Fibre Bundles and Spin Structures

Part 3: Differential Forms

David O'Connell

Okinawa Institute of Science and Technology

1919-1 Tancha, Onna-son, Okinawa 904-0495, Japan

On the surface it may seem as though differential forms are just another type of tensor. However, it turns out that they have an interesting application in topology, differential geometry, and physics. In this section we will discuss several enjoyable properties of differential forms. Specifically, we will see that

- Differential forms comprise an algebra,
- Differential forms are the correct object to be integrated on a manifold,
- Differential forms admit a generalisation of the differential df, and
- Differential forms can induce a series of algebraic structures that can be used to test the underlying topology of M.

We will explore all of these ideas, and more. Again, in two hours of lectures per week we do not have enough time to cover all of the amazing facts surrounding differential forms. As such, I've tried to include several more details in this note, so that you can dive down the differential rabbit hole in your spare time.

Contents

1 The Exterior Algebra

Last week we introduced the dual space V^* of a vector space V as the collection linear maps from V into the underlying field $\mathbb{R}$. We can generalise this idea into the notion of a tensor. In simple terms, tensors are machines that turn collections of vector spaces into numbers, in such a way that the linear structure of all of the vector spaces are somehow respected. Formally, a rank (r, s) tensor T is a multilinear map

$$T : \underbrace{V^* \times \cdots \times V^*}_{r-\text{many}} \times \underbrace{V \times \cdots \times V}_{s-\text{many}} \to \mathbb{R}.$$

By "multilinear", we mean that a tensor T is linear in any of its arguments, for example in the first argument we demand that

$$T(\bigwedge \omega_1, \omega_2, \cdots, \omega_r, v_1, \cdots, v_s) = \bigwedge T(\omega_1, \omega_2, \cdots, \omega_r, v_1, \cdots, v_s)$$

for any real number $\bigwedge$, and

$$\begin{aligned} T(\omega_1 + \eta, \omega_2, \cdots, \omega_r, v_1, \cdots, v_s) = \bigwedge T(\omega_1, \omega_2, \cdots, \omega_r, v_1, \cdots, v_s) \\ + T(\eta, \omega_2, \cdots, \omega_r, v_1, \cdots, v_s) \end{aligned}$$

for any other $\eta \in V^*$. We will denote the space of all rank (r, s) tensors by $T^{(r,s)}(V)$. We can make the following observations:

$$T^{(0,0)}(V) = \mathbb{R},\ T^{(0,1)}(V) = V^*,\ \text{and } T^{(1,0)}(V) = V.$$

The first and second items should be clear. The third equality comes from the fact that $(V^*)^* = V$.

There is an important operation that allows us to describe complicated tensors in terms of simpler ones. This operation is known as the *tensor product*, and is denoted by $\otimes$. For a pair of tensors $T \in T^{(r,s)}(V)$ and $T' \in T^{(r',s')}(V)$, the tensor product $T \otimes T'$ is a member of the space $T^{(r+r',s+s')}$, and is defined as follows:

$$(T \otimes T')(\omega_1, ..., \omega_{r+r'}, v_1, ..., v_{s+s'})$$
$$= T(\omega_1, ..., \omega_r, v_1, ..., v_s)T'(\omega_{r+1}, ..., \omega_{r+r'}, v_{s+1}, ..., v_{s+s'}),$$

where we use the product of real numbers on the right-hand side.

Denote by $T(V)$ the space of all tensors we can build from V. Together with the tensor product $\otimes$, the space $T(V)$ has the structure of an infinite-dimensional algebra. In fact, this algebra is so large and so general that many other interesting algebras can be constructed as subalgebras of $T(V)$.

1.1 Forms as Totally Antisymmetric Tensors

We will now restrict our attention to the rank $(0, k)$ tensors, that is, tensors of the form

$$T : \underbrace{V \times \cdots \times V}_{k-\text{many}} \to \mathbb{R}.$$

The tensor product of a rank $(0, k)$ and a rank $(0, l)$ tensor will be a rank $(0, k + l)$ tensor. Therefore, the collection of all such tensors forms a subalgebra of the full tensor algebra $T(V)$.

We say that a tensor T in $T^{(0,k)}(V)$ is antisymmetric if the interchange of any two arguments carries an overall minus sign. The simplest non-trivial case would be the rank $(0, 2)$-tensors. An antisymmetric $(0, 2)$ tensor would satisfy:

$$T(v, w) = -T(w, v).$$

The general formula is slightly more complicated, but mirrors the same idea. We say that a rank $(0, k)$-tensor is totally-antisymmetric if the equality

$$T(v_1, ..., v_k) = \text{sgn}(\pi)T(v_{\pi(1)}, ..., T_{\pi(k)})$$

holds, where here π is some permutation of the set $\{1, 2, ..., k\}$, and $\mathrm{sgn}(\pi)$ counts the sign of the permutation.[1] We will denote the space of all totally-antisymmetric rank $(0, k)$-tensors by $\bigwedge^k(V)$. It is also common to refer to $\bigwedge^k(V)$ as the space of k-forms.

