Alternating Tensor Fields in Differential Geometry: Differential Forms and Multivector Fields

McKenzie Lamb

September 13, 2007

One of the first hurdles faced by beginning students in differential geometry is understanding tensor fields–in particular, differential forms and multivector fields. The notion of a tensor field is somewhat different in flavor from most of the ideas students typically see as undergraduates. Furthermore, the notation used for differential forms and multivector fields, while ultimately practical, can be terribly confusing to someone who is encountering the subject for the first time. While there are a number of textbooks on elementary differential geometry (e.g., [1], [3], [4]), many students find that the detailed, airtight style of a textbook renders this subject, in particular, a confusing mess.

The goal of this paper is to provide a simple, intuitive explanation of differential forms and multivector fields as well as a brief description of tensor fields in general. I have included a number of examples and explanations that seem to be missing from many textbooks. The intended audience is beginning graduate students in mathematics. It is assumed that the reader understands the basics of smooth manifolds and tangent bundles. These prerequisites aside, this text is (in theory) more or less self-contained. There are several additional comments about Poisson and symplectic geometry which may or may not be meaningful to people who have not studied this subject.

Note: Throughout this paper, I will use angle brackets to denote the pairing between a vector space and its dual. That is, if V is a vector space and V^* is its dual, for any $v \in V$, $\phi \in V^*$, instead of writing $\phi(v)$ for the result of applying the function ϕ to v, I will write $\langle \phi, v \rangle$. The motivation for this notation is that, given an inner product $\langle \cdot, \cdot \rangle$ on V, every element $\phi \in V^*$ can be written in the form $v \mapsto \langle v_0, \cdot \rangle$ for some vector $v_o \in V$. The dual space V^* can therefore be identified with V (via $\phi \mapsto v_0$). The reason we use the term, *pairing*, is that duality of vector spaces is symmetric. By definition, ϕ is a linear function on V, but v is also a linear function on V^*. So instead of thinking of applying ϕ to v, we can just as well consider v to be acting on ϕ. The result is the same in both cases. For this reason, it makes more sense to think of a *pairing* between V and V^*: If you pair an element from a vector space with an element from its dual, you get a real number.

1 Multilinear Algebra

1.1 Tensor Products

Tensor products are described in all sorts of creative and horrible ways. If you read about tensor products in an algebra book, for instance, you will get a very general definition, possibly involving universal mapping principles and other abstract nastiness. Fortunately, if we restrict our attention to vector spaces–which is often possible in differential geometry–the situation is fairly simple. All of the vector spaces in this discussion are finite-dimensional with base field $\mathbb{R}$.

The following is one way to think about tensor products. Suppose V and W are vector spaces with bases $\{v_1, \ldots, v_n\}$ and $\{w_1, \ldots, w_m\}$, respectively. Then $V \otimes W$ is the formal span over $\mathbb{R}$ of the elements $\{v_i \otimes w_j \,|\, 1 \leq i \leq n, 1 \leq j \leq m\}$. In other words, an element of $V \otimes W$ has the form

$$\sum_{1 \leq i \leq n,\, 1 \leq j \leq m} c_{ij}(v_i \otimes w_j),$$

for some constants c_{ij}. That's it, and that's all. This point of view is obviously basis dependent, and there are situations in which other perspectives are more practical, but this is a very concrete way to think about tensor products of vector spaces. For example, if $V = \mathrm{span}\{dx_1, dx_2\}$ and $W = \mathrm{span}\{dy_1, dy_2\}$ then the elements of $V \otimes W$ have the form

$$c_{11}dx_1 \otimes dy_1 + c_{12}dx_1 \otimes dy_2 + c_{21}dx_2 \otimes dy_1 + c_{22}dx_2 \otimes dy_2,$$

for some constants c_{ij}, $1 \leq i, j \leq 2$.

Another approach (which is described on Wikipedia, for example) is to define $V \otimes W$ to be the set of elements $\{v \otimes w \,|\, v \in V,\, w \in W\}$ modulo the following equivalence relations:

$$\begin{aligned} v \otimes (w_1 + w_2) &= v \otimes w_1 + v \otimes w_2 \\ (v_1 + v_2) \otimes w &= v_1 \otimes w + v_2 \otimes w \\ (cv) \otimes w = v \otimes (cw) &= c(v \otimes w). \end{aligned}$$

It is easy to see that this quotient space is a vector space, and that, for any bases $\{v_1, \ldots, v_n\}$ and $\{w_1, \ldots, w_m\}$, for V and W, $\{v_i \otimes w_j \,|\, 1 \leq i \leq n, 1 \leq j \leq m\}$ is a basis for $V \otimes W$.

