MORSE THEORY AND HANDLE DECOMPOSITIONS

REMY BOHM

Abstract. We construct a handle decomposition of a smooth manifold from a Morse function on that manifold. We then use handle decompositions to prove Poincaré duality for smooth manifolds.

Contents

Introduction

The goal of this paper is to provide a relatively self-contained introduction to handle decompositions of manifolds. In particular, we will prove the theorem that a handle decomposition exists for every compact smooth manifold using techniques from Morse theory. Sections 1 through 3 are devoted to building up the necessary machinery to discuss the proof of this fact, and the proof itself is in Section 4. In Section 5, we discuss an application of handle decompositions to algebraic topology, namely Poincaré duality.

We assume familiarity with some real analysis, linear algebra, and multivariable calculus. Several theorems in this paper rely heavily on commonplace results in these other areas of mathematics, and so in many cases, references are provided in lieu of a proof. This choice was made in order to avoid getting bogged down in difficult proofs that are not directly related to geometric and differential topology, as well as to make this paper as accessible as possible.

Before we begin, we introduce a motivating example to consider through this paper. Imagine a torus, standing up on its end, behind a curtain, and what the torus would look like as the curtain is slowly lifted. The pictures in Figure 1 show the portions of the torus that are visible at different moments as the curtain is lifted. A closer look will reveal that during this unveiling process, the topology of the revealed portion changes; at first it is simply a disk, then a tube, then a torus

This paper was written during the second half of University of Chicago Mathematics Research Experience for Undergraduates program, in the late summer of 2019.

with one boundary component, and finally, a whole torus. This paper aims to provide an explanation for how the topology of the torus changes as it is unveiled, as well as how that informs other studies within topology.

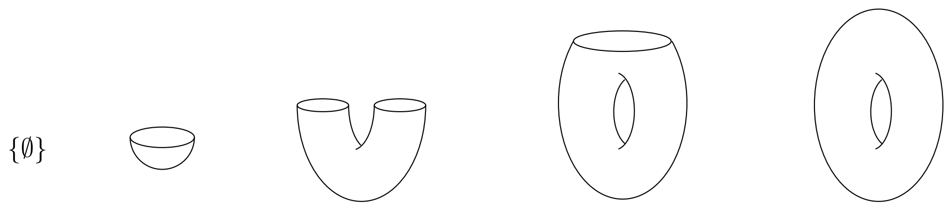

FIGURE 1. Unveiling a torus.

1. Smooth Manifolds and Handles

We begin by defining topological manifolds.

Definition 1.1. A *topological manifold* M is a second countable, Hausdorff topological space such that for all points p in M, there exists an open neighborhood N_p of p such that N_p is homeomorphic to the Euclidean open n-ball, $B^n := \{x \in \mathbb{R}^n \mid |x| < 1\}$.

It will be standard notation throughout this paper to use M^n to denote an n-dimensional manifold when the dimension of the manifold is relevant, after which the manifold may be simply referred to as M.

Definition 1.2. A *manifold with boundary* M is a second countable, Hausdorff topological space such that for all points p in M, there exists an open neighborhood N_p of p such that N_p is homeomorphic to either the Euclidean open n-ball $\{x \in \mathbb{R}^n : |x| < 1\}$ or the Euclidean open half-n-ball $\{x \in \mathbb{R}^n_+ : |x| < 1\}$.

Two points p and q in a manifold M may have neighborhoods that overlap, but are both homeomorphic to Euclidean balls B^n. We therefore introduce the idea of a transition map on the intersection.

Definition 1.3. Let M be a manifold, and let U, V be open subsets of M^n with homeomorphisms $P_U : U \to B^n$ and $P_V : V \to B^n$ such that $U \cap V \neq \emptyset$. The map $\phi\colon \mathbb{R}^n \to \mathbb{R}^n$ sending $P_U(U \cap V)$ to $U \cap V$ and then to $P_V(U \cap V)$ is called the *transition map* on $U \cap V$.

Transition maps are important in the study of manifolds, since they allow one to patch together local coordinate systems on manifolds to form globally defined structures.

We now proceed to definitions pertaining to smooth manifolds.

Definition 1.4. A function $f\colon \mathbb{R}^n \to \mathbb{R}$ is said to be C^∞ or *smooth* if it is infinitely differentiable.

Definition 1.5. A *smooth manifold* is a manifold M^n such that all of its transition maps are C^∞.

Definition 1.6. If M and N are smooth manifolds, then $f : U \to V$ is a *diffeomorphism* if it is a homeomorphism and if f and f^{-1} are differentiable.

For our purposes, as is common in the literature, we will take diffeomorphisms to be infinitely differentiable, to match our infinitely differentiable manifolds.

Definition 1.7. A *manifold with corners* M is a second countable, Hausdorff topological space such that for all points p in M, there exists an open neighborhood N_p of p such that N_p is homeomorphic to one of the following:
(i) the Euclidean open n-ball $\{x \in \mathbb{R}^n \colon |x| < 1\}$
(ii) the Euclidean open half-n-ball $\{x \in \mathbb{R}^{n-1} \times \mathbb{R}_+ \colon |x| < 1\}$
(iii) other subsets of the Euclidean n-ball where more than one coordinate is restricted positive $\{x \in \mathbb{R}^{n-m} \times \mathbb{R}^m_{+^m} \colon |x| < 1\}$

Note that manifolds with corners are homeomorphic to manifolds with boundaries, but not necessarily diffeomorphic to them.

Definition 1.8. For each point p in a smooth manifold M, let N_p be a coordinate neighborhood with a local homeomorphism $\phi \colon N_p \to U \subset \mathbb{R}^n$. Consider the equivalence classes $[\gamma]$ of curves $\gamma \colon [-1, 1] \to N_p$ passing through p such that $\gamma(0) = p$, under the equivalence relation $\gamma_1 \equiv \gamma_2$ if $\frac{\partial(\phi \circ \gamma_1)}{\partial t} = \frac{\partial(\phi \circ \gamma_2)}{\partial t}$ as maps $\phi \circ \gamma \colon [-1, 1] \to U \subset \mathbb{R}^n$. We say that an equivalence class of such local paths v is a *tangent vector* to M at p, and that the vector space spanned by all such v is the *tangent space* of M at p, denoted T_pM.

The union of all the tangent spaces over M is called the *tangent bundle* on M, denoted TM.

Note that all topological manifolds have tangent spaces, but they do not necessarily patch together in a way that will be useful for our purposes without a smooth structure already in place. The tangent bundle on smooth manifolds is a major object of study in geometric and differential topology, and it comes with a lot of interesting structure. However, we will only need its definition in this paper, where it appears in the definition of a vector field on a manifold. For more information, see [5].

The main theme of this paper is to understand smooth manifolds by breaking them up into smaller, topologically trivial chunks called *handles*.

Definition 1.9. An *n-dimensional k-handle* is a contractible smooth manifold $D^k \times D^{n-k}$.

We specify the construction of k-handles as $D^k \times D^{n-k}$ so that we can denote with k the region of h^k along which we "glue" it to another topological space of the same dimension n. We make the notion of "gluing" precise below.

Definition 1.10. Let X, Y be topological spaces, and let $K \subset X$ and $L \subset Y$ be subspaces such that there exists a homeomorphism $\phi \colon K \to L$. We obtain a new space, which we call *X glued to Y along ϕ* by taking ${}^{X \sqcup Y}/_{x \sim \phi(x)}$. We call ϕ the *attaching map*.

With the goal of gluing handles to other topological spaces in mind, we now define some useful parts of a k-handle.

Definition 1.11. There are five subsets of a k-handle which will be of interest to us. They are:
(i) the *attaching region*, defined to be $\partial D^k \times D^{n-k}$. In this paper, it is shown in bold in figures. Note that the attaching map of a k-handle is a homeomorphism of

the attaching region into a subset of the space being glued to.
(ii) the *attaching sphere*, denoted A^k: $\partial D^k \times \{0\}$,
(iii) the *core*, denoted C^k: $D^k \times \{0\}$,
(iv) the *belt sphere*, denoted B^k: $\{0\} \times \partial D^{n-k}$,
(v) the *co-core*, denoted K^k: $\{0\} \times D^{n-k}$.

Envisioning D^k and D^{n-k} both as products of the unit interval, we draw the following diagram of a handle in Figure 2.

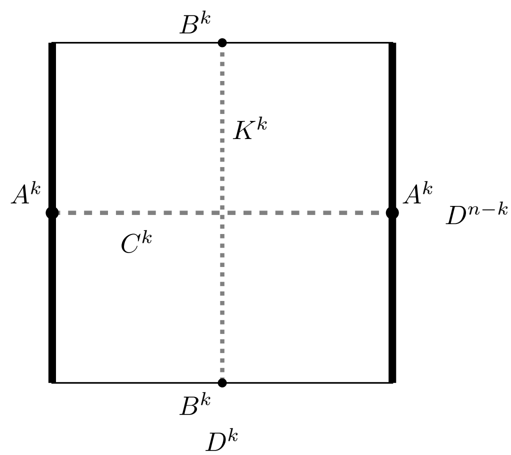

FIGURE 2. Anatomy of a k-handle, with attaching region shown in bold.