1.2 The Wedge Product

The full tensor algebra has the tensor product $\otimes$ that can send pairs of tensors into a higher ranked one. There is a similar operation for k-forms, and it is defined to be the antisymmetrization of the tensor product. In the simplest case, wedge product of two 1-forms ω and η will be a 2-form defined by:

$$\omega \wedge \eta = \omega \otimes \eta - \eta \otimes \omega.$$

An immediate consequence is the following.

$$\omega \wedge \eta(v_1, v_2) = \omega(v_1)\eta(v_2) - \omega(v_2)\eta(v_1).$$

Clearly an exchange of arguments v_i will carry an overall minus sign in the wedge product.

As you might expect, the general formula for the wedge product of arbitrary forms is more complicated, and again involves permutations. We can use the previous formula to describe the wedge product of $\omega \in \bigwedge^k(V)$ and $\eta \in \bigwedge^l(V)$ as:

$$\omega \wedge \eta(v_1, ..., v_{k+l}) = \frac{1}{k!l!} \sum_{\pi} sgn(\pi) \omega \otimes \eta(v_{\pi(1)}, ..., v_{\pi(k+l)}), \tag{1}$$

where we sum over all permutations of the set of numbers $\{1, ..., k + l\}$. The strange coefficient is needed to avoid overcounting. To see this in action, let's

[1] The permutation can always be broken down into a series of permutations of pairs of indices – called a transposition. The sign of a permutation π is either $+1$ or -1 depending on whether the number of total transpositions is even or odd, respectively.

consider an example. Suppose that we have a two-form ω and a one-form η, with V a vector space of dimension at least 3. The product $\omega \wedge \eta$ is a 3-form:

$$\begin{aligned}
\omega \wedge \eta(v_1, v_2, v_3) &= \frac{1}{2!1!} \sum_{\pi} sgn(\pi) \omega \otimes \eta(v_{\pi(1)}, ..., v_{\pi(r+s)}) \\
&= \frac{1}{2}\Big(\omega(v_1, v_2)\eta(v_3) - \omega(v_1, v_3)\eta(v_2) - \omega(v_2, v_1)\eta(v_3) \\
&\quad - \omega(v_3, v_2)\eta(v_1) + \omega(v_2, v_3)\eta(v_1) + \omega(v_3, v_1)\eta(v_2)\Big) \\
&= \frac{1}{2}\Big(\omega(v_1, v_2)\eta(v_3) - \omega(v_1, v_3)\eta(v_2) + \omega(v_1, v_2)\eta(v_3) \\
&\quad + \omega(v_2, v_3)\eta(v_1) + \omega(v_2, v_3)\eta(v_1) - \omega(v_1, v_3)\eta(v_2)\Big) \\
&= \frac{1}{2}\Big(2\omega(v_1, v_2)\eta(v_3) - 2\omega(v_1, v_3)\eta(v_2) + 2\omega(v_3, v_2)\eta(v_2)\Big) \\
&= \omega(v_1, v_2)\eta(v_3) - \omega(v_1, v_3)\eta(v_2) + \omega(v_3, v_2)\eta(v_2)
\end{aligned}$$

There are many cool properties of the wedge product. Here is a summary of some of the facts that we will use throughout the course.

Theorem 1.1. *Let V be a finite dimensional vector space, and let ω and η be k and l-forms on V, respectively.*

1. $\wedge : \bigwedge^k(V) \times \bigwedge^l(V) \to \bigwedge^{k+l}(V)$ *is bilinear.*
2. $\omega \wedge \eta = -\eta \wedge \omega$.
3. $\omega \wedge \omega = 0$.
4. *If* $e_1, ..., e_n$ *is a basis of* V, *then* $\{e_{i_1} \wedge \cdots \wedge e_{i_k} \mid 1 \leq i_1, ..., i_k \leq n\}$ *is a basis of* $\bigwedge^k(V)$.
5. *The dimension of* $\bigwedge^k(V)$ *is* $\binom{n}{k} = \frac{n!}{k!(n-k)!}$, *where* $n = dim(V)$.

The proofs of these properties all follow in a routine manner from the definition of the wedge product and the linearity of the forms themselves.

1.3 Constructions of the Exterior Algebra

There are two ways to construct the exterior algeba:

1. a "bottom-up" construction, where we add layers of k-forms one-by-one, and
2. a "top-down" approach, where we quotient part of the Tensor algebra $T(V)$.

The first is simple enough – given our finite-dimensional vector space V, we may write $\Omega^k(V)$ for the space of totally-symmetric rank $(0, k)$ tensors on V. The full exterior algebra can be defined as

$$\bigwedge(V) := \bigoplus_{i=1}^{n} \bigwedge^{i}(V).$$

The $\oplus$ operation is known as the *direct sum* of vector spaces. It acts as a Cartesian product, where scalar multiplication and addition are performed componentwise.[2]

As for the second construction, things are not so simple. In the first week of lectures we saw that a manifold can always be constructed intrinsically by "gluing" together open charts along their transition maps. Formally, this gluing was achieved by specifying an equivalence relation, which was required to satsify certain properties in order to guarantee that the resulting space had the desired manifold structure.