Now suppose V is a vector space and V^* is its dual. The elements of $V^* \otimes V^*$ can be identified with real-valued functions on $V \times V$ as follows: If $\phi \otimes \psi \in V^* \otimes V^*$ and $v, w \in V$, then

$$\langle \phi \otimes \psi, (v, w) \rangle := \langle \phi, v \rangle \cdot \langle \psi, w \rangle.$$

The relations above imply that $\phi \otimes \psi$ is a *multilinear*–i.e., linear in each slot–map on $V \times V$. In fact, it's easy to see that $V^* \otimes V^*$ is the space of *all* multilinear maps on $V \times V$. Alternatively, $V^* \otimes V^*$ can be thought of as the space of *linear* maps on $V \otimes V$. It amounts to the same thing. Therefore, $V \otimes V$ and $V^* \otimes V^*$ are dual as vector spaces.

More generally, one can consider *mixed* tensors, such as $V \otimes V^* \otimes V$. An element $v_1 \otimes \phi_1 \otimes w_1$ of this space acts on the space $V^* \otimes V \otimes V^*$ in the obvious way: For $\phi_2, \psi_2 \in V$ and $v_2 \in V^*$,

$$\langle v_1 \otimes \phi_1 \otimes w_1, \phi_2 \otimes v_2 \otimes \psi_2 \rangle := \langle v_1, \phi_2 \rangle \cdot \langle \phi_1, v_2 \rangle \cdot \langle w_1, \psi_2 \rangle.$$

Thus, $V^* \otimes V \otimes V^*$ and $V \otimes V^* \otimes V$ are dual as vector spaces.

One simple application of mixed tensors that seems to be fairly useful is the identification of $V^* \otimes V$ with the space $\mathrm{End}(V)$ of linear transformations on V. Given an element $\sum_i \phi_i \otimes v_i \in V^* \otimes V$, we can produce a linear transformation $f : V \mapsto V$ by defining

$$f(v) := \sum_i \phi_i(v) \cdot v_i.$$

The simplest way to show that every linear transformation on V can be produced in this way is to choose a basis for V. The details are left to the reader.

1.2 Wedge Products

Of particular interest in differential geometry is the subspace $V^* \wedge V^* \subset V^* \otimes V^*$ of so-called *alternating tensors*. For any $\phi, \psi \in V^*$, define the element $\phi \wedge \psi \in V^* \otimes V^*$ by its action on elements of the form $v \otimes w \in V \otimes V$:

$$\langle \phi \wedge \psi, v \otimes w \rangle := \det \begin{pmatrix} \langle \phi, v \rangle & \langle \psi, v \rangle \\ \langle \phi, w \rangle & \langle \psi, w \rangle \end{pmatrix} = \langle \phi, v \rangle \cdot \langle \psi, w \rangle - \langle \phi, w \rangle \cdot \langle \psi, v \rangle.$$

Of course, typical elements of $V \otimes V$ have the form $\sum_i v_i \otimes w_i$, and similarly for $V^* \otimes V^*$, so the above definition must be extended by linearity.

Notice that interchanging v and w interchanges the rows and interchanging ϕ and ψ interchanges the columns of the above matrix. Since we are taking the determinant, this has the effect of multiplying the result by -1. Linear maps on $V \otimes V$ with this property are called *alternating 2-tensors*. Of course, $\phi \wedge \psi$ is clearly a linear map, again by properties of the determinant. One can prove that the space of all alternating 2-tensors on V is the vector space $V^* \wedge V^* = \{\phi \wedge \psi \,|\, \phi, \psi \in V^*\}$. It can also be shown that if $\{\phi_1, \ldots, \phi_n\}$ is a basis for V^*, then $\{\phi_i \wedge \phi_j \,|\, 1 < i \leq j \leq n\}$ is a basis for $V^* \wedge V^*$. (For a given $1 \leq i \neq j \leq n$, since $\phi_j \wedge \phi_i = -\phi_i \wedge \phi_j$, only one of $\phi_i \wedge \phi_j$ and $\phi_j \wedge \phi_i$ should be included in a basis–hence the condition, $i < j$). Thus, if $\{\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z}\}$ is a basis for V, then the elements of $V \wedge V$ have the form

$$c_1 \frac{\partial}{\partial x} \wedge \frac{\partial}{\partial y} + c_2 \frac{\partial}{\partial x} \wedge \frac{\partial}{\partial z} + c_3 \frac{\partial}{\partial y} \wedge \frac{\partial}{\partial z},$$

where c_1, c_2, c_3 are real numbers. Furthermore, $V^* \wedge V^*$ is dual to $V \wedge V$. That is, we have defined $\phi \wedge \psi$ as an alternating, linear function on $V \otimes V$, but it could just as well be thought of as a linear function on $V \wedge V$. In fact, $\phi \wedge \psi$ can also be thought of as an alternating, multilinear function on $V \times V$.

