THE DOLD-THOM THEOREM

THOMAS BARNET-LAMB

ABSTRACT. We give a proof of the Dold-Thom theorem, following Hatcher [3]. This states that for a connected CW complex X, $H_i(X) \cong \pi_i \circ SP(X)$. Here SP is a functor which we introduce, which arises naturally from both a geometric and categorical perspective, as we discuss. A brief discussion as to the significance of the result for deeper developments is given.

1. INTRODUCTION

If challenged to give a description of what the subject of algebraic topology 'is about', then (at least at an elementary level) one could do a lot worse than:

> Algebraic topology is the study of functors F (both covariant and contravariant) from **Top** (the category of topological spaces and continuous maps) to some 'algebraic category', normally the category **Grp** of groups, which have the key property of *homotopy invariance*; if $f, g : X \to Y$ are parallel morphisms which are homotopic, then $Ff = Fg$.

It is then a commonplace (again, at least at the elementary level) that there are two major families of such functors, which are as different as chalk and cheese. On the one hand there are the *homotopical* functors $\pi_0, \pi_1, \pi_2, \ldots$ which generalise the fundamental group π_1: easy to define, and apparently encoding much deep information about the space, but still (after 50 years of study) almost impossible to calculate in any real generality. On the other hand, the homological functors $H_0, H_1, H_2, \ldots$ and cohomological functors $H^0, H^1, H^2, \ldots$, which while trickier to define than the homotopical functors, and possibly yielding slightly less information about the deep structure of the space concerned, have the enormous practical advantage of being computable, through the paraphernalia of long exact sequences (attached to pairs and of the Mayer-Vietoris kind) and excision isomorphisms.

Unsurprisingly, considerable interest has grown in connections between these twin pillars of the theory. There is, of course, the Hurcewicz theorem, which tells us that, if the first nonzero homotopy group of positive degree d has $d > 1$, then it is isomorphic to the homology group of the same degree. One can also give a construction of cohomology using homotopical ideas: rather than looking at homotopy classes of maps *from* a fixed space (normally a sphere) into a varying space X (which gives $\pi_*(X)$), one can fix some space Y_0, and for a space X consider homotopy classes of maps from X to Y_0. Under appropriate conditions on Y_0, this will be a group, and it turns out that for a particular natural family of Y_0s (the $K(G, i)$), one recovers the 'usual' cohomology groups. (It also turns out, after significantly more work, that *every* cohomology theory arises in this way, for a particular family of Y_0s).

There is also the beautiful theory that has grown up around the Freudenthal suspension theorem. This allows one to define a process called 'stabilisation', which makes homotopy groups more well-behaved, and it turns out that the theory you get by taking homotopy and applying this process, so-called 'stable homotopy' is actually a homology theory (albeit one which is still not computable in generality).

One thing that is missing from this list is a way of reducing homology to homotopy. The Dold-Thom theorem, which is the subject of this essay, provides just such a connection: and it turns out that the connection which emerges is surprisingly beautiful. To state it, we must introduce the functor SP. For the present, we will just state this is a functor $\mathbf{Top}_* \to \mathbf{Top}_*$, which arises naturally from both geometric and category theoretic perspectives. (A full discussion will be given in section 2). Then for any connected CW-complex X, we have

Theorem 1 (The Dold-Thom theorem). *There is an isomorphism $H_i(X) \cong \pi_i \circ SP(X)$. (To apply SP to X we must give it a basepoint; but any basepoint will do.)*

The Dold-Thom is not, however, merely of interest as a technical curio connecting homology and homotopy theory. Recent developments have given it a much deeper significance. Let us pause for a while to examine what this significance is. To start with, it turns out that, using this theorem as a centerpiece, pretty much all of traditional homology theory can be set up without recourse to any of the standard foundations (singular, simplicial, cellular, Cech, Alexander, sheaves and derived functors, etc.): thus homotopy theory, via the Dold-Thom theorem, forms another possible foundation for homology theory (see [1] for a complete exposition of the theory, including more advanced topics like K-theory, from this viewpoint). This, of course, is not an earth-shattering discovery *per se*; homology theory already has an embarrassment of potential foundations (we've just listed six limiting ourselves to ones that might be regarded as 'standard'). Indeed a reasonable amount of work in the subject is devoted to establishing the various theorems that show that these disparate foundations in fact calculate the same object, at least for nice spaces. (When the first few equivalent foundations were discovered, it was very interesting that these seemingly-completely-different constructions were calculating the same thing, but now 'yet another foundation' is not news.)

The key point, which renders the homotopical construction of homology more than 'just another' construction, is its relevance to algebraic geometry. There has, for some time, been a desire to allow one to use the tools of algebraic topology in the study of algebraic geometry: if and when they can be made to apply, they often provide important insights into what is going on, and tools for proving other important results. However, most of the 'foundations' of algebraic topology are not at all applicable to the algebraic geometric case. Of the 'standard' foundations, only the the (sheaf theoretic) Cech and derived functor approaches have any analogue, and while study of these analogues is an absolutely central tool in the study of algebraic geometry, the analogy is not so close as one might hope. For instance, a curve of genus g, over $\mathbb{C}$, can also be realised as a Riemann surface, which is a topological object whose cohomology groups are of considerable interest. But alas, if you try and compute the analogous groups using the algebraic analogue of the theory, all the groups turn out to be zero, and therefore less interesting! (Indeed, this led to the construction of *étale cohomology*, a foundation for cohomology in algebraic geometry which has no analogues, in terms of how it is constructed, on the topological side, but which at least gives 'the right answer', roughly speaking, compared to $\mathbb{C}$, for the cohomology of curves.

Thus there is considerable interest in other foundations for homology *which have algebraic analogues*; and it turns out that the homotopical foundation, building from Dold-Thom, has a strong algebraic analogue. Voevodsky (with others) has translated many of the concepts of algebraic topology into this context. This allowed Voevodsky to prove the Milnor conjecture, which concerns a certain relationship between the Galois cohomology groups of a field F and Milnor's K theory groups of F. For more information, see the introduction to [1]. While we will not explore these connections in this essay, focusing on the topological context, it is hoped that explaining this connection gives some motivation for the result.

We turn now to some brief notes on the literature. Dold and Thom originally published their result in German ([2]). I have managed to track down only two secondary sources. First is Hatcher's omnibus work on Algebraic Topology ([3]) which covers the Dold-Thom theorem as part of its aim to be 'a background reference for many additional topics' which do not fit into a normal Algebraic Topology course. Second is the work *Algebraic Topology from a Homotopical*

Viewpoint, [1]. As we have already mentioned, this work actually carries out, completely explicitly, the construction of all the concepts of homology (and cohomology) from the standpoint of homotopy theory and SP, and along the way, essentially proves the Dold-Thom theorem[1]. It is somewhat surprising that there are so few references for such an important theorem, but I am not the only one to have searched for other references in vain: the authors of [1] mention that, at the time of publication (2002—presumably just before Hatcher's book, published the same year!) their exposition is, so far as the authors know, the only one in the literature apart from the original research announcement.

The rest of this essay is devoted to giving a proof of the Dold-Thom theorem 1, in the form stated above. (Thus we do not phrase the result as [1] does, as a fully-worked out alternative foundation for homology theory, since to do this properly, one really needs to explicate quite a substantial portion of homology theory in the new context, which we do not have room for in this essay!) In the proof itself, we shall essentially follow Hatcher, although we will occasionally incorporate elements from [1]. There are a number of technical lemmas which are used in the proof; Hatcher proves these first (which leaves one wondering for some time where the whole thing is going), I will try the alternate approach of beginning with an analysis of the Dold-Thom theorem, from which it will emerge that we could prove the theorem if we had these various technical lemmas (and that, in some sense, they are the natural technical lemmas to prove and use). Thus motivated as to their importance, we will proceed to prove the lemmas, and so finally get a proof of the theorem.