The case of quotienting an algebraic structure is similar. We have a desired structure – addition, scalar multiplication, the zero vector, and the product operation. We want to somehow "glue" or "collapse" the algebra down into a smaller space. Formally, this is done by quotienting by something called an ideal. We will not delve into the definition of an ideal, so for now

[2] To make this more concrete: $V \oplus W$ has vectors of the form (v, w), where $v \in V$ and $w \in W$. The addition acts as $(v_1, w_1) + (v_2, w_2) = (v_1 +_V v_2, w_1 +_W w_2)$, with scalar multiplication defined similarly.

you may read "ideal" as "the thing that is needed to properly define quotients of algebras". We can define the exterior algebra $\bigwedge(V)$ as the quotient

$$\bigwedge(V) := {\tilde{T}(V)}\big/{I}$$

where $\tilde{T}(V)$ is the subalgebra of $T(V)$ consisting of only the rank $(0, k)$ tensors (for all k), and I is the ideal of $\tilde{T}(V)$ that is generated by elements of the form $v \otimes v$. Intuitively, the ideal I consists of the symmetric part of $\tilde{T}(V)$, and we are collapsing it to zero. The equivalence relation on $\tilde{T}(V)$ is defined by saying that two tensors T and S are equivalent if their difference $T - S$ is in I. This turns out to be an equivalence relation, and thus the quotient space is well-defined. The collection of equivalence classes inherit a linear structure from $\tilde{T}(V)$, by defining:

$$[T] + [S] := [T + S], \; r[T] := [rT], \text{ and } 0_{\bigwedge(V)} := [0].$$

The fact that all of these operations are well-defined[3] follows from the fact that I is an ideal. The tensor product $\otimes$ can also be transferred into the quotient space, by using $[x \otimes y]$. With a simple computation, we can see that the transferred tensor product becomes antisymmetric on equivalence classes:

$$\begin{aligned} 0 = [(T + S) \otimes (T + S)] &= [(T \otimes T) + (T \otimes S) + (S \otimes T) + (S \otimes +S)] \\ &= [T \otimes T] + [T \otimes S] + [S \otimes T] + [S \otimes +S] \\ &= [T \otimes S] + [S \otimes T], \end{aligned}$$

which implies that $[T \otimes S] = -[S \otimes T]$. This agrees with our previous definition of the wedge product as the antisymmetrization of the tensor product.

2 Tensors on Manifolds

Before discussing differential forms on manifolds, we will first need to talk about the general idea of tensors on manifolds. There are two situtations

[3]That is, independent on the particular choice of representative of each class

that we need to consider: tensors defined pointwise, and *fields* of tensors across the whole manifold.

2.1 Tensors

We can apply the general linear-algebraic constructions seen above to the case when V equals the tangent space T_pM. Therefore, we may obtain a hierarchy of tensors built on a manifold. As one might expect, we can use the coordinate bases of both T_pM and T_p^*M to induce a basis of the rank (r,s) tensors. Given a local chart (U,φ) with coordinates x^i, the collection of tensors of the form

$$dx^{i_1} \otimes \cdots \otimes dx^{i_r} \otimes \frac{\partial}{\partial x^{j_1}} \otimes \cdots \otimes \frac{\partial}{\partial x^{j_s}}$$

will give a basis for the space $T_p^{(r,s)}M$. This means that in local coordinates, a tensor will be expressed as

$$T = T_{i_1,\dots,i_r}^{j_1,\cdots,j_s} dx^{i_1} \otimes \cdots \otimes dx^{i_r} \otimes \frac{\partial}{\partial x^{j_1}} \otimes \cdots \otimes \frac{\partial}{\partial x^{j_s}},$$

where $T_{i_1,\dots,i_r}^{j_1,\cdots,j_s}$ is some real number that has dependence on p.

We can run through some arguments similar to those seen last week to determine the transformation properties of tensors. As one might expect, the components of a tensor will transform according to some combination of the transformation properties for tangent and cotangent vectors. For a pair of charts (U,φ) and (V,ψ) with coordinates x^i and $\tilde{x}^i$ respectively, the transformation properties of the components of a rank (r,s) tensor will carry r-many inverse Jacobians for the covector parts, and s-many Jacobians for the vector parts:

$$\tilde{T}_{i_1,\dots,i_r}^{j_1,\cdots,j_s} = \frac{\partial x^{k_1}}{\partial \tilde{x}^{i_1}} \cdots \frac{\partial x^{k_r}}{\partial \tilde{x}^{i_r}} \frac{\partial \tilde{x}^{j_1}}{\partial x^{l_1}} \cdots \frac{\partial \tilde{x}^{j_s}}{\partial x^{l_s}} T_{k_1,\dots,k_r}^{l_1,\cdots,l_s}.$$

Note that in the above expression we have suppressed the dependence on p and $\varphi(p)$ for the sake of simplification.