A k-handle in dimensions higher than 2 is impossible to draw, but to give a sense of how to interpret Figure 2, we show the attaching of a 2-dimensional 1-handle to a surface in Figure 3.

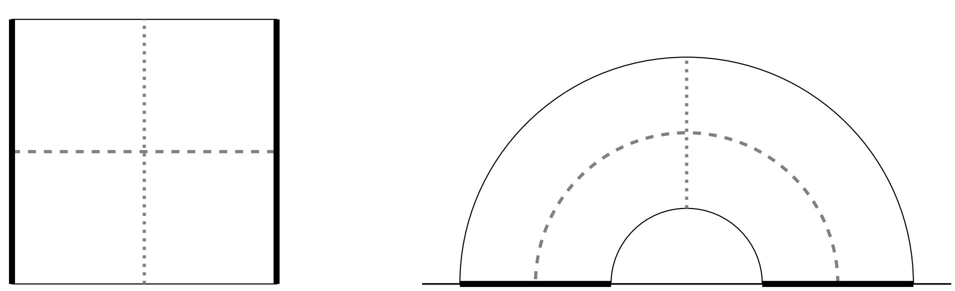

FIGURE 3. Gluing a k-handle.

Definition 1.12. A *handle decomposition* of a compact manifold M is a finite sequence of manifolds $W_0, \dots, W_l$ such that:
(i) $W_0 = \emptyset$,
(ii) W_l is diffeomorphic to M,
(iii) W_i is obtained from W_{i-1} by attaching a handle.
A *handlebody* is a compact manifold expressed as the union of handles.

Handle decompositions allow one to construct a manifold piece by piece, attaching one k-handle at a time. An example of a handle decomposition of a torus is shown below in Figure 4. The reader should take a moment to convince themselves that the final attachment of a 2-handle in the figure really does produce a torus.

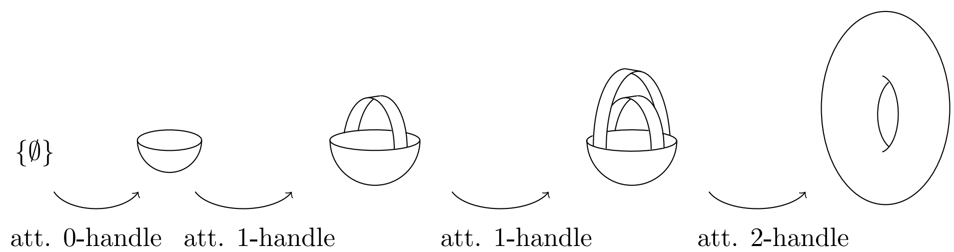

FIGURE 4. A handle decomposition of a torus.

It is also important to note that a handle decomposition of a given manifold is not unique. For instance, below are two decompositions of the unit sphere S^2.

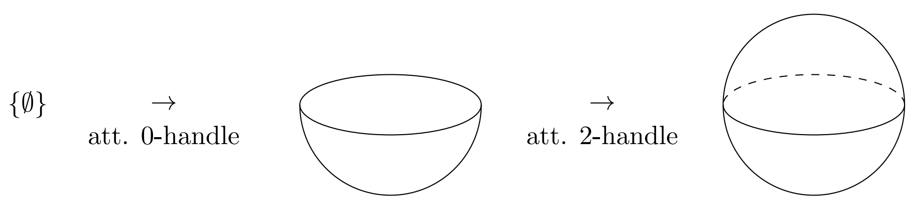

FIGURE 5. One decomposition of S^2.

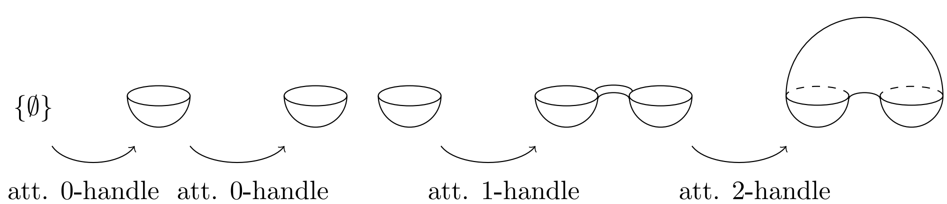

FIGURE 6. Another decomposition of S^2.

Even though any given manifold has many different handle decompositions, handle decompositions are nevertheless very useful tools for understanding the topology of manifolds, as they provide a “manual” of sorts for building a manifold piece by piece. In particular, all closed smooth manifolds admit handle decompositions, allowing many problems in topology to be studied purely in the context of handlebodies. The proof of this fact requires an understanding of some basic Morse theory, which we will now discuss.

2. Morse Functions

The basic idea of Morse theory is to understand manifolds by studying certain real-valued maps, called Morse functions, on them. We begin by introducing some properties of smooth functions on manifolds.

Definition 2.1. The *gradient vector field* of a function f is the vector field on the domain of f that takes the value $(\frac{\partial f}{\partial x_1}, \ldots, \frac{\partial f}{\partial x_n})$ at each point. We denote this vector field ∇f and its value at a point p as $\nabla f|_p$.

Definition 2.2. A *critical point* of a function $f\colon \mathbb{R}^n \to \mathbb{R}$ is a point $p \in \mathbb{R}^n$ such that $\nabla f|_p = 0$. Similarly, a *critical value* of f is a value $c \in \mathbb{R}$ such that $f(p) = c$ for p a critical point of f.

Definition 2.3. The *Hessian* of a function $f\colon M \to \mathbb{R}$, denoted $\mathcal{H}_f$, is the matrix of mixed second order partial derivatives of f:

$$\mathcal{H}_f = \begin{pmatrix} \frac{\partial^2 f}{\partial x_1^2} & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1} & \cdots & \frac{\partial^2 f}{\partial x_n^2} \end{pmatrix}$$

The Hessian evaluated at a point p is written $\mathcal{H}_f(p)$.

Definition 2.4. A critical point p of a continuous function f is called *degenerate* if $\det(\mathcal{H}_f(p)) = 0$.

We can now define a Morse function.

Definition 2.5. Given a smooth manifold M and a smooth function $f\colon M \to \mathbb{R}$, we say that f is *Morse* if f has no degenerate critical points on M.

The prototypical example of a Morse function on a manifold is a height function on a surface. That is, imagine your favorite closed surface floating in space above a plane. Then let your Morse function simply measure the height of level sets of the surface above the plane. A visual of this example is shown in Figure 7.

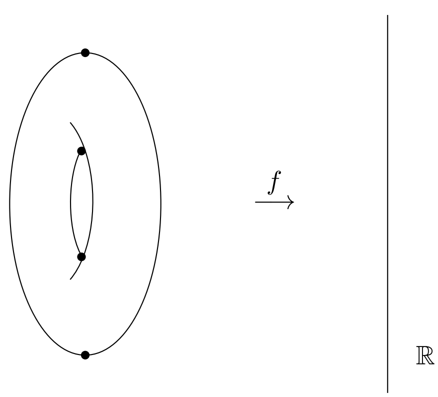

Figure 7. A Morse function that measures height on a torus, with critical points shown in bold.

The point of studying Morse functions is that if a function has only nondegenerate critical points, the function's local behavior in the neighborhood of its critical points can be further studied and classified, as is shown in the following definition.

Definition 2.6. Let M be a smooth manifold, $f\colon M \to \mathbb{R}$ be smooth, and p be a nondegenerate critical point of f. Then the *index* of f and p is defined to be the number of negative eigenvalues of the Hessian $\mathcal{H}_f$ evaluated at p.

Heuristically, the index of the Hessian tells us how many directions f is decreasing on. It will be the key to understanding how Morse functions relate to the actual attachment of handles to a manifold.

Proposition 2.7. *The nondegeneracy and the index of a function f at a critical point p do not depend on choice of local coordinates.*

Proof. We appeal to Sylvester's Law, which states that the number of negative eigenvalues of the Hessian is independent of the way it is diagonalized. Since diagonalization of a matrix corresponds to changing the basis of the source vector space so that the basis vectors are the eigenvectors of the matrix, this means that the number of negative eigenvalues of the Hessian is invariant under coordinate transformation. □

To make Morse functions effective tools in general, we must prove that they exist on all compact smooth manifolds. The proof of this fact is usually stated in the literature as the theorem that the set of Morse functions on a smooth, closed manifold M is dense in $C^\infty(M)$. In this treatment, we prove that one can always find a "very similar" function, or a (C^2, ε)-approximation, of any function such that the approximation function is Morse. Even with this modification, this is a rather involved proof requiring two fundamental lemmas dealing in real analysis. We therefore provide intuitive outlines for the proofs below, rather than fully rigorous ones. A more thorough treatment can be found in [6], from which these proofs are adapted.

