Theorem 6.1. (Regular neighbourhoods of submanifolds) *Let L be a submanifold of pl manifold M. Then $\mathcal{N}(L)$ is the total space of a fibre bundle over L, with fibre a disc D^n, and with the inclusion $L \to \mathcal{N}(L)$ being a section.*

In this course, we will only consider very simple bundles. We therefore only give the briefest outline of their theory. Normally bundles are dealt with in the smooth category, but there is of course a pl version. This is less satisfactory in high dimensions, but in dimension three, it works well.

Definition. A map $p: B \to M$ is a *fibre bundle over M with total space B and fibre F* (or an *F-bundle*) if M has an open cover $\{U_\alpha\}$ such that

- the closure $\overline{U}_\alpha$ of each U_α is simplicial, and
- each $p^{-1}(\overline{U}_\alpha)$ is (pl) homeomorphic to $F \times \overline{U}_\alpha$ so that the following diagram commutes:

$$\begin{array}{ccc} p^{-1}(\overline{U}_\alpha) & \xrightarrow{\cong} & F \times \overline{U}_\alpha \\ \big\downarrow{\scriptstyle p} & & \big\downarrow{\scriptstyle \text{projection onto 2nd factor}} \\ \overline{U}_\alpha & = & \overline{U}_\alpha \end{array}$$

If U_α and U_β intersect, then there are two maps

$$p^{-1}(\overline{U}_\alpha \cap \overline{U}_\beta) \to F \times (\overline{U}_\alpha \cap \overline{U}_\beta),$$

one given via U_α, one via U_β. Hence, we obtain a map $g_{\beta\alpha}: F \times (\overline{U}_\alpha \cap \overline{U}_\beta) \to F \times (\overline{U}_\alpha \cap \overline{U}_\beta)$, such $g_{\beta\alpha}|_{F\times\{x\}}$ is a homeomorphism onto $F \times \{x\}$ for each $x \in \overline{U}_\alpha \cap \overline{U}_\beta$. These maps $g_{\beta\alpha}$ are known as the *transition maps*, and satisfy the following conditions:

1. $g_{\alpha\alpha} = \text{id}$,
2. $g_{\beta\alpha} = g_{\alpha\beta}^{-1}$,
3. $g_{\gamma\beta} \circ g_{\beta\alpha} = g_{\gamma\alpha}$.

Usually, one insists that, for each α and β and each $x \in \overline{U}_\alpha \cap \overline{U}_\beta$, $g_{\alpha\beta}|_{F\times\{x\}}$ should lie in some specified subgroup of $\text{Homeo}(F, F)$, known as the *structure group* of the bundle. In this case, all we insist is that these homeomorphisms be pl.

Note that a fibre bundle over M with fibre F can be specified by an open cover $\{U_\alpha\}$ of M (with each $\overline{U}_\alpha$ simplicial), together with transition maps satisfying the above three conditions.

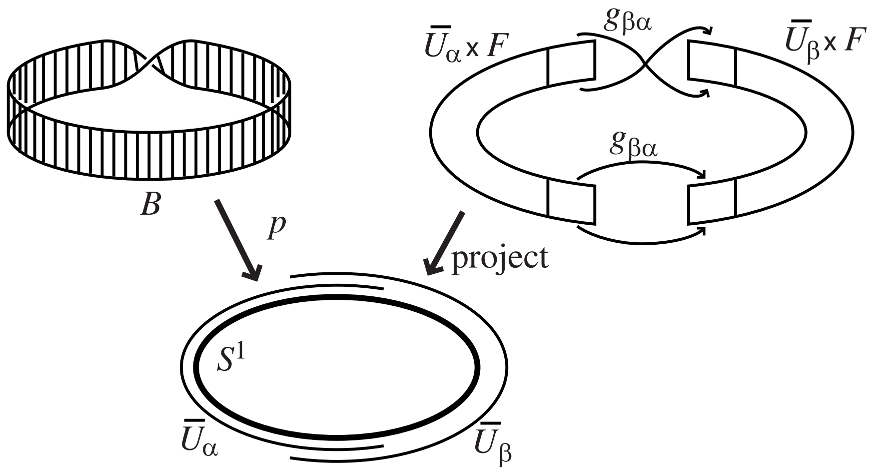

Figure 21.

