 We are going to continue to look at stuff in space. We have been working with triple integrals and seeing how to set them up in all sorts of coordinate systems. And the next topic we will be looking at are vector fields in space. And so, in particular, we will be learning about flux and work. So, just for a change, we will be starting with flux first. And we will do work, actually, after Thanksgiving. Just to remind you, a vector field in space is just the same thing as in the plane. At every point you have a vector, and the components of this vector depend on the coordinates x, y and z. Let's say the components might be P, Q, R, or your favorite three letters, where each of these things is a function of coordinates x, y, z. You have seen that in the plane it is already pretty hard to draw a vector field. Usually, in space, we won't really try too hard. But it is still useful to try to have a general idea for what the vectors in there are doing, whether they are all going in the same direction, whether they may be all vertical or horizontal, pointing away from the origin, towards it, things like that. But, generally-speaking, we won't really bother with trying to draw a picture because that is going to be quite hard. Just to give you examples, well, the same kinds of examples as the plane, you can think of force fields. For example, the gravitational attraction -- -- of a solid mass, let's call this mass big M, at the origin on a mass M at point x, y, z. That would be given by a vector field that points toward the origin and whose magnitude is inversely proportional to the square of a distance from the origin. Such a field would be directed towards the origin and its magnitude would be of the order of a constant over pho squared where pho is the distance from the origin. The picture, if I really wanted to draw a picture, would be everywhere it is a field that points towards the origin. And if I am further away then it gets smaller. And, of course, I am not going to try to draw all these vectors in there. If I wanted to give a formula for that -- A formula for that might be something of a form minus c times x, y, z over pho cubed. Let's see. Well, the direction of this vector, this vector is proportional to negative x, y, z. is the vector that goes from the origin to your point. The negative goes towards the origin. Then the magnitude of this guy, well, the magnitude of x, y, z is just the distance from the origin rho. So the magnitude of this thing is one over rho cubed times some constant factor. That would be an example of a vector field that comes up in physics. Well, other examples would be electric fields. Actually, if you look at the electric field generated by a charged particle at the origin, it is given by exactly the same kind of formula, and there are magnetic fields and so on. Another example comes from velocity fields. If you have a fluid flow, for example, if you want to study wind patterns in the atmosphere. Well, wind, most of the time, is kind of horizontal, but maybe it depends on the altitude. At high altitude you have jet streams, and the wind velocity is not the same at all altitudes. And, just to give you more examples, in math we have seen that the gradient of a function of three variables gives you a vector field. If you have a function u of x, y, z then its gradient field has just components, u sub x, u sub y and u sub z. And, of course, the cases are not mutually exclusive. For example, the electric field or gravitational field is given by the gradient of the gravitational or electric potential. So, these are not like different cases. There is overlap. Anyway, hopefully, you are kind of convinced that you should learn about vector fields. What are we going to do with them? Well, let's start with flux. Remember not so long ago we looked at flux of a two-dimensional field of a curve. We had a curve in the plane and we had a vector field. And we looked at the component of a vector field in the direction that was normal to the curve. We formed the flux integral that was a line integral F dot n ds. And that measured how much the vector field was going across the curve. If you were thinking of a velocity field, that would measure how much fluid is passing through the curve in unit time. Now let's say that we were in space. Well, we cannot really think of flux as a line integral. Because, if you have a curve in space and say that you have wind or something like that, you cannot really ask how much air is flowing through the curve. See, to have a flow through something you need a surface. If you have a net maybe then you can ask how much stuff is passing through that surface. There is going to be a big difference here. In the three-dimensional space, flux will be measured through a surface. And so it will be a surface integral, not a line integral anymore. That means we will be integrating, we will be summing over all the pieces of a surface in space. Because a surface is a two-dimensional object, that will end up being a double integral. But, of course, we will have to