CHAPTER 1

Principal bundles and connections

1. Motivation: Gauge theory

The simplest example of a gauge theory in physics is electromagnetism. Recall Maxwell's equations for an electromagentic field,

$$\vec{\nabla} \cdot \vec{B} = 0, \quad \vec{\nabla} \times \vec{E} = -\frac{\partial}{\partial t}\vec{B} \quad \text{(homogeneous equations)} \tag{1}$$

and

$$\vec{\nabla} \cdot \vec{E} = \rho, \quad \vec{\nabla} \times \vec{B} = \frac{\partial}{\partial t}\vec{E} + \vec{j} \quad \text{(inhomogeneous equations)}\ . \tag{2}$$

Here we have chosen units so that ϵ_0, μ_0 and c are equal to 1. $\vec{E}$ is the electric field, $\vec{B}$ the magnetic field, $\vec{j}$ the current density, and ρ the electric charge density. All of these fields are functions on space-time $\mathbb{R}^4$, with coordinates $x^0 = t$ and $(x^1, x^2, x^3) = \vec{x}$. One can re-write Maxwell's equations in more concise, coordinate free form. Let the *electro-magnetic field strength* F be the following 2-form on $\mathbb{R}^4$,

$$F = \mathrm{d}x^0 \wedge (E^1 \mathrm{d}x^1 + E^2 \mathrm{d}x^2 + E^3 \mathrm{d}x^3) - B^1 \mathrm{d}x^2 \wedge \mathrm{d}x^3 + B^2 \mathrm{d}x^3 \wedge \mathrm{d}x^1 + B^3 \mathrm{d}x^1 \wedge \mathrm{d}x^2,$$

and define a charge-current density 3-form by

$$J = \rho \mathrm{d}x^1 \wedge \mathrm{d}x^2 \wedge \mathrm{d}x^3 - \mathrm{d}x^0 \wedge (j^1 \mathrm{d}x^2 \wedge \mathrm{d}x^3 + j^2 \mathrm{d}x^3 \wedge \mathrm{d}x^1 + j^3 \mathrm{d}x^1 \wedge \mathrm{d}x^2).$$

Also $* : \Omega^k(\mathbb{R}^4) \to \Omega^{4-k}(\mathbb{R}^4)$ be the star operator with respect to the standard orientation and the Minkowski metric on $\mathbb{R}^4$ [1] Then Maxwell's equations are equivalent to

$$\mathrm{d}F = 0 \quad \text{(homogeneous equations)} \tag{3}$$

and

$$\mathrm{d} * F = J \quad \text{(inhomogeneous equations)}. \tag{4}$$

The inhomogeneous equations require the integrability condition $\mathrm{d}J = 0$, which is easily recognized as *charge conservation*. Note that the Equations (4) are not quite as "natural" as the homogenous Maxwell-equations: The homogeneous Maxwell equations make sense

[1] For any pseudo-Riemannian manifold (M, g) of dimension n, g extends to a metric on $\wedge^k T^*M$. If M is oriented, one therefore obtains a unique oriented volume form Λ such that $||\Lambda||^2 = \pm 1$. The Hodge star operator $* : \Omega^k(M) \to \Omega^{n-k}(M)$ is defined by the equation $(\alpha, *\beta)\Lambda = \alpha \wedge \beta$. For $\mathbb{R}^4$ with Minkowski metric, one has $\Lambda = \mathrm{d}x^0 \wedge \mathrm{d}x^1 \wedge \mathrm{d}x^2 \wedge \mathrm{d}x^3$, thus e.g. $*(\mathrm{d}x^0 \wedge \mathrm{d}x^1) = \mathrm{d}x^2 \wedge \mathrm{d}x^3$.

on any 4-manifold M while the inhomogeneous ones involve the choice of a (pseudo-Riemannian) metric g. More precisely, since multiplying g by a positive function does not change the star operator *on 2-forms* in $\mathbb{R}^4$, they depend on the conformal structure of g.

The homogeneous Maxwell equations just say that F is a closed 2-form. Of course on $\mathbb{R}^4$, any closed 2-form is exact, so $F = \mathrm{d}A$ for some 1-form $A \in \Omega^1(\mathbb{R}^4)$ called the electro-magnetic potential. The 1-form A is defined up to a closed 1-form. On $\mathbb{R}^4$, any closed 1-form is exact so that any two potentials A, A' for F are related by $A' = A + \mathrm{d}f$ for some function $f \in \Omega^0(\mathbb{R}^4)$. One says that f defines a *gauge transformation* of the potential A. At this stage, the potential A on $\mathbb{R}^4$ itself does not have "physical meaning": Only $F = \mathrm{d}A$, i.e. the gauge equivalence class of A, can be interpreted in terms of $\vec{B}$ and $\vec{E}$ and can be measured in experiments. [2]

Using A the inhomogeneous Maxwell equation(s) can be obtained as Euler-Lagrange equations for a Lagrange density (a 4-form)

$$\mathcal{L} = \tfrac{1}{2}\mathrm{d}A \wedge *\mathrm{d}A + A \wedge J.$$

The first term in this expression can also be written $\frac{1}{2}||F_A||^2 \mathrm{d}x_0 \wedge \mathrm{d}x_1 \wedge \mathrm{d}x_2 \wedge \mathrm{d}x_3$, by definition of the star operator. In terms of $\vec{B}$ and $\vec{E}$, $||F_A||^2 = ||\vec{B}||^2 - ||\vec{E}||^2$. The functional

$$A \mapsto \tfrac{1}{2}\int_{\mathbb{R}^4} \mathrm{d}A \wedge *\mathrm{d}A$$

(defined on potentials A which decrease sufficiently fast at infinity) is a special case of the *Yang-Mills functional.*

Gauge transformations become more interesting if the electromagnetic field is coupled to some particle field. Consider for example an electron described according to Dirac's theory by a wave function $\psi : \mathbb{R}^4 \to \mathbb{C}^4$. (Here $\mathbb{C}^4$ accomodates both the electron (as a spin $\frac{1}{2}$ particle) and its anti-particle, the positron). Dirac's equation for a free electron reads, $(i\sum_{\mu=0}^3 \gamma^\mu \partial_\mu + m)\psi = 0$, where the gamma-matrices γ^μ are 4×4 matrices giving a representation of the Clifford algebra of $\mathbb{R}^4$. The equation for ψ in a (non-quantized) electromagnetic field $A = \sum_{\mu=0}^3 A_\mu \mathrm{d}x^\mu$ are obtained by "minimal coupling", i.e. replacing derivatives $\frac{1}{i}\partial_\mu$ by *covariant derivatives* $\frac{1}{i}\partial_\mu - eA_\mu$:

$$(\sum_{\mu=0}^{3} \gamma^\mu (i\partial_\mu + eA_\mu) + m)\psi = 0. \tag{5}$$

This equation is invariant under the gauge transformations,

$$(\psi, A) \mapsto (e^{if}\psi, A + \mathrm{d}f),$$

[2] Note that on a non-simply connected space time M, it is possibly for two 1-forms A, A' to have the same differential without being gauge equivalent. Thus gauge equivalence of A is a finer invariant than field strength $F_A = \mathrm{d}A$. The famous Aharanov-Bohm experiment shows that the gauge equivalence class *does* have physical meaning.

for $f : \mathbb{R}^4 \to \mathbb{R}$. Note that only $g = e^{if} : \mathbb{R}^4 \to \mathrm{U}(1)$ really enters the gauge transformations. In terms of g,

$$(\psi, A) \mapsto (g\,\psi, A - i\mathrm{d}gg^{-1}).$$

In this sense, the theory of an electromagnetic field is a gauge theory with gauge group $\mathcal{G} = \{g : \mathbb{R}^4 \to \mathrm{U}(1)$. Note that the theory of an electron without an EM field only has a global U(1) gauge symmetry: Electromagnetism is necessary to turn the global U(1) gauge symmetry into a local gauge symmetry.

The idea of Yang-Mills theory is to replace the abelian gauge group U(1) by non-commutative Lie groups G. The gauge fields A are now 1-forms with values in the Lie algebra $\mathfrak{g}$ of G. Again there is a notion of curvature 2-form F_A, which is interpreted as a field strength. The coupling of this gauge field to a particle field (with values in some complex inner product space V) depends on the choice of a unitary representation of G on V, and replacing derivatives by covariant derivatives in such a way that the coupled theory becomes gauge invariant.

