A Survey on Morse Theory and its Extensions

Carfora Giulia, Muñiz-Brea Álvaro and Giannini Riccardo

December 3, 2021

Contents

Introduction

Every mathematician has a secret weapon.

Mine is Morse theory.

-Raoul Bott

In 1924 Marston Morse published the paper "A fundamental class of geodesics on any closed surface of genus greater than one" in the Transactions of the American Mathematical Society, laying the foundations for the theory named after him. The key insight from Morse was that a typical smooth function on a manifold can reflect its topology quite directly. More specifically, one can study the topology of a smooth manifold through the critical points of a suitable smooth function. These critical points turn out to encode information about the topology of the manifold. In fact, when Morse first started developing the theory the field was known as critical point theory.

For example, it is a standard result in topology that any continuous function on a compact manifold must achieve a maximum and a minimum. Here, the topology of the manifold is giving restrictions on the functions it admits. Morse theory allows us to consider an "inverse" question: if a manifold admits a function with a unique maximum and minimum, can we say anything about its topology? It turns out we can: in this case, the manifold is homeomorphic to a sphere (Theorem 1.6). Morse theory was later developed into a homology theory, allowing us to obtain a topological invariant. Following these ideas, Morse homology was extended to many different settings and its applications nowadays are diverse and striking.

Over time, Morse theory evolved in many ways and its ideas were carried to other settings, spanning over apparently unrelated areas of mathematics. The examples we are going to delve into in this project are Lagrangian Floer homology and discrete Morse theory, but further topics can be found in literature.

For instance, in the 1950's Stephen Smale enabled us to understand finite dimensional manifolds in terms of handle decomposition. This is at the core of the h-cobordism theorem and the topological Poincaré conjecture. Furthermore, one can extend this formalism to categories other than just smooth finite dimensional manifolds, such as CW-complexes and infinite-dimensional spaces. One of the most famous examples of the latter comes from considering the energy functional on the space $\mathcal{L}M$ of loops in M, defined by

$$E(\gamma) = \int_{S^1} |\dot{\gamma}|^2, \quad \gamma \in \mathcal{L}M.$$

Using an infinite-dimensional analogue of Morse theory, the functional E can be used to prove existence theorems for closed geodesics, such as the celebrated result that any metric on S^2 has at least three geodesics [20, 4] or, more generally, that any two points on S^n are joined by an infinite number of geodesics, regardless of the chosen metric [23]. In a similar vein, these ideas can be applied to CW-complexes. A discrete adaptation of the theory was firstly introduced by Robin Forman and developed with a geometric group theory perspective by Mladen Bestvina and Noel Brady. Inspired by the classical approach of John Milnor, they built a well-behaved machinery as useful as its smooth counterpart.

This survey will focus on the original smooth finite-dimensional Morse theory, Lagrangian Floer homology and discrete Morse theory. The structure is as follows. Section 1 introduces smooth finite-dimensional Morse theory. We will begin with the definition of a Morse function, and then present some of the most important results in the field, justifying why it is sensible to look at critical points. The rest of the section will then be devoted to introducing the most recent approach to the study of Morse theory. We firstly introduce pseudo-gradient fields adapted to a Morse function, which allow us to define the stable and unstable manifolds of a critical point. We then define Morse homology and show that it is independent of the choice of Morse function and pseudo-gradient field. Finally, we conclude the section by stating that Morse homology agrees with singular homology. In Section 2 we will introduce the theory of Lagrangian Floer (co)homology, used to study intersection properties of Lagrangian submanifolds. The section will begin with Arnold's conjecture - which can be thought of as the motivation for the subject - and some introduction to complex geometry. We will then give an idea of the construction: although the idea is not too complicated, the proper definition will turn out to be very technical, with a lot of things to take care of. The main part of this section is the rest of subsection 2.2, where we analyze all the technical difficulties to finally arrive to a correct definition. The section concludes by showing how to phrase the theory in Morse terms and an application of the theory in the context of displaceability. Lastly, Section 3 is entirely devoted to give the basics for the above-mentioned Bestvina-Brady construction, the analog formulation of Morse Theory established in the contest on CW-complexes. Their main result shows a way to build a group with prescribed finiteness properties out of a simplicial complex. More specifically, they define a functor from the family of finite flags[1] to the category of groups. It turns out that the homotopy type of the simplicial complex gives finiteness conditions to the group associated. We conclude this section with one of the many application their main theorem provides.

1 Smooth finite-dimensional Morse theory

This section will provide a brief introduction to the fundamental ideas of smooth finite-dimensional Morse theory discussed in the introduction. We will start our discussion by introducing Morse functions and related key results, following the approach in [18]. The remaining part of the section will then be devoted to setting up Morse homology, as in [2, 8, 17, 12].

1.1 Morse functions

Throughout this section, we will assume M is a closed compact smooth manifold of finite dimension n. Consider a smooth function $f : M \to \mathbb{R}$.

Definition 1.1. Suppose p is a point on M and consider the local coordinates $(x_1, \dots, x_n)$ of a neighbourhood of p, such that p can be regarded as the origin. We say that p is a *critical point* of f if

$$\frac{\partial f}{\partial x_i}(p) = 0, \quad i = 1, \dots, n.$$

Definition 1.2. Suppose $f : M \to \mathbb{R}$ is a smooth function on the finite n-manifold M and let p be a point on M. Let $(x_1, \dots, x_n)$ be local coordinates near p. We define the *Hessian* of f at p to be the matrix of second partial derivatives of f. The determinant of such matrix is denoted by $H_f(p)$, namely

$$H_f(p) = \det\left(\frac{\partial^2 f}{\partial x_i \partial x_j}(p)\right).$$

[1] We will see that a finite flag is a finite simplcial complex uniquely determined by its 1-skeleton.

Note that in this definition we chose local coordinates $(x_1, \ldots, x_n)$. We could easily choose a different set of coordinates, say $(y_1, \ldots, y_n)$. Then in this case the Hessian $H'_f(p)$ is related to $H_f(p)$ from our definition, via a change of basis matrix P, which expresses the new set of coordinates in terms of the original ones, as follows [18]:

$$\left(\frac{\partial^2 f}{\partial x_i \partial x_j}(p) \right) = P^\intercal \left(\frac{\partial^2 f}{\partial y_i \partial y_j}(p) \right) P.$$

Definition 1.3. Suppose p is a critical point of $f : M \to \mathbb{R}$. We say that p is a *nondegenerate critical point* if $H_f(p) \neq 0$.

This condition turns out to be well-defined, since it does not depend on the coordinate system chosen near p. Following our previous observation, in particular we note that the sign $H_f(p)$ is independent of the coordinate system [18]. The notion of a nondegenerate critical point is fundamental to introduce the concept of a Morse function.

Definition 1.4. A smooth function $f : M \to \mathbb{R}$ is a *Morse function* if each of its critical points is nondegenerate.

Example 1.1. We may now introduce one of the most simple and well-known examples of a Morse function. Consider the 2-sphere S^2 in $\mathbb{R}^3$. We define the *height function* $h : S^2 \to \mathbb{R}$ as $h(x, y, z) = z$.

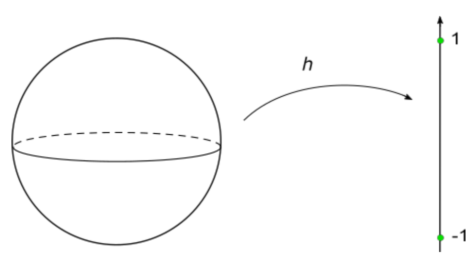

Figure 1: The height function on the sphere S^2 in $\mathbb{R}^3$.

Intuitively, it is clear from the picture that the height function on S^2 only has two critical points. This can be made clearer by looking at the level sets of h. Indeed,

$$h^{-1}(c) = \begin{cases} \varnothing & \text{if } c < -1, \\ \text{a point} & \text{if } c = -1, \\ \text{a circle} & \text{if } -1 < c < 1, \\ \text{a point} & \text{if } c = 1, \\ \varnothing & \text{if } c > 1. \end{cases}$$

Indeed, note that there is a well-defined topology on each level set and it changes each time we pass through a critical point [15]. Hence, it is clear that S^2 has two critical points, namely the south pole and the north pole. We may now concern ourselves with verifying whether h is a Morse function by going through the necessary calculations, following the example from [18]. The upper hemisphere of S^2 can be parametrised via $(x, y, \sqrt{1 - x^2 - y^2})$ for (x, y) in the unit disc in $\mathbb{R}^2$, hence $h(x, y) = \sqrt{1 - x^2 - y^2}$. It can be easily computed that

$$\frac{\partial h}{\partial x} = \frac{-x}{\sqrt{1 - x^2 - y^2}} \quad \text{and} \quad \frac{\partial h}{\partial y} = \frac{-y}{\sqrt{1 - x^2 - y^2}},$$

and hence the only critical point in the upper hemisphere of S^2 is the north pole. The determinant of the Hessian of h can thus be computed to be

$$H_h(x, y) = \det \begin{pmatrix} \frac{y^2 - 1}{(1 - x^2 - y^2)^{\frac{3}{2}}} & \frac{-xy}{(1 - x^2 - y^2)^{\frac{3}{2}}} \\ \frac{-xy}{(1 - x^2 - y^2)^{\frac{3}{2}}} & \frac{x^2 - 1}{(1 - x^2 - y^2)^{\frac{3}{2}}} \end{pmatrix} = \frac{1}{(1 - x^2 - y^2)^2}.$$

It follows that the north pole is a nondegenerate critical point of h, since clearly $H_h(0, 0) \neq 0$. In a completely analogous fashion it is possible to parametrise the southern hemisphere and find out that the south pole is another nondegenerate critical point. By parametrising the eastern and western hemispheres as well we can then verify that there are no more critical points. Hence we have verified that h is a Morse function.

Before proceeding with our discussion, we introduce the concept of a sublevel set. This idea is part of the classical approach to Morse theory, and is used to study how the topology of a manifold changes as we move along it using a Morse function. It is very useful to prove that manifolds are diffeomorphic by understanding them in terms of attaching handles [17]. For the remainder of the section, however, we will focus on a newer approach to Morse theory using pseudo-gradient fields, as it can be extended to the case of infinite-dimensional manifolds, where the classical approach turns out to be useless [17]. The idea of a sublevel set, however, will appear again in Sec. 3, where we introduce discrete Morse theory.

Definition 1.5. Let $f : M \to \mathbb{R}$ be a smooth function on the n-manifold M. For $a \in \mathbb{R}$ the *sublevel set* $M_{(-\infty,a]}$ is defined as

$$M_{(-\infty,a]} = f^{-1}(-\infty, a] = \{x \in M : f(x) \leq a\}.$$

So we can now look back at our example. For $a < -1$, $S^2_{(-\infty,a]} = \varnothing$. $S^2_{(-\infty,-1]} = (0,0,-1)$, the south pole, while for $-1 < a < 1$, $S^2_{(-\infty,a]}$ is a disc centred at $(0,0,-1)$, which we may note is topologically contractible to a point. Finally, $S^2_{(-\infty,1]}$ is the whole sphere S^2. The important thing to observe by looking at sublevel sets is the shift in topology every time we pass a critical point of the Morse function on our manifold. We may now look at sublevel sets in another important example concerning the torus.

Example 1.2. Consider the torus T^2 in $\mathbb{R}^3$ and the height function $h : T^2 \to \mathbb{R}$. We can prove h is Morse doing calculations similar to the ones in the previous example. It is important to note that we are looking at a torus standing vertically: if we were to consider the torus lying flat together with the height function h we would have a circle of maxima, and hence h would not be Morse. Looking at Fig. 2 we notice there are four critical points of h,

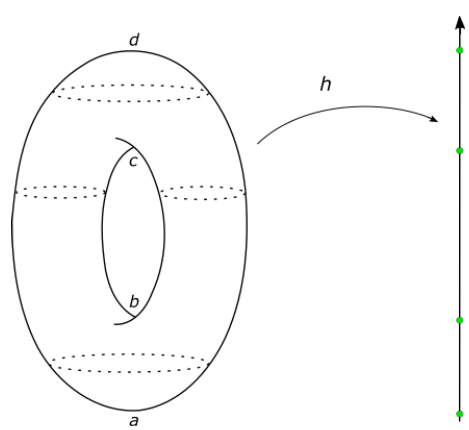

Figure 2: The height function on the torus T^2 in $\mathbb{R}^3$.

denoted by a, b, c and d. We now want to examine the sublevel sets, which are shown in Fig 3. Assuming the torus to have height 0 at the point a, we can see that $T^2_{(-\infty,0]}$ only contains the point a. For a point p on the real line such that $p < h(b)$, $T^2_{(-\infty,p]}$ is a disc centred at a. Now, for $h(b) < p < h(c)$ we have that $T^2_{(-\infty,p]}$ is a cylinder. Similarly, for $h(c) < p < h(d)$, $T^2_{(-\infty,p]}$ is a torus with a disc removed, therefore we have a significant change in topology. Finally, for $p = h(d)$, we obtain that $T^2_{(-\infty,p]}$ is the whole torus. Thus, we see

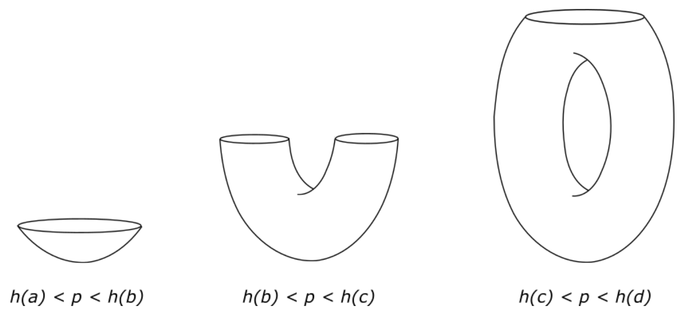

Figure 3: Sublevel sets for the height function on the vertical torus.

that near a the surface looks like a paraboloid opening upwards. Near b and c it looks like a saddle with two orthogonal directions in which h decreases or increases. Finally, near d the surface once again looks like a paraboloid, but opening downwards [18].

As previously mentioned, the torus is a key example in our understanding of Morse theory and we will return to it later on in the present section. An important observation that can be deduced from this example is that, given a manifold M and a Morse function f on it, if p and q are two points on the real line with no critical values of f between them, then M_p and M_q are homeomorphic. In particular, M_a is a deformation retract of M_b. More on this result can be found in [22]. This means there is an impressive amount of information about the topology of a manifold M that is stored by the critical point of a Morse function associated to it. The following theorem is a famous consequence of this fact.

Theorem 1.6 (Reeb's Theorem). *Let M be a compact n-manifold and that $f : M \to \mathbb{R}$ is a Morse function with exactly two critical points. Then M is homeomorphic to the sphere S^m.*

1.2 Key results

Example 1.1 illustrates the fact that near the north pole, the height function can be described by $h(x,y) = 1 - x^2 - y^2$. This representation relies on the fact that by Taylor's theorem smooth functions can be approximated by a Taylor expansion. The following theorem, known as the Morse lemma, states that locally in an appropriate chart, often referred to as a *Morse chart*, a Morse function can always be represented as above.

Theorem 1.7 (Morse Lemma). *Suppose p is a nondegenerate critical point of the Morse function $f : M \to \mathbb{R}$. Then there is a local coordinate system $(x_1, \dots, x_n)$ in a neighbourhood U of p, where p corresponds to the origin, such that on U f can be represented as*

$$f = f(p) - x_1^2 - x_2^2 - \dots - x_i^2 + x_{i+1}^2 + \dots + x_n^2.$$

A proof of this theorem will not be provided here, as it is outside the scope of the section, but it can be found in [18]. The Morse lemma also prompts us to define the index of a critical point.

