Homology Theory, Morse Theory, and the Morse Homology Theorem

Alex Stephanus

June 8, 2016

Abstract

This thesis develops the Morse Homology Theorem, first starting by motivation for and developing of the notion of homology groups of a manifold. Following this, we introduce some introductory Morse Theory, specifically the process of building a CW complex which is homotopy equivalent to a given manifold. Finally, we introduce the machinery required for the proof of the Morse Homology Theorem, which relates the homology groups generated by the n-cells of a manifold to the homology groups generated by critical points of certain functions on the manifold.

Contents

1 Introduction **1**

2 Motivation for Homology Groups **2**

3 Homology **5**

4 Morse Theory: Constructing CW Complexes **13**

5 Gradient Flows, Stable and Unstable Manifolds, and Morse-Smale functions **21**

6 Corollaries to the λ-lemma **24**

7 The Morse Homology Theorem **26**

8 Acknowledgements **36**

1 Introduction

Classification of manifolds has been a central goal of topology since the field has existed, and even today it remains a key focus of the field. In the last hundred years or so algebraic topology has sought to ascribe groups to manifolds and, in doing so, apply insights from algebra to the topological problem of manifiold classification. Within the past 50 years even more disciplines have been pulled in to help understand this problem. In this thesis we build up the theory of Morse Homology, which links the algebraic structures of homology with the more geometric flow structure that particular real-valued functions induce on a manifold.

We begin by developing some of the machinery of homology theory, which seeks to understand a manifold's different-dimensional "holes" in an algebraic sense. We show some important properties of homology groups, for instance that they are a homotopy invariant of a space, and explore the relative homology groups of a space, which provide a way to ignore certain subspaces of a topological space. The other important result that we prove in this section is that, for a CW complex, the homology which arises from the cell structure is identical to the more traditional singular homology.

Following the equivalence of cellular and singular homology we turn to Morse Theory to construct a way to put a cell structure on a manifold. As it turns out, a smooth real-valued function with some nice properties is all that we need to determine a cell structure on a manifold, and in this section we flesh out the process for determining a cell structure given an appropriate function on a manifold.

Finally, we bring these two theories together to prove the Morse Homology Theorem. This theorem, which is essentially the capstone result of this thesis, states that that singular homology groups of a manifold coincide exactly with the homology groups generated by the manifold's critical points.

This thesis is meant to be largely self-contained, if the reader has a fairly good footing in basic abstract algebra, differential topology, and Riemannian geometry then everything should roll out smoothly. However, due to the need to build up multiple disciplines some proofs are omitted for the sake of brevity and readability. In general, if a result is stated without proof it is likely because the proof is both long and not terribly insightful for what we're building to.

2 Motivation for Homology Groups

Rather than just dive into homology theory, we begin with the motivating example of homotopy groups of a space. These groups are very intuitively defined and can give insight into how a space is structured. However, homotopy groups are also problematic in some ways. Some of these problems provide motivation for the development of homology theory which, althouth not as intuitively defined, is generally much better-behaved.

2.1 The Fundamental Group

One of the simplest tools for studying manifolds is the fundamental group, π_1. The fundamental group of a manifold X is the group of based loops in X, i.e.

maps $\alpha : [0,1] \to X$ such that $\alpha(0) = \alpha(1)$. subject to the relation that two loops are considered equivalent if they are homotopically identical. The group operation is concatenation of loops, specifically concatenating two paths α and β to get a path $\alpha \circ \beta$, defined as follows:

$$(\alpha \circ \beta)(l) = \begin{cases} \alpha(2l) & \text{if } 0 \leq l \leq .5 \\ \beta(2l-1) & \text{if } .5 < l \leq 1 \end{cases}$$

The set of maps combined with this concatenation operation satisfies the group axioms as follows.

1. Closure: If one concatenates two loops α and β which both begin and end at some point x_0 in X then their concatenation is also clearly a loop in X based at x_0.

2. Identity: The identity is given by the path which simply stays at x_0

3. Associativity: Given 3 paths α, β, and γ the order of composition does not matter.

4. Inversion: Given a loop α based at x_0, let α^{-1} be defined as $\alpha^{-1}(l) = \alpha(1-l)$. The homotopy given by

$$(f_t)(l) = \begin{cases} \alpha(2l) & \text{if } 0 \leq l \leq \frac{t}{2} \\ \alpha(t) & \text{if } \frac{t}{2} \leq l \leq 1 - \frac{t}{2} \\ \alpha^{-1}(2l-1) & \text{if } 1 - \frac{t}{2} \leq l \leq 1 \end{cases}$$

   is smooth with $f_0(l) = x_0$ and $f_1(l) = (\alpha \circ \alpha^{-1})(l)$, so $\alpha \circ \alpha^{-1} \simeq \beta$, where β is the identity loop.

The fundamental group is useful for a variety of reasons, but one of the most important is its ease of computation. Van Kampen's Theorem gives a relatively simple way to calculate the fundamental group of a space by splitting it up into smaller spaces. A simplified statement of the theorem is as follows:

Theorem 2.1 (van Kampen's Theorem) *Let X be a topological space, and let U_1 and U_2 be two open, path-connected subspaces of X. If $U_1 \cap U_2$ is path connected and nonempty then X is path connected and the inclusion morphisms draw the commutative pushout below:*

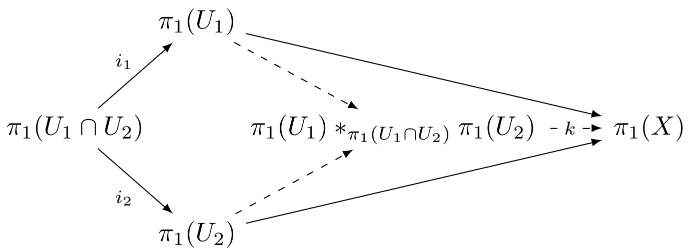

The natural morphism k is an isomorphism, that is to say that the fundamental group X is the free product of the fundamental groups of 1 *and* 2 *with amalgamation of the fundamental group of their intersection.*

The proof of van Kampen's Theorem is omitted here, as it is fairly involved and not particularly relevant to us. For a full proof see pg. 53 of [1].

Van Kampen's Theorem, as it describes spaces which are built out of smaller spaces, leads naturally to the question of whether there is a convenient way to construct spaces out of smaller spaces. As it turns out there is and these spaces, known as CW complexes, are central to the study of algebraic topology.

2.2 CW Complexes

In algebraic topology, it is often helpful to look at topological spaces which are endowed with a useful structure. One such set of spaces considered to be particularly nice are called CW complexes. A CW complex is constructed by inductively gluing together disks of increasing dimension along their boundaries. The definition of a CW complex is below, note that e^n is taken to be the open disk of dimension n. When necessary, we will refer to the boundary of this disk as $\hat{e}^n$.

1. Start with a discrete set X^0, regarding the points of this set to be 0-cells.
2. Inductively build the n-skeleton, X^n by attaching n-cells e^n_α to X^{n-1} via maps $\phi_\alpha : S^{n-1} \to X^{n-1}$.
3. This process can either terminate at some finite n, in which case $X = X^n$, or it can go on infinitely, in which case $X = \bigcup_i X^i$

CW Complexes are particularly well-behaved spaces, with this inherent inductive structure supplying many useful properties. For instance, in a CW complex X, the fundamental group is the free group with generators coming from the 1-skeleton of X and relations coming from the 2-skeleton of X.

2.3 Homotopy Groups

The fundamental group allows us to characterize manifolds in a low-dimensional manner. Anywhere a loop cannot be contracted to a point it means that there is some sort of hole in the manifold. For instance, the circle S^1 has fundamental group isomorphic to $\mathbb{Z}$, which is indicative of the area in the center of the circle which a loop cannot be contracted through. However, S^n is simply connected for $n > 1$, that is to say it has a trivial fundamental group. These higher-dimensional spheres still have holes in a similar manner to S^1, but that information is not captured by the fundamental group. This makes sense, as the fundamental group considers low-dimensional maps and homotopies. In fact, in a CW complex the fundamental group can be calculated knowing only the 2-skeleton of the complex. This lack of higher-dimensional information provided by the fundamental group leads naturally to the question of whether higher-dimensional analogs exist. The fundamental group considers maps of S^1 into spaces and homotopies of those maps, and similarly the n^{th} homotopy group is constructed in a similar manner, only one considers maps and homotopies of maps of S^n. The main issue with homotopy groups is that they cannot generally be computed as easily as the fundamental group. Van Kampen's Theorem gives a nice way to calculate the fundamental group of a space, but it does not generalize into higher dimensions.

3 Homology

Homotopy groups, while useful for examining a space's structure, have the drawback that they are not terribly easy to use and they are not always algebraically simple. For instance, the fundamental group of a space is not necessarily abelian, a property that we would like to have. So, in an effort to avoid some of the computational difficulties associated with homotopy groups, we develop homology theory. The definition is much less intuitive at first, but it has a similar goal of putting an algebraic structure on the n-dimensional holes of a topological space.

3.1 Simplices

The theory of homology is developed primarily by considering simplices, so before we develop homology theory, it is appropriate to first define a simplex. An n-dimensional simplex, denoted as Δ^n, is a generalization of a 2-d triangle into higher dimensions. n-simplices are defined by sets of points $[v_0, \ldots, v_n]$ subject to the condition that no set of k points lies in a hyperplane of dimension less than $k-1$, for $2 < k \leq n+1$. Simplices are meant to generalize the notion of a triangle into higher dimensions, so this condition is a generalization of the idea that one cannot have a triangle with three collinear points.

$$\Delta^n = \{(v_0 t_0, \ldots, v_n t_n) \in \mathbb{R}^n \mid \sum_i t_i = 1 \text{ and } t_i \geq 0 \text{ for every } i\}$$

The order of a simplex's vertices is relevant when considering boundary maps later, so we keep the order given by the subscripts of the vertices.

Additionally, we consider the boundary of Δ^n to be the union of the simplices of dimension $n-1$ formed when deleting a vertex from the original n-simplex, i.e. the simplices $[v_0, \ldots, \hat{v}_i, \ldots, v_n]$, with $\hat{v}_i$ indicating that the i^{th} vertex has been removed. These $(n-1)$-simplices are called the faces of Δ^n. One important convention to take is that in the faces of Δ^n, or rather any subsimplex Δ^k of Δ^n, the vertices of Δ^k are ordered as they are in Δ^n. Additionally, we orient the edges of Δ^n, $[v_i, v_j]$ towards the higher-ordered vertex. This orientation of edges carries over into all subsimplices of Δ^n.

3.2 Singular Homology

Singular homology stems from so-called 'singular' mappings of simplices into a space X. These maps need only be continuous, so the term singular captures the idea that a map need not necessarily resemble a simplex in its image. So, we define a singular n-simplex to be a map $\sigma : \Delta^n \to X$. To start with, we define the group $C_n(X)$ to be the free abelian group whose basis is the set of singular n-simplices in X. The elements of $C_n(X)$, n-chains, are finite sums of the form $\sum_i n_i\sigma_i$, with n_i in $\mathbb{Z}$ and σ_i singular n-simplices of X.

Homology concerns itself with how the n-simplices of X attach to the $n-1$-simplices of X, and in order to examine that, we define the following boundary map $\partial_n : C_n(X) \to C_{n-1}(X)$, below.

$$\partial_n(\sigma) = \sum_i (-1)^i \sigma \mid [v_0, \dots, \hat{v_i}, \dots, v_n]$$

Restricting σ to the set $[v_0, \dots, \hat{v_i}, \dots, v_n]$ gives us a new $(n-1)$-chain, with the $(-1)^i$ precipitating out of the orientation of the edges which are removed. An important property of this boundary map follows from the definition:

Lemma 3.1 *For any n-simplex,* $\partial_{n-1}\partial_n = 0$

Proof The proof follows from simply composing ∂_n with $\partial n - 1$

$$\partial_n(\sigma) = \sum_i (-1)^i \sigma \mid [v_0, \dots, \hat{v_i}, \dots, v_n] \text{ so, we get that}$$

$$\partial_{n-1}\partial_n(\sigma) = \sum_{j<i} (-1)^i(-1)^j \sigma \mid [v_0, \dots, \hat{v_j}, \dots, \hat{v_i}, \dots, v_n] +$$
$$\sum_{j>i} (-1)^i(-1)^{j-1} \sigma \mid [v_0, \dots, \hat{v_i}, \dots, \hat{v_j}, \dots, v_n]$$

As j ranges through all values other than i, switching j and i in the second sum cancels out the first sum, so the sum is zero when taken over all possible values of i and j.

