Cobordism and Exotic Spheres

A Senior Honors Thesis Submitted to
The Department of Mathematics, Harvard University
Cambridge, MA 02138
April 5, 1999

Joshua B. Plotkin
975 Memorial Drive #910
Cambridge, MA 02138
(617) 864-6702

Preface

Jules Henri Poincaré may rightly be considered the father of modern topology (Leonhard Euler and the Königsberg bridges notwithstanding). It is fitting, therefore, that many of the questions explored in this thesis originated with Poincaré.

Poincaré's famous conjecture – which has driven so much of twentieth-century topology – arose, in fact, as the successor of an earlier conjecture. Poincaré originally conjectured that any 3-manifold with the *homology* of S^3 is homeomorphic to S^3. But he soon found a counterexample, known today as the Poincaré homology sphere (obtained from S^3 by surgery along the right trefoil with framing 1). As a result, Poincaré re-formulated his conjecture into the now-famous statement that any compact, connected, oriented 3-manifold with trivial fundamental group is homeomorphic to S^3.

Although yet unresolved in dimension three, there is a natural generalization of Poincaré's conjecture stating that any compact, oriented n-manifold with the homotopy type of S^n, $n \geq 4$, is homeomorphic to S^n. Stephen Smale resolved the Generalized Poincaré Conjecture affirmatively for $n \geq 5$ in 1961. Finally, in 1982, Michael Freedman provided a proof in dimension four.

Smale approached the Generalized Poincaré Conjecture by investigating the properties of manifolds equivalent up to *cobordism.* As we shall see, cobordism theory is properly seen as a generalization of singular homology theory. It is ironic, therefore, in light of Poincaré's original homology conjecture, that the Generalized Poincaré Conjecture was resolved via cobordism theory.

As it turns out, however, cobordism theory can address questions even more subtle than the Generalized Poincaré Conjecture. In fact, cobordism and surgery may be used together to distinguish between smooth n-manifolds which are homeomorphic but not *diffeomorphic* to S^n. Such "exotic spheres" (or exotic differential structures on any manifold), were presumed nonexistent until John Milnor constructed the first one in 1956. The systematic characterization of exotic spheres, completed in 1963 by Milnor and Kervaire, relies heavily on cobordism, surgery, and the Hirzebruch signature formula. It is remarkable, in retrospect, that the cobordism tools first used by Smale to approach Poincaré's questions are also useful in the more subtle setting of exotic differential structures.

Contents

1 Introduction

In order to demonstrate that a smooth manifold M is an exotic sphere, there are clearly two distinct tasks required: (1) to show that M is homeomorphic to a sphere, and (2) to show that M is not diffeomorphic to a sphere. The former task is usually accomplished via the h-cobordism theorem, and the latter by the Hirzebruch signature formula. Unfortunately, the proof of the signature formula requires, more than anything else, a strong stomach for characteristic classes and standard algebraic topology. The proof of the h-cobordism theorem, on the other hand, is more topological in nature. Therefore, in the spirit of topology as an inherently visual subject, this thesis will focus on the h-cobordism theorem and its application to exotic spheres. In particular, this thesis will trace through the development of cobordism theory (and its relationship with Morse theory, handle theory, and surgery theory), prove the h-cobordism theorem, and use these tools to investigate exotic spheres. We will present a detailed construction of some exotic 7-spheres, and we will discuss higher-dimensional exotic spheres in so far as space (and time!) permits.

Aside from Smale's original paper, [S], there is already a detailed proof of the h-cobordism theorem found in [M2]. Fortunately (for this thesis), Smale and Milnor adopt different approaches to the theorem. Smale emphasizes handle theory, while Milnor utilizes Morse theory and gradient-like vector fields. Of course, when seen properly, these two approaches are entirely equivalent. In an attempt to reflect the important relationship between handle theory and Morse theory, the presentation of the h-cobordism theorem in this thesis is somewhat of a hybrid between Smale's and Milnor's approaches. Moreover, this thesis introduces and utilizes *cellular* homology as the algebraic framework within which to express handle cancellation.

Along the path towards the h-cobordism theorem, we will also encounter a small assortment of related results. Among these we will find an intuitively-pleasing proof of Poincaré Duality, as well as the result that every compact manifold has the homotopy type of a CW-complex.

Despite my efforts at synthesis, much of the material in this thesis draws heavily from excellent sources. I have tried to be careful to reference results accordingly throughout. Nevertheless, in the interest of academic honesty, I must mention in advance the three sources from which most of the material has been derived: Milnor's excellent notes on cobordism ([M2]) and the original papers of Milnor and Kervaire ([M3] and [KM]).

There has been much effort to make the material in this thesis almost entirely self-contained; only standard algebraic topology (Eilenberg-Steenrod homology) and bundle-theory is presumed. The only glaring departure from this framework is the use, without proof, of the Hirzebruch signature formula. In addition, in order to apply the h-cobordism theorem, the theory of characteristic classes is required. All of the necessary definitions and results about characteristic classes have been summarized in Appendix A. In the final section on higher-dimensional homotopy spheres, however, I make liberal use of results by Bott ([B]), Pontrjagin ([P]), and Serre ([Se]).

There is a large collection of people – faculty and students at Harvard and beyond – who have guided my exploration of the mathematics comprising this thesis. In particular, I wish to thank sincerely Raoul Bott for his inspirational flair, Mak Trifkovic for teaching me Lagrange's theorem on subgroups, and Alice Chen for her eagle-eyed pooh-reading. I especially would like to thank my advisor, Peter Kronheimer, for all of his encouragement and explanations during the past months. Finally, I am most grateful to John Roe for his patient, thoughtful guidance during the whole of last year.

Most importantly, I must thank my loving parents who, if not directly related to the completion of this thesis, have certainly provided a fair contribution.

2 Generalities

Unless otherwise stated, all (co)homology is taken with integer coefficients. Similarly, all manifolds are C^∞ manifolds and are assumed compact and oriented.

Cobordisms

A *cobordism* (M, V_0, V_1) is given by M, a smooth manifold with boundary, such that ∂M is the disjoint union of V_0 and V_1. More precisely, a cobordism is the equivalence class of such objects under diffeomorphisms preserving the decomposition of ∂M. In this case, V_0 is said to be *cobordant* to V_1. Thus, the study of manifolds up to cobordism equivalence is, in a sense, a generalization of singular homology.

Given a cobordism from V_0 to V_1 and from V_1 to V_2, it is clear how to form the composition cobordism from V_0 to V_2. (We have ignored the technical questions about well-definedness. See [K] for a careful discussion of these matters.)

A cobordism (M, V_0, V_1) is an *h-cobordism* if, in addition, V_0 and V_1 are both deformation retracts of M. The h-cobordism theorem states that a simply-connected h-cobordism is always a product cobordism – *i.e.* that such a cobordism is diffeomorphic to $V_0 \times [0, 1]$.

If a manifold M comes equipped with an orientation, then ∂M acquires an induced orientation. Oriented manifolds V_0 and V_1 are in the same *oriented cobordism class* if there is some manifold with boundary M and a orientation-preserving diffeomorphism from ∂M to $V_0 \coprod V_1$. By the Collar Neighborhood Theorem (see [MS]), the relation of oriented cobordism is reflexive, symmetric, and transitive.

Surgery and Handles

Let N be an n-manifold with an embedded sphere S^q having a trivial normal bundle. Then, by the tubular neighborhood theorem, we obtain an embedding $D^p \times S^q \subset N$ as a tubular neighborhood. We say that N' is obtained by *surgery* from N when we remove this $D^p \times S^q$ and glue in $S^{p-1} \times D^{q+1}$ instead, along their common boundary $S^{p-1} \times S^q$.

Now we discuss the notion of adding a handle to a manifold with boundary. We will follow the conventions set forth in [Roe]. To be precise, however, we must first introduce the notion of a manifold with corners. Such an object is, by definition, locally modelled on open subsets of $(\mathbb{R}^+)^n$. Specifically, a n-manifold with corners, M, has boundaries $\partial_S M$, possibly empty, for each subset $S \subset \{1, 2, \ldots, n\}$. The space $\partial_S M$ is locally modelled upon $\{(x_1, \ldots, x_n) \in \mathbb{R}^n | x_i = 0, i \in S\}$.

Unfortunately, introducing such objects is a nuisance in the smooth category. In the topological or piecewise-linear category, of course, manifolds with corners are nothing more than manifolds with boundary. In the smooth category, however, they are distinct objects, and the process of going from the former to the later is called *smoothing*. For our purposes, however, we will only use manifolds with corners whose boundaries $\partial_S M = \emptyset$ for $|S| \geq 3$. Aside from this section we will not belabor ourselves with the technical details of smoothing such corners. Roughly speaking, however, a manifold with second-order corners may be smoothed by excising tubular neighborhoods of the corners, doubling the angles (thus reducing to first-order corners), and re-attaching. This process will be called *unbending the corners*. Similarly, given a manifold-with-boundary M and a codimension zero submanifold-with-boundary $N \subset \partial M$, then we can *bend* M along ∂N to form a manifold with second-order corners, M_c, such that $\partial_1 M_c = N$ and $\partial_2 M_c = \partial M \setminus \operatorname{Int} N$. See the figure below.

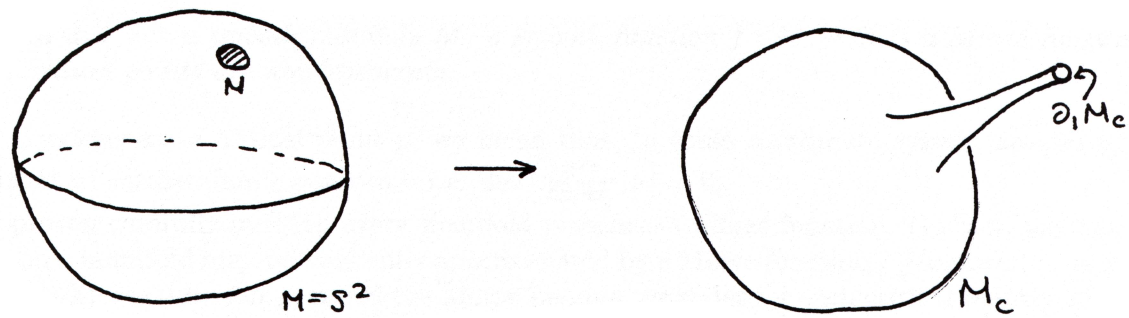

Now we can describe the process of attaching a handle to an n-manifold-with-boundary M. In essence, we simply perform surgery on ∂M. More specifically, given an embedded sphere S^q in ∂M with trivial normal bundle, we first obtain the tubular neighborhood $N = D^p \times S^q$ (where $p + q = n - 1$). N is a codimension zero submanifold of ∂M. Next, as described above, bend M along ∂N to form a manifold with corners M_c such that $\partial_1 M_c = D^p \times S^q$. Notice that a *handle*, $H = D^p \times D^{q+1}$, may naturally be considered as a manifold with corners such that $\partial_1 H = D^p \times S^q$ and $\partial_2 H = S^{p-1} \times D^{q+1}$. In the natural way, glue the handle H to M_c along their diffeomorphic ∂_1-boundaries, obtaining a new manifold with boundary $M' = M \cup_N H$.

Definition 2.1 *In the situation above, the manifold M' is said to be obtained from M by attaching a $(q+1)$-handle to $S^q \subset \partial M$.*

By convention, when M' is obtained by attaching a $(q+1)$-handle to M, as above, we call $S^q \subset \partial M$ the *attaching sphere*. On the other hand, D^p is called the *belt disc* of the handle. Dually, S^{p-1} is called the *belt sphere*, and D^{q+1} the *attaching disc* or *core disc*.

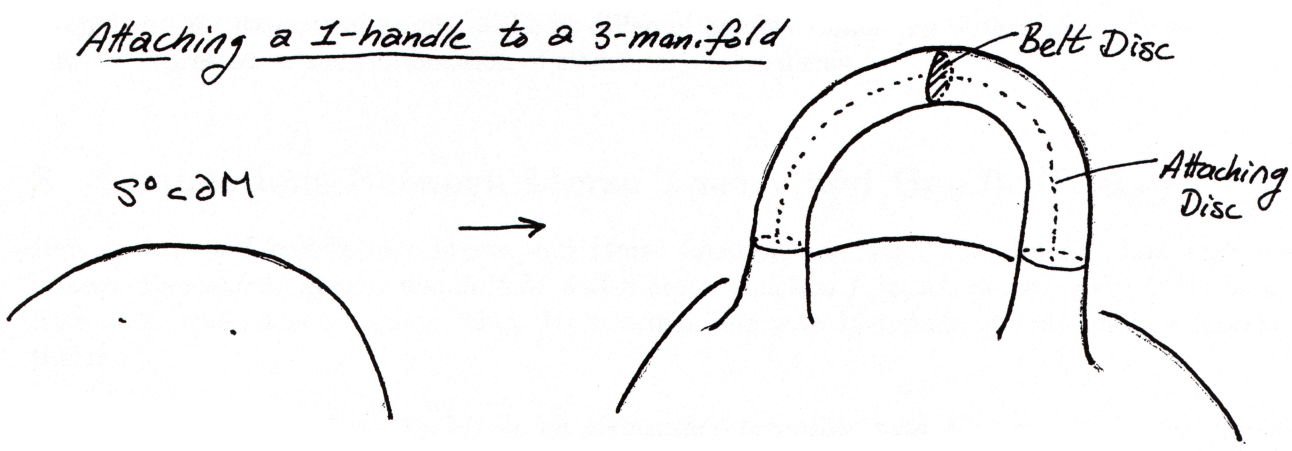

The following lemma is apparent from the definition of handle attachment:

Lemma 2.1 *If M' is obtained from M by attaching a handle to S^q in ∂M, then $\partial M'$ is obtained from ∂M by performing surgery on the sphere S^q.*

Notice that on the level of homotopy type, attaching handles is the same as attaching cells. By contracting the belt disc D^p to a point, we see that attaching a $(q+1)$-handle to S^q is equivalent (up to homotopy) to attaching a $(q+1)$-cell via the attaching map.

Morse Functions

Definition 2.2 *For a smooth manifold M, a smooth function $f : M \to \mathbb{R}$ is a Morse function if all of its critical points are nondegenerate.*

By a nondegenerate critical point p, we mean that, in some coordinate system around p, the determinant of the hessian is nonzero – *i.e.* $\det\left(\frac{\partial^2 f}{\partial x_i \partial x_j}|p\right) \neq 0$.

As is proven carefully in [M1], every manifold possesses a Morse function. (In fact, any smooth function on a manifold may be uniformly approximated by a Morse function.) Moreover, a standard fact about a Morse function f, called the Morse Lemma, provides for each critical point p a system of coordinates in which f takes the following form:

$$f(x_1, x_2, \ldots, x_n) = -x_1^2 - \ldots - x_\lambda^2 + x_{\lambda+1}^2 + \ldots + x_n^2 + f(p),$$

where λ, the *index* of the critical point, is the the number of negative eigenvalues of the Hessian (with multiplicity). A proof of the Morse lemma is little more than an exercise in diagonalization of matrices (see [M1]). Such coordinates around a critical point p will be called *Morse coordinates*.

Notice that, according to the Morse Lemma, the critical points of a Morse function are isolated, and thus there are finitely many of them for M compact.

By definition, a Morse function for a cobordism (M, V_0, V_1) is a smooth, nondegenerate function $f : M \to [0,1]$ such that $V_0 = f^{-1}(0)$ and $V_1 = f^{-1}(1)$. Furthermore, the critical points of f are required to lie on the interior of M. As with manifolds, cobordisms always possess Morse functions.

