Morse and Cerf Theory

Source: https://euclidlab.org/research/morse-and-cerf-theory (October 27, 2020)

2012-01-09

Morse Theory studies the topology of smooth manifolds by looking at generic smooth maps (Morse functions) from smooth manifolds to the reals (or, sometimes, to the circle), and investigating the critical points of such maps and their indices, interactions, the gradient flow lines for the maps, etc. **Cerf Theory** studies smooth 1-parameter families of functions connecting different Morse functions on the same smooth manifold. Although there are important infinite-dimensional versions of these theories, I will focus on the finite-dimensional setting and, especially, on low dimensions like 2, 3, 4 and 5. I will focus less on the foundational analytic technicalities and more on the applications. Some of the most important results I want to get to are (not necessarily in this order, and not necessarily in full detail):

1. The classification of surfaces.
2. The existence of Heegaard splittings of 3-manifolds.
3. Generators and relations for the mapping class groups of surfaces.
4. Handlebody decompositions for 4-manifolds and surgery diagrams for 3-manifolds.
5. The Kirby calculus for 3-manifolds and 4-manifolds.
6. The isotopy versus pseudo-isotopy problem as studied by Cerf and Hatcher-Wagoner.
7. Understanding Morse 2-functions (generic maps to 2-manifolds).

We begin with a problem: Consider the embedding of S^2 in $\mathbb{R}^3$ shown at left. This is supposed to be rotationally symmetric about the red z-axis. Let Σ denote this particular submanifold of $\mathbb{R}^3$ (diffeomorphic to S^2). Now let $f : \Sigma \to \mathbb{R}$ be orthogonal projection onto the z-axis. Note that f has two critical points, indicated, a maximum and a minimum. Now note that, for any unit vector $v \in S^2$, we can project Σ orthogonally onto the oriented line spanned by v to get another function $f_v : \Sigma \to \mathbb{R}$. Thus our original f is $f_{(0,0,1)}$, and $-f$ is $f_{(0,0,-1)}$. (We are assuming the origin is in the middle of the picture at left.) For another example, $f_{(1,0,0)}$ will have six critical points: two maxima, two minima and two saddles, as in the figure below:

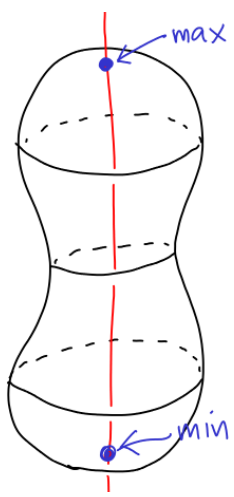

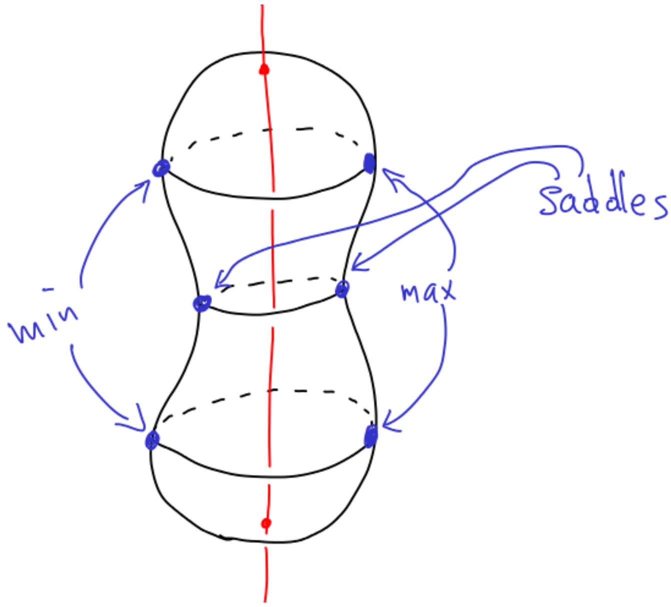

Thus we have a 2-parameter family of functions f_v, parameterized by the 2-dimensional parameter space S^2. The problem is to somehow depict with a diagram the behavior of the critical values of f_v and their indices as v ranges over S^2. (The index of a minimum is 0, the index of a saddle is 1, the index of a maximum is 2, more on indices later.)

(Recall: for a smooth $f : X \to Y$, a point $p \in X$ is a critical point if $Df_p : T_pX \to T_{f(p)}Y$ does not have maximal rank. If $p \in X$ is a critical point, then $f(p) \in Y$ is a critical value. When Y is 1-dimensional, not having maximal rank simply means being equal to 0.)

One possible way to depict this is to draw a diagram in $S^2 \times \mathbb{R}$, where in each $v \times \mathbb{R}$ you draw the critical values of f_v, labelled with their indices. Since the critical values are usually isolated, you should have some kind of surface in $S^2 \times \mathbb{R}$, possibly with some interesting singularities. A convenient picture for $S^2 \times \mathbb{R}$ is an open shell between two concentric spheres (identifying $\mathbb{R}$ with an open interval).

This picture drawn in $S^2 \times \mathbb{R}$ can also be thought of as the set of critical values of the function $F : S^2 \times \Sigma \to S^2 \times \mathbb{R}$ given by $(v, p) \mapsto (v, f_v(p))$.

We are just about to define a Morse function properly, given the above preamble, but first we mention one famous problem that was studied using Morse and Cerf theory:

An *isotopy* between two maps $f_0, f_1 : X \to Y$ can be define as a diffeomorphism $F : [0,1] \times X \to [0,1] \times Y$ such that $F(0,p) = (0, f_0(p))$, $F(1,p) = (1, f_1(p))$ and F is "level-preserving", i.e. F is the identity on the $[0,1]$ component, or $F(t,p) = (t, f_t(p))$. A *pseudo-isotopy* between f_0 and f_1 is a diffeomorphism $F : [0,1] \times X \to [0,1] \times Y$ satisfying the first two criteria but not necessarily level-preserving. Note than when gluing manifolds together along boundaries, the resulting manifold is determined up to diffeomorphism by the pseudo-isotopy class of the gluing map (this is a good exercise to prove), but in other contexts the difference between isotopy and pseudo-isotopy is very important. Cerf studied 1-parameter families of functions connecting Morse functions in order to understand this problem, and hopefully we'll get to that later.

Now:

Preliminary definition: A *Morse function* is a smooth map f from an n-manifold X to a 1-manifold Y such that, for every critical point $p \in X$, there exist local coordinates $(x_1, \ldots, x_n)$ about p and a coordinate y about $f(p)$ with respect to which $f(x_1, \ldots, x_n) = -x_1^2 - \ldots - x_k^2 + x_{k+1}^2 + \ldots x_n^2$. The integer k is called the *index* of p.

We need to see that Morse functions exist, that the index is independent of the coordinates, that there are lots of Morse functions, that the property of being Morse is stable (doesn't change under small perturbations), and many other foundational facts. But first some examples:

When $n = 2$, we have minima $f(x_1, x_2) = x_1^2 + x_2^2$ with index $k = 0$, saddles $f(x_1, x_2) = -x_1^2 + x_2^2$ with index $k = 1$ and maxima $f(x_1, x_2) = -x_1^2 - x_2^2$ with index $k = 2$.

When $n = 1$ we have minima $f(x) = x^2$ with index $k = 0$ and maxima $f(x) = -x^2$ with index $k = 1$.

Note that $f(x) = x^3$ is not Morse, and that there is an interesting perturbation $f_t(x) = x^3 + tx$. When $t < 0$, f_t is Morse with one min and one max, and when $t > 0$, f_t is Morse with no critical points. We'll discuss these phenomena more carefully soon, and they do arise in the problem that we began with.

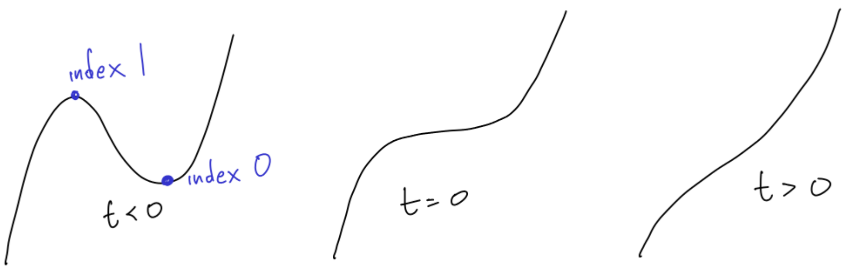

(YouTube link)

2012-01-11

Addendum to the problem from last time:

In the exercise from last time, since the embedded sphere Σ is rotationally symmetric about the z-axis, the critical values of f_v, as a function of $v \in S^2$, will be invariant under rotation of $v \in S^2$ about the z-axis, and thus really we might as well think of the parameter v as in S^1, or even just in an interval from north pole to south pole in S^2.

So do the problem as stated, but note that your answer is rotation-invariant. Then do it again but this time using the surface Σ at right:

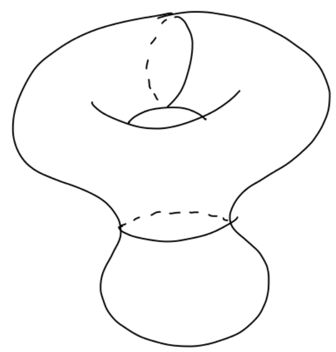

Continuation of lecture:

Last lecture's definition of a Morse function was called *preliminary* because it did not discuss boundary behavior and compactness. Here is the full definition:

Definition: A function $f : X^n \to Y^1$ is *Morse* if the following conditions are satisfied:

1. For each critical point $p \in X$ there are coordinates around p and $f(p)$ with respect to which $f(x_1, \ldots, x_n) = -x_1^2 - \ldots - x_k^2 + x_{k+1}^2 + \ldots x_n^2$.
2. X is compact.
3. $f^{-1}(\partial Y) = \partial X$

(In class I said that either X is closed and Y is anything or X is a cobordism (see drawing) from M_0 to M_1 and $Y = [0, 1]$ with $f^{-1}(0) = M_0$ and $f^{-1}(1) = M_1$. The way I've said it above is only slightly more general, and the cobordism case is generally the most important case to consider.)

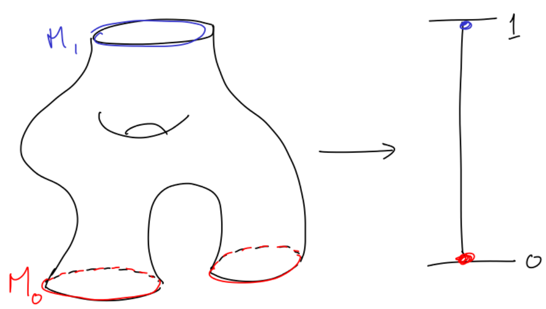

(To say that X is a *cobordism* from M_0 to M_1 means that X is compact and $\partial X = M_0 \amalg M_1$. To say that X is *closed* means that X is compact with $\partial X = \emptyset$.)

Here are two important results about Morse functions, the proofs of which we will defer till later in the interest of getting quickly to the topological applications:

Theorem: For any compact X there exists a Morse function on X. More precisely, for any compact n-manifold X and any 1-manifold Y, and given any function $f : \partial X \to \partial Y$, there exists an extension of f to a Morse function $f : X \to Y$. (And furthermore, Morse function are generic, i.e. there are lots of them, and any given function can be perturbed in an arbitrarily small way to be Morse, more on this later.)

Thus we have Morse functions when we need them.

Theorem: If $p \in \mathbb{R}^n$ is a critical point of $f : \mathbb{R}^n \to \mathbb{R}$ such that the Hession $Hf(p)$ (the $n \times n$ matrix of second order partial derivatives) is non-degenerate as a bilinear form, then f is *locally Morse* at p, i.e. there are coordinate around p and $f(p)$ with respect to which $= -x_1^2 - \ldots - x_k^2 + x_{k+1}^2 + \ldots x_n^2$. Furthermore, the index k of p is precisely the index of $Hf(p)$ as a bilinear form, the number of negative diagonal entries of $Hf(p)$ after diagonalizing.

Thus the index of a Morse critical point is independent of the coordinate system.

Topology from Morse functions

Our first example of recovering topological information from a Morse function is the case of a Morse function with no critical points.

Theorem: If X is a cobordism from M_0 to M_1 with a Morse function $f : X \to [0, 1]$, if f has no critical points then X is diffeomorphic to $[0, 1] \times M_0$.

