

Fiber Bundles

In the discussion of topological manifolds, one often comes across the useful concept of starting with two manifolds M_1 and M_2 , and building from them a new manifold, using the product topology: $M_1 \times M_2$. A fiber bundle is a natural and useful generalization of this concept.

4.1 Product Manifolds: A Visual Picture

One way to interpret a product manifold is to place a copy of M_1 at each point of M_2 . Alternatively, we could be placing a copy of M_2 at each point of M_1 . Looking at the specific example of $\mathbb{R}^2 = \mathbb{R} \times \mathbb{R}$, we take a line \mathbb{R} as our base, and place another line at each point of the base, forming a plane. For another example, take $M_1 = S^1$, and $M_2 =$ the line segment $(0,1)$. The product topology here just gives us a piece of a cylinder, as we find when we place a circle at each point of a line segment, or place a line segment at each point of a circle.




Figure 4.1 The cylinder can be built from a line segment and a circle.

A More General Concept

A fiber bundle is an object closely related to this idea. In any local neighborhood, a fiber bundle looks like $M_1 \times M_2$. Globally, however, a fiber bundle is generally not a product manifold. The prototype example for our discussion will be the möbius band, as it is the simplest example of a nontrivial fiber bundle. We can create the möbius band by starting with the circle S^1 and (similarly to the case with the cylinder) at each point on the circle attaching a copy of the open interval $(0,1)$, but in a nontrivial manner. Instead of just attaching a bunch of parallel intervals to the circle, our intervals perform a 180° twist as we go around. This gives the manifold a much more interesting geometry.

We look at the object we've formed, and note that *locally* it is indistinguishable from the cylinder. That is, the “twist” in the möbius band is not located at any particular point on the band; it is entirely a global property of the manifold. Motivated by this example, we seek to generalize the language of product spaces, to include objects like the möbius band which are only locally a product space. This generalization is what we will come to know as a fiber bundle.



Figure 4.2 The möbius band can also be built from a circle and line segment, but using a more general method of construction.

4.2 Fiber Bundles: The Informal Description

When we build up the language to describe a fiber bundle, we want to regard a fiber bundle as “locally a product space” in the same sense that a manifold is “locally euclidean”. For this reason, the language describing fiber bundles will mimic the language of manifolds quite closely.

The manifold we are creating will be called the total space, E . It will be constructed from a base manifold M , and a fiber F . In our examples of the cylinder and the möbius band, S^1 is the base manifold, and the interval $(0,1)$ is the fiber. However, since globally the cylinder and möbius band differ, we're definitely going to need some additional data to distinguish them.

We shall find that for a general point q in E , we can directly associate this with a point p in the base manifold, M . However, we cannot generally associate q with a point in the fiber, F . This is an important fact which makes the fiber bundle a more general object than a product space; M and F are fundamentally different objects, and it can be seen readily in the case of the möbius band:

Say we want to parameterize the möbius band by a point θ on the circle, and a real number f in the interval $(0,1)$. Concretely, say $\theta = 0$ and $f = \frac{3}{4}$. Now, we transport the point around the möbius band by increasing θ and keeping f fixed at $\frac{3}{4}$. When $\theta \rightarrow 2\pi$, we should return to the same point q in E , since it corresponds to the same θ and f , and hence the same q . However, because of the inescapable twist in the band, the point we return to is associated with $\theta = 0, f = \frac{1}{4}$. Our parameterization for F somehow “flips” when we move one turn around the möbius band.




Figure 4.3 Fiber parameterization on the möbius band is not global.

What you should take away from this is that the parameterization for F only works in a local sense; it does not extend globally. You may be tempted to think “Okay, this parameterization certainly won't do, but perhaps we're just not looking hard enough. How do we know that there is *no* global fiber parameterization for the möbius band?” Let's say there existed such a parameterization. Then the set of points $f = \frac{3}{4}$ would give us a closed loop going around the möbius band. Similarly, the points with fiber-value $f = \frac{1}{4}$ would constitute another closed loop around the band. However, you can check that any two such loops around the möbius band must cross at some point, and at this point our fiber parameterization would not be well-defined. Fiber parameterization is not global. In this way, it is a great deal like coordinate parameterization on a manifold. In a neighborhood of a point, we can parameterize points in E by (p, f) , but when we go to another neighborhood, we use a different parameterization, (p, f') .

How to express this mathematically? First, we associate each point in E with a point in M , which we *can* globally do. This association can be accomplished by a projection map

$$\pi : E \rightarrow M \tag{4.1}$$

which projects points q in E to their associated point p in M . This map is generally not one-to-one, of course; we want it to map entire fibers F_p to points p , capturing the fact that we're

attaching a copy of the fiber F to each point p . We can enforce this condition by the following requirement on the preimage:

$$\pi^{-1}(p) \approx F \tag{4.2}$$

for each p in M . We still want to locally parameterize these fibers, but leave the definition open to include different parameterizations of F for different neighborhoods. This is the part which may be familiar from the standpoint of coordinate parameterization.

When one defines coordinates on manifolds, it is generally accomplished by an open covering $\{U_i\}$ and a homeomorphism Φ_i for each U_i associating it with an open set of \mathbb{R}^n . We will do something very similar in the case of fiber bundles. We take an open covering of M , $\{U_i\}$, and a set of smooth homeomorphisms, $\{\Phi_i\}$, associating an open set of E given by the preimage $\pi^{-1}(U_i)$ with a product space. Formally,

$$\Phi_i: U_i \times F \rightarrow \pi^{-1}(U_i) \tag{4.3}$$

Since we have a map Φ_i between $\pi^{-1}(U_i)$ and $U_i \times F$, we can locally express points in E using points in $U_i \times F$. We first project the point q in E to a point p in M , using $p = \pi(q)$, then find an open neighborhood U_j of p , then there is a corresponding map Φ_j which associates q with a point in $U_j \times F$, i.e. a pair of points (p, f) .