2. Principal bundles and connections

In this section we review basic material on principal bundles, connections, curvature and parallel transport.

2.1. Fiber bundles. Let F be a given manifold. A (smooth) fiber bundle with *standard fiber* F is a smooth map $\pi : E \to B$ from a manifold E (the *total space*) to a manifold B (called *base*) with the following property, called *local triviality*: There exists an open covering U_α of B and diffeomorphisms

$$\phi_\alpha : \pi^{-1}(U_\alpha) \to U_\alpha \times F$$

such that $\mathrm{pr}_{U_\alpha}(\phi(x)) = \pi(x)$. (That is, ϕ intertwines π with projection to the first factor.) We will denote the fibers by $E_b = \pi^{-1}(b)$.

A morphism of fiber bundles $\pi_j : E_j \to B_j$ with fibers F_j is a smooth map $\phi : E_1 \to E_2$ taking fibers to fibers. In particular, ϕ induces a map $B_1 \to B_2$ on the base. Conversely, suppose $E \to B$ is a fiber bundle with fiber F, and $\psi : Y \to B$ is a smooth map. Then one can define a pull-back bundle $\psi^* E \to Y$ as a fibered product

$$\psi^* E = \{(y, x) \in Y \times E| \ \ \psi(y) = \pi(x)\}.$$

By construction one has a bundle map $\psi^* E \to E$. The bundle E is called trivializable if there exists a bundle map $\phi :\ E \to B \times F$; the choice of ϕ is called a trivialization.

A special case of this construction is the *fibered product* of two fiber bundles $E_1, E_2 \to B$ over the same base: Let $E_1 \times E_2 \to B \times B$ be the direct product, and $\mathrm{diag}_B : B \to B \times B$ the diagonal map. One defines

$$E_1 \times^B E_2 = \mathrm{diag}_B^*(E_1 \times E_2).$$

Thus $E_1 \times^B E_2 \to B$ has fibers

$$(E_1 \times^B E_2)_b = (E_1)_b \times (E_2)_b.$$

Additional structures on F give rise to special types of fiber bundles:

- If $F = V$ is a vector space, one defines a *vector bundle* with standard fiber V to be a fiber bundle $\pi : E \to B$ where all fibers $\pi^{-1}(b)$ are vector spaces and the local trivializations ϕ_α can be chosen to be fiberwise linear. A homomorphism of two vector bundles is a fiber bundle homomorphism that is fiberwise linear. The fibered product of vector bundles E_1, E_2 is a vector bundle (also called *Whitney sum* and denoted $E_1 \oplus E_2$). One has natural bundle maps $E \oplus E \to E$ (fiberwise addition) and a map $\mathbb{R} \times E \to E$ (fiberwise scalar multiplication).
- If $F = G$ has the structure of a Lie group, one defines a *group bundle* $\mathcal{G} \to B$ with standard fiber G to be a fiber bundle where all fibers carry group structures and the local trivializations can be chosen to be fiberwise group homomorphisms. A group bundle homomorphism is a fiber bundle homomorphism which is fiberwise a group homomorphism. The fibered product of group bundles is a group bundle. One has natural bundle maps $\mathcal{G} \times^B \mathcal{G} \to \mathcal{G}$ (fiberwise group multiplication) and $\mathcal{G} \to \mathcal{G}$ (fiberwise inversion).

Similarly, one defines algebra bundles, Lie algebra bundles, ... An action of a group bundle $\mathcal{G} \to B$ on a fiber bundle $E \to B$ is a smooth map $\mathcal{G} \times^B E \to E$ preserving fibers and defining a group action fiberwise. Similarly one can define fiberwise linear actions of group or algebra bundles on vector bundles, ...

Example 2.1. Given a vector bundle $\pi : E \to B$, one obtains a group bundle $\mathrm{GL}(E) \to B$ with fiber over b the invertible transformations (automorphisms) of the vector space $E_b = \pi^{-1}(b)$, and an algebra bundle $\mathrm{End}(E) \to B$ with fibers the endomorphisms of E_b. There are natural fiberwise linear maps $\mathrm{GL}(E) \times^B E \to E$ and $\mathrm{End}(E) \times^B E \to E$.

2.2. Principal bundles. Let G be a Lie group. A *principal homogeneous G-space* is a manifold X together with a transitive, free G-action. For example, G itself with action $g.a = ag^{-1}$ is a principal homogeneous space. (We could also use the left-action, of course.) Any principal homogeneous G-space X is isomorphic to this example: Given $x_0 \in X$ there is a unique G-equivariant isomorphism $X \to G$ taking x_0 to e. The only reason for introducing the pedantic notion of principal homogeneous G-space is that in general, there is no distinguished choice of x_0. Note that the group $\mathrm{Diff}(X)^G$ of G-equivariant diffeomorphisms of X to itself is diffeomorphic to G (but not canonically so), and that the space of G-equivariant diffeomorphisms $\mathrm{Diff}(X, G)^G$ is a principal homogeneous G-space (with G acting on the target G by left multiplication) which is canonically isomorphic to X itself.

Example 2.2. If V is a real vector space of dimension n, the space $\mathrm{Fr}(V)$ of linear isomorphisms ("frames") $V \to \mathbb{R}^n$ is a principal homogeneous $\mathrm{GL}(n, \mathbb{R})$-space. The abstract vector space V can be recovered from X as a quotient,

$$V = (\mathrm{Fr}(V) \times \mathbb{R}^n) / \mathrm{GL}(n, \mathbb{R})$$

by the diagonal action. In fact, all the "natural" (that is, $\mathrm{GL}(n,\mathbb{R})$-equivariant) constructions with $\mathbb{R}^n$ give rise to the corresponding constructions of V by a similar quotient: For instance, $P(V) = (\mathrm{Fr}(V) \times \mathbb{R}P^{n-1})/\,\mathrm{GL}(n,\mathbb{R})$ is the projectivization of V, $\wedge^k(V) = (\mathrm{Fr}(V) \times \wedge^k(\mathbb{R}^n))/\,\mathrm{GL}(n,\mathbb{R})$ is the exterior algebra, and so on.

A *principal G-bundle* is a fiber bundle $\pi : \mathcal{P} \to B$ where each fiber $\mathcal{P}_b$ carries the structure of a principal homogeneous G-space, and such that the local trivializiations ϕ_α can be chosen to be fiberwise G-equivariant.

EXAMPLES 2.3. a. Let $S^1 = \mathbb{R}/\mathbb{Z}$ the standard circle. The k-fold covering $\pi : S^1 \to S^1$ is a principal bundle with structure group $\mathbb{Z}_k = \mathbb{Z}/k\mathbb{Z}$.

b. Let $\pi : E \to B$ be a real vector bundle of rank n. Then the *frame bundle* $\mathrm{Fr}(E)$ with fibers $\mathrm{Fr}(E)_b = \mathrm{Fr}(E_b)$ is a principal $\mathrm{GL}(n,\mathbb{R})$ bundle.

c. If in addition the fibers E_b carry inner products and orientations, depending smoothly on $b \in B$, one can define the bundle $\mathrm{Fr}_{SO}(E)_b = \mathrm{Fr}_{SO}(E_b)$ of special orthogonal frames, which is a principal $\mathrm{SO}(n)$ bundle.

A *homomorphism of principal G-bundles* $\pi_j : \mathcal{P}_j \to B_j$ is a G-equivariant smooth map $\phi : \mathcal{P}_1 \to \mathcal{P}_2$ taking fibers to fibers. Any such map induces a smooth map $\psi : B_1 \to B_2$ on the base. Conversely, given a smooth map $\psi : Y \to B$ and given a principal G-bundle $\pi : \mathcal{P} \to B$, one can form the *pull-back* bundle $\psi^*\mathcal{P} \to Y$ using the fibered product of ψ and π: That is,

$$\psi^*\mathcal{P} = \{(y,p) \in Y \times \mathcal{P} |\, \psi(y) = \pi(p)\}.$$

The bundle projection $\psi^*\mathcal{P} \to Y$ is induced by projection $Y \times \mathcal{P} \to Y$ to the first factor, and projection to the second factor $Y \times \mathcal{P} \to \mathcal{P}$ gives rise to a principal bundle homomorphism $\psi^*P \to \mathcal{P}$. In fact, any principal bundle homomorphism $\phi : \mathcal{P}_1 \to \mathcal{P}_2$ identifies $\mathcal{P}_1$ with the pull-back of $\mathcal{P}_2$ by the map $B_1 \to B_2$ on the base.

