Differential forms and vector calculus

Sean Fitzpatrick

November 11, 2008

This is a brief introduction to the language of differential forms in $\mathbb{R}^3$ and how they relate to vector calculus. You may ask, why bother to re-write everything when we already have perfectly good formulas in terms of vectors? For one, I hope you'll agree that these formulas look a lot nicer, and are more easily remembered, in terms of forms. Another reason is the need to do vector calculus in spaces other than $\mathbb{R}^3$. For example, electrodynamics takes place in 4D (3D space + time), and general relativity requires curved space. Our vector calculus formulas don't work in these settings (try defining the cross product of two four-dimensional vectors!), but differential forms work just fine. So let's get started with the translation!

1 Differential forms on $\mathbb{R}^3$

Let $f(x, y, z)$ be a smooth function, and let $\vec{F}(x, y, z) = P(x, y, z)\hat{\imath} + Q(x, y, z)\hat{\jmath} + R(x, y, z)\hat{k}$ be a smooth vector field. (For the purposes of this note, "smooth" will mean that any partial derivatives we need exist and are continuous.) We can express the four types of differential form in $\mathbb{R}^3$ in terms of f and $\vec{F}$ as follows.

- **0-forms** A 0-form is just a smooth function, such as $f(x, y, z)$.
- **1-forms** To each vector field $\vec{F}$ we associate a 1-form

  $$W_{\vec{F}} = P(x, y, z)\,dx + Q(x, y, z)\,dy + R(x, y, z)\,dz,$$

  called the *work form* associated to $\vec{F}$. Every 1-form is an expression of this type.
- **2-forms** We can also associate $\vec{F}$ to a 2-form $\Phi_{\vec{F}}$, called the *flux form* associated to $\vec{F}$, and given by

  $$\Phi_{\vec{F}} = P(x, y, z)\,dy \wedge dz + Q(x, y, z)\,dz \wedge dx + R(x, y, z)\,dx \wedge dy.$$
- **3-forms** Finally, our function $f(x, y, z)$ can be associated to a 3-form, given by

  $$\rho_f = f(x, y, z)\,dx \wedge dy \wedge dz,$$

  which we'll call the "volume form" associated to f. (Technically, to be considered a volume form we should demand that $f(x, y, z) > 0$ everywhere, but we'll abuse the language.)

For each k, we denote the space of k-forms by $\Omega^k(\mathbb{R}^3)$. We remark that every 1-form is the work form of some vector field, every 2-form is the flux form of some vector field, and every 3-form is the volume form of some function.

OK, so now that we have this notation, what can we do with it? We first need to define two operations on differential forms: the wedge product appearing in $\Phi_{\vec{F}}$ and ρ_f, and the differential d.

The *wedge* product $\wedge$ is a multiplication that takes a k-form α and an l-form β and gives a $k+l$-form $\alpha \wedge \beta$ such that:

1. For any smooth functions f and g, and differential forms α, β and γ, we have
$$\alpha \wedge (f(x,y,z)\beta + g(x,y,z)\gamma) = f(x,y,z)\alpha \wedge \beta + g(x,y,z)\alpha \wedge \gamma.$$
2. If α is a k-form, and β is an l-form, then
$$\alpha \wedge \beta = (-1)^{kl}\beta \wedge \alpha.$$

In particular, this tells us that $dx \wedge dy = -dy \wedge dx$, and that $dx \wedge dx = dy \wedge dy = dz \wedge dz = 0$. This also tells us why there are no 4-forms, 5-forms, etc: we have, for example, $dy \wedge (dx \wedge dy \wedge dz) = -(dy \wedge dy) \wedge (dx \wedge dz) = 0$.

Recall that the differential of a continuously differentiable function $f(x,y,z)$ is given by
$$df(x,y,z) = f_x(x,y,z)dx + f_y(x,y,z)dy + f_z(x,y,z)dz.$$
This defines the differential d on 0-forms; note that the result is a 1-form. We also get the identity
$$df = W_{\nabla f}.$$
That is, the differential of f is the work form of its gradient. Notice that this gives an alternative answer to the question "When is $\vec{F}$ a gradient?" - we have $\vec{F} = \nabla f$ if and only if $W_{\vec{F}} = df$. That is, $\vec{F}$ is a conservative vector field if its work form is the differential of some function.