set it up properly because the surface that is in space, and we will probably have x, y and z to deal with at the same time, and we will have to somehow get rid of one variable so that we can set up and evaluate a double integral. So conceptually it is very similar to line integrals. In the line integral in the plane, you had two variables that you reduced to one by figuring out what the curve was. Here you have three variables that you will reduce to two by figuring out what the surface is. Let me give you a definition of flux in 3D. Let's say that we have a vector field and s, a surface in space. Let me draw some kind of a picture. I have my surface and I have my vector field F. Well, at every point it changes with a point. Well, I want to figure out how much my vector field is going across that surface. That means I want to figure out the normal component of my vector field, so I will use, as in the plane case, the unit normal vector to s. I take my point on the surface and build a unit vector that is standing on it perpendicularly. Now, we have to decide which way it is standing. We can build our normal vector to go this way or to go the other way around. There are two choices. Basically, whenever you want to set up a flux integral you have to choose one side of the surface. And you will count positively what flows toward that side and negatively what flows towards the other side. There are two choices for n. We need to choose a side of the surface. In the case of curves, we made that choice by deciding that because we were going along some direction on the curve we could choose one side by saying let's rotate clockwise from the tangent vector. And, in a way, what we were doing was really it was a recipe to choose for us one of the two sides. Here we don't have a notion of orienting the surface other than by precisely choosing one of the two possible normal vectors. So, in fact, this is called choosing an orientation of a surface. When you are saying you are orienting the surface that really means you are deciding which side is which. Let's call that orientation. Now, there is no set convention that will work forever. But the usually traditional settings would be to take your normal vector pointing maybe out of the solid region because then you will be looking at flux that is coming out of that region of space. Or, if you have a surface that is not like closed or anything but maybe you will want the flux going up through the region. Or, there are various conventions. Concretely, on problem sets it will either say which choice you have to make or you get to choose which one you want to make. And, of course, if you choose the other one then the sign becomes the opposite. Now, once we have made a choice then we can define the flux integral. It will just be the double integral over a surface of F dot n dS. Now I am using a big dS. That stands for the surface area element on this surface. I am using dS rather than dA because I still want to think of dA as maybe the area in one of the coordinate planes like the one we had in double integrals. You will see later where this comes in. But conceptually it is very similar. Concretely what this means is I cut my surface into little pieces. Each of them has area delta S. And, for each piece, I take my vector field, I take my normal vector, I dot them and I multiply by this surface area and sum all these things together. That is what a double integral means. In particular, an easy case where you know you can get away without computing anything is, of course, if your vector field is tangent to the surface because then you know that there is no flux. Flux is going to be zero because nothing passes through the surface. Otherwise, we have to figure out how to compute these things. That is what we are going to learn now. Well, maybe I should box this formula. I have noticed that some of you seem to like it when I box the important formulas. (APPLAUSE) By the way, a piece of notation before I move on, sometimes you will also see the notation vector dS. What is vector dS? Vector dS is this guy n dS put together. Vector dS is a vector which points perpendicular to the surface and whose length corresponds to the surface element. And the reason for having this shortcut notation, well, it is not only laziness like saving one n, but it is because this guy is very often easier to compute than it is to set up n and dS separately. Actually, if you remember in the plane, we have seen that vector n little ds can be written directly as dy, - dx. That was easier than finding n and ds separately. And here the same is going to be true in many cases. Well, any questions before we do examples? No. OK. Let's do examples. The first example for today is we are going to look at the flux of vector field xi yj xk through the sphere of radius a -- -- centered at the origin. What does the picture look like? We have a sphere of radius a. I have my vector field. Well, , see, that is a vector field that is equal to the vector from the origin to the point where I am, so it is pointing