2.2 Tensor Fields

Throughout this course we will need to talk about *fields* of tensors. I have briefly mentioned these in the lectures, and it is worth expanding on some of the details. Before getting into it, we should note – I am being intentionally vague with some of the details here, the exact formulation of fields requires vector bundles, so we will discuss them properly in a few weeks time. For now, take everything with a grain of salt – the information below is still technically correct, but is very much a simplification.

A field of tensors will smoothly assign a tensor of fixed rank to each point in the manifold. The tensor in question needs to be one that is defined at the same point, e.g. for a vector field we assign something in T_pM to p, for every point. There are many similarities between the space of fields of a given type, and the space of vectors of a given type. Here is a brief summary of the algebraic properties of fields.

- A field of objects can be added and scaled. The addition acts pointwise in an obvious way, but the scaling is slightly more sophisticated.
- At every point, a tensor has components $T^{r_1\cdots}_{s_1\cdots}$ which are real numbers. A field carries these components across the whole manifold M. So, at every point p we associate some real number, and this association is smooth. We have already met such objects several times already – they are smooth functions, i.e. members of $C^\infty(M)$. This means that the components of a tensor *field* are really more like a smooth function, i.e. a real number that carries a smooth dependence on p. Put differently, we scale tensor fields by smooth functions instead of fixed real numbers.
- Since we can add tensor fields and scale them, it is extremely tempting to call them a vector space. However, this is not strictly true, because in order for a structure to be a vector space it's scalars need to form

an algebraic field.[4] The real numbers are an algebraic field, but the space $C^\infty(M)$ is not – it is merely a ring. Therefore the space of rank (r, s)-tensor fields has a different structure: it is called a *module* over $C^\infty(M)$.

- We can always see tensor fields as vector spaces over $\mathbb{R}$ by changing the scalars from $C^\infty(M)$ to scalars in $\mathbb{R}$. This can be seen as stipulating that we only scale by constant functions instead of arbitrary ones. Although this allows us to view tensor fields as real-valued vector spaces, there is a trade-off: there are no finite local bases anymore.

- It is not always the case that the module of tensor fields will admit a globally-defined basis – i.e. a collection of fields that form a basis in each pointwise space. However locally they always will, since we can use local charts to force the fields into $\mathbb{R}^n$, where things are well-behaved.

Aside from this, we should also remark that it is not always the case that tensor fields are non-zero everywhere. It is sometimes unavoidable that certain fields will have to vanish at least somewhere (cf. the Hairy Ball theorem). This is due to global structural issues of vector bundles that we will see later on in the course.

Note that we can always use the local charts (U, φ) to create a local basis for fields. In the case of vector fields, say, we would have a globally-defined vector v that can be represented locally as

$$v = v^i \frac{\partial}{\partial x^i},$$

where $v^i : U \to \mathbb{R}$ are smooth. This representation is valid on all of U.

[4] This in unfortunate terminology, but is standard. The term "field" used in algebra is distinct from the sense of the word "field" used in physics/differential geometry. After some research, it seems that the word "field" was first used in algebra in an 1893 paper by E. H. Moore. Maxwell was describing fields at least 30 years prior.

3 Differential Forms

In Section 1 we saw that any finite-dimensional vector space V admits an exterior algebra, which consists of k-forms. In the manifold setting, we can define an analogous system at each point by taking $V = T_pM$. There is thus an exterior algebra $\bigwedge T_pM$ at every point in the manifold. A differential k-form is defined to be a *field* of k-forms on M.

3.1 Basic Properties

We define a differential form as a field of k-forms on M. In distinction to the vector space setting, we will denote the space of k-form fields as $\Omega^k(M)$. A pointwise k-form $\omega_p \in \bigwedge_p^k(M)$ will eat k-many tangent vectors and deliver a real number. In contrast, a differential k-form ω will eat k-many vector *fields* and deliver a smooth function. At low ranks, we have that

- the space $\Omega^0(M)$ is equal to $C^\infty(M)$, and
- the space $\Omega^1(M)$ is equal to the space of covector fields on M.

Locally, ω can be written in the coordinate basis

$$\omega = \omega_{i_1 i_2 \ldots i_p} dx^{i_1} \wedge dx^{i_2} \wedge \cdots \wedge dx^{i_p} = \omega_I dx^I,$$

where we write dx^I as shorthand for the wedge product $dx_1^i \wedge \cdots \wedge dx_p^i$. As with covectors, the coefficients ω_I can be determined by their action on the coordinate basis vectors:

$$\omega_I = \omega\left(\frac{\partial}{\partial x^{i_1}}, \cdots, \frac{\partial}{\partial x^{i_p}}\right).$$

3.2 Pullbacks and Pushforwards

The differential of a smooth map $f : M \to N$ sends tangent vectors in T_pM to tangent vectors in $T_{f(p)}N$. If this is the case, we can say that the vector $df_p(v)$ is the *pushforward* of the vector v. There is an analogous notion for covectors, except that the map goes in the opposite direction – if we have a smooth map covectors of the image can be *pulled back* to covectors of the domain. Let $f : M \to N$ be a smooth map, and let ω_p be a member of $T^*_{f(p)}N$. Then we may define the pullback $f^*(\omega) : T_pM \to \mathbb{R}$ by

$$f^*(\omega)_p(v) := \omega(df_p(v)).$$

This is well-defined and linear, therefore $f^*\omega$ is a member of T_pM.