The *Morse number* of a cobordism (M, V_0, V_1) is the minimum over all Morse functions f of the number of critical points of f. A cobordism with Morse number one is called an *elementary cobordism*. As is demonstrated carefully in [M2], a Morse function on a cobordism may easily be altered so that its critical points all lie at different heights. Thus, we obtain that any cobordism may be expressed as the composition of elementary cobordisms.

3 Connections Between Morse Theory and Handle Theory

Having discussed handle attachment and Morse functions, we shall now describe how they are related. Henceforth, given a manifold M with a Morse function f, let M_a denote space $f^{-1}(-\infty, a]$. Note that, when a is a regular value, M_a is a manifold with boundary (by the implicit function theorem).

Proposition 3.1 *Let (M^n, V_0, V_1) be an elementary cobordism with $f(p) = a \in \mathbb{R}$, the critical point, of index λ. Then $M_{a+\epsilon}$ may be obtained from $M_{a-\epsilon}$ by attaching a λ-handle. Moreover, a cobordism with Morse number zero is a product cobordism.*

PROOF: First we consider the cobordism with no critical points. Endow M with a Riemannian metric and consider the vector field dual to df – $i.e$ the gradient vector field ∇f such that $< \eta, \nabla f >= \eta(f)$ for any other vector field η. Since f is non-critical on M, the flow lines of ∇f provide us with a diffeomorphism from $V_0 \times [0,1]$ to M. (Note: This proof rests, in essence, on the existence and uniqueness of solutions to ordinary elliptic differential equations.)

Now consider the elementary cobordism under the Morse coordinates in a neighborhood of p. Notice (see picture below) that to pass from $M_{a-\epsilon}$ to $M_{a+\epsilon}$ we must add a handle $D^\lambda \times D^{n-\lambda}$. Moreover, we have a natural disc $D^{n-\lambda}$ and sphere $S^{\lambda-1}$ embedded in the $(n-1)$-manifold $\partial M_{a-\epsilon}$ which play the roles of the belt disc and the attaching sphere. ■

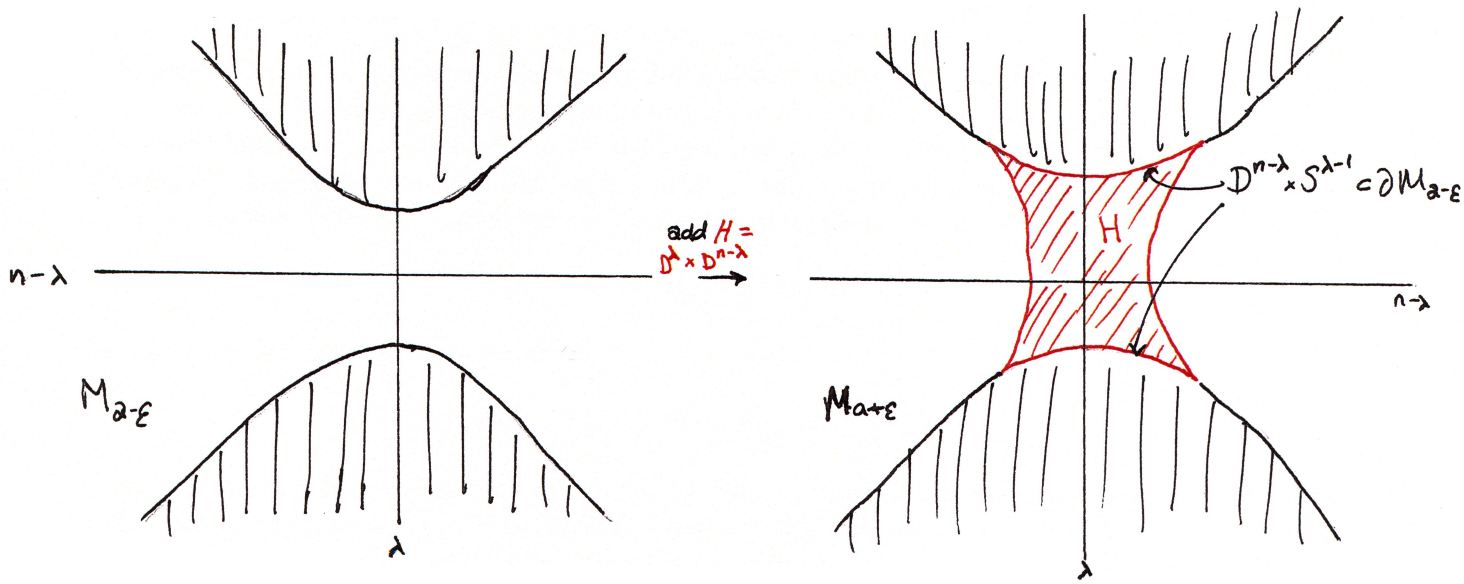

According to this proposition, specifying a Morse function on M is equivalent to specifying a *handle decomposition* of M.

The observation that a cobordism is a product if and only if it has Morse number zero forms the inspiration for Morse's proof of the h-cobordism theorem. By "cancelling critical points," we will show that an arbitrary Morse function on a h-cobordism may be altered so as to have no critical points.

Because we already know the effect of handle attachment on homotopy, we have the following immediate corollary of Proposition 3.1.

Corollary 3.1 *In the first situation described in Proposition 3.1, the manifold $M_{a+\epsilon}$ has the homotopy type of $M_{a-\epsilon}$ with a λ-cell attached.*

Many proofs of the corollary above employ the flow lines of ∇f in order to find an explicit retraction from $M_{a+\epsilon}$ to $M_{a-\epsilon} \cup D^\lambda$. Such techniques, which were championed by Marston Morse himself, will be useful for our purposes (in addition to handle-theoretic tools of Stephen Smale), and we will explore them in the next section.

This section has revealed the equivalence of Morse-decompositions and handle-decompositions up to homotopy type. Nevertheless, we should mention (without proof) that this equivalence extends also to the homeomorphism type. In other words, given a handle decomposition of M, we can always find an "equivalent" Morse function on M whose critical points have indices which agree with the corresponding handles.

Gradient-like Vector Fields

We now present the main tool used by Morse in his proof of the h-cobordism theorem: gradient-like vector fields. Near critical points, such vector fields mimic the behavior of the actual gradient of a Morse function, and they are nonzero elsewhere. More precisely,

Definition 3.1 *Given a cobordism (M, V_0, V_1) and a Morse function f, a vector field ξ on M is a gradient-like vector field for f provided that (1) $\xi(f) > 0$ away from the critical points, and (2) at each critical point there are coordinates $(\overline{x}, \overline{y}) = (x_1, \ldots, x_\lambda, \ldots, x_n)$ such that $f = f(p) - |x|^2 + |y|^2$ and ξ is given by $(-x_1, \ldots, -x_\lambda, x_{\lambda+1}, \ldots, x_n)$.*

It is well known that gradient-like vector fields always exist for Morse functions. Moreover, their trajectories provide us with a alternative definitions of the attaching and belt discs. These new definitions holds even outside of the local Morse coordinates. Given a manifold M^n with Morse function f, associated gradient-like vector field ξ, and a critical point p of index λ, let $\gamma(x,t)$ describe the flow lines of ξ. Then we can define the attatching disc as

$$D^\lambda = \{x \in M | \lim_{t \to \infty} \gamma(x,t) = p\}.$$

Similarly, we can define the belt disc as

$$D^{n-\lambda} = \{x \in M | \lim_{t \to -\infty} \gamma(x,t) = p\}.$$

By considering these definitions within the Morse coordinates, they are clearly equivalent to the definitions given previously. To be careful, however, this assumes that in the Morse-coordinates, the Riemannian structure on M agrees with the standard structure on R^n – which is not, in fact, absolutely correct. There is, however, always a homeomorphism of M which brings the flow of the Riemannian structure of M in-line with the standard flow in Morse coordinates.

Because of the equivalence of the handle-theoretic and Morse-theoretic definitions, we will often use the terminology *ascending disc* for belt disc, and *descending disc* for attaching disc. Similarly, the ascending and descending spheres are the boundaries of these discs. (See the figure below.)

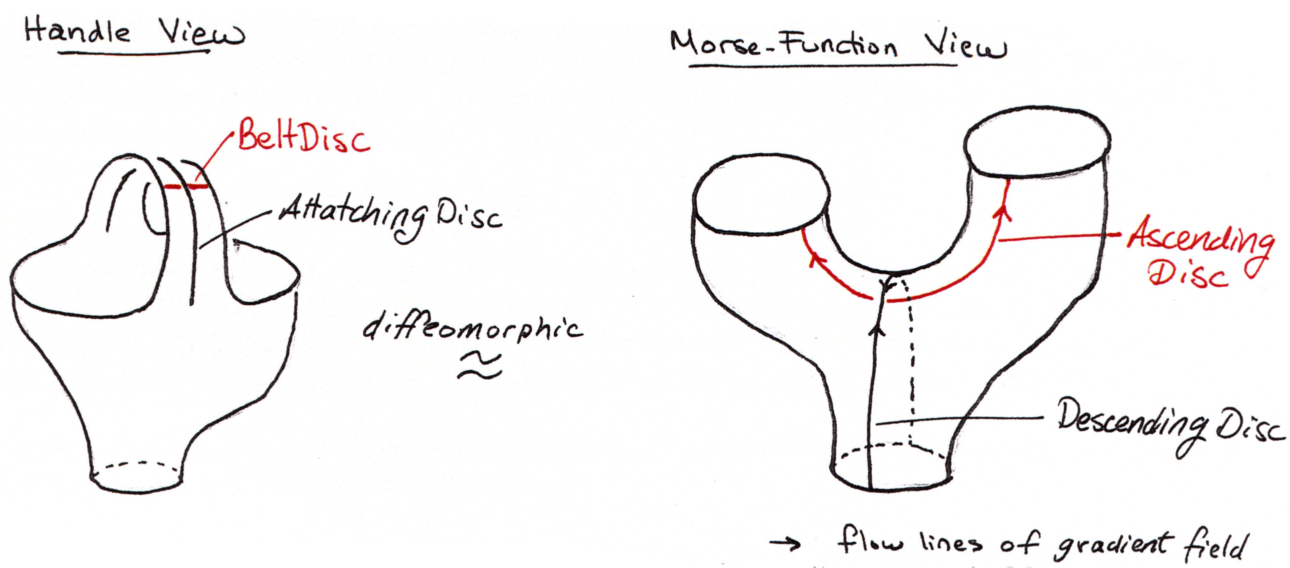

Our first step towards a proof of the h-cobordism is to alter a generic Morse function on (M, V_0, V_1) so that it becomes *self-ordered*. In other words, we seek a Morse function f such that for any two critical points p and q, if $\text{index}(p) > \text{index}(q)$ then $f(p) > f(q)$.

Proposition 3.2 *Given a Morse function f on M, we may alter f to a new Morse function which is self-ordered and which has the same critical points as f.*

PROOF: (Sketch) Consider two critical points of f, p and q of index λ and μ respectively, with $\lambda < \mu$, $f(p) > f(q)$, and no other critical points of intermediate height. The ascending sphere from q and descending sphere from p have codimension one in the $(n-1)$-dimensional level surface $L = f^{-1}(\frac{f(p)+f(q)}{2})$. Therefore, generically these spheres do not intersect in L. Hence, by slightly altering the gradient-like vector field for f, we may isotope p's descending sphere off of the handle associated with q. Once accomplished, it is evident that the p- and q-handles may be attached in any order desired. Attaching the p-handle first is equivalent to altering f so that $f(p) < f(q)$. ■

Aside from initiating our proof of the h-cobordism theorem, this proposition reveals that a compact smooth manifold has the homotopy type of a CW-complex.

Cellular Homology

In this section we introduce the tools of *cellular homology* for CW-complexes. We will use this flavor of homology theory throughout the proof of the h-cobordism theorem.

For a CW-complex X, let X^k denote the k-skeleton. Let Σ_k denote the set of k-dimensional cells of X. Although we have also used the notation D^k for discs, we shall now denote k-cells by σ^k. Arbitrarily fix orientations for each cell in Σ_k. Then notice the following obvious fact about the relative homology groups:

$$H_q(X^k, X^{k-1}) = H_q\left(\bigvee_{\sigma\in\Sigma_k} S^k_\sigma\right) = \begin{cases} 0 & \text{for } q \neq k \\ C_k(X) & \text{for } q = k \end{cases}$$

where $C_k(X)$ is the free abelian group on the elements of Σ_k. (The rightmost equality may easily be proven by induction and the relation $(\bigvee_{i\in I} D^n_i)/(\bigvee_{i\in I} \partial D^n_i) = \bigvee_{i\in I} S^n_i$. Alternatively, see [MS] p. 261). When there is no chance for confusion, we will denote $C_k(X)$ simply by C_k.

For our cellular differential complex we use the groups $H_q(X^k, X^{k-1}) = C_k$. We define the differential operator $\partial : C_k \to C_{k-1}$ by composing the maps δ and π_* from the exact sequences for the relative pairs (X^k, X^{k-1}) and (X^{k-1}, X^{k-2}):

$$H_k(X^k, X^{k-1}) \xrightarrow{\delta} H_{k-1}(X^{k-1}) \xrightarrow{\pi_*} H_{k-1}(X^{k-1}, X^{k-2}).$$

The composition above, denoted by ∂, is often called the boundary operator for the triple (X^k, X^{k-1}, X^{k-2}). It is immediately apparent, by the exactness of the sequences involved, that ∂ thus defined satisfies $\partial^2 = 0$. Hence, we have constructed a cellular differential complex (C_k, ∂). Of course, the dual construction $C^k = \mathrm{Hom}(C_k, \mathbb{Z})$ yields an associated dual complex.

Proposition 3.3 *The homology of the complex $(C_k(X), \partial)$ is canonically isomorphic to the singular homology of X (and similarly for cohomology).*

Proof: First we make a simple observation about the singular homology of the skeleta. Using induction and our calculation of $H_q(X^k, X^{k-1})$, it follows that $H_q X^k$ is zero for $q > k$ and isomorphic to $H_q X$ for $q < k$ (consider the exact sequence for the pair (X^k, X^{k-1}) and induct on k).

Now we shall inspect $\partial_k : C_k \to C_{k-1}$, in order to show that $\ker\partial_k/\operatorname{im}\partial_{k+1}$ is canonically isomorphic to $H_k(X)$. Consider the following commutative diagram

$$\begin{array}{ccccccc} & & 0 & & & & \\ & & \downarrow & & & & \\ C_{k+1} & \to & H_k(X^k, X^{k-2}) & \xrightarrow{\alpha} & H_k(X^{k+1}, X^{k-2}) & \to & 0 \\ & & \downarrow & & & & \\ & & C_k & & & & \\ & & \downarrow & & & & \\ & & C_{k-1} & & & & \end{array}$$

where the horizontal line is from the homology exact sequence of the triple (X^{k+1}, X^k, X^{k-2}), and the vertical line from the sequence corresponding to (X^k, X^{k-1}, X^{k-2}). The zero on top follows because $H_k(X^{k-1}, X^{k-2}) = 0$, and the one on the right because $H_k(X^{k+1}, X^k) = 0$.

The vertical line reveals that

$$\ker\partial_k \cong H_k(X^k, X^{k-2}).$$

And the horizontal line shows that

$$H_k(X^k, X^{k-2})/\ker\alpha \cong H_k(X^{k+1}, X^{k-2}).$$

Therefore, by commutativity, $\ker\partial_k/\operatorname{im}\partial_{k+1} \cong H_k(X^{k+1}, X^{k-2})$. But, using our observation at the beginning of the proof (and the exact sequence for the pair (X^{k+1}, X^{k-2})), we see that

$$H_k(X^{k+1}, X^{k-2}) \cong H_k(X^{k+1}) \cong H_k(X),$$

as desired.