radially away from the origin. My vector field is really sticking out everywhere away from the origin. Now I have to find the normal vector to the sphere if I want to set up double integral over the sphere of F dot vector ds, or if you want F dot n dS. What does the normal vector to the sphere look like? Well, it depends, of course, whether I choose it pointing out or in. Let's say I am choosing it pointing out then it will be sticking straight out of a sphere as well. Hopefully, you can see that if I take a normal vector to the sphere it is actually pointing radially out away from the origin. In fact, our vector field and our normal vector are parallel to each other. Let's think a bit more about what a normal vector looks like. I said it is sticking straight out. It is proportional to this vector field. Maybe I should start by writing because that is the vector that goes from the origin to my point so it points radially away from the origin. Now there is a small problem with that. It is not a unit vector. So what is its length? Well, its length is square root of x^2 y^2 z^2. But, if I am on the sphere, then that length is just equal to a because distance from the origin is a. In fact, I get my normal vector by scaling this guy down by a factor of a. And let me write it down just in case you are still unsure. This is unit because square root of x^2 y^2 z^2 is equal to a on the sphere. OK. Any questions about this? No. It looks OK? I see a lot of blank faces. That physics test must have been hard. Yes? I could have put a rho but I want to emphasize the fact that here it is going to be a constant. I mean rho has this connotation of being a variable that I will need to then maybe integrate over or do something with. Yes, it would be correct to put rho but I then later will want to replace it by its actual value which is a number. And the number is a. It is not going to actually change from point to point. For example, if this was the unit sphere then I would just put x, y, z. I wouldn't divide by anything. Now let's figure out F dot n. Let's do things one at a time. Well, F and n are parallel to each other. F dot n, the normal component of F, is actually equal to the length of F. Well, times the length of n if you want, but that is going to be a one since F and n are parallel to each other. And what is the magnitude of F if I am on the sphere? Well, the magnitude of F in general is square root of x^2 y^2 z^2 on the sphere that is going be a. The other way to do it, if you don't want to think geometrically like that, is to just to do the dot product x, y, z doted with x over a, y over a, z over a. You will be x^2 y^2 z^2 divided by a. That will simplify to a because we are on the sphere. See, we are already using here the relation between x, y and z. We are not letting x, y and z be completely arbitrary. But the slogan is everything happens on the surface where we are doing the integral. We are not looking at anything inside or outside. We are just on the surface. Now what do I do with that? Well, I have turned my integral into the double integral of a dS. And a is just a constant, so I am very lucky here. I can just say this will be a times the double integral of dS. And, of course, some day I will have to learn how to tackle that beast, but for now I don't actually need to because the double integral of dS just means I am summing the area of each little piece of the sphere. I am just going to get the total area of the sphere which I know to be 4pi a2. This guy here is going to be the area of S. I know that to be 4pi a^2. So I will get 4pi a^3. That one was relatively painless. That was too easy. Let's do a second example with the same sphere. But now my vector field is going to be just z times k. Well, let me give it a different name. Let me call it H instead of f or something like that just so that it is not called F anymore. Well, the initial part of the setup is still the same. The normal vector is still the same. What changes is, of course, my vector field is no longer sticking straight out so I cannot use this easy geometric argument. It looks like I will have to compute F dot n and then figure out how to integrate that with dS. Let's do that. We still have that n is /a. That tells us that H dot n will be dot / a. It looks like I will be left with z^2 over a. H dot n is z^2 over a. The double integral for flux now becomes double integral on the sphere of z^2 over a dS. Well, we can take out one over a, that is fine, but it looks like we will have to integrate z^2 on the surface of the sphere. How do we do that? Well, we have to figure out what is dS in terms of our favorite set of two variables that we will use to integrate. Now, what is the best way to figure out where you are on the sphere? Well, you could try to use maybe theta and z. If you know how high you are and where you are around, in principle you know where you are on the sphere. But since spherical coordinates we have actually learned about something much more interesting, namely spherical coordinates. It looks like