Proof: We will construct a vector field V on X such that $df(V) \equiv 1$ (which is the same thing as saying that $Df(V) = \partial/\partial y$, where I'm using d for the exterial differential and D for the derivative). Using this we will flow forward along V from M_0 to construct the dif-

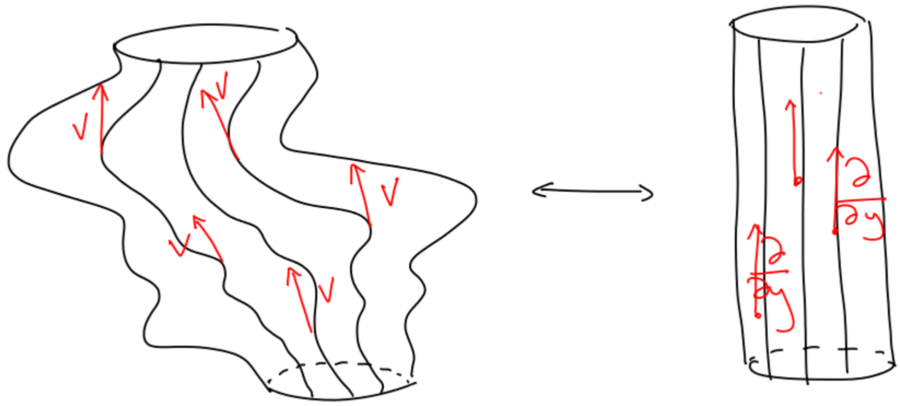

feomorphism.

To get V we need a Riemannian metric (there are other more direct ways using a partition of unity to directly patch together such V's on coordinate charts, but using a Riemannian metric has some advantages and is at the very least an important idea). A *Riemannian metric* g on X is a choice of an inner product g_p on T_pX for each $p \in X$, varying smoothly in p. (*Varying smoothly in* p just means that, when g is written as an $n \times n$ matrix in local coordinates, the entries of the matrix are all smooth functions of p.) Next time we will use a partition of unity to show that Riemannian metrics exist and then show how to use such a metric to get V.

2012-01-13

Further addendum to homework problem: As a warmup, do the following one-lower-dimensional version: Let Σ be the embedding of S^1 in $\mathbb{R}^2$ shown at right and now, for each $v \in S^1$, define the analogous $f_v : S^1 \to \mathbb{R}$. Then draw a graph in $S^1 \times \mathbb{R}$ of the critical values and their indices as a function of v. Now the indices will only be 0 and 1. The critical events to note are births and deaths of pairs of critical points and crossings of critical values (one critical value rising above or below another one).

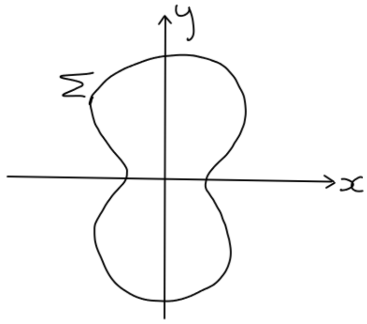

Now continuing our proof: We need a Riemannian metric on X, so here is the quick proof that Riemannian metrics exist. Cover X with coordinate charts $\{U_i\}$ with a corresponding partition of unity $\{\mu_i\}$. In each coordinate chart choose the standard Euclidean inner product g_i. Then let $g = \sum \mu_i g_i$. This works because convex combinations of positive definite symmetric matrices are positive definite symmetric matrices.

Now note that a metric g, at each point p, is a non-degenerate bilinear form $g_p : T_pX \times T_pX \to \mathbb{R}$ and can thus equivalently be thought of as an isomorphism $g_p : T_pX \to T_p^*X$. Then using this isomorpism, we construct a vector field W by $W_p = g_p^{-1}(df_p)$. Because f has no critical points, W is nowhere 0. This is the *gradient vector field* for f with respect to the metric g, denoted $\nabla_g f$. As a basic exercise you should verify that, when g is the standard inner product on $\mathbb{R}^n$ and $f : \mathbb{R}^n \to \mathbb{R}$, then $\nabla_g f$ is the usual gradient ∇f.

Now because W is never 0, $df(W) = g(W, W)$ is never 0, so we can let $V = (1/df(W))W$, so that $df(V) \equiv 1$. This is the vector field we wanted. Now we construct a diffeomorphism $\phi : [0,1] \times M_0 \to X$ by making $\phi(t,p)$ equal to the point q you get to by starting at $p \in M_0$ and flowing forward along V for time t. The fact that $df(V) \equiv 1$ means that $f(q) = t$, and from this and the existence and uniqueness of solutions to ordinary differential equations shows that ϕ is a diffeomorphism. $\square$

So now what if there are critical points?

Suppose that X is a cobordism from M_0 to M_1 with a Morse function $f : X \to [0,1]$ with one single critical point $p \in X$ of index k, as in the picture at right. We will again use a gradient vector field $W = \nabla_g f$ to understand the topology of X in terms of the topology of W_0 but now, because $df_p = 0$, we cannot rescale W to get a vector field V with $df(V) \equiv 1$ on all of X. So instead we will divide X into four parts, on three of which we will rescale W. But before we do this we need to construct our metric g a little more carefully: We want there to be a coordinate chart U around p with respect to which g is the standard Euclidean inner product and f is the standard Morse local model $\sum \pm x_i^2$, so that $W = \sum \pm 2x_i \partial_{x_i}$. This is possible because we can start with a standard Morse chart around p as one of the charts in our partition of unity construction and then arrange that, in a ball neighborhood around p, one of the μ_i's is identically 1 and all the others are 0.

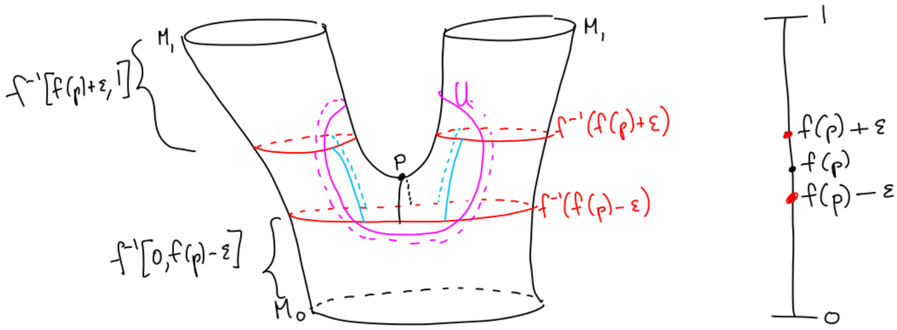

So now, assuming that g, f and W are standard inside a neighborhood U of p, we draw a picture of U with the level sets of f and the flow lines of V to the right. We choose an $\epsilon > 0$ so that $f^{-1}(f(p) - \epsilon)$ and $f^{-1}(f(p) + \epsilon)$ intersect U as shown. Then our four pieces of X, which we will study more carefully next time, are:

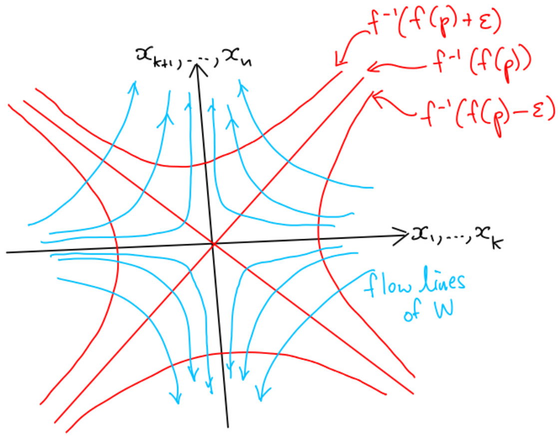

1. $f^{-1}[0, f(p) - \epsilon]$, which is diffeomorphic to $[0, f(p) - \epsilon] \times M_0$ using flow along $(1/df(W))W$.
2. $f^{-1}([f(p) + \epsilon, 1]$, which is diffeomorphic to $[f(p) + \epsilon, 1] \times M_1$ using backward flow along $(1/df(W))W$.
3. The intersection of $f^{-1}[f(p) - \epsilon, f(p) + \epsilon]$ with the closure of the union of all flow lines for W which start in some tubular neighborhood of $x_1^2 + \ldots + x_k^2$ in $f^{-1}(f(p) - \epsilon)$. This is the "mystery piece" that we will understand better soon.
4. The rest of $f^{-1}[f(p) - \epsilon, f(p) + \epsilon]$, which is a product that we will also discuss next time.

(YouTube link)

2012-01-18

Exercise: Let us say that a Riemannian metric g is *adapted* to a Morse function f if, for each critical point p of f, there exist coordinates around p with respect to which $f = \sum \pm x_i^2$ and g is the standard inner product. Show that the space of metrics adapted to a fixed Morse function is connected. I.e. if g_0 and g_1 are adapted to f then they are connected by a smooth family g_t, adapted to f for each t. It might be helpful to show that any two coordinate charts near p, with the same orientation, for which f is standard can be connected by a smooth path of such coordinate charts. (Thanks to Bruce Bartlett and Eric Burgess for pointing out the importance of the orientation here, since $O(n)$ is disconnected.) It is also helpful to show that the space of inner products on $\mathbb{R}^n$ is connected.

Now back to the main thread: We are thinking about the situation where X is a cobordism from M_0 to M_1 with a Morse function $f : X \to [0, 1]$ with a single critical point p of index k, and we want to understand what this says about the topology of X. Refer again to the figure at right.

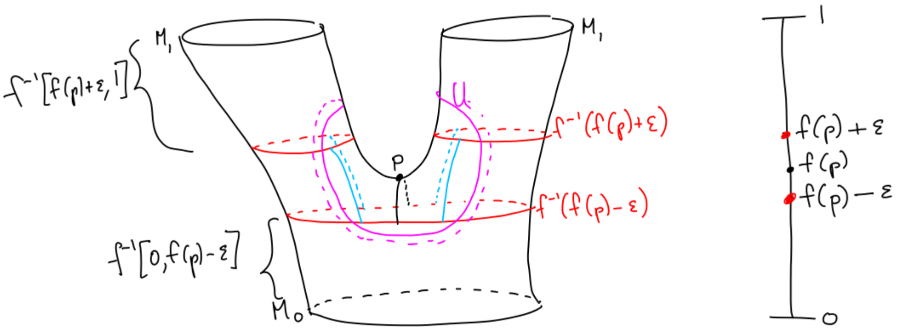

Where we are going is: we want to describe X as built as a product on M_0 at the bottom, with some kind of "handle" attached going over the critical point p, followed by another product on M_1 at the top.

For our first approach to making this precise, we break X into four pieces: $f^{-1}[0, f(p) - \epsilon]$, $f^{-1}[f(p) + \epsilon, 1]$ (both of which are products) and two pieces making up $f^{-1}[f(p) - \epsilon, f(p) + \epsilon]$. The small $\epsilon > 0$ is chosen so that there is a coordinate chart U around p making f standard, with coordinates $(x_1, \ldots, x_n)$, such that the closed ball $\sum x_i^2 \leq \epsilon$ is contained in U. Then, for some $\epsilon' > \epsilon$, we can take our coordinate chart U to be the open ball $\sum x_i^2 < \epsilon'$, and U and f look like the figure at right. Note that $f^{-1}(f(p) - \epsilon) \cap \{x_{k+1} = \ldots = x_n = 0\}$ is a sphere S^{k-1}. Pick some small $\delta > 0$ so that the

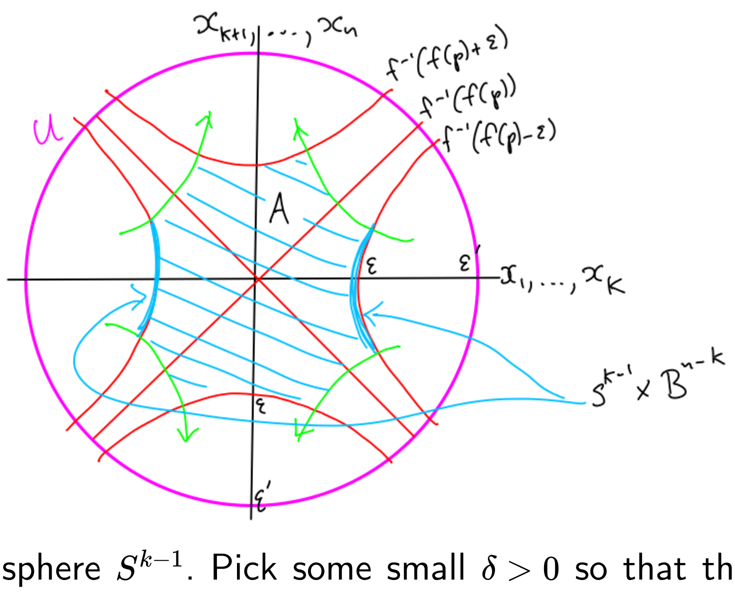

δ-neighborhood $S^{k-1} \times B^{n-k}$ of this S^{k-1} in $f^{-1}(f(p)-\epsilon)$ is contained in U, and then let A be the closure of the union of the flow lines for $\nabla_g f$ which pass through this $S^{k-1} \times B^{n-k}$, intersected with $f^{-1}[f(p)-\epsilon, f(p)+\epsilon]$. This is the region shaded in blue. We want to think of A as a cobordism from $S^{k-1} \times B^{n-k}$ to $B^k \times S^{n-k-1}$, where the $B^k \times S^{n-k-1}$ is $A \cap f^{-1}(f(p)+\epsilon)$, which is a δ-neighborhood of $f^{-1}(f(p)+\epsilon) \cap \{x_1 = \ldots = x_k = 0\}$ in $f^{-1}(f(p)+\epsilon)$. In the preceding figure of the whole cobordism X, the region A is also outlined in blue.