Alternatively, consider the spectral sequence arising from the filtration of X by skeleta. Since $H_q(X^k, X^{k-1})$ is zero for $q \neq k$, $E_2^{p,q} = 0$ if $q \neq 0$ and $E_2^{p,0} = H_p(X)$. Since either the target or the range of d_3 is zero, we obtain $E_3 = E_\infty$. Finally, in E_∞ each line $p + q = k$ contains only one non-zero group, $E_\infty^{k,0} = E_2^{k,0}$. Therefore $E_\infty^{k,0} = H_k(X)$. ■

Now we give a more geometric interpretation of the cellular differential operator ∂. In fact, the efficacy of cellular homology arises from this geometric interpretation of ∂. Let I index the set Σ_k, and J index the set Σ_{k-1}. Let the differential $\partial_k : \mathbb{Z}^{|I|} \to \mathbb{Z}^{|J|}$ be expressed by the integer matrix (A_{ij}). Then, as an inspection of ∂_k's definition shows, the matrix element A_{ij} is simply the degree of the composition $S^{k-1} \to X \to S^{k-1}$, where the first map is the cellular gluing map from $\partial\sigma_i^k$ to X, and the second map contracts $X \setminus \sigma_j^{k-1}$ to a point. In other words, the entry A_{ij} equals the generic number of inverse images (counted with sign), under the attaching map $f_i : \partial\sigma_i^k \to X^{k-1}$, of a point in the interior of σ_j^{k-1}.

Cellular Homology and Morse Functions

Now let us apply the techniques of cellular homology to the CW-complex associated with a Morse function on M^n (or, equivalently, with a handle decomposition of M^n). We assume that our Morse function is self-ordered – *i.e* that the associated handle decomposition yields a sequence of co-dimension zero (in M) submanifolds-with-boundary, $\emptyset = M_{-1} \subset M_0 \subset \ldots \subset M_n = M$. In this sequence, M_k is obtained by attaching k-handles to M_{k-1}. Arbitrarily orient the attaching disc of each handle. Note that the attaching sphere $\alpha = S^{k-1}$ of a k-handle and the belt sphere $\beta = S^{n-k}$ of a $(k-1)$-handle are both submanifolds of ∂M_{k-1} whose dimensions sum to $(n-1) = \dim \partial M_{k-1}$. Therefore, these spheres generically intersect in a finite number of points. Moreover α receives an induced orientation from the disc it bounds, and the normal bundle to β is oriented as well. (Note that the normal bundle of β is simply the attaching disc of the $(k-1)$-handle.) Therefore, we can define an intersection number $[\alpha : \beta] \in \mathbb{Z}$ by summing over the intersection points counted with sign.

In this situation, we can alternatively consider the complex whose chain groups, $\mathfrak{C}_k$, are the free abelian groups on the k-handles themselves. The boundary map is given by the matrix $\mathfrak{M}_k$ whose (ij)'th entry equals the intersection number $[\alpha_i : \beta_j]$ of the i'th k-handle's attaching sphere with the j'th $(k-1)$-handle's belt sphere. In this context we have the following important result.

Proposition 3.4 (Smale) *In the situation above, the homology of $(\mathfrak{C}_k, \mathfrak{M}_k)$ is canonically isomorphic to the homology of M, or, more descriptively, to the cellular homology of the CW-complex associated with the handle decomposition of M.*

Proof: Consider the j'th $(k-1)$-handle $H = D^{n-k+1} \times D^{k-1}$ as a manifold with corners. Note that $\partial_2 H = \beta_j \times D^{k-1}$ lies in ∂M_{k-1}, and $\partial_1 H$ in ∂M_{k-2}. According to our definition of cellular homology and the CW-complex associated with a handle decomposition, the matrix entry $A_{ij} = \partial_k$ is the generic number of preimages of a point p in the interior of D^{k-1} under the map $\alpha_i \to S^{k-1}$ which attaches the k-handle, then shrinks the belt disc D^{n-k-1} to a point. But clearly this number equals the intersection number of α_i with the belt sphere $S^{n-k} \times \{p\}$. Thus the boundary operator for cellular homology agrees with the matrix $\mathfrak{M}_k$. ■

Nota Bene: According to the equivalence of Morse- and handle-decompositions, the complex $\mathfrak{C}_k$ may also be described as the free group on the critical points of index k. In this context, the (ij)'th element of $\mathfrak{C}_k$ denotes the intersection number of the descending sphere from p_i (of index k) with the ascending sphere of q_j (of index $k-1$). Thus, the dynamics of the Morse function's flow alone provides the homology of M.

As an aside, we note that the classical *Morse inequalities* follow easily from Proposition 3.4. Let b_k denote the k'th Betti number of M. Consider the complex of Proposition 3.4 but with rational coefficients. Since c_k, the number of critical points of index k, equals the dimension of the chain group $\mathfrak{C}_k$ (by our note above), elementary linear algebra yields the inequalities

$$\begin{aligned} b_0 &\leq c_0 \\ b_1 - b_0 &\leq c_1 - c_0 \\ b_2 - b_1 + b_0 &\leq c_2 - c_1 + c_0 \\ &\vdots \\ b_n - \ldots \pm b_0 &= c_n - \ldots \pm c_0. \end{aligned}$$

An Aside: The Dual Decomposition

Given a manifold M with Morse function f, we can also consider $-f$ as a Morse function with the same critical points (visually, we have simply turned M upside down). Notice that a critical point for f of index k becomes a critical points for $-f$ of index $n-k$. The handle-decomposition associated with $-f$ is called the *dual decomposition* to the one associated with f; the k-handles of one are the $(n-k)$ handles of the other, and the roles of the attaching and belt discs have been reversed. Moreover, when M is oriented, an orientation for the attaching disc of a handle determines a natural orientation for its belt disc (prescribing that the intersection point have positive sign). Clearly the matrices $\mathfrak{M}_k$ for the "upside-right" complex $\{\mathfrak{C}_k\}$ are replaced by their transposes in the "upside-down" complex $\{\overline{\mathfrak{C}}_k\}$. In other words, there is an isomorphism

$$\overline{\mathfrak{C}}_*(M) \cong \mathrm{Hom}(C_{n-*}(M), \mathbb{Z}) = \mathfrak{C}^{n-*}(M).$$

But these complexes compute the homology and cohomology of M, and thus we have found an intuitive proof of Poincaré duality.

4 Altering Handle Presentations

As we saw in the previous section, the matrices $\mathfrak{M}_k$ completely determine the homology of a M. In this section we shall investigate geometric maneuvers which effect standard algebraic manipulations of these boundary-operator matrices.

The following cancellation result of Smale forms the true backbone of the h-cobordism theorem.

Proposition 4.1 *Let M' be obtained from M by the successive attachment of a λ and $(\lambda+1)$-handle to ∂M. Suppose further that the attaching sphere of the $(\lambda+1)$-handle intersects the belt sphere of the λ-handle transversely in one point. Then M' is diffeomorphic to M.*

PROOF: (Sketch) Smale's proof of this theorem amounts to showing that the two handles together form $D^1 \times D^\lambda \times D^{n-\lambda-1}$. Without a detailed discussion, however, it is difficult to comprehend Smale's reasoning. Therefore, instead we will sketch Morse's proof of the result. Morse demonstrates the equivalent statement that a cobordism (M, V_0, V_1) whose Morse function, f, has two critical points p and q of index λ and $(\lambda+1)$ is, under the assumption about their ascending and descending spheres, a product cobordism.

More specifically, given f and a gradient-like vector field ξ for f, we will alter ξ near a trajectory T from p to q producing a nonzero vector field ξ' whose trajectories all flow from V_0 to V_1 – thus giving the cobordism a product structure. Moreover, ξ' will be a gradient-like vector field for a new Morse function f' without critical points, which agrees with f near ∂M.

Note first that, by our assumption on the ascending and descending spheres, there is a *unique* trajectory T from p to q. Next, we find coordinates near T such that

(1) p and q are given by $(0, \ldots, 0)$ and $(1, 0, \ldots, 0)$.
(2) $\xi(x)$ is given by $(v(x_1), -x_2, \ldots, -x_\lambda, -x_{\lambda+1}, x_{\lambda+2}, \ldots, x_n)$, where
(3) $v(x_1)$ is a smooth function, positive on $(0,1)$, zero at $0, 1$, and negative elsewhere, with $\left|\frac{\partial v}{\partial x_1}(x_1)\right| = 1$ near $x_1 = 0, 1$.

Now we may alter ξ to form ξ' which, in these coordinates, is expressed by,

(1) $\xi'(\vec{x}) = (w(x_1, \rho(x)), -x_2, \ldots, x_n)$ where $\rho = (x_2^2 + \ldots x_n^2)^{\frac{1}{2}}$ and
(2) $w(x_1, \rho) = v(x_1)$ away from T, and $w(x_1, 0)$ is always negative.

A picture of ξ and ξ' will be useful in understanding these conditions.

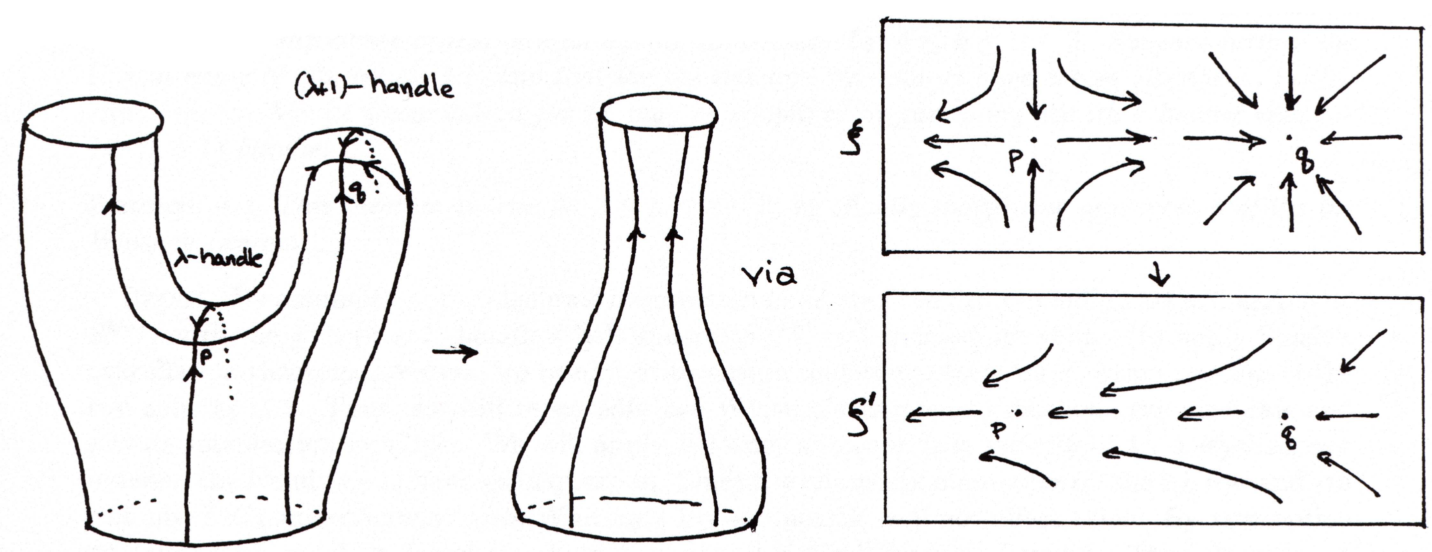

It is straightforward to verify that, as defined above, ξ' satisfies the desired properties of the proposition. (The diagrams above should serve to convey how the critical points are cancelled with each other.) The technical difficulties of Morse's proof, therefore, lie in constructing the canonical local coordinates near T. See [M2] (whence this sketch is distilled) for the details of this argument. ■

Notice that Smale's cancellation theorem requires that relevant attaching and belt spheres intersect transversely in one point. This condition is rather stringent (at least to an algebraic topologist). Before continuing with our proof of the h-cobordism theorem, we must first state a technical lemma due to Whitney which weakens Smale's requirements. The proof of this lemma, which employs the most sophisticated and interesting techniques presented in this paper so far, will be delayed until after the h-cobordism result.

Consider two submanifolds ${N_1}^{k_1}$ and ${N_2}^{k_2}$ of complementary dimensions intersecting transversely inside of M^n, an n-manifold without boundary. Furthermore, suppose that N_1 is oriented, as is the normal bundle of N_2 in M. In this case, the actual, setwise number of intersection points of N_1 and N_2 may be larger than the algebraic intersection number $[N_1 : N_2]$; often two intersection points of opposite sign may cancel algebraically. The Whitney Lemma illuminates when we can deform the submanifolds so that the the actual, setwise number of intersection points equals the topologically-necessary, algebraic intersection number.

Lemma 4.1 (Whitney) *In the situation described above, suppose further that $k_1, k_2 \geq 3$. Let P and P' be two intersection points of N_1 and N_2 having opposite signs. Suppose there exists paths γ_1 and γ_2 from P to P' in N_1 and N_2, respectively, such that $\gamma_1^{-1}\gamma_2$ is nullhomotopic in M. Then there is an ambient isotopy of N_1 into a submanifold N_1' transverse to N_2 such that*

$$N_1' \cap N_2 = N_1 \cap N_2 \setminus \{P, P'\}.$$

Eliminating Middle-Dimensional Handles

Now we shall re-interpret Smale's cancellation theorem in terms of a handle presentation for M^n, with intersection matrices $\{\mathfrak{M}_k\}$:

Lemma 4.2 *Assume that the cobordism M has no 0, 1, n or $n-1$ handles. Suppose also that M and the two components of ∂M are all simply connected. Let $4 \leq k \leq n-3$. Suppose further the (i, j)th entry of $\mathfrak{M}_k$ equals ± 1, and that the ith row and jth column have zeroes elsewhere. In this situation, the handle presentation for M may be simplified by cancelling the ith k-handle with the jth $(k-1)$ handle.*

Remark 4.1 *This theorem is true for $3 \leq k \leq n-3$, by slightly sharpening our version of the the Whitney Lemma.*

Proof: By assumption, the algebraic intersection number of the i'th k-handle's attaching sphere, S^{k-1}, with the j'th $(k-1)$-handle's belt sphere, S^{n-k}, is 1 in absolute value. To apply Smale's cancellation theorem, however, we need first to ensure that the *setwise* intersection number of the two spheres is 1. Thus, we will repeatedly use Whitney's Lemma to reconcile the algebraic and setwise intersection numbers. We will apply Whitney's isotopy in a vicinity of V, a level surface between the k and $(k-1)$ handles in question. Since the intersection numbers on the ith row and jth column of $\mathfrak{M}_k$ are otherwise zero, Whitney's isotopy near V will not affect them. By assumption we know that $k-1 \geq 3$ and $n-k \geq 3$, as required for Whitney's Lemma. Thus, in order to complete the proof, we must show that V is simply connected – so as to find the contractible loop required in Whitney's hypotheses. But by Van Kampen's theorem (and the bound on k), $\pi_1(V) = \pi_1(D^k \cup V \cup D^{n-k+1})$, where the discs are the attaching and belt discs of k and $(k-1)$-handles in question. But, again by Van Kampen's, if M and its boundaries are simply connected, then so too is the sub-cobordism consisting of the ith k-handle and jth $(k-1)$-handle. But this sub-cobordism is homotopy equivalent to $D^k \cup V \cup D^{n-k+1}$. ■

Now consider an h-cobordism (M, V_0, V_1). By assumption we know that $H_*(M, V_0) = 0$. In this case, we can use a relative version of the cellular homology described above. The triviality of this relative homology implies that, for each k, the (relative) differential operator is given by $\mathfrak{M}_k$, a square and invertible matrix. Fix some basis for $H_k(M, V_0)$. Then the process of passing from this fixed basis to any other basis is equivalent to performing row and column operations on $\mathfrak{M}_k$. But any elementary row or column operation on $\mathfrak{M}_k$ may be obtained by a corresponding shift in handle presentation (this important result relies on geometric maneuvers such as handle sliding; see [M2] Theorem 7.6 or [RS] for a detailed proof). Thus, we may change the handle presentation so as to ensure that the attaching discs of the new k-handles represent (in singular theory) the given basis of $H_k(M, V_0)$. In other words, we can ensure that each $\mathfrak{M}_k$ is the identity matrix. Using this we obtain

Corollary 4.1 *Suppose (M, V_0, V_1) is an h-cobordism of dimension $n \geq 7$ possessing a handle decomposition with no 0, 1, $n-1$, or n-handles. Then M is a product cobordism.*

PROOF: By assumption we know that M, V_0, and V_1 are all simply connected (hence orientable) and that $H_*(M, V_0) = 0$. Choose an ordered handle presentation of M. In the notation of cellular homology, we have a complex

$$\mathfrak{C}_{n-2} \xrightarrow{\partial} \mathfrak{C}_{n-3} \xrightarrow{\partial} \dots \xrightarrow{\partial} \mathfrak{C}_2$$

which forms an exact sequence. Alter the presentation of M so that each ∂_k is given by the identity matrix . Then every pair (i, i) satisfies the hypotheses of Lemma 4.2. Thus, for each k, we may cancel every k-handle with a corresponding $(k-1)$ handle. When all handles have been cancelled, M is manifestly a product. ■

Eliminating Low Dimensional Handles

In light of Corollary 4.1, we may finish a proof of the h-cobordism theorem by removing the 0, 1, $n-1$, and n-handles.