longitude / latitude is the way to go when trying to figure out where you are on a sphere. We are going to use phi and theta. And, of course, we have to figure out how to express dS in terms of d phi and d theta. Well, if you were paying really, really close attention last time, you will notice that we have actually already done that. Last time we saw that if I have a sphere of radius a and I take a little piece of it that corresponds to small changes in phi and theta then we said that -- Well, we argued that this side here, the one that is going east-west was a piece of the circle that has a radius a sin phi because that is r, so that side is a sin phi delta theta. And the side that goes north-south is a piece of the circle of radius a corresponding to angle delta phi, so it is a delta phi. And so, just to get to the answer, we got dS equals a^2 sin phi d phi d theta. When we set up a surface integral on the surface of a sphere, most likely we will be using phi and theta as our two variables of integration and dS will become this. Now, it is OK to think of them as spherical coordinates, but I would like to encourage you not to think of them as spherical coordinates. Spherical coordinates are a way of describing points in space in terms of three variables. Here it is more like we are parameterizing the sphere. We are finding a parametric equation for the sphere using two variables phi and theta which happen to be part of the spherical coordinate system. But, see, there is no rho involved in here. I am not using any rho ever, and I am not going to in this calculation. I have two variable phi and theta. That is it. It is basically in the same way as when you parameterize a line integral in the circle, we use theta as the parameter variable and never think about r. That being said, well, we are going to use phi and theta. We know what dS is. We still need to figure out what z is. There we want to think a tiny bit about spherical coordinates again. And we will know that z is just a cos phi. In case you don't quite see it, let me draw a diagram. Phi is the angle down from the positive z axes, this distance is a, so this distance here is a cos phi. Now I have everything I need to compute my double integral. z^2 over a dS will become a double integral. z^2 becomes a^2 cos^2 phi over a times, ds becomes, a^2 sin phi d phi d theta. Now I need to set up bounds. Well, what are the bounds? Phi goes all the way from zero to pi because we go all the way from the north pole to the south pole, and theta goes from zero to 2pi. And, of course, I can get rid of some a's in there and take them out. Let's look at what number we get. First of all, we can take out all those a's and get a^3. Second, in the inner integral, we are integrating cos^2 phi sin phi d phi. I claim that integrates to cos3 up to some factor, and that factor should be negative one-third. If you look at cos3 phi and you take its derivative, you will get that guy with a negative three in front between zero and pi. And, while integrating over theta, we will just multiply things by 2pi. Let me add the 2pi in front. Now, if I evaluate this guy between zero and pi, well, at pi cos^3 is negative one, at zero it is one, I will get two-thirds out of this. I end up with four-thirds pi a^3. Sorry I didn't write very much because I am trying to save blackboard space. Yes? That is a very natural question. That looks a lot like somebody we know, like the volume of a sphere. And ultimately it will be. Wait until next class when we talk about the divergence theorem. I mean the question was is this related to the volume of a sphere, and ultimately it is, but for now it is just some coincidence. Yes? The question is there is a way to do it M dx plus N dy plus stuff like that? The answer is unfortunately no because it is not a line integral. It is a surface integral, so we need to have to variables in there. In a way you would end up with things like some dx dy maybe and so on. I mean it is not practical to do it directly that way because you would have then to compute Jacobians to switch from dx dy to something else. We are going to see various ways of computing it. Unfortunately, it is not quite as simple as with line integrals. But it is not much harder. It is the same spirit. We just use two variables and set up everything in terms of these two variables. Any other questions? No. OK. By the way, just some food for thought. Never mind. Conclusion of looking at these two examples is that sometimes we can use geometric. The first example, we didn't actually have to compute an integral. But most of the time we need to learn how to set up double integrals. Use geometry or you need to set up for double integral of a surface. And so we are going to learn how to do that in general. As I said, we need to have two parameters on the surface and express everything in terms of these. Let's look at various examples. We are going to see various situations where we can do things. Well, let's start with an easy one. Let's call that number zero. Say that my surface S is a