If we look at all of these maps for a given space X

$$\dots \xrightarrow{\partial_{n+1}} C_n(X) \xrightarrow{\partial_n} C_{n-1}(X) \xrightarrow{\partial_{n-1}} \dots \xrightarrow{\partial_1} C_0(X) \xrightarrow{\partial_0} 0$$

we get what is called a chain complex on the space, a series of maps between abelian groups with the property $\partial_n\partial_{n+1} = 0$.

The composition $\partial_n\partial_{n+1}$ being identically zero is equivalent to the statement that the image of ∂_{n+1} is contained in the kernel of ∂_n. As $\mathrm{Im}(\partial_{n+1})$ is a (not necessarily proper) subgroup of $\mathrm{Ker}(\partial_n)$, the quotient $\mathrm{Ker}(\partial_n)/\mathrm{Im}(\partial_{n+1})$ exists. We define the n^{th} singular homology group via this quotient.

$$H_n(X) = \mathrm{Ker}(\partial_n)/\mathrm{Im}(\partial_{n-1})$$

We can think of $\mathrm{Ker}(\partial_n)$ as the group of cycles of $n-$simplices, and $\mathrm{Im}(\partial_{n+1})$ to be the group of boundaries (this group contains precisely the chains of n-simplices which form the boundary of some $(n+1)$-chain). In this sense, homology groups are obtained by taking cycles mod boundaries.

3.3 Homotopy Invariance of Homology Groups

One particularly important property of singular homology groups is that they are invariant under homotopy, that is to say that if two spaces X and Y are homotopy equivalent then they have isomorphic homology groups. This section contains the proof of that fact, which stems from the homomorphism of homology groups which is induced by a map between two spaces.

This homomorphism comes from exactly where one would expect it to, taking a map $f : X \to Y$ and composing each $\sigma : \Delta^n \to X$ with f. This composition of maps induces a homomorphism $f_* : H_n(X) \to H_n(Y)$ in the following manner. If we denote the composition $f\sigma$ as $f_\#(\sigma)$, then we get that each $f_\#(\sigma)$ is a map from Δ^n into Y. We can then extend this map linearly to chains in Y by defining $f_\#(\Sigma_i n_i \sigma_i) = \Sigma_i n_i f_\#(\sigma_i)$ The crucial detail of this composition is the following chain of equalities.

$$\begin{aligned} f_\#\partial(\sigma) &= f_\#(\Sigma_i(-1)^i \sigma|[v_0, \dots, \hat{v_i}, \dots, v_n]) \\ &= \Sigma_i(-1)^i f\sigma|[v_0, \dots, \hat{v_i}, \dots, v_n] = \partial f_\#(\sigma) \end{aligned}$$

Where ∂ is the same boundary operator defined previously. Because ∂ and f commute, we can construct the following commutative diagram.

$$\begin{array}{ccccccccc} \dots & \longrightarrow & C_{n+1}(X) & \xrightarrow{\partial} & C_n(X) & \xrightarrow{\partial} & C_{n-1}(X) & \longrightarrow & \dots \\ & & \downarrow f_\# & & \downarrow f_\# & & \downarrow f_\# & & \\ \dots & \longrightarrow & C_{n+1}(Y) & \xrightarrow{\partial} & C_n(Y) & \xrightarrow{\partial} & C_{n-1}(Y) & \longrightarrow & \dots \end{array}$$

These $f's$ define a chain map from the chain complex of X into that of Y. We can see that $f_\#$ takes cycles to cycles and boundaries to boundaries by observing that $\partial\alpha = 0$ gives $\partial(f_\#\alpha) = f_\#(\partial\alpha) = 0$, and also that $f_\#(\partial\beta) = \partial(f_\#\beta)$. So, $f_\#$ induces a homomorphism $f_* : H_n(X) \to H_n(Y)$. The following theorem is one of the more important results we will be using from homology theory, and it will show up throughout this thesis.

Theorem 3.2 *If two maps $f, g : X \to Y$ are homotopic, then they induce the same homomorphism $f_* = g_* : H_n(X) \to H_n(Y)$*

Proof They key to this proof is the following way to divide $\Delta^n \times I$ into simplices. Let $\Delta^n \times 0$ be given by $[v_0, \dots, v_n t]$ and $\Delta^n \times 1$ by $[w_0, \dots, w_n]$ subject to the condition that v_i and w_i have the same image under the natural projection $\Delta^n \times I \to \Delta^n$. We can then go from $[v_0, \dots, v_n]$ to $[w_0, \dots, w_n]$ by stitching them together with a sequence of n-simplices. These simplices are of the form $[v_0, \dots, v_i, w_{i+1}, \dots, w_n]$ for every i with $-1 \leq i \leq n$. Between each successive pair of these simplices, i.e. between $[v_0, \dots, v_i, w_{i+1}, \dots, w_n]$ and $[v_0, \dots, v_{i-1}, w_i, w_{i+1}, \dots, w_n]$, is an $(n+1)$-simplex. The union of all of these $(n+1)$-simplices spans the region $\Delta^n \times I$.

Following this we define a new operator, called a prism operator, using a homotopy $F : X \times I \to Y$ and a simplex $\Delta^n \to X$. Composing these two maps, we get a map $F \circ (\sigma \times I) : \Delta^n \times I \to X \times I \to Y$. Then, we define the prism operator $P : C_n(X) \to C_{n+1}(Y)$ as follows

$$P(\sigma) = \sum_I (-1)^i F \circ (\sigma \times 1)|[v_0, \ldots, v_i, w_i, \ldots, w_n]$$

These operators have the property that $\partial P = g_\# - f_\# - P\partial$ (for verification of this property see the proof of 2.10 in [1]). Using this information, we know that if $\alpha \in C_n(X)$ is a cycle then we have $g_\#(\alpha) - f_\#(\alpha) = \partial P(\alpha) + P\partial(\alpha) = \partial P(\alpha)$. The last equality comes from the fact that $\partial\alpha = 0$ (as it is a cycle). Thus, $g_\# - f_\#$ is a boundary, which we know from before means that $g_\#$ and $f_\#$ determine the same homology class. As such, $g_\#$ and $f_\#$ are equal on the homology class of α, and we are done.

This statement implies immediately that for a homotopy equivalence $f : X \to Y$ the induced homomorphisms $f* : H_n(X) \to H_n(Y)$ are in fact isomorphisms.

This is one of the most important results we will take from homology theory, and we will combine with soon-to-come results about the homology of a CW complex to do a lot of heavy lifting later.

3.4 Exact Sequences, Relative Homology, and the Mayer-Vietoris Sequence

A question that comes up naturally when studying spaces is how subspaces fit into and define the larger space. Spaces can often be constructed inductively from smaller spaces, so understanding the behavior of a larger space relative to its subspaces is typically very useful information to have. Homology is no exception to this rule, and the concept of a relative homology group attempts to understand these relationships. Relative homology groups essentially convey how the homology of a space changes when we ignore one of its subspaces. However, before we rigorize this idea we will define something called an exact sequence, which is a structure that is quite important for the understanding of relative homology groups. We have already defined a chain complex, that is a series of homomorphisms from the chain groups of a space into each other:

$$\ldots \xrightarrow{\partial_{n+2}} C_{n+1}(X) \xrightarrow{\partial_{n+1}} C_n(X) \xrightarrow{\partial_n} C_{n-1}(X) \xrightarrow{\partial_{n-1}} \ldots$$

with the additional information that $\partial_n \partial_{n+1} = 0$ for every n, and we also know that this condition is equivalent to the inclusion $\mathrm{Im}(\partial_{n+1}) \subseteq \mathrm{Ker}(\partial_n)$. The notion of an exact sequence strengthens this condition from an inclusion to an equivalence. That is to say that an exact sequence is a sequence of homomorphisms between abelian groups

$$\ldots \xrightarrow{\alpha_{n+2}} A_{n+1} \xrightarrow{\alpha_{n+1}} A_n \xrightarrow{\alpha_n} A_{n-1} \xrightarrow{\alpha_{n-1}} \ldots$$

with the additional constraint that $\text{Im}(\alpha_{n+1}) = \text{Ker}(\alpha_n)$. The inclusions necessary for this sequence to be a chain complex are of course satisfied, and because the two groups are identical it means that $\text{Ker}(\partial_n)/\text{Im}(\partial_{n+1}) = H_n(X)$ is trivial.

Here we will just list a few properties of exact sequences that follow immediately from the definition of an exact sequence. They come up often enough that it is worth listing them out.

1. $0 \to A \xrightarrow{\alpha} B$ is exact if and only if $\text{Ker}(\alpha) = 0$, so if α is injective.

2. $A \xrightarrow{\alpha} B \to 0$ is exact iff $\text{Im}(\alpha) = B$, so if α is surjective. It follows naturally then that

3. $0 \to A \xrightarrow{\alpha} B \to 0$ is exact iff α is an isomorphism.

4. $0 \to A \xrightarrow{\alpha} B \xrightarrow{\beta} C \to 0$ is exact if and only if α is injective, β is surjective, and $\text{Im}(\alpha) = \text{Ker}(\beta)$. In this case we get that C is isomorphic to $B/\text{Im}(\alpha)$. This property is not as obvious as the other three, but it will be relevant later.

One particularly important result concerning exact sequences is the following theorem.

Theorem 3.3 *If A is a nonempty closed subspace of a space X, and is also a deformation retract of a neighborhood in X, then there exists an exact sequence*

$$\dots \tilde{H}_n(A) \xrightarrow{i_*} \tilde{H}_n(X) \xrightarrow{j_*} \tilde{H}_n(X/A) \xrightarrow{\partial}$$

$$\tilde{H}_{n-1}(A) \xrightarrow{i_*} \tilde{H}_{n-1}(X) \xrightarrow{j_*} \dots \xrightarrow{j_*} \tilde{H}_0(X/A) \to 0$$

where i_* is the homomorphism induced by the inclusion $i : A \hookrightarrow X$ and j_* is induced by the quotient map $j : X \to X/A$.

It should be noted that the groups $\tilde{H}$ are the reduced homology groups of a nonempty space X. These groups are the homology groups of the chain complex

$$\dots \to C_2(X) \xrightarrow{\partial} C_1(X) \xrightarrow{\partial} C_0(X) \xrightarrow{\epsilon} \mathbb{Z} \to 0$$

with the map $\epsilon(\sum_i n_i\sigma_i) = \sum_i n_i$. $H_n(x) \approx \tilde{H}_n(X)$ for $n > 0$, and $H_0(X) \approx \tilde{H}_0(X) \oplus \mathbb{Z}$.

A proof of this theorem and a construction of the boundary map ∂ can be found in Chapter 2 of [1] (Theorem 2.13 and proof). An important bit of terminology to remember is that if a space X and subspace $A \subseteq X$ satisfy the theorem's conditions then we call (X, A) a good pair.

We have already posed the question of how to understand how the homology of X changes when you "ignore" one of its subspaces, and one natural idea is to examine the homology of the space when you quotient out the subspace. Relative homology groups do just that, and are defined in the following way so as to effectively ignore the homology of X generated by the subspace.

Given a space X and a subspace $A \subseteq X$, we define $C_n(X, A)$ to be the quotient $C_n(X)/C_n(A)$ so any chain in A is quotiented out. Chains in A have boundaries in A, so the boundary map ∂ takes $C_n(A)$ into $C_{n-1}(A)$. This means that the boundary map induces a quotient map $\partial : C_n(X, A) \to C_{n-1}(X, A)$. So, we get a similar chain complex

$$\dots \xrightarrow{\partial_{n+1}} C_n(X, A) \xrightarrow{\partial_n} C_n(X, A) \xrightarrow{\partial_n n-1} \dots$$

We know the relation $\partial_n \partial_{n+1} = 0$ holds in $C_n(X)$, so it holds for the quotient-induced operator as well. The homology groups of this chain complex are defined to be the relative homology groups $H_n(X, A)$. Elements of the group $H_n(X, A)$ are represented by relative cycles, which are n-chains whose boundaries lie completely in A.