2. The functor SP

We begin with a description of the functor SP. Suppose (X, x_0) is a space-with-basepoint. We construct the n-fold symmetric product of X (which we denote SP_nX), as the quotient of the space X^n (the n-fold power of X) by the action of the group S_n, acting by permutation of the factors. We can embed X^n in X^{n+1} by sending $(p_1, \ldots, p_n) \mapsto (x_0, p_1, \ldots, p_n)$; it is clear that this descends to give an embedding of SP_n in SP_{n+1}; by taking the direct limit of these spaces (with the direct limit topology), we end up with the space $SP(X)$. It has a natural basepoint: the point with all components equal to x_0

Thus, as a set, $SP(X)$ is the same as the space you get by the following construction: take the product of countably many copies of X; quotient by the action of the symmetric group on countably many generators, and then take the subspace in which all but finitely many components are equal to x_0. Unfortunately, this does not quite give the right topology on $SP(X)$, but if one carries out the standard construction of passing to the corresponding compactly generated topology, then you do get the right topology, at least in the case of X finite.

It is clear that a map of spaces $f : X \to Y$ gives a map on $SP_kX \to SP_kY$ for all k in a natural way; these are then compatible, and give a map $SP(f) : SP(X) \to SP(Y)$. Moreover, a homotopy of maps $f, g : (X, x_0) \to (Y, y_0)$ gives rise to a homotopy of the corresponding maps on SP_k and SP. Finally, it is clear that the construction is functorial; if $f : X \to Y$ and $g : Y \to Z$ are maps, then the map induced on SP's by the composition $g \circ f$ is the same as the composition $SP(g) \circ SP(g)$. We summarise all this by saying that SP is a homotopy functor (a functor sending homotopic maps to homotopic maps).

Another way of looking at the space $SP(X)$, and one that perhaps gives a clearer picture of what it is goes as follows. Recall the notion of the free group on a set. This, roughly speaking, takes a set X and 'turns it into a group' in the 'simplest' way possible. More formally, 'taking

[1]Of course, since their whole project is to set up homology theory without mentioning the usual construction of $\tilde{H}_i$, they do not prove a theorem of the form $\tilde{H}_* = \pi_* \circ SP$, rather taking this as the definition of $\tilde{H}_*$, at least for connected complexes (there is a cunning way to extend to non-connected complexes, as we will see in due course). But they do establish that $\pi_* \circ SP$ satisfies all the usual axioms for a reduced homology theory, which then by general theory automatically shows it coincides with any of the other frameworks on CW complexes.

the free group' is the unique functor F from **Grp** to **Set** which is *left adjoint* to the natural forgetful functor U going the other way, in that for all groups G and sets S,

$$\mathbf{Set}(S, UG) \cong \mathbf{Grp}(FS, G)$$

naturally in S and G, where the LHS is the collection of set maps from S to UG, and the RHS is the collection of homomorphisms from FS to G.

Similarly, there is a natural forgetful functor U from **AbTopMonoid**, the category of abelian topological monoids, to $\mathbf{Top}_*$, the category of pointed topological spaces and basepoint preserving maps (the basepoint comes from the identity of the monoid). (A monoid is a like a group except one drops the requirement that inverses exist: and similarly, topological monoids are like topological groups without inverses.) It is possible to show that there is a functor F going the other way which is left adjoint to this forgetful functor (indeed, by the end of this discussion it should be clear how one might do this). This functor 'takes the free abelian topological monoid' on a given topological space.

What does this free abelian topological monoid FX look like, for a pointed space (X, x_0). Let's first consider it merely as a set. Well, by analogy with free groups, an element $\alpha \in FX$ looks like formal linear sums of elements of X, say $x_1+x_2+\cdots+x_n$ (we write it as a sum since order doesn't matter). We then topologise it by saying that 'small changes in any of the x_i makes a small change in the sum'; more formally, we use the finest possible topology consistent with the map $(x_1, x_2, \ldots, x_n) \mapsto x_1+x_2+\cdots+x_n$ being continuous. Finally, the monoid on FX structure comes from juxtaposition: $(x_1+x_2+\cdots+x_n)+(x'_1+x'_2+\cdots+x'_m) = x_1+x_2+\cdots+x_n+x'_1+x'_2+\cdots+x'_m$.

Thus the subset of FX that can be written using fewer than n terms looks exactly like SP_nX, at least as a topological space; taking the direct limit, we get that FX's underlying topological space is precisely $SP\,X$; that is, $UFX = SP\,X$. Thus another way of looking at SP is that it is precisely the adjunction $UF : \mathbf{Top}_* \to \mathbf{Top}_*$ arising from the adjoint pair U, F. Since free/forgetful pairs are very natural objects, this adjunction is very natural too. Now, an upshot of this is that $SP\,X$ has a natural structure as an abelian topological monoid (the structure we forgot about using U). This structure will turn out to be very useful to us, and so we will sometimes want to think of $SP\,X$ as a monoid. We will thus abuse notation and sometimes treat it as one.

We close this section with an example: we calculate, quite explicitly, the homotopy type of $SP(S^1)$.

Theorem 2. *We have that $SP(S^1) \simeq S^1$.*

Proof: It is well known that $S^1 \simeq \mathbb{C}\backslash\{0\}$, so that $SP(S^1) \simeq SP(\mathbb{C}\backslash\{0\})$, and so it suffices to show that $SP(\mathbb{C}\backslash\{0\}) \simeq S^1$ (since SP is a homotopy functor). We first consider $SP_k(\mathbb{C}\backslash\{0\})$. This consists of k-tuples of nonzero complex numbers, disregarding order. Now, it is well known that the space of k-tuples of complex numbers, disregarding order, but without the 'nonzero' restriction, is isomorphic (even topologically isomorphic) to $\mathbb{C}^k$, the isomorphism being given by

$$\mathbb{C}^k/S^k \to \mathbb{C}^k : (x_1, \ldots, x_k) \mapsto (s_1(x_1, \ldots, x_k), \ldots, s_1(x_1, \ldots, x_k)) = (r_1, \ldots, r_k) \quad \text{(say)}$$

where the s_i are the elementary symmetric polynomials, so for example

$$s_1(x_1, \ldots, x_k) = x_1 + \cdots + x_k; \quad s_2(x_1, \ldots, x_k) = \sum_{i<j} x_i x_j; \text{ and } s_k(x_1, \ldots, x_k) = x_1 \ldots x_k$$

(It is clear that the map is continuous, open, and it is a standard result of elementary algebra that it has fibres precisely the orbits of S_n). Now, it is clear that $r_k = 0$ iff at least one of the x_i is zero. Thus we can restrict to give a homeomorphism between the space $SP_k(\mathbb{C}\backslash\{0\})$ of k-tuples of non-zero complex numbers, disregarding order, with the space of k tuples of complex numbers the last of which is not zero, and we have:

$$SP_k(\mathbb{C}\backslash\{0\}) \cong \mathbb{C}^{k-1} \times \mathbb{C}\backslash\{0\}$$

Now, the right hand side deformation retracts onto a subspace which is an S^1 (say take the subspace $(0,\ldots,0,t), |t|=1$ and use as deformation retraction a linear homotopy to the retraction $(x_1,\ldots,x_n)\mapsto(0,\ldots,0,x_n/|x_n|)$. So we deduce that $SP_k(\mathbb{C}\backslash\{0\})\simeq S^1$.

Now, we would like to then say 'these deformation retractions can be made compatible with the inclusions $SP_k(\mathbb{C}\backslash\{0\})\to SP_{k+1}(\mathbb{C}\backslash\{0\})$, so we can pass to the limit and get a deformation retraction of $SP\,X$ to a subspace homeomorphic to S^1'. This is more or less all Hatcher and [1] say, but the matter involves a little subtlety (which makes it non-obvious that a consistent choice can be made), so we shall be slightly more explicit.

