))$. We observe that $E_{k,l}^1 = H_k(C_{*,l}(\mathcal{Q}^\alpha), C_{*,l-1}(\mathcal{Q}^\alpha)) = HM_k(f^\alpha, g^\alpha, \sigma^\alpha, N_l) \cong HI(S_l, \phi)$, and the homology of the original complex is known by

$$HI_*(M, \phi) \cong H_*(C_*(f^\alpha)) \cong HM_*(f^\alpha, g^\alpha, \sigma^\alpha) \cong H_*(M).$$

The spectral sequence collapses after at most $\min(m, n)$ steps with $\bigoplus_l E_{k,l}^\infty \cong H_k(M)$.

Moreover, by the naturality of the relative long exact sequences, if $\mathcal{Q}^\beta = (f^\beta, g^\beta, \sigma^\beta)$ is another choice of datum as above, there is a commutative diagram

$$\begin{array}{ccccccc} \dots & \rightarrow & H_*(C_{*,l-1}(\mathcal{Q}^\alpha)) & \xrightarrow{i} & H_*(C_{*,l}(\mathcal{Q}^\alpha)) & \xrightarrow{j} & H_*(C_{*,l}(\mathcal{Q}^\alpha), C_{*,l-1}(\mathcal{Q}^\alpha)) \xrightarrow{k} \dots \\ & & \downarrow \Phi_*^{\beta\alpha} & & \downarrow \Phi_*^{\beta\alpha} & & \downarrow \Phi_*^{\beta\alpha} \\ \dots & \rightarrow & H_*(C_{*,l-1}(\mathcal{Q}^\alpha)) & \xrightarrow{i} & H_*(C_{*,l}(\mathcal{Q}^\alpha)) & \xrightarrow{j} & H_*(C_{*,l}(\mathcal{Q}^\alpha), C_{*,l-1}(\mathcal{Q}^\alpha)) \xrightarrow{k} \dots \end{array}$$

Where the vertical maps are induced by continuation as usual. This implies that there is a spectral sequence for the Morse-Conley-Floer homology $E_{k,l}^r = \varprojlim E_{k,l}^r(\mathcal{Q}^\alpha)$, with $E_{k,l}^1 = HI_k(S_l, \phi)$ and $\bigoplus_l E_{k,l}^\infty \cong H_k(M)$. \square

5.2 The Morse-Conley relations via the spectral sequence

The existence of the spectral sequence gives the Morse-Conley relations of Theorem 2.6.4, as we prove now. We have the following elementary lemma.

Lemma 5.2.1. *Suppose C_k is a chain complex, finitely generated in each degree, with $C_k = 0$ for $k < 0$. Denote by H_k the homology. Then there is an equality of formal polynomials*

$$\sum \text{rank } C_k t^k = \sum \text{rank } H_k t^k + (1+t)Q_t$$

where Q_t has non-negative coefficients.

Proof. We have short exact sequences

$$0 \rightarrow \ker \partial_k \rightarrow C_k \rightarrow \text{im } \partial_k \rightarrow 0,$$

and

$$0 \rightarrow \text{im } \partial_{k+1} \rightarrow \ker \partial_k \rightarrow H_k \rightarrow 0.$$

from which it follows that $\text{rank } C_k = \text{rank im } \partial_k + \text{rank ker } \partial_k$ and $\text{rank ker } \partial_k = \text{rank } H_k + \text{rank im } \partial_{k+1}$. We compute

$$\begin{aligned} \sum_{k=0} (\text{rank } C_k - \text{rank } H_k) t^k &= \sum_{k=0} (\text{rank im } \partial_k + \text{rank ker } \partial_k \\ &\quad - \text{rank ker } \partial_k + \text{rank im } \partial_{k+1}) t^k \\ &= (1+t) \sum_{k=0} \text{rank im } \partial_{k+1} t^k \end{aligned}$$

where we used the fact that $\text{rank im } \partial_0 = 0$. □

A Lyapunov function for a Morse decomposition $\{S_l\}$ is a function that decreases along the flow outside the Morse sets S_l , and is locally constant on the Morse sets. Given a Morse decomposition $\{S_l\}$ of M , we can coarsen the partial order by introducing extra relations. We obtain a linearly ordered Morse decomposition. A Lyapunov function f_ϕ associated to such a linear order, satisfies the assumptions of Section 5.1. The existence of the spectral sequence in Theorem 5.1.1 a reinterpretation of the Morse-Conley relations.

Of Morse-Conley relations of Theorem 2.6.4. Write $P_t(E^r) = \sum_k \text{rank } \bigoplus_l E_{k,l}^r t^k$. Then we get, successively applying the previous lemma, using the fact that each page is the homology of the previous one

$$\begin{aligned} \sum_l P_t(S_l) &= P_t(E^1) \\ &= P_t(E^2) + (1+t)Q_t^1 \\ &= P_t(E^3) + (1+t)(Q_t^1 + Q_t^2) \\ &= \dots \\ &= P_t(E^\infty) + (1+t)Q_t^\infty, \\ &= P_t(M) + (1+t)Q_t^\infty, \end{aligned}$$

where Q_t^∞ is the sum of all the Q_t^i . □

Remark 5.2.2. In the construction of the spectral sequences we have thrown away some information. A Morse decomposition of a flow is ordered by an arbitrary poset, not necessarily a linear one. Generically, choosing a Lyapunov function for the Morse decomposition *coarsens* the order by introducing extra relations. For each such coarsening we obtain a spectral sequence, and we have maps between the spectral sequences defined through continuation. Because these coarsenings must always respect the original partial order we know a lot a priori about the spectral sequences. This seems to be related to the notion of a spectral system [Mat13], which arises when a chain complex is filtered by an arbitrary poset, and warrants further study. The relation with the connection matrix, cf. [Fra89], also requires investigation. In the case of linearly ordered Morse decompositions some work has been done, cf. [Bar07, Bar09, Bar05].

5.3 Interpretation of the spectral sequence for Morse-Bott functions

The first page of the spectral sequence in Theorem 5.1.1 can be computed explicitly in terms of critical point data for Morse-Bott functions, satisfying an orientability requirement.

Morse-Bott functions admit larger critical sets than isolated critical points. Morse-Bott functions appear in situations with symmetry, for example Lie groups acting on manifolds. They still have good critical sets, in the sense that the critical sets are “Morse in the normal direction”. Recall that the Hessian of a function on a Riemannian manifold is defined via $\text{Hess}_g f(V, W) = g(\nabla_V \nabla_g f, W)$. At a critical point of f this is independent of the choice of metric, since the difference of two connections is a tensor. If B is a connected component of $\text{Crit } f$ that is a submanifold, the Hessian descends to the normal bundle NB of B inside M . We call such a connected component a critical manifold. Denote the Hessian in the normal direction by $\text{Hess}^N f : NB \times NB \rightarrow \mathbb{R}$. A function $f_\phi : M \rightarrow \mathbb{R}$ is a *Morse-Bott function* if all connected components of $\text{Crit } f_\phi$ are submanifolds, and the Hessians in the normal directions $\text{Hess}^N f_\phi|_B$ are non-degenerate. Then there are splittings $NB = N^-B \oplus N^+B$ into the negative N^-B and positive N^+B eigenspaces of Hess^N with respect to the metric. The number $|B| := \dim N^-B$ is the *Morse-Bott index* of B , which does not depend on the choice of metric.

Proposition 5.3.1. *Let B be a critical manifold of a Morse-Bott function. Suppose*

5.3. Interpretation of the spectral sequence for Morse-Bott functions

that the negative normal bundle $N^-B \rightarrow B$ is orientable. Then

$$HI_k(B, \phi) \cong H_{k-|B|}(B).$$

Proof. The proof is basically the Morse homological interpretation of the Thom isomorphism, cf. [AS10, CS09], on the bundle where Hess^N is negative definite. We first consider a slightly more general construction. Consider a vector bundle $E \rightarrow B$. A metric on E gives a horizontal-vertical splitting:

$$T_p E \cong T_p^h E \oplus T_p^v E \cong T_{\pi(p)} B \oplus \pi^{-1}(\pi(p)) \quad (5.1)$$

via the connection induced by the metric. Assume now that $E = F \oplus G$ is the direct sum of two vector bundles with F orientable, and write $\pi_F : E \rightarrow F$ and $\pi_G : E \rightarrow G$ for the projections. Choose a metric h on E compatible with the splitting. Fix a Morse-Smale pair $(f_B, g_B, \mathfrak{o}_B)$ on the base B . Then

$$f(x) = f_B(\pi(x)) - |\pi_F(x)|_h^2 + |\pi_G(x)|_h^2$$

is seen to be a Morse function. Moreover the critical points are located at the zero section and we identify

$$C_k(f_B) \cong C_{k+\dim F}(f). \quad (5.2)$$

Since F is assumed to be orientable, we can choose a coherent orientation \mathfrak{o}_F for the fibers of $F_{\pi(x)}$. We orient all unstable manifolds of critical points of f via $\mathfrak{o}_B \pi(x) \oplus \mathfrak{o}_F$. Gradient flow lines of f project to gradient flowlines of f_B , cf. (5.1). The differential ∂_E agrees with the differential ∂_B under the identification (5.2). Hence we get isomorphisms

$$HM_k(f_B, g_B, \mathfrak{o}_B) \cong HM_{k+\dim F}(f, g, \mathfrak{o}, N).$$

where N is any isolating neighborhood of the zero section B . Now we return to our situation. Let f_ϕ be a Morse function with critical manifold B , and $E = NB$ and $F = NB^-$ and $F = NB^+$. The Morse-Bott lemma, cf. [BH09, Lemma 5], states that f_ϕ is locally of the form $f = f_\phi|_B - |\pi_F(x)|^2 + |\pi_G(x)|^2$. For a Morse function f_B, f , and $\epsilon > 0$ small isolated homotopic to

$$f = f_\phi(0) + \epsilon f_B(\pi(x)) - |\pi_F(x)|_h^2 + |\pi_G(x)|_h^2,$$

and by the previous argument $HI_k(B, \phi) \cong HM_k(f, g, \mathfrak{o}, N) \cong HM_{k-|B|}(f_B, g_B, \mathfrak{o}_B) \cong HM_{k-|B|}(B)$. □

If f_ϕ is a Morse-Smale function whose critical manifolds have orientable negative bundles, and such that the values of the critical manifolds are distinct, we get a spectral sequence converging to the homology of M , whose E^1 page is given by the homology groups $H_{k-|B_i|}(B_i)$. An more elementary expression of this fact are the Morse-Bott relations. For a Morse-Bott function with critical manifolds B_i , we define the Morse-Bott polynomial by $MB_t(f) := \sum_i P_t(B_i)t^{|B_i|}$. Theorem 2.6.4, which follows from the spectral sequence and Proposition 5.3.1, now gives the Morse-Bott relations.

Theorem 5.3.2 (Morse-Bott relations). *Let $f : M \rightarrow \mathbb{R}$ be a Morse Bott function on a closed manifold, all of whose critical manifolds have orientable negative normal bundle NB^- . Then*

$$MB_t(f) = P_t(M) + (1 + t)Q_t,$$

for some polynomial Q_t with non-negative coefficients.

5.4 A counter example to misstated Morse-Bott inequalities.

Orientation assumptions for the Morse-Bott relations are necessary. The assumptions in Theorem 5.3.2 are sufficient. In [Hur13, BH04, BH09] the following theorem is stated and in [Jia99] no orientation assumptions are made.

Theorem 5.4.1 (False Morse-Bott relations). *Let $f : M \rightarrow \mathbb{R}$ be a Morse-Bott function on a closed orientable manifold. Assume that all critical manifolds are orientable. Then*

$$MB_t(f) = P_t(M) + (1 + t)Q_t$$

where Q_t has non-negative coefficients.

We have the following counterexample.

Proposition 5.4.2. *The function $f : \mathbb{R}P^5 \rightarrow \mathbb{R}$ given in homogeneous coordinates by*

$$f([x_0, \dots, x_5]) = \frac{-x_4^2 + x_5^2}{x_0^2 + \dots + x_5^2}$$

is Morse-Bott with orientable critical manifolds. The Morse-Bott inequalities of Theorem 5.4.1 are violated.

Proof. Note that $\mathbb{R}P^{2n+1}$ is orientable. The function f is Morse-Bott with three critical manifolds: one minimum at $[0, 0, 0, 0, 1, 0]$, one maximum at $[0, 0, 0, 0, 0, 1]$ and all other critical points form an $\mathbb{R}P^3 =$

5.4. A counter example to misstated Morse-Bott inequalities.

$\{[x_0, x_1, x_2, x_3, 0, 0] \mid x_i \in \mathbb{R}\}$ of Morse-Bott index 1. The Morse-Bott polynomial of f with \mathbb{Z} coefficients is

$$MB_t(f) = P_t(\text{pt})t^0 + P_t(\mathbb{R}\mathbb{P}^3)t^1 + P_t(\text{pt})t^5.$$

Using the fact that $P_t(\mathbb{R}\mathbb{P}^{2n+1}) = 1 + t^{2n+1}$ the Morse-Bott inequalities state that

$$1 + t + t^4 + t^5 = 1 + t^5 + (1 + t)Q_t$$

for a polynomial Q_t with non-negative coefficients, which is absurd. \square

Of course the Morse-Bott polynomial of f in the previous proposition with $\mathbb{Z}/2\mathbb{Z}$ coefficients does satisfy the Morse-Bott inequalities. In this case

$$P_t(\mathbb{R}\mathbb{P}^{2n+1}, \mathbb{Z}/2\mathbb{Z}) = 1 + t + \dots + t^{2n} + t^{2n+1}$$

and

$$\begin{aligned} MB_t(f) &= P_t(\min, \mathbb{Z}/2\mathbb{Z})t^0 + P_t(\mathbb{R}\mathbb{P}^3, \mathbb{Z}/2\mathbb{Z})t^1 + P_t(\max, \mathbb{Z}/2\mathbb{Z})t^5 \\ &= 1 + t + t^2 + t^3 + t^4 + t^5, \end{aligned}$$

hence $Q_t = 0$. One can also define Morse homology with local coefficients [KM07] through which it is possible to obtain a more refined version of the Morse-Bott inequalities. This is also the statement that appears in Bott's original paper [Bot54], and in Bismut's paper [Bis86].

Remark 5.4.3. The orientability assumption of Theorem 5.4.1 probably comes from the following exact sequence for the normal bundle of a critical manifold

$$0 \rightarrow TB \rightarrow TM \rightarrow NB \rightarrow 0.$$

Assuming B and M are oriented, this exact sequence orients the normal bundle NB . Unfortunately this does not orient neither NB^- or NB^+ , as can be seen from the example above, or more explicitly for an embedded open Moebius strip in \mathbb{R}^3 .




Figure 5.1: A sketch of the counterexample of Proposition 5.4.2. We depicted $\mathbb{R}P^5$ as a solid ball, where the opposite endpoints of the boundary sphere are identified. The function has three critical manifolds. Two points, the maximum and the minimum, and a critical $B \cong \mathbb{P}^3$, which is the vertical line in the middle. The negative N^-B bundle and the positive bundle N^+B are one-dimensional and are not orientable. The normal bundle $NB = N^-B \oplus N^+B$ is.

Part II

Periodic Orbits on Non-Compact Hypersurfaces

Chapter 6

Periodic orbits in cotangent bundles of non-compact manifolds

Please don't ask me anything
about this paper. I write to forget.

Michael Hutchings

6.1 Introduction

The question of existence of periodic orbits of a Hamiltonian vector field X_H on a given regular energy level, i.e. a level set $\Sigma = H^{-1}(0)$ of the Hamiltonian function H , with $dH \neq 0$ on Σ , has been a central question in Hamiltonian dynamics and symplectic topology which has generated some of the most interesting recent developments in those areas, see Section 1.10 and [Pas12] for a discussion. Currently not much is known about the existence of periodic orbits on *non-compact* energy hypersurfaces. In [vdBPV09] the authors proved an existence result for periodic orbits on non-compact hypersurfaces in \mathbb{R}^{2n} of mechanical hamiltonians. In this chapter we follow their method of proof and generalize this existence result to a class of cotangent bundles.

6.1.1 Main result

A Riemannian manifold is said to have *flat ends* if the curvature tensor vanishes outside a compact set. The main theorem of this paper is the following existence result.

Theorem 6.1.1. *Let $H : T^*M \rightarrow \mathbb{R}$ be a mechanical Hamiltonian, i.e. $H(q, \theta_q) = \frac{1}{2}g_q^*(\theta_q, \theta_q) + V(q)$ and assume*

- (i) *(M, g) is an n -dimensional complete orientable Riemannian manifold with flat ends*
- (ii) *The hypersurface $\Sigma = H^{-1}(0)$ is regular, i.e. $dH \neq 0$ on Σ*
- (iii) *The potential V is asymptotically regular, i.e. there exist a compact K and constant $V_\infty > 0$ such that*

$$|\text{grad } V(q)| \geq V_\infty, \text{ for } q \in M \setminus K \text{ and } \frac{\|\text{Hess } V(q)\|}{|\text{grad } V(q)|} \rightarrow 0, \text{ as } d(q, K) \rightarrow \infty.$$

- (iv) *Assume moreover that there exists an integer $0 \leq k \leq n - 1$ such that*
 - $H_{k+1}(\Lambda M) = 0$ and $H_{k+2}(\Lambda M) = 0$, and
 - $H_{k+n}(\Sigma) \neq 0$.

*Then Σ has a periodic orbit which is contractible in T^*M .*

Here ΛM denotes the free loop space of H^1 loops into M . The proof of Theorem 6.1.1 follows the scheme of the proof of the existence result for non-compact hypersurfaces in \mathbb{R}^{2n} presented in [vdBPV09]. We regard periodic orbits as critical points of a suitable action functional \mathcal{A} . The functional does not satisfy the Palais-Smale condition. Therefore we introduce a sequence of approximating functionals \mathcal{A}_ϵ , for $\epsilon > 0$, which do satisfy the Palais-Smale condition. Critical points of \mathcal{A}_ϵ satisfying certain bounds converge to critical points of \mathcal{A} as $\epsilon \rightarrow 0$. Next, based on the assumptions on the topology of Σ and M , we construct linking sets in M and lift these to linking sets in the free loop space, where we apply a linking argument to produce critical points for the approximating functionals satisfying the appropriate bounds. These critical points then converge to a critical point of \mathcal{A} as $\epsilon \rightarrow 0$. Because we construct the linking sets in the component of the loop space containing the contractible loops, this critical point corresponds to a contractible loop.




Figure 6.1: The manifold M viewed as two sheets of \mathbb{R}^6 connected by a tube $S^5 \times [0, 1]$. The projection N of Σ to the base manifold M is shown in red. In the picture ∂N is depicted as two copies of S^0 , but is diffeomorphic to two copies of S^1 .

One of the main issues in cotangent bundles (as opposed to \mathbb{R}^{2n}) is that the linking arguments get more involved due to the topology of M . Another difficulty is that curvature terms appear in the analysis of the functional, which require some care.

Theorem 6.1.1 directly generalizes the results of [vdBPV09]. In [vdBPV09] examples are given that show that both topological and geometric assumptions on Σ are necessary. Theorem 6.1.1 also improves the result in the \mathbb{R}^{2n} case, as it requires weaker assumptions on the metric.

Recently a preprint by Suhr and Zehmisch [SZ13] appeared where it is shown that some technical hypotheses of Theorem 6.1.1 can be removed. In particular the estimate in Lemma 6.4.7 has been shown to hold under the weaker assumption of bounded geometry of M , instead of flat ends. Moreover due to a different minimax construction, no assumptions on the homology of the loop space are necessary, see also the example below.

6.1.2 Examples

We give an example of a manifold with non-trivial topology that satisfies the conditions of Theorem 6.1.1. Let $M = \mathbb{R}^6 - \{\text{pt}\} \cong S^5 \times \mathbb{R}$. Let (ϕ, r) be coordinates on $S^5 \times \mathbb{R}$, and define the metric on M via: $g_M = f(r)g_{S^5} + dr^2$, with f equal to r^2 outside $[-1, 1]$ and positive everywhere, and g_{S^5} is the round

metric on the 5-sphere. This manifold is asymptotically conformally flat and admits an asymptotically flat and complete representative via g_M itself. This manifold can also be visualized as two copies of \mathbb{R}^6 , with a ball of radius 1 cut out. These are then connected via a tube $\mathbb{S}^5 \times [0, 1]$, cf. Figure 6.1.2. In these coordinates we take the potential function

$$V(q) = \frac{\rho(q)}{2}(q_1^2 + q_2^2 - q_3^2 - q_4^2 - q_5^2 - q_6^2) - C,$$

with ρ equal to zero within the ball of radius 1 and equal to 1 outside the ball of radius 2. For $C > 0$ chosen large enough, we can study the topology of Σ . The topology of the projection of Σ to the base manifold, and its boundary

$$N = \pi(\Sigma) = \{q \in M \mid V(q) \leq 0\}, \quad \text{and} \quad \partial N = \{q \in M \mid V(q) = 0\},$$

are easily visualized, N deformation retracts to \mathbb{S}^5 , and the boundary ∂N deformation retracts to two disjoint circles \mathbb{S}^1 . From this, using Proposition 6.7.2, we can compute the the homology of Σ . From the long exact sequence of the pair we read off (for $k = 1$) that $H_2(N, \partial N) = \mathbb{Z}^2$, and hence $H_7(\Sigma) \neq 0$. The free loop space of M is homotopic to the free loop space of \mathbb{S}^5 . This is well studied, and from the path and loop space fibrations, and a spectral sequence argument, we see that $H_2(\Lambda M) = H_3(\Lambda M) = 0$. Thus Σ contains a contractible closed characteristic.

We now give an example that show that some assumptions are too strong, see also [SZ13]. Consider the question of existence of periodic orbits of the Hamiltonian system on the cylinder $M = \mathbb{S}^1 \times \mathbb{R}$, with an asymptotically regular potential V , which is negative on N , a set homeomorphic to $\mathbb{S}^1 \times [0, 1] \cup [0, \frac{1}{2}] \times \mathbb{R}_{\geq 0}$, where the corners are sufficiently smoothed, cf. Figure 6.1.2. The boundary ∂N consists of two components, one diffeomorphic to \mathbb{R} and one component diffeomorphic to \mathbb{S}^1 . Using Proposition 6.7.2 we find that $H_{k+n}(\Sigma) \neq 0$, for $k = 0$. The free loop space of M is homotopic to $\mathbb{S}^1 \times \mathbb{Z}$, and we see that $H_1(\Lambda M) \neq 0$. It seems Theorem 6.1.1 is not applicable. However it is not hard to find an embedding of Σ into $T^*\mathbb{R}^2$, cf. Figure 6.1.2, by constructing an appropriate asymptotically regular potential, which shows the existence of closed characteristics on Σ .

6.2 Contact type conditions

In this section we show that mechanical hypersurfaces are always of contact type. We give two proofs. One proof does not provide an explicit contact




Figure 6.2: Admissible embeddings of N in two different manifolds M . On the left hand side we embed Σ in T^*M with $M = \mathbb{S}^1 \times \mathbb{R}$, and on the right hand we embed Σ in $M = \mathbb{R}^2$.

form. If the the potential is asymptotically regular, it is possible to write down an explicit contact form which has good asymptotic properties at infinity. This is the content of Proposition 6.2.3.

