Spiros Adams - Flor
30/01/09

Geometry Club Talk: The h-cobordism theorem

Before stating the h-cobordism theorem, it will be useful to recap a few concepts

Defⁿ: Let M^n, N^n be two oriented n-dimensional manifolds. We say that M is cobordant to N if there exists an oriented $(n+1)$-dimensional manifold W^{n+1}, whose boundary is the disjoint union $\bar{M} \amalg N$. $\bar{M}$ denotes M with the opposite orientation. W is then called a cobordism between M and N.

Example: A cobordism between S^1 and $S^1 \amalg S^1$:

Defⁿ: We call $M \times [0,1]$ the trivial cobordism:

Defⁿ: A cobordism (W, M, M') is called an h-cobordism if the inclusions $M \hookrightarrow W$, $M' \hookrightarrow W$ are homotopy equivalences. In the case that all of W, M, M' are simply connected, this is equivalent to the condition $H_*(W, M; \mathbb{Z}) = 0$.

Defⁿ: Let $M, N \subset W$ be submanifolds. We say that M intersects N transversely written $M \pitchfork N$, if at all points $x \in M \cap N$ we have $T_xM \oplus T_xN = T_xW$ in the case where $\dim M + \dim N = \dim W$, and $M \cap N = \emptyset$ if $\dim M + \dim N \neq \dim$

Examples: transverse intersection   non-transverse intersection:

$W = \mathbb{R}^3$

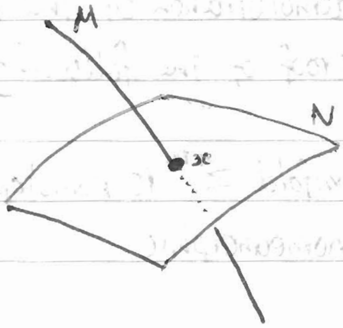

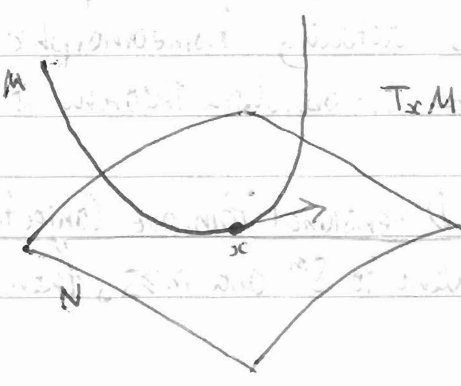

$T_xM \subset T_xN$

Def^n: Let $M, N \subset W$ be submanifolds such that $M \pitchfork N$, and further assume M, N are compact, then we may define an <u>intersection number</u> $M.N$ as follows: To each $x \in M \cap N$ assign a sign $\varepsilon(x) \in \{\pm 1\}$ and then set $M.N := \sum_{x \in M \cap N} \varepsilon(x)$. $\varepsilon(x)$ is defined by choosing orientations for M, N and W, the orientations for M and N at $x \in M \cap N$ give us an orientation there for W and we set $\varepsilon(x) = 1$ if this agrees with our chosen orientation for W and -1 otherwise.

Example:

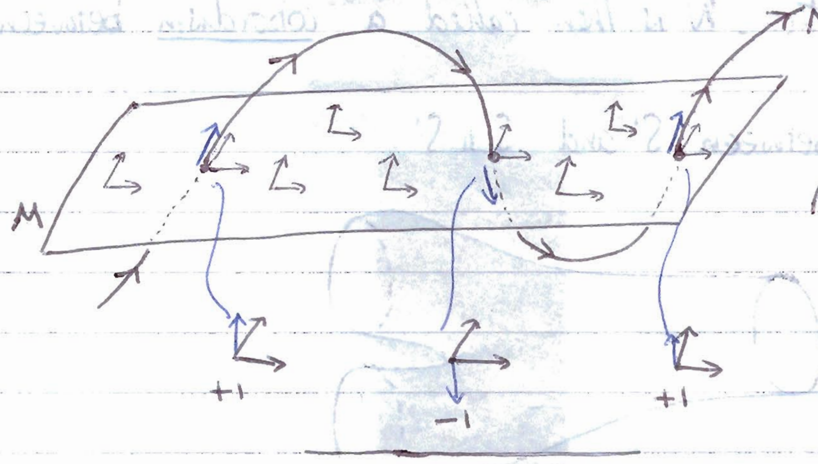

$M.N = +1 -1 +1 = +1$

Now we are in a position to state the theorem:

h-cobordism theorem: Let M^m and N^m be compact, simply connected, oriented m-dimensional manifolds which are h-cobordant through the simply connected $(m+1)$-dimensional manifold W^{m+1}. If $m \geq 5$, then there is a diffeomorphism $W \cong M \times [0,1]$, which may be chosen to be the identity from $M \subset W$ to $M \times \{0\} \subset M \times [0,1]$. In particular M and N are diffeomorphic.

This is a fundamental result in high dimensional manifold theory, proven in 1961 by Stephen Smale for which he was awarded a Fields Medal. This theorem helps us to answer the important question "When are two homotopy equivalent spaces actually homeomorphic/diffeomorphic?". A demonstration of the power of the h-cobordism theorem is its almost immediate proof of the following:

High Dimensional Poincaré Conjecture: If a smooth m-manifold Σ^m is homotopy equivalent to S^m and $m \geq 5$, then Σ^m and S^m must be homeomorphic.

<u>Proof</u>: There are two parts to the proof, it is possible to prove for $m \geq 6$ one way and for $m=5$ another way. So first suppose that $m \geq 6$.

We have our manifold $\Sigma^m \simeq S^m$. Cutting two m-disks D^m out of Σ^m we are left with an h-cobordism between two copies of S^{m-1} (the boundaries of the two disks we cut out). Hence by the h-cobordism theorem there exists a diffeomorphism $f: \Sigma \setminus (D_1^m \cup D_2^m) \longrightarrow S^{m-1} \times [0,1]$ such that the restriction to ∂D_1^m is the identity map to $S^{m-1} \times \{0\}$ say.

We now attempt to extend our diffeomorphism over the two disks we removed to obtain a global diffeomorphism $\Sigma^m \cong S^m$.

Since $f|_{\partial D_1^m} = id$ we may extend f to $\Sigma \setminus D_2^m$ by taking the identity map on D_1^m and patching f with id we obtain another diffeomorphism. So now we need to extend to D_2^m. By the h-cobordism theorem we know that $f|_{\partial D_2^m}$ is a diffeomorphism onto $S^{m-1} \times \{1\}$, so the problem is whether we may extend a diffeomorphism $S^{m-1} \longrightarrow S^{m-1}$ to a diffeomorphism $D^m \longrightarrow D^m$. Sadly this cannot always be done, however simply extending radially by

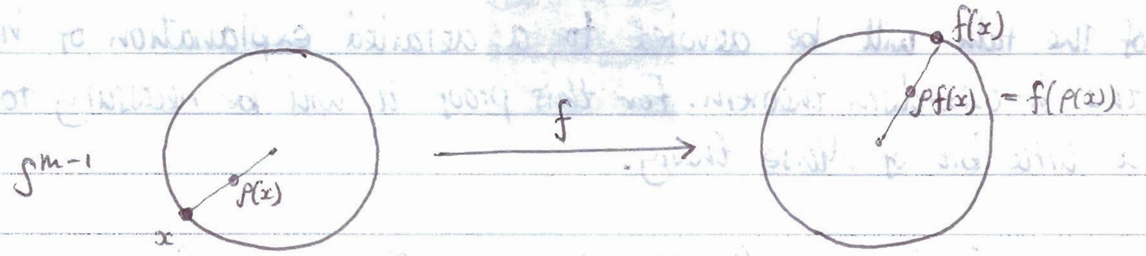

gives us a diffeomorphism except at the centre of D^m, however this extension gives a homeomorphism, hence patching with f we get a homeomorphism of Σ with S^m.

