MIT OpenCourseWare
http://ocw.mit.edu

18.02 Multivariable Calculus, Fall 2007

Please use the following citation format:

Denis Auroux. *18.02 Multivariable Calculus, Fall 2007*. (Massachusetts
Institute of Technology: MIT OpenCourseWare). http://ocw.mit.edu
(accessed MM DD, YYYY). License: Creative Commons Attribution-
Noncommercial-Share Alike.

Note: Please use the actual date you accessed this material in your citation.

For more information about citing these materials or our Terms of Use, visit:
http://ocw.mit.edu/terms

Let's get going. Yesterday we learned about flux and we have seen the first few examples of how to set up and compute integrals for a flux of a vector field for a surface. Remember the flux of a vector field F through the surface S is defined by taking the double integral on the surface of F dot n dS where n is the unit normal to the surface and dS is the area element on the surface. As we have seen, for various surfaces, we have various formulas telling us what the normal vector is and what the area element becomes. For example, on spheres we typically integrate with respect to phi and theta for latitude and longitude angles.

On a horizontal plane, we would just end up degrading dx, dy and so on. At the end of lecture we saw a formula. A lot of you asked me how we got it. Well, we didn't get it yet. We are going to try to explain where it comes from and why it works. The case we want to look at is if S is the graph of a function, it is given by z equals some function in terms of x and y. Our surface is out here. Z is a function of x and y. And x and y will range over some domain in the x, y plane, namely the region that is the shadow of the surface on the x, y plane.

I said that we will have a formula for n dS which will end up being plus/minus minus f sub x, minus f sub y, one dxdy, so that we will set up and evaluate the integral in terms of x and y. Every time we see z we will replace it by f of xy, whatever the formula for f might be. Actually, if we look at a very easy case where this is just a horizontal plane, z equals constant, the function is just a constant, well, the partial derivatives become just zero. You get <0, 0, 1> dx, dy.

That is what you would expect for a horizontal plane just from common sense. This is more interesting, of course, if a function is more interesting. How do we get that? Where does this come from? We need to figure out, for a small piece of our surface, what will be n delta S. Let's say that we take a small rectangle in here corresponding to sides delta x and delta y and we look at the piece of surface that is above that. Well, the question we have now is what is the area of this little piece of surface and what is its normal vector? Observe this little piece up here. If it is small enough, it will look like a parallelogram.

I mean it might be slightly curvy, but roughly it looks like a parallelogram in space. And so we have seen how to find the area of a parallelogram in space using cross-product. If we can figure out what are the vectors for this side and that side then taking that cross-product and taking the magnitude of the cross-product will give us the area. Moreover, the cross-product also gives us the normal direction. In fact, the cross-product gives us two in one.

It gives us the normal direction and the area element. And that is why I said that we will have an easy formula for n dS while n and dS taken separately are more complicated because you would have to actually take the length of a direction of this guy. Let's carry out this problem. Let's say I am going to look at a small piece of the x, y plane. Here I have delta x, here I have delta y, and I am starting at some point

(x, y). Now, above that I will have a parallelogram on my surface. This point here, the point where I start, I know what it is. It is just (x, y). And, well, z is f(x, y).

Now what I want to find, actually, is what are these two vectors, let's call them U and V, that correspond to moving a bit in the x direction or in the y direction? And then U cross V will be, well, in terms of the magnitude of this guy will just be the little piece of surface area, delta S. And, in terms of direction, it will be normal to the surface. Actually, I will get just delta S times my normal vector. Well, up to sign because, depending on whether I do U V or V U, I might get the normal vector in the direction I want or in the opposite direction. But we will take care of that later. Let's find U and V. And, in case you have trouble with that small picture, I have a better one here.

Let's keep it just in case this one gets really too cluttered. It really represents the same thing. Let's try to figure out these vectors U and V. Vector U starts at the point x, y, f of x, y and it goes to -- Whereas, its head, well, I will have moved x by delta x. So, x plus delta x and y doesn't change. And, of course, the z coordinate has to change. It becomes f of x plus delta x and y. Now, how does f change if I change x a little bit? Well, we have seen that it is given by the partial derivative f sub x. This is approximately equal to f of x, y plus delta x times f sub x at the given point x, y.

