Lecture 1: Connections on principal fibre bundles

The beauty and profundity of the geometry of fibre bundles were to a large extent brought forth by the (early) work of Chern. I must admit, however, that the appreciation of this beauty came to physicists only in recent years.
— CN Yang, 1979

In this lecture we introduce the notion of a principal fibre bundle and of a connection on it, but we start with some motivation.

1.1 Motivation: the Dirac monopole

It is only fitting to start this course, which takes place at the JCMB, with a solution of Maxwell's equations. The magnetic field $\mathbf{B}$ of a magnetic monopole sitting at the origin in $\mathbb{R}^3$ is given by

$$\mathbf{B}(\boldsymbol{x}) = \frac{\boldsymbol{x}}{4\pi r^3} ,$$

where $r = |\boldsymbol{x}|$. This satisfies $\operatorname{div}\mathbf{B} = 0$ and hence is a solution of Maxwell's equations in $\mathbb{R}^3 \setminus \{\mathbf{0}\}$.

> Maxwell's equations are linear and hence this solution requires a source of magnetic field, namely the monopole which sits at the origin. We will see later in this course that there are other monopoles which do not require sources and which extend smoothly to the 'origin'.

In modern language, the vector field $\mathbf{B}$ is understood as the 2-form $\mathrm{F} \in \Omega^2(\mathbb{R}^3 \setminus \{\mathbf{0}\})$ given by

$$\mathrm{F} = \frac{1}{2\pi r^3}\left(x_1\, dx_2 \wedge dx_3 + \text{cyclic}\right)$$

in the cartesian coordinates of $\mathbb{R}^3$. Maxwell's equations say that $d\mathrm{F} = 0$. This is perhaps more evident in spherical coordinates ($x_1 = r\sin\theta\cos\phi, \ldots$), where

$$\mathrm{F} = \frac{1}{4\pi}\sin\theta\, d\theta \wedge d\phi .$$

Since $d\mathrm{F} = 0$ we may hope to find a one-form A such that $\mathrm{F} = d\mathrm{A}$. For example,

$$\mathrm{A} = -\frac{1}{4\pi}\cos\theta\, d\phi .$$

This is not regular over all of $\mathbb{R}^3 \setminus \{\mathbf{0}\}$, however. This should not come as a surprise, since after all spherical coordinates are singular on the x_3-axis. Rewriting A in cartesian coordinates

$$\mathrm{A} = \frac{x_3}{4\pi r}\,\frac{x_2\, dx_1 - x_1\, dx_2}{x_1^2 + x_2^2} ,$$

clearly displays the singularity at $x_1 = x_2 = 0$.

> In the old Physics literature this singularity is known as the "Dirac string," a language we shall distance ourselves from in this course.

This singularity will afflict any A we come up with. Indeed, notice that F restricts to the (normalised) area form on the unit sphere $\mathrm{S} \subset \mathbb{R}^3$, whence

$$\int_{\mathrm{S}} \mathrm{F} = 1 .$$

If $\mathrm{F} = d\mathrm{A}$ for a smooth one-form A, then Stokes's theorem would have forced $\int_{\mathrm{S}} \mathrm{F} = 0$.

The principal aim of the first couple of lectures is to develop the geometric framework to which F (and A) belong: the theory of connections on principal fibre bundles, to which we now turn.

1.2 Principal fibre bundles

A **principal fibre bundle** consists of the following data:

- a manifold P, called the **total space**;
- a Lie group G acting freely on P on the right:

$$\begin{aligned} P \times G &\to P \\ (p, g) &\mapsto pg \quad \text{(or sometimes } R_g p) \end{aligned}$$

  where by a free action we mean that the stabilizer of every point is trivial, or paraphrasing, that every element G (except the identity) moves every point in P. We will also assume that the space of orbits $M = P/G$ is a manifold (called the **base**) and the natural map $\pi : P \to M$ taking a point to its orbit is a smooth surjection. For every $m \in M$, the submanifold $\pi^{-1}(m) \subset P$ is called the **fibre over** m.

Further, this data will be subject to the condition of **local triviality**: that M admits an open cover $\{U_\alpha\}$ and G-equivariant diffeomorphisms $\psi_\alpha : \pi^{-1}U_\alpha \to U_\alpha \times G$ such that the following diagram commutes

$$\begin{tikzcd} \pi^{-1}U_\alpha \arrow[rr, "\psi_\alpha"] \arrow[dr, "\pi"'] & & U_\alpha \times G \arrow[dl, "pr_1"] \\ & U_\alpha & \end{tikzcd}$$

This means that $\psi_\alpha(p) = (\pi(p), g_\alpha(p))$, for some G-equivariant map $g_\alpha : \pi^{-1}U_\alpha \to G$ which is a fibrewise diffeomorphism. Equivariance means that

$$g_\alpha(pg) = g_\alpha(p)g \; .$$

One often abbreviates the above data by saying that

$$\begin{tikzcd} P \arrow[d, "\pi"] \\ M \end{tikzcd} \quad \text{or} \quad \begin{tikzcd} G \arrow[r] & P \arrow[d, "\pi"] \\ & M \end{tikzcd} \quad \text{is a principal G-bundle,}$$

but be aware that abuses of language are rife in this topic. Don't be surprised by statements such as "Let P be a principal bundle..."