Relative homology groups share a similar property to Theorem 3.3, in that they fit into the following exact sequence:

$$\dots H_n(A) \xrightarrow{i_*} H_n(X) \xrightarrow{j_*} H_n(X, A) \xrightarrow{\partial}$$

$$H_{n-1}(A) \xrightarrow{i_*} H_{n-1}(X) \xrightarrow{j_*} \dots \xrightarrow{j_*} H_0(X, A) \to 0$$

where i_* and j_* are the homomorphisms induced by the inclusions $A \hookrightarrow X$ and $(X, \emptyset) \hookrightarrow (X, A)$. The boundary operator $\partial : C_n(X, A) \to C_{n-1}(X, A)$ is defined by considering the boundaries of relative cycles which represent elements of $H_n(X, A)$. If a relative cycle α represents an element $[\alpha]$ then the boundary map sends $[\alpha]$ to the class of the boundary of α in $H_n(A)$. We refer to the operator ∂ in this sequence as a connecting homomorphism, a definition which will be important later. One important property of this map ∂ is that it is a natural transformation, which is to say that for a map $f : (X, A) \to (Y, B)$, the diagram below is commutative.

$$\begin{array}{ccccccccccc}
\dots & \longrightarrow & H_n(A) & \xrightarrow{i_*} & H_n(X) & \xrightarrow{j_*} & H_n(X, A) & \xrightarrow{\partial} & H_{n-1}(A) & \longrightarrow & \dots \\
 & & \downarrow f_* & & \downarrow f_* & & \downarrow f_* & & \downarrow f_* & & \\
\dots & \longrightarrow & H_n(B) & \xrightarrow{i_*} & H_n(Y) & \xrightarrow{j_*} & H_n(Y, B) & \xrightarrow{\partial} & H_{n-1}(B) & \longrightarrow & \dots
\end{array}$$

A more thorough discussion of naturality can be found in Chapter 2 of [1].

One additional result that we will use once later is the Excision Theorem. This theorem lays out conditions under which a a subset $Z \subset A \subset X$ can be deleted without affecting the relative homology groups $H_n(X, A)$. The statement of the theorem is as follows.

Theorem 3.4 *Given subspaces $Z \subset A \subset X$ such that $\overline{Z}$ is contained in the interior of A, then the inclusion $(X - Z, A - Z) \hookrightarrow (X, A)$ induces isomorphisms $H_N(X - Z, A - Z) \to H_n(X, A)$ for all n. Equivalently, for subspaces $A, B \subset X$ whose interiors cover X, the inclusion $(B, A \cap B) \hookrightarrow (X, A)$ induces isomorphisms $H_n(B, A \cap B) \to H_n(X, A)$ for all n.*

We only use this theorem once, and the proof is quite laborous, so it is omitted here. For a fully fleshed-out proof of the excision theorem the reader is directed to Theorem 2.20 in [1].

As we mentioned earlier, one thing we would very much like to have is a way to build up the homology of a space X by considering it as the union of smaller spaces. The fundamental group has this property in van Kampen's Theorem, but this does not generalize effectively to higher dimensions, one problem with homotopy groups. As it turns out, homology groups have a largely identical property, the Mayer-Vietoris sequence. This is a long exact sequence that exists for a space X and two subspaces $A, B \subseteq X$ such that X is the union of the interiors of A and B. The sequence is as follows:

$$\ldots \to H_n(A \cap B) \xrightarrow{(i_*, j_*)} H_n(A) \oplus H_n(B) \xrightarrow{k_* - l_*} H_n(X) \xrightarrow{\partial_*} H_{n-1}(A \cap B) \to$$

$$\ldots \to H_0(X) \to 0$$

Here the homomorphisms are induced by the maps $i : A \cap B \hookrightarrow A, j : A \cap B \hookrightarrow B, k : A \hookrightarrow X$, and $l : B \hookrightarrow X$. We will not use Mayer-Vietoris sequences at all, but as an analogue for van Kampen's theorem they they are a property of homology groups worth mentioning. One can find a derivation of the Mayer-Vietoris sequence in Chapter 2 of [1].

3.5 Cellular Homology

We already know that a cell structure on a space is a very nice thing to have, and we get many useful properties from that. That niceness extends to the homology groups of a space–it turns out that we can compute cellular homology groups directly from the cell structure of a space, and these homology groups are identical to the singular homology groups introduced before. As such, they carry all of the properties from earlier as well.

To begin, there are a few relevant facts which are worth establishing about the homology groups of a CW complex. They are as follows:

If a space X is a CW complex then the following three properties are true.

1. 1. $H_k(X^n, X^{n-1})$ is zero for $k \neq n$ and for $k = n$ is free abelian with a basis generated by the n-cells of X.
2. 2. $H_k(X^n) = 0$ for $k > n$ Particularly, if X is of finite dimension then $H_k(X) = 0$ for $k >$dim(X).
3. 3. The map $H_k(X^n) \to H_k(X)$ arising naturally from the inclusion $X^n \hookrightarrow X$ is an isomorphism for $k < n$ and is surjective for $k = n$.

The proof of these statements can be found in [1], proposition 2.34. We will not reproduce it here as they are not central results, but rather useful pieces of information.

Theorem 3.5 *$H_n^{CW}(X) \approx H_n(X)$, where $H_n(X)$ is the n^{th} singular homology group, and $H_n^{CW}(X)$ is the n^{th} cellular homology group, which we will now define.*

Let X be a CW complex, then by stitching together pieces of the exact sequences for the relative homology groups we arrive at the following diagram:

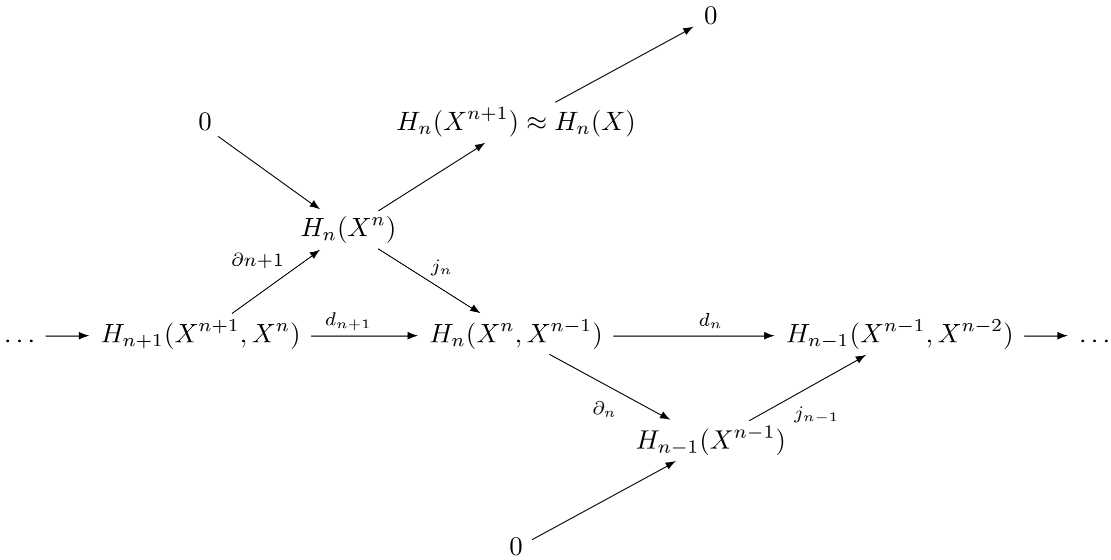

Note that the isomorphism $H_n(X^{n+1}) \approx H_n(X)$ comes from fact 3 from before.

In this diagram the maps d_n and d_{n+1} are defined simply by the compositions $j_{n-1}\partial_n$ and $j_n\partial_{n-1}$ in the diagram. It we look at $d_n d_{n+1}$, we can see that it is the composition $j_{n-1}\partial_n j_n \partial_{n+1}$. which contains the composition $\partial_n j_n$, which is contained in the exact sequence for (X^n, X^{n-1}). As such, the composition $d_n d_{n+1}$ is zero, so the sequence traversed by the maps d_k is a chain complex ($H_n(X^n, X^{n-1})$ is a free abelian group generated by the n-cells of X, so the rest of the requirements of something to be a chain complex are satisfied). We call the homology groups of this complex the cellular homology groups of X, and write them as $H_n^{CW}(X)$. The proof that these are isomorphic to the singular homology groups $H_n(X)$ comes from considering the relationships between maps in the above diagram.

Proof We know from property 4 of exact sequences listed earlier that $H_n(X)$ is isomorphic to $H_n(X^n)/\text{Im}(\partial_{n+1})$. Because j_n is injective, it gives rise to an isomorphism between $\text{Im}(\partial_{n+1}) \subseteq H_n(X^n)$ and $\text{Im}(j_n\partial_{n+1}) \subseteq H_n(X^n, X^{n-1})$. Note that $\text{Im}(j_n\partial_{n+1}) = \text{Im}(d_{n+1})$, so this means that $\text{Im}(\partial_{n+1}) \approx \text{Im}(d_n)$. Additionally, because j_{n-1} is also injective we know by the same argument that $\text{Ker}(d_n) \approx \text{Ker}(\partial_n)$.

Because $\text{Im}(\partial_{n+1}) \approx \text{Im}(d_{n+1})$ and $\text{Ker}(d_n) \approx \text{Ker}(\partial_n)$, we know that the quotients are also isomorphic, so that we get $H_n(X) \approx H_n(X^n)/\text{Im}(\partial_{n+1}) \approx \text{Ker}(d_n)/\text{Im}(d_{n+1}) = H_n^{CW}(X)$.

This equivalence is quite important to us, as the cellular homology groups of a space X can be computed directly from a cell structure on that space. Singular homology groups, on the other hand, are rooted in monstrous free abelian groups, often with uncountably many generators. Computation becomes much more straightforward then as the groups required to compute cellular homologies are so much more manageable than those required for singular homology.

4 Morse Theory: Constructing CW Complexes

In light of the equivalence of cellular homology and singular homology, it is natural that we would like to know how reliably we can apply the techniques of cellular homology to manifolds in general. In order to define the cellular homology groups of a manifold the manifold has to have a cell structure on it, so which manifolds can be given a cell structure is therefore an important question for us to try and answer. Morse theory gives us with an answer to that question by providing a way to obtain a cell structure on a manifold M using a function on the manifold with certain important properties.

Before we begin proving any results in this section, it is worth taking a little bit of time to state some definitions which we will use here and throughout the rest of this thesis. For starters, any time we consider a manifold M it will be a smooth manifold, "smooth" here meaning C^∞. If we consider a smooth function $f : M \to \mathbb{R}$, a point p in M is called a critical point of f if the map $f_* : TM_p \to T\mathbb{R}_{f(p)}$ is zero. The image of p is called a critical value of f. If a point $q \in \mathbb{R}$ is not a critical value we call it a regular value.

A bit of notation we will use quite often is M^a, defined as

$$M^a = \{x \in M | f(x) \leq a\}.$$

If a is a regular value, then M^a is a smooth manifold with boundary.

We call a critical point p non-degenerate if the matrix

$$(\frac{\partial^2 f}{\partial x^i \partial x^j}(p))$$

is nonsingular, and we call the functional associated to this matrix the Hessian of f at p. Obviously, if p is a non-degenerate critical point of f then the above matrix must be of full rank at p. An additional piece of terminology we will use is calling a smooth function $f : M \to \mathbb{R}$ a Morse function if it has no degenerate critical points.

One last important definition is that of the index of a bilinear functional. For a functional H, the index of the functional is the maximum dimension of a subspace V on which H is negative definite. Specifically, we will call the index of the Hessian on TM_p the index of f at p. The index of a critical point is a crucial concept throughout this section, as the construction of the CW complex associated to a manifold M stems precisely from the indices of the critical points of a function on M. The following lemma, which we will prove half of, lays the groundwork for this construction as well as gives a more tangible meaning to the index of a critical point.

Lemma 4.1 *If p is a non-degenerate critical point of f (from the previous lemma), then in a neighborhood U of p there exists a local coordinate system $(y^1, \ldots, y^n)$ with $y^i(p) = 0$ for all i and such that the identity*

$$f = f(p) - (y^1)^2 - \ldots - (y^\lambda)^2 + (y^{\lambda+1})^2 + \ldots + (y^n)^2$$

holds through U. In this case, λ is the index of f at p.

Proof The part of this lemma that we will prove is that if such an f exists, then λ must be the index of f at p. If we take the coordinate system $(z^1, \ldots, z_n)$ at p such that

$$f(q) = f(p) - z^1(q)^2 - \ldots - z^\lambda(q)^2 + z^{\lambda+1}(q)^2 + \ldots + z^n(q)^2$$

then we know that

$$\frac{\partial^2 f}{\partial z^i \partial z^j}(p) = \begin{cases} -2 & \text{if } i = j \leq \lambda \\ 2 & \text{if } i = j > \lambda \\ 0 & \text{if } i \neq j \end{cases}$$

This is equivalent to the matrix of the Hessian of f with respect to the basis generated by the coordinate functions at p being diagonal, with λ (-2)'s followed by $(n - \lambda)$ 2's along the diagonal of the matrix. So, there is a λ-dimensional subspace W of TM_p where the Hessian is negative definite, and an $(n - \lambda)$-dimensional subspace V where the Hessian is positive definite. Each of these subspaces is the largest of their respective types, as if there were a negative definite subspace of dimension greater than λ then it would intersect with V, which cannot be the case as V is positive definite. A similar argument shows that V is also the positive definite subspace of maximal dimension. So, λ is the index of the Hessian at p by the definition of index.