I am not going to add it because the notation is already long enough. That means my vector U, well, approximately because I am using this linear approximation, <delta x, 0, f sub x times delta x>. Is that OK with everyone? Good. Now, what about V? Well, V works the same way so I am not going to do all the details. When I move from here to here x doesn't change and y changes by delta y. X component nothing happens. Y component changes by delta y. What about the z component? Well, f changes by f sub y times delta y. That is how f changes if I increase y by delta y.

I have my two sides. Now I can take that cross-product. Well, maybe I will first factor something out. See, I can rewrite this as one, zero, f sub x times delta x. And this one I will rewrite as zero, one, f sub y delta y. And so now the cross-product, n hat delta S up to sign is going to be U cross V. We will have to do the cross-product, and we will have a delta x, delta y coming out. I am just saving myself the trouble of writing a lot of delta x's and delta y's, but if you prefer you can just do directly this cross-product.

Let's compute this cross-product. Well, the i component is zero minus f sub x. The y component is going to be, well, f sub y minus zero but with the minus sign in front of everything, so negative f sub y. And the z component will be just one times delta x delta y. Does that make sense? Yes. Very good. And so now we shrink this rectangle, we shrink delta x and delta y to zero, that is how we get this formula for n dS equals negative fx, negative fy, one, dxdy. Well, plus/minus because it is up to us to choose whether we want to take the normal vector point up or down.

See, if you take this convention then the z component of n dS is positive. That corresponds to normal vector pointing up. If you take the opposite signs then the z component will be negative. That means your normal vector points down. This one is with n pointing up. I mean when I say up, of course it is still perpendicular to the surface. If the surface really has a big slope then it is not really going to go all that much up, but more up than down. OK. That is how we get the formula. Any questions? No. OK. That is a really useful formula.

You don't really need to remember all the details of how we got it, but please remember that formula. Let's do an example, actually. Let's say we want to find the flux of the vector field z times k, so it is a vertical vector field, through the portion of the paraboloid z equals x^2 y^2 that lives above the unit disk. What does that mean? z = x^2 y^2. We have seen it many times. It is this parabola and is pointing up. Above the unit disk means I don't care about this infinite surface. I will actually stop when I hit a radius of one away from the z-axis. And so now I have my vector field which is going to point overall up because, well, it is z times k. The more z is positive, the more your vector field goes up.

Of course, if z were negative then it would point down, but it will live above. Actually, a quick opinion poll. What do you think the flux should be? Should it be positive, zero, negative or we don't know? I see some I don't know, I see some negative and I see some positive. Of course, I didn't tell you which way I am orienting my paraboloid. So far both answers are correct. The only one that is probably not correct is zero because, no matter which way you choose to orient it you should get something. It is not looking like it will be zero. Let's say that I am going to do it with the normal pointing upwards.

Second chance. I see some people changing back and forth from one and two. Let's draw a picture. Which one is pointing upwards? Well, let's look at the bottom point. The normal vector pointing up, here we know what it means. It is this guy. If you continue to follow your normal vector, see, they are actually pointing up and into the paraboloid. And I claim that the answer should be positive because the vector field is crossing our paraboliod going upwards, going from the outside out and below to the inside and upside. So, in the direction that we are counting positively.

We will see how it turns out when we do the calculation. We have to compute the integral for flux. Double integral over a surface of F dot n dS is going to be -- What are we going to do? Well, F we said is <0, 0, z>. What is n dS. Well, let's use our brand new formula. It says negative f sub x, negative f sub y, one, dxdy. What does little f in here? It is x^2 y^2. When we are using this formula, we need to know what little x stands for.