We say that the bundle is **trivial** if there exists a diffeomorphism $\Psi : P \to M \times G$ such that $\Psi(p) = (\pi(p), \psi(p))$ and such that $\psi(pg) = \psi(p)g$. This last condition is simply the G-equivariance of Ψ.

A **section** is a (smooth) map $s : M \to P$ such that $\pi \circ s = \text{id}$. In other words, it is a smooth assignment to each point in the base of a point in the fibre over it. Sections are rare. Indeed, one has

Done? ❑ **Exercise 1.1.** Show that a principal fibre bundle admits a section if and only if it is trivial. (This is in sharp contrast with, say, vector bundles, which always admit sections.)

Nevertheless, since P is locally trivial, local sections do exist. In fact, there are local sections $s_\alpha : U_\alpha \to \pi^{-1}U_\alpha$ canonically associated to the trivialization, defined so that for every $m \in U_\alpha$, $\psi_\alpha(s_\alpha(m)) = (m, e)$, where $e \in G$ is the identity element. In other words, $g_\alpha \circ s_\alpha : U_\alpha \to G$ is the constant function sending every point to the identity. Conversely, a local section s_α allows us to identify the fibre over m with G. Indeed, given any $p \in \pi^{-1}(m)$, there is a unique group element $g_\alpha(p) \in G$ such that $p = s_\alpha(m)g_\alpha(p)$.

On nonempty overlaps $U_{\alpha\beta} := U_\alpha \cap U_\beta$, we have two ways of trivialising the bundle:

$$\begin{tikzcd} U_{\alpha\beta} \times G \arrow[dr, "pr_1"'] & \pi^{-1}U_{\alpha\beta} \arrow[l, "\psi_\beta"'] \arrow[r, "\psi_\alpha"] \arrow[d, "\pi"] & U_{\alpha\beta} \times G \arrow[dl, "pr_1"] \\ & U_{\alpha\beta} & \end{tikzcd}$$

For $m \in U_{\alpha\beta}$ and $p \in \pi^{-1}(m)$, we have $\psi_\alpha(p) = (m, g_\alpha(p))$ and $\psi_\beta(p) = (m, g_\beta(p))$, whence there is $\bar{g}_{\alpha\beta}(p) \in G$ such that $g_\alpha(p) = \bar{g}_{\alpha\beta}(p) g_\beta(p)$. In other words,

$$\bar{g}_{\alpha\beta}(p) = g_\alpha(p) g_\beta(p)^{-1} . \tag{1}$$

In fact, $\bar{g}_{\alpha\beta}(p)$ is constant on each fibre:

$$\begin{aligned} \bar{g}_{\alpha\beta}(pg) &= g_\alpha(pg) g_\beta(pg)^{-1} \\ &= g_\alpha(p) g g^{-1} g_\beta(p) && \text{(by equivariance of } g_\alpha, g_\beta) \\ &= \bar{g}_{\alpha\beta}(p) . \end{aligned}$$

In other words, $\bar{g}_{\alpha\beta}(p) = g_{\alpha\beta}(\pi(p))$ for some function

$$g_{\alpha\beta} : U_{\alpha\beta} \to G .$$

From equation (1), it follows that these **transition functions** obey the following **cocycle conditions**:

$$g_{\alpha\beta}(m) g_{\beta\alpha}(m) = e \qquad \text{for every } m \in U_{\alpha\beta}, \text{ and} \tag{2}$$

$$g_{\alpha\beta}(m) g_{\beta\gamma}(m) g_{\gamma\alpha}(m) = e \qquad \text{for every } m \in U_{\alpha\beta\gamma}, \tag{3}$$

where $U_{\alpha\beta\gamma} = U_\alpha \cap U_\beta \cap U_\gamma$.

Done? ❑ **Exercise 1.2.** Show that on double overlaps, the canonical sections s_α are related by

$$s_\beta(m) = s_\alpha(m) g_{\alpha\beta}(m) \qquad \text{for every } m \in U_{\alpha\beta}.$$

One can reconstruct the bundle from an open cover $\{U_\alpha\}$ and transition functions $\{g_{\alpha\beta}\}$ obeying the cocycle conditions (2) as follows:

$$P = \bigsqcup_\alpha (U_\alpha \times G) \Big/ \sim ,$$

where $(m, g) \sim (m, g_{\alpha\beta}(m) g)$ for all $m \in U_{\alpha\beta}$ and $g \in G$. Notice that π is induced by the projection onto the first factor and the action of G on P is induced by right multiplication on G, both of which are preserved by the equivalence relation, which uses *left* multiplication by the transition functions. (Associativity of group multiplication guarantees that right and left multiplications commute.)

