Integration on surfaces

Just as you can integrate a differential 1-form (the ordinary kind without the wedge product) on an oriented curve, so you can integrate a **differential 2-form** (two 1-forms multiplied together by the wedge product or an expression built out of such products) on an oriented surface. (Similarly, you can integrate a differential 3-form on an oriented region of space, and so on for higher rank forms in spaces of higher dimension, but we're not doing any of that except for the volume integrals that we've already covered.)

Similarly, just as you can integrate a vector field along an oriented curve by taking a dot product with $\mathrm{d}\mathbf{r}$ to get a differential 1-form and you can also integrate a vector field across a pseudooriented curve by taking a cross product with $\mathrm{d}\mathbf{r}$ to get a differential pseudo-1-form (and then reintepreting this as an honest differential 1-form on an oriented curve), so you can integrate a vector field across a pseudooriented surface by taking a dot product with $\mathrm{d}\mathbf{S}$ to get a differential pseudo-2-form (and then reintepreting this as an honest differential 2-form on an oriented surface).

So now I need to explain what all of this means.

Parametrizing surfaces

Just as you use 1 parameter (often called t) to parametrize a curve, so you use 2 variables (often called u and v) to parametrize a surface. For example, on the surface of the unit sphere (the sphere of radius 1 centred at $(x, y, z) = (0, 0, 0)$), we can use spherical coordinates with $\rho = 1$, so that

$$\begin{aligned} x &= r\cos\theta = \rho\sin\phi\cos\theta = \sin\phi\cos\theta, \\ y &= r\sin\theta = \rho\sin\phi\sin\theta = \sin\phi\sin\theta, \text{ and} \\ z &= \rho\cos\theta = \cos\theta. \end{aligned}$$

That is, ϕ and θ are the parameters. (You can call them u and v instead, but it's convenient to call them by more familiar names when possible.) Strictly speaking, the parametrization should also indicate the range of values taken by the parameters; in this case,

$$0 \leq \phi\pi,\ 0 \leq \theta \leq 2\pi.$$

Now I have made this sphere into a **parametrized surface** (in 3-dimensional space).

In general, you can use ϕ and θ as parameters whenever the surface can be described by giving ρ as a function of ϕ and θ. (In the example above, that function was the constant function with value 1.) Besides using spherical coordinates, cylindrical coordinates are also often useful for parametrization. Most often, you'll use r and θ as the parameters, but sometimes you'll use z and θ; in any case, you'll need a way to express the other variable as a function of the two that you're using as parameters. Then using $x = r\cos\theta$ and $y = r\sin\theta$, you have x, y, and z all given as functions of the parameters. Finally, if you can express z as a function of x and y, then you can use x and y themselves as the parameters. (You could also use x and z or y and z, as long as the missing variable is given as a function of the two that you use.)

While most examples will use familiar coordinates as the parameters, in general, so long as you have $P = (x, y, z)$ given as a point-valued function of two variables u and v, then the range of this function is a **parametrized surface**. For purposes of integrals, this function should ideally be one-to-one, but as long as the overlap is contained within a few lines in the (u, v)-plane, then it won't affect the value of any integrals. (This is the same condition as for change of variables in a double integral.) In the case of cylindrical coordinates, the overlap is when θ is 0 or 2π, or (if r is being used as a parameter) when $r = 0$; but these are contained within lines. In the case of spherical coordinates, the overlap is when θ is 0 or 2π again, when $\phi = 0$, or when $\phi = \pi$; again, these are contained within lines. So cylindrical and spherical coordinates are always acceptable for integrals. (With rectangular coordinates, there is no overlap, so they are definitely acceptable.)

Orienting surfaces

In the case of a curve, there are two ways to go along the curve, giving two orientations. In the case of a surface, there are many ways to go along it, but if you start going in some direction, then you can *turn* from that direction in one way or the other; these give the two **orientations** of the surface. (Actually, not every surface can be oriented; a Möbius strip is a famous example of a surface that cannot be orinted continuously everywhere. However, any parametrized surface can be broken into pieces on which it can be oriented, so it is possible to do some integrals on unorientable surfaces, as long as they are integrals whose values don't depend on the orientation. Surface area and other integrals of scalar fields, discussed below, are examples of these.)

A differential form such as $\mathrm{d}u \wedge \mathrm{d}v$ *matches* the orientation of a surface if moving in the direction in which u increases and then turning in the direction in which v increases matches the surface's orientation. For example, the (x, y)-plane can be oriented clockwise or counterclockwise; $\mathrm{d}x \wedge \mathrm{d}y$ matches the counterclockwise orientation (if (x, y) is a counterclockwise coordinate system as usual), while $\mathrm{d}y \wedge \mathrm{d}x$ matches the clockwise orientation.

