

Topology and Physics 2019 - lecture 4

Marcel Vonk

February 26, 2019

2.1 Gauge fields or field strengths?

2.1.1 Gauge fields vs field strengths

In lecture 2, we have seen that it is very useful to write Maxwell's theory of electromagnetism in terms of differential forms. This not only allows us to study electromagnetism in a more efficient notation; it also helps us understand how the topological aspects of space-time play a role in electromagnetism – a role which is rather unclear if we describe Maxwell's theory in terms of electric and magnetic fields. Two examples of the role that topology plays in physics have already been discussed in the exercises: the Dirac monopole in exercise 2.3, and the Aharonov-Bohm effect in exercise 3.1.

However, if you have done these exercises, you may also feel slightly uneasy. Two differential forms have played an important role in our description of Maxwell theory:

- the **field strength**, F , and
- the **gauge potential** (gauge field), A .

As long as space-time is topologically trivial, say Minkowski space, these two objects contain exactly the same information: Maxwell's equations tell us that $dF = 0$, and then the Poincaré lemma tells us that this implies $F = dA$. Thus, every physically possible field strength arises from some gauge potential. Of course, there are then still multiple gauge fields that can lead to the same field strength F , but using the Poincaré lemma once again (see exercise 3. part (ib)), we know that two such gauge fields are always related by a gauge transformation, $A_1 - A_2 = d\Lambda$. Therefore, on a topologically trivial space, there are two fully equivalent ways to describe field configurations: as

The set of closed field strengths F : $\Omega_{cl}^2(M)$

and as

The set of gauge fields A up to gauge transformations: $\Omega^1(M)/d\Omega^0(M)$

So far, so good, but from the two examples we have seen in the exercises, it has become clear that on a topologically *nontrivial* space, these two descriptions are not quite equivalent. And so the question arises: which of the two descriptions above is the physically correct one? The answer will be: neither! In lecture 3 it was already mentioned that the fully correct way to describe gauge fields is in terms of connections on principle bundles, and this reasoning will be continued from a mathematical perspective in the next lecture. For now, however, we want to see exactly what goes wrong when we stick to the language of differential forms and cohomology on space-time itself. This turns out to be a very useful exercise, which clearly shows the virtues and limitations of the purely cohomological description of gauge theories.

2.1.2 A useful sequence

To see that neither one-forms nor two forms are a perfect description of electromagnetic field configurations, let us look at the exact sequence of maps that was introduced in exercise 3.1:

$$0 \xrightarrow{f_0} H_{\text{dR}}^1(M) \xrightarrow{f_1} \Omega^1(M)/d\Omega^0(M) \xrightarrow{f_2} \Omega_{\text{cl}}^2(M) \xrightarrow{f_3} H_{\text{dR}}^2(M) \xrightarrow{f_4} 0. \quad (2.1)$$

Here, we have the following maps:

- f_0 is the inclusion map,
- f_1 is also inclusion: the equivalence class $[A]$ maps to $[A]$ for any closed 1-form A ,
- f_2 is induced by the exterior derivative: $[A]$ maps to dA ,
- f_3 is the standard map from closed forms to cohomology: F maps to $[F]$,
- f_4 is the projection to 0.

In exercise 3.1, you have convinced yourself that all of these maps are well-defined, and that the sequence is *exact*: $\text{im}(f_n) = \ker(f_{n+1})$ for all pairs of consecutive maps.

Let us now have a closer look at the above sequence for the three cases we are interested in: topologically trivial space, the Aharonov-Bohm case, and the Dirac monopole.

