The wedge product

The *wedge product* of differential forms is kind of like the cross product of vectors; however, instead of trying to interpret it as another vector (or a scalar), we view it as another differential form of higher 'rank' than the original forms. (Just as the operation that produces the cross product may be called outer multiplication of vectors, so the operation that produces the wedge product may be called *exterior multiplication* of differential forms, but the term 'wedge product' is much more common.) The ordinary differential forms that we're used to have rank 1, and they can be evaluted at a point and a vector; to evaluate a differential form of rank 2, you need a point and 2 vectors. If you keep going with more wedge products, then you get differential forms of even higher rank; to evaluate a differential form of rank p, you need a point and p vectors.

The wedge product also involves subtracting one thing from another (again like the cross product); if α and β are 1-forms (differential forms of rank 1, as we've been using so far), P_0 is a point, and $\mathbf{v}_1$ and $\mathbf{v}_2$ are vectors, then

$$(\alpha \wedge \beta)|_{\substack{P=P_0,\\ \mathrm{d}P=\mathbf{v}_1,\mathbf{v}_2}} = \alpha|_{\substack{P=P_0,\\ \mathrm{d}P=\mathbf{v}_1}} \beta|_{\substack{P=P_0,\\ \mathrm{d}P=\mathbf{v}_2}} - \alpha|_{\substack{P=P_0,\\ \mathrm{d}P=\mathbf{v}_2}} \beta|_{\substack{P=P_0,\\ \mathrm{d}P=\mathbf{v}_1}}.$$

For example, if $\alpha = x^2\,\mathrm{d}x + xy\,\mathrm{d}y$, $\beta = y^2\,\mathrm{d}x - xy\,\mathrm{d}y$, $P_0 = (2,3)$, $\mathbf{v}_1 = \langle 0.01, 0.04\rangle$, and $\mathbf{v}_2 = \langle -0.01, 0\rangle$, then

$$\begin{aligned}
&\Big((x^2\,\mathrm{d}x + xy\,\mathrm{d}y) \wedge (y^2\,\mathrm{d}x - xy\,\mathrm{d}x)\Big)\Big|_{\substack{(x,y)=(2,3),\\ \mathrm{d}(x,y)=\langle 0.01,0.04\rangle,\langle -0.01,0\rangle}} \\
&\quad = (x^2\,\mathrm{d}x + xy\,\mathrm{d}y)|_{\substack{(x,y)=(2,3),\\ \langle \mathrm{d}x,\mathrm{d}y\rangle=\langle 0.01,0.04\rangle}} (y^2\,\mathrm{d}x - xy\,\mathrm{d}y)|_{\substack{(x,y)=(2,3),\\ \langle \mathrm{d}x,\mathrm{d}y\rangle=\langle -0.01,0\rangle}} \\
&\qquad - (x^2\,\mathrm{d}x + xy\,\mathrm{d}y)|_{\substack{(x,y)=(2,3),\\ \langle \mathrm{d}x,\mathrm{d}y\rangle=\langle -0.01,0\rangle}} (y^2\,\mathrm{d}x - xy\,\mathrm{d}y)|_{\substack{(x,y)=(2,3),\\ \langle \mathrm{d}x,\mathrm{d}y\rangle=\langle 0.01,0.04\rangle}} \\
&\quad = \Big((2)^2(0.01) + (2)(3)(0.04)\Big)\Big((3)^2(-0.01) - (2)(3)(0)\Big) \\
&\qquad - \Big((2)^2(-0.01) + (2)(3)(0)\Big)\Big((3)^2(0.01) - (2)(3)(0.04)\Big) \\
&\quad = (0.28)(-0.09) - (-0.04)(-0.15) = -0.0312.
\end{aligned}$$

A few basic properties of the wedge product follow immediately (where α, β, γ are 1-forms and u is a 0-form, that is an ordinary non-differential quantity):

$$\begin{gathered}
\alpha \wedge (u\beta) = (u\alpha) \wedge \beta = u(\alpha \wedge \beta); \\
(\alpha + \beta) \wedge \gamma = \alpha \wedge \gamma + \beta \wedge \gamma; \\
\alpha \wedge (\beta + \gamma) = \alpha \wedge \beta + \alpha \wedge \gamma; \\
\alpha \wedge \beta = -\beta \wedge \alpha; \\
\alpha \wedge \alpha = 0.
\end{gathered}$$

(What these equations technically mean is that if you evaluate each side at the same point and vectors, then you'll get the same result on both sides, assuming that the operations appearing in the expressions are defined.) So if you treat the wedge product as a kind of multiplication, then you can use the ordinary rules of algebra, so long as you keep track of the order of multiplication in the wedge product and throw in a minus sign whenever you reverse the order of multiplication of two 1-forms.