Seeing A as a cobordism between manifolds with boundary means enlarging the definition of a cobordism to include manifolds with boundary and corners, with the corners separating "vertical" boundary (which is a product) and "horizontal" boundary (the top and the bottom). If we allow this, and the definitions are natural, then the complement of A $f^{-1}[f(p)-\epsilon, f(p)+\epsilon]$ is also a cobordism, but this time a product. Thus we can characterize X as follows: X is built from M_0 by first constructing a product $[0,1] \times M_0$ (we replace $[0, f(p)-\epsilon]$ with $[0,1]$ for simplicity). Then we attach A to $\{1\} \times M_0$ via an embedding of $S^{k-1} \times B^{n-k}$ into $\{1\} \times M_0$. At this point we do not have a smooth manifold but we make it smooth by also attaching a product cobordism to the complement of this embedding, and gluing the sides of the product cobordism to the sides of A. Finally we complete with another product cobordism $[0,1] \times M_1$, but since this doesn't "do anything" we can just ignore that step.

Next time I hope to say this a little more carefully, so I'll leave out the rest of my waffle from this lecture and clarify in my next post. Here's the video (thanks to Eddie Beck):

(YouTube link)

2012-01-20

I belabored some subtleties about handles last time and will continue to do so a little bit more here. The upshot of the story should be that, when X is an n-dimensional cobordism from M_0 to M_1 with a Morse function $f : M \to [0,1]$ with a single critical point of index k, then X is diffeomorphic to $[0,1] \times M_0$ with an n-dimensional k-handle attached to $\{1\} \times M_0$. Here are 3 different approaches to defining a handle and what it means to attach a handle, and hence making sense of the preceding sentence:

1. The most standard thing is to say that an n-dimensional k-handle is $H = H^n_k = B^k \times B^{n-k}$. This is glued to the top M of a cobordism X via an embedding $\phi : S^{k-1} \times B^{n-k} \hookrightarrow M$. The boundary of H is divided into two parts: the *attaching region* $(\partial B^k) \times B^{n-k} = S^{k-1} \times B^{n-k}$ and the *free region* $B^k \times \partial B^{n-k} = B^K \times S^{n-k-1}$. Note first that H is not a smooth manifold, but a manifold with corners, and that, after attaching such a handle, we get a manifold with corners, which need to be smoothed. All this is illustrated in the figure below. There are subtleties one could discuss about what it means precisely to smooth corners and about the fact that any reasonable way of smoothing the corners produces the same smooth manifold. Some of these details are dealt with in Kosinski, *Differential Manifolds*.

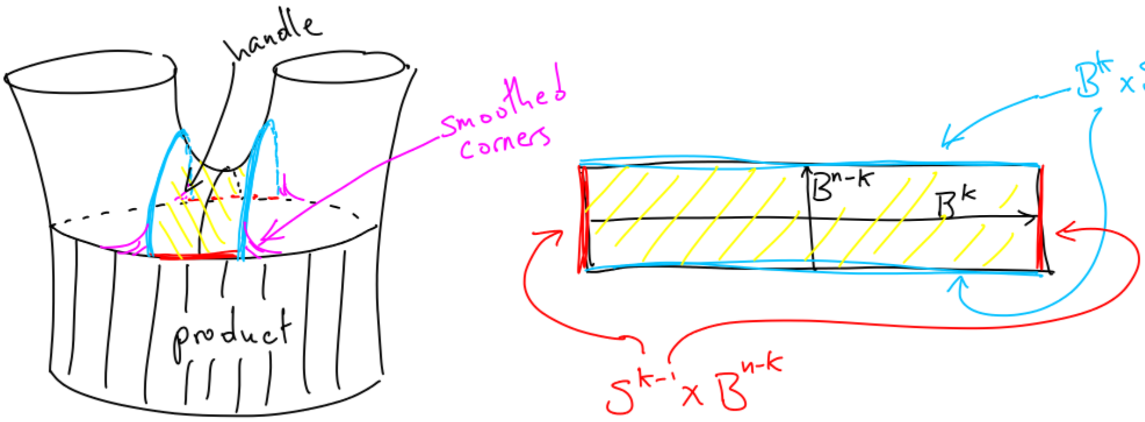

2. The second approach is the approach discussed first in the preceding lecture, in which the handle is itself a cobordism, but a cobordism with corners between manifolds with boundary. More precisely, $H = H^n_k$ is a subset of $\mathbb{R}^n$ defined as follows: Let $f = -x_1^2 - \ldots - x_k^2 + x_{k+1}^2 + \ldots + x_n^2$. Let $\partial_- H = f^{-1}(-1) \cap \{x_{k+1}^2 + \ldots x_n^2 \leq 1\}$. Then let H be the closure of the intersection of $f^{-1}[-1,1]$ with the union of all flow lines for ∇f starting on $\partial_- H$. Note that $\partial_- H \cong S^{k-1} \times B^{n-k}$ and that H is a cobordism from $\partial_- H$ to $\partial_+ H = f^{-1}(1) \cap \{x_1^2 + \ldots + x_k^2 \leq 1\} \cong B^k \times S^{n-k-1}$. Furthermore, H has "product sides" $\partial_0 H \cong [0,1] \times S^{k-1} \times S^{n-k-1}$, the part of ∂H consisting of flow lines starting at $\partial(partial_- H)$. In this case, what it means to attach H is to choose an

embedding $\phi : S^{k-1} \times B^k \hookrightarrow M$, glue H using this attaching map, and then attach a product $[0,1] \times (M \setminus \phi(S^{k-1} \times B^k))$, glueing the sides of this product to the sides of H, resulting in a smooth manifold. This is illustrated below.

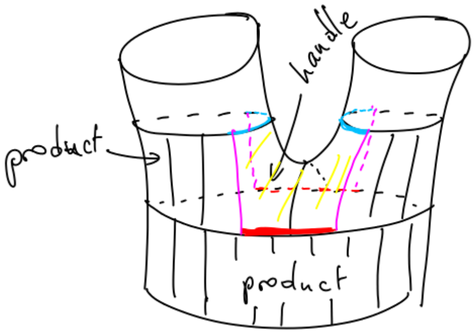

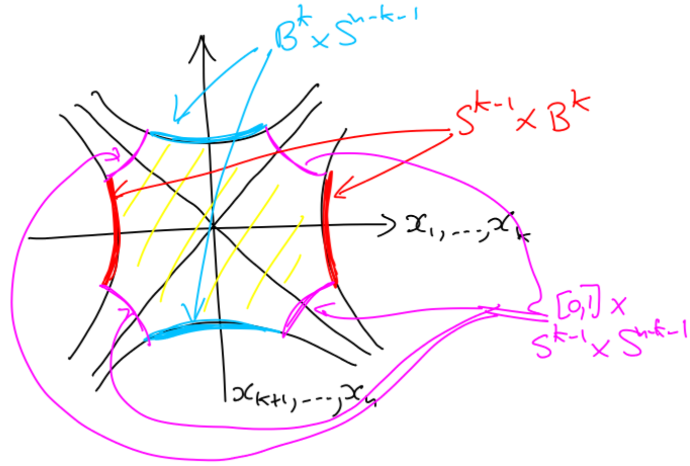

3. The third approach is to describe the handle as something that immediately produces a smooth manifold after being attached; in this case it should have "flanges" instead of corners. Milnor, in his *Morse Theory*, does this by comparing the Morse function f and a small perturbation of f supported in a neighborhood of the critical point. Here is another way: H^n_k is a subset of $\mathbb{R}^n$ described as follows: Let $f : \mathbb{R}^n \to \mathbb{R}$ and $\partial_- H$ be as in the preceding construction. Now let $\tau : \partial_- H \to (0, \infty]$ be the time it takes to flow from a point on $\partial_- H \subset f^{-1}(-1)$ to $f^{-1}(1)$. Thus $\tau = \infty$ along $S^{k-1} \times \{0\} \subset S^{k-1} \times B^{n-k} = \partial_- H$. Choose a bump function $\mu : B^{n-k} \to [0,1]$ and consider the function $\mu\tau : S^{k-1} \times B^{n-k} = \partial_- H \to (0, \infty]$. Then let H be the closure of the union of flow lines for ∇f starting at $\partial_- H$ and flowing forward for time $\mu\tau$. This is illustrated below and, once again, is attached via an embedding of $S^{k-1} \times B^{n-k}$. The resulting manifold is immediately smooth and the new boundary is obtained from the old boundary by "replacing" the image of the embedding of $S^{k-1} \times B^{n-k}$ with the other part of the boundary of H, which is diffeomorphic to $B^k \times S^{n-k-1}$.

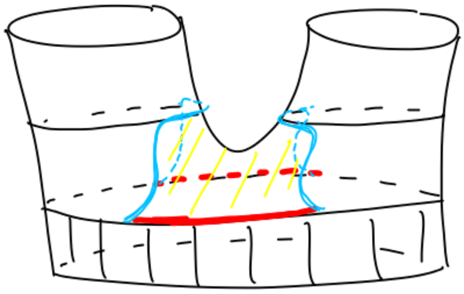

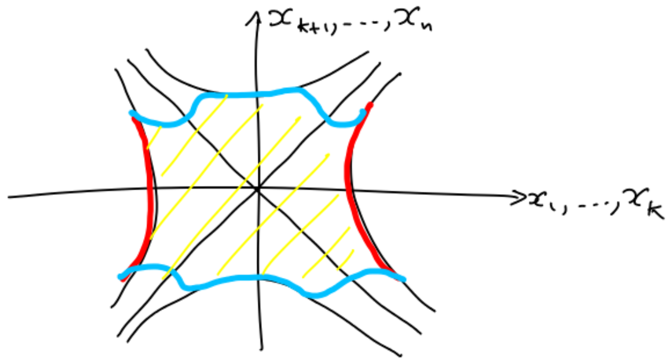

Henceforth we will use the simplest $B^k \times B^{n-k}$ model, but I wanted to discuss these subtleties because, in some contexts, in can become important. For example, if one wants to build manifolds with certain additional structures (symplectic or metric structures, for example), one would like to be very careful about extending such structures across handles, and then one may need to be quite careful with the rounded corners.

In the rest of the lecture I went through examples of handles in dimensions one and two, and discussed the cases $k = 0$ and $k = n$. I'll save that writeup for the next blog post. Here is the video:

(YouTube link)

2012-01-23

Here we focus on examples of handles. Recall that an n-dimensional k-handle is $H = H_k^n = B^k \times B^{n-k}$ with ∂H divided into two regions, the *attaching region* $S^{k-1} \times B^{n-k}$ and the *free region* $B^k \times S^{n-k-1}$, and a handle is attached to a pre-existing cobordism X from M_0 to M_1 via an embedding of the attaching region into M_0, producing a new cobordsim X' from M_0 to M_1', containing X. Note that we have not yet discussed carefully how M_1' is obtained from M_1, but when we do discuss this, the general method will be known as *surgery*.

First we note that all of this even makes sense when $k = 0$ or $k = n$, with the convention that B^0 is a point and $S^{-1} = \emptyset$. For example, S^1 is built with a 0-handle $B^0 \times B^1$, attached along an embedding of $\emptyset$, i.e. not attached to anything, following by a 1-handle $B^1 \times B^0$ attached along an embedding of $S^0 \times B^0$, as in the figure to the right. This picture obviously generalizes to S^n built with a 0-handles and a n-handle.

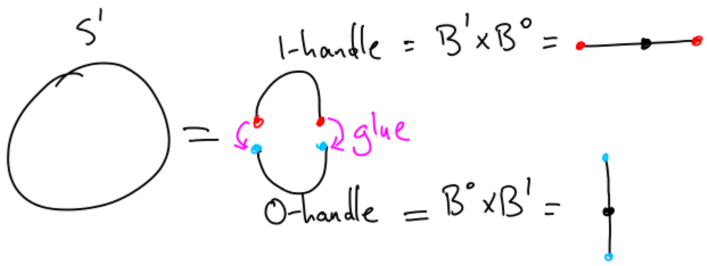

Our main results thus far concern Morse functions with zero or one critical points, but these immediately imply the following general result:

Corollary: Every cobordism decomposes into a sequence of products and handles. In particular, every closed n-manifold can be built starting with a 0-handle, then attaching some number of other handles of index $0 \leq k \leq n$, and then capping off with a n-handle.

(In fact we can always arrange that, if the manifold is connected, we only need one 0-handle and one n-handle, but this fact is not entirely trivial and we will try to prove it carefully later.)