We begin with the definition of a (C^2, ε)-approximation:

Definition 2.8. A function $f\colon K \subset \mathbb{R}^n \to \mathbb{R}$ is said to be a (C^2, ε)-*approximation* of a function $g\colon K \to \mathbb{R}$ if the following inequalities hold for all points $p \in K$:

$$|f(p) - g(p)| < \varepsilon$$

$$\left|\frac{\partial f}{\partial x_i}(p) - \frac{\partial g}{\partial x_i}(p)\right| < \varepsilon \qquad i = 1, \ldots, n$$

$$\left|\frac{\partial^2 f}{\partial x_i \partial x_j}(p) - \frac{\partial^2 g}{\partial x_i \partial x_j}(p)\right| < \varepsilon \qquad i, j = 1, \ldots, n$$

We can now move on to the requisite lemmas from analysis.

Lemma 2.9. *Let U be an open subset of $\mathbb{R}^n$ with coordinates $\{x_1, \ldots, x_n\}$ and let $f\colon U \to \mathbb{R}$ be a smooth function. Then there exist real numbers $\{a_i\}$ such that $f(x_1, \ldots, x_n) - (a_1 x_1 + \cdots + a_n x_n)$ is Morse on U. Moreover, for all $\varepsilon > 0$, each a_i can be chosen such that $|a_i| < \varepsilon$.*

Proof. The proof of this lemma is dependent on Sard's theorem, which states that the set of critical values of a continuous function $g\colon U \subset \mathbb{R}^n \to \mathbb{R}$ has measure 0 in $\mathbb{R}$. This result has very powerful applications in differential topology, but its proof is analytical, and so we refer the reader to Appendix C of [1] for a proof.

We begin with a function $f\colon U \to \mathbb{R}$ that may or may not have degenerate critical points in U. Let $h\colon U \to \mathbb{R}^n$ send $p \in U$ to $\nabla f(p)$. Then the matrix of partial derivatives of h is precisely the Hessian $\mathcal{H}_f$ at each point $p \in U$. Thus, critical points of h are precisely the degenerate critical points of f (points p where $\det(\mathcal{H}_f(p)) = 0$).

By Sard's Theorem, we can choose a point $a = (a_1, \dots, a_n) \in \mathbb{R}^n$ such that a is arbitrarily close to 0 but a is not a critical value of h.

We now claim that $\bar{f} := f - (a_1x_1 + \cdots + a_nx_n)$ is Morse on U. To see this, let p be a critical point of $\bar{f}$. Then $h(p) = a$ since $\frac{\partial \bar{f}}{\partial x_i}(p) = \frac{\partial f}{\partial x_i}(p) - a_i = 0$. But since a was chosen to not be a critical value of h, p must not be a critical point of h, and hence $\det(\mathcal{H}_f(p)) \neq 0$. Furthermore, $\mathcal{H}_f = \mathcal{H}_{\bar{f}}$ since f and $\bar{f}$ differ only by linear terms which vanish under second derivatives. Conclude $\det(\mathcal{H}_{\bar{f}}(p)) \neq 0$, and so p is nondegenerate. □

The upshot of this lemma is that we only ever need to modify a smooth function on an open subset of a manifold by some arbitrarily small linear term to make it Morse.

Lemma 2.10. *Let K be a compact subset of a manifold M, and suppose that $g\colon M \to \mathbb{R}$ has no degenerate critical points in K. Then for sufficiently small $\varepsilon > 0$, any (C^2, ε)-approximation f of g has no degenerate critical point in K.*

Proof. Let $\{U_i\}$ be a finite cover by open coordinate neighborhoods of K. For any function f to have no degenerate critical points in a given U_i, it must have no points where all of its partial derivatives and the determinant of its Hessian matrix with respect to the coordinates $\{x_1, \dots, x_n\}$ on U_i are all 0. Equivalently, it must satisfy the following inequality:

$$\left|\frac{\partial f}{\partial x_1}\right| + \cdots + \left|\frac{\partial f}{\partial x_n}\right| + |\det(\mathcal{H}_f)| > 0$$

But if f is a (C^2, ε)-approximation of g, which we know satisfies the above inequality, then we have that:

$$\left|\frac{\partial f}{\partial x_i} - \frac{\partial g}{\partial x_i}\right| < \varepsilon \qquad i = 1, \dots, n$$

$$\left|\frac{\partial^2 f}{\partial x_i \partial x_j} - \frac{\partial^2 g}{\partial x_i \partial x_j}\right| < \varepsilon \qquad i, j = 1, \dots, n$$

Therefore, for sufficiently small $\varepsilon > 0$, we have that f satisfies the desired inequality, and therefore has no degenerate critical points on U_i.

If we repeat this process on all the U_i, we get that f has no degenerate critical points on all of K as desired. □

Together, these lemmas allow us to perturb continuous functions on open subsets of a manifold to make them Morse, as well as ensure that these perturbations have minimal effect outside of the subsets on which they are defined. The content of the existence theorem, then, is stitching these *local* perturbations together to form a function that is *globally* Morse.

Theorem 2.11 (Existence of Morse functions)**.** *Let M be a compact manifold and $f_0\colon M \to \mathbb{R}$ be smooth. Then there exists a Morse function f on M that is an arbitrarily close approximation of f_0.*

Proof. Let $\{U_l\}_{1 \leq l \leq k}$ be a finite open cover of M such that for each U_l, there exists a compact subset K_l of U_l with $\{K_l\}$ a cover of M by compact sets. We begin with some smooth function f_0 on M that may have degenerate critical points. The idea of this proof is to inductively define functions f_l on M such that f_l is Morse on $\bigcup_{j=1}^{l} K_j$, denoted C_l for brevity. When $l = k$, we will have f_k Morse on $C_k = M$.

Our base case for induction will be to let $K_0 := \{\emptyset\}$ with f_0 our base function.

For our inductive hypothesis, suppose that we already have $f_{l-1}\colon M \to \mathbb{R}$ such that f_{l-1} is Morse on C_{l-1}. We want now to show that there exists a function f_l that is Morse on $C_{l-1} \cup K_l = C_l$.

To do this, let $\{x_1, \dots, x_n\}$ be local coordinates on U_l. Lemma 2.9 then tells us that there exist real numbers $\{a_i\}$ such that $f_{l-1}(x_1, \dots, x_n) - (a_1x_1 + \cdots + a_nx_n)$ is Morse on U_l.

We cannot simply set f_l to be this modified version of f_{l+1}, however, since the coordinates $\{x_1, \dots, x_n\}$ are local to U_l. To fix this, we introduce a smooth bump function on $h_l\colon U_l \to [0,1]$ such that $h_l = 1$ on an open neighborhood V_l of K_l contained in U_l, but $h_l = 0$ outside of a compact neighborhood $\bar{V}_l$. This is a lot of sets to keep track of, so a picture is shown below.

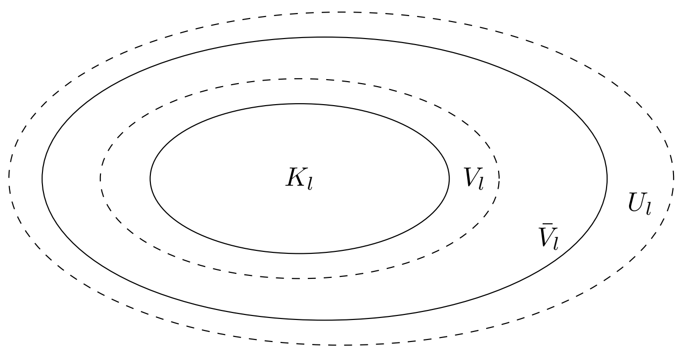

FIGURE 8. Relevant sets.

We can now define f_l on all of M as follows:

$$f_l(p) = \begin{cases} f_{l-1}(p) - h_l(p) \cdot (a_1x_1 + \cdots + a_nx_n) & p \in \bar{V}_l \\ f_{l-1}(p) & p \notin \bar{V}_l \end{cases}$$

All that remains is to check that f_l is a (C^2, ε)-approximation of f_{l-1}. Inside K_l, we can simply calculate the following inequalities:

$$|f_l - f_{l-1}| = |(a_1x_1 + \cdots + a_nx_n)|(h_l)$$

$$\left|\frac{\partial f_l}{\partial x_i} - \frac{\partial f_{l-1}}{\partial x_i}\right| = \left|a_ih_l + (a_1x_1 + \cdots + a_nx_n)\frac{\partial h_l}{\partial x_i}\right| \qquad i = 1, \dots, n$$

$$\left|\frac{\partial^2 f_l}{\partial x_i \partial x_j} - \frac{\partial^2 f_{l-1}}{\partial x_i \partial x_j}\right| = \left|a_i\frac{\partial h_l}{\partial x_j} + a_j\frac{\partial h_l}{\partial x_i} + (a_1x_1 + \cdots + a_nx_n)\frac{\partial^2 h_l}{\partial x_i \partial x_j}\right| \qquad i, j = 1, \dots, n$$

We know that h_l is bounded on $\bar{V}_l$, which is compact, and 0 elsewhere, and so $|\frac{\partial h_l}{\partial x_i}|$ and $|\frac{\partial^2 h_l}{\partial x_i \partial x_j}|$ must also be bounded on $\bar{V}_l$. Therefore for all $\varepsilon > 0$, by choosing each a_i small enough, we can make $|f_l - f_{l-1}|$, $|\frac{\partial f_l}{\partial x_i} - \frac{\partial f_{l-1}}{\partial x_i}|$, and $|\frac{\partial^2 f_l}{\partial x_i \partial x_j} - \frac{\partial^2 f_{l-1}}{\partial x_i \partial x_j}|$ all less than ε. Therefore, on K_l, f_l is a (C^2, ε)-approximation of f_{l-1}.