To see how this works, revisit the example above where $\alpha = x^2\,\mathrm{d}x + xy\,\mathrm{d}y$ and $\beta = y^2\,\mathrm{d}x - xy\,\mathrm{d}y$, and consider $\alpha \wedge \beta$. This can be simplified as follows:

$$\begin{aligned}
\alpha \wedge \beta &= (x^2\,\mathrm{d}x + xy\,\mathrm{d}y) \wedge (y^2\,\mathrm{d}x - xy\,\mathrm{d}y) && * \\
&= (x^2\,\mathrm{d}x) \wedge (y^2\,\mathrm{d}x) + (x^2\,\mathrm{d}x) \wedge (-xy\,\mathrm{d}y) + (xy\,\mathrm{d}y) \wedge (y^2\,\mathrm{d}x) + (xy\,\mathrm{d}y) \wedge (-xy\,\mathrm{d}y) \\
&= (x^2)(y^2)(\mathrm{d}x \wedge \mathrm{d}x) + (x^2)(-xy)(\mathrm{d}x \wedge \mathrm{d}y) + (xy)(y^2)(\mathrm{d}y \wedge \mathrm{d}x) + (xy)(-xy)(\mathrm{d}y \wedge \mathrm{d}y) \\
&= x^2y^2(0) - x^3y\,\mathrm{d}x \wedge \mathrm{d}y + xy^3(-\mathrm{d}x \wedge \mathrm{d}y) - x^2y^2(0) && * \\
&= (-x^3y - xy^3)\,\mathrm{d}x \wedge \mathrm{d}y = -xy(x^2 + y^2)\,\mathrm{d}x \wedge \mathrm{d}y. && *
\end{aligned}$$

I've written this out in detail so that each step uses only one of the basic algebraic properties of the wedge product; but with a little practice, you should only need to write down the lines with asterisks after them. When you multiply the expressions (think FOIL), make sure to keep track of the order in which you multiply the differentials; if you multiply a differential by itself (such as $\mathrm{d}x \wedge \mathrm{d}x$), then you get zero, and if you multiply differentials in an order different from the order that you prefer (such as $\mathrm{d}y \wedge \mathrm{d}x$ instead of $\mathrm{d}x \wedge \mathrm{d}y$ if you prefer alphabetical order), then you can rearrange the order if you throw in a minus sign whenever two differentials switch places. In this way, you can go from the first line in the calculation above to the next line with an asterisk, skipping over the lines in between. (With a little more practice, you can even skip that line and go straight from the first line to the last line.)

To check that this simplification of $\alpha \wedge \beta$ is correct, evaluate it again at $P_0 = (2, 3)$, $\mathbf{v}_1 = \langle 0.01, 0.04\rangle$, and $\mathbf{v}_2 = \langle -0.01, 0\rangle$. You should get

$$\begin{aligned}(-xy(x^2+y^2)\,\mathrm{d}x \wedge \mathrm{d}y)|_{\substack{(x,y)=(2,3),\\ \mathrm{d}(x,y)=\langle 0.01,0.04\rangle,\langle -0.01,0\rangle}} &= \Big(-xy(x^2+y^2)\Big)\Big|_{(x,y)=(2,3)} (\mathrm{d}x \wedge \mathrm{d}y)|_{\langle \mathrm{d}x,\mathrm{d}y\rangle=\langle 0.01,0.04\rangle,\langle -0.01,0\rangle} \\ &= -(2)(3)\Big((2)^2+(3)^2\Big)\Big((0.01)(0)-(0.04)(-0.01)\Big) = -0.0312,\end{aligned}$$

the same result as before. (What makes the original and simplified versions of $\alpha \wedge \beta$ equal to each other is precisely that you will get the same result for each as long as you use the same values of P_0, $\mathbf{v}_1$, and $\mathbf{v}_2$, no matter what those values are.)