Now, it is possible that we could have chosen a different U_k about p , thus a different map Φ_k associating it with a different point (p, f') in $U_k \times F$. This is fine, but we need to understand the relationship between f and f' . In other words, to distinguish the bundle properly, we need to know about all possible choices of fiber parameterization. In the case of the cylinder, there was only one fiber parameterization, because the space was globally a product manifold. In the case of the möbius band, there are two possible parameterizations, and we can make the transformation explicit by

$$f' = 1 - f \tag{4.4}$$

Neither parameterization f nor f' works globally, but we can cover the circle with two overlapping segments, and choose one parameterization for one segment, and the opposite for the other segment.

Changes in the parameterization of the fiber are known as transition functions. These are written formally as

$$t_{ij} = \Phi_i \circ \Phi_j^{-1} : (U_i \cap U_j) \times F \rightarrow (U_i \cap U_j) \times F \tag{4.5}$$

so they smoothly carry a point from one product space to another, in the overlapping region $U_i \cap U_j$. However, there are more enlightening ways to look at the transition functions.

First of all, note that they carry the point p to itself:

$$(p, f) \mapsto (p, f') \tag{4.6}$$

So, it may be more enlightening to think of this as a set of maps from F to itself for each point p in the overlap. Symbolically,

$$t_{ij}(p): F \rightarrow F \tag{4.7}$$

Now, notice that $\{t_{ij}\}$ satisfy group axioms:

$$\begin{aligned} t_{ij}(p)[t_{jk}(p)] &= t_{ik}(p) && (\text{closure}) \\ t_{ii}(p)[f] &= f && (\text{identity element}) \\ t_{ji}(p) &= t_{ij}^{-1}(p) && (\text{inverses}) \end{aligned} \tag{4.8}$$

We can now think of $\{t_{ij}(p)\}$ as a group. Specifically, since we are interested in smooth transition functions, $\{t_{ij}(p)\}$ is a lie group. This group is denoted G , and is called the structure group of the fiber bundle. Of course, we can't forget that it is also a map from the fiber to itself. This is a realization of the structure group. This group action is also required to be smooth, since that was an original requirement on the transition functions. To summarize our repackaging, t_{ij} is a map

$$t_{ij}: U_i \cap U_j \rightarrow G \tag{4.9}$$

into the structure group G which acts smoothly on the fiber F . The transition functions characterize the fiber bundle. In the case of the cylinder, the structure group is just the trivial group of one element. In the möbius band, the structure group is the group of two elements, Z_2 , given by $\{1, x\}$, where $x^2 = 1$. In other words, we only have two parameterizations, and thus only one transition function other than the identity, which is its own inverse:

$$f'' = 1 - (1 - f) = f \tag{4.10}$$

In these two cases, the structure group is a discrete, finite set of elements, and therefore the dimensionality of the lie group is zero. Keep in mind that these are very simple cases of fiber bundles, and generally the lie group consists of a continuous spectrum of transition functions. In other words, we call this a lie group for a reason. Also keep in mind that we have a choice when determining our structure group, since we don't have to use all of the elements. For example, if the fiber bundle is trivial, like the cylinder, we can use any group G we want, but let it act trivially on the fiber. However, it generally makes the most sense to choose the smallest group that is convenient for our purposes.

4.3 Fiber Bundles: The Bloated Description

We have finally laid out all the pieces we need to describe a fiber bundle. Let's give a preliminary brute-force formal definition, before eventually refining it to look a bit nicer.

A Differential Fiber Bundle (E, π, M, F, G) consists of the following:

1. A differential manifold E called the total space
2. A differential manifold M called the base manifold
3. A differential manifold F called the fiber
4. A surjective map $\pi: E \rightarrow M$ called the projection, such that $\pi^{-1}(p) = F_p \cong F$
5. An open covering $\{U_i\}$ of M with diffeomorphisms

$$\Phi_i: U_i \times F \rightarrow \pi^{-1}(U_i) \tag{4.11}$$

called local trivializations, with

$$\pi(\Phi_i(p, f)) = p \tag{4.12}$$

6. A lie group, G , known as the structure group, which acts on the fiber, F .

Finally, there is the requirement that the transition functions,

$$\Phi_i \circ \Phi_j^{-1} = t_{ij}(p) \tag{4.13}$$

are smooth and live in G , the structure group.

$$\Phi_j(f) = \Phi_i(t_{ij}(p) \cdot f) \tag{4.14}$$

The base manifold and the fiber tell you exactly what the bundle looks like locally. At the level of the base manifold, M , open neighborhoods just look like pieces of \mathbb{R}^n , and M 's transition functions tell us precisely how they are sewn together. At the level of the bundle, open neighborhoods are just pieces of $\mathbb{R}^n \times F$, and there is an additional sewing operation. We need to glue the fiber F_p over p from the patch U_i to the same fiber F_p from the patch U_j . The structure group gives you the additional information required to tell you how to “glue” the fibers together.




Figure 4.4 The transition functions specify precisely how the fibers are sewn together.

4.4 Special Types of Fiber Bundles

In the general case of fiber bundles, F can be any differential manifold, and G can be any lie group that acts on F . By adding further requirements, we can define certain special bundles.