Example 1.1 (Möbius band)**.** The boundary of the Möbius band is an example of a nontrivial principal $\mathbb{Z}_2$-bundle. This can be described as follows. Let $S^1 \subset \mathbb{C}$ denote the complex numbers of unit modulus and let $\pi : S^1 \to S^1$ be the map defined by $z \mapsto z^2$. Then the fibre $\pi^{-1}(z^2) = \{\pm z\}$ consists of two points. A global section would correspond to choosing a square-root function smoothly on the unit circle. This does not exist, however, since any definition of $z^{1/2}$ always has a branch cut from the origin out to the point at infinity. Therefore the bundle is not trivial. In fact, if the bundle were trivial, the total space would be disconnected, being two disjoint copies of the circle. However building a paper model of the Möbius band one quickly sees that its boundary is connected.

We can understand this bundle in terms of the local data as follows. Cover the circle by two overlapping open sets: U_1 and U_2. Their intersection is the disjoint union of two intervals in the circle: $V_1 \sqcup V_2$. Let $g_i : V_i \to \mathbb{Z}_2$ denote the transition functions, which are actually constant since V_i are connected and $\mathbb{Z}_2$ is discrete, so we can think of $g_i \in \mathbb{Z}_2$. There are no triple overlaps, so the cocycle condition is vacuously satisfied. It is an easy exercise to check that the resulting bundle is trivial if and only if $g_1 = g_2$ and nontrivial otherwise.

1.3 Connections

The push-forward and the pull-back

Let $f : M \to N$ be a smooth map between manifolds. The **push-forward**

$$f_* : TM \to TN \qquad \text{(also written } Tf \text{ when in a categorical mood)}$$

is the collection of fibre-wise linear maps $f_* : T_m M \to T_{f(m)} N$ defined as follows. Let $v \in T_m M$ be represented as the velocity of a curve $t \mapsto \gamma(t)$ through m; that is, $\gamma(0) = m$ and $\gamma'(0) = v$. Then $f_*(v) \in T_{f(m)} N$ is the velocity at $f(m)$ of the curve $t \mapsto f(\gamma(t))$; that is, $f_* \gamma'(0) = (f \circ \gamma)'(0)$. If $g : N \to Q$ is another smooth map between manifolds, then so is their composition $g \circ f : M \to Q$. The **chain rule** is then simply the "functoriality of the push-forward": $(g \circ f)_* = g_* \circ f_*$. Dual to the push-forward, there is the **pull-back** $f^* : T^*N \to T^*M$, defined for a one-form α by $(f^*\alpha)(v) = \alpha(f_* v)$. It is also functorial, but now reversing the order $(g \circ f)^* = f^* \circ g^*$.

Let $\pi : P \to M$ be a principal G-bundle and let $m \in M$ and $p \in \pi^{-1}(m)$. The **vertical** subspace $V_p \subset T_p P$ consists of those vectors tangent to the fibre at p; in other words, $V_p = \ker \pi_* : T_p P \to T_m M$. A vector field $v \in \mathscr{X}(P)$ is **vertical** if $v(p) \in V_p$ for all p. The Lie bracket of two vertical vector fields is again vertical. The vertical subspaces define a G-invariant distribution $V \subset TP$: indeed, since $\pi \circ R_g = \pi$, we have that $(R_g)_* V_p = V_{pg}$. In the absence of any extra structure, there is no natural complement to V_p in $T_p P$. This is in a sense what a connection provides.

1.3.1 Connections as horizontal distributions

A **connection** on P is a smooth choice of **horizontal** subspaces $H_p \subset T_p P$ complementary to V_p:

$$T_p P = V_p \oplus H_p$$

and such that $(R_g)_* H_p = H_{pg}$. In other words, a connection is a G-invariant distribution $H \subset TP$ complementary to V.

For example, a G-invariant riemannian metric on P gives rise to a connection, simply by defining $H_p = V_p^\perp$. This simple observation underlies the Kałuża–Klein programme relating gravity on P to gauge theory on M. It also underlies many geometric constructions, since it is often the case that 'nice' metrics will give rise to 'nice' connections and viceversa.

We will give two more characterisations of connections on P, but first, a little revision.

Some Lie group technology

A Lie group is a manifold with two smooth operations: a multiplication $G \times G \to G$, and an inverse $G \to G$ obeying the group axioms. If $g \in G$ we define diffeomorphisms

$$\begin{aligned} L_g : G &\to G \\ x &\mapsto gx \end{aligned} \qquad \text{and} \qquad \begin{aligned} R_g : G &\to G \\ x &\mapsto xg \end{aligned}$$

called **left** and **right** multiplications by g, respectively.
A vector field $v \in \mathscr{X}(G)$ is **left-invariant** if $(L_g)_* v = v$ for all $g \in G$. In other words, $v(g) = (L_g)_* v(e)$ for all $g \in G$, where e is the identity. The Lie bracket of two left-invariant vector fields is left-invariant. The vector space of left-invariant vector fields defines the **Lie algebra** $\mathfrak{g}$ of G. A left-invariant vector field is uniquely determined by its value at the identity, whence $\mathfrak{g} \cong T_e G$.
The **(left-invariant) Maurer–Cartan form** is the $\mathfrak{g}$-valued 1-form θ on G defined by