It is whatever the formula is for z as a function of x and y. We take x^2 y^2 and we take the partial derivatives with minus signs. We get negative 2x, negative 2y and one, dxdy. Well, of course here it didn't really matter because we are going to dot them with zero. Actually, even if we had made a mistake we somehow wouldn't have had to pay the price. But still. We will end up with double integral on S of z dxdy. Now, what do we do with that? Well, we have too many things. We have to get rid of z. Let's use z equals x^2 y^2 once more.

That becomes double integral of x^2 y^2 dxdy. And here, see, we are using the fact that we are only looking at things that are on the surface. It is not like in a triple integral. You could never do that because z, x and y are independent. Here they are related by the equation of a surface. If I sound like I am ranting, but I know from experience this is where one of the most sticky and tricky points is.

OK. How will we actually integrate that? Well, now that we have just x and y, we should figure out what is the range for x and y. Well, the range for x and y is going to be the shadow of our region. It is going to be this unit disk. I can just do that for now. And this is finally where I have left the world of surface integrals to go back to

a usual double integral. And now I have to set it up. Well, I can do it this way with dxdy, but it looks like there is a smarter thing to do. I am going to use polar coordinates. In fact, I am going to say this is double integral of r^2 times r dr d theta. I am on the unit disk so r goes zero to one, theta goes zero to 2pi. And, if you do the calculation, you will find that this is going to be pi over two.

Any questions about the example. Yes? How did I get this negative 2x and negative 2y? I want to use my formula for n dS. My surface is given by the graph of a function. It is the graph of a function x^2 y^2. I will use this formula that is up here. I will take the function x^2 y^2 and I will take its partial derivatives. If I take the partial of f, so x^2 y^2 with respect to x, I get 2x, so I put negative 2x. And then the same thing, negative 2y, one, dxdy. Yes?

Which k hat? Oh, you mean the vector field. It is a different part of the story. Whenever you do a surface integral for flux you have two parts of the story. One is the vector field whose flux you are taking. The other one is the surface for which you will be taking flux. The vector field only comes as this f in the notation, and everything else, the bounds in the double integral and the n dS, all come from the surface that we are looking at. Basically, in all of this calculation, this is coming from f equals zk. Everything else comes from the information paraboloid z = x^2 y^2 above the unit disk.

In particular, if we wanted to now find the flux of any other vector field for the same paraboloid, well, all we would have to do is just replace this guy by whatever the new vector field is. We have learned how to set up flux integrals for this paraboloid. Not that you should remember this one by heart. I mean there are many paraboloids in life and other surfaces, too. It is better to remember the general method. Any other questions? No. OK. Let's see more ways of taking flux integrals. But, just to reassure you, at this point we have seen the most important ones. 90% of the problems that we will be looking at we can do with what we have seen so far in less time and this formula.

Let's look a little bit at a more general situation. Let's say that my surface is so complicated that I cannot actually express z as a function of x and y, but let's say that I know how to parametize it. I have a parametric equation for my surface. That means I can express x, y and z in terms of any two parameter variables that might be relevant for me. If you want, this one here is a special case where you can parameterize things in terms of x and y as your two variables. How would you do it in the fully general case? In a way, that will answer your question that, I think one of you, I forgot, asked yesterday how would I do it in general? Is there a formula like M dx plus N dy? Well, that is going to be the general formula. And you will see that it is a little bit too complicated, so the really useful ones are actually the special ones.

Let's say that we are given a parametric description -- -- of a surface S. That means we can describe S by formulas saying x is some function of two parameter variables. I am going to call them u and v. I hope you don't mind. You can call them t1 and t2. You can call them whatever you want. One of the basic properties of a surface is because I have only two independent directions to move on. I should be able to express x, y and z in terms of two variables. Now, let's say that I know how to do that. Or, maybe I should instead think of it in terms of a position vector if it helps you. That is just a vector with components <x, y, z> is given as a function of u and v.

It works like a parametric curve but with two parameters. Now, how would we actually set up a flux integral on such a surface. Well, because we are locating ourselves in terms of u and v, we will end up with an integral du dv. We need to figure out how to express n dS in terms of du and dv. N dS should be something du dv. How do we do that? Well, we can use the same method that we have actually used over here. Because, if you think for a second, here we used, of course, a rectangle in the x, y plane and we lifted it to a parallelogram and so on. But more generally you can think what happens if I change u by delta u keeping v constant or the other way around?