For the 0-handles we can apply Smale's theorem almost directly.

Lemma 4.3 *Let (M^n, V_0, V_1) be a connected cobordism with V_0 nonempty. Then M has a presentation with no 0-handles. If V_0 is empty, then it has a presentation with exactly one 0-handle.*

PROOF: The second statement is a consequence of the first; given M with V_0 empty, simply remove a small D^n from a M (*i.e.* remove a 0-handle) to reduce to the first case. To prove the first statement, consider a presentation of M with a minimal number of 0-handles. Since M is connected, so too is M_1 (the submanifold of all 1- and 0-handles). Therefore, if M_1 has a 0-handle it must also have a 1-handle connecting it to somewhere else. But the attaching sphere of this 1-handle is S^0, with one point in the 0-handle, and one point elsewhere. Therefore, by Smale's theorem we may cancel the 1-handle and the 0-handle. (See the picture below.) ∎

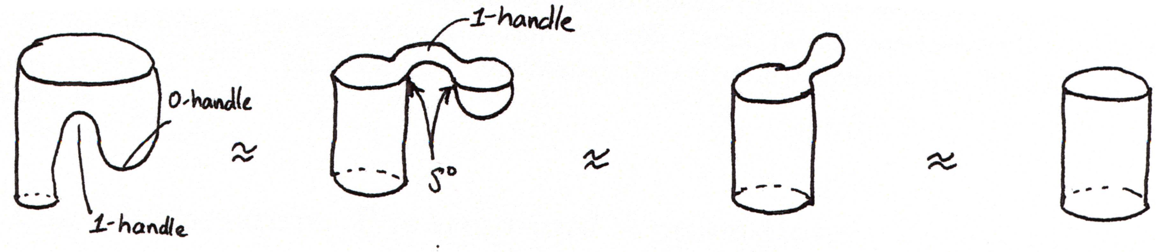

Following the 0-handles, we now face the 1-handles. As it turns out, these can be eliminated at the cost of introducing new 3-handles.

Lemma 4.4 *Let (M, V_0, V_1) be a simply connected h-cobordism with $n \geq 5$. Furthermore, let M be given by a presentation without* 0 *handles. Then we may alter the presentation of M so that it has no* 0 *or* 1*-handles, but the same number of handles as before in dimension* 4 *or greater.*

Proof: Let us denote V_0 by $\partial_- M$ and V_1 by $\partial_+ M$, and similarly for each submanifold $M_k \subset M$ in the handle presentation. By assumption, V_0 is connected. Consider a presentation with a minimal number of 1-handles. Choose a 1-handle H^1 attached to $\partial_+ M_0$. Let Γ_1 be a simple closed curve in $\partial_+ M$ which intersects the belt sphere of H^1 transversely in one point, and returns through $\partial_+ M_0$.

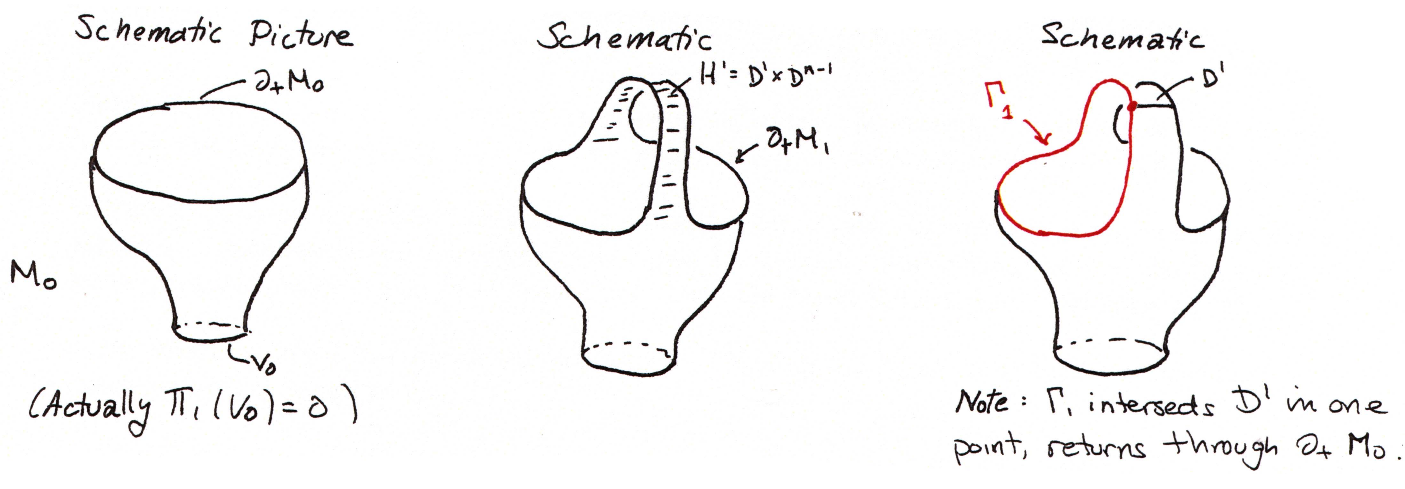

Because $n \geq 5$, we may assume by transversality that Γ_1 is disjoint from that attaching spheres of all 2-handles. As a result, Γ_1 lies in $\partial_+ M_2$. Since Γ_1 is nullhomotopic in M, and M is obtained by adding handles to ∂M_2 of index 3 or more, Γ_1 is, in fact, nullhomotopic in $\partial_+ M_2$.

Schematic

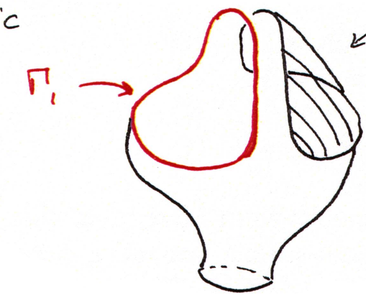

Any 2-handles attach here, do not hit Γ_1. Therefore $\Gamma_1 \subset \partial_+ M_2$

(Note: "Γ_1 nullhomotopic" is difficult to see in this schematic picture.)

Out of the way of all other handles, attach to M_1 a 2-handle H^2 and a cancelling 3-handle H^3. Of course this attachment does not alter M_1, and (as argued above) the attaching sphere of H^2 is a nullhomotopic curve Γ_2 in $\partial_+ M_2$.

Since $n \geq 5$, we may use Whitney's standard embedding theorem to find an embedded cylinder with boundary $\Gamma_1 \coprod \Gamma_2$. Pushing along this cylinder yields an isotopy between these two curves. We use this isotopy to move the trivial handle pair so that H^2 attaches along Γ_1. But, as a result, $H^1 \cup H^2$ is now trivial, by Smale's cancellation theorem (recall Γ_1 has one intersection point with the belt sphere of H^1). Thus we have eliminated the offending 1-handle by "trading" it for a 3-handle. ■

5 The *h*-cobordism Theorem and Applications

Lemmas 4.3 and 4.4 allow us finally to prove our desired result.

Theorem 5.1 *Suppose that (M^n, V_0, V_1) is a cobordism such that*

(1) $M, V_0,$ *and* V_1 *are simply connected*
(2) $H_*(M, V_0) = 0$
(3) $\dim M \geq 7$

Then M is diffeomorphic to $V_0 \times [0,1]$.

Remark 5.1 *This theorem is true for $n \geq 6$, by slightly sharpening our version of the the Whitney Lemma.*

Remark 5.2 *Certainly an h-cobordism satisfies (2) above. Moreover, it is true (although we will not use this fact) that (1) and (2) together imply that M is an h-cobordism (see [M2]).*

Proof: Choose an ordered handle decomposition for M (or, equivalently, a self-indexing Morse function). Eliminate 0 and 1-handles by Lemmas 4.3 and 4.4. Considering the dual decomposition, we can also eliminate n and $(n-1)$-handles. Then the result follows from Corollary 4.1. ■

We will now list a few major applications of Theorem 5.1. The most direct consequence is a characterization of the n-disc.

Theorem 5.2 *Let M^n be a compact, simply connected, smooth manifold, $n \geq 6$, with simply connected (nonempty) boundary. Then the following are equivalent*

(1) M *is diffeomorphic to* D^n
(2) M *is homeomorphic to* D^n
(3) M *is contractible*
(4) M *has the integral homology of a point*

PROOF: The only nontrivial implication is (4) $\Rightarrow$ (1). Consider a small n-disc D_0 embedded in $\operatorname{Int} M$. We have a natural simply-connected cobordism $(M - \operatorname{Int} D_0, \partial D_0, \partial M)$. By excision $H_*(M - \operatorname{Int} D_0, \partial D_0) \cong H_*(M, D_0) = 0$. Therefore we may apply Theorem 5.1 to deduce that $(M, \partial M)$ is a composition of $(D_0, \emptyset, \partial D_0)$ with the product cobordism $(M - \operatorname{Int} D_0, \partial D_0, \partial M) \approx \partial D_0 \times [0, 1]$. Since there is a unique, compatible differential structure on the composition of two smooth cobordisms, M is diffeomorphic to D_0. ■

The Generalized Poincaré Conjecture also follows easily.

Theorem 5.3 *If M^n, $n \geq 6$ is a closed, simply-connected smooth manifold with the integral homology of the n-sphere, then M is homeomorphic to S^n.*

Corollary 5.1 (Smale) *If a closed smooth manifold M^n has the homotopy type of S^n, $n \geq 5$, then it is homeomorphic to S^n.*

Remark 5.3 *This result is also true for $n = 5$ [KM].* [1]

PROOF: Let $D_0 \subset M$ be a smooth n-disc. The we have the following basic homology calculation:

$$\begin{aligned} H_i(M - \operatorname{Int} D_0) &\cong H^{n-i}(M - \operatorname{Int} D_0, \partial D_0) \text{ (Poincaré Duality)} \\ &\cong H^{n-i}(M, D_0) \text{ (excision)} \\ &\cong \begin{cases} 0 & \text{if } i > 0 \\ \mathbb{Z} & \text{if } i = 0 \end{cases} \text{ (exact relative sequence)} \end{aligned} \tag{1}$$

Therefore, by our characterization of the n-disc, $M = (M - \operatorname{Int} D_0) \cup D_0$ is diffeomorphic to a union of two discs, glued along some diffeomorphism of their boundaries.

Such a manifold, called a "twisted sphere," is clearly homeomorphic to S^n. (Alternatively, note that such a twisted sphere has Morse-number 2, and apply the fundamental result of Reeb [Re]). ■

Other corollaries of the h-cobordism theorem, which will not directly concern us, include the differentiable Schoenfliess theorem for dimensions ≥ 5 [M2].

[1] In fact, Michael Freedman proved a 5-dimensional topological h-cobordism theorem in 1982. This result, in turn, extends the Generalized Poincaré Conjecture affirmatively to dimension four. Freedman's general approach was to specialize the Whitney Lemma, which fails in dimension four. In dimension four, the Whitney Lemma would require self-intersections – a problem which Freedman resolved (in essence) by the clever, infinitely-repeated application of the Whitney Lemma itself.

6 The Proof of the Whitney Lemma

Finally, we will now address the proof of Whitney's Lemma. Notice that Whitney's Lemma was used rather extensively within the proof of the h-cobordism theorem. Moreover, the application of this lemma imposed the dimension restriction on the h-cobordism theorem. For the reader's convenience, we will restate Whitney's Lemma before proving it.

Consider two submanifolds $N_1^{k_1}$ and $N_2^{k_2}$ of complementary dimensions intersecting transversely inside of M^n, an n-manifold without boundary. Furthermore, suppose that N_1 is oriented, as well as the normal bundle of N_2 in M. In this case, the actual, setwise number of intersection points of N_1 and N_2 may be larger than the algebraic intersection number $[N_1 : N_2]$.

Lemma 6.1 (Whitney) *In the situation described above, suppose further that $k_1, k_2 \geq 3$. Let P and P' be two intersection points of N_1 and N_2 having opposite signs. Suppose there exists paths γ_1 and γ_2 from P to P' in N_1 and N_2, respectively, such that $\gamma_1^{-1}\gamma_2$ is nullhomotopic in M. Then there is an ambient isotopy of N_1 into a submanifold N_1' transverse to N_2 such that*

$$N_1' \cap N_2 = N_1 \cap N_2 \setminus \{P, P'\}.$$

Proof: First, by slightly deforming our paths, we may assume (using transversality) that γ_1 and γ_2 do not intersect any points of $N_1 \cap N_2$ except for P and P'.

From here, the idea of the proof is quite simple. We simply embed a "standard model" within which the required isotopy is easy to write down. For our standard model we choose two curves C_1 and C_2 in $\mathbb{R}^2$ intersecting transversely at points Q and Q', and enclosing a disc D. Choose an embedding $\phi_1 : C_1 \cup C_2 \to N_1 \cup N_2$ mapping the C_i to the γ_i.