Next, we define the differential of a 1-form $\alpha = P\,dx + Q\,dy + R\,dz$ by
$$dW_{\vec{F}} = dP \wedge dx + dQ \wedge dy + dR \wedge dz,$$
where dP, dQ, and dR are the usual differentials. If we write out the differentials in full, we can use the properties of the wedge product to simplify this expression. For example, we have
$$\begin{aligned} dP \wedge dx &= (P_x\,dx + P_y\,dy + P_z\,dz) \wedge dx \\ &= P_z\,dz \wedge dx - P_y\,dx \wedge dy, \end{aligned}$$
since $dy \wedge dx = -dx \wedge dy$ and $dx \wedge dx = 0$. If we do the same for $dQ \wedge dy$ and $dR \wedge dz$, and put everything together, we get
$$dW_{\vec{F}} = \left(\frac{\partial R}{\partial y} - \frac{\partial Q}{\partial z}\right) dy \wedge dz + \left(\frac{\partial P}{\partial z} - \frac{\partial R}{\partial x}\right) dz \wedge dx + \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y}\right) dx \wedge dy.$$
Now, notice that the functions appearing here are the components of $\nabla \times \vec{F}$ - we get the identity
$$dW_{\vec{F}} = \Phi_{\nabla \times \vec{F}}.$$

In other words, applying the differential d to a 1-form has the same effect as taking the curl of a vector field.

Now, let's apply the differential d to a 2-form $\Phi_{\vec{F}}$. The pattern is the same: we take the differentials of the component functions, and then work out the wedge products. We have:

$$d\Phi_{\vec{F}} = dP \wedge dy \wedge dz + dQ \wedge dz \wedge dx + dR \wedge dx \wedge dy.$$

In this case if we look at one of the terms, say $dQ \wedge dz \wedge dx$, since dz and dx are already there, only the term $Q_y dy$ is going to survive, and we end up with $Q_y\, dy \wedge dz \wedge dx = Q_y\, dx \wedge dy \wedge dz$. (Here there's no sign change because we make an even number of swaps: $dy \wedge dz \wedge dx = -dy \wedge dx \wedge dz = -(-dx \wedge dy \wedge dz)$. Putting everything together, we get

$$d\Phi_{\vec{F}} = (P_x + Q_y + R_z)\, dx \wedge dy \wedge dz = \nabla \cdot \vec{F}\, dx \wedge dy \wedge dz = \rho_{\nabla \cdot \vec{F}}.$$

Summing up, we get a sequence of maps

$$\Omega^0(\mathbb{R}^3) \xrightarrow{d} \Omega^1(\mathbb{R}^3) \xrightarrow{d} \Omega^2(\mathbb{R}^3) \xrightarrow{d} \Omega^3(\mathbb{R}^3)$$

such that $df = W_{\nabla f}$, $dW_{\vec{F}} = \Phi_{\nabla \times \vec{F}}$, and $d\Phi_{\vec{F}} = \rho_{\nabla \cdot \vec{F}}$. Of course, applying the differential d to a 3-form gives zero, since there are no 4-forms on $\mathbb{R}^3$.

Finally, the differential d has the following useful property:

$$d^2 = d \circ d = 0.$$

In other words, if we apply d twice in a row, we always get zero. This gives us the two separate identities we know from vector calculus:

$$d^2 f = d(df) = d(W_{\nabla f}) = \Phi_{\nabla \times \nabla f} = 0,$$

which tells us that $\nabla \times \nabla f = 0$, and

$$d^2 W_{\vec{F}} = d(dW_{\vec{F}}) = d(\Phi_{\nabla \times \vec{F}}) = \rho_{\nabla \cdot (\nabla \times \vec{F})} = 0,$$

or $\nabla \cdot (\nabla \times \vec{F}) = 0$.

