Math 3204: Calculus IV[1] Mikhail Lavrov

Lecture 25: Stokes' theorem and conservative vector fields

November 12, 2024 *Kennesaw State University*

1 Stokes' theorem and differential forms

When taking the flux integral of $\mathbf{F} = M\,\mathbf{i} + N\,\mathbf{j} + P\,\mathbf{k}$ across a surface S, we have used two forms of notation:

$$\iint_S \mathbf{F} \cdot \mathbf{n}\,\mathrm{d}A = \iint_S M\,\mathrm{d}y \wedge \mathrm{d}z + N\,\mathrm{d}z \wedge \mathrm{d}x + P\,\mathrm{d}x \wedge \mathrm{d}y.$$

We can also think of this as the integral of the 2-form (differential form of degree 2) given by $M\,\mathrm{d}y \wedge \mathrm{d}z + N\,\mathrm{d}z \wedge \mathrm{d}x + P\,\mathrm{d}x \wedge \mathrm{d}y$. A differential form of degree 2 is the most general thing you can integrate over a 2-dimensional object like the surface S.

How does this interact with the curl of a vector field? Well, it turns out that the curl operator $\nabla\times$ has an equivalent in the world of differential forms. This equivalent is the **exterior derivative**, which we've already seen once when working in two dimensions.

For this, we'll have to take $\mathbf{F} = M\,\mathbf{i} + N\,\mathbf{j} + P\,\mathbf{k}$ and turn it into a differential form in a different way: we'll turn it into the 1-form $M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z$. The exterior derivative of this 1-form is defined as

$$\mathrm{d}(M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z) = \mathrm{d}M \wedge \mathrm{d}x + \mathrm{d}N \wedge \mathrm{d}y + \mathrm{d}P \wedge \mathrm{d}z.$$

For example, suppose we start with the 1-form $xyz\,\mathrm{d}x + e^y\,\mathrm{d}y + \sin x\,\mathrm{d}z$. Then we get

$$\begin{aligned}
\mathrm{d}(xyz\,\mathrm{d}x + e^y\,\mathrm{d}y + \sin x\,\mathrm{d}z) &= \mathrm{d}(xyz) \wedge \mathrm{d}x + \mathrm{d}(e^y) \wedge \mathrm{d}y + \mathrm{d}(\sin x) \wedge \mathrm{d}z \\
&= (yz\,\mathrm{d}x + xz\,\mathrm{d}y + xy\,\mathrm{d}z) \wedge \mathrm{d}x + e^y\,\mathrm{d}y \wedge \mathrm{d}y + \cos x\,\mathrm{d}x \wedge \mathrm{d}z \\
&= xz\,\mathrm{d}y \wedge \mathrm{d}x + xy\,\mathrm{d}z \wedge \mathrm{d}x + \cos x\,\mathrm{d}x \wedge \mathrm{d}z \\
&= (xy - \cos x)\,\mathrm{d}z \wedge \mathrm{d}x - xz\,\mathrm{d}x \wedge \mathrm{d}y.
\end{aligned}$$

As before, remember the two rules for simplifying wedge products:

1. For any variable u, $\mathrm{d}u \wedge \mathrm{d}u = 0$.
2. For any two variables u, v, $\mathrm{d}u \wedge \mathrm{d}v = -(\mathrm{d}v \wedge \mathrm{d}u)$.

We begin by distributing, then use the first rule to cancel anything that simplifies to 0. Finally, to collect like terms, we use the second rule to write the answer in terms of $\mathrm{d}y \wedge \mathrm{d}z$, $\mathrm{d}z \wedge \mathrm{d}x$, and $\mathrm{d}x \wedge \mathrm{d}y$.

If you simplify the general exterior derivative $\mathrm{d}(M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z)$, you will see a surprising thing. The exterior derivative operator is the same as taking the curl! More precisely, the result, $U\,\mathrm{d}y \wedge \mathrm{d}z + v\,\mathrm{d}z \wedge \mathrm{d}x + W\,\mathrm{d}x \wedge \mathrm{d}y$, will satisfy

$$\nabla \times (M\,\mathbf{i} + N\,\mathbf{j} + P\,\mathbf{k}) = U\,\mathbf{i} + V\,\mathbf{j} + W\,\mathbf{k}.$$

[1]This document comes from an archive of the Math 3204 course webpage: http://misha.fish/archive/3204-fall-2024