Definition. A *section* of a fibre bundle $p: B \to M$ is a map $s: M \to B$ such that $p \circ s = \mathrm{id}_M$.

Sketch proof of Theorem 6.1. Pick a triangulation of M in which L is simplicial. This induces handle structures on L and M. Each i-handle of L is contained in an i-handle of M. The union of these handles of M containing L forms $\mathcal{N}(L)$. Careful choice of product structures on the handles (starting with the highest index handles and working downwards) can be used to define the bundle map $p: \mathcal{N}(L) \to L$. Each U_α is (a small extension) of a handle of L. $\square$

Definition. Two bundles $p_1: B_1 \to M$ and $p_2: B_2 \to M$ are *equivalent* if there is a homeomorphism $h: B_1 \to B_2$ so that the following commutes:

$$\begin{array}{ccc} B_1 & \xrightarrow{h} & B_2 \\ \Big\downarrow{\scriptstyle p_1} & & \Big\downarrow{\scriptstyle p_2} \\ M & = & M \end{array}$$

Definition. If $p: B \to M$ is a fibre bundle and $f: M' \to M$ is any map, then there is a bundle over M', known as the *pull-back bundle*. It is constructed by taking the open cover $\{U_\alpha\}$ via which M is defined, and letting $\{f^{-1}(U_\alpha)\}$ be the open cover for M'. If $g_{\alpha\beta}: F \times (\overline{U}_\alpha \cap \overline{U}_\beta) \to F \times (\overline{U}_\alpha \cap \overline{U}_\beta)$ is a transition map then the transition map at a point x of $f^{-1}(\overline{U}_\alpha) \cap f^{-1}(\overline{U}_\beta)$ is given by $g_{\alpha\beta}|_{F \times \{f(x)\}}$.

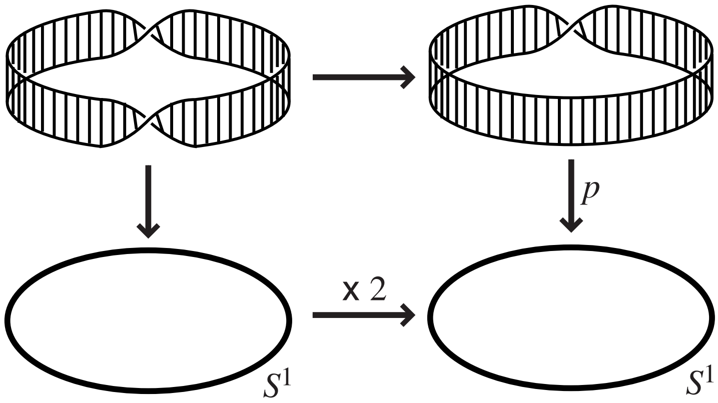

Figure 22.

Examples. Let B be any bundle over M. If $i: M' \to M$ is an inclusion map, then the pull-back bundle is the restriction of the bundle to M'. The pull-back of B with respect to id_M is the same bundle as B. The pull-back with respect to a constant map $M' \to M$ is a product bundle.

The following important result is not very difficult. Its proof can be found in Husemoller's book 'Fibre Bundles'.

Theorem 6.2. *Let M be compact, and let $p: B \to M \times [0,1]$ be a fibre-bundle. Then the associated bundles over $M \times \{0\}$ and $M \times \{1\}$ are equivalent.*

Corollary 6.3. *A bundle over a contractible space M is a product bundle.*

Proof. Since M is contractible, there is a homotopy $M \times [0,1] \to M$ between id_M and a constant map. Pull back the bundle over M to a bundle over $M \times [0,1]$. The bundle over $M \times \{0\}$ is the original bundle. The bundle over $M \times \{1\}$ is the product bundle. They are equivalent by Theorem 6.2. $\square$.

Lemma 6.4. *For each $n \in \mathbb{N}$, there are precisely two D^n-bundles over S^1 up to bundle equivalence.*

Proof. The two D^n-bundles over S^1 are constructed as follows. Start with the product bundle $D^n \times [0,1]$ over $[0,1]$, and glue $D^n \times \{0\}$ to $D^n \times \{1\}$ via some homeomorphism. The result is a D^n-bundle over S^1. It is easy to see that isotopic gluing homeomorphisms give equivalent bundles. By Proposition 4.6, there are two isotopy classes of such homeomorphisms. To see that the bundles are inequivalent, note that their underlying spaces are not homeomorphic: one is orientable and one is not.