So here is a sequence of examples:

Dimension 1: S^1 again, but with more handles; notice the different possible ways you might cancel pairs of 0- and 1-handles:

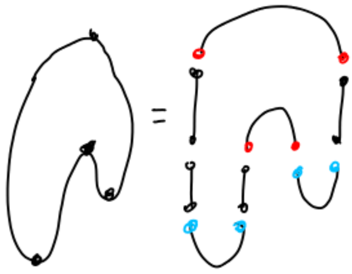

Dimension 2**:** Here is a standard picture of a torus decomposed into a 0-handle, two 1-handles and a 2-handle, with products in between. In terms of this decomposition into elementary cobordisms, the second 1-handle is attached to the top of the product above the first 1-handle. However, we can let the attaching map for the second 1-handle flow down along the gradient vector field through the product and past the first 1-handle (as long as we are not unlucky and don't get sucked into the first 1-handle's critical point - this is again an issue to be discussed more carefully soon), and then see both 1-handle attached simultaneously to the boundary of the 0-handle.

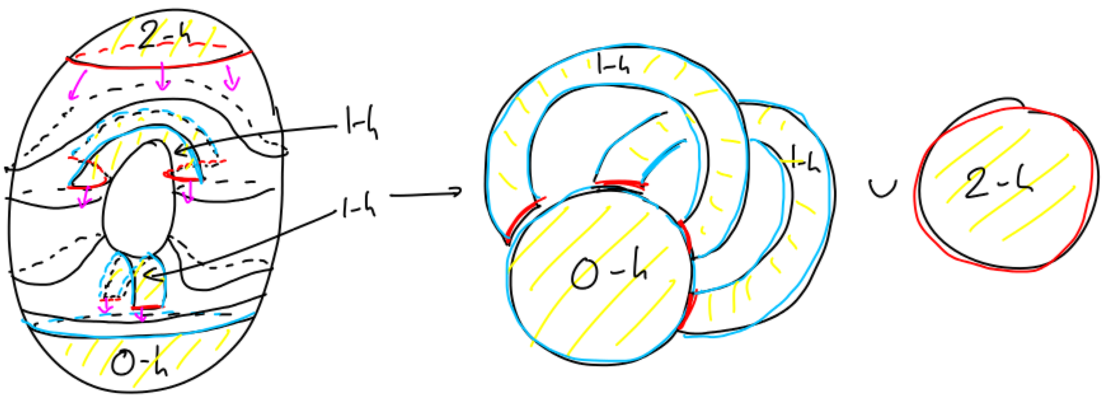

Dimension 3**:** Now it gets interesting. First, we must abandon hope of embedding the 3-manifold in $\mathbb{R}^3$ and seeing the Morse function as the height function. So instead we will just draw some handle decompositions and, perhaps, some level sets of the Morse functions. First recall the three kinds of handles in dimension 3:

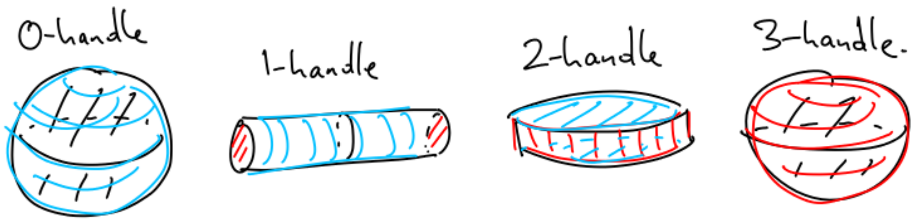

We can put these together as follows, for a simple example:

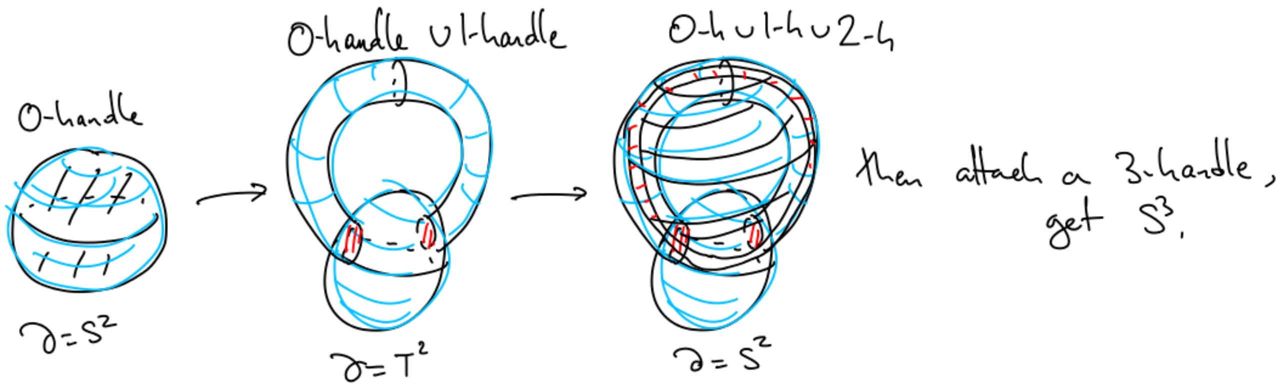

Note that, after attaching the 2-handle, we have a ball again, so we might as well not have attached the 1- and the 2-handle at all. I.e. these two handles can be *cancelled* in a way that will be made precise in due time. (This is the convertible roof, see end of Lecture 5 video for the hand gestures,) This picture is hard to look at so we can flatten it and draw only the images of the attaching maps in the boundary of the 0-handle (identifying S^2 with $\mathbb{R}^2 \cup \infty$) as follows:

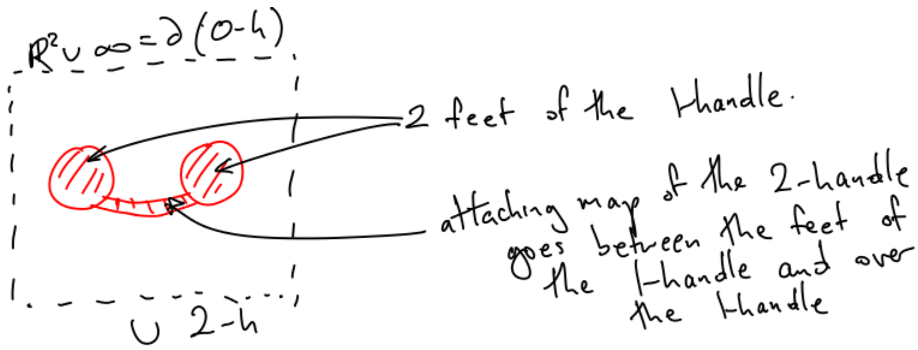

Note that we could have many 1-handles, so some labelling of the feet is appropriate, and note that we only need to draw the core of the attaching map of each 2-handle (the image of $S^1 \times \{0\} \subset S^1 \times B^1$) to specify the isotopy class of the attaching map. In fact, the same could be said for 1-handles (assuming everything is oriented) but it is visually convenient to draw the whole disk. So here is another example:

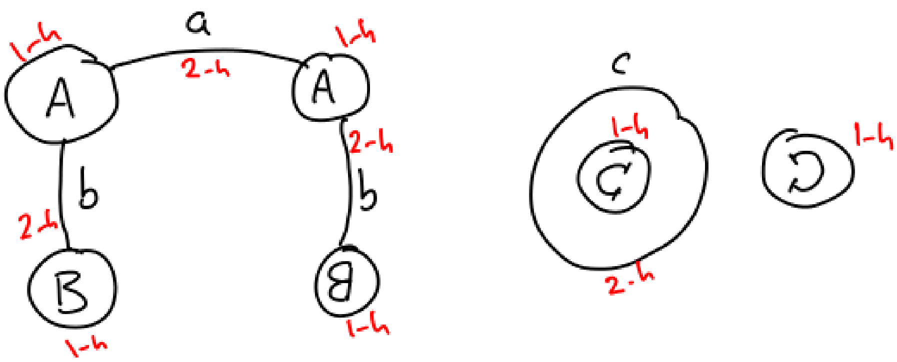

So this has three 1-handles, labelled A, B and C, and three 2-handles, labelled a, b and c, and, of course, a 0-handle that is the "background" to this picture and a 3-handle that caps everything off. First, to see that you can cap it off with a 3-handle you need to verify that the boundary is S^2.

Exercise: show that this manifold is $S^1 \times S^2$.

I'll end here, although in the lecture I then discussed Heegaard splittings. I'll write that up next time. Here's the video:

(YouTube link)

2012-01-25

Exercise: Let Σ be a torus with two boundary components embedded in $\mathbb{R}^3$ as in the top picture at right, and let $f : \Sigma \to [0, 1]$ be projection onto the vertical direction, a Morse function with two critical points of index 1 and with $a \cup b$ being a level set. Identify Σ with the square-with-opposite-sides glued minus 2 disks shown below the embedded picture, so that the indicated curves a, b and c match up. On the square picture, we have an automorphism $\phi : \Sigma \to \Sigma$ obtained by rotating the square 90°. Let $f_0 = f$ and let $f_1 = f \circ \phi$. Find a generic homotopy f_t from f_0 to f_1, constant on $f^{-1}(0)$ and $f^{-1}(1)$. Generic means that f_t is Morse for all but finitely many times t, and when not Morse, we have a simple birth or death of a pair of cancelling critical points of successive index. Also, if you don't consider distinct critical points with the same critical value as being Morse, then you should also allow finitely many times when two critical points cross. Draw the *Cerf graphic* for this homotopy, i.e. the graph in $[0, 1] \times [0, 1]$ of the critical values with their indices for all $t \in [0, 1]$.

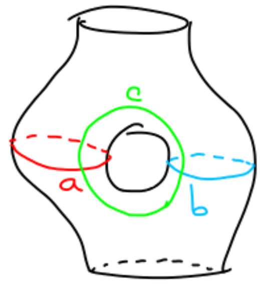

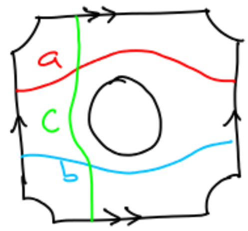

Heegaard splittings and Heegaard diagrams: We need three facts here that will be proved later:

1. For any closed connected n-manifold X there is a Morse function on X with exactly one index 0 critical point (minimum) and one index n critical point (maximum).
2. Given any Morse function $f : X \to \mathbb{R}$, there is a homotopy of maps $f_t : M \to \mathbb{R}$, with $f_0 = f$, and f_t Morse for all t, such that the critical values of f_1 are ordered by index. In other words, if the indices of critical point p and q are i and j, resp., and if $i < j$, then $f(p) < f(q)$. (Note that there if f does not satisfy this property then there will necessarily be times t where f_t has two critical points with the same critical value, and these times should be isolated. We call these times "critical value crossing times".)
3. Given a Morse function $f : X \to \mathbb{R}$ with critical values ordered by index, a generic choice of an adapted metric allows us to assume that all the k-handles are attached simultaneously to the boundary of the union of the handles of index $< k$. In other words, if p and q are critical points of index k with $f(p) < f(q)$ and no critical values in $(f(p), f(q))$, we can use the gradient flow to compare the free region of the boundary of the handle for p and the attaching region of the boundary of the handle for q inside an intermediate level set $f^{-1}(y)$ for $y \in (f(p), f(q))$. When the metric is chosen generically these will be disjoint, and thus we can flow from the attaching re-

gion of the handle for q down along the gradient field past the handle for p and see them as both attached to a level set below $f(p)$.

(As a consequence of these results, one sometimes works with *self-indexing Morse functions*, functions with the property that, for a critical point p, the index of p equals $f(p)$. Note that these are not quite generic because distinct critical points do not have distinct critical values, so some might consider these not to be Morse functions, but that might be being too nitpicky. I will not use that terminology much, but you see it a lot in 3-manifold topology.)

Now consider a closed connected oriented 3-manifold X^3 with a Morse function $f : X \to \mathbb{R}$ and corresponding handle decomposition as in 3 above. Let y be a regular value between the index 1 and index 2 critical values, and let $A = f^{-1}(-\infty, y]$, $B = f^{-1}[y, \infty)$ and $\Sigma = f^{-1}(y)$. Then A is the result of attaching some number g of 1-handles to a ball (0-handle) and $\Sigma = \partial A$ is a genus g surface. On B, consider the Morse function $-f$; the index 2 critical points of f become index 1 criitical points for $-f$ and thus B is the result of attaching g' 1-handles to a ball and $\Sigma = \partial B$ is a genus g' surface. Therefore $g = g'$, so f has the same number of index 1 and 2 critical points, and both A and B are diffeomorphic to the standard genus g handlebody (the solid object in $\mathbb{R}^3$ bounded by the standard embedding of a genus g surface in $\mathbb{R}^3$). This decomposition of X into two solid handlebodies is called a *Heegaard splitting* of X.