This means that if we stick to differential form notation throughout, the statement of Stokes' theorem becomes:

$$\int_C M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z = \iint_S \mathrm{d}(M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z)$$

where S is an oriented surface in $\mathbb{R}^3$ and C is its compatibly oriented boundary. Or, if we define $\omega = M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z$, we get

$$\int_C \omega = \iint_S \mathrm{d}\omega.$$

2 A return to the component test

Let's look at Stokes' theorem together with the fundamental theorem of line integrals. If you recall that theorem, it says that if a vector field $\mathbf{F}$ is really the gradient field ∇f of a scalar function, then we can take line integrals of $\mathbf{F}$ simply by evaluating f. That is, if C is any curve that starts at point $\mathbf{a}$ and ends at point $\mathbf{b}$, then

$$\int_C \nabla f \cdot \mathrm{d}\mathbf{r} = f(\mathbf{b}) - f(\mathbf{a}).$$

In particular, if C is a closed curve, we can think of it as starting and ending at the same point $\mathbf{a}$, and the circulation of $\mathbf{F} = \nabla f$ around C is $f(\mathbf{a}) - f(\mathbf{a}) = 0$.

In the language of differential forms, $\nabla f = \frac{\partial f}{\partial x}\,\mathbf{i} + \frac{\partial f}{\partial y}\,\mathbf{j} + \frac{\partial f}{\partial z}\,\mathbf{k}$ corresponds to

$$\mathrm{d}f = \frac{\partial f}{\partial x}\,\mathrm{d}x + \frac{\partial f}{\partial y}\,\mathrm{d}y + \frac{\partial f}{\partial z}\,\mathrm{d}z,$$

the exterior derivative of the function (or "0-form") f. So in that form, the fundamental theorem of line integrals says that

$$\int_C \mathrm{d}f = f(\mathbf{b}) - f(\mathbf{a}).$$

A differential form which is an exterior derivative of another, lower-degree differential form is called **exact**. So, for example, $yz\,\mathrm{d}x + xz\,\mathrm{d}y + xy\,\mathrm{d}z$ is an exact 1-form, because it is equal to $\mathrm{d}(xyz)$. Also, $(xy - \cos x)\,\mathrm{d}z \wedge \mathrm{d}x - xz\,\mathrm{d}x \wedge \mathrm{d}y$ is an exact 2-form, because (as we saw in the previous section) it is equal to $\mathrm{d}(xyz\,\mathrm{d}x + e^y\,\mathrm{d}y + \sin x\,\mathrm{d}z)$.

How do you test whether a differential form is exact? Well, there is a short one-way test. A differential form ω is called **closed** if $\mathrm{d}\omega = 0$. This is relevant because:

Theorem 2.1. *All exact differential forms are closed. That is, for any differential form ω, taking its exterior derivative twice yields $\mathrm{d}(\mathrm{d}\omega) = 0$.*

Proof. There is a very general argument here that works for any degree of differential forms in any number of dimensions. The idea is that when our differential form contains a scalar function f somewhere inside it, taking the exterior derivative once turns f into

$$\mathrm{d}f = \sum_{i=1}^{n} \frac{\partial f}{\partial x_i}\,\mathrm{d}x_i$$

and taking the exterior derivative again turns it into

$$\mathrm{d}(\mathrm{d}f) = \sum_{i=1}^{n}\sum_{j=1}^{n} \frac{\partial^2 f}{\partial x_i\,\partial x_j}\,\mathrm{d}x_i \wedge \mathrm{d}x_j.$$

All "diagonal" terms (terms with $i = j$) simplify to 0, and all other terms pair up into pairs of the form:

$$\frac{\partial^2 f}{\partial x_i\,\partial x_j}\,\mathrm{d}x_i \wedge \mathrm{d}x_j + \frac{\partial^2 f}{\partial x_j\,\partial x_i}\,\mathrm{d}x_j \wedge \mathrm{d}x_i.$$

Here, the partial derivatives are equal[2], but $\mathrm{d}x_i \wedge \mathrm{d}x_j = -(\mathrm{d}x_j \wedge \mathrm{d}x_i)$, so the two terms in each pair cancel. This tells us that $\mathrm{d}(\mathrm{d}f) = 0$ for any scalar function f. For a more complicated differential form, this argument tells us that every scalar function in every term is annihilated by applying d twice, and so the result is still 0. □

This result about exterior derivatives corresponds to a statement about curl and gradient: for any scalar function f, $\boldsymbol{\nabla} \times \boldsymbol{\nabla} f = \mathbf{0}$. Both of these are also equivalent to what we called the "component test" for conservative vector fields, in the past.