Let us fix some notation. We will use vectors and covectors on the base manifold, as well as vectors and covectors on the cotangent bundle. In an attempt to reduce possible confusion, we denote elements of these bundles by different fonts. We use lowercase roman for vectors, thus a vector on M is denoted by $x = (q, x_q) \in TM$, where $x_q \in T_qM$. A covector on M is denoted in lowercase greek, i.e. $\theta = (q, \theta_q) \in T^*M$ with $\theta_q \in T_q^*M$. Vectors on the cotangent bundle, e.g. the Hamiltonian vector field X_H , are denoted in uppercase roman. In uppercase greek we denote covectors on T^*M , e.g. the Liouville form Λ . The Riemannian metric g on M gives the musical isomorphism $\flat : TM \rightarrow T^*M$, which is defined by

$$\flat(q, x_q) = (q, g_q(x_q, \cdot)),$$

where $g_q(x_q, \cdot) \in T_q^*M$. Its inverse is $\sharp : T^*M \rightarrow TM$. We also write \flat and \sharp for the fiber-wise isomorphisms. The Riemannian metric g induces a non-degenerate symmetric bilinear form g^* on T^*M via

$$g_q^*(\theta_q, \theta'_q) = g_q(x_q, x'_q),$$

where $(q, x_q) = \sharp(q, \theta_q)$ and $(q, x'_q) = \sharp(q, \theta'_q)$. By abuse of notation, both for $(q, x_q), (q, x'_q) \in TM$ and for $(q, \theta_q), (q, \theta'_q) \in T^*M$ we write

$$\begin{aligned} |x_q|^2 &= g_q(x_q, x_q), & \langle x_q, x'_q \rangle &= g_q(x_q, x'_q) \\ |\theta_q|^2 &= g_q^*(\theta_q, \theta_q), & \langle \theta_q, \theta'_q \rangle &= g_q^*(\theta_q, \theta'_q). \end{aligned}$$

We need some standard constructions in symplectic geometry, which can be found for example in Hofer and Zehnder [HZ11]. The tangent map of the bundle projection $\pi : T^*M \rightarrow M$ is a mapping $T\pi : TT^*M \rightarrow TM$. The Liouville form Λ , a one-form on T^*M , is defined by

$$\Lambda_{(q,\theta_q)} = \theta_q \circ T_{(q,\theta_q)}\pi. \quad (6.1)$$

The cotangent bundle T^*M has a canonical symplectic structure via the Liouville form given by the 2-form $\Omega = d\Lambda$. Recall the definition of a contact type hypersurface.

Definition 6.2.1. A hypersurface $\iota : \Sigma \rightarrow T^*M$ is of *contact type* if there exists a 1-form Θ on Σ such that

- (i) $d\Theta = \iota^*\Omega$,
- (ii) $\Theta(X) \neq 0$, for all $X \in \mathcal{L}_\Sigma$ with $X \neq 0$,

where \mathcal{L}_Σ is the characteristic line bundle for Σ in (T^*M, Ω) , defined by $\mathcal{L}_\Sigma = \ker \iota^*\Omega$.

The Hamiltonian vector field X_H on Σ , is defined by the relation

$$i_{X_H}\Omega = -dH,$$

and is a smooth non-vanishing section in the bundle \mathcal{L}_Σ . The proof of the following theorem is an adaptation of [AVDBV07, Lemma 3.3], which is in turn inspired by a contact type theorem for mechanical hypersurfaces in \mathbb{R}^{2n} in [HZ11].

Theorem 6.2.2. *A regular hypersurface $\Sigma = H^{-1}(0)$ of a mechanical hamiltonian is of contact type, irrespective of the compactness of Σ .*

Proof. The restriction of the Liouville form $\iota^*(\Lambda)$ to Σ vanishes on

$$S = \{(q, \theta_q) \in \Sigma \mid \theta_q = 0\},$$

which is a submanifold of Σ of dimension $n - 1$. The form $\iota^*\Lambda$ satisfies condition (i) of Definition 6.2.1, but does not satisfy condition (ii). Any exact perturbation $\Theta = \iota^*\Lambda + df$, for functions $f : \Sigma \rightarrow \mathbb{R}$, obviously satisfies condition (i) as well. We now construct an f such that the condition (ii) is also satisfied. This is first done locally. We observe that $i_X \iota^*\Lambda \geq 0$ for all positive multiples X of X_H . We therefore only need to prove that $i_{X_H} \Theta > 0$.

Let $x \in S$. Because Σ is a regular hypersurface, that is $dH \neq 0$ on Σ , the Hamiltonian vector field is transverse to S , i.e. $X_H \notin TS$. Thus there exists a chart $U \ni x$, with coordinates $y = (y_1, \dots, y_{2n-1})$, with $S \cap U \subset \{(y_1, \dots, y_{2n-2}, 0)\}$, and $X_H = \frac{\partial}{\partial y_{2n-1}}$. Construct a smooth and positive bump function h around x , supported in U , which is a product of different bump functions of the form

$$h(y) = h_1(y_{2n-1})h_2(y_1, \dots, y_{2n-2}).$$

Assume h is equal to 1 on a neighborhood $W \subset U$ of x . The function $f_x : \Sigma \rightarrow \mathbb{R}$ defined by

$$f_x = \begin{cases} y_{2n-2}h(y) & \text{for } y \in U \\ 0 & \text{otherwise,} \end{cases}$$

then satisfies

$$i_{X_H}df_x = h + y_{2n-1} \frac{\partial h_1}{\partial y_{2n-1}} h_2 \quad \text{on } U.$$

Because $y_{2n-1} = 0$ on S , this shows that $i_{X_H}df_x \geq 0$ on S , and equal to 1 at x , hence positive on a small neighborhood around x . By the product structure of the bump function $i_{X_H}df_x < 0$ only on a compact set outside a neighborhood of S . The quantity

$$A_x = \sup_{\Sigma \setminus S} \frac{\max(-i_{X_H}df_x, 0)}{i_{X_H}t^*\Lambda},$$

is therefore finite and non-zero.

Now we patch the local constructions together. Take a countable and dense collection of points $\{x_k\}_{k=1}^\infty$ in S , and construct f_{x_k} as above. Define $f : \Sigma \rightarrow \mathbb{R}$ via

$$f = \frac{1}{2} \sum_{k=1}^\infty 2^{-k} \frac{f_{x_k}}{A_k + \|f_{x_k}\|_{C^k}}.$$

This series converges in C^∞ to a smooth function. On S we have $i_{X_H}df_{x_k} \geq 0$ for all k , and for any $x \in S$, there exists a k such that $i_{X_H}f_{x_k}(x) > 0$ by the density of the sequence $\{x_k\}_{k=1}^\infty$. Thus $i_{X_H}f|_S > 0$. We compute this quantity

outside of S

$$\begin{aligned}
 -i_{X_H}df &= \frac{1}{2} \sum_{k=1}^{\infty} 2^{-k} \frac{-i_{X_H}df_{x_k}}{A_k + \|f_{x_k}\|_{C^k}}, \\
 &\leq \frac{1}{2} \sum_{k=1}^{\infty} 2^{-k} \frac{\max(-i_{X_H}df_{x_k}, 0)}{A_k}, \\
 &\leq \frac{1}{2} \sum_{k=1}^{\infty} 2^{-k} i_{X_H}\Lambda < \frac{1}{2} i_{X_H}\Lambda.
 \end{aligned}$$

On $\Sigma \setminus S$, we therefore have that $i_{X_H}df \geq -\frac{1}{2}i_{X_H}\Lambda$. From this we see that $i_{X_H}\Theta = i_{X_H}(i^*\Lambda + df) \geq 0$ on Σ . Thus Θ is a contact form. \square

Theorem 6.2.2 shows that the contact type condition is a general property of hypersurfaces of mechanical Hamiltonians regardless of asymptotic behavior at infinity. In the case that the potential is asymptotically regular, an explicit contact form can be constructed and a stronger contact type condition holds. Consider the vector field:

$$v(q) = -\frac{\text{grad } V(q)}{1 + |\text{grad } V(q)|^2}, \quad (6.2)$$

and the function $f : T^*M \rightarrow \mathbb{R}$ defined by $f(x) = \theta_q(v(q))$, for all $x = (q, \theta_q) \in T^*M$. For $\kappa > 0$, define the 1-form $\Theta = \Lambda + \kappa df$ on T^*M . Clearly, $d\Theta = \Omega$. Define the energy surfaces $\Sigma_\epsilon = \{x \in T^*M \mid H(x) = \epsilon\}$.

Proposition 6.2.3. *Let $\Sigma = H^{-1}(0)$ be a regular hypersurface of a mechanical Hamiltonian and assume that the potential is asymptotically regular. There exists $\epsilon_0, \kappa_0 > 0$, such that for every $-\epsilon_0 < \epsilon < \epsilon_0$, $\Theta = \Lambda + \kappa df$ restricts to a contact form on Σ_ϵ , for all $0 < \kappa \leq \kappa_0$. Moreover, for every κ , there exists a constant $a_\kappa > 0$ such that*

$$\Theta(X_H) \geq a_\kappa > 0, \quad \text{for all } x \in \Sigma_\epsilon \quad \text{and for all } -\epsilon_0 < \epsilon < \epsilon_0.$$

The energy surfaces Σ_ϵ are said to be of uniform contact type.

Proof. A tedious, but straightforward, computation, which we have moved to an appendix, Section 6.12, reveals that

$$X_H(f)(q, \theta_q) = \frac{|\operatorname{grad} V(q)|^2}{1 + |\operatorname{grad} V(q)|^2} - \frac{\operatorname{Hess} V(q)(\#\theta_q, \#\theta_q)}{1 + |\operatorname{grad} V(q)|^2} + \frac{2\theta_q(\operatorname{grad} V(q)) \operatorname{Hess} V(q)(\operatorname{grad} V, \#\theta_q)}{(1 + |\operatorname{grad} V(q)|^2)^2}. \quad (6.3)$$

The reverse triangle inequality, and Cauchy-Schwarz directly give

$$\begin{aligned} X_H(f) &\geq \frac{|\operatorname{grad} V(q)|^2}{1 + |\operatorname{grad} V(q)|^2} \\ &\quad - \left(\frac{\|\operatorname{Hess} V(q)\| |\theta_q|^2}{1 + |\operatorname{grad} V(q)|^2} + \frac{2\|\operatorname{Hess} V(q)\| |\operatorname{grad} V(q)|^2 |\theta_q|^2}{(1 + |\operatorname{grad} V(q)|^2)^2} \right) \\ &\geq \frac{|\operatorname{grad} V(q)|^2}{1 + |\operatorname{grad} V(q)|^2} - \frac{3\|\operatorname{Hess} V(q)\| |\theta_q|^2}{1 + |\operatorname{grad} V(q)|^2}. \end{aligned}$$

By asymptotic regularity there exists a constant C such that $\frac{3\|\operatorname{Hess} V(q)\|}{1 + |\operatorname{grad} V(q)|^2} \leq C$, hence

$$X_H(f) \geq \frac{|\operatorname{grad} V(q)|^2}{1 + |\operatorname{grad} V(q)|^2} - C |\theta_q|^2.$$

This yields the following global estimate

$$\begin{aligned} \Theta_x(X_H)(q, \theta_q) &\geq (1 - \kappa C) |\theta_q|^2 + \kappa \frac{|\operatorname{grad} V(q)|^2}{1 + |\operatorname{grad} V(q)|^2} \\ &\geq \frac{1}{2} |\theta_q|^2 + \kappa \frac{|\operatorname{grad} V(q)|^2}{1 + |\operatorname{grad} V(q)|^2} > 0, \quad \text{for all } x \in T^*M, \end{aligned}$$

for all $0 < \kappa \leq \kappa_0 = 1/2C$. The final step is to establish a uniform positive lower bound on a_κ for $(q, \theta_q) \in \Sigma_\epsilon$, independent of (q, θ_q) and ϵ .

If $d_M(q, K) \geq R$ is sufficiently large, then asymptotic regularity gives that $|\operatorname{grad} V(q)| > V_\infty$. Thus, in this region,

$$\frac{1}{2} |\theta_q|^2 + \kappa \frac{|\operatorname{grad} V(q)|^2}{1 + |\operatorname{grad} V(q)|^2} \geq \kappa \frac{|\operatorname{grad} V(q)|^2}{1 + |\operatorname{grad} V(q)|^2} \geq \frac{\kappa V_\infty^2}{1 + V_\infty^2}.$$

On $d_M(q, K) < R$ we can use standard compactness arguments. For $(q, \theta_q) \in \Sigma_\epsilon$, we have the energy identity

$$\frac{1}{2}|\theta_q|^2 + V(q) = \epsilon.$$

Suppose that $\frac{1}{2}|\theta_q|^2 < \epsilon_0$, then $|V(q)| < \epsilon + \epsilon_0$. If ϵ_0 is sufficiently small, this implies that $|\text{grad } V(q)| \geq V_0 > 0$ for some constant V_0 , because $\text{grad } V \neq 0$ at $V(q) = 0$. Therefore in this case

$$\frac{1}{2}|\theta_q|^2 + \kappa \frac{|\text{grad } V(q)|^2}{1 + |\text{grad } V(q)|^2} > \frac{\kappa V_0^2}{1 + V_0^2}.$$

If $|\theta_q|^2 > \epsilon_0$, then obviously

$$\frac{1}{2}|\theta_q|^2 + \kappa \frac{|\text{grad } V(q)|^2}{1 + |\text{grad } V(q)|^2} > \epsilon_0.$$

We have exhausted all possibilities and

$$a_\kappa = \min\left(\frac{\kappa V_\infty^2}{1 + V_\infty^2}, \frac{\kappa V_0^2}{1 + V_0^2}, \epsilon_0\right) > 0,$$

is a uniform lower bound for $i_{X_H} \Theta$. □

6.3 The variational setting

Closed characteristics on M can be regarded as critical points of the action functional

$$\mathcal{B}(q, T) = \int_0^T \left\{ \frac{1}{2} |q'(t)|^2 - V(q(t)) \right\} dt,$$

for mappings $q : [0, T] \rightarrow M$. Via the coordinate transformation

$$(q(t), T) \mapsto (c(s), \tau) = (q(sT), \log(T)).$$

we obtain the rescaled action functional

$$\mathcal{A}(c, \tau) = \frac{e^{-\tau}}{2} \int_0^1 |c'(s)|^2 ds - e^\tau \int_0^1 V(c(s)) ds,$$

for mappings $c : [0, 1] \rightarrow M$ and $\tau \in \mathbb{R}$. For closed loops we impose the boundary condition $c(0) = c(1)$. By defining the parametrized circle $\mathbb{S}^1 = [0, 1] / \{0, 1\}$, the mappings $c : \mathbb{S}^1 \rightarrow M$ satisfy the appropriate boundary condition.

In order to treat closed characteristics as critical points of the action functional \mathcal{A} we need an appropriate functional analytic setting. We briefly recall this setting. Details can be found in the books of Klingenberg [Kli78] and [Kli95]. A map $c : \mathbb{S}^1 \rightarrow M$ is called H^1 if it is absolutely continuous and the derivative is square integrable with respect to the Riemannian metric g on M , i.e. $\int_0^1 |c'(s)|^2 ds < \infty$. The space of H^1 maps is denoted by $H^1(\mathbb{S}^1, M)$, or ΛM , the free loop space of H^1 -loops. An equivalent way to define ΛM is to consider continuous curves $c : \mathbb{S}^1 \rightarrow M$ such that for any chart (U, φ) of M , the function $\varphi \circ c : J = c^{-1}(U) \rightarrow \mathbb{R}^n$ is in $H^1(J, \mathbb{R}^n)$. There is a natural sequence of continuous inclusions

$$C^\infty(\mathbb{S}^1, M) \subset H^1(\mathbb{S}^1, M) \subset C^0(\mathbb{S}^1, M),$$

and the free loop space $C^0(\mathbb{S}^1, M)$ is complete with respect to the metric $d_{C^0}(c, e) = \sup_{s \in \mathbb{S}^1} d_M(c(s), \tilde{e}(s))$, where d_M is the metric on M induced by g . The smooth loops $C^\infty(\mathbb{S}^1, M)$ are dense in $C^0(\mathbb{S}^1, M)$. The free loop space $\Lambda M = H^1(\mathbb{S}^1, M)$ can be given the structure of a Hilbert manifold. Let $c \in C^\infty(\mathbb{S}^1, M)$ and denote by $C^\infty(c^*TM)$ the space of smooth sections in the pull-back bundle $c^*\pi : c^*TM \rightarrow \mathbb{S}^1$. For sections $\zeta, \eta \in C^\infty(c^*TM)$ consider the norms and inner products:

$$\begin{aligned} \|\zeta\|_{C^0} &= \sup_{s \in \mathbb{S}^1} |\zeta(s)|, \\ \langle \zeta, \eta \rangle_{L^2} &= \int_0^1 \langle \zeta(s), \eta(s) \rangle ds, \\ \langle \zeta, \eta \rangle_{H^1} &= \int_0^1 \langle \zeta(s), \eta(s) \rangle ds + \int_0^1 \langle \nabla_s \zeta(s), \nabla_s \eta(s) \rangle ds, \end{aligned}$$

with ∇_s the induced connection on the pull-back bundle (from the Levi-Civita connection on M). The spaces $C^0(c^*TM)$, $L^2(c^*TM)$ and $H^1(c^*TM)$ are the completions of $C^\infty(c^*TM)$ with respect to the norms $\|\cdot\|_{C^0}$, $\|\cdot\|_{L^2}$ and $\|\cdot\|_{H^1}$, respectively.

The loop space ΛM is a smooth Hilbert manifold locally modeled over the Hilbert spaces $H^1(c^*TM)$, with c any smooth loop in M . For each smooth loop c , the space $H^1(c^*TM)$ is a separable Hilbert space, hence these are all

isomorphic. The tangent space $T_c\Lambda M$ at a smooth loop c consists of H^1 vector fields along the loop and is canonically isomorphic to $H^1(c^*TM)$.

For fixed $\tau \in \mathbb{R}$ the functional is well-defined on the loop space ΛM . The kinetic energy term $\mathcal{E}(c) := \frac{1}{2} \int_0^1 |c'(s)|^2 ds$ is well-defined and continuous. The embedding of ΛM into $C^0(\mathbb{S}^1, M)$ and the continuity of the potential V imply that the potential energy is also well-defined and continuous. The latter implies the continuity of the functional $\mathcal{A} : \Lambda M \times \mathbb{R} \rightarrow \mathbb{R}$. We now discuss differentiability and the first variation formula. The details are in the book of Klingenberg [Kli95].

We also denote the weak derivative of $c \in \Lambda M$ by $\partial c = c'$. It is possible to extend the Levi-Civita connection on ΛM to differentiate $\xi \in H^1(c^*TM)$ with respect to $\eta \in L^2(c^*TM)$. We denote this connection with the same symbol $\nabla_\eta \xi$, and we write, for $\xi \in T_c\Lambda M$,

$$\nabla \xi = \nabla_{\partial c} \xi.$$

In general $\nabla \xi \in L^2(c^*TM)$. If ξ and c are smooth, then $\nabla \xi(s) = \nabla_s \xi(s)$. For $c \in \Lambda M$, and $\xi, \eta \in T_c\Lambda M$, the metric on ΛM can be written as

$$\langle \xi, \eta \rangle = \langle \xi, \eta \rangle_{L^2} + \langle \nabla \xi, \nabla \eta \rangle_{L^2}.$$

Observe that the kinetic energy is given by $\mathcal{E}(c) = \frac{1}{2} \|\partial c\|_{L^2}^2$. Let $\xi \in T_c\Lambda M$ and $\gamma : (-\epsilon, \epsilon) \rightarrow \Lambda M$ a smooth curve with $\gamma(0) = c$ and $\gamma'(0) = \xi$. Then,

$$\left. \frac{d}{dt} \mathcal{E}(\gamma(t)) \right|_{t=0} = \langle \partial \gamma(t), \nabla_{\partial \gamma(t)} \frac{\partial \gamma(t)}{\partial t} \rangle_{L^2} \Big|_{t=0} = \langle \partial c, \nabla \xi \rangle_{L^2}.$$

For the function $\mathcal{W}(c) = \int_0^1 V(c(s)) ds$ a similar calculation gives:

$$\left. \frac{d}{dt} \mathcal{W}(\gamma(t)) \right|_{t=0} = \langle \text{grad } V \circ c, \xi \rangle_{L^2}.$$

For convenience we write $\mathcal{A}(c, \tau) = e^{-\tau} \mathcal{E}(c) - e^\tau \mathcal{W}(c)$.

Lemma 6.3.1. *The action $\mathcal{A} : \Lambda M \times \mathbb{R} \rightarrow \mathbb{R}$ is continuously differentiable. For any $(\xi, \sigma) \in T_{(c, \sigma)} \Lambda M \times \mathbb{R}$ the first variation is given by*

$$\begin{aligned} d\mathcal{A}(c, \tau)(\xi, \sigma) &= e^{-\tau} \int_0^1 \langle \partial c(s), \nabla \xi(s) \rangle ds - e^\tau \int_0^1 \langle \text{grad } V(c(s)), \xi(s) \rangle ds \\ &\quad - \int_0^1 \left[\frac{e^{-\tau}}{2} |\partial c(s)|^2 + e^\tau V(c(s)) \right] \sigma ds \\ &= d_c \mathcal{A}(c, \tau) \xi - \left(e^{-\tau} \mathcal{E}(c) + e^\tau \mathcal{W}(c) \right) \sigma, \end{aligned}$$

where the gradient $\text{grad } V$ is taken with respect to the metric g on M .

Proof. The first term in \mathcal{A} is $e^{-\tau} \mathcal{E}(c)$ of which the derivative is given by

$$-e^{-\tau} \mathcal{E}(c) \sigma + e^{-\tau} d\mathcal{E}(c) \xi = -e^{-\tau} \mathcal{E}(c) \sigma + e^{-\tau} \langle \partial c, \nabla \xi \rangle_{L^2}.$$

Writing out the L^2 -metric gives the desired result. The second and third term follow directly from the definition of derivative. \square

A *critical point* of \mathcal{A} is a pair $(c, \tau) \in \Lambda M \times \mathbb{R}$ such that $d\mathcal{A}(c, \tau) = 0$. Critical points are contained in $C^\infty(\mathbb{S}^1, M) \times \mathbb{R}$ and satisfy the second order equation $e^{-\tau} \nabla_s c'(s) + e^\tau \text{grad } V(c(s)) = 0$. One can see this by taking variations of the form $(\xi, 0)$. Consider the expression $H(s) = \frac{e^{-2\tau}}{2} |c'(s)|^2 + V(c(s))$, then

$$\frac{d}{ds} H(s) = e^{-2\tau} \langle c'(s), \nabla_s c'(s) \rangle + \langle \text{grad } V(c(s)), c'(s) \rangle = 0,$$

which implies that $H(s)$ is constant. From the first variation formula, with variations $(0, \sigma)$ it follows that $\int_0^1 H(s) ds = 0$, and therefore $H \equiv 0$ for critical points of \mathcal{A} .

Since $\Lambda M \times \mathbb{R}$ has a Riemannian metric we can define the gradient of \mathcal{A} . The metric on $\Lambda M \times \mathbb{R}$ is denoted by $\langle \cdot, \cdot \rangle_{H^1 \times \mathbb{R}} = \langle \cdot, \cdot \rangle_{H^1} + \langle \cdot, \cdot \rangle_{\mathbb{R}}$. The gradient $\text{grad } \mathcal{A}(c, \tau)$ is the unique vector such that

$$\langle \text{grad } \mathcal{A}(c, \tau), (\xi, \sigma) \rangle_{H^1 \times \mathbb{R}} = d\mathcal{A}(c, \tau)(\xi, \sigma),$$

for all $(\xi, \sigma) \in T_{c, \tau} \Lambda M \times \mathbb{R}$, and $\text{grad } \mathcal{A}$ defines a vector field on $\Lambda M \times \mathbb{R}$.