2.3. Associated bundles. Let $\pi : \mathcal{P} \to B$ be a *principal G-bundle.* Given a G-manifold F, one defines the *associated fiber bundle* by

$$F(\mathcal{P}) \equiv \mathcal{P} \times_G F := (\mathcal{P} \times F)/G.$$

The space $\mathcal{P} \times_G F$ is a fiber bundle over $B = \mathcal{P}/G$ with standard fiber F. The sections $\Gamma^\infty(B, \mathcal{P} \times_G F)$ of this fiber bundle are naturally identified with the space $C^\infty(\mathcal{P}, F)^G$ of equivariant maps $\mathcal{P} \to F$.

EXAMPLES 2.4. a. If V is a vector space on which G acts linearly, then $\mathcal{P} \times_G V$ is a vector bundle. Taking $V = \mathfrak{g}$ with the adjoint representation one obtains the *adjoint bundle* $\mathfrak{g}(P) := \mathcal{P} \times_G \mathfrak{g}$.

b. If K is a Lie group on which G acts by automorphisms, $\mathcal{P} \times_G K$ is a group bundle. Taking $K = G$ with G acting by the adjoint action, one obtains a group bundle $G(P) := \mathcal{P} \times_G G$ which is also called the adjoint bundle. It has $\mathfrak{g}(P)$ as its Lie algebra bundle.

c. If F is a principal homogeneous H-space on which G acts H-equivariantly, the associated bundle $\mathcal{P} \times_G F$ is a principal H-bundle. For example, if $F = H = G$ with G acting by left multiplication one recovers $\mathcal{P}$ itself.

d. Let E be a real vector bundle of rank n, and $\mathrm{Fr}(E)$ its $\mathrm{Gl}(n)$ frame bundle. The bundle associated to the defining representation of $\mathrm{GL}(n)$ on $\mathbb{R}^n$ is E itself:

$$\mathrm{Fr}(E) \times_{\mathrm{GL}(n)} \mathbb{R}^n = E.$$

The representation on $\wedge^k \mathbb{R}^n$ resp. $S^k \mathbb{R}^n$ gives the anti-symmetric resp. symmetric powers $\wedge^k E$ and $S^k E$. The bundle associated to the conjugation action of $GL(n)$ on itself is the group bundle $\mathrm{GL}(E)$. For $s \in \mathbb{R}$, the real line bundle $\mathrm{Fr}(E) \times_{\mathrm{GL}(n)} \mathbb{R}$ defined by the representation $A \mapsto |\det(A)|^s$ of $\mathrm{GL}(n, \mathbb{R})$ is a real line bundle called the bundle of s-densities. The line bundle for the representation $A \mapsto \det(A)$ is the determinant line bundle $\det(E)$. The bundle to the contragredient representation $A \mapsto (A^{-1})^t$ on $\mathbb{R}^n$ is the dual bundle E^*.

G-equivariant maps $F_1 \to F_2$ give rise to fiber bundle homomorphisms $F_1(\mathcal{P}) \to F_2(\mathcal{P})$. For instance, if V is a G-representation, the map

$$G \times V \to V,\ (a, v) \mapsto a.v$$

is equivariant for the G-action $g.(a, v) = (gag^{-1}, g.v)$ on $G \times V$. It follows that one obtains an action of the group bundle $G(\mathcal{P})$ on $V(\mathcal{P})$.

2.4. Sections. Let $\pi : E \to B$ be a fiber bundle with fiber F. A smooth map $\sigma : B \to E$ is called a section of E if $\pi \circ \sigma = \mathrm{Id}_B$. The set of sections will be denoted $\Gamma^\infty(B, E)$. For a vector bundle, it is a vector space, for an algebra bundle it is an algebra, for a group bundle it is a group. Let $\mathcal{P} \to B$ be a principal G-bundle, and F a G-manifold. The pull-back bundle $\pi^* F(\mathcal{P})$ is canonically isomorphic to the trivial bundle:

$$\pi^* F(\mathcal{P}) \cong \mathcal{P} \times F.$$

Hence, pull-back of sections gives a canonical isomorphism,

$$\Gamma^\infty(B, F(\mathcal{P})) \cong C^\infty(\mathcal{P}, F)^G.$$

For sections $\sigma \in \Gamma^\infty(B, F(\mathcal{P}))$ we will denote by $\tilde{\sigma} \in C^\infty(\mathcal{P}, F)^G$ the corresponding equivariant function.

Group bundles and vector bundles always have a distinguished section, the identity section resp. zero section. In general, fiber bundles need not admit any sections at all. A principal fiber bundle has a section if and only if it is trivializable, and the choice of a section is equivalent to a choice of trivialization. Indeed, given $\phi : \mathcal{P} \to B \times G$ one can define a section $\sigma(b) = \phi^{-1}(b, e)$. Conversely, given σ one defines $\phi^{-1} : B \times G \to \mathcal{P}$ by $\phi^{-1}(b, g) = g.\sigma(b)$.

For vector bundles $E \to B$, one can define a space of E-valued differential forms

$$\Omega^k(B, E) = \Gamma^\infty(B, \wedge^k T^* M \otimes E).$$

It is a module for the algebra $\Omega^\star(B)$ but does not carry a natural differential. Suppose $E = \mathcal{P} \times_G V$ is an associated bundle. Then there is a canonical identification,

$$\Omega^k(B, E) \cong (\Omega^k_{\mathrm{hor}}(\mathcal{P}) \otimes V)^G =: \Omega^k_{\mathrm{basic}}(\mathcal{P}, V).$$

The isomorphism is obtained as follows: By construction we have a bundle homomorphism $\mathcal{P} \times V \to E, \quad (p, v) \mapsto G.(p, v)$ covering the bundle projection $\pi : \mathcal{P} \to B$. This identifies the pull-back $\pi^*E \to \mathcal{P}$ with the trivial bundle $\mathcal{P} \times V$. It follows that π^* induces a map

$$\pi^* : \ \Omega^k(B, E) \to \Omega^k(\mathcal{P}, \pi^*E) = \Omega^k(\mathcal{P}) \otimes V.$$

The image takes values in horizontal forms (since the pull-back of a form under π vanishes on vectors tangent to the fibers) which are G-invariant (since π is G-invariant). It is easy to check that this map is 1-1 and onto $\Omega^k_{\mathrm{basic}}(\mathcal{P}, V)$.

We will often use this isomorphism, since it is usually easier to work with vector-space valued forms rather than vector-bundle valued forms. If $\alpha \in \Omega^\star(B, V(\mathcal{P}))$, we will denote by $\tilde{\alpha}$ the corresponding form in $\Omega^\star_{\mathrm{hor}}(\mathcal{P}, V)^G$.

2.5. Gauge transformations. A principal bundle automorphism is a G-equivariant diffeomorphism $\phi : \mathcal{P} \to \mathcal{P}$ taking fibers to fibers. The group of principal bundle automorphisms will be denoted $\mathrm{Aut}(\mathcal{P})$. The *gauge group* $\mathrm{Gau}(\mathcal{P}) \subset \mathrm{Aut}(\mathcal{P})$ consists of automorphisms $\phi : \mathcal{P} \to \mathcal{P}$ inducing the identity map on the base B. That is, $\mathrm{Gau}(\mathcal{P})$ is defined by an exact sequence of groups

$$1 \longrightarrow \mathrm{Gau}(\mathcal{P}) \to \mathrm{Aut}(\mathcal{P}) \to \mathrm{Diff}(B). \tag{6}$$

In general the last map is not onto.