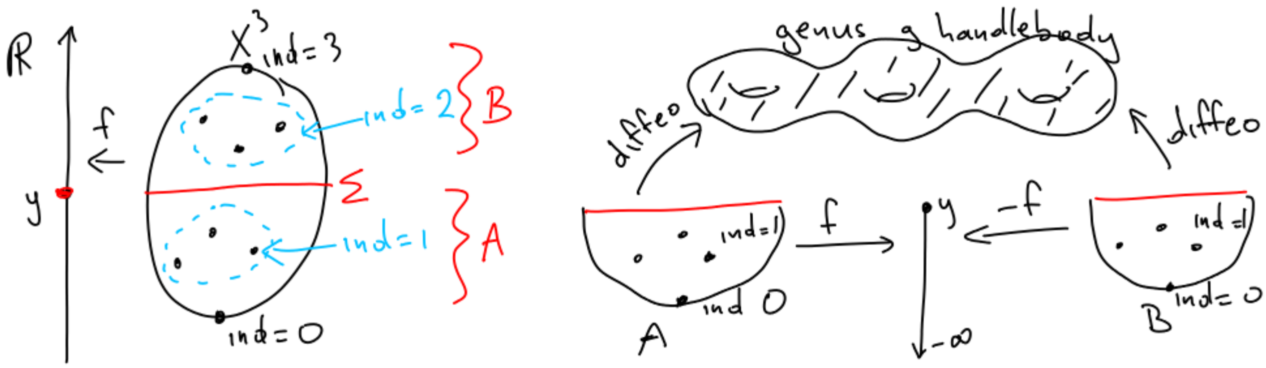

Thinking now of constructions of manifolds, rather than decompositions of manifolds, we get the related notion of a *Heegaard diagram*. In the above figure, noting that Σ is diffeomorphic to the standard genus g surface Σ_g, we can now instead consider the Morse functions $-f$ on A and f on B, in which case we see both A and B as built by attaching g 2-handles and a 3-handle to Σ_g. In other words, X is built (or, rather, a 3-manifold diffeomorphic to X is built) by starting with $[-1, 1] \times \Sigma_g$ and attaching g 2-handles and a 3-handle to $1 \times \Sigma_g$ (producing B) and then turning things upside down and attaching g more 2-handles and a 3-handle to $-1 \times \Sigma_g$ (producing A). This construction is completely determined by the $2g$ attaching circles (simple closed curves), often labelled $\alpha_1, \ldots, \alpha_g$ for A and $\beta_1, \ldots, \beta_g$ for B. The α curves must be mutually disjoint and their complement in Σ_g must be a $2g$-punctured sphere; ditto for the β curves. Any such collection of simple closed curves in Σ_g determines a closed 3-manifold; this is a Heegaard diagram. The example below is $S^1 \times S^2$, again:

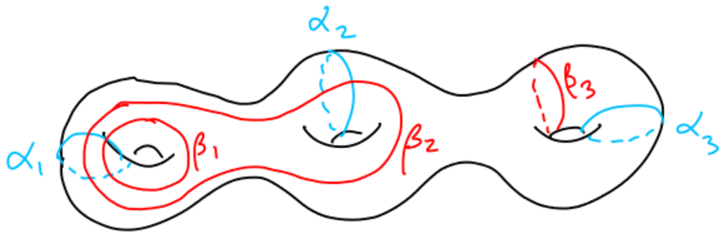

This is the data that is used to compute the famous Heegaard-Floer invariants of 3-manifolds, but that story is beyond the scope of this course.

Surgery: The comment above that, in a Heegaard diagram, the complement of the α curves (resp. β curves) should be a $2g$-punctured sphere merits further discussion. This is a condition that guarantees that, after attaching the 2-handles along these curves, the new boundary is S^2 and so we can cap off with a 3-handle. *Surgery* is the process by which the boundary of a manifold changes when one attaches a handle along that boundary. If X^n is a cobordism from M_0 to M_1 and we attach an n-dimensional k-handle along an embedding $\phi : S^{k-1} \times B^{n-k} \hookrightarrow M_1$, we get a new cobordism from M_0 to M_2, and M_2 is diffeomorphic to $(M_1 \setminus \phi(S^{k-1} \times B^{n-k})) \cup_\phi B^k \times S^{n-k-1}$, because we "cover up" the image of the attaching region of the handle and "expose" the free region of the handle. The free region $B^k \times S^{n-k-1}$ is glued via $\phi : S^{k-1} \times S^{n-k-1} \hookrightarrow M_1$. This is $(n-1)$-dimensional $(k-1)$-surgery.

In the Heegaard diagram case, we have $n = 3$ and $k = 2$, so that we are doing 2-dimensional 1-surgery by removing a $S^1 \times B^1$ and replacing with $B^2 \times S^0$. In other words, we cut open Σ along the attaching curve, leaving two new boundary components, and then we cap off each component with a disk.

We just touched on 4-manifolds in this lecture, but in the notes I'll start that in the next post. Here's the video (thanks Eddie)

(YouTube link)

2012-01-27

(At the beginning of class, Eric Burgess presented a solution to the problem of seeing the 1-parameter family of Morse functions associated to projecting a particular embedding of S^1 to various directions, see the video at the end of this post.)

4-**manifolds**: Now we extend the ideas from dimension 3 to dimension 4 to do our best to draw pictures of 4-manifolds built from handles. We start with a single 0-handle. It's boundary is $S^3 = \mathbb{R}^3 \cup \infty$, which we draw as $\mathbb{R}^3$, the ambient space in which the rest of the drawing will happen. This is just the background of our drawing - i.e. we don't really "draw" anything, we just imply it. We will assume no handles except the final 4-handle will have attaching maps which hit ∞.

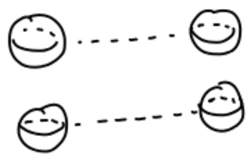

A 4-dimensional 1-handle is $B^1 \times B^3$ attached along $S^0 \times B^3$, a pair of balls. We draw these balls as small balls in $\mathbb{R}^3$, with dotted lines connecting the "feet" of each 1-handle so we know which ball goes with which ball. Thus the illustration at right simply describes a 0-handle with two 1-handles attached. Note that the new boundary, after attaching these 1-handles, is not S^3 because the interior of each ball is no longer in the boundary of the 4-manifold, and each $S^0 \times B^3$ has been replaced with $B^1 \times S^2$. We should visualize each pair of balls as indicating "worm-holes" allowing us to tunnel from one region of space to another along an interval's worth of S^2's. In particular, this new boundary here is exactly the connected sum of two copies of $S^1 \times S^2$. Thus we could not attach a 4-handle at this point.

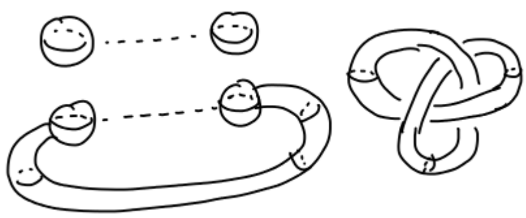

A 4-dimensional 2-handle is $B^2 \times B^2$ attached along $S^1 \times B^2$, i.e. a solid torus. They can go over 1-handles or not, as show in the diagram at right, involving two 1-handles and two 2-handles, one of which goes over one 1-handle once and one of which does not go over any 1-handles. Unfortunately just drawing the images of the attaching maps now no longer determines the attaching maps up to isotopy. To understand this better we now consider framings.

Framings: Consider an embedding $\phi : S^{k-1} \hookrightarrow M^{n-1}$. Under what conditions does this extend to an embedding $\Phi : S^{k-1} \times B^{n-k} \hookrightarrow M$? An obvious necessary condition is that the image K of the embedding ϕ should have a trivial normal bundle, and this is a sufficient condition by the tubular neighborhood theorem. When $k-1=0$ we always have triviality. When $k-1=1$ we may not have triviality if M is nonorientable (e.g. the core of a Mobius band). But if M is orientable then S^1's always have trivial normal bundles. Thing get more interesting when $k-1=1$ and $n-1 \geq 4$, but we will stay away from such high dimensions for now.

Now, assuming K has trivial normal bundle ν, there may be more than one trivialization of ν (isomorphism $\nu \to K \times B^{n-k}$). The standard proof of the tubular neighborhood theorem extends to show that isotopy classes of trivializations of ν are in one-to-one correspondence with isotopy classes of extensions Φ of ϕ. A *framing* of K is precisely an isotopy class of trivializations of the normal bundle ν. We will discuss framings in full generality next time, but for now, consider two cases:

First, $n-1=0$ and $k-1=1$. So K is a pair of points in a 1-manifold. We claim that, up to isotopy, there are exactly four framings of K. These are illustrated below:

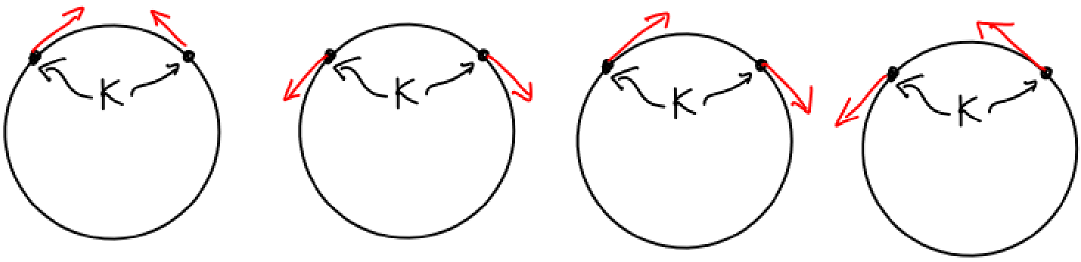

Next we show how each of these framings specifies a handle attachment:

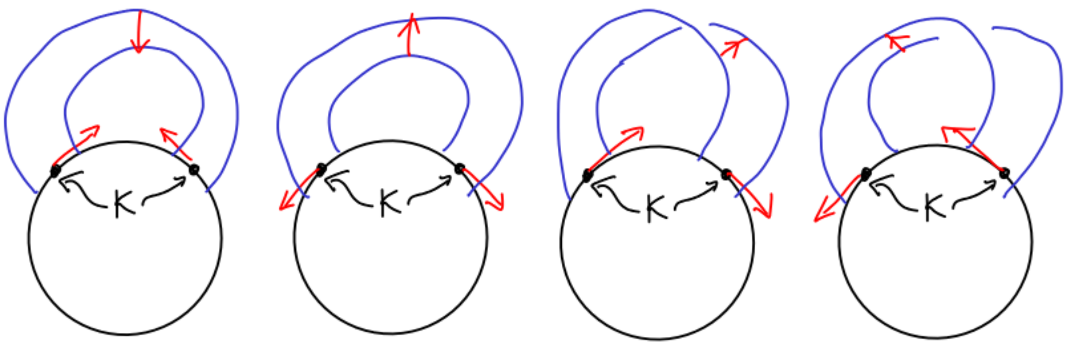

Note that the two on the left are diffeomorphic, and the two on the right are diffeomorphic, and the distinction is orientability. This example generalizes to higher dimensional 0-handles, with the upshot being that, if we agree that we are only working with orientable manifolds, we do not need to specify the framing of the attaching map of a 0-handle and need simply draw the images the attaching maps (pairs of balls). Note that we could draw points rather than balls but that the use of balls allows us to distinguish different curves going over a single 1-handle more easily than points would.

Next, consider $n-1=1$ and $k-1=2$. Now we are framing a simple closed curve in a surface, and a moment's thought shows that there are only two framings, but that handle attachment with either framing produces diffeomorphic manifolds. In fact, this phenomenon where two distinct framings produce the same manifold will persist in all dimensions, so we should immediately mod out by it. We leave this detail for the reader to sort out.

Finally for this post, consider $n-1=1$ and $k-1=3$. Now we are framing a knot K in a 3-manifold M. A framing is a pair of linearly independent normal vectors to K at each point along K (varying smoothly along K of course). To mod out by the issue mentioned in the preceding paragraph, we will assume that K is oriented, M is oriented, and that the tangent to K followed by the two normal vectors is an oriented basis, so that, up to isotopy, we do not need to specify the second normal vector. Thus the framing is just given by a single nowhere zero vector field along K, normal to K. After identifying ν with a tubular neighborhood of K, this can be seen as a parallel copy of K on the boundary of a tubular neighborhood, as shown at right.

Intuitively, this is characterized up to isotopy by the number of times it twists around K. I.e. somehow framings of K can be identified with $\mathbb{Z}$. The problem is that, in general, there is no preferred 0-framing. The best we can say in general is that the set of framings of K is a $\mathbb{Z}$-torsor, or an affine space for $\mathbb{Z}$; it looks like $\mathbb{Z}$ but we don't know where 0 is. Equivalently, given any framing, we can add or subtract 1 to it in a consistent way to produce a new framing. We use the right-hand rule for the sign convention; an example of adding $+1$ to the previous drawing is shown at right.

