

[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: OK. Great. So at this point, we basically have all of the most important-- well, we now have the full understanding of how to work with tensors. I haven't really done too much physics with them at this point, but we've very carefully-- one might even argue excessively carefully-- laid out this mathematical structure that is going to be-- we're going to use it to contain the geometric objects that will describe the physics that we're going to study over the course of the entire semester.

So there was a lot of twiddly detail in the preceding lecture, but I would say the two most important things I want you to take out of this is the idea that we now have tensors as a general class of geometric objects which map one-forms-- or dual vectors, if you prefer-- a combination of one-forms and vectors to the Lorentz invariant real numbers.

And further, the distinction between one-forms, and vectors, and what is being mapped to what is not so important, because I can always use the metric to either raise or lower indices on tensors or vectors or one forms in order to convert from one to another. So if I raise an index-- if I have a tensor that's got m indices in the upstairs position and n in the downstairs, I raise an index, and I've got $m + 1$ upstairs, $n - 1$ downstairs. And likewise, I can lower index-- put one of those ones in the upstairs, down to the downstairs.

We wanted to go through all that stuff with great care, because we're going to need the foundations of that to make a lot of what we talk about later rigorous. It is really overkill for where we're starting, but it's worthwhile to have that scaffolding in place. I'm sick of it, though, so today I'm going to try to do some things that are more physics, because, well, not all of you are necessarily physics students, but I'm a physics professor and I'm a little tired of doing math.

So let's think about-- when we're studying physics, generally we are interested in the

behavior of bodies and fields. Let's move this to the relativity framework. We're looking at the behavior of bodies and fields in spacetime. So far, if I think about the quantities that we have introduced-- and by the way, I apologize for the blackboards. A previous lecturer today used chalk that cannot be erased very well, so they're sort of grayish boards today, unfortunately. But we'll manage.

Anyhow, so far, the quantities that we have introduced that have solid physics content are good for describing particles-- individual particles-- maybe a handful of them at once. So the two that I want to focus on in particular are the 4-velocity-- so if I have a body with a 4-velocity u , a particular Lorentz observer will describe that as having a timelike component γ -- special relativistic γ -- and a spacelike component γv .

This is defined such that $u \cdot u$ equals minus 1. It is great for describing-- it is the tool we use to describe the trajectory of a material body that follows a timelike trajectory through spacetime. Notice that its norm is minus 1. The fact that it's 1 means that it's normalized, and the minus tells us that it is timelike.

By the way, because this is timelike, the 4-velocity is not going to be a useful tool for us when we want to describe the behavior of light. If we want to talk about the motion of a photon in our spacetime, we can't use a 4-velocity for it. The trajectory of a photon is null. Whatever quantity that is going to describe its trajectory-- I take its dot product with itself, I have to get 0.

And you can intuitively get a sense what's going on with that. If I take the way a particular observer interprets these things and says, what would this turn into if v equals c , well, you have infinity and infinity, because your gammas are diverging there. So it's a singular limit. It doesn't behave well. We will overcome this difficulty.

The other quantity that we have used-- let me see if I can clean this a little bit better. All right. The other quantity that we have defined which is good for particles is the 4-momentum, which is just that 4-velocity multiplied by a rest mass.

So this is defined-- you can clearly, by trimming together a bunch of definitions-- you have $p \cdot p$ equals minus m squared. We're actually going to interpret this, though. So if a particular Lorentz observer makes a measurement of this, they will

call the timelike component E and the spatial component p .

And so you put all these things together, and this tells me that E^2 minus magnitude of p^2 equals m^2 . Notice this is perfectly well-behaved for a massless particle. This is actually going to be the trick by which, when I want to describe the trajectory of photons or of light in a spacetime-- I can't use 4-velocity. I can use 4-momentum.

It works perfectly well. It's 0 times a 4-velocity, but if I use this intuition that this thing is diverging when I take the speed of light limit, I sort of-- I guess I cheat, but I'm basically getting 0 times infinity. And it behaves.

We're going to do it a little bit more rigorously than that, but if you want a little bit of intuition, that's essentially why it works. So the last thing which I will say here is that this is good for $m \neq 0$. OK. And when you do have $m = 0$, it's often convenient to write this as $\hbar \omega$.

So use the usual formula for the energy of a photon-- let me move this so that you can see-- and then I just put a unit vector into the spatial direction that defines the direction in which this is moving.

OK. So this is all fine as far as it goes, but it's kind of restrictive. We want to a little bit more with our physics than just worry about the kinetics of individual particles. There's a lot we can do with that, but one of the things that we like to do when we're studying gravity is to have a description of continual matter, like things that gravitate. You want to be able to build a description of something like a star, or an exotic star-- a neutron star-- something like that.