Remark 1.1. Note that if V and W are different vector spaces, the tensor product of V and W makes sense, but the wedge product does not.

In general, for $\phi_1, \ldots, \phi_n \in V^*$ and $v_1, \ldots, v_n \in V$, we define $\phi_1 \wedge \cdots \wedge \phi_n$ by its action on $\otimes^n V$, the n-fold tensor product of V with itself:

$$\langle \phi_1 \wedge \cdots \wedge \phi_n, (v_1, \ldots v_n) \rangle := \det\left(\langle \phi_j, v_i \rangle\right) = \det \begin{pmatrix} \langle \phi_1, v_1 \rangle, & \ldots & \langle \phi_n, v_n \rangle \\ \vdots & & \vdots \\ \langle \phi_1, v_n \rangle & \ldots & \langle \phi_n, v_n \rangle \end{pmatrix}.$$

Again, this action is extended by linearity. As before, interchanging v_i with v_j or ϕ_i and ϕ_j changes the sign of the result. More generally, any permutation of the ϕ's or the v's can be decomposed into a product of transpositions. For a given permutation, the parity of the number of transpositions in this decomposition is called the *sign* of the permutation. Thus, an even permutation of the ϕ's or v's leaves the result unchanged, and an odd permutation changes the sign of the result. This is what it means to be alternating in the case of n-fold tensors. For example, suppose $n = 3$, and let $\phi_1, \phi_2, \phi, 3 \in V^*$, $v_1, v_2, v_3 \in V$. Then

$$\langle \phi_1 \wedge \phi_2 \wedge \phi_3, v_1 \otimes v_3 \otimes v_2 \rangle = -\langle \phi_1 \wedge \phi_2 \wedge \phi_3, v_1 \otimes v_2 \otimes v_3 \rangle$$

because $v_1 \otimes v_3 \otimes v_2$ and $v_1 \otimes v_2 \otimes v_3$ differ by a single transposition. On the other hand,

$$\langle \phi_1 \wedge \phi_2 \wedge \phi_3, v_3 \otimes v_1 \otimes v_2 \rangle = \langle \phi_1 \wedge \phi_2 \wedge \phi_3, v_1 \otimes v_2 \otimes v_3 \rangle$$

since $v_3 \otimes v_1 \otimes v_2$ and $v_1 \otimes v_2 \otimes v_3$ differ by a pair of transpositions. As noted, the results would be the same if the ϕ's were permuted instead of the v's. As in the case of 2-fold products, an element of $\wedge^n(V^*)$ can be thought of as (1) an alternating, multilinear function on the n-fold Cartesian product of V with itself, (2) an alternating linear function on $\otimes^n V$, or (3) a linear function on $\wedge^n(V)$.

2 Multivector Fields and Differential Forms

Given a manifold M, associated to each point $p \in M$ is a vector space, the tangent space at p, denoted T_pM. The cotangent space at p is the dual of T_pM, denoted T_p^*M. I will refer to the collection of all tangent spaces as the tangent bundle, TM, and the collection of all cotangent spaces as the cotangent bundle, T^*M. Suppose $(x_1, \ldots, x_n)$ are coordinates on an open set $U \subset M$. Then for any $p \in U$, $\left\{ \left(\frac{\partial}{\partial x_1}\right)_p, \ldots, \left(\frac{\partial}{\partial x_n}\right)_p \right\}$ is a basis for T_pM and $\{(dx_1)_p, \ldots (dx_n)_p\}$ is a basis for T_p^*M.