<u>Idea</u>:

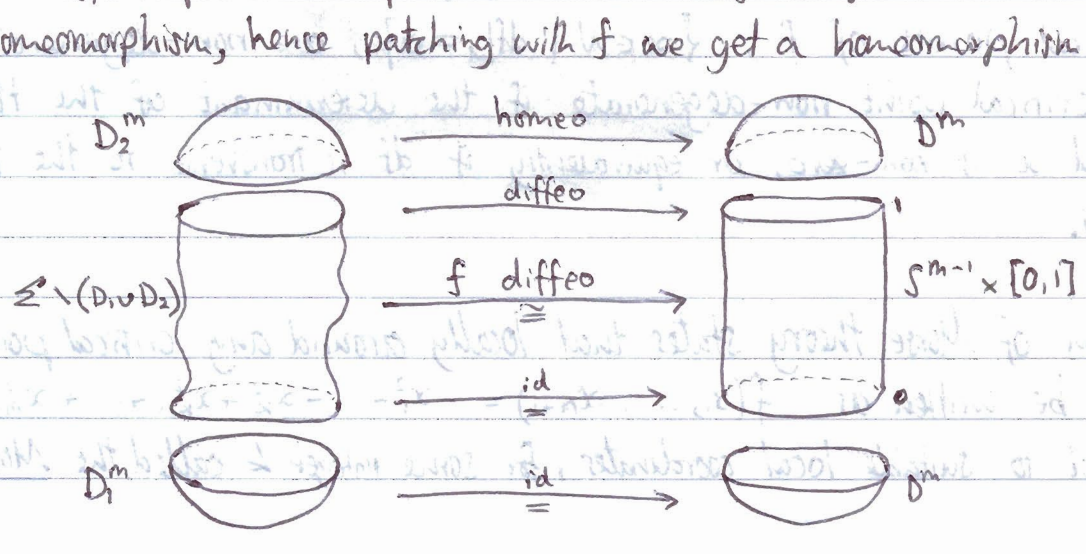

Note that this approach does not work in the case $m=5$ since after cutting out two D^5s we are left with an h-cobordism between S^4 and S^4 and the h-cobordism theorem is for h-cobordant 5-manifolds or higher. For this reason we adopt a different strategy.

Let $m=5$, $\Sigma^5 \simeq S^5$. All 5-manifolds necessarily bound a 6-manifold (provided of course that the 5-manifold is closed). Hence $\exists V^6$ s.t $\partial V = \Sigma$. Further V is necessarily contractible since $\Sigma \simeq S^5$. That is to say, we may extend the homotopy equivalence $\partial V \simeq \partial D^6$ to an homotopy equivalence $V \simeq D^6 \simeq *$. Now, cutting out a D^6 from the interior of V we are left with an h-cobordism between $\partial D^6 = S^5$ and Σ^5, hence we may apply the h-cobordism theorem (everything is simply connected!) and we deduce that in fact Σ must be diffeomorphic to S^5. We have proven something much stronger than we set out to; a corollary of this proof is that there do not exist any exotic 5-spheres!

QED

The rest of the talk will be devoted to a detailed explanation of how to prove the h-cobordism theorem. For this proof it will be necessary to introduce a little bit of Morse theory.

Def^n: A Morse function on (W, M, M') is a differentiable function $f : W \to [a,b]$ such that $f^{-1}(a) = M$, $f^{-1}(b) = M'$ and such that all the critical points of f, $\{x \in W \mid df|_x = 0\}$, are non-degenerate. We call a critical point non-degenerate if the determinant of the Hessian matrix at x is non-zero, or equivalently if df is transverse to the zero section of T^*W.

A theorem of Morse theory states that locally around any critical point x, f may be written as $f(x_1, \dots, x_{m+1}) = -x_1^2 - \dots - x_k^2 + x_{k+1}^2 + \dots + x_{m+1}^2 + \text{const}$ with respect to suitable local coordinates, for some integer k called the Morse index of x.

Another way to think of the morse index is as the dimension of the negative eigenspace of the Hessian of f at x. Intuitively, the index is the number of linearly independant directions out of x along which the value of f decreases.

Examples: Firstly, an example of degenerate critical points is a cylinder on its side with f being the height function (projection onto z-axis):

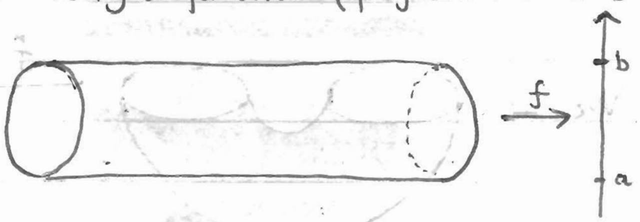

$f^{-1}(b)$ is the top line of the cylinder, $f^{-1}(a)$ the bottom line. All points in $f^{-1}(a) \cup f^{-1}(b)$ are degenerate critical points.

Next the canonical example of a morse function, the height function on a torus:

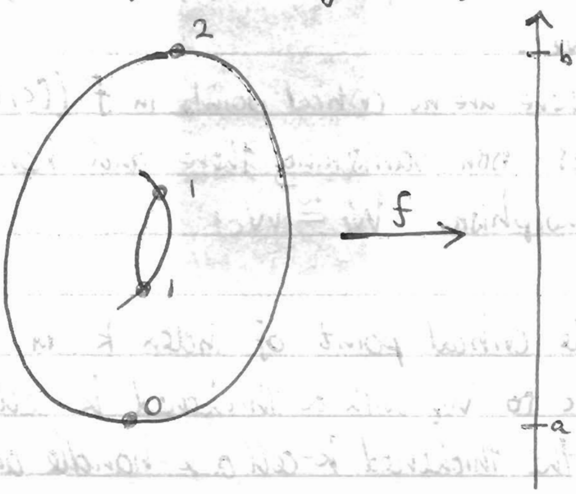

Four (non-degenerate) critical points are marked on the torus with their morse indices next to them, these are all the critical points for this function. Note, a local maximum has index $= \dim(W) = 2$ in this case since the value of f decreases in all directions, similarly a local minimum has index zero. Above the two points of index 1 are saddle points, f increases in one direction, and decreases in the other orthogonal direction.

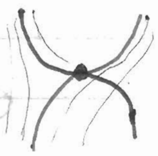

<u>Def^n</u>: Let f be a Morse function on (W, M, M') with $f^{-1}(a) = M$, $f^{-1}(b) = M'$, then we define the <u>ascending cobordism</u> to be $W_c = f^{-1}([a,c])$ for $c \in [a,b]$. W_c is a cobordism between $f^{-1}(a) = M$ and $f^{-1}(c)$

<u>Example</u>: the torus & height function:

$W_c =$

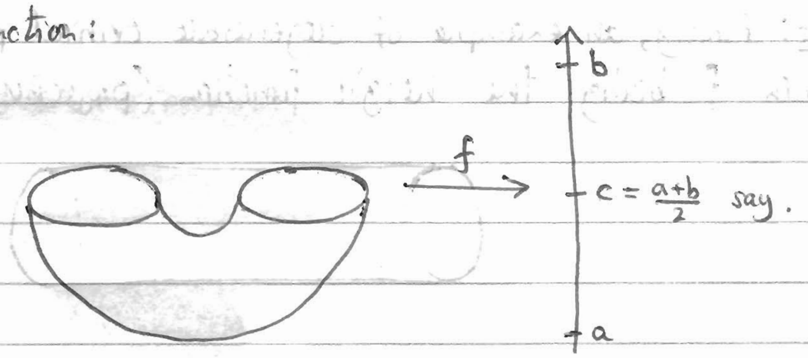

The ascending cobordism has some very nice topological properties:

1) If there are no critical points in $f^{-1}([c,c'])$, then W_c is diffeomorphic to $W_{c'}$.
<u>Justification</u>: Since there are no critical points in $f^{-1}([c,c'])$, the gradient vector field of f is non-vanishing there and hence may be integrated to yield a diffeomorphism $W_c \cong W_{c'}$

2) If there is a single critical point of index k in $f^{-1}([c,c'])$, then $W_{c'}$ is diffeomorphic to W_c with a thickened k-cell $D^k \times D^{(m+1)-k}$ attached. We call the thickened k-cell a k-handle and say that we obtain $W_{c'}$ from W_c by attaching a k-handle.
<u>Justification</u>: This is a famous result from Morse theory - see any good book on the subject!