$$\theta_g = (L_{g^{-1}})_* : T_g G \to T_e G = \mathfrak{g} .$$

If v is a left-invariant vector field, then $\theta(v) = v(e)$, whence θ_e is the natural identification between $T_e G$ and $\mathfrak{g}$. For a matrix group, $\theta_g = g^{-1} dg$, from where it follows that θ is left-invariant and satisfies the **structure equation**:

$$d\theta = -\tfrac{1}{2}[\theta, \theta] ,$$

where the bracket in the RHS denotes both the Lie bracket in $\mathfrak{g}$ and the wedge product of 1-forms.
Every $g \in G$ defines a smooth map $\mathrm{Ad}_g : G \to G$ by $\mathrm{Ad}_g = L_g \circ R_g^{-1}$; that is,

$$\mathrm{Ad}_g\, h = g h g^{-1} .$$

This map preserves the identity, whence its derivative there defines a linear representation of the group on the Lie algebra known as the **adjoint representation** $\mathrm{ad}_g := (Ad_g)_* : \mathfrak{g} \to \mathfrak{g}$, defined explicitly by

$$\mathrm{ad}_g X = \frac{d}{dt}\left(g e^{tX} g^{-1}\right)\Big|_{t=0} .$$

For G a matrix group, $\mathrm{ad}_g(X) = gXg^{-1}$. Finally, notice that $R_g^* \theta = \mathrm{ad}_{g^{-1}} \circ \theta$.

The action of G on P defines a map $\sigma : \mathfrak{g} \to \mathscr{X}(P)$ assigning to every $X \in \mathfrak{g}$, the **fundamental vector field** $\sigma(X)$ whose value at p is given by

$$\sigma_p(X) = \frac{d}{dt}\left(p e^{tX}\right)\Big|_{t=0} .$$

Notice that

$$\pi_* \sigma_p(X) = \frac{d}{dt}\pi\left(p e^{tX}\right)\Big|_{t=0} = \frac{d}{dt}\pi\left(p\right)\Big|_{t=0} = 0 ,$$

whence $\sigma(X)$ is a vertical vector field. In fact, since G acts freely, the map $X \mapsto \sigma_p(X)$ is an isomorphism $\sigma_p : \mathfrak{g} \xrightarrow{\cong} V_p$ for every p.

Lemma 1.1.

$$(R_g)_* \sigma(X) = \sigma\left(\mathrm{ad}_{g^{-1}} X\right) .$$

Proof. By definition, at $p \in \mathrm{P}$, we have

$$\begin{aligned}(\mathrm{R}_g)_* \sigma_p(\mathrm{X}) &= \frac{d}{dt} \mathrm{R}_g \left(p e^{t\mathrm{X}}\right)\Big|_{t=0} \\ &= \frac{d}{dt} \left(p e^{t\mathrm{X}} g\right)\Big|_{t=0} \\ &= \frac{d}{dt} \left(p g g^{-1} e^{t\mathrm{X}} g\right)\Big|_{t=0} \\ &= \frac{d}{dt} \left(p g e^{t \operatorname{ad}_{g^{-1}} \mathrm{X}}\right)\Big|_{t=0} \\ &= \sigma_{pg}\left(\operatorname{ad}_{g^{-1}} \mathrm{X}\right) .\end{aligned}$$

□

Done? ❑ **Exercise 1.3.** Let $g_\alpha : \pi^{-1}\mathrm{U}_\alpha \to \mathrm{G}$ be the maps defined by the local trivialisation. Show that $(g_\alpha)_* \sigma_p(\mathrm{X}) = (\mathrm{L}_{g_\alpha(p)})_* \mathrm{X}$.

1.3.2 The connection one-form

The horizontal subspace $\mathrm{H}_p \subset \mathrm{T}_p\mathrm{P}$, being a linear subspace, is cut out by $k = \dim \mathrm{G}$ linear equations $\mathrm{T}_p\mathrm{P} \to \mathbb{R}$. In other words, H_p is the kernel of k one-forms at p, the components of a one-form ω at p with values in a k-dimensional vector space. There is a natural such vector space, namely the Lie algebra $\mathfrak{g}$ of G, and since ω annhilates horizontal vectors it is defined by what it does to the vertical vectors, and we do have a natural map $\mathrm{V}_p \to \mathfrak{g}$ given by the inverse of σ_p. This prompts the following definition.