The Trivial Bundle

Almost unnecessary to include, except that it draws the important connection which shows that bundles are generalizations of product manifolds. Simply put, a trivial bundle is a product manifold. The base and fiber are interchangeable, and the structure group is the trivial group. The trivial bundle can be covered with one patch; the entire base manifold. The “local trivialization” Φ is really a global trivialization, since Φ covers all of the base. The trivial bundle can clearly be formed using any two manifolds as base and fiber.

Vector Bundles

A vector bundle is defined by two restrictions: first, the requirement that F be isomorphic to \mathbb{R}^k , i.e. that F is a vector space. Second, that the structure group acts linearly on the vector space. Because of these two requirements, the transition functions have a k -dimensional representation on the fibers. In other words, the transition functions are subgroups of $GL_k\mathbb{R}$.

Special Case: The Tangent Bundle

Take the base manifold to be M , and the fiber at p to be the tangent space T_pM , which is indeed a vector space. The projection operator sends $T_pM \rightarrow p$. For the open covering, we can use the same coordinate patches $\{U_i\}$ that we use to define the manifold structure of M .




Figure 4.5 The tangent bundle is an example of a fiber bundle with which we are already familiar.

This space is known as the tangent bundle of M, $E = TM$. We have a local trivialization in any given patch, given by the coordinate representation of the vectors V^i in T_pM . In other words, the coordinate charts not only give us a local parameterization for M, they also give us a local parameterization for TM, in that they give us the vector components.

In the neighborhood U_A we use coordinates $\{x^i\}$ and

$$V_p \in TM = V^i \frac{\partial}{\partial x^i} \Big|_p \quad (4.15)$$

and in the neighborhood U_B we use coordinates $\{y^j\}$ and

$$V_p \in TM = \tilde{V}^j \frac{\partial}{\partial y^j} \Big|_p \quad (4.16)$$

At the level of the manifold, the transition functions are given by $x(y)$ and $y(x)$, but at the level of the bundle, we see that

$$V^i = \tilde{V}^j \frac{\partial x^i}{\partial y^j} \quad (4.17)$$

The transition functions on the bundle are the partial derivatives of the manifold transition functions. These can be thought of as elements of $GL_n\mathbb{R}$ mapping the vector components in one patch to vector components in another patch. Since we can arbitrarily write down smooth coordinate transformations on M, we can locally construct any set of matrices to produce the transition functions. Thus, the structure group of the tangent bundle of an n-dimensional manifold is simply $GL_n\mathbb{R}$. The transition functions which the manifold itself together generate the maps which glue the fibers together.

Vector bundles in general are quite useful, due to their concrete nature. We will tend to think of them as generalizations of the tangent bundle. In this light, we can use familiar formalisms for defining parallel transport and curvature in the more general setting of vector bundles, as we shall soon see.

Principal Bundles

A principal bundle is a fiber bundle in which the fiber over any point is a copy of the bundle's structure group, $F = G$. Since G is a lie group, it is a manifold by definition. The group G can act on itself by left multiplication. Principal bundles will play a central role in constructing a geometric physical theory of everything.

The Frame Bundle

To construct the frame bundle, we start with a manifold M, and let the fiber over each point p be the space of all ordered bases $\{e_i\}$ for T_pM . An ordered basis provides a frame at p. So, we are looking at all frames $\{e_i\}$ at each $p \in M$. The reason we choose an "ordered" basis is so that we don't distinguish between two sets of bases $\{e_i\}$ and $\{f_j\}$ where the $\{e_i\}$ are merely a permutation of the $\{f_j\}$. This is not quite a vector bundle, because a given element of E is a set of n *linearly independent* vectors. The independence condition prevents the fiber

from being a vector space. For example, there is no “zero” element in the frame bundle. Note that given any initial basis for T_pM , one can find any other by operating on this basis by a suitable element of $GL_n\mathbb{R}$:

$$f_j = M_{ij} e_i \tag{4.18}$$




Figure 4.6 Any frame can be replaced by the element of the general linear group which gets you there from the reference frame.

Since the $\{f_j\}$ are a linear combination of the $\{e_j\}$, and they are linearly independent since g is nonsingular, the $\{f_j\}$ do indeed form a basis for T_pM . Moreover, this exhausts the space of all frames. If you provide me with any frame $\{h_j\}$, it is always possible to express each h_j as a sum of vectors in the original frame. This is equivalent to writing $h_j = M_{ij} e_i$, where M is an invertible matrix.

By starting with any fiducial or “point-of-reference” basis, we can get all other bases by acting with elements of $GL_n\mathbb{R}$. Then we can label a given frame by the element of $GL_n\mathbb{R}$ that got us there from the fiducial frame. In this way, we have shown that the fiber of all frames is nothing but the group $GL_n\mathbb{R}$! In other words, the frame bundle is equivalent to a $GL_n\mathbb{R}$ principal bundle. It will be left as an exercise to make this equivalence more explicit.

The equivalence to which we've just alluded seems useful. Is there any way of naturally going the other direction? That is, can we produce some kind of useful fiber bundle from a principal bundle? The answer is yes, from a principal bundle we can build associated k -dimensional vector bundles, provided that G has a k -dimensional representation. This is useful, because vector bundles are much more concretely-described fiber bundles.

Associated Vector Bundles

The basic idea of constructing associated vector bundles is as follows: Rip out the copy of G at each $p \in M$. Replace this group fiber by a vector space V of dimension k . Choose the k -dimensional representation ρ_k of G . Then choose the transition functions to be $\rho_k(t_{ij}) = k \times k$ matrices acting on V .

