Math 208H

Topics since the third exam

(Parametrized) surfaces and surface area

Just as curves can be represented as the image of a function from an interval into 3-space, a surface Σ in space can be parametrized as a function (of two variables), or really three functions (x, y, and z) from a region R in the u,v-plane;

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

For example, the graph of a function $f : R \to \mathbb{R}$ can be parametrized by $x = u$, $y = v$, $z = f(u,v)$ for (u,v) in R. For such a surface Σ we can formulate an integral which will compute its *surface area*; as usual, we start by approximating the surface by things whose area we can compute, in this case, by parallelograms. We worked out the basic technology when we developed change of variables formulas; at any point P of Σ, $P = T(u_0, v_0)$, a small rectangle in u,v-space is sent by T to something approximated by the parallelogram spanned by the vectors $T_u = (x_u, y_u, z_u)$ and $T_v = (x_v, y_v, z_v)$ with tails at P. The area of Σ can therefore be approximated by a sum looking like $\sum ||T_u \times T_v|| \Delta u \Delta v$, since these are the areas of the parallelograms. But this in turn is an approximation of an integral over the region R in $u.v$-space! SO we define

$$\text{Area of } \Sigma \text{ - } \int\int_R ||T_u \times T_v|| \ du \ dv$$

As an example, using sphereical coordinates to parametrixze a sphere (just set the radius to be a constant), we find that the area of a sphere of radius r is $4\pi r^2$.

Because we could interpret the above integral as the area of Σ, we obtain the important observation that this integral does not 'really' depend on the parametrization, but only on the surface being parametrized.

Flux Integrals

More generally, we can integrate a function f, whose domain contains the surface σ, over σ. The basic idea is that surface area is essentially the integral of the function 1 over Σ, and so the proper formulation of the integral of f should be

$$\int\int_\Sigma f \ dA = \int\int_R f(x(u,v), y(u,v), z(u,v)) \ ||T_u \times T_v|| \ du \ dv$$

Again, we can see that this is independent of the parametrization of Σ, since it can be thought of as approximated by (Riemann) sums of the value of f times the area of small pieces of Σ, representing an 'average' value of f over the surface.

Perhaps more importantly (the concept, at least, appears throughout the sciences), we can integrate *vector fields* (in 3-space) over a surface Σ. 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) 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, typically as a parametrized surface. 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 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 often choose the *outward pointing normal*, pointing away from the region that our surface surrounds. Other times, we choose the *upward pointing normal*, the one with positive z-coordinate. 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 $F \bullet \vec{n} = 1$, and so our flux integral will just compute the area of the surface S.

Computing using a parametrization of a surface

For computations, what we need is some approaches to calculating $\vec{n}\ dA$. The most general approach to this is to use a parametrization of the surface S, as we did for surface area. The unit normal is given by

$\dfrac{T_u \times T_v}{||T_u \times T_v||}$ and dA 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 parametrizations we know, 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

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

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. Computing the vector field G can take some work, though; the general technique can be found listed under 'Helmholtz decomposition', if you're interested...

The Divergence Theorem

If W is a region in 3-space, its 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 at the center of the 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.