2.1.3 Topologically trivial space

In the topologically trivial case, both cohomology groups that appear in our exact sequence are vanishing, and so the sequence simply reduces to

$$0 \xrightarrow{f_1} \Omega^1(M)/d\Omega^0(M) \xrightarrow{f_2} \Omega_{\text{cl}}^2(M) \xrightarrow{f_3} 0. \quad (2.2)$$

If you're not used to working with exact sequences, it may be useful to spell out what exactness means at the different nodes:

- The image of f_1 must equal the kernel of f_2 . Now the image of f_1 is clearly $\{0\}$, so the kernel of f_2 is $\{0\}$. That is: no two different elements map to the same element: f_2 is *injective*.
- The image of f_2 must equal the kernel of f_3 . Now the kernel of f_3 is clearly the full $\Omega_{\text{cl}}^2(M)$, so the image of f_2 is $\Omega_{\text{cl}}^2(M)$. That is: every element is in the image: f_2 is *surjective*.

If a map (in this case f_2 is both injective and surjective, it is of course bijective: the two spaces under consideration are isomorphic. That is, they are really “the same space”:

$$\Omega^1(M)/d\Omega^0(M) \cong \Omega_{\text{cl}}^2(M) \quad (2.3)$$

This is nothing but what we said in the introduction to this lecture: we may just as well describe field configurations on a topologically trivial space using either construction: as gauge fields up to gauge transformations, $[A] \in \Omega^1(M)/d\Omega^0(M)$ or as closed field strengths, $F \in \Omega_{\text{cl}}^2(M)$.

2.1.4 The Aharonov-Bohm configuration

Next, let us recall the situation of exercise 3.1: the Aharonov-Bohm effect. Here, we are in a situation where $H_{\text{dR}}^1(M)$ is nontrivial. This can be achieved for example by introducing into our space a long (in theory: infinitely long) solenoid with circular cross section of radius R . Inside the solenoid a magnetic field is present. Let us say the solenoid extends in the z -direction, and inside the solenoid there is a constant magnetic field in that same direction, B^z . Outside the solenoid, all electric and magnetic fields are zero. Thus, we have a field strength

$$\begin{aligned} r \leq R : \quad F &= -B^z dx \wedge dy, \\ r > R : \quad F &= 0. \end{aligned} \quad (2.4)$$

As a simple exercise, you may convince yourself that if we use polar coordinates (r, θ) instead of (x, y) , the field strength for $r \leq R$ may be written as $F = dA$ with

$$r \leq R : \quad A = -\frac{1}{2}B^z r^2 d\theta. \quad (2.5)$$

As a result, it makes sense to choose

$$r > R : \quad A = -\frac{1}{2}B^z R^2 d\theta, \quad (2.6)$$

which is constant, so that indeed $F = dA = 0$ for $r > R$. As another consistency check, note that for a 1-cycle γ surrounding the solenoid, we have

$$\int_{\gamma} A = -\pi R^2 B^z = \int_D F \quad (2.7)$$

where D is the disk-shaped two-cycle that is bounded by γ , and we can compute using either (2.6) or (2.4). That is: our setup is (of course) consistent with Stokes' theorem.

So far, our space is not topologically nontrivial yet, but now we can note that as long as we are only interested in the *outside* of the solenoid, we may just as well shrink it to zero size, and consider a gauge field of the form (2.6) for any $(x, y) \neq (0, 0)$. Of course, at the line $(x, y) = (0, 0)$ we now have a problem, as $A \sim d\theta$ is ill-defined there. This can be resolved by removing this line from our space altogether, so that A becomes well-defined again. Of course, the “cost” is that we have now made our space topologically nontrivial: there is now a closed 1-cycle γ which is no longer the boundary of any 2-cycle, and therefore by De Rham's theorem, $H_{\text{dR}}^1(M) \neq 0$. In fact, it is not hard to guess what the nontrivial element of $H_{\text{dR}}^1(M)$ corresponding to γ is: it is simply the one-form $d\theta$, which is indeed closed but not exact. (Note that it can not be written as d acting on the 0-form “ θ ”, as the latter is not a well-defined, single-valued function.)