Definition 1.8. The number i of minus signs appearing in the representation of f in Theorem (1.7) is called the *index* of the critical point p, and it is such that $0 \leq i \leq n$.

It follows from the definition that the index i is actually the number of negative eigenvalues of the Hessian of f. Looking back at Example 1.2, we see that a has index 0, b and c have index 1, and d has index 2. In general, it can be observed that a local maximum has full index n and a local minumum has index 0, while saddle points have index $0 < i < n$ [12]. We now present two important consequences of the Morse lemma.

Corollary 1.9. *The critical points of a Morse function are isolated.*

Proof. Consider a Morse function $f : M \to \mathbb{R}$ and one of its critical points p. By the Morse lemma then f can be described as $f = f(p) - x_1^2 - \dots - x_i^2 + x_{i+1}^2 + \dots + x_n^2$ in a neighbourhood of p. By carrying out a simple calculation we then find out that there are no other critical points in this neighbourhood [18]. □

Corollary 1.10. *A Morse function f on a compact manifold M has only finitely many critical points.*

Proof. This statement is proved by contradiction. We suppose there are infinitely many critical points of f:

$$p_1, p_2, p_3, p_4, \dots$$

Now, as mentioned at the beginning of this section, we assume M to be a compact manifold. Then by compactness of M there must exist a convergent subsequence of critical points, say

$$p_{i_1}, p_{i_2}, p_{i_3}, p_{i_4}, \dots$$

lying in a neighbourhood of the limit point p. The partial derivatives $\frac{\partial f}{\partial x_i}$ vanish at the p_{i_k}, and by continuity they vanish at p. This implies that p is a critical point of f with a neighbourhood in which there are infinitely many other critical points. This contradicts Corollary 1.9, hence there are only finitely many critical points [18]. □

Up to this point, we have discussed Morse functions on a manifold and started getting a glimpse into why they are so powerful. However, this point of view may seem restrictive, as we do not know that we can always find a Morse function on a manifold in hope of better understanding its topology. The following result tells us that Morse functions on a manifold exist and there are actually many of them. A proof and more complete discussion on the topic can be found in [22, 12].

Lemma 1.11. *Suppose M is an n-dimensional manifold. Then for almost any $p \in \mathbb{R}$ the function $f : M \to \mathbb{R}$ defined by*

$$x \mapsto \|x - p\|^2$$

is a Morse function.

Together with Whitney's embedding theorem this implies that on smooth manifolds Morse functions always exist. Furthermore, the next result tells us that Morse functions are generic, that is that any smooth function on a manifold can be approximated by a Morse function [12].

Theorem 1.12. *Suppose M is a manifold and $f : M \to \mathbb{R}$ is a smooth function defined on it. Let $k \in \mathbb{N}$. Then on any compact subset of M, f can be approximated by a Morse function in C^k-norm.*

1.3 Pseudo-gradients

In this subsection we continue our discussion of Morse theory using a newer approach, that is by considering pseudo-gradients on a manifold that will then help us connect the critical points of a Morse function. Recall that the gradient of a smooth function $f : M \to \mathbb{R}$ is a vector field defined as

$$\text{grad } f = \left(\frac{\partial f}{\partial x_1}, \ldots, \frac{\partial f}{\partial x_n} \right).$$

For vector fields Y on $\mathbb{R}^n$ the gradient vector field of f can equivalently be defined by

$$g(\text{grad } f, Y) = df(Y),$$

where g is a Euclidean metric [12]. This idea can be generalized to a Riemannian manifold, prompting the next definition from [12]. Recall that a Riemannian manifold is a pair made of a smooth manifold M and a Riemannian metric on it g.

Definition 1.13. Let (M, g) be a Riemannian manifold and suppose $f : M \to \mathbb{R}$ is smooth. The *gradient* of f is the vector field defined by

$$g(\text{grad } f, Y) = df(Y)$$

for all vector fields Y.

We can observe from this definition that the gradient vanishes if and only if $df = 0$, that is on critical points. Moreover, f decreases along integral curves of f [12]. However, for our purposes we might not always precisely require the gradient of a function [8, 2], so using the aforementioned properties we construct pseudo-gradient fields, whose integral curves connect critical points of Morse functions [12].

Definition 1.14. Let $f : M \to \mathbb{R}$ be a Morse function on M and let X be a vector field on M. Then X is a *pseudo-gradient adapted to* f if

1. For every $p \in M$ we have that $(df)_p \cdot X_p \leq 0$, and equality holds if and only if p is a critical point of f;

2. If p is a critical point of f of index i, then p has a sufficiently small coordinate neighbourhood U such that f has a standard form

$$f = f(p) - x_1^2 - x_2^2 - \cdots - x_i^2 + x_{i+1}^2 + \cdots + x_n^2$$

   and X can be written as $-grad\ f$.

Thus we observe that if (M, g) is a Riemannian manifold, then the vector field $-\text{grad } f$ defined using g is a pseudo-gradient [8]. Once again, we might wonder if pseudo-gradients exist for a given manifold, and that turns out to be the case.

Theorem 1.15. *If $f : M \to \mathbb{R}$ is a Morse function on a compact smooth manifold M, then there exists a pseudo-gradient adapted to f.*

Proof. The following proof is adapted from [8, 12]. An alternative approach would be to make use of the existence of Riemannian metrics on manifolds. Consider the critical points of f, $p_1, p_2, \dots, p_r$, which are finitely many, as previously proved. Let $(U_1, h_1), (U_2, h_2), \dots, (U_r, h_r)$ be Morse charts in the neighbourhoods of each of these critical points whose images $\{\Omega_i\}_{1 \le i \le r}$ are disjoint. We can add more charts to extend them to a finite atlas which can be refined in a way such that each p_i is contained in exactly one Ω_i.
Now consider X_i, which is the pushforward of the negative gradient of f in U_i by h_i. Moreover, let $\{\varphi_i\}_{1 \le i \le m}$ be a partition of unity subordinate to $\{\Omega_i\}_{1 \le i \le m}$ such that $\varphi_i(p_i) = 1$ for all i. Now define the vector fields

$$\tilde{X}_i = \begin{cases} \varphi_i(x) X_i & \text{if } x \in \Omega_i, \\ 0 & \text{otherwise.} \end{cases}$$

and

$$X = \sum_{i=1}^{m} \tilde{X}_i.$$

We want to verify that X is a pseudo-gradient adapted to f.
Indeed, note that for $x \in M$

$$df_x \cdot X_x = \sum_{i=1}^{m} df_x \cdot \tilde{X}_{i,x} \le 0,$$

follows from the definition of $\tilde{X}_i$. Furthermore, we see that this quantity is zero if and only if $\varphi_i(x) X_i = 0$ for all i. This means that either x is a critical point of f, or $\varphi_i(x) = 0$ for all i. However, the latter is impossible, since we previously defined $\{\varphi_i\}_{1 \le i \le m}$ to be a partition of unity, so we must have that x is a critical point. Finally, suppose p_i is a critical point and recall that X was constructed to be the negative gradient of f with the canonical metric over $U_i \cap (\bigcup_{i \ne j} U_j)$, which also contains a neighbourhood of p_i. □

We now continue our investigation on how to connect critical points. Suppose we have a Morse function $f : M \to \mathbb{R}$ and a pseudo-gradient field X. We want to look at vector flows of X, often called *trajectories* of X, and denoted φ^t [12]. The flow actually defines a one-parameter family of diffeomorphisms $\varphi^t : \mathbb{R} \times M \to M$ for $t \in \mathbb{R}$ such that $\varphi^0 = \text{id}$ and $\frac{\partial \varphi^t}{\partial t} = X$ [17]. This enables us to define stable and unstable manifolds, which are collections of trajectories tending to or moving away from a critical point.

Definition 1.16. Suppose p is a critical point of a Morse function $f : M \to \mathbb{R}$. The *stable manifold* of p is defined as

$$W^s(p) = \{x \in M : \lim_{t\to\infty} \varphi^t(x) = p\}.$$

Similarly, the *unstable manifold* of p is

$$W^u(p) = \{x \in M : \lim_{t\to-\infty} \varphi^t(x) = p\}.$$

Often in the literature they are also referred to as the *ascending* and *descending manifolds* of p, respectively.

The stable and unstable manifolds of a critical point can actually be shown to be submanifolds of M. Moreover, given a critical point p, $W^u(p)$ is diffeomorphic to an open disc and

$$\dim W^u(p) = \text{index}(p),$$

thus

$$\dim W^s(p) = \text{codim}\, W^u(p) = \text{index}(p).$$

Loosely speaking, this implies that the trajectories belonging to stable and unstable manifolds describe critical points. Furthermore, it can be proved that all trajectories belong to such manifolds [12].

Proposition 1.17. *Suppose M is a compact manifold and let $\varphi^t(x)$ be a trajectory of a pseudo-gradient field X of f. Then there are critical points c and d of f, such that*

$$\lim_{t\to\infty} \varphi^t(x) = c \text{ and } \lim_{t\to-\infty} \varphi^t(x) = d.$$

An outline of a proof for this statement can be found in [12].
To make good use of the idea of connecting critical points we need to introduce a condition on the stable and unstable manifolds. This condition in turn relies on the idea of transversality.

Definition 1.18. Let M be a manifold and suppose A and B are smooth submanifolds of M. A and B are said to *intersect transversally* if for each point $x \in A \cap B$ we have that

$$T_x M = T_x A + T_x B.$$

If A and B intersect transversally, we write $A \pitchfork B$.

If two submanifolds A and B intersect transversally it can be shown that $A \cap B$ is a submanifold of M and $\dim(M) = \dim(A) + \dim(B) - \dim(A \cap B)$ [8].

Definition 1.19. Let $f : M \to \mathbb{R}$ be a Morse function. A pseudo-gradient X adapted to f is said to satisfy the *Smale condition* if all stable and unstable manifolds intersect transversally. That is, for any two critical points c and d, $W^u(c) \pitchfork W^s(d)$. When this is the case, often in the literature (f, X) are said to be a *Morse-Smale pair*.

It follows from this definition that if we have a pseudo-gradient field on an n-manifold M satisfying the Smale condition, then for any two critical points c, d we have

$$\begin{aligned}\dim(W^u(c) \cap W^s(d)) &= n - \operatorname{codim}(W^u(c) \cap W^s(d)) \\ &= n - (\operatorname{codim}(W^u(c)) + \operatorname{codim}(W^s(d))) \\ &= n - (n - \operatorname{index}(c) + \operatorname{index}(d)) \\ &= \operatorname{index}(c) - \operatorname{index}(d).\end{aligned}$$

Thus we can retrieve the dimension of a submanifold via the difference of indexes of two critical points. Said submanifold contains the trajectories of the pseudo-gradient that connects c to d, and is denoted by

$$\mathcal{M}(c, d) \coloneqq W^u(c) \cap W^s(d) = \{x \in M : \lim_{t\to\infty} \varphi^t(x) = d \text{ and } \lim_{t\to-\infty} \varphi^t(x) = c\}.$$

If this submanifold is non-empty then it contains at least one trajectory, hence it has dimension at least one. Note as well that indices of critical points always decrease as we go along a trajectory [12]. We may note that $\mathbb{R}$ as a Lie group acts on $\mathcal{M}(c, d)$ by translations in time [12], that is we have:

$$t \cdot x = \varphi^t(x).$$

It turns out that this action of $\mathbb{R}$ is smooth, proper, and free, more on which can be found in [12], prompting our next definition.

Definition 1.20. We define the *moduli space of trajectories* from c to d as

$$\mathcal{L}(c,d) = \mathcal{M}(c,d)/\mathbb{R},$$

where $\mathbb{R}$ acts on $\mathcal{M}(c,d)$ by the flow $\{\varphi^t\}$.

Moreover, $\mathcal{L}(c,d)$ is actually a manifold, and as a consequence of the Morse-Smale condition we have that

$$\dim \mathcal{L}(c,d) = \text{index}(c) - \text{index}(d) - 1.$$

In particular, when $\text{index}(c) - \text{index}(d) = 1$, $\dim \mathcal{L}(c,d) = 0$, and so $\mathcal{L}(c,d)$ is a discrete set. However, $\mathcal{L}(c,d)$ is also compact, as we will discuss, so it must be finite [8].

1.4 Morse Homology

We now aim to continue the work started in the previous subsection by using the pseudo-gradient associated to a Morse function satisfying the Smale condition to define the Morse complex and a differential on it. A simpler initial approach to Morse homology can be seen by studying Morse homology modulo 2, which allows us to ignore issues related to orientation. A discussion about this approach can be found in [8, 2, 12]. However, to keep this brief introduction to Morse homology concise, we will be discussing integral Morse homology directly.

Hence we must start by introducing an orientation on $\mathcal{L}(c,d)$, for critical points $c \neq d$. We do so by first choosing an orientation for the unstable manifold $W^u(c)$. At any point of of a flow line γ from c to d, there is a canonical isomorphism at the level of orientations [17]:

$$\begin{aligned} TW^u(c) &\cong T(W^u(c) \cap W^s(d)) \oplus TM/TW^s(d) \\ &\cong T_\gamma \mathcal{L}(c,d) \oplus T_\gamma \oplus T_d(W^u(d)). \end{aligned}$$

In the above expression, the first isomorphism comes from the Smale condition. For the second one, note that $T(W^u(c) \cap W^s(d)) \cong T_\gamma \mathcal{L}(c,d) \oplus T_\gamma$, since $\dim \mathcal{L}(c,d) = \text{index}(c) - \text{index}(d) - 1$. Finally, if we translate the subspace $T_d(W^u(d)) \subset T_d M$ along γ keeping it complementary to $TW^s(d)$, it follows that $TM/TW^s(d) \cong T_d(W^u(d))$ [18]. Now, by assuming this isomorphism is orientation preserving, we have induced an orientation on $\mathcal{L}(c,d)$ [17].

As previously mentioned, we are mostly interested in the case when index$(c)-$ index$(d) = 1$, as then $\mathcal{L}(c,d)$ has dimension zero, and thus it consists of finite points. We want to be able to count these points, and in order to do this we need to know that our moduli space of trajectories is compact [17]. This can be deduced using the next definition and the following theorem from [25].

Definition 1.21. A *smooth manifold with corners* is a second countable Hausdorff space such that each point has a neighbourhood homeomorphic to $\mathbb{R}^{n-k} \times [0,\infty)^k$ for some $k \in \mathbb{N}$, and such that transition maps are smooth.