As an example consider again the torus, starting with $f^{-1}(a) =$ point we attach a thickened 0-cell, i.e. a $D^0 \times D^2$, this gives us a disk

The ascending cobordism remains topologically the same up to the first saddle point x_1 where a 1-handle $D^1 \times D^1$ is attached:

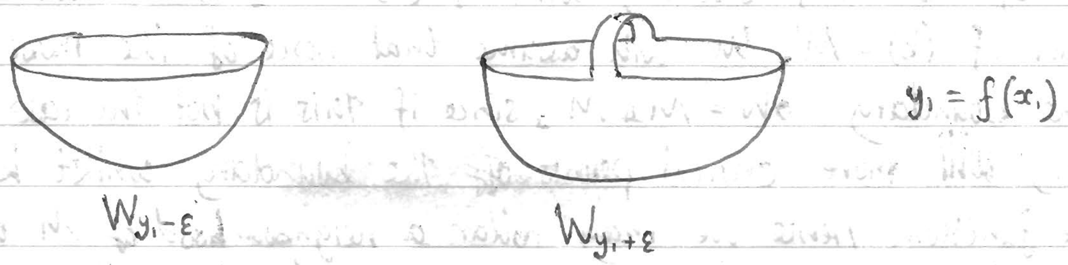

Again the ascending cobordism remains topologically unchanged up to the second saddle point x_2 where another 1-handle is attached:

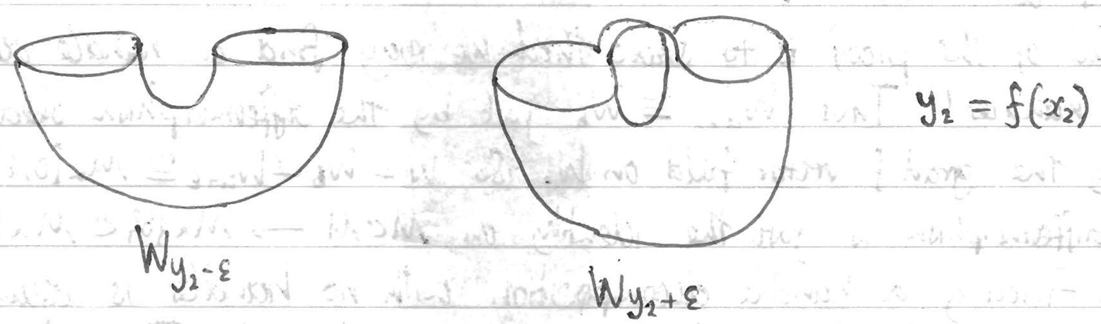

Finally at the last critical point we attach a 2-handle $D^2 \times D^0$

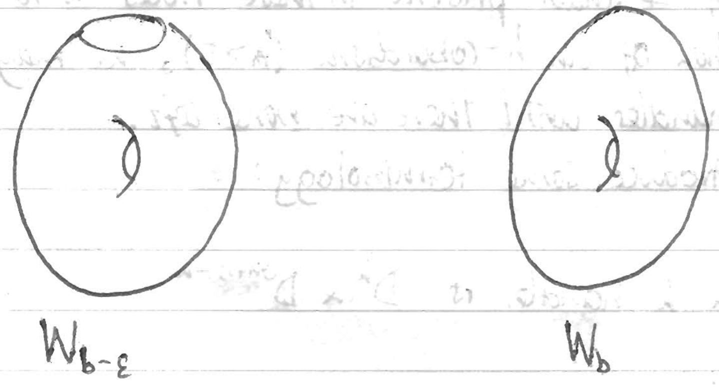

One final important fact we will use from Morse theory is the fact that Morse functions exist, or rather given a cobordism (W, M, M') there exists a Morse function on W such that $f^{-1}(a) = M$, $f^{-1}(b) = M'$. In fact Morse functions form an open and dense subset of the space of all functions on W. This will allow us to perturb a Morse function and still have a Morse function.

So given an h-cobordism (W, M, M'), we may find a Morse function f on W. Then considering the ascending cobordisms of f, the Morse function gives us a recipe of how to, starting with $f^{-1}(a) = M$, attach handles until we obtain $f^{-1}(b) = M'$. We will assume that none of the morse critical points lie on the boundary $\partial W = M \amalg M'$, since if this is not the case perturbing f slightly will move critical points off the boundary whilst keeping f as a morse function. Hence we may collar a neighbourhood of M with no critical points in it, diffeomorphic to $M \times [0, \varepsilon]$. We then attach a series of handles to $M \times [0, \varepsilon]$ to obtain W. We call this a <u>handle decomposition</u> of W. Note there will be an i-handle in the handle decomposition for every critical point of f with index i.

The goal of the proof is to show that we may find a handle decomposition with no handles! Thus $W_{a+\varepsilon} \cong W_b$ just by the diffeomorphism obtained from integrating the grad f vector field on W. So $W \cong W_b \cong W_{a+\varepsilon} \cong M \times [0, \varepsilon] \cong M \times [0, 1]$ and our diffeomorphism is just the identity on $M \subset M \longrightarrow M \times \{0\} \subset M \times [0, 1]$ as required. Finding a handle decomposition with no handles is equivalent to finding a Morse function on W with no critical points. The latter approach is the one taken by Milnor in "Lectures on the h-cobordism theorem".

The technique of proof I shall present in these notes is to show that given a handle decomposition of an h-cobordism $(m \geq 5)$, we may reorder, slide, exchange and cancel handles until there are none left.

For this I must introduce some terminology:

In $(m+1)$-dimensions, a k-handle is $D^k \times D^{(m+1)-k}$

We call $S^{k-1} \times \{0\}$ the attaching sphere
We call $\{0\} \times S^{m-k}$ the belt sphere
We call $D^k \times \{0\}$ the core

e.g. a 1-handle in $\mathbb{R}^3$: $D^1 \times D^2$

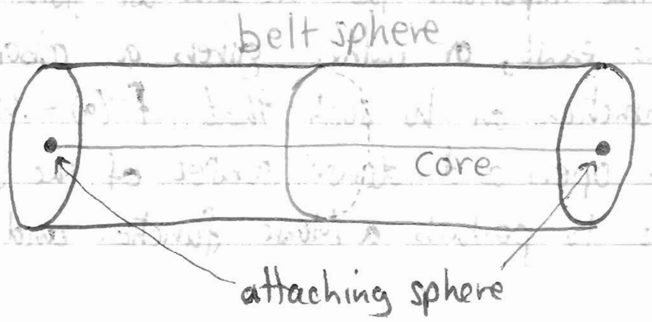

The process of attaching a handle looks as follows

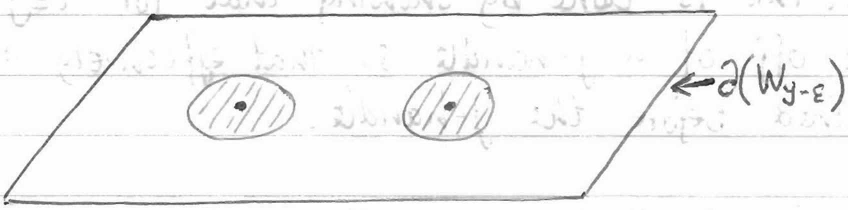

Find an embedded $S^{k-1} \times D^{(m+1)-k}$ in $\partial(W_{y-\varepsilon})$ to which we attach our k-handle $D^k \times D^{(m+1)-k}$. Note $\partial(D^k \times D^{(m+1)-k}) = S^{k-1} \times D^{(m+1)-k} \cup D^k \times S^{m-k}$. So we glue the k-handle by identifying the first boundary component with our embedded $S^{k-1} \times D^{(m+1)-k}$. The result of this procedure is to replace $S^{k-1} \times D^{(m+1)-k}$ in $\partial(W_{y-\varepsilon})$ with $D^k \times S^{m-k}$:

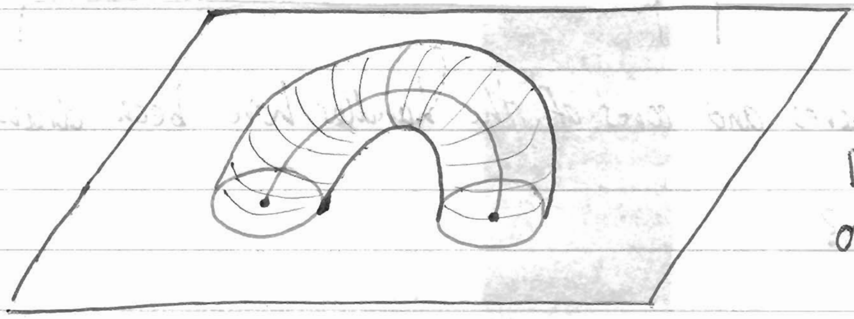

The process of cutting out $S^{k-1} \times D^{(m+1)-k}$ and gluing in $D^k \times S^{m-k}$ is called k-surgery on $\partial(W_{y-\varepsilon})$

We have a schematic picture of the handle decomposition as follows

lots of handles attached ontop of M and ontop of each other

Top boundary component $= f^{-1}(b) = M'$

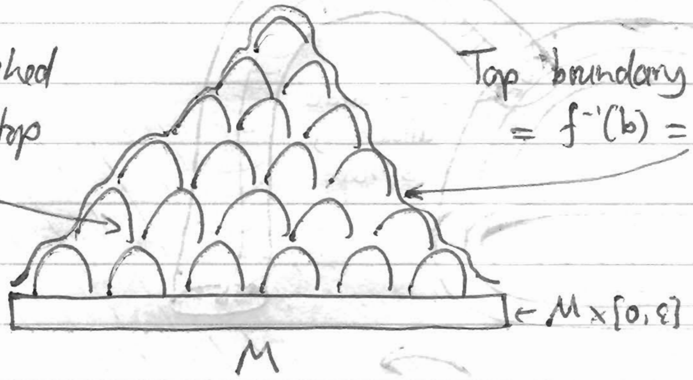

Each k-handle represents a k-surgery on the top boundary component, so the handle decomposition provides us with instructions for how to obtain M' from M via a sequence of surgeries.