Any $\phi \in \mathrm{Gau}(\mathcal{P})$ gives rise to a map $\tilde{\phi} : \mathcal{P} \to G$ by

$$\phi(p) = \tilde{\phi}(p)^{-1}.p. \tag{7}$$

The map $\tilde{\phi}$ is equivariant since

$$g\tilde{\phi}(p)^{-1}.p = g.\phi(p) = \phi(g.p) = \tilde{\phi}(g.p)^{-1}g.p$$

Conversely, given any equivariant map $\tilde{\phi} : P \to G$ the map ϕ defined by (7) $\phi(p) = \tilde{\phi}(p)^{-1}.p$ is a gauge transformation. The map $\phi \mapsto \tilde{\phi}$ is a group homomorphism since

$$(\phi_1\phi_2)(p) = \phi_1(\tilde{\phi}_2(p)^{-1}.p) = \tilde{\phi}_2(p)^{-1}.\phi_1(p) = \tilde{\phi}_2(p)^{-1}\tilde{\phi}_1(p)^{-1}.p = (\tilde{\phi}_1\tilde{\phi}_2)(p)^{-1}.p$$

It therefore defines a canonical isomorphism

$$\mathrm{Gau}(\mathcal{P}) \cong C^\infty(\mathcal{P}, G)^G \cong \Gamma^\infty(B, G(\mathcal{P})).$$

The Lie algebra of the group $\mathrm{Aut}(\mathcal{P})$ is the Lie algebra $\mathrm{aut}(\mathcal{P}) = \mathfrak{X}(\mathcal{P})^G$ of G-invariant vector fields on $\mathcal{P}$, and the Lie algebra of the gauge group is the subspace $\mathrm{gau}(\mathcal{P})$ of vertical invariant vector fields. One has an exact sequence of Lie algebras

$$0 \longrightarrow \mathrm{gau}(\mathcal{P}) \longrightarrow \mathrm{aut}(\mathcal{P}) \longrightarrow \mathfrak{X}(B) \to 0. \tag{8}$$

Here surjectivity of the map $\mathrm{aut}(\mathcal{P}) \longrightarrow \mathfrak{X}(B)$ can be proved using a trivializing open cover U_α of B and a partition of unity subordinate to that cover. $\mathrm{gau}(\mathcal{P})$ can be identified with sections of $\mathfrak{g}(\mathcal{P})$:

$$\mathrm{gau}(\mathcal{P}) \cong \Omega^0(B, \mathfrak{g}(\mathcal{P})).$$

The group $\mathrm{Gau}(\mathcal{P})$ acts on all associated bundles $F(\mathcal{P})$, since $\phi \times \mathrm{id} : \mathcal{P} \times F \to \mathcal{P} \times F$ descends to quotients by the G-action. Furthermore it acts on the space of sections $\Gamma^\infty(B, F(\mathcal{P}))$ by $(\phi.\sigma)(b) = \phi.(\sigma(b))$. One has

$$\widetilde{\phi.\sigma} = \tilde{\phi}.\tilde{\sigma} \quad .$$

If $F = V$ is a vector space on which G acts linearly, this action extends uniquely to $\Omega^\star(B, V(\mathcal{P}))$ in such a way that $\phi.(\alpha \wedge \beta) = \alpha \wedge \phi.\beta$ for $\alpha \in \Omega^\star(B)$ and $\beta \in \Omega^\star(B, V(\mathcal{P}))$. One has

$$\widetilde{\phi.\alpha} = \tilde{\phi}.\tilde{\alpha}, \quad \alpha \in \Omega^\star(B, V(\mathcal{P})).$$

Infinitesimally one has a Lie algebra action of $\mathrm{gau}(\mathcal{P}) = \Omega^0(B, \mathfrak{g}(\mathcal{P}))$; this action extends to an action of the graded Lie algebra $\Omega^\star(B, \mathfrak{g}(\mathcal{P}))$.

3. Connections

3.1. Connections on fiber bundles. For any fiber bundle $\pi : E \to B$ the tangent bundle TE of the total space has a distinguished subbundle, the vertical bundle $VE \hookrightarrow TE$. The fiber V_xE for $\pi(x) = b$ is the image of $T_x(F_b)$ under the natural inclusion $TF_b \hookrightarrow TE$. An (Ehresmann) connection on E is the choice of a complementary *horizontal subbundle* HE such that $TE = VE \oplus HE$. Equivalently, a connection is a bundle projection $TE \to VE$ which is left-inverse to the inclusion $VE \to TE$; one defines HE as the kernel of this projection. Ehresmann connections always exist: For instance, one may take HE to be the orthogonal complement of VE for some Riemannian metric on the total space of E. If $\psi : Y \to B$ is a smooth map and $HE \subset TE$ a given connection, one obtains a connection on the pull-back bundle $\psi^*E \to Y$ by defining

$$H\psi^*E = (\phi_*)^{-1}HE.$$

Here $\phi : \psi^*E \to E$ is the bundle map covering ψ, and $\phi_* : T\psi^*E \to TE$ its differential.

For any smooth path $\gamma : [t_0, t_1] \to B$ from $b_0 = \gamma(t_0)$ to $b_1 = \gamma(t_1)$, the connection defines a *parallel transport*

$$\Pi^\gamma : E_{b_0} \to E_{b_1}$$

as follows: Given $p_0 \in E_{b_0}$ let $\tilde{\gamma} : [t_0, t_1] \to E$ be the unique path with $\pi \circ \tilde{\gamma} = \gamma$ and $\frac{\mathrm{d}}{\mathrm{d}t}\tilde{\gamma}(t) \in H_{\tilde{\gamma}(t)}E$. Put $\Pi^\gamma(p_0) := \tilde{\gamma}(t_1)$. The definition extends uniquely to piecewise smooth paths, in such a way that for any two paths γ_1, γ_2 such that the end point of γ_1 equals the initial point of γ_2,

$$\Pi^{\gamma_2 * \gamma_1} = \Pi^{\gamma_2} \circ \Pi^{\gamma_1}$$

where $\gamma_2 * \gamma_1$ is the concatenation of γ_2 and γ_1. (This can be nicely phrased in terms of "groupoid homomorphisms" but we won't do this here.) If one fixes $b = b_0$ and considers only loops γ based at b, one obtains a map from the group $L(M,b)$ of piecewise smooth loops based at b into $\mathrm{Diff}(b)$, called holonomy of the connection:

$$\mathrm{Hol}: L(M,b) \to \mathrm{Diff}(E_b).$$

If B is connected, the holonomy subgroups with respect to different base points b, b' are isomorphic, each choice of a path from b, b' defines an isomorphism. A connection is called *flat* if the holonomy for any contractible loop is trivial.

3.2. Principal connections. If F has additional structure, one is interested in connections such that parallel transport preserves that structure. For example, if E is a vector bundle, each Π^γ should be a linear map, for group bundles it should be a fiberwise group homomorphism and so on.

To construct such connections it is most convenient to work with the corresponding principal bundle.

For a principal G-bundle $\pi: \mathcal{P} \to B$, the map

$$\mathcal{P} \times \mathfrak{g} \to V\mathcal{P}, \quad (p,\xi) \mapsto \xi_P(p) := \frac{\mathrm{d}}{\mathrm{d}t}|_{t=0} \exp(-t\xi).p$$

defines a canonical trivialization of the vertival bundle $V\mathcal{P}$. Since $g_*(\xi_P(p)) = (\mathrm{Ad}_g\,\xi)_{\mathcal{P}}(g.p)$, the trivialization is G-equivariant for the adjoint action of G on its Lie algebra.

Hence an Ehresmann connection on $\mathcal{P}$ is equivalent to a vector bundle homomorphism $T\mathcal{P} \to \mathcal{P} \times \mathfrak{g}$ taking $\xi_P(p)$ to ξ. Equivalently, a connection is a Lie-algebra valued 1-form $\theta \in \Omega^1(\mathcal{P}, \mathfrak{g})$ with

$$\iota(\xi_P)\theta = \xi \text{ for all } \xi \in \mathfrak{g}. \tag{9}$$

The connection is G-equivariant if and only if the projection map $T\mathcal{P} \to P \times \mathfrak{g}$ is G-equivariant, that is,

$$g^*\theta = \mathrm{Ad}_g\,\theta \text{ for all } g \in G. \tag{10}$$

Definition 3.1. A (principal) connection on a fiber bundle is an equivariant Lie-algebra valued 1-form $\theta \in \Omega^1(\mathcal{P}, \mathfrak{g})^G$ such that $\iota(\xi_P)\theta = \xi$ for all $\xi \in \mathfrak{g}$. The space of principal connections will be denoted $\mathcal{A}(\mathcal{P})$.