Outside of $\bar{V}_l$, $f_l = f_{l-1}$, but we know that some compact sets K_j must intersect K_l since all the K_j together cover M. We therefore must check that f_l is a (C^2, ε)-approximation on the overlaps $K_l \cap K_j$, $j \neq l$, that is to say, with respect to the coordinates on U_j for those U_j which intersect U_l. Fortunately, because M is smooth, all of its transition maps between overlapping open sets are C^∞, and so the composition of any transition map from U_l to U_j with f_l differs from f_{l-1} on U_j by a bounded term. Therefore, for all $\varepsilon > 0$, we can adjust the a_i to be even smaller such that f_l is a (C^2, ε)-approximation of f_{l-1} on the overlaps $K_l \cap K_j$. Outside of $\bar{V}_l$, $f_l = f_{l-1}$, and so we conclude that f_l is a (C^2, ε)-approximation of f_{l-1} on all of M.

By Lemma 2.10, we can now say that if f_{l-1} had no degenerate critical points in any K_j for $j \neq l$, then f_l must also have no degenerate critical points in any of those sets. Thus, after inducting on l until $l = k$, we have f_k Morse on M. $\square$

Now that we have familiarized ourselves with the idea and existence of Morse functions, we can proceed to the first result of Morse theory: the use of the index of a critical point to define new coordinate systems on neighborhoods of critical points. This observation is made rigorous in the Morse lemma below. Proving the Morse lemma is a key step in the construction of handlebody decompositions, as it allows us to reduce the behavior of a Morse function near a critical point to simply telling us how many coordinates f is increasing on, and how many it is decreasing on. If f is a height function, as in our examples, this is equivalent to walking along inside our manifold at p and noting how many directions one could walk in to go "down" and how many one could walk in to go "up".

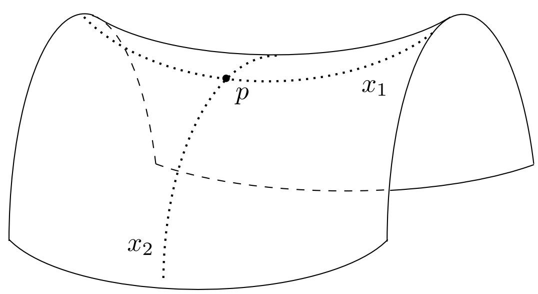

FIGURE 9. An index 1 critical point p on a surface, with local coordinates $\{x_1, x_2\}$.

This is shown in Figure 9. Note that in this figure, if one were walking along the surface at p, one could walk along the x_1 axis to walk "up" or along the x_2 axis to walk "down". Figure 9 therefore corresponds to an index 1 critical point.

Before we can get to the actual proof of the Morse lemma, we will need a lemma from multivariable calculus.

Lemma 2.12. *Let $f : \mathbb{R}^n \to \mathbb{R}$ be C^∞ on a convex neighborhood $U \subset \mathbb{R}^n$ containing the origin, and suppose that $f(0, \ldots, 0) = 0$. Then there exist C^∞ functions $\{g_i\}_{1 \leq i \leq n}$ defined on U such that:*

$$f = \sum_{i=1}^{n} x_i g_i$$

with:

$$g_i(0,\dots,0) = \frac{\partial f}{\partial x_i}(0,\dots,0)$$

We will omit the proof of this lemma for brevity, as it is an application of the multivariable chain rule. For a concise proof, see Part I, Chapter 2 of [7], or [4].

We are finally ready to prove the Morse lemma. There are many proofs of the Morse lemma out there, all in varying levels of detail. Here we provide a general idea of the proof that is palatable for those not in the mood to do lots of coordinate transformations, with a particular emphasis on aspects of the proof that are enlightening for its later use in handlebody decompositions. Of course, Milnor in [7] has a proof. For a more fleshed-out version, however, we recommend [4].

Theorem 2.13 (Morse lemma)**.** *Let f be a Morse function on a manifold M and p be a nondegenerate critical point of f. Then there exist local coordinates $\{x_1,\cdots,x_n\}$ on a neighborhood N_p such that on N_p, $f(x_1,\dots,x_n)$ has the form:*

$$f(x_1,\cdots,x_n) = -x_1^2 - \cdots - x_k^2 + x_{k+1}^2 + \cdots + x_n^2$$

Where k is the index of f at p, and p corresponds to the origin of this coordinate system.

Proof. We begin by letting $\{y_1,\dots,y_n\}$ be local coordinates on N_p and considering $f(y_1,\dots,y_n)$. To rearrange the coordinates such that f takes a quadratic form on a neighborhood of p, we would like to rewrite f in such a way that allows us to see the behavior of its second partial derivatives at $p=(0,\dots,0)$.

To do this, we use Lemma 2.12 twice; one iteration of the lemma applied to f defines functions $g_i\colon \mathbb{R}^n \to \mathbb{R}$ such that $f(y_1,\dots,y_n) = \sum_{i=1}^n y_i g_i(y_1,\dots,y_n)$ and $g_i(0) = \frac{\partial f}{\partial y_i}(0) = 0$, and the second iteration applies the lemma to each g_i to define functions $h_{i,j}\colon \mathbb{R}^n \to \mathbb{R}$ such that $h_{i,j}(0) = \frac{\partial g_i}{\partial x_j}(0)$. The details of this calculation are not enlightening, but it is necessary to obtain the form below for f:

$$f(y_1,\dots,y_n) = \sum_{i=1}^n \sum_{j=1}^n y_i y_j h_{i,j}(y_1,\dots,y_n)$$

Because all partial derivatives of f are assumed to exist on N_p, $h_{i,j} = h_{j,i}$. Furthermore, if we compute the 2^{nd} partial derivatives of f at $p=(0,\dots,0)$ in terms of $\{h_{i,j}\}$, we see that:

$$\frac{\partial^2 f}{\partial y_i \partial y_j}(p) = \begin{cases} 2h_{i,j}(p) & i = j \\ h_{i,j}(p) & i \neq j \end{cases}$$

Crucially, this observation means that $2h_{i,i}(p)$ is equal to the i^{th} diagonal entry of $\mathcal{H}_f(p)$ in terms of coordinates $\{y_1,\dots,y_n\}$, and that $h_{i,j}$ is equal to the $(i,j)^{\text{th}}$ entry for $i \neq j$. Our goal, therefore, is to diagonalize $\mathcal{H}_f(p)$ in order to get a quadratic form for f.

Now because second partial derivatives commute, the Hessian is a symmetric matrix with real entries and is therefore diagonalizable. There therefore exists a coordinate basis $\{\bar{x}_1,\dots,\bar{x}_n\}$ on N_p such that in this basis, $\mathcal{H}_f(p)$ is diagonal. If we let λ_i be the i^{th} diagonal entry of $\mathcal{H}_f(p)$, then we know that in this basis, f

takes the following form:

$$f(\bar{x}_1, \ldots, \bar{x}_n) = \sum_{i=0}^{n} \frac{\lambda_i}{2} \bar{x}_i^2$$

We perform one more coordinate transformation J on N_p to get rid of the $\frac{\lambda_i}{2}$ coefficients, being careful to keep their sign:

$$x_i = J(\bar{x}_i) := \operatorname{sign}(\lambda_i) \cdot \sqrt{\frac{|\lambda_i|}{2}} \bar{x}_i$$

Thus we obtain a quadratic form for f in the coordinates $\{x_1, \ldots, x_n\}$:

$$f(x_1, \ldots, x_n) = \operatorname{sign}(\lambda_1) x_1 + \cdots + \operatorname{sign}(\lambda_n) x_n$$

If we wish, we can now permute the coordinates to group them by the signs of their coefficients to achieve the desired form for f.

$$f(x_1, \cdots, x_n) = -x_1^2 - \cdots - x_k^2 + x_{k+1}^2 + \cdots + x_n^2$$

□

If f can be locally modeled by quadratics in each coordinate, then because quadratics have no critical points that cannot be isolated from their extrema by open neighborhoods (in fact, they have no other critical points at all), f must not have any critical points too near to any other. We therefore obtain the following result:

Corollary 2.14. *Nondegenerate critical points on any manifold can be isolated by open neighborhoods.*

This is an important corollary, as it allows us to consider nondegenerate critical points one by one.