It's often easier to think of a **pseudoorientation** of a surface, which (in a 3-dimensional space) is a direction *across* the surface. The textbook never refers directly to orientations of surfaces, but only to pseudoorientations, which it (confusingly) calls 'orientations'. However, you can switch between orientations and pseudoorientations using the right-hand rule: if you curl the fingers of your right hand in the direction of turning indicated by an orientation, then your thumb will point in the direction of crossing indicated by the corresponding pseudoorientation. So the textbook applies this right-hand rule whenever it needs an orientation but really has a pseudoorientation.

Defining surface integrals

As with other definitions of integrals, people never use this directly if they can help it, and you'll never need to use it to solve any of the problems. But for the record, here it is.

So, suppose that you have a differential 2-form α written using the variables $P = (x, y, z)$ and their differentials, and an oriented surface in (x, y, z)-space, given by some parametrization function S (so that $P = (x, y, z) = S(u, v)$ on the surface) whose domain is a compact region R. Then we can try to integrate α along the surface, by defining the integral

$$\int_{P=S(u,v)} \alpha.$$

To form a Riemann sum to approximate this integral, dividing the region R into n triangles, pick one vertex of each triangle, and let $\mathbf{v}_k$ and $\mathbf{w}_k$ (where $k = 1, 2, \ldots n$ counts the triangles) the vectors (in the ambient (x, y, z)-space) from that vertex to the other two vertices; select which is $\mathbf{v}_k$ and which is $\mathbf{w}_k$ so that, when you turn from $\mathbf{v}_k$ to $\mathbf{w}_k$, this matches the orientation of the surface. Finally, tag this partition with a point c_k within each triangle. The **Riemann sum** is

$$\sum_{k=1}^{n} \alpha\big|_{\substack{P=S(c_k),\\ \mathrm{d}P=\mathbf{u}_k,\mathbf{v}_k}}.$$

If you require that the areas of the triangles to all be less than δ and take the limit of the Riemann sums as $\delta \to 0^+$, then the value of the integral is defined to be this limit, if it exists.

There is a theorem that this limit does exist, at least if α is piecewise continuous and S is piecewise continuously differentiable (and sometimes otherwise); I don't know a nice proof of this directly, but you can prove that it exists because the practical calculation method in the next section works. Similarly, there is now a theorem that the value of this integral does not depend on the parametrization of the surface, only the orientation.

Calculating integrals

The practical method of evaluating an integral along a surface is to pick any convenient parametrization (preferably one that is continuously differentiable) and put everything in terms of those parameters.

For example, I'll integrate $z\,\mathrm{d}x \wedge \mathrm{d}y$ on the top half of the unit sphere, oriented to turn clockwise when viewed from above the sphere. I'll use the parametrization given earlier using spherical coordinates:

$$\begin{aligned} x &= \sin\phi\cos\theta, \\ y &= \sin\phi\sin\theta, \\ z &= \cos\phi. \end{aligned}$$

Since I only want the top half of the sphere, I use

$$0 \leq \phi \leq \frac{\pi}{2},\ 0 \leq \theta \leq 2\pi.$$

Now I differentiate the parametrization:

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

Then

$$\mathrm{d}x \wedge \mathrm{d}y = \cos\phi\sin\phi\cos^2\theta\,\mathrm{d}\phi \wedge \mathrm{d}\theta - \sin\phi\cos\phi\sin^2\theta\,\mathrm{d}\theta \wedge \mathrm{d}\phi = \sin\phi\cos\phi\,\mathrm{d}\phi \wedge \mathrm{d}\theta.$$

(Remember that $\mathrm{d}\phi \wedge \mathrm{d}\phi$ and $\mathrm{d}\theta \wedge \mathrm{d}\theta$ are 0, so that half of the terms immediately vanish, and that $\mathrm{d}\theta \wedge \mathrm{d}\phi = -\mathrm{d}\phi \wedge \mathrm{d}\theta$, so that the other two terms can be combined into one.) Finally,

$$z\,\mathrm{d}x \wedge \mathrm{d}y = \sin\phi\cos^2\phi\,\mathrm{d}\phi \wedge \mathrm{d}\theta.$$

So, I am basically looking at

$$\int_{\substack{0\leq\phi\leq\pi/2,\\ 0\leq\theta\leq 2\pi}} \sin\phi\cos^2\phi\,\mathrm{d}\phi\,\mathrm{d}\theta,$$

but I still need to think about the orientation. I really have $\mathrm{d}\phi \wedge \mathrm{d}\theta$ rather than $\mathrm{d}\phi\,\mathrm{d}\theta$, and this matches an orientation in which I turn from a direction in which ϕ increases to a direction in which θ increases. But this appears counterclockwise from above, while the orientation of the surface is clockwise from above. To fix this, I could rewrite the form to use $\mathrm{d}\theta \wedge \mathrm{d}\phi$, or equivalently put in a minus sign wherever $\mathrm{d}\phi \wedge \mathrm{d}\theta$ appears. So my real integral is

$$\int_{\theta=0}^{2\pi}\int_{\phi=0}^{\pi/2}(-\sin\phi\cos^2\phi)\,\mathrm{d}\phi\,\mathrm{d}\theta = \int_{\theta=0}^{2\pi}\left(-\frac{1}{3}\right)\mathrm{d}\theta = -\frac{2}{3}\pi.$$

