

3.2. Differential forms. The line integral of a 1-form over a curve is a very nice kind of integral in several respects. In the spirit of differential geometry, it does not require any additional structure, such as a metric. In practice, it is relatively simple to compute. If two curves c and c' are very close together, then $\int_c \alpha \approx \int_{c'} \alpha$. The line integral appears in a nice integral theorem.

This contrasts with the integral $\int_c f ds$ of a function with respect to length. That is defined with a metric. The explicit formula involves a square root, which often leads to difficult integrals. Even if two curves are extremely close together, the integrals can differ drastically.

We are interested in other nice integrals and other integral theorems. The generalization of a smooth curve is a submanifold. A p -dimensional *submanifold* $\Sigma \subseteq M$ is a subset which looks locally like $\mathbb{R}^p \subseteq \mathbb{R}^n$.

We would like to make sense of an expression like

$$\int_{\Sigma} \omega.$$

What should ω be? It cannot be a scalar function. If we could integrate functions, then we could compute volumes, but without some additional structure, the concept of “the volume of Σ ” is meaningless.

In order to compute an integral like this, we should first use a coordinate system to chop Σ up into many tiny p -dimensional parallelepipeds (think cubes). Each of these is described by p vectors which give the edges touching one corner. So, from p vectors v_1, \dots, v_p at a point, ω should give a number $\omega(v_1, \dots, v_p)$ which is something like “the volume according to ω ”. If one of these is the sum of two vectors, then the “volumes” add:

$$\omega(u + w, v_2, \dots) = \omega(u, v_2, \dots) + \omega(w, v_2, \dots).$$

This should be true even if u and w are not tangent to Σ , because adding a small bump to Σ should not effect the integral significantly.




FIGURE 2. A parallelepiped determined by 3 vectors and the effect of adding vectors

This implies that $\omega(v_1, \dots, v_p)$ is a linear function of each of the vectors v_1, \dots, v_p . If any two of the vectors are equal, then the parallelepiped collapses and its “volume” is 0:

$$\omega(w, w, v_3, \dots, v_p) = 0.$$

This implies that ω is (completely) *skew symmetric*, i.e., if any two of these vectors are interchanged then the sign of the result changes. For example,

$$0 = \omega(v_1 + v_2, v_1 + v_2, v_3, \dots) = \omega(v_1, v_2, v_3, \dots) + \omega(v_2, v_1, v_3, \dots).$$

If we use p vector *fields*, then $\omega(v_1, \dots, v_n) \in C^\infty(M)$ is $C^\infty(M)$ -linear in each argument.

This leads to the following definition. A p -form (or *differential form of degree p*) is a $C^\infty(M)$ -multilinear, skew symmetric map $\omega : \mathcal{X}(M)^p \rightarrow C^\infty(M)$.

This has a simple meaning in terms of tensor notation. A p -form is just a skew symmetric tensor $\omega_{i_1 \dots i_p}$ with p subscript indices. These definitions are related by

$$\omega(v_1, \dots, v_p) = \omega_{i_1 \dots i_p} v_1^{i_1} \dots v_p^{i_p}.$$

The set of p -forms on M is denoted $\Omega^p(M)$. The set of all differential forms on M is denoted $\Omega^\bullet(M)$. Note that 0-forms are just functions.

If $\deg \omega > n$ then $\omega = 0$. This is because any sequenced of more than n integers $1, \dots, n$ must have a repetition.

In general, how many p -forms are there? The nonvanishing components of a p -form are labeled by sequences of p different integers between 1 and n . Up to a sign, the order of these numbers does not matter. So, the *independent* components are labeled by sets of p numbers between 1 and n . Therefore, the dimension of the space of p -forms at a point of an n -dimensional manifold is $\binom{n}{p} = \frac{n!}{p!(n-p)!}$. In particular, there are as many p -forms as $n - p$ -forms.