We conclude with some basic inequalities for the various metrics which will be used in this analysis. The proofs are in Klingenberg [Kli95]. Let $c, e \in \Lambda M$, then

$$d_M(c(s), c(s')) \leq \sqrt{|s - s'|} \sqrt{2\mathcal{E}(c)}, \quad (6.4)$$

$$d_{C^0}(c, e) \leq \sqrt{2} d_{\Lambda M}(c, e), \quad (6.5)$$

where $d_{\Lambda M}$ is the metric induced by the Riemannian metric on ΛM , and $d_{C^0}(c, e) = \sup_{s \in \mathbb{S}^1} d_M(c(s), e(s))$. Furthermore, for $\xi \in T_c \Lambda M$, we have

$$\|\xi\|_{L^2} \leq \|\xi\|_{C^0} \leq \sqrt{2} \|\xi\|_{H^1}. \quad (6.6)$$

6.4 The Palais-Smale condition

The functional \mathcal{A} does not satisfy the Palais-Smale condition. We therefore approximate this functional by functionals \mathcal{A}_ϵ , and show that the approximating functionals do satisfy PS. We then find critical points of the approximating functionals using a linking argument. Finally we show that these critical points converge to a critical point of \mathcal{A} as $\epsilon \rightarrow 0$. The approximating, or penalized, functionals are defined by

$$\mathcal{A}_\epsilon(c, \tau) = \mathcal{A}(c, \tau) + \epsilon(e^{-\tau} + e^{\tau/2}).$$

The term $\epsilon e^{-\tau}$ penalizes orbits with short periods, and $\epsilon e^{\tau/2}$ penalizes orbits with long periods. Recall, that for $\epsilon > 0$ fixed, a sequence $\{(c_n, \tau_n)\} \in \Lambda M \times \mathbb{R}$ is called a *Palais-Smale sequence* for \mathcal{A}_ϵ , if:

- (i) there exist constants $a_1, a_2 > 0$ such that $a_1 \leq \mathcal{A}_\epsilon(c_n, \tau_n) \leq a_2$;
- (ii) $d\mathcal{A}_\epsilon(c_n, \tau_n) \rightarrow 0$ as n tends to ∞ .

Condition (ii) can be equivalently rewritten as

$$d\mathcal{A}_\epsilon(c_n, \tau_n)(\xi, \sigma) = \langle \text{grad } \mathcal{A}_\epsilon(c_n, \tau_n), (\xi, \sigma) \rangle_{H^1 \times \mathbb{R}} = o(1)(\|\xi\|_{H^1} + |\sigma|), \quad (6.7)$$

as $n \rightarrow \infty$ and $(\xi, \sigma) \in T_{(c_n, \tau_n)}\Lambda M \times \mathbb{R}$. Condition (i) implies that, by passing to a subsequence if necessary, $\mathcal{A}_\epsilon(c_n, \tau_n) \rightarrow a_\epsilon$, with $a_1 \leq a_\epsilon \leq a_2$. In what follows we tacitly assume we have passed to such a subsequence.

Remark 6.4.1. We will only consider Palais-Smale sequences that are positive, thus $a_1 > 0$. The functionals \mathcal{A}_ϵ satisfy the Palais-Smale condition for critical levels $a_\epsilon > a_1 > 0$.

The relation between \mathcal{A} and \mathcal{A}_ϵ gives:

$$d\mathcal{A}_\epsilon(c_n, \tau_n)(\xi, \sigma) = d\mathcal{A}(c_n, \tau_n)(\xi, \sigma) - \epsilon\left(e^{-\tau} - \frac{1}{2}e^{\tau/2}\right)\sigma.$$

Lemma 6.4.2. *A Palais-Smale sequence (c_n, τ_n) satisfies*

$$2e^{-\tau_n} \mathcal{E}(c_n) + \epsilon\left(2e^{-\tau_n} + \frac{1}{2}e^{\tau_n/2}\right) = a_\epsilon + o(1), \quad (6.8)$$

$$e^{\tau_n} \mathcal{W}(c_n) - \epsilon \frac{3}{4}e^{\tau_n/2} = -\frac{a_\epsilon}{2} + o(1), \quad \text{as } n \rightarrow \infty. \quad (6.9)$$

Proof. Consider variations of the form $(\xi, \sigma) = (0, 1)$. From the variation formula and (6.7) we then derive that

$$e^{-\tau_n} \mathcal{E}(c_n) + e^{\tau_n} \mathcal{W}(c_n) = -\epsilon \left(e^{-\tau_n} - \frac{1}{2} e^{\tau_n/2} \right) + o(1)$$

as $n \rightarrow \infty$. On the other hand, $\mathcal{A}_\epsilon(c_n, \tau_n) \rightarrow a_\epsilon$ means that

$$e^{-\tau_n} \mathcal{E}(c_n) - e^{\tau_n} \mathcal{W}(c_n) = -\epsilon \left(e^{-\tau_n} + e^{\tau_n/2} \right) + a_\epsilon + o(1).$$

Combining these two estimates completes the proof. \square

We obtain the following a priori bounds on τ_n .

Lemma 6.4.3. *Let (c_n, τ_n) be a Palais-Smale sequence. There are constants $T_0 < T_1$ (depending on ϵ) such that $T_0 \leq \tau_n \leq T_1$ for sufficiently large n .*

Proof. From Equation (6.8) it follows that $\epsilon(2e^{-\tau_n} + \frac{1}{2}e^{\tau_n/2}) \leq a_\epsilon + 1$ for sufficiently large n , which proves the lemma. \square

From this we also obtain a priori bounds on the energy.

Lemma 6.4.4. *Let (c_n, τ_n) be a Palais-Smale sequence. Then $\|\partial c_n\|_{L^2}^2 = 2\mathcal{E}(c_n) \leq C(\epsilon)$, independent of n .*

Proof. The a priori bounds on τ_n can be used in Equation (6.8), which yields

$$2\mathcal{E}(c_n) = e^{\tau_n} \left\{ C_\epsilon - \epsilon \left(2e^{-\tau_n} + \frac{1}{2} e^{\tau_n/2} \right) + o(1) \right\} \leq C(\epsilon),$$

which proves the lemma. \square

The following proposition establishes the Palais-Smale condition for the action \mathcal{A}_ϵ , $\epsilon > 0$.

Proposition 6.4.5. *Let (c_n, τ_n) be a Palais-Smale sequence for \mathcal{A}_ϵ . Then (c_n, τ_n) has an accumulation point $(c_\epsilon, \tau_\epsilon) \in \Lambda M \times \mathbb{R}$ that is a critical point, i.e. $d\mathcal{A}_\epsilon(c_\epsilon, \tau_\epsilon) = 0$ and the action is bounded $0 < a_1 \leq \mathcal{A}_\epsilon(c_\epsilon, \tau_\epsilon) = a_\epsilon \leq a_2$.*

6. PERIODIC ORBITS

Proof. From Lemmas 6.4.3 and 6.4.4 we have that $\mathcal{E}(c_n) \leq C$ and $|\tau_n| \leq C'$, with the constants $C, C' > 0$ depending only on ϵ . Fix $s_0 \in \mathbb{S}^1$, then by Eq. (6.4) we have $d_M(c_n(s), c_n(s_0)) \leq \sqrt{|s - s_0|} \sqrt{2C} \leq \sqrt{2C}$, and therefore $c_n(s) \in B_{\sqrt{2C}}(c_n(s_0))$, for all $s \in \mathbb{S}^1$ and all n . We argue that $c_n(s_0)$ must remain in a compact set as $n \rightarrow \infty$. From this, and the energy bound on c_n , it follows that $c_n(s)$ remains in a compact set for all $s \in \mathbb{S}^1$ as $n \rightarrow \infty$. The sequence c_n is equicontinuous and the generalized Arzela-Ascoli Theorem shows that c_n must have a convergent subsequence in the metric d_{C_0} . We then argue that this is actually an accumulation point for the metric $d_{\Lambda}M$, which shows that \mathcal{A}_ϵ satisfies the Palais-Smale condition.

We argue by contradiction. Suppose $d_M(c_n(s_0), K) \rightarrow \infty$ as $n \rightarrow \infty$. Since $K \subset M$ is compact, its diameter is finite and there exists a constant C_K such that $d_M(q, \tilde{q}) \leq C_K$, for all $q, \tilde{q} \in K$. If $K \cap B_{\sqrt{2C}}(c_n(s_0)) \neq \emptyset$ for all n , then there exist points $q_n \in K \cap B_{\sqrt{2C}}(c_n(s_0))$, with $d_M(q, q_n) \leq C_K$ for all $q \in K$ and all n . This implies:

$$C_K + \sqrt{2C} \geq d_M(q, q_n) + d_M(q_n, c_n(s_0)) \geq d_M(q, c_n(s_0)) \geq d_M(c_n(s_0), K) \rightarrow \infty,$$

as $n \rightarrow \infty$, which is a contradiction, and thus there exists an N such that

$$K \cap B_{\sqrt{2C}}(c_n(s_0)) = \emptyset, \quad \text{for } n \geq N.$$

As a consequence $c_n(s) \in M \setminus K$, for all $s \in \mathbb{S}^1$ and all $n \geq N$. Next we show it is impossible that $d_M(c_n(s_0), K) \rightarrow \infty$. Suppose there does not exist a constant C'' such that $d_M(c_n(s_0), K) \leq C''$, i.e. $d_M(c_n(s_0), K) \rightarrow \infty$ as $n \rightarrow \infty$. By the previous considerations there exists an N such that $c_n(s) \in M \setminus K$, for all $s \in \mathbb{S}^1$ and all $n \geq N$, and thus, using the asymptotic regularity, we have that

$$|\text{grad } V(c_n(s))|_{c_n(s)} \geq V_\infty > 0, \quad \text{for all } s \in \mathbb{S}^1, \text{ and for all } n \geq N.$$

For $q \in M \setminus K$ we can define the smooth vector field on M by

$$x(q) = -\frac{\text{grad } V(q)}{|\text{grad } V(q)|^2}.$$

For a loop $c \in C^\infty(\mathbb{S}^1, M)$ the vector field along c is given by

$$\tilde{\zeta}(s) = x(c(s)) = -\frac{\text{grad } V(c(s))}{|\text{grad } V(c(s))|^2},$$

which is a smooth section in the pull-back bundle $C^\infty(c^*TM)$. We now investigate the H^1 -norm of ξ in $H^1(c^*TM)$. For a vector field y on TM we have that

$$\begin{aligned} |\nabla_y x(q)| &= \left| -\frac{\nabla_y \text{grad } V(q)}{|\text{grad } V(q)|^2} - 2\frac{\langle \nabla_y \text{grad } V(q), \text{grad } V(q) \rangle}{|\text{grad } V(q)|^4} \text{grad } V(q) \right| \\ &\leq 3\frac{|\nabla_y \text{grad } V(q)|}{|\text{grad } V(q)|^2} \leq 3\frac{\|\text{Hess } V(q)\|}{|\text{grad } V(q)|^2} |y|, \end{aligned}$$

where we used the identity $|\nabla_y \text{grad } V(q)|^2 = \text{Hess } V(q)(y, \nabla_y \text{grad } V(q))$. Now set $\xi_n(s) = x(c_n(s))$. Using the asymptotic regularity of V , we obtain that

$$|\nabla_s \xi_n(s)| \leq 3\frac{\|\text{Hess } V(c(s))\|}{|\text{grad } V(c(s))|^2} |\partial c(s)| = o(1)|\partial c(s)|,$$

and $|\xi_n(s)| \leq \frac{1}{V_\infty}$ as $d_M(c_n(s_0), K) \rightarrow \infty$. Thus the H^1 norm of ξ_n is

$$|\xi_n|_{H^1} = \frac{1}{V_\infty} + o(1)|\partial c|_{L^2}.$$

By the first variation formula in Lemma 6.3.1

$$\begin{aligned} d\mathcal{A}_\epsilon(c_n, \tau_n)(\xi_n, 0) &= e^{-\tau_n} \langle \partial c_n, \nabla \xi_n \rangle_{L^2} - e^{\tau_n} \langle \text{grad } V \circ c_n, \xi_n \rangle_{L^2} \\ &= o(1)\|\partial c_n\|_{L^2}^2 + e^{\tau_n} = o(1) + e^{\tau_n}. \end{aligned} \quad (6.10)$$

On the other hand, since $c_n(s)$ is a Palais-Smale sequence, we have

$$d\mathcal{A}_\epsilon(c_n, \tau_n)(\xi_n, 0) = o(1)\|\xi_n\|_{H^1} = o(1)\left(\frac{1}{V_\infty} + Co(1)\right) = o(1),$$

which contradicts (6.10), since $|\tau_n|$ is bounded by Lemma 6.4.3. This now shows that $d_M(c_n(s_0), K) \leq C''$ and therefore there exists an $0 < R < \infty$ such that $c_n(s) \in B_R(K)$ for all $s \in \mathbb{S}^1$ and all $n \geq N$. Since (M, g) is complete, the Hopf-Rinow Theorem implies that $B_R(K)$ is compact and thus $\{c_n(s)\} \subset M$ is pre-compact for any fixed $s \in \mathbb{S}^1$. The sequence $\{c_n(s)\}$ is point wise relatively compact and equicontinuous by Eq. (6.4). Therefore, by the generalized version of the Arzela-Ascoli Theorem [Mun75] there exists a subsequence c_{n_k} converging in $C^0(\mathbb{S}^1, M)$ (uniformly) to a continuous limit $c_\epsilon \in C^0(\mathbb{S}^1, M)$. It remains to show that c_ϵ is an accumulation point in ΛM , thus in H^1 sense.

Due to the above convergence in $C^0(S^1, M)$, the sequence $\{c_n\}$ can be assumed to be contained in a fixed chart $(\mathcal{U}(c_0), \exp_c^{-1})$, for a fixed $c_0 \in C^\infty(S^1, M)$. Following [Kli78] it suffices to show that $\exp_{c_0}^{-1} c_n$ is a Cauchy sequence in $T_{c_0} \Lambda M = H^1(c_0^* TM)$. This final technical argument is identical to Theorem 1.4.7 in [Kli78], which proves that $\{c_n\}$ has an accumulation point in $(c_\epsilon, \tau_\epsilon) \in \Lambda M \times \mathbb{R}$, proving the Palais-Smale condition for \mathcal{A}_ϵ . The limit points satisfy $d\mathcal{A}_\epsilon(c_\epsilon, \tau_\epsilon) = 0$, and $\mathcal{A}_\epsilon(c_\epsilon, \tau_\epsilon) = a_\epsilon$. \square

For critical points of \mathcal{A}_ϵ we prove additional a priori estimates on τ_ϵ . The latter imply a priori estimates on c_ϵ . This allows us to pass to the limit as $\epsilon \rightarrow 0$.

We start with pointing out that critical points of the penalized action \mathcal{A}_ϵ satisfy the following Hamiltonian identity

$$H_\epsilon(s) = \frac{e^{-2\tau_\epsilon}}{2} |c'_\epsilon(s)|^2 + V(c_\epsilon(s)) \equiv \epsilon \left(-e^{-2\tau_\epsilon} + \frac{1}{2} e^{-\tau_\epsilon/2} \right) = \tilde{\epsilon}. \quad (6.11)$$

Thus the critical point $(c_\epsilon, \tau_\epsilon)$ corresponds to a closed characteristic on $\Sigma_{\tilde{\epsilon}}$. Via the transformation $q_\epsilon(t) = c_\epsilon(te^{-\tau})$ and the Legendre transform of $(q_\epsilon, q'_\epsilon)$ to a curve γ_ϵ on the cotangent bundle we see that the Hamiltonian action is

$$\mathcal{A}_\epsilon^H(\gamma_\epsilon, \tau_\epsilon) = \int_{\gamma_\epsilon} \Lambda + \epsilon(e^{-\tau_\epsilon} + e^{\tau_\epsilon/2}),$$

The following a priori bounds are due to the uniform contact type of Σ , cf. Lemma 6.2.3.

Lemma 6.4.6. *Let $(c_\epsilon, \tau_\epsilon)$ be critical points of \mathcal{A}_ϵ , with $0 < a_1 \leq \mathcal{A}_\epsilon(c_\epsilon, \tau_\epsilon) \leq a_2$. Then there is a constant T_2 , independent of ϵ , such that $\tau_\epsilon \leq T_2$ for sufficiently small ϵ .*

Proof. We start with the case $\tau_\epsilon \geq 0$. The Hamiltonian action satisfies

$$\mathcal{A}_\epsilon^H(\gamma_\epsilon, \tau_\epsilon) = \int_{\gamma_\epsilon} \Lambda + \epsilon(e^{-\tau_\epsilon} + e^{\tau_\epsilon/2}) \leq a_2,$$

and thus $\int_{\gamma_\epsilon} \Lambda \leq a_2$. Since Σ is of uniform contact type it holds for $\gamma_\epsilon \subset \Sigma_{\tilde{\epsilon}}$, with $\tilde{\epsilon} \leq \epsilon \leq \epsilon_0$, that

$$a_2 \geq \int_{\gamma_\epsilon} \Lambda = \int_{\gamma_\epsilon} \Theta = \int_0^{e^{\tau_\epsilon}} \alpha_{\gamma_\epsilon}(X_H) \geq a_\kappa e^{\tau_\epsilon}.$$

We conclude that always $\tau_\epsilon \leq \max\{0, \log(a_2/a_\kappa)\}$, which proves the lemma. \square

We can also establish a lower bound on τ_ϵ under the condition that M is asymptotically flat. This is the only estimate that requires *flat* ends. All other estimates carry through under the weaker assumption of bounded geometry.

Lemma 6.4.7. *Let $(c_\epsilon, \tau_\epsilon)$ be critical points of \mathcal{A}_ϵ , with $0 < a_1 \leq \mathcal{A}_\epsilon(c_\epsilon, \tau_\epsilon) \leq a_2$. If (M, g) is asymptotically flat, then there is a constant T_3 , independent of ϵ , such that $\tau_\epsilon \geq T_3$ for sufficiently small ϵ .*

Proof. Assume, by contradiction that $\tau_\epsilon \rightarrow -\infty$ as $\epsilon \rightarrow 0$. Then Equation (6.8) gives

$$2\mathcal{E}(c_\epsilon) = e^{\tau_\epsilon} a_\epsilon - 2\epsilon - \frac{\epsilon}{2} e^{3\tau_\epsilon/2} \rightarrow 0, \quad \text{as } \epsilon \rightarrow 0. \quad (6.12)$$

Fix $s_0 \in \mathbb{S}^1$. Then the previous equation implies, using Equation 6.4, that $c_\epsilon(s) \in B_{\epsilon'}(c_\epsilon(s_0))$, where $\epsilon' = \sqrt{e^{\tau_\epsilon} a_2 - 2\epsilon - \frac{\epsilon}{2} e^{3\tau_\epsilon/2}}$. We distinguish two cases:

(i) There exists an $R > 0$ such that $d_M(c_\epsilon(s_0), K) \leq R$ for all ϵ . Then $c_\epsilon(s) \in B_{\epsilon'+R}(K)$, and therefore $|V(c_\epsilon(s))| \leq C$ for all $s \in \mathbb{S}^1$ and all $\epsilon > 0$. This implies $\int_0^1 e^{\tau_\epsilon} V(c_\epsilon(s)) ds \rightarrow 0$, which contradicts (6.9), as $a_\epsilon > 0$, and thus $\tau_\epsilon \geq T_3$.

(ii) Now we assume no such $R > 0$ exists, and assume thus that $d_M(c_\epsilon(s_0), K) \rightarrow \infty$ as $\epsilon \rightarrow 0$ to derive a contradiction. By bounded geometry of M , every point $q \in M$ has a normal charts $(\mathcal{U}_q, \exp_q^{-1})$ and constants $\rho_0, R_0 > 0$ such that $B_{\rho_0}(q) \subset \mathcal{U}_q$ and $|\partial^\ell \Gamma_{ij}^k(q)| \leq R_0$. This implies that $c_\epsilon(s) \in \mathcal{U}_{c_\epsilon(s_0)}$ for sufficiently small ϵ . We assume M has flat ends, and since $d(c_\epsilon(s_0), K) \rightarrow \infty$ the metric on the charts $\mathcal{U}_{c_\epsilon(s_0)}$ is flat. We identify these charts with open subsets of \mathbb{R}^n henceforth. The differential equation c_ϵ satisfies is

$$e^{-2\tau_\epsilon} \nabla_s c'_\epsilon(s) + \text{grad } V(c_\epsilon(s)) = 0. \quad (6.13)$$

Take the unique geodesic γ from $c_\epsilon(s_0)$ to $c_\epsilon(s)$ parameterized by arc length, i.e.

$$\gamma(0) = c_\epsilon(s_0), \quad \gamma(d_M(c_\epsilon(s_0), c_\epsilon(s))) = c_\epsilon(s), \quad \text{and} \quad |\gamma'(t)| = 1.$$

6. PERIODIC ORBITS

Then, by asymptotic regularity, $\|\text{Hess } V(\gamma(t))\| \leq C|\text{grad } V(\gamma(t))|$ for some constant $C > 0$, and

$$\begin{aligned} \frac{d}{dt} |\text{grad } V(\gamma(t))|^2 &= 2 \text{Hess } V(\gamma(t))(\text{grad } V(\gamma(t)), \gamma'(t)) \\ &\leq 2\|\text{Hess } V(\gamma(t))\| |\text{grad } V(\gamma(t))| \leq 2C |\text{grad } V(\gamma(t))|^2. \end{aligned}$$

Gronwall's inequality therefore implies that $|\text{grad } V(\gamma(t))| \leq |\text{grad } V(\gamma(0))|e^{Ct}$. We identify $U_{c_\epsilon(s_0)}$ with an open subset of \mathbb{R}^n , and we write

$$\text{grad } V(\gamma(t)) = \text{grad } V(\gamma(0)) + \int_0^t \frac{d}{d\sigma} \text{grad } V(\gamma(\sigma)) d\sigma.$$

Hence

$$\begin{aligned} |\text{grad } V(\gamma(t)) - \text{grad } V(\gamma(0))| &\leq \int_0^t \|\text{Hess } V(\gamma(\sigma))\| d\sigma \\ &\leq C \int_0^t |\text{grad } V(\gamma(\sigma))| d\sigma \quad (6.14) \\ &\leq |\text{grad } V(\gamma(0))|(e^{Ct} - 1). \end{aligned}$$

For any solution to Equation (6.13), we compute

$$\begin{aligned} \frac{d}{ds} e^{2\tau_\epsilon} \langle \text{grad } V(c_\epsilon(s_0)), c'_\epsilon(s) \rangle &= e^{2\tau_\epsilon} \langle \text{grad } V(c_\epsilon(s_0)), \nabla_s c'_\epsilon(s) \rangle \\ &= -\langle \text{grad } V(c_\epsilon(s_0)), \text{grad } V(c_\epsilon(s)) \rangle \\ &= -\langle \text{grad } V(c_\epsilon(s_0)), \text{grad } V(c_\epsilon(s_0)) \rangle \\ &\quad - \langle \text{grad } V(c_\epsilon(s_0)), \text{grad } V(c_\epsilon(s)) - \text{grad } V(c_\epsilon(s_0)) \rangle \end{aligned}$$

By asymptotic regularity and Estimate (6.14), we find that

$$\frac{d}{ds} e^{2\tau_\epsilon} \langle \text{grad } V(c_\epsilon(s_0)), c'_\epsilon(s) \rangle \leq -V_\infty^2 + V_\infty^2 (e^{Cd_M(c_\epsilon(s_0), c_\epsilon(s))} - 1).$$

We see that as $\epsilon \rightarrow 0$ that $e^{2\tau_\epsilon} \langle \text{grad } V(c_\epsilon(s_0)), c'_\epsilon(s) \rangle$ is monotonically decreasing in s . We conclude that c_ϵ cannot be periodic. This is a contradiction, therefore there exists a constant T_3 such that $\tau_\epsilon \geq T_3$, for sufficiently small ϵ . \square

Proposition 6.4.8. *Let $(c_\epsilon, \tau_\epsilon)$, $\epsilon \rightarrow 0$ be a sequence satisfying $d\mathcal{A}_\epsilon(c_\epsilon, \tau_\epsilon) = 0$, and $0 < a_1 \leq \mathcal{A}_\epsilon(c_\epsilon, \tau_\epsilon) \leq a_2$. If (M, g) has flat ends, then there exists a convergent subsequence $(c_{\epsilon'}, \tau_{\epsilon'}) \rightarrow (c, \tau)$ in $\Lambda M \times \mathbb{R}$, $\epsilon' \rightarrow 0$. The limit satisfies $d\mathcal{A}(c, \tau) = 0$, and $0 < a_1 \leq \mathcal{A}(c, \tau) \leq a_2$.*

Proof. From Lemmas 6.4.6 and 6.4.7 we obtain uniform bounds on τ_ϵ . We can now repeat the arguments of the proof of Proposition 6.4.5 on the sequence $\{c_\epsilon\}$, from which we draw the desired conclusion. \square

6.5 Minimax characterization

We prove that bounds of the action functional on certain linking homology classes give rise to critical values. Set $E = \Lambda M \times \mathbb{R}$. For $d \in \mathbb{R}$ define the sublevel sets

$$\mathcal{A}_\epsilon^d = \{(c, \tau) \in E \mid \mathcal{A}_\epsilon(c, \tau) \leq d\}.$$

Lemma 6.5.1. *Suppose we have subsets $A, B \subset E$ such that*

- (i) *The morphism induced by inclusion $i_{k+1} : H_{k+1}(A) \rightarrow H_{k+1}(E - B)$ is non-trivial.*
- (ii) *The subset A is compact, and the action functional satisfies the bounds*

$$\mathcal{A}|_B \geq b > 0 \quad \mathcal{A}|_A \leq a < b.$$

- (iii) *The homology of the loop space vanishes in the $(k + 1)$ -st degree, and therefore $H_{k+1}(E) = 0$.*

Then there exist $\bar{a} > 0$ and $\epsilon^ > 0$ such that for all $0 < \epsilon \leq \epsilon^*$, there exist a non-trivial $[y_\epsilon^{\bar{a}}] \in H_{k+2}(E, \mathcal{A}_\epsilon^{\bar{a}})$ and*

$$C_\epsilon = \inf_{y_\epsilon^{\bar{a}} \in [y_\epsilon^{\bar{a}}]} \max_{|y_\epsilon^{\bar{a}}|} \mathcal{A}_\epsilon,$$

satisfies the estimates $\bar{a} < C_\epsilon < C$ for a finite constant C , independent of ϵ .