We will not prove the second half of this lemma, that such a coordinate system exists, as the proof revolves around manipulating and transforming different coordinate systems and is not helpful to understanding the rest of this section. The proof of the first half is helpful as it helps ground the idea of the index of a critical point as well as demonstrates its importance.

The process for building a CW complex for a manifold M boils down to traversing the manifold by considering M^a for increasing values of a and attaching cells at critical points of f. The following theorems detail the process of attaching these cells, culminating eventually with a theorem which describes the cell structure associated to a manifold. To begin, we show that this structure depends entirely on the critical points of f.

Theorem 4.2 *Suppose f is a smooth real-valued function on a manifold M. Assume also that a and b are numbers in $\mathbb{R}$ with $a < b$, and that the set $f^{-1}[a, b]$*

is compact and contains no critical points of f. It follows that M^a is diffeomorphic to M^b. In fact, M^a is a deformation retract of M^b, so the inclusion map $M^a \hookrightarrow M^b$ is a homotopy equivalence.

Before we proceed with the proof of this theorem, we must introduce another couple of important definitions.

To start, we define the gradient vector field of a function on a Riemannian manifold. Consider a Riemannian manifold (M, g), and let $< X, Y >$ denote the inner product of two tangent vectors provided by the metric. For a smooth function $f : M \to \mathbb{R}$, the gradient vector field of f with respect to the metric g is the smooth vector field ∇f such that for every smooth vector field V on M, $< \nabla f, V >= df(V) = V \cdot f$. It should be noted that ∇f is unique.

Additionally, we must define a 1-parameter group of diffeomorphisms on a manifold M. A 1-parameter group of diffeomorphisms on a manifold M is a continuous action of the real numbers on M given by the map $\phi : \mathbb{R} \times M \to M$ with the following two properties.

1. For every $t \in \mathbb{R}$ the map $\phi_t : M \to M$ defined as $\phi_t p = \phi(t, p)$ is a diffeomorphism of M onto M.
2. $\phi_t \circ \phi_s = \phi_{t+s}$ for all pairs $t, s \in \mathbb{R}$

An important fact about 1-parameter groups which will be used in the following proof is as follows:

Lemma 4.3 *A smooth vector field on M which vanishes outside of a compact set $K \subset M$ generates a unique 1-parameter group of diffeomorphisms for M.*

For a proof of this result, the reader is directed to Section I.2 of [2] (Lemma 2.4 and proof). With these two definitions in mind, we proceed with the proof of the theorem.

Proof Essentially, what we want to do is push M^b down to M^a orthogonally through the level sets $f^{-1}(c)$ for c in $[a, b]$. This should deform M^b smoothly into M^a and provide us with a suitable deformation retraction.
In order to do this, we take ρ to be smooth function $\rho : M \to \mathbb{R}$ such that

$$\rho = \frac{1}{< \nabla f, \nabla f >}$$

within the set $f^{-1}[a, b]$, and vanishes outside of a compact neighborhood of this set. We then let X be the vector field defined as

$$X_q = \rho(q)(\nabla f)_q$$

So, this vector field points in the same direction as ∇f, but with inverse magnitude. This set satisfies the conditions of Lemma 4.3, so it generates a unique 1-parameter group of diffeomorphisms $\phi_t : M \to M$. Now, we fix some q in M and, if it is contained in the set $f^{-1}[a, b]$ then we know that

$$\frac{df(\phi_t(q))}{dt} \;=\; < \frac{d\phi_t(q)}{dt}, \nabla f > \;=\; < X, \nabla f > \;=\; 1.$$

So, we know that the map $t \to f(\phi_t(q))$ is then linear with derivative equal to 1 when $f(\phi_t(q))$ is between a and b. Because of this, ϕ_{b-a} maps M^a diffeomorphically onto M^b. We then define the 1-parameter family of maps $r_t : M^b \to M^b$ as

$$r_t(q) = \begin{cases} q & \text{if } f(q) \leq a \text{ (i.e. if } q \text{ is in } M^a \\ \phi_{t(a-f(q))}(q) & \text{if } a \leq f(q) \leq b \end{cases}$$

For this family of maps, r_0 is the identity map on M^b, and the homotopy-equivalent r_1 retracts M^b onto M^a. This is the definition of a deformation retraction, so we are done.

Theorem 4.4 *Let $f : M \to \mathbb{R}$ be a smooth function, and suppose p in M is a non-degenerate critical point of index λ. Say $f(p) = c$, and assume that for some $\epsilon > 0$ $f^{-1}[c-\epsilon, c+\epsilon]$ is compact and contains no other critical points of f. Then, for all ϵ small enough, $M^{c+\epsilon}$ is homotopy equivalent to $M^{c-\epsilon}$ with a λ-cell attached to it.*

Proof The soul of this proof is contained in two separate parts. The first step is to select a neighborhood H of p in M and show that $M^{c-\epsilon} \cup e^\lambda$ is a deformation retract of $M^c \cup H$. Following this, we use the preceding theorem to show that $M^{c-\epsilon} \cup H$ is a deformation retract of $M^{c+\epsilon}$. Combining these 2 retractions then gives the result we're looking for. To begin, we look for the retraction from $M^{c-\epsilon} \cup H$ to $M^{c-\epsilon} \cup e^\lambda$. We want to choose a coordinate system $u^1, \ldots, u^n$ in a neighborhood U of p where the following property holds

$$f = f(p) - (u^1)^2 - \ldots - (u^\lambda)^2 + (u^\lambda + 1)^2 + \ldots + (u^n)^2$$

By the Morse Lemma we know that such a coordinate system exists, as p is non-degenerate. We also know that $u^k(p) = 0$ for all $1 \leq k \leq n$. Following this, we choose a small enough $\epsilon > 0$ such that the following two properties hold.

1. $f^{-1}[c-\epsilon, c+\epsilon]$ is compact and only contains p as a critical point.
2. U's image in $\mathbb{R}^n$ coming from this coordinate system contains the closed ball
$$\{(u^1, \ldots, u^n) | \sum_i (u^i)^2 \leq 2\epsilon\}$$

We then define e^λ to be the set of points in U where

$$(u^1)^2 + \ldots + (u^\lambda)^2 \leq \epsilon \text{ and } u^{\lambda+1}, \ldots, u^n = 0$$

It is worth noting that the intersection of $e^\lambda \cap M^{c-\epsilon}$ is the boundary $\hat{e}^\lambda$. Now that we have established exactly the λ-cell we are using, we proceed to establish the two deformation retractions we are looking for. To do so, we construct a new function $F : M \to \mathbb{R}$ as follows. First we take a function $\mu : \mathbb{R} \to \mathbb{R}$ with the following properties.

1. $\mu(0) > \epsilon$
2. $\mu(r) = 0$ for $r \geq 2\epsilon$
3. $-1 \leq \frac{d\mu}{dr} \leq 0$ for all r. We refer to $\frac{d_\mu}{d_r}$ as μ'.

We now use μ to define F by slightly perturbing f in the region U. That is, we define F as

$$F = \begin{cases} f & \text{for } x \text{ outside of the neighborhood } U \\ f - \mu((u^1)^2 + \ldots + (u^\lambda)^2 + 2(u^{\lambda+1})^2 & + \ldots + 2(u^n)^2) \text{ for } x \text{ contained in } U \end{cases}$$

For the following steps, it simplifies computations if we define two functions $\xi, \eta : U \to [0, \infty)$ as follows:

$$\xi(u^1, \ldots, u^n) = (u^1)^2 + \ldots + (u^\lambda)^2$$
$$\eta(u^1, \ldots, u^n) = (u^{\lambda+1})^2 + \ldots + (u^n)^2$$

Intuitively, they represent the magnitude of a point $(u^1, \ldots, u^n)$'s presence in the negative definite and positive definite subspaces of TM_p. Following immediately from these definitions, we know that $f = c + \xi - \eta$ and that $F = c - \xi + \eta - \mu(\xi + 2\eta)$.

To start, we know that the region $F^{-1}(-\infty, c+\epsilon]$ is identical to $f^{-1}(\infty, c+\epsilon] = M^{c+\epsilon}$. This is because when $\xi + 2\eta \leq 2\epsilon$ we know that $\mu(\xi + 2\eta) = 0$ When $\xi + 2\eta \leq 2\epsilon$ we know the following:

$$F \leq f = c - \xi + \eta \leq c \leq \frac{\xi}{2} \leq c + \epsilon$$

We also know that F has the same critical point set as f on M. The functions are equal outside of U, so we look in the region U. Examining the partials

$$\frac{\partial F}{\partial \xi} = -1 - \mu'(\xi + 2\eta)$$

$$\frac{\partial F}{\partial \eta} = 1 - 2\mu'(\xi + 2\eta)$$

and observing that

$$dF = \frac{\partial F}{\partial \xi} d\xi + \frac{\partial F}{\partial \eta} d\eta$$

we can see that $dF = 0$ only when $d\xi$ and $d\eta$ are zero. However, this happens only at the origin, so the origin can be the only critical point of F in U. So, the critical points of f and F are identical.

Now, we look at the region $F^{-1}[c - \epsilon, c + \epsilon]$. We know from previously that $F^{-1}(-\infty, c + \epsilon]$ is identical to $M^{c+\epsilon}$, and because $F \leq f$ in all of M, we know that $M^{c-\epsilon} \subseteq F^{-1}(-\infty, c - \epsilon]$. From these two observations, we know that $F^{-1}[c-\epsilon, c+\epsilon] \subseteq f^{-1}[c-\epsilon, c+\epsilon]$. As $F^{-1}[c-\epsilon, c+\epsilon]$ is closed and a subset of compact $f^{-1}[c - \epsilon, c + \epsilon]$, it is compact as well. As a subset of $f^{-1}[c - \epsilon, c + \epsilon]$, it also contains no critical points of F other than p. However, we know that $F(p) = c - \mu(0) < c - \epsilon$, so $F^{-1}[c-\epsilon, c+\epsilon]$ does not contain p as a critical point either. As $F^{-1}[c-\epsilon, c+\epsilon]$ contains no critical points, we know from the previous theorem that $F^{-1}(-\infty, c-\epsilon]$ is a deformation retract of $F^{-1}(-\infty, c+\epsilon] = M^{c+\epsilon}$. For convenience's sake, we will refer to $F^{-1}(-\infty, c-\epsilon]$ as $M^{c-\epsilon} \cup H$, where we define H to be the closure of $(F^{-1}(-\infty, c-\epsilon] - M^{c-\epsilon})$.

We now consider the cell e^λ which consists of all points q with $\xi(q) \leq \epsilon$, and $\eta(q) = 0$. We know that $\frac{\partial F}{\partial \xi} < 0$ so for all $q \in e^\lambda$ we have that $F(q) \leq F(p) < c - \epsilon$. However, $f(q) \geq c - \epsilon$ for $q \in e^\lambda$, so e^λ is in the region H.

Now, we seek to prove that $M^{c-\epsilon} \cup e^\lambda$ is a deformation retract of $M^{c-\epsilon} \cup H$. We proceed simply by defining the deformation retraction that we want. We denote this deformation retraction $r_t : M^{c-\epsilon} \cup H \to M^{c-\epsilon} \cup H$, and it will be defined differently for different parts of this space. For starters, we can define r_t to be the identity outside of U. Within U, the definition is a bit more complicated, taking on the form of three cases depending on the value of ξ.

The first case is when $\xi \leq \epsilon$, in which case we let r_t be the transformation

$$(u^1, \ldots, u^n) \to (u^1, \ldots, u^\lambda, tu^{\lambda+1}, \ldots, tu^n).$$

So, r_1 is simply the identity map on this region and r_0 maps $(u^1, \ldots, u^n)$ to $(u^1, \ldots, u^\lambda, 0, 0, \ldots, 0)$, with the last $(n - \lambda)$ coordinates equal to zero. This is precisely the region e^λ.

The second case is for the region $\epsilon \leq \xi \leq \eta + \epsilon$. In this region, we define r_t as mapping

$$(u^1, \ldots, u^n) \to (u^1, \ldots, u^\lambda, s_t u^{\lambda+1}, \ldots, s_t u^n).$$

We define the number s_t as

$$s_t = t + (1-t)(\frac{\xi - \epsilon}{\eta})^{\frac{1}{2}}.$$

We have s_1 is simply 1, and s_0 is $(\frac{\xi-\epsilon}{\eta})^{\frac{1}{2}}$. So, r_1 is the identity map, and r_0 takes $(u^1, \ldots, u^n)$ to $(u^1, \ldots, u^\lambda, (\frac{\xi-\epsilon}{\eta})^{\frac{1}{2}} u^{\lambda+1}, \ldots, (\frac{\xi-\epsilon}{\eta})^{\frac{1}{2}} u^n)$. Taking $f(r_0(u^1, \ldots, u^n))$, we get $f = c - \xi + \eta(\frac{\xi-\epsilon}{\eta}) = c - \epsilon$. So, r_0 takes all of region 2 into $f^{-1}(c - \epsilon)$.