3. Flows on Manifolds

We now take a brief detour into another topic in differential topology and geometry: flows. Morally, a flow is a group of diffeomorphisms associated to a smooth vector field X on a manifold M that sends a point on M in the direction of the vector associated to that point by X. We make these two definitions precise below:

Definition 3.1. A *smooth vector field* X on a manifold M is a smooth map from M into its tangent bundle TM that assigns to each p in M a vector v_p in T_pM.

Definition 3.2. A *flow* on a manifold M is a map $\Phi\colon \mathbb{R} \times M \to M$ with the following properties:
(i) $\varphi_t(p) := \Phi(t, p)$ is a diffeomorphism of M,
(ii) $\varphi_0(p)$ is the identity diffeomorphism of M,
(iii) For all $s, t \in \mathbb{R}$, $\varphi_{s+t} = \varphi_s \circ \varphi_t$.

The *trajectory* of a point p in M of a flow is a map $\psi_p\colon \mathbb{R} \to M$ that sends $t \in \mathbb{R}$ to $\varphi_t(p)$ such that $\psi_p(0) = \varphi_0(p) = p$.

Note that the definition of a flow is equivalent to that of a smooth group action of $\mathbb{R}$ on M by diffeomorphisms.

One can think of a flow as being generated by a smooth vector field X on M by defining X such that $\frac{\partial \Phi(t,p)}{\partial t} = X \circ \varphi_t(p)$. In this framework, the trajectory of a point is equivalent to an integral curve of the vector field.

For our purposes, we will want to understand when this generation of flows by smooth vector fields is unique. It turns out that the restriction we put on smooth vector fields on M to have them generate unique flows on M is the property of being *compactly supported*, or taking on a value of 0 outside of some compact subset of M.

The following proof of this fact is adapted from [7], but it can be found in a variety of different texts in differential topology in varying forms. A particularly thorough treatment can be found in Chapter 12 of [5].

Theorem 3.3. *Let X be a smooth vector field on a manifold M and suppose that X is compactly supported on $K \subset M$. Then X generates a unique flow Φ on M.*

Proof. Given X, consider the set of differential equations on t parametrized by points p in K given by

$$\frac{\partial \Phi(t,p)}{\partial t} = X \circ \varphi_t(p)$$

with the initial condition $\varphi_0(p) = p$ at all points p in K.

Existence and uniqueness of Φ both, then, are a consequence of the common result in ordinary differential equations that guarantees that an ordinary differential equation with an initial condition has a unique, smooth solution that depends smoothly on the initial condition. So for any given point $p \in M$, there exists an open neighborhood N_p of p with a unique $\Phi(t,p)$ defined on it that satisfies the above differential equation for $t \in (-\varepsilon, \varepsilon)$. Furthermore, because X was smooth, for a given set of solutions $\{\varphi_{t,\alpha}\}$ such that each $\varphi_{t,\alpha}$ is defined on the same open set $(-\varepsilon', \varepsilon') \subset \mathbb{R}$, we can patch these local solutions on M together to define φ_t globally on M within $(-\varepsilon', \varepsilon')$.

It now remains to show that we can find such an open set in $\mathbb{R}$ that all φ_t are defined on. To do this, note that because K is compact, we can restrict a cover of K by neighborhoods N_p of individual points p_i indexed by $i \in I \subset \mathbb{N}$ to finitely many open neighborhoods $\{N_{p_i}\}$. Let ε_0 be the smallest of the ε_i corresponding to these neighborhoods N_{p_i}. Note that ε_0 is well-defined and nonzero because there are only finitely many neighborhoods N_{p_i}.

If $\varphi_t(q) = q$ for all $q \notin K$ and for all $t \in \mathbb{R}$, then we know that φ_t is defined for all of M for $t \in (-\varepsilon_0, \varepsilon_0)$. Furthermore, as each φ_t is generated by X, $\varphi_t \circ \varphi_s = \varphi_{t+s}$ for all $t, s \leq \varepsilon_0$ and so we can simply iterate φ_t on itself to generate $\varphi_{t'}$ for all t' such that $|t'| > \varepsilon_0$. Thus we obtain $\Phi(t,p)$ defined globally on M and for all $t \in \mathbb{R}$. □

This theorem will be essential in showing the existence of a diffeomorphism between submanifolds of M whose boundaries' image under f do not include a critical point.

4. From Morse Functions to Handle Decompositions

We have now constructed enough machinery to obtain a handlebody decomposition of any smooth, compact manifold. This theorem will follow from the following two intermediate results about the local behavior of a Morse function on a manifold, Theorem 4.2 and Theorem 4.4.

Definition 4.1. The *sublevel set* of a Morse function on M at a point $a \in \mathbb{R}$ is $\{p \in M \mid f(p) \leq a\}$. It is denoted M_a.

Theorem 4.2. *Let M be a compact manifold and $f : M \to \mathbb{R}$ be a Morse function. Suppose that $a, b \in \mathbb{R}$ are such that $f^{-1}[a, b]$ is nonempty. If $f^{-1}[a, b]$ does not contain a critical point of f, then M_a is diffeomorphic to M_b.*

Proof. The idea of this proof is to find a vector field that we can associate to f and use that vector field to generate a flow that will give us a diffeomorphism from M_a to M_b.

Choose a Riemannian metric on M with inner product $\langle\ ,\ \rangle$ and let $||\ ||$ denote the induced norm on M. Note that we can choose such a metric because M is assumed to be smooth. For more information, see Chapter 8 of [5].

To construct a satisfactory flow, we need our vector field to have only unit vectors on $f^{-1}[a, b]$. To do this, define a new function $g : M \to \mathbb{R}$ such that $g = \frac{1}{||\nabla f||^2}$ on $f^{-1}[a, b]$ and g vanishes outside of a compact neighborhood of $f^{-1}[a, b]$. Note that the existence of such a g is a consequence of the existence of bump functions on manifolds; see [5].

We then define a vector field X on M by:

$$X(p) := g(p) \cdot \nabla f(p)$$

So on $f^{-1}[a, b]$, X takes the form:

$$X(p) = \frac{\nabla f}{||\nabla f||^2}(p)$$

However, X is compactly supported, since g was defined to vanish outside of a compact neighborhood of $f^{-1}[a, b]$. This allows us to apply Theorem 3.3 to generate a flow Φ on M.

We want to show that Φ contains a diffeomorphism that sends M_a to M_b. To do this, let $F\colon \mathbb{R} \to \mathbb{R}$ be defined by $F(t) = f \circ \varphi_t(p)$. We now calculate the derivative of F with respect to t.

$$\frac{\partial F}{\partial t} = \left\langle \frac{\partial \Phi}{\partial t}, \nabla f \right\rangle = \langle X, \nabla f \rangle = 1$$

This tells us that F, as a function from $\mathbb{R}$ to $\mathbb{R}$, is linear with slope 1. Therefore, φ_0 is the identity diffeomorphism on M_a, and $\varphi_{b-a}(M_a) = M_b$.

We have found a diffeomorphism from M_a to M_b, thus completing the proof. $\square$

We now deal with the case where $f^{-1}[a, b]$ contains a critical point. We will need to use the following lemma.

Lemma 4.3. *Let M be a smooth manifold with corners. Then there exists a smooth manifold M' without corners that is homeomorphic to M and diffeomorphic to M outside of a neighborhood of the corner points. Furthermore, M' is unique up to diffeomorphism.*

We refer the reader to [8] for a proof.

Theorem 4.4. *Let M, f, and a, b be as in Theorem 4.2. If $f^{-1}[a, b]$ contains one critical point of f with index k, then M_b is diffeomorphic to the union of M_a with a k-handle.*

Proof. Let p be the critical point in $f^{-1}(a, b)$, denote its image under f as c, and let k be the index of p. Because of Corollary 2.14, we can assume, up to adjusting a and b to decrease $|a - c|$ and $|b - c|$, that there exists a coordinate neighborhood N_p that intersects both preimages $f^{-1}(a)$ and $f^{-1}(b)$ and that contains no other critical points of f. By the Morse lemma, we can alter the coordinates on N_p such that f takes the form $f(x_1, \ldots, x_n) = -x_1^2 - \cdots - x_k^2 + x_{k+1}^2 + \cdots + x_n^2$. With respect to these coordinates, we are able to draw a contour map of f on N_p as in Figure 10.

