Math 208H

Topics since the third exam

Vector fields

A vector field is a field of vectors, i.e., a choice of vector $F(x, y)$ (or $F(x, y, z)$) in the plane for every point in some part of the plane (the domain of F), and similarly in 3-space. We can think of F as $F(x, y) = (F_1(x, y), F_2(x, y))$; each coordinate of F is a function of several variables. We can represent a vector field pictorially by place the vector $F(x, y)$ in the plane with its *tail* at the point (x, y) . A vector field is therefore a choice of a direction (and magmitude) at each point in the plane (or 3-space...). Such objects naturally occur in many disciplines, e.g., a vector field may represent the wind velocity at each point in the plane, or the direction and magnitude of the current in a river.

One of the most important class of vector fields that we will encounter are the *gradient vector fields*. If we have an (ordinary) function $f(x, y, z)$ of several variables, then for each point (x, y, z), $\nabla(f)$ can be thought of as a *vector*, which we have in fact already taken to drawing with its tail at the point (x, y, z) (so that, for example, we can use it as a normal vector for the tangent plane to the graph of f). Many vector fields are gradient vector fields, e.g., $(y, x) = \nabla(xy)$; one of the questions we will need to answer is: 'How do you tell when a vector field is a gradient vector field?'. We shall see several answers to this question shortly.

Line Integrals

The basic idea

We introduced vector fields $F(x, y)$ in large part because these are the objects that we can most naturally integrate over a (parametrized) curve. The reason for this is that along a curve we have the notion of a velocity vector $\vec{v}$ at each point, and we can *compare* these two vectors, by taking their dot product. This tells us the extent to which F points in the direction of $\vec{v}$. Integration is all about taking averages, and so we can think if the integral of F over the curved C as measuring the *average* extent to which F points in the same direction as C.

We can set this up as we have all other integrals, as a limit of sums. Picking points $\vec{c_i}$ strung along the curve C, we can add together the dot products $F(\vec{c_i}) \bullet (\vec{c_{i+1}} - \vec{c_i})$, and then take a limit as the lengths of the vectors $\vec{c_{i+1}} - \vec{c_i}$ between consecutive points along the curve goes to 0. We denote this number by

$$\int_C F \bullet d\vec{r}$$

Such a quantity can be interpreted in several ways; we will mostly focus on the notion of *work*. If we interpret F as measuring the amount of force being applied to an object at each point (e.g., the pull due to gravity), then $\displaystyle\int_C F \bullet d\vec{r}$ measures the amount of work done by F as we move along C. In other words, it measures the amount that the force field F *helped* us move along C (since moving in the same direction, it helps push us along, while when moving opposite to it, it would slow us down).

In the case that F measures the current in a river or lake or ocean, and C is a *closed* curve (meaning it begins and ends at the same point), we interpret the integral of F along C as the *circulation* around C, since it measures the extent to which the current would *push* you around the curve C.

Of course, as usual, we would never want to *compute* a line integral by taking a limit! But if we use a parametrization of C, we can interpret $\displaystyle\int_C F \bullet d\vec{r}$ as an 'ordinary' integral. The idea is that if we use a parametrization $\vec{r}(t)$ for C then $F(\vec{c_i}) \bullet (\vec{c_{i+1}} - \vec{c_i})$ becomes

$$F(\vec{r}(t_i)) \bullet (\vec{r}(t_{i+1}) - \vec{r}(t_i))$$

But using tangent lines, we can approximate $\vec{r}(t_{i+1}) - \vec{r}(t_i)$ by $\vec{r}'(t_i)(t_{i+1} - t_i) = \vec{r}'(t_i)\Delta y$. So we can instead compute our line integral as

$$\int_C F \bullet d\vec{r} = \int_a^b F(\vec{r}(t)) \bullet \vec{r}'(t)\ dt$$

where $\vec{r}$ parametrizes C with $a \leq t \leq b$.