It is possible to simplify the schematic picture above by reordering handles.

Claim: We may rearrange handles so that they are attached in order of ascending index. This is done by showing that for $i \leq j$ we may slide an i-handle off of a j-handle so that effectively it may be considered as attached before the j-handle.

Example: Sliding a 1-handle off of a 2-handle in 3-dimensions:

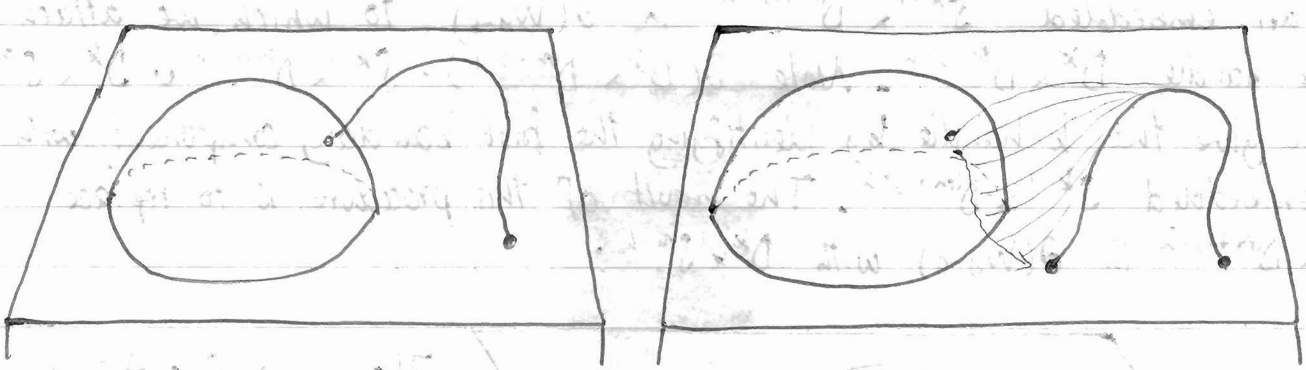

(Here only the attaching spheres and cores of the handles have been drawn)

So, when can we do this?

Consider a j-handle attached first and then an i-handle ontop:

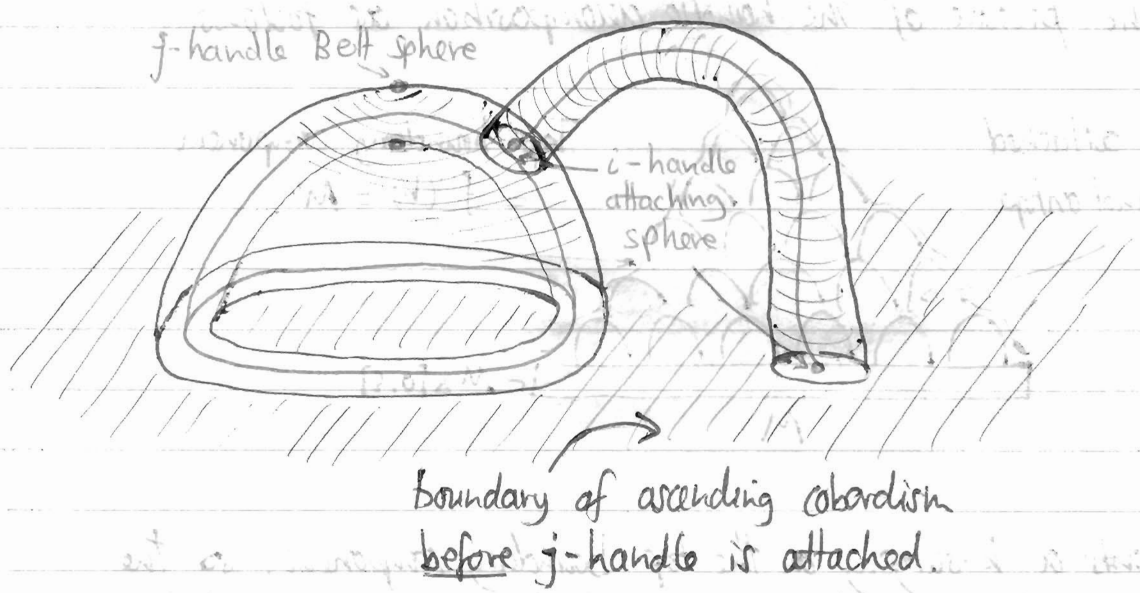

After attaching the j-handle, we get a boundary component $D^j \times S^{m-j}$ to which the i-handle is attached (or at least part of the i-handle attaching sphere sits in this $D^j \times S^{m-j}$).

If the i-attaching sphere does not intersect the j-belt sphere, then we may push the i-handle off of the j-handle radially.

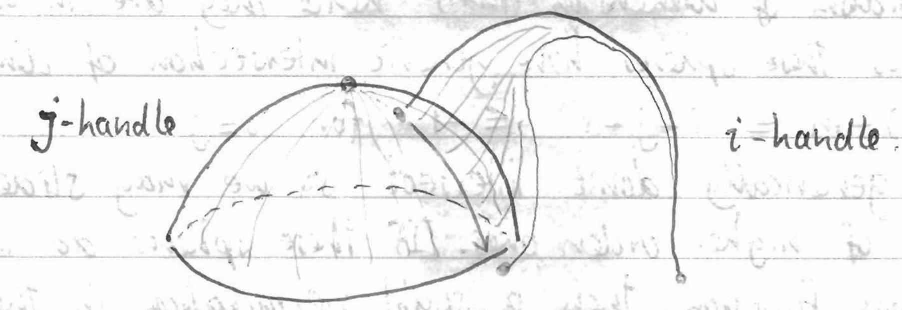

We have to be slightly careful: it is possible that the attaching sphere of the i-handle and the belt sphere of the j-handle do not intersect but that $(S^{i-1} \times D^{m+1-i}) \cap (\{0\} \times S^{m-j}) \neq \emptyset$, i.e. that the fattened attaching sphere intersects the belt sphere:

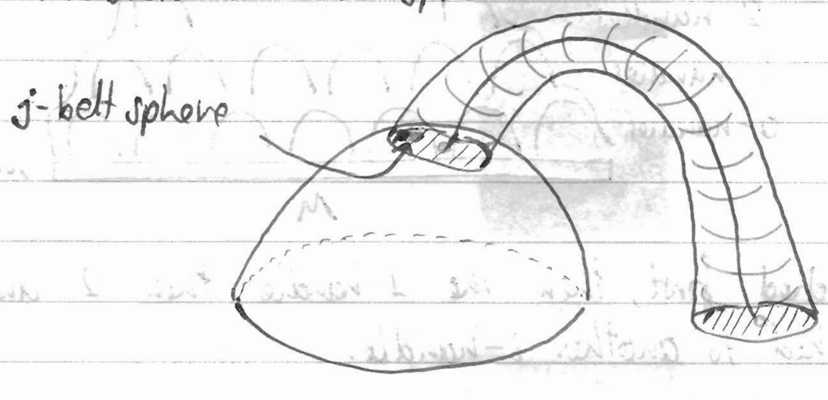

If this is the case we perturb the i-handle making it thinner until it no longer intersects the belt sphere of the j-handle:

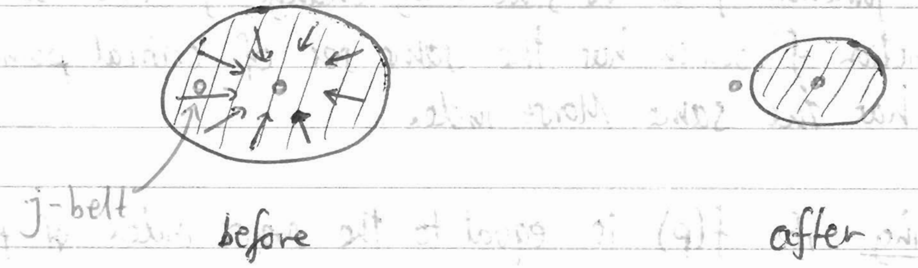

After this is done once again we may push the whole thing off radially. Hence, whether we may push the i-handle off is determined by whether the attaching S^{i-1} intersects the belt S^{m-j}. In what dimensions can we guarantee that these spheres generically do not intersect?