It will be important for us to understand whether or not we can transfers fields of objects similarly. The short story is that we cannot pushforward vector fields since

- the map f might not be injective, meaning that two vectors from two different tangent spaces get mapped to the same tangent space in N, and
- f might not be surjective, meaning that there are points in N that do not have an assigned tangent vector.

Interestingly, covector fields can always be pulled back. This is mostly due to the definition of a function – the domain of a smooth map f is the whole of M, and any single point p in M is mapped to a *unique* point in N. We can define the pullback pointwise just as we did before. It is beyond the scope of our current discussion to prove that the pullback of a covector field is smooth, but this is indeed true.[5]

3.3 Orientations

Suppose we have a vector space V. We can change the basis of V by using a linear map $A : V \to V$ that is invertible. Geometrically speaking, we can use the determinant of A to see how a unit volume transforms under A. In particular, the sign of $det(A)$ tells us whether the vector space V has been

[5]See Lee, Proposition 11.26.

"flipped over", that is, whether or not the orientation of V has changed. We therefore say that

$$A \text{ is } \begin{cases} \text{orientation preserving if } det(A) > 0 \\ \text{orientation reversing if } det(A) < 0 \end{cases}.$$

Note that these are the only two options, since any invertible matrix has non-zero determinant.

The same idea can be carried across to manifolds. Suppose that M is a smooth manifold and (U, φ) and (V, ψ) are two charts of M. As we saw last week, we can use the coordinate bases of T_pM with respect to both U and V, to obtain a coordinate transformation

$$\frac{\partial}{\partial \tilde{x}^i} = \frac{\partial x^j}{\partial \tilde{x}^i} \frac{\partial}{\partial x^j}$$

of bases, where we have simplified the expression by supressing the p and $\varphi(p)$ that occurred previously. In a direct analogue to the setting of arbitrary vector spaces, we make the following definition.

Definition 3.1. *A smooth manifold M is called oriented if for any pair of charts (U, φ) and (V, ψ) at a point p, the Jacobian $\frac{\partial x^j}{\partial \tilde{x}^i}(\varphi(p))$ has positive determinant.*

In general, the atlas of a smooth manifold M is too large to be oriented. What usually happens is that we restrict the atlas down to a smaller set of charts which makes M oriented. We refer to M as *orientable* if can be turned into an oriented manifold under some restriction of charts.

The above definition of orientability might seem a little abstract. There is a more-intuitive definition of orientation in terms of differential forms. As a consequence of Theorem 1.1 there are no differential forms of degree higher than $dim(M)$.[6] This means that the exterior power $\bigwedge^n M$ is the last non-trivial collection of forms. According to Theorem 1.1, this space of highest degree forms is one-dimensional. A differential $dim(M)$-form has a special name – it is called a *top form*. The orientability of M can be characterised by specifying a non-vanishing top form – for the full details see Lee Chapter 15.

4 Integration of Forms

In high school we are introduced to the definite integral

$$\int_a^b f(x)dx.$$

This is usually interpreted as integration of the function f with respect to x along the interval $[a, b]$. Once you grow a bit, you learn that this definite integral can be approximated by computing discrete strips of area. We can split the interval $[a, b]$ into N-many linearly-ordered pieces $[x_i, x_{i+1}]$, and then

$$\int_{[a,b]} f(x)dx \approx \sum_i^{N-1} f(x_i)\Delta(x_i),$$

where $\Delta(\cdot)$ computes the difference between adjacent x_i. The overall integral is then some limit of the above expression, when $N \to \infty$. Although this picture is somewhat correct, the more honest interpretation is that $f(x)dx$ is a coordinate expression for a field of 1-forms in $\mathbb{R}$, and these are the real object that is integrated over.

The geometric intuition here is that the one-form dx is a field of infinitesimal, oriented line elements. These are infinitely-summed over to yield a

[6] If this is not clear to you, consider the following: a k-form for $k = dim(M) + 1$ would have as a basis $dx' \wedge dx^I$. The form dx' can itself be represented as a linear combination of the dx^I, so if we expand out, the linearity of the wedge product guarantees that we will get $dx^i \wedge dx^i$ somewhere.

finite result. Moreover, the sign of the one-form can indicate the orientation of the integral:

$$\int_a^b f(x)dx = \int_b^a -(f(x)dx) = -\int_b^a f(x)dx.$$

The case for higher-dimensional integrals is similar – a 2-form is a field of infinitesimal, signed area elements, a 3-form is a field of infinitesimal, signed volume elements, and so on. A more-technical reason that we integrate over differential forms is that the resulting integral is independent of any choice of coordinates, and thus is well-defined.