For later reference let us note the infinitesimal version of the invariance condition:

$$\mathcal{L}_{\xi_{\mathcal{P}}}\theta = -[\xi, \theta].$$

Proposition 3.2. *The space $\mathcal{A}(\mathcal{P})$ of principal connections has a natural affine structure, with underlying vector space the space $\Omega^1(B, \mathfrak{g}(\mathcal{P}))$ of 1-forms on B with values in the adjoint bundle.*

Proof. We first show that $\mathcal{A}(\mathcal{P})$ is non-empty. This is obvious if $\mathcal{P}$ is a trivial bundle $\mathcal{P} = B \times G$. If $\mathcal{P}$ is non-trivial, choose an open cover U_α of B with trivializations $\phi_\alpha : \pi^{-1}(U_\alpha) \cong U_\alpha \times G$, and a subordinate partition of unity ρ_α. Over each U_α, th trivialization defines a principal connection θ_α. A global connection is given by $\theta =$

$\sum_\alpha (\pi^*\rho_\alpha)\theta_\alpha$. The difference $\gamma = \theta_1 - \theta_2$ between any two connections satisfies $\iota(\xi_P)(\theta_1 - \theta_2) = 0$. It is thus a 1-form in $(\Omega^1_{\mathrm{hor}}(\mathcal{P}) \otimes \mathfrak{g})^G \cong \Omega^1(B, \mathfrak{g}(\mathcal{P}))$. Conversely, adding such a form γ to a principal connection produces a new principal connection. □

By construction, parallel transport $\Pi\gamma$ with respect to a principal connection is a G-equivariant map. That is, the holonomy at $b \in B$ becomes a group homomorphism

$$\mathrm{Hol}:\ L(B,b) \to \mathrm{Diff}(\mathcal{P}_b)^G.$$

Any choice of a base point $p \in \mathcal{P}_b$ identifies $\mathcal{P}_b \cong G$ and therefore $\mathrm{Diff}(\mathcal{P}_b)^G \cong G$.

3.3. Connections on associated bundles. If F is any G-manifold, the principal connection on $\mathcal{P}$ defines a connection on $\mathcal{P} \times_G F$ as follows: Let $q : \mathcal{P} \times F \to \mathcal{P} \times_G F$ the quotient map, and note that

$$H\mathcal{P} \oplus TF \oplus \ker(q_*) = T(\mathcal{P} \times F).$$

Hence $q_*(TF) = V(\mathcal{P} \times_G F)$, and

$$H(\mathcal{P} \times_G F) = q_*(H\mathcal{P})$$

defines a complementary subbundle. If G acts by automorphisms of a given structure on F, parallel transport on the associated bundle preserves that structure.

(Say F is a vector space V with G acting linearly. View $\mathcal{P} \times V$ as a G-equivariant fiber bundle over $\mathcal{P}$. The pull-back of the connection on $\mathcal{P} \times_G V$ is given by the horizontal subbundle $H(\mathcal{P} \times V) = H\mathcal{P} \oplus \ker(q_*)$ (the horizontal subspace of $\mathcal{P}$ together with orbit directions). Vectorfields taking values in $H\mathcal{P}$ preserve the linear structure, as do generating vector field for the diagonal G-action. By equivariance, parallel transport on $\mathcal{P} \times V$ descends to parallel transport on the associated bundle.)

3.4. Covariant derivative. A form $\alpha \in \Omega^k(\mathcal{P})$ is called *horizontal* if $\iota(\xi_{\mathcal{P}})\alpha = 0$ for all ξ, and *basic* if in addition α is G-invariant. Denote the space of horizontal k-forms by $\Omega^k_{\mathrm{hor}}(\mathcal{P})$ and the space of basic k-forms by $\Omega^k_{\mathrm{basic}}(\mathcal{P})$. Pull-back induces an isomorphism $\pi^* :\ \Omega^k(B) \cong \Omega^k_{\mathrm{basic}}(\mathcal{P})$. A principal connection gives rise to a G-equivariant projection operator

$$\mathrm{Hor}^\theta :\ \Omega^k(\mathcal{P}) \to \Omega^k_{\mathrm{hor}}(\mathcal{P}).$$

The covariant derivative defined by θ is the composition,

$$\mathrm{d}^\theta = \mathrm{Hor}^\theta \circ \mathrm{d} :\ \Omega^k(\mathcal{P}) \to \Omega^{k+1}_{\mathrm{hor}}(\mathcal{P}).$$

More generally, let V be a G-representation, and $E = \mathcal{P} \times_G V$ the corresponding associated vector bundle. As we remarked earlier, pull-back defines an isomorphism

$$\Omega^k(B, E) \cong \Omega^k_{\mathrm{hor}}(\mathcal{P}, \pi^* E)^G = \Omega^k_{\mathrm{hor}}(\mathcal{P}, V)^G$$

since $\pi^* E = \mathcal{P} \times V$ canonically. The differential d^θ (extended to V-valued forms) preserves this space, so it defines an operator

$$\mathrm{d}^\theta :\ \Omega^k(B, V(\mathcal{P})) \to \Omega^{k+1}(B, V(\mathcal{P})).$$

Often one denotes this operator by ∇^θ or simply ∇. It respects the $\Omega^\star(B)$-module structure, in the sense that

$$\mathrm{d}^\theta(\alpha \wedge \beta) = \mathrm{d}\alpha \wedge \beta + (-1)^k \alpha \wedge \mathrm{d}^\theta \beta,$$

for all $\alpha \in \Omega^k(B)$, $\beta \in \Omega^l(B, V(\mathcal{P}))$.

LEMMA 3.3. *For all* $\alpha \in \Omega^\star_{\mathrm{hor}}(\mathcal{P}, V)^G$,

$$d^\theta \tilde{\alpha} = d\tilde{\alpha} + \theta.\tilde{\alpha}$$

Here the "." denotes the Lie algebra representation of $\mathfrak{g}$ *on* V.

PROOF. The right hand side is horizontal, since

$$\iota(\xi_{\mathcal{P}})(\mathrm{d}\tilde{\alpha} + \theta.\tilde{\alpha}) = L_{\xi_{\mathcal{P}}}\tilde{\alpha} - \mathrm{d}\iota(\xi_P)\tilde{\alpha} + \xi.\tilde{\alpha} = 0.$$

Here we have used that $\iota(\xi_{\mathcal{P}})\tilde{\alpha} = 0$ by horizontality and $L_{\xi_{\mathcal{P}}}\tilde{\alpha} + \xi.\tilde{\alpha} = 0$ by invariance. It remains to show that the two sides agree on horizontal vectors. But this is clear since $\theta.\tilde{\alpha}$ vanishes on any $k+1$ horizontal vector fields, and $\mathrm{d}^\theta \tilde{\alpha} = \mathrm{Hor}^\theta \mathrm{d}\tilde{\alpha}$ and $\mathrm{d}\tilde{\alpha}$ agree on any $k+1$ horizontal vector fields, by definition. □

3.5. Curvature. Let θ be a principal connection on $\pi : \mathcal{P} \to B$. The *curvature* of θ is the $\mathfrak{g}$-valued 2-form

$$\tilde{F}^\theta = \mathrm{d}^\theta \theta.$$

By definition, $\tilde{F}^\theta$ is G-invariant and horizontal, so it can also be viewed as a 2-form $F^\theta \in \Omega^2(B, \mathfrak{g}(\mathcal{P}))$. One has the alternative expression

$$\tilde{F}^\theta = \mathrm{d}\theta + \tfrac{1}{2}[\theta, \theta].$$

To see this, note that $\mathrm{d}\theta + \frac{1}{2}[\theta,\theta]$ agrees with $\mathrm{d}\theta$ on horizontal vectors, and that it is horizontal since

$$\iota(\xi_{\mathcal{P}})\mathrm{d}\theta = L_{\xi_{\mathcal{P}}}\theta = -[\xi, \theta] = -\tfrac{1}{2}\iota(\xi_{\mathcal{P}})[\theta, \theta].$$

PROPOSITION 3.4. *The curvature satisfies the* Bianchi identity

$$d^\theta F^\theta = 0.$$

PROOF. Since $\tilde{F}^\theta \in \Omega^2_{\mathrm{hor}}(\mathcal{P}, \mathfrak{g})^G$, Lemma 3.3 applies, and we can calculate

$$\mathrm{d}^\theta \tilde{F}^\theta = \mathrm{d}\tilde{F}^\theta + [\theta, \tilde{F}^\theta] = \tfrac{1}{2}\mathrm{d}[\theta,\theta] + [\theta, \mathrm{d}\theta] + \tfrac{1}{2}[\theta, [\theta, \theta]].$$

The first two terms cancel and the last vanishes by the Jacobi identity for $\mathfrak{g}$. □

There are many different interpretations of what curvature "measures". The following Proposition says that curvature measures the extent to which the bracket of horizontal vector fields fails to be horizontal. Equivalently (by Frobenius' theorem) curvature measures the extent to which the horizontal distribution $\ker(\theta) \subset T\mathcal{P}$ fails to be integrable.