More formally, let V be a k -dimensional vector space on which G acts via a k -dimensional representation, ρ_k . Then, given a principal bundle P , define the associated vector bundle $P \times_{\rho} V$ by starting with $P \times V$ and imposing the equivalence relation

$$(u, v) \sim (u \cdot g^{-1}, \rho(g)v) \tag{4.19}$$

In a given local neighborhood of u , we can write this equivalence relation in coordinate-form:

$$(x, h, v) \sim (x, h \cdot g^{-1}, \rho(g)v) \tag{4.20}$$




Figure 4.7 A principal bundle has an associated vector bundle for each of its representations.

Problem 4.1 Show that every point in this space, expressed in any coordinate system, is equivalent to a point at the same location on the base, but at the identity point on the G -fiber. Since each point in the space can be expressed in this manner, we have collapsed the G -fiber to a point, while the V -fiber persists. Show that the transition functions for the persisting V -fiber are

$$\tilde{t}_{ij} = \rho_k(t_{ij}) \tag{4.21}$$

The large-scale picture you should have in mind is of a single principal bundle, underneath which sits a multitude of associated vector bundles. For every matrix representation of the structure group G , there exists another associated vector bundle.

4.5 Fiber Bundles: The More Elegant Description

As mathematicians, we are inclined to rigorously define the tools that we use. Specifically, where do they live, and what distinguishes them? For a fiber bundle, we have not yet explained this. Is it the total space? Is it the collection of spaces? What is the specific object we are calling the “bundle”, and how does it specify all the underlying structure? This is a delicate question, which is why it has been put off until we could get a clearer conceptual picture.

Upon close inspection, you may notice that the fiber bundle is entirely specified by the projection map, π , subject to a rigorous series of requirements. All the other objects are defined through π . E and M are its domain and range, and it is required that they are both differential manifolds. Since it is required that E is a differential manifold, it is assumed that its differential structure is already fixed, but this structure is subject to all the requirements in the definition. F is given by π 's preimage of a point in M . Local trivializations are required to exist and be compatible, but they play a similar role to that of coordinate charts on M . Since we required E has a specific topology and differential structure, the local trivializations are just all possible maps that are compatible with this structure. Once all trivializations are given, this implicitly defines the set of all transition functions, and hence the structure group, G . Thus, all the pieces are truly given by just the projection map, and for this reason, it is the projection map itself which is often called the “fiber bundle”.

It will sometimes be useful to deem two different bundles π_1 and π_2 to be “equivalent”. We have already done this for the case of the frame bundle and a $GL_n\mathbb{R}$ principal bundle in the previous section (although we did not give much justification). To do so, we need to be sure that the two total spaces E_1 and E_2 are equivalent differential manifolds, i.e. there exists a diffeomorphism

$$\Theta: E_1 \rightarrow E_2 \tag{4.22}$$

However, there must be an equivalence of projection maps as well, so that the base manifolds are defined in the same way. In other words, for all $u \in E_1$,

$$\Theta(\pi_1(u)) = \pi_2(\Theta(u)) \tag{4.23}$$

Since the map specifies the bundle, this diffeomorphism equivocates the two bundles. To get the more general notion of a “bundle map”, we remove the invertibility condition; that is, we require that Θ be a smooth map between the two bundles, but it need not be bijective.

Problem 4.2 Show that the frame bundle is equivalent to a $GL_n\mathbb{R}$ principal bundle, as was stated in section 4.4.

Problem 4.3 Show that the associated vector bundles described in section 4.4 are given by a smooth bundle map from a principal bundle to a vector bundle.

Sections

A section σ of a fiber bundle is a map from the base manifold into the total space, picking out a point on the fiber over each point p on the base. Since sections pick out points in the total space that lie above the point on the manifold they're mapping from, we can project back down and recover the original point:

$$\pi(\sigma(p)) = p \tag{4.24}$$

For example, a vector field is a section of a tangent bundle.

4.6 Curvature: Generalizing past the Tangent Bundle

Curvature is defined on a manifold by transporting vectors from one tangent space to another. Thus, curvature is best thought of as a quantity associated with the tangent bundle of a manifold. In this light, we can now generalize this concept to *other* vector bundles.

In order to make this jump, another conceptual reformulation must occur. When one normally approaches the concept of a connection over the tangent bundle, Γ_{jk}^i , it is done in terms of the transportation properties of *coordinate* basis vectors. Before we move into more general territory, we need to reformulate our definition of a connection and do so in the context of general basis vectors, not necessarily tied to any coordinate system. We define a connection to tell us how our basis transforms as we move along the manifold. Specifically, if we have a general basis $\{e_\mu(p)\}$ for the tangent space T_pM , and a basis $\{e^{*\nu}(p)\}$ for the cotangent space T_p^*M as well, we define the connection ω by the rate at which the basis transforms as we move in a given direction away from p . This derivative is a vector of course, so we can express it as a linear combination of basis vectors:

$$\nabla e_\mu = \omega_\mu^\sigma \otimes e_\sigma \quad (4.25)$$

The notation here may look somewhat unusual, but it is consistent with the Γ_{jk}^i formulation. The two-indexed object ω_μ^σ is a one-form for each value of μ and σ .

$$\omega_\mu^\sigma = \omega_{\nu\mu}^\sigma e^{*\nu} \quad (4.26)$$

In other words, ω can be thought of as an $n \times n$ matrix of one-forms. The one-form component of ω tells you in which direction you're transporting in the manifold.