4.1 Line Integrals

Suppose that we have a smooth curve with endpoints $\gamma : [a, b] \to M$, and consider a 1-form ω that is defined on M. We can define the line integral of ω in M by pulling ω back to a one-form in $\mathbb{R}$:

$$\int_\gamma \omega := \int_a^b \gamma^* \omega.$$

Note that we need to assume that the covector field ω is non-zero on the image of γ.

Example 4.1. *Consider the differential form* $\omega := \frac{xdy - ydx}{x^2+y^2}$*, defined on* $\mathbb{R}^2$ *minus the origin. We will compute the integral of* ω *around the closed curve* $\gamma : S^1 \to \mathbb{R}^2\backslash\{0\}$*. Suppose that* γ *has parameter* t*, and consider the circle given by* $x = \cos(t)$ *and* $y = \sin(t)$*. It follows that* $dx = -\sin(t)$ *and* $dy = \cos(t)$*. The pullback* $\gamma^*\omega$ *is then given by*

$$\gamma^* \omega = \frac{\cos(t)(\cos(t)dx) - \sin(t)(-\sin(t)dt)}{\cos^2(t) + \sin^2(t)} = dt.$$

We can then compute the integral of ω *around* γ *as*

$$\int_\gamma \omega = \int_{S^1} \gamma^* \omega = \int_{S^1} dt = 2\pi.$$

4.2 Integration of Top Forms

The multivariable case of integration is slightly more complicated. Instead of a curve, we will use a full-dimensional domain of integration $D \subset M$, and a top-form ω of M.[7] In order to yield a finite answer for the eventual integral, we will need the top form ω to be zero everywhere outside of D.[8] We will consider two levels of generality:

- the case where D is contained in a single chart, and
- the case where D is compact, but not contained in a single chart.

If the domain of integration is contained within a single chart, then we can simply transfer our integral to Euclidean space by using the pullback

$$\int_{D \subset U} \omega = \int_{\varphi(D)} (\varphi^{-1})^* \omega,$$

where the second expression is a regular multivariable integral over $\mathbb{R}^n$.

In the case that our domain of integration D is not contained within a single chart, we can cover D with finitely-many open charts and compute parts of the integral using the above expression. We can then glue these results together via a partition of unity to define the total integral over D. For the details, see Lee Chapter 16.

[7] Note that this is not strictly required – we can also integrate k-forms over k-dimensional submanifolds of M, but we will not need this in this course.

[8] There is a slightly more sophisticated version of this: we define the support of ω to be $supp(\omega) = \{p \in M \mid \omega_p \neq \emptyset\}$. We can then show that the topological closure of this set can be contained within a compact domain of integration.

5 Simplicial Homology

In this section we will introduce a basic form of homology known as *simplicial homology*. The idea is to turn a manifold M into a system of higher-dimensional triangles (called simplices), and to determine an algebraic system for dealing with the interrelation of the triangles. The goal is to express topological holes in terms of linear-algebraic data, with an eye towards actual computations. The building blocks for simplicial homology are the following:

1. Create a triangulation Δ_M of M.

2. Create a chain complex $C_k(M)$, which is a vector space generated from the individual k-simplices. Geometrically, sums of basis vectors correspond to chaining simplices of the same dimension together.

3. The map δ which sends each chain in $C_k(M)$ to its topological boundary in $C_{k-1}(M)$ can be seen as a linear map. The chain complex then looks like

$$C_0 \xleftarrow{\delta_1} C_1 \xleftarrow{\delta_2} \cdots \xleftarrow{\delta_k} C_k \xleftarrow{\delta_{k+1}} \cdots \xleftarrow{\delta_n} C_n$$

   where here we have indexed the δ map for simplicity.

4. Since the boundary of a boundary is zero, the boundary map δ satisfies $\delta^2 = 0$.

5. Holes are characterised by n-dimensional loops which do not bound an $n+1$-dimensional disk. In simplicial terms, holes are characterised by closed chains (called cycles) that have empty boundary.

6. The k^{th} simplicial homology group is defined to be the k-cycles modulo the k-chains that are boundaries of something in $C_{k+1}(M)$. Formally

$$H_k^S(M) := \frac{Ker(\delta_k)}{Im(\delta_{k+1})}.$$

At some point in the future we will need to understand the formal details of simplicial homology. Before we need it, I will include a more thorough introduction to Simplicial Homology in this section. Watch this space!

5.1 Simplicial Complexes and Triangulations

5.2 Chains, Cycles and Boundaries

5.3 The Homology Groups

5.4 Some Worked Examples

6 De Rham Cohomology

Differential forms have a deep relationship with the topological structure on which they are defined. This can go in two directions:

1. The space of differential forms can be turned into a topological invariant that can used to tell two topologies apart, and

2. the topology of the manifold can pose as an obstruction for the guaranteed existence of particular types of differential forms.

In this section we will begin to explore this interesting relationship. We will start by exploring a special derivative operator that can be defined on differential forms, and then we will use this to construct a series of algebraic objects known as the *de Rham cohomology groups*.

6.1 The Exterior Derivative

So far we have only defined the differential of a smooth map at a single point. As a matter of fact, we can interpret df to be a smooth covector field – we simply assign the pointwise differential df_p to each point p in M.