Theorem 1.22. *If M is a closed manifold, $f : M \to \mathbb{R}$ is a Morse function, X is a pseudo-gradient field such that (f,X) is Morse-Smale, then for any two critical points c, d, the moduli space of trajectories $\mathcal{L}(c,d)$ admits a natural compactification to a smooth manifold with corners $\overline{\mathcal{L}(c,d)}$, whose codimension k stratum is*

$$\overline{\mathcal{L}(c,d)}_k = \bigcup_{r_1,\dots,r_k \in Crit(f)} \mathcal{L}(c,r_1) \times \mathcal{L}(r_1,r_2) \times \cdots \times \mathcal{L}(r_{k-1},r_k) \times \mathcal{L}(r_k,d),$$

where the r_i are critical points of f distinct from each other and from c and d. We refer to $\mathcal{L}(c,d)$ as the space of broken trajectories from c to d.
When $k = 1$, as oriented manifolds we have

$$\partial\overline{\mathcal{L}(c,d)} = \bigcup_{r \neq c,d} (-1)^{\text{index}(c)+\text{index}(d)+1} \mathcal{L}(c,r) \times \mathcal{L}(r,d).$$

Intuitively, this theorem is telling us that we can somehow split a flow line from c to d into flow lines passing through intermediate critical points. Looking at some specific cases, if index$(c) = i$ and index$(d) = i-1$, then $\mathcal{L}(c,d)$ is compact. If index$(d) = i-2$ then there is a compactification of $\mathcal{L}(c,d)$ which is a 1-manifold whose boundary is

$$\partial\overline{\mathcal{L}(c,d)} = \bigcup_{r \in Crit_{i-1}(f)} \mathcal{L}(c,r) \times \mathcal{L}(r,d). \tag{1}$$

In this case the critical point r can only have index $i-1$. This is because $\mathcal{L}(c,r)$ is non-empty, therefore this implies that index$(r) \leq i-1$. Similarly, $\mathcal{L}(r,d)$ is non-empty, thus index$(r) \geq i-1$ and hence index$(r) = i-1$ [17].
Following the discussion in [17], we will briefly present an outline of the proof of the above theorem, which consists of two main parts, a compactness result

and a gluing theorem. The first part focuses on showing that a sequence of flow lines $\mathcal{L}(c,d)$ has a subsequence converging to a broken flow line in $\overline{\mathcal{L}(c,d)_k}$ for some $k \geq 0$. The second part shows that it is possible to perturb any broken flow line in $\overline{\mathcal{L}(c,d)_k}$ to obtain an integral flow line in $\mathcal{L}(c,d)$. This is done via perturbations that can by parametrised by special gluing parameters which can be taken to infinity and correspond to breaking the flow line at one of the intermediate critical points. Orientations also need to be taken into account [17]. The intuition provided by the theorem is that it is possible to compactify moduli spaces of flow lines into compact manifolds with corners via the addition of suitably "broken" flow lines [17]. The importance of this theorem also comes from the fact that it can be generalised to an infinite dimensional version, becoming a key result in Floer theory, which is discussed later on in this paper.

We are now able to proceed and define the Morse complex. Suppose M is a smooth compact closed manifold with a Riemannian metric g. Given a Morse function f on M, let $Crit_i(f)$ denote the critical points of f of index i. This is a finite set generating a free $\mathbb{Z}$-module, which is the chain module denoted by C_i:

$$C_i(f,g) := \mathbb{Z}\, Crit_i(f).$$

The boundary map, or differential, $\partial : C_i \to C_{i-1}$ then counts gradient flow lines [17]. Given $c \in Crit(f)$, it is defined as

$$\partial(c) := \sum_{d \in Crit_{i-1}(f)} \#\mathcal{L}(c,d) \cdot d.$$

The term $\#\mathcal{L}(c,d)$ is defined making use of the orientation on $\mathcal{L}(c,d)$: to each point in the moduli space there is an orientation induced by that of $W^u(c)$. We say this orientation is positive, and denote it by +1 if it agrees with that in $W^u(c)$, and −1 otherwise [18]. To determine $\#\mathcal{L}(c,d)$ we take the sum of these signs.

Definition 1.23. The *Morse complex* $\mathbb{M}_*$ is the chain complex

$$\cdots \to C_{i+1} \xrightarrow{\partial_{i+1}} C_i \xrightarrow{\partial_i} C_{i-1} \to \dots.$$

We now need to show that the Morse complex is indeed a complex, that is that $\partial^2 = \partial \circ \partial = 0$. This is one of the main results of this section.

Lemma 1.24. *For ∂ defined as above, $\partial^2 = \partial \circ \partial = 0$.*

Proof. This fact follows from Equation (1), previously discussed. Indeed, given two critical points whose indexes differ by 2, the boundary of the 2-cell they determine is an oriented compact 1-manifold. Such a manifold has zero points when counted with sign. To be more specific, suppose x is an element of the Morse complex, and that y is a basis element of it. We may denote by $\langle x, y\rangle$ the coefficient of y in the expression of x with respect to the basis. Now, for two critical points c and d of index i and $i-2$ respectively, we have:

$$\begin{aligned}
\langle \partial_{i-1}(\partial_i(c)), d\rangle &= \sum_{r \in Crit_{i-1}(f)} \langle \partial_i(c), r\rangle \langle \partial_{i-1}(r), d\rangle \\
&= \# \bigcup_{r \in Crit_{i-1}(f)} \mathcal{L}(c, r) \times \mathcal{L}(r, d) \\
&= \# \partial \overline{\mathcal{L}(c, d)} \\
&= 0.
\end{aligned}$$

The proof presented is an adaptation of the one in [17, 18]. ☐

Definition 1.25. The homology of the Morse complex $\mathbb{M}_*$ is known as the *Morse homology $MH_*(f, g)$*.

As it can be seen from the definition, we are keeping track of a dependance on the choice of Morse function f and the Riemannian metric g in our notation.

Example 1.3. We now compute the Morse homology of the torus T^2 with the height function h previously shown in Fig. 2. First of all we must note that the Smale condition is not satisfied since the two saddles b and c are joined by two trajectories [12]. However, we can slightly tilt the torus using the function h' (or equivalently perturb the metric g on it), so that it looks like the one in Fig. 4. We can check that the Morse complex is

$$\cdots \to 0 \to \mathbb{Z} \to \mathbb{Z}^2 \to \mathbb{Z} \to \ldots.$$

The next step is to determine the boundary maps by taking the signed count of trajectories between critical points. Working through the example, we may note that the trajectories in $\mathcal{L}(d, c)$ cancel out as they have opposite signs.

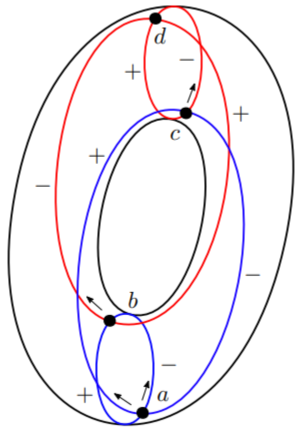

Figure 4: Signed Morse complex on the tilted torus T^2. Figure from [12].

This turns out to be the case for all the trajectories, since for each saddle point there are two trajectories going to a maximum which cancel out, and two trajectories going to a minimum cancelling out as well [17]. Hence all the boundary maps are trivial and the Morse Homology is

$$MH_i(h', g) = \begin{cases} \mathbb{Z} & i = 0, 2 \\ \mathbb{Z}^2 & i = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Note that Morse homology on the torus agrees with singular homology. We will soon find out that this is always the case.

An attentive reader might have noticed that we have used a specific Morse function and choice of pseudo-gradient field to compute the Morse homology of the torus in the example above. Thus it is natural to ask what would happen if we chose a different Morse function. An interesting aspect about Morse homology is that it turns out to be independent of the choice of Morse function and of the pseudo-gradient field used to define the boundary map [8], therefore it only depends on the smooth manifold we are working with.

Suppose we have two Morse functions f_0 and f_1 defined on a smooth manifold M. We can define a homotopy $F : M \times I \to \mathbb{R}$ from f_0 to f_1. In general, this function is not Morse, for example when the number of critical points changes, but we can overcome this problem if we can construct a morphism of complexes which induces isomorphisms on homology [12]. That is, we want to prove the following theorem.

Theorem 1.26. *Suppose f_0, $f_1 : M \to \mathbb{R}$ are two Morse functions on the manifold M. Let X_0, X_1 be pseudo-gradients adapted to f_0 and f_1 such that the Smale condition is satisfied. Then, there exists is a morphism of complexes $\Phi_* : (C_*(f_0), X_0) \to (C_*(f_1), X_1)$ which induces isomorphisms on homology.*

Proof. As this proof is quite long, this survey will only present a brief overview based on [12], but full details can be found in [8]. We begin by considering the homotopy $F : M \times [0,1] \to \mathbb{R}$, which is the smooth function defined as

$$F|_{t\in[0,\frac{1}{3}]} = f_0, \qquad F|_{t\in[\frac{2}{3},1]} = f_1.$$

F is referred to as an *end-constant interpolation.* It is possible to define the category $\mathbf{EndConstInt}(M)$ of Morse-Smale pairs on M, where the morphisms are given by the equivalence classes of end-constant interpolations. We may also introduce the category $\mathbf{MoCplx}(M)$ of Morse complexes for a Morse-Smale pair, with morphisms given by chain maps [12].

Let F and G be end-constant interpolations between (f_0, X_0) and (f_1, X_1), and between (f_1, X_1) and (f_0, X_0), respectively. Moreover, suppose Φ is a functor $\Phi : \mathbf{EndConstInt}(M) \to \mathbf{MoCplx}(M)$. It follows that $\Phi^F \circ \Phi^G$ and $\Phi^G \circ \Phi^F$ are each the identity map on the respective complexes, so Φ^F induces an isomorphism on homology. Therefore to prove the theorem all we need is to construct the functor Φ [12]. The proof then consists of two main steps:

1. Use F to construct $\Phi^F : C_*(f_0) \to C_*(f_1)$ and show it only depends on the equivalence class of F.
2. Prove functoriality in homology. This consists of proving two results. The first one is that if I is the constant homotopy from a map to itself, then $\Phi^I = Id(C_*(f_0))$. The second one is to prove that for end-constant interpolations F, G, H defined as above, with F going to f_0 to f_1, G from to f_1 to f_2, and H from f_0 to f_2, we have that $\Phi^G \circ \Phi^F = \Phi^H$.

We will briefly discuss how to construct Φ, but will not focus on the second part of the proof, which can be found in [2, 8]. The map F defined above can be extended to $\left[-\frac{1}{3}, \frac{4}{3}\right]$ by keeping the ends constant. F can be used to define a Morse function $g : \left[-\frac{1}{3}, \frac{4}{3}\right] \to \mathbb{R}$, with only two critical points, 0, a maximum, and 1, a minimum. We want g to decrease rapidly enough in $[0,1]$, that is for all $x \in M$ and $t \in (0,1)$ we require

$$\frac{\partial F}{\partial t}(x,t) + g'(t) < 0.$$

This condition can be achieved by compactness of M and by letting the critical value of 0 be very large [17]. From here we can define $F + g : M \times \left[-\frac{1}{3}, \frac{4}{3}\right] \to \mathbb{R}$, which is Morse with critical points

$$Crit(f_0) \times \{0\} \bigcup Crit(f_1) \times \{1\}.$$

Now, for all $a \in Crit(f_0)$, $(a,0)$ has index $\text{ind}(a) + 1$. Similarly, for all $b \in Crit(f_1)$, $(b,1)$ has index $\text{ind}(b)$. We may now take a pseudo-gradient field X adapted to $F + g$ such that:

- It coincides with $X_0 - \operatorname{grad} g$ on $M \times \left[-\frac{1}{3}, \frac{1}{3}\right]$.
- It coincides with $X_1 - \operatorname{grad} g$ on $M \times \left[-\frac{2}{3}, \frac{4}{3}\right]$.
- It satisfies the Smale condition.

The last condition is actually not needed, as we can always find an approximation of a pseudo-gradient field which satisfies the Smale condition. Now, it can be observed that

$$C_*(F + g|_{M \times [-\frac{1}{3}, \frac{1}{3}]}, X) = C_{*+1}(f_0, X_0),$$
$$C_*(F + g|_{M \times [\frac{2}{3}, \frac{4}{3}]}, X) = C_*(f_1, X_1).$$

Thus we obtain a decomposition of the chain complex $C_{i+1}(F + g, X) = C_i(f_0, X_0) + C_{i+1}(f_1, X_1)$ and of the boundary map ∂_X as

$$\partial_X = \begin{pmatrix} \partial_{X_0} & 0 \\ \Phi^F & \partial_{X_1} \end{pmatrix}.$$

The map $\Phi^F : C_*(f_0) \to C_*(f_1)$ is defined over the generators of $C_i(f_0)$ as

$$\Phi(c) := \sum_{d \in Crit_i(f_1)} n_X(c,d) \cdot d.$$

Here $n_X(c,d)$ denotes the number of trajectories of X that connect a critical point $c \in M \times \{0\} \cap Crit(F + g)$ to $d \in M \times \{1\} \cap Crit(F + g)$ [2]. The proof then continues by showing that Φ^F is the desired map and that $F \mapsto \Phi^F$ is a functor. □

We have just seen that Morse homology is independent of the Morse function defined on a manifold, so there is actually no need to specify Morse function and pseudo-gradient field considered in our notation. However, there is another result which shows us why Morse theory is so powerful and Morse homology in particular is an extremely useful tool. Earlier in Example 1.3 we noticed that the Morse homology computed agreed with singular homology on the torus. This turns out to always be the case, proving that Morse homology is not only independent of the Morse-Smale pair used to define it, but it actually only depends on the topological structure of the manifold [12], leading us to state the following fundamental fact about finite-dimensional Morse theory.

Theorem 1.27. *Let M be a smooth compact closed finite-dimensional manifold and let (f, X) be a Morse-Smale pair defined on it. Then there is a canonical isomorphism*

$$MH_*(f, X) \cong H_*(M).$$

We will not present a proof of this fact in this survey due to time constraints, but it is possible to find one in [8]. A well-known consequence of this theorem is that the Betti numbers of a manifold can be expressed using Morse homology, thus obtaining an equivalent definition as in the case of singular homology [12]. From here it is then possible to derive both the strong and weak Morse inequalities, which relate the numbers of critical points on a manifold to the Betti numbers. A derivation of such inequalities can be found in [8, 12]. Using the Morse inequalities it is then possible to prove the h-cobordism theorem and the smooth Poincaré conjecture for dimensions greater than 6. This sections served as simple and self-contained introduction to Morse theory highliting how far-reaching it is.

2 Lagrangian Floer homology

In this section we present the theory of Lagrangian Floer homology, and show it can be seen as an infinite-dimensional analogue of Morse homology. We will give the basic definitions but will assume some standard results from symplectic topology - these are covered in any introductory text such as [9] or [21]. The exposition is based on [3].

2.1 Introduction

2.1.1 Motivation

Lagrangian Floer homology was developed by Andreas Floer in the late 1980s to study the intersection properties of Lagrangian submanifolds. It lead to the solution of a particular case of Arnold's conjecture, which gives a lower bound for the number of fixed points of a symplectomorphism.

More precisely, recall that a *symplectic manifold* is an even-dimensional manifold M^{2n} equipped with a closed, non-degenerate 2-form $\omega \in \Omega^2(M)$. Given a symplectic manifold (M, ω) and a smooth family of functions $H_t \in C^\infty(M, \mathbb{R})$, non-degeneracy of ω gives a time-dependent vector field uniquely defined by $\omega(\cdot, X_t) = dH_t$. For a compact manifold M we can integrate the family X_t over $t \in [0, 1]$ to obtain a diffeomorphism $\varphi : M \to M$. We say φ is a *Hamiltonian diffeomorphism*, and we write $\mathrm{Ham}(X, \omega)$ for the space of Hamiltonian diffeomorphism of M. Then Arnold's conjecture states the following:

Conjecture 2.1 (Arnold's conjecture [1]). *Let $\psi \in \mathrm{Ham}(M, \omega)$ be a Hamiltonian diffeomorphism and $\mathrm{Fix}\,\psi$ its fixed points. Assume that for every fixed point $p \in \mathrm{Fix}\,\psi$, the linear map $d\psi(p)$ does not have 1 as an eigenvalue. Then*

$$\#\,\mathrm{Fix}\,\psi \geq \sum_i \dim H^i(M; \mathbb{Z}_2)$$

Note the fixed points of any diffeomorphism can be seen as the intersection points of its graph and the diagonal. In the context of symplectic topology, this is the intersection of two Lagrangian submanifolds.[2] Therefore, the study of fixed points is intimately related to the intersection of Lagrangian submanifolds.