A vector field on M is a function X which picks out from each tangent space T_pM a vector X_p in such a way that the vectors X_p "vary smoothly" as you move around the manifold. For simplicity of notation suppose M is a 3-manifold. For a given point $p \in U$ a vector in T_pM looks like

$$c_1 \left(\frac{\partial}{\partial x}\right)_p + c_2 \left(\frac{\partial}{\partial y}\right)_p + c_3 \left(\frac{\partial}{\partial z}\right)_p,$$

where c_1, c_2, and c_3 are some constants. A vector field on M has the same form, but the coefficients vary from point to point. Thus, in local coordinates on $U \subset M$, a vector field on M has the form

$$f_1(p) \left(\frac{\partial}{\partial x}\right)_p + f_2(p) \left(\frac{\partial}{\partial y}\right)_p + f_3(p) \left(\frac{\partial}{\partial z}\right)_p,$$

where f_1, f_2, f_3 are functions from U to $\mathbb{R}$. In coordinates, "varying smoothly" just means that the coefficient functions f_1, f_2, and f_3 are smooth on U. In practice, the p subscripts are often omitted, and the f_i's are written in x, y, z coordinates. For example,

$$V = y\frac{\partial}{\partial x} - x\frac{\partial}{\partial y} + 2\frac{\partial}{\partial z}$$

is a vector field on $\mathbb{R}^3$.

Remark 2.1. Note that the vector field V is tangent to any cylinder centered at the z-axis. It could therefore be considered as a vector field on any of these submanifolds. This kind of situation can lead to a great deal of confusion. Let M be the cylinder of radius 1 centered at the z-axis. The coordinates x, y, z are not local coordinates on M. They can't be, of course, because M is 2-dimensional, not 3-dimensional. The vectors $\frac{\partial}{\partial x}$, $\frac{\partial}{\partial y}$, and $\frac{\partial}{\partial z}$ do not, in general, live in the tangent space to M at a given point, but any tangent vector can be written as a linear combination of $\frac{\partial}{\partial x}$, $\frac{\partial}{\partial y}$, and $\frac{\partial}{\partial z}$. For many purposes, it is convenient to write vector fields in *ambient* coordinates, as V is here, but it is important to remember that ambient coordinates are not local coordinates.

Instead of choosing a vector in each tangent space, a *bivector field* chooses an element of $\wedge^2(T_pM)$ at each point $p \in M$. If M is 3-dimensional, for instance, a bivector at a point p has the form

$$c_1\frac{\partial}{\partial x} \wedge \frac{\partial}{\partial y} + c_2\frac{\partial}{\partial x} \wedge \frac{\partial}{\partial z} + c_3\frac{\partial}{\partial y} \wedge \frac{\partial}{\partial z}$$

in local coordinates, where c_1, c_2, c_3 are constants. In local coordinates, a bivector field has the form

$$f_1(p)\left(\frac{\partial}{\partial x} \wedge \frac{\partial}{\partial y}\right)_p + f_2(p)\left(\frac{\partial}{\partial x} \wedge \frac{\partial}{\partial z}\right)_p + f_3(p)\left(\frac{\partial}{\partial y} \wedge \frac{\partial}{\partial z}\right)_p,$$

where f_1, f_2, f_3 are real-valued functions on (an open subset of) M. Again, the p subscripts are usually omitted. For example,

$$\pi = z\frac{\partial}{\partial x} \wedge \frac{\partial}{\partial y} - y\frac{\partial}{\partial x} \wedge \frac{\partial}{\partial z} + x\frac{\partial}{\partial y} \wedge \frac{\partial}{\partial z}$$

is a bivector field on R^3 (or on an open subset of a 3-manifold).

In general, a k-vector field on a manifold chooses an element of $\wedge^k(T_pM)$ at each point $p \in M$. In local coordinates, a k-vector field has the form

$$\sum_{i_1 < \cdots < i_k} f_{i_1,\ldots,i_k}\frac{\partial}{\partial x_{i_1}} \wedge \cdots \wedge \frac{\partial}{\partial x_{i_k}},$$

where each $f_{i_1,\ldots,i_k}$ is a function on M.

Remark 2.2. A vector field on a manifold might also be described as a (smooth) *section* of the tangent bundle. The notion of a section is nicely captured by Figure 1 (see below). Similarly, a k-vector field is a section of $\wedge^k TM$, and a k-form is a section of $\wedge^k(T^*M)$. In general, the space of sections of a vector bundle E is often denoted by $\Gamma(V)$. The space of 3-forms, for instance, on a manifold M is denoted by $\Gamma(\wedge^3(T^*M))$.