Now we must show that every D^n-bundle over S^1 is equivalent to one of these. Pick a point $x \in S^1$. Then, restricting to the bundle over $S^1 - \mathrm{int}(\mathcal{N}(x))$ is a bundle over the interval, which by Corollary 6.3 is a product. Hence, our bundle is constructed as above. □

We now give a characterisation of whether a manifold is orientable.

Proposition 6.5. *An n-manifold M is orientable if and only if it contains no embedded copy of the total space of the non-orientable D^{n-1}-bundle over S^1.*

Proof. If such a bundle embeds in M, then some triangulation of M is non-orientable, and hence M is non-orientable.

Conversely, suppose that M contains no such bundle. Pick an orientation on some n-simplex of M. This specifies unique compatible orientations on its neighbouring n-simplices. Repeat with these simplices. In this way, we orient M, unless at some stage we return to an n-simplex and assign it an orientation the opposite from its original orientation. This specifies a loop, running between the n-simplices through the $(n-1)$-dimensional faces. We may take this loop ℓ to be embedded. Then $\mathcal{N}(\ell)$ is the required non-orientable D^{n-1}-bundle over S^1. □

The total space of the non-orientable D^{n-1}-bundle over S^1 is the Möbius band for $n = 2$ and the solid Klein bottle for $n = 3$.

Proposition 6.6. *Let S be a surface properly embedded in a compact orientable 3-manifold M. Then S is orientable if and only if $\mathcal{N}(S)$ is homeomorphic to $S \times I$.*

Proof. It suffices to consider the case where S is connected. Let $p\colon \mathcal{N}(S) \to S$ be the I-bundle over S from Theorem 6.1. Suppose first that S has non-empty boundary. Then there is a collection A of disjoint properly embedded arcs in S, such that cutting S along A gives a disc D. Then, by Corollary 6.3, the restriction of p to $p^{-1}(D)$ is a product I-bundle. Now identify arcs in ∂D in pairs to give S. These arcs inherit an orientation from some orientation on ∂D.

If two arcs α_1 and α_2 are glued so that their orientations agree, then S contains an embedded Möbius band and so is non-orientable. When $\alpha_1 \times I$ is glued to $\alpha_2 \times I$, the orientations of the I factors must be reversed (otherwise M would contain a solid Klein bottle and hence be non-orientable). Hence, the ∂I-bundle over S is connected, and therefore $\mathcal{N}(S)$ is not homeomorphic to $S \times I$.

Suppose therefore that each pair of arcs α_1 and α_2 in ∂D are identified in a way that reverses orientation. Then S is orientable. Also, the gluing map between $\alpha_1 \times I$ and $\alpha_2 \times I$ preserves the orientation of the I-factor, otherwise M would contain a solid Klein bottle. Hence, after an isotopy of the gluing maps, we may assume that it is the identity in the I-factors. Hence, $\mathcal{N}(S)$ is a product I-bundle.

Now consider the case where S is closed. Remove the interior of a small disc D to give a surface S'. Then S is orientable if and only if S' is. If $\mathcal{N}(S)$ is product I-bundle, then its restriction to S' is. Conversely, if its restriction to S' is a product I-bundle, then we may extend the product structure over $p^{-1}(D)$ to give a product structure on $\mathcal{N}(S)$. $\square$

A codimension one submanifold X of a manifold is known as *two-sided* if $\mathcal{N}(X)$ is a product I-bundle. The existence of a product neighbourhood for a properly embedded orientable surface S in an orientable 3-manifold M is very important. For example, it is vital in the proof of Theorem 3.3, which asserts that S is incompressible if and only if it is π_1-injective. This can in fact fail for non-orientable surfaces. For example, there is a non-orientable incompressible embedded surface in some lens space which is not π_1-injective.

§7. Homology of 3-manifolds

Definition. For $i \in \mathbb{Z}_{\geq 0}$, the i^{th} *Betti number* $\beta_i(M)$ of a space M is the dimension of $H_i(M; \mathbb{Q})$ viewed as a vector space over $\mathbb{Q}$.