In the third case, we consider the region $\eta + \epsilon \leq \xi$, we just define r_t to be the identity. It should be noted that when $\xi = \eta + \epsilon$ in case 2 we get that s_t is equal to 1 so r_t is the identity map. Thus, we have defined a deformation retraction from $F^{-1}(-\infty, c+\epsilon]$ to $M^{c-\epsilon} \cup e^\lambda$. Composing this with the deformation retraction from $M^{c+\epsilon}$ to $F^{-1}(-\infty, c+\epsilon]$ we have a deformation retraction from $M^{c+\epsilon}$ to $M^{c-\epsilon} \cup e^\lambda$ and are done.

Now that we know how to attach cells to critical points of f, we can build up a cell structure on a manifold M using the following theorem.

Theorem 4.5 *If f is a Morse function on a manifold M and M^a is compact for each a, then M is homotopy-equivalent to a CW complex, with each critical point of index λ corresponding to a different λ-cell.*

Before we begin the proof of this theorem, we first prove two additional lemmas which the proof hinges on. The lemmas and their proofs are as follows.

Lemma 4.6 *Let ϕ_0 and ϕ_1 be homotopic maps from the sphere $\hat{e}^\lambda$ to a space X. Then, the identity map of X extends to a homotopy equivalence*

$$k : X \cup_{\phi_0} e^\lambda \to X \cup_{\phi_1} e^\lambda$$

Proof We prove this simply by defining the homotopy equivalence k as follows.

$$\begin{cases} k(x) = x & \text{for } x \text{ in } X \\ k(tu) = 2tu & \text{for } 0 \le t \le \frac{1}{2} \text{ and } u \text{ in } \hat{e}^\lambda \\ k(tu) = \phi_{2-2t}\phi(u) & \text{for } \frac{1}{2} \le t \le 1 \text{ and } u \text{ in } \hat{e}^\lambda \end{cases}$$

ϕ_t indicates a homotopy map between ϕ_0 and ϕ_1 (we know this exists from the conditions of the lemma). tu also indicates multiplying a unit vector in $\hat{e}^\lambda$ with $t \in [0, 1]$.

Lemma 4.7 *If we let $\phi : \hat{e}^\lambda \to X$ be an attaching map then any homotopy equivalence $f : X \to Y$ extends to a homotopy equivalence*

$$F : X \cup_\phi e^\lambda \to Y \cup_{f\phi} e^\lambda$$

Proof Let us start off by defining a function F as follows:

$$F = \begin{cases} f \text{ for } x \text{ in } X \\ \text{identity for } x \text{ in } e^\lambda \end{cases}$$

Let $g : Y \to X$ be a homotopy inverse to f and define then $G : Y \cup_{f\phi} e^\lambda \to X \cup_{gf\phi} e^\lambda$ in a similar manner to F. That is, set $G|Y = g$ and $G|e^\lambda$ is the identity map.

We know then that $gf\phi$ is homotopic to ϕ, so it follows from the previous lemma that there is a homotopy equivalence

$$k : X \cup_{gf\phi} e^\lambda \to X \cup_\phi e^\lambda$$

Firstly, we prove that the triple-composition

$$kGF : X \cup_\phi e^\lambda \to X \cup_\phi e^\lambda$$

is homotopic to the identity map on $X \cup_\phi e^\lambda$. Let h_t be a homotopy between gf and the identity function. Then, from the definitions of k, F, and G we get that

$$\begin{cases} kGF(x) = gf(x) & \text{for } x \text{ in } X \\ kGF(tu) = 2tu & \text{for } 0 \le t \le \frac{1}{2} \text{ and } u \text{ in } \hat{e}^\lambda \\ kGF(tu) = h_{2-2t}\phi(u) & \text{for } \frac{1}{2} \le t \le 1 \text{ and } u \text{ in } \hat{e}^\lambda \end{cases}$$

From this set of equations we can define an appropriate homotopy

$$q_\tau : X \cup_\phi e^\lambda \to X \cup_\phi e^\lambda$$

This homotopy is as follows:

$$\begin{cases} q_\tau(x) = h_\tau(x) & \text{for } x \text{ in } X \\ q_\tau(tu) = \frac{2}{1+\tau} & \text{for } 0 \leq t \leq \frac{1+\tau}{2} \text{ and } u \text{ in } \hat{e}^\lambda \\ q_\tau(tu) = h_{2-2t+\tau}\phi(u) & \text{for } \frac{1+\tau}{2} \leq t \leq 1 \text{ and } u \text{ in } \hat{e}^\lambda \end{cases}$$

In order to complete the proof of this lemma, it is important to prove the following statement about left and right homotopy inverses.

Aside: If a map F has a left homotopy inverse L and a right homotopy inverse R, then F is a homotopy equivalence and R and L are two-sided inverses.

The proof is relatively simple. $LF \simeq$ identity and $FR \simeq$ identity give that $L \simeq L(FR) = (LF)R \simeq R$, so $RF \simeq LF \simeq$ identity. This shows that R is a 2-sided inverse, and a similar proof shows that L is as well, so we are done.

With this in mind, we examine the functions we have defined for this proof. We know that $kGF \simeq$ identity, so we know that F has a left homotopy inverse. However, were we to construct G via a function g such that f were a homotopy inverse of g, we would have that $kFG \simeq$ identity, so G has a left homotopy inverse.

From Lemma 4.6, we know that k is a homotopy equivalence, so that means that $k(GF) \simeq$ identity implies that $(GF)k \simeq$ identity. This means that G has a right homotopy inverse, which from the previous aside indicates that G is a 2-sided homotopy inverse of (Fk). So, we know that $(Fk)G \simeq$ identity. From this and the knowledge that F has a left homotopy inverse, we get that F is a homotopy equivalence and we are done.

With these two lemmas proved, we continue on to the proof of the main theorem.

Proof Let $c_1 < c_2 < c_3 < \ldots$ be the critical values of a smooth function $f : M \to \mathbb{R}$. The sequence has no accumulation point since we know that each M^a is compact. Additionally, if $\mathbb{R} \ni a < c_1$ then we know that M^a is empty. So, suppose that $a \neq c_n$ for any n, and that M^a is homotopy equivalent to a CW complex. Let us take c to be the smallest $c_i > a$. By the preceding theorems, we know that $M^{c+\epsilon}$ is homotopic to $M^{c-\epsilon} \cup_{\phi_1} e^{\lambda_1} \cup_{\phi_2} \ldots \cup_{\phi_{j(c)}} e^{\lambda_{j(c)}}$ for attachment maps $\phi_1, \ldots, \phi_{j(c)}$. We also know that there exists a homotopy equivalence $h : M^{c-\epsilon} \to M^a$, and we have assumed that a homotopy equivalence $h' : M^a \to K$ exists, where K is a CW complex.

In this case, we have that for each i, $h' \circ h \circ \phi_i$ is homotopic to a map ψ_i, which attaches the cell e^{λ_i} onto the $(i-1)$-skeleton of K. So, we have that $K \cup_{\psi_1} e^{\lambda_1} \cup_{\psi_2} \ldots \cup_{\psi_{j(c)}} e^{\lambda_{j(c)}}$ is again a CW complex, and is homotopy equivalent to $M^{c+\epsilon}$ by Lemmas 4.6 and 4.7.

Inducting, we know that $M^{a'}$ is homotopy equivalent to a CW complex for every a'. If M has all its critical points contained in a compact set M^a then we are done, as even if M is not compact, if all its critical points are compact then our result follows from Theorem 4.2. If M has its critical points contained in some noncompact set then we get an infinite chain

$$\begin{array}{ccccccc} M^{a_1} & \subset & M^{a_2} & \subset & M^{a_3} & \subset & \dots \\ \downarrow & & \downarrow & & \downarrow & & \\ K_1 & \subset & K_2 & \subset & K_3 & \subset & \dots \end{array}$$

where each downward arrow is a homotopy equivalence coming from earlier in the proof. Define K to be the union of the K_i in the limit topology, and let $g : M \to K$ be the limiting map. As g is the limit of homotopy equivalences, g induces isomorphisms of homotopy groups of every dimension. This completes the proof for the case of a noncompact interval being needed to encompass all the critical points of M, and thus completes the proof of the Theorem.

So, given a Morse function on a manifold M, we can use that function to build up a CW complex which is homotopy-equivalent to M. In light of the homotopy-invariance of the homology groups of a space, this means that we can apply the techniques of cellular homology to any smooth manifold that we can find a Morse function on. However, the need for a Morse function on M raises the question of how often we can find such a function. We get a very well-behaved structure on a manifold if there is a Morse function on it, but it would be a shame if Morse functions were too uncommon to be really useful.

Luckily, Morse functions are very easy to find for any smooth manifold M. We will not detail the process here, but in Section I.6 of [1], Milnor shows that for any smooth manifold M there exists a Morse function for which M^a is compact for every a in $\mathbb{R}$. So, for any smooth manifold M we can find a Morse function on M, which means that every smooth manifold is homotopy-equivalent to a CW complex.

5 Gradient Flows, Stable and Unstable Manifolds, and Morse-Smale functions

In this section we will lay much of the groundwork for the Morse Homology Theorem. We begin by defining gradient flows on a manifold and stating some of their properties, and following that we explore stable and unstable manifolds, which are an important way of characterizing a flow on a manifold. Specifically, for a flow on a manifold, the stable and unstable manifolds of a critical point rigorize the idea of spaces flowing into and out of that point. With stable and unstable manifolds developed we will define Morse-Smale functions, which are Morse functions that satisfy an additional constraint on the intersections of their stable and unstable manifolds. We then detail some important results concerning flow lines between critical points which we will use extensively in the proof of the Morse Homology Theorem.

5.1 Gradient Flows

Before we proved Theorem 4.2 earlier, we defined the gradient vector field of a function on a manifold, and we also defined a 1-parameter group of diffeomorphisms on a manifold. In this section, we will combine the two concepts to create a 1-parameter group of diffeomorphisms which is generated by the gradient vector field of a function f. Specifically, we define a gradient flow on a manifold as follows: given a smooth function f on a manifold M, we take ϕ_t to be the 1-parameter group of diffeomorphisms generated by the $-\nabla f$. This is the group with

$$\begin{aligned} \phi_0(x) = x \text{ and} \\ \tfrac{d}{dt}\phi_t(x) = -(\nabla f)(\phi_t(x)). \end{aligned}$$

Given this group ϕ_t, we can define the curve $\gamma_x : (a, b) \to M$ by

$$\lambda_x(t) = \phi_t(x)$$

For every $x \in M$. This curve is called a gradient flow line. This is due to the following 2 results, Propositions 3.18 and 3.19 from [3], presented here without proof.

Lemma 5.1 *Every smooth function $f : M \to \mathbb{R}$ on a finite-dimensional smooth Riemannian manifold (M, g) decreases along its gradient flow lines.*

Lemma 5.2 *Let $f : M \to \mathbb{R}$ be a Morse function on a finite dimensional compact smooth Riemannian manifold (M, g). Then every gradient flow line f begins and ends at a critical point of M.*

5.2 Stable and unstable manifolds

Another important definition which we will use is that of stable and unstable manifolds. These rigorize the ideas of flows into and out of a point. These manifolds are defined for any non-degenerate critical point $p \in M$ of a smooth function f.

The stable manifold of p is the set $W^s(p) = \{x \in M | \lim_{t\to\infty} \phi_t(x) = p\}$

The unstable manifold of p is the set $W^u(p) = \{x \in M | \lim_{t\to-\infty} \phi_t(x) = p\}$

Essentially, the stable manifold of the point p is the set of points which flow into p under the 1-parameter group ϕ_t, and the unstable manifold is the set of points which flow out of p under ϕ_t. There is an extremely important theorem to do with the stable and unstable manifolds of critical points, whose proof is covered through most of Chapter 4 of [3].

Theorem 5.3 *Let $f : M \to \mathbb{R}$ be a Morse function on a compact smooth Riemannian manifold (M, g) of dimension $m < \infty$. If $p \in M$ is a critical point of f, then the tangent space at p splits as*

$$T_pM = T_p^sM \bigoplus T_p^uM$$

where the Hessian is positive definite on T_p^sM and negative definite on T_p^uM. Additionally, the stable and stable manifolds are surjective images of smooth embeddings

$$E^s : T_p^sM \to W^s(p) \subseteq M$$
$$E^u : T_p^uM \to W^u(p) \subseteq M$$

So, $W^s(p)$ is a smoothly embedded open disk of dimension $m - \lambda_p$, and $W^u(p)$ is a smoothly embedded open disk of dimension λ_p, where λ_p is the index of p.