And so just having the behavior of individual particles is not good enough. I want to be able to describe things like fluids. I want to be able to have a continuum that describes things. So what I'm going to start doing today is begin to introduce the mathematical tools that we will use.

Pardon me. I seem to have twisted this around myself in a crazy way. We're going to describe the mathematical tools that are useful for dealing with continuous matter, rather than just particles.

So let's call this more interesting matter. So the simplest continuum form of matter

which we're going to talk about is-- we're going to call it dust. OK. When you're here at the chalkboard, you can't get away with it, so you have an idea of what it means.

Physically, what I mean by dust is that it is a collection of particles that do not interact with each other. So they pass through each other. There's no pressure that is generated.

So they have energy density associated with them, but no interaction. So it's the most boring kind of matter that you like. You can think of it as essentially just particles, but it's a ton of them, and they're smearing out into a continuum.

In particular, one thing I'm going to say is, so I can imagine, if I take my two erasers here, I have a nice dust field in front of me now, which I helpfully just inhaled. And you can think of this thing as a field of little dust elements that are all moving around.

If I go in and I track an individual element of that dust-- I go in and I make a little cube that's a nanometer on each side-- every little element has its own rest frame. In that cloud-- I'm not going to do that again-- but as I made that cloud, different elements within the cloud have different rest frames, but for each individual element, I can define a rest frame.

So we're going to use the properties of the dust in that rest frame as a way of beginning to normalize and get a grasp on how we're going to describe this. But I just want to emphasize again-- in a given cloud, different elements may have different rest frames.

OK. But let's suppose that I clap my erasers, I make my little cloud, and I go and I zoom in on one particular nanometer by nanometer by nanometer chunk of that thing, and I go into the rest frame of that particular little bit of dust. So how am I going to characterize that?

So presumably, each particle has its own mass associated with it. Each little bit of dust that's in that thing has its own mass. We're going to worry about how that comes into the picture a little bit later. Let's just begin by saying, one of the things it's going to be interested in knowing is, how much dust is there in that little cubic nanometer of dust in front of me there?

So the first thing we're going to focus on is just counting how many bits of dust are in this element. In particular, I'm going to want to know, how many elements per unit volume? In other words, the number density.

So I claim that's a good thing for us to get started on. And it looks like they haven't quite poisoned this board as badly. Good. So let's call n_0 the number density in the rest frame of that element.

OK. So this will just be some number per unit volume. OK. So it's a quantity that has dimensions 1 over volume. Now, in the rest frame, like I said, we're going to talk about how I treat the mass associated with each element, and how much energy is in that thing, a little bit later. There's a few other tools I want to introduce first before I get to there.

But in the rest frame, if I'm not worrying about the mass and the energy, this is it. This is really the only thing I can say about this. So what do I'm going to do now is say, OK, well, if I have a general cloud here-- as I described, when we made this little turbulent cloud of dust in front of us, it was all swirling around and doing its own thing.

You're generally not in the rest frame. OK. So you also want to know how to characterize the dust out of the rest frame. So let's move out of the rest frame.

When I do this, two things happen. So first of all, let's continue to focus on-- so we've got a particular nanometer cubed of dust that we are looking at here. So we're very attached to that particular thing. So I now boost into a frame where I'm going $3/4$ of the speed of light or something like that, but I'm going to pay attention to that-- I'm very attached to the dust in that little cube, so I'm going to keep myself focused on that.

Boosting into frame can't change the total number, so the amount of dust has to be the same. But of course, there's going to be a Lorentz contraction, and so the volume will get smaller. So the number in a particular volume stays the same while the volume Lorentz contracts.

And so let's call n the number density in this new frame. This is just going to be the

Lorentz gamma times n_0 . So again, I'm assuming you're all perfectly fluent with special relativity. So all I did there was say it's going to contract along that direction. The lengths are smaller by a factor of gamma, therefore the volume is larger by a factor of gamma.

But there's a second thing that happens as well. When I'm in this new frame, there is now a flow associated with the dust. So this will now be flowing through space. In particular, what you can do is define a flux that describes the number of particles crossing unit area in unit time.

So let's define n 3-vector to be this. And you use a little bit of intuition, and just think about the flow of these particles as they kind of go translating past you. This can only be the number density that you measure in that frame times the 3-velocity.

OK? Think about that a little bit. It's very much like when you learn about current density in basic E&M. Notice it's n and not n_0 here. So when you look at these two things for a moment or two-- let me just write down one more step. I can write this as $\gamma n_0 v$.

These are screaming to be put together to make a 4-vector. They're looking at you and saying-- this bit's kind of like, I'm a timelike piece of a 4-vector. Spatial piece, I love you. Let's get together. They clearly are just screaming to go together.