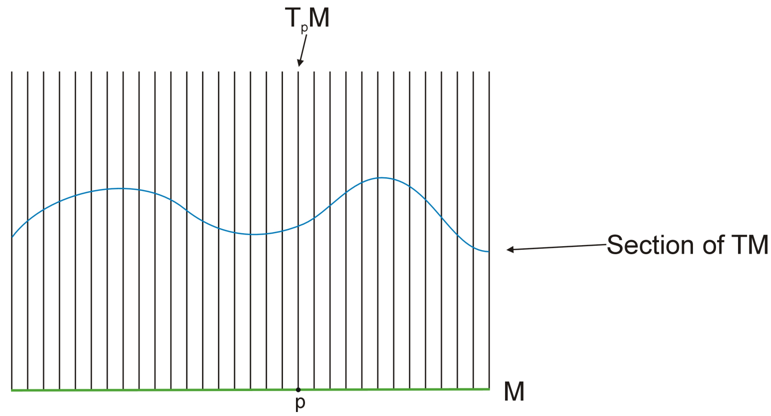

Figure 1: A section of the tangent bundle of a manifold.

2.1 Sharp and Flat Maps

Let V be a vector space, and let $\pi \in \wedge^2 V$. For the moment, it will be useful to think of π as an alternating, multilinear map on $V^* \times V^*$. Fix $\phi \in V^*$. Then for any $\psi \in V^*$, $\langle \pi, (\phi, \psi)$ is a real number. Therefore, thinking of ψ as a variable,

$$\psi \mapsto \langle \pi, (\phi, \psi) \rangle$$

is a linear map from V^* to $\mathbb{R}$. This map is more commonly denoted by

$$\langle \pi, (\phi, \cdot) \rangle$$

and is described as the *contraction* of π with ϕ. Changing perspectives, we can also think of ϕ as a variable. Then for any $\phi \in V^*$, $\langle \pi, (\phi, \cdot) \rangle$ is a linear map from V^* to $\mathbb{R}$–i.e., an element of V. Thus, π can be thought of as a linear map from V^* to V, defined by

$$\phi \mapsto \langle \pi, (\phi, \cdot) \rangle.$$

The same thing works at the bundle level. A bivector field π (a section of $\wedge^2 TM$) induces a map that takes in a 1-form (a section of the cotangent bundle) and produces a vector field (a section of the tangent bundle)–i.e., a map from $\Gamma(T^*M)$ to $\Gamma(T^M)$. At each point p on the manifold, π_p is a linear map from T_p^*M to T_pM, as defined above. Given a 1-form η, the map defined by the bivector field π takes the covector η_p to the vector $\langle \pi_p, (\eta_p, \cdot) \rangle$ for each $p \in M$. This map is often denoted by $\pi^\#$. (Note: In Poisson geometry, the convention is to contract in the second slot, which changes the sign of the result. For this paper, I will contract in the first slot.)

For example, consider the bivector field

$$\pi = z\frac{\partial}{\partial x} \wedge \frac{\partial}{\partial y} - y\frac{\partial}{\partial x} \wedge \frac{\partial}{\partial z} + x\frac{\partial}{\partial y} \wedge \frac{\partial}{\partial z}$$

on $\mathbb{R}^3$. Let $\eta = zdx + xdy + ydz$. Then

$$\begin{aligned} \pi^\#(\eta) &= z\left(z\frac{\partial}{\partial y} - y\frac{\partial}{\partial z}\right) + x\left(-z\frac{\partial}{\partial x} + x\frac{\partial}{\partial z}\right) + y\left(y\frac{\partial}{\partial x} - x\frac{\partial}{\partial y}\right) \\ &= (y^2 - xz)\frac{\partial}{\partial x} + (z^2 - xy)\frac{\partial}{\partial y} + (x^2 - zy)\frac{\partial}{\partial z}. \end{aligned}$$

Similarly, a 2-form ω on M defines a map $\omega^\flat$ which takes in a vector field and produces a 1-form by contraction.

In fact, a bivector field (resp. 2-form) can be defined by its # (resp. $\flat$) map. For example, suppose π is a bivector field such that $\pi_p^\#$ is invertible at each point $p \in M$ (a bivector field with this property is called *nondegenerate*). Then the map

$$\omega^\flat : \Gamma(TM) \to \Gamma(T^*M),$$

defined at each point to be the inverse of $\pi_p^\#$, determines a 2-form ω. Conversely, if ω is a 2-form such that $\omega^\flat$ is invertible at each point, then the inverse of $\omega^\flat$ defines a map from $\Gamma(T^*M)$ to $\Gamma(TM)$, which in turn determines a bivector field.

Remark 2.3. If ω is a symplectic form (a closed, nondegenerate 2-form), then the bivector field π determined by the inverse of $\omega^\flat$ is Poisson. The Jacobi identity for π is equivalent to the closed condition on ω. On the other hand, given a *nondegenerate* Poisson bivector field π, the inverse of $\pi^\#$ defines a symplectic form. However, Poisson bivectors are not necessarily nondegenerate. Therefore, a symplectic manifold is a special case of a Poisson manifold.