You will get some sort of mesh grid on your surface and you will look at a little parallelogram that is an elementary piece of that mesh and figure out what is its area and normal vector. Well, that will again be given by the cross-product of the two sides. Let's think a little bit about what happens when I move a little bit on my surface. I am taking this grid on my surface given by the u and v directions.

And, if I take a piece of that corresponding to small changes delta u and delta v, what is going to be going on here? Well, I have to deal with two vectors, one corresponding to changing u, the other one corresponding to changing v. If I change u, how does my point change? Well, it is given by the derivative of this with respect to u. This vector here I will call, so the sides are given by, let me say, partial r over partial u times delta u. If you prefer, maybe I should write it as partial x over partial u times delta u. Well, it is just too boring to write.

And so on. It means if I change u a little bit, keeping v constant, then how x changes is, given by partial x over partial u times delta u, same thing with y, same thing with z, and I am just using vector notation to do it this way. That is the analog of when I said delta r for line integrals along a curve, vector delta r is the velocity vector dr dt times delta t. Now, if I look at the other side --

Let me start again. I ran out of space. One side is partial r over partial u times delta u. And the other one would be partial r over partial v times delta v. Because that is how the position of your point changes if you just change u or v and not the other one. To find the surface element together with a normal vector, I would just take the cross-product between these guys. If you prefer, that is the cross-product of partial r over partial u with partial r over partial v, delta u delta v. And so n dS is this cross-product times du dv up to sign.

It depends on which choice I make for my normal vector, of course. That, of course, is a slightly confusing equation to think of. A good exercise, if you want to really understand what is going on, try this in two good examples to look at. One good example to look at is the previous one. What is it? It is when u and v are just x and y. The parametric equations are just x equals x, y equals y and z is f of x, y. You should end up with the same formula that we had over there. And you should see why because both of them are given by a cross-product. The other case you can look at just to convince yourselves even further. We don't need to do that because we have seen the formula before, but in the case of a sphere we have seen the formula for n and for dS separately.

We know what n dS are in terms of d phi, d theta. Well, you could parametize a sphere in terms of phi and theta. Namely, the formulas would be x equals a sine phi cosine theta, y equals a sign phi sine theta, z equals a cosine phi. The formulas for circle coordinates setting Ro equals a . That is a parametric equation for the sphere.

And then, if you try to use this formula here, you should end up with the same things we have already seen for n dS, just with a lot more pain to actually get there because cross-product is going to be a bit complicated. But we are seeing all of these formulas all fitting together. Somehow it is always the same question. We just have different angles of attack on this general problem. Questions?

No. OK. Let's look at yet another last way of finding n dS. And then I promise we will switch to something else because I can feel that you are getting a bit overwhelmed for all these formulas for n dS. What happens very often is we don't actually know how to parametize our surface. Maybe we don't know how to solve for z as a function of x and y, but our surface is given by some equation. And so what that means is actually maybe what we know is not really these kinds of formulas, but maybe we know a normal vector.

And I am going to call this one capital N because I don't even need it to be a unit vector. You will see. It can be a normal vector of any length you want to the surfaces. Why would we ever know a normal vector? Well, for example, if our surface is a plane, a slanted plane given by some equation, ax by cz = d. Well, you know the normal vector. It is <a, b, c>. Of course, you could solve for z and then go back to that case, which is why I said that one is very useful. But you can also just stay with a normal vector. Why else would you know a normal vector? Well, let's say that you know an equation that is of a form g of x, y, z equals zero. Well, then you know that the gradient of g is perpendicular to the level surface. Let me just give you two examples.

If you have a plane, ax by cz = d, then the normal vector would just be <a, b, c>. If you have a surface S given by an equation, g(x, y, z) = 0, then you can take a normal vector to be the gradient of g. We have seen that the gradient is perpendicular to the level surface. Now, of course, we don't necessarily have to follow what is going to come. Because, if we could solve for z, then we might be better off doing what we did over there. But let's say that we want to do it this. What can we do? Well, I am going to give you another way to think geometrically about n dS.