Definition. The *Euler characteristic* $\chi(M)$ of a compact triangulable space M is

$$\sum_i (-1)^i \beta_i(M).$$

Theorem 7.1. *Pick any triangulation of a compact space M, and let σ_i be the number of i-simplices in this triangulation. Then $\chi(M) = \sum_i (-1)^i \sigma_i$.*

Remark. If $H_i(M) \cong \mathbb{Z}^a \oplus T$, where each element of T has finite order, then $\beta_i(M) = a$.

The following result, which we quote without proof, is one of the cornerstones of manifold theory.

Theorem 7.2. (Poincaré duality) *Let M be a compact connected orientable n-manifold. Then for each i, $H_i(M, \partial M; \mathbb{Q}) \cong H_{n-i}(M; \mathbb{Q})$.*

Remark. The corresponding statements for coefficients in $\mathbb{Z}$ is not true.

Corollary 7.3. *Let M be a closed orientable m-manifold, with m odd. Then $\chi(M) = 0$.*

Corollary 7.4. *For a compact orientable m-manifold M, with m odd, $\chi(M) = (1/2)\chi(\partial M)$.*

Proof. Let DM be two copies of M glued along ∂M, via the 'identity' map. Then a triangulation of M induces one for DM. Counting i-simplices gives $0 = \chi(DM) = 2\chi(M) - \chi(\partial M)$. □

Theorem 7.5. *Let M be a compact orientable 3-manifold, with at least one component of ∂M not a 2-sphere. Then there is an element of $H_1(\partial M)$ which has infinite order in $H_1(M)$.*

Proof. Let $\hat{M}$ be the 3-manifold obtained by attaching a 3-ball to each 2-sphere component of ∂M. Then $H_1(\hat{M}) \cong H_1(M)$. Since $\hat{M}$ is not closed, $H_3(\hat{M}) = 0$ and so $\beta_3(\hat{M}) = 0$. Since M is orientable, so is $\hat{M}$ and $\partial \hat{M}$. Since $\partial \hat{M}$ contains no 2-spheres, $\chi(\partial \hat{M}) \leq 0$. Corollary 7.4 implies that $\chi(\hat{M}) \leq 0$. But $\chi(\hat{M}) = \beta_0(\hat{M}) - \beta_1(\hat{M}) + \beta_2(\hat{M}) - \beta_3(\hat{M}) = 1 - \beta_1(\hat{M}) + \beta_2(\hat{M}) \leq 0$. So, $\beta_1(\hat{M}) > \beta_2(\hat{M})$. Therefore, in the long exact sequence of the pair $(\hat{M}, \partial \hat{M})$, the map $H_1(\hat{M}; \mathbb{Q}) \rightarrow H_1(\hat{M}, \partial \hat{M}; \mathbb{Q})$ has non-trivial kernel. Hence, there is an element of $H_1(\hat{M}; \mathbb{Q})$ in the image of $H_1(\partial \hat{M}; \mathbb{Q})$. Clearing denominators from the coefficients gives an infinite order element of $H_1(\hat{M})$ in the image of $H_1(\partial \hat{M})$. The following diagram commutes, where each map is induced by inclusion.

$$\begin{array}{ccc} H_1(\partial \hat{M}) & \longrightarrow & H_1(\hat{M}) \\ \Big\downarrow \cong & & \Big\uparrow \cong \\ H_1(\partial M) & \rightarrow & H_1(M) \end{array}$$

This proves the theorem. □

We introduce some standard terminology.

Definition. A 3-manifold M is *irreducible* if any embedded 2-sphere bounds a 3-ball in M.

By Proposition 3.5, a 3-manifold is irreducible if and only if it is prime and not $S^2 \times S^1$.