3.2.1. Exterior products. The *exterior product* (or *wedge product*) of p 1-forms is a p -form. It is defined by multiplying them (to give a tensor) and summing over all $p!$ permutations with a $-$ sign on the odd permutations. The exterior product is denoted by \wedge . For example

$$(\alpha \wedge \beta)(u, v) = \alpha(u) \beta(v) - \beta(u) \alpha(v)$$

$$(\alpha \wedge \beta)_{ij} = \alpha_i \beta_j - \beta_i \alpha_j$$

$$\begin{aligned} (\alpha \wedge \beta \wedge \gamma)_{ijk} &= \alpha_i \beta_j \gamma_k - \beta_i \alpha_j \gamma_k + \beta_i \gamma_j \alpha_k - \gamma_i \beta_j \alpha_k + \gamma_i \alpha_j \beta_k - \alpha_i \gamma_j \beta_k \\ &= 6\alpha_{[i} \beta_j \gamma_{k]} \end{aligned}$$

where the bracket means skew symmetrization of indices.

The relationship between tensor notation and p -form notation is

$$\omega = \frac{1}{p!} \omega_{i_1 \dots i_p} dx^{i_1} \wedge \dots \wedge dx^{i_p}. \quad (3.4)$$

So, any differential form is a sum of exterior products of 1-forms. From this, we can define the exterior product of differential forms by requiring the properties of associativity and distributivity:

$$\omega \wedge (\sigma \wedge \lambda) = (\omega \wedge \sigma) \wedge \lambda$$

$$\omega \wedge (\sigma + \lambda) = \omega \wedge \sigma + \omega \wedge \lambda.$$

The exterior product with a function is just multiplication by that function. The exterior product is also *anticommutative*:

$$\omega \wedge \sigma = (-1)^{\deg \omega \deg \sigma} \sigma \wedge \omega.$$

In tensor notation, the exterior product of a p -form and a q -form is

$$(\omega \wedge \sigma)_{i_1 \dots i_{p+q}} = \frac{(p+q)!}{p!q!} \omega_{[i_1 \dots i_p} \sigma_{i_{p+1} \dots i_{p+q}]}$$

Wald gives this formula as the definition of the exterior product, although that makes the numerical factor seem capricious. The index notation is more common in physics texts.

Example. Note that in 3 dimensions, the dimensions of the spaces of differential forms at a point are 1, 3, 3, and 1. In the standard geometry of vectors on Euclidean \mathbb{R}^3 , we identify 1-forms with vectors. We can also identify 2-forms with vectors and 3-forms with scalars. So, any differential form is identified with a scalar or vector.

In this way, the exterior products of forms turn out to be familiar operations. The exterior product of two 1-forms is the cross product, $\mathbf{v} \times \mathbf{w}$. The exterior product of a 1-form and a 2-form is the dot product, $\mathbf{v} \cdot \mathbf{w}$. The exterior product of three 1-forms is the triple product of vectors, $\mathbf{u} \cdot (\mathbf{v} \times \mathbf{w})$. The exterior product of two 2-forms is 0.

Remark. Two 1-forms α and β are parallel if and only if $\alpha \wedge \beta = 0$. In general, if $\deg \sigma$ is odd, then $\sigma \wedge \sigma = -\sigma \wedge \sigma = 0$. On the other hand, if $\deg \omega$ is even, then $\omega \wedge \omega$ need not be 0.

3.2.2. Interior products. The *interior product* of a vector v and a p -form ω is a $p-1$ -form $v \lrcorner \omega$. If we think of a p -form as a multilinear map, then $v \lrcorner \omega = \omega(v, \dots)$, i.e., v is inserted as the first argument in ω . In tensor notation

$$(v \lrcorner \omega)_{i_1 \dots i_p} = v^j \omega_{ji_1 \dots i_p}.$$

The interior product of a vector and a 0-form (scalar) is understood to be 0, since there are no (-1) -forms. The interior product of a vector and a 1-form is just another way of writing the pairing, $v \lrcorner \alpha = \langle \alpha, v \rangle = \alpha(v)$.

Although the interior product involves no derivatives, it satisfies a Leibniz rule

$$v \lrcorner (\omega \wedge \sigma) = (v \lrcorner \omega) \wedge \sigma + (-1)^{\deg \omega} \omega \wedge (v \lrcorner \sigma). \quad (3.5)$$

3.2.3. Orientation and volume forms. The space of n -forms at a point of an n -dimensional manifold is only 1-dimensional. Locally, there are only two kinds of nonvanishing n -forms. There is an ambiguity in the notion of integration; nothing tells us *a priori* which n -forms should integrate to positive numbers and which should give negative numbers. An *orientation* of an n -dimensional manifold is a choice of which n -forms to consider positive.