It will be useful to determine how the components of ω transform when we use a different set of basis vectors, $\{f_\mu(p)\}$. Each basis vector will be expressible as a linear combination of the old basis:

$$f_\beta = \Lambda_\beta^\alpha e_\alpha \quad (4.27)$$

Problem 4.4 Show that the connection ω changes in the following manner under such a basis transformation:

$$\tilde{\omega}_\mu^\sigma = \Lambda_\mu^\beta \omega_\beta^\alpha \Lambda^{-1\sigma}_\alpha + d\Lambda_\mu^\alpha \Lambda^{-1\sigma}_\alpha \quad (4.28)$$

Problem 4.5 Show that this expression is consistent with the following transformation law for Γ given by change of coordinates:

$$\tilde{\Gamma}_{jk}^i = \frac{\partial x^i}{\partial y^l} \left(\frac{\partial y^m}{\partial x^j} \frac{\partial y^n}{\partial x^k} \Gamma_{mn}^l + \frac{\partial^2 y^l}{\partial x^j \partial x^k} \right) \quad (4.29)$$

Remember that you must also perform a coordinate transformation on the index associated with the basis covector that's been contracted with ω .

The key point to understand is we are no longer changing coordinates on M ; we are just changing the basis we are using for the tangent space. This is the important transition that we are required to make before moving to the more general language of vector bundles.

Connections on Vector Bundles

Consider a vector bundle V , with vector space fiber $F \cong \mathbb{R}^k$. Let $\{h_\mu(p)\}$ be a basis for F_p , $\{e_\sigma(p)\}$ a basis for T_pM , and $\{e^{*\nu}(p)\}$ a basis for T_p^*M , where μ runs from one to k , and ν and σ from one to n . Define $\nabla_\nu h_\mu$ to be the rate of change of the μ th basis vector of the fiber F in the ν th direction in the base M . This can be expressed as a linear combination of the h_ρ 's:

$$\nabla_\nu h_\mu = \omega_{\nu\mu}^\rho h_\rho \quad (4.30)$$

Once again, we can think of the ∇ operator as a base-manifold one-form, by contracting its lower index with the basis covector $e^{*\nu}(p)$:

$$\nabla h_\mu = \omega_{\mu}^\rho h_\rho \quad (4.31)$$

where

$$\omega_{\mu}^\rho = \omega_{\nu\mu}^\rho e^{*\nu}(p) \quad (4.32)$$

This may seem repetitive, but only because we've set up the notation so that we're basically copying the equations above (4.25 - 4.27) for connections on the tangent bundle. Note that the μ and ρ components of ω refer to directions in the vector-space fiber, while the ν component (which we tend to repress) refers to directions in the manifold. If we now change fiber basis

$$g_\beta = \Lambda_\beta^\alpha h_\alpha \quad (4.33)$$

we get the same transformation law (4.28), which we will often express in abstract matrix form:

$$\tilde{\omega} = \Lambda \cdot \omega \cdot \Lambda^{-1} + d \Lambda \cdot \Lambda^{-1} \quad (4.34)$$

ω is a matrix-valued one-form, meaning each of its n^2 components sits in the cotangent space of the manifold. The matrix itself also lives in an interesting space, as we will now explore.

Let's say we parallel-transport a basis vector along some coordinate direction. How do we find out what new vector results? We have seen in the past that we can make a macroscopic translation by performing a large number of infinitesimal translations. This, we have seen multiple times, results in an exponential map. We cut to the chase, taking the exponential map of the covariant derivative, and acting on a basis vector:

$$\tilde{h}_\mu = [\exp\{\Delta x^\alpha \nabla_\alpha\}] h_\mu = h_\mu + \Delta x \nabla h_\mu + \frac{1}{2} \Delta x^2 \nabla^2 h_\mu + \dots \quad (4.35)$$

We're omitting the notation for the sum $\Delta x^\alpha \nabla_\alpha$ because it's clear that's what we're doing, and adding more indices will only complicate the expression. We are simply contracting the vector Δx with the one-form component of the covariant derivative ∇ .

Problem 4.6 Show that the transported vector is found by taking the matrix exponential of the connection:

$$\tilde{h}_\mu = [\exp\{\Delta x \omega\}]^\rho_\mu h_\rho \quad (4.36)$$

From problem 4.6 we find that $(\exp\{\Delta x \omega\})$ is a transformation matrix that gives the value of a parallel-transported basis vector, given an initial basis vector. We can imagine setting up a coordinate system in which the basis $\{g_\mu\}$ is parallel-transported along this curve. In this case, the transition function from the h -basis to the g -basis will just be the

matrix $(\exp\{\Delta x \omega\})^\rho_\mu$.

What we are attempting to demonstrate here is that $\exp\{\omega\}$ is a matrix which sits in the structure group of the vector bundle,

$$\exp\{\omega\} \in G \tag{4.37}$$

which means that ω must live in the lie algebra of the structure group,

$$\omega \in \mathfrak{g} \tag{4.38}$$

To summarize, a connection ω over a vector bundle is specified by a lie-algebra-valued one-form over the base manifold.