Another useful way to draw a framed knot is to adopt the convention that we always use the *blackboard framing*, the framing where the pushoff is in the plane of the surface on which the surface is drawn. Here are some examples; note that, to acheive extra twists we introduce small kinks in K:

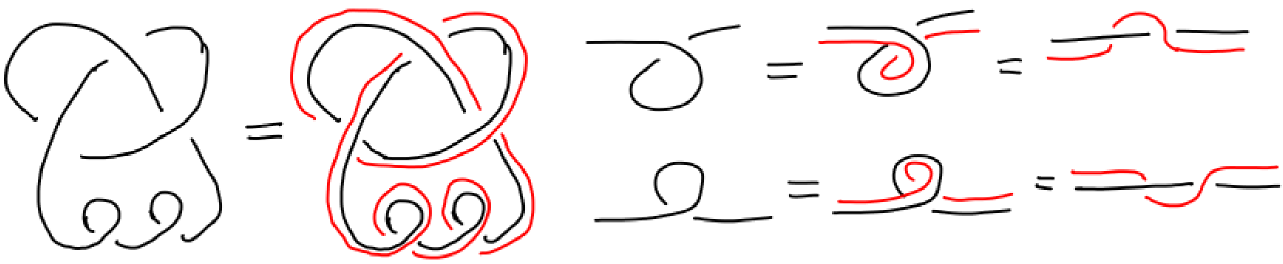

Here's the video:

([YouTube link]())

2012-01-30

Let's discuss framings more carefully: A framing of an embedded $(k-1)$-sphere K in a $(n-1)$-manifold M is an isotopy class of trivializations of the normal bundle $\nu K \cong K \times \mathbb{R}^{n-k}$. (So K needs a trivial normal bundle to begin with.) The difference between two *trivializations* is a map $K \to GL(n-k)$, but since we only care about trivializations up to isotopy, we only care about this map up to homotopy, so the difference between two *framings* is an element of $\pi_{k-1}(O(n-k))$. (Here we use that $GL(n-k)$ deformation retracts onto $O(n-k)$. Also, if $k-1 \geq 1$, we really mean $GL^+(n-k)$ and $SO(n-k)$.)

Thus the difference two framings of a knot in a 3-manifold lies in $\pi_1(SO(2)) = \mathbb{Z}$. This explains the comment last time that the set of framings is a $\mathbb{Z}$-torsor. This is worth comparing to the case of 1-knots (embedded S^1's) in a 4-manifold, in which case we have $\pi_1(SO(3)) = \mathbb{Z}/2\mathbb{Z}$, so there are two different ways to attach a 2-handle to a 5-manifold along a fixed embedding of S^1 in the 4-manifold boundary.

Now there is a special case when the set of framings of a knot $K \subset M^3$ *can* be canonically indentified with $\mathbb{Z}$, when $[K] = 0 \in H_1(M;\mathbb{Z})$. (In particular, when $M = S^3$ this holds.) In this case $K = \partial\Sigma$ for some oriented, compact surface $\Sigma \subset M$, and we use this *Seifert surface* Σ to define a preferred 0-framing. (Note that algebraic topology just tells us that K is the boundary of a singular 2-chain, but not that the 2-chain can be realized by a smooth surface. Seifert gave an explicit algorithm for constructing such a surface from a knot diagram when $M = S^3$, but there is also a general smooth topology argument using a map from $M \setminus K$ to S^1 and pulling back a regular value.) We arbitrarily orient K and then orient Σ so that the boundary orientation agrees with the given orientation on K. Then a framing of K gives a parallel push-off K' of K, with orientation coming from K. Make K' transverse to Σ and count intersections with sign ($+$ means K' goes from the negative side of Σ to the positive side and $-$ means the opposite). This associates an integer to the framing; this integer is otherwise known as the linking number $lk(K, K')$. (See Rolfsen, *Knots and Links*, for a beautiful discussion of ten different definitions of linking number.) Here are some pictures to illustrate this:

OK, now back to building 4-manifolds. We know how to draw pictures of a 0-handle with some 1-handles, and now we can add the 2-handles by drawing framed knots, some of which may go over the 1-handles. To the right is a particularly nice example, called the Mazur manifold, involving one 0-handle, one 1-handle and one 2-handle. It is contractible, roughly because the attaching circle for the 2-handle has a clasp which, if undone, could slide off the 1-handle leaving the 2-handle only going once over the 1-handle, and thus cancelling the 1-handle (c.f. convertible roof). But it is not diffeomorphic to B^4, and it's boundary is a 3-manifold with the same homology as S^3 (a *homology sphere*) but nontrivial π_1. In particular, we cannot complete this 4-manifold to a closed 4-manifold by attaching a 4-handle; to do so we would need the boundary to be S^3. Thus we see that we need to think more carefully about surgery on 3-manifolds in order to understand the boundaries of 4-dimensional handlebodies and know whether we can cap them off to make closed 4-manifolds.

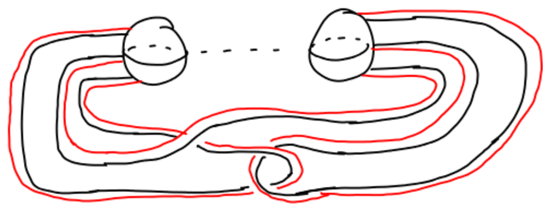

So what is surgery along a knot in a 3-manifold? If a $M^3 = \partial X^4$ and we attach a 2-handle along $\phi : S^1 \times B^2 \hookrightarrow M$ to produce a new 4-manifold X' with $\partial X' = M'$, then M' can be described precisely as:

$$M' = (M \setminus \phi(S^1 \times B^2)) \cup_\phi (B^2 \times S^1)$$

Here the gluing map is labelled ϕ, by which I mean $\phi|_{S^1 \times S^1}$. How do we see this pictorially? Let $K = \phi(S^1 \times \{0\})$ and let $\nu = \phi(S^1 \times B^2)$, a knot and a tubular neighborhood in M. Let $\mu = \phi(\{p\} \times S^1)$, the *meridian curve*, and let $\lambda = \phi(S^1 \times \{p\})$, the *longitude curve*. (Here p is just some arbitrary point in S^1.) The meridian μ is characterized up to isotopy by being homologically nontrivial in $\partial\nu$ but bounding a disk in ν. On the other hand the longitude λ depends on the framing of K, and is in fact exactly the parallel push-off we have been discussing above. (λ is partially characterized by intersecting μ once, but that's not enough.) Then M' can be described as the result of removing ν from M and glueing back in a solid torus so that, after gluing it back in, λ now bounds a disk (the B^2 in $B^2 \times S^1$) while μ does not. This is illustrated below:

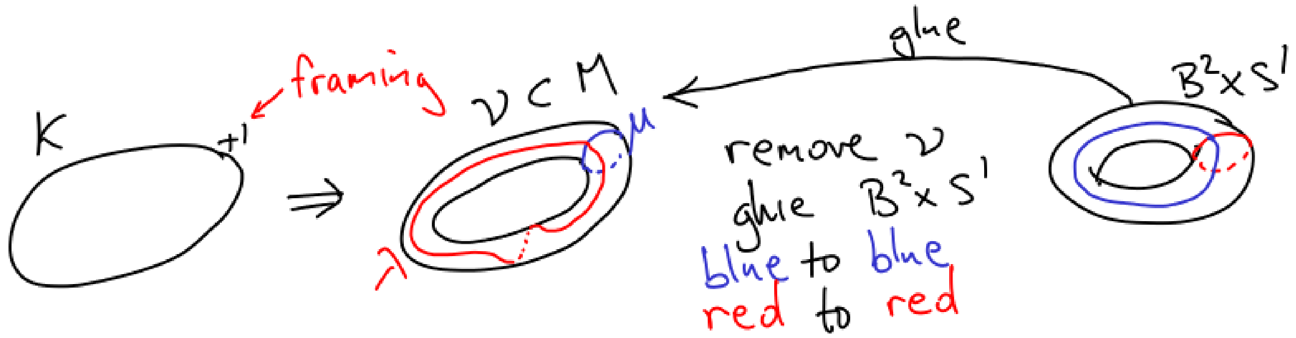

Basic examples (exercises)

1. 0-framed surgery on the unknot in S^3 gives $S^1 \times S^2$. Thus if we build a 4-manifold with a 0-handle and a 2-handle attached along the 0-framed unknot, we cannot cap it off to a closed 4-manifold by just attaching a 4-handle.
2. ± 1-framed surgery on the unknot gives S^3. Thus we can build a closed 4-manifold with one 0-, one 2- and one 4-handle, with the 2-handle attached along the ± 1-framed unknot. In fact these 4-manifolds are $\mathbb{CP}^2$ (for $+1$ framing) and $\overline{\mathbb{CP}^2}$ (for -1 framing).

I ended with an attempt to describe S^2's inside $\mathbb{CP}^2$ and $\overline{\mathbb{CP}^2}$ with self-intersection ± 1, but that's best left for the blog until the next post.

Here's the movie:

(YouTube link)

2012-02-01

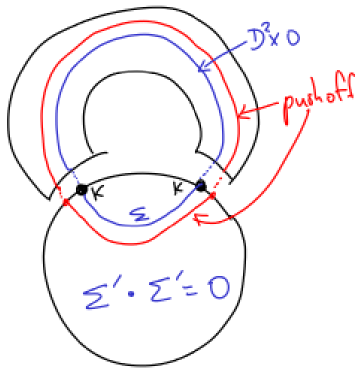

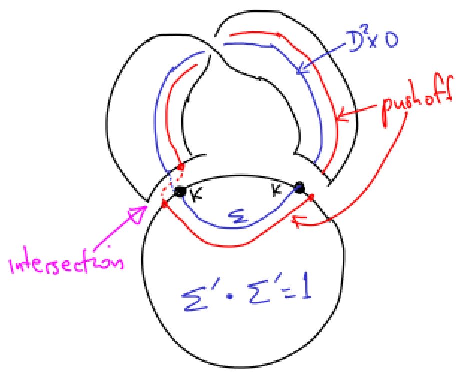

Regarding the handle decompositions given last lecture for $\mathbb{CP}^2$ and $\overline{\mathbb{CP}^2}$, we observe that attaching a 2-handle to a 0-handle with framing ± 1 to an unknot immediately yields an embedded sphere with self-intersection ± 1. We see this, and a more general version, as follows (and illustrated at right in the half-dimensional version): If we attach a 2-handle $B^2 \times B^2$ to a 4-manifold X along a *null-homologous* knot $K \subset \partial X$ with framing n, and $K = \partial \Sigma$ for some surface $\Sigma \subset \partial X$, then the interior of Σ can be pushed in to the interior of X so as to meet ∂X transversely along K. Then gluing the core $B^2 \times \{0\}$ of the handle to Σ we get a smooth closed surface $\Sigma' \subset X'$, where X' is the result of attaching the handle to X.

The self-intersection of a surface in a 4-manifold is the intersection of the surface with a nearby parallel copy of itself, counted with signs. We can push Σ off itself in X without self-intersections, and we can push $B^2 \times \{0\}$ off itself in $B^2 \times B^2$ without self-intersections, each restricting to a parallel push-off of K in ∂X. The difference between these push-offs (literally their intersection number in the boundary of a tubular neighborhood of K) is exactly the framing n of K (yes, the signs do work out right). Thus the self-intersection of Σ', denoted $\Sigma' \cdot \Sigma'$, is exactly equal to n. In the case of an unknot, Σ is a disk, so Σ' is a sphere.

Similarly, if one attaches two 2-handles along nullhomologous knots, each one produces a closed surface, and the intersection number between the two surfaces is the linking number between the knots.

Ascending and Descending manifolds: We have been drawing handle diagrams as if all the handles were attached "at once", i.e. as if the corresponding critical points had the same value. But certainly in some instances this cannot be done; as in our examples where 2-handles run over 1-handles. To understand better when we can or cannot do this, we need to see larger handles than we have seen so far. In our argument that a Morse function on a manifold gives a decomposition of the manifold into handles, each handle is contained in a small neighborhood of the critical point, and "most" of the manifold is made up of products.

Instead, consider a gradient-like vector field V for a Morse function f and a critical point p. Then the **ascending manifold for** p, A_p, is the union of $\{p\}$ and all flow lines whose backward-time limit is p, while the **descending manifold for** p, D_p, is the union of $\{p\}$ and all flow lines whose forward-time limit is p. In local coordinates near p in which $f = -x_1^2 - \ldots - x_k^2 + \ldots x_{k+1}^2 + \ldots x_n^2$, D_p is the $x_1, \ldots, x_k$-space and A_p is the $x_{k+1}, \ldots, x_n$-space. Near every other point in D_p (resp. A_p), forward (resp. backward) flow along V for some time gives a local diffeomorphism into such a local coordinate chart around p, and hence we see that D_p is a smooth k-dimensional submanifold and A_p is a smooth $(n-k)$-dimensional submanifold. They intersect transversely at p. If X is closed, so that there is no boundary to run into, they are diffeomorphic to $\mathbb{R}^k$ and $\mathbb{R}^{n-k}$.