Some notation that we will occasionally use: If the vector field $F = (P, Q, R)$ and $\vec{r}(t) = (x(t), y(t), z(t))$, then $d\vec{r} = (dx, dy, dz)$, so $F \bullet d\vec{r} = Pdx + Qdy + Rdz$. So we can write $\displaystyle\int_C F \bullet d\vec{r} = \int_a^b Pdx + Qdy + Rdz$

Gradient fields and path independence

In general, the computation of a line integral can be quite cumbersome, in part because we need to evaluate the vector field F *at the point* $\vec{r}(t)$, which can yield quite complicated formulas. But there is one class of vector fields that cause a lot less trouble to integrate: gradient vector fields. This is because we can compute:

if $F = \nabla(f)$, then $\displaystyle F(\vec{r}(t)) \bullet \vec{r}'(t) = \frac{\partial f}{\partial x}\frac{dx}{dt} + \frac{\partial f}{\partial y}\frac{dy}{dt} + \frac{\partial f}{\partial z}\frac{dz}{dt} = \frac{d}{dt}(f(\vec{r}(t)))$

so $\displaystyle\int_C F \bullet d\vec{r} = \int_a^b F(\vec{r}(t)) \bullet \vec{r}'(t)\ dt = \int_a^b \frac{d}{dt}(f(\vec{r}(t)))\ dt = f(\vec{r}(b)) - f(\vec{r}(a))$. We call this the *Fundamental Theorem of Calculus for Line Integrals.*

We say that a vector field f is *path-independent* (or *conservative*) if the value of a line integral over a curve C depends only on what the endpoints P, Q of C are, i.e., the integral would be the same for any *other* curve running from P to Q. Our result right above can then be interpreted as saying that gradient vector fields are conservative. What is amazing is that it turns out that every conservative vector field F is the gradient vector field for some function f. We can alctually write down the function, too (by stealing an idea from the Fundamental Theorem of Calculus...), as

$\displaystyle f(x, y) = \int_C F \bullet d\vec{r}$, where C is *any* curve from (0,0) to (x, y).

Green's Theorem

All of which is very nice, but far too theoretical for practical purposes. What we need are simple ways to tell that a vector field is conservative, and to build the function f when it is. Luckily, this is not too hard!

First, a slight reinterpretation: a vector field F is path-independent if $\int_C F \bullet d\vec{r}$=0 for every *closed* curve C.

If F is conservative, then $F = (F_1, F_2) = (f_x, f_y)$ for some function f. But then by using the equality of mixed partials for f, we can then conclude that we *must* have $(F_1)_y = (F_2)_x$. In fact, this is enough to *guarantee* that F is conservative; this is because of *Green's Theorem*: defining the *curl* of F to be $(F_2)_x - (F_1)_y$, we have

If R is a region in the plane, and C is the boundary of R, parametrized so that we travel *counterclockwise* around R, then

$$\int_C F \bullet d\vec{r} = \int_R \text{curl}(F)\ dA$$

In particular, if the curl is 0, then the integral of F along C is always 0 for every closed curve, so F is conservative.

We can actually use this result to evaluate line integrals *or* double integrals, whichever we wish. For example, we can compute the area of a region R as a line integral, by integrating the function 1 over R, and then using a vector field around the boundary whose curl is 1, such as $(0, x)$ or $(-y, 0)$ or $(y, 2x)$ or ....

This allows us to spot conservative vector fields quite easily, but doesn't tell us how to compute the function it is the gradient *of* (called its *potential function*). But in practice this is not tootough; we simply write down a function f with $\frac{\partial f}{\partial x}$=$F_1$ (e.g., $f(x, y) = \int F_1(x, y)dx$). This is actually a *family* of functions, because we have the constant of integration to worry about, which we should *really* think of as a *function of* y (because we integrated a function of two variables, dx). To figure out *which* function of y, simply take $\frac{\partial}{\partial y}$ of your function, and compare with $F_2 = \frac{\partial f}{\partial y}$; just adjust the constant of integration accordingly.

Finally, there is a similar result for vector fields in dimension 3; for $F = (F_1, F_2, F_3)$, we can define curl(F) = "$\nabla \times F$" $= ((F_3)_y - (F_2)_z, -((F_3)_x - (F_1)_z), (F_2)_x - (F_1)_y)$

Then $F = \nabla f$ exactly when curl(F) = (0,0,0) ; and we can actually *construct* f using a procedure analogous to the one we came up with for vector fields with two variables.