The **connection one-form** of a connection $\mathrm{H} \subset \mathrm{TP}$ is the $\mathfrak{g}$-valued one-form $\omega \in \Omega^1(\mathrm{P};\mathfrak{g})$ defined by

$$\omega(v) = \begin{cases} \mathrm{X} & \text{if } v = \sigma(\mathrm{X}) \\ 0 & \text{if } v \text{ is horizontal.} \end{cases}$$

Proposition 1.2. *The connection one-form obeys*

$$\mathrm{R}_g^* \omega = \operatorname{ad}_{g^{-1}} \circ \omega .$$

Proof. Let $v \in \mathrm{H}_p$, so that $\omega(v) = 0$. By the G-invariance of H, $(\mathrm{R}_g)_* v \in \mathrm{H}_{pg}$, whence $\mathrm{R}_g^*\omega$ also annhilates v and the identity is trivially satisfied. Now let $v = \sigma_p(\mathrm{X})$ for some $\mathrm{X} \in \mathfrak{g}$. Then, using Lemma 1.1,

$$\mathrm{R}_g^* \omega(\sigma(\mathrm{X})) = \omega\left((\mathrm{R}_g)_* \sigma(\mathrm{X})\right) = \omega\left(\sigma\left(\operatorname{ad}_{g^{-1}} \mathrm{X}\right)\right) = \operatorname{ad}_{g^{-1}} \mathrm{X} .$$

□

Conversely, given a one-form $\omega \in \Omega^1(\mathrm{P};\mathfrak{g})$ satisfying the identity in Proposition 1.2 and such that $\omega(\sigma(\mathrm{X})) = \mathrm{X}$, the distribution $\mathrm{H} = \ker \omega$ defines a connection on P.

We say that a form on P is **horizontal** if it annihilates the vertical vectors. Notice that if ω and ω' are connection one-forms for two connections H and H′ on P, their difference $\omega - \omega' \in \Omega^1(\mathrm{P};\mathfrak{g})$ is horizontal. We will see later that this means that it defines a section through a bundle on M associated to P.

1.3.3 Gauge fields

Finally, as advertised, we make contact with the more familiar notion of gauge fields as used in Physics, which live on M instead of P.

Recall that we have local sections $s_\alpha : \mathrm{U}_\alpha \to \pi^{-1}\mathrm{U}_\alpha$ associated canonically to the trivialisation of the bundle, along which we can pull-back the connection one-form ω, defining in the process the following $\mathfrak{g}$-valued one-forms on U_α:

$$\mathrm{A}_\alpha := s_\alpha^* \omega \in \Omega^1(\mathrm{U}_\alpha;\mathfrak{g}) .$$

Proposition 1.3. *The restriction of the connection one-form* ω *to* $\pi^{-1}U_\alpha$ *agrees with*

$$\omega_\alpha = \mathrm{ad}_{g_\alpha^{-1}} \circ \pi^* A_\alpha + g_\alpha^* \theta ,$$

where θ *is the Maurer–Cartan one-form.*

Proof. We will prove this result in two steps.

1. First we show that ω_α and ω agree on the image of s_α. Indeed, let $m \in U_\alpha$ and $p = s_\alpha(m)$. We have a direct sum decomposition
$$T_p P = \mathrm{im}(s_\alpha \circ \pi)_* \oplus V_p ,$$
so that every $v \in T_p P$ can be written uniquely as $v = (s_\alpha)_* \pi_*(v) + \bar{v}$, for a unique vertical vector $\bar{v}$. Applying ω_α on v, we obtain (since $g_\alpha(s_\alpha(m)) = e$)
$$\begin{aligned} \omega_\alpha(v) &= \left(\pi^* s_\alpha^* \omega\right)(v) + \left(g_\alpha^* \theta_e\right)(v) \\ &= \omega\left((s_\alpha)_* \pi_* v\right) + \theta_e\left((g_\alpha)_* v\right) \\ \text{(since } (g_\alpha \circ s_\alpha)_* = 0\text{)} \qquad &= \omega\left((s_\alpha)_* \pi_* v\right) + \theta_e\left((g_\alpha)_* \bar{v}\right) \\ &= \omega\left((s_\alpha)_* \pi_* v\right) + \omega(\bar{v}) \\ &= \omega(v) . \end{aligned}$$

2. Next we show that they transform in the same way under the right action of G:
$$\begin{aligned} R_g^*(\omega_\alpha)_{pg} &= \mathrm{ad}_{g_\alpha(pg)^{-1}} \circ R_g^* \pi^* s_\alpha^* \omega + R_g^* g_\alpha^* \theta \\ \text{(equivariance of } g_\alpha\text{)} \qquad &= \mathrm{ad}_{(g_\alpha(p)g)^{-1}} \circ R_g^* \pi^* s_\alpha^* \omega + g_\alpha^* R_g^* \theta \\ \text{(since } \pi \circ R_g = \pi\text{)} \qquad &= \mathrm{ad}_{g^{-1} g_\alpha(p)^{-1}} \circ \pi^* s_\alpha^* \omega + g_\alpha^* (\mathrm{ad}_{g^{-1}} \circ \theta) \\ &= \mathrm{ad}_{g^{-1}} \circ \left(\mathrm{ad}_{g_\alpha(p)^{-1}} \circ \pi^* s_\alpha^* \omega + g_\alpha^* \theta\right) \\ &= \mathrm{ad}_{g^{-1}} \circ (\omega_\alpha)_p . \end{aligned}$$