Proof. Choose \bar{a} and \bar{b} such that $a < \bar{a} < \bar{b} < b$, with $\bar{a} > 0$. Define $\epsilon^* = \frac{1}{2} \inf_{(c, \tau) \in A} \frac{\bar{a} - a}{e^{-\tau} + e^{\frac{\tau}{2}}}$, which is finite and positive because A is compact. For all $0 < \epsilon \leq \epsilon^*$, and $(c, \tau) \in A$, the following estimate holds

$$\mathcal{A}_\epsilon(c, \tau) = \mathcal{A}(c, \tau) + \epsilon \left(e^{-\tau} + e^{\frac{\tau}{2}} \right) < a + \frac{a - \bar{a}}{2} < \bar{a},$$

For $(c, \tau) \in B$ we obtain

$$\mathcal{A}_\epsilon(c, \tau) = \mathcal{A}(c, \tau) + \epsilon \left(e^{-\tau} + e^{\frac{\tau}{2}} \right) > b > \bar{b}.$$

6. PERIODIC ORBITS

For the remainder of this proof, we assume $\epsilon \leq \epsilon^*$. One immediately verifies the following inclusions of sublevel sets

$$A \longrightarrow \mathcal{A}_{\epsilon^*}^{\bar{a}} \longrightarrow \mathcal{A}_{\epsilon}^{\bar{a}} \longrightarrow \mathcal{A}_{\epsilon}^{\bar{b}} \longrightarrow E \setminus B.$$

When we pass to homology, the sequence becomes

$$H_{k+1}(A) \rightarrow H_{k+1}(\mathcal{A}_{\epsilon^*}^{\bar{a}}) \rightarrow H_{k+1}(\mathcal{A}_{\epsilon}^{\bar{a}}) \rightarrow H_{k+1}(\mathcal{A}_{\epsilon}^{\bar{b}}) \rightarrow H_{k+1}(E \setminus B), \quad (6.15)$$

By assumption (i), the morphism induced by the inclusion is non-trivial, and it factors through this sequence. Therefore all the homology groups in Sequence (6.15) are non-trivial. Consider the following part of the long exact sequence of the pair $(E, \mathcal{A}_{\epsilon}^{\bar{b}})$

$$H_{k+2}(E, \mathcal{A}_{\epsilon}^{\bar{b}}) \rightarrow H_{k+1}(\mathcal{A}_{\epsilon}^{\bar{b}}) \rightarrow H_{k+1}(E).$$

By assumption $H_{k+1}(E) = 0$, hence $H_{k+2}(E, \mathcal{A}_{\epsilon}^{\bar{b}}) \rightarrow H_{k+1}(\mathcal{A}_{\epsilon}^{\bar{b}})$ is surjective. Because $H_{k+1}(\mathcal{A}_{\epsilon}^{\bar{b}}) \neq 0$, we conclude $H_{k+2}(E, \mathcal{A}_{\epsilon}^{\bar{b}}) \neq 0$. The same argument shows that the homology groups $H_{k+2}(E, \mathcal{A}_{\epsilon}^{\bar{a}})$ are non-zero. Naturality of the boundary map with respect to the inclusions of pairs shows that the following diagram commutes

$$\begin{array}{ccccccc} H_{k+1}(A) & \longrightarrow & H_{k+1}(\mathcal{A}_{\epsilon^*}^{\bar{a}}) & \longrightarrow & H_{k+1}(\mathcal{A}_{\epsilon}^{\bar{a}}) & \longrightarrow & H_{k+1}(\mathcal{A}_{\epsilon}^{\bar{b}}) \longrightarrow H_{k+1}(E \setminus B) . \\ & & \uparrow & & \uparrow & & \uparrow \\ & & H_{k+2}(E, \mathcal{A}_{\epsilon^*}^{\bar{a}}) & \longrightarrow & H_{k+2}(E, \mathcal{A}_{\epsilon}^{\bar{a}}) & \longrightarrow & H_{k+2}(E, \mathcal{A}_{\epsilon}^{\bar{b}}) \end{array} \quad (6.16)$$

The vertical maps are the boundary homomorphisms of the long exact sequences of the pairs. Choose $[x] \in H_{k+1}(A)$ such that $i_{k+1}([x]) \neq 0$. Because the map i_{k+1} factors through Sequence (6.15), we get non-zero elements in those groups. We can lift the non-zero $[y^{\bar{a}}] \in H_{k+1}(\mathcal{A}_{\epsilon^*}^{\bar{a}})$ to some non-zero $[y_{\epsilon^*}^{\bar{a}}] \in H_{k+2}(E, \mathcal{A}_{\epsilon^*}^{\bar{a}})$ by the surjectivity of the boundary homomorphism. The element $[x]$ is mapped by the inclusions in equation (6.15) as follows

$$\begin{array}{ccccccc} [x] & \longmapsto & [x_{\epsilon^*}^{\bar{a}}] & \longmapsto & [x_{\epsilon}^{\bar{a}}] & \longmapsto & [x_{\epsilon}^{\bar{b}}] \longmapsto i_{k+1}([x]) . \\ & & \uparrow \downarrow & & \uparrow & & \uparrow \\ & & [y_{\epsilon^*}^{\bar{a}}] & \longmapsto & [y_{\epsilon}^{\bar{a}}] & \longmapsto & [y_{\epsilon}^{\bar{b}}] \end{array}$$

This diagram defines the undefined elements. All elements in the above diagram are non-zero. Now define

$$C_\epsilon = \inf_{\mathbf{y}_\epsilon^{\bar{a}} \in [\mathbf{y}_\epsilon^{\bar{a}}]} \max_{|\mathbf{y}_\epsilon^{\bar{a}}|} \mathcal{A}_\epsilon.$$

Here we abuse notation and write $\mathbf{y}_\epsilon^{\bar{a}}$ for the representatives of $[\mathbf{y}_\epsilon^{\bar{a}}]$. The infimum runs over all elements representing the class $[\mathbf{y}_\epsilon^{\bar{a}}]$. The support, or image, of $\mathbf{y}_\epsilon^{\bar{a}}$ is denoted by $|\mathbf{y}_\epsilon^{\bar{a}}|$. The morphism induced by inclusion maps $\mathbf{y}_\epsilon^{\bar{a}} \mapsto \mathbf{y}_\epsilon^{\bar{b}}$, an element in $C_{k+2}(E, \mathcal{A}_\epsilon^{\bar{b}})$, which represents the class $[\mathbf{y}_\epsilon^{\bar{b}}]$. The support does not change, hence

$$\max_{|\mathbf{y}_\epsilon^{\bar{a}}|} \mathcal{A}_\epsilon = \max_{|\mathbf{y}_\epsilon^{\bar{b}}|} \mathcal{A}_\epsilon > \bar{b}.$$

This follows from the following observation. If $\max_{|\mathbf{y}_\epsilon^{\bar{b}}|} \mathcal{A}_\epsilon \leq \bar{b}$, then $|\mathbf{y}_\epsilon^{\bar{a}}| \subset \mathcal{A}_\epsilon^{\bar{b}}$, hence $[\mathbf{y}_\epsilon^{\bar{b}}] = 0$ in $H_{k+2}(E, \mathcal{A}_\epsilon^{\bar{b}})$, which is a contradiction. We therefore conclude $C_\epsilon \geq \bar{b} > \bar{a}$. If we now pick a representative $\mathbf{y}_{\epsilon^*}^{\bar{a}} \in [\mathbf{y}_{\epsilon^*}^{\bar{a}}]$, which is mapped to $\mathbf{y}_\epsilon^{\bar{a}}$, we have that

$$\inf_{\mathbf{y}_\epsilon^{\bar{a}} \in [\mathbf{y}_\epsilon^{\bar{a}}]} \max_{|\mathbf{y}_\epsilon^{\bar{a}}|} \mathcal{A}_\epsilon \leq \max_{|\mathbf{y}_{\epsilon^*}^{\bar{a}}|} \mathcal{A}_\epsilon \leq \max_{|\mathbf{y}_{\epsilon^*}^{\bar{a}}|} \mathcal{A}_{\epsilon^*} =: C. \quad (6.17)$$

We conclude that we have, for each $\epsilon \leq \epsilon^*$, a non-trivial class $[\mathbf{y}_\epsilon^{\bar{a}}] \in H_{k+2}(E, \mathcal{A}_\epsilon^{\bar{a}})$, such that

$$\bar{a} < C_\epsilon < C. \quad (6.18)$$

The constants that bound C_ϵ are independent of ϵ . \square

A minimax argument now gives, if the functionals satisfy (PS), a critical point of \mathcal{A} .

Proposition 6.5.2. *Assume the hypotheses of Lemma 6.5.1 are met. There exists constants $0 < a_1 < a_2 < \infty$, such that, for all $\epsilon > 0$ sufficiently small, \mathcal{A}_ϵ has a critical point $(c_\epsilon, \tau_\epsilon)$, satisfying $a_1 < \mathcal{A}(c_\epsilon, \tau_\epsilon) < a_2$.*

Proof. The homology class $[\mathbf{y}_\epsilon^{\bar{a}}]$ is homotopy invariant, and the minimax value C_ϵ over this class is finite and greater than zero, by Lemma 6.5.1. According to the Minimax principle, cf. [Cha93], this gives rise to a Palais-Smale sequence (c_n, τ_n) for \mathcal{A}_ϵ , with $a_1 = \bar{a}$ and $a_2 = C$. Proposition 6.4.5 states that \mathcal{A}_ϵ satisfies the Palais-Smale condition, hence this produces, for each $\epsilon > 0$ sufficiently small, a critical point $(c_\epsilon, \tau_\epsilon)$ of \mathcal{A}_ϵ . \square

6.6 Some remarks on the geometry and topology of the loop space

The following observation is a simple relation between the homology of the loop space and the base manifold.

Proposition 6.6.1. *There exists an isomorphism*

$$H_*(\Lambda M) \cong H_*(M) \oplus H_*(\Lambda M, M). \quad (6.19)$$

Proof. We show that M is a retract of ΛM . Recall that the topology on ΛM coincides with the compact open topology. Define the inclusion map $\iota : M \rightarrow \Lambda M$, by sending a point $q \in M$ to the constant loop c_q , with $c_q(s) = q$ for all $s \in \mathbb{S}^1$. This map is an embedding [Kli78, Theorem 1.4.6]. Also the evaluation map $\text{ev} : \mathbb{S}^1 \times \Lambda M \rightarrow M$, defined by $\text{ev}_s(c) = c(s)$ is continuous [Mun75, Theorem 46.10]. Pick a fixed $s_0 \in \mathbb{S}^1$. Then $\text{ev}_{s_0} : \Lambda M \rightarrow M$ is continuous, and $\text{ev}_{s_0} \circ \iota = \text{id}$ is the identity on M . Hence M is a retract (but in general not a deformation retract) of ΛM and the Splitting Lemma [Hat02] gives the desired result. \square

The goal is to construct a link in the function space with right bounds for \mathcal{A} on the linking sets. We will in fact construct this link in a tubular neighborhood of M in ΛM . To construct the linking sets in the loop space, we need to have a geometric understanding of the submanifold of constant loops. By Proposition 6.6.1 the inclusion $\iota : M \hookrightarrow \Lambda M$ sending a point to the constant loop at that point is an embedding. By the assumption of bounded geometry, we can construct a well behaved tubular neighborhood. Let NM be the normal bundle of $\iota(M)$ in ΛM . Recall that we denote the constant loop at $q \in M$, by c_q , thus $c_q(s) = q$ for all $s \in \mathbb{S}$. Elements $\xi \in N_{c_q}M$ are characterized by the fact that $\int_{\mathbb{S}^1} \langle \xi(s), \eta_0 \rangle ds = 0$, for all $\eta_0 \in T_q M$.

Proposition 6.6.2. *Assume that M is of bounded geometry. Then there exists an open neighborhood \mathcal{V} of $\iota(M)$ in ΛM and a diffeomorphism $\phi : NM \rightarrow \mathcal{V}$, with the property that it maps $\xi \in NM$ with $\|\xi\|_{H^1} \leq \frac{\text{inj} M}{2}$ to $\phi(\xi) \in \Lambda M$ with $d_{H^1}(c_q, \phi(\xi)) = \|\xi\|_{H^1}$.*

Proof. Let $k : NM \rightarrow NM$ be a smooth injective radial fiber-wise contraction such that

$$k(\xi) = \begin{cases} \xi & \text{for } \|\xi\|_{H^1} < \frac{\text{inj} M}{2} \\ \frac{\text{inj} M}{1 + \|\xi\|_{H^1}} \xi & \text{for } \|\xi\|_{H^1} \text{ large} \end{cases}.$$

6.6. Some remarks on the geometry and topology of the loop space

Thus $\|k(\xi)\|_{H^1} < \text{inj } M$ for all $\xi \in NM$, and k is the identity on the disc bundle of radius $\text{inj } M/2$. Define the map $\phi : NM \rightarrow \Lambda M$ by

$$\phi(\xi) = \exp_{c_q} k(\xi),$$

for $\xi \in N_{c_q}M$. This is a diffeomorphism onto an open subset $\mathcal{V} \subset \Lambda M$. We use the fact that k is injective, and along with bounded geometry this shows that \exp_{c_q} is injective for $\|\xi\|_{H^1} < \text{inj } M$. We use bounded geometry to globalize this statement to the whole of M . \square

The inclusion of the zero section in the normal bundle is denoted by $\zeta : M \rightarrow NM$. The zero section of the normal bundle is mapped diffeomorphically into $\iota(M) \subset \Lambda M$ by ϕ . To form the required linking sets, we need a different, but equivalent, Riemannian structure on the normal bundle.

Proposition 6.6.3. *On the normal bundle NM , the norm $\|\cdot\|_{\perp}$ defined by*

$$\|\xi\|_{\perp} = \int_{S^1} \langle \nabla \xi(s), \nabla \xi(s) \rangle ds, \quad (6.20)$$

is equivalent to the norm $\|\cdot\|_{H^1}$. To be precise the following estimate holds

$$\|\xi\|_{\perp} \leq \|\xi\|_{H^1} \leq \sqrt{2}\|\xi\|_{\perp}. \quad (6.21)$$

Proof. It is obvious that $\|\xi\|_{\perp} \leq \|\xi\|_{H^1}$. We need to show the remaining inequality. We do this using Fourier expansions. For $\xi \in T_{c_q}\Lambda M$, a vector field along a constant loop c_q , we can write $\xi(s) = \sum_k \lambda_k e^{2\pi i k s}$, with $\lambda_k \in T_q M \otimes \mathbb{C}$. Then

$$\|\xi\|_{H^1}^2 = \sum_k \langle \lambda_k, \bar{\lambda}_k \rangle + \sum_k 4\pi^2 k^2 \langle \lambda_k, \bar{\lambda}_k \rangle.$$

If $\xi \in NM$ then $\lambda_0 = 0$, by the characterization of the normal bundle given above, hence $\sum_k \langle \lambda_k, \bar{\lambda}_k \rangle \leq \sum_k 4\pi^2 k^2 \langle \lambda_k, \bar{\lambda}_k \rangle$, which gives that $\|\xi\|_{H^1} \leq \sqrt{2}\|\xi\|_{\perp}$. \square

We will use the next proposition to show that the linking sets we construct in the normal bundle persist in the loop space.

Proposition 6.6.4. *Let \mathcal{V} be the tubular neighborhood of M inside ΛM , constructed in Proposition 6.6.2. Let T be any subspace of \mathcal{V} whose closure is contained in \mathcal{V} . Assume $H_{k+2}(\Lambda M) = 0$. Then the morphism*

$$H_{k+1}(\mathcal{V} \setminus T) \rightarrow H_{k+1}(\Lambda M \setminus T), \quad (6.22)$$

induced by inclusion is injective, and

$$H_{k+2}(\mathcal{V}\setminus T) \rightarrow H_{k+2}(\Lambda M\setminus T), \quad (6.23)$$

is surjective.

Proof. We know from Proposition 6.6.1 that

$$H_{k+1}(\Lambda M) \cong H_{k+1}(\Lambda M, M) \oplus H_{k+1}(M).$$

We assume $H_{k+2}(\Lambda M) = 0$, thus $H_{k+2}(\Lambda M, M) = 0$. The tubular neighborhood \mathcal{V} deformation retracts to M . If we apply the Five-Lemma to the long exact sequences of the pairs $(\Lambda M, M)$ and $(\Lambda M, \mathcal{V})$, where the vertical maps are induced by the deformation retraction,

$$\begin{array}{ccccccccc} H_{k+2}(\mathcal{V}) & \longrightarrow & H_{k+2}(\Lambda M) & \longrightarrow & H_{k+2}(\Lambda M, \mathcal{V}) & \longrightarrow & H_{k+1}(\mathcal{V}) & \longrightarrow & H_{k+1}(\Lambda M) \\ \downarrow \cong & & \downarrow \cong & & \downarrow & & \downarrow \cong & & \downarrow \cong \\ H_{k+2}(M) & \longrightarrow & H_{k+2}(\Lambda M) & \longrightarrow & H_{k+2}(\Lambda M, M) & \longrightarrow & H_{k+1}(M) & \longrightarrow & H_{k+1}(\Lambda M) \end{array}$$

we see that $H_{k+2}(\Lambda M, \mathcal{V}) \cong H_{k+2}(\Lambda M, M)$. Since T is contained in the closure of \mathcal{V} , we can excise T . This gives an isomorphism $H_{k+2}(\Lambda M\setminus T, \mathcal{V}\setminus T) \cong H_{k+2}(\Lambda M, \mathcal{V}) \cong 0$. The long exact sequence of the pair $(\Lambda M\setminus T, \mathcal{V}\setminus T)$ is

$$H_{k+2}(\mathcal{V}\setminus T) \rightarrow H_{k+2}(\Lambda M\setminus T) \rightarrow H_{k+2}(\Lambda M\setminus T, \mathcal{V}\setminus T) \rightarrow H_{k+1}(\mathcal{V}\setminus T) \rightarrow H_{k+1}(\Lambda M\setminus T)$$

The homology group $H_{k+2}(\Lambda M\setminus T, \mathcal{V}\setminus T)$ is zero by the preceding argument, thus the map on the right is injective and the map on the left is surjective. \square

6.7 The topology of the hypersurface and its shadow

We investigate the relation between the topology of Σ and its projection $N = \pi(\Sigma)$ to the base manifold. In the case N does not have a boundary, this relation is expressed by the classical Gysin sequence for sphere bundles. Recall that we assume H to be mechanical and that the hypersurface $\Sigma = H^{-1}(0)$ is regular. Thus N and its boundary ∂N are given by

$$N = \{q \in M \mid V(q) \leq 0\}, \quad \text{and} \quad \partial N = \{q \in M \mid V(q) = 0\},$$

and ∂N is smooth. Let us consider a basic example, the harmonic oscillator on \mathbb{R} . The potential is $V(q) = \frac{k}{2}q^2 - V_0$ for constants $k, V_0 > 0$. One directly




Figure 6.3: The relation between the topology of Σ and $N = \pi(\Sigma)$ for the harmonic oscillator. The fibers of Σ over the interior of N are copies of S^0 that get smaller at the boundary, where they are collapsed to a point.

verifies that $N = \left[-\sqrt{\frac{2V_0}{k}}, \sqrt{\frac{2V_0}{k}}\right]$ is an interval and $\Sigma = \{(q, \theta) \in \mathbb{R}^2 \mid \frac{1}{2}\theta^2 + \frac{k}{2}q^2 = V_0\}$ is an ellipse. The fibers over the interior of N are copies of S^0 , whose size is determined by the distance to the boundary ∂N . Close to the boundary the fibers get smaller and at the boundary the fibers are collapsed to a point. Topologically the collapsing process can be understood by the gluing of an interval at the boundary of N . The hypersurface Σ is thus homeomorphic to $N \times S^0 \cup_{\partial N \times S^0} \partial N \times D^1$, which Figure 6.3 illustrates. This picture generalizes to arbitrary hypersurfaces Σ . We have the topological characterization

$$\Sigma \cong ST^*N \bigcup_{ST^*N|_{\partial N}} DT^*N|_{\partial N}. \quad (6.24)$$

The characterization is given in terms of the sphere bundle ST^*N and the disc bundle DT^*N in the cotangent bundle of N . The vertical bars denote the restriction of the bundles to the boundary. This topological characterization gives a relation between the homology of Σ and N . In this section we identify Σ with this characterization.