In [10] Floer developed a homology theory to study the intersection properties of Lagrangian submanifolds. Roughly speaking, he considered a chain complex $CF(L_0, L_1)$ freely generated by the intersection points of L_0 and L_1 and equipped it with a differential $\partial : CF(L_0, L_1) \to CF(L_0, L_1)$ such that:

1. $\partial^2 = 0$, so the *Lagrangian Floer homology* $HF(L_0, L_1)$ of the complex is defined.
2. for $\varphi \in \mathrm{Ham}(M, \omega)$ there is an isomorphism $HF(L_0, L_1) \cong HF(L_0, \varphi(L_1))$.

[2]A submanifold $L \subset M$ is said to be a *Lagrangian* submanifold if $\omega|_L = 0$ and $\dim L = n$.

3. for any Lagrangian L there is an isomorphism $HF(L,L) \cong H^*(L)$ (with suitable coefficients).

The theory is nowadays known as Lagrangian Floer homology, and with it he was able to solve Arnold's conjecture for a particular class of closed symplectic manifolds. His result is in fact more general and stated the following:

Theorem 2.2 (Floer [10])**.** *Assume that any disk on M with boundary on L has area zero, i.e. $[\omega]\cdot\pi_2(M,L) = 0$. Let $\varphi \in \mathrm{Ham}(M,\omega)$ be such that L and $\varphi(L)$ intersect transversely. Then*

$$\#(L \cap \varphi(L)) \geq \sum_i \dim H^i(L;\mathbb{Z}_2) \tag{2}$$

Under the assumption of the theorem, Arnold's conjecture is a corollary for the case when L is the diagonal $\Delta \subset M \times M$ and $\varphi = \mathrm{id} \times \psi$.

Ignoring coefficients, Theorem 2.2 can be deduced from the three properties of Lagrangian Floer homology stated above. Indeed, the right hand side of (2) is the dimension of the singular cohomology of L, which by the third and second properties coincides with the dimension of $HF(L,\varphi(L))$. This is bounded above by the number of generators in the chain complex, which is precisely the left hand side of (2).

The intersection properties of Lagrangians submanifolds are extremely important in symplectic topology and many questions remain unsolved. As an example, the following conjecture - in the line of Theorem 2.2 - has only been solved for specific cases and is still an open problem:

Conjecture 2.3 (Arnold-Givental conjecture)**.** *Let $\varphi \in \mathrm{Ham}(M,\omega)$ and $L \subset M$ be a compact Lagrangian submanifold that is the fixed point set of an antisymplectic involution.*[3] *Assume that L and $\varphi(L)$ intersect transversely. Then*

$$\#(L \cap \varphi(L)) \geq \sum_i \dim H^i(L;\mathbb{Z}_2).$$

It is worth noting that it implies Arnold's conjecture, since the diagonal is the fixed point set of the antisymplectic involution of $(M \times M, \omega \oplus -\omega)$ given by swapping the factors.

[3]An *antisymplectic involution* is a diffeomorphism $\varphi : M \to M$ such that $\varphi^2 = \mathrm{id}_M$ and $\varphi^*\omega = -\omega$.

2.1.2 Preliminaries on complex geometry

In modern symplectic topology one often studies symplectic manifolds by analyzing maps from Riemann surfaces satisfying a certain PDE. These are called (pseudo)-holomorphic maps, and can be seen as a generalization of holomorphic curves in complex geometry. In this section we will briefly introduce the necessary background to work with pseudo-holomorphic maps.

An *almost complex structure* on a manifold M is a bundle automorphism $J : TM \to TM$ such that $J^2 = -\mathrm{Id}$. The standard example comes from complex manifolds, i.e. even-dimensional manifolds whose transition functions are biholomorphisms. In this case, if $(x_1 + iy_1, \dots, x_n + iy_n) \in \mathbb{C}^n \cong \mathbb{R}^{2n}$ are local coordinates, the natural automorphism $J : TM \to TM$ locally defined as

$$J\left(\frac{\partial}{\partial x_i}\right) = \frac{\partial}{\partial y_i}, \quad J\left(\frac{\partial}{\partial y_i}\right) = -\frac{\partial}{\partial x_i},$$

defines an almost complex structure on M.

Given two almost complex manifolds (M, J) and (M, J'), we have a natural endomorphism of the space of bundle automorphisms $TM \to TM'$ given by

$$\begin{aligned} \Psi : \mathrm{Hom}(TM, TM') &\to \mathrm{Hom}(TM, TM') \\ F &\mapsto J' \circ F \circ J. \end{aligned}$$

It is immediate that $\Psi^2 = \mathrm{Id}$, so that the eigenvalues of Ψ are ± 1. This gives a splitting

$$\mathrm{Hom}(TM, TM') = \mathrm{Hom}(TM, TM')^{0,1} \oplus \mathrm{Hom}(TM, TM')^{1,0},$$

as a direct sum of the eigenspaces corresponding to 1 and -1 respectively. We say a smooth map $\varphi : M \to M'$ is *holomorphic* if its differential $d\varphi \in \mathrm{Hom}(TM, TM')$ satisfies $(d\varphi)^{0,1} = 0$, where $(d\varphi)^{0,1}$ is the projection of $d\varphi$ to the first factor. In other words, $d\varphi$ satisfies $d\varphi = -J' \circ d\varphi \circ J$, i.e.

$$J' \circ d\varphi = d\varphi \circ J.$$

Thus, holomorphic maps are those respecting the complex structures. More generally, if $d\varphi$ satisfies the inhomogenous equation $(d\varphi)^{0,1} = \nu$ for some *perturbation* term ν, we say φ is *pseudo-holomorphic*. We will often use the words holomorphic and pseudo-holomorphic indistinctly.

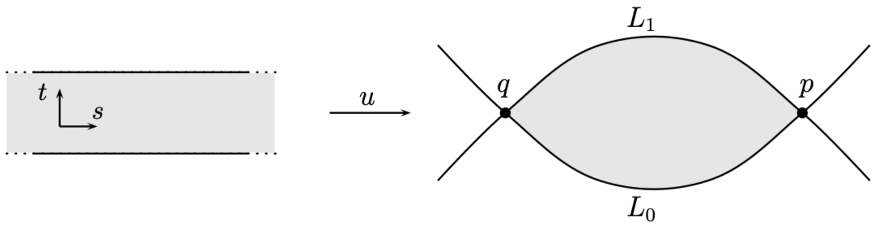

Figure 5: Schematic picture of the type of holomorphic maps that we consider in the differential.

Lastly, in the case of a symplectic manifold (M, ω), we say that an almost complex structure J on M is *compatible* with ω if $g(\cdot,\cdot) := \omega(\cdot, J\cdot)$ defines a Riemannian metric. It is a classical result that the space of compatible almost complex structure on any symplectic manifold is contractible and non-empty [21], so the choice of almost complex structure is essentially unique. Thus, we will pick *any* compatible almost-complex structure when neccessary.

2.2 Construction of Lagrangian Floer homology

The idea of Lagrangian Floer homology is the following. Let L_0 and L_1 be compact Lagrangians, which we will assume for now to be transverse (this condition will be relaxed later). Recall that the idea behind the Morse complex was:

- $C = \bigoplus_i C_i$ is freely generated by critical points;
- $\partial : C_i \to C_{i-1}$ is obtained by considering the moduli-space of gradient flow lines between critical points.

In an analogy with this idea of studying geometric configurations between generators, we could try to define the Floer complex $(CF(L_0, L_1), \partial)$ as:

- $CF(L_0, L_1)$ is freely generated by the intersection points of L_0 and L_1;
- $\partial : CF(L_0, L_1) \to CF(L_0, L_1)$ is obtained by considering the moduli-space of *pseudo-holomorphic strips* with boundary in L_0 and L_1 (see Figure 5). More precisely, given intersection points $p, q \in L_0 \cap L_1$, we will consider the space of smooth maps $u : \mathbb{R}_s \times [0,1]_t \subset \mathbb{C} \to M$ satisfying

$$\frac{\partial u}{\partial s} + J(u)\frac{\partial u}{\partial t} = 0 \tag{3}$$

and with boundary conditions

- $u(s,i) \in L_i$ for $i = 0, 1$ and all $s \in \mathbb{R}$;
- $\lim_{s\to-\infty} u(s,t) = q$ and $\lim_{s\to\infty} u(s,t) = p$ for all $t \in [0,1]$.

This is indeed the idea behind the construction, but there are several technical difficulties. These will be discussed in the following sections.

2.2.1 Gradings

First of all, in Morse theory we have a grading on the chain complex, i.e. we can assign degrees (or indexes) to the generators. This is neccessary to work only with zero-dimensional moduli-spaces. Equipping the Lagrangian Floer complex with a grading is a much more subtle matter, and in fact it is not necessary for the theory to work. Instead, we will work with a "grading" on the moduli-space itself. We need the following two definitions:

- Denote by $LGr(n)$ the Grassmannian of Lagrangian n-planes in $(\mathbb{R}^{2n}, \omega_0)$. The unitary group $\mathrm{U}(n) \subset \mathrm{GL}(n,\mathbb{C}) \subset \mathrm{GL}(n,\mathbb{R}^{2n})$ acts transitively on Lagrangian planes with stabilizer $\mathrm{O}(n)$, and thus $LGr(n) \cong \mathrm{U}(n)/\mathrm{O}(n)$. Composing any loop in $LGr(n)$ with the map $\det^2 : \mathrm{U}(n)/(n) \to S^1$ gives a map $S^1 \to S^1$. We define the *Maslov index* of any loop γ to be the degree of the class $[\det^2 \circ \gamma] \in \pi_1(S^1) \cong \mathbb{Z}$.
- A classical result in symplectic topology states that $\mathrm{Sp}(2n,\mathbb{R})$ - the space of symplectomorphisms of $(\mathbb{R}^{2n}, \omega_0)$ - acts transitively on pairs of transverse Lagrangian planes. Thus, given transverse Lagrangian n-planes $l_0, l_1 \in LGr(n)$, there exists $A \in \mathrm{Sp}(2n,\mathbb{R})$ such that $A(l_0) = \mathbb{R}^n$ and $A(l_1) = (i\mathbb{R})^n$, where we have made the usual identification $\mathbb{R}^{2n} \cong \mathbb{C}^n$, The "clockwise path" $A^{-1}((e^{-i\pi t/2}\mathbb{R})^n)$ from l_0 to l_1 is called the *canonnical short path.*

Given a pseudo-holomorphic strip u, denote by γ_p the canonical short path from T_pL_0 to T_pL_1 and by γ_q that from T_qL_0 to T_qL_1. For $i \in \{0,1\}$, let γ_i be the path u^*TL_i, oriented with s going from $+\infty$ to $-\infty$. Fixing a trivialisation of u^*TM, we can view all this paths as paths in $LGr(n)$.

Definition 2.4. The *Maslov index* $\mathrm{ind}(u)$ of a pseudo-holomorphic strip u is the Maslov index of the loop in $LGr(n)$ given by following $-\gamma_0, \gamma_p, \gamma_1$ and $-\gamma_q$.

With this definition, the Maslov index of u depends only on its homotopy class (where the homotopy goes through maps $(\mathbb{R}\times[0,1],\mathbb{R}\times\{0,1\})\to(M,L_0\cup L_1)$ with appropriate boundary conditions). Furthermore, under some transversality assumptions the moduli-space $\widehat{\mathcal{M}}(p,q;[u])$ of strips whose homotopy class is that of u is a manifold of dimension $\operatorname{ind}([u])$. It follows that the moduli-space $\mathcal{M}(p,q,[u]) := \widehat{\mathcal{M}}(p,q;[u])/\mathbb{R}$ of unparametrized strips (where the $\mathbb{R}$-action is $a\cdot u := u(\cdot+a,\cdot)$) is a manifold of dimension $\operatorname{ind}([u])-1$. If in addition $\mathcal{M}(p,q,[u])$ is compact, when $\operatorname{ind}([u])=1$ we get a well-defined number $\#\mathcal{M}(p,q,[u])\in\mathbb{Z}_2$ as the mod 2 count of its points. In the next sections we will deal with these transversality and compactness issues.

2.2.2 Transversality

Transversality is needed to ensure the moduli-spaces under consideration are smooth manifolds of the expected dimension. Using an infinite-dimensional analogue of the implicit function theorem, this would follow if the linearization of the operator mapping a strip u to $(du)^{0,1}$ was surjective at every solution. Although this is not always the case, it can be achieved by replacing J with a t-dependent family J_t of compatible almost-complex structures.

A more basic issue is how to define the Lagrangian Floer complex when the two Lagrangians do not intersect transversely, since the intersection points need not be finite in this case (in particular, we would like to define the Lagrangian Floer homology $HF(L,L)$ of a Lagrangian with itself). In view of the desired property of Hamiltonian isotopy invariance, the natural solution is to use a Hamiltonian diffeomorphism $\varphi\in\operatorname{Ham}(M,\omega)$ to perturb L_1 so that it intersects L_0 transversely, and then define $CF(L_0,L_1) := CF(L_0,\varphi(L_1))$.

An equivalent solution is to consider *perturbed* pseudo-holomorphic equations. To understand this, suppose $H\in C^\infty([0,1]\times M,\mathbb{R})$ is the Hamiltonian generating the Hamiltonian isotopy φ^t and $\tilde{u}:\mathbb{R}\times[0,1]\to M$ is a pseudo-holomorphic strip between intersection points $p,q\in L_0\cap\varphi(L_1)$ and with boundary in L_0 and $\varphi(L_1)$. The strip $u(s,t) := \varphi^t(\tilde{u}(s,t))$ has boundary in L_0 and L_1, but in general it will not be pseudo-holomorphic. Instead, since

$$\frac{\partial u}{\partial t} = (\varphi^t)_*\left(\frac{\partial\tilde{u}}{\partial t}\right) + X_H$$

we see that

$$\frac{\partial u}{\partial s} + J_t' \left(\frac{\partial u}{\partial t} - X_H \right) = 0, \tag{4}$$

where $J_t' := (\varphi^t)_* J_t$. In this case, u no longer converges to intersection points as $s \to \pm\infty$, but instead to flow lines of X_H.

Therefore, for non-transverse Lagrangians we can define the Floer complex by studying the moduli-space of strips satisfying the perturbed equation (4), mapping the boundaries of the strip to L_0 and L_1 and converging to flow lines of X_H between points in L_0 and L_1.

2.2.3 Compactness

Compactness of the moduli-space is one of the most complicated matters in Lagrangian Floer theory. It is required to ensure that we can actually count the points in the zero-dimensional moduli-space $\mathcal{M}(p, q, [u])$ when $\text{ind}([u]) = 1$, as well as to show that the differential squares to zero.

Compactness of the moduli-spaces (or, more precisely and as in Morse theory, the moduli-spaces with some extra points) is based on Gromov's compactness theorem, which states that any sequence of pseudo-holomorphic strips has a subsequence converging, up to reparametrization, to a *nodal tree* of pseudo-holomorphic strips. For a sequence of pseudo-holomorphic strips $u_n : \mathbb{R} \times [0, 1] \to M$ with boundary in L_0 and L_1, there are three possible scenarios:

1. *strip breaking*: in this case, for a suitable sequence $a_n \to \pm\infty$ the translated strips $u_n(s - a_n, t)$ converge (on compact sets) to non-constant limit strips;

2. *disc bubbling*: in this case, suitable rescalings of u_n converge to a pseudo-holomorphic disc with boundary fully contained in a single Lagrangian;

3. *sphere bubbling*: lastly, this occurs when suitable rescalings of u_n converge to a pseudo-holomorphic sphere in M.