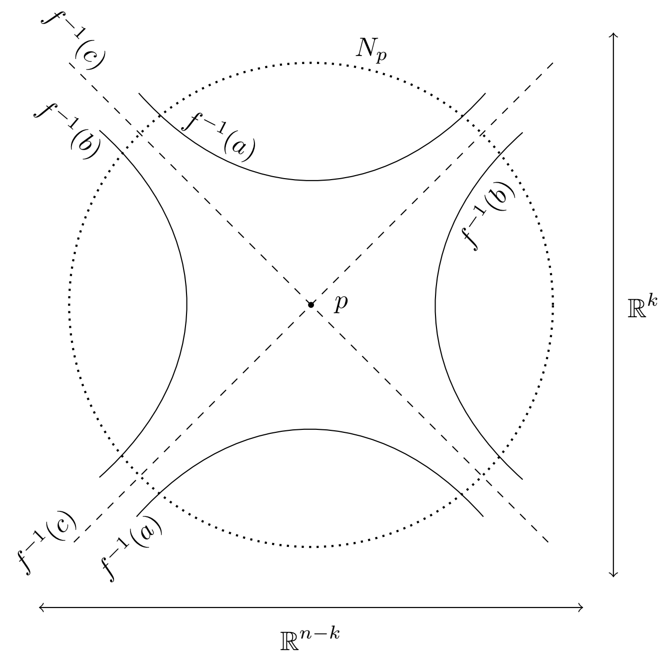

FIGURE 10. A neighborhood of a nondegenerate critical point p.

The Morse lemma allows us to split the n dimensions of our manifold into two subspaces, one of dimension k on which f takes values less than c, and the other of dimension $n-k$ on which f takes values greater than c. Note that the level set $f^{-1}(a)$ in N_p intersects the coordinate axes of $\{x_1, \ldots, x_k\}$, and that it does not intersect the coordinate axes of $\{x_{k+1}, \ldots, x_n\}$. To see this, imagine standing at p and noting that walking along any of the axes x_i for $i \leq k$ will lead you down towards $f^{-1}(a)$, while walking along any of the axes x_i for $i > k$ will lead you "up". The same logic allows us to represent $f^{-1}(b)$ as intersecting the axes $\{x_{k+1}, \ldots, x_n\}$ and avoiding the others.

We will be interested in the intersections of M_a and M_b with N_p, since the local behavior of f at p occurs within N_p. The intersections $N_p \cap M_b$ and $(N_p \cap M_b) - M_a$ are shown below.

Let H be the subset $(N_p \cap M_b) - M_a$. In fact, H is, as a topological space, a $k-$handle, just with an unfamiliar shape. Recalling that our goal is to show a diffeomorphism between $M_a \cup H$ and M_b, we apply Lemma 4.3 to round out the corners of H where it does not meet M_a. This is shown in Figure 13. H then is diffeomorphic, as a manifold with corners, to $D^k \times D^{n-k}$.

It now remains to show that $M_a \cup H$ is diffeomorphic to M_b. Rather than construct such a map explicitly, we appeal to Theorem 4.2. The existence of Morse functions (Theorem 2.11) guarantees the existence of another Morse function g such

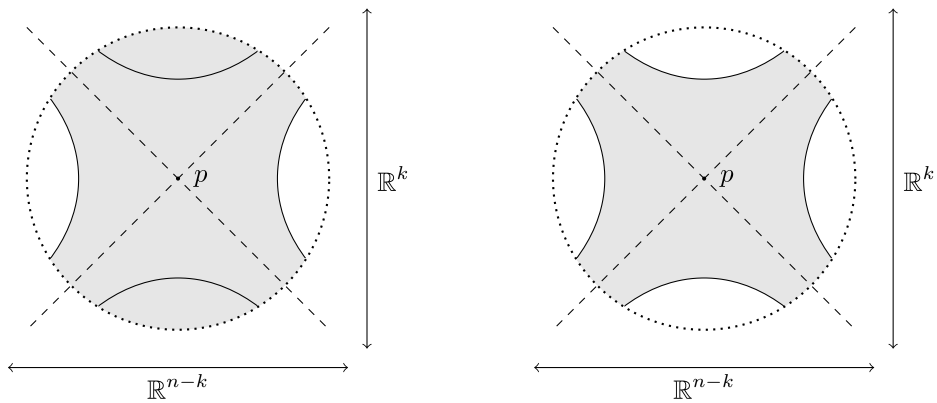

FIGURE 11. Left: $N_p \cap M_b$. Right: $(N_p \cap M_b) - M_a$.

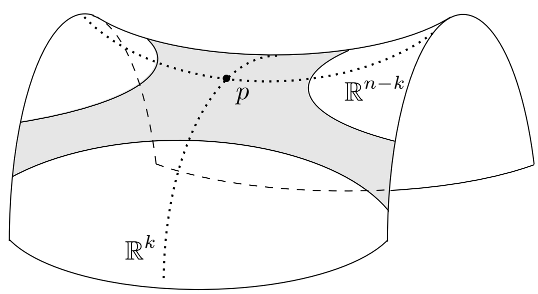

FIGURE 12. H on a torus

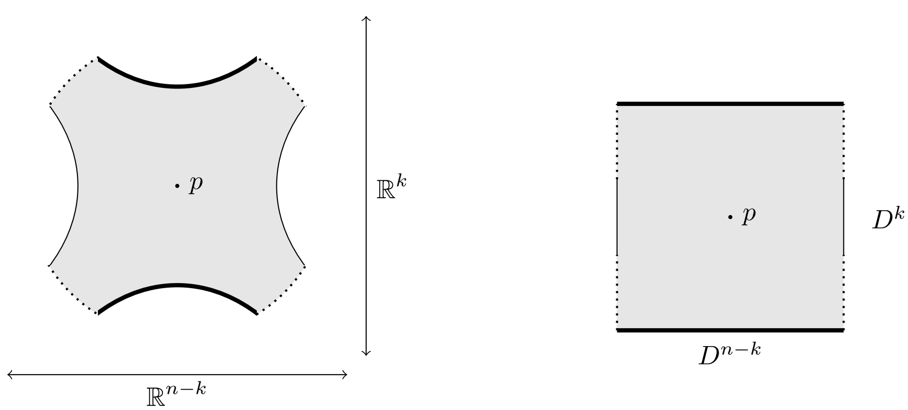

FIGURE 13. Smoothing corners of H, with attaching region of H to M_a shown in bold.

that for some $a' \in \mathbb{R}$, $M_{g=a'} = M_a \cup H$ and $M_{g=b} = M_b$. In particular, g can be chosen such that M does not contain a critical point of g in $g^{-1}[a', b]$ (since f does not by assumption). So by Theorem 4.2, we have that $M_{g=a'} = M_a \cup H$ is diffeomorphic to $M_{g=b} = M_b$. □

We need one more lemma in order to prove the existence of handle decompositions for all compact smooth manifolds.

Lemma 4.5. *Let M be a smooth manifold with f a Morse function on M. Then if p and q are both critical points of f such that $f(p) = f(q)$, then there exists a smooth manifold M' that is diffeomorphic to M such that $f(p) \neq f(q)$.*

We refer the reader to [7] for a proof.

Theorem 4.6 (Existence of handle decompositions)**.** *There exists a handle decomposition for every compact smooth manifold.*

Proof. Let M be a compact smooth manifold. By Theorem 2.11, there exists a Morse function f on M. Because M is compact, we know that there exist $A, B \in \mathbb{R}$ such that $M_A = \{\emptyset\}$ and $M_B = M$. Compactness also guarantees us that there are only finitely many critical points p_i of f, and Lemma 4.5 guarantees that we can adjust M by diffeomorphism such that $f(p_i) \neq f(p_j)$ for all $i \neq j$. Our goal now is to use these critical points to build a handle decomposition of M.

Let L be the total number of critical points on M. Index each critical point p_i such that if $i < j$, then $f(p_i) < f(p_j)$. This way, p_1 is the lowest critical point, and p_L is the highest. For each pair p_i, p_{i+1} for $i = 1, \ldots, L-1$, let $a_i = \frac{f(p_i)+f(p_{i+1})}{2}$. Note that a_i is defined such that the only critical points that the sublevel set M_{a_i} contains are $p_1, \ldots, p_i$. For notation, set $M_0 = \{\emptyset\}$, and set $M_L = M$.

If we compare M_{a_i} with $M_{a_{i+1}}$, we see that $f^{-1}[a_i, a_{i+1}]$ contains exactly one critical point, p_{i+1}. By Theorem 4.4, we have that $M_{a_{i+1}}$ is diffeomorphic to M_{a_i} with the attachment of a k-handle, where k is the index of a_{i+1}. Furthermore, by Theorem 4.2, we have that the topology of two sublevel sets M_{a_i} and M_b for $b > a_i$ only differs when $b > f(p_{i+1})$. Therefore, the sequence $\{M_0, M_1, \ldots, M_{L-1}, M_L\}$ is a handle decomposition for M. □

5. Handlebodies in Algebraic Topology

Our goal for this section is to illustrate an application of handlebodies to another area of the study of manifolds, namely Poincaré duality. Note that Poincaré duality deals in the homology and cohomology of manifolds, and therefore is not sensitive to differential structure, only the homotopy type of a manifold. To accomodate this, our first step in proving Poincaré duality is to come up with a way to view a handlebody as a CW complex.

Proposition 5.1. *A handlebody decomposition of a manifold M defines a CW complex X that is homotopy equivalent to M.*

Proof. The idea here is to recognize that up to homotopy equivalence, the structure of a k-handle can be reduced to that of a k-cell where the attaching sphere of the k-handle becomes the boundary of the k-cell glued on. This is made precise by noting that the key information contained in the gluing map of a k-handle can be reduced to the dimension of the attaching sphere A^k and its placement on the manifold.