Flux Integrals

The basic idea

The basic idea is that we can also integrate vector fields (in 3-space) over a *surface.* The interpretation we will use is that we are measuring the amount of fluid flowing through a surface (e.g., a cell membrane) immersed in the fluid.

We can think of a wire-frame surface sitting in a river; we would like to compute the amount of water flowing (each second, perhaps) flowing through the surface. (Or, you can think of computing the amount of rain falling on the surface of your body...)

Our input is a (velocity) vector field F, and a surface S, described in some fashion (e.g., as the graph of a function of two variables). The idea is that a piece of surface which is tilted with respect to the vector field will not contribute much to the total. In other words, the amount flowing through the surface is related to the extent to which the (**unit**) *normal vector* for the surface is pointing in the same direction as F. We measure this with the dot product, $F \bullet \vec{n}$. This amount is also clearly proportional to the *size* of the surface; twice as much surface will give twice as much flow. This leads us to believe that what we need to add up in order to compute the flow through the surface is $F \bullet \vec{n}\ dA$ (to take into account tilt and size). So we define the *flux integral* of a vector field F over a surface S to be

$$\int_S \vec{F} \bullet d\vec{A} = \int_S (\vec{F} \bullet \vec{n})\ dA$$

Now at every point of the surface S, we actually have *two* choices of unit normal vector $\vec{n}$; we will see in the next section how to make a more or less 'obvious' consistent choice of normal, the *outward pointing normal.* For example, if S is a sphere of radius R, centered at (0,0,0), the outward unit normal at (x, y, z) is just $(x/R, y/R, z/R)$. If we choose F to be this *same* vector, then it is easy to see that $F \bullet \vec{n} = 1$, and so our flux integral will just compute the area of the surface S.

Computing using graphs, cylindrical, and spherical coordinates

Of course, we don't want to compute flux integrals as limits of sums, either! What we need is some approaches to calculating $\vec{n}\ dA$. The most general approach to this is to *parametrize* the surface S, in a manner just like our parametrizations of curves. We think of the surface S (in xyz-space) as the image of a region R (in uv-space) under a change of variables

$$x = x(u, v)\ ,\ y = y(u, v)\ ,\ z = z(u, v)$$

A small rectangle $du \times dv$ is carried by the parametrization to a small parallelogram with sides

$T_u = (x_u, y_u, z_u)\ du$ and $T_v = (x_v, y_v, z_v)\ dv$

This is the piece of the surface that we work with above. Its unit normal is given by

$\dfrac{T_u \times T_v}{||T_u \times T_v||}$

and dA (taking a cue from our change of variables formula!) is $||T_u \times T_v||\ du\ dv$. So $\vec{n}\ dA = T_u \times T_v\ du\ dv$, and so

$\int\int_S (\vec{F} \bullet \vec{n})\ dA = \int\int_R \vec{F} \bullet (T_u \times T_v)\ du\ dv$

In particular, there are three cases we can consider separately, coming from our standard coordinate systems:

Suppose S is the graph of a function f, having domain R in the plane. What we would really like to do is to compute the flux integral as the integral of a function over R. To do this, we note that the vector $v = (-f_x, -f_y, 1)$ is normal to the graph of f; it's the normal vector we used to express the tangent plane to the graph of f. It just so happens that $v = (1, 0, f_x) \times (0, 1, f_y)$, and so its length is equal to the area of the parallelogram that these two vectors span. But!, these are exactly the paralleograms we would use to approximate the graph, i.e., this length is also dA. So, $\vec{n}\ dA = (-f_x, -f_y, 1)$, and so

$$\int_S F \bullet \vec{n}\ dA = \int_R F(x, y, f(x, y)) \bullet (-f_x, -f_y, 1)\ dx\ dy\ dz$$