Some further thought shows that in our example we still have that $H_{\text{dR}}^2(M) = 0$, so that now our exact sequence (2.1) becomes

$$0 \xrightarrow{f_0} H_{\text{dR}}^1(M) \xrightarrow{f_1} \Omega^1(M)/d\Omega^0(M) \xrightarrow{f_2} \Omega_{\text{cl}}^2(M) \xrightarrow{f_3} 0. \quad (2.8)$$

So now we have an exact sequence containing *three* nontrivial terms. Such a sequence is called a *short* exact sequence in mathematics, and it is a construction one encounters over and over again. Again, if you have never encountered such sequences, it is useful to spell out what exactness of such a sequence means. Here, in particular, it is useful to recall that for any linear map f between vector spaces, we have that the image is isomorphic to the domain modulo the kernel:

$$\text{im}(f) \cong \text{dom}(f)/\ker(f) \quad (2.9)$$

In fact, this statement is not only true for vector spaces; it holds in many different categories of objects, which is why short exact sequences are so useful in mathematics. Let us now apply this statement to the map f_2 , and also use exactness of our sequence:

- $\text{im}(f_2) = \ker(f_3) = \Omega_{\text{cl}}^2(M)$, and
- $\ker(f_2) = \text{im}(f_1)$, which is isomorphic to the full $H_{\text{dR}}^1(M)$ by yet another (easier) application of (2.9), this time on f_1 .

Thus, using exactness, (2.9) for f_2 reduces to

$$\Omega_{\text{cl}}^2(M) \cong \frac{\Omega^1(M)/d\Omega^0(M)}{H_{\text{dR}}^1(M)}. \quad (2.10)$$

In words: the space of closed two-forms F is strictly *smaller* than the space of gauge fields up to gauge transformations: it is exactly gauge fields of the form $A \sim d\theta$ that the field strength “does not see”. Yet the presence of such gauge fields can be measured (for example

by measuring the phase shift of light passing along both sides of the solenoid), and so it seems that our description of nature in terms of 1-forms A is richer and more physical than a description in terms of 2-forms F .

While the general spirit of this conclusion is correct, it is not the whole story – as will become clear by studying our final example, the Dirac monopole.

2.1.5 The Dirac monopole

In the case of the Dirac monopole, we are essentially in the opposite situation: here, we removed a single *point* (as opposed to a full *line*) from three-dimensional space¹. As a result, now any closed 1-cycle can be viewed as the boundary of a 2-cycle (one can always draw a disk-like shape that “misses” the excluded point), but now there is a closed 2-cycle that can not be viewed as the boundary of a 3-cycle: the sphere (or anything homotopic to it) surrounding the removed point.

Thus, from De Rham’s theorem, we now know that $H_{\text{dR}}^1(M) = 0$, but that $H_{\text{dR}}^2(M)$ is nontrivial. In fact, there is only one nontrivial closed 2-cycle, and therefore $H_{\text{dR}}^2(M) = \mathbb{R}$: it is generated exactly by the two-form

$$F = \sin \theta d\theta \wedge d\phi \tag{2.11}$$

that we studied in exercise 1.3.

If we once again write down our favorite exact sequence (2.1), we therefore have

$$0 \xrightarrow{f_1} \Omega^1(M)/d\Omega^0(M) \xrightarrow{f_2} \Omega_{\text{cl}}^2(M) \xrightarrow{f_3} H_{\text{dR}}^2(M) \xrightarrow{f_4} 0. \tag{2.12}$$

Again, we find a *short* exact sequence, and therefore we can immediately use the same result as in the previous example for this sequence: we find

$$H_{\text{dR}}^2(M) \cong \frac{\Omega_{\text{cl}}^2(M)}{\Omega^1(M)/d\Omega^0(M)}. \tag{2.13}$$

As the left hand side is nonzero, we see that now the set of closed 2-forms F is strictly *larger* than the space of gauge fields up to gauge transformations. It seems that now F contains more information than A !

The reason for this conclusion is clear: since our space is topologically nontrivial, and since the nontriviality shows up exactly in the cohomology of 2-forms, Maxwell’s equation $dF = 0$ no longer implies that $F = dA$. As a result, we should no longer be all that interested in the space of all gauge fields up to gauge transformations: these only parameterize part of

¹Note that in our entire discussion, we can essentially forget about the time-direction of space-time, as all of our field configurations are constant in time.

the allowed field configurations. Somehow, we should consider a more broad definition of “gauge field” that also incorporates the monopole-like field configurations that now exist.