Not all manifolds are orientable. The Möbius strip and Klein bottle are simple non-orientable examples.

A coordinate system (on an oriented manifold) is *oriented* if $dx^1 \wedge \dots \wedge dx^n$ is positive. A transformation between oriented coordinate systems has positive Jacobian. An orientation can be defined by an atlas of oriented charts.

A positive n -form is often called a *volume form*, because it can be used to define a notion of volume for open subsets of M .

Example. If M is an oriented manifold with metric tensor g_{ij} , then *the* volume form is defined in any oriented coordinate system by

$$\epsilon := \sqrt{|\det g|} dx^1 \wedge \cdots \wedge dx^n.$$

Given a volume form ϵ , we can identify a vector v with the $n-1$ -form $v \lrcorner \epsilon$. Some vectors (namely, *currents*) are most naturally thought of as $n-1$ -forms.

Exercise 3.5. Using (3.5), prove that for $v \in \mathcal{X}(M)$, $\alpha \in \Omega^1(M)$ and $\epsilon \in \Omega^n(M)$,

$$\alpha \wedge (v \lrcorner \epsilon) = \langle \alpha, v \rangle \epsilon.$$

Remark. If M is an oriented manifold with a metric, then we can identify p -forms and $n-p$ -forms. Starting with some $\sigma \in \Omega^p(M)$, first raise its indices with the metric, and then define the *Hodge dual* $*\sigma \in \Omega^{n-p}(M)$ by

$$*\sigma_{i_1 \dots i_{n-p}} = \frac{1}{p!} \sigma^{j_1 \dots j_p} \epsilon_{j_1 \dots j_p i_1 \dots i_{n-p}}.$$

3.2.4. Exterior derivatives. The exterior derivative of a p -form σ is a $p+1$ -form $d\sigma$. Exterior differentiation is a linear map $d : \Omega^p(M) \rightarrow \Omega^{p+1}(M)$. On functions, it is just the gradient. It satisfies the Leibniz rule

$$d(\omega \wedge \sigma) = d\omega \wedge \sigma + (-1)^{\deg \omega} \omega \wedge d\sigma$$

and

$$dd\sigma = 0.$$

From these properties, we can compute the exterior derivative of the product of a function and p gradients. Applying this to eq. (3.4), we get

$$d\sigma = \frac{1}{p!} d(\sigma_{i_1 \dots i_p}) dx^{i_1} \wedge \cdots \wedge dx^{i_p} = \frac{1}{p!} \sigma_{i_1 \dots i_p, j} dx^j \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_p}.$$

This gives the explicit formula in a coordinate system,

$$d\sigma_{j i_1 \dots i_p} = (p+1) \sigma_{[i_1 \dots i_p, j]}, \quad (3.6a)$$

or using a covariant derivative

$$d\sigma_{j i_1 \dots i_p} = (p+1) \sigma_{[i_1 \dots i_p | j]}. \quad (3.6b)$$

Exercise 3.6. Show that the exterior derivative, as given by eq. (3.6a), has the properties (Liebniz and $d^2 = 0$) postulated above.

Example. In Euclidean \mathbb{R}^3 , the exterior derivative of a 1-form is equivalent to the curl. The exterior derivative of a 2-form is equivalent to the divergence.

$$\begin{array}{ccccccc} \text{0-form} & \xrightarrow{d} & \text{1-form} & \xrightarrow{d} & \text{2-form} & \xrightarrow{d} & \text{3-form} \\ \text{scalar} & \xrightarrow{\text{gradient}} & \text{vector} & \xrightarrow{\text{curl}} & \text{vector} & \xrightarrow{\text{divergence}} & \text{scalar} \end{array}$$

More generally, if ϵ is the volume form on a Riemannian manifold, then

$$d(v \lrcorner \epsilon) = (\nabla \cdot v) \epsilon$$

where $\nabla \cdot v = v^i_{|i}$ is the covariant divergence. In fact, a volume form alone (without a metric) defines a notion of divergence.