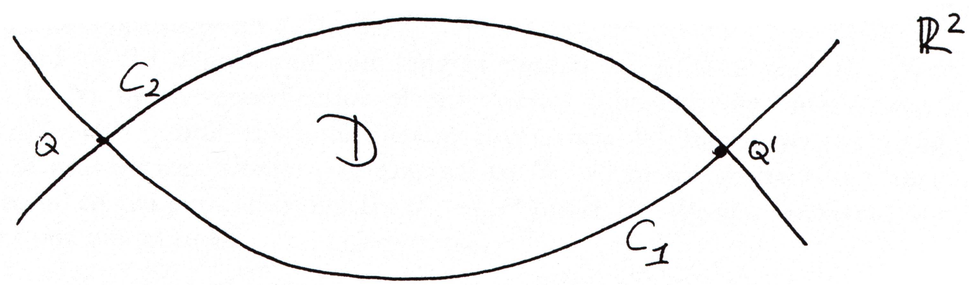

Whitney's result will follow, rather simply, once we show that the embedding ϕ_1 can be extended to the full standard model:

Lemma 6.2 *For some neighborhood U of D in the plane, we can extend ϕ_1 to an embedding $\phi : U \times \mathbb{R}^{k_1-1} \times \mathbb{R}^{k_2-1} \to M$ such that N_1 corresponds to $(U \cap C_1) \times \mathbb{R}^{k_1-1} \times 0$ and N_2 corresponds to $(U \cap C_2) \times 0 \times \mathbb{R}^{k_2-1}$.*

Assuming Lemma 6.2 for now, we quickly show how to finish the proof the Whitney Lemma. We will find an isotopy F_t of M from the identity-map to F_1 such that $N_1' = F_1(N_1)$ meets the required conditions. In fact, F_t will be the identity outside our embedded standard model. To define F_t, we define, instead, an isotopy G_t of $U \times \mathbb{R}^{k_1-1} \times \mathbb{R}^{k_2-1}$, and then pushforword via ϕ. G_t is defined as the identity on the second two coordinates. On U, however, G_t is defined so at to "pull" C_1 out of the way of C_2 (and as the identity near the boundary of $\overline{U} - U$):

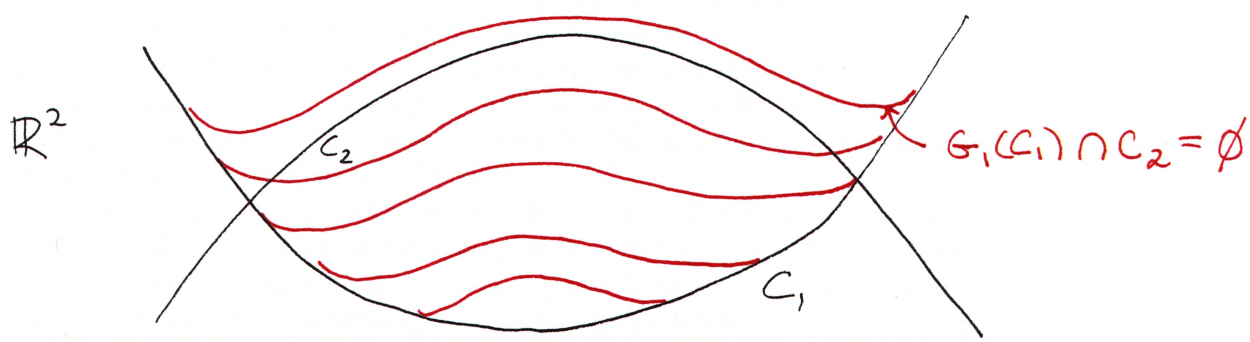

As defined, it it clear that F_t successfully pulls N_1 away from N_2 so that the intersection points P and P' are removed. ■

We have left to prove the crucial Lemma 6.2. This proof, which relies on general techniques in homotopy and obstruction theory, will occupy the remainder of this section. Our general approach to finding the desired embedding is geometric. We will endow M with a metric, construct vector fields on the model, and embed using the exponential map (as in the proof of the tubular neighborhood theorem).

First we endow M with a Riemannian metric which is Euclidean near P and P', and under which N_1 and N_2 are totally geodesic submanifolds. (Constructing such a metric is easy, and is left as an exercise for the industrious reader.) Our general plan is to use normal bundles and the tubular neighborhood theorem to extend ϕ_1 to ϕ.

Let $\tau_2(P)$ and $\tau_2(P')$ denote the unit tangent vectors to γ_2 at P and P'. Note that $\tau_2(P)$ is orthogonal to N_1 by the specification of our metric. Consider the bundle over γ_1 of vectors orthogonal to N_1; this bundle is trivial since γ_1 is contractible. Hence, we can extend $\tau_2(P)$ to a smooth field of unit vectors along γ_1 orthogonal to N_1, equal to the parallel translates of $\tau_2(P)$ near P, and equal to the parallel translates of $-\tau_2(P')$ near P'. We also construct a corresponding vector field on our model in $\mathbb{R}^2$.

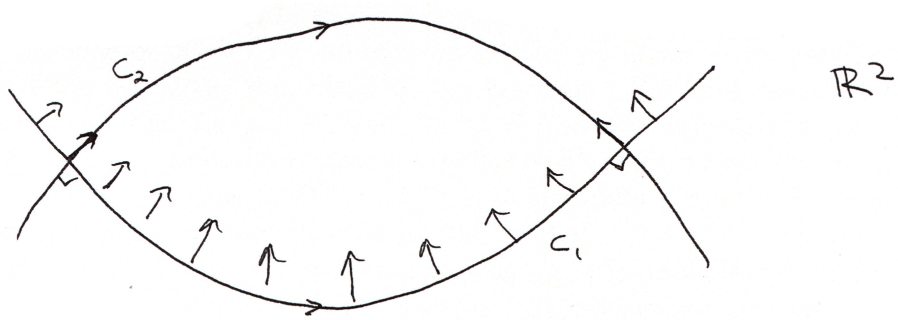

Now, by the standard properties of the exponential map (a local diffeomorphism) we find a neighborhood of C_1 in the plane and an extension of $\phi_1|_{C_1}$ to an embedding of this neighborhood into M.

Let $\tau_1(P)$ and $\tau_1(P')$ denote the unit tangent vectors to γ_1 at P and P'. Again, we may extend $\phi_1|_{C_2}$ to an embedding of a neighborhood of C_2 (using a field of unit vectors along C_2 orthogonal to N_2 *etc.* as for we did C_1.)

Since our specified metric is Euclidean near P and P', we see that the two embeddings agree at these points, and thus combine to form an embedding $\phi_2 : A \to M$, where A is an annular neighborhood of $C_1 \cup C_2$ in the plane.

Our next task is to extend ϕ_2 to a neighborhood, U, of the entire disc D. Let S denote the inner boundary of the annulus A. Since $\gamma_1^{-1}\gamma_2$ is homotopic to $\phi_2(S)$, $\phi_2(S)$ is also contractible in M. Moreover, by Whitney's standard embedding theorem there is a (homotopy class of) map $\phi_3 : U \to M$ realizing the null-homotopy of S (because $n \geq 5 = 2 \cdot 2 + 1$). Moreover the embedded disc $\phi_3(D)$ may be assumed disjoint from N_1 and N_2 (by transversality, because the codimensions of these submanifolds are both at least three).

It remains, finally, to extend ϕ_3 to $U \times \mathbb{R}^{k_1-1} \times \mathbb{R}^{k_2-1}$. We will find an obstruction to this extension in general, which will be obviated by our assumption that the signs of P and P' are opposite. Once again, our approach is to construct appropriate vector fields and use the exponential map.

In particular, we use the following intermediate lemma whose statement and proof are due to Milnor. Henceforth let U' denote $\phi_3(U)$, and let γ_1, γ_2, C_1 and C_2 denote $U' \cap \gamma_1$, $U' \cap \gamma_2$, $U' \cap C_1$, and $U' \cap C_2$.

Lemma 6.3 *There exist smooth vector fields $\xi_1, \ldots, \xi_{k_1-1}, \eta_1, \ldots, \eta_{k_2-1}$ on U' which satisfy (1) below, and such that the ξ's satisfy (2) while the η's satisfy (3)*

(1) are orthonormal and orthogonal to U'
(2) are tangent to N_1 along γ_1
(3) are tangent to N_2 along γ_2

Proof: We construct the ξ's in steps: first along γ_1, then extending to $\gamma_1 \cup \gamma_2$, and then finally to all of U'. Each of the two extensions will require a bundle argument.

Let τ_1 and τ_2 be the normalized velocity vector fields along γ_1 and γ_2. Let β_2 be the field of unit vectors along γ_2 which are tangent to $U' \subset M$ and inward orthogonal to γ_2. Finally, let $\nu(N_2)$

denote the normal bundle of $N_2 \subset M$. Note that $\beta_2(P) = \tau_1(P)$ and $\beta_2(P') = -\tau_1(P')$.

Choose $k_1 - 1$ vectors $\xi_1(P), \ldots, \xi_{k_1-1}(P)$ which are tangent to N_1 at P, orthogonal to U', and such that the k_1-frame $\tau_1(P), \xi_1(P), \ldots, \xi_{k_1-1}(P)$ is positively oriented in $T_P N_1$. We parallel translate these vectors to define the ξ's along γ_1. These vectors automatically satisfy (2) because parallel translation along a curve in a a totally geodesic submanifold sends tangent vectors (to N_1) to tangent vectors. By continuity, the k_1 frame constructed is positively oriented in TN_1 along γ_1.

In small neighborhoods of P and P', we can extend the ξ's along γ_2 by parallel translation. We wish to extend the ξ's along the whole of γ_2, however. We have assumed that the intersection numbers of N_1 and N_2 are $+1$ and -1 at P and P', respectively. In other words, $\tau_1(P), \xi_1(P), \ldots, \xi_{k_1-1}(P)$ is positively oriented in $\nu(N_2)$ at P, while negatively oriented in $\nu(N_2)$ at P'. But, since $\beta_2(P) = \tau_1(P)$ and $\beta_2(P') = -\tau_1(P')$, at all points of γ_2 near P and P' the frames $\beta_2, \xi_1, \ldots, \xi_{k_1-1}$ are positively oriented in N_2's normal bundle.

We wish to extend the ξ's to γ_2 – *i.e.* to find a moving $(k_1 - 1)$-frame over γ_2 agreeing with the frame already defined over γ_1. Instead of looking for $(k_1 - 1)$ independent sections over γ_2, we look for a single (nonzero) cross section of the frame-bundle of (k_1-1)-frames $(\zeta_1, \ldots, \zeta_{k_1-1})$, orthogonal to N_2 and to U', and such that $\beta_2, \zeta_1, \ldots, \zeta_{k_1-1}$ is positively oriented in N_2's normal bundle. This frame bundle is trivial with fiber $SO(n - k_2 - 1) = SO(k_1 - 1)$. Since the fiber is connected (and we are trying to extend over a 1-dimensional manifold γ_2), we can extend $\xi_1, \ldots, \xi_{k_1-1}$ to a smooth field of $(k_1 - 1)$-frames on $\gamma_1 \cup \gamma_2$ satisfying conditions (1) and (2).

Aiming to extend to all of U', we must consider the frame-bundle over U' of orthonormal (k_1-1)-frames orthogonal to U' – which is trivial with fiber $O(k_1 + k_2 - 2)/O(k_2 - 1) = V_{k_1-1}(\mathbb{R}^{k_1+k_2-2})$, the Stiefel manifold of orthonormal $(k_1 - 1)$-frames in $\mathbb{R}^{k_1+k_2-2}$. We already have defined a smooth section of this bundle over $\gamma_1 \cup \gamma_2$. Composing this section with projection onto the fiber gives us a smooth map of $S^1 = \gamma_1 \cup \gamma_2$ into $O(k_1 + k_2 - 2)/O(k_2 - 1)$. Thus, the obstruction to extending this section lies in $\pi_1(V_{k_1-1}(\mathbb{R}^{k_1+k_2-2}))$ which, as $k_2 \geq 3$ is trivial according to [St]. Thus we can perform the required extension to all of U', satisfying (1) and (2).

It is interesting to note that the triviality of $\pi_1(V_{k_1-1}(\mathbb{R}^{k_1+k_2-2}))$ is itself demonstrated by using general position arguments. Thus, throughout the entire proof of Whitney's Lemma so far, we have relied almost entirely upon general position.

Finally, to define the η's, consider the bundle over U' of orthonormal frames $\eta_1, \ldots, \eta_{k_2-1}$ in TM such that each η_i is orthogonal to U' and to the ξ's. This bundle is trivial by the contractibility of U'. Hence we can find the field of frames $\eta_1, \ldots, \eta_{k_2-1}$, which, together with the ξ's satisfy the conditions (1), (2), and (3) of Lemma 6.3. Of course, condition (3) is satisfied by the η's because they are orthogonal to the ξ's, which were constructed to be orthogonal to N_2 along γ_2. ■

Having finished the proof of Lemma 6.3, we now complete the proof of Lemma 6.2, and thus also of the Whitney Lemma. We shall use the vector fields constructed to define a map $\phi_4 : U \times \mathbb{R}^{k_1-1} \times \mathbb{R}^{k_2-1} \to M$ by

$$\phi_4 : (u, x_1, \ldots, x_{k_1-1}, y_1, \ldots, y_{k_2-1}) \mapsto \exp\left[\sum_{i=1}^{k_1-1} x_i \xi_i(\phi_3(u)) + \sum_{j=1}^{k_2-1} y_j \eta_j(\phi_3(u))\right]$$

(Notice that this map is a slightly altered version of the embedding used to prove the tubular neighborhood theorem.) By the standard properties of the exponential map, we know that ϕ_4, near the origin in $\mathbb{R}^{k_1-1} \times \mathbb{R}^{k_2-1}$, is an embedding. But a neighborhood of the origin of $\mathbb{R}^{k_1-1} \times \mathbb{R}^{k_2-1}$ is diffeomorphic to $\mathbb{R}^{k_1-1} \times \mathbb{R}^{k_2-1}$ itself, and so we define $\phi : U \times \mathbb{R}^{k_1-1} \times \mathbb{R}^{k_2-1} \to M$ to be this diffeomorphism composed with ϕ_4. Then $\phi(C_1 \times \mathbb{R}^{k_1-1} \times 0) \subset N_1$ (This follows because N_1 is a totally geodesic submanifold by choice of metric; on $C_1 \times \mathbb{R}^{k_1-1} \times 0$, ϕ only deals with vectors $v \in T_p(N_1) \subset T_p(M)$, and $\exp(tv)$ is a geodesic in M, tangent to N_1 at $t = 0$, and thus entirely within N_1). Similarly $\phi(C_2 \times 0 \times \mathbb{R}^{k_2-1}) \subset N_2$. Moreover, since $\phi(U \times 0) = U'$ intersects N_1 and N_2 transversely in γ_1 and γ_2, it follows that $\phi^{-1}(N_1) = C_1 \times \mathbb{R}^{k_1-1} \times 0$ and $\phi^{-1}(N_2) = C_2 \times 0 \times \mathbb{R}^{k_2-1}$, as desired. ∎

7 Signatures and Connected Sums

Before turning to the construction of some exotic 7-spheres, we must first recall the notions of signature and connected sum.

Let M by a $4n$-dimensional, oriented manifold with (possibly empty) boundary. Denote its orientation class by $\mu_M \in H_{4n}(M, \partial M)$. Then Poincaré Duality defines a nondegenerate, quadratic form

$$F : H^{2n}(M, \partial M)/\text{torsion} \to \mathbb{Z}$$

given by $\alpha \mapsto < \mu_M, \alpha \cup \alpha >$. Such forms are classified by their signatures – the number of positive minus negative eigenvalues over $\mathbb{Q}$. We define $\sigma(M)$, *the signature of* M, to be the signature of its this form F.

There is a remarkable connection, due to Hirzebruch, between a manifold's signature and its Pontrjagin classes. We will only state this theorem for dimension eight.

