

Bundles

Michael Weiss

March 28, 2015

This is a brief note on fiber bundles and principal bundles. The sole purpose is to clarify the relation between a principal bundle and its associated bundles. No formal proofs.

First we work through the case of the Möbius band, then the tangent bundle of the 2-sphere. We describe how to go from the fiber bundle to the principal bundle, and back again.

Fiber Bundle to Principal Bundle

Let (X, B, π, F, G) be the band, where:

1. X is the band (a topological space).
2. B is the base space, which is the circle S^1 ; we can think of B as the centerline of the band.
3. $\pi : X \rightarrow B$ is the projection map.
4. F is the fiber, which is the interval $[-1, 1]$. So for each $b \in B$, $\pi^{-1}(b)$ is homeomorphic to F , but *not canonically*.

5. G is the structure group, namely $\mathbb{Z}_2 = \{1, -1\}$, which acts on F in the obvious way: 1 is the identity mapping, and -1 sends s to $-s$. We write gs for this action ($g \in G, s \in F$). We make G into a topological group by giving it the discrete topology.

Next: the principal bundle (P, B, τ, G) , where:


1. P is the total space of bundle.
2. B is the base space.
3. τ is the projection map.
4. G is both the group and the fiber.

The fibers $\tau^{-1}(b)$ are all homeomorphic to G (but *not canonically*), so P is a double cover of B . The idea is to make P look as much like X as possible, given the different fiber. There are exactly two double covers of B : the trivial one and the one given by $z \mapsto z^2$ on the unit circle. The latter is the one we want; we can think of it as the edge of the band, which wraps twice around the centerline.

To construct P , we imagine that the band X is constructed by taking two rectangular strips of paper and gluing them together in the right way (i.e., with a twist: see fig.1). P is constructed almost as a byproduct, since each strip contains a top edge and a bottom edge, and these four edges get glued into a single circle as the two strips are glued together to make the band.

More formally, we cover B with open sets $\{U_\alpha : \alpha = 1, 2\}$ which “trivialize” the bundle X , i.e., for which $\pi^{-1}(U_\alpha)$ is homeomorphic to $U_\alpha \times F$. (So those are the two strips.) Also, these homeomorphisms preserve fibers: if $\varphi_\alpha : \pi^{-1}(U_\alpha) \rightarrow U_\alpha \times F$ is one of the homeomorphisms, then φ_α maps $\pi^{-1}(b)$ to the fiber over b in $U_\alpha \times F$, i.e., $\{b\} \times F$.

Figure 1: Möbius Band



Möbius Band obtained by gluing two rectangular strips together. The inner two vertical edges are glued together without a twist, the outer two vertical edges with a twist. In the formal treatment, not just vertical edges but vertical open “zones” are glued (identified), since U_α and U_β are open sets, and the α and β representations of $\{b\} \times [-1, 1]$ are identified for any $b \in U_\alpha \cap U_\beta$.

So if $b \in U_\alpha \cap U_\beta$ ($\alpha \neq \beta$), then φ_α and φ_β both make $\pi^{-1}(b)$ homeomorphic to F . But in two different ways: $\varphi_\beta \varphi_\alpha^{-1}$ induces a homeomorphism of F to itself, and this composition is either the identity or the map $s \mapsto -s$. In other words, it's given by one of the actions of G on F . (We say “induces” because technically, $\varphi_\beta \varphi_\alpha^{-1}$ maps $\{b\} \times F$ to itself, a fine point we will ignore from now on.)

This formal description assumes we already have X in hand, otherwise we can't speak of $\pi^{-1}(U_\alpha)$. Let's make the math follow the arts and crafts more closely. Start with the spaces $\{U_\alpha \times F : \alpha = 1, 2\}$, i.e., the two strips of paper shown in fig.1. Take the disjoint union. For each $b \in U_\alpha \cap U_\beta$, we chose an element of G , call it $g_{\alpha\beta}(b)$. Moreover, let's make $g_{\alpha\beta}$ continuous. Then $g_{\alpha\beta}$ tells us how to glue the strips together. Formally, we identify (b, s) as an element of $U_\alpha \times F$ with $(b, g_{\alpha\beta}(b)s)$ as an element of $U_\beta \times F$. (Recall that gs stands for the action of G on F .) The $g_{\alpha\beta}$ work as gluing instructions, though they are called transition functions.

Now for the punchline: we can do the same thing if we replace the fiber F with G as a fiber. We start off with the disjoint union of spaces $\{U_\alpha \times G : \alpha = 1, 2\}$ — that's the four edges. The transition functions $g_{\alpha\beta}$ restrict to transition functions for G , so we know how to glue the four edges together.