Exercise 3.7. Show that the exterior derivative of a 1-form is given by the formula,

$$d\alpha(u, v) = u[\alpha(v)] - v[\alpha(u)] - \alpha([u, v]).$$

(It is sufficient to prove this for $\alpha = f dg$, since any 1-form is a sum of such expressions.) Similar expressions hold in higher degrees.

3.2.5. *Closed forms.* A differential form is called *closed* if its exterior derivative is 0. It is called *exact* if it is the exterior derivative of another form. The exterior derivative of an exterior derivative is always 0 (succinctly, $d^2 = 0$) so exact forms are always closed.

Conversely, the *Poincaré lemma* states that over a contractable region (such as all of \mathbb{R}^n) any closed form is exact. So, closed forms are “locally exact”.

This is not true globally. For example, any n -form on an n -dimensional manifold must be closed. However, the volume form on the sphere S^2 of radius 1 is not exact. The vector potential in Exercise 3.3 is a closed but not exact 1-form.

Example. In standard vector calculus on \mathbb{R}^3 , a vector field is a gradient iff its curl vanishes. It is a curl iff its divergence vanishes.

Example. If we take the exterior derivative of the first law of thermodynamics (3.1), then $ddE = 0$ implies that

$$dT \wedge dS = dp \wedge dV.$$

The *Maxwell relations* can easily be derived from this equation. For example, if we write dT and dV in terms of dS and dp , and insert these into the above equation, then this immediately gives

$$\left(\frac{\partial T}{\partial p}\right)_S = \left(\frac{\partial V}{\partial S}\right)_p$$

(times $dS \wedge dp$).

Example. Suppose that $\alpha \in \Omega^1(M)$ is normal to a family of hypersurfaces. That is, through any point of M , there is exactly 1 of these hypersurfaces, and α is normal to it there. In some neighborhood of any point, there exists a coordinate system in which the hypersurfaces are the level sets of a coordinate (call it x). In such a coordinate system, α must be of the form

$$\alpha = f dx.$$

So, $d\alpha = df \wedge dx$ and

$$\alpha \wedge d\alpha = 0.$$

This equation is the necessary and sufficient condition for a 1-form to be locally “hypersurface orthogonal”.

3.2.6. *Pull-backs.* We previously described pull-backs of functions. If $\varphi : M \rightarrow N$ is a smooth map and $f \in C^\infty(N)$ a smooth function, then the pull-back of f is

$$\varphi^*f := f \circ \varphi. \tag{3.7}$$

The concept of pull-back extends to differential forms. A map $\varphi^* : \Omega^\bullet(N) \rightarrow \Omega^\bullet(M)$ is defined by three properties: For functions, it is defined by eq. (3.7); it is compatible with the exterior derivative

$$\varphi^*d\omega = d(\varphi^*\omega)$$

and exterior product

$$\varphi^*(\sigma \wedge \omega) = \varphi^*\sigma \wedge \varphi^*\omega.$$

In coordinates, it is very easy to compute a pull-back. This is just a generalization of changing coordinates in a 1-form. If $\{y^i\}$ are the coordinates on N , then write ω in terms of dy^i 's. The map φ gives the y^i 's as functions of coordinates on M . The pull-back $\varphi^*\omega$ is just given by inserting these formulas into the expression for ω .

If $i : \Sigma \hookrightarrow M$ is the inclusion of a submanifold, then the pull-back of functions by i is just restriction. For differential forms, it is a little more subtle. If $\alpha \in \Omega^1(M)$ is normal to $\Sigma \subset M$, then $i^*\alpha = 0$. For example, if f is constant along Σ , then df is normal to Σ and

$$i^*df = d(i^*f) = 0.$$

So, the pull-back i^* restricts differential forms to Σ and throws away components normal to Σ .

Exercise 3.8. The graph of $z = x^2 + y^2$ is a submanifold of \mathbb{R}^3 . Using x and y as coordinates on this, compute the pull-back of

$$\omega = (x \, dy - y \, dx) \wedge dz.$$