Curvature of Vector Bundles

We are finally ready to define curvature on vector bundles. At this point, we've set things up to be a fairly straightforward generalization from what we've presumably already seen. First we look back at the more familiar case, where $V = TM$. After introducing the Christoffel Connection, we parallel-transport a vector along two paths. Comparing V_{path1} V_{path2} gives us the curvature. In component form, the Riemann curvature can be expressed in terms of the Christoffel connection:

$$R^i_{jkl} = \partial_k \Gamma^i_{jl} + \partial_l \Gamma^i_{jk} + \Gamma^i_{km} \Gamma^m_{jl} - \Gamma^i_{lm} \Gamma^m_{jk} \tag{4.39}$$

Since R is antisymmetric in its last two indices, we can think of this portion of the tensor as a two-form:

$$R = \frac{1}{2} R^\alpha_{\beta\gamma\delta} dx^\beta \wedge dx^\gamma = \text{an } n \times n \text{ matrix of two-forms} \tag{4.40}$$

Problem 4.7 This viewpoint becomes very natural when noting the formula for R is much simpler. Show:

$$R = d\omega - \omega \wedge \omega \tag{4.41}$$

Where we are representing the Christoffel connection with our more general form of the connection, a lie-algebra-valued one-form. (Remember that since ω is a matrix of one-forms, the wedge product implies matrix multiplication as well, meaning it is a nontrivial combination of one-forms; if the matrix size were 1×1 , ω would just be a single one-form, and $\omega \wedge \omega = 0$, by antisymmetry of the wedge product.)




Figure 4.8 Curvature of a general vector bundle is the difference found by transporting a vector along two different paths.

The connection ω is a matrix of one-forms, and the curvature $R = d\omega - \omega \wedge \omega$ is a matrix of two-forms. We generalize this to the language of vector bundles, basically by replacing the letter R with the symbol Ω . The curvature two-form on a vector bundle V is given by:

$$\Omega = \frac{1}{2} \Omega_{\beta ij}^{\alpha} dx^i \wedge dx^j = d\omega - \omega \wedge \omega \tag{4.42}$$

Clearly this reduces to the Riemann curvature tensor for $V = TM$. Note that α and β run from one to k, and i and j run from one to n. In other words, these indices refer refer to vectors in different spaces; α and β represent vectors in the fiber, and i and j represent tangent vectors. Since Ω is a tensor, it is possible to view it as a multilinear map

$$\Omega: T_p M \otimes T_p M \otimes F_p M \rightarrow F_p M \tag{4.43}$$

It rests on basically the same concept as the Riemann curvature; we transport a vector in the fiber through the manifold on two different paths, specified by two tangent vectors. The difference in the result of the two transports will be found by contracting Ω with the two directions and the original vector.

Problem 4.8 When we perform a change of basis in the fiber given by (4.33), show that the curvature transforms in the following manner:

$$\tilde{\Omega} = \Lambda^{-1} \cdot \Omega \cdot \Lambda \tag{4.44}$$

That is, it transforms as a (1,1) tensor in $F_p^* M \otimes F_p M$.

You may be wondering why we care so much how things transform under changes of fiber bases. We'd like to be able to construct quantities whose components transform, but which can be identified with invariant objects. Currently, the curvature transforms like a (1,1) fiber-valued tensor, and thus we can identify it as such, but the transformation law (4.34) for the connection doesn't lend itself to anything of this nature. To put it more plainly, ω 's components are one-forms over the base manifold, but they are not one-forms over the total space. When we deal with principal bundles, we will be able to define an invariant connection which will indeed be associated with a one-form on the total space, but for now we have to settle with this half-defined object.

Problem 4.9 Recall that the Λ 's all sit in the bundle's structure group, G. If Λ can be found by taking the exponential of a lie algebra element, κ ,

$$\Lambda_{\beta}^{\alpha} = [\exp\{\kappa\}]_{\beta}^{\alpha} \tag{4.45}$$

Show that the transformation law for the connection reduces to

$$\tilde{\omega} = \exp\{\kappa\} \cdot \omega \cdot \exp\{-\kappa\} + d\kappa \tag{4.46}$$

Note, for small values of κ , this looks like

$$\tilde{\omega} = \omega + d\kappa + [\kappa, \omega] \tag{4.47}$$

highlighting the fact that ω and κ are both lie algebra elements, and that the lie algebra is closed under addition and commutation.

4.7 Parallel Transport on General Fiber Bundles

In the case of a general fiber bundle, we are no longer able to abstract the tangent

bundle's vector language any further, but we can still perform many of the same calculations, though a bit less concretely. We cannot abstract the quantitative notion of curvature to a general fiber bundle, but we *can* do so the special case of principal bundles. We start by revisiting parallel transport, for the case of a general fiber bundle.

The need for a specification of parallel transport is built into the language of fiber bundles. The projection map $\pi: E \rightarrow M$ allows us to definitively say over which point we sit in the base, but we cannot say definitively at which point we sit in the fiber. F provides a local coordinate system for points in the neighborhood of p , but these coordinates in themselves have no substantial meaning, as we can change them at will, thus they can tell us nothing about parallel transport. We need to provide additional information which tells us whether we are moving up or down the fiber as we move along a path in E . This additional data is our final notion of a "connection".

Start with a fiber bundle $\pi: E \rightarrow M$, with fiber F . Imagine studying the tangent bundle of the total space, TE . This is a fairly complex object; it's the tangent bundle of the total space of a fiber bundle. It's worth taking a moment to get your head around this concept. Now, imagine we look at a specific tangent space $T_u E$ at an arbitrary point $u \in E$. Further imagine that we have a means of decomposing the vector space $T_u E$ into a combination of two vector spaces:

$$T_u E = V_u E \oplus H_u E \tag{4.48}$$

where $V_u E \cong T_p F$ is the "vertical" subspace, consisting of directions along the fiber F , and $H_u E \cong T_p M$ is the "horizontal" subspace, providing directions along the base M (for some f and p corresponding to the point u in E).