2.2 Pushing Forward (Multi)Vector Fields

Given two manifolds, M and N, let φ be a smooth map from M to N. At each point $p \in M$, φ induces a map φ_{*p} from T_pM to $T_{\varphi(p)}N$. If tangent vectors are identified with equivalence classes of curves through p, then φ_{*p} is defined by

$$[\gamma] \mapsto [\varphi \circ \gamma],$$

where γ is a curve in M such that $\gamma(0) = p$. Equivalently, if tangent vectors are identified with derivations on smooth, real-valued functions, then φ_{*p} is defined by

$$(\varphi_{*p}(X))(f)_p := X(f \circ \varphi)_p,$$

where X is a derivation on M and f is a smooth function on N. If φ is injective, the disjoint union $\coprod \varphi_{*p}$ defines a map φ_* from $\Gamma(TM)$ to $\Gamma(TN)$. If φ is not injective, then the map φ_* will generally not be well-defined on vector fields: If $X \in \Gamma(TM)$ and $q \in N$ has two inverse images, p_1 and p_2, then $\varphi_{*p_1}(X_{p_1}$ and $\varphi_{*p_2}(X_{p_2}$ may not be equal, in which case $\varphi_*(X)$ is not well-defined.

Now let $x_1, \ldots, x_n$ be local coordinates on a neighborhood U of p. Writing $X_p = \sum_i c_i \frac{\partial}{\partial x_i}$, $\varphi_*(X_p)$ can be computed by applying the Jacobian of ϕ to the column vectors associated to the vectors $\frac{\partial}{\partial x_i}$. That is,

$$\varphi_{*p}(X_p) = \sum_i c_i D\varphi_p \left(\frac{\partial}{\partial x_i}\right),$$

where $D\varphi_p$ is the Jacobian matrix of φ at p. If φ is injective, for a vector field $X = \sum_i f_i \frac{\partial}{\partial x_i}$, we have

$$\varphi_*(X) = \sum_i (f_i \circ \varphi^{-1}) D\varphi \left(\frac{\partial}{\partial x_i}\right).$$

The map φ_{*p} can be extended to k-vector fields by defining

$$\varphi_{*p}(X_p^1 \wedge \cdots \wedge X_p^k) := \varphi_{*p}(X_p^1) \wedge \cdots \wedge \varphi_{*p}(X_p^k).$$

If φ is injective, the map φ_* is a well-defined map on multivector fields. For example, consider the bivector field $\pi = z\frac{\partial}{\partial x} \wedge \frac{\partial}{\partial y} - y\frac{\partial}{\partial x} \wedge \frac{\partial}{\partial z} + x\frac{\partial}{\partial y} \wedge \frac{\partial}{\partial z}$ on $\mathbb{R}^3$. Since π is just the standard volume form on S^2 with $d(\cdot)$ replaced by $\frac{\partial}{\partial(\cdot)}$, it's easy to see that π is tangent to S^2 (that is to say, it lives in the wedge product of the tangent bundle to the sphere with itself). We should, therefore, be able to write π in stereographic coordinates. The computation is as follows: We are using stereographic projection from the north pole, given by

$$\varphi : S^2 \setminus \{\text{north pole}\} \to \mathbb{R}^2 : (x, y, z) \mapsto \left(\frac{x}{1-z}, \frac{y}{1-z}\right) = (u, v),$$

and

$$\varphi^{-1}(u, v) = \left(\frac{2u}{1+u^2+v^2}, \frac{2v}{1+u^2+v^2}, \frac{u^2+v^2-1}{1+u^2+v^2}\right) = (x, y, z).$$

For brevity of notation, let $l = 1 + u^2 + v^2$, and note that $1 - z = \frac{2}{l}$. Then

$$D\varphi = \begin{pmatrix} \frac{1}{1-z} & 0 & \frac{x}{(1-z)^2} \\ 0 & \frac{1}{1-z} & \frac{y}{(1-z)^2} \end{pmatrix},$$

which gives

$$\begin{aligned}
\varphi_*\left(\frac{\partial}{\partial x}\right) &= \frac{1}{1-z}\frac{\partial}{\partial u} = \frac{l}{2}\frac{\partial}{\partial u}, \\
\varphi_*\left(\frac{\partial}{\partial y}\right) &= \frac{1}{1-z}\frac{\partial}{\partial v} = \frac{l}{2}\frac{\partial}{\partial v}, \\
\varphi_*\left(\frac{\partial}{\partial z}\right) &= \frac{x}{(1-z)^2}\frac{\partial}{\partial u} + \frac{y}{(1-z)^2}\frac{\partial}{\partial v} = \frac{ul}{2}\frac{\partial}{\partial u} + \frac{vl}{2}\frac{\partial}{\partial v}.
\end{aligned}$$