And when you do that, you say, hmm, let's see what happens here. Two great tastes that taste great together. OK. That's nice. Let's take advantage of the fact that both of these have a really simple form in terms of the density in the dust's own rest frame.

And you look at this and go, holy moly. That is nothing more than the number density times a 4-velocity. So that's pretty cool. If I have this stuff in front of me-- you know, again, what the heck. This time I'm going to actually do my clap again.

So I do this. Every little element-- I can think of it as having some trajectory through spacetime. I can attach a 4-velocity to that. And if I know what the density of dust is in that rest frame, I put it together-- now I've got a geometric object that describes this thing.

And one of the reasons why this is powerful is, geometric objects-- as I have

emphasized repeatedly-- have a geometric meaning that transcends a particular coordinate representation. Another way to think about this is that all observers-- I don't care whether you're at rest in this classroom, or you've had too much coffee and you're dashing through at $3/4$ the speed of light.

We all agree on N -- capital N vector-- but we choose different ways of splitting spacetime into space and time. And so I will have a different n from you, I will have a different n_v from you. We all have different ways of splitting this spacetime into space and time, and so we have different ways of splitting these two things up.

But this is a geometric object we all agree on. It's something that we can hang a lot of our physics on, and build frame-independent powerful geometric representations. So before going on to do a few things with this, let's just explore a couple of the properties of this.

All right. So whenever you've got a 4-vector, at a certain point, you just sort of say to yourself, well, I've got this thing. Let's take the dot product with itself. Why not? So in this case, $N \cdot N$ equals minus n_0 squared.

Or, turning this around, this tells you that if someone comes by and goes, psst, here's a number 4-vector, you can figure out density in the rest frame associated with each element by the following operations. So that's nice. OK. In other words, the number density, in a sense, tells you what the normalization of this vector is.

A few moments ago, I defined N as being kind of a flux of these things. It's a flux in the spatial directions. We want to think about flux in a more general kind of way. And so another thing which we're going to do-- and this is where 1-forms are going to play an important role for us-- is a nice systematic way to pick out a flux across a surface.

So here's the bit where, when I was talking about it in Tuesday's lecture, I was getting the blank stares. And sorry, I'm coming right back to this again, so we're going to harp on this concept a little bit more. So recall we talked about the fact that if I have a particular coordinate system-- I have a particular set of coordinates x^α that I use to represent spacetime.

If I make 1-forms that are the gradients of those coordinates, I can represent them

as essentially, level surfaces at unit ticks of the coordinate x^α . OK. And basically, the way to think of it-- I kept emphasizing that the basis 1-forms are dual to basis vectors.

And if you have this intuition that your basis vector is a little arrow pointing like so, then the thing that is dual to it is a surface that is everywhere but pointing along like so. That's kind of the way I want you to think about it. And those graphics that are scans from MTW I put on the web give better pictures of that than I am able to draw.

So putting these concepts together, it basically tells me that I can use these basis 1-forms as a way of defining in an abstract form. So if I want to know the flux of N in the-- let's write it this way-- in the x^α direction.

So remember, the basis 1-forms-- the α in that case is not like a coordinate index. It's labeling a particular member of a set. And so what you would do is, you would construct this by saying, OK, let's take the β component of a basis 1-form α .

Contract it like so. And that tells me about the flux-- how much of this stuff is flowing in the direction associated with α . So this is a tool that we're going to use from time to time.

One little thing which I want to emphasize-- actually, two things I want to emphasize. So if I'm working in an intelligent coordinate system, remember, this is basically just the identity matrix, right? And so this ends up being a very simple thing that I actually work out. And it says, if I want to know the timelike component, then this is 1 in the timelike direction, 0 everywhere else, and I just pick out $N_{\text{sub } t}$ -- in other words, the density itself in that frame.

That is the other thing which I want to emphasize. So when you do this, if I pick out the timelike piece of this, this just gives me the 0 component of this thing, which is the density that I measure in this frame. This is very, very simple in terms of the calculation, but I want you to stop and think about what it means physically.

So flux in the timelike direction-- what does it mean for a flux to be in the timelike direction, right? If I take my water-- actually, I'm kind of thirsty-- so I want a little bit of flux of water going down my throat. You have an intuition about what that means.

The water's actually flowing in a particular direction.

If it's flowing through time-- there it goes. That basically means it's just sitting there. It's not doing anything, but just moving in time, as we all are doing. One time I remember my wife saying about our daughter-- she's like, she's growing up so fast. I was like, eh, she's growing up at 1 second per second. And that's just the way things go. It's just sitting there and it's living its life at 1 second per second, as we all do.

That is density. So we described the flux of a thing in time as just being its density. That's another way-- when we do a lot of our calculations, and we talk about the flow of things in the timelike direction-- that tends to be just the simple density associated with stuff. And then if I did this-- if I pulled out the x direction of this thing, I would pull out the x component of velocity times that. And that's sort of the flux of something in the x direction like you are probably used to from other classes.