Recall that a map is proper if preimages of compact sets are compact. In the proof of the next proposition, compactly supported cohomology $H_c^*(M)$ is used, which is contravariant with respect to proper maps. In singular (co)homology, homotopic maps induce the same maps in (co)homology. For

compactly supported cohomology, maps that are homotopic via a homotopy of proper maps, induce the same maps in cohomology. If $\partial N = \emptyset$ the following proposition directly follows from the Gysin sequence.

Proposition 6.7.1. *There exist isomorphisms $H_c^i(\Sigma) \cong H_c^i(N)$ for all $0 \leq i \leq n - 2$.*

Proof. Let C be the closure of a collar of ∂N in N , which always exists, cf. [Hir94]. Thus C deformation retracts via a proper homotopy onto ∂N . Now $\pi^{-1}(C)$ is the closure of a collar of $ST^*N|_{\partial N} = \partial ST^*N$ in ST^*N , and therefore it deformation retracts via a proper homotopy onto $ST^*N|_{\partial N}$. Define $D \subset \Sigma$ by

$$D = \pi^{-1}(C) \bigcup_{ST^*N|_{\partial N}} DT^*N|_{\partial N}.$$

This is a slight enlargement of the disc bundle of M restricted to the boundary ∂N , which Figure 6.4 clarifies. By construction D deformation retracts properly to $DT^*N|_{\partial N}$, which in turn deformation retracts properly to ∂N . This induces an isomorphism

$$H_c^*(D) \cong H_c^*(\partial N). \tag{6.25}$$

Let $S = ST^*N$. The intersection $D \cap S$ deformation retracts properly to $ST^*N|_{\partial N}$. Thus the isomorphism

$$H_c^*(D \cap S) \cong H_c^*(ST^*N|_{\partial N}), \tag{6.26}$$

holds. The inclusions in the diagram

$$\begin{array}{ccccc} & & & S & \\ & & \nearrow^{i_1} & & \searrow^{j_1} \\ S \cap D & & & & \Sigma \\ & & \searrow^{i_2} & & \nearrow^{j_2} \\ & & & D & \end{array}$$

are proper maps, because the domains are all closed subspaces of the codomains. This gives rise to the contravariant Mayer-Vietoris sequence of compactly supported cohomology of the triad (Σ, S, D)

$$\rightarrow H_c^i(\Sigma) \xrightarrow{(j_1^i, -j_2^i)} H_c^i(S) \oplus H_c^i(D) \xrightarrow{i_1^i + i_2^i} H_c^i(S \cap D) \rightarrow H_c^{i+1}(\Sigma) \rightarrow \tag{6.27}$$

The map i_2^i is an isomorphism: this can be seen from the Gysin sequence for compactly supported cohomology as follows. Recall that from any vector bundle $E \rightarrow B$ of rank n over a locally compact space B , we can construct

6.7. The topology of the hypersurface and its shadow

a sphere bundle $SE \rightarrow B$. The Gysin sequence relates the cohomology of SE and B ,

$$\dots \rightarrow H_c^{i-n}(B) \xrightarrow{\epsilon^i} H_c^i(B) \xrightarrow{\pi^i} H_c^i(SE) \xrightarrow{\delta} H_c^{i-n+1}(B) \rightarrow \dots \quad (6.28)$$

The map ϵ^i is the cup product with the Euler class of the sphere bundle. We apply this sequence to the sphere bundle in T^*N restricted to ∂N . For dimensional reasons the sequence breaks down into short exact sequences

$$0 \rightarrow H_c^i(\partial N) \xrightarrow{\pi^i} H_c^i(ST^*N|_{\partial N}) \rightarrow 0 \quad \text{for} \quad 0 \leq i \leq n-2. \quad (6.29)$$

The diagram

$$\begin{array}{ccc} H_c^i(D) & \xrightarrow{i_2^i} & H_c^i(S \cap D) \\ \cong \downarrow & & \downarrow \cong \\ H_c^i(\partial N) & \xrightarrow{\pi^i} & H_c^i(ST^*N|_{\partial N}) \end{array}$$

commutes. This shows that i_2^i is an isomorphism for $0 \leq i \leq n-2$. The map $i_1^i + i_2^i$ in the Mayer-Vietoris sequence, Equation (6.27), is surjective, and the sequence breaks down into short exact sequences

$$0 \longrightarrow H_c^i(\Sigma) \longrightarrow H_c^i(S) \oplus H_c^i(D) \xrightarrow[p]{i_1^i + i_2^i} H_c^i(S \cap D) \longrightarrow 0. \quad (6.30)$$

More is true, since the sequence actually splits by the map $p = (0, (i_2^i)^{-1})$. If we study the Gysin sequence for N and S we see that

$$0 \longrightarrow H_c^i(N) \xrightarrow{\pi^i} H_c^i(S) \longrightarrow 0, \quad \text{for} \quad 0 \leq i \leq n-2. \quad (6.31)$$

The isomorphisms (6.31), (6.25), (6.26), and (6.29) can be applied to the sequence in Equation (6.30), which becomes

$$0 \rightarrow H_c^i(\Sigma) \rightarrow H_c^i(N) \oplus H_c^i(\partial N) \xrightarrow{\pi^i} H_c^i(\partial N) \rightarrow 0, \quad \text{for} \quad 0 \leq i \leq n-2.$$

Because the sequence is split the stated isomorphism holds. □




Figure 6.4: A sketch of the spaces D , and S . In the picture N is a half-line, hence ∂N is a point. The topology of the energy hypersurface can be recovered from its projection N .

Proposition 6.7.2. *For all $2 \leq i \leq n$ there is an isomorphism*

$$H_i(N, \partial N) \cong H_{i+n-1}(\Sigma). \quad (6.32)$$

Proof. This is a double application of Poincaré duality for non-compact manifolds with boundary. The dimension of N is n , and therefore Poincaré duality gives $H_i(N, \partial N) \cong H_c^{n-i}(N)$. The boundary of Σ is empty, and its dimension equals $2n - 1$, thus $H_{n+i-1}(\Sigma) \cong H_c^{n-i}(\Sigma)$. By Proposition 6.7.1 we have $H_c^{n-i}(N) \cong H_c^{n-i}(\Sigma)$, for all $2 \leq i \leq n$. The isomorphism stated in the proposition is the composition of the isomorphisms. \square

We would also like the previous proposition to be true if $i = 1$. This is the case if the bundle ST^*N is trivial, but in general this is not true. However, the following result is sufficient for our needs.

Proposition 6.7.3. *If $H_n(\Sigma) \neq 0$ and $H_n(M) = 0$, then $H_1(N, \partial N) \neq 0$.*

Proof. We will show that a non-zero element in $H_c^{n-1}(\Sigma)$ gives rise to a non-zero element in $H_c^{n-1}(N)$. A double application of Poincaré duality, as in the previous proposition, will give the desired result. We will use the same notation as in the proof of Proposition 6.7.1. The Gysin sequence, Equation (6.28), for the sphere bundle $ST^*N|_{\partial N}$ over ∂N breaks down to the short exact

shows a naturality property of the Gysin sequence.

$$\begin{array}{ccccccc}
 0 & \longrightarrow & H_c^{n-1}(\partial N) & \longrightarrow & H_c^{n-1}(ST^*N|_{\partial N}) & \xrightarrow{\delta} & H_c^0(\partial N) & \longrightarrow & 0 \\
 & & \uparrow \cong & & \uparrow = & & \uparrow \Phi^{-1} \cong & & \\
 0 & \longrightarrow & H_c^{n-1}(DT^*N|_{\partial N}) & \longrightarrow & H_c^{n-1}(ST^*N|_{\partial N}) & \xrightarrow{\delta} & H_c^n(DT^*N|_{\partial N}, ST^*N|_{\partial N}) & \longrightarrow & 0 \\
 & & \uparrow & & \uparrow i_1^{n-1} & & \uparrow & & \\
 0 & \longrightarrow & H_c^{n-1}(DT^*N) & \longrightarrow & H_c^{n-1}(ST^*N) & \xrightarrow{\delta} & H_c^n(DT^*N, ST^*N) & \longrightarrow & 0 \\
 & & \uparrow \cong & & \uparrow = & & \uparrow \Phi \cong & & \\
 0 & \longrightarrow & H_c^{n-1}(N) & \longrightarrow & H_c^{n-1}(ST^*N) & \xrightarrow{\delta} & H_c^0(N) & \longrightarrow & 0
 \end{array}$$

The top and bottom rows are the Gysin sequences of $(\partial N, ST^*N|_{\partial N})$ and $(N, ST^*N|_N)$ respectively. The vertical maps between the middle rows are the pullback maps of the inclusion of pairs $(DT^*N|_{\partial N}, ST^*N|_{\partial N}) \rightarrow (DT^*N, ST^*N|_N)$. The vertical maps Φ are the Thom isomorphisms. The map i_1^{n-1} in the diagram is the same as the map induced by $i_1 : S \cap D \rightarrow S$, under the isomorphism induced by the homotopy equivalence $S \cap D \cong ST^*N|_{\partial N}$, which we therefore denote by the same symbol.

We want to show that $\delta j_1^{n-1}(y) = 0$ for all $y \in H_c^{n-1}(\Sigma)$. For this we argue as follows. Recall that $H_c^0(N)$ consists of constant functions of compact support, and therefore is generated by the number of compact components of N . If the vertical map in the third column, from $H_c^0(N) \rightarrow H_c^0(\partial N)$ is not injective, then N has a compact component without boundary. This implies that M must have a compact component without boundary. But we assume that $H_n(M) = 0$, therefore M does not have orientable compact components and $H_c^0(N) \rightarrow H_c^0(\partial N)$ is injective. Let $[y] \in H_c^0(\Sigma)$ be non-zero. Obviously in $H_c^{n-1}(ST^*N|_{\partial N})$ we have the equality $i_1^{n-1} j_1^{n-1}([y]) = i_2^{n-1} j_2^{n-1}([y])$, and from the definition of the boundary map in the long exact sequence of the pair in the second row of the diagram, we obtain

$$\begin{aligned}
 \delta i_1^{n-1} j_1^{n-1}([y]) &= \delta i_2^{n-1} j_2^{n-1}([y]) = [p^{-1} d(i_2^{n-1})^{-1} i_2^{n-1} j_2^{n-1} y] \\
 &= [p^{-1} d j_2^{n-1} y] = [p^{-1} j_2^{n-1} d y] = 0.
 \end{aligned}$$

where p is the projection map in the defining short exact sequence. By the injectivity of the map $H_c^0(N) \rightarrow H_c^0(\partial N)$, and the commutativity of the dia-

gram we must have that $\delta j_1^{n-1}([\mathbf{y}]) = 0 \in H_c^0(N)$. The exactness of the bottom row now shows that there must be an element in $H_c^{n-1}(N)$ which is mapped to $j_1^{n-1}([\mathbf{y}])$, because it $j_1^{n-1}([\mathbf{y}])$ is in the kernel of δ . Poincaré duality for non-compact manifolds with boundary states that $H_n(\Sigma) \cong H_c^{n-1}(\Sigma)$, and $H_c^{n-1}(N) \cong H_1(N, \partial N)$. Thus, by the preceding argument we get a non-zero class in $H_1(N, \partial N)$. \square

Proposition 6.7.4. *Suppose that $H_{k+n}(\Sigma) \neq 0$ and $H_{k+1}(M) = 0$, for some $0 \leq k \leq n - 1$. Then there exists a non-zero class in $H_k(M \setminus N)$ which is mapped to zero in $H_k(M)$ by the morphism induced by the inclusion.*

Proof. Consider the long exact sequence of the pair $(M, M \setminus N)$

$$H_{k+1}(M) \rightarrow H_{k+1}(M, M \setminus N) \rightarrow H_k(M \setminus N) \rightarrow H_k(M).$$

The homology group $H_{k+1}(M)$ is zero by assumption, thus the middle map is injective. If we can find a non-zero element in $H_{k+1}(M, M \setminus N)$, then we see it is mapped to a non-zero element of $H_k(M \setminus N)$, which in turn is mapped to zero in $H_k(M)$ by exactness of the sequence. Using excision, we will now show that $H_{k+1}(M, M \setminus N)$ is isomorphic to $H_{k+1}(N, \partial N)$. The latter group is non-zero, by Propositions 6.7.2 and 6.7.3. Take a neighborhood V of N , which strongly deformation retracts to N . This is possible because N is a submanifold with boundary. The complement of V in M is closed, and $M \setminus V \subset M \setminus N$. The latter space is open, hence we can excise $M - V$. Thus

$$H_{k+1}(M, M \setminus N) \cong H_{k+1}(V, V \setminus N). \tag{6.34}$$

The pair $(V, V \setminus N)$ deformation retracts to the pair $(N, \partial N)$ by construction. By the Five Lemma applied to the long exact sequences of the pairs, we have $H_{k+1}(M, M \setminus N) \cong H_{k+1}(N, \partial N)$. Propositions 6.7.2 and 6.7.3 show that this homology group is non-zero. Thus there is a non-zero element of $H_k(M \setminus N)$ which is mapped to zero in $H_k(M)$ by the inclusion for $0 \leq k \leq n - 1$. \square

Remark 6.7.5. In our setting, the assumption $H_{k+1}(M) = 0$ is automatically satisfied. This follows from the assumption $H_{k+1}(\Lambda M) = 0$ on the topology of the loop space, and Lemma 6.6.1.




Figure 6.5: The projection $N = \pi(\Sigma)$ is shrunk to N_ν using the gradient flow of the function f defined in Equation (6.8.1). The boundary ∂N is shown in red, and the boundary ∂N_ν is grey.

6.8 The link

6.8.1 The parameter ν

For analytical reasons, we need to shrink the set $N = \pi(\Sigma) = \{q \in M \mid V(q) \leq 0\}$ to

$$N_\nu = \{q \in M \mid V(q) \leq -\nu\sqrt{1 + |\text{grad } V(q)|^2}\}. \quad (6.35)$$

For small ν this can be done diffeomorphically. On the modified set N_ν we estimate the potential V uniformly. A sketch is given in Figure 6.5.

Lemma 6.8.1. *There exist $\nu > 0$ sufficiently small, such that*

- *The spaces N and N_ν are diffeomorphic, and $M \setminus N$ and $M \setminus N_\nu$ are diffeomorphic.*
- *If $H_{k+n}(\Sigma) \neq 0$ and $H_{k+1}(M) = 0$ for some k , there exists a non-zero class in $H_k(M - N_\nu)$ which is mapped to zero in $H_k(M)$ by the morphism induced by the inclusion.*
- *There exists a $\rho_\nu > 0$ such that, for all $q \in N_\nu$,*

$$V(\tilde{q}) \leq -\frac{\nu}{2}, \quad \text{for all } \tilde{q} \in B_{\rho_\nu}(q). \quad (6.36)$$

Proof. Consider the function $f : M \rightarrow \mathbb{R}$ defined by

$$f(q) = \frac{V(q)}{\sqrt{1 + |\text{grad } V(q)|^2}}.$$

The gradient flow of this function induces the diffeomorphism. Because N is non-compact, the standard Morse Lemma does not apply. We show that this function satisfies the Palais-Smale condition and that it has no critical values between 0 and $-\nu$. Because M is assumed to be complete, a theorem of Palais [Pal63, Theorem 10.2] shows that N and N_ν are diffeomorphic.

We first show that f satisfies the Palais-Smale condition. Let q_n be a Palais-Smale sequence for f . That is, we assume there exists a constant a such that the sequence satisfies

$$|f(q_n)| < a \quad \text{and} \quad |\text{grad } f(q_n)| \rightarrow 0 \quad \text{as} \quad n \rightarrow \infty.$$

Sequences for which q_n stays in a compact set have a convergent subsequence, which must be a critical point of f . Thus we assume that $d(q_n, K) \rightarrow \infty$, and show that we get a contradiction. We compute the gradient of f

$$\text{grad } f(q) = \frac{\text{grad } V(q)}{\sqrt{1 + |\text{grad } V(q)|^2}} - \frac{1}{2} f(q) \frac{\text{grad } |\text{grad } V(q)|^2}{1 + |\text{grad } V(q)|^2},$$

and estimate

$$|\text{grad } f(q_n)| \geq \frac{|\text{grad } V(q_n)|}{\sqrt{1 + |\text{grad } V(q_n)|^2}} - \frac{1}{2} |f(q_n)| \frac{|\text{grad } |\text{grad } V(q_n)|^2|}{1 + |\text{grad } V(q_n)|^2}.$$

The first term is bounded from below by $V_\infty / \sqrt{1 + V_\infty^2}$ for large n , by asymptotic regularity. In the second term, $|f(q_n)|$ is also bounded because q_n is a Palais-Smale sequence, and a calculation reveals that $d|\text{grad } V(q_n)|^2 = 2 \text{Hess } V(q_n)(\text{grad } V(q_n), -)$. We find

$$|\text{grad } f(q_n)| \geq \frac{V_\infty}{\sqrt{1 + V_\infty^2}} - a \frac{\|\text{Hess } V(q_n)\|}{|\text{grad } V(q_n)|} \geq \frac{1}{2} \frac{V_\infty}{\sqrt{1 + V_\infty^2}},$$

for n sufficiently large, since asymptotic regularity gives that $\frac{\|\text{Hess } V(q_n)\|}{|\text{grad } V(q_n)|} \rightarrow 0$ as $n \rightarrow \infty$. This is a contradiction for n large. Hence $d(q_n, K)$ is bounded as $n \rightarrow \infty$. Since M is complete, q_n is contained in a compact set and therefore it

contains a convergent subsequence. We conclude that f satisfies the condition of Palais and Smale.

We use a similar argument to show that f does not have critical values between 0 and $-\nu$, provided ν is small enough. Let $q \in N \setminus N_\nu$, then we have the following estimates with $C_1, C_2 > 0$, independent of q ,

$$|\text{grad } V(q)| > C_1 \quad \text{and} \quad \|\text{Hess } V(q)\| < C_2 |\text{grad } V(q)|.$$

For $d(q, K) > 0$ this directly follows from asymptotic regularity. For $q \in K$ we argue as follows. By regularity of the energy surface $\text{grad } V \neq 0$ on ∂N . Hence the gradient does not vanish on a tubular neighborhood of ∂N . For ν small the strip $(N \setminus N_\nu) \cap K$ lies inside this tubular neighborhood, by compactness. Thus $|\text{grad } V(q)|$ does not vanish on $(N \setminus N_\nu) \cap K$ for ν small. On the compact set $\{q \mid d(q, K) \leq R\}$, $\|\text{Hess } V\|$ is bounded, thus we can find a C_2 such that the second estimate holds.

Hence for all $q \in N - N_\nu$, the norm of $\text{grad } f(q)$ is bounded from below

$$\begin{aligned} |\text{grad } f(q)| &\geq \frac{|\text{grad } V(q)|}{\sqrt{1 + |\text{grad } V(q)|^2}} - \frac{1}{2} |f(q)| \frac{|\text{grad } |\text{grad } V(q)|^2|}{1 + |\text{grad } V(q)|^2}, \\ &\geq \frac{C_1}{\sqrt{1 + C_1^2}} - \nu \frac{|\text{Hess } V(\text{grad } V(q), -)|}{1 + |\text{grad } V(q)|^2}, \\ &\geq \frac{C_1}{\sqrt{1 + C_1^2}} - \nu C_2 > 0. \end{aligned} \tag{6.37}$$

provided $\nu > 0$ is chosen small enough. Thus f does not have critical values between 0 and $-\nu$.

This shows that N is diffeomorphic to N_ν . By considering $-f$ we also see that $M - N_\nu$ is diffeomorphic to $M - N$. The following diagram commutes, by homotopy invariance

$$\begin{array}{ccc} H_k(M \setminus N) & \longrightarrow & H_k(M). \\ \downarrow & \nearrow & \\ H_k(M \setminus N_\nu) & & \end{array} \tag{6.38}$$

Proposition 6.7.4 therefore shows that there exists a non-zero class in $H_{k+1}(M - N_\nu)$ that is mapped to zero in $H_k(M)$ by the morphism induced by the inclusion.

We now estimate V uniformly on balls of radius ρ_ν around points of N_ν . By continuity and compactness, there exists a $\rho_\nu > 0$ such that for all $q \in N_\nu$ with $d(q, K) < 1$, and all $\tilde{q} \in B_{\rho_\nu}(q)$, the estimate $V(\tilde{q}) \leq -\frac{\nu}{2}$ holds. Away from K , i.e. $q \in N_\nu$ with $d(q, K) \geq 1$, the following argument works. Take $\rho_\nu > 0$ small. By the reversed triangle inequality, then for all $\tilde{q} \in B_{\rho_\nu}(q)$, $d(\tilde{q}, K) > 0$. By asymptotic regularity there exists a $C > 0$, independent of q , such that

$$\|\text{Hess } V(\tilde{q})\| \leq C |\text{grad } V(\tilde{q})|. \quad (6.39)$$

If ρ_ν is less than the injectivity radius $\text{inj } M$ of M , there exists a unique geodesic c from q to \tilde{q} parameterized by arc length, with length $\rho'_\nu < \rho_\nu$. One then has:

$$\begin{aligned} \frac{d}{ds} |\text{grad } V(c(s))|^2 &= 2 |\text{Hess } V(\text{grad } V(c(s)), c'(s))|, \\ &\leq 2 \|\text{Hess } V(c(s))\| |\text{grad } V(c(s))| |c'(s)| \leq 2C |\text{grad } V(c(s))|^2. \end{aligned}$$

The above estimate uses Estimate (6.39), and the fact that c is parameterized by arc length. Gronwall's inequality now implies

$$|\text{grad } V(c(s))|^2 \leq |\text{grad } V(c(0))|^2 e^{2Cs}.$$

Taking the square root of the previous equation gives the inequality

$$|\text{grad } V(c(s))| \leq |\text{grad } V(q)| e^{Cs}.$$

This allows us to estimate V along the geodesic c . We compute

$$\begin{aligned} V(\tilde{q}) &= V(q) + \int_0^{\rho'_\nu} \frac{d}{ds} V(c(s)) ds, \\ &= V(q) + \int_0^{\rho'_\nu} \langle \text{grad } V(c(s)), c'(s) \rangle ds, \\ &\leq V(q) + |\text{grad } V(q)| \frac{e^{C\rho'_\nu} - 1}{C}, \\ &\leq -\nu \sqrt{1 + |\text{grad } V(q)|^2} + |\text{grad } V(q)| \frac{e^{C\rho_\nu} - 1}{C}. \end{aligned} \quad (6.40)$$

The function $x \mapsto -\nu \sqrt{1 + x^2} + \frac{e^{C\rho_\nu} - 1}{C} x$ has the maximum $-\sqrt{\nu^2 - \left(\frac{e^{2C\rho_\nu} - 1}{C}\right)^2}$ for $\frac{e^{2C\rho_\nu} - 1}{C} \leq \nu$. We can find $\rho_\nu > 0$ small such that $V(\tilde{q}) \leq -\frac{\nu}{2}$. This is independent of q , because C is. \square




Figure 6.6: Sketches of the domain $\Lambda M \times \mathbb{R}$ of the functional \mathcal{A} . The manifold M is embedded in ΛM by the map sending $q \in M$ to the constant loop $c_q(s) = q$, and hence is embedded in $\Lambda M \times \mathbb{R}$. On the left, the (k) -link in the base manifold, between W and N_v is shown. This link obviously does not persist in the loop space. However, it is possible to *lift* the link to $\Lambda M \times \mathbb{R}$, depicted on the right, to the sets A and B , which $(k + 1)$ -link in $\Lambda M \times \mathbb{R}$.

6.8.2 Constructing linking sets.

We will use the topological assumptions in Theorem 6.1.1, to construct linking subspaces of the loop space. These are in turn used to find candidate critical values of the functional \mathcal{A} . The minimax characterization, cf. Section 6.5, shows that these candidate critical values contain critical points.

