Math 208H
Topics for the third exam

(Technically, everything covered on the first two exams, *plus*)

Triple Integrals

Triple integrals are just like double integrals, only more so. We can define them as a limit of a huge sum; here the terms in the sum would be the value of the function f times the volume of a tiny rectangular box. The usual interpretation of a triple integral arises by thinking of the function f as giving the density of the matter at each point of a solid region W in 3-space. Since density times volume is mass, the integral of f over the region W would compute the mass of the solid object occupying the region W. In the special case that f is the function 1, the integral will compute the volume of the region W.

Again, as with double integrals, the way we really compute a triple integral is as a (triply) iterated integral. You pick a direction to slice (x=constant, y=constant, or z=constant) W up, and compute the integral of f over each slice. Each of these is a double integral (computed as an iterated integral), whose value *depends* on the variable you sliced along. To compute the integral over W, you integrate these double integrals over the last variable, getting three iterated integrals.

Put slightly differently, you can evaluate a triple integral by integrating out each variable, one at a time. Typically, we start with z, since our region W is usually described as the region lying between the graphs of two functions, given as z=blah and z=bleh . The idea is to first, for each fixed value of x and y, integrate the function f, dz, from blah to bleh. (Their resulting values depend on x and y, i.e., are a *function* of x and y.) Then we integrate over the region, R, in the plane consisting of the points (x, y) such that the vertical line *hits* the region W. We usually call this region R the *shadow* of W in the x-y plane. In symbols

$\int\int\int_W f\ dV = \int\int_R (\int_{a(x,y)}^{e(x,y)} f(x, y, z)\ dz)\ dA$

For example, the integral of a function over the region lying above the x-y plane and inside the sphere of radius 2, centered at the origin, would be computed as

$\int\int_R (\int_0^{\sqrt{4-x^2-y^2}} f(x, y, z)\ dz)\ dA = \int_{-2}^{2} (\int_{-\sqrt{4-x^2}}^{\sqrt{4-x^2}} (\int_0^{\sqrt{4-x^2-y^2}} f(x, y, z)\ dz)\ dy)\ dx$

where R is the shadow of W (in this case, the disk of radius 2, centered at the origin, in the x-y plane).

Change of variables for Triple Integrals

As with double integrals, we can carry out a change of coordinates for 3 variables; we then write

$$x = x(s, t, u),\ y = y(s, t, u), \text{ and } z = z(s, t, u)$$

A little box with sides of length ds, dt, and du gets carried to a little parallelopiped, with sides the vectors

$$(x_s, y_s, z_s)\ ds,\ (x_t, y_t, z_t)\ dt, \text{ and } (x_u, y_u, z_u)\ du$$

(call these V_s, V_t, and V_u). This has volume $|V_s \bullet (V_t \times V_u)|$, which is the Jacobian of this change of variables, and serves as the necessary "fudge factor" to express an integral in terms of s, t, and u.

As before, this usually works best when our change of variables functions x, y, z carry what we would consider a 'nice' region in (s, t, u)-space to the region R in (x, y, z)-space that we find ourselves needing to integrate over. Good choices are rectangular boxes or regions lying under the graph of a function $u = u(s, t)$ over some nice region in the (s, t)-plane. As before, we often use change of variables to make the region we integrate over nicer, at the expense of making the quantity we are integrating less nice (in distinction with the usual philosophy of u-substitution in one-variable calculus!).

Triple integrals with spherical and cylindrical coordinates

It turns out that we can impose *two* new coordinate systems on 3-space, analogous to polar coordinates in 2-space; each can sometimes be used render an integration problem more tractible, usually by making the *region* we integrate over more 'routine'.

With cylindrical coordinates, we simply replace (x, y, z) with (r, θ, z), i.e., use polar coordinates in the xy-plane. In the new coordinate system, $dV = (r\ dr\ d\theta)\ dz$, since that will be the volume of a small 'cylinder' of height dz lying over the small sector in the xy-plane that we use to compute dA above.

Usually, we will actually integrate in cylindrical coordinates in the order $dz\ dr\ d\theta$, since this coordinate system is most useful when the cross-sections z=constant of our region are *disks* (so the limits of integrations for z will depend only on r).