As we saw in Section 3, the functions $C^\infty(M)$ can be identified with $\Omega^0(M)$. The space of covector fields can be identified with $\Omega^1(M)$. Therefore we can think of the differential as a map $d : \Omega^0(M) \to \Omega^1(M)$. In this section we will extend the notion of the differential to a map which raises the degree of k-forms by one. This map is known as the *exterior derivative*.

Suppose that we have a k-form ω. Then in a coordinate basis induced from a chart, we can locally represent ω as the linear combination

$$\omega = \omega_{i_1 \cdots i_k} dx^{i_1} \wedge \cdots \wedge dx^{i_k} = \omega_I dx^I,$$

where the component $w_{i_1 \cdots i_k}$ is some function in $C^\infty(U)$. We can define the exterior derivative locally by

$$d(\omega_I dx^I) = (d\omega_I) \wedge dx^I.$$

We will now expand on some of the properties of the exterior derivative.

Theorem 6.1. *The exterior derivative d on $\Omega(M)$ satisfies the following properties.*

1. *d is $\mathbb{R}$-linear.*

2. *d satisfies the graded Leibniz law: $d(\omega \wedge \eta) = (d\omega) \wedge \eta + (-1)^k \omega \wedge (d\eta)$, where $\omega \in \Omega^k(M)$.*

3. *$d(d\omega) = 0$.*

The proofs of these theorems are generally quite routine. You can find the full arguments in Lee. Just to convince you that some of these do actually hold, we will argue locally. Suppose that we have $\omega = \omega_i dx^i$ and $\eta = \eta_j dx^j$. Then:

$$\begin{aligned}
d(\omega \wedge \eta) &= d\left((\omega_i dx^i)(\eta_j dx^j)\right) \\
&= d\left(\omega_i \eta_j dx^i \wedge dx^j)\right) \\
&= d(\omega_i \eta_j) \wedge dx^i \wedge dx^j \\
&= (d(\omega_i)\eta_j + \omega_i(d\eta_j)) \wedge dx^i \wedge dx^j \\
&= (d(\omega_i)\eta_j) \wedge dx^i \wedge dx^j + (\omega_i(d\eta_j)) \wedge dx^i \wedge dx^j \\
&= \eta_j d\omega_i \wedge dx^i \wedge dx^j + \omega_i d\eta_j \wedge dx^i \wedge dx^j \\
&= (d\omega_i \wedge dx^i) \wedge (\eta_j dx^j) + (-1)^k (\omega_i dx^i) \wedge (d\eta_j \wedge dx^j) \\
&= (d\omega) \wedge \eta + (-1)^k \omega \wedge d\eta
\end{aligned}$$

You can construct a similar expansion to show that $d^2 = 0$, the key idea in this case is that second-order partial derivatives commute, but the differentials dx^i anti-commute, and thus there is pairwise cancellation in the local expansion of $d(d\omega)$.

We can use the exterior derivative to identify two important types of differential form.

Definition 6.2. *A differential k-form is called*

1. *closed if* $d\omega = 0$, *and*
2. *exact if* $\omega = d\eta$.

According to Theorem 6.1 every exact form is closed, but the converse is not always true.

6.1.1 Stoke's Theorem

Recall the fundamental theorem of calculus:

$$\int_{[a,b]} f(x)dx = F(b) - F(a).$$

The function F is the antiderivative of f. This means that $\frac{dF}{dx} = f(x)$, or put differently, that $dF = f(x)dx$. We can paraphrase the fundamental theorem of calculus as:

$$\int_{[a,b]} dF = \int_{[a,b]} f(x)dx = F(b) - F(a) = \int_{\{a,b\}} F = \int_{\partial([a,b])} F.$$

As we can see, we can replace an exterior derivative with a boundary to get the same answer for the integral. There is a generalisation of this theorem to arbitrary differential forms, known as the generalised Stoke's theorem, which reads:

$$\int_D d\omega = \int_{\partial D} \omega.$$

You can find the proof in Lee Chapter 16. The theorem has several nice consequences, including the following.

Proposition 6.3. *Let* M *be a compact oriented manifold without boundary. If* ω *is an exact top form, then the integral* $\int_M \omega$ *vanishes.*

Proof. Suppose that $\omega = d\eta$. Then by Stoke's theorem we have

$$\int_M \omega = \int_M d\eta = \int_{\partial M} \eta = 0,$$

since M has no boundary. □

According to the above result, the two-form defined in Example 4.1 is not exact, since its integral around a closed loop is non-zero.

6.2 de Rham Cohomology

As we saw in Section 5, homology can be seen as algebraic theories for counting holes in a topological space. In simplicial homology, we turn our space into a system of triangles, and then formally count the inter-relatedness of the triangles to eventually obtain a collection of vector spaces $H_k^S(M)$. In the case of de Rham cohomology we will do something similar, but we will use differential forms together with d as our starting point instead of simplices and the boundary operator ∂.