Thus,

$$\begin{aligned}
\varphi_*(\pi) &= \frac{l-2}{2}\left(\frac{l^2}{4}\frac{\partial}{\partial u}\right) \wedge \frac{\partial}{\partial v} - \frac{2v}{l}\left(\left(\frac{l}{2}\frac{\partial}{\partial u}\right) \wedge \left(\frac{ul}{2}\frac{\partial}{\partial u} + \frac{vl}{2}\frac{\partial}{\partial v}\right)\right) \\
&\quad + \frac{2u}{l}\left(\left(\frac{l}{2}\frac{\partial}{\partial v}\right) \wedge \left(\frac{ul}{2}\frac{\partial}{\partial u} + \frac{vl}{2}\frac{\partial}{\partial v}\right)\right) \\
&= \frac{1}{4}(l^2 - 2l - 2v^2 l - 2u^2 l)\frac{\partial}{\partial u} \wedge \frac{\partial}{\partial v} \\
&= -\frac{1}{4}(1 + u^2 + v^2)\frac{\partial}{\partial u} \wedge \frac{\partial}{\partial v}.
\end{aligned}$$

Note: This example is not artificial. The bivector π is a *Poisson structure* on $\mathbb{R}^3$, which induces a foliation of $\mathbb{R}^3$ into *symplectic leaves*. The submanifold S^2 is one of these leaves. The restriction of π to S^2 is in fact (the inverse of) a symplectic form.

2.3 Pulling Back Differential Forms

If $\varphi : M \longrightarrow N$ is smooth, and ω is a k-form on N, then we can define a form $\varphi^*\omega$ on M by

$$\varphi^*\omega(v_1, \ldots, v_k) := \omega(\varphi_*(v_1), \ldots, \varphi_*(v_k)).$$

The form $\varphi^*\omega$ is called the *pullback* of ω. If x and y are coordinate systems on M and N, respectively, then the following properties can be used to compute $\varphi^*\omega$.

1. $\varphi^*(\omega_1 + \omega_2) = \varphi^*\omega_1 + \varphi^*\omega_2$
2. $\varphi^*(\omega_1 \wedge \omega_2) = \varphi^*\omega_1 \wedge \varphi^*\omega_2$
3. $\varphi^*(dy^i) = \displaystyle\sum_{j=1}^{n} \frac{\partial(y^i \circ \varphi)}{\partial x^j} dx^j.$
4. If g is a 0-form, then $\varphi^* g = g \circ \varphi$.

Note: I'm using the convention that for a 0-form (a smooth function on M) g,

$$g \wedge \omega = g \cdot \omega.$$

Example: Let $M = N = \{(x, y) \in \mathbb{R} \,|\, x, y > 0\}$ and let φ be the identity map. Let (x, y) and (u, v) be coordinate charts on M where $u = x^2$ and $v = 2x + y$. That is,

$$\varphi(x, y) = (x^2, 2x + y) = (u, v),$$

and

$$\varphi^{-1}(u, v) = (\sqrt{u}, v - 2\sqrt{u}) = (x, y).$$

Let $\omega = xy^2 dx \wedge dy$. Then

$$\begin{aligned}
\varphi^*\omega &= \varphi^*(xy^2)\varphi^*(dx) \wedge \varphi^*(dy) \\
&= (x \circ \varphi)(y \circ \varphi)^2 d(x \circ \varphi) \wedge d(y \circ \varphi) \\
&= (\sqrt{u})(v - 2\sqrt{u})^2 d(\sqrt{u}) \wedge d(v - 2\sqrt{u}) \\
&= (\sqrt{u})(v - 2\sqrt{u})^2 \left(-\frac{1}{2\sqrt{u}} du\right) \wedge (dv - \left(\frac{1}{\sqrt{u}}\right) du) \\
&= (\sqrt{u})(v - 2\sqrt{u})^2 \left(-\frac{1}{2\sqrt{u}}\right)\left(-\frac{1}{\sqrt{u}}\right) du \wedge dv \\
&= \left(\frac{v\sqrt{u} - 2u)^2}{2u}\right) du \wedge dv.
\end{aligned}$$

2.4 Tensor Fields

The tangent bundle and the cotangent bundle of a manifold M are examples of *vector bundles*. Roughly speaking, a vector bundle is a manifold together with a collection of vector spaces of the same dimension, each associated to one point on the manifold. The manifold is the *base space*, and each vector space can be thought of as sitting over the corresponding point on the manifold. There is also a local triviality condition which ensures that things are locally reasonably nice. Other examples of vector bundles are $\wedge^k(TM)$ and $\wedge^k(T^*M)$ for $k > 1$.