Recall from Chapter 1 that the attaching sphere of a k-handle was defined to be $\partial D^k \times \{0\}$, which is homeomorphic to S^{k-1}. It is the boundary of the core $D^k \times \{0\}$. Since h^k can be viewed as the direct product of its core with D^{n-k}, which is contractible, h^k deformation retracts onto its core.

This deformation retraction shrinks the attaching region to the attaching sphere A^k. However, this does not change the topology of the attaching map, since the attaching region is just $A^k \times D^{n-k}$. Thus, deformation retracting h^k to its core induces a deformation retraction on the space h^k is attached to, but does not change the topology of either. We therefore say that the k-cell e^k associated to h^k is the core of h^k.

Note that the attaching map of the handle h^k restricted to A^k is now precisely the attaching map of the cell e^k. □

A visual of the handle-to-cell homotopy is shown in Figure 14.

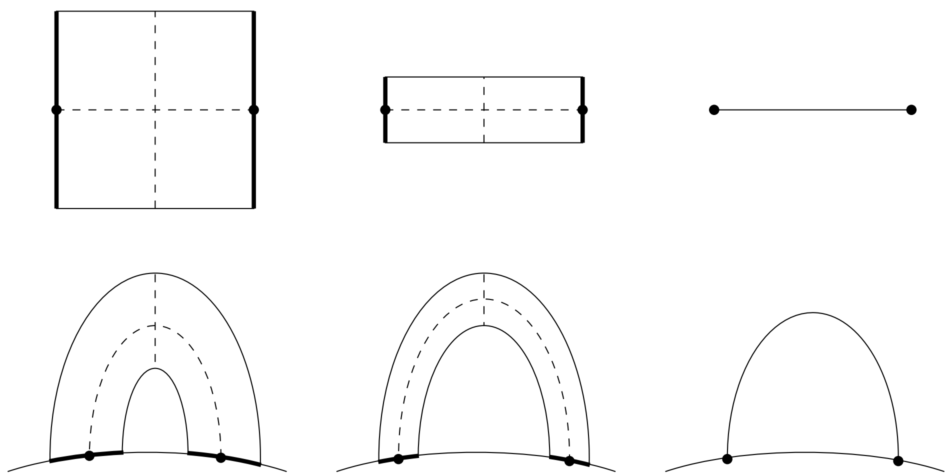

FIGURE 14. Shrinking a k-handle to its core.

Armed with a cellular description of M, we can now proceed to construct its cellular homology and cohomology. Below, we provide the reader with a brief overview of these theories. For more information, we recommend a combination of [3] for geometric intuition around homology and [2] for cellular homology specifically. It is important to note that this is not intended to be a sufficiently thorough introduction without prior familiarity with homology and cohomology.

Definition 5.2. The *cellular chain complex* of a CW complex X is the chain complex associated to a space X where the $k-$dimensional chain groups $C_k(X)$ are defined to be the free abelian groups generated by the set of k-cells in X.

In a cellular chain complex, the boundary map sends a k-cell e^k to the formal sum of the (k-1)-cells in the image of the attaching map of e^k. This is formalized in the *cellular boundary formula*:

$$\partial_k(e_i^k) = \sum_{j \in J} d_{ij} e_j^{k-1}$$

where d_{ij} is the degree of the composition of the following three maps: the attaching map of e_i^k sending $\partial e_i^k \cong S_i^{k-1}$ into X^{k-1}, the quotient map sending X^{k-1} to

${}^{X^{k-1}}/{}_{X^{k-2}}$, and the collapsing map sending all the copies of S^{k-1} in ${}^{X^{k-1}}/{}_{X^{k-2}}$ to a single S_j^{k-1}. The composition of all of these maps defines a single map from S_i^{k-1} to S_j^{k-1}, the degree of which is well-defined.

More information on cellular homology, as well as a proof of the cellular boundary formula from the typical construction of cellular homology, can be found in [2].

To each chain complex we can associate the dual complex, called the cochain complex. For the cellular case, it is defined as follows:

Definition 5.3. The *cochain complex* associated to a cellular chain complex is the chain complex in which the $k-$dimensional chains $C^k(X)$ are defined to be $C^k(X) := \mathrm{Hom}(C_k(X), \mathbb{Z})$.

For a given $f \in C^k(X)$, the composition $f \circ \partial_{k+1}$ defines a homomorphism from $C_{k+1}(X)$ to $\mathbb{Z}$, which is precisely an element of $\mathrm{Hom}(C_{k+1}(X), \mathbb{Z})$, or $C^{k+1}(X)$. We therefore define the *cellular coboundary map* $\delta^k \colon C^k(X) \to C^{k+1}(X)$ as follows:

$$\delta^k(f) = f \circ \partial_{k+1}$$

The chain and cochain complexes of a CW decomposition X of a compact, orientable smooth manifold M can be related using Morse functions on M! The following proposition constructs the foundation for this relationship.

Proposition 5.4. *If $f \colon M^n \to \mathbb{R}$ is a Morse function, then $-f \colon M^n \to \mathbb{R}$ is also a Morse function with the same critical points as f.*

Furthermore, if p is an index k critical point of f, then p is an index $n-k$ critical point of $-f$.

Proof. If f is smooth, then $-f$ is smooth, as it is the composition of f with the map $\mathbb{R} \to \mathbb{R}$ sending x to $-x$.

Let p be a critical point of f. Then on a neighborhood of p with local coordinates $\{x_i\}$, $\frac{\partial f}{\partial x_i}(p) = 0$. Hence $\frac{\partial(-f)}{\partial x_i}(p) = -\frac{\partial f}{\partial x_i}(p) = 0$, and so p is a critical point of $-f$. Furthermore, by our assumption that f is Morse, $\det(\mathcal{H}_f(p)) \neq 0$. But $\mathcal{H}_{-f}(p) = -\mathcal{H}_f(p)$, and so $\det(\mathcal{H}_{-f}(p)) = -\det(\mathcal{H}_f(p)) \neq 0$. So p is a nondegenerate critical point of $-f$. This completes the proof that $-f$ is Morse.

As for the index of a critical point p of $-f$, note that multiplication by -1 of a matrix $\mathcal{H}_f(p)$ flips the sign of all of its eigenvalues. Therefore, since $\mathcal{H}_{-f}(p) = -\mathcal{H}_f(p)$, the number of negative eigenvalues of the $n \times n$ matrix $\mathcal{H}_{-f}(p)$ is $n-$ the number of negative eigenvalues of $\mathcal{H}_f(p)$. □

This lemma leads us to the following key corollary, which forms the foundation for Poincaré duality.

Corollary 5.5. *Up to diffeomorphism, the k-handles of the decomposition associated to f are equal as submanifolds to the $(n-k)$-handles of the decomposition associated to $-f$.*

We are now ready to prove Poincaré duality.

Theorem 5.6 (Poincaré duality)**.** *Let M^n be a closed, orientable manifold. Then for all $0 \leq k \leq n$, $H^{n-k}(M) \cong H_k(M)$.*

Proof. To begin, let f be a Morse function on M. Then by Proposition 5.4, $-f$ is a Morse function on M.

Let $\{M_r\}$ be the handle decomposition of M obtained from f, and let X be the CW complex obtained from $\{M_r\}$. Let $\{W_s\}$ be the handle decomposition of M obtained from $-f$, and let Y be the CW complex obtained from $\{W_s\}$. Note that M is homotopy equivalent to both X and Y.

We will refer to $n-k$-handles of $\{M_r\}$ as x_i^{n-k}, and k-handles of $\{W_s\}$ as y_i^k. Note that we may index these both with the same variable because Corollary 5.5 guarantees us a map from $n-k$-handles of $\{M_r\}$ to k-handles of $\{W_s\}$. In fact, they are equal as submanifolds. Therefore, for a given handle H_i (a $n-k$-handle when viewed in $\{M_r\}$ and a k-handle when viewed in $\{W_s\}$), we define new variables for its core and co-core, since those terms are no longer well defined when switching between decompositions. Define α_i to be the core of H_i seen as y_i^k, or equivalently, the co-core of H_i seen as x_i^{n-k}. Similarly, define β to be the co-core of H_i seen as y_i^k, or the core of H_i seen as x_i^{n-k}. This can be seen in Figure 15.

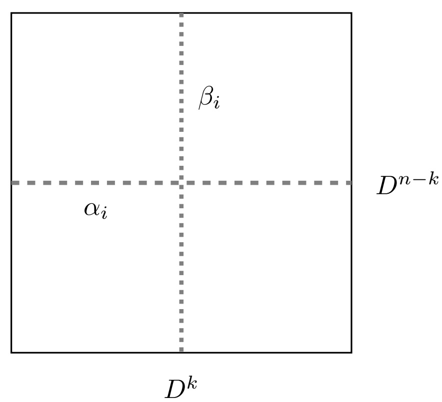

FIGURE 15. Anatomy of H_i. Note that when H_i is viewed as a k-handle y^k, its core is α_i, but when it is viewed as a $(n-k)$-handle, its core is β_i.