In exercise 1.3, we saw a somewhat cumbersome way to describe such “more general gauge fields”: we cut our space in two (topologically trivial) halves, studied gauge fields on the individual halves, and then glued the halves together – but in such a way that the glueing itself could be done up to gauge transformations. In the rest of this lecture and in particular the next one, we will formalize this idea more. As we will see, this will require some extra technical equipment: we must stop viewing the gauge field A as a (globally defined) one-form, and we must start using the language of *fibre bundles*. These fibre bundles will be equipped with gauge-field-like objects called *connections*, and when the dust settles, we will discover that those connections are the “correct” way to think about the gauge field A , in such a way that one can now describe all situations that have appeared in this section. (As well as the even more general situation where *both* $H_{\text{dR}}^1(M)$ and $H_{\text{dR}}^2(M)$ are nonzero.)

In the next lecture, we will introduce fibre bundles and connections with some mathematical rigor. As a warm-up, in the rest of this lecture we want to get a better feeling for how these objects arise in physics. It will turn out that for this purpose, our simple example of Maxwell theory is in fact a bit too simple: we will need to study a broader class of theories called Yang-Mills theories.

2.2 Yang-Mills theory

2.2.1 Gauge fields and covariant derivatives

To understand the broader definition of our gauge field, let us discuss a seemingly rather different way in which gauge fields can arise in physics. As a toy model, we consider a quantum field theory where the field is a complex scalar field $\phi(x)$. What kind of action could such a field have? Of course, we are free to choose any functional as the action, but in practice most field theories follow the “kinetic energy minus potential energy” recipe that we have seen before. The kinetic energy term would contain two factors of the time derivative of $\phi(x)$ (kinetic energy is proportional to velocity squared); since moreover we want energy to be a real quantity, and we want our theory to be relativistic (so that time derivatives appear on an equal footing with spatial derivatives) we naturally arrive at something like

$$S \sim \int \eta^{\mu\nu} \partial_\mu \phi^*(x) \partial_\nu \phi(x), \quad (2.14)$$

where $*$ denotes complex conjugation. As for the potential energy term, we of course have a lot of freedom, though the requirement that the energy is real remains. Let us take the simplest nontrivial case, where $V(x) \sim \phi^*(x)\phi(x)$ and consider

$$S \sim \int \left(\eta^{\mu\nu} \partial_\mu \phi^*(x) \partial_\nu \phi(x) - m \phi^*(x) \phi(x) \right) d^4x, \quad (2.15)$$

with m some constant which for our purposes is irrelevant. (It turns out to be related to the *mass* of the excitations of the quantum field.) In fact, we could have added any function of $\phi^*\phi$; all arguments below will remain unchanged.

To get a grip on a quantum field theory, it is always useful to start by finding the *symmetries* of the model: operations we can apply to the field configurations such that the action remains unchanged. For our toy action, such a symmetry is for example the global phase rotation

$$\phi(x) \rightarrow e^{i\Lambda}\phi(x) \quad (2.16)$$

for any $\Lambda \in \mathbb{R}$. We call this a *global* phase rotation because the phase Λ is independent of the position x . What happens if we instead consider a *local* phase rotation? This would mean we would change the field as

$$\phi(x) \rightarrow e^{i\Lambda(x)}\phi(x). \quad (2.17)$$

Clearly, the combination $\phi^*\phi$ is still invariant under this operation, but this is not the case for the derivative terms. In fact, we have

$$\partial_\mu\phi(x) \rightarrow e^{i\Lambda(x)}\left(\partial_\mu\phi(x) + i\partial_\mu\Lambda(x)\phi(x)\right). \quad (2.18)$$