By Proposition 6.7.4 and Lemma 6.8.1 there exists a non-zero $[w] \in H_k(M - N_v)$ such that $i_k([w]) = 0$ in $H_k(M)$. In this formula i is the inclusion $i : M - N_v \rightarrow M$, and i_k the induced map in homology of degree k . Because $i_k[w] = 0$, there exists a $u \in C_{k+1}(M)$ such that $\partial u = w$. We disregard any connected component of u that does not intersect w . Set $W = |w|$ and $U = |u|$ where $|\cdot|$ denotes the support of a cycle. Both are compact subspaces of M . The inclusion $H_k(W) \rightarrow H_k(M - N_v)$ is non-trivial by construction. We say that W (k) -links N_v in M . The linking sets discussed above will be used to construct linking sets in the loop space, satisfying appropriate bounds, cf. Proposition 6.9.1. A major part of this construction is carried out by the “hedgehog” function. This function is constructed in section 6.11, and is a continuous map $h : [0, 1] \times U \rightarrow \Lambda M$ with the following properties

- (i) $h_0(U) \subset \mathcal{V}$, with the tubular neighborhood \mathcal{V} defined in Proposition 6.6.2.
- (ii) The restriction $h_t|_W$ is the inclusion of W in the constant loops in ΛM ,




Figure 6.7: The normal bundle of M . The set S is a sphere bundle over N_v .

(iii) Only W is mapped to constant loops. Thus $h_t(q) \in \iota(M)$ if and only if $q \in W$.

(iv) $\int_0^1 V(h_1(q)(s)) ds > 0$ for all $q \in U$.

The idea is that, for $t = 0$, a point $q \in U$ is mapped to a loop close (in H^1 sense) to the constant loop $c_q(s) = q$. Points on the boundary W are mapped to constant loops, but other points are never mapped to a constant loop. This ensures the first three properties (for $t = 0$). The construction shows that there are a finite number of points, such that the loops stay at these points for almost all time. These points are then homotoped to points where the potential is positive. This ensures the last property, using compactness of U . Properties (i) and (ii) are used to lift the link of M to a link in ΛM . The remaining properties are used to deform the link to sets where the functional satisfies appropriate bounds, and show that the link is not destroyed during the homotopy.

Because U is compact, and N_v is closed, $\iota(U \cap N_v)$ is compact. Moreover, it does not intersect $h_t(U)$ for any t , by property (iii). Hence $d_{\Lambda M}(h_{[0,1]}(U), \iota(U \cap N_v)) > 0$. Set $0 < \rho < \min(\frac{\text{inj } M}{2}, \frac{\rho_v}{2})$ such that

$$\rho < \frac{1}{2} d_{\Lambda M}(h_{[0,1]}(U), \iota(U \cap N_v)). \tag{6.41}$$

Define $f : U \rightarrow NM$ by the equation $f(q) = \phi^{-1}h_0(q)$, where $\phi : NM \rightarrow \mathcal{V}$ is defined in Proposition 6.6.2. The restriction of f to W is the inclusion of W




Figure 6.8: The projected normal bundle $\hat{\pi}(NM) = M \times \mathbb{R}$. The set Z and $N_\nu \times \{\rho\}$ $(k + 1)$ link in $M \times \mathbb{R}$.

into the zero section of NM by Property (ii), see also Figure 6.7. Recall that the normal bundle comes equipped with the equivalent norm $\|\cdot\|_\perp$, cf. Proposition 6.6.3. Define the map $\hat{\pi} : NM \rightarrow M \times \mathbb{R}$ by

$$\hat{\pi}(q, \zeta) = (q, \|\zeta\|_\perp).$$

Define $S = \hat{\pi}^{-1}(N_\nu \times \{\rho\})$. This is a sphere sub-bundle of radius ρ in the normal bundle over N_ν . Recall that the inclusion of M as the zero section in NM is denoted by $\zeta : M \rightarrow NM$. Set

$$Z = \hat{\pi}(\zeta(U) \cup f(U)) = U \times \{0\} \cup \hat{\pi}(f(U)).$$

The sets are depicted in Figure 6.8. Because W (k) -links N_ν in M , the set Z $(k + 1)$ -links $\hat{\pi}(S) = N_\nu \times \{\rho\}$ in $M \times \mathbb{R}$, as we prove below.

Lemma 6.8.2. *Suppose $H_k(W) \rightarrow H_k(M \setminus N_\nu)$ is non-trivial. Then*

$$H_{k+1}(Z) \rightarrow H_{k+1}(M \times \mathbb{R} \setminus N_\nu \times \{\rho\}),$$

is non-trivial.

Proof. Recall that $W = |\mathbf{w}|$ and $\mathbf{w} = \partial \mathbf{u}$ with $\mathbf{u} \in C_{k+1}(M)$ a $(k + 1)$ -cycle. Define the cycle $\mathbf{x} \in C_{k+1}(Z)$ by

$$\mathbf{x} = \hat{\pi}_{k+1} \zeta_{k+1}(\mathbf{u}) - \hat{\pi}_{k+1} f_{k+1}(\mathbf{u}),$$




Figure 6.9: To show that the sets link, we need to apply a homotopy to Figure 6.8.

This cycle is closed, because

$$\begin{aligned} \partial x &= \hat{\pi}_k \zeta_k(\partial u) - \hat{\pi}_k f_k(\partial u) \\ &= \hat{\pi}_k \zeta_k(\mathbf{w}) - \hat{\pi}_k f_k(\mathbf{w}) = 0. \end{aligned}$$

In the last step we used that $f|_W = \zeta|_W$. Hence $[x] \in H_{k+1}(Z)$. We show that this class is mapped to a non-trivial element in $H_{k+1}(M \times \mathbb{R} \setminus N_v \times \{\rho\})$

For technical reasons we need to modify Z and $N_v \times \{\rho\}$. Define the set \tilde{Z} , which is depicted in Figure 6.9, by

$$\tilde{Z} = U \times \{0\} \cup W \times [0, \rho] \cup T_\rho(\hat{\pi}f(U)),$$

where $T_\rho : M \times \mathbb{R} \rightarrow M \times \mathbb{R}$ is the translation over ρ in the \mathbb{R} direction, i.e. $T_\rho(q, r) = (q, r + \rho)$. Denote by I_ρ the interval $(\frac{\rho}{3}, \frac{2\rho}{3})$. There exists a homotopy $m_t : M \times \mathbb{R} \rightarrow M \times \mathbb{R}$, with the following properties:

- (i) $m_0 = \text{id}$,
- (ii) $m_t(\tilde{Z}) \cap m_t(N_v \times I_\rho) = \emptyset$, for all t ,
- (iii) $m_1(\tilde{Z}) = Z$,
- (iv) $m_1(N_v \times I_\rho) = N_v \times \{\rho\}$.

These properties ensure that Z $(k + 1)$ -links $N_v \times \{\rho\}$ if and only if \tilde{Z} $(k + 1)$ -links $N_v \times I_\rho$. Define $[\tilde{x}] = (m_1)_{k+1}^{-1}[x] \in H_{k+1}(\tilde{Z})$. This is well defined

6. PERIODIC ORBITS

because $(m_1)_{k+1}$ is an isomorphism. We will reason that this class includes non-trivially in $H_{k+1}(M \times \mathbb{R} \setminus N_\nu \times I_\rho)$. For this we apply Mayer-Vietoris to the triad (\tilde{Z}, U_1, U_2) , with

$$\begin{aligned} U_1 &= U \times \{0\} \cup W \times [0, \frac{2\rho}{3}), \\ U_2 &= W \times (\frac{\rho}{3}, \rho] \cup T_\rho \hat{\pi}f(U). \end{aligned}$$

Note that $U_1 \cap U_2 = W \times I_\rho$. From the Mayer-Vietoris sequence for the triad we get the boundary map

$$H_{k+1}(\tilde{Z}) \xrightarrow{\delta} H_k(W \times I_\rho).$$

By definition of the boundary map δ in the Mayer-Vietoris sequence, we have that $\delta[\tilde{\mathbf{x}}] = (m_1)_k^{-1} \hat{\pi}_k \zeta_k[\mathbf{w}]$. Now we consider a second Mayer-Vietoris sequence, the Mayer-Vietoris sequence of the triad

$$\left(M \times \mathbb{R} \setminus N_\nu \times I_\rho, M \times \mathbb{R}_{>\frac{\rho}{3}} \setminus N_\nu \times I_\rho, M \times \mathbb{R}_{<\frac{2\rho}{3}} \setminus N_\nu \times I_\rho \right).$$

By naturality of Mayer-Vietoris sequences, the following diagram commutes

$$\begin{array}{ccc} H_{k+1}(\tilde{Z}) & \xrightarrow{\delta} & H_k(W \times I_\rho) \\ i_{k+1} \downarrow & & \downarrow i_k \\ H_{k+1}(M \times \mathbb{R} \setminus N_\nu \times I_\rho) & \xrightarrow{\delta} & H_k(M \times I_\rho \setminus N_\nu \times I_\rho). \end{array}$$

We argued that $\delta[\tilde{\mathbf{x}}] = (m_1)_k^{-1} \hat{\pi}_k \zeta_k[\mathbf{w}]$. We have that $i_k(m_1)_k^{-1} \hat{\pi}_k \zeta_k[\mathbf{w}] \neq 0$ by assumption. By the commutativity of the above diagram we conclude that $i_{k+1}[\tilde{\mathbf{x}}] \neq 0$. Thus \tilde{Z} $(k+1)$ -links $N_\nu \times I_\rho$ in $M \times \mathbb{R}$, which implies that Z $(k+1)$ -links $N_\nu \times \{\rho\}$ in $M \times \mathbb{R}$. \square

The previous lemma lifted the link in the base manifold to a link in $M \times \mathbb{R}$. We now lift this link to the full normal bundle.

Lemma 6.8.3. *The fact that $H_{k+1}(Z) \rightarrow H_{k+1}(M \times \mathbb{R} \setminus N_\nu \times \{\rho\})$ is non-trivial implies that $H_{k+1}(\zeta(U) \cup f(U)) \rightarrow H_{k+1}(NM \setminus S)$ is non-trivial.*

Proof. The following diagram commutes

$$\begin{array}{ccc}
 H_{k+1}(\zeta(U) \cup f(U)) & \xrightarrow{\hat{\pi}_{k+1}} & H_{k+1}(Z) \\
 \downarrow i_{k+1} & & \downarrow i_{k+1} \\
 H_{k+1}(NM \setminus S) & \xrightarrow{\hat{\pi}_{k+1}} & H_{k+1}(M \times \mathbb{R} \setminus N_\nu \times \{0\}).
 \end{array}$$

Define $[\mathbf{y}] = \zeta_{k+1}[\mathbf{u}] - f_{k+1}[\mathbf{u}]$. Recall that $[\mathbf{x}] = \hat{\pi}_{k+1}[\mathbf{y}]$ includes non-trivially in $H_{k+1}(M \times \mathbb{R} \setminus N_\nu \times \{\rho\})$ by the construction in lemma 6.8.2. By the commutativity of the above diagram $\pi_{k+1}i_{k+1}[\mathbf{y}] = i_{k+1}\pi_{k+1}[\mathbf{y}] \neq 0$. Thus $i_*[\mathbf{y}] \neq 0$. The inclusion $H_{k+1}(\zeta(U) \cup f(U)) \rightarrow H_{k+1}(NM \setminus S)$ is non-trivial. \square

The domain of \mathcal{A} is not the free loop space ΛM , but $\Lambda M \times \mathbb{R}$. The extra parameter keeps track of the period of the candidate periodic solutions. Thus we need once more to lift the link to a bigger space. In this process we also globalize the link, moving it from the normal bundle to the full free loop space. Recall that we write $E = \Lambda M \times \mathbb{R}$. The subsets $A^t = A_I \cup A_{II} \cup A_{III}^t$ are defined by

$$\begin{aligned}
 A_I &= \phi(\zeta(U)) \times \{\sigma_1\} \\
 A_{II} &= \phi(\zeta(W)) \times [\sigma_1, \sigma_2] \\
 A_{III}^t &= h_t(U) \times \{\sigma_2\}
 \end{aligned}$$

The constants $\sigma_1 < \sigma_2$ will be specified in Proposition 6.9.1. Finally we define the sets $A, B \subset E$ by

$$A = A^1 \quad \text{and} \quad B = \phi(S) \times \mathbb{R}. \tag{6.42}$$

Figure 6.10 depicts the sets A , and B .

Lemma 6.8.4. *Assume that $H_{k+2}(\Lambda M) = 0$. The fact that $H_{k+1}(\zeta(U) \cup f(U)) \rightarrow H_{k+1}(NM \setminus S)$ is non-trivial implies that the inclusion $H_{k+1}(A) \rightarrow H_{k+1}(E \setminus B)$ is non-trivial.*

Proof. By Lemma 6.8.3 the morphism induced by the inclusion $H_{k+1}(\zeta(U) \cup f(U)) \rightarrow H_{k+1}(NM \setminus S)$ is non-trivial. By applying the diffeomorphism ϕ , we see therefore that

$$H_{k+1}(\phi(\zeta(U)) \cup \phi(f(U))) \rightarrow H_{k+1}(\mathcal{V} \setminus \phi(S)),$$

is non-trivial. Proposition 6.6.4 shows that the map

$$H_{k+1}(\mathcal{V}\setminus\phi(S)) \rightarrow H_{k+1}(\Lambda M\setminus\phi(S)),$$

is injective, because we assumed $H_{k+2}(\Lambda M) = 0$, and $\phi(S)$ is closed in \mathcal{V} . It follows that

$$H_{k+1}(\phi(\zeta(U) \cup f(U))) \rightarrow H_{k+1}(\Lambda M\setminus\phi(S)),$$

is non-trivial. Let $\pi_1 : \Lambda M \times \mathbb{R} \rightarrow \Lambda M$ be the projection to the first factor. Because of the choice of ρ , cf. Equation (6.41) the set $\pi_1(A^t)$ never intersects $\pi_1(B)$. By the construction of the sets A^t and B , the map π_1 induces a homotopy equivalence between A^t and $\pi_1(A^t)$ and between $E\setminus B$ and $\Lambda M\setminus\pi_1(B)$, so that the diagram

$$\begin{array}{ccc} H_{k+1}(A^t) & \longrightarrow & H_{k+1}(E\setminus B) \\ (\pi_1)_{k+1} \downarrow & & \downarrow (\pi_1)_{k+1} \\ H_{k+1}(\pi_1(A^t)) & \longrightarrow & H_{k+1}(\Lambda M\setminus\pi_1(B)), \end{array}$$

commutes. We see that $H_{k+1}(A^t) \rightarrow H_{k+1}(E\setminus B)$ is non-trivial if and only if $H_{k+1}(\pi_1(A^t)) \rightarrow H_{k+1}(\Lambda M\setminus\pi_1(B))$ is non-trivial. For all $t \in [0, 1]$ the induced maps are the same, because of homotopy invariance. For $t = 0$ we have that $\pi_1(A^0) = \phi(\zeta(U) \cup f(U))$, and $\Lambda M\setminus\pi_1(B) = \Lambda M\setminus\phi(S)$. We conclude that A $(k+1)$ -links B in ΛM . \square

6.9 Estimates

We need to estimate \mathcal{A} on the sets $A, B \subset E$, defined in Equation (6.42).

Proposition 6.9.1. *If ν and ρ are sufficiently small, then there exist constants $\sigma_1 < \sigma_2$ and $0 < a < b$, such that*

$$\mathcal{A}|_A \leq a \quad \text{and} \quad \mathcal{A}|_B > b. \tag{6.43}$$

Proof. We first estimate \mathcal{A} on $B = \phi(S) \times \mathbb{R}$. Let $(c_1, \tau) \in \phi(S) \times \mathbb{R}$. Then $c_1 = \phi(\xi) = \exp_{c_0}(\xi)$ where ξ is a vector field along a constant loop c_0 at $q \in N_\nu$, for which $\|\xi\|_\perp = \|\nabla \xi\|_{L^2} = \rho$. From the Gauss lemma, and the following estimate, cf. Equations (6.6) and (6.21),

$$\|\xi\|_{C^0} \leq \sqrt{2}\|\xi\|_{H^1} \leq 2\|\xi\|_\perp$$




Figure 6.10: A sketch of the linking sets $A = A_I \cup A_{II} \cup A_{III}^1$ and B . For graphical reasons the base manifold M and the loop space ΛM are depicted one dimensional. This graphical representation therefore does not agree with the dimensionality of Figure 6.7, where the normal bundle of $\iota(M)$ in ΛM was depicted as two dimensional.

we see that $\sup_{s \in \mathbb{S}^1} d_M(c_0(s), c_1(s)) \leq 2\rho$. Recall that we assumed $\rho \leq \frac{\rho_V}{2}$. Hence for all $s \in \mathbb{S}^1$, we have $c_1(s) \in B_{\rho_V}(q)$ and therefore $V(c_1(s)) \leq -\frac{\nu}{2}$, by Lemma 6.8.1. We use this to estimate the second term of

$$\mathcal{A}(c_1, \tau) = \frac{e^{-\tau}}{2} \int_0^1 |\partial c_1(s)|^2 ds - e^\tau \int_0^1 V(c_1(s)) ds. \quad (6.44)$$

Let us now concentrate on the first term. We construct the geodesic from c_0 to c_1 in the loop space, namely

$$c_t(s) = \exp_{c_0}(t\zeta(s)).$$

This can also be seen as a singular surface in M , cf. [Kli95]. Now we apply Taylor's formula with remainder to $t \mapsto \mathcal{E}(c_t)$. There exists a $0 \leq \tilde{t} \leq 1$ such that

$$\mathcal{E}(c_1) = \mathcal{E}(c_0) + \frac{d}{dt} \mathcal{E}(c_t) \Big|_{t=0} + \frac{1}{2} \frac{d^2}{dt^2} \mathcal{E}(c_t) \Big|_{t=0} + \frac{1}{6} \frac{d^3}{dt^3} \mathcal{E}(c_t) \Big|_{t=\tilde{t}}. \quad (6.45)$$

It is obvious that $\mathcal{E}(c_0) = 0$, since c_0 is a constant loop. Because $t \mapsto c_t$ is a geodesic $\frac{d}{dt} \Big|_{t=0} \mathcal{E}(c_t) = 0$. The second order neighborhood of a closed geodesic

is well studied [Kli95, Lemma 2.5.1]. We see that c_0 is a (constant) closed geodesic, therefore

$$\frac{d^2}{dt^2} \Big|_{t=0} \mathcal{E}(c_t) = D^2 \mathcal{E}(c_0)(\zeta, \zeta) = \|\zeta\|_{\perp}^2 = \rho^2.$$

The curvature term in the second variation vanishes at $t = 0$, because c_0 is a constant loop. The third derivative of the energy functional can be bounded in terms of the curvature tensor and its first covariant derivative times a third power of $\|\zeta\|_{\perp}$. By the assumption of bounded geometry, we can therefore uniformly bound $\mathcal{E}(c_1)$. The main point is that for ρ sufficiently small, $\mathcal{E}(c_1) \geq C\rho^2$, for some constant $C > 0$. We now can estimate \mathcal{A} on B .

$$\mathcal{A}(c_1, \tau) \geq \frac{e^{-\tau}}{2} C\rho^2 + \frac{e^{\tau}}{2} \nu \geq \sqrt{C\nu} \rho \quad (6.46)$$

Set $b = \sqrt{C\nu} \rho$, then $\mathcal{A}|_B > b$. It remains to estimate \mathcal{A} on the set $A = A_I \cup A_{II} \cup A_{III}^1$. Let $(c, \sigma_1) \in A_I = \phi(\zeta(U)) \times \{\sigma_1\}$. Recall that U is compact, hence $V_{\max} = \sup_{q \in U} -V(q) < \infty$. Because c is a constant loop, we find

$$\mathcal{A}(c, \sigma_1) = -e^{\sigma_1} \int_0^1 V(c(s)) ds \leq e^{\sigma_1} V_{\max}. \quad (6.47)$$

By choosing $\sigma_1 \leq \log(\frac{b}{2V_{\max}})$ we get $\mathcal{A}|_{A_I} \leq b/2$. On $A_{II} = \phi(W) \times [\sigma_1, \sigma_2]$ all the loops are constants as well, moreover their image is contained in W . The potential is positive on W hence $\mathcal{A}|_{A_{II}} < 0 < \frac{b}{2}$. It remains to estimate \mathcal{A} on $A_{III}^1 = h_1(U) \times \{\sigma_2\}$. Recall that we constructed h in such a way that for any $q \in U$ we have $\int_0^1 V(h_1(q)(s)) ds > 0$. This gives

$$\begin{aligned} \mathcal{A}(c, \sigma_2) &= \frac{e^{-\sigma_2}}{2} \int_0^1 |\partial h_1(q)(s)|^2 ds - e^{\sigma_2} \int_0^1 V(h_1(q)(s)) ds \\ &\leq \frac{e^{-\sigma_2}}{2} \int_0^1 |\partial h_1(q)(s)|^2 ds. \end{aligned} \quad (6.48)$$

Because h is continuous and U is compact, $\mathcal{E}_{\max} = \sup_{q \in U} \mathcal{E}(h_1(q)) < \infty$. And therefore

$$\mathcal{A}(c, \sigma_2) \leq \frac{e^{-\sigma_2}}{2} \mathcal{E}_{\max}. \quad (6.49)$$

By setting $\sigma_2 > \max(\log(\frac{\mathcal{E}_{\max}}{b}), \sigma_1)$ we get $\mathcal{A}|_{A_{III}^1} \leq \frac{b}{2}$. Now set $a = b/2$, and we see that $\mathcal{A}|_A < a < b$. \square

6.10 Proof of the main theorem

Proof of Theorem 6.1.1. From the assumptions $H_{k+1}(\Sigma) \neq 0$ and $H_{k+1}(\Lambda M) = H_{k+2}(\Lambda M) = 0$, we are able to construct linking sets A and B in the loop space, cf. Lemma 6.8.4. We estimate \mathcal{A} on A and B in Proposition 6.9.1. This gives rise to estimates for the penalized functionals \mathcal{A}_ϵ and this in turn gives the existence of critical points, cf. Proposition 6.5.2. Under the assumption of flat ends, Proposition 6.4.8 produces a critical point of \mathcal{A} , by taking the limit $\epsilon \rightarrow 0$. The critical point of \mathcal{A} corresponds to a closed characteristic on Σ . \square

6.11 Appendix: construction of the hedgehog

In this section we construct the hedgehog function $h : [0, 1] \times U \rightarrow \Lambda M$. Recall that this function needs to satisfy the following properties:

- (i) $h_0(U) \subset \mathcal{V}$, with the tubular neighborhood \mathcal{V} defined in Proposition 6.6.2.
- (ii) The restriction $h_t|_W$ is the inclusion of W in the constant loops in ΛM ,
- (iii) Only W is mapped to constant loops. Thus $h_t(q) \in \iota(M)$ if and only if $q \in W$.
- (iv) $\int_0^1 V(h_1(q)(s)) ds > 0$ for all $q \in U$.

The precise construction of the hedgehog is technically involved, but the idea behind it is simple. It is based on a similar construction in [vdBPV09], but we have to construct the map locally, which adds some subtleties. Points in W are mapped to constant loops, and, for $t = 1$, points in $U \setminus W$ are mapped to non-constant loops which remain for most time at points where the potential is positive.