Therefore they agree everywhere on $\pi^{-1}U_\alpha$. □

Now since ω is defined globally, we have that $\omega_\alpha = \omega_\beta$ on $\pi^{-1}U_{\alpha\beta}$. This allows us to relate A_α and A_β on $U_{\alpha\beta}$. Indeed, on $U_{\alpha\beta}$,

$$\begin{aligned} A_\alpha = s_\alpha^* \omega_\alpha &= s_\alpha^* \omega_\beta \\ &= s_\alpha^* \left(\mathrm{ad}_{g_\beta(s_\alpha)^{-1}} \circ \pi^* A_\beta + g_\beta^* \theta\right) \\ \text{(using } g_\beta \circ s_\alpha = g_{\beta\alpha}\text{)} \qquad &= \mathrm{ad}_{g_{\alpha\beta}} \circ A_\beta + g_{\beta\alpha}^* \theta . \end{aligned}$$

In summary,

$$A_\alpha = \mathrm{ad}_{g_{\alpha\beta}} \circ A_\beta + g_{\beta\alpha}^* \theta , \tag{4}$$

or, equivalently,

$$A_\alpha = \mathrm{ad}_{g_{\alpha\beta}} \circ \left(A_\beta - g_{\alpha\beta}^* \theta\right) ,$$

where we have used the result of the following

Done? ❑ **Exercise 1.4.** Show that

$$\mathrm{ad}_{g_{\alpha\beta}} \circ g_{\alpha\beta}^* \theta = -g_{\beta\alpha}^* \theta .$$

For matrix groups this becomes the more familiar

$$\boxed{A_\alpha = g_{\alpha\beta} A_\beta g_{\alpha\beta}^{-1} - d g_{\alpha\beta} g_{\alpha\beta}^{-1} .} \tag{5}$$

Conversely, given a family of one-forms $A_\alpha \in \Omega^1(U_\alpha;\mathfrak{g})$ satisfying equation (4) on overlaps $U_{\alpha\beta}$, we can construct a globally defined $\omega \in \Omega^1(P;\mathfrak{g})$ by the formula in Proposition 1.3. Then ω is the connection one-form of a connection on P.

In summary, we have three equivalent descriptions of a connection on P:

1. a G-invariant horizontal distribution $H \subset TP$,
2. a one-form $\omega \in \Omega^1(P;\mathfrak{g})$ satisfying $\omega(\sigma(X)) = X$ and the identity in Proposition 1.2, and
3. a family of one-forms $A_\alpha \in \Omega^1(U_\alpha;\mathfrak{g})$ satisfying equation (4) on overlaps.

Each description has its virtue and we're lucky to have all three!

1.4 The space of connections

Connections exist! This is a fact which we are not going to prove in this course. The proof can be found in [KN63, § II.2]. What we will prove is that the space of connections is an (infinite-dimensional) affine space. In fact, we have already seen this. Indeed, we saw that if ω and ω' are the connection one-forms of two connections H and H′, their difference $\tau = \omega - \omega'$ is a horizontal $\mathfrak{g}$-valued one-form on P satisfying the equivariance condition $R_g^*\tau = \mathrm{ad}_{g^{-1}} \circ \tau$. Let us see what this means on M. Let $\tau_\alpha \in \Omega^1(U_\alpha;\mathfrak{g})$ be the pull-back of τ by the local sections:

$$\tau_\alpha = s_\alpha^* \tau = s_\alpha^*(\omega - \omega') = A_\alpha - A'_\alpha \ .$$

Then equation (4) on $U_{\alpha\beta}$, says that

$$\tau_\alpha = g_{\alpha\beta}\tau_\beta g_{\alpha\beta}^{-1} = \mathrm{ad}_{g_{\alpha\beta}} \circ \tau_\beta \ . \tag{6}$$

We claim that the $\{\tau_\alpha\}$ define a section of a vector bundle associated to P.

Associated fibre bundles

Let G act on a space F via automorphisms and let $\varrho : G \to \mathrm{Aut}(F)$ be the corresponding representation. For example, F could be a vector space and $\mathrm{Aut}(F) = GL(F)$, or F could be a manifold and $\mathrm{Aut}(F) = \mathrm{Diff}(F)$.

The data defining the principal fibre bundle $P \to M$ allows to define a fibre bundle over M as follows. Consider the quotient

$$P \times_G F := (P \times F)/G$$

by the G-action $(p, f)g = (pg, \varrho(g^{-1})f)$. Since G acts freely on P, it acts freely on $P \times F$ and since P/G is a smooth manifold, so is $P \times_G F$. Moreover the projection $\pi : P \to M$ induces a projection $\pi_F : P \times_G F \to M$, by $\pi_F(p, f) = \pi(p)$, which is well-defined because $\pi(pg) = \pi(p)$. The data $\pi_F : P \times_G F \to M$ defines a fibre bundle **associated to** P **via** ϱ. For example, taking the adjoint representations $\mathrm{ad} : G \to GL(\mathfrak{g})$ and $\mathrm{Ad} : G \to \mathrm{Diff}(G)$ in turn, we arrive at the associated vector bundle $\mathrm{ad}\,P = P \times_G \mathfrak{g}$ and the associated fibre bundle $\mathrm{Ad}\,P = P \times_G G$.