To illustrate the first two phenomena, consider inside the cylinder $M = \mathbb{R} \times S^1$ two closed Lagrangians intersecting at two points, L_0 being non-contractible and L_1 contractible (see Figure 6 left). Choosing a chart $U \cong \mathbb{C}$ containing L_1 and such that L_0 corresponds to the real axis and L_1 to the unit circle,

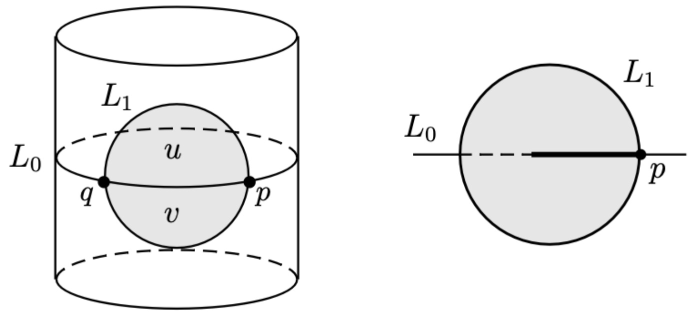

Figure 6: Example where strip breaking and disc bubbling appear.

we define holomorphic maps $u_n : D^2 \setminus \{\pm 1\} \subset \mathbb{C} \to \mathbb{C}$ given by

$$u_n(z) = \frac{z^2 + \alpha_n}{1 + \alpha_n z^2}, \quad \alpha_n \in (-1, 1).$$

(Note that $D^2 \setminus \{\pm 1\}$ is biholomorphic to $\mathbb{R} \times [0, 1]$ by the Riemann mapping theorem: composing with the biholomorphism we get a holomorphic map from the strip $\mathbb{R} \times [0, 1]$ as usual.) The maps u_n take the top boundary of $D^2 \setminus \{\pm 1\}$ to L_0 and the bottom to L_1, and α_n is the endpoint of the slit (see Figure 6 right).

If we consider a sequence $\alpha_n \to -1$, the strips u_n converge to a broken strip with components u and v. On the other hand, for a sequence $\alpha_n \to 1$ the strips u_n converge to a constant strip at p and a disc bubble whose boundary is precisely L_0.

Bubbling is such a big issue that, in general, the differential will *not* square to zero, and the Lagrangian Floer homology will not be defined. However, if we can guarantee that disc and sphere bubbling do not occur, we get a well-defined chain complex with $\partial^2 = 0$. This can be achieved under some topological assumptions like $\pi_2(M, L_i) = 0$, which is the technical hypothesis in Theorem 2.7.

2.2.4 Coefficients

Up to now we have argued that when $\text{ind}([u]) = 1$ the moduli-spaces $\mathcal{M}(p, q, [u])$ are compact manifolds of dimension zero. We could then try to define the differential as

$$\partial(p) = \sum_{\substack{q \in L_0 \cap L_1 \\ [u] \mid \operatorname{ind}([u])=1}} \#\mathcal{M}(p, q, [u]) \cdot q. \tag{5}$$

for $p \in L_0 \cap L_1$. However, the sum in the right hand side of (5) might not be well-defined: although the number $\#\mathcal{M}(p, q, [u]) \in \mathbb{Z}_2$ is well-defined, there might be infinitely many different homotopy classes $[u]$ with $\operatorname{ind}([u]) = 1$. This issue is solved by changing the coefficients of the chain complex.

Given a pseudo-holomorphic strip u, we define its *energy* to be

$$E(u) := \int_{\mathbb{R}\times[0,1]} u^*\omega = \int_{\mathbb{R}\times[0,1]} \omega\left(\frac{\partial u}{\partial s}, \frac{\partial u}{\partial t}\right) dsdt = \int_{\mathbb{R}\times[0,1]} \left|\frac{\partial u}{\partial s}\right|^2 dsdt.$$

Remark 2.5. *Note that, by Stoke's theorem, the energy of u depends only on its homotopy class: if u and u′ are homotopic through maps* $(\mathbb{R} \times [0,1], \mathbb{R} \times \{0,1\}) \to (M, L_0 \cup L_1)$*, writing G for the homotopy we get*

$$\begin{aligned}
\int_{\mathbb{R}\times[0,1]} (u^*\omega - u'^*\omega) &= \int_{\mathbb{R}\times[0,1]} (u^*\omega - u'^*\omega) + \int_{\mathbb{R}\times\{0,1\}\times[0,1]} G^*\omega && (\omega|_{L_i} = 0) \\
&= \int_{\partial(\mathbb{R}\times[0,1]\times[0,1])} G^*\omega \\
&= \int_{\mathbb{R}\times[0,1]\times[0,1]} d(G^*\omega) && \textit{(Stoke's)} \\
&= 0 && (d\omega = 0)
\end{aligned}$$

and thus $E(u) = E(u') =: E([u])$*.*

Although strips can have arbitrarily high energy, Gromov's compactness theorem ensures that, given any energy bound $E_0 \in \mathbb{R}$, there exist finitely many homotopy classes $[u]$ with energy $E([u]) \leq E_0$. In particular, there are only finitely many strips with a given energy, and the possible energie values are either finite or tend to infinity. Thus, by using the Novikov field over $\mathbb{Z}_2$

$$\Lambda := \left\{ \sum_{i=0}^{\infty} a_i T^{\alpha_i} \mid a_i \in \mathbb{Z}_2, \alpha_i \in \mathbb{R}, \lim_{i\to\infty} \alpha_i = \infty \right\},$$

we can finally get a well-defined Lagrangian Floer chain complex.

Definition 2.6. Given transverse compact Lagrangians L_0 and L_1, the *Lagrangian Floer chain complex* $(CF(L_0, L_1), \partial)$ consists of:

- $CF(L_0, L_1) \coloneqq \bigoplus_{p \in L_0 \cap L_1} \Lambda \cdot p$ is the free Λ-module generated by the intersection points $p \in L_0 \cap L_1$;
- $\partial : CF(L_0, L_1) \to CF(L_0, L_1)$ is the Λ-linear map defined on generators $p \in L_0 \cap L_1$ as

$$\partial(p) \coloneqq \sum_{\substack{q \in L_0 \cap L_1 \\ [u] \mid \mathrm{ind}([u])=1}} \#\mathcal{M}(p, q, [u]) T^{E([u])} \cdot q.$$

Note that by Gromov's compactness theorem the formal sum

$$\sum_{[u] \mid \mathrm{ind}([u])=1} \#\mathcal{M}(p, q, [u]) T^{E([u])}$$

belongs to the Novikov field for any $q \in L_0 \cap L_1$.

After dealing with all the technical difficulties, we arrive at the main Theorem of this section, originally due to Floer:

Theorem 2.7. *Assume that all discs with boundary in L_0 and L_1 have symplectic area zero, i.e. $[\omega] \cdot \pi_2(M, L_0) = [\omega] \cdot \pi_2(M, L_1) = 0$. Then the Floer differential is well-defined, satisfies $\partial^2 = 0$, and, up to isomorphism, the Floer cohomology $HF(L_0, L_1)$ is independent of the chosen almost-complex structure J and invariant under Hamiltonian isotopies of L_0 or L_1*

Proof. First we prove that the Floer homology is well-defined, i.e. that $\partial^2 = 0$. The flavor of this proof is the same as that in Lemma 1.24.

Suppose that L_0 and L_1 are transverse (otherwise perturb L_1 by a Hamiltonian diffeomorphism; alternatively, as discussed previously, consider a perturbed flow equation). Given generators $p, q \in L_0 \cap L_1$, the coefficient of q in $\partial^2(p)$ comes from studying the moduli-space

$$\bigsqcup_{r \in L_0 \cap L_1} \mathcal{M}(p, r, [u']) \times \mathcal{M}(r, q, [u'']), \quad \mathrm{ind}([u']) = \mathrm{ind}([u'']) = 1. \tag{6}$$

The claim is that this is the boundary of a compact 1-manifold with boundary; in particular, it has an even number of points. To show this, consider the compactification $\bar{\mathcal{M}}(p, q, [u])$ of the one-dimensional moduli-space $\mathcal{M}(p, q, [u])$ of strips of index 2 with homotopy class $[u]$. Gromov's compactness implies that $\bar{\mathcal{M}}(p, q, [u])$ is obtained from $\mathcal{M}(p, q, [u])$ by adding broken strips connecting p to q and representing the total class $[u]$ (here it is crucial that there is no sphere or disk bubbling, otherwise these configurations should be added too). Since $\operatorname{ind}([u']) + \operatorname{ind}([u'']) = \operatorname{ind}([u]) = 2$ and any strip must have index at least one,[4] the only possibility is that $\operatorname{ind}([u']) = \operatorname{ind}([u'']) = 1$. Therefore, only terms of the form of those in (6) can appear in the boundary of $\bar{\mathcal{M}}(p,q,[u])$.

To finish we must prove that all such terms appear. This is guaranteed by a *gluing* theorem, which states that any broken strip is locally the limit of a unique family of honest index 2 strips. That is, Gromov's compactness and the gluing theorem state that when $\operatorname{ind}([u]) = 2$:

$$\partial\bar{\mathcal{M}}(p,q,[u]) = \bigsqcup_{\substack{r\in L_0\cap L_1\\ [u']+[u'']=[u]\\ \operatorname{ind}([u'])=\operatorname{ind}([u''])=1}} \mathcal{M}(p,r,[u']) \times \mathcal{M}(r,q,[u''])$$

It follows that, with $\mathbb{Z}_2$ coefficients,

$$\sum_{\substack{r\in L_0\cap L_1\\ [u']+[u'']=[u]\\ \operatorname{ind}([u'])=\operatorname{ind}([u''])=1}} \#\mathcal{M}(p,r,[u'])\#\mathcal{M}(r,q,[u''])T^{E([u']+E([u''])} = 0,$$

where we have used that $E([u']) + E([u'']) = E([u])$. Summing over all possible homotopy classes $[u]$ with index 2, the result is precisely the coefficient of q in $\partial^2(p)$. This proves that $\partial^2 = 0$.

Lastly, we show that Lagrangian Floer cohomology does not depend on the choice of almost-complex structure J or on the Hamiltonian perturbation H. Suppose we have two choices (J,H) and (J',H') for which transverality holds.[5] Choose a *generic* smooth deformation $(J_\tau, H_\tau), \tau \in [0,1]$ between both choices (this is always possible since both spaces are contractible). We define a *continuation map* $F : CF(L_0,L_1;H,J) \to CF(L_0,L_1;H',J')$ by counting solutions of

$$\frac{\partial u}{\partial s} + J_{g(s)}(t,u)\left(\frac{\partial u}{\partial t} - X_{H_{g(s)}}(t,u)\right) = 0 \tag{7}$$

[4]Although we have not discussed it here, the index of a solution u can alternatively be defined by considering the linearization at u of the operator mapping u to $(du)^{0,1}$, which is a Fredholm operator. Then one defines the index of u to be the index of the corresponding Fredholm operator. With this definition, the last statement follows.

[5]Thus, both the almost-complex structure and the Hamiltonian might be t-dependent.

where $g(s)$ is a smooth function which equals 1 for $s << 0$ and 0 for $s >> 0$. Given generators $p \in CF(L_0, L_1; H, J)$ and $p' \in CF(L_0, L_1; H', J')$, the coefficient of p' in $F(p)$ will be be the number of index-zero strips converging to p as $s \to \infty$ and to p' as $s \to -\infty$, weighted by energy as usual. (Note that (7) is no longer s-invariant, so we work with $\widehat{\mathcal{M}}$ instead of $\mathcal{M}$.)

To see that F is a chain map, i.e. $\partial' \circ F = F \circ \partial$, we study the spaces of index 1 solutions of (7). Just as before, these are 1-dimensional manifolds, whose boundary consists of an index 0 solution of (7) either preceded by an index 1 J-holomorphic strip with perturbation data H or followed by an index 1 J'-holomorphic strip with perturbation data H'. The map $F \circ \partial$ counts the former and the map $\partial' \circ F$ the latter, and being the boundary of a compact 1-manifold, these two numbers must be equal mod 2. Thus F is a chain map.

Using a function g with the opposite behavior (i.e. $g(s) = 1$ for $s >> 0$ and $g(s) = 0$ for $s << 0$) the same argument gives a continuation map $F' : CF(L_0, L_1; H', J') \to CF(L_0, L_1; H, J)$. The claim is that these maps are quasi-inverses, i.e. that $F' \circ F$ is chain homotopic to $\mathrm{id}_{CF(L_0,L_1,H)}$ and similarly for $F \circ F'$. To see this, consider the perturbed equation (7) with $g(s)$ equal to 0 for $|s| >> 0$ and non-zero on a certain interval $(-a, a)$. We would like to define a map $G : CF(L_0, L_1, H) \to CF(L_0, L_1, H)$ by counting index -1 solutions of (7); however, in general there will be no such solutions. Although we will not give details, we sketch the argument, as it is not a simple generalization but a new idea.

We consider a homotopy $\{g_\lambda\}_{\lambda \in [0,1]}$ between $g_0 = g$ and the constant function $g_1 \equiv 0$ (note the latter reduces (7) to the equation of J-holomorphic strips with perturbation H). There exist no index -1 solutions for g_0 and g_1, but there might be *accidental solutions*, i.e. values of λ for which equation (7) has solutions of index -1.[6] It turns out that counting solutions for all accidental values is well-defined, and this will be the chain homotopy. □

This theorem completes the construction of Lagrangian Floer homology. Some important topics we have not covered include how to give a grading to the chain complex, the proof that $HF(L, L) \cong H^*(L; \Lambda)$ or how to work with fields other than $\mathbb{Z}_2$, where orientation comes into play; see [3] for such matters.

[6]These will be values of λ for which $(J_{g_\lambda}, H_{g_\lambda})$ is no longer transverse.

2.3 Formulation with the action functional

Lagrangian Floer theory can alternatively be introduced as an infinite dimensional analogue of Morse theory. That is, we want generators to be critical points of a suitable function and the differential to count flow lines between critical points. The idea is based on that in [24] for Floer homology.

Consider the path space $\mathcal{P}(L_0, L_1) := \{\gamma : [0,1] \to M \mid \gamma(0) \in L_0, \gamma(1) \in L_1\}$. Its universal cover $\tilde{P}(L_0, L_1)$ can be thought of as tuples $(\gamma, [\Gamma])$, where $\gamma \in \mathcal{P}(L_0, L_1)$ is a path from L_0 to L_1 and $[\Gamma]$ is the equivalence class of a homotopy $\Gamma : [0,1] \times [0,1] \to M$ between γ and a fixed base point in its connected component.[7] On it we define the functional

$$\begin{aligned} \mathcal{A} : \tilde{P}(L_0, L_1) &\to \mathbb{R} \\ (\gamma, [\Gamma]) &\mapsto -\int_\Gamma \omega. \end{aligned}$$

Remark 2.8. *For $\mathcal{A}$ to be well-defined (i.e. not dependent on the choice of representative for $[\Gamma]$) we need some extra hypothesis. More precisely, we should ask that M is aspherical, i.e. $[\omega]\cdot\pi_2(M) = 0$. In this case, if $\Gamma' \in [\Gamma]$ is another homotopy from a fixed base path η to γ, we can restrict a homotopy $G : [0,1]^2 \times [0,1] \to M$ between $\Gamma = G(\cdot,\cdot,0)$ and $\Gamma' = G(\cdot,\cdot,1) = \Gamma$ to the boundary $\partial([0,1]^2 \times [0,1])$ to get a map $G|_{\partial([0,1]^2\times[0,1])} : S^2 \to M$. Now note that the bottom and top faces of the cube are Γ and Γ', two of the lateral faces map to L_0 and L_1 (where $\omega = 0$), while the remaining two stay fixed at γ and γ' ($G(t,0,u) = \eta(t)$ and $G(t,1,u) = \gamma(t)$ for all $u \in [0,1]$). Therefore, the area of this sphere is precisely $\int_\Gamma \omega - \int_{\Gamma'} \omega$, and this vanishes by assumption.*

We claim that the critical points of $\mathcal{A}$ correspond to constant paths at intersection points, and that flow lines of the gradient of $\mathcal{A}$ (with respect to a suitable metric) are pseudo-holomorphic strips bounded by L_0 and L_1.