PROPOSITION 3.5. *If* $X, Y \in \mathfrak{X}(\mathcal{P})$ *are horizontal vector fields on* $\mathcal{P}$,

$$\theta([X, Y]) = -\tilde{F}^\theta(X, Y).$$

PROOF. Using that $\iota_X\theta = \iota_Y\theta = 0$,

$$\theta([X,Y]) = \theta(L_X Y) = L_X(\iota_Y\theta) - \iota_Y L_X\theta = -\iota_Y L_X\theta = -\iota_Y\iota_X \mathrm{d}\theta = -\tilde{F}^\theta(X,Y).$$

□

One can look at this result from a slightly different angle. The connection θ determines a *horizontal lift* of vector fields $\mathrm{Lift}^\theta : \mathfrak{X}(B) \to \mathfrak{X}(\mathcal{P})$ such $\mathrm{Lift}^\theta(X)$ is horizontal and projects down to X (i.e. is π-related to X). In general, the lift is not a Lie algebra homomorphism. However, since with any two pairs of π-related vector fields their bracket is also π-related, the difference $\mathrm{Lift}^\theta([X,Y]) - [\mathrm{Lift}^\theta(X), \mathrm{Lift}^\theta(Y)]$ is vertical. By the proposition, if X, Y are any two vector fields on B,

$$\tilde{F}^\theta(\mathrm{Lift}^\theta(X), \mathrm{Lift}^\theta(Y)) = \theta\big(\mathrm{Lift}^\theta([X,Y]) - [\mathrm{Lift}^\theta(X), \mathrm{Lift}^\theta(Y)]\big).$$

Thus curvature measures the extent to which the horizontal lift Lift^θ fails to be a Lie algebra homomorphism. Finally, curvature measures the extent to which the covariant differential d^θ fails to define a differential:

LEMMA 3.6. *Let V be a linear G-representation. In the notation of Lemma 3.3, if $\alpha \in \Omega^k(B, \mathcal{P} \times_G V)$,*

$$(d^\theta)^2\alpha = F^\theta.\alpha.$$

PROOF. Using Lemma 3.3, $(\mathrm{d}^\theta)^2\tilde{\alpha} = (\mathrm{d}+\theta.)^2\tilde{\alpha} = (\mathrm{d}\theta).\tilde{\alpha} + \theta.\theta.\tilde{\alpha} = (\mathrm{d}\theta + \frac{1}{2}[\theta,\theta]).\tilde{\alpha}$ □

3.6. Gauge transformations of connections. The group of automorphisms $\phi \in \mathrm{Aut}(\mathcal{P})$ acts on the space $\mathcal{A}(\mathcal{P})$ of principal connections by pull-back by the inverse. For the action of the subgroup $\mathrm{Gau}(\mathcal{P})$, there is an alternative formula for this action which we derive below. Let $\theta^L, \theta^R \in \Omega^1(G, \mathfrak{g})$ the left- and right-invariant Maurer Cartan forms. Thus if ξ^L, ξ^R are the left- and right-invariant vector fields on G equal to $\xi \in \mathfrak{g} = T_eG$ at the group unit, $\iota(\xi^L)\theta^L = \xi$ and $\iota(\xi^R)\theta^R = \xi$. In a matrix representation of G, $\theta^L = g^{-1}\mathrm{d}g$ and $\theta^R = \mathrm{d}g g^{-1}$. One has $\theta^R = \mathrm{Ad}_g(\theta^L)$. Under the inversion map $\iota : G \to G$, $g \mapsto g^{-1}$ one has $\iota^*\theta^L = -\theta^R$.

PROPOSITION 3.7. *Let $\phi \in \mathrm{Gau}(\mathcal{P})$ and $\tilde{\phi} \in C^\infty(\mathcal{P}, G)^G$ the corresponding equivariant function. For all $X \in \mathrm{aut}(\mathcal{P}) = \mathfrak{X}(\mathcal{P})^G$,*

$$\phi_* X - X \in \mathrm{gau}(\mathcal{P}) = \mathfrak{X}_{\mathrm{vert}}(\mathcal{P})^G.$$

The corresponding section $\zeta \in \Omega^0(B, \mathfrak{g}(\mathcal{P})) \cong \mathrm{gau}(\mathcal{P})$ is given by

$$\tilde{\zeta} = \iota_X(\tilde{\phi}^*\theta^R).$$

PROOF. The vector field $\phi_* X$ is invariant since ϕ commutes with the G-action, and under π_* projects to the same vector field as X since $\pi \circ \phi = \pi$. This shows that

$\phi_* X - X \in \mathfrak{X}_{\mathrm{vert}}(\mathcal{P})^G$. Let Ψ_t denote the flow of X. Then

$$\begin{aligned}(\phi_* X)_p &= \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0}\phi(\Psi_t(p)) \\ &= \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0}\tilde{\phi}(\Psi_t(p))^{-1}.\Psi_t(p) \\ &= \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0}\tilde{\phi}(p)^{-1}.\Psi_t(p) + \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0}\tilde{\phi}(\Psi_t(p))^{-1}.p.\end{aligned}$$

The first term is $X_{\phi(p)}$ since

$$\tilde{\phi}(p)^{-1}.\Psi_t(p) = \Psi_t(\tilde{\phi}(p)^{-1}.p) = \Psi_t(\phi(p)).$$

For the second term, define a curve on G starting at e by

$$g_t = \tilde{\phi}(\Psi_t(p))^{-1}\tilde{\phi}(p).$$

Then

$$\frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0}\tilde{\phi}(\Psi_t(p))^{-1}.p = \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0} g_t\tilde{\phi}(p)^{-1}.p = \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0} g_t.\phi(p) = (\xi_P)_{\phi(p)},$$

with

$$\xi = -\frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0} g_t = \iota(\tilde{\phi}_* X_p)(\theta^L_{\tilde{\phi}(p)}) = \big(\iota_X(\tilde{\phi}^*\theta^L)\big)_p = \big(\iota_X(\tilde{\phi}^*\theta^R)\big)_{\phi(p)}.$$

□

PROPOSITION 3.8. *The natural action of* $\mathrm{Gau}(\mathcal{P})$ *on* $\mathcal{A}(\mathcal{P})$ *is given by the formula,*

$$\phi.\theta = (\phi^{-1})^*\theta = \mathrm{Ad}_{\tilde{\phi}}\,\theta - \tilde{\phi}^*\theta^R.$$

PROOF. We will prove the equivalent property $\phi^*\theta = \mathrm{Ad}_{\tilde{\phi}^{-1}}\,\theta + \tilde{\phi}^*\theta^L$. It suffices to check this identity on invariant vector fields $X \in \mathfrak{X}(\mathcal{P})^G$. By Proposition 3.7, and since $\phi^* f = \mathrm{Ad}_{\tilde{\phi}^{-1}} f$ for all $f \in C^\infty(\mathcal{P}, \mathfrak{g})^G$,

$$\iota(X)\phi^*\theta = \phi^*(\iota(\phi_* X)\theta) = \mathrm{Ad}_{\tilde{\phi}^{-1}}(\iota(\phi_* X)\theta) = \iota(X)(\mathrm{Ad}_{\tilde{\phi}^{-1}}\,\theta + \tilde{\phi}^*\theta^L).$$

as desired. □

Let us (informally) regard $\mathcal{A}(\mathcal{P})$ as an infinite dimensional manifold, equipped with an action of an infinite-dimensional Lie group. What are the generating vector fields, and what are the stabilizer subgroupsΓ Given $\zeta \in \Omega^0(\mathcal{P}, \mathfrak{g}(\mathcal{P})) \cong \mathrm{gau}(\mathcal{P})$, and identifying $T_\theta\mathcal{A}(\mathcal{P}) \cong \Omega^1(\mathcal{P}, \mathfrak{g}(\mathcal{P}))$, one has:

COROLLARY 3.9. *The generating vector field for the action of* ζ *on* $\mathcal{A}(\mathcal{P})$ *is given by the covariant differential:*

$$\zeta_{\mathcal{A}(\mathcal{P})}(\theta) = d^\theta\zeta.$$

PROOF. Let $\tilde{\zeta} \in \Omega^0(\mathcal{P}, \mathfrak{g})^G$ be the equivariant function corresponding to ζ. Then

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0}\exp(-t\zeta).\theta &= \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0}\mathrm{Ad}_{\exp(-t\tilde{\zeta})}\,\theta - \exp(-t\tilde{\zeta})^*\theta^R \\ &= -[\tilde{\zeta}, \theta] + \mathrm{d}\tilde{\zeta} = \mathrm{d}^\theta\tilde{\zeta}\end{aligned}$$

where we have used Lemma 3.3. □

The stabilizer subgroups $\mathrm{Gau}(\mathcal{P})_\theta$ can be characterized in terms of the holonomy. Let us first describe how gauge transformations act on parallel transport. For any path $\gamma : [t_0, t_1] \to B$, let

$$\Pi^\theta_\gamma : \mathcal{P}_{\gamma(t_0)} \to \mathcal{P}_{\gamma(t_1)}$$

denote parallel transport with respect to θ. We will need:

LEMMA 3.10. *For all $\phi \in \mathrm{Gau}(\mathcal{P})$,*

$$\Pi^{\phi.\theta}_\gamma = \phi(\gamma(t_1)) \circ \Pi^\theta_\gamma \circ \phi(\gamma(t_0))^{-1}.$$

We leave the proof as an exercise. (Hint: Pulling back $\mathcal{P}$ under γ, we may assume B is 1-dimensional.)

Suppose B is connected, and let $b \in B$. Recall the holonomy homomorphism

$$\mathrm{Hol}^\theta : \; L(B, b) \to \mathrm{Aut}(\mathcal{P}_b), \quad \gamma \mapsto \Pi^\theta_\gamma$$

from the group of loops based at b into the group $\mathrm{Aut}(\mathcal{P}_b) \cong G$ of automorphisms of the fiber. Let the *holonomy subgroup* $G^\theta \subset \mathrm{Aut}(\mathcal{P}_b)$ be the image of the holonomy map, and the *restricted holonomy subgroup* G^θ_0 the image of the subgroup $L_0(B, b)$ of contractible loops. G^θ_0 is a closed subgroup (hence, Lie subgroup) of $\mathrm{Aut}(\mathcal{P}_b)$, while G^θ need not be a closed subgroup.

PROPOSITION 3.11. *Let $\mathcal{P} \to B$ be a principal G-bundle over a connected base, $b \in B$, and $\theta \in \mathcal{A}(\mathcal{P})$. The evaluation map $\mathrm{Gau}(\mathcal{P})_\theta \to \mathrm{Aut}(\mathcal{P}_b)$, $\phi \mapsto \phi(b)$ is injective, and defines an isomorphism of $\mathrm{Gau}(\mathcal{P})_\theta$ with the centralizer in $\mathrm{Aut}(\mathcal{P}_b)$ of the holonomy group:*

$$\mathrm{Gau}(\mathcal{P})_\theta \cong Z_{\mathrm{Aut}(\mathcal{P}_b)}(G^\theta).$$

PROOF. A gauge transformation $\phi \in \mathrm{Gau}(\mathcal{P})$ preserves θ if and only if it preserves the horizontal subbundle $H\mathcal{P} \subset T\mathcal{P}$. Equivalently, by the Lemma (with ϕ^{-1} in place of ϕ) $\phi \in \mathrm{Gau}(\mathcal{P})_\theta$ if and only if for all paths $\gamma : [t_0, t_1] \to B$, ϕ commutes with parallel transport:

$$\phi(\gamma(t_1)) = \Pi_\gamma \circ \phi(\gamma(t_0)) \circ \Pi_\gamma^{-1}. \tag{11}$$

Taking paths γ with $\gamma(t_0) = b$, it follows that if ϕ is trivial at b then ϕ is trivial everywhere. That is, the evaluation map $\phi \mapsto \phi(b)$ is injective. Taking $\gamma \in L(B, b)$, (11) shows that $\phi(b)$ centralizes G^θ. Conversely, if $\phi(b) \in Z(G^\theta)$, then we can define a gauge transformation $\phi \in \mathrm{Gau}(\mathcal{P})_\theta$ by putting

$$\phi(b') = \Pi_\gamma \circ \phi(b) \circ \Pi_\gamma^{-1},$$

where γ is any path from b to b'; the condition $\phi \in \mathrm{Gau}(\mathcal{P})_\theta$ guarantees that the right hand side does not depend on the choice of γ. □

The proposition shows in particular that all stabilizer subgroups are finite dimensional Lie groups. Moreover:

COROLLARY 3.12. *The* based gauge group $\mathrm{Gau}(\mathcal{P}, b) = \{\phi \in \mathrm{Gau}(\mathcal{P}) | \, \phi(b) = e\}$ *acts freely on* $\mathcal{A}(\mathcal{P})$.

COROLLARY 3.13. *If* G *is abelian, every principal connection has stabilizer equal to* G*, viewed as constant gauge transformations in* $\mathrm{Gau}(\mathcal{P}) = \Gamma^\infty(B, G(\mathcal{P}))^G = C^\infty(B, G)$.

The quotient $\mathcal{A}(\mathcal{P})/\,\mathrm{Gau}(\mathcal{P})$ is called the moduli space of connections. It is still infinite-dimensional. To obtain finite dimensional moduli spaces, one has to impose additional (gauge-invariant) constraints on θ: E.g. that θ is a flat connection, or more generally a Yang-Mills connection. (See below.)

3.7. Reducible connections. Let $\mathcal{P} \to B$ be a principal G-bundle.

DEFINITION 3.14. A *reduction of the structure group* from G to a Lie subgroup $H \subset G$ is a principal H-bundle $\mathcal{Q} \to B$ together with an H-equivariant fiber bundle map $\iota : \mathcal{Q} \hookrightarrow \mathcal{P}$ making $\mathcal{Q}$ into an H-invariant submanifold of $\mathcal{P}$. A principal connection θ on $\mathcal{P}$ is called reducible to H if the horizontal distribution $H\mathcal{P}$ for θ is tangent to $\mathcal{Q}$. A connection on $\mathcal{P} \to B$ is called *irreducible* if it is not reducible to a subgroup H of G.

Thus, θ is a reducible connection if and only if its pull-back $\iota^*\theta$ takes values in $\mathfrak{h}$, and is a principal connection for $\mathcal{Q}$. The principal bundle $\mathcal{P}$ can be recovered from $\mathcal{Q}$ as an associated bundle $\mathcal{P} \cong \mathcal{Q} \times_H G$ where H acts on G from the left. The reducible connections on $\mathcal{P}$ are those which come from connections on $\mathcal{Q}$, using the construction for associated bundles.

EXAMPLES 3.15. A choice of an inner product on a real vector bundle is *equivalent* to a reduction of the structure group of the frame bundle $\mathrm{Fr}(E)$ from $\mathrm{GL}(n, \mathbb{R})$ to $O(n)$, a choice of an orientation is equivalent to a reduction of the structure group to the identity component $\mathrm{GL}^+(n, \mathbb{R})$. Similarly the choice of a complex structure, symplectic structure, Hermitian structure, etc. are equivalent to various reductions of the structure group. Connections which are reducible to these subgroups have parallel transports preserving these extra structures.

Reducible connections can be recognized from their holonomy groups. Indeed, if θ is reducible to an $H \subset G$ connection on $\mathcal{Q} \subset \mathcal{P}$, then its holonomy subgroup is contained in $\mathrm{Aut}(\mathcal{Q}_b) \subset \mathrm{Aut}(\mathcal{P}_b)$. Conversely, if the holonomy subgroup $G^\theta \subset \mathrm{Aut}(\mathcal{P}_b)$ at $b \in B$ is a proper subgroup of $\mathrm{Aut}(\mathcal{P}_b)$, pick $p \in \mathcal{P}_b$ and let $\mathcal{Q}$ be the set of all points in $\mathcal{P}$ that can be reached from p by parallel transport. Then $\mathcal{Q}$ is a principal H-bundle for the subgroup H corresponding to G^θ under the identification $\mathrm{Aut}(\mathcal{P}_b) \cong G$ given by the choice of p.