The associated bundle $P \times_G F$ can also be constructed locally from the local data defining P, namely the open cover $\{U_\alpha\}$ and the transition functions $\{g_{\alpha\beta}\}$ on double overlaps. Indeed, we have that

$$P \times_G F = \bigsqcup_\alpha (U_\alpha \times F)\Big/ \sim \ ,$$

where $(m, f) \sim (m, \varrho(g_{\alpha\beta}(m))f)$ for all $m \in U_{\alpha\beta}$ and $f \in F$.

Sections of $P \times_G F$ are represented by functions $f : P \to F$ with the equivariance condition:

$$R_g^* f = \varrho(g^{-1}) \circ f \ ,$$

or, equivalently, by a family of functions $f_\alpha : U_\alpha \to F$ such that

$$f_\alpha(m) = \varrho(g_{\alpha\beta}(m)) f_\beta(m) \qquad \text{for all } m \in U_{\alpha\beta}.$$

We therefore interpret equation (6) as saying that the family of one-forms $\tau_\alpha \in \Omega^1(U_\alpha;\mathfrak{g})$ defines a one-form with values in the adjoint bundle $\operatorname{ad} P$. The space $\Omega^1(M;\operatorname{ad} P)$ of such one-forms is an (infinite-dimensional) vector space, whence the space $\mathscr{A}$ of connections on P is an infinite-dimensional affine space modelled on $\Omega^1(M;\operatorname{ad} P)$. It follows that $\mathscr{A}$ is contractible. In particular, the tangent space $T_A\mathscr{A}$ to $\mathscr{A}$ at a connection A is naturally identified with $\Omega^1(M;\operatorname{ad} P)$.

1.5 Gauge transformations

Every geometrical object has a natural notion of automorphism and principal fibre bundles are no exception. A **gauge transformation** of a principal fibre bundle $\pi : P \to M$ is a G-equivariant diffeomorphism $\Phi : P \to P$ making the following diagram commute

$$\begin{tikzcd} P \arrow[rr, "\Phi"] \arrow[dr, "\pi"'] & & P \arrow[dl, "\pi"] \\ & M & \end{tikzcd}$$

In particular, Φ maps fibres to themselves and equivariance means that $\Phi(pg) = \Phi(p)g$. Composition makes gauge transformations into a group, which we will denote $\mathscr{G}$.

We can describe $\mathscr{G}$ in terms of a trivialisation. Since it maps fibres to themselves, a gauge transformation Φ restricts to a gauge transformation of the trivial bundle $\pi^{-1}U_\alpha$ over U_α. Applying the trivialisation map $\psi_\alpha(\Phi(p)) = (\pi(p), g_\alpha(\Phi(p)))$, which lets us define $\bar\phi_\alpha : \pi^{-1}U_\alpha \to G$ by

$$\bar\phi_\alpha(p) = g_\alpha(\Phi(p))g_\alpha(p)^{-1} .$$

Equivariace of g_α and of Φ means that

$$\bar\phi_\alpha(pg) = \bar\phi_\alpha(p) ,$$

whence $\bar\phi_\alpha(p) = \phi_\alpha(\pi(p))$ for some function

$$\phi_\alpha : U_\alpha \to G .$$

For $m \in U_{\alpha\beta}$, and letting $p \in \pi^{-1}(m)$, we have

$$\begin{aligned} \phi_\alpha(m) &= g_\alpha(\Phi(p))g_\alpha(p)^{-1} \\ &= g_\alpha(\Phi(p))g_\beta(\Phi(p))^{-1}g_\beta(\Phi(p))g_\beta(p)^{-1}g_\beta(p)g_\alpha(p)^{-1} \\ \text{(since } \pi(\Phi(p)) = m\text{)} \quad &= g_{\alpha\beta}(m)\phi_\beta(m)g_{\alpha\beta}(m)^{-1} \\ &= \operatorname{Ad}_{g_{\alpha\beta}(m)}\phi_\beta(m) , \end{aligned}$$

whence the $\{\phi_\alpha\}$ define a section of the associated fibre bundle $\operatorname{Ad} P$. Since the $\{\phi_\alpha\}$ determine Φ uniquely (and viceversa), we see that $\mathscr{G} = C^\infty(M;\operatorname{Ad} P)$.

1.6 The action of $\mathscr{G}$ on $\mathscr{A}$

The group $\mathscr{G}$ of gauge transformations acts naturally on the space $\mathscr{A}$ of connections. We can see this in several different ways.