Theorem 7.6. *Let M be a compact irreducible 3-manifold with $H_1(M)$ infinite. Then M contains a connected 2-sided non-separating properly embedded incompressible surface S, which is not a 2-sphere. Furthermore, if there is an infinite order element of $H_1(M)$ in the image of $H_1(\partial M)$, then we may guarantee that ∂S has non-zero signed intersection number with some loop in ∂M.*

Lemma 7.7. *Let M be a compact connected 3-manifold and let X be a space with $\pi_2(X) = 0$. Then, for any basepoints $m \in M$ and $x \in X$, any homomorphism $\pi_1(M, m) \to \pi_1(X, x)$ is induced by a map $M \to X$.*

Proof. Pick a triangulation of M with m a 0-simplex. The 0-simplices and 1-simplices form a graph in M. Pick a maximal tree T in this graph and map it to x. For each remaining 1-simplex σ_1 of M, there is a unique path in T joining the endpoints of σ_1. The union of this path with σ_1 forms a loop which (when oriented) represents an element of $\pi_1(M, m)$. The given homomorphism $\pi_1(M, m) \to \pi_1(X, x)$ determines a loop in X (up to homotopy). Send σ_1 to this loop.

Let σ_2 be any 2-simplex of M. Its three boundary 1-simplices $\partial\sigma_2$ have been mapped into X. Since $\partial\sigma_2$ is homotopically trivial in M and group homomorphisms send the identity element to the identity element, the image of $\partial\sigma_2$ is homotopically trivial in X. Using this homotopy, we may extend our map over σ_2.

Now, let σ_3 be any 3-simplex of M. We have mapped $\partial\sigma_3$ to a 2-sphere in

X. Since $\pi_2(X) = 0$, this extends to a map of the 3-ball into X. Hence, we may extend over each 3-simplex. □

Lemma 7.8. *Let M be a compact irreducible 3-manifold, and let X be a pl k-manifold containing a properly embedded 2-sided $(k-1)$-submanifold Y. Suppose that* $\ker(\pi_1(Y) \to \pi_1(X)) = 1$ *and* $\pi_2(Y) = \pi_2(X - Y) = \pi_3(X) = 0$. *Then any map $f \colon M \to X$ is homotopic to a map g such that*

(i) each component of $g^{-1}(Y)$ is a properly embedded 2-sided incompressible surface in M,

(ii) no component of $g^{-1}(Y)$ is a 2-sphere, and

(iii) for properly chosen product neighbourhoods $\mathcal{N}(Y)$ and $\mathcal{N}(g^{-1}(Y))$, the map $g|_{\mathcal{N}(g^{-1}(Y))}$ sends fibres homeomorphically onto fibres.

Proof. Since Y is a pl submanifold of X, there is a triangulation of X in which Y is a union of simplices. By assumption, $\mathcal{N}(Y)$ is a product I-bundle. Hence, we may alter the triangulation of X, by replacing each simplex σ of Y with the standard triangulation of the product $\sigma \times [-1,1]$. Then $Y = Y \times \{0\}$ embeds in X transversely to the triangulation. Using the Simplicial Approximation Theorem, we may subdivide a given triangulation of M and perform a homotopy to f so that afterwards it is simplicial.

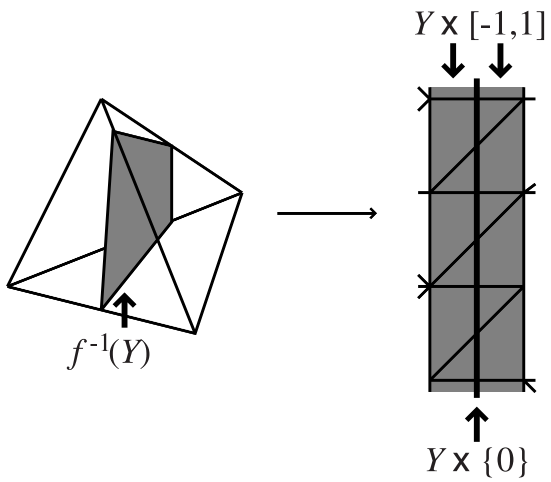

Figure 23.

Then each component of $f^{-1}(Y)$ is a properly embedded 2-sided surface, satisfying condition (iii) relative to $Y \times [-1/2, 1/2]$ and $f^{-1}(Y \times [-1/2, 1/2])$. If $f^{-1}(Y)$ is incompressible, and no component is a 2-sphere, we are done.