We have shown above that the stabilizer group of θ under gauge transformations is isomorphic to the centralizer of G^θ. Hence, for an irreducible connection it is isomorphic to the center $Z(G)$. The converse is not true (cf. Corollary 3.13).

Reducibility can also be detected from the curvature, using that the holonomy group and restricted holonomy group coincide in this case, together with:

THEOREM 3.16 (Ambrose-Singer theorem). *Let $\mathcal{P} \to B$ be a principal bundle over a connected base B, and θ a principal connection. Given $p \in \mathcal{P}$ let $Q \subset B$ be the set of all points which can be connected to p by parallel transport. Then the Lie algebra of the holonomy subgroup relative to p is equal to the set of all $\tilde{F}^\theta(X,Y)_q \in \mathfrak{g}$ with $q \in Q$.*

For a proof, see e.g. Kobayashi-Nomizu I, p. 89.

3.8. The universal connection. Consider $\mathcal{A}(\mathcal{P}) \times B$ as an infinite-dimensional manifold, which is the base for a universal principal bundle

$$\mathcal{A}(\mathcal{P}) \times \mathcal{P} \to \mathcal{A}(\mathcal{P}) \times B$$

(not, of course, to be confused with the classifying bundle $EG \to BG$, which often is also called universal bundle). The tangent space at any point (θ, p) is the direct sum of $\Omega^1(B, \mathfrak{g}(\mathcal{P}))$ and of $T_p\mathcal{P}$. Let $\Theta \in \Omega^1(\mathcal{A}(\mathcal{P}) \times \mathcal{P}, \mathfrak{g})$ be the 1-form such that $\Theta_{(\theta,p)}$ vanishes on $\Omega^1(B, \mathfrak{g}(\mathcal{P}))$ and equals θ_p on $T_p\mathcal{P}$. Clearly, Θ is a connection 1-form. It is called the *universal connection*. Let us calculate its curvature $F^\Theta = \mathrm{d}\Theta + \frac{1}{2}[\Theta, \Theta]$ of this connection. For $a, b \in \Omega^1(B, \mathfrak{g}(\mathcal{P}))$ and $X \in T_p\mathcal{P}$ we have

$$\iota(a)\iota(b)(\mathrm{d}\Theta)_{(\theta,p)} = \iota(a)\frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0} \Theta_{(\theta+tb,p)} = 0,$$

$$\iota(X)\iota(a)(\mathrm{d}\Theta)_{(\theta,p)} = \iota(X)\frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0} \Theta_{(\theta+ta,p)} = \iota(X)a_p$$

and finally $\iota(X)\iota(Y)\mathrm{d}\Theta = \iota(X)\iota(Y)\mathrm{d}\theta$. Thus

$$F^\Theta(a,b) = 0, \quad F^\Theta(X,Y) = F^\theta(X,Y), \quad F^\Theta(X,a) = a(X).$$

The group $\mathrm{Gau}(\mathcal{P})$ acts on $\mathcal{A}(\mathcal{P}) \times \mathcal{P}$ by automorphisms, preserving the universal connection Θ, hence also preserving the curvature F^Θ.

3.9. Yang-Mills connections. Suppose $\pi : \mathcal{P} \to B$ is a principal G-bundle over a compact, oriented, Riemannian manifold B. The inner product on TB gives rise to an inner product on T^*M and on all $\wedge^k T^*M$. Taking the inner product of differential form, followed by integration over B with respect to the Riemannian volume form, defines an inner product on $\Omega^\star(B)$. In terms of the Hodge star operator,

$$\langle \alpha, \beta \rangle = \int_B \alpha \wedge *\beta.$$

This inner product extends to $\mathfrak{g}(\mathcal{O})$-valued forms, using the inner product on the vector bundle $\mathfrak{g}(\mathcal{P})$:

$$\langle \alpha, \beta \rangle = \int_B (\alpha, *\beta).$$

Let $||\cdot||$ be the norm corresponding to $\langle \cdot, \cdot \rangle$. The Yang-Mills functional on $\mathcal{A}(\mathcal{P})$ is the functional

$$\mathrm{YM}(\theta) = ||F^\theta||^2 = \int_B (F^\theta, *F^\theta).$$

It is invariant under the action of the gauge group

$$\mathrm{YM}(\theta) = \mathrm{YM}(\phi.\theta),$$

hence all its critical points (called Yang-Mills connections) are invariant as well. If $\mathcal{P}$ admits flat connections, then $\mathcal{A}_{\text{flat}}(\mathcal{P})$ is the absolute minimum of YM.

PROPOSITION 3.17. *A connection θ is a critical point of the Yang-Mills functional if and only if it satisfies the Yang-Mills equation,*

$$d^\theta * F^\theta = 0.$$

PROOF. We use the formula $\tilde{F}^\theta = \mathrm{d}\theta + \frac{1}{2}[\theta, \theta]$. For any $\eta \in \Omega^1_{\text{hor}}(\mathcal{P}, \mathfrak{g})^G$,

$$\tilde{F}^{\theta+\eta} = \tilde{F}^\theta + \mathrm{d}\eta + \tfrac{1}{2}[\eta, \eta] + [\theta, \eta] = \tilde{F}^\theta + \mathrm{d}^\theta\eta + \tfrac{1}{2}[\eta, \eta].$$

Hence

$$\begin{aligned} \mathrm{YM}(\theta + \eta) - \mathrm{YM}(\theta) &= \int_B (\mathrm{d}^\theta\eta, *F^\theta) + \int_B (F^\theta, *\mathrm{d}^\theta\eta) + \dots \\ &= 2\int_B (\eta, \mathrm{d}^\theta * F^\theta) + \dots \end{aligned}$$

where ... denotes terms which are at least quadratic in η. Thus θ is a stationary point if and only if $\int_B (\eta, \mathrm{d}^\theta * F^\theta) = 0$ for all η, that is, $\mathrm{d}^\theta * F^\theta = 0$. □

The quotient of the space of Yang-Mills connections by the action of the gauge group is called the Yang-Mills moduli space.

The Yang Mills-equations depend upon the Riemannian metric on B only via the star operator on $\Omega^2(B)$. The case $\dim B = 4$ is special in that $*$ takes $\Omega^2(B)$ to itself, since $4 - 2 = 2$. We mentioned already that in this case, the Yang-Mills equations are conformally invariant: Multiplying the metric by a positive function does not change the star operator in middle dimension, hence does not change the Yang-Mills equations. A special type of Yang-Mills connections in 4 dimensions are those satisfying one of the equations

$$*F^\theta = F^\theta \quad \text{or} \quad *F^\theta = -F^\theta$$

(self-duality resp. anti-self-duality) because for such connections, the Yang-Mills equations are a consequence of the Bianchy identity $\mathrm{d}^\theta F^\theta = 0$. A change of orientation of B changes the sign of the $*$ operator, and therefore exchanges the notion of duality and anti-self duality. The value of the Yang-Mills functional for an ASD connection is

$$\mathrm{YM}(\theta) = -\int_B (F^\theta, F^\theta).$$

From the theory of charactersitic classes, one knows that the right hand side is independent of θ (it is a multiple of the second Chern number $c_2(\mathcal{P})$). Using this fact, and decomposing the curvature of a connection into its self-dual ann ASD part, one finds that for $c_2(\mathcal{P}) \geq 0$, ASD connections give the absolute minimum of the Yang-Mills functional.

Anti-self dual connections over S^4 are also called *instantons.* Up to gauge transformation, they can be viewed as connections on the trivial bundle over $\mathbb{R}^4$ which are flat outside a compact set. The famous ADHM construction (Atiyah-Drinfeld-Hitchin-Manin, 1978) gives a complete description of such instantons.

The moduli space for anti-self dual YM-connections for $G = \mathrm{SU}(2)$ is the starting point for Donaldson theory of 4-manifolds. As realized by Donaldson, they contain information not only about the topology but also the differentiable structure of 4-manifolds.

In this course, we will be concerned with the very different case $\dim B = 2$, and in fact mostly with flat connections.