Let's start by thinking about the slanted plane. Let's say that my surface is just a slanted plane. My normal vector would be maybe somewhere here. And let's say that I am going to try -- I need to get some handle on how to set up my integrals, so maybe I am going to express things in terms of x and y. I have my coordinates, and I will try to use x and y. Then I would like to relate delta S or dS to the area in the x y plane. That means I want maybe to look at the projection of this guy onto a horizontal plane.

Let's squish it horizontally. Then you have here another area. The guy on the slanted plane, let's call that delta S. And let's call this guy down here delta A. And delta A would become ultimately maybe delta x, delta y or something like that. The question is how do we find the conversion rate between these two areas? I mean they are not the same. Visually, I hope it is clear to you that if my plane is actually horizontal then, of course, they are the same. But the more slanted it becomes the more delta A becomes smaller than delta S. If you buy land and it is on the side of a cliff, well, whether you look at it on a map or whether you look at it on the actual cliff, the area is going to be very different.

I am not sure if that is a wise thing to do if you want to build a house there, but I bet you can get really cheap land. Anyway, delta S versus delta A depends on how slanted things are. And let's try to make that more precise by looking at the angel that our plane makes with the horizontal direction. Let's call this angle alpha, the angle that our plane makes with the horizontal direction. See, it is all coming together. The first unit about cross-products, normal vectors and so on is actually useful now.

I claim that the surface element is related to the area in the plane by delta A equals delta S times the cosine of alpha. Why is that? Well, let's look at this small rectangle with one horizontal side and one slanted side. When you project this side does not change, but this side gets shortened by a factor of cosine alpha. Whatever this length was, this length here is that one times cosine alpha.

That is why the area gets shrunk by cosine alpha. In one direction nothing happens. In the other direction you squish by cosine alpha. What that means is that, well, we will have to deal with this. And, of course, the one we will care about actually is delta S expressed in terms of delta A. But what are we going to do with this cosine? It is not very convenient to have a cosine left in here. Remember, the angle between two planes is the same thing as the angle between the normal vectors. If you want to see this angle alpha elsewhere, what you can do is you can just take the vertical direction. Let's take k. Then here we have our angle alpha again.

In particular, cosine of alpha, I can get, well, we know how to find the angle between two vectors. If we have our normal vector N, we will do N dot k, and we will divide by length N, length k. Well, length k is one. That is one easy guy. That is how we find the angle. Now I am going to say, well, delta S is going to be one over cosine alpha delta A. And I can rewrite that as length of N divided by N dot k times delta A.

Now, let's multiply that by the unit normal vector. Because what we are about is not so much dS but actually n dS. N delta S will be, I am just going to multiply by N. Well, let's think for a second. What happens if I take a unit normal N and I multiply it by the length of my other normal big N? Well, I get big N again. This is a normal vector of the same length as N, well, up to sign. The only thing I don't know is whether this guy will be going in the same direction as big N or in the opposite direction. Say that, for example, my capital N has, I don't know, length three for example. Then the normal unit vector might be this guy, in which case indeed three times little n will be big n.

Or it might be this one in which case three times little n will be negative big N. But up to sign it is N. And then I will have N over N dot k delta A. And so the final formula, the one that we care about in case you don't really like my explanations of how we get there, is that N dS is plus or minus N over N dot k dx dy. That one is actually kind of useful so let's box it. Now, just in case you are wondering, of course, if you didn't want to project to x, y, you would have maybe preferred to project to say the plane of a blackboard, y, z, well, you can do the same thing. To express n dS in terms of dy dz you do the same argument.

Simply, the only thing that changes, instead of using the vertical vector k, you use the normal vector i. So you would be doing N over N dot i dy dz. The same thing. So just keep an open mind that this also works with other variables. Anyway, that is how you can basically project the vectors of this area element onto the x, y plane in

a way. Let's look at the special case just to see how this fits with stuff we have seen before.