In Proposition 5.1, we saw that to consider a handlebody as a CW complex, each handle was shrunk to its core. We can therefore say that for every handle H_i in $\{M_r\}$ and $\{W_s\}$, β_i, its core as a $n-k$-handle in $\{M_r\}$ is a generator of $C_{n-k}(x)$. Similarly, α_i, its core as a k-handle in $\{W_s\}$, is a generator of $C_k(Y)$. We will use this duality to show that the cellular cochain complex of X is isomorphic to the cellular chain complex of Y.

We begin by defining a homomorphism $\psi_k\colon C_k(Y) \to C^{n-k}(X)$ given by:

$$\psi_k(\alpha_i) = c_i^{n-k}$$

where c_i^{n-k} denote the element of $C^{n-k}(X)$ that maps β_i to $1 \in \mathbb{Z}$ and all other β_j to 0 for $j \neq i$. Note that because $\{\beta_i\}$ generate $C_{n-k}(X)$, the maps c_i which send β_i to 1 and all other β_j to 0 generate $C^{n-k}(X)$. Furthermore, $\text{rank}(C_{n-k}(X)) = \text{rank}(C^{n-k}(X))$ as free abelian groups.

The equivalence of ranks of these chain and co-chain groups guarantees that ψ_k is bijective for all k. Every handle H_i has exactly one core/co-core α_i and one co-core/core β_i. Furthermore, for every β_i in $C_{n-k}(X)$, there is exactly one c_i in $C^{n-k}(X)$.

To extend the maps ψ_k to an isomorphism of chain complexes, we must show that they commute with the boundary maps of each complex. Specifically, we want to show that the following diagram commutes:

$$\begin{array}{ccc} C_k(Y) & \xrightarrow{{}_Y\partial_k} & C_{k-1}(Y) \\ \downarrow{\scriptstyle \psi_k} & & \downarrow{\scriptstyle \psi_{k-1}} \\ C^{n-k}(X) & \xrightarrow[{}_X\,\delta^{n-k}]{} & C^{n-k+1}(X) \end{array}$$

To avoid confusion, we take a moment to give names to the elements of these chain groups:

Elements of $C_k(Y)$ are denoted α_i.
Elements of $C_{k-1}(Y)$ are denoted $\bar{\alpha}_j$.
Elements of $C_{n-k}(X)$ are denoted β_i.
Elements of $C_{n-k+1}(X)$ are denoted $\bar{\beta}_j$.
Elements of $C^{n-k}(X)$ are denoted c_i.
Elements of $C^{n-k+1}(X)$ are denoted $\bar{c}_j$.

Consider first $\psi_{k-1} \circ {}_Y\partial_k$.

The map ${}_Y\partial_k$ sends the core of a k-handle y_i^k to a formal sum of cores of $(k-1)$-handles. The image of α_i, the core of y_i^k, under ${}_Y\partial_k$ is then:

$$\sum_{j=1}^{\operatorname{rank}(C_{k-1}(Y))} A_{i,j}\bar{\alpha}_j$$

Where $\bar{\alpha}_j$ is the core of y_j^{k-1}, and thus a generator of $C_{k-1}(Y)$.

Formally, $A_{i,j}$ is the degree of the attaching map of y_i^k. Geometrically, the coefficients $A_{i,j}$ represent the number of times the attaching region of y_i^k "wraps around" the core of each y_j^{k-1}, with sign determined by the orientation of the cores. However, because the core and co-core of any handle intersect transversely exactly once, the signed number of times y_i^k "wraps around" each y_j^{k-1} is precisely the signed transverse intersection number of α_i with the co-cores of each y_j^{k-1}. Note that these co-cores, which we denote $\bar{\beta}_j$, are the generators of $C_{n-k+1}(X)$.

If we now apply ψ_{k-1} to $\sum_{j=1}^{\operatorname{rank}(C_{k-1}(Y))} A_{i,j}\bar{\alpha}_j$, we obtain the following:

$$\psi_{k-1} \circ {}_Y\partial_k(\alpha_i) = \sum_{j=1}^{\operatorname{rank}(C^{n-k+1}(X))} A_{i,j}\bar{c}_j$$

We now examine ${}_X\,\delta^{n-k} \circ \psi_k$.

Recall that ${}_X\,\delta^{n-k}$ was defined so that ${}_X\,\delta^{n-k}(c_i) = c_i \circ {}_X\partial_{n-k+1}$. ${}_X\partial_{n-k+1}$ sends the core of a $(n-k+1)$-handle x^{n-k+1} to a formal sum of cores of $(n-k)$-handles. We can therefore denote the image of $\bar{\beta}_j$ under ${}_X\partial_{n-k+1}$ as follows:

$$\sum_{i=1}^{\operatorname{rank}(C_{n-k}(X))} B_{i,j}\beta_i$$

As before, note that geometrically the coefficients $B_{i,j}$ represent the number of times the attaching region of x_j^{n-k+1} "wraps around" the core of each x_i^{n-k}, with signs determined by orientation. And again, the core and co-core of each x_i^{n-k} intersect transversely exactly once, and so $B_{i,j}$ is equivalent to the signed transverse intersection number of $\bar{\beta}_j$ with the co-cores of each x_i^{n-k}. But the co-core of x_i^{n-k} is precisely α_i! So $B_{i,j} =$ the signed transverse intersection number of α_i with $\bar{\beta}_j = A_{i,j}$. Note that we implicitly used orientability of our manifold here to ensure that the orientations chosen for α_i and β_i are consistent with their boundary components under both ${}_Y\partial_k$ and ${}_X\partial_{n-k+1}$, thus giving a well-defined signed transverse intersection number.

With this in mind, we can rewrite the image of $\bar{\beta}_j$ under ${}_X\partial_{n-k+1}$ as:

$$\sum_{i=1}^{\operatorname{rank}(C_{n-k}(X))} A_{i,j}\beta_i$$

Applying c_i to this sum, we obtain the following form for ${}_X\,\delta^{n-k}$:

$${}_X\,\delta^{n-k}(c_i) = \sum_{j=1}^{\operatorname{rank}(C^{n-k+1}(X))} A_{i,j}\bar{c}_j$$

If we precompose this map with ψ_k, since ψ_k is an isomorphism sending α_i to c_i, we obtain:

$${}_X\,\delta^{n-k} \circ \psi_k(\alpha_i) = \sum_{j=1}^{\operatorname{rank}(C^{n-k+1}(X))} A_{i,j}\bar{c}_j$$

Thus we have:

$$\psi_{k-1} \circ {}_Y\partial_k = {}_X\,\delta^{n-k} \circ \psi_k$$

Therefore, the chain complex of Y is isomorphic to the co-chain complex of X.

Isomorphic chain complexes have isomorphic homology groups, and so we can conclude that $H_k(Y) \cong H^{n-k}(X)$. But recall that X and Y were merely two CW structures on the same manifold, and since the homology of a manifold is independent of its CW structure, we know that $H_k(Y) \cong H_k(X)$.

We can therefore conclude that $H_k(X) \cong H^{n-k}(X)$. □

There are numerous other uses of handlebodies in other areas of topology in addition to this proof of Poincaré duality. The idea of a handle decomposition, as well as the theory of rearranging handles in a manifold known as handle trading, is the key step in the proof of the h-cobordism theorem. Handlebodies can also be used to prove a classification of compact surfaces, and they are also closely related to Heegard splittings of 3-manifolds. The reader is referred to the references for further reading on these topics.

Acknowledgments

I would first like to thank my mentor at the University of Chicago, Hana Jia Kong, for her advice and support. I would also like to thank J. Peter May, who allowed me to participate in the University of Chicago Mathematics REU despite the unusual circumstances.

Additionally, I must thank John Etnyre for his superb lectures at PCMI 2019 which inspired this project, as well as Paul Melvin for his advice and support during PCMI.

References

[1] Glen E. Bredon. Topology and Geometry. Springer-Verlag. 1993.
[2] Paolo Degiorgi. Cellular Homology and the Cellular Boundary Formula. www.math.uchicago.edu/ may/REU2016/REUPapers/Degiorgi.pdf. 2016.
[3] Allen Hatcher. Algebraic Topology. Cambridge University Press. 2002.
[4] Amy Hua. An Introductory Treatment of Morse Theory on Manifolds. www.math.uchicago.edu/ may/VIGRE/VIGRE2010/REUPapers/Hua.pdf. 2010.
[5] John M. Lee. Introduction to Smooth Manifolds. Springer-Verlag. 2003.
[6] Yukio Matsumoto. Translation by Kiki Hudson and Masahico Saito. An Introduction to Morse Theory. American Mathematical Society. 2002.
[7] John Milnor. Morse Theory. Princeton University Press. 1963.
[8] C.T.C. Wall. Differential Topology. Cambridge University Press. 1961.