When a space admits a triangulation, simplicial and singular homology coincide. We will exploit this fact by constructing our map simplex by simplex in simplicial homology. It is possible to find a smooth triangulation of M , which restricts to a smooth triangulation of N_ν [Mun66, Theorem 10.4], and by subdivision we can make this triangulation as fine as we want, i.e. we assume that each simplex is contained in a single chart, and that the diameter (measured with respect to the Riemannian distance) of each simplex is bounded uniformly, by $0 < \frac{\epsilon_1}{2} < \text{inj } M$. Recall that U and W are images of cycles, and as such are triangulated and compact.

The hedgehog is constructed by the following procedure. First we construct h_0 . Each point in U is mapped to a loop whose image is in the simplex it is contained in, or in small spines emanating from the corners. This is done in such a manner for most times s , the loop is locally constant with values at the corner points and the spines. By starting with a fine triangulation this ensures that this map satisfies property (i). The points on W are mapped to constant loops, which ensures property (ii). We also ensure that these are the only points which are sent to constant loops, which enforces (iii). Finally we apply a homotopy $f_t : M \rightarrow M$, which is the identity outside a neighborhood of U , and maps the 0-simplices in the triangulation that are not in W to points in $M - (N \cup W)$. The potential V is positive on this set. We then define $h_t = f_t \circ h_0$, and this, together with the property that the loops are for most times at these corner points, makes sure we satisfy property (iv).

For the construction we need some spaces which closely resemble the standard simplex.

Definition 6.11.1. The *standard m -simplex* Δ_m is defined by

$$\Delta_m = \left\{ y = (y_0, \dots, y_m) \in \mathbb{R}_{\geq 0}^{m+1} \mid \sum_{i=0}^m y_i = 1 \right\}, \quad (6.50)$$

the *extended m -simplex* \square_m by

$$\square_m = \left\{ y = (y_0, \dots, y_m) \in \mathbb{R}_{\geq 0}^{m+1} \mid 1 \leq \sum_{i=0}^m y_i \leq 2 \right\}, \quad (6.51)$$

and the *m -simplex with spines* by

$$\triangle_m = \Delta_m \bigcup_{j=0}^m \left\{ y \in \mathbb{R}_{\geq 0}^{m+1} \mid 1 \leq y_j \leq 2 \text{ and } y_{i \neq j} = 0 \right\}. \quad (6.52)$$

These simplices are sketched in figure 6.11. We can now make the previous discussion more precise. Denote by $\mathcal{T}_U = \{l_m : \Delta_m \rightarrow M\}$ the triangulation of U . Because U is a $(k+1)$ -cycle, \mathcal{T}_U only contains simplices of dimension less or equal to $k+1$. To make consistent choices in the interpolations to come, we need to fix a total order $<$ on the zero simplices. All the other simplices respect this order, in the sense that for any simplex $l_k : \Delta_k \rightarrow M$ we have

$$l_k(1, 0, \dots, 0) < l_k(0, 1, 0, \dots, 0) < \dots < l_k(0, \dots, 0, 1).$$




Figure 6.11: The three types of simplices are depicted for $m = 2$.

Now we start with some fixed top-dimensional, i.e. $(k + 1)$ -dimensional, simplex in the triangulation \mathcal{T}_U , say $l_{k+1} : \Delta_m \rightarrow M$. We define the map \tilde{h}_{k+1} , which constructs initial non-constant loops on the vertices not in W , and interpolates between them. The images of these loops lie in the extended simplex. The loops produced in this manner are not smooth, only of class L^2 . Therefore we need to smooth these loops, by taking the convolution \star with a standard mollifier ζ_{ϵ} , which lands us in $C^\infty(\mathbb{S}^1, \square_m)$. This is possible because the extended simplex is a convex set. The map Φ contracts the extended simplex to the simplex with spines. By post-composing the loops with this contraction, denoted $\Phi \cdot -$, we get loops in the simplex with spines. This simplex with spines is mapped to the free loop space of M using l_{k+1} . This is a slight modification of the simplex. The spines are mapped to paths close-by the corner points. This finalizes the construction of h_0 restricted to the simplex. After this we homotope M moving all the 0-simplices of the triangulation to points with positive potential, and not in W . The construction is such that on intersection of simplices in the triangulation of U the maps coincide. Therefore these maps can be patched together to make a continuous map from $U \rightarrow \Lambda M$. For reference, the points in the simplex l_{k+1} will eventually be mapped to ΛM by the following sequence

$$\begin{array}{lcl}
 l_{k+1}(\Delta_{k+1}) & \xrightarrow{l_{k+1}^{-1}} & \Delta_{k+1} \xrightarrow{\tilde{h}_{k+1}} L^2(\mathbb{S}^1, \square_{k+1}) \xrightarrow{\zeta_\epsilon \star} C^\infty(\mathbb{S}^1, \square_{k+1}) \\
 & & \xrightarrow{\Phi \cdot -} H^1(\mathbb{S}^1, \odot_{k+1}) \xrightarrow{\tilde{l}_{k+1} \cdot -} \Lambda M \xrightarrow{f_{i \cdot -}} \Lambda M.
 \end{array}$$

The precise details on the construction of the maps in this diagram follows.

The interpolation operator I_{Δ_m}

We will construct the loops inductively. We initially set loops at the vertices in the next paragraph. We want to interpolate them to the 1 dimensional faces. These are then in turn interpolated to the 2 dimensional faces, etc. Thus we need to construct a continuous interpolation operator $I_{\Delta_m} : C(\partial\Delta_m, L^2(\mathbb{S}^1, \square_m)) \rightarrow C(\Delta_m, L^2(\mathbb{S}^1, \square_m))$. We define the loop at the center of the simplex to be the concatenation of loops at the corner points. Each interior point on the simplex lies on a unique line segment from the barycenter to a boundary point. We can parameterize this line segment by a parameter $0 \leq \lambda \leq 1$. A point on the line segment corresponding to the parameter λ is mapped to the loop which follows the boundary loop up to time $s = \lambda$ and then follows the loop of the barycenter of the simplex, for time $\lambda < s < 1$.

Fix m and consider the barycenter $\mathcal{M} = (\frac{1}{m+1}, \dots, \frac{1}{m+1})$ of the m -simplex. Set, for $y = (y_1, \dots, y_m) \in \Delta_m - \mathcal{M}$,

$$\lambda_y = \min_{\substack{0 \leq i \leq m \\ y_i < 1/(m+1)}} \frac{1}{1 - (m+1)y_i}. \quad (6.53)$$

This is the smallest positive value λ_y such that $\mathcal{M} + \lambda_y(y - \mathcal{M})$ is on the boundary, cf. Figure 6.12. Now assume we have a continuous map $k : \partial\Delta_m \rightarrow L^2(\mathbb{S}^1, \square_m)$. Interpolate this map to the barycenter by

$$(I_{\Delta_m} k)(\mathcal{M})(s) = \begin{cases} k(1, 0, \dots, 0) & \text{if } 0 \leq s \leq \frac{1}{m} \\ k(0, 1, 0, \dots, 0) & \frac{1}{m} \leq s \leq \frac{2}{m} \\ \dots & \dots \\ k(0, \dots, 0, 1) & \frac{m-1}{m} \leq s \leq 1 \end{cases}, \quad (6.54)$$

and interpolate it to the full simplex Δ_m by:

$$(I_{\Delta_m}(k))(y)(s) = \begin{cases} k(\mathcal{M} + \lambda_y(y - \mathcal{M}))(s) & \text{if } 0 \leq s \leq 1/\lambda_y \\ k(\mathcal{M})(s) & 1/\lambda_y < s \leq 1, \end{cases} \quad (6.55)$$

for all y in the interior of Δ_m but not equal to \mathcal{M} . If y is on the boundary, then $I_{\Delta_m}(k)(y)(s) = k(y)(s)$. If k was continuous, we have that $I_{\Delta_m}(k) \in C(\Delta_m, L^2(\mathbb{S}^1, \square_m))$.




Figure 6.12: The interpolation operator interpolates between the barycenter \mathcal{M} and the boundary.

The map \tilde{h}_{k+1}

Recall that we have chosen a fixed top dimensional, that is $(k + 1)$ -simplex $l_{k+1} : \Delta_{k+1} \rightarrow M$ in \mathcal{T}_U . For each lower dimensional face of Δ_{k+1} the image of its interior is either disjoint from W or it is contained in W . We say that the face is *in* W in the latter case. To construct \tilde{h}_{k+1} we assign non-constant initial loops to the 0-dimensional faces not in W . The 0-dimensional faces in W are mapped to constant loops. Then we interpolate, using I_{Δ_1} , on all the 1-dimensional faces not in W . On the 1-dimensional faces in W we will have constant loops. This process can be done on higher dimensional faces by induction. The resulting map is $\tilde{h}_{k+1} : \Delta_{k+1} \rightarrow L^2(\mathbb{S}^1, \square_{m+1})$. Let us now formalize the discussion.

To start, let us define \tilde{h}_{k+1} on the vertices in Δ_{k+1} . For a vertex $y = (0, \dots, 1, \dots, 0)$ set

$$\tilde{h}_{k+1}(y)(s) = \begin{cases} y & \text{if } y \text{ is in } W \text{ or } 0 \leq s \leq \frac{1}{2} \\ 2y & \text{if } y \text{ is not in } W \text{ and } \frac{1}{2} \leq s \leq 1 \end{cases}. \quad (6.56)$$

Now suppose \tilde{h}_{k+1} is defined for all m dimensional faces with $m < k + 1$. We now define \tilde{h}_{k+1} on the $(m + 1)$ -dimensional faces. If the $(m + 1)$ -dimensional face is in W set $\tilde{h}_{k+1}(y)(s) = y$ for y in this simplex. For each $(m + 1)$ -dimensional face not in W we will use the interpolation operator $I_{\Delta_{m+1}}$ in the following manner.

We have standard projections $\pi : \mathbb{R}^{(k+1)+1} \rightarrow \mathbb{R}^{(m+1)+1}$, which drops $k - m$ variables, and inclusions $\iota : \mathbb{R}^{(m+1)+1} \rightarrow \mathbb{R}^{(k+1)+1}$, which set $k - m$

coordinates to zero. There are many such maps, and we need them all, but we denote them by the same symbol. These maps are linear maps that respect the total order $<$. Now if \tilde{h} has been defined on the m -dimensional faces of Δ_{k+1} , then for an $(m + 1)$ -dimensional face

$$\pi \circ \tilde{h} \circ \iota|_{\partial\Delta_{m+1}} \in C(\partial\Delta_{m+1}, L^2(\mathbb{S}^1, \square_{m+1})), \quad (6.57)$$

is a map we can apply $I_{\Delta_{m+1}}$ to. This gives a continuous map in $C(\Delta_{m+1}, L^2(\mathbb{S}^1, \square_{m+1}))$. By post composing with ι we can map this back to Δ_{k+1} . Doing this to all $(m + 1)$ -dimensional faces, gives us a continuous map on the $(m + 1)$ -dimensional faces. Because the interpolation is a continuous operator, we can proceed by induction to obtain the full map $\tilde{h}_{k+1} : \Delta_{k+1} \rightarrow L^2(\mathbb{S}^1, \square_{k+1})$.

The loops constructed in this way are locally constant a.e. with values at the special points $(0, \dots, 0, 1, 0, \dots, 0)$ and $(0, \dots, 0, 2, 0, \dots, 0)$, or points that are mapped to W . Furthermore, if $y \in \Delta_{k+1}$ is not in W , then $\tilde{h}_{k+1}(y)$ attains at least two different values on sets with non-zero measure.

The convolution $\zeta_{\hat{\epsilon}} \star$

The set \square_{k+1} is a convex subset of \mathbb{R}^{n+1} so we can apply to each loop we got from \tilde{h}_{k+1} the convolution with a standard mollifier $\zeta_{\hat{\epsilon}}$ with parameter $\hat{\epsilon}$ to obtain smooth loops in \square_{k+1} . The smooth loops obtained in this manner are constant at the corner points, on a set of measure at least $1 - 2(k + 1)\hat{\epsilon}$. The loops that were constant before mollifying, i.e. those loops that arose from points in W , are left unchanged by this process. No non-constant loop is mapped to a constant loop, if $\hat{\epsilon}$ is chosen small enough.


The retraction Φ

We now construct a retraction from \square_{k+1} to \triangle_{k+1} . A sketch of Φ is depicted in Figure 6.13. Let $y \in \square_{k+1}$. Denote by $2\text{ndmax}(y)$ the second largest component of y . Define

$$\mu_y = \min(2\text{ndmax}(y), \frac{1}{k+2} (\sum_{i=0}^{k+1} y_i - 1)).$$

Then the map $\Phi : \square_{k+1} \rightarrow \triangle_{k+1}$ is defined component-wise by

$$\Phi^i(y) = \max(0, y_i - \mu_y).$$


 Figure 6.13: The map Φ retracts \square_m to \triangle_m .

Note, that if $y \in \square_{k+1}$ then $\Phi(y) \in \triangle_{k+1}$. Post composing an H^1 -loop in \square_{k+1} with Φ produces an H^1 loop with image in \triangle_{k+1} .

The spine maps \tilde{l}_{k+1}

For each 0-simplex l_0^i in $\mathcal{T}_U - \mathcal{T}_W$, with image x_i choose a point \tilde{x}_i close by. To be precise assume that $d_M(x_i, \tilde{x}_i) < \frac{\epsilon_1}{2}$. For each i there exists a unique geodesic γ_{x_i} in the simplex with $\gamma_{x_i}(0) = x_i$ and $\gamma_{x_i}(1) = \tilde{x}_i$. Without loss of generality we can choose \tilde{x}_i such that the geodesics γ_{x_i} don't intersect one another. For the 0-simplices in W , i.e. $l_0^i \in \mathcal{T}_W$, with images x_i we choose the points $\tilde{x}_i = x_i$ and $\gamma_{x_i}(s) = x_i$ is the constant geodesic at x_i .

We have fixed a top dimensional simplex $l_{k+1} : \triangle_{k+1} \rightarrow M$ in $\mathcal{T}_U - \mathcal{T}_W$. The map \tilde{l}_{k+1} is defined by

$$\tilde{l}_{k+1}(y) = \begin{cases} l_{k+1}(y) & \text{if } y \in \triangle_{k+1} \\ \gamma_{l_{k+1}(0, \dots, 1, \dots, 0)}(y_i - 1) & \text{if } y = (0, \dots, y_i, \dots, 0) \end{cases} \quad (6.58)$$

On the simplex itself \tilde{l}_{k+1} is just l_{k+1} , but on the spines it follows the geodesics $\gamma_{l_{k+1}(0, \dots, 1, \dots, 0)}$ emanating from the images of the corners of the $(k+1)$ -simplex constructed above. This map is smooth in the following sense. The restriction to the interior, or the restriction to the interior of \triangle_{k+1} is smooth, because we started with a smooth triangulation. The restriction to the spines is also smooth, because the geodesics are.




Figure 6.14: The spine map \tilde{l}_{k+1} . The spines of Δ_{k+1} are mapped to geodesics in U

Patching

If we apply the above maps after each other we obtain a map $l_{k+1}(\Delta_{k+1}) \rightarrow \Lambda M$. If we have another simplex l'_{k+1} , the construction coincides on common lower dimensional faces, for which we use the total order $<$. The simplices are closed, hence the maps patch together to a continuous map $h_0 : U \rightarrow \Lambda M$. The image of a point $q \in U$ is in H^1 , because Φ is the only H^1 map and all other maps in the construction are piecewise smooth.

Tubular neighborhood

We need to argue that h_0 lies in the tubular neighborhood \mathcal{V} of M in ΛM . For this it is sufficient to show that the H^1 distance of a loop $c_1 := h_0(q) \rightarrow \Lambda M$ and a constant loop $c_0 \in \iota(M)$ is as small as we want, see proposition 6.6.2. We consider the fixed simplex l_{k+1} , and a point q in its interior, and denote by c_0 the constant loop at q . Because the image of c_1 is contained in the simplex l_{k+1} and the geodesics γ_{x_i} emanating from its corners points (which are maximally $\epsilon_1 = \frac{\epsilon_1}{2} + \frac{\epsilon_1}{2}$ away from q , we can estimate that $d_\infty(c_0, c_1) \leq \epsilon_1$. Because the loops are only non-constant on a set of measure at most $2(k+1)\hat{\epsilon}$, and the derivate of the loop can be estimated in terms of ϵ_1 we argue that the energy of c_1 is also bounded from above, by ϵ_2 , which is small if ϵ_1 is small. The following lemma shows that the image of h_0 is contained in \mathcal{V} .

Lemma 6.11.2. *Suppose c_0 is a constant loop and $c_1 \in \Lambda M$ is such that $d_\infty(c_0, c_1) < \epsilon_1 < \text{inj } M$ and $\mathcal{E}(c_1) < \epsilon_2$, then $\|\tilde{\zeta}\|_{H^1} \leq \sqrt{\epsilon_1^2 + 2\epsilon_2}$.*

Proof. Because $d_\infty(c_0, c_1) < \epsilon_1 < \text{inj } M$ we can write $c_1(s) = \exp_{c_0} \zeta(s)$ with $\zeta \in T_{c_0} \Lambda M$. We expand $\mathcal{E}(c_t)$ along the geodesic $c_t = \exp_{c_0}(t\zeta)$, cf. Equation (6.45). Because c_0 is constant, $\mathcal{E}(c_0) = 0$, and because c_t is a geodesic on the loop space $\frac{d}{dt}\big|_{t=0} \mathcal{E}(c_t) = 0$. Thus

$$\mathcal{E}(c_1) = D^2E(\zeta, \zeta) + o(\|\zeta\|_{H^1}^2).$$

We see that for small enough loops $D^2E(\zeta, \zeta) \leq 2\epsilon_2$. But at a constant loop $D^2E(\zeta, \zeta) = \|\nabla \zeta\|_{L^2}^2$. Thus, by the inequality $\|\zeta\|_{L^2} \leq \|\zeta\|_\infty$, Equation (6.6), we have that

$$\|\zeta\|_{H^1}^2 = \|\zeta\|_{L^2}^2 + \|\nabla \zeta\|_{L^2}^2 \leq \epsilon_1^2 + 2\epsilon_2.$$

□

Because we control the size of ϵ_1 , and therefore of ϵ_2 we can ensure h_0 lands in \mathcal{V} .

The homotopy f_t

For each 0-simplex l_0^i in $\mathcal{T}_U \setminus \mathcal{T}_W$ construct paths starting at x_i and \tilde{x}_i , and ending at points in $M \setminus (N \cup W)$, making sure that these paths do not intersect. Define $f_t : M \rightarrow M$, to be the identity outside small enough balls around the x_i 's, and \tilde{x}_i 's. For all points inside these balls, trace a small tube along the path constructed, cf. Figure 6.15, making sure that $f_1(x_i) \neq f_1(\tilde{x}_i)$ and $f_1(x_i), f_1(\tilde{x}_i) \in M \setminus N \cup W$. There are no topological obstructions to this construction. To see this, recall that we assumed that all connected components of U intersect W , and we could construct the paths lying in U .

Estimates

Define $h_t = f_t \circ h_0$. The set U is compact, h_1 is a continuous function hence there exists a constant such that $V(g_1(u)(s)) > -C$ for some $C > 0$. Define \tilde{C} such that $V(f_1(x_i)), V(f_1(\tilde{x}_i)) > \tilde{C} > 0$. If the mollifying parameter $\hat{\epsilon} > 0$ has been chosen sufficiently small the loop $g_1(u)$ is locally constant on a set of measure at least $1 - 2(k+1)\hat{\epsilon}$, with values in $\bigcup_i V(f_1(x_i)) \cup V(f_1(\tilde{x}_i))$. Thus we can bound

$$\int_0^1 V(g_1(u)(s)) ds \geq (1 - 2(k+1)\hat{\epsilon})\tilde{C} - 2(k+1)\hat{\epsilon}C > 0, \quad (6.59)$$

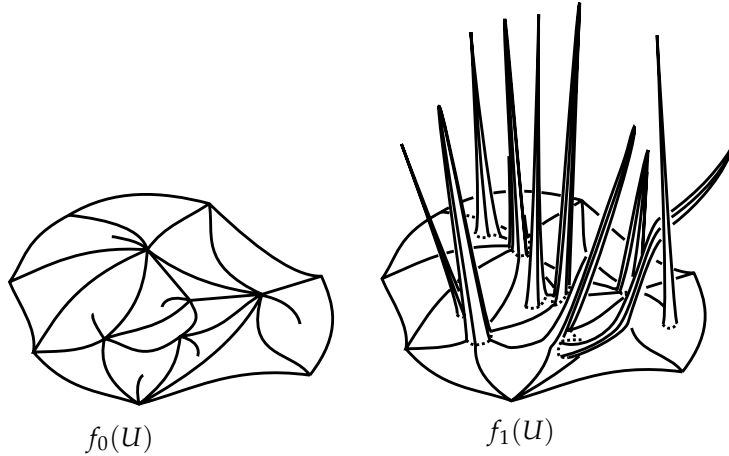


Figure 6.15: The homotopy f_t is depicted for $t = 0$, where it is the identity, and $t = 1$. The zero simplices x_i and the endpoints of the spines \tilde{x}_i have been dragged to points where the potential is positive.

if $\hat{\epsilon}$ is chosen sufficiently small. If $u \in \mathcal{T}_W$ then $g_1(u)(s) = u \in W$ and $V|_W > 0$. Therefore h_t satisfies property (iv). Remark that $h_t(u)$ attains at least two different values if $u \notin \mathcal{T}_W$, and is thus never constant, so that we satisfy (iii). This completes the construction of the hedgehog function.

6.12 Appendix: local computations

In this section we show that equation (6.3) is valid. The computation is done in local coordinates. We use Einsteins summation convention, and we denote by ∂_i the i -th partial derivative with respect to the coordinates on the base manifold. Similarly the the metric tensor g^{ij} and the Christoffel symbols $\Gamma_{ij}^k = \frac{1}{2}g^{kl}(\partial_i g_{kj} + \partial_j g_{ik} - \partial_l g_{ij})$ refer to those on the base manifold M . We omit the subscript q in θ_q and write $(q, \theta) \in T_q M$. In local coordinates the Hamiltonian is expressed as

$$H(q, \theta) = \frac{1}{2}g^{ij}\theta_i\theta_j + V(q),$$

and the Hamiltonian vector field is

$$X_H = \left(-\partial_i + \Gamma_{il}^k g^{lj} \theta_k \theta_j \right) \frac{\partial}{\partial \theta_i} + g^{ij} \theta_j \frac{\partial}{\partial q^i}.$$

To relieve notational burden we introduce the function $h = -\frac{1}{1+|\text{grad } V|^2}$ so

$$f = \theta_i v^i = h \theta_i g^{ij} \partial_j V.$$

Then

$$X_H(f) = h g^{ij} \partial_j V \left(-\partial_i V + \Gamma_{il}^k g^{lj} \theta_k \theta_j \right) + \theta_i g^{kl} \theta_l \partial_k \left(h g^{ij} \partial_j V \right). \quad (6.60)$$

If we rewrite this, relabeling dummy indices, we arrive at

$$X_H(f) = -h g^{ij} \partial_i V \partial_j V + \quad (6.61)$$

$$h \theta_i \theta_j g^{lj} \left(\Gamma_{kl}^i g^{km} \partial_m V + \partial_l (g^{ik} \partial_k V) \right) + \quad (6.62)$$

$$\theta_i \theta_j g^{il} \partial_l V g^{ki} \partial_k h. \quad (6.63)$$

The first term (6.61) is equal to $\frac{|\text{grad } V|^2}{1+|\text{grad } V|^2}$, while the third term (6.63) can be expressed as follows

$$\begin{aligned} \theta_i \theta_j g^{il} \partial_l V g^{ki} \partial_k h &= \theta(\text{grad } V) \theta(\text{grad } h) \\ &= \theta(\text{grad } V) \theta(h^2 \text{grad } |\text{grad } V|^2) \\ &= \frac{2\theta(\text{grad } V) \text{Hess } V(\text{grad } V, \#\theta)}{(1 + |\text{grad } V|^2)^2}. \end{aligned}$$

We now focus on the second term, Equation (6.62). If we write out the derivative, and we relabel some dummy indices, we obtain

$$\begin{aligned} h \theta_i \theta_j g^{lj} \left(\Gamma_{kl}^i g^{km} \partial_m V + \partial_l (g^{ik} \partial_k V) \right) &= h g^{ik} \theta_i g^{lj} \theta_j (\partial_l \partial_k V - \Gamma_{lk}^m \partial_m V) \\ &= h \text{Hess } V(\#\theta, \#\theta) = \frac{-\text{Hess } V(\#\theta, \#\theta)}{1 + |\text{grad } V|^2}. \end{aligned}$$

Combining all the terms gives Equation (6.3).