You should be able to visualize this example geometrically well enough to see that the answer would have to be negative. Since $\mathrm{d}x \wedge \mathrm{d}y$ matches an orientation in which you turn from a direction in which x increases to a direction in which y increases, which appears counterclockwise from above, while the orientation is supposed to be clockwise from above, the factor $\mathrm{d}x \wedge \mathrm{d}y$ will always contribute something negative. The factor z, on the other hand, will always contribute something positive, since z is always positive on the top half of the sphere. So, the product $z\,\mathrm{d}x \wedge \mathrm{d}y$ will always be negative, so the overall integral must also be negative.

In this way, you can integrate any continuous differential 2-form on any surface with a continuously differentiable parametrization, because this process will always leave you with a continuous double integral to do.

Integrating vector fields

In the textbook, you'll never be directly given differential forms to integrate (other that 1-forms to integrate along curves). In some of Section 15.6 and much of Sections 15.7 and 15.8, you integrate a vector field across a surface; to integrate the vector field $\mathbf{F}$, you integrate the differential form $\mathbf{F}(x,y,z) \cdot đ\mathbf{S}$, where $đ\mathbf{S}$ is the **oriented surface element**

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

(People often write $đ\mathbf{S}$ 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

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

(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 $đ\mathbf{S}$.)

The middle formula for $đ\mathbf{S}$ (the one without P or $\mathbf{r}$ in it) requires the use of the right-hand rule for the cross product. This is because $đ\mathbf{S}$ is really a **pseudoform**, meaning that it changes sign if you switch between right-hand and lett-hand rules. (Recall that multiplying vectors with the cross product similarly results in a pseudovector, also called an axial vector.) 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.

(The textbook never writes $đ\mathbf{S}$ or even $\mathrm{d}\mathbf{S}$; instead, it writes $\mathbf{n}\,đ\sigma$, or rather $\mathbf{n}\,\mathrm{d}\sigma$. But $đ\sigma$ is just $\|đ\mathbf{S}\|$, the magnitude of $đ\mathbf{S}$; and $\mathbf{n}$ is just $\widehat{đ\mathbf{S}}$, 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. So $\mathbf{n}\,\mathrm{d}\sigma$ is really just a complicated way of saying $đ\mathbf{S}$. To actually calculate $\mathbf{n}$ and $đ\sigma$ is a waste of time if $đ\mathbf{S}$ is all that you really want.)

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

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

on that hemisphere oriented clockwise when viewed from above, because turning the fingers of your right hand clockwise results in your thumb pointing downwards. Above, I calculated this integral to be $-2/3\,\pi$), and that is exactly how I would finish this problem.

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. (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 $đ\mathbf{S}$, and the final result would be the same.)

Integrating scalar fields

In Section 15.5 and some of Section 15.6, you integrate a scalar field (that is a functions of 3 variables) on a surface; to integrate the scalar field f, you integrate the differential form $f(x,y,z)\,đ\sigma$, where

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

Because the differentials only appear inside a vector magnitude, square, or absolute value (depending on which version you look at), orientation is irrelevant; instead, simply make sure that all parameters are increasing in the iterated integral.

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

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

on that hemisphere with either orientation. To work out that expression using the parameters ϕ and θ, I can use $\mathrm{d}x \wedge \mathrm{d}y = \sin\phi\cos\phi\,\mathrm{d}\phi \wedge \mathrm{d}\theta$ from earlier, but I also need to find $\mathrm{d}y \wedge \mathrm{d}z$ and $\mathrm{d}z \wedge \mathrm{d}x$. I already have the individual differentials from page 3, so

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

and

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

Therefore, I am integrating

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

(Here, 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

$$\int_{\theta=0}^{2\pi}\int_{\phi=0}^{\pi/2} \sin\phi\cos\phi\,\mathrm{d}\phi\,\mathrm{d}\theta = \int_{\theta=0}^{2\pi} \frac{1}{2}\,\mathrm{d}\theta = \pi.$$

(You should expect the integral to be positive, since z is always positive on the top hemisphere.)

If instead I simply want the area of this surface, then I can simply integrate $đ\sigma$ itself, which gives

$$\int_{\theta=0}^{2\pi}\int_{\phi=0}^{\pi/2} \sin\phi\,\mathrm{d}\phi\,\mathrm{d}\theta = \int_{\theta=0}^{2\pi} \mathrm{d}\theta = 2\pi.$$

(And that is indeed the area of a hemisphere of radius 1.)