Theorem 7.1 (Hirzebruch) *Let M be an 8-manifold with orientation class ν. Then its signature obeys the following equation:*

$$\sigma(M) = < \nu, \frac{1}{45}(7p_2(M) - p_1^2(M)) > .$$

Hirzebruch's signature formula – and, in particular, the (somewhat trivial) number-theoretic aspects of the coefficients involved – will play a crucial role in discovering which spheres are exotic. The proof of Hirzebruch's theorem, which is less strictly-topological than the h-cobordism theorem, may be found in [MS].

Given two n-manifolds M_1 and M_2, we require a formal definition of their connected sum. Let $D^n \subset \mathbb{R}^n$ be the unit n-disc, and let $2D^n$ denote the disc with radius 2. Choose an orientation preserving embedding $i_1 : 2D^n \to M_1$, and an orientation reversing embedding $i_2 : 2D^n \to M_2$. Then we define the *connected sum* of M_1 and M_2 as the quotient of $M_1 \coprod M_2$ formed by deleting $i_1(\operatorname{Int} D^n)$ and $i_2(\operatorname{Int} D^n)$, and identifying $i_1(S^{n-1})$ with $i_2(S^{n-1})$ in the natural way. We denote the connected sum by $M_1 \# M_2$.

We should remark that we have used $2D^n$ to ensure that the connected sum has a well-defined manifold structure on the embedded $(n-1)$ spheres. It is well know that $M_1 \# M_2$ does not depend upon the choice of i_1 and i_2 (see [K]). The connected sum inherits a natural orientation and smooth structure. Moreover, the connected-sum is obviously a commutative and associative operation for which S^n acts as an identity.

Now we will generalize this notion to the connected sums of bundles. We shall only employ the notion of a bundle connected sum over a sphere (although our definition will be easy to generalize). Let ξ and ξ' be two S^{n-1} bundles (or associated $\mathbb{R}^n$, D^n, or SO_n bundles) over $B = S^k$. Consider $S^k \bigvee S^k$ as $S^k \coprod S^k$ with the two north poles identified. There is an obvious bundle $\xi \bigvee \xi'$ defined over $S^k \bigvee S^k$ which equals ξ and ξ' over their respective base spheres: simply identify the fibers over the north poles using an orientation-preserving isomorphism. Let $f : S^k = S^k \# S^k \to S^k \bigvee S^k$ be the obvious map which takes $i_1(S^{k-1}) = i_2(S^{k-1})$ to the common north pole. Then we have the following

Definition 7.1 *The bundle connected sum $\xi \# \xi'$ over S^k is defined as $f^*(\xi \bigvee \xi)$.*

For ξ and ξ' both S^{n-1} bundles over S^k we have the following

Proposition 7.1 *For any k, $e(\xi \# \xi') = e(\xi) + e(\xi')$. And for $k = 4i$, $p_i(\xi \# \xi') = p_i(\xi) + p_i(\xi')$.*

Proof: (I learned this proof from M. Thaddeus) Let $f : S^k = S^k \# S^k \to S^k \bigvee S^k$ be defined as it was for the bundle connected sum. Let $g : S^k \coprod S^k \to S^k \bigvee S^k$ be the map identifying north poles. We know from [K] that g^* induces an isomorphism in cohomology. Let us consider the map $f^*(g^*)^{-1} : H^k(S^k \coprod S^k) \to H^k(S^k \# S^k) = H^k(S^k)$ in light of the standard isomorphisms $H^k(S^k \coprod S^k) \cong H^k(S^k \bigvee S^k) \cong \mathbb{Z} \oplus \mathbb{Z}$, and $H^k(S^k) \cong \mathbb{Z}$. We shall inspect $f^*(g^*)^{-1}$ on the generators for $H^k(S^k \coprod S^k)$ by looking at their effect on the orientation class $\mu \in H_k(S^k)$:

$$< \mu, f^*(g^*)^{-1}(0,1) >=< f_*\mu, (g^*)^{-1}(0,1) >=< (\mu,\mu), (0,1) >= 1$$
$$< \mu, f^*(g^*)^{-1}(1,0) >=< f_*\mu, (g^*)^{-1}(1,0) >=< (\mu,\mu), (1,0) >= 1$$

Thus, by linearity, $f^*(g^*)^{-1}$ is simply addition. But by naturality $e(\xi \# \xi') = f^*(e(\xi \bigvee \xi')) = f^*(g^*)^{-1} e(\xi \coprod \xi') = f^*(g^*)^{-1}((e(\xi), e(\xi')) = e(\xi) + e(\xi')$. The proof for Pontrjagin classes is exactly analogous. ■

8 Exotic 7-Spheres

With the h-cobordism theorem in hand, we are finally able to approach the construction of the first exotic spheres following [M3]. As a pool of candidate exotic spheres we will consider the set of S^3-bundles over S^4. Notice that such bundles can be trivialized over either hemisphere of S^4, and thus they are determined by the homotopy class of a map from the equator $S^3 \subset S^4$ to SO_4. In other words, such bundles are parametrized by $\pi_3(\mathrm{SO}_4) = \mathbb{Z} \oplus \mathbb{Z}$. In fact, an explicit isomorphism is provided by [St]. Consider $u \in S^3$ and $v \in \mathbb{R}^4$ as quaternions, and to $(h,j) \in \mathbb{Z} \oplus \mathbb{Z}$ we associate the map $f_{hj} : S^3 \to \mathrm{SO}_4$ defined by $f_{hj}(u) \cdot v = u^h v u^j$. Let ξ_{hj} denote the S^3-bundle determined by f_{hj}, and let E_{hj} denote its total space. Similarly, let B_{hj} denote the total space the associated D^4-bundle. The manifold B_{hj} receives a natural differentiable structure which restricts to a smooth structure on $E_{hj} = \partial B_{hj}$. We will define the orientations on these manifolds somewhat later.

In some cases it is easy to characterize E_{hj} and B_{hj}.

Proposition 8.1 *$B_{00} = D^4 \times S^4$ and $E_{00} = S^3 \times S^4$.*

Proof: Note that on the equator of S^4, $f_{00}(u) = \mathrm{id} : \mathbb{R}^4 \to \mathbb{R}^4$. ■

Proposition 8.2 *B_{10} is the quaternionic projective space, $\mathbb{HP}^2$, minus an open disk. E_{10} is S^7.*

PROOF: Recall that $\mathbb{HP}^n$ is the quotient of $\mathbb{H}^{n+1} - \{0\}$ under the identification $(u, v, w) \sim (xu, xb, xw)$, for $x \in \mathbb{H}^*$. There is a natural injection of $\mathbb{HP}^1 \hookrightarrow \mathbb{HP}^2$ given by $[u, v] \mapsto [u, v, 0]$. There is also a natural fibration $\mathbb{H} \hookrightarrow \mathbb{HP}^2 - \{[0, 0, 1]\} \xrightarrow{\pi} \mathbb{HP}^1$ given by $\pi^{-1}([u, v]) = \{[u, v, w] \mid w \in \mathbb{H}\}$. Notice that $\mathbb{HP}^1$, may be decomposed into two 4-discs,

$$D_1 = \{[u, 1] \mid \|u\| \leq 1\}$$
$$D_2 = \{[1, v] \mid \|v\| \leq 1\}$$

sewn together via the reflection map $[u, 1] \mapsto [1, u^{-1}]$. Therefore, the fibration is simply an $\mathbb{H}$- or $\mathbb{R}^4$-bundle over S^4. In fact, we can write down local trivializations of the bundle over D_1 and D_2 as

$$\phi_1 : D_1 \times \mathbb{H} \to \pi^{-1}(D_1),\ \phi_1([u, 1], w) = [u, 1, w] \text{ and}$$
$$\phi_2 : D_2 \times \mathbb{H} \to \pi^{-1}(D_2),\ \phi_2([v, 1], w') = [1, v, w'].$$

The bundle's transition function acts on the equator (*i.e* where $[u, 1] = [1, v]$). Over this equatorial S^3, we have

$$\phi_2^{-1}\phi_1([u, 1], w) = \phi_2^{-1}([u, 1, w]) = \phi_2^{-1}([1, u^{-1}, u^{-1}w]) = ([1, v], vw),$$

and so $\phi_2^{-1}\phi_1$ corresponds to f_{10}.

From the total space $\mathbb{HP}^2 - \{[0, 0, 1]\}$ we wish to remove the open 8-disc $\{[u, v, 1] \mid \|u\|^2 + \|v\|^2 < 1\}$ centered at $[0, 0, 1]$. In other words, we must restrict the fiber over $[u, v]$ to the set $\{[u, v, w] \mid \|w\|^2 \leq \|u\|^2 + \|v\|^2\}$. But, for each fixed $[u, v]$, then, the fiber is homeomorphic to D^4, and the resulting D^4-bundle over S^4 also has transition function f_{10}. Therefore the disc-bundle B_{10} is none other than $\mathbb{HP}^2$ minus an open 8-disc. Moreover, $\partial B_{10} = E_{10}$ is homeomorphic to S^7, the boundary of the extracted 8-disc. ■

Calculating e and p_1

Using our characterization of B_{10}, we will be able to derive formulas for the characteristic classes of all ξ_{hj}, up to sign. Let ι be the standard generator of $H^4(S^4)$. Then we have the following

Proposition 8.3 *The characteristic classes of ξ_{hj} are given by*

$$e(\xi_{hj}) = \pm(h + j)\iota,$$
$$p_1(\xi_{hj}) = \pm 2(h - j)\iota.$$

PROOF: First, by considering the bundle connected sum, we will show that these characteristic classes are linear in h and j. Consider two bundles ξ_{hj} and $\xi_{h'j'}$ as well as the connected sum $\xi_{hj}\#\xi_{h'j'}$. Since we may choose the embedded discs $i_1(D^4)$ and $i_2(D^4)$ so as to intersect the equatorial S^3, by naturality the bundle connected sum adds according the group operation in $\pi_3(\mathrm{SO}_4)$, *i.e.* $\xi_{hj}\#\xi_{h'j'} = \xi_{h+h',j+j'}$. Then, by Proposition 7.1, we have

$$e(\xi_{h+h',j+j'}) = e(\xi_{hj}\#\xi_{h'j'}) = e(\xi_{hj}) + e(\xi_{h'j'}) \text{ and}$$
$$p_1(\xi_{h+h',j+j'}) = p_1(\xi_{hj}\#\xi_{h'j'}) = p_1(\xi_{hj}) + p_1(\xi_{h'j'}).$$

Thus these characteristic classes are linear in h and j. Next consider the effect of reversing the fiber orientation. Interpreting S^3 as the unit quaternions, this is equivalent to conjugation by the map $v \overset{g}{\mapsto} v^{-1}$. But $g^{-1}(f_{hj}(g(v))) = (u^h v^{-1} h^j)^{-1} = (u^{-j} v u^{-h})$, and so ξ_{hj} becomes ξ_{-j-h}. But reversing the fiber orientation is detected by e and not by p_1, and so $e(\xi_{hj}) = -e(\xi_{-j-h})$ while $p_1(\xi_{hj}) = p_1(\xi_{-j-h})$. Thus $e(\xi_{hj}) = k_1(h+j)\iota$ and $p_1(\xi_{hj}) = k_2(h-j)\iota$ for some yet undetermined constants k_i.

But we might as well evaluate these constants by considering the familiar bundle ξ_{10}. We apply the Gysin sequence to this bundle, yielding an exact sequence

$$H^3(E_{10}) \to H^0(S^4) \xrightarrow{m} H^4(S^4) \to H^4(E_{10})$$

where m is given by multiplication by $e(\xi_{10})$. But we already have shown that $E_{10} = S^7$, and so the first and last groups above are zero. Thus $e(\xi_{10})$ is a generator, and so $k_1 = \pm 1$, as desired.

Now we calculate k_2. Recall that $B_{10} = \mathbb{HP}^2 - D^8$. Notice that the map $i : B_{10} \hookrightarrow \mathbb{HP}^2$ induces an isomorphism $i^* : H^4(\mathbb{HP}^2) \to H^4(B_{10})$ by the exact cohomology sequence and excision (applied to $(\mathbb{HP}^2, B_{10})$). Similarly, notice that $\pi : B_{10} \to S^4$ induces an isomorphism in H^4 because it is a retraction onto the zero section. Let α denote the generator of $H^4(B_{10})$ and β of $H^4(\mathbb{HP}^2)$. We consider the tangent bundle TB_{10}, which is naturally the Whitney sum of the sub-bundles of those vectors parallel to the fiber and of those parallel to the 0-section. But the former bundle is induced via π from ξ_{10}, while the latter is induced from TS^4. Thus we have $TB_{10} \cong \pi^*(\xi_{10}) \oplus \pi^*(TS^4)$ and π lifts to a bundle map $\tilde{\pi}$:

$$\begin{array}{ccc} TB_{10} & \xrightarrow{\tilde{\pi}} & \xi_{10} \oplus TS^4 \\ \downarrow & & \downarrow \\ B_{10} & \xrightarrow{\pi} & S^4 \end{array}$$

(Note that $\pi^*(\xi_{10})$ really denotes the associated $\mathbb{R}^4$-bundle.) Because π may be covered by a bundle map, we may compute p_1 as

$$\begin{aligned} p_1(TB_{10}) &= \pi^*(p_1(\xi_{10} \oplus TS^4)) \\ &= \pi^*(p_1(\xi_{10}) + p_1(TS^4)) \quad \text{(Whitney Product Theorem)} \\ &= \pi^*(p_1(\xi_{10} + 0)). \end{aligned}$$

Therefore, because π^* and i^* are isomorphisms in H^4,

$$\begin{aligned} p_1(\xi_{10}) &= {\pi^*}^{-1}(p_1(TB_{10})) \\ &= {\pi^*}^{-1}(i^*(p_1(T\mathbb{HP}^2)) \\ &= {\pi^*}^{-1}(i^*(2\beta)) \\ &= {\pi^*}^{-1}(\pm 2\alpha) \\ &= \pm 2\iota, \end{aligned}$$

which reveals the value of k_2. Note that we have used the well-known calculations of $p_1(S^4)$ and $p_1(\mathbb{HP}^2)$ (see [MS]). ■

The Homeomorphism Class of E_{hj}

Having calculated the characteristic classes of E_{hj}, we now show that many of these bundles are homeomorphic to S^7. Our strategy, in light of the Generalized Poicaré Conjecture, is to compute the homology of these bundles.

Proposition 8.4 *If $h + j = 1$, then E_{hj} is homeomorphic to S^7.*

PROOF: According to Theorem 5.3, we need only show that such E_{hj} have the same homology as S^7. But we can simply apply the Gysin sequence of ξ_{hj},

$$\ldots \to H^i(S^4) \xrightarrow{m} H^{i+4}(S^4) \to H^{i+4}(E_{hj}) \to H^{i+1}(S^4) \to \ldots$$

from which it follows that

$$H^i(E_{hj}) = \begin{cases} 0 \text{ for } i \neq 0, 3, 4, 7 \\ \mathbb{Z} \text{ for } i = 0, 7. \end{cases}$$

For $i = 3, 4$ we obtain the sequence

$$0 \to H^3(E_{hj}) \to H^0(S^4) \xrightarrow{m} H^4(S^4) \to H^4(E_{hj}) \to 0$$

where m is multiplication by $e(\xi_{hj})$. But, according to Proposition 8.3, $e_{hj} = \pm\iota$ generates $H^4(S^4)$, and so $H^3(E_{hj}) = H^4(E_{hj}) = 0$. ■

Remark 8.1 *Our use of the Generalized Poincaré Conjecture in the proof above was, in fact, unnecessary. Milnor originally proved Proposition 8.4 by a simple application of Morse Theory. Nevertheless, Milnor's technique was somewhat ad hoc, and it does not generalize to higher dimensional exotic spheres. (To be fair to Milnor, in 1956 the h-cobordism theorem had not yet been proven!)*

A Smooth Invariant

We now concentrate our focus on those bundles E_{hj} where $h + j = 1$. We have already completed half the task of demonstrating that these manifolds are exotic spheres – *i.e.* we have shown that they are, in fact, spheres. As for the "exotic" part, we will use an invariant of the differential structure on such manifolds to distinguish some of them from the "real" S^7.