Clearly because of the presence of the second term in this transformation, the term $\eta^{\mu\nu}\partial_\mu\phi^*(x)\partial_\nu\phi(x)$ is not invariant. There is, however, a way to solve this issue. That way is to introduce the *covariant derivative*,

$$D_\mu\phi(x) \equiv \partial_\mu\phi(x) - iA_\mu(x)\phi(x), \quad (2.19)$$

where at this point $A_\mu(x)$ is a new field about which we have not specified anything other than the fact that it has a lower μ -index. (Of course, the notation strongly suggests the interpretation of A_μ we are working towards.) If we apply our local phase rotation to the above term, we find that

$$D_\mu\phi(x) \rightarrow e^{i\Lambda(x)}\left(\partial_\mu\phi(x) + i\partial_\mu\Lambda(x)\phi(x) - i\tilde{A}_\mu(x)\phi(x)\right). \quad (2.20)$$

In this expression, we have written \tilde{A}_μ , since we have not specified yet how our would-be symmetry acts on $A_\mu(x)$ – we simply gave the field after the transformation the name \tilde{A}_μ . However, now we see that if we impose that A_μ transforms as

$$A_\mu(x) \rightarrow \tilde{A}_\mu(x) \equiv A_\mu(x) + \partial_\mu\Lambda(x), \quad (2.21)$$

we get that

$$D_\mu\phi(x) \rightarrow e^{i\Lambda(x)}D_\mu\phi(x), \quad (2.22)$$

and therefore

$$\eta^{\mu\nu}D_\mu\phi^*(x)D_\nu\phi(x) \quad (2.23)$$

is invariant. The transformation (2.21) looks familiar: it is exactly a gauge transformation! That is, if we interpret $A = A_\mu(x)dx^\mu$ as a 1-form, the gauge transformation is simply

$$A \rightarrow A + d\Lambda \tag{2.24}$$

where we interpret the function $\Lambda(x)$ as a 0-form.

The above observation allows us to construct a broad class of interesting theories. Of course, now that we have discovered that it is interesting to “couple” the complex scalar field $\phi(x)$ to a gauge field $A_\mu(x)$, we may as well study theories where $A_\mu(x)$ itself has a kinetic term so that it becomes a “dynamical” field. Of course, this kinetic term for A must be gauge invariant as well, and so the natural choice is to construct the kinetic term out of $F = dA$. As we mentioned before, the term $F \wedge F$ does not lead to any interesting physics, and so we are led to consider a kinetic term of the form $F \wedge \star F$ – that is, we study field theories with an action²

$$S[A, \phi] = \int \left(F \wedge \star F + \eta^{\mu\nu} D_\mu \phi^* D_\nu \phi + V(\phi^* \phi) \right). \tag{2.25}$$

In a very different way, we have found back the action for Maxwell theory for A , now “coupled” to an action for a scalar field $\phi(x)$. It turns out that in nature, this is exactly how the electromagnetic field interacts with other, charged fields: through covariant derivatives.

2.2.2 Covariant derivatives and parallel transport

The fact that a gauge field A shows up in the covariant derivative leads us to a very different way of interpreting gauge fields: as being related to *parallel transport*. This will be made more explicit from a mathematical perspective in the next lecture, but let us already give the basic ideas here. To simplify notation, let us now look at a space which is simply \mathbb{R} , parameterized by a single coordinate x . To begin with, recall that the ordinary derivative can be used as a “generator” for translation of functions: for a function $\phi(x)$, using a Taylor series, we have that

$$\begin{aligned} \phi(x + \epsilon) &= \phi(x) + \epsilon \frac{d\phi}{dx} + \frac{1}{2} \epsilon^2 \frac{d^2\phi}{dx^2} + \dots \\ &= e^{\epsilon \frac{d}{dx}} \phi(x). \end{aligned} \tag{2.26}$$

Thus, the derivative “transports” the function value. Is there a similar interpretation for the covariant derivative?