Figure 4.9 The specification of a horizontal subspace defines horizontal transport on a general fiber bundle.

Clearly we haven't been very rigorous, but the idea here is that we are providing a definitive split between “horizontal” and “vertical” transport. Given a curve γ through E , we say that $\gamma(t)$ is parallel-transported if the velocity sits in the horizontal subspace:

$$\frac{d\gamma}{dt} \in H_{\gamma(t)}E \quad (4.49)$$

Thus, if we can provide data which specifies a way of splitting $T_uE = V_uE \oplus H_uE$, we have defined parallel transport in the fiber bundle.

The Connection on a Principal Bundle

In the special case that π is a principal bundle, $P \rightarrow M$, with fiber G given by the structure group, this additional data can, in fact, be expressed as a lie-algebra-valued one-form over the total space, $\omega \in \mathfrak{g} \otimes T^*P$. Note that we previously defined an ω which lived in the cotangent space of the base manifold, T^*M , but under the condition that its components transformed under a change of fiber basis. This new ω lives in the cotangent space of the total space, $\omega \in T^*P$.

To find the appropriate lie-algebra-valued one-form, we first choose bundle coordinates (x, g) where $x \in M$, $g \in G$. In these coordinates, any one-form can be expressed as:

$$\omega = \alpha_\mu dx^\mu + \beta_{ij} dg^{ij} \quad (4.50)$$

Problem 4.10 We can narrow down the form of our connection by imposing a basic requirement. When the connection acts on a tangent vector that's just moving vertically along the fiber, we want it to return the value of the lie algebra element associated with it. This restricts the set of lie-algebra-valued one-forms we can choose from. A path passing through (x, g) whose tangent vector is the lie algebra element A can be written

$$P(t) = (x, g \cdot e^{At}) \quad (4.51)$$

Thus we have the following requirement:

$$\omega\left(\frac{dP}{dt}\right) = A \quad (4.52)$$

Show that this implies:

$$\beta_{ij} = g_{ij}^{-1} \quad (4.53)$$

Thus we can always find parameterization of the fiber such that ω has the following form:

$$\omega = \alpha_\mu dx^\mu + g_{ij}^{-1} dg^{ij} \quad (4.54)$$

α_μ contains all of the information about horizontal transport. If we rewrite this expression, we can make contact with our earlier definitions for the connection:

$$\omega = g^{-1} \cdot A_\mu \cdot g dx^\mu + g_{ij}^{-1} dg^{ij} \quad (4.55)$$

where A is in the lie algebra of G .

We treat A_μ like our old definition of a connection. It transforms annoyingly under fiber reparameterizations (4.34), but the expression (4.55) has just the right form to cancel terms picked up by A 's unusual transformation law. We change our notation again, redefining what we mean by the connection ω :

$$\omega = g^{-1} \cdot A \cdot g + g_{ij}^{-1} dg^{ij}, \text{ where} \quad (4.56)$$

$$A = A_{\mu}^a(x) T^a dx^{\mu} \quad (4.57)$$

Here, $\{T^a\}$ is a basis for the lie algebra, and $\{A_{\mu}^a\}$ is a collection of coefficients for the $\{T^a\}$.

We will now show that ω so defined provides a split of $T_u P$ into $V_u P \oplus H_u P$.

First, we write down a basis for $H_u P$:

$$\left\{ \frac{\partial}{\partial x^{\mu}} + C_{\mu}^{ij} \frac{\partial}{\partial g^{ij}} \right\} \in H_u P \quad (4.58)$$

The $\{\partial/\partial x^{\mu}\}$ are, of course, directions in the base manifold, while the $\{\partial/\partial g^{ij}\}$ correspond to directions along the fiber (in a given parameterization). $\{C_{\mu}^{ij}\}$ are coefficients specified as input, selecting a definition of "horizontal". Given a lie-algebra-valued one-form ω , we can determine a set of C 's, via the following equation:

$$H_u P = \{V \in T_u P \mid \omega(V) = 0\}$$

Problem 4.11 We write a general vector $V \in T_u P$ as

$$V = \left(\alpha^{ij} \frac{\partial}{\partial g^{ij}} \right) + \beta^{\mu} \left(\frac{\partial}{\partial x^{\mu}} + C_{\mu}^{ij} \frac{\partial}{\partial g^{ij}} \right) \quad (4.59)$$

When $\alpha_{ij} = 0$, $V \in H_u P$, meaning

$$\alpha^{ij} = 0 \Rightarrow \omega(V) = 0 \text{ for all } \{\beta^{\mu}\} \quad (4.60)$$

Use (4.56) and (4.57) to compute the coefficients fixing the horizontal subspace:

$$C_{\mu}^{ij} = -A_{\mu}^a(x) [T^a]^{ik} g^{kj} \quad (4.61)$$

The $\{A_{\mu}^a(x)\}$ is data supplied by ω , which tells us the C_{μ}^{ij} , which nails down the horizontal subspace $H_u P$.

$$H_u P = \left\{ \beta^{\mu} \left(\frac{\partial}{\partial x^{\mu}} - A_{\mu}^a(x) [T^a]^{ik} g^{kj} \frac{\partial}{\partial g^{ij}} \right) \right\} \quad (4.62)$$

Parallel Transport on a Principal Bundle

It will now be possible to write down an equation for parallel transport in a principal bundle. Let $\Upsilon(t)$ be a path in the base M .