Time to generalize. Let (X, B, π, F, G) be any old fiber bundle. So we have an open cover $\{U_\alpha\}$ of B , with trivializing homeomorphisms $\varphi_\alpha : \pi^{-1}(U_\alpha) \rightarrow U_\alpha \times F$ preserving fibers. Moreover each transition function $\varphi_\beta \varphi_\alpha^{-1} : F \rightarrow F$ is given by the action of G on F , according to the formula

$$\varphi_\beta \varphi_\alpha^{-1}(s) = g_{\alpha\beta}(b)s, \quad b = \pi(s)$$

where $g_{\alpha\beta} : U_\alpha \cap U_\beta \rightarrow G$ is a continuous function. So we can form the disjoint union of the $\{U_\alpha \times F\}$, and then use the $g_{\alpha\beta}$ to glue together (i.e., identify) various fibers.

Next, the principal bundle: we form the disjoint union of the $\{U_\alpha \times G\}$. Potential problem: there's no reason to believe that G will be a subset of F . So we can't "restrict" the $g_{\alpha\beta}$ to G . But this problem dissolves when you realize that $g_{\alpha\beta}(b)$ is an element of G , so it acts on G by multiplication (on the left).

Let's do another example: the tangent bundle of S^2 . The fibers $T_b(S^2)$ are all homeomorphic to $\mathbb{R}^2$; indeed they are linearly isomorphic. So set $F = \mathbb{R}^2$. For each U_α we flatten the open set diffeomorphically onto the plane; this gives linear isomorphisms $\varphi_\alpha(b) : T_b(S^2) \rightarrow \mathbb{R}^2$ for each $b \in U_\alpha$ (the derivative maps). (As before, we will be cavalier about the distinction between $T_b(S^2) \rightarrow \mathbb{R}^2$ and $T_b(S^2) \rightarrow \{b\} \times \mathbb{R}^2$.)

It's pretty clear that the transition functions $\varphi_\beta \varphi_\alpha^{-1}$ are elements of $GL_{\mathbb{R}}(2)$, thus $g_{\alpha\beta} : U_\alpha \cap U_\beta \rightarrow GL_{\mathbb{R}}(2)$. In fact, the $g_{\alpha\beta}$ are all smooth, but let's not get into that.

Now we form the principal bundle. This is often called the frame bundle. Here's the idea: if $[\hat{x}, \hat{y}]$ is the canonical basis of $\mathbb{R}^2$, then we have a canonical 1–1 correspondence between $GL_{\mathbb{R}}(2)$ and the set of bases of $\mathbb{R}^2$, namely $g \leftrightarrow$

$[g\hat{x}, g\hat{y}]$. (We use brackets to indicate an ordered basis.) So the elements of the fiber can be viewed as “frames”.

Let's try to picture an element t of the principal bundle above $b \in S^2$. From our definitions, it's an equivalence class: for each $U_\alpha \ni b$, we have a “representative” of t in $U_\alpha \times GL_{\mathbb{R}}(2)$, call it $[\hat{x}_\alpha, \hat{y}_\alpha]$. Representatives are related by the transition functions, and t is the equivalence class of all the representatives.

However, we can give a more concrete description. We have a linear isomorphism $\varphi_\alpha(b) : T_b(S^2) \rightarrow \mathbb{R}^2$. That sets up a 1–1 correspondence between bases in $T_b(S^2)$ and bases in $\mathbb{R}^2$. The same goes, *mutatis mutandis*, for φ_β . Now $g_{\alpha\beta}(b)$ is given by a composition:

$$[\hat{x}_\alpha, \hat{y}_\alpha] \xrightarrow{\varphi_\alpha^{-1}(b)} [\hat{x}_t, \hat{y}_t] \xrightarrow{\varphi_\beta(b)} [\hat{x}_\beta, \hat{y}_\beta]$$

where $[\hat{x}_t, \hat{y}_t] \subset T_b(S^2)$ is a basis “living on the sphere”. Each equivalence class is represented in a natural way by a basis of $T_b(S^2)$.

Obviously we have been way too specific, and this discussion works for any differential n -manifold. If the manifold has more structure, we can replace $GL(n)$ with other groups. For example, if it's a Riemannian manifold, then we can restrict ourselves to orthonormal bases, and replace $GL(n)$ with $O(n)$, the orthogonal group.

Principal Bundle to Fiber Bundle

Suppose we have a principal bundle (P, B, τ, G) . How can we replace the fiber G with another fiber F ? Answer: if we have a left action of G on F , say $(g, s) \mapsto gs$, then we can define the so-called associated bundle (X, B, π, F, G) .

The definition is rather simple: just let $\tilde{X} = P \times F$, and identify:

$$(p, s) \equiv (pg^{-1}, gs) \quad \forall g \in G$$

and let X be $\tilde{X}/\equiv$. I omit the rest of the construction.