The subtlety arises because the inclusion $SP_k(\mathbb{C}\backslash\{0\})\to SP_{k+1}(\mathbb{C}\backslash\{0\})$, when translated into an inclusion $\mathbb{C}^{k-1}\times\mathbb{C}\backslash\{0\}\to\mathbb{C}^k\times\mathbb{C}\backslash\{0\}$ turns out to be a slightly more complicated map than one might hope. In particular, it is the map

$$\iota_k:(r_1,\ldots,r_k)\mapsto(1+r_1,r_1+r_2,\ldots,r_{k-1}+r_k,r_k)$$

as can easily be checked[2]. Thus, for instance, if we make, as above, the natural choice of $(0,\ldots,0,t), |t|=1$ as our subspace homeomorphic to S^1 in $\mathbb{C}^{k-1}\times\mathbb{C}\backslash\{0\}$ onto which we want it to retract, for each k, then we have a problem. For our subspace $(0,\ldots,0,t), |t|=1$ in $\mathbb{C}^{k-1}\times\mathbb{C}\backslash\{0\}$ maps into $\mathbb{C}^k\times\mathbb{C}\backslash\{0\}$ as $(1,0,\ldots,0,t,t), |t|=1$; so it does not map onto $(0,0,\ldots,0,0,t), |t|=1$, our chosen subset of $\mathbb{C}^k\times\mathbb{C}\backslash\{0\}$, as it should.

The problem turns out not to be too serious, however: we just need to be slightly careful in our choices. Let us pick as our embedded S^1 in $\mathbb{C}^{k-1}\times\mathbb{C}\backslash\{0\}$ the subspace E_k given by:

$$\left(\tbinom{k-1}{1}+\tbinom{k-1}{0}t,\tbinom{k-1}{2}+\tbinom{k-1}{1}t,\ldots,\tbinom{k-1}{k-2}+\tbinom{k-1}{k-3}t,\tbinom{k-1}{k-1}+\tbinom{k-1}{k-2}t,\tbinom{k-1}{k-1}t\right):t\in\mathbb{C},|t|=1$$

and pick as our deformation retraction a linear homotopy onto the retraction

$$\varphi_k:\mathbb{C}^{k-1}\times\mathbb{C}\backslash\{0\}\to E_k:(x_1,\ldots,x_k)\mapsto(\left(\tbinom{k-1}{1}+\tbinom{k-1}{0}\frac{x_k}{|x_k|},\ldots,\tbinom{k-1}{k-1}+\tbinom{k-1}{k-2}\frac{x_k}{|x_k|},\tbinom{k-1}{k-1}\frac{x_k}{|x_k|}\right)$$

Then it is easy to check that the maps φ_k are compatible with the inclusions ι_k, in that $\iota_k\circ\varphi_k=\varphi_{k+1}\circ\iota_k$, using standard binomial coefficient identities. In particular the coice of the E_k as our embedded S^1 is compatible with the ι_k. Now, since the ι_k are affine, and we pick homotopies which are linear, the compatibility of the retractions φ_k with the ι_k then implies that the *deformation* retractions are also compatible with the ι_k. So we do indeed get, passing to the limit, a homotopy equivalence $SP(\mathbb{C}\backslash\{0\})\simeq\lim_{\longrightarrow}E_k=S^1$ $\square$

3. Proving Dold-Thom modulo some technical lemmas

The usual way of showing that some construction agrees with the usual homology functors is to show that it satisfies the full Eilenberg-Steenrod axioms for a reduced homology theory, including the dimension axiom; it then follows from a general theorem (which shows that the axioms completely determine the homology of CW complexes) that the construction does just give the 'usual' groups. Let us briefly remind ourselves what these axioms are. For a theory h, we need:

(1) The h_i are functors $\mathbf{Top}_*\to\mathbf{Grp}$, and satisfy homotopy invariance: if $f\simeq g$, then $h_*f=h_*g$.
(2) There is a boundary map $\partial:h_i(X/A)\to h_{i-1}(A)$, defined for each pair (X,A) where A is a subcomplex of X, such that there is an exact sequence

$$\cdots\to h_i(A)\to h_i(X)\to h_i(X/A)\to h_{i-1}(A)\to\ldots$$

and satisfying appropriate naturality properties.

[2]Suppose we start with a tuple $(x_1,\ldots,x_k)\in\mathbb{C}^k/S_k$ with corresponding tuple $(r_1,\ldots,r_k)\in\mathbb{C}^{k-1}\times\mathbb{C}\backslash\{0\}$. This maps to the tuple $(x_1,\ldots,x_k,1)\in\mathbb{C}^{k+1}/S_{k+1}$, whose corresponding element $(r'_1,\ldots,r'_{k+1})\in\mathbb{C}^k\times\mathbb{C}\backslash\{0\}$ satisfies $r'_i=s_i(x_1,\ldots,x_k,1)=s_i(x_1,\ldots,x_k)+s_{i-1}(x_1,\ldots,x_k)=r_i+r_{i-1}$ for $i\neq k+1$ and $r'_{k+1}=s_{k+1}(x_1,\ldots,x_k,1)=x_1\ldots x_k=s_k(x_1,\ldots,x_k)=r_k$.

(3) We have $h_i(\bigwedge_\alpha X_\alpha) \cong \bigoplus_\alpha h_i(X_\alpha)$
(4) We have $h_0(S^0) = \mathbb{Z}$ while $h_i(S^0) = 0$ for $i > 0$.

This is what we will do, but with an additional wrinkle. The wrinkle arises because the *Dold-Thom theorem only holds for connected complexes* X, whereas the axioms above need a construction which works on all spaces X. Luckily, we can cunningly dodge this difficulty. The idea that makes this possible is that for any based topological space X, the reduced suspension ΣX is always a connected complex. Then, if Dold-Thom holds, we have $\pi_{i+1} \circ SP\Sigma X = H_{i+1}\Sigma X = H_i X$, so if we define $h'_i = \pi_{i+1} \circ SP \circ \Sigma$, then these ought to be homology functors, so ought to satisfy the axioms. Turning this round, we can make one of our subgoals towards proving Dold-Thom proving:

Lemma 3. *The functors $h'_i = \pi_{i+1} \circ SP \circ \Sigma$ satisfy axioms 1-4. Thus, $h'_i \cong H_i$ for all i.*

Before we turn to the proof of lemma 3, let us see how we can use this to finish off the proof of Dold-Thom. The other key ingredient is the following, which is a kind of version of axiom 2, applied to the functors $\pi_* \circ SP$; but restricted so it only talks about connected complexes (and therefore has a chance of being true). This claim is also the natural statement one proves when establishing the 'axiom 2' part of lemma 3, so in some sense it does not 'come out of a hat' as much as it seems. (We shall soon see how this statement is used in the proof of lemma 3.):

Claim 4. *There is a boundary map $\partial : \pi_i \circ SP(X/A) \to \pi_{i-1} \circ SP(A)$, defined for each pair (X, A) where A is a connected subcomplex of X, which is also connected, such that there is an exact sequence*

$$\cdots \to \pi_i \circ SP(A) \to \pi_i \circ SP(X) \to \pi_i \circ SP(X/A) \to \pi_{i-1} \circ SP(A) \to \dots$$

and satisfying appropriate naturality properties.

Now, suppose Y is connected. Then setting $(X, A) = (CY, Y)$ we have $X/A = \Sigma X$, and we certainly satisfy the connectivity properties to apply claim 4; applying it, we have an exact sequence

$$\cdots \to \pi_{i+1} \circ SP(CY) \to \pi_{i+1} \circ SP(\Sigma Y) \to \pi_i \circ SP(Y) \to \pi_i \circ SP(CY) \to \dots$$

Now, for each k,

$$\begin{aligned} \pi_k \circ SP(CY) &= \pi_k \circ SP(pt) && \text{(since } CY \simeq pt\text{, so } SP(CY) \simeq SP(pt)) \\ &= \pi_k(pt) && \text{(since } SP\, pt = pt) \\ &= 0 \end{aligned}$$

so the groups at each end of the portion of exact sequence above are zero, and we have an isomorphism $\pi_i \circ SP(Y) \cong \pi_{i+1} \circ SP(\Sigma Y)$, but then $\pi_{i+1} \circ SP(\Sigma Y) = h'_i(Y) = H_i(Y)$, so $\pi_i \circ SP(X) = H_i(X)$, which establishes Dold-Thom.