To see this, we first analyze the tangent space of these spaces. Given a path $\gamma \in \mathcal{P}(L_0, L_1)$, we can think of the tangent space $T_\gamma\mathcal{P}(L_0, L_1)$ as a suitable space of sections of the pullback bundle γ^*TM. Namely, a tangent vector $X \in T_\gamma\mathcal{P}(L_0, L_1)$ will be a section $X : [0,1] \to \gamma^*TM$ such that $X(0) \in T_{\gamma(0)}L_0$ and $X(1) \in T_{\gamma(1)}L_1$. Furthermore, note that we have natural isomorphisms

[7] The path space need not be connected. Thus, to be precise we would have to consider the universal cover of each connected component.

$T_{(\gamma,[\Gamma])}\tilde{\mathcal{P}}(L_0, L_1) \cong T_\gamma P(L_0, L_1)$, so that we can think of tangent vectors to $\tilde{\mathcal{P}}(L_0, L_1)$ at a point $(\gamma, (\Gamma])$ just as before.

To define flow lines we also need the notion of a gradient, and thus a metric on $\tilde{\mathcal{P}}(L_0, L_1)$. We can do this as follows. Choosing a compatible almost complex structure $J \in \mathrm{End}(TM, TM)$, we define an inner product on $T\tilde{\mathcal{P}}(L_0, L_1)$ as

$$\langle X, Y\rangle := \int_{[0,1]} \omega_{\gamma(t)}(X(t), JY(t))dt, \quad X, Y \in T_{(\gamma,[\Gamma])}\tilde{\mathcal{P}}(L_0, L_1).$$

An easy calculation then shows that $d\mathcal{A}_{(\gamma,[\Gamma])} \cdot X = \langle J\dot{\gamma}, X\rangle$, and thus

$$\operatorname{grad} \mathcal{A}_{(\gamma,[\Gamma])} = J\dot{\gamma}. \tag{8}$$

From this we deduce:

1. Since J is an automorphism (in particular, injective), critical points are paths with $\dot{\gamma} = 0$, i.e. constant paths which necessarily live at intersection points.
2. By definition, gradient flow lines are smooth maps $u : \mathbb{R} \to \tilde{\mathcal{P}}(L_0, L_1)$ such that $\frac{\partial u}{\partial s} = -\operatorname{grad} \mathcal{A}$. Thinking of u as a map $u : \mathbb{R} \times [0, 1] \to M$ (with the obvious boundary conditions), using (8) we can describe flow lines by the equation
$$\frac{\partial u}{\partial s} + J(u)\frac{\partial u}{\partial t} = 0,$$
which is precisely equation (3) of a pseudo-holomorphic strip.

Thus, we can think of the Lagrangian Floer complex as the infinite-dimensional Morse complex associated to the functional $\mathcal{A}$. However, the difficulties of infinite-dimensional analysis make this approach much harder. We will instead stick to the geometric study of moduli-spaces.

2.4 Application

In this section we will use this machinery to prove a displaceability result. Recall two subsets $N_1, N_2 \subset M$ are said to be *(Hamiltonian) displaceable* if there exists $\varphi \in \mathrm{Ham}(M, \omega)$ such that $N_1 \cap \varphi(N_2) = \varnothing$. We want to characterize which closed embedded Lagrangians of the cylinder $M = \mathbb{R} \times S^1$ are displaceable.[8]

[8] Note that since $\dim M = 2$ Lagrangian submanifolds are the same as one-dimensional submanifolds.

Proposition 2.9. *Given two non-contractible closed curves $L_0, L_1 \subset M$, they are displaceable if and only if $HF(L_0, L_1) = 0$.*

Proof. One of the directions is immediate: if they are displaceable by $\varphi \in \text{Ham}(M, \omega)$, then $CF(L_0, \varphi(L_1)) = 0$ and thus $HF(L_0, L_1) = HF(L_0, \varphi(L_1)) = 0$.[9]

To see the other direction, first note that a non-contractible closed curve will wrap around the cylinder exactly once (in other words, the embedding $\gamma : S^1 \to M$ defining the curve induces an isomorphism on π_1). The main part of the proof consists of the following Lemma:

Lemma 2.10. *If L_0 and L_1 are homotopic through $H : S^1 \times [0,1] \to M$ and*

$$\int_{S^1 \times [0,1]} H^* \omega = 0$$

then L_0 is Hamiltonian isotopic to L_1.

In other words, if L_0 and L_1 bound a 2-chain of vanishing area (e.g. L and $\psi(L)$ in Figure 7), then they are Hamiltonian isotopic. We omit the proof of this Lemma as it is a matter of symplectic topology and not Lagrangian Floer theory. An immediate corollary is:

Corollary 2.11. *Every non-contractible closed curve $L \subset M$ is Hamiltonian isotopic to a (unique) horizontal circle $C_z := \{z\} \times S^1$.*

With this in mind, consider horizontal circles C_{z_0} and C_{z_1} Hamiltonian isotopic to L_0 and L_1. If L_0 and L_1 are non-displaceable, then $z_0 = z_1 = z$. We must prove that $HF(C_z, C_z) \neq 0$.

Choose a perturbation of C_z as that of Figure 7. This can be achieved through a Hamiltonian on $M = S^1 \times \mathbb{R}$ of the form $H(p,t) = h(p)$, where $h : S^1 \to \mathbb{R}$ has a unique maximum and minimum at p and q. In this case the Lagrangian Floer chain complex $CF(C_z, \psi(C_z))$ is generated by p and q, and there are only two index 1 holomorphic strips (shaded in the figure). By our convention they both go from q to p, thus $\partial(q) = 0$. On the other hand,

[9] In fact, this is a general result: if two Lagrangians are displaceable then their Lagrangian Floer homology vanishes.

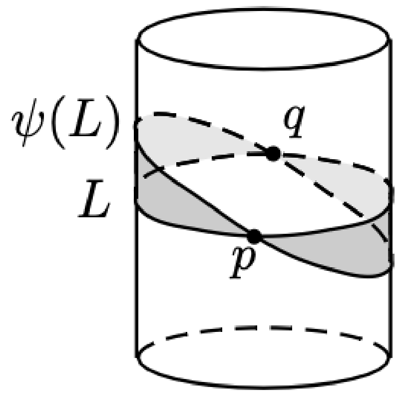

Figure 7: Chosen Hamiltonian perturbation of $L = C_z$ in the proof of Theorem 2.9

the two strips from q to p bound the same area and thus have equal energy, which implies that $\partial(p) = 0$. Thus the complex takes the form

$$\Lambda \cdot p \underset{0}{\overset{0}{\rightleftarrows}} \Lambda \cdot q$$

and $HF(C_z, C_z) = HF(C_z, \psi(C_z)) = \Lambda \cdot p \oplus \Lambda \cdot q \neq 0$. □

Note that, following the discussion about compactness in subsection 2.2.3, this result might not even make sense if one of the Lagrangians is a contractible curve: disk bubbling cannot be excluded, the differential does not square to zero and thus the Lagrangian Floer homology is not defined.

3 Discrete Morse Theory

The concepts we explored in the previous sections deal with smooth real-valued functions $f : M \to \mathbb{R}$ defined over smooth manifolds. In 1997 Mladen Bestvina and Noel Brady, inspired by the classical approach of John Milnor [22], developed a discrete counterpart as useful as the one we have already seen, with applications in geometric group theory. After the proof of their main theorem, we will be able to answer the following question, posed in [19, Problem 7.4]. Recall that a VF group, namely a virtually finite group, is a group with a finite-index subgroup which is of finite type (see Definition 3.17).

Question 3.1. *Is there a group of type VF with infinitely many conjugacy classes of finite subgroups?*

3.1 Discrete Morse functions

From now on, the domain of a discrete Morse function is no more a smooth manifold. Alternatively, we will consider a specific family of CW-complexes.

Definition 3.2. An *affine CW-complex* X is a CW-complex equipped with an integer $n \in \mathbb{Z}$ and a pair (C_e, χ_e) for each cell e of X, where

1. C_e is a convex polyedral cell embedded in $\mathbb{R}^n$;
2. $\chi_e : C_e \to e$ is an embedding, called the characteristic function associated to e, with the property that the restriction to each face of C_e is again a characteristic function, up to precomposition with an affine homeomorphism.

Definition 3.3. Let X be an affine CW-complex. A function $f : X \to \mathbb{R}$ is a *discrete Morse function* provided

1. $f\chi_e$ is the restriction of an affine map;
2. if f is constant on a cell e, then $\dim(e) = 0$;
3. the image of the 0-skeleton of X is a discrete set.

We observe that, without loss of generality, we can always picture $f\chi_e$ as a height function. As a matter of fact, we only need to correctly re-position C_e in $\mathbb{R}^n$ i.e. precomposing $f\chi_e$ with an affine homeomorphism.

Lemma 3.4. *Let $f : X \to \mathbb{R}$ be a discrete Morse function. Then every $f\chi_e$ can be seen as a height function on C_e.*

Proof. By the first property listed in Definition 3.3, the map $f\chi_e$ is the restriction of an affine function on $\mathbb{R}^n$. Therefore, we can think of it as an affine real-valued map from $\mathbb{R}^n$. Recall that an affine map is a composition of a translation and a linear map. Without loss of generality we can think of $f\chi_e$ as a vector $x \in \mathbb{R}^n$ up to precompose it with an appropriate translation.

If we suppose $f\chi_e$ not to be the zero map (i.e. the cell e has dimension at least 1) we conclude the proof by recalling that the group $\mathrm{GL}_n(\mathbb{R})$ acts transitively on $\mathbb{R}^n \setminus \{0\}$. In other words, we can think of x as the last element in the canonical basis of $\mathbb{R}^m$. □

Hence, the second property stated in Definition 3.3 can be shortly rephrased by saying that no discrete Morse function has horizontal cells. We might also want to highlight the role that the 0-skeleton of X plays in this setup. Shortly, it can be thought as the set of critical points for discrete Morse functions. The following appears in [22, Theorem 3.1] and it should be compared with its discrete analogue (see Proposition 3.7 below).

Theorem 3.5. *Let $f : M \to \mathbb{R}$ be a Morse function on a smooth manifold. Suppose that f is proper and $f^{-1}([a,b])$ contains no critical points of f. Then, the fiber $f^{-1}(b)$ is diffeomorphic to $f^{-1}(a)$.*

As for the classical Morse function theory, we have the following definition.

Definition 3.6. Given a discrete Morse function $f : X \to \mathbb{R}$ and an interval I of $\mathbb{R}$, the notation X_I stands for the set $f^{-1}(I)$, while X_t will denote the fiber $f^{-1}(t)$.

Proposition 3.7. *Let $f : X \to \mathbb{R}$ be a discrete Morse function. Suppose the interval $(a,b]$ is disjoint from $f(X^{(0)})$. Then, X_I deformation retracts to X_a.*

Proof. We begin by defining A_{-1} to be X_a and $A_i = X_a \cup (X_I \cap X^{(i)})$ for each $i \in \mathbb{N}_0$. The idea is to show that each A_i deformation retracts to A_{i-1}. Note that the set-difference $A_i \setminus A_{i-1}$ consists of the subsets of i-cells in X with values via f in I. Suppose we are given with an i-cell of X, identified with C_e via χ_e. In Lemma 3.4, we proved that C_e can be arranged such that f is a height function on e. Therefore, $e \cap X_I$ is a convex polytope whose bottom face A maps to a and the top one B maps to b. Clearly, the clousure of $\partial(e \cap X_I) \setminus B$ is homeomorphic to the disk $\mathbb{D}^{\dim(e)}$ and deformation retracts to A. This deformetion retraction extends to $e \cap X_I$ (and then recursively on the whole A_i) since every CW-complex has the HEP property. □

Question 3.8. *What happen if X_I hits the* 0*-skeleton of X?*

If we look at the literature in the classical Morse theory setup we can find the following result ([22, Theorem 3.2]).

Theorem 3.9. *Let $f : M \to \mathbb{R}$ be a Morse function on a smooth manifold. Suppose that f is proper and p is the only critical point in $f^{-1}([a, b])$. If λ is the index of p, then $f^{-1}(b)$ has the same homotopy type of $f^{-1}(a)$ with a λ-cell attached.*

A reformulation of this result in our discrete setup is possible: we only need to look for a discrete candidate that can replace λ-cells.

Definition 3.10. Let (X, v) be an affine CW-complex with a vertex v. We say that a cell e containing v as a vertex attains its minimum at v if $f|_e$ attains its minimum at v. The *descending link* $Lk_{\downarrow}(X, v)$ associated to the pair (X, v) is the link of v in the union of all descenting cells. The *ascending link* $Lk_{\uparrow}(X, v)$ can be defined similarly by replacing minimum with maximum.

Perhaps, some examples can help to have a better understanding of how a descending/ascending link looks like.

Example 3.1. Consider the following affine CW-complexes in Figure 3.1, where the discrete Morse function is just the height function. The ascending and descending links are pictures below.

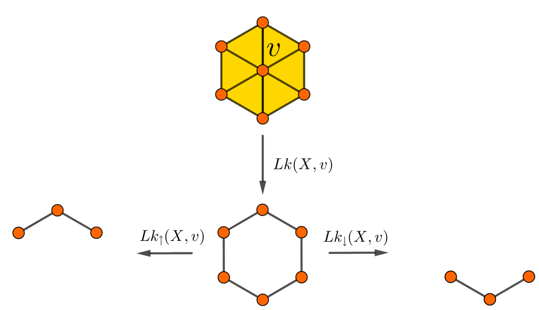

Figure 8: An example in which the descending and ascending link of X are isomorphic by a central symmetry.

Proposition 3.11. *Suppose $f : X \to \mathbb{R}$ is a discrete Morse function such* *that $X_{[a,b]}$ contains one vertex v in X_b. Then $X_{[a,b]}$ is homotopy equivalent to X_a with the cone of $Lk_{\downarrow}(X,v)$ attached.*

Proof. For each descending cell e containing v, consider the cone of $Lk_{\downarrow}(e,v)$ i.e. the set $S_e := e \cap X_{[a,b]}$ where v is the cone point and X_a contains the base. The idea is to use the same idea as Proposition 3.7 and define a chain

$$A_{-1} \subset A_0 \subset \ldots A_n = X_{[a,b]}$$

where each $A_i + 1$ deformation retracts to A_i. If A_{-1} is the union between X_a and $\bigcup S_e$, where the last union is taken among all descending cells e, and A_i is defined recursively as $A_{-1} \cup (X_{[a,b]}) \cap X^{(i)}$, the claim follows as in Proposition 3.7. The main difference lies in the definition of A_{-1}: it is X_a together with the cone on $Lk_{\downarrow}(X,v)$ attached. □

3.2 Finiteness properties of groups

This subsection is exclusively devoted to explain the finiteness relations of groups appearing in [5, Main Theorem]. Let us start defining the weakest among these properties.

Definition 3.12. A group H is of *type FP_n* if there exists an exact sequence of finitely generate and projective $\mathbb{Z}H$-modules $(P_i)_{i=0}^n$ as follows:

$$P_n \to P_{n-1} \to \ldots \to P_0 \to R \to 0, \tag{9}$$

where $\mathbb{Z}$ is considered with the structure of $\mathbb{Z}H$-module given by the trivial action of H on it. In this case, we write $H \in FP_n$.

Definition 3.13. A group H is of *type FP* if there exists an exact sequence of finitely generate and projective $\mathbb{Z}H$-modules $(P_i)_{i=0}^n$ as follows

$$0 \to P_n \to P_{n-1} \to \ldots \to P_0 \to R \to 0, \tag{10}$$

where $\mathbb{Z}$ is considered again as a $\mathbb{Z}H$-module. In this case, we write $H \in FP$.