Let $H \subset TP$ be a connection and let $\Phi : P \to P$ be a gauge transformation. Define $H^\Phi := \Phi_* H$. This is also a connection on P. Indeed, the equivariance of Φ makes $H^\Phi \subset TP$ into a G-invariant distribution:

$$\begin{aligned} (R_g)_* H^\Phi_{\Phi(p)} &= (R_g)_*\Phi_* H_p \\ \text{(equivariance of } \Phi\text{)} \quad &= \Phi_*(R_g)_* H_p \\ \text{(invariance of H)} \quad &= \Phi_* H_{pg} \\ \text{(definition of } H^\Phi\text{)} \quad &= H^\Phi_{\Phi(pg)} \\ \text{(equivariance of } \Phi\text{)} \quad &= H^\Phi_{\Phi(p)g} . \end{aligned}$$

Moreover, H^Φ is still complementary to V because Φ_* is an isomorphism which preserves the vertical subspace.

Done? ❑ **Exercise 1.5.** Show that the fundamental vector fields $\sigma(X)$ of the G-action are **gauge invariant**; that is, $\Phi_*\sigma(X) = \sigma(X)$ for every $\Phi \in \mathscr{G}$. Deduce that if ω is the connection one-form for a connection H then $\omega^\Phi := (\Phi^*)^{-1}\omega$ is the connection one-form for the gauge-transformed connection H^Φ.

Finally we work out the effect of gauge transformations on a gauge field. Let $m \in U_\alpha$ and $p \in \pi^{-1}(m)$. Let A_α and A_α^Φ be the gauge fields on U_α corresponding to the connections H and H^Φ. By Proposition 1.3, the connection one-forms ω and ω^Φ are given at p by

$$\begin{aligned} \omega_p &= \mathrm{ad}_{g_\alpha(p)^{-1}} \circ \pi^* A_\alpha + g_\alpha^*\theta \\ \omega_p^\Phi &= \mathrm{ad}_{g_\alpha(p)^{-1}} \circ \pi^* A_\alpha^\Phi + g_\alpha^*\theta\,. \end{aligned} \tag{7}$$

On the other hand, $\omega^\Phi = (\Phi^{-1})^*\omega$, from where we can obtain a relation between A_α and A_α^Φ. Indeed, letting $q = \Phi^{-1}(p)$, we have

$$\begin{aligned} \omega_p^\Phi = (\Phi^{-1})^*\omega_q &= \mathrm{ad}_{g_\alpha(q)^{-1}} \circ (\Phi^{-1})^*\pi^* A_\alpha + (\Phi^{-1})^* g_\alpha^*\theta \\ \text{(functoriality of pull-back)} \quad &= \mathrm{ad}_{g_\alpha(q)^{-1}} \circ (\pi\circ\Phi^{-1})^* A_\alpha + (g_\alpha\circ\Phi^{-1})^*\theta \\ \text{(since } \pi\circ\Phi^{-1} = \pi) \quad &= \mathrm{ad}_{g_\alpha(q)^{-1}} \circ \pi^* A_\alpha + (g_\alpha\circ\Phi^{-1})^*\theta \\ \text{(since } g_\alpha(p) = \bar\phi_\alpha(p) g_\alpha(q)) \quad &= \mathrm{ad}_{g_\alpha(p)^{-1}\bar\phi_\alpha(p)} \circ \pi^* A_\alpha + (g_\alpha\circ\Phi^{-1})^*\theta\,. \end{aligned}$$

Now, $(g_\alpha\circ\Phi^{-1})(p) = g_\alpha(q) = \bar\phi_\alpha(p)^{-1} g_\alpha(p)$, whence

$$(g_\alpha\circ\Phi^{-1})^*\theta = g_\alpha^*\theta - \mathrm{ad}_{g_\alpha(p)^{-1}\bar\phi_\alpha(p)}\,\bar\phi_\alpha^*\theta\,.$$

This identity is easier to prove for matrix groups, since

$$(g_\alpha\circ\Phi^{-1})^*\theta = g_\alpha(p)^{-1}\bar\phi(p)\,d(\bar\phi(p)^{-1} g_\alpha(p))\,.$$

Now we put everything together using that $\bar\phi_\alpha = \phi_\alpha\circ\pi$ to arrive at

$$\omega_p^\Phi = \mathrm{ad}_{g_\alpha(p)^{-1}\phi_\alpha(m)} \circ \pi^*\left(A_\alpha - \phi_\alpha^*\theta\right) + g_\alpha^*\theta\,,$$

whence comparing with the second equation in (7), we conclude that

$$A_\alpha^\Phi = \mathrm{ad}_{\phi_\alpha} \circ \left(A_\alpha - \phi_\alpha^*\theta\right)\,,$$

or for matrix groups,

$$\boxed{A_\alpha^\Phi = \phi_\alpha A_\alpha \phi_\alpha^{-1} - d\phi_\alpha\,\phi_\alpha^{-1}\,.} \tag{8}$$

Comparing with equation (4), we see that in overlaps gauge fields change by a local gauge transformation defined on the overlap. This means that any gauge-invariant object which is constructed out of the gauge fields will be well-defined globally on M.