It remains to prove lemma 3 using claim 4, and finally to turn to the proof of claim 4.

Proof of lemma 3: For axiom 1, we simply observe that Σ is a homotopy functor (recall we are using the *reduced* suspension); thus $\pi_i \circ SP \circ \Sigma$, being the composition of three functors, is a functor: and then if $f \simeq g$, $\Sigma f \simeq \Sigma g$, so $SP\,\Sigma f \simeq SP\,\Sigma g$, so $\pi_k SP\,\Sigma f = \pi_k SP\,\Sigma f$ for all k (that gives $h'_{i-1} f = h'_{i-1} g$).

For axiom 2, we note that for A a subcomplex of B, ΣA is a connected subcomplex of ΣB, so we can apply claim 4 to $(\Sigma A, \Sigma X)$. We deduce that we have an exact sequence

$$\cdots \to \pi_{i+1} \circ SP(\Sigma A) \to \pi_{i+1} \circ SP(\Sigma X) \to \pi_{i+1} \circ SP(\Sigma(X/A)) \to \pi_i \circ SP(\Sigma A) \to \dots$$

(using the fact that $\Sigma(X/A) = \Sigma X/\Sigma A$). That is

$$\cdots \to h'_i(A) \to h'_i(X) \to h'_i(X/A) \to h'_{i-1}(A) \to \dots$$

which gives the required exact sequence, with $\partial'_k : h'_k(X/A) \to h'_{k-1}(A)$ coming from the map $\partial_{k+1} : \pi_{k+1} \circ SP(\Sigma X/\Sigma A) \to \pi_k \circ SP(\Sigma A)$ we got from the claim. Naturality of that ∂ tells us we have, for every map $g : (X', A') \to (Y', B')$, a diagram:

$$\begin{array}{ccc} \pi_{k+1} \circ SP(X'/A') & \xrightarrow{\partial_{k+1}} & \pi_k \circ SP(A') \\ {\scriptstyle \pi_{k+1}\circ SP(g)}\downarrow & & \downarrow{\scriptstyle \pi_k\circ SP(g)} \\ \pi_{k+1} \circ SP(Y'/B') & \xrightarrow{\partial_{k+1}} & \pi_k \circ SP(B') \end{array}$$

Applying this to the map $\Sigma f : (\Sigma X, \Sigma A) \to (\Sigma Y, \Sigma B)$, where $f : (X, A) \to (Y, B)$ is arbitrary, we get

$$\begin{array}{ccc} \pi_{k+1} \circ SP(\Sigma(X/A)) & \xrightarrow{\partial_{k+1}} & \pi_k \circ SP(\Sigma A) \\ {\scriptstyle \pi_{k+1}\circ SP(\Sigma f)}\downarrow & & \downarrow{\scriptstyle \pi_k\circ SP(\Sigma f)} \\ \pi_{k+1} \circ SP(\Sigma(Y/B)) & \xrightarrow{\partial_{k+1}} & \pi_k \circ SP(\Sigma B) \end{array}$$

That is,

$$\begin{array}{ccc} h'_k(X/A) & \xrightarrow{\partial'_k} & h'_{k-1}(A) \\ {\scriptstyle f_*}\downarrow & & \downarrow{\scriptstyle f_*} \\ h'_k(Y/B) & \xrightarrow{\partial'_k} & h'_{k-1}(B) \end{array}$$

Which establishes the naturality we require.

For axiom 3, we have that that

$$SP \bigwedge_\alpha X_\alpha = \prod_\alpha^\circ SP\, X_\alpha$$

where by $\prod^\circ$ we mean the weak product[3]. To see this, for each finite set of the αs, A say, and each map f of A to the integers, consider the space $X_{A,f} = \prod_{\alpha \in A} SP_{f(\alpha)} X_\alpha$. Now, we may partially order such (A, f) by saying $(A, f) \leqslant (A', f')$ if $A \subseteq A'$ and $f(\alpha) \leqslant f'(\alpha)$ for all α where this makes sense. Clearly this forms a direct system, and if $(A, f) \leqslant (A', f')$ then we can include $X_{A,f}$ in $X_{A',f'}$. So we can consider the direct limit of the $X_{A,f}$, say X. Now consider first taking the limit over possible f, for given A, then taking the limit over A; after a little thought, one sees that one gets $\prod_\alpha^\circ SP\, X_\alpha$. Since taking the limit in two pieces like this does not affect the answer, $\prod_\alpha^\circ SP\, X_\alpha = X$. On the other hand, we can define n as the sum $\sum_{\alpha \in A} f(a)$, and consider taking the limit first over pairs (X, A) with fixed n, then over increasing n; doing this, one gets $SP \bigwedge_\alpha X_\alpha$. Thus $SP \bigwedge_\alpha X_\alpha = \prod_\alpha^\circ SP\, X_\alpha$, both equalling X.

We also have that, for a collection X_α of basepointed spaces, that $\pi_i \prod^\circ X_\alpha = \bigoplus \pi_i X_\alpha$. We can see this as follows.

First, for each of the X_α, pick a neighbourhood U_α of the basepoint which deformation retracts onto the basepoint, and pick a deformation retraction θ_α. Now, we can consider an element $x \in \prod^\circ X_\alpha$ as a tuple (x_α), where each $x_\alpha \in X_\alpha$ and where all but finitely many of the x_α are just the basepoint of their corresponding X_α. Define, for each finite subset S of the α, a subset U_S of $\prod^\circ X_\alpha$ as;

$$U_S = \{x \in \prod^\circ X_\alpha | x_\alpha \in U_\alpha \text{ for all } \alpha \notin S\}$$

[3]The *weak product* is the direct limit of the products of finite subsets of the factors, where for example we include $X_1 \times \ldots \times X_n$ into $X_1 \times \ldots \times X_n \times X_{n+1}$ via $(x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_n, e)$ where e is the basepoint of X_{n+1}.

this is open (since it's intersection with a finite product $X_{\alpha_1} \times \cdots \times X_{\alpha_k}$ is open: in particular, it looks like $T_{\alpha_1} \times \cdots \times T_{\alpha_k}$ where each T_{α_i} is either X_{α_i} or U_{α_i}, and so in either case is open in X_{α_i}). The open neighbourhoods U_S cover X, since X is simply the union of the spaces $\prod_{\alpha\in S} X_\alpha$ where S is a finite set of the α (considered as living in X by setting all opther coordinates equal to the basepoint); and $\prod_{\alpha\in S} X_\alpha \subset U_S$ for all S.

Thus, given an element χ of a π_k, and picking a representative map φ, we have, since φ has compact image, that its image lies in a union $\bigcup U_{S_i}$ for some finite sets S_i; then if we set $S = \bigcup S_i$, we have that the image of φ lies in U_S. We can deformation retract U_S into $\prod_{\alpha\in S} X_\alpha$ by using maps F_t which are always the identity on all the factors corresponding to indices in S, but which carry out the homotopies θ_α on the other factors. Post composing φ with this retraction gives a homotopy of φ into another map φ', which of course still represents χ, but which now maps into $\prod_{\alpha\in S} X_\alpha$.