We can also use cylindrical and spherical coordinates, in special cases. If S is a piece of a cylinder, given by $r = r_0$, for θ and z in some range of values R, then the outward normal at r_0, θ, z is $(\cos\theta, \sin\theta, 0)$, while $dA = r_0\ d\theta\ dz$, so

$$\int_S F \bullet \vec{n}\ dA = \int_R F(r_0\cos\theta, r_0\sin\theta, z) \bullet (\cos\theta, \sin\theta, 0) r_0\ d\theta\ dz$$

If S is a piece of sphere, given by $\rho = \rho_0$ for θ and ϕ in some range R of values, then the outward normal is $(\cos\theta\sin\phi, \sin\theta\sin\phi, \cos\phi)$ while dA is $\rho_0^2 \sin\phi\ d\theta\ d\phi$, so

$$\int_S F \bullet \vec{n}\ dA =$$
$$\int_R F(\rho_0\cos\theta\sin\phi, \rho_0\sin\theta\sin\phi, \rho_0\cos\phi) \bullet (\cos\theta\sin\phi, \sin\theta\sin\phi, \cos\phi)\ \rho_0^2\sin\phi\ d\theta\ d\phi$$

The divergence of a vector field

In terms of the coordinates $\vec{F} = (F_1, F_2, F_3)$ of a vector field, the divergence is

$$\text{div}(F) = (F_1)_x + (F_2)_y + (F_3)_z$$

It can be identified with the *flux density* of the vector field $\vec{F}$ at a point P: this should be though of as the flux integral of F through a tiny box around the point P. This measures the extent to which the vector field is 'expanding', at each point.

A vector field F is *divergence-free* if $\text{div}(F) = 0$. For example, $F = (y, z, x)$ is divergence free, but $F = (x, y, z)$ is not; $\text{div}(F) = 3$.

Some formulas that can help to calculate divergence:

$$\text{div}(fF) = (\nabla f) \bullet F + f \cdot (\text{div}F)$$
$$\text{div}(F \times G) = (\text{curl } F) \bullet G - F \bullet (\text{curl}G) \qquad \text{in 3-space}$$
$$\text{div}(\text{curl}(\vec{F})) = 0 \qquad \text{in 3-space}$$

It turns out that this last result works the other way; a vector field F , defined over an entire box, which is divergence-free, is the curl of *some* other vector field G.

The Divergence Theorem

If W is a region in 3-space, it boundary is a surface S. (S might actually consist of several pieces; this won't really effect our discussion.) We can choose normal vectors for each piece of S by insisting that $\vec{n}$ alway points *out* of W. Then we have, for any vector field F *which is defined everywhere in* W:

The Divergence Theorem: $\displaystyle\int_S \vec{F} \bullet d\vec{A} = \int_W (\text{div } F)\, dV$

In other words, we can compute flux integrals over a surface S that forms the boundary of a region W, by computing the integral of a *different* function over W. This is especially useful when the vector field is divergence-free; for example if the region W has two surfaces for boundary and F is divergence-free, then the flux integral of F over one surface, with normals pointing out of W, is *equal* to the flux integral of F over the *other* surface, with normals pointing *into* W. Even if F is not divergence-free, we can compute the flux integral of one as the flux integral of the other *plus* the triple integral of the divergence over W.

The curl of a vector field

We have already met the curl of a vector field $\vec{F} = (F_1, F_2)$ in 2-space; there is a similar definition for a vector field $\vec{F} = (F_1, F_2, F_3)$ in 3-space, except that it is a vector. In terms of coordinates:

$$\text{curl}(\vec{F}) = ((F_3)_y - (F_2)_z, -((F_3)_x - (F_1)_z), (F_2)_x - (F_1)_y)$$