$$\Upsilon(t) = (x^1(t), x^2(t), \dots, x^n(t)) \quad (4.63)$$

we want to perform parallel translation in the bundle, lifting $\Upsilon(t): \mathbb{R} \rightarrow M$ to $\Gamma(t): \mathbb{R} \rightarrow P$, where $\Gamma(t)$ always lies over the point $\Upsilon(t)$ in the base manifold:

$$\pi \circ \Gamma = \Upsilon \quad (4.64)$$

Parallel transport requires that the velocity of Γ lives in the horizontal subspace:

$$\frac{d\Gamma}{dt} \in H_{\Gamma(t)} P, \text{ for all } t \quad (4.65)$$

More explicitly,

$$\frac{d}{dt} = \frac{dx^{\mu}}{dt} \frac{\partial}{\partial x^{\mu}} + \frac{dg^{ij}}{dt} \frac{\partial}{\partial g^{ij}} = \beta^{\mu} \left(\frac{\partial}{\partial x^{\mu}} - A_{\mu}^a(x) [T^a \cdot g]^{ij} \frac{\partial}{\partial g^{ij}} \right) \quad (4.66)$$

Problem 4.12 Use equation (4.66) to derive the equation of motion for $g(t)$:

$$\frac{dg^{ij}}{dt} + \frac{dx^\mu}{dt} A_\mu^a(x) [T^a \cdot g]^{ij} = 0 \quad (4.67)$$

Given a connection ω specified by A , we simply solve this equation for $g(t)$ to determine parallel transport on a principal bundle. How now do we determine a rule for assigning *curvature* to a principal bundle? We cannot simply proceed by analogy, for there is no obvious way of comparing group elements, as there is for vectors. In order to motivate our search for curvature on principal bundles, we take a second look at vector bundles, now in a more specific context.

4.8 Connections on Associated Vector Bundles

Recall that given a principal bundle, we can build associated vector bundles, when given a representation of the structure group. It should come as no surprise that a connection on a principal bundle induces vector-bundle-connections on all of its associated vector bundles. Given an associated vector bundle with representation ρ , we can perform parallel transport on a vector V at a point p along a curve $\Upsilon(t)$ in the base by the following:

Look at the point $(p = \Upsilon(0), e)$ in the principal bundle over M , where e is the identity element of G . Using the connection ω , parallel-transport the identity element along the curve $\Upsilon(t) \subset M$, producing the curve $\Gamma(t) \subset P$. Act with the representation of the parallel-transported group element on the vector V_p to get the vector $V_{\Upsilon(t)}$, the parallel-transported vector in the vector bundle. Explicitly,

$$V_{\Upsilon(t)} = \rho(\Gamma(t)) \cdot V_{\Upsilon(0)} \quad (4.68)$$

Problem 4.13 Recall that we already have an equation of motion (4.67) for $\Gamma(t) = g(t)$, so these two formulas now yield an equation of motion for V :

$$\frac{dV}{dt} = - \frac{dx^\mu}{dt} A_\mu^a(x) \rho(T^a) \cdot V \quad (4.69)$$

Remember that a representation of a lie group induces a representation on its lie algebra. This is what is meant by $\rho(T^a)$.

We can rewrite this in terms of covariant differentiation, by

$$\frac{dV}{dt} = \frac{dx^\mu}{dt} \frac{\partial V}{\partial x^\mu} \quad (4.70)$$

so that (4.69) takes the following form:

$$\left(\frac{\partial}{\partial x^\mu} + A_\mu^a(x) \rho(T^a) \right) V = 0 \quad (4.71)$$

This is simply the covariant derivative of V . The connection ω on a principal bundle induces a connection $A = A^a(x) \rho(T^a)$ on its associated vector bundles. This is the same kind of connection we previously defined for vector bundles. The curvature of an associated vector bundle can be directly given by

$$F = dA - A \wedge A \quad (4.72)$$

When we constructed ω , it had a piece corresponding to A , which by itself was not a one-form on the total space, but it was a one-form defined on the base with a connection-like transformation law. As it turns out, this piece *is* the connection over the associated vector bundles of the principal bundle, when evaluated in the given representation.




Figure 4.10 A connection on a principal bundle induces connections on all its associated vector bundles.

Curvature on Principal Bundles

Finally, we are in a position to define curvature on principal bundles. Curvature over a principal bundle will be associated with the curvature that it induces over its associated vector bundles. We state the answer and show that it is consistent:

$$\Omega = d\omega - \omega \wedge \omega \quad (4.73)$$

where now ω is a lie-algebra-valued one-form over the total space, which can be expressed in a given fiber parameterization like so:

$$\omega = g^{-1} A g + g^{-1} dg \quad (4.74)$$

Problem 4.14 Show, in this fiber parameterization, the curvature is given by the following:

$$\Omega = g^{-1} (dA - A \wedge A) g \quad (4.75)$$

Remembering that the curvature of a vector bundle transforms like (4.44), the result of problem 4.14 is motivation enough to define the curvature of a principal bundle in this way. To summarize this picture, the curvature over a principal bundle is found by the connection form ω , a lie-algebra-valued one-form over the total space, P . The curvature Ω is a lie-algebra-valued two-form. For each representation of the structure group G , there exists

an associated vector bundle, with vector bundle connection given by A , which is a lie-algebra-valued one-form over the base, and curvature $F = dA - A \wedge A$, a lie-algebra-valued two-form over the base, and $(1,1)$ tensor over the vector space fiber.