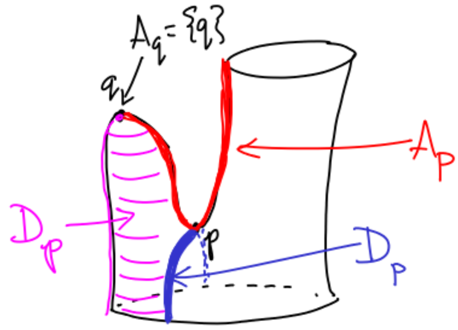

Now consider a cobordism X from M_0 to M_1 with Morse $f : X \to [0,1]$ and critical points $p_1, \ldots, p_m$, and suppose that all the ascending and descending manifolds of distinct critical values are disjoint. (Descending manifolds are always disjoing, and likewise ascending manifolds are always disjoint, but a point can be on the ascending manifold of one critical point and the descending manifold of another.) In this case, the attaching maps for the handles corresponding to the critical points all flow all the way down to M_0, and so we can see X as built from $[0,1] \times M_0$ with all the handles attached simultaneously to $\{1\} \times M_0$. In fact, the handle $B^k \times B^{n-k}$ corresponding to critical point p_i is precisely a small neighborhood of $D_{p_i} \cup A_{p_i}$, with D_{p_i} being the core $B^k \times \{0\}$ and A_{p_i} being the *co-core* $\{0\} \times B^{n-k}$.

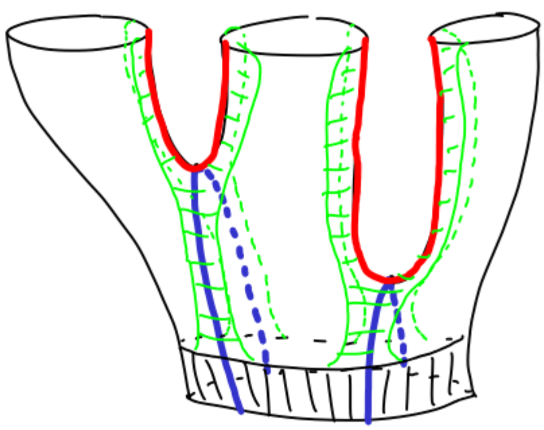

In the next post we will investigate circumstances under which we can arrange for this disjointness to occur.

([YouTube link])

2012-02-03

Consider two critical points $p, q \in X$ of a Morse function f on X^n, of indices k and l, respectively, with $f(p) < f(q)$ and no critical values in between. We want to investigate conditions under which we can assume that their ascending and descending manifolds, A_p and D_q, can be assumed to be disjoint. Of course, A_p and D_q are not defined until we choose a metric g or, at least, a gradient-like vector field V.

Note first that, in between $f^{-1}(f(p))$ and $f^{-1}(f(q))$, everything is a product and is determined by behavior in $f^{-1}(y)$ for some regular $y \in (f(p), f(q))$. Thus we look at $A_p^y = A_p \cap f^{-1}(y) \cong S^{n-k-1}$ and $D_q^y = D_q \cap f^{-1}(y) \cong S^{l-1}$, all inside the $(n-1)$-dimensional manifold $f^{-1}(y)$.

Next we note that any isotopy of A_p^y (resp. D_q^y) can be realized by homotoping the vector field, and thus the metric g, inside $f^{-1}[y-2\epsilon, y-\epsilon]$ (resp. $f^{-1}[y+\epsilon, y+2\epsilon]$). Again, this uses the product structure on $f^{-1}[y-2\epsilon, y-\epsilon]$ and just spreads the isotopy out across this product. Furthermore, any homotopy of g or V moves A_p^y and D_q^y by (independent) isotopies in $f^{-1}(y)$. Thus we can apply the transversality theorem to say that g (or V) can be homotoped to make $A_p^y \cap D_q^y$ transverse in $f^{-1}(y)$ and that, if they are transverse, a small perturbation of g or V will keep them transverse.

So now we assume that A_p^y and D_q^y intersect transversely in $f^{-1}(y)$ and now we count the dimension of their intersection. Recall that, for transverse intersections, the mantra is *"codimensions add"*. A_p^y has dimension $n-k-1$ in the $(n-1)$-manifold, hence codimension k. D_q^y has dimension $l-1$, hence codimension $n-l$. Thus $A_p^y \cap D_q^y$ has codimension $n+k-l$, hence dimension $n-1-(n+k-l) = l-k-1$. This is negative if $l < k+1$ or $l \leq k$. Thus we can assume that $A_p \cap D_q = \emptyset$ as long as $l \leq k$.

As a corollary, if a cobordism X has a Morse function with all critical points of the same index, then X can be built as a handlebody with all the handles attached at once to the bottom level.

Note that if $l = k+1$ then $A_p^y \cap D_q^y$ has dimension 0, i.e. points, in which case we have isolated flow lines from q down to p. In terms of handles, the handle for q *"goes over"* the handle for p; we have seen many examples of this when $k = 1$ and $l = 2$.

Now we want to consider what sorts of intersections between A_p and D_q to expect as we move through a 1-parameter family of Morse functions and gradient-like vector fields. The first case to consider is where f stays fixed, but the vector field varies as V_t, $t \in [0,1]$. Now consider $\mathcal{A}_p$ and $\mathcal{D}_q$ in $[0,1] \times X$, defined by $\mathcal{A}_p \cap \{t\} \times X = \{t\} \times \mathcal{A}_{p,t}$, where $\mathcal{A}_{p,t}$

is the ascending manifold for p with respect to V_t (and similary for $\mathcal{D}_q$. A similar argument to the preceding case shows that, if we want to move $\mathcal{A}_p$ through an isotopy in $[0,1] \times f^{-1}(y)$ (remaining transverse to the slices $\{t\} \times X$), we can do this by homotoping the homotopy V_t in a slab $f^{-1}[y-2\epsilon, y-\epsilon]$ (and comparable statement for $\mathcal{D}_q$. And, similarly, any homotopy of the homotopy V_t moves these manifolds by isotopies. Thus, again, transversality applies and we can assume $\mathcal{A}_p \cap \mathcal{D}_q$ is transverse.

Now when we count dimensions we discover that this intersection, if transverse, should be empty if $k < l$. Thus, for example, in a 1-parameter family we do not expect a critical point of index 1 to suddenly develop a flow line down to a critical point of index 2. But, if $k = l$, then we expect $\mathcal{A}_p \cap \mathcal{D}_q \cap ([0,1] \times f^{-1}(y))$ to have dimension 0. This means that at isolated times, there will be a single point of intersection between $A^y_{p,t}$ and $D^y_{q,t}$ or, equivalently, a single flow line from q down to p. Such events are called **handle slides**. Below is a simple example that justifies this term; we will discuss handle slides more carefully next time.

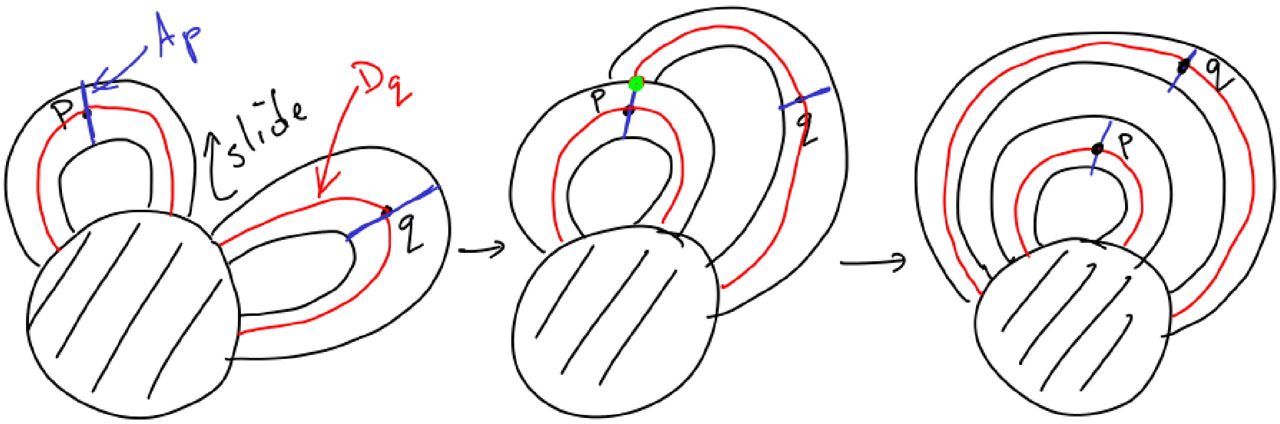

(YouTube link)

2012-02-06

We spent today's class with students presenting solutions and/or half-baked ideas about exercises. We had a complete solution to all of the various $S^1 \times S^2$ problems. The problem of showing that the space of metrics adapted to a fixed Morse function is connected (by which I really meant path-connected) was reduced to the following question:

Let f be a standard Morse model function $f = \sum \pm x_i^2$ on $\mathbb{R}^n$ and let $\phi : \mathbb{R}^n \to \mathbb{R}^n$ be any orientation-preserving diffeomorphism sending 0 to 0 and respecting f, i.e. $f \circ \phi = f$. Show that ϕ is isotopic to the identity through a 1-parameter family of maps f_t with $f_t(0) = 0$ and $f_t \circ \phi = f_t$.

A suggestion for showing that $f(x) = x^2$ is stable was to use the fact that any function (in particular, a 1-parameter family f_t) can be approximated by polynomials. Another approach suggested was to show that f_t, for small t, has "the same kind of singularity" that f has, where "same kind" means $f'(x_t) = 0$ and $f''(x_t) > 0$.

That's it.

2012-02-08

We have seen that, in a 1-parameter family V_t of gradient-like vector fields for a fixed Morse function f, we should generically expect isolated times at which ascending and descending manifolds of critical points of the same index intersect in intermediate regular levels. It is not hard to generalize this to the case where the function varies as well, in which case we have a pair (f_t, V_t). As long as f_t remains Morse and critical values do not cross, we can apply all the same transversality arguments from before, letting $\mathcal{A}_p$ be the descending manifold in $[0,1] \times X$ of an arc of critical points labelled p.

We also identified these isolated times as **handle slides** and showed one example where the total dimension is $n = 2$ and the critical points have index $k = 1$. We want to investigate this more generally.

The first point to make is that, in calling these events "handle slides", we are really describing a particulation operation on handle attaching maps (framed embedded spheres), and claiming that this operation is exactly how the handle attaching maps change from before one of these isolated time events to after the event. So first I will attempt to describe this operation.

Exercise: Generalize the following examples to an operation that makes sense for any dimension $n > 2$ and any index k with $1 < k < n$. We exclude 0 and n because you need some ascending and some descending manifold to get the discussion started. We exclude 1 because we have already discussed it and because it is hard to make sense of many smooth operations on 0-manifolds; e.g. what is the connected sum of two S^0's? The operation we are looking for should take two framed S^{k-1}'s, K_p and K_q, in a $(n-1)$-manifold, and produce a new framed S^{k-1} K_q' which results from sliding q over p.

Example: n=3, k=2 : Here K_p and K_q are framed S^1's in a surface M^2, in which case there is only one framing so we ignore the framing completely. The resulting K_q' is an embedded S^1 in M such that $K_p \cup K_q \cup K_q'$ together bound a pair of pants. This is illustrated below:

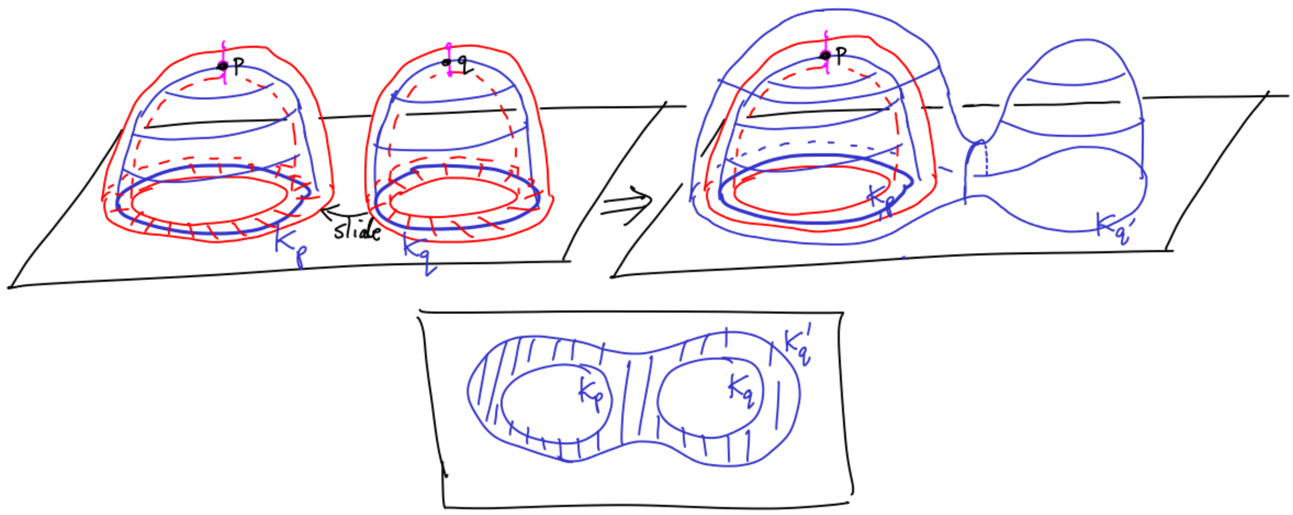

In this lecture I then proceeded to describe the 4-dimensional version $n = 4$, $k = 2$, but once again the exposition improved with the review in the next lecture, so I'll save it for the next post.