We have shown that every element of $\pi_i \prod^\circ X_\alpha$ arises from the inclusion of $\pi_i \prod_{\alpha\in S} X_\alpha$ for some S. This is equivalent to saying that the natural map[4]

$$\varinjlim_{S \text{ finite}} \pi_i \prod_{\alpha\in S} X_\alpha \to \pi_i \overset{\circ}{\prod} X_\alpha$$

is surjective. A similar argument to the one above applied to homotopies between maps then shows it is injective. Thus

$$\begin{aligned}
\pi_i \overset{\circ}{\prod} X_\alpha &= \varinjlim_{S \text{ finite}} \pi_i \prod_{\alpha\in S} X_\alpha \\
&= \varinjlim_{S \text{ finite}} \prod_{\alpha\in S} \pi_i X_\alpha \\
&= \varinjlim_{S \text{ finite}} \bigoplus_{\alpha\in S} \pi_i X_\alpha \\
&\qquad \text{(since } S \text{ is finite)} \\
&= \bigoplus_{\alpha} \pi_i X_\alpha
\end{aligned}$$

Then using our results that $SP \bigwedge_\alpha X_\alpha = \prod_\alpha^\circ SP\, X_\alpha$ and $\pi_i \prod^\circ X_\alpha = \bigoplus \pi_i X_\alpha$ we have

$$\begin{aligned}
h_i'(\bigwedge_\alpha X_\alpha) &= \pi_{i+1} \circ SP\Sigma \bigwedge_\alpha X_\alpha \\
&= \pi_{i+1} \circ SP \bigwedge_\alpha \Sigma X_\alpha
\end{aligned}$$

[4] This is the map induced by the compatible collection of maps

$$(\iota_{S*} : \pi_i \prod_{\alpha\in S} X_\alpha \to \pi_i \overset{\circ}{\prod_\alpha} X_\alpha)_{S \text{ finite}}$$

where ι_{S*} is the map on homotopy induced by the natural inclusion

$$\iota_S : \prod_{\alpha\in S} X_\alpha \to \overset{\circ}{\prod_\alpha} X_\alpha$$

coming from the very definition of $\prod^\circ$ as a limit. We can see compatibility of the ι_S easily: if $S \subset S'$, then $\iota_S = \iota_{S'} \circ \iota_{S,S'}$, where $\iota_{S,S'} : \prod_{\alpha\in S} X_\alpha \to \prod_{\alpha\in S'} X_\alpha$ is the inclusion; so $\iota_{S*} = \iota_{S'*} \circ (\iota_{S,S'})_*$

(we can swap Σ and $\bigwedge$ for reduced suspensions)

$$\begin{aligned} &= \pi_{i+1} \prod_{\alpha}^{\circ} SP\,\Sigma X_\alpha \\ &= \bigoplus_{\alpha} \pi_{i+1} SP\,\Sigma X_\alpha \\ &= \bigoplus_{\alpha} h'_i(X_\alpha) \end{aligned}$$

For axiom 4, recall $SP\,S^1 \simeq S^1$, so $\pi_1\,SP\,S^1 = \pi_1\,S^1$; then

$$\begin{aligned} h'_0(S^0) &= \pi_1\,SP\,\Sigma S^0 = \pi_1\,SP\,S^1 \\ &= \pi_1 S^1 = \mathbb{Z} \end{aligned}$$

This establishes all the axioms, and so completes the proof. □

We now turn to a consideration of the key claim 4. We want an exact sequence

$$\cdots \to \pi_i(SP(A)) \to \pi_i(SP(X)) \to \pi_i(SP(X/A)) \to \pi_i(SP(A)) \to \dots$$

Now, we have the usual exact sequence

$$\cdots \to \pi_i(SP(A)) \to \pi_i(SP(X)) \to \pi_i(SP(X), SP(A)) \to \pi_i(SP(A)) \to \dots$$

so comparing, what we want is for the quotient map $q : SP(X) \to SP(X/A)$ to induce an isomorphism $q_* : \pi_i(SP(X), SP(A)) \to \pi_i(SP(X/A))$; that is, we wish for $q_* : \pi_i(SP(X), q^{-1}(b)) \to \pi_i(SP(X/A))$ to be an isomorphism. Establishing this fact is the crux of the Dold-Thom theorem, and turns out to be a little tricky. As a starting point, let us give a name to maps $p : X \to Y$ such that $p_* : \pi_k(X, p^{-1}(b)) \to \pi_k(Y, b)$ is an isomorphism for all b: *quasifibrations*. What we want, therefore, is to show q is a quasifibration. (The name is chosen since the standard theorem one proves about fibrations which gives them their homotopical usefulness is that they have this quasi-fibration property: so 'a quasi-fibration is something that has the useful homotopical property fibrations have'.)

Now, quite a few results in topology are proved by first showing they hold for certain particularly nice subsets, then patching this together to show it holds for larger subsets, and finally using some kind of limiting argument to extend this to the whole space (for example, Poincaré duality is proved in this way). We shall use this kind of approach here. We will therefore need the following criteria for being a quasifibration, (whose proof we defer to the next section) which allow us to patch together the property of being a quasifibration from smaller to larger spaces, and which allow us to pass the property of being a quasifibration over to limits.

Patching criterion for quasifibrations: If $p : X \to Y$ is a continuous map with connected fibers and U, V are sets with $Y = U \cup V$, and assuming $U, V, p^{-1}(U), p^{-1}(V), X, Y$ are connected, then if the three maps

$$p|_{p^{-1}(U)} : p^{-1}(U) \to U, \quad p|_{p^{-1}(V)} : p^{-1}(V) \to V \quad \text{and } p|_{p^{-1}(U\cap V)} : p^{-1}(U \cap V) \to U \cap V$$

are quasifibrations, then so is p.

Limiting criterion for quasifibrations: Suppose $p : X \to Y$ is a continuous map and $Y = \bigcup_{i=0}^{\infty} Y_i$, with the union (direct limit) topology. Suppose further that both X and Y are Hausdorff. Then if $p|_{p^{-1}(Y_i)} : p^{-1}(Y_i) \to Y_i$ is a quasifibration for all i, then p is a quasifibration.

Another technique which is often useful in showing something has a particular property is that the property is resilient to some kind of smooth change. We will need a technical result like this for quasifibrations (whose proof we again defer). The particular kind of smooth change that turns out to be helpful here is a deformation of a subspace into a smaller subspace (this is sometimes called a 'deformation in the weak sense' to distinguish it from a deformation retraction; but we will not use this terminology, since it takes too many words.)

So what is this strange 'deformation' concept? Let $A \subset X \subset Y$; a *deformation of X into A* is a homotopy F_t of maps $Y \to Y$ which starts with the identity on Y, and for which F_1 maps X completely into A. We must also have that F_t, for all t, maps A into itself and X into itself. (Note that we do *not* require that F_t fixes every point in A, as we would require for $F_t|_X$ to be a deformation retraction.) Then we have:

Homotopy criterion for quasifibrations: If $p : X \to Y$ is a continuous map with path-connected fibers and we are given a deformation of X into a subspace X' (F_t say), covering a deformation G_t of Y into a subspace Y' (with X, Y, X', Y' all path connected), and such that $p|_{X'} : X' \to Y'$ is a quasifibration, and (finally) such that $F_1 : p^{-1}(b) \to p^{-1}(G_1(b))$ is a weak homotopy equivalence for all b, then p is a quasifibration.

(By saying that F_t covers G_t, we just mean that $p \circ F_t = G_t \circ p$, for all t.)

Armed with these criteria, we can go forward to prove the key claim by showing that $q : SP(X) \to SP(X/A)$ (as defined above) is a quasifibration. We will do this by showing (inductively) that for each n, $p|_{p^{-1}(B_i)}$ is a quasifibration, where $B_i = SP_i(X/A)$. (Since $SP(X/A)$ is the union of the B_i with the union topology, we are then done by the limiting criterion.) For $n = 0$, B_0 is a point, and the result is trivial. So let us turn to the inductive case, $n > 0$. We shall find an open neighbourhood U of B_{n-1} over which q is a quasifibration (by which we mean that $q|_{q^{-1}(U)}$ is a quasifibration). We shall also show that q is a quasifibration over $V = B_n \backslash B_{n-1}$ (which is clearly open) and $U \cup V$, and then we'll be done by the patching criterion. Before we wade into doing this, however, we will first note that we might as well replace X with the mapping cylinder of the inclusion of A into X (which is homotopy equivalent to X). This will turn out to make a few things slightly easier in the rest of the proof.