For convenience, we will re-index E_{hj}, $h+j = 1$ by the odd integers. For each odd k let $h = \frac{1+k}{2}$ and $j = \frac{1-k}{2}$ so that $h + j = 1$ and $h - j = k$. Let us redefine $\xi_k = \xi_{hj}$ by these formulae. In these terms, the important results so far may be summarized by $\xi_k = \pm\iota$, $p_1(\xi_k) = \pm 2k\iota$, and E_k is homeomorphic to S^7.

We will define an invariant λ for those smooth 7-manifolds E, homeomorphic to S^7, which bound 8-manifolds B. Of course, our constructions $E_k = \partial B_k$ clearly qualify. (In fact, all 7-manifolds bound 8-manifolds, according to [T].) We will define $\lambda(E)$ in terms of the signature $\sigma(B)$ and the first Pontrjagin class $p_1(B)$.

For the manifolds E under consideration, H^3 and H^4 vanish. Thus, by the exact cohomology sequence,

$$i^* : H^4(B, E) \to H^4(B)$$

is an isomorphism. Therefore ${i^*}^{-1}(p_1(B))$ is a well-defined cohomology class in in $H^4(B, E)$. Using the orientation class $\mu_B \in H^8(B, E)$, we define

$$q(B) = F({i^*}^{-1}(p_1(B)) =< \mu_B, \left({i^*}^{-1}(p_1(B))\right)^2 >,$$

where F is the intersection form (and the cup-product is understood). Finally, we define our invariant $\lambda(E)$ to be the residue class modulo 7 of $2q(B) - \sigma(B)$. Because σ and q are diffeomorphism invariants, so too is λ. Moreover, λ is well-defined with respect to choice of B, as we learn from

Proposition 8.5 *Let B and B' be two 8-manifolds such that $\partial B = \partial B' = E$. Then*

$$2q(B) - \sigma(B) \equiv 2q(B') - \sigma(B') \pmod 7.$$

Proof: We begin by forming the smooth manifold $C = B \cup B'$ by gluing along E. We choose an orientation μ_C consistent with μ_B, and therefore with $-\mu_{B'}$. Consider the following diagram in which the isomorphisms in the columns are derived from the exact cohomology sequences, and the isomorphism on the bottom row from the Mayer-Vietoris sequence:

$$\begin{array}{ccc} H^4(B,E) \oplus H^4(B',E) & \stackrel{h}{\leftarrow} & H^4(C,E) \\ \downarrow i^* \oplus i'^* & & \downarrow j \\ H^4(B) \oplus H^4(B') & \stackrel{k}{\leftarrow} & H^4(C) \end{array}$$

Since the other three arrows are isomorphisms, so too must be h. Let $\alpha \in H^4(B,E)$ and $\alpha' \in H^4(B',E)$ and define $\beta = jh^{-1}(\alpha \oplus \alpha') \in H^4(C)$. Then

$$< \mu_C, \beta^2 >=< \mu_C, jh^{-1}(\alpha^2 \oplus \alpha'^2) >=< \mu_B \oplus \mu_{B'}, \alpha^2 \oplus \alpha'^2) >=< \mu_B, \alpha^2 > - < \mu_{B'}, \alpha'^2 > . \quad (2)$$

In other words, the quadratic form of C is the direct sum of the form of B and minus the form of B'. Thus their signatures clearly satisfy

$$\sigma(C) = \sigma(B) - \sigma(B'). \quad (3)$$

Now we specify $\alpha = {i^*}^{-1}(p_1(B))$ and $\alpha' = (i'^*)^{-1}(p_1(B'))$. By naturality of the Pontrjagin class, $k^*(p_1(C)) = p_1(B) \oplus p_1(B')$, which implies that

$$jh^{-1}(\alpha \oplus \alpha') = p_1(C).$$

But then Equation 2 reads

$$< \mu_C, p_1^2(C) >=< \mu_B, \alpha^2 > - < \mu_{B'}, \alpha'^2 >,$$

or, equivalently,

$$q(C) = q(B) - q(B'). \quad (4)$$

Combining Equations 3 and 4 yields

$$(2q(B) - \sigma(B)) - (2q(B') - \sigma(B')) = 2q(C) - \sigma(C).$$

Hence, we will be finished once we show that $2q(C) - \sigma(C) \equiv 0 \pmod 7$. But, according the Hirzebruch signature formula,

$$\sigma(C) =< \mu_C, \frac{1}{45}(7p_2(C) - p_1^2(C)) >$$

and so by bilinearity

$$45\sigma(C) + q(C) = 7 < \mu_C, p_2(C) > \equiv 0 \pmod 7$$

or, equivalently,

$$2q(C) - \sigma(C) \equiv 0 \pmod 7. \blacksquare$$

Thus λ is indeed a well-defined, smooth manifold invariant. Hopefully we can use λ to distinguish between some of the E_k and S^7. To that end we will prove

Proposition 8.6 $\lambda(E_k) \equiv k^2 - 1 \pmod 7$.

Proof: The bundle E_k is the boundary of the associated disc bundle $B_k \xrightarrow{\pi} S^4$. Notice that, similar to our proof of Proposition 8.3, $\pi^* : H^4(S^4)) \to H^4(B_k)$ and $i^* : H^4(B_k, E_k) \to H^4(B_k)$ are both isomorphisms; the former because π is retraction onto the 0-section, and the latter by the exact cohomology sequence. Let us denote the generator of $H^4(B_k, E_k)$ by $\alpha = (i^*)^{-1}\pi^*(\iota)$. Choose an orientation μ for B_k such that $F(\alpha) =< \mu, \alpha \cup \alpha >= 1$. Then $\sigma(B_k) = 1$.

Now, exactly as in the proof of Proposition 8.3, π lifts to a bundle map $\tilde{\pi} : TB_k \to \xi_k \oplus TS^4$. Thus we may calculate $p_1(B_k) = \pi^*(\pm 2k\iota + 0) = \pm 2k\pi^*(\iota)$ and $q(B_k) = F((i^*)^{-1}(\pm 2k\pi^*(\iota))) = F(\pm 2k\alpha) = 4k^2$. Then $\lambda(E_k) = 2q(B_k) - \sigma(B_k) = 8k^2 - 1 \equiv k^2 - 1 \pmod 7$. ■

According to this calculation, $\lambda(S^7) = \lambda(E_1) = 0$. Therefore, in order to find an exotic 7-sphere we simply choose an odd k such that $k^2 - 1$ is not congruent to 0, modulo 7. Therefore, this construction furnishes us with three clearly distint exotic 7-spheres: $\lambda(E_3) = 1$, $\lambda(E_5) = 3$, and $\lambda(E_7) = 6$.

We should pause briefly to make a quick remark about characteristic classes and Milnor's construction of 7-spheres. Notice that characteristic classes are used *both* to show that E_k is homeomorphic to S^7, and to show that some of the E_k are not diffeomorphic to S^7. The Euler class performs the first task (along with the h-cobordism theorem), and the Pontrjagin class performs the second task (along with the Hirzebruch signature formula). In some sense, therefore, the precise factor we have used to distinguish some spheres as exotic is the different effect which reversing the fiber-orientation has on e and p_1, as seen in the proof of Proposition 8.3.

9 Higher Dimensional Homotopy Spheres

In this section we will outline some of the results which generalize Milnor's original discovery of exotic differential structures on S^7.

Groups of Homotopy Spheres

Definition 9.1 *Let θ^n denote the set of (equivalence classes of) oriented manifolds with the homotopy type of S^n, up to h-cobordism.*

Nota Bene: Henceforth, we only concern ourselves with $n \geq 5$. In these dimensions, we know that a homotopy n-sphere is homeomorphic to S^n. We also know that manifolds which are h-cobordant are diffeomorphic. Thus θ^n is naturally the set of distinct differentiable structures on the topological n-sphere. Therefore, by our previous results on exotic 7-spheres, we already know that $|\theta^7| \geq 4$.

The h-cobordism theorem not only allows us to recognize θ^n as the set of differentiable structures on S^n, but it also allows us to endow θ^n with a natural group structure.

Theorem 9.1 *θ^n is an abelian group under the connected-sum operation.*

Proof: Clearly the connected-sum of two homotopy n-spheres is again a homotopy n-sphere. The operation # is well-defined via a result of Cerf [C] (that the group of manifold diffeomorphisms acts transitively). It is clear that # is commutative and associative. It is likewise obvious that $S^n \in \theta^n$ acts as the identity. Let $E \in \theta^n$ be a homotopy n-sphere. We claim that $-E$, the homotopy sphere E with the opposite orientation, provides an inverse under the operation #. Notice that $E\# -E$ bounds the contractible manifold $M = (E - D^n) \times [0,1]$. But then ∂M is a homotopy n-sphere, and, in particular, it is simply-connected. Hence by Theorem 5.2, M is diffeomorphic to D^{n+1} and thus $E\# -E = S^n$. ■

In their paper [KM], Milnor and Kervaire undertake the task of computing the group θ^n. In the remainder of this section, we will summarize their general approach and results. The mathematics found in [KM] is a *tour de force* of cobordism and surgery theory. Although the details of their entire paper are beyond the scope of this thesis, we will at least summarize, in some detail, a large part of their analysis. For the record, all of the results in this section are drawn from [KM].

Definition 9.2 *A manifold is parallelizable if its tangent bundle is trivial.*

To start their analysis, Milnor and Kervaire focus on the important subgroup $bP_{n+1} \subset \theta^n$ consisting of those homotopy n-spheres which bound parallelizable manifolds. Notice that $S^n \in bP_{n+1}$. In fact, any homotopy sphere $E \in \theta^n$ which bounds a contractible manifold is, by Theorem 5.2, diffeomorphic to S^n (and is thus in bP_{n+1}.)

We should pause to note that bP_{n+1} is, in fact, a subgroup of θ^n. If two homotopy spheres $E_1 = \partial M_1$ and $E_2 = \partial M_2$ bound parallelizable manifolds, then $E_1 \# E_2$ bounds the (parallelizable) manifold obtained by attaching a 0-handle to $M_1 \coprod M_2$.

In light of the subgroup bP_{n+1}, there are two distinct reasons why a homotopy sphere $E \in \theta^n$ may be exotic (*i.e.* why it may not be diffeomorphic to S^n). First, E may not bound any parallelizable manifold whatsoever. To quantify this possibility, Milnor and Kervaire investigate the group θ^n/bP_{n+1}. Alternatively, E may be in bP_{n+1}, but it may not bound a contractible manifold. This possibility is analyzed via surgery and signatures. For the remainder of the thesis we will focus on the first of these two possibilities.

An Analysis of θ^n/bP_{n+1}

We will now summarize the method by which Milnor and Kervaire inspect the group θ^n/bP_{n+1}. In essence, they embed a homotopy sphere E into a higher-dimensional sphere, and they show (via stable-parallelizability) that its normal bundle is trivial. Then, they apply the Pontrjagin-Thom construction to realize θ^n/bP_{n+1} as a quotient of the stable homotopy group of spheres Π_n.

Definition 9.3 *A manifold M with tangent bundle $\tau(M)$ is stably parallelizable if $\tau \oplus \epsilon^1$ is a trivial bundle.*

Theorem 9.2 *Homotopy spheres are stably parallelizable.*

Proof (Sketch) We can trivialize $\tau(E)$ on each hemisphere of a homotopy sphere $E \in \theta^n$. The overlap map is thus a map $f : S^{n-1} \to \mathrm{SO}_n$. Therefore, the bundle $\tau(M) \oplus \epsilon^1$ will be trivial if the composition

$$S^{n-1} \xrightarrow{f} \mathrm{SO}_n \hookrightarrow \mathrm{SO}_{n+1}$$

is homotopic to a constant map. By the homotopy exact sequence, however, $\pi_{n-1}(\mathrm{SO}_{n+1}) = \pi_{n-1}(\mathrm{SO}_{n+k})$, $k \geq 1$. Therefore, by Bott's wonderful theorem on the stable homotopy of the classical groups [B], there is no obstruction to trivialization when $n \equiv 3, 5, 6, 7$ modulo 8.

For $n \equiv 1, 2$ modulo 8, Milnor and Kervaire demonstrate directly that there is no obstruction (by appealing to sophisticated results of Rohlin and Adams.) Finally, for dimensions congruent to 0 and 4 modulo 8, Milnor and Kervaire inspect the composition above to show that the obstruction is a multiple of the signature $\sigma(E)$, which is certainly zero. ■

Given a homotopy n-sphere (or any n-manifold) $E \in \theta^n$, we may always embed E into a high-dimensional sphere S^{n+k}, $k > n+1$. This embedding is unique up to smooth isotopy of S^{n+k}. Since we know that E is stably parallelizable, we might expect that the normal bundle of E in S^{n+k} has nice properties. In fact, we have the following

Lemma 9.1 *When E is embedded in S^{n+k}, $k > n+1$, the normal bundle of E is trivial.*

The proof rests on the following elementary lemma.

Lemma 9.2 *Let ξ be a vector bundle with fibre dimension k over an n-dimensional base space, $k > n$. If $\xi \oplus \epsilon^r$ is trivial, then so too is ξ itself.*

PROOF: By induction, we need only consider $\xi \oplus \epsilon^1$. In this case the isomorphism $\xi \oplus \epsilon^1 \cong \epsilon^{k+1}$ induces a bundle map f from ξ to the bundle ψ^k of oriented k-planes in $(k+1)$-space. The base space of ξ has dimension n while the base space of ψ^k is S^k, $k > n$. Thus f is homotopic to a constant map, and ξ is trivial. ■.

Returning to the proof of Lemma 9.1, let ν denote the normal bundle of $E \subset S^{n+k}$, and τ its tangent bundle. Since E is embedded in S^{n+k}, $\tau \oplus \nu$ is trivial. Clearly, then, $(\tau \oplus \epsilon^1) \oplus \nu$ is trivial. Since E is stably parallelizable, $\tau \oplus \epsilon^1$ is trivial, and we may apply Lemma 9.2 to deduce that ν is trivial as well. (Notice that the fibre dimension of ν equals $k > n$, as required by Lemma 9.2.) ■

Because of Theorem 9.2 and Lemma 9.1, once we embed E into S^{n+k}, $k > n+1$, we know that there exists a framing ϕ for its normal bundle. Given (E, ϕ), the Pontrjagin-Thom construction furnishes a map

$$p(E, \phi) : S^{n+k} \to S^k.$$

Moreover, the homotopy class of $p(M, \phi)$ is a well defined element in the stable homotopy group $\Pi_n = \pi_{n+k}(S^k)$. In fact, according to [P], there is a one-to-one correspondence between framed cobordism classes (E, ϕ) and elements $p(E, \phi) \in \Pi_n$. And this correspondence p is a homomorphism with respect to the disjoint union of framed cobordism classes, and the group operation in Π_n.