(YouTube link)

2012-02-10

Today's goal is to present the handle slide move in dimension 4, for 2-handles, and then justify this move in both dimensions 3 and 4. I claim that, in both the cases $n = 3, k = 2$ and $n = 4, k = 2$, the move can be described as follows (bearing in mind that the framing is irrelevant when $n = 3$): Let $K_p, K_q \subset M^{n-1}$ be the framed descending S^1's for index 2 critical points p and q, resp., with $f(p) < f(q)$, and M a level set below p and q. Then the result of sliding q over p is that K_q is replaced by K_q', where $K_p \cup K_q \cup K_q' = \partial\Sigma$ and $\Sigma \subset M$ is an embedded pair of pants *realizing the given framings of* K_p *and* K_q. The framing of K_q' that results from the slide is exactly the framing coming from Σ. Below is an example when $n = 4$, so that the drawing takes place in a 3-manifold:

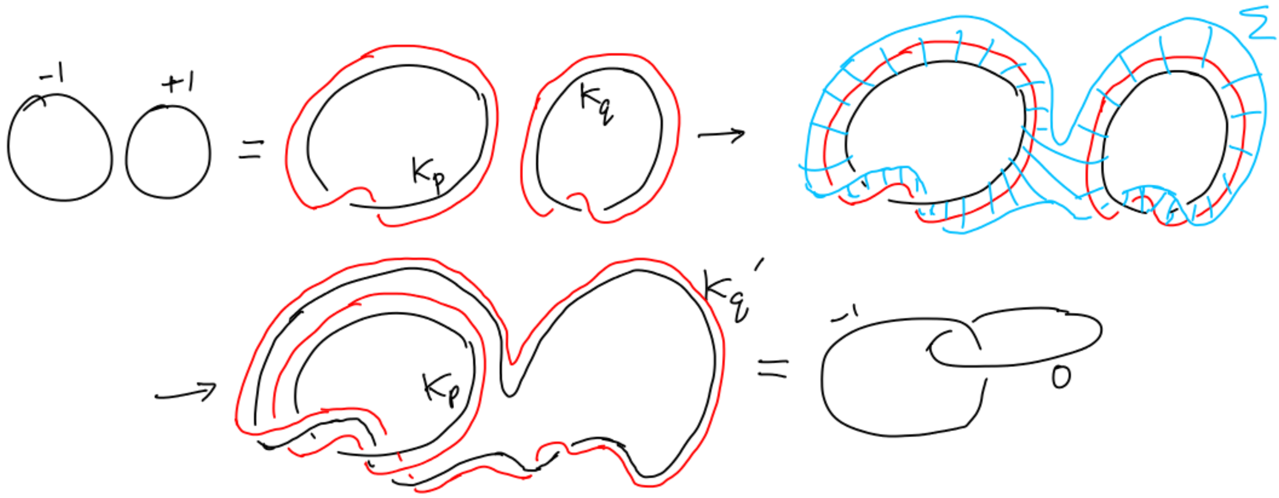

Now we want to simultaneously justify the following two statements: (1) If two handle diagrams are related by a handle slide then they describe diffeomorphic manifolds. (2) If two Morse functions (with gradient-like vector fields) are related by a homotopy in which the function remains Morse, then their corresponding handle diagrams are related by handle slides. To do this, consider two critical points p and q of the same index k with $f(p) < f(q)$ and let us follow their ascending and descending manifolds in two different regular level sets: $f^{-1}(y_0)$ and $f^{-1}(y_1)$, with $y_0 < f(p) < y_1 < f(q)$. We focus on times just before and just after a time t_0 at which there is a single point of intersection between D_q and A_p in $f^{-1}(y_1)$ (a transverse intersection between $\mathcal{D}_q$ and $\mathcal{A}_p$). In $f^{-1}(y_1)$, D_q^y moves around by an isotopy, crossing A_p^y (transversely in time) at time t_0. But in $f^{-1}(1)$, D_q makes a discrete jump somehow from before t_0 to after. As D_q^y crosses A_p^y, it sweeps out an annulus punctured once by A_p^y. Removing a disk neighborhood of this puncture from the annulus, we get a pair of pants $\Sigma \subset f^{-1}(y_1) \setminus A_p^y$ with boundary the union of D_q^y before t_0, D_q^y after t_0, and the boundary of the disk we removed from the annulus. Since Σ is disjoint from A_p, it can flow down to $f^{-1}(y_0)$. The "before" and "after"versions of D_q^y

flow down to become "before" and "after" attaching spheres for q (framed by Σ) while the boundary of the disk we removed flows down to become a parallel push-off of the attaching sphere for p.

This demonstrates directly that the singular event in the Morse function movie corresponds to a handle slide in the handle diagram. To go the other way, note that the previous paragraph also provides a construction of a Morse function movie that corresponds to a given hande slide; this, together with the fact that handle diagrams uniquely determine manifolds up to diffeomorphism, shows that, when two diagrams are related by handle slides, then they describe diffeomorphic manifolds.

I claim that the above argument also works in higher dimensions and different indices, but leave that to the reader to sort out. Also, as a suggestion, it might be useful to construct the pair of pants and its higher-dimensional generalizations as a framed cobordism in $[0,1] \times f^{-1}(y_0)$, with a single critical point for the Morse function arising from projection to the $[0,1]$ factor.

(YouTube link)

2012-02-13

We have been a bit unclear, when discussing handle slides, about whether we are thinking of families (f_t, V_t) of Morse functions paired with gradient-like vector fields or of a fixed Morse function f with a family of vector fields V_t. I claim that this distinction is not important due to the following fact:

Lemma: If $f_t : X \to [0,1]$ is Morse for all t (with distinct critical values for all t), with $t \in [0,1]$, then there exist isotopies $\phi_t : X \to X$ and $\psi_t : [0,1] \to [0,1]$, with ϕ_0 and ψ_0 identity maps, such that $f_t = \psi_t \circ f_0 \circ \phi_t$.

Thus, as long as there are no crossings of critical values, we may pull everything back by ϕ_t and ψ_t to treat f_t as constant in t. When there are crossings, in which case we have no hope of making the function constant in time, we can arrange that no handle slides occur in a short time interval around the crossing.

It is obvious that two Morse functions cannot in general be connected by a path of Morse functions, since critical points remain discrete and therefore the number of critical points would need to be constant. However, they can be connected by a *generic homotopy*, which we define as follows:

Definition: a generic homotopy is a homotopy $f_t : X \to [0,1]$, $t \in [0,1]$, between Morse functions f_0 and f_1 such that, near every point $p \in X$ and time t_0, there exist coordinates τ around $t_0 \in [0,1]$ and τ-dependent coordinates $x_1^\tau, \ldots, x_n^\tau$ around $p \in X$, and τ-dependent coordinates y^τ around $f_{t_0}(p)$, with respect to which f_t has one of the following three local models (in which we suppress the dependence of the x_i's and y on τ):

1. $(x_1, \ldots, x_n) \mapsto x_1$; i.e. there is no singularity here.
2. $(x_1, \ldots, x_n) \mapsto -x_1^2 - \ldots - x_k^2 + x_{k+1}^2 + \ldots + x_n^2$; i.e. p is an index k Morse singularity for f_{τ_0} and there is a path p_t, $t \in (t_0 - \epsilon, t_0 + \epsilon)$, with $p = p_0$ such that p_t is an index k Morse singularity for f_t.
3. $(x_1, \ldots, x_n) \mapsto -x_1^2 - \ldots - x_k^2 + x_k + 1^3 - \tau x_{k+1} + x_{k+2}^2 + \ldots x_n^2$; i.e. a birth or death of a pair of critical points of index k and $k+1$ occurs at p at time t_0.

The **claim**, which we offer without proof or perhaps defer to a later date, is that generic homotopies are *generic and stable.* (This emphasizes, of course, the poor choice of the adjective "generic" in the definition.) Here stable means that a generic homotopy is still a generic homotopy after a small perturbation, and generic means that any homotopy can be perturbed to generic by an arbitrarily small perturbation.

Thus, between any two Morse functions, we can always find a homotopy which remains Morse except at isolated times when a birth or death occurs (and, if we insist that Morse functions have discrete critical values then we also count critical value crossings as special isolated non-Morse events).

Immediately after a birth has occurred, producing critical points p and q with indices k and $k+1$, respectively, then $f(p) < f(q)$ and there are no critical values in $(f(p), f(q))$. Also, there exists a gradient-like vector field with respect to which there is a single flow line from q down to p; i.e. $A_p^y \cap D_q^y$ is a single point in $f^{-1}(y)$ for some $y \in (f(p), f(q))$. In terms of handles, this means that the handle attaching sphere S^k for q "goes over" the handle $B^k \times B^{n-k}$ for p once, intersecting the *belt sphere* $\{0\} \times S^{n-k-1}$ at one point. We have already examples of such diagrams.

In the next post I'll explain the converse of this, namely that if a $(k+1)$-handle goes over a k-handle once then we can cancel them; Gompf and Stipsicz give a non-Morse theory proof which is more basic, but we will construct a generic homotopy cancelling the two critical points.

(YouTube link)

2012-02-15

Note that the assertion that A_p^y and D_q^y, the ascending and descending spheres for critical points p and q in an intermediate regular level y, meet at a single point, is equivalent to the assertion that there is a unique gradient flow line from q down to p.

In this entry we want to sketch the proof of the following:

Theorem: If $f : X \to [0, 1]$ is Morse with exactly two critical points p and q with a unique gradient flow line from q down to p, with ascending and descending manifolds meeting transversely, then there is a generic homotopy f_t from $f_0 = f$ to f_1 which cancels p and q.

Sketch of proof: Here's what I send in lecture, *but actually it's subtly* ***wrong***: Find an arc $A \cong [0, 1]$ embedded in X containing this unique flow line as its middle third $[1/3, 2/3]$, with q at $1/3$ and p at $2/3$, and with $[0, 2/3)$ contained in the descending manifold for q and $(1/3, 1]$ contained in the ascending manifold for p. Thus, up to reparametrization, $f|_A$ looks like $f(x) = x^3 - x$. Now we claim that there is a tubular neighborhood ν of A, with coordinate x_{k+1} on A and coordinates $x_1, \ldots, x_k, x_{k+2}, \ldots, x_n$, where k is the index of p, such that $f|_\nu = -x_1^2 - \ldots - x_k^2 + x_{k+1}^3 - x_{k+1} + x_{k+2}^2 + \ldots + x_n^2$. The idea is that, along the given flow line, A_p and D_q intersect transversely, so that the descending coordinates $x_1, \ldots, x_k$ come from D_q and the ascending coordinates $x_{k+2}, \ldots, x_n$ come from A_p. Once we have this local model, we can cancel the critical points using $x_{k+1}^3 - t x_{k+1}$. This is illustrated below:

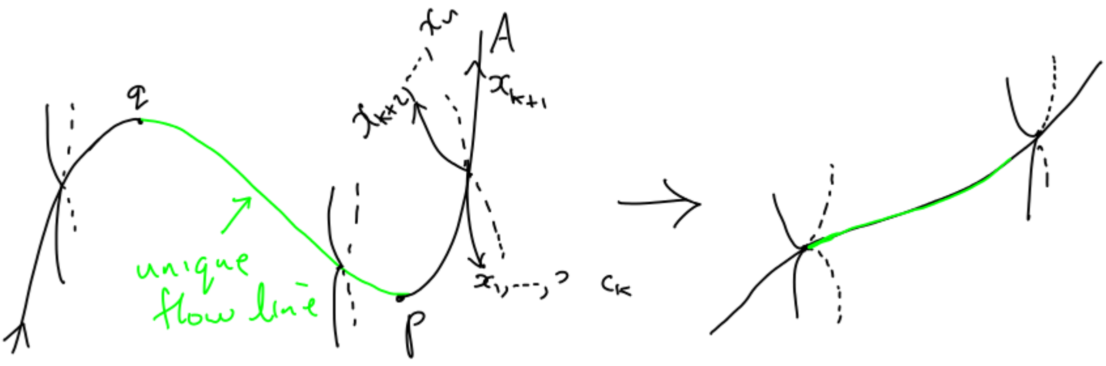

So what is **wrong** with this argument? The first problem is that, yes, one may find a local patch (in this case a tubular neighborhood of an arc) in which there is a certain local model (that much is correct in the above argument), but then one cannot blithely apply a polynomial perturbation because polynomials are not compactly supported, and we should be constructing a homotopy which is constant outside the given patch. Thus one should cut if off with a bump function. But then, when cutting things off with a bump function, one has the potential to accidentally create new critical points, as illustrated in this picture:

So I owe a proper sketch of this proof - the point is that one really does need to work with the full descending manifold for q and the full ascending manifold for p. $\square$