Chapter 7

Nederlandse samenvatting

Schandere opmerking:

' $E = mc^2$ '

Ja, Albert Einstein -

Die kende zijn vak

En zijn betoog was heel

Argumentatierijk

(Blijft hier achterwege

Voor het gemak)

Formule - Marjolein Kool & Drs. P

7.1 Overzicht

In dit proefschrift worden twee ogenschouwelijk onafhankelijke problemen aangepakt. De samenhang tussen de onderwerpen is Morse theorie. In Deel I wordt een nieuwe homologietheorie ontwikkeld voor stromingen op eindigdimensionale variëteiten, en in Deel II wordt de existentie van periodieke banen op niet-compact energie hyperoppervlakken onderzocht.

We gaan nu wat dieper in op de resultaten.

7.2 Deel I: Morse-Conley-Floer homologie

Een *stroming* is een afbeelding $\phi : \mathbb{R} \times M \rightarrow M$ zodat $\phi(t, \phi(s, x)) = \phi(t + s, x)$ en $\phi(x, 0) = x$ voor alle $x \in M$ en $t, s \in \mathbb{R}$. Stromingen komen natuurlijk voor:

Onder milde technische voorwaarden zijn dit precies de oplossingen van de autonome differentiaalvergelijkingen $x' = X(x)$ voor vectorvelden X . Een belangrijk studieobject voor stromingen zijn de *invariante verzamelingen*. Dit zijn verzamelingen $S \subset M$ zodat $\phi(S, t) = S$ voor alle $t \in \mathbb{R}$. Voorbeelden van zulke verzamelingen zijn bijvoorbeeld evenwichten, periodieke banen en (chaotische) attractoren.

In dit proefschrift wordt een homologietheorie –Morse-Conley-Floer homologie– ontwikkeld voor *geïsoleerde* invariante verzamelingen. Dit zijn invariante verzamelingen die gevangen worden in omgevingen die geen grotere invariante verzamelingen bevatten. Een geïsoleerde invariante verzameling laat altijd een *Lyapunov functie* toe. Zo'n functie daalt langs de oplossingen buiten de invariante verzameling. De Morse-Conley-Floer homologie is grofweg gedefinieerd als de *Morse* homologie van een kleine verstoring van een Lyapunov functie.

Een hoofdresultaat van dit proefschrift, beschreven in Hoofdstuk 2, is dat dit een goede definitie is. Het hangt –op uniek isomorfisme na– niet af van de keuzes die gemaakt zijn om de Morse-Conley-Floer homologie te definiëren. De homologie is verder stabiel in de zin dat het invariant is onder continuaties. Deze stabiliteit kan gebruikt worden om de homologie uit te rekenen in niet-triviale situaties.

In Hoofdstuk 3 beschrijven de functoriële eigenschappen van de homologie theorie. Natuurlijke afbeeldingen tussen variëteiten met stromingen zijn de zogenaamde equivariante afbeeldingen. Deze afbeeldingen verstrengelen de beide stromingen en induceren afbeeldingen in Morse-Conley-Floer homologie. Dit kan gebruikt worden om formele berekeningen te doen over het mogelijke bestaan van bepaalde stromingen, equivariante afbeeldingen en geïsoleerde invariante verzamelingen.

In Hoofdstuk 5 bestuderen we gedegenererde *gradient-systemen* en *Morse decomposities* van algemene stromingen. Het hoofdresultaat is hier het bestaan van een *spectraalrij* wat een relatie geeft tussen Morse-Conley-Floer homologie van de geïsoleerde invariante verzamelingen en de globale homologie van de onderliggende ruimte.

7.3 Deel II: Periodieke banen op niet compacte hyperoppervlakken

In Deel II behandelen we de existentie van periodieke banen voor een klasse van Hamiltoniaanse dynamische systemen. Hamiltoniaanse systemen

beschrijven klassieke mechanische bewegingen zonder wrijving. De zoektocht naar het bestaan van zulke periodieke oplossingen is altijd een drijfveer geweest in het onderzoek in de symplectische en contact meetkunde. Bijna alle resultaten die behaald zijn gaan over systemen waar de energieoppervlakken begrensd (compact) zijn. In Hoofdstuk 6 geven we een klasse van Hamiltoniaanse systemen waar de energie oppervlakken niet compact hoeven te zijn, maar waar er toch altijd een periodieke baan aanwezig is. De methode die gebruikt wordt om dit te bewijzen is klassiek. Het probleem wordt geformuleerd als een vraag naar de existentie van kritieke punten van de Lagrangiaanse actiefunctie op de lussenruimte van de onderliggende ruimte. Dat deze functionaal niet triviale kritieke punten heeft wordt bewezen met behulp van een verstrengelings methode. In de lussenruimte worden verzamelingen gevonden die verstrengeld zijn, en waar de functionaal een bepaald gedrag vertoont. Dit forceert kritieke punten van de functionaal, en daarmee de existentie van periodieke banen op het gevraagde energieoppervlak.

Bibliography

- [AB83] M F Atiyah and R Bott. The Yang-Mills equations over Riemann surfaces. *Philosophical Transactions of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 308(1505):523–615, 1983.
- [AB95] D M Austin and P J Braam. Morse-Bott theory and equivariant cohomology. In *The Floer memorial volume*, pages 123–183. Birkhäuser, Basel, 1995.
- [Abb97] Alberto Abbondandolo. A new cohomology for the Morse theory of strongly indefinite functionals on Hilbert spaces. *Topological Methods in Nonlinear Analysis*, 9(2):325–382, 1997.
- [AD14] Michèle Audin and Mihai Damian. *Morse theory and Floer homology*. Universitext. Springer, London; EDP Sciences, Les Ulis, 2014.
- [AM01] Alberto Abbondandolo and Pietro Majer. Morse homology on Hilbert spaces. *Communications on Pure and Applied Mathematics*, 54(6):689–760, 2001.
- [AM04] Alberto Abbondandolo and Pietro Majer. When the Morse index is infinite. *International Mathematics Research Notices*, (71):3839–3854, 2004.
- [AM05] Alberto Abbondandolo and Pietro Majer. A Morse complex for infinite dimensional manifolds. I. *Advances in Mathematics*, 197(2):321–410, 2005.
- [AM06] Alberto Abbondandolo and Pietro Majer. Lectures on the Morse complex for infinite-dimensional manifolds. In *Morse theoretic*

- methods in nonlinear analysis and in symplectic topology*, pages 1–74. Springer, Dordrecht, 2006.
- [AM08] Alberto Abbondandolo and Pietro Majer. A Morse complex for Lorentzian geodesics. *Asian Journal of Mathematics*, 12(3):299–319, 2008.
- [Arn78] V I Arnold. *Mathematical methods of classical mechanics*. Springer-Verlag, New York-Heidelberg, 1978.
- [AS10] Alberto Abbondandolo and Matthias Schwarz. Floer homology of cotangent bundles and the loop product. *Geometry and Topology*, 14(3):1569–1722, 2010.
- [AVDBV07] S B Angenent, J B Van Den Berg, and R C A M Vandervorst. Contact and non-contact type Hamiltonian systems generated by second-order Lagrangians. *Ergodic Theory and Dynamical Systems*, 27(1):23–49, 2007.
- [AZ08] Avraham Aizenbud and Frol Zapolsky. Functoriality in Morse theory on closed manifolds. *arXiv*, math.DG, May 2008.
- [Bar05] Piotr Bartłomiejczyk. The Conley index and spectral sequences. *Topological Methods in Nonlinear Analysis*, 25(1):195–203, 2005.
- [Bar07] Piotr Bartłomiejczyk. Simple connection matrices. *Annales Polonici Mathematici*, 90(1):77–87, 2007.
- [Bar09] Piotr Bartłomiejczyk. Spectral sequences and detailed connection matrices. *Topological Methods in Nonlinear Analysis*, 34(1):187–200, 2009.
- [BEM14] Matthew Strom Borman, Yakov Eliashberg, and Emmy Murphy. Existence and classification of overtwisted contact structures in all dimensions. *arXiv.org*, April 2014.
- [Ben84] V Benci. Closed geodesics for the Jacobi metric and periodic solutions of prescribed energy of natural Hamiltonian systems. *Annales de l’Institut Henri Poincaré. Analyse Non Linéaire*, 1(5):401–412, 1984.

-
- [BH04] Augustin Banyaga and David Hurtubise. *Lectures on Morse homology*, volume 29 of *Kluwer Texts in the Mathematical Sciences*. Kluwer Academic Publishers Group, Dordrecht, 2004.
- [BH09] Augustin Banyaga and David E Hurtubise. The Morse-Bott inequalities via a dynamical systems approach. *Ergodic Theory and Dynamical Systems*, 29(6):1693–1703, 2009.
- [Bis86] Jean-Michel Bismut. The Witten complex and the degenerate Morse inequalities. *J. Differential Geom.*, 23(3):207–240, 1986.
- [Bol78] S V Bolotin. Libration motions of natural dynamical systems. *Vestnik Moskovskogo Universiteta. Seriya I. Matematika, Mekhanika*, 1(6):72–77, 1978.
- [Bot54] Raoul Bott. Nondegenerate critical manifolds. *Annals of Mathematics. Second Series*, 60:248–261, 1954.
- [Bot56] Raoul Bott. An application of the Morse theory to the topology of Lie-groups. *Bulletin de la Société Mathématique de France*, 84:251–281, 1956.
- [Bot82] Raoul Bott. Lectures on Morse theory, old and new. *American Mathematical Society. Bulletin. New Series*, 7(2):331–358, 1982.
- [Bot88] Raoul Bott. Morse theory indomitable. *Inst. Hautes Études Sci. Publ. Math.*, (68):99–114 (1989), 1988.
- [Bre97] Glen E Bredon. *Topology and geometry*, volume 139 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.
- [BT82] Raoul Bott and Loring W Tu. *Differential forms in algebraic topology*, volume 82 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1982.
- [Cha93] Kung-ching Chang. *Infinite-dimensional Morse theory and multiple solution problems*. Progress in Nonlinear Differential Equations and their Applications, 6. Birkhäuser Boston, Inc., Boston, MA, 1993.
- [Con78] Charles Conley. *Isolated invariant sets and the Morse index*, volume 38 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, R.I., 1978.

BIBLIOGRAPHY

- [CS09] Ralph L Cohen and Matthias Schwarz. A Morse theoretic description of string topology. In *New perspectives and challenges in symplectic field theory*, pages 147–172. Amer. Math. Soc., Providence, RI, 2009.
- [CZ84] Charles Conley and Eduard Zehnder. Morse-type index theory for flows and periodic solutions for Hamiltonian equations. *Communications on Pure and Applied Mathematics*, 37(2):207–253, 1984.
- [FHW94] A Floer, H Hofer, and K Wysocki. Applications of symplectic homology. I. *Mathematische Zeitschrift*, 217(4):577–606, 1994.
- [Flo88a] Andreas Floer. A relative Morse index for the symplectic action. *Communications on Pure and Applied Mathematics*, 41(4):393–407, 1988.
- [Flo88b] Andreas Floer. Morse theory for Lagrangian intersections. *J. Differential Geom.*, 28(3):513–547, 1988.
- [Flo88c] Andreas Floer. The unregularized gradient flow of the symplectic action. *Communications on Pure and Applied Mathematics*, 41(6):775–813, 1988.
- [Flo89a] Andreas Floer. Symplectic fixed points and holomorphic spheres. *Communications in Mathematical Physics*, 120(4):575–611, 1989.
- [Flo89b] Andreas Floer. Witten’s complex and infinite-dimensional Morse theory. *J. Differential Geom.*, 30(1):207–221, 1989.
- [Fra89] Robert D Franzosa. The connection matrix theory for Morse decompositions. *Transactions of the American Mathematical Society*, 311(2):561–592, 1989.
- [Fuk96] Kenji Fukaya. Floer homology of connected sum of homology 3-spheres. *Topology. An International Journal of Mathematics*, 35(1):89–136, 1996.
- [GG03] Viktor L Ginzburg and Basak Z Gürel. A C^2 -smooth counterexample to the Hamiltonian Seifert conjecture in \mathbb{R}^4 . *Annals of Mathematics. Second Series*, 158(3):953–976, 2003.

-
- [Gin95] Viktor L Ginzburg. An embedding $S^{2n-1} \rightarrow \mathbb{R}^{2n}$, $2n - 1 \geq 7$, whose Hamiltonian flow has no periodic trajectories. *International Mathematics Research Notices*, (2):83–97 (electronic), 1995.
- [GZ83] Herman Gluck and Wolfgang Ziller. Existence of periodic motions of conservative systems. In *Seminar on minimal submanifolds*, pages 65–98. Princeton Univ. Press, Princeton, NJ, 1983.
- [Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [Hil88] Peter Hilton. A brief, subjective history of homology and homotopy theory in this century. *Mathematics Magazine*, 61(5):282–291, 1988.
- [Hir94] Morris W Hirsch. *Differential topology*, volume 33 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994.
- [Hur13] David E Hurtubise. Three approaches to Morse-Bott homology. *African Diaspora Journal of Mathematics*, 14(2):145–177, 2013.
- [HV88] H Hofer and C Viterbo. The Weinstein conjecture in cotangent bundles and related results. *Annali della Scuola Normale Superiore di Pisa. Classe di Scienze. Serie IV*, 15(3):411–445 (1989), 1988.
- [HZ11] Helmut Hofer and Eduard Zehnder. *Symplectic invariants and Hamiltonian dynamics*. Modern Birkhäuser Classics. Birkhäuser Verlag, Basel, 2011.
- [Izy01] Marek Izydorek. A cohomological Conley index in Hilbert spaces and applications to strongly indefinite problems. *Journal of Differential Equations*, 170(1):22–50, 2001.
- [Izy06] Marek Izydorek. The $\mathcal{L}\mathcal{S}$ -index: a survey. In *Morse theoretic methods in nonlinear analysis and in symplectic topology*, pages 277–320. Springer, Dordrecht, 2006.
- [Jia99] Mei-Yue Jiang. Morse homology and degenerate Morse inequalities. *Topological Methods in Nonlinear Analysis*, 13(1):147–161, 1999.
- [Kli78] Wilhelm Klingenberg. *Lectures on closed geodesics*. Springer-Verlag, Berlin-New York, 1978.

BIBLIOGRAPHY

- [Kli95] Wilhelm P A Klingenberg. *Riemannian geometry*, volume 1 of *de Gruyter Studies in Mathematics*. Walter de Gruyter and Co., Berlin, second edition, 1995.
- [KM07] Peter Kronheimer and Tomasz Mrowka. *Monopoles and three-manifolds*, volume 10 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2007.
- [KMV13] William D Kalies, Konstantin Mischaikow, and Robert C A M Vandervorst. Lattice Structures for Attractors I. *arXiv.org*, July 2013.
- [Kup63] Ivan Kupka. Contribution à la théorie des champs génériques. *Contributions to Differential Equations*, 2:457–484, 1963.
- [Kup94] Krystyna Kuperberg. A smooth counterexample to the Seifert conjecture. *Annals of Mathematics. Second Series*, 140(3):723–732, 1994.
- [LF51] L A Lyusternik and A I Fet. Variational problems on closed manifolds. *Doklady Akad. Nauk SSSR (N.S.)*, 81:17–18, 1951.
- [Man03] Ciprian Manolescu. Seiberg-Witten-Floer stable homotopy type of three-manifolds with $b_1=0$. *Geometry and Topology*, 7:889–932 (electronic), 2003.
- [Mat02] Yukio Matsumoto. *An introduction to Morse theory*, volume 208 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 2002.
- [Mat13] Benjamin Matschke. Successive Spectral Sequences. *arXiv.org*, August 2013.
- [McC88] Christopher McCord. Mappings and homological properties in the Conley index theory. *Ergodic Theory and Dynamical Systems*, 8*(Charles Conley Memorial Issue):175–198, 1988.
- [McC92] Christopher McCord. Poincaré-Lefschetz duality for the homology Conley index. *Transactions of the American Mathematical Society*, 329(1):233–252, 1992.

-
- [McC01] John McCleary. *A user's guide to spectral sequences*, volume 58 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2001.
- [Mil63] John Milnor. *Morse theory*. Based on lecture notes by M. Spivak and R. Wells. *Annals of Mathematics Studies*, No. 51. Princeton University Press, Princeton, N.J., 1963.
- [Mil65] John Milnor. *Lectures on the h-cobordism theorem*. Notes by L. Siebenmann and J. Sondow. Princeton University Press, Princeton, N.J., 1965.
- [Mis95] Konstantin Mischaikow. Conley index theory. In *Dynamical systems (Montecatini Terme, 1994)*, pages 119–207. Springer, Berlin, 1995.
- [Mis99] Konstantin Mischaikow. The Conley index theory: a brief introduction. In *Conley index theory (Warsaw, 1997)*, pages 9–19. Polish Acad. Sci., Warsaw, 1999.
- [MM02] Konstantin Mischaikow and Marian Mrozek. Conley index. In *Handbook of dynamical systems, Vol. 2*, pages 393–460. North-Holland, Amsterdam, 2002.
- [Mor25] Marston Morse. Relations between the critical points of a real function of n independent variables. *Transactions of the American Mathematical Society*, 27(3):345–396, 1925.
- [Mun66] James R Munkres. *Elementary differential topology*, volume 1961 of *Lectures given at Massachusetts Institute of Technology, Fall*. Princeton University Press, Princeton, N.J., 1966.
- [Mun75] James R Munkres. *Topology: a first course*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1975.
- [Nic11] Liviu Nicolaescu. *An invitation to Morse theory*. Universitext. Springer, New York, second edition, 2011.
- [Nor95] Douglas E Norton. The fundamental theorem of dynamical systems. *Commentationes Mathematicae Universitatis Carolinae*, 36(3):585–597, 1995.

BIBLIOGRAPHY

- [Pal63] Richard S Palais. Morse theory on Hilbert manifolds. *Topology. An International Journal of Mathematics*, 2:299–340, 1963.
- [Pas12] Federica Pasquotto. A short history of the Weinstein conjecture. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 114(3):119–130, 2012.
- [PdM82] Jacob Palis, Jr and Welington de Melo. *Geometric theory of dynamical systems*. Springer-Verlag, New York-Berlin, 1982.
- [Poi10] Henri Poincaré. *Papers on topology*, volume 37 of *History of Mathematics*. American Mathematical Society, Providence, RI; London Mathematical Society, London, 2010.
- [PS64] R S Palais and S Smale. A generalized Morse theory. *Bulletin of the American Mathematical Society*, 70:165–172, 1964.
- [Rab78] Paul H Rabinowitz. Periodic solutions of Hamiltonian systems. *Communications on Pure and Applied Mathematics*, 31(2):157–184, 1978.
- [Rab79] Paul H Rabinowitz. Periodic solutions of a Hamiltonian system on a prescribed energy surface. *Journal of Differential Equations*, 33(3):336–352, 1979.
- [RS92] Joel W Robbin and Dietmar A Salamon. Lyapunov maps, simplicial complexes and the Stone functor. *Ergodic Theory and Dynamical Systems*, 12(1):153–183, 1992.
- [RV14] T O Rot and R C A M Vandervorst. Morse–Conley–Floer homology. *Journal of Topology and Analysis*, 6(3):305–338, 2014.
- [Sal90] Dietmar Salamon. Morse theory, the Conley index and Floer homology. *The Bulletin of the London Mathematical Society*, 22(2):113–140, 1990.
- [Sch93] Matthias Schwarz. *Morse homology*, volume 111 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1993.
- [Sma61] Stephen Smale. On gradient dynamical systems. *Annals of Mathematics. Second Series*, 74:199–206, 1961.

-
- [Sma63] S Smale. Stable manifolds for differential equations and diffeomorphisms. *Ann. Scuola Norm. Sup. Pisa (3)*, 17:97–116, 1963.
- [Sma67] S Smale. Differentiable dynamical systems. *Bulletin of the American Mathematical Society*, 73:747–817, 1967.
- [Sta14] Maciej Starostka. Morse cohomology in a Hilbert space via the Conley Index. *arXiv.org*, January 2014.
- [SZ13] Stefan Suhr and Kai Zehmisch. Linking and closed orbits. *arXiv.org*, May 2013.
- [Tau09a] Clifford Henry Taubes. Notes on the Seiberg-Witten equations, the Weinstein conjecture and embedded contact homology. In *Current developments in Mathematics*, pages 221–245. Int. Press, Somerville, MA, 2009.
- [Tau09b] Clifford Henry Taubes. The Seiberg-Witten equations and the Weinstein conjecture. II. More closed integral curves of the Reeb vector field. *Geometry and Topology*, 13(3):1337–1417, 2009.
- [VDBPRV13] J B Van Den Berg, F Pasquotto, T O Rot, and R C A M Vandervorst. Closed characteristics on non-compact mechanical contact manifolds. *arXiv*, math.SG, March 2013.
- [vdBPV09] Jan Bouwe van den Berg, Federica Pasquotto, and Robert C Vandervorst. Closed characteristics on non-compact hypersurfaces in \mathbb{R}^{2n} . *Mathematische Annalen*, 343(2):247–284, 2009.
- [Vit87] Claude Viterbo. A proof of Weinstein’s conjecture in \mathbb{R}^{2n} . *Annales de l’Institut Henri Poincaré. Analyse Non Linéaire*, 4(4):337–356, 1987.
- [Vit97] Claude Viterbo. Exact Lagrange submanifolds, periodic orbits and the cohomology of free loop spaces. *J. Differential Geom.*, 47(3):420–468, 1997.
- [Vit99] C Viterbo. Functors and computations in Floer homology with applications. I. *Geometric and Functional Analysis*, 9(5):985–1033, 1999.
- [VKM13] Vandervorst, Kalies, and Mischaikow. *Conley Index Theory*. Preprint, August 2013.

BIBLIOGRAPHY

- [Web06] Joa Weber. The Morse-Witten complex via dynamical systems. *Expositiones Mathematicae*, 24(2):127–159, 2006.
- [Wei78] Alan Weinstein. Periodic orbits for convex Hamiltonian systems. *Annals of Mathematics. Second Series*, 108(3):507–518, 1978.
- [Wei79] Alan Weinstein. On the hypotheses of Rabinowitz' periodic orbit theorems. *Journal of Differential Equations*, 33(3):353–358, 1979.
- [Wit82] Edward Witten. Supersymmetry and Morse theory. *J. Differential Geom.*, 17(4):661–692 (1983), 1982.
- [WY73] Jr Wilson and James A Yorke. Lyapunov functions and isolating blocks. *Journal of Differential Equations*, 13:106–123, 1973.