The construction itself is reliant on the enjoyable properties of the exterior derivative listed in Theorem 6.1. In analogy to the chain complex of simplicial homology, we can construct the following series of spaces.

$$\Omega^0(M) \xrightarrow{d_0} \Omega^1(M) \xrightarrow{d_1} \cdots \xrightarrow{d_{k-1}} \Omega^k(M) \xrightarrow{d_k} \cdots \xrightarrow{d_{n-1}} \Omega^n(M)$$

This is now called a *cochain* complex, since the arrows are raising degrees instead of lowering them.

Since the exterior derivative operator is a linear map between the spaces $\Omega^k(M)$, we can meaningfully consider the image and kernel of d. Indeed, in a direct analogy to Section 5, we can define

$$H^k_{dR}(M) := \frac{Z^k(M)}{B^k(M)} = \frac{Ker(d_k)}{Im(d_{k-1})} = \frac{\{\text{Closed} \quad k-\text{forms}\}}{\{\text{Exact} \quad k-\text{forms}\}}.$$

Formally, this quotient is described by an equivalence relation $\cong$, where $\omega \cong \eta$ if the difference $\omega - \eta$ is exact. By construction this means that $[\omega] \in H^k_{dR}(M)$ can be described as

$$[\omega] = \{\eta \in Z^k(M) \mid \omega = \eta + d\theta\},$$

that is, two closed k-forms lie in the same cohomology class if they differ by an exact form.

6.3 Relationship to Simplicial Homology

The above construction gives us a collection of cohomology spaces $H^k_{dR}(M)$ induced from the differential forms of M. However, it is not immediately clear that the de Rham cohomology groups are related to the topology of M at all. We will now argue that there is in fact a deep connection to the underlying topology of M. We will do this by relating de Rham cohomology to the simplicial homology introduced in Section 5. If seen for the first time, it seems that this is a surprising connection. However, we have already seen clues to the connection in the form of Stoke's theorem:

$$\int_C d\omega = \int_{\partial C} \omega.$$

We can rewrite this in a more suggestive manner by considering the integral as a function:

$$\int : C_k(M) \times \Omega^k(M) \to \mathbb{R},$$

Where $C_k(M)$ is the space of k-chains induced from the triangulation of M, and $\Omega^k(M)$ is the space of differential k-forms on M. This is known as a *pairing*, and very much mirrors the properties of an inner product. In particular, this pairing is bilinear. In this notation, we may write Stoke's theorem as

$$\langle C, d\omega\rangle = \langle \partial C, \omega\rangle.$$

The content of de Rham's theorem states that this map can be converted into a map between homology and cohomology, that is, if we restrict this pairing by integration down to the relevant equivalent classes, then we will obtain a non-degenerate bilinear form that can be used to induce an isomorphism of vector spaces. Let us define a map $\mathcal{I} : H^S_k(M) \times H^k_{dR}(M) \to \mathbb{R}$ by

$$\mathcal{I}([C], [\omega]) := \langle C, \omega\rangle = \int_C \omega.$$

In order for this map to be well-defined, it should be independent on the choice of representative in each equivalence class. However, this turns out to be the case. Suppose that we have a cycle $\tilde{C}$ in $[C]$ and a closed form $\tilde{\omega}$

in $[\omega]$. Then by construction, $\tilde{C} = C + \partial B$ where B is a $(k+1)$-chain, and $\tilde{\omega} = \omega + d\eta$, where η is a $(k-1)$-form. We can then use Stoke's theorem to conclude that:

$$\begin{aligned}
\langle \tilde{C}, \tilde{\omega} \rangle &= \langle C + \partial B, \omega + d\eta \rangle \\
&= \langle C, \omega \rangle + \langle C, d\eta \rangle + \langle \partial B, \omega \rangle + \langle \partial B, d\eta \rangle \\
&= \langle C, \omega \rangle + \langle \partial C, \eta \rangle + \langle B, d\omega \rangle + \langle B, d^2\eta \rangle \\
&= \langle C, \omega \rangle + \langle \partial C, \eta \rangle + \langle B, d\omega \rangle + \langle B, d^2\eta \rangle \\
&= \langle C, \omega \rangle + \langle 0, \eta \rangle + \langle B, 0 \rangle + \langle B, 0 \rangle \\
&= \langle C, \omega \rangle.
\end{aligned}$$

It follows that our induced integral map $\mathcal{I}$ is indeed well-defined. The amazing connection between de Rham and simplicial can now be stated properly.

Theorem 6.4 (de Rham's Theorem)**.** *If M is compact then $H^k_{dR}(M)$ and $H^S_k(M)$ are both finite-dimensional. Moreover, the pairing*

$$\mathcal{I} : H^S_k(M) \times H^k_{dR}(M) \to \mathbb{R}$$

is bilinear and non-degenerate.

It follows from this theorem that $H^S_k(M)$ is isomorphic to $H^k_{dR}(M)$, and moreover $H^k_d R(M)$ is the dual vector space of $H^S_k(M)$. Under this reading, we can interpret the exterior derivative as the dual to the topological boundary.