Clearly, a group of type FP is of type FP_n for every $n \in \mathbb{N}$.

Definition 3.14. A group is said to be of *type FH_n* if it acts cellularly on a CW-complex X such that

1. the action is free, properly discontinuous and cocompact;

2. the reduced group homology $\tilde{H}_i(X, R) = 0$ for all $i \leq n-1$.

In this case, we write $H \in FH_n$.

Definition 3.15. A group is said to be of *type* FH if it acts cellularly, freely, properly and cocompact on an acyclic CW-complex X. In this case, we write $H \in FH_n$.

Similarly as before, each FH groups is indeed an FH_n for each $n \in \mathbb{N}$.

Definition 3.16. A group H is said of *type* F_n if it acts cellularly, freely, properly discontinuously and cocompactly on an n-connected CW-complex. In this case, we write $H \in F_n$.

Definition 3.17. A group H is said to be a of *type* F if there exists a finite $K(H, 1)$ CW-complex.

Again, a group of type F is always a group of type F_n for every $n \in \mathbb{N}$. The relations between these properties are sketched below and proved in the following proposition:

$$F_n \Rightarrow FH_n \Rightarrow FP_n$$
$$F \Rightarrow FH \Rightarrow FP.$$

Proposition 3.18. *Any group H of type F_n is of type FH_n. Moreover, a group of type FH_n acting on CW-complex X is also a group of type FP_n. The same holds for groups of type F, FH and FP.*

Proof. The first statement comes straightforward from [16, Theorem 4.32], since any n-connected space has vanishing reduced homology groups. In order to show that any FH_n group is a FP_n, we only need to consider every P_i as the ith element of the reduced chain complex of X. They all are free $\mathbb{Z}H$-modules. In particular, they are all projective. From Definition 3.14 we obtain that the chain is exact until the $(n-1)$-th term. We only need to show that these $\mathbb{Z}H$-modules are all finitely generated. But we notice that the action of H on X is cocompact i.e. the quotient space X/H is compact and therefore it has finitely many cells. Hence, the reduced chain complex of X is finitely generated as a $\mathbb{Z}H$-module. The respective impication for groups of type F, FH and FP can be proved in the same way. □

3.3 The Bestvina-Brady theorem

Let us give the first definitions we are goingto need for the statement of Bestvina-Brady theorem [5, Main Theorem].

Definition 3.19. A simplicial complex L is a *flag* if there is no larger simplicial complex with the same 1-skeleton.

Definition 3.20. The *right angled Arting group* associated to a finite flag L is the group G_L generated by elements $\{g_1, \dots, g_N\}$ in bijective correspondence with the vertex set $L^{(0)} = \{v_1, \dots, v_N\}$ of L, modulo the following relations

$$[g_i, g_j] = 1 \text{ for all edges } \{v_i, v_j\} \text{ in } L^{(1)}. \tag{11}$$

Next, consider the epimorphism $\phi : G_L \to \mathbb{Z}^N$ that takes each generator g_i to different standard basis elements in $\mathbb{Z}^N$. We denote by ϕ_L the map obtained by postcomposing ϕ with the epimorphism $\mathbb{Z}^N \to \mathbb{Z}$ sending every $(x_1, \dots, x_N)$ to the sum $\sum_{i=1}^N x_i$. Throughout this section, we will always denote by H_L the kernel of the map ϕ_L.

We are now able to state the main theorem. The idea is to build a group with finiteness properties related to the homotopy type of the simplex it arises from. This machinery gives a particularly useful way that provides us with examples of groups with a chosen finiteness type. The techniques Bestvina and Brady adopted rely on the definition of a ϕ_L-equivariant discrete Morse function.

Theorem 3.21. *Let L be a non-empty finite flag. If $n \geq 0$, then*

1. *H_L is of type FP_{n+1} if and only if L is homologically n-connected;*
2. *H is of type FP if and only if L is acyclic;*
3. *H is finitely presented if and only if L is simply connected.*

The following corollary is an easy consequence.

Corollary 3.22. *Let L be a non-empty finite flag. If $n \geq 1$, then*

1. *H_L is of type F_{n+1} if and only if L is n-connected;*
2. *H_L is of type F if and only if L is contractible.*

Proof. Suppose H_L is of type F_{n+1}. In particular, if it is finitely presented (or, equivalently, it is of type F_2) and of type FP_{n+1}. Therefore, L must be simply connected and homologically n-connected. By the Hurewicz theorems we conclude that it is also n-connected. Conversely, if L is n-connected then H_L is of both types F_2 and FP_{n+1}. From [7, Chapter VIII, Section 7] we can conclude is it also of type F_{n+1}. Moreover, if H_L is of type F then L must be weakly contractible. Since it is a CW-complex, it is also contractible. Conversely, if L is contractible then H_L is of type F_n for every $n \in \mathbb{N}$. Note that this does not imply automatically that H_L is of type F. Later in the dissertation, we will see that H_L acts on a contractible finitely dimensional CW-complex. This is enough to conclude the claim (see e.g.[13, Proposition 7.2.13] together with [7, Chapter VIII, Section 6] explaining Wall's finiteness obstruction). □

Example 3.2. Let us consider L to be the $(n-1)$-sphere given by the canonical triangulation as the join of n pairs of points. From Bestvina-Brady theorem we obtain that

1. there are groups of type F_{n-1} but not of type F_n for every $n \geq 1$;
2. there are groups of type FP_{n-1} but not of type FP_n for every $n \geq 1$.

Moreover, an acyclic non simply-connected finite flag complex provides us with a group H_L that is FP but not finitely presented.

Definition 3.23. A *piecewise euclidean cubic complex* (or briefly PE cubical complex) is a CW-complex built by gluing faces of a finite family of disjoint regular cubes via isometries.

Next, we are going to define a an affine CW-complex X that is build from PE cubical complexes. This will be the domain of a discrete Morse function. More specifically, we want it to be a universal cover with a well-behaved metric. Such classes of metrics are know as CAT(0) metrics. The following definitions has been taken from [6, Chapter II,1]. Recall that a geodesic segment joining two points p, q is the image of a path of length $d(p, q)$ joining p to q.

Definition 3.24. A *geodesic triangle* Δ in a metric space is a subspace defined by three points $\{p, q, r\} \subset X$ and a choice of geodesics segments (namely $[p, q], [q, r]$ and $[r, p]$) connecting each corresponding pair. A *comparison triangle* $\overline{\Delta}$ for Δ in $\mathbb{R}^2$ is a triangle with vertices $\overline{p}, \overline{q}, \overline{r}$ and edges $[\overline{p}, \overline{q}], [\overline{q}, \overline{r}]$ and

$[\overline{r}, \overline{p}]$ of the same length as the geodesics segments in Δ. A point $\overline{x} \in [\overline{p}, \overline{q}]$ is a *comparison point* for $x \in [p, q]$ if $d(q, x) = d(\overline{q}, \overline{x})$.

Definition 3.25. A metric space X is said to satisfy the *CAT(0) inequality* if the following holds:

1. X is a complete metric space;
2. for each pair of points there exists a geodesic connecting them;
3. for each pair of points x, y in a geodesic triangle Δ and each pair of comparison points $\overline{x}, \overline{y}$ in a comparison triangle $\overline{\Delta}$ we have

$$d(x, y) \leq d(\overline{x}, \overline{y}).$$

A similar notion can be given by looking at the unique complete simply-connected Riemannian surface M with curvature $k \in \mathbb{Z}$ and by replacing comparison triangles in $\mathbb{R}^2$ with comparison triangles in M with perimeter less that $2\,\mathrm{diam}(M)$. These metric spaces, usually called CAT(k) spaces, have contractible balls of radius less than $\mathrm{diam}(M)$ ([6, Proposition 1.4]). For $k = 0$ we deduce that each CAT(0) space is indeed contractible.

Definition 3.26. A PE cubical complex is said to be *nonpositively curved* if its universal cover is a CAT(0) metric space.

In particular, the universal cover of a nonpositively curved PE cubical complex is contractible. The following is a criterion we will use to check that PE cubical complexes are nonpositively curved. Note that the set of unit tangent vectors at a vertex v, pointing into the cubical complex, can be considered as a union of simplicial complexes.

Proposition 3.27. *If the set of unit tangent vectors at any vertex v pointing into a PE cubical complex X is a flag, then X is nonpositively curved.*

Proof. Omitted. See [14, pag.120] for details. □

We are now ready to describe the Bestvina-Brady construction. Our initial datum is a finite flag complex L.

Let $N \in \mathbb{N}$ be the number of vertices of L. By mapping each vertex $v_i \in L^{(0)}$ to the endpoint of the canonical basis vector e_i of $\mathbb{R}^N$, and each n-simplex $\{v_{i_0}, \ldots, v_{i_n}\}$ to the convex hull in $\mathbb{R}^N$ of $\{e_{i_0}, \ldots, e_{i_n}\}$, we have defined what is usually called the geometric realization of L. We associate to every simplex $\sigma = \{v_{i_0}, \dots, v_{i_n}\}$ the regular $(n+1)$-cube based at origin with edges $\{e_{i_0}, \dots, e_{i_n}\}$ and denote this cube by $\square_\sigma$. The image of $\bigcup_\sigma \square_\sigma$ under the universal covering map of the torus $q : \mathbb{R}^N \to \mathbb{R}^N/\mathbb{Z}^N$ is a CW-complex with just one vertex. We denote it by Q_L. Intuitively, it can be viewed as a CW–complex as follows. Its 0–skeleton consists of a single point. The 1–skeleton consists of a wedge of circles in one-to-one correspondence with the generators of G_L. The 2-skeleton is obtained by attaching a 2–torus by $g_1 g_2 g_1^{-1} g_2^{-1}$ for each edge $\{g_1, g_2\} \in L$. The 3-skeleton is obtained by attaching a 3–torus for each triangle in L, and so on. Let us give a formal description of it.

Remark 3.28. *The complex Q_L is finite. Let $T = \prod_{\omega \in L^{(0)}} S^1_\omega$ be the product of S^1 for each vertex of L, seen as a CW-complex with the usual construction with a 0-cell and a 1-cell. Note that T is a torus with the product CW-complex structure. Clearly, each simplex of L identifies a subtorus of T. Then Q_L is the union of all such subtori.*

Definition 3.29. Given a k-dimensional simplex of L, the pre-image in X of the associated torus consists of pairwise disjoint copies of $\mathbb{R}^{k+1}$. These are called *sheets*.

In particular, this description gives X a structure of PE cubical complex.

Lemma 3.30. *The complex Q_L is nonpositively curved.*

Proof. We would like to show that the set of unit tangent vectors based at v and pointing into X is a flag. We will denote it by U_v. If we restrict it to a sheet of X, it gives a canonical triangulation of S^k, inductively defined as the join $S^k \star S^0$. Furthermore, we can associate each vertex in U_v to a vertex of L and this correspondence gives a simplicial map

$$\psi : U_v \to L.$$

The pre-image of each n-simplex in $\mathrm{Star}_L(\psi(v'))$ consists of n-simplices and their faces. Since L is a flag, also U_v must be a flag. □

We can conclude that the universals cover of Q_L is contractible. Moreover, if $\pi_1(Q_L) = G_L$ the PE cubical complex Q_L is a finite $K(G_L, 1)$ space.

Proposition 3.31. *The fundamental group of Q_L is isomorphic to G_L.*

Proof. Recall that Q_L is a CW-complex with just one vertex. The fundamental group for such (oriented) CW-complexes can be easily computed by looking exclusively at its 2-skeleton. As a matter of fact, the generators correspond to the 1-cells of Q_L i.e. the vertices of L, while the relations depend only on the 2-cells and their boundaries. This gives exactly the relations in (11). □

Theorem 3.32. *Let L be a finite flag complex. Then, there exists a map $l : Q_L \to S^1$ with the following properties:*

1. *l induces an epimorphism $\phi_L : G_L \to \mathbb{Z}$ on fundamental groups;*
2. *the lift of l to the universal cover is a ϕ_L-equivariant[10] discrete Morse function $f : X \to \mathbb{R}$;*

Proof. The linear map

$$\begin{aligned} h_N : \mathbb{R}^N &\longrightarrow \mathbb{R} \\ (x_1, \cdots, x_N) &\mapsto x_1 + \cdots + x_N \end{aligned}$$

descends to a continuous map $\mathbb{R}^N/\mathbb{Z}^N \to S^1$ since the S^1-valued map is constant on the equivalence classes. We claim we can choose l to be the restriction of it on Q_L. Let us denote by $\phi_L : G_L \to \mathbb{Z}$ the map induced by l on the fundamental groups. As l sends each canonical basis vector of $\mathbb{R}^N$ to the unit interval in $\mathbb{R}$, the homomorphism ϕ_L maps each generator of $\pi(Q_L) = G_L$ to $1 \in \mathbb{Z}$. The lifting of l to the universal covers gives a map $f : X \to \mathbb{R}$, that is ϕ_L-equivariant. We are left to show that the map is a discrete Morse function. Since l sends each edge of Q_L homeomorphically onto S^1, we deduce that f is non–constant on edges of X, and hence f is also non–constant on higher dimensional cells. We see that, if e is an m-dimensional cell of X and τ is an integer translation, we can write $f \circ \chi_e = \tau \circ h_m$ up to precomposing the attaching map χ_e by an isometry of the regular m-cube $\square_m$. Here, the map h_m is defined in the same way we defined h_N, but on $\mathbb{R}^m$. □

Remark 3.33. *The kernel H_L acts on the level sets X_t since f is ϕ_L-equivariant. In particular, it acts properly, cellularly, freely and cocompactly.*

[10]With respect to the G_L-action ox X by deck transformations and the $\mathbb{Z}$-action by translations on $\mathbb{R}$.

To see this, we only need to show that it acts cocompactly: the other properties come straightforward from the G_L-action on X that satify them. Firstly, we notice that, for $x \in X$ given, there exists $g_x \in G$ such that $\phi(g_x) = t - f(x)$ (ϕ is an epimorphism). Hence, $g_x x$ is in the level set X_t and, if there is another $g'_x \in G_L$ such that $\phi(g'_x) = t - f(x)$, then $g_x^{-1} g'_x \in H_L$. The map

$$\begin{aligned} X/G_L &\to X_t/H_L \\ Gx &\mapsto (g_x \cdot x)H_L \end{aligned}$$

is therefore well-defined and X_t/H_L is compact since X/G_L is.

Proposition 3.34. *Let $f : X \to \mathbb{R}$ be the above-defined discrete Morse function. Then, all the ascending and descending links of X are isomorphic to L.*

Proof. The local picture of X at a vertex can be embedded in $\mathbb{R}^N$, where the vertex corresponds to the origin. Hence, the discrete Morse function f is given locally at this vertex as the linear map h_N defined above in Theorem 3.32. It is easy to see that the ascending and descending links in this case are isomorphic to L (Figure 9). □

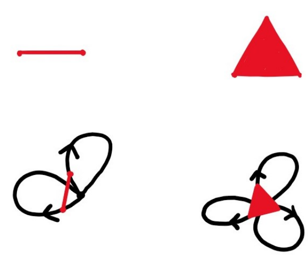

Figure 9: The flags (above, in red), together with the respective 1-skeletons of Q_L (black) and the .

We are finally able to show how Proposition 3.7 and Proposition 3.11 turn out to be useful in the proof of main theorem.