Its physical interpretation is as the direction where the *circulation density* of the vector field $\vec{F}$, at the point P, is the *largest.* The circulation density measures the extent to which objects caught up in a (velocity) vector field 'want' to rotate with their axis pointing in the direction of a (unit) vector $\vec{n}$, and is computed as the limit, as the side lengths go to 0, of the line integral of $\vec{F}$ around the boundary of a little square around P and *perpendicular* to $\vec{n}$, divided by the area of the square. In terms of the curl, it can be computed as

$$\text{circ}_{\vec{n}}(\vec{F}) = \text{curl}(\vec{F}) \bullet \vec{n}$$

We have already used the curl to detect conservative vector fields; this stems from the formula

$$\text{curl}(\nabla\vec{F}) = (0,0,0)$$

A vector field $\vec{F}$ is *curl-free* if $\text{curl}\vec{F} = (0{,}0{,}0)$. This means that in any *box* in which $\vec{F}$ is defined, $\vec{F}$ is a gradient vector field (although it is possible that $\vec{F}$ cannot be expressed as the gradient of a function everywhere that $\vec{F}$ is defined *at the same time*; the standard example of this is the vector field

$$\vec{F} = (\frac{-y}{x^2+y^2}, \frac{x}{x^2+y^2}, 0)$$

$\vec{F}$ is curl-free, but it is not a gradient vector field, since (as you can check) the line integral of $\vec{F}$ around the circle of radius one in the x-y plane with center (0,0,0) is 2π. Green's Theorem does not work, because $\vec{F}$ (and so its curl) is not defined on the entire disk with boundary the circle.)

Stokes' Theorem

If S is a surface in 3-space, with a normal orientation $\vec{n}$, the boundary of S is a colection of paramatrized curves (there can easily be more than one, e.g, if S is a cylinder). We can orient each curve using a *right-hand rule*; if we stand on the curve and walk along it the chosen orientation with our heads pointing in the direction of $\vec{N}$, then the surface S dshould always be on our left. Then Stokes' Theorem says that, for any vector field $\vec{F}$ defined everywhere on S:

$$\int_C \vec{F} \bullet d\vec{r} = \int_S (\text{curl}\vec{F}) \bullet d\vec{A}$$

This allows us to compute line integrals as flux integrals, and, with a little work, flux integrals as line integrals.

For example, it says that the line integral of a curl-free vector field $\vec{F}$ around a closed curve is always 0, *so long as* the curve is the boundary of a surface contained entirely in the domain of $\vec{F}$.

We say that a vector field $\vec{F}$ is a *curl field* if $\vec{F} = \text{curl}(\vec{G})$ for some vector field $\vec{G}$. $\vec{G}$ is called a *vector potential* of $\vec{F}$. Then Stokes' Theorem says that, for any surface S in the domain of $\vec{F}$ with boundary C,

$$\int_S \vec{F} \bullet d\vec{A} = \int_S \text{curl}\vec{G} \bullet d\vec{A} = \int_C \vec{G} \bullet d\vec{r}$$

So, for example, for a curl field $\vec{F}$ and *two* surfaces S_1 and S_2 with the *same* boundary C, we have

$$\int_{S_1} \vec{F} \bullet d\vec{A} = \int_{S_2} \vec{F} \bullet d\vec{A}$$

So the flux integral of a curl field *really* depends just on the boundary of the surface, not on the surface.

We can test for whether or not $\vec{F}$ is a curl field, using the divergence, since $\text{div}(\text{curl}(\vec{G})) = 0$, so a curl field must be divergence-free. (The opposite, as we have seen, is *almost* true; it is true, for example, if the vector field is defined in a big box.)

The whole idea behind these three theorems (Green's, Divergence, and Stokes') is that the integral of one kind of function over one kind of region can be computed instead as the integral of *another* kind of function over the *boundary* of the region.

Green's: Integral of the vector field $\vec{F}$ over a closed curve in the plane equals integral of its curl of $\vec{F}$ over the region in the plane that the curve bounds.
Divergence: The flux integral of a vector field $\vec{F}$ through the boundary of a region in 3-space equals the integral of the divergence of $\vec{F}$ over the region in 3-space.
Stokes': The line integral of the vector field $\vec{F}$ over a closed curve C in 3-space equals the flux integral of the curl of $\vec{F}$ over any surface S that has C as its boundary.

Note that Green's Theorem is really just a special case of Stokes' (where the curve C lies in the plane, and the third coordinate of $\vec{F}$ just happens to be 0). All of these, like the Fundamental Theorem of Line Integrals, are really a kind of Fundamental Theorem of Calculus, where we are computing a kind of integral by instead computing something else across the boundary of the region we are interested in.