To define a wedge product between forms of higher rank, you have to add and subtract all possible permutations of the possible orders in which to write the vectors at which the result is evaluated. Keeping track of all of this in a general formula is complicated, but the important point for our calculations is that the rules above continue to apply, and additionally we have an associative law for wedge products:

$$(\alpha \wedge \beta) \wedge \gamma = \alpha \wedge (\beta \wedge \gamma).$$

We will not actually need to evaluate these higher-rank forms in this course; what's necessary is to work with them algebraically. In other words, the only calculation in these notes so far that you really need to know how to do is the one towards the bottom of Page 1.

An imporant example that we've looked at before is the transformation between rectangular and polar coordinates. Given

$$\begin{aligned} x &= r\cos\theta, \\ y &= r\sin\theta, \end{aligned}$$

we differentiate to get

$$\begin{aligned} \mathrm{d}x &= \cos\theta\,\mathrm{d}r - r\sin\theta\,\mathrm{d}\theta, \\ \mathrm{d}y &= \sin\theta\,\mathrm{d}r + r\cos\theta\,\mathrm{d}\theta. \end{aligned}$$

Given this, the algebra of the wedge product determines this calculation:

$$\begin{aligned} \mathrm{d}x \wedge \mathrm{d}y &= (\cos\theta\,\mathrm{d}r - r\sin\theta\,\mathrm{d}\theta) \wedge (\sin\theta\,\mathrm{d}r + r\cos\theta\,\mathrm{d}\theta) \\ &= \cos\theta\sin\theta\,(0) + r\cos^2\theta\,(\mathrm{d}r \wedge \mathrm{d}\theta) - r\sin^2\theta\,(-\mathrm{d}r \wedge \mathrm{d}\theta) - r^2\sin\theta\cos\theta\,(0) \\ &= (r\cos^2\theta + r\sin^2\theta)\,\mathrm{d}r \wedge \mathrm{d}\theta = r\,\mathrm{d}r \wedge \mathrm{d}\theta. \end{aligned}$$

I did this before to calculate the area element

$$đA = |\mathrm{d}x \wedge \mathrm{d}y| = |r\,\mathrm{d}r \wedge \mathrm{d}\theta| = r\,|\mathrm{d}r \wedge \mathrm{d}\theta|,$$

assuming that $r \geq 0$.

Besides changing coordinates in the plane, you can also use two variables to parametrize a surface in three-dimensional space. For example, on the surface of the unit sphere (the sphere of radius 1 centred at $(x, y, z) = (0, 0, 0)$), if we write x and y using r and θ above, then we can further write

$$r = 1 \sin \phi,$$
$$z = 1 \cos \phi,$$

where the 1 indicates the radius of the sphere and the angle ϕ varies from 0 to π. (In other words, I'm using spherical coordinates with $\rho = 1$.) Differentiating,

$$\mathrm{d}r = \cos \phi \, \mathrm{d}\phi,$$
$$\mathrm{d}z = -\sin \phi \, \mathrm{d}\phi.$$

Thus,

$$\begin{aligned} \mathrm{d}x \wedge \mathrm{d}y &= r \, \mathrm{d}r \wedge \mathrm{d}\theta = \sin \phi \cos \phi \, \mathrm{d}\phi \wedge \mathrm{d}\theta, \\ \mathrm{d}x \wedge \mathrm{d}z &= (\cos \theta \cos \phi \, \mathrm{d}\phi - \sin \phi \sin \theta \, \mathrm{d}\theta) \wedge (-\sin \phi \, \mathrm{d}\phi) = 0 + \sin^2 \phi \sin \theta \, \mathrm{d}\theta \wedge \mathrm{d}\phi, \\ \mathrm{d}y \wedge \mathrm{d}z &= (\sin \theta \cos \phi \, \mathrm{d}\phi + \sin \phi \cos \theta \, \mathrm{d}\theta) \wedge (-\sin \phi \, \mathrm{d}\phi) = 0 - \sin^2 \phi \cos \theta \, \mathrm{d}\theta \wedge \mathrm{d}\phi. \end{aligned}$$

However,

$$\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z = (\sin \phi \cos \phi \, \mathrm{d}\phi \wedge \mathrm{d}\theta) \wedge (-\sin \phi \, \mathrm{d}\phi) = \sin^2 \phi \cos \phi \, \mathrm{d}\phi \wedge \mathrm{d}\phi \wedge \mathrm{d}\theta = 0,$$

because $\mathrm{d}\phi \wedge \mathrm{d}\phi = 0$. This makes sense if $\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z$ represents something like a volume, since the volume of the *surface* of a sphere is zero.