2 Integration of differential forms

We now come to the use of the language of forms in the integration theorems of vector calculus. There are three advantages to using forms. The first is that differential forms behave nicely under change of variables (we get the change of variables formula, Jacobians, etc. almost for free!). The second is that all the different theorems (FTC for line integrals, Green's theorem, Stokes' theorem, divergence theorem) we covered are special cases of one general result (also called Stokes' theorem). The third is that differential forms keep track of orientation: we get a minus sign if we switch from right-handed coordinates to left-handed, courtesy of the wedge product.

The basic principle for integrating a differential form is this: a k-form is something you integrate over a k-dimensional space. That is, a 1-form is integrated along a curve (defined by one parameter), a 2-form is integrated over a surface (defined by two parameters), and a 3-form is integrated over a region in space (given by three parameters). (A 0-form is "integrated" along a zero-dimensional space; that is, evaluated at a point.)

First, let's sum up the integrals we get when we integrate differential forms:

$$\begin{aligned}\int_C W_{\vec{F}} &= \int_C P\,dx + Q\,dy + R\,dz := \int_C (\vec{F}\cdot\vec{T})\,ds\\ \iint\limits_S \Phi_{\vec{F}} &= \iint\limits_S P\,dy\wedge dz + Q\,dz\wedge dx + R\,dx\wedge dy := \iint\limits_S (\vec{F}\cdot\vec{n})\,dS\\ \iiint\limits_T \rho_f &= \iiint\limits_T f(x,y,z)\,dx\wedge dy\wedge dz := \iiint\limits_T f(x,y,z)dV.\end{aligned}$$

So far, we've written our usual vector calculus integrals as integrals of differential forms. It seems like we're missing integrals of functions with respect to arc length and surface area ($\int_C f(x,y,z)ds$ and $\iint\limits_S f(x,y,z)dS$, respectively). However, since $\vec{T}$ and $\vec{n}$ are unit vectors, we can write

$$\int_C f(x,y,z)ds = \int_C (f(x,y,z)\vec{T})\cdot\vec{T}ds = \int_C W_{f\vec{T}} = \int_C fW_{\vec{T}},$$

and

$$\iint\limits_S f(x,y,z)\,dS = \iint\limits_S (f(x,y,z)\vec{n})\cdot\vec{n}\,dS = \iint\limits_S \Phi_{f\vec{n}} = \iint\limits_S f\Phi_{\vec{n}}.$$

Now, we need to be able to evaluate these integrals, which requires a parameterization of our curve, surface or space region. Once we define a parameterization, the rest will just be simple substitution.

Let's work everything out for a 1-form. Suppose our curve C is given by the parameterization $\vec{r}(t) = \langle x(t), y(t), z(t)\rangle$, for $t\in[a,b]$. Then we have

$$\begin{aligned}W_{\vec{F}} &= P(x,y,z)\,dx + Q(x,y,z)\,dy + R(x,y,z)\,dz\\ &= P(\vec{r}(t))x'(t)\,dt + Q(\vec{r}(t))y'(t)\,dt + R(\vec{r}(t))z'(t)\,dt\\ &= \vec{F}(\vec{r}(t))\cdot\vec{r}'(t)\,dt.\end{aligned}$$

The case of a 2-form is similar, but it takes a bit more work. Suppose we're given a parameterization $\vec{r}(u,v) = \langle x(u,v), y(u,v), z(u,v)\rangle$. Then we have, for example,

$$\begin{aligned}dx\wedge dy &= (x_u\,du + x_v\,dv)\wedge(y_u\,du + y_v\,dv)\\ &= (x_uy_v - x_vy_u)\,du\wedge dv\\ &= \frac{\partial(x,y)}{\partial(u,v)}\,du\wedge dv.\end{aligned}$$

Similar calculations give $dz\wedge dx = \frac{\partial(z,x)}{\partial(u,v)}\,du\wedge dv$, and $dy\wedge dz = \frac{\partial(y,z)}{\partial(u,v)}\,du\wedge dv$. Plugging everything in, we recover the formula

$$\iint\limits_S \Phi_{\vec{F}} = \int_c^d\int_a^b \left(P(\vec{r}(u,v))\frac{\partial(y,z)}{\partial(u,v)} + Q(\vec{r}(u,v))\frac{\partial(z,x)}{\partial(u,v)} + R(\vec{r}(u,v))\frac{\partial(x,y)}{\partial(u,v)}\right)\,du\,dv.$$

Finally, if $(x, y, z) = T(u, v, w)$ is a smooth transformation, a lengthy (but straightforward) calculation lets us show that

$$f(x, y, z)\, dx \wedge dy \wedge dz = f(T(u, v, w))J_T(u, v, w)\, du \wedge dv \wedge dw,$$

which lets us recover the change of variables formula.