To answer this question, let us look at *constant* functions in particular: functions for which

$$\frac{d}{dx} \phi(x) = 0 \tag{2.27}$$

²Note the slightly awkward notation here, where in the first term we integrate a 4-form (which can be done without making a choice of coordinates), whereas the integration of the other two terms should really contain a d^4x and *is* coordinate-dependent in general.

for every x . We want to figure out what the analogous statement $D_\mu\phi(x) = 0$ for the covariant derivative means.

To this end, we must realize that there are two different ways to think of a function $\phi : \mathbb{R} \rightarrow \mathbb{C}$. The first one is the usual way: there is a complex plane \mathbb{C} , and every point $x \in \mathbb{R}$ is mapped to some point $y = \phi(x) \in \mathbb{C}$. But there is a slightly different way to think about functions: we can think of ϕ as being defined by a subset

$$\Sigma \subset \mathbb{R} \times \mathbb{C} \tag{2.28}$$

such that for every $x \in \mathbb{R}$ there is exactly one element of the form $(x, y) \in \Sigma$, and then we denote $y \equiv \phi(x)$. Of course, this definition of a function is fully equivalent to the previous one, but note the “moral” difference: every point $x \in \mathbb{R}$ now has “its own” complex plane \mathbb{C}_x attached to it (consisting of all points of the form (x, z) for some z), and the point x is mapped to one point in that particular plane \mathbb{C}_x . This construction, where over every point in some *base space* (here: \mathbb{R}) we have a copy of some given *fibre* (here: \mathbb{C}) can be generalized into the concept of a *fibre bundle* that we began to discuss in the previous lecture.

For now, let us note that we can interpret our (ordinary) constant function $\phi(x)$ in this language as follows. Say we start with a single point $(x_0, y_0) \in \mathbb{R} \times \mathbb{C}$. The set Σ for the constant function $f(x) = y_0$ now contains all points (x, y_0) in Σ , and therefore it tells us how to *parallel transport* the point (x_0, y_0) along the \mathbb{R} -direction to points (x, y_0) in different fibers \mathbb{C}_x .

However, recall that we want to study topology, and therefore we want to make all of our constructions depend on specific choices of coordinates as little as possible. For example, instead of using a variable z_x to parameterize a certain \mathbb{C}_x , we may want to use a different variable $z'_x = e^{i\Lambda(x)}z_x$. Then, of course, our function $f(x)$, which was constant in the variables z_x , is no longer constant in the variables z'_x .

The upshot of this remark is the following: the related notions of being “constant” and of “parallel transporting”, require an extra piece of information: specific functions are only constant when we use a specific coordinate system on our fibers. The requirement

$$\frac{d}{dx}f(x) = 0 \tag{2.29}$$

is therefore a very coordinate dependent statement! However, as we have seen, for the requirement

$$D_x\phi(x) = 0 \tag{2.30}$$

where we have used a *covariant* derivative, this is not the case: if we change fibre coordinates by a coordinate dependent phase $e^{i\Lambda(x)}$, the left hand side only changes by an overall phase, and so this statement is still true in the new coordinates. This is the interpretation of the gauge field $A_\mu(x)$ that we were looking for: it gives us a way to define parallel transport

and constant functions (and, later on, other constant quantities like vectors, spinors, and so on) in a way which does not depend on local phase rotations of the fibre. An object which does this is called a *connection* in mathematics; the general concept of connections will be introduced in the next lecture.

Before proceeding, let us make a final remark about the above interpretation of the gauge field. The reader may wonder: why all this hassle? Couldn't we just have stayed with our original concept where a function maps \mathbb{R} into a single space like \mathbb{C} , with fixed coordinates chosen once and for all? The reason that this turns out not to be sufficient for physics purposes is that fields in physics are not always ordinary functions. Often, we may parallel transport a field value at a particular point around some topologically nontrivial closed loop in space-time, and *not* end up at the point where we started. As a toy model, our space could be a circle S^1 , and our field could take values in a fibre \mathbb{R}_x at each $x \in S^1$, but if we parallel transport (x, y) around the circle we might end up in the point $(x, -y)$. The fields satisfying this condition may be viewed as continuous “functions” on $[0, 1] \times \mathbb{R}$, but with $(0, y)$ identified with $(1, -y)$. That is, our “functions” (called *sections* in proper mathematic language) are subsets of a Möbius strip, which clearly is *not* topologically equivalent to $S^1 \times \mathbb{R}$. This situation occurs in physics quite often: field configurations are sections of “bundles” which are not simply a product of the base space with a given fibre, and as soon as this happens, we need to carefully determine what parallel transport means, and so we need a connection: a gauge field.