With stable and unstable manifolds defined, we can use them to define a Morse-Smale function on M, which we will be basing the proof of the Morse Homology Theorem on.

5.3 Morse-Smale Functions

We have already dealt with Morse functions on a manifold M. That is, smooth functions $f : M \to \mathbb{R}$ which have no degenerate critical points. However, a stronger condition on the function, called the Morse-Smale transversality condition, places additional constraints on the stable and unstable manifolds of the critical points of f. This condition is as follows:

A Morse function on a finite-dimensional smooth Riemannian Manifold (M,g) is said to satisfy the Morse-Smale transversality condition if the stable and unstable manifolds of f intersect transversally, that is to say that

$$W^u(q) \pitchfork W^s(p)$$

for all critical points p, q of f. If a Morse function satisfies this condition then we call if a Morse-Smale function.

Stemming directly from this transversality condition is a categorization of intersections of the stable and unstable manifolds of different critical points.

Lemma 5.4 *Let $f : M \to \mathbb{R}$ be Morse-Smale function on a finite-dimensional compact smooth manifold (M, g). If p and q are critical points of f with non-trivial intersection, then $W^u(q) \cap W^s(p)$ is an embedded $(\lambda_q - \lambda_p)$-dimensional submanifold of M.*

Proof The proof comes simply from the fact that $W^u(q)$ is a λ_q-dimensional submanifold of M, and $W^s(p)$ is a $(m - \lambda)$-dimensional submanifold. So, their intersection is then a submanifold whose dimension $\dim(W^u(q) \cap W^s(p))$ is equal to

$$\dim W^u(q) + \dim W^s(p) - m = \lambda_q + (m - \lambda_p) - m = \lambda_q - \lambda_p$$

For the future, we will refer to this intersection $W^u(q) \cap W^s(p)$ as $W(q,p)$. Thus, $W(q,p)$ denotes the flow out of a critical point q into a different critical point p.

An immediate consequence of this is the following corollary:

Corollary 5.5 *IF $F : M \to \mathbb{R}$ is a Morse-Smale function on a finite-dimensional compact smooth Riemannian manifold (M,g), then the index of critical points decreases strictly along gradient flow lines. Specifically, if p and q are critical points of f with $W(q,p) \neq$ then $\lambda_q > \lambda_p$.*

The proof of this comes simply from observing that if $W(q,p) \neq$, then ther must be at least one flow line from q into p. Flow lines have dimension 1, so $\dim W(q,p) = \lambda_q - \lambda_p \geq 1$.

One more result about Morse-Smale functions that is worth noting is the following Kupka-Smale Theorem, concerning how many functions on a manifold satisfy the Morse-Smale transversality condition.

Theorem 5.6 *If (M,g) is a finite-dimensional compact smooth Riemannian manifold, then the set of Morse-Smale gradient vector fields of class C^r is a generic subset of the set of all C^r gradient vector fields on M for $1 \leq r \leq \infty$.*

For a proof of the Kupka-Smale Theorem the reader is directed to Section 6.1 of [3]. The reason we mention it here is that almost all of the remaining results of this thesis will be rooted in Morse-Smale functions. As such, it is comforting to know that they are so universal.

6 Corollaries to the λ-lemma

In this section we list some important results for the proof of the Morse Homology Theorem. For proofs of these results, we direct the reader to section 6.3 of [3]. The reason we do not go through the proofs here is that they all depend on the $\lambda-$lemma, a result whose statement and proof are very intricate and not particularly illuminatory for us. Rather, the meat of the λ-lemma for us are the corollaries which follow it. The results are as follows:

Lemma 6.1 *Let M be a smooth manifold, and suppose that p and q are hyperbolic fixed points of ϕ, a diffeomorphism of M. If $W^u(q)$ and $W^s(p)$ have a point of transverse intersection, then*

$$\overline{W^u(q)} \supseteq W^u$$

As $W^u(q)$ and $W^u(p)$ can be thought of as the flows out of q and p respectively, this inclusion effectively means that all of the flow out of p is contained in the closure of q's outflow. The next result has to do with the flows from one point into a second and then into a third.

Lemma 6.2 *Let M be a smooth manifold, and let p_1, p_2, and p_3 be three hyperbolic fixed points of ϕ, a diffeomorphism from M to itself. If $W^u(p_3$ and $W^s(p2)$ have non-null transverse intersection, and the same is true for $W^u(p_2)$ and $W^s(p_1)$ then $W^u(p_3)$ and $W^s(p_1)$ have a point of transverse intersection, and the inclusion*

$$\overline{W(p_3, p_1)} \subseteq W(p_3, p_2) \bigcup W(p_2, p_1) \bigcup \{p_1, p_2, p_3\}$$

This is effectively a transitive property for the outflow from a fixed point p_3, through p_2, and into p_1. In fact, when the diffeomorphism in this preceding lemma is generated by a gradient vector field of a Morse-Smale function f on M, then this result holds for any trio of critical points of the function f. With this in mind, we can define a partial ordering of the critical points of f. For any two critical points p and q for which $W(q, p) \neq \emptyset$ holds, we say that $q \preceq p$.

From this partial ordering we get the following two inclusions:

Corollary 6.3 *For any critical point q of $f : M \to \mathbb{R}$ the inclusion below holds.*

$$\overline{W^u(q)} \supseteq \bigcup_{q \succeq p} W^u(p)$$

Additionally, if p and q are critical points of f such that $q \succeq p$, then

$$\overline{W(q, p)} \supseteq \bigcup_{q \preceq \tilde{q} \preceq \tilde{p} \preceq p} W(\tilde{q}, \tilde{p})$$

The following lemmas show that the opposite inclusions also hold, that is to say that the inclusions in the above lemmas and corollary are actually equalities.

Lemma 6.4 *If $\overline{W^u(q)} \cap W^u(p) \neq$, then $p \in \overline{W^u(q)}$*

Lemma 6.5 *If $p \neq q$ is a critical point of $f : M \to \mathbb{R}$ such that $p \in \overline{W^u(q)}$, then $\overline{W^u(q)}$ intersects $W^s(p) - p$.*

Corollary 6.6 *Suppose that q and p are critical points of $f : M \to \mathbb{R}$, and $\overline{W^u(q)} \cap W^s(p) \neq$. Then*

$$\overline{W^u(q)} \subseteq \bigcup_{q \preceq p} W^u(p).$$

That is to say that $q \preceq p$.

Corollary 6.7 *For any critical point q of $f : M \to \mathbb{R}$ we have that*

$$\overline{W^u(q)} = \bigcup_{q \preceq p} W^u(p)$$

Additionally,

$$\overline{W^s(q)} = \bigcup_{r \preceq q} W^s(r)$$

Corollary 6.8 *If p and q are critical points of $f : M \to \mathbb{R}$ and $q \preceq p$ then*

$$\overline{W9q,p)} = \overline{W^u(q)} \cap \overline{W^s(p)} = \bigcup_{q \preceq \tilde{q} \preceq \tilde{p} \preceq p} W(\tilde{q}, \tilde{p})$$

where this union is taken over all critical points between p and q under the partial order taken earlier.

Corollary 6.9 *If p and q are critical points of relative index one ($\lambda_q = \lambda_p = 1$), then*

$$\overline{W(q,p)} = W(q,p) \cup p, q$$

Additionally, $W(q,p)$ has finitely many components, that is to say that there are a finite number of gradient flow lines from q to p.

With these results under our belt we can proceed to the culminating result of this thesis, the Morse Homology Theorem.

7 The Morse Homology Theorem

The result we have been building to is the Morse Homology Theorem, which identifies the singular homology groups of a manifold M with the homology groups generated by the critical points of a Morse-Smale function $f : M \to \mathbb{R}$. However, in order to do this we must define the homology groups generated by the critical points of f. To do this, we introduce the Morse-Smale-Witten chain complex of a function f on a manifold .

Definition Let $f : M \to \mathbb{R}$ be a Morse-Smale function on a compact smooth oriented Riemannian manifold M of dimension $M < \infty$, and assume that orientations for the unstable manifolds of f have been chosen. Let $C_k(f)$ be the free abelian group generated by $Cr_k(f)$, and define

$$C_*(f) = \bigoplus_{k=0}^{m} C_k(f).$$

We now define a homomorphism $\partial_k : C_k(f) \to C_{k-1}(f)$ by

$$\partial_j(q) = \sum_{p \in Cr_{k-1}(f)} n(q,p)p.$$

where $n(q,p)$ is defined below. We call this the Morse-Smale-Witten boundary operator, and the pair $(C_*(f), \partial_*)$ is the Morse-Smale-Witten chain complex of f.

The Morse Homology Theorem states that the Morse-Smale-Witten chain complex is in fact a chain complex and, furthermore, that the homology groups it generates are isomorphic to the singular homology groups $H_n(M)$

7.1 Orientations and $n(q,p)$

We have defined the Morse-Smale-Witten boundary operator already and involved in the summation is the number $n(q,p)$. However, we have not explicitly defined what this number is. Effectively, this number $n(q,p)$ is a way of counting the gradient flow lines from q to p while accounting for the orientations of everything[ELABORATE ON TH IS]

Before we can do this, though, we must first establish some preliminary conventions for how we are orient

So, we take q and p to be critical points of relative index one. Let's say that $\lambda_q = k$ and $\lambda_p = k-1$. Let us assume that $q \preceq p$ (If q does not precede p then there are no gradient flows between q and p, so $n(q,p) = 0$), and let $\gamma : \mathbb{R} \to M$ be a gradient flow line from q to p. At any point x in $\gamma(\mathbb{R})$, we can complete $(-\nabla f)(x)$ to a positive basis of $T_x W^u(q), ((-\nabla f)(x), \hat{B}^u_x)$. If we take a positive basis B^s_x of $T_x W^s(p)$, then $(B^x_s, \hat{B}^u_x)$ is a basis for $T_x M$ (as $\hat{B}^u_x$ has $k-1$ elements, and B^s_x has $(m-k+1)$ elements). If $(B^s_x, \hat{B}^u_x)$ is a positively-oriented basis for $T_x M$ then we assign +1 to this flow γ. If not, then we assign -1 to γ.

7.2 Index Pairs

To start, we let ϕ_t be a 1-parameter group of diffeomorphisms on a locally compact metric space M. A subset S of M is called an invariant subset if $\phi_t(S) = S$ for every $t \in \mathbb{R}$. Following this, we can define the maximal invariant subset $I(N)$ of a subset N of M as follows:

$$I(N) = \{x \in N | \phi_t(x) \in N \text{ for every } t \in \mathbb{R}\}$$

Furthermore, a compact invariant subset S is called isolated if there exists a compact neighborhood N of S such that $I(N) = S$. Following this, we can define index pairs as follows. Given an isolated compact invariant subset S, an index pair (N, L) for S is a pair of compact sets $L \subset N$ such that the following properties hold.

1. $S = I(\overline{(N-L)}) \subseteq int(N-L)$
2. $x \in L$ and $x \cdot [0,t] \subseteq N$ implies that $x \cdot [0,t] \subseteq L$. We denote this property by saying that L is positively invariant in N.
3. If an orbit of a point leaves N, it has to go through L first. More specifically, if for a point x in N we have that $x \cdot t \notin N$ for some $t > 0$, then there must be some $t' \in [0,t]$ with the property that $x \cdot [0,t] \subseteq N$ and $x \cdot t' \in L$.

A couple of key results regarding index groups are the following theorems (Theorems 7.14 and 7.15 in [3]).

Theorem 7.1 *Every isolated compact invariant set admits an index pair*

Theorem 7.2 *If S is an isolated compact invariant set and (N, L) and $(\tilde{N}, \tilde{L})$ are two index pairs for S, then N/L and $\tilde{N}/\tilde{L}$ are homotopy equivalent as pointed spaces via maps induced by the flow ϕ.*

We will not prove these theorems here, but a proof of the first theorem can be found in Section 4.1 of [4] and Section 7.5 of [3] is dedicated to proving the second theorem.

An important property that an index pair (N, L) can have is for the inclusion map $L \hookrightarrow N$ to be a cofibration. If that is the case, then we say that the index pair is regular.