Finally, we state the general Stokes' theorem. If α is a k-form, and σ is a region defined by $(k + 1)$-parameters, then

$$\int_\sigma d\alpha = \int_{\partial\sigma} \alpha,$$

where $\partial\sigma$ denotes the boundary of σ. (We've written a single integral sign to stand for single, double, or triple integrals, depending on the context.)

Let's translate this theorem back into vector calculus, case-by-case. Let's start with a 0-form $\alpha = f(x, y, z)$. Then σ is a curve C, and $\partial\sigma$ is given by the two end-points of the curve; let's denote these by A and B. We write this as $\partial C = B - A$ to indicate that the left endpoint is assigned the opposite orientation (this is a part of the discussion I've skipped). Then we get:

$$\int_C df = \int_{B-A} f = f(B) - f(A),$$

which is the Fundamental Theorem for line integrals.

Now, if $\alpha = W_{\vec{F}}$ is a 1-form, and S is a surface, we get

$$\iint\limits_S dW_{\vec{F}} = \int\limits_{\partial S} W_{\vec{F}}.$$

The boundary of S is a closed curve C, and $dW_{\vec{F}} = \Phi_{\nabla\times\vec{F}}$, so in vector form, this equation becomes

$$\iint\limits_S (\nabla \times \vec{F}) \cdot \vec{n}\, dS = \oint_C (\vec{F} \cdot \vec{T})\, ds,$$

which is the classical Stokes' theorem.

Finally, if $\alpha = \Phi_{\vec{F}}$ is a 2-form, and T is a region in space, we get

$$\iiint\limits_T d\Phi_{\vec{F}} = \iint\limits_{\partial T} \Phi_{\vec{F}}.$$

The boundary of T is a closed surface S, and $d\Phi_{\vec{F}} = \rho_{\nabla\cdot\vec{F}}$, so in vector form, this equation reads

$$\iiint\limits_T \nabla \cdot \vec{F}\, dV = \oiint\limits_S \vec{F} \cdot \vec{n}\, dS,$$

which is the divergence theorem.

Let us end with a mention of the geometric interpretation of differential forms. The 1-forms dx, dy and dz can be thought of as the linear functions on $\mathbb{R}^3$ that give the x, y, z components, respectively, of a vector $\vec{v} \in \mathbb{R}^3$: if $\vec{v} = \langle a, b, c\rangle$, then $dx(\vec{v}) = a$, $dy(\vec{v}) = b$, and $dz(\vec{v}) = c$.

Now what about a 2-forms? Well this should be something that eats 2 vectors and returns a number. If α and β are 1-forms, we define

$$\alpha \wedge \beta(\vec{u}, \vec{v}) = \alpha(\vec{u})\beta(\vec{v}) - \alpha(\vec{v})\beta(\vec{u}).$$

In particular, let's look at the basic 2-forms such as $dx \wedge dy$. If $\vec{u} = \langle u_1, u_2, u_3\rangle$ and $\vec{v} = \langle v_1, v_2, v_3\rangle$, we find

$$\begin{aligned} dx \wedge dy(\vec{u}, \vec{v}) &= dx(\vec{u})dy(\vec{v}) - dx(\vec{v})dy(\vec{u}) \\ &= u_1 v_2 - v_1 u_2 = \begin{vmatrix} u_1 & v_1 \\ u_2 & v_2 \end{vmatrix}, \end{aligned}$$

which is what we get if we project $\vec{u}$ and $\vec{v}$ onto the xy-plane, and then compute the area of the parallelogram spanned by these projections. In other words, $dx \wedge dy$ determines the "xy-component" of the area spanned by two vectors. Similarly, $dz \wedge dx$ and $dy \wedge dz$ give the areas determined by projecting onto the xz or yz planes, respectively.