To see how $\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z$ indeed represents something like a volume, I should explain how to integrate higher-rank differential forms. (Of course, I need to explain that anyway if theses notes are going to apply to Sections 15.5&15.6 of the textbook, which are all about integrals!) You typically integrate a differential form over a shape (or 'manifold') whose dimension (as given by the number of parameters used to parametrize it) matches the rank of the form. We have already seen this with rank-1 forms integrated over 1-dimensional curves, which can be parametrized by 1 parameter t. In general, to approximate the integral of a rank-p form over a p-dimensional manifold (one parametrized by p parameters), you divide the manifold up into pieces along level curves (or level surfaces) of the parameters, for each piece evaluate the differential form at a point in that piece and at the vectors across the piece along level curves through that point (evaluating $\mathrm{d}u \wedge \mathrm{d}v$ first at the vector along which u increases and then at the vector along which v increases), multiply by ± 1 according to the orientation of the manifold (see the next paragraph), and then add these pieces up. If the limit of the result of this process, no matter which level curves (or level surfaces) you pick and no matter which points in each piece you pick, exists as long as the size of the largest piece goes to zero, then that limit is the value of the integral. As with other Riemann integrals, this is guaranteed to exist if you're integrating a continuous differential form on a manifold with a continuously differentiable parametrization that is compact (closed and bounded).

I still need to explain the **orientation** of the manifold; this indicates directions along it. In the case of a curve, there are two ways to go along the curve, giving two orientations. In the case of a surface, if we start going in some direction, then we can turn from that direction in one way or the other. In particular, the coordinate plane can be oriented clockwise or counterclockwise. Ordinary three-dimensional space has right-handed and left-handed orientations. In general, every small piece of a manifold has two orientations, no matter what the manifold's dimension. A differential form such as $\mathrm{d}u \wedge \mathrm{d}v$ matches the orientation of a manifold if moving in the direction in which u increases and then turning in the direction in which v increases matches the manifold's orientation; if not, then we must use -1 for that piece when forming a Riemann sum.

As usual, you don't need to evaluate an integral as a limit of Riemann sums as long as you have nice formulas for everything; instead, you evaluate it as an iterated integral in the parameters. To do this, you simply set up bounds of integration over the values that the parameters can take and write down an iterated integral that makes sense, inserting a factor of -1 if the orientation of the differential form is opposite that of the manifold. (In principle, whether the orientations match may depend on where you are, and then you'll need to divide the manifold into several pieces, multiplying some of them by -1 but not others. In practice, however, this is rarely necessary.) For example, to integrate $\mathrm{d}x \wedge \mathrm{d}y = \sin\phi\cos\phi\,\mathrm{d}\phi \wedge \mathrm{d}\theta$ on the top half of the unit sphere, oriented to turn clockwise when viewed from above the sphere, we start with

$$\int_{\theta=0}^{2\pi}\left(\int_{\phi=0}^{\pi/2}\sin\phi\cos\phi\,\mathrm{d}\phi\right)\mathrm{d}\theta = \int_{\theta=0}^{2\pi}\frac{1}{2}\,\mathrm{d}\theta = \pi;$$

but then, because we turn *counterclockwise* to move from a direction in which ϕ increases to a direction in which θ increases (which is opposite the specified orientation of the surface), the actual value is $-\pi$.