Both spheres are thought of as living in a boundary component of some ascending cobordism of dimension $(m+1)$ hence they are in a manifold of dimension m. Thus these spheres have generic intersection of dimension

$$(i-1) + (m-j) - m = i-j-1 \leq -1 \quad \text{for } i \leq j$$

So for $i \leq j$ they generically don't intersect so we may slide lower order handles off of higher order ones. If these spheres do intersect for our chosen morse function, then a slight perturbation in the function will remove the intersection.

So we may assume now that our schematic looks like

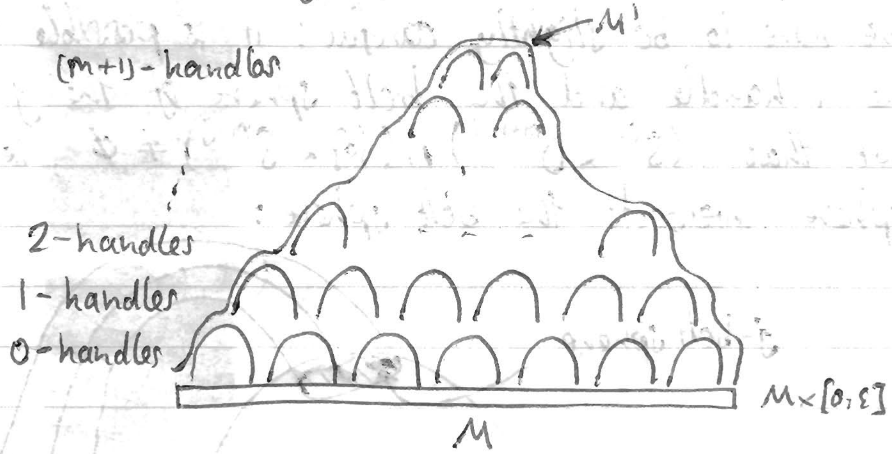

The 0-handles are all attached first, then the 1-handles, then 2 and so on, and no i-handle is attached to another i-handle.

It is possible to make a different argument to achieve the same result by using the following theorem from Morse Theory:

<u>Theorem</u>: Given any Morse function f on W, we may change f into a self-indexing Morse function f' which has the same set of critical points and each critical point has the same Morse index.

<u>Def^n</u>: f is <u>self indexing</u> if $f(p)$ is equal to the morse index of f at p for all critical points p.

Using the theorem we obtain a self-indexing Morse function, and then the ascending cobordism of this yields the required handle decomposition as in the schematic diagram..

<u>Note</u>: We have handle orders ranging from 0 to $m+1$; a zero-handle is $D^0 \times D^{m+1}$ attached via $S^{-1} \times D^{m+1}$ where we adopt the convention $S^{-1} := \emptyset$. So in some sense a 0-handle is not attached to $M \times [0,\varepsilon]$, but rather just disjoint unioned with it

$M \times [0,\varepsilon]$

$M \times [0,\varepsilon]$ with 'attached' 0-handle

An $(m+1)$-handle attached via $S^m \times D^0$ is filling in a hole whose boundary is that S^m.

~

Given our handle decomposition, we may pass from the geometry to algebra, by considering the handle chain complex:

let $\quad C_k = \mathbb{Z}\{k\text{-handles } h_\alpha^k\}$ $\quad$ free abelian group generated by the k-handles

define $\quad \partial_k : C_k \to C_{k-1}$

by $\quad h_\alpha^k \longmapsto \sum_\beta \langle h_\alpha^k | h_\beta^{k-1} \rangle h_\beta^{k-1}$

where $\langle h_\alpha^k | h_\beta^{k-1} \rangle$ is the intersection number of the attaching sphere of h_α^k with the belt sphere of h_β^{k-1}.

The homology of this chain complex, $H_k(C_*) = \dfrac{\ker \partial_k}{\operatorname{Im} \partial_{k+1}}$ is naturally isomorphic to $H_k(W, M, \mathbb{Z})$.

A good way to see this is to think in terms of cellular homology, a handle decomposition is precisely a cell decomposition where the cells have been fattened, also the boundary map ∂_k can be seen to correspond exactly to the boundary map of the cellular chain complex and this is also the easiest way to see why $\partial_{k-1} \circ \partial_k = 0$

With respect to the basis of handles, the boundary maps ∂_k are all matrices with integer entries (Since intersection numbers are integers). Explicitly ∂_k is a $rk(C_{k-1}) \times rk(C_k)$ matrix.

Now I introduce the procedures by which we may alter the handle decomposition until there are no handles left.

1) <u>Handle Creation / Cancellation</u>

A $(k-1)$-handle when attached may create a hole which is later filled by a k-handle. In this case the handles can be eliminated. Similarly we may create two such cancelling handles and add them to an existing handle decomposition. This idea can easily be interpreted in terms of Morse functions

Again consider a height function:

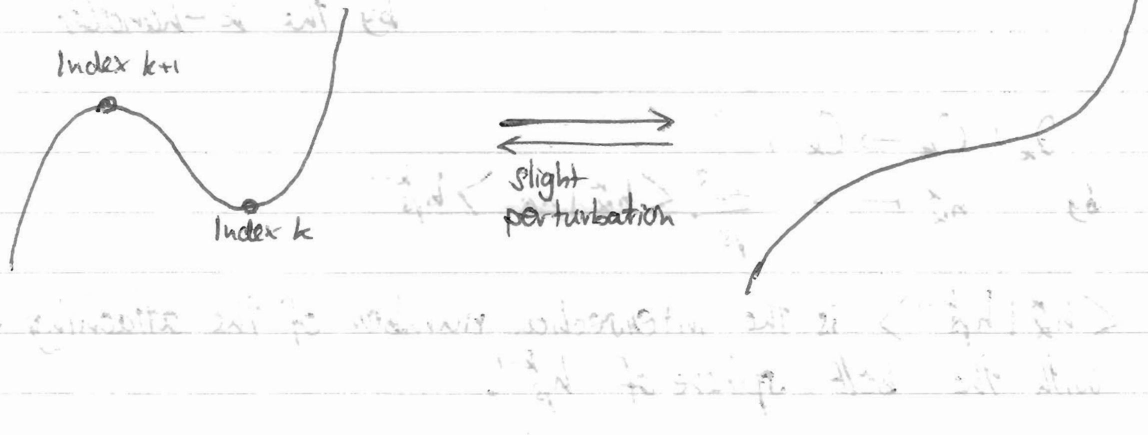

Examples of handle cancellation

Take $M \times [0, \varepsilon]$ , attach a 0-handle to get a disjoint $D^0 \times D^{m+1}$:

Now attach a 1-handle connecting the two connected components

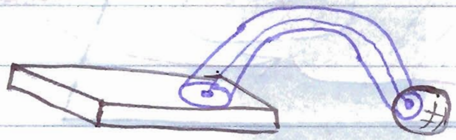

Now we may deform smoothly back to $M \times [0, \varepsilon]$:

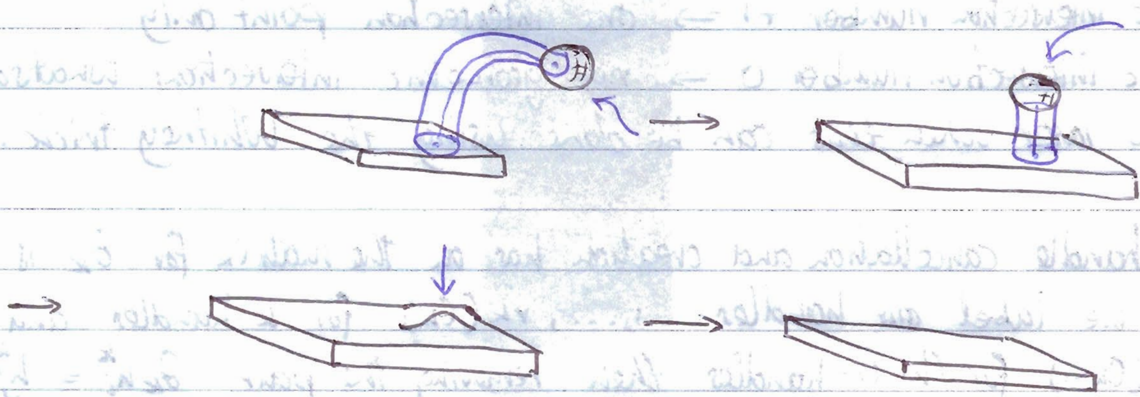

Similarly with attaching a 1-handle followed by a 2-handle

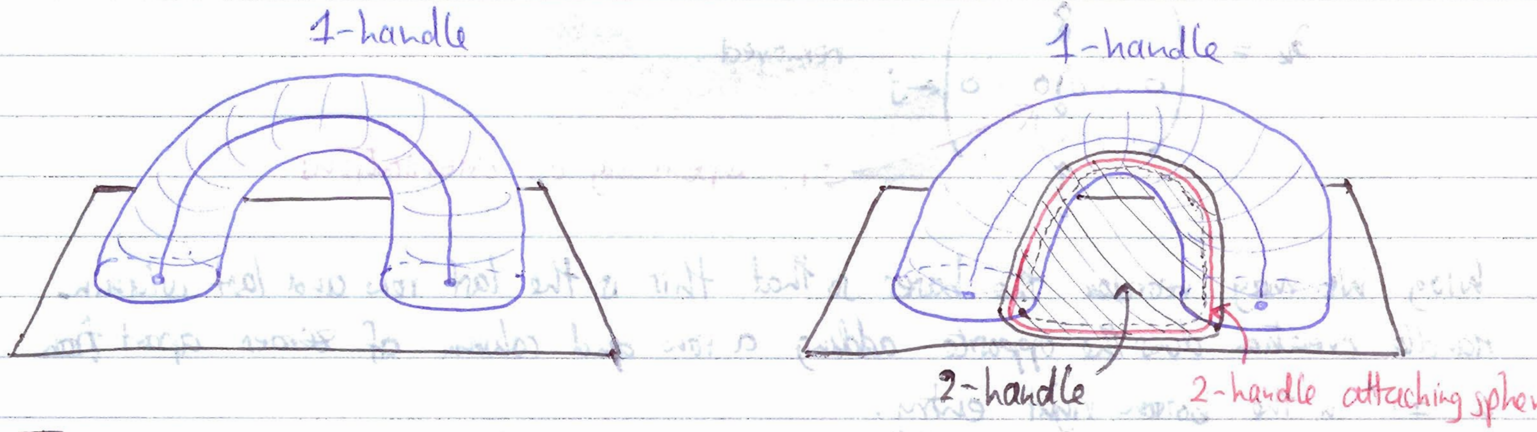

This is now just a smooth bump which can be flattened.

In order for the handles to cancel we require that $\partial_k h^k_\alpha = \pm h^{k-1}_\beta$, in which case h^k_α cancels with h^{k-1}_β. This condition is saying that the attaching sphere of h^k_α has ± 1 algebraic intersection with the belt sphere of h^{k-1}_β, and zero algebraic intersection with the belt spheres of all other $(k-1)$-handles.

Example

$$\partial_2 h^2_\beta = h^1_\alpha$$

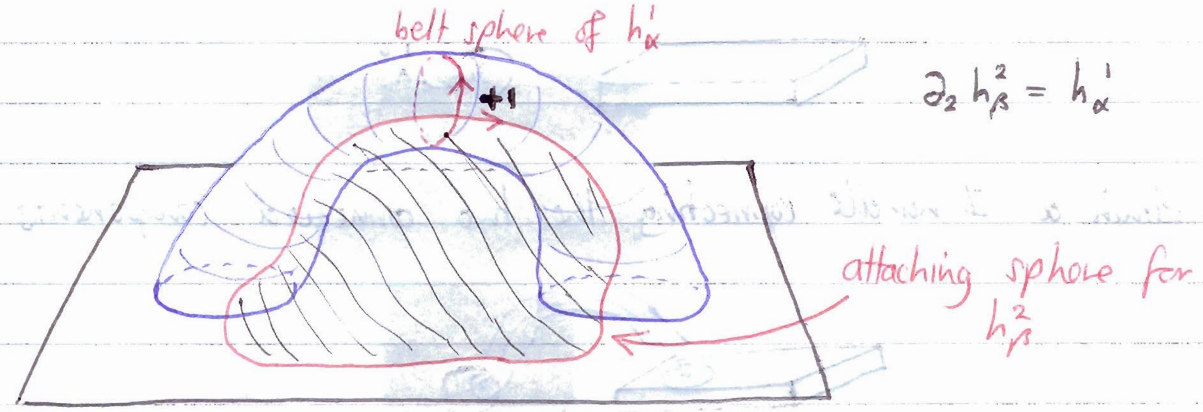

In order to cancel, the algebraic intersections need to be realised geometrically
ie algebraic intersection number $+1 \Rightarrow$ one intersection point only
algebraic intersection number $0 \Rightarrow$ no geometric intersections whatsoever.
We will show later when this can be done using the Whitney trick.

The effect handle cancellation and creation has on the matrix for ∂_k is as follows. If we label our handles $1, \dots, rk(C_k)$ for k-handles and $1, 2, \dots, rk(C_{k-1})$ for $(k-1)$-handles then removing the pair $\partial_k h^k_i = h^{k-1}_j$ has the effect of removing the j^{th} row and the i^{th} column of ∂_k and remove a generator for C_k and one for C_{k-1}.

$$\partial_k = \begin{array}{cc} \left(\begin{array}{ccccccc} & & & 0 & & & \\ & & & 0 & & & \\ & & & \vdots & & & \\ 0 & \cdots & 0 & \pm 1 & 0 & \cdots & 0 \\ & & & 0 & & & \\ & & & \vdots & & & \end{array} \right) & \leftarrow j \\ \uparrow i & \end{array} \quad \text{removed.}$$

± 1 depending on orientations

Wlog we may reorder the bases so that this is the last row and last column. Handle creation does the opposite adding a row and column of zeroes apart from a ± 1 in the bottom right entry.

2) <u>Handle Sliding</u>: In the discussion earlier we showed that for $i \leq j$ we may slide an i-handle over a j-handle, so certainly we may slide i over i.

<u>Examples</u>: Sliding a 1-handle:

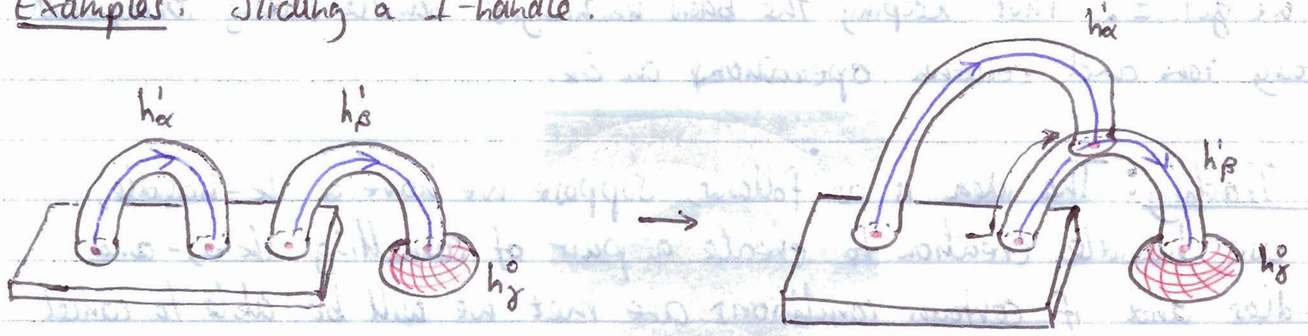

$\langle h^1_\alpha, h^0_\gamma \rangle = 0$, $\langle h^1_\beta, h^0_\gamma \rangle = +1$

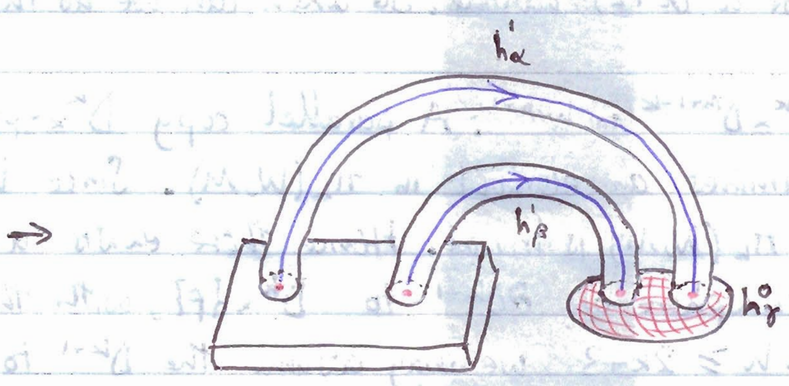

$\langle h^1_\alpha, h^0_\gamma \rangle = +1$, $\langle h^1_\beta, h^0_\gamma \rangle = +1$

Sliding a 2-handle:

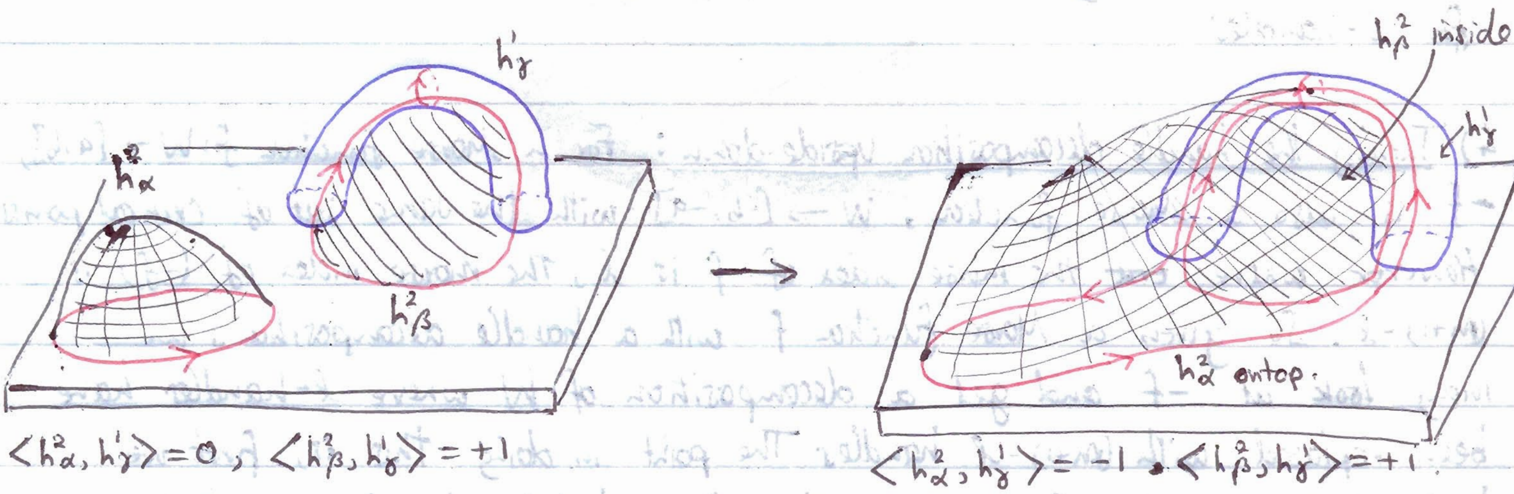

$\langle h^2_\alpha, h^1_\gamma \rangle = 0$, $\langle h^2_\beta, h^1_\gamma \rangle = +1$ $\qquad$ $\langle h^2_\alpha, h^1_\gamma \rangle = -1$, $\langle h^2_\beta, h^1_\gamma \rangle = +1$

The algebraic effect of sliding a k-handle over another k-handle is to change the boundary matrix $\partial_k : C_k \to C_{k-1}$. If h_α^k slides over h_β^k this has the effect of replacing h_α^k by $h_\alpha^k \pm h_\beta^k$ in the basis for C_k, depending on the orientation we get $\pm$. Thus keeping the basis unchanged handle sliding corresponds to elementary row and column operations on ∂_k.

3) <u>Handle Trading</u>: The idea is as follows, suppose we have a k-handle, we may use handle creation to create a pair of cancelling $(k+1)$- and $(k+2)$-handles and if certain conditions are met we will be able to cancel the k- and $(k+1)$-handles leaving behind just the $(k+2)$-handle, so in effect we trade a k-handle for a $(k+2)$-handle. So when can we do this?

Let h_α^k be a k-handle $D^k \times D^{m+1-k}$ in W^{m+1}. A parallel copy $D^k \times \{p\}$ ($p \in S^{m-k}$) of the core $D^k \times \{0\}$ determines an element in $\pi_k(W, M)$. Since W is an h-cobordism we know that $\pi_k(W, M)$ is trivial. Hence there exists a map $D^{k+1} \to W$ mapping a hemisphere of ∂D^{k+1} to $D^k \times \{p\}$, with the other hemisphere in W. For $\dim W \geq 2k+3$ we may assume the D^{k+1} to be embedded. Given an embedding, D^{k+1} may be fattened into a cancelling pair of $(k+1)$ and $(k+2)$ handles, and the $(k+1)$-handle cancels the handle h_α^k by construction. So the limiting factor is k, the order of the handle. In this proof we are assuming $m \geq 5 \Rightarrow \dim W^{m+1} \geq 6$, so $2k+3 \leq 6$ for $k = 0, 1$. This tells us we may trade all 0-handles for 2-handles and all 1-handles for 3-handles

4) <u>Turning the handle decomposition upside-down</u>: For a Morse function $f: W \to [a,b]$, $-f$ is also a Morse function, $W \to [-b, -a]$ with the same set of critical points. However where ~~crit~~ the morse index of f is ℓ, the morse index of $(-f)$ is $(m+1) - \ell$. So given a Morse function f with a handle decomposition, we may look at $-f$ and get a decomposition of W where ℓ-handles have been replaced with $(m+1)-\ell$ handles. The point in doing this is, first we trade all 0,1-handles for 2,3-handles, then looking at $-f$ we obtain a

handle decomposition with no m and $(m+1)$-handles. In this decomposition we may trade away all 0 and 1-handles to get a decomposition with no $0, 1, m$ or $(m+1)$ handles, then we turn the decomposition back right-side up and we still have no $0, 1, m$, or $(m+1)$-handles. Effectively we have cleverly traded $(m+1)$-handles for $(m-1)$-handles and m-handles for $(m-2)$-handles.

Note: A k-handle is $D^k \times D^{m+1-k}$, turning the decomposition upside-down corresponds to interchanging the two factors to get $D^{m+1-k} \times D^k$, a $(m+1)-k$-handle

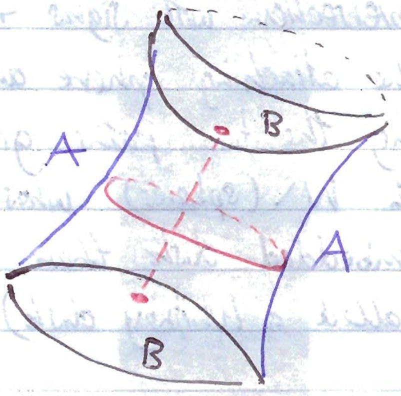

A 1-handle attached to B is a 2-handle attached to A

$\sim$

Now we return to the somewhat delicate matter of handle cancellation. We want handles which cancel 'algebraically' to cancel 'geometrically'. That is we want to change

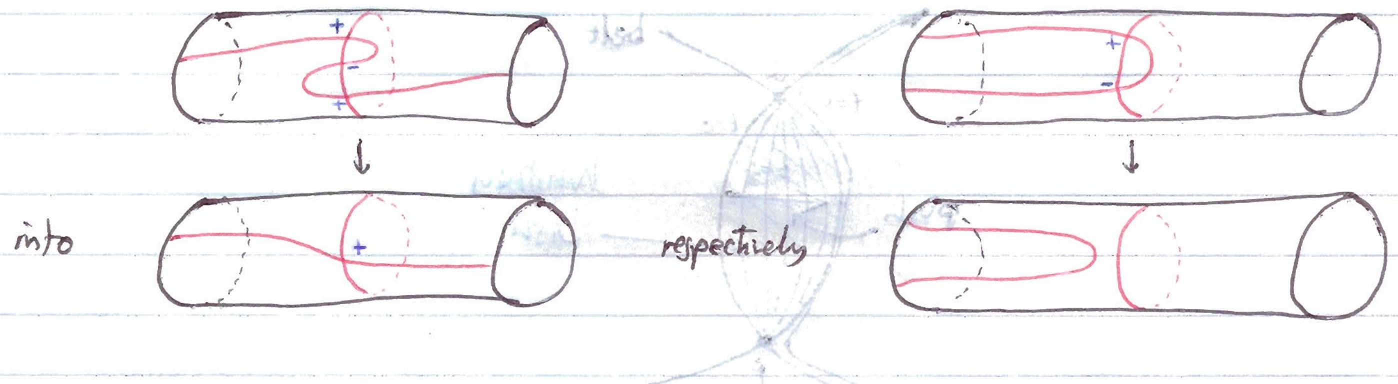

For this we use the Whitney Trick!

Question: When can we remove algebraically cancelling pairs of intersections?