2.2.3 Field strength and curvature

Our example in the previous section, where we only have one “space”-dimension \mathbb{R} , is really rather simple. If we have two spatial dimensions, say parameterized by x^1 and x^2 , we can consider parallel transport in either direction. Now of course, these operations need *not* commute. Compare this to the transport of a tangent vector along, say, a sphere. A little experimenting shows that because the sphere is curved, transporting in one direction and then another need not give the same result as transporting along the same directions in the opposite order. In fact, if you know general relativity, you may know that this difference is exactly what the Riemann curvature tensor $R^\mu{}_{\nu\rho\sigma}$ measures: if the two directions in which we transport are themselves given by tangent vectors X^ρ and Y^σ , then the matrix

$$A^\mu{}_\nu = R^\mu{}_{\nu\rho\sigma} X^\rho Y^\sigma \tag{2.31}$$

gives (to first order in the length of the vectors) the rotation that relates a vector Z^ν parallel transported along X and then along Y to the same vector parallel transported along Y and then along X .

Something very similar turns out to happen in the case we studied above. Let us now consider two different spatial directions, labeled by x^μ and x^ν . Then the covariant derivative D_μ generates parallel transport in the x^ν -direction, and similarly for D_ν in the x^μ -direction. If we want to transport in both directions, we can choose two orders, and the difference

between those operations is of course $[D_\mu, D_\nu]\phi(x)$. A short computation (see exercise 4.1) now shows that

$$[D_\mu, D_\nu]\phi(x) = -iF_{\mu\nu}\phi(x) \quad (2.32)$$

That is, the field strength F measures the analogue of “curvature” for the parallel transport defined by the gauge field (connection) A_μ . This observation will play a role in the final part of this lecture, where we introduce Yang-Mills theories.

2.2.4 Covariant derivatives for more general gauge groups

To recover the Maxwell gauge field, we studied a complex scalar field $\phi(x)$ and its covariant derivative under $U(1)$ phase rotations $e^{i\Lambda(x)}$. This is a simple case of a much more generic situation: we can have a vector of fields, $\phi^i(x)$, on which a more general (matrix) Lie group acts – think of $SO(N)$, $U(N)$, $Sp(2N)$, etc. This is a situation we encounter a lot in physics: for example, the fundamental forces act on fields which transform as representations of $SU(3)$ for the strong nuclear force, $SU(2)$ for the weak nuclear force, and – as we have seen – $U(1)$ for the electromagnetic force.

Therefore, it is an interesting question whether we can also introduce covariant derivatives for this more general case. That is, suppose we have transformations of the form

$$\phi^i(x) \rightarrow (e^{ic^a(x)T_a})^i_j \phi^j(x) \quad (2.33)$$

where T_a are the generators of the Lie algebra for the group we are interested in, and $c^a(x)$ are position-dependent factors³. The entire computation is now exactly equivalent to the $U(1)$ case, and we find that the covariant derivative

$$D_\mu \phi^i(x) = \partial_\mu \phi^i(x) - iA_\mu^a(x)(T_a)^i_j \phi^j(x) \quad (2.34)$$

transforms in the same way as the field itself,

$$D_\mu \phi^i(x) \rightarrow (e^{ic^a(x)T_a})^i_j D_\mu \phi^j(x), \quad (2.35)$$

provided $A_\mu^a(x)$ transforms as

$$A_\mu^a(x) \rightarrow g(x)A_\mu^a(x)g^{-1}(x) + \partial_\mu g(x) \cdot g^{-1}(x), \quad (2.36)$$

where $g(x) = e^{ic^a(x)T_a}$ is the group element that also appears in (2.33).