Suppose we have a vector bundle E with base space M, and for each $p \in M$, let V_p be the corresponding vector space. A *tensor field* is a choice, at each point p, of an element of $\left(\otimes^k V_p\right) \otimes \left(\otimes^l V_p^*\right)$ for some k, l. Equivalently, a tensor field is a section of the vector bundle $\left(\otimes^k E\right) \otimes \left(\otimes^l E^*\right)$. Thus, k-forms and k-vector fields are tensor fields. These are alternating tensor fields because the vectors associated to each point are alternating–i.e., they live in the kth wedge power of a vector space.

One can also consider symmetric tensor fields. In this case, the vector space associated to each point in the base space is a symmetric tensor power of, for example, the tangent space. A Riemannian metric is an example of a symmetric tensor field.

2.5 Covariant and Contravariant Tensor Fields

As we have seen, multivector fields map forward and differential forms pull back. Tensor fields that map forward are called *contravariant*, and tensor fields that pull back are called covariant. This is, of course, completely ridiculous: the terms should be reversed. This unfortunate convention has been in place so long, however, that it will be almost impossible to change.

For computational purposes, it is usually easier to pull back forms than to push forward multivector fields. Suppose that we have a bivector field on $\mathbb{R}^n = \{(x_1, \ldots, x_n\}$, $\pi = f\frac{\partial}{\partial x_i} \wedge \frac{\partial}{\partial x_j}$, which we would like to push forward under a diffeomorphism $\phi : \mathbb{R}^n \to \mathbb{R}^n$. On the image of ϕ, we will use coordinates $(y_1, \ldots, y_n)$. The bivector field π seems to contain the same information as the 2-form $\omega = f dx_i \wedge dj$, and pulling back ω under ϕ^{-1} transfers this information from the first copy of $\mathbb{R}^3$ to the second. Is pulling back ω by ϕ^{-1} the same as pushing forward π by ϕ? Well, not quite. In coordinates,

$$\left(\phi^{-1}\right)^* \omega = (f \circ \phi^{-1}) \sum_{1 \le k < l \le n} \det\left(\frac{\partial(x_i, x_j)}{\partial(y_k, y_l)}\right) dy_k \wedge dy_l,$$

and

$$\phi_* \omega = (f \circ \phi^{-1}) \sum_{1 \le k < l \le n} \det\left(\frac{\partial(y_k, y_l)}{\partial(x_i, x_j)}\right) \frac{\partial}{\partial y_k} \wedge \frac{\partial}{\partial y_l},$$

where

$$\left(\frac{\partial(y_k, y_l)}{\partial(x_i, x_j)}\right) = \begin{pmatrix} \frac{\partial y_k}{\partial x_i} & \frac{\partial y_k}{\partial x_j} \\ \frac{\partial y_l}{\partial x_i} & \frac{\partial y_l}{\partial x_j} \end{pmatrix}.$$

(These computations generalize in the obvious way to higher degree forms/multivector fields.) Using the fact that

$$\det\left(\frac{\partial(y_k, y_l)}{\partial(x_i, x_j)}\right) = \left(\det\left(\frac{\partial(x_i, x_j)}{\partial(y_k, y_l)}\right)\right)^{-1},$$

it is therefore possible to push forward a multivector field by (1) changing it into a form by replacing the dx_i's by $\frac{\partial}{\partial x_i}$'s, (2)pulling back this form under the inverse map, (3) taking reciprocals of the appropriate parts of the coefficients of the pullback, and (4) changing the resulting form back into a multivector field.

There is much more to say about covariant and contravariant tensor fields. A more extensive description can be found in [2], for example.

References

[1] L. Conlon. *Differentiable Manifolds*. Birkhauser, 2001.

[2] J.M. Lee. *Riemannian manifolds*. Springer-Verlag, New York, 1997.

[3] J.M. Lee. *Introduction to Smooth Manifolds*. Springer, 2003.

[4] M. Spivak. *A Comprehensive Introduction to Differential Geometry*. 1999.