Now, we begin by picking a neighborhood W of A in X with a deformation of W into A. Since X is a mapping cylinder, this can be chosen to just slide points along the mapping cylinder, compressing the part of the mapping cylinder near A to the A end, while stretching the bit near the X end. This then has the additional property that the deformation actually fixes all points in A at all times t.

We then define U as the set of points in $B_n = SP_n(X/A)$ which have at least one factor in W/A; this is clearly a neighbourhood of e. Now, let f_t be our deformation of W into A; this clearly gives a deformation $\bar{f}_t$ of W/A onto $b = A/A$ (just postcompose with the quotient map $W \to W/A$; the resulting composite clearly factors through the quotient $W \to W/A$). Now, define a map

$$F_t : SP(X) \to SP(X) : x_1 + \cdots + x_k \mapsto f_t(x_1) + \cdots + f_t(x_k)$$

(this is well defined, since f_t fixes the basepoint e, as it's in A, so it doesn't matter if we add extra es in the expression in the LHS). We can similarly define a map $\bar{F}_t$. Clearly F_t covers $\bar{F}_t$; and $\bar{F}_t$ restricts to give a deformation of U into B_{n-1}. Similarly, F_t restricts to give a retraction of $p^{-1}(U)$ into $p^{-1}(B_n)$. We inductively know q is a quasifibration over B_{n-1}; so we'll be done by the homotopy condition for quasifibrations if we can show F_1 induces a weak homotopy equivalence $p^{-1}(b) \to p^{-1}(\bar{F}(b))$ for all b. So fix a $b = \bar{x}_1 + \cdots + \bar{x}_n$, where the $\bar{x}_i$ are points in X/A. Now, we may as well assume none of the $\bar{x}_i$ is the basepoint A/A; which means each has a unique lift to a point x_i of $X \backslash A$. Then it is clear that $p^{-1}(b)$ is homeomorphic to $SP(A)$ via

$$\psi : SP(A) \to p^{-1}(b) : P \mapsto P + x_1 + \cdots + x_n$$

Now, when we apply f_1 to the x_i, some will be mapped into A, while others will remain in $X \backslash A$. We may wlog assume that we have reordered such that $f_1(x_1), \ldots, f_1(x_k) \in X \backslash A$ while $f_1(x_{k+1}), \ldots, f_1(x_n) \in A$. Then it is clear that $p^{-1}(\bar{f}(b))$ is homeomorphic to $SP(A)$ via

$$\varphi : SP(A) \to p^{-1}(\bar{f}(b)) : P \mapsto P + f_1(x_1) + \cdots + f_1(x_k)$$

Then we have, for $P \in SP(A)$

$$\begin{aligned} F_1(\psi(P)) &= F_1(P + x_1 + \cdots + x_n) \\ &= F_1(P) + f_1(x_1) + \dots f_1(x_n) \\ &= \varphi(F_1(P)) + f_1(x_{k+1}) + \cdots + f_1(x_n) \\ &= \varphi(F_1(P)) + Q \qquad \text{(say, where } Q = f_1(x_{k+1}) + \cdots + f_1(x_n)) \end{aligned}$$

Now, since A is path connected, so is $SP(A)$, and we can find a path Q_t from Q to the basepoint. Then $F_1 \circ \psi : P \mapsto \varphi(F_1(P)) + Q$ is homotopic via $\mu_t : P \mapsto \varphi(F_{1-t}(P)) + Q_t$ to $P \mapsto \varphi(F_0(P)) = \varphi(P)$. We deduce F_1 satisfies

$$F_1|_{p^{-1}(b)} = F_1 \circ \psi \circ \psi^{-1} \simeq \varphi \circ \psi^{-1}$$

and so, being homotopic to a homeomorphism, is a homotopy equivalence (since φ and ψ are both homeomorphisms, so is $\varphi \circ \psi^{-1}$). This is as we require.

Now, let us turn to the space $V = B_n \backslash B_{n-1}$. It is clear that $V = SP_n((X/A)\backslash\{e\})$. Now, there is homeomorphism $(X/A)\backslash\{e\} \to X\backslash A$, which gives a homeomorphism $SP_n((X/A)\backslash\{e\}) \cong SP_n(X\backslash A)$. Let $\theta : V \to SP_n(X\backslash A)$ be this homeomorphism. It is clear that $SP_n(X\backslash A)$ includes into $SP(X)$; thus we can think of θ mapping to $SP(X)$. It is then clear that $p \circ \theta$ is the identity on V (θ gives a section for p), and that each $P \in p^{-1}(V)$ may be written as $P = \theta(p(P)) + Q$ for some Q: a little thought shows $Q \in SP(A)$. Now, it is an easy excercise that the 'subtraction' map on SP is continuous on compact sets where defined, so if we define $Q : p^{-1}(V) \to SP(A)$ so $P = \theta(p(P)) + Q(P)$, then Q is continuous on compact sets. Clearly $p \circ Q$ is identically equal to e. Putting this together, and pretending Q were continuous for a moment, we'd have a homeomorphsm:

$$\begin{aligned} p^{-1}(V) &\cong V \times SP(A) \\ P &\to (p(P), Q(P)) \\ \theta(v) + Q &\leftarrow (v, Q) \end{aligned}$$

Then considering p as a map on $V \times SP(A)$ rather than $p^{-1}(V)$, it would be simply projection onto the second factor. Thus it would be a fiber bundle (indeed, a trivial bundle), and so a fibration, and so a quasifibration. But seeing as we're interested in homotopy, and maps representing homotopy elements always have compact image, the fact that Q is only continuous on compact subsets does not prevent the above argument going through.

(The same argument gives a map $p^{-1}(V \cap U) \to (V \cap U) \times SP(A)$ which is a homeomorphism on compact sets: so we get that $p|_{p^{-1}(U\cap V)}$ is also a quasifibration.)

4. Establishing the criteria for quasifibrations

We'll first look at the patching criterion, and we'll find that we need a technical tool in the argument; the proof of this tool will be deferred to the next section (and that's the *last* time a proof is deferred in this essay, I promise!) We'll call it 'homotopical Mayer-Vietoris', since it allows us to deduce facts about what a map does on the homotopy groups of a whole space from what it does on two subsets the union of whose interiors is the whole space.

Technical tool: Suppose $f : (X; A, B) \to (Y; C, D)$ is a continuous map (that is a map $X \to Y$ sending A to C and B to D), where X is the union of the interiors of A and B and similarly for Y, C and D. Suppose further that the induced maps $\pi_i(A, A \cap B) \to \pi_i(C, C \cap D)$ and $\pi_i(B, A \cap B) \to \pi_i(D, C \cap D)$ are surjections for $i < k$ and bijections for $i = k$, then the same can be said for the induced maps $\pi_i(X, A) \to \pi_i(Y, C)$ (and, symmetrically, for the induced maps $\pi_i(X, B) \to \pi_i(Y, D)$ also).

Theorem 5. *The patching criterion for being a quasifibration holds. That is, if $p : X \to Y$ is a continuous map with path-connected fibers and U, V are sets with $Y = U \cup V$, and assuming $U, V, p^{-1}(U), p^{-1}(V), X, Y$ are connected, then if the three maps*

$$p|_{p^{-1}(U)} : p^{-1}(U) \to U, \quad p|_{p^{-1}(V)} : p^{-1}(V) \to V \quad \text{and } p|_{p^{-1}(U\cap V)} : p^{-1}(U \cap V) \to U \cap V$$

are quasifibrations, then so is p.

Proof: For convenience, write $\bar{U}$ for $p^{-1}(U)$ and similarly $\bar{V}$ for $p^{-1}(V)$. Consider the map induced by p from the long exact sequence of the triple $(\bar{U}, \bar{U} \cap \bar{V}, p^{-1}(b))$ to that of $(U, U \cap V, b)$. The maps $\pi_k(\bar{U}, p^{-1}(b)) \to \pi_k(U, b)$ are isomorphisms by assumption that p is a quasifibration over U; similarly, since it is a quasifibration over $U \cap V$, the maps $\pi_k(\bar{U} \cap \bar{V}, p^{-1}(b)) \to \pi_k(U \cap V, b)$ are isomorphisms. Thus, by the five lemma[5], we deduce the maps $\pi_k(\bar{U}, \bar{U} \cap \bar{V}) \to \pi_k(U, U \cap V)$ are isomorphisms. A similar argument shows the same holds for the maps $\pi_k(\bar{V}, \bar{U} \cap \bar{V}) \to \pi_k(V, U \cap V)$.