horizontal plane, say z equals a. When I say a horizontal plane, it doesn't have to be the entire horizontal plane. It could be a small piece of it. It could even be, to trick you, maybe an ellipse in there or a triangle in there or something like that. What you have to recognize is my surface is a piece of just a flat plane, so I shouldn't worry too much about what part of a plane it is. Well, it will become important when I set up bounds for integration. But, when it comes to looking for the normal vector, be rest assured that the normal vector to a horizontal plane is just vertical. It is going to be either k or negative k depending on whether I have chosen to orient it pointing up or down. And which one I choose might depend on what I am going to try to do. The normal vector is just sticking straight up or straight down. Now, what about dS? Well, it is just going to be the area element in a horizontal plane. It just looks like it should be dx dy. I mean if I am moving on a horizontal plane, to know where I am, I should know x and y. So dS will be dx dy. If I play the game that way, I have my vector field F. I do F dot n. That just gives me the z component which might involve x, y and z. x and y I am very happy with. They will stay as my variables. Whenever I see z, well, I want to get rid of it. That is easy because z is just equal to a. I just plug that value and I am left with only x and y, and I am integrating that dx dy. It is actually ending up being just a usual double integral in x, y coordinates. And, of course, once it is set up anything is fair game. I might want to switch to polar coordinates or something like that. Or, I can set it up dx dy or dy dx. All the usual stuff applies. But, for the initial setup, we are just going to use these and express everything in terms of x and y. A small variation on that. Let's say that we take vertical planes that are parallel to maybe the blackboard plane, so parallel to the yz plane. That might be something like x equals some constant. Well, what would I do then? It could be pretty much the same. The normal vector for this guy would be sticking straight out towards me or away from me. Let's say I am having it come to the front. The normal vector would be plus/minus i. And the variables that I would be using, to find out my position on this guy, would be y and z. In terms of those, the surface element is just dy dz. Similarly for planes parallel to the xz plane. You can figure that one out. These are somehow the easiest ones, because those we already know how to compute without too much trouble. What if it is a more complicated plane? We will come back to that next time. Let's explore some other situations first. Number one on the list. Let's say that I gave you a sphere of radius a centered at the origin, or maybe just half of that sphere or some portion of it. Well, we have already seen how to do things. Namely, we will be saying the normal vector is x, y, z over a, plus or minus depending on whether we want it pointing in or out. And dS will be a^2 sin phi d phi d theta. In fact, we will express everything in terms of phi and theta. If I wanted to I could tell you what the formulas are for x, y, z in terms of phi and theta. You know them. But it is actually better to wait a little bit. It is better to do F dot n, because F is also going to have a bunch of x's, y's and z's. And if there is any kind of symmetry to the problem then you might end up with things like x^2 y^2 z^2 or things that have more symmetry that are easier to express in terms of phi and theta. The advice would be first do the dot product with F, and then see what you get and then turn it into phi and theta. That is one we have seen. Let's say that I have -- It is a close cousin. Let's say I have a cylinder of radius a centered on the z-axis. What does that look like? And, again, when I say cylinder, it could be a piece of cylinder. First of all, what does the normal vector to a cylinder look like? Well, it is sticking straight out, but sticking straight out in a slightly different way from what happens with a sphere. See, the sides of a cylinder are vertical. If you imagine that you have this big cylindrical type in front of you, hopefully you can see that a normal vector is going to always be horizontal. It is sticking straight out in the horizontal directions. It doesn't have any z component. I claim the normal vector for the cylinder, if you have a point here at (x, y, z), it would be pointing straight out away from the central axis. My normal vector, well, if I am taking it two points outwards, will be going straight away from the central axis. If I look at it from above, maybe it is easier if I look at it from above, look at x, y, then my cylinder looks like a circle and the normal vector just points straight out. It is the same situation as when we had a circle in the 2D case. The normal vector for that is just going to be x, y and 0 in the z component. Well, plus/minus, depending on whether you want it sticking in or out. We said in our cylinder normal vector is