Proposition 3.35. *Let $f : X \to \mathbb{R}$ be a discrete Morse function and $I' \subset I \subset \mathbb{R}$ connected sets.*

1. *If each descending and ascending link of X is n-connected, then the inclusion $X_{I'} \hookrightarrow X_I$ induces isomorphisms on $\tilde{H}_i$ for $i \leq n$ and an epimorphism on $\tilde{H}_{n+1}$;*

2. *if each descending and ascending link of X is simply connected, then the inclusion $X_{I'} \hookrightarrow X_I$ induces an isomorphisms on π_1;*

3. *if each descending and ascending link of X is connected, then the inclusion $X_{I'} \hookrightarrow X_I$ induces an epimorphism on π_1.*

Proof. Recall that $f(X^{(0)})$ is discrete. Therefore, we can proceed by induction. We only need Proposition 3.7 and Proposition 3.11, together with the Mayer-Vietoris and Seifert-Van Kampen Theorems with respect to $X_{I'}$ and the cones on the descending/ascending links of the vertices $v \in f^{-1}(I \smallsetminus I')$. □

Corollary 3.36. *Let $f : X \to \mathbb{R}$ be as above.*

1. *If all descending and ascending links are n-homologically connected, then $H_L \in FP_{n+1}$;*

2. *If all descending and ascending links acyclic, then $H_L \in FP$:*

3. *If all descending and ascending links are simply-connected, then H_L is finitely presented.*

Proof. 1. By the previous Proposition, we have that the inclusion $X_{(-\infty,t]} \hookrightarrow X_{(-\infty,s]}$ induces isomorphisms of $\tilde{H}_i$ for $i \leq n$ and and epimorphism for $i = n+1$, if $t < s$ holds. We also know that X can be written as $\bigcup_{t \in \mathbb{Z}} X_{(-\infty,t]}$ and in particular it is contractible. We deduce that $\tilde{H}(X_{(-\infty,t]})$ vanish for $i \leq n$ and for all t. To see this, we use the followig chain of equalities for $i \leq n$:

$$H_i(X_{(-\infty,t]}) = \varinjlim H_i(X_{(-\infty,t]}) \cong H_i(\varinjlim X_{(-\infty,t]}) = H_i(X) = 0,$$

where the first equality holds since all $H_i(X_{(-\infty,t]})$ are isomorphic to each other and the second one follows for [16, Proposition 3.33]. Similarly, $\tilde{H}(X_{[t,\infty)})$ vanish for $i \leq n$ and for all t. It follows from the Mayer-Vietoris sequence that $X_t = X_{(-\infty,t]} \cap X_{[t,\infty)}$ is homologically n-connected. To prove that H_L is of type FP_{n+1}, it is enough to look at the exact cellular chain complex

$$C_{n+1}(X_t) \to C_n(X_t) \to \ldots \to C_1(X_t) \to C_0(X_t) \to \mathbb{Z} \to 0.$$

Since H_L acts on X_t freely and cocompactly, we deduce that each group $C_i(X_t)$ is finitely generated as a $\mathbb{Z}H_L$-module.

2. If all descending and ascending links are acyclic, we can proceed as above to show that X_t is acyclic for any t. Observe that X is finite dimensional. Hence, in the cellular chain complex above, $C_i(X_t)$ are non-zero for only finitely many i's. Furthermore, since X_t is in particular acyclic, we deduce that this cellular chain is exact and hence we obtain a finite resolution of $\mathbb{Z}$ by finitely generated $\mathbb{Z}H_L$-modules.

3. If all ascending and descending links are simply-connected, again by the previous Proposition, the inclusion $X_t \hookrightarrow X$ induces an isomorphism on π_1. So $\pi_1(X_t) \cong \pi_1(X) \cong 1$. Also, by part (1) of this corollary, X_t is homologically 1-connected and so X_t is connected. Thus, X_t is simply-connected. Now observe that the orbit space of X_t can be made into a $K(H_L, 1)$ space by adding cells of dimension 3 and higher to X_t in order to turn it into an acyclic space. Hence, we can conclude that H_L is of type F_2 since X_t/H_L is then compact and has finitely many 1-cells.

□

Proof of (⇐) in Theorem 3.21. It is a straightforward application of Corollary 3.36 and Proposition 3.34. □

We are left to prove the forwarding implications of Theorem 3.21. They all rely on the following theorem. It can be found in [5, Corollary 7.2].

Proposition 3.37. *For each $i \in \mathbb{N}_0$ there exists an isomorphism of $\mathbb{Z}H_L$-modules*

$$\tilde{H}_i(X_t) \cong \oplus_{v \notin X_t} \tilde{H}_i(L). \tag{12}$$

Proof (⇒) of 1) and 2) in Theorem 3.21. 1. Let $H_L \in \mathrm{FP}_{n+1}$. Set m to be $\min\{i : \widetilde{H}_i(L) \neq 0\}$. Let us assume that $m \leq n$ and arrive at a contradiction. By (12) for $i = m$, we have that $\widetilde{H}_m(X_t) \cong \oplus_{v \notin X_t} \widetilde{H}_m(L)$. Since f is onto $\mathbb{Z}$, we must have infinitely many vertices not in X_t. So, $\widetilde{H}_m(X_t)$ is not finitely generated as a $\mathbb{Z}H_L$-module. Consider the cellular chain complex

$$C_m(X_t) \to C_{m-1}(X_t) \to \ldots \to C_1(X_t) \to C_0(X_t) \to \mathbb{Z} \to 0.$$

Since H_L acts cocompactly on X_t, we deduce that each $C_i(X_t)$ is a finitely generated free $\mathbb{Z}H$-module. Furthermore, from (12), we have

$\widetilde{H}_i(X_t) \cong \oplus_{v \notin X_t} \widetilde{H}_i(L) \cong 0$ for all $0 \leq i \leq m-1$, and hence the above chain is a partial projective resolution of finite type of length m. By [7, Theorem 4.3] we have that $K = \ker\{C_m(X_t) \to C_{m-1}(X_t)\}$ is finitely generated (to apply the theorem, we use the fact that $H_L \in \mathrm{FP}_{n+1}$). But this implies that $\widetilde{H}_m(X_t)$ is a quotient of K and hence must be finitely generated, giving us the desired contradiction. Thus, L is homologically n-connected.

2. Suppose $H_L \in \mathrm{FP}$. Thus, there exists a finite resolution

$$0 \to P_n \to \ldots \to P_1 \to P_0 \to \mathbb{Z} \to 0$$

by finitely generated projective $\mathbb{Z}H_L$-modules. Hence, for all m, there exists a partial resolution of $\mathbb{Z}$ of length m by finitely generated projective $\mathbb{Z}H$-modules. So, for all m, we have $H \in \mathrm{FP}_{m+1}$ and therefore by part (1) of the Bestvina–Brady theorem, L is homologically m-connected for all m. Thus, L is acyclic.

□

The proof of the third part of the Bestvina–Brady Theorem is more technical and we will skip some details.

Sketch of the proof for (⇒) *of 3) in Theorem 3.21.* If L is not connected, then L is not homologically 1-connected and by part 1 of the Bestvina–Brady Theorem we have that H cannot be of type FP_2 and so it cannot be of type F_2. If we suppose that L is connected but not simply connected, than $\pi_1(X_t)$ is generated by H_L-translates of finitely many loops ([5, Proposition 3.9]). Since all this loops are null-homotopic in X, we can shrink X to $X_{[t-T,t+T]}$ for some small T and suppose they are null-homotopic in it. Moreover, H acts by translation and H-orbits of these loops are then null-homotopic in $X_{[t-T,t+T]}$. The inclusion $X_t \hookrightarrow X_{[t-T,t+T]}$ induces an epimorphism on the fundamental groups and therefore $X_{[t-T,t+T]}$ is simply connected. This cannot happen if H is finitely presented. □

3.4 About some VF groups

Our goal is to show that the following result (due to Brown [7, Lemma 13.2]) cannot be generalized to finite subgroups.

Proposition 3.38. *Let G be a group with a finite index subgroup of type F. Then, it can contain only finitely many conjugacy classes of subgroups of prime power order.*

In particular, we would like to exhibit a group G together with a finite index subgroup H of type F, such that G has a finite subgroup having infinitely many conjugates. The idea is to use Bestvina-Brady construction and define our group G to be the semidirect product $H_L \rtimes Q$, where Q is a finite group of automorphism of L.

We observe that such Q induces a well-defined action on G_L. Let $g_v \in G_L$ be a generator associated to a vertex $v \in L$. We define $q \cdot g_v$ to be $g_{q(v)}$. Cleary, if $[g_v, g_w] = 1$ for some v and w L vertices, then $[g_{q(v)}, g_{q(w)}] = 1$ since Q acts by automorphisms on L and, in particular, it sends simplices to simplices.

Let us fix a point $x_0 \in X$. Then, there exists an Q-action on X: if an n-cell in X is indexed by $(g \cdot x_0, (v_1, \dots, v_{n+1}))$, and element $q \in Q$ send it to $((q \cdot g) \cdot x_0, (q \cdot v_1, \dots, q \cdot v_{n+1}))$. Note that x_0 is the only fixed point of Q on X. As G_L acts by deck transformations on X, we can extend the Q-action on X by a $G_L \rtimes Q$ action on X.

Definition 3.39. An automorphism of L is said to be *admissible* if the setwise and pointwise stabilizer of every simplex coincide. A subgroup $Q \leq \mathrm{Aut}(L)$ is admissible if every element in it is admissible.

Proposition 3.40. *Let L be a non-empty finite flag complex together with a finite admissible group Q acting on L by automorphisms. Suppose that Q do not fix any element in L. Then, there are infinitely many conjugacy classes of Q in $H_L \rtimes Q$.*

Proof. Let us fix x_0 to be the only fixed point of X by Q. Hence, the stabilizer of $g \cdot x_0$ in $G_L \rtimes Q$ is exactly the conjugate Q^g. We also notice that, as $f : X \to \mathbb{R}$ is G_L-equivariant, it is also $G_L \rtimes Q$-equivariant.

Let us denote by X^P the fixed point of a finite subgroup P of $G_L \rtimes Q$. We claim that, since $G_L \rtimes Q$ acts by isometries, the subspace X^P of X is CAT(0).

To see this, we first notice that G_L acts by deck transformation on X and so we only need to show that Q acts by isometries on X. But this is trivial since it acts by automorphisms on L, preserving the distances. Hence, it preserves the distances on elements of the form $g \cdot x_0$. Now, since X is CAT(0), there exists a unique geodesic γ between two arbitrary points $x_1, x_2 \in X^P$. If $g \in P$, then the endpoints of γ are the same as the endpoints of $g \cdot \gamma$. Since g is an isometry, we have that $g \cdot \gamma$ is a geodesic and that $g \cdot \gamma$ is exactly γ pointwise. To prove that X^P is complete, we only need to show that it is closed in X. Since X is complete, the claim will follow. As a matter of fact, it is the finite union of the sets $\{x \in X \mid p \cdot x = x\}$, indexed by $p \in P$.

The admissibility of Q on L implies that the action is cellular, and hence X^P is a subcomplex. In particular, the subcomplex X^P is contractible. Since Q do not fix any element, we conclude that X^Q cannot have 1-cells and so it consists of a single point $\{g \cdot x_0\}$, with image by f the point $\phi_L(g)$. If $P = Q^h$ for some $h \in H_L \rtimes Q$, then $f(X^P)$ consists of the point $\phi_L(hg) = \phi_L(g)$. In other words, conjugation by an element in $H_L \rtimes Q$ preserves the height.

Let us consider $t \in G_L$ such that $\phi_L(t) = 1$. Then, all the subgroups Q^{t^i} cannot be conjugated in $H_L \rtimes Q$, for every $i \in \mathbb{Z}$. □

The following is a straightforward application of Bestvina-Brady main theorem.

Corollary 3.41. *Together with the same hypothesis as in Proposition 3.40, suppose furthermore that L is contractible. Then, the group $H_L \rtimes Q$ is of type VF and contains infinitely many conjugacy classes of the finite group Q.*

In order to answer Question 3.1, one only needs to exhibit a contractible finite flag L together with a finite automorphism with no fixed points. The admissibility can be omitted as we can always consider the barycentric subdivision L' of L: in this case the action on L' is admissible and L' has the same realization as L. It turns out that a fixed-point free action on a contractible flag complex is not easy to define. Nevertheless, an example exists. In [11] we can find an example of a 2-dimensional fixed-point free action on a finite acyclic flag. Then, the group $G \times \mathbb{Z}/2\mathbb{Z}$ acts without fixed point on the suspension of X, that is contractible.

References

[1] Arnold, Vladimir. "Sur une propriétés topologique des applications globalment canonique de la méchanique classique." *CR. Acad. Sci. Paris* 261 (1965): 3719-3722.

[2] Audin Michéle and Damian Mihai. *Morse Theory and Floer Homology*, Springer, 2013.

[3] Auroux, Denis. "A beginner's introduction to Fukaya categories." *Contact and symplectic topology*. Springer, Cham, 2014. 85-136.

[4] Ballmann, Werner. *Der satz von Lusternik und Schnirelmann*. No. 1. Mathematisches Institut der Universität, 1978.

[5] Bestvina, Mladen, and Noel Brady. *Morse theory and finiteness properties of groups.*, Inventiones mathematicae 129.3, 445-470 (1997).

[6] Bridson, Martin R., and André Haefliger. "Metric spaces of non-positive curvature" *Springer Science & Business Media*, 2013.

[7] Brown, Kenneth S. "Cohomology of groups", *Springer Science & Business Media*, 2012.

[8] Brugués Mora, Joaquim. "Morse Theory and Floer Homology", *Universitat Politècnica de Catalunya - BarcelonaTech* (2019).

[9] Da Silva, Ana Cannas. *Lectures on symplectic geometry*. Berlin: Springer, 2008.

[10] Floer, Andreas. "Morse theory for Lagrangian intersections." *Journal of differential geometry* 28.3 (1988): 513-547.

[11] Floyd, E.E., Richardson, R.W. *An action of a finite group on an n-cell without stationary points.* Bull. Am. Math. Soc. 65, 73–76 (1959).

[12] Fushida-Hardy, Shintaro. "Morse Theory", *Stanford University* (2020).

[13] Geoghegan, Ross. "Topological methods in group theory." Vol. 243. *Springer Science & Business Media*, (2007).

[14] Gromov, Mikhael. "Hyperbolic groups. Essays in group theory." *Springer*, New York, NY, 1987. 75-263.

[15] Gutierrez Abed, Mario. "Morse Homology: A Brief Introduction", *CUNY Hunter College* (2016).

[16] Hatcher, Allen. "Algebraic topology", *Cambridge University Press*, (2002).

[17] Hutchings, Michael. "Lecture notes on Morse homology (with an eye towards Floer theory and pseudoholomorphic curves)", *University of California, Berkeley* (2002).

[18] Knudson, Kevin. *Morse Theory: Smooth and Discrete*, World Scientific, 2015.

[19] Lück, Wolfgang, *The type of the classifying space for a family of subgroups*, Journal of Pure and Applied Algebra, 149, 2, 177-203 (2000)

[20] Lusternik, L., and Lev Schnirelmann. "Sur le problème de trois géodésiques fermées sur les surfaces de genre 0." *CR Acad. Sci. Paris* 189.269 (1929): 271.

[21] McDuff, Dusa, and Dietmar Salamon. *Introduction to symplectic topology.* Vol. 27. Oxford University Press, 2017.

[22] Milnor, John. *Morse Theory*, Princeton University Press, 1963.

[23] Morse, Harold Marston. "A fundamental class of geodesics on any closed surface of genus greater than one". Trans. Amer. Math. Soc. 26 (1): 25–60.

[24] Salamon, Dietmar."Lectures on Floer homology" *Symplectic geometry and topology (Park City, UT, 1997)* 7 (1999): 143-229.

[25] Schwarz, Matthias. *Morse Homology*, Birkhäuser, 1993.