In the textbook, you'll never be directly given differential forms to integrate in more than 1 dimension. In some of Section 15.6 and much of Sections 15.7 and 15.8, you integrate vector fields through surfaces; to integrate the vector field $\mathbf{F}$, you integrate the differential form $\mathbf{F}(x,y,z)\cdot \text{đ}\mathbf{S}$, where $\text{đ}\mathbf{S}$ is the **oriented surface element**

$$\text{đ}\mathbf{S} = \frac{1}{2}\,\mathrm{d}P \hat{\times} \mathrm{d}P = \langle \mathrm{d}y \wedge \mathrm{d}z, \mathrm{d}z \wedge \mathrm{d}x, \mathrm{d}x \wedge \mathrm{d}y\rangle = \frac{\partial P}{\partial u} \times \frac{\partial P}{\partial v}\,\mathrm{d}u \wedge \mathrm{d}v.$$

(People often write $\text{đ}\mathbf{S}$ as simply $\mathrm{d}\mathbf{S}$, although there is no quantity $\mathbf{S}$ that it is the differential of.) Here, $P = (x,y,z)$ as usual; the book prefers $\mathbf{r} = \langle x,y,z\rangle$, but since $\mathrm{d}P = \mathrm{d}\mathbf{r}$, partial derivatives of P and of $\mathbf{r}$ are the same, so we can equally well write

$$\text{đ}\mathbf{S} = \frac{1}{2}\,\mathrm{d}\mathbf{r} \hat{\times} \mathrm{d}\mathbf{r} = \langle \mathrm{d}y \wedge \mathrm{d}z, \mathrm{d}z \wedge \mathrm{d}x, \mathrm{d}x \wedge \mathrm{d}y\rangle = \frac{\partial \mathbf{r}}{\partial u} \times \frac{\partial \mathbf{r}}{\partial v}\,\mathrm{d}u \wedge \mathrm{d}v.$$

(When I write $\hat{\times}$ between vector-valued differential forms, I mean to multiply them as vectors using the cross product and as differential forms using the wedge product. Note that you get two minus signs when switching the order of multiplication, so the result of multiplying $\mathrm{d}P = \mathrm{d}\mathbf{r}$ by itself is not zero but rather twice something, and that something is what we mean by $\text{đ}\mathbf{S}$.)

The middle formula for $\text{đ}\mathbf{S}$ (the one without P or $\mathbf{r}$ in it) requires the use of the right-hand rule for the cross product; in other words, it involves an orientation of the ambient three-dimensional space (not to be confused with an orientation of the surface itself). Properly, $\text{đ}\mathbf{S}$ is a **pseudoform**, meaning that it must be given with an orientation of the ambient space and changes sign if that orientation reverses. (Recall that multiplying vectors with the cross product similarly results in a pseudovector, or axial vector.) When using this pseudoform given with the right-handed orientation of space, we accordingly use the right-hand rule to convert between a direction through the surface (which is a **pseudoorientation**) and an orientation on the surface. In this way, it makes sense to integrate a vector field through a pseudooriented surface; if you consistently use the left-hand rule instead of the right-hand rule, then the final result will be the same.

So for example, integrating the constant vector field $\mathbf{F}(x,y,z) = \langle 0,0,1\rangle = \mathbf{k}$ through the top half of the unit sphere pseudooriented downwards is the same as integrating the rank-2 differential form

$$\mathbf{F}(x,y,z)\cdot \text{đ}\mathbf{S} = \langle 0,0,1\rangle \cdot \langle \mathrm{d}y \wedge \mathrm{d}z, \mathrm{d}z \wedge \mathrm{d}x, \mathrm{d}x \wedge \mathrm{d}y\rangle = 0 + 0 + \mathrm{d}x \wedge \mathrm{d}y = \mathrm{d}x \wedge \mathrm{d}y$$

on that hemisphere oriented clockwise when viewed from above (which I calculated above to be $-\pi$), because turning the fingers of your right hand clockwise results in your thumb pointing downwards. (If you used the left-hand rule instead, then you'd turn the fingers of your left hand counterclockwise to make your left thumb point downwards, but you'd also use $\langle \mathrm{d}z \wedge \mathrm{d}y, \mathrm{d}x \wedge \mathrm{d}z, \mathrm{d}y \wedge \mathrm{d}x\rangle$ for $\text{đ}\mathbf{S}$, and the final result would be the same.) Since the vector field that we integrated points upwards while the surface through which we integrated is pseudooriented downwards, you should expect the final result to be negative; guessing the sign of the integral ahead of time like this can help you to avoid mistakes with orientation.