Notice that the element $p(E, \phi) \in \Pi_n$ depends upon a choice of framing ϕ for $\nu(E)$ in S^{n+k}. Allowing the possible framings to vary, we define

$$p(E) = \{p(E, \phi) \mid \phi \text{ is a framing for } \nu(E)\} \subset \Pi_n.$$

So defined, $p(E)$ will contain the zero element of Π_n if and only if E bounds a parallelizable manifold. Moreover the map p is well-behaved with respect to connected sums:

Lemma 9.3 *For $E_1, E_2 \in \theta^n$, we have the inclusion $p(E_1) + p(E_2) \subset p(E_1 \# E_2)$.*

PROOF: Note that $E_1 \coprod E_2$ is cobordant to $E_1 \# E_2$. A cobordism M is given by the boundary connected-sum

$$M = (E_1 \times [0,1], E_1 \times \{1\}) \;\#\; (E_2 \times [0,1], E_2 \times \{0\}).$$

Here M has the homotopy type of $E_1 \vee E_2$.

We may embed this entire cobordism M into $S^{n+k} \times [0,1]$, k large, such that $M \cap (S^{n+k} \times \{0\}) = E_1 \coprod E_2$, and $M \cap (S^{n+k} \times \{1\}) = E_1 \# E_2$. Given any normal framings ϕ_1 and ϕ_2 of $\nu(E_1)$ and $\nu(E_2)$, these framings extend naturally to all of $\nu(M)$. Thus we obtain (by restriction) a framing on $E_1 \# E_2$ denoted $\phi_1 \# \phi_2$. Hence, by the properties of the Pontrjagin-Thom construction listed above,

$$p(E_1 \# E_2, \phi_1 \# \phi_2) = p(E_1 \coprod E_2, \phi_1 \coprod \phi_2) = p(E_1, \phi_1) + p(E_2, \phi_2).$$

The result follows because the equality above holds for all framings ϕ_1 and ϕ_2. ■

Finally, we arrive at the following pleasing result.

Proposition 9.1 *The set $p(S^n) \subset \Pi_n$ is a subgroup. For any $E \in \theta^n$, $p(E)$ is a coset of $p(S^n)$. Thus the map $E \mapsto p(E)$ defines a homomorphism from θ^n to $\Pi_n / p(S^n)$.*

PROOF: We will apply Lemma 9.3 to three very trivial identities. First, we know that $S^n \# S^n = S^n$, and hence $p(S^n) + p(S^n) \subset p(S^n)$. In other words, $p(S^n)$ is a veritable subgroup of Π_n. Second, we know that for any $E \in \theta^n$, $S^n \# E = E$, and hence $p(S^n) + p(E) \subset p(E)$. In other words $p(E)$ is the union of cosets of $p(S^n)$. Finally, we know that $E \# -E = S^n$, and hence $p(E) + p(-E) \subset p(S^n)$. But the last inclusion reveals that $p(E)$ contains exactly one coset of $p(S^n)$ – because otherwise $p(S^n)$ would contain elements from two of its own cosets. ■

By this proposition, we have found a map $\overline{p} : \theta^n \to \Pi_n / p(S^n)$, whose kernel, as mentioned before, consists exactly of those homotopy spheres which bound parallelizable manifolds. In other words, we have recognized θ^n / bP_{n+1} as a quotient of Π_n,

$$\theta^n / bP_{n+1} \cong \Pi_n / p(S^n).$$

This gives us an immediate bound on the size of θ^n / bP_{n+1}. We know from Serre [Se] that the stable homotopy group of spheres is finite. Therefore, we have arrived at the following important result:

Theorem 9.3 *The group θ^n / bP_{n+1} is finite. Thus the number of exotic spheres which do not bound parallelizable manifolds is finite.*

Kervaire and Milnor proceed to calculate the exact size of θ^n / bP_{n+1} (in so far as $|\Pi_n|$ is known) by recognizing the subgroup $p(S^n)$ as the image of the Hopf-Whitehead J-homomorphism. By using surgery, then, Milnor and Kervaire are able to relate the size of coker J to the Bernouli numbers. This analysis eventually furnishes sharp answers for $|bP_{n+1}|$ (see below).

θ^n In The Large

In the second-half of their paper, Kervaire and Milnor investigate our "second reason" why a sphere may be exotic – *i.e.* they investigate those exotic spheres within bP_{n+1} itself. By use of surgery and signature analysis, they count the number of homotopy spheres which bound parallelizable, but not contractible manifolds.

In essence, the following approach is used. A manifold M is called *almost-parallelizable* if $M \backslash \{\text{point}\}$ is parallelizable. It is possible to characterize the signatures of smooth, almost parallelizable manifolds via their Pontrjagin classes and the Hirzebruch signature formula. Then, given $E = \partial M \in bP_{n+1}$, one constructs the closed *topological* manifold $M^* = M \cup_h D^{n+1}$, where $h : \partial D^{n+1} \to E$ is an orientation-preserving *homeo*morphism. Now, if E were in fact *diffeo*morphic to S^n, then M^* would be an almost-parallelizable *smooth* manifold, with a predictable signature. Therefore, in order to construct and count exotic homotopy n-spheres within bP_{n+1}, Milnor and Kervaire construct examples of M^* with other signatures (*i.e.* signatures other than the ones predicted).

Instead of attempting a hasty description of their actual surgery construction, we will simply report the orders of θ^n and bP_{n+1} as computed by Milnor and Kervaire via this method.

n	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
$\lvert bP_{n+1}\rvert$	1	1	?	1	1	1	28	1	2	1	992	1	1	1	8128	1	2	1
$\lvert\theta^n\rvert$	1	1	?	1	1	1	28	2	8	6	992	1	3	2	16256	2	16	16

Despite our previous construction of exotic 7-spheres as boundaries of disc-bundles with $p_1 \neq 0$, the table above shows that all homotopy 7-spheres, in fact, bound parallelizable manifolds. In other words, $\theta^7 = bP_8$. Milnor and Kervaire also demonstrate that bP_{n+1} is always finite cyclic ($n \neq 3$). In particular, this implies that all 28 homotopy 7-spheres may be obtained by repeated connected-sums with some generator exotic sphere. Unfortunately, I do not yet know which (if any) of the exotic spheres constructed in Section 8 generates θ^7.

We will conclude this thesis (as we began it) with a few remarks on the Poincaré Conjecture. Notice that the ? in the table above denotes our current ignorance as to how many h-cobordism classes of homotopy 3-spheres exist. If the Poincaré Conjecture were true, however, there would exist only one homotopy 3-sphere, S^3 itself. Thus, the Poincaré Conjecture would imply that $? = 1$ in the table above. Moreover, according to Siebenmann and Kirby, the PL, topological, and smooth categories are all equivalent in dimension three. Thus the Poincaré Conjecture would imply a unique homeomorphism class of homotopy 3-spheres, as well as a unique diffeomorphism class.

Ironically, the sophisticated surgery and cobordism theory which has untangled the mysteries of higher-dimensional differentiable structures has failed to answer simple questions about homotopy and homeomorphism in dimension three. It seems that Poincaré's original inquiry – despite the quantity of successful mathematics it has inspired – must await new tools for its resolution.

A Characteristic Classes

In general, we have assumed basic knowledge of characteristic classes. In this appendix, we will highlight the results which we have required for applications to exotic spheres. Most of this presentation will follow [MS].

In this section ξ will denote a vector bundle over base space B, and E will denote its total space. The number n will denote the fiber dimension of ξ. Throughout, however, there will be analogous results for the associated S^{n-1}, D^n, and SO_n bundles.

Given a vector bundle ξ, let E_0 denote E with its zero section removed. (For the associated D^n bundle, E_0 denotes the S^{n-1} bundle ∂E).

We start with the following absolutely fundamental result of Thom. Let E be the total space of an oriented n-plane bundle over B. In other words, for each fiber F we have chosen a preferred generator $u_F \in H^n(F, F_0; \mathbb{Z})$. Then we have the following result (with coefficients taken in $\mathbb{Z}$) which we will state without proof.

Theorem A.1 (Thom) *$H^i(E, E_0) = 0$ for $0 < i < n$, and $H^n(E, E_0)$ is generated by the unique cohomology class u whose restriction to each fiber (F, F_0) agrees with u_F. Moreover, the map $y \mapsto y \cup u$ maps $H^j(E)$ isomorphically onto $H^{j+n}(E, E_0)$ for each integer j.*

The cohomology class u will be called the *Thom class* of the vector bundle ξ. Alternatively, as a result of this theorem, we may redefine a bundle orientation as a choice of generator $u \in H^n(E, E_0)$.

Nota bene: On the other hand, the bundle projection $\pi : E \to B$ induces an isomorphism $\pi^* : H^j(B) \cong H^j(E)$, since the zero section embeds B as a deformation retract of E. Thus we obtain the *Thom isomorphism* $H^j(B) \xrightarrow{\cong} H^{j+n}(E, E_0)$ which is the composition:

$$H^j(B) \xrightarrow{\pi^*} H^j(E) \xrightarrow{\cup u} H^{j+n}(E, E_0).$$

Definition A.1 *The euler class of a vector bundle ξ is the cohomology class $e(\xi) \in H^n(B, \mathbb{Z})$ which corresponds to $u|_E$ under the canonical isomorphism $\pi^* : H^j(B) \cong H^j(E)$.*

In other words, given the restriction map $\rho : E \to (E, E_0)$, $e(\xi)$ is defined to be $(\pi^*)^{-1}\rho^*(u)$. In such circumstances, we often use the notation $|_E$ for ρ^*.

There are a few elementary properties of the euler class whose proofs will will also omit. First, it is natural with respect to orientation-preserving bundle maps (and thus it is zero on trivial bundles). Second, it changes sign when the orientation of ξ is flipped. Finally, $e(\xi_1 \coprod \xi_2) = e(\xi_1) \oplus e(\xi_2)$ when the disjoint union has a compatible orientation.

The most useful application of the euler class will be to derive the Gysin sequence. As usual, let E be the total space of a vector bundle ξ over B.

Proposition A.1 *There is an exact sequence*

$$\ldots \to H^i(B) \xrightarrow{\cup e(\xi)} H^{i+n}(B) \xrightarrow{\pi_0^*} H^{i+n}(E_0) \to H^{i+1}(B) \to \ldots$$

where the first map denotes cup product with $e(\xi)$, and π_0 denotes the restriction of the projection π to $E_0 \subset E$.

Proof: We start with the cohomology exact sequence for the pair (E, E_0)

$$H^i(E, E_0) \to H^i(E) \to H^i(E_0) \xrightarrow{\delta} H^{i+1}(E, E_0) \to \ldots$$

Using the isomorphism $\cup u$ of Thom, however, we substitute $H^{i-n}(E)$ for $H^i(E, E_0)$ to obtain

$$H^{i-n}(E) \xrightarrow{g} H^i(E) \to H^i(E_0) \to H^{i+1-n}(E) \to \ldots,$$

where $g(x) = (x \cup u)|_E = x \cup (u|_E)$. Now we use the isomorphism π^* to substitute $H^*(B)$ for $H^*(E)$. But, under this substitution, the cohomology class $u|_E$ corresponds to the Euler class in $H^n(B)$, yielding

$$H^{i-n}(B) \xrightarrow{\cup e(\xi)} H^i(B) \to H^i(E_0) \to H^{i+1-n}(B) \to \ldots \blacksquare$$

Now we provide definitions of Chern and Pontrjagin classes. We start with a complex n-plane bundle ω with total space E over a base space B. We construct a complex $(n-1)$-plane bundle ω_0 over the space $E_0 = E \setminus \{\text{zero section}\}$. Given a point in E_0, specified by a fiber F and nonzero vector $v \in F$, we define the covering fiber of ω_0 to be $F/\mathbb{C}v$ (or, when a Hermitian metric is defined, the orthogonal complement of v).

Now we apply the Gysin sequence to the real vector bundle underlying ω_0 to obtain the exact sequence

$$\ldots \to H^{i-2n}(B) \xrightarrow{\cup e(\omega)} H^i(B) \xrightarrow{\pi_0^*} H^i(E_0) \to H^{i-2n+1}(B) \to \ldots.$$

Since the groups $H^j(B)$ are zero for $j < 0$, we find the isomorphism $\pi_0^* : H^i(B) \to H^i(E_0)$ for $i < 2n-1$. Using this isomorphism we can make the following

Definition A.2 *The Chern classes $c_i(\omega) \in H^{2i}(B)$ are defined by induction on the complex dimension n of ω. The top Chern class, $c_n(\omega)$ equals the Euler class of the underlying real vector bundle. For $i < n$ we define*

$$c_i(\omega) = \pi_0^* c_i(\omega_0).$$

For $i > n$, we define $c_i(\omega)$ to equal zero.

The formal sum $c(\omega) = 1 + \sum_{0 \le i \le n} c_i(\omega)$ is called the *total Chern class* of ω. Such formal sums may be multiplied as are polynomials (or any graded commutative ring).

Aside from naturality, the Chern class satisfies the following elementary property.

Proposition A.2 *Let ω and ϵ^k be complex bundles over B, then $c(\omega \oplus \epsilon^k) = c(\omega)$.*

We define the Pontrjagin classes of a real vector bundle by using the Chern classes of its complexification.

Definition A.3 *For a real n-plane bundle ξ, the ith Pontrjagin class $p_i(\xi) \in H^{4i}(B; \mathbb{Z})$ is defined to be the cohomology class $(-1)^i c_{2i}(\xi \otimes_{\mathbb{R}} \mathbb{C})$.*

Similarly, the total Pontrjagin class is defined to be the sum

$$p(\xi) = 1 + p_1(\xi) + \ldots + p_{[\frac{n}{2}]}(\xi).$$

As do their Chern counterparts, the Pontrjagin classes are natural with respect to bundle maps. The Pontrjagin classes also are unaffected by adding trivial dimensions to a bundle. Moreover, modulo elements of order two, the total Pontrjagin class is multiplicative with respect to the Whitney sum of vector bundles.

References

[B] Bott, R. *The stable homotopy of the classical groups*, Ann. of Math. **70** (1959) 313-337.

[C] Cerf, J. *Topologie de certains espaces de plongements*, Bull. Soc. Math. France **89** (1961) 227-380.

[KM] Kervaire, M. and J. Milnor. *Groups of homotopy spheres I*, Ann. of Math. **77** (1963) 504-537.

[K] Kosinkski, A. (1993) *Differential Manifolds*. Academic Press, New York.

[MS] Milnor, J. and J. Stasheff. (1974) *Characteristic Classes*. Princeton University Press. Princeton.

[M1] Milnor, J. (1969) *Morse Theory*. Princeton University Press, Princeton.

[M2] Milnor, J. (1965) *Lectures on the h-cobordism Theorem*. Princeton University Press, Princeton.

[M3] Milnor, J. *On manifolds homeomorphic to the 7-sphere*, Ann. of Math. **64** (1956) 399-405.

[P] Pontrjagin, L. *Smooth manifolds and their applications in homotopy theory*, A.M.S. Translations, series 2, 11, 1-114.

[Re] Reeb, G. *Sur certain propriétés topologiques des variétés feuilletées*, Actual. Sci. Industr. 1183, Paris (1952) 91-154.

[Roe] Roe, J. (1997) *Notes on Surgery Theory*. mimeo

[RS] Rourke, C. and B. Sanderson. (1972) *Introduction to piecewise-linear Topology*. Springer Verlag. Berlin.

[Se] Serre, J.-P., *Homologie singulière des despaces fibrés*, Ann. of Math. **54** (1951), 425-505.

[S] Smale, S. *Generalized Poincaré Conjecture in dimensions greater than four*, Ann. of Math. **74** (1961) 391-406.

[St] Steenrod, N. (1951) *Topology of Fibre Bundles*. Princeton University Press, Princeton.

[T] Thom, R. *Quelques propriétés globale des variétés différentiables*, Comment. Math. Helv. **28** (1954) 17-86.