We will begin our examination of index pairs by constructing them for the critical points of a Morse function f. We know for any $q \in Cr_k(f)$ the tangent space T_qM splits into $T_q^sM \oplus T_q^uM$, with $T_q^uM = T_qW^u(q)$ and $T_q^sM = T_qW^s(q)$. We also know that $W^s(q)$ and $W^u(q)$ intersect transversally at q. So, we get a coordinate chart $\phi : U \to T_qM$ around q that maps $W^s(q) \cap U$ into T_q^sM and $W^u(q) \cap U$ into T_q^uM. Now, we take $\epsilon > 0$ and define

$$D_\epsilon^s = \{v \in T_q^sM | \|v\| \leq \epsilon\} \text{ and}$$

$$D_\epsilon^u = \{v \in T_q^uM | \|v\| \leq \epsilon\}$$

Then, let $N_q = \phi_q^{-1}(D_\epsilon^s \times D_\epsilon^u)$ and $L_q = \phi_q^{-1}(D_\epsilon^s \times \partial D_\epsilon^u)$. N_q is an isolating neighborhood for $\{q\}$, a compact invariant set, and (N_q, L_q) is an index pair for this set. Now, we can see that the index pair (N_q, L_q) is regular, and we have that

$$(N_q, L_q) \approx (D_\epsilon^s \times D_\epsilon^u, D_\epsilon^s \times \partial D_\epsilon^u) \simeq (D_\epsilon^u, \partial D_\epsilon),$$

So N_q/L_q is a pointed sphere of dimension λ_q. So, $H_j(N_q, L_q) \approx H_j(D_\epsilon^u, \partial D_\epsilon^u) \approx H_j(D_\epsilon^u/\partial D_\epsilon^u, *) \approx H_j(N_q/L_q, *)$ for every j. Additionally, because an orientation of T_q^uM determines a generator of $H_k(N_q, L_q) \approx H_k(D_\epsilon^u, \partial D_u^\epsilon) \approx \mathbb{Z}$. Therefore, we get an identification

$$C_k(f) \approx \bigoplus_{q \in Cr_k(f)} H_k(N_q, L_q).$$

This identification is quite important to us for proving the Morse Homology theorem.

One particular set of index pairs we will be relying quite heavily on is the following sequence, which provides index pairs for the flows on M determined by $-\nabla f$. To start, we define

$$W(k, j) = \bigcup_{j \leq \lambda_p \leq \lambda_q \leq k} W(q, p)$$

where $0 \leq j \leq k \leq m$. We know that these spaces are compact as they are the images of closed disks in a compact manifold. So, let us take a compact neighborhood N of $W(k, j)$ such that $Cr(f) \cap N = Cr(f) \cap W(k, j)$, then we know that

$$I(N) = \{x \in N | \phi_t(x) \in N \text{ for all } t \in \mathbb{R}\} = W(k, j).$$

So, we know that $W(k,j)$ is an isolated compact invariant set for $0 \leq j \leq k \leq m$. We know from Corollary 6.7 that the sets

$$W_j^s = \bigcup_{j \leq \lambda_p} W^s(p) \text{ and}$$

$$W_j^u = \bigcup_{\lambda_p \leq j} W^u(p)$$

are compact for all $j = 0, \ldots, m$. Set $N_m = M$, and then choose a cofibered compact neighborhood N_{m-1} of W_{m-1}^u in M which is positively invariant in N_m and fulfills the property $N_{m-1} \cap W_m^s = \emptyset$. One such neighborhood is

$$N_m - \bigcup_{q \in Cr_m(f)} intN_q.$$

With N_{m-1} fulfilling these conditions (N_m, N_{m-1}) is a regular index pair for $Cr_m(f)$. Similarly, we then choose a cofibered compact neighborhood N_{m-2} of W_{m-2}^u in N_{m-1} which is positively invariant in N_{m-1} and satisfies $N_{m-2} \cap W_{m-1}^s = \emptyset$. We then know that (N_{m-1}, N_{m-2}) is an index pair for $Cr_{m-1}(f)$ and (N_m, N_{m-2}) is an index pair for $W(m, m-1)$. Repeating this choosing of cofibered compact neighborhoods establishes a sequence

$$\emptyset = N_{-1} \subseteq N_0 \subseteq N_1 \subseteq \ldots \subseteq N_m = M$$

with the property that (N_k, N_{j-1}) is a regular index pair for $W(k,j)$ for all values $0 \leq j \leq k \leq m$.

Before we proceed, we would first like to define an important tool in proving the Morse Homology Theorem. Given the homology exact sequence of the triple $N_0 \subseteq N_1 \subseteq N_2$

$$\begin{array}{ccccccccc}
0 & \longrightarrow & C_k(N_1,N_0) & \xrightarrow{i} & C_k(N_2,N_0) & \xrightarrow{j} & C_k(N_2,N_1) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow \overline{\partial}_k & & \downarrow & & \\
0 & \longrightarrow & C_{k-1}(N_1,N_0) & \xrightarrow{i} & C_{k-1}(N_2,N_0) & \xrightarrow{j} & C_{k-1}(N_2,N_1) & \longrightarrow & 0
\end{array}$$

we define the connecting homomorphism

$$\delta_* : H_k(N_2, N_1) \to H_{k-1}(N_1, N_0)$$

as follows. Let $a \in C_k(N_2, N_1)$ be a $k-cycle$, and let $\overline{a} \in C_k(N_2, N_0)$ be the lift $j^{-1}(a)$. In this case, $\overline{\partial}_k(\overline{a})$ is a cycle in $C_{k-1}(N_2, N_0)$ which comes from a unique cycle $\hat{a}$ in $C_{k-1}(N_1, N_0)$. We define $\delta_*[a] = [\hat{a}]$. It is worth noting that $\delta_* = j_* \circ \delta_k$ where δ_k is the connecting homomorphism for the homology exact sequence of (N_2, N_1) and $j : (N_1, \emptyset) \to (N_1, N_0)$ is the inclusion map. An important peoperty of this connecting homomorphism is that it is a natural transformation, which is to say that if $g : (N_2, N_1, N_0) \to (\tilde{N}_2, \tilde{N}_1, \tilde{N}_0)$ is a map of triples then the diagram below is commutative for every k :

$$\begin{array}{ccc} H_k(N_2, N_1) & \xrightarrow{\delta_*} & H_{k-1}(N_1, N_0) \\ \downarrow g_* & & g_* \downarrow \\ H_k(\tilde{N}_2, \tilde{N}_1) & \xrightarrow{\delta_*} & H_{k-1}(\tilde{N}_1, \tilde{N}_0) \end{array}$$

The homotopy equivalences we get from Theorem 7.2 are flow-induced, so the diagram commutes for regular index pairs.

The last step before we start proving the Morse Homology Theorem is constructing an important group homomorphism Δ_k. To start, we consider some Morse-Smale function $f : M \to \mathbb{R}$ on a finite dimensional smooth compact Riemannian manifold M. Let $q \in Cr_k(f)$ and $p \in Cr_{k-1}$. We know that $S = W(q, p) \cup \{p, q\}$ is an isolated compact invariant set, so by Theorem 7.1 it admits an index pair (N_2, N_0). Let $N_1 = N_0 \cup (N_2 \cap M^c)$, where $f(p) < c < f(q)$. Then, (N_2, N_1) is an index pair for q and (N_1, N_0) is an index pair for p. Assume that these index pairs are regular(Which we can do according to Section 5.1 of [4]).

Let (N_q, L_q) be a regular index pair for q and (N_p, L_p) be a regular index pair for p. Let us define the homomorphism $\Delta_k(q, p) : H_k(N_q, L_q) \to H_{k-1}(N_p, L_p)$ via the composition of maps

$$H_k(N_q, L_q) \xrightarrow{\approx} H_k(N_2, N_1) \xrightarrow{\delta_*} H_{k-1}(N_1, N_0) \xrightarrow{\approx} H_{k-1}(N_p, L_p).$$

The first and third isomorphism come from the homotopy equivalence of Theorem 7.1, and δ_* is the connecting homomorphism in the exact sequence of (N_2, N_1, N_0). Thus, we get the homomorphism

$$\Delta_k : \bigoplus_{q \in Cr_k(f)} H_k(N_q, L_q) \to \bigoplus_{p \in Cr_{k-1}(f)} H_{k-1}(N_p, L_p)$$

7.3 Proving the Morse Homology Theorem

With the preceding definitions taken care of, we can proceed with the Morse Homology Theorem. The proof of this theorem depends strongly on the following lemma, which allows us to construct the commutative diagram on which the Morse Homology Theorem's proof hinges. The lemma is as follows:

Lemma 7.3 *The diagram below commutes.*

$$
\begin{array}{ccc}
C_k(f) & \xrightarrow{\partial_k} & C_{k-1}(f) \\
\updownarrow \approx & & \updownarrow \approx \\
\bigoplus_{q \in Cr_k(f)} H_k(N_q, L_q) & \xrightarrow{\Delta_k} & \bigoplus_{p \in Cr_{k-1}(f)} H_{n-1}(N_p, L_p) \\
\updownarrow \approx & & \updownarrow \approx \\
H_k(N_k, N_{k-1}) & \xrightarrow{\delta_*} & H_{k-1}(N_{k-1}, N_{k-2})
\end{array}
$$

In this diagram, (N_q, L_q) indicates a regular index pair for $q \in Cr_k(f)$, (N_p, L_p) is a regular index pair for $p \in Cr_{k-1}(f)$, ∂_k is the Morse-Smale-Witten boundary operator, Δ_k is the homomorphism defined earlier, and δ_* is a connecting homomorphism.

Proof We know that the bottom square commutes because the connecting homomorphism δ_* is a natural transformation and Δ_k is constructed using δ_*, so we proceed with the top square. We will restrict ourselves to the case where $q \in Cr_k(f)$ and $p \in Cr_{k-1}(f)$ are the only two critical points in $f^{-1}[a, b]$, with $f(p) = a$ and $f(q) = b$. If this is not the case then we need only alter f outside an isolating neighborhood of $S = W(q, p) \cup \{q, p\}$. This alteration affects neither Δ_k nor ∂_k, so we are fine.

We first start off by defining some sets which will lead to convenient index pairs. We have already defined the set $M^c = \{x \in M | f(x) \leq c\}$, and we define M_c to be the corresponding set $M_c = \{x \in M | f(x) \geq c\}$. Let $\epsilon > 0$ be very small, and let $T \gg 0$.

$$
\begin{aligned}
N_q &= \{x \in M_c | f(\phi_{-T}(x)) \leq b + \epsilon\} \\
L_q &= \{x \in N_q | f(x) = c\} \\
N_p &= \{x \in M^c | f(\phi_T(x)) \geq a - \epsilon\} \\
L_p &= \{x \in N_p | f(\phi_T(x)) = a - \epsilon\}
\end{aligned}
\tag{1}
$$

We also define $C = N_p \cup N_q, B = N_p \cup L_q$, and $A = L_p \cup \overline{(L_q - N_p)}$.

The definitions of N_q, N_p, L_q, and L_p are a bit obtuse, so it helps to have an intuitive way of thinking about them. One can think of N_q as the set of points in M which map to a value $\geq c$ and did not flow from too far above q. L_q is just the intersection of this set with $f^{-1}(c)$. N_p are the the points with value $\leq c$ and which don't flow too far past p, and L_p is simply the set of these points which start at c.

(N_q, L_q) and (C, B) are index pairs for q, (N_p, L_p) and (B, A) are index pairs for p, and (C, A) are index pairs for $S = W(q, p) \cup \{q, p\}$.

We know that, for every x in $W^u(q)$, $f(\phi_{-T}) \leq f(q) = b$, so $W^u(q) \cap M_c$ is contained in N_q. Additionally, with $\epsilon > 0$ fixed, we get that N_q contracts to $W^u \cap M_c$ as $T \to \infty$. So, (N_q, L_q) contracts to

$$(W^u(q) \cap M_c, W^u(q) \cap f^{-1}(c)) \approx (D^k, \partial D^k).$$

As the inclusion $\partial D^n \hookrightarrow D^n$ is a cofibration, we then know that (N_q, L_q) is a regular index pair.

Similarly, N_p is a tubular neighborhood of $W^s(p) \cap M^c$. This contracts to $W^s(p) \cap M^c$ as $T \to \infty$ in a similar manner, so we know that

$$(N_p, L_p) \approx (D^{k-1} \times D^{m-k+1}, \partial D^{k-1} \times D^{m-k+1})$$

This is also a cofibration, so we get that (N_p, L_p) is a regular index pair.

Now, we note that N_p is a neighborhood of the stable closed disk $W^s(p) \cap M^c$, and note that the unstable sphere $W^u(q) \cap f^{-1}(c)$ is contained in L_q. We know from corollary 6.9 that $W(q,p)$ has finitely many components, so know that $N_p \cap S^u(q) = N_p \cap W^u(q) \cap f^{-1}(c)$ has finitely many components, each with a single point of intersection with $W(q,p)$ for each component. We denote these components $V_1, \ldots, V_j$ and the points of intersection $V_j \cap W(q,p)$ as x_j.