In Section 15.5 and some of Section 15.6, you integrate scalar fields on surfaces; to integrate the scalar field f, you integrate the differential form $f(x,y,z)\,đ\sigma$, where

$$đ\sigma = |đ\mathbf{S}| = \sqrt{(\mathrm{d}y \wedge \mathrm{d}z)^2 + (\mathrm{d}z \wedge \mathrm{d}x)^2 + (\mathrm{d}x \wedge \mathrm{d}y)^2} = \left|\frac{\partial \mathbf{r}}{\partial u} \times \frac{\partial \mathbf{r}}{\partial v}\right| |\mathrm{d}u \wedge \mathrm{d}v|.$$

Now orientation is irrelevant, instead, simply make sure that all parameters are increasing in the iterated integral. It's possible to write $\mathbf{n}\,đ\sigma$ for $đ\mathbf{S}$, where $\mathbf{n}$ is $\widehat{đ\mathbf{S}}$, a unit vector in the direction of $đ\mathbf{S}$, that is a unit vector perpendicular to the surface pointing in the direction given by its pseudoorientation. This is how the book writes it, but actually calculating $\mathbf{n}$ and $đ\sigma$ is a waste of time if $đ\mathbf{S}$ is all that you really want.

So for example, integrating the constant scalar field $f(x,y,z) = 1$ on the top half of the unit sphere is the same as integrating the rank-2 differential form

$$f(x,y,z)\,đ\sigma = 1\sqrt{(\mathrm{d}y \wedge \mathrm{d}z)^2 + (\mathrm{d}z \wedge \mathrm{d}x)^2 + (\mathrm{d}x \wedge \mathrm{d}y)^2}$$

on that hemisphere with either orientation. In terms of the parameters ϕ and θ, this differential form is

$$\begin{aligned}&\sqrt{\sin^2\phi\cos^2\phi\,(\mathrm{d}\phi \wedge \mathrm{d}\theta)^2 + \sin^4\phi\sin^2\theta\,(\mathrm{d}\theta \wedge \mathrm{d}\phi)^2 + \sin^4\phi\cos^2\theta\,(\mathrm{d}\theta \wedge \mathrm{d}\phi)^2} \\ &\quad = \sqrt{\sin^2\phi\cos^2\phi\,(\mathrm{d}\phi \wedge \mathrm{d}\theta)^2 + \sin^4\phi\,(\mathrm{d}\phi \wedge \mathrm{d}\theta)^2} = \sqrt{\sin^2\phi\,(\mathrm{d}\phi \wedge \mathrm{d}\theta)^2} = \sin\phi\,|\mathrm{d}\phi \wedge \mathrm{d}\theta|.\end{aligned}$$

Here, I switched from $\mathrm{d}\theta \wedge \mathrm{d}\phi$ to $\mathrm{d}\phi \wedge \mathrm{d}\theta$ at one point, which should introduce a minus sign, but since this was inside a square, the minus sign was irrelevant. Also, I simplified $\sqrt{\sin^2\phi}$ to $\sin\phi$ rather than to $|\sin\phi|$, since $0 \leq \phi \leq \pi$ (so that $\sin\phi \geq 0$). The value of the integral is now

$$\int_{\theta=0}^{2\pi}\left(\int_{\phi=0}^{\pi/2} \sin\phi\,\mathrm{d}\phi\right)\mathrm{d}\theta = \int_{\theta=0}^{2\pi} 1\,\mathrm{d}\theta = 2\pi.$$

(This is the area of the surface in question, which should always be what you get when you integrate $đ\sigma$ itself, or equivalently when you integrate the constant scalar field with the value 1.)

As for Chapter 14, here we are integrating scalar fields on the flat surface of the plane, using

$$đA = |\mathrm{d}x \wedge \mathrm{d}y|$$

and on all of three-dimensional space, using

$$đV = |\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z|.$$

You can also think of $đA$ as the pseudoform $\mathrm{d}x \wedge \mathrm{d}y$, giving the plane its counterclockwise orientation, and similarly think of $đV$ as the pseudoform $\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z$, giving space its right-handed orientation. This will be useful for some purposes later on, but for purposes of calculation, it's easier to use the absolute values so that you don't have to think about orientation.