What type of object does this define? Note that we can remove all indices of A by considering

$$A(x) \equiv A_\mu^a(x)T_a dx^\mu. \quad (2.37)$$

³We use “physics notation” here, where Lie group elements are obtained by exponentiating Lie algebra elements with a factor of i , analogous to the $U(1)$ example.

That is: A is a 1-form, but the coefficients of the different dx^μ are no longer x -dependent numbers: they have become x -dependent Lie algebra elements. We therefore say that A is a *Lie algebra valued 1-form*. In the next lecture, we will see that again, these objects play an important role in the theory of fibre bundles, where they describe the connections that define parallel transport.

2.2.5 Yang-Mills theory

Using covariant derivatives, we can once again construct field theory actions that are invariant under local Lie group symmetries. For example, suppose there is some bilinear form

$$Q(\phi, \chi) \equiv G_{ab} \phi^a \chi^b \quad (2.38)$$

which is invariant if we act on both ϕ and χ with elements of our Lie group:

$$Q(U\phi, U\chi) = Q(\phi, \chi) \quad (2.39)$$

In fact, many Lie groups can be defined through the existence of such a form: $SO(N)$, for example, consists of all linear transformations U such that the bilinear form with $G_{ab} = \delta_{ab}$ is invariant under U .

If we have such a bilinear form, described by some matrix G_{ab} , we can easily construct symmetric actions: any action of the form

$$\int d^4x \left(G_{ab} \eta^{\mu\nu} D_\mu \phi^a D_\nu \phi^b - V(G_{ab} \phi^a \phi^b) \right) \quad (2.40)$$

is now invariant under local Lie group transformations.

Can we also construct a kinetic term for the gauge field $A(x)$ itself in this case, just like we recovered Maxwell's action from the $U(1)$ covariant derivative? The answer is yes, and the easiest way to construct such a kinetic term is by recalling that in the $U(1)$ case, the commutator of two covariant derivatives gave us the field strength that we could build the action from. It turns out the same thing is true here, but that the field strength looks slightly different: one can compute (see exercise 4.1) that again

$$[D_\mu, D_\nu] \phi^i(x) = -i F_{\mu\nu}^a(x) T_a \phi(x) \quad (2.41)$$

but that now the field strength F has an extra term:

$$F_{\mu\nu}^a(x) = \partial_\mu A_\nu^a(x) - \partial_\nu A_\mu^a(x) - i f_{bc}^a A_\mu^b(x) A_\nu^c(x), \quad (2.42)$$

where f_{bc}^a are the structure constants of the Lie algebra, defined from the commutators of its generators:

$$[T_b, T_c] = f_{bc}^a T_a. \quad (2.43)$$

Again, we can write F as a Lie-algebra valued two-form $F(x) = F_{\mu\nu}^a(x)T_a dx^\mu \wedge dx^\nu$, which (exercise 4.1) makes (2.42) look much nicer:

$$F = dA - iA \wedge A \tag{2.44}$$

Here, the wedge product is extended to also include the multiplication of the generators T_a ; this is the reason that $A \wedge A$ is no longer zero (as it was for ordinary 1-forms), as two such generators generically do not commute.

Constructing an action for the gauge field now goes along exactly the same lines as for the Maxwell field: the natural action one can write down is

$$S \sim \int \text{Tr} (F \wedge \star F) \tag{2.45}$$

or, in terms of components,

$$S \sim \int \text{Tr} (F^{\mu\nu} F_{\mu\nu}). \tag{2.46}$$

The only difference now is that we need a trace operator (which for matrix groups is the ordinary matrix trace) to turn the Lie algebra-valued forms in the action into ordinary numbers. This action is called the *Yang-Mills action*; it plays an essential role in particle physics, as it describes the behaviour of the gauge fields for all fundamental forces between particles except for gravity.