The idea is as follows, consider the intersecting belt sphere and attaching sphere, a ~~p~~ pair of algebraically cancelling intersections looks something like

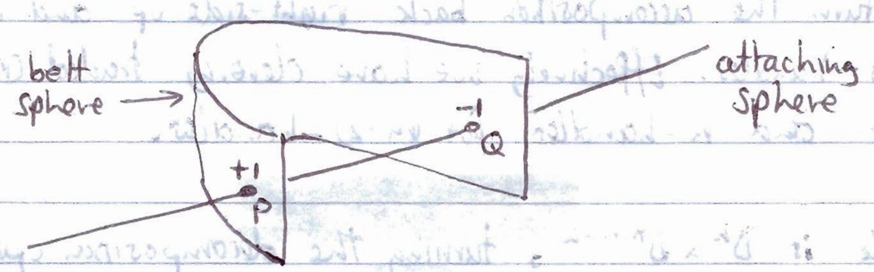

Let P and Q be the points of intersection with signs + and − respectively. Take a path from P to Q in the attaching sphere and a path from Q to P in the belt sphere. Composing the two paths gives us a loop in W. We seek an embedded disk in W \ (spheres) whose boundary is this loop. Provided we may find such an embedded disk then we may smoothly push one sphere along the disk (called a Whitney disk) until it no longer intersects the other sphere:

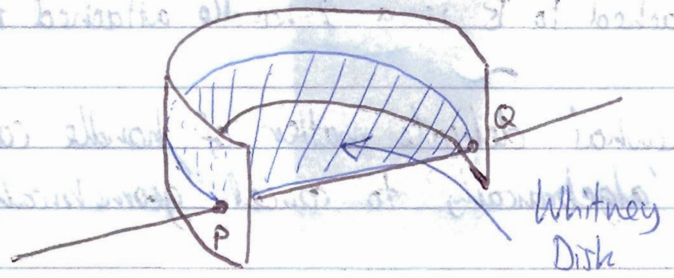

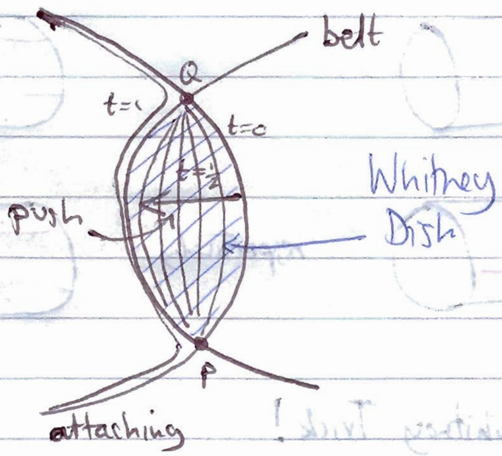

Provided we can show W with the two spheres cut out remains simply connected, then we know any loop in W \ (spheres) bounds an immersed disk, however embeddings are dense in the space of all maps $A^n \longrightarrow B^{2n+1}$, and since $m \geq 5$ we have that embeddings of disks are dense, so given an immersion we may perturb to an embedding and hence smoothly pull the two spheres apart.

So when can we guarantee that W remains simply connected after cutting out our two spheres?

The situation we consider is that of a $(k+1)$-handle attached to a k-handle. We are looking at the intersection of the attaching S^k with the belt S^{m-k}. These spheres both lie in the upper boundary of an ascending cobordism, this boundary $\widetilde{W}$ has dimension m so the spheres intersect generically in zero dimensions. We have removed all $0, 1, m$ and $m+1$ handles, hence we need only consider the cases $k = 2, 3, \ldots, m-2$.

First we look at the cases $k \in \{3, 4, \ldots, \lceil \frac{m+1}{2} \rceil\}$, removing an $m-k$ sphere does not affect simply connectedness of $\widetilde{W}$ for $k \geq 3$ as above ie codim ≥ 3, then subsequently removing a k sphere also does not affect simply-connectedness. The case we need to take care with is $k=2$, where we remove a codimension 2 sphere. Let M_1 be the upper boundary of the ascending cobordism before attaching the k-handle, and let M_2 be the upper boundary after attaching. We contemplate removing S^{m-k} from M_2

$$\text{Since } (D^k \setminus \{0\}) \times S^{m-k} \simeq S^{k-1} \times S^{m-k} \simeq S^{k-1} \times (D^{m-k+1} \setminus \{0\})$$

we have that $M_2 \setminus S^{m-k}$ is diffeomorphic to $M_1 \setminus S^{k-1}$

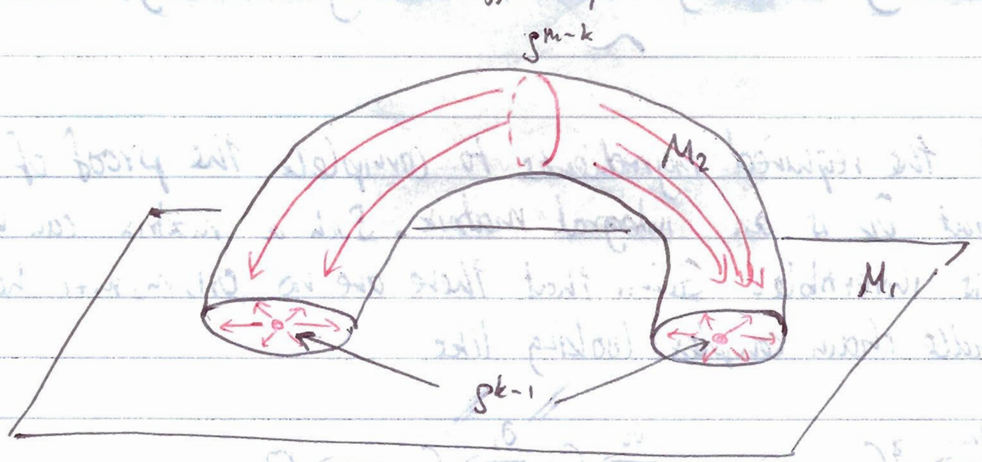

For $k=2$, we are removing S^1 from M_1^m, for $m \geq 5$ removing S^1 does not destroy simply-connectedness, hence $M_1 \setminus S^1$ simply connected and diffeomorphic to $M_2 \setminus S^{m-2}$ $\Rightarrow$ cutting out the codim 2 sphere we are left with something that is still simply connected. After this cutting out the attaching 2-sphere is also not a problem.

So, we have just shown that $k = 2, 3, \ldots, \lceil \frac{m+1}{2} \rceil$ are fine. By turning the handle decomposition upside down we get that $\lceil \frac{m+1}{2} \rceil + 1, \ldots, m-2$ are also fine. Thus for $k = 2, \ldots, m-2$ we may find an immersed disk and hence an embedded disk as claimed. This is of course provided $m \geq 5$.

If $m = 4$, then we encounter several problems, firstly the codimension two spheres, when removed, may destroy simply-connectedness, and secondly, given an immersed disk we cannot necessarily find an embedded disk!

P and Q must have opposite signs otherwise we run the risk of introducing more intersections of the spheres after pushing along the whitney disk.

So we have just shown for all handles left in the decomposition, we may geometrically remove algebraically cancelling intersections so that intersection number = 0 means no intersection, and 1 means a single intersection.

This means algebraically cancelling handles, cancel geometrically.

~

Now we have all the required ingredients to complete the proof of the theorem. Recall that ∂_k is an integral matrix. Such a matrix can be diagonalised if it is invertible. Given that there are no $0, 1, m, m+1$ handles we have the handle chain complex looking like

$$0 \to C_{m-1} \xrightarrow{\partial_{m-1}} C_{m-2} \to \cdots \xrightarrow{\partial_4} C_3 \xrightarrow{\partial_3} C_2 \to 0.$$

Since $H_*(W, M; \mathbb{Z}) = 0$, the sequence is exact, $\mathrm{Im}\, \partial_{k+1} = \ker \partial_k$. Also as the C_k are free we may split the sequence as follows

$$\begin{array}{ccccccc}
\longrightarrow C_4 & \longrightarrow & C_3 & \xrightarrow{\partial_3} & C_2 & \longrightarrow & 0 \\
 & & \| & & \| & & \\
\| & & \mathrm{Coker}\, \partial_4 & \xrightarrow[\cong]{\partial_3} & \mathrm{Im}\, \partial_3 & & \\
 & & \oplus & & & & \\
\mathrm{Coker}\, \partial_5 & \xrightarrow[\cong]{\partial_4} & \mathrm{Im}\, \partial_4 & & & & \\
\oplus & & & & & & \\
\vdots & & & & & &
\end{array}$$

Each ∂_k restricts to an invertible morphism $\mathrm{Coker}\, \partial_{k+1} \to \mathrm{Im}\, \partial_k$, thus we can diagonalise ∂_k via the elementary row and column operations and by stabilising.

The diagonalisation of all the ∂_k means that all the handles are paired with a single other handle. Switching orientations of handles we may assume that all diagonal entries are non-negative, since the matrices are invertible the diagonal entries must be 1, thus all the paired up handles cancel with eachother algebraically and hence geometrically yielding a decomposition with no handles just as we sought

<u>End Remarks</u>: Before Smale proved the theorem in 1961, mathematicians had got stuck understanding manifolds of dimension 3 & 4 assuming higher dimensions to be harder, but the h-cobordism theorem showed at least for simply connected manifolds that in fact dims ≥ 5 are easier.

The proof of the theorem hinges on the fact that $m \geq 5$, in order to use the Whitney trick and to embed disks. In dimension 4 things are not so easy...