The fiber will be F . I like to look at it this way: we start with a point $b \in B$. For each $p \in P$ in the τ -fiber over b , we have a *coordinate system* for the π -fiber over b . Coordinates are elements of F ; the coordinate of $(p, s)/\equiv$ is just s , using the coordinate system associated with p .

Let's look at the Möbius band. The principal bundle looks like the edge, a double covering of the centerline. The τ -fiber over b is a pair of points. Now we want to construct the band. Here the fibers will be homeomorphic to $F = [-1, 1]$. Given a $b \in B$, we have two coordinate systems for $\pi^{-1}(b)$, as we can lay out $[-1, 1]$ along $\pi^{-1}(b)$ going either way. A coordinate s using one system becomes $-s$ using the other system. Thus $-1 \in G$ will take us from one edge to the other, and will adjust the coordinate accordingly.

The tangent space example is also not too difficult. The elements of the τ -fiber can be viewed as bases of $T_b(S^2)$, as noted. If t is such a basis, then we have a coordinate system for $T_b(S^2)$ associated with it. Coordinates are elements of the fiber $\mathbb{R}^2$, i.e., pairs of real numbers. A change of basis results in a change of coordinates, and “contravariantly”. For example, if we go from $[\hat{x}, \hat{y}]$ to $[2\hat{x}, \hat{y}]$, then the coordinates (x, y) go to $(\frac{1}{2}x, y)$. That's why we have (pg^{-1}, gs) in the definition of $\equiv$.

Torsors and the Principal Bundle

When constructing the principal bundle, we started with the disjoint union of the $U_\alpha \times G$, and then identified elements. In the Möbius band, these $U_\alpha \times G$ were two pairs of edges, which glued together became the edge of the band.

There's something a bit odd about this: the group $G = \{1, -1\}$ has a distinguished element, namely the identity. However, the two edges of a pair “look the same”. To reflect this, we should really start with the disjoint union of $U_\alpha \times T$, where T has the same cardinality as G , but the bijection between T and G is *not canonical*.

The notion needed is a *torsor*. T is a torsor for G if:

1. There is a right action of G on T : $(t, g) \mapsto tg$.
2. The action is transitive: for any $s, t \in T$, there is a g such that $sg = t$.
3. The action is free: except for $1 \in G$, the action is without fixed points.
So $(\exists t)tg = t \Rightarrow g = 1$.

It follows that as soon as we pick an element $t_0 \in T$, we have a 1–1 correspondence between G and T , $g \mapsto t_0g$.

Accordingly, we write (P, B, τ, G, T) for a principal bundle, replacing G with T as the “generic fiber”.

The prototypical example of a torsor is an oriented circle: the rotation group $U(1)$ acts on it. In contrast, the sphere S^2 is not a torsor for the rotation group $SO(3)$, since the action is not free.

Any set with two elements is a torsor for $\mathbb{Z}_2$, with a uniquely determined action. This is the torsor we need for the Möbius band.

Let V be a finite-dimensional real vector space. The set of bases for V (call it T) is a torsor for $GL(V)$. It is slightly less obvious, but T is also a torsor for $GL_{\mathbb{R}}(n)$, if $\dim V = n$. How do we define the action of $g \in GL_{\mathbb{R}}(n)$ on a basis $[v_1, \dots, v_n]$? Answer: $[v_1, \dots, v_n]$ determines a linear isomorphism between V and $\mathbb{R}^n$, call it $q : V \rightarrow \mathbb{R}^n$. Then $q^{-1}gq$ is in $GL(V)$, so we can apply it to $[v_1, \dots, v_n]$. In matrix terms: g has a canonical representation in $\text{Mat}_{\mathbb{R}}(n \times n)$ (just use the standard basis), and we can apply the same matrix to $[v_1, \dots, v_n]$.

Example: if $g \in GL_{\mathbb{R}}(2)$ is defined by $[\hat{x}, \hat{y}]g = [2\hat{x}, \hat{y}]$, then for any basis $[v_1, v_2]$ of V , we let $[v_1, v_2]g = [2v_1, v_2]$.

This highlights a key property of principal bundles, not shared with more general bundles: we have an action of G on the entire space P . Not so for

(X, B, π, G, F) . We have an action of G on the generic fiber F , but no way of applying this “across the board”, since the homeomorphisms between $\pi^{-1}(b)$ and F are not canonical.

For example, compare the tangent with the frame bundle on S^2 . $G = GL_{\mathbb{R}}(2)$. As noted, we know how to apply $g \in G$ to any basis of $T_p(S^2)$. But if we are looking at a vector $v \in T_p(S^2)$, we don’t generally know how to apply g to v , as illustrated by $[v_1, v_2] \mapsto [2v_1, v_2]$.