We have already noted that N_p is a tubular neighborhood of $W^s(p) \cap M^c$, so as $W^s(p) \cap M^c \approx D^{m-k+1}$ we get a diffeomorphism

$$\Psi_p : N_p \to D^{k-1} \times D^{m-k+1}.$$

with the following properties:

$$\Psi_p(L_p) = \partial D^{k-1} \times D^{m-k+1}$$

$$\Psi_p(N_p \cap W^s(p)) = \{0\} \times D^{M-k+1}$$

$$\Psi_p(V_j) = D^{k-1} \times \{\theta_j\} \text{ for } \theta_j \in \partial D^{m-k+1}$$

So, V_j is a $(k-1)$-dimensional manifold which is diffeomorphic to D^{k-1} via the maps $\Psi_{p,j} = \pi_1 \circ \Psi_p|_{v_j}$, whose composition takes V_j to D^{k-1} as follows:

$$V_j \xrightarrow{\Psi_p|_{V_j}} D^{k-1} \times D^{m-k+1} \xrightarrow{\pi_1} D^{k-1}$$

where the map π_1 simply projects a point (x, y) in $D^{k-1} \times D^{m-k+1}$ onto its first coordinate x in D^{k-1}. This map takes L_p into ∂D^{k-1} and therefore induces the isomorphism

$$H_{k-1}(N_p, L_p) \overset{(\Psi_{p,1})_*}{\approx} H_{k-1}(D^{k-1}, \partial D^{k-1}) \overset{(\Psi_j^{-1})_*}{\approx} H_{k-1}(V_j, \partial V_j) \approx \mathbb{Z}.$$

The orientation of $T_p^u M$ determines a generator α of $H_{k-1}(N_p, L_p) \approx \mathbb{Z}$, and a chain representing α also determines a generator for $H_{k-1}(B, A) \approx \mathbb{Z}$. Therefore, we can identify $H_{k-1}(N_p, L_p)$ with $H_{k-1}(B, A)$. This generator α is mapped onto a generator α_j of $H_{k-1}(V_j, \partial V_j)$ by $(\Psi_j^- 1)_* \circ (\Psi_{p,1})_*$. So, the homology class α_j

is determined by the orientation which $T_{x_j}V_j$ inherits from the orientation of T_p^uM by way of the isomorphism $T_{x_j}V_j \to T_p^uM$ which is determined by the flow ϕ_t. We get another orientation of $T_{x_j}V_j$ coming from $W^u(q)$, as

$$T_{x_j}V_j = (-\nabla f)(x_j)^{\perp} \cap T_{x_j}W^u(q) \subseteq T_{x_j}W^u(q)$$

where $-\nabla f(x_j)$ is the first vector in a positive basis. Let n_j be either +1 or -1 depending on whether or not these two orientations for T_{x_j} agree or not. So, n_j is the sign associated to the gradient flow line containing x_j.

We know that $S^u(q) = W^u(q) \cap f^{-1}(c)$ is the same region as $w^u(q) \cap L_q$. So, we consider the diagram below. The vertical maps arise from inclusions and the horizontal maps are connecting homomorphisms.

$$\begin{array}{ccc} H_k(W^u(q) \cap N_q, S^u(q)) & \xrightarrow{\delta_*} & H_{k-1}(S^u(q), \overline{(S^u(q) - \coprod_{j=1}^{n} V_j)}) \\ \downarrow \approx & & \downarrow s_* \\ H_k(C,B) & \xrightarrow{\quad\quad\delta_*\quad\quad} & H_{k-1}(B,A) \end{array}$$

Connecting homomorphisms are natural transformations, so this diagram is commutative. It is important to observe that

$$H_{k-1}(S^u(q), \overline{S^u(q) - \coprod_{j=1}^{n} V_j}) \approx \bigoplus_{j=1}^{n} H_{k-1}(S^u(q), \overline{S^u(q) - V_j}) \approx \mathbb{Z}.$$

Additionally, define $\delta_{*,j}$ to be the j^{th} component of the connecting homomorphism

$$\delta_* : H_k(W^u(q) \cap N_q, S^u(q)) \to H_{k-1}(S^u(q), \overline{S^u(q) - \coprod_{j=1}^{n} V_j})$$

under this identification. Therefore, for any $\beta \in H_k(W^u(q) \cap N_q, S^u(q))$ we have $\delta_*(\beta) = \delta_{*,1}(\beta) + \ldots + \delta_{*,\beta}$ and

$$s_*(\delta_*(\beta)) = \sum_{j=1}^{n} s_*(\delta_{*,j}(\beta)) \in H_{k-1}(B,A) \approx \mathbb{Z}$$

We know that N_q is a thickening of the unstable closed disk $W^u(q) \cap M_c = W^u(q) \cap N_q$, so we have the following relationship.

$$(N_q, L_q) \simeq (W^u(q) \cap N_q, W^u(q) \cap L_q) = (W^u(q) \cap N_q, S^u(q)).$$

Therefore, $H_k(W^u(q) \cap N_q, S^u(q)) \approx H_k(N_q, L_q)$. We use the following method to get a generator β for this group. We start with triangulations of the $(k-1)$-dimensional closed disks $V_j \subseteq S^u(q)$ for all $j = 1, \ldots, n$, and extend these triangulations to a triangulation of the k-dimensional manifold with boundary $W^u(q) \cap N_q \approx D^k$. This triangulation together with the orientation of $W^u(q)$ gives us the generator β. Similarly to α earlier, β also determines a generator for $H_k(C, B) \approx H_k(N_q, L_q)$, so we can identify

$$H_k(C, B) = H_k(N_q, L_q) = H_k(W^u(q) \cap N_q, S^u(q)).$$

By the excision theorem, we know that $H_{k-1}(S^u(q), \overline{S^u(q) - V_j}) \approx H_{k-1}(V_j, \partial V_j) \approx \mathbb{Z}$ for all $j = 1, \ldots, n$. In this isomorphism the imgae of the homology class $\delta_{*,j}(\beta)$ is represented by the original triangulation of V_j together with the original orientation, so it corresponds to $n_j \alpha_j$. We know that $\alpha_j \in H_{k-1}(V_j, \partial V_j)$ corresponds to α under the isomorphism $H_{k-1}(N_p, L_p) \approx H_{k-1}(V_j, \partial V_j)$, and so we get that

$$s_*(\delta_*(\beta)) = \sum_{j=1}^{n} n_j \alpha_j = n(q, p)\alpha \in H_{k-1}(B, A) = H_{k-1}(N_p, L_p)$$

Identifying β with q and α with p we then get the formula

$$\Delta_k(q) = \sum_{p \in Cr_{k-1}(f)} n(q, p)p = \partial_k(q).$$

Theorem 7.4 (Morse Homology Theorem) *The Morse-Smale-Witten chain complex $(C_*(f), \partial_*)$ is a chain complex, and its homology is isomorphic to the singular homology $H_n(M)$.*

Proof We begin with the commutative diagram below:

$$\begin{array}{ccccc}
C_{k+1}(f) & \xrightarrow{\partial_{k+1}} & C_k(f) & \xrightarrow{\partial_k} & C_{k-1}(f) \\
\updownarrow \approx & & \updownarrow \approx & & \updownarrow \approx \\
H_{k+1}(N_{k+1}, N_k) & \xrightarrow{\delta_*} & H_k(N_k, N_{k-1}) & \xrightarrow{\delta_*} & H_{k-1}(N_{k-1}, N_{k-2})
\end{array}$$

where $\delta*$ is the connecting homomorphism. We know that such a commutative diagram exists from the previous lemma, so we know that the Morse-Smale-Witten boundary operator has the property that $\partial_k \partial_{k+1} = 0$.

Following this, we examine the index pair (N_k, N_{k-1}). We know that $W(k, k) = Cr_k(f)$, and we use the constructions N_q and L_q from the example of gradient flows to produce a new index pair for $Cr_k(f)$

$$(\tilde{N}_k, \tilde{L}_k), \text{ setting } \tilde{N}_k = \bigcup_{q \in Cr_k(f)} N_q \text{ and } \tilde{L}_k = \bigcup_{q \in Cr_k(f)} L_q$$

By Theorem 7.2 we know that

$$N_k / N_{k-1} \simeq \tilde{N}_k / \tilde{L}_k \simeq \bigvee_{q \in Cr_k(f)} S^k_q$$

As N_k/N_{k-1} is homotopic to a wedge sum of k-spheres, we know that $H_j(N_k, N_{k-1}) = 0$ for $j \neq k$. We can then plug this into the exact sequence for (N_k, N_{k-1})

$$\ldots \to H_{j+1}(N_k, N_{k-1}) \to H_j(N_{k-1}) \to H_j(N_k) \to H_j(N_k, N_{k-1}), \to \ldots$$

to yield a short exact sequence of the form

$$0 \to H_j(N_{k-1}) \xrightarrow{i_*} H_j(N_k) \to 0$$

for all values of $j \neq k, k-1$. By the properties listed earlier for exact sequences we know then that the map i_* is an isomorphism for all $j \neq k, k-1$. So, we know that the inclusion map $H_j(N_k) \to H_j(M)$ is an isomorphism when $j < k$.

We now wish to show that $H_j(N_k)$ is trivial for $j > k$. To do so we induct on k. Starting at $k = 0$ we know that $H_j(N_0) = 0$ for $j > 0$ because $(N_0, \emptyset) = (N_0, N_{-1})$ is an index pair for $Cr_0(f)$.

For the inductive step we assume that $H_j(N_{k-1}) = 0$ for all $j > k-1$ and consider the exact sequence associated to (N_k, N_{k-1}) :

$$\ldots \to H_j(N_{k-1}) \to H_j(N_k) \to H_j(N_k, N_{k-1}) \to \ldots$$

We know that $H_j(N_k, N_{k-1}) = 0$ for $j > k$ from earlier, and we know that $H_j(N_{k-1}) = 0$ from the induction hypothesis, so from the properties of exact sequences we know that $H_j(N_{k-1})$ maps surjectively onto $H_j(N_k)$. However, $H_j(N_{k-1})$ is trivial, so $H_j(N_k)$ must be as well.

Now that we know that $H_j(N_k) = 0$ for all $j > k$ we have the following commitative diagram:

$$\begin{array}{ccccccc}
 & & & & & & H_{k-1}(N_{k-2} = 0 \\
 & & & & & & \downarrow \\
0 = H_k(N_{k-1}) & \longrightarrow & H_k(N_k) & \longrightarrow & H_k(N_k, N_{k-1}) & \xrightarrow{\delta_k} & H_{k-1}(N_{k-1}) \\
 & & & & & \searrow^{\delta_*} & \downarrow j_* \\
 & & & & & & H_{k-1}(N_{k-1}, N_{k-2})
\end{array}$$

The horizontal sequence is exact as it comes from the exact sequence for (N_k, N_{k-1}), and the vertical sequence is comes from the exact sequence for the pair (N_{k-1}, N_{k-2}). As the vertical sequence is exact, we know from the properties of exact sequences that j_* is injective.

We then know that the diagram below is commutative:

$$\begin{array}{ccccccc}
H_{k+1}(N_{k+1}, N_k) & \xrightarrow{\delta_{k+1}} & H_k(N_k) & \longrightarrow & H_k(N_{k+1}) & \longrightarrow & 0 \\
\updownarrow \approx & & \updownarrow \approx & & \updownarrow \approx & & \\
C_{k+1}(f) & \xrightarrow{\partial_{k+1}} & Ker(\partial_k) & \longrightarrow & H_k(M) & \longrightarrow & 0
\end{array}$$

From this diagram, we know that the bottom row is exact as the top row comes from the exact sequence for (N_{k+1}, N_k). So, we know by property (iv) of exact sequences that $H_k(M) \approx \text{Ker}(\partial_k)/\text{Im}(\partial_(k+1)$. As that quotient is the n^{th} homology group of $(C_*(f), \partial_*)$, we get an isomorphism between the homology groups of $(C_*(f), \partial_*)$ and the singular homology groups $H_*(M)$. This completes the proof of the Morse Homology Theorem.

8 Acknowledgements

I would like to thank my advisor, Professor Ralph Cohen, for his suggestion of Morse Homology as well as the help and guidance he gave me through this entire process. I would like to also thank Dan Berwick-Evans, whose classes first sparked my interest in topology. And I am of course forever grateful to my parents, whose encouragement and support have been crucial not only through this thesis, but throughout my entire collegiate career.

References

[1] Allen Hatcher. *Algebraic Topology.* Allen Hatcher, 2001.

[2] John Milnor *Morse Theory.* Princeton University Press, Princeton, N.J., 1963.

[3] Agustin Banyaga and David Hurtubise *Lectures on Morse Homology.* Springer Science and Business Media, Dordrecht, Netherlands, 2004.

[4] Dietmar Salamon *Connected simple systems and the Conley index of isolated irwariant sets.* Transactions of the American Mathematical Society 291 (1985).