Then the technical tool tells us the maps $\pi_k(X, \bar{U}) \to \pi_k(Y, U)$ are isomorphisms. And we know the maps $\pi_k(\bar{U}, p^{-1}(b)) \to \pi_k(U, b)$ are isomorphisms. Then considering the map induced by p from the long exact sequence of the triple $(X, \bar{U}, p^{-1}(b))$ to that of (Y, U, b), the five lemma tells us that the maps $\pi_k(X, p^{-1}(b)) \to \pi_k(Y, b)$ are isomorphisms, which is as we require. $\square$

Before we turn to the limiting criterion, it will be useful to prove a lemma.

Lemma 6. *Suppose Y is a Hausdorff space such that $Y = \bigcup_{i=0}^{\infty} Y_i$, with the union (direct limit) topology, for subspaces Y_i. Then a compact subset K of Y actually lies in Y_i for some i.*

Proof: For each i, pick an element $y_i \in Y$ such that $y_i \notin Y_i$. For each k, we claim the set $S_k = \{y_k, y_{k+1}, \dots\}$ is closed. For this, it suffices to prove that for each n, the intersection $S_k \cap Y_n$ is closed. But $S_k \cap Y_n = \{y_k, \dots, y_{n-1}\}$ (or is empty if $k \geqslant n$, so certainly closed); this is a finite set, so (as we're in a Hausdorff space) certainly closed. Thus our claim is proved.

Now, the $S_k \subset K$, and every finite intersection $\bigcap S_{k_i}$ of the S_k is nonempty (being S_N for $N = \max k_i$); but the intersection of all the S_k is empty (since $S_k \subset Y \backslash Y_k$, so $\bigcap S_k \subset Y \backslash \bigcup Y_k = Y \backslash Y = \{\}$). This contradicts the compactness of K. $\square$

Theorem 7. *The limiting criterion for being a quasifibration holds: suppose $p : X \to Y$ is a continuous map and $Y = \bigcup_{i=0}^{\infty} Y_i$, with the union (direct limit) topology. Suppose further both X and Y are Hausdorff. Then if $p|_{p^{-1}(Y_i)} : p^{-1}(Y_i) \to Y_i$ is a quasifibration for all i, then p is a quasifibration.*

Proof: Fix a basepoint b in Y. Suppose χ is an element of π_k, and φ is a representative map. Then, since the image of φ is compact, it lies in some Y_i; we may write $\varphi = \iota_i \circ \varphi'$, where $\iota_i : Y_i \to Y$ is inclusion and φ' is the map φ considered as a map to Y_i. Then χ is $= \iota_{i*}([\varphi'])$ and so lies in the image of ι_{i*}. This tells us that the natural map $\lim_{\to} \pi_k(Y_i) \to \pi_k(Y)$ is surjective (compare the establishment of axiom 3 in the proof of lemma 3); then (as usual) a similar argument applied to homotopies gives that the map is injective: and we conclude $\lim_{\to} \pi_k(Y_i) \cong \pi_k(Y)$. A similar argument gives that $\lim_{\to} \pi_k(p^{-1}(Y_i), p^{-1}(b)) \cong \pi_k(X, p^{-1}(b))$ (here we only take the limit where it makes sense: i.e. over those i large enough that $b \in Y_i$). (We begin this 'similar argument' by saying, for a map φ to X representing a homotopy element, that since the image of $p \circ \varphi$ is compact, it lies in some Y_i; so the image of φ lies in $p^{-1}(Y_i)$.)

Now let us turn to the map $p_* : \pi_k(X, p^{-1}(b)) \to \pi_k(Y, b)$. We have

$$\begin{aligned} \pi_k(X, p^{-1}(b)) &\cong \lim_{\to} \pi_k(p^{-1}(Y_i), p^{-1}(b)) \\ &\cong \lim_{\to} \pi_k(Y_i) \end{aligned}$$

[5]We note that the bit at the end of the long exact sequence, where everything stops being a group, has everything zero anyway, since everything in sight is path connected. It is then an easy exercise that we can deduce the relative π_1 is zero, even though the usual (group) 5 lemma does not apply.

(since by assumption $\pi_k(p^{-1}(Y_i), p^{-1}(b) \cong \pi_k(Y_i)$ via the map induced by $p|_{p^{-1}(Y_i)}$; these isomorphisms are compatible as they all restrict p_*)

$$\cong \pi_k(Y, b)$$

The isomorphism is clearly induced by p_* (since on each $\pi_k(X, p^{-1}(b))$ it restricts to give $p_*|_{\pi_k(p^{-1}(Y_i),p^{-1}(b))}$); this is as required. $\square$

Theorem 8. *The homotopy criterion for being a quasifibration holds: if $p : X \to Y$ is a continuous map with path-connected fibers and we are given a deformation of X into a subspace X' (F_t say), covering a deformation G_t of Y into a subspace Y' (with X, Y, X', Y' all path connected), and such that $p|_{X'} : X' \to Y'$ is a quasifibration, and (finally) such that $F_1 : p^{-1}(b) \to p^{-1}(G_1(b))$ is a weak homotopy equivalence for all b, then p is a quasifibration.*

Proof: Fix a point b in Y. Since F_1 covers G_1, we have $p \circ F_1 = G_1 \circ p$; this induces the following commutative diagram on homotopy groups:

$$\begin{tikzcd}
\pi_k(X, p^{-1}(b)) \arrow[r, "p_*"] \arrow[d, "F_{1*}"'] & \pi_k(Y, b) \arrow[d, "G_{1*}"'] \\
\pi_k(X', p^{-1}(G_1(b))) \arrow[r, "p_*"] & \pi_k(Y', G_1(b))
\end{tikzcd}$$

Our aim is to prove the map across the top is an isomorphism for all k, so it suffices to prove the other three maps are isomorphisms for all k. The fact that $p|_{X'}$ is a quasifibration immediately gives this for the map across the bottom.

To show the left vertical map is an isomorphism, it suffices (by the long exact sequence for relative homotopy, and the five lemma[6]) to show that $F_{1*} : \pi_k(X) \to \pi_k(X')$ and $F_{1*} : \pi_k(p^{-1}(b)) \to \pi_k(p^{-1}(G_1(b)))$ are isomorphisms. But the former is true since F_1 is homotopic via F_t to the identity, and the latter is true by the hypothesis that $F_1 : p^{-1}(b) \to p^{-1}(G_1(b))$ is a weak homotopy equivalence for all b.

Similarly, to show that the right map is an isomorphism, we just need $G_{1*} : \pi_k(Y) \to \pi_k(Y')$ is an isomorphism (which it is since G_1 is homotopic to the identity), and $G_{1*} : \pi_k(\{b\}) \to \pi_k(\{G_1(b)\})$ is an isomorphism (which it is since both the groups $\pi_k(\{b\})$ and $\pi_k(\{G_1(b)\})$ are trivial, being homotopy groups of a point. $\square$

5. Establishing 'homotopical Mayer-Vietoris'

This section is one of the most directly geometric of the entire essay. We will working explicitly with maps representing homotopy group elements and homotopies between them. To save a little work, we will prove the following lemma, which for an inclusion $(X, A) \to (Y, C)$ gives a single geometric condition which is equivalent both to the injectivity of the map on π_{k-1} and the surjectivity of the map of π_k. (It should not be altogether surprising that these both reduce to a single condition, seeing as the injectivity involves the existence of a homotopy of a given map of S^{k-1} to zero—a k dimensional object—while the surjectivity involves the construction of a map of S^k, which is again a k dimensional object.)