Are all closed differential forms exact? The answer to that depends on the domain. It is true if the differential form is defined on all of $\mathbb{R}^n$ (for whichever n makes sense; in this class, n can be 2 or 3). We've already seen an example of a 2-dimensional differential form that is undefined at $\mathbf{0}$, and is closed but not exact.

Let's stop being fully general, and limit our attention to 1-forms $\omega = M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z$ in $\mathbb{R}^3$. Here, we are interested in whether ω is exact, because if it is, then integrals of ω are path-independent by the fundamental theorem of line integrals. Equivalently, when we ask, "Is ω exact?" we really want to know, "Does ω integrate to 0 around any closed curve?"

With the help of Stokes' theorem, we can now give a very general answer. We call a region $D \subseteq \mathbb{R}^3$ **simply connected** if it has no "holes" in the following sense: every closed curve contained in D is the boundary of some surface contained in D.

Theorem 2.2. *Suppose $\omega = M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z$ is a **closed** differential 1-form defined on some simply connected domain D. Then for every closed curve C contained in D, we have*

$$\int_C \omega = 0.$$

Proof. Because D is simply connected, there is a surface S contained in D whose boundary is C. Therefore we can apply Stokes' theorem:

$$\int_C \omega = \iint_S \mathrm{d}\omega.$$

Because ω is closed, $\mathrm{d}\omega = 0$, so integrating $\mathrm{d}\omega$ over S also yields 0. This proves the theorem. □

[2] This requires some niceness assumptions on the function f, but we make those assumptions by default when working with differential forms.

3 Two different types of holes

The theorems we proved in the previous section are very similar to theorems we proved when working over $\mathbb{R}^2$. However, there is a sense in which they are more powerful, because the nature of simply connected regions changes when we have 3 or more dimensions.

If a 2-dimensional differential form $M\,\mathrm{d}x + N\,\mathrm{d}y$ is defined everywhere except the point $(0, 0)$, then its domain is not simply connected. A closed curve around the origin is not the boundary of any bounded 2-dimensional region that avoids $(0, 0)$.

If a 3-dimensional differential form $M\,\mathrm{d}x+N\,\mathrm{d}y+P\,\mathrm{d}z$ is defined everywhere except the point $(0, 0, 0)$, there is no trouble. This domain is still simply connected. For any closed curve C that avoids the origin, we can draw a surface whose boundary is C that still avoids the origin. It would take an infinite set of undefined points to give a non-simply-connected domain in $\mathbb{R}^3$.

To see this in action, let's look at two differential forms that gave us some amount of trouble in two dimensions:

$$\omega_1 = \frac{-y\,\mathrm{d}x + x\,\mathrm{d}y}{x^2 + y^2} \qquad \text{and} \qquad \omega_2 = -\frac{x\,\mathrm{d}x + y\,\mathrm{d}y}{(x^2 + y^2)^{3/2}}.$$

Both of these are closed: $\mathrm{d}\omega_1 = \mathrm{d}\omega_2 = 0$. Both of these are undefined at the origin, however. So if C is a closed loop around the origin in $\mathbb{R}^2$, we cannot assume that the integral of ω_1 or ω_2 over C is 0.

In fact, we've already seen that ω_1 is irretrievably bad: its integral around the unit circle is 2π, not 0. But ω_2 can be rescued, and one way to rescue it is to think "outside the plane" and go to $\mathbb{R}^3$.

Define the 3-dimensional differential form ω_3 to be

$$\omega_3 = -\frac{x\,\mathrm{d}x + y\,\mathrm{d}y + z\,\mathrm{d}z}{(x^2 + y^2 + z^2)^{3/2}}.$$

This extension of ω_2 is "backwards compatible" in the sense that when we set $z = 0$, ω_3 turns back into ω_2. In particular, if C is any closed curve in the plane that avoids the origin, then

$$\int_C \omega_2 = \int_C \omega_3.$$

Also, it is true (though it takes some tedious algebra to check) that ω_3 is still closed: $\mathrm{d}\omega_3 = 0$.

Finally, ω_3 is defined everywhere except the point $(0, 0, 0)$. This is a simply connected subset of $\mathbb{R}^3$. Therefore, by Theorem 2.2, we have

$$\int_C \omega_3 = 0$$

for every closed curve C in $\mathbb{R}^3$ that avoids the origin. In particular, since this is true for the xy-plane, we conclude that the differential form ω_2 also has this property, even though Theorem 2.2 does not apply to it directly.