Spherical coordinates are much like polar coordinates; we describe a point (x, y, z) by distance (which we call ρ and direction, except we need to use *two* angles to completely specify the direction; first, the angle θ that $(x, y, 0)$ makes with the x-axis in the xy-plane, and then the angle ϕ that the line through our point makes with the (positive) z-axis (which we can always assume lies between 0 and π). A little trigonometry leads us to the formulas

$$(x, y, z) = (\rho \cos\theta \sin\phi, \rho \sin\theta \sin\phi, \rho \cos\phi)$$

Again, the idea is that regions troublesome to describe in rectangular coordinates can be far more routine to describe spherically; for example, the inside of a sphere of radius R_0 can be described as the *rectangle* $0 \leq \rho \leq R_0$, $o \leq \theta \leq 2\pi$, and $0 \leq \phi \leq \pi$.

It is a bit more work to compute what dV is in spherical coordinates; computing the Jacobian, we find that it is

$$dv = \rho^2 \sin\phi\ d\rho\ d\theta\ d\phi$$

.

So the 'change of variables formula' for spherical coordinates reads:

$$\int\int\int_W f(x, y, z)\ dV = \int\int\int_R f(\rho\cos\theta\sin\phi, \rho\sin\theta\sin\phi, \rho\cos\phi)\ \rho^2 \sin\phi\ d\rho\ d\theta\ d\phi$$

So, for example, the integral of the function $f(x, y, z) = xz$ over the top half of a ball of radius 5 could be computed as

$$\int_0^5 \int_0^{2\pi} \int_0^{\pi/2} (\rho\cos\theta\sin\phi)(\rho\cos\phi)\ (\rho^2\sin\phi)\ d\phi\ d\theta\ d\rho$$

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 magnitude) 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, or the direction and magnitude of a force acting at each point.

One of the most important classes 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). Gradient vector fields point 'uphill' with respect to their parent function f, that is, in the direction of greatest increase of f. Many vector fields are gradient vector fields, e.g., $(y,x) = \nabla(xy)$; one of the questions we ?will find it useful 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

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 $\int_C F \bullet d\vec{r}$ 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 as F, it helps push us along, while when moving opposite to it, it would hinder us).

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 $\int_C F \bullet d\vec{r}$ 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 $\vec{r}'(t_i)(t_{i+1} - t_i) = \vec{r}'(t_i)\Delta t$. 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 $a \le t \le b$.

An important point is that the value of the line integral is <u>independent</u> of the parametrization of the curve (so long as we traverse γ in the same direction); this follows from our original description which did not really <u>use</u> a parametrization, or directly (via u-substitution) by considering a change of parametrization (as a change of variable, u for t).

Some notation that we will occasionally use: If the vector field $F = (M,N,P)$ and $\vec{r}(t) = (x(t), y(t), z(t))$, then $d\vec{r} = (dx, dy, dz)$, so $F \bullet d\vec{r} = Mdx + Ndy + Pdz$. So we can write

$$\int_C F \bullet d\vec{r} = \int_a^b M dx + N dy + P dz = \int_a^b (M\frac{dx}{dt} + N\frac{dy}{dt} + P\frac{dz}{dt})\ dt$$

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

so $\int_C F \bullet d\vec{r} = \int_a^b F(\vec{r}(t)) \bullet \vec{r}'(t)\ dt = \int_a^b \frac{d}{dt}(f(\vec{r}(t)))\ dt = f(\vec{r}(b)) - f(\vec{r}(a))$. 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

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

Showing that F is actually the gradient of this function f is, in the end, a relatively straightforward computation.

Green's Theorem

All of which is very nice, but far too theoretical for practical purposes. What we need are quicker ways to tell us that a vector field is conservative, and to build the function f when it is.

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. This really is a remarkable result, although once you suspect it is true, it can essentially be verified by a straighforward computation (as in our handout).

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 quickly (their curl must be 0), but it doesn't tell us how to compute the function it is the gradient *of* (called its *potential* *function*). But in practice this is not too tough; we 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. For $\vec{F} = (M, N, P)$, if curl($\vec{F}$)=$\vec{0}$, then $\vec{F} = \nabla f$ for $f(x,y,z) = \int M \ dx = m(x,y,z) + c(y,z)$, so $c_y = f_y - m_y = N - m_y$, so $c(y,z) = \int N - m_y \ dy = n(y,z) + k(z)$, so $f = m + n + k$, so $k'(z) = f_z - m_z - n_z = P - m_z - n_z$, so $k(z) = \int P - m_z - n_z \ dz$. At every step, we know the function we are integrating, so we can, in principle, carry these integrations out. Note that $c_y = N - m_y$ must be a function of y and z alone, and $k'(z) = P - m_z - n_z$ must be a function of z alone; if we find this not to be true, it is a clue that the original vector field was not a gradient field!