Let's do a special example where our surface is given by the equation z minus f of x, y equals zero. That is a strange way to write the equation. z equals f of x, y. That we saw before. But now it looks like some function of x, y, z equals zero. Let's try to use this new method. Let's call this guy g(x, y, z). Well, now let's look at the normal vector. The normal vector would be the gradient of g, you see. What is the gradient of this function? The gradient of g --

Well, partial g, partial x, that is just negative partial f, partial x. The y component, partial g, partial y is going to be negative f sub y, and g sub z is just one. Now, if you take N over N dot k dx dy, well, it looks like it is going to be negative f sub x, negative f sub y, one divided by -- Well, what is N dot k? If you dot that with k you will get just one, so I am not going to write it, dx dy. See, that is again our favorite formula. This one is actually more general because you don't need to solve for z, but if you cannot solve for z then it is the same as before.

I think that is enough formulas for n dS. After spending a lot of time telling you how to compute surface integrals, now I am going to try to tell you how to avoid computing them. And that is called the divergence theorem. And we will see the proof and everything and applications on Tuesday, but I want to at least the theorem and see how it works in one example. It is also known as the Gauss-Green theorem or just the Gauss theorem, depending in who you talk to.

The Green here is the same Green as in Green's theorem, because somehow that is a space version of Green's theorem. What does it say? It is 3D analog of Green for flux. What it says is if S is a closed surface -- Remember, it is the same as with Green's theorem, we need to have something that is completely enclosed. You have a surface and there is somehow no gaps in it. There is no boundary to it. It is really completely enclosing a region in space that I will call D.

And I need to choose my orientation. The orientation that will work for this theorem is choosing the normal vector to point outwards. N needs to be outwards. That is one part of the puzzle. The other part is a vector field. I need to have a vector field that is defined and differentiable -- -- everywhere in D, so same instructions as usual. Then I don't have actually to compute the flux integral. Double integral of f dot n dS of a closed surface S.

I am going to put a circle just to remind you it is has got to be a closed surface. It is just a notation to remind us closed surface. I can replace that by the triple integral of a region inside of divergence of F dV. Now, I need to tell you what the divergence of a 3D vector field is. Well, you will see that it is not much harder than in the 2D case. What you do is just -- Say that your vector field has components P, Q and R. Then you will take P sub x Q sub y R sub z. That is the definition. It is pretty easy to remember. You take the x component partial respect to S plus partial respect to y over y component plus partial respect to z of the z component.

For example, last time we saw that the flux of the vector field zk through a sphere of radius a was four-thirds pi a cubed by computing the surface integral. Well, if we do it more efficiently now by Green's theorem, we are going to use Green's theorem for this sphere because we are doing the whole sphere. It is fine. It is a closed surface. We couldn't do it for, say, the hemisphere or something like that. Well, for a

hemisphere we would need to add maybe the flat face of a bottom or something like that.

Green's theorem says that our flux integral can actually be replaced by the triple integral over the solid bowl of radius a of the divergence of zk dV. But now what is the divergence of this field? Well, you have zero, zero, z so you get zero plus zero plus one. It looks like it will be one. If you do the triple integral of 1dV, you will get just the volume -- -- of the region inside, which is four-thirds by a cubed. And so it was no accident. In fact, before that we looked at also xi yj zk and we found three times the volume.

That is because the divergence of that field was actually three. Very quickly, let me just say what this means physically. Physically, see, this guy on the left is the total amount of stuff that goes out of the region per unit time. I want to figure out how much stuff comes out of there. What does the divergence mean? The divergence means it measures how much the flow is expanding things. It measures how much, I said that probably when we were trying to understand 2D divergence.

It measures the amount of sources or sinks that you have inside your fluid. Now it becomes commonsense. If you take a region of space, the total amount of water that flows out of it is the total amount of sources that you have in there minus the sinks. I mean, in spite of this commonsense explanation, we are going to see how to prove this. And we will see how it works and what it says.