[6]The same remarks as made during the proof of the patching criterion, about the non-group bits of the sequence vanishing, hold here.

Lemma 9. *Suppose we have an inclusion $(X, A) \to (Y, C)$; then the following conditions are equivalent:*

(1) *The induced map on π_{k-1} is injective and the induced map on π_k is surjective.*
(2) *Let us write the surface ∂D^k of the k dimensional disk as the union of hemispheres $\partial_+ D^k$ and $\partial_- D^k$ meeting in S^{k-2}. Suppose we are given a map of $D^k \times \{0\} \cup \partial_+ D^k \times I$ to Y which sends*
   - *$\partial_- D^k \times \{0\}$ and $S^{k-2} \times I$ to A*
   - *$\partial_+ D^k \times \{1\}$ to X*
   - *$S^{k-2} \times \{1\}$ to C*

   Then it extends to a map $D^k \times I \to Y$ sending
   - *$\partial_- D^k \times I$ to A*
   - *$D^k \times \{1\}$ to X*
   - *$\partial_- D^k \times \{1\}$ to C*
(3) *Condition (2) weakened by only requiring that we can extend maps which are independent of the I coordinate on the $\partial_+ D^k \times I$ part.*

Proof: Clearly (3) $\Rightarrow$ (2). Let us turn to the implication (3) $\Rightarrow$ (1). So suppose (3) holds. Let us first show that we have injectivity on π_{k-1}. Given an element of the kernel of $\iota_* : \pi_{k-1}(X, A) \to \pi_{k-1}(Y, C)$, we can represent it by a map $\psi : (D^{k-1}, \partial D^{k-1}) \to (X, A)$, or equivalently a map $(\partial_+ D^k \times \{1\}, S^{k-2} \times \{1\}) \to (X, A)$. We can extend this to $\partial_+ D^k \times I$ using the requirement that the map be independent of the I coordinate. We can then extend across $D^k \times \{0\}$ using the the fact that ψ is homotopic to a constant as a map to (Y, C). Then using (3), we can extend to a map on all of $D^k \times I$. Restricting this do $D^k \times \{1\}$ shows ψ is homotopic to a constant map as a map to (X, A) too. Hence the kernel is trivial.

Now let us show surjectivity on π_k. Given an element of $\pi_k(Y, C)$, we can represent it by an map $D^k \times \{0\}$ to Y sending $\partial_+ D^k \times \{0\}$ to the basepoint and $\partial_- D^k \times \{0\}$ to C. Again we extend to $\partial_+ D^k \times I$ by the requirement that the map be independent of the I coordinate. Then we can use (3) to extend to a map on $D^k \times I$. The restriction of this to $D^k \times \{1\}$ gives an element of $\pi_k(X, A)$; then the whole map on $D^k \times I$ gives a homotopy which shows that once we map into $\pi_k(Y, C)$, this is equal to our original element.

We now turn to (1) $\Rightarrow$ (2). Given a map f as in (2), we may consider the map on $(\partial_+ D^k \times \{1\}, S^{k-2} \times \{1\})$ as representing an element θ of $\pi_{k-1}(X, A)$. Once we map this into $\pi_{k-1}(Y, C)$, this is conjugate to the element represented by the map on $(\partial_+ D^k \times \{0\}, S^{k-2} \times \{0\})$ (the map on $\partial_+ D^k \times I$ gives the homotopy to show this), which is zero (by consideration of the map f restricted to $D^k \times \{0\}$. Thus by the injectivity of (1), θ is zero, which means we can extend f across $D^k \times \{1\}$, to get f_1 (say). Now, choose a small k-disk E^n in $\partial_- D^k \times I$, which intersects $D^n \times \{1\}$ in a hemisphere $\partial_+ E^n$ of its boundary. By a small deformation of f_1, we may assume that $\partial_+ E^n$ is sent identically to some fixed point (call the modified map f_2). The surjectivity part of (1) means that we can extend our map f_2 to cover all of $D^k \times I$, such that everything is sent to the places as required by (2), except that E_n gets sent to X and $\partial_- E^n$ gets sent to A. We then do a small deformation to 'pivot' E_n into $D^k \times \{1\}$, and we're done. □

Theorem 10. *Suppose $f : (X; A, B) \to (Y; C, D)$ is a continuous map (that is a map $X \to Y$ sending A to C and B to D). Suppose further that the induced maps $\pi_i(A, A \cap B) \to \pi_i(C, C \cap D)$ and $\pi_i(B, A \cap B) \to \pi_i(D, C \cap D)$ are surjections for $i < k$ and bijections for $i = k$, then the same can be said for the induced maps $\pi_i(X, A) \to \pi_i(Y, C)$ (and, symmetrically, for the induced maps $\pi_i(X, B) \to \pi_i(Y, D)$ also).*

Proof: Our first goal is to replace our arbitrary map f with an inclusion using mapping cylinders. We can replace Y with the mapping cylinder of $f : X \to Y$; we then replace C with $M_{f|A} \cup (f^{-1}(C) \times (1/2, 1])$ where by $Xf^{-1}(C) \times (1/2, 1]$ we really mean its image in M_f. Similarly replace D with $M_{f|B} \cup (f^{-1}(D) \times (1/2, 1])$. (The reason we add the $(f^{-1}(C) \times (1/2, 1])$ bits, which do not change the homotopy types, is to ensure the new Y is still the union of the interiors of C and D.) Then it is clear that it suffices to prove the result for the inclusion of (X, A, B) into this new complex.

In view of the lemma, it suffices to assume that property (2) holds for the inclusions $(A, A \cap B) \to (C, C \cap D)$ and $(B, A \cap B) \to (D, C \cap D)$ and prove that property (3) holds for the inclusion $(X, A) \to (Y, C)$. The argument upon which we are about to embark will make much use of subdivision of disks into smaller disks. This is a little easier, from the standpoint of bookkeeping, if we use cubes I^n rather than D^n; so let us choose an identification of D^n with I^n which sends one face of the cube to $\partial_- D^k$ and sends the disc formed by all the remaining faces to to $\partial_+ D^k$.

Now, suppose we are given our function f defined on $I^k \times \{0\} \cup \partial_+ I^k \times I$. We first split I^k into small cubes, each small enough so that each of the small cubes is either sent completely into C or D. In constructing our extension, we will preserve the following key property:

- For K a small cube or indeed any face of such a cube, then if $K \times \{0\}$ is sent to C or D, then $(K \times I, K \times \{1\})$ is sent to (C, A) or (B, D) respectively.

We may assume this condition holds initially if we suppose $A = X \cup C$ and $B = C \cup D$, which certainly holds for the mapping cylinder construction we performed above.

We then slowly extend our map to the small cubes and their lower dimensional faces (working up through the dimensions). Let K be such a subcube (or lower dimensional face), and suppose we have already extended f to $\partial_+ K \times I$. If f maps $\partial_- K \times \{0\}$ to $C \cap D$, then using property (2) for either the inclusion $(A, A \cap B) \to (C, C \cap D)$ (if $K \times \{0\}$ is sent to C) or the inclusion $(B, A \cap B) \to (D, C \cap D)$ (if $K \times \{0\}$ is sent to D), we can make the extension required, preserving the key property. Otherwise, the given f takes $(K \times \{0\}, \partial_- K \times \{0\})$ to either (C, C) or (D, D), and it is easy to construct the required extension with our bare hands, by simply picking a retraction of $K \times I$ to $\partial_+ K \times I \cup K \times \{0\}$ which sends $K \times \{1\}$ to $\partial_+ K \times \{1\}$, then simply defining the new f to be the old f postcomposed with this retraction. $\square$

References

[1] M Aguilar, S Gitler, C Prieto, *Algebraic Topology from a Homotopical Viewpoint,* New York: Springer-Verlag, 2002

[2] A Dold, R Thom, *Quasifaserungen und unendliche symmetrische Produkte,* Ann. of Math. 67 (1958), 239–281

[3] A Hatcher, *Algebraic Topology,* Cambridge: Cambridge University Press, 2002