Suppose now that D is a compressing disc for $f^{-1}(Y)$. Choose a regular neighbourhood $\mathcal{N}(D)$ in M such that $A = \mathcal{N}(D) \cap f^{-1}(Y)$ is an annulus properly embedded in $\mathcal{N}(D)$. Let D_1 and D_2 be disjoint discs properly embedded in $\mathcal{N}(D)$ such that $\partial D_1 \cup \partial D_2 = \partial A$. Define $f_1 \colon M \to X$ as follows. Put $f_1|_{M-\mathrm{int}(\mathcal{N}(D))} = f|_{M-\mathrm{int}(\mathcal{N}(D))}$. The map $f|_{D_i}$ is a trivialising homotopy for the curve $f|_{\partial D_i}$. Since $\ker(\pi_1(Y) \to \pi_1(X)) = 1$, we may extend $f_1|_{\partial D_i}$ to a map $f_1|_{D_i}$ into Y. Extend f_1 over a small neighbourhood $\mathcal{N}(D_i)$ of D_i using the product structure of $\mathcal{N}(Y)$. Then $\mathcal{N}(D) - (\mathrm{int}(\mathcal{N}(D_1 \cup D_2))$ is three 3-balls. On their boundaries, f_1 is already defined, mapping into $Y - X$. Since $\pi_2(Y - X) = 0$, we may extend f_1 over all of $\mathcal{N}(D)$, avoiding Y. Then $f_1^{-1}(Y) = f^{-1}(Y) \cup D_1 \cup D_2 - \mathrm{int}(A)$. Thus, $f_1^{-1}(Y)$ is obtained from $f^{-1}(Y)$ via a compression. It therefore reduces the complexity of the surface, defined in §3. Note that f and f_1 differ only within a 3-ball, and therefore they are homotopic, since $\pi_3(X) = 0$.

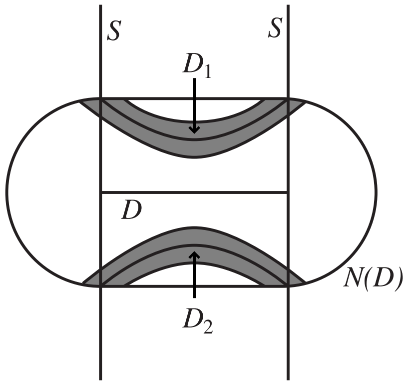

Figure 24.

If some component of $f^{-1}(Y)$ is a 2-sphere, then it bounds a 3-ball B in M. We define a map $f_1 \colon M \to X$ as follows. Let $f|_{M-\mathrm{int}(B)} = f_1|_{M-\mathrm{int}(B)}$. Using that $\pi_2(Y) = 0$, we may extend $f|_B$ to a map $f_1|_B \colon B \to Y$. Then use the product structure on $\mathcal{N}(Y)$ to define a small homotopy so that $f_1(B) \cap Y = \emptyset$, removing the 2-sphere component of $f^{-1}(Y)$. This leaves the complexity of the surface unchanged, but it reduces the number of components. Hence, we eventually obtained the map g as required. □

Proof of Theorem 7.6. Since $H_1(M)$ is infinite but finitely generated, it has $\mathbb{Z}$ as a summand. Hence, there is a surjective homomorphism $H_1(M) \to \mathbb{Z}$. If there is an infinite order element of $H_1(M)$ in the image of $H_1(\partial M)$, we may assume that the composition $H_1(\partial M) \to H_1(M) \to \mathbb{Z}$ is surjective.

Now, there is a surjective homomorphism $\pi_1(M) \to H_1(M)$ which sends a based oriented loop in M to a sum of oriented 1-simplices representing that loop. Hence, there is a surjection $\pi_1(M) \to \mathbb{Z}$. In the case where there is an infinite order element of $H_1(M)$ in the image of $H_1(\partial M)$, we may take $\pi_1(\partial M) \to \pi_1(M) \to \mathbb{Z}$ to be surjective. The map $\pi_1(M) \to \mathbb{Z}$ is induced by a map $M \to S^1$, by Lemma 7.7. Apply Lemma 7.8 to a point Y in S^1. Then some component of $g^{-1}(Y)$ is a 2-sided non-separating incompressible surface S in M that is not a 2-sphere. If $\pi_1(\partial M) \to \pi_1(M) \to \mathbb{Z}$ is surjective, a loop in ∂M mapping to $1 \in \mathbb{Z}$ must have odd signed intersection number with ∂S. $\square$