plus or minus x, y, zero over a. What about the surface element? Before we ask that, maybe we should first figure out what coordinates are we going to use to locate ourselves in a cylinder. Well, yes, we probably want to use part of a cylindrical coordinate, except for, well, we don't want r because r doesn't change, it is not a variable here. Indeed, you probably want to use z to tell how high you are and theta to tell you where you are around. dS should be in terms of dz d theta. Now, what is the constant? Well, let's look at a small piece of our cylinder corresponding to a small angle delta theta and a small height delta z. Well, the height, as I said, is going to be delta z. What about the width? It is going to be a piece of a circle of radius a corresponding to the angle delta theta, so this side will be a delta theta. Delta S is a delta theta delta z. DS is just a dz d theta or d theta dz. It doesn't matter which way you do it. And so when we set up the flux integral, we will take first the dot product of f with this normal vector. Then we will stick in this dS. And then, of course, we will get rid of any x and y that are left by expressing them in terms of theta. Maybe x becomes a cos theta, y becomes a sin theta. These various formulas, you should try to remember them because they are really useful, for the sphere, for the cylinder. And, hopefully, those for the planes you kind of know already intuitively. What about marginals or faces? Not everything in life is made out of cylinders and spheres. I mean it is a good try. Let's look at a marginal kind of surface. Let's say I give you a graph of a function z equals f of x, y. This guy has nothing to do with the integrand. It is not what we are integrating. We are just integrating a vector field that has nothing to do with that. This is how I want to describe the surface on which I will be integrating. My surface is given by z as a function of x, y. Well, I would need to tell you what n is and what dS is. That is going to be slightly annoying. I mean, I don't want to tell them separately because you see they are pretty hard. Instead, I am going to tell you that in this case, well, let's see. What variables do we want? I am going to tell you a formula for n dS. What variables do we want to express this in terms of? Well, most likely x and y because we know how to express z in terms of x and y. This is an invitation to get rid of any z that might be left and set everything up in terms of dx dy. The formula that we are going to see, I think we are going to see the details of why it works tomorrow, is that you can take negative partial f partial x, negative partial f partial y, one, dx dy. Plus/minus depending on which way you want it to go. If you really want to know what dS is, well, dS is the magnitude of this vector times dx dy. There will be a square root and some squares and some stuff. What is the normal vector? Well, you take this vector and you scale it down to unit length. Just to emphasize it, this guy here is not n and this guy here is not dS. Each of them is more complicated than that, but the combination somehow simplifies nicely. And that is good news for us. Now, concretely, one way you can think about it is this tells you how to reduce things to an integral of x and y. And, of course, you will have to figure out what are the bounds on x and y. That means you will need to know what does the shadow of your surface look like in the x, y plane. To set up bounds on whatever you will get dx dy, well, of course you can switch to dy dx or anything you would like, but you will need to look at the shadow of S in the x y plane. But only do that after you gotten rid of all the z. When you no longer have z then you can figure out what the bounds are for x and y. Any questions about that? Yes? For the cylinder. OK. Let me re-explain quickly how I got a normal vector for the cylinder. If you know what a cylinder looks like, you probably can see that the normal vector sticks straight out of it horizontally. That means the z component of n is going to be zero. And then the x, y components you get by looking at it from above. One last thing I want to say. What about the geometric interpretation and how to prove it? Well, if your vector field F is a velocity field then the flux is the amount of matter that crosses the surface that passes through S per unit time. And the way that you would prove it would be similar to the picture that I drew when we did it in the plane. Namely, you would consider a small element of a surface delta S. And you would try to figure out what is the stuff that flows through it in a second. Well, it is the stuff that lives in a small box whose base is that piece of surface and whose other side is given by the vector field. And then the volume of that is given by base times height, and the height is F dot n. It is the same argument as what we saw in the plane. OK. Next time we will see more formulas. We will first see how to prove this, more ways to do it, more examples. And then we will get to the divergence theorem.