Now, a place where this turns out-- so doing it when I'm just using the basis 1-forms is kind of trivial, [INAUDIBLE] sucks. More generally, what you can do is define a surface as being the solution of some scalar function in spacetime.

So let's say I do something like $\psi(t, x, y) = \text{constant}$. OK. This is very abstract, so let me just make it a little bit more concrete. Suppose my function ψ were square root of $x^2 + y^2 + z^2$, and my constant was 5.

OK. Well then, my scalar field would be picking out a sphere of radius 5. OK. You can make a little bit more complicated than that, and people often do. You can define the unit 1-form that is associated with the normal to this thing.

Or rather, it's the 1-form associated with the surface. If you translate this to a vector-- you raise the index-- with the vector, it would be the vector normal to the surface.

You might need to normalize it. Let's imagine that we normalize this thing. And then we would just need to contract it along this, and that tells me about the flux through this particular surface. OK?

This is one of the things that-- this idea of 1-forms being level surfaces-- it tends to be useful for things like that. We're going to come back to a similar sort of picture in just a moment, because we'll have to start talking about integration.

OK. In prep for that discussion, now that I'm talking about things a little bit more complicated than just the kinematics of simple particles, we're going to want to have some laws. And we want to have geometric ways of describing those laws that essentially act as constraints on what those particles can do.

OK. So the form in which we're going to express them-- we're going to tend to put things into the form of conservation laws. So suppose, here's my little element that's got dust in it, and it's embedded in an environment with a bunch of dust around it.

And over some time interval, some dust flows in, some does flows out. The density can go up, the density can go down. The total number-- it may vary depending on how things are flying. The spatial flux out of the sides must come-- and I'm going to say it's the flux out-- so let's say that that comes at the expense of the density of dust already there.

OK. And so if you were to just to intuitively write down what kind of conservation law you would expect to see, you would write it as something like this, based on simple Euclidean intuition. What's kind of nice-- you look at that for a second.

You go, ooh, if I think about this as being the time complement of my 4-vector. This is the space component of my 4-vector. This has a very obvious form when I write it in a geometric framework. This whole thing can be rewritten as a conservation law that looks like this.

OK. So I'll remind you, $d_{\text{sub } \alpha} = d \text{ by } dx_{\alpha}$. I'm going to talk a little bit about some of the derivatives a little bit later, because there are a few subtle points that can get introduced. But for now, we just know that $d_{\text{downstairs } t} \text{ is } d \text{ by } dt$. $d_{\text{downstairs } x} \text{ is } d \text{ by } dx$. $d_{\text{downstairs } y}$, et cetera.

So this is really nice. One thing which I want to emphasize-- again, coming back to what I said over there just a moment ago. When I write down this conservation law, I'm assuming that someone has defined what time means and someone has defined what space means.

This is a form that's covariant, right? All observers agree that this goes together. When they actually make their own coordinate systems, they're going to have their

own time coordinate and their own x-coordinate, but they're all going to be different flavors-- just different ways of instantiating what this is.

Now, pardon me for just one moment. Depending on the pace of the course, we're about to switch over to a different set of lecture notes, and I want to make sure I smoothly go from one to the other.

OK. So in many of your physics classes, you have learned, when you get a conservation law, both a differential form like this-- the rate of change of N in a box is related to the amount of N flowing out of that box and its sides. And you also learn an integral form of the conservation law.

So without proof, let me just say that it should be intuitively clear that what I've written down over here is equivalent to-- OK, so it looks like this. You know what? Let me just fix up my notation a little bit. Let me call this lowercase n .

OK. A few symbols that I've introduced here-- so V^3 is some volume in three-dimensional space, and dV like this-- this is a symbol that means the boundary of that 3-volume. And that's basically-- it's a form of Gauss's law. That's what I've written down there.

OK? So you've all seen things like that. Again, let me emphasize that when I write down a formula like that, I can only do that having assumed a particular Lorentz frame. That t is the t of some observer. That volume is the volume of that particular observer who is using t .

You jump into a different Lorentz frame, their volumes will not be the same. Their times will not be the same. I'm going to make some coordinates up. So an integral form like this, as I've written it there, only works in one given Lorentz frame.

Nonetheless, we are going to find it useful, even though in some sense, when you do it in an integral form like this, you're saying things in the framework of some particular observer. Sometimes you want to know in the framework of some particular observer. It could be you, right? And you might care about these sorts of things.

And that's good, but the way I've written it here-- first of all, it's in language that-- it's in a mathematical formulation that's not easy to generalize as I take things up to

higher dimensions. And so what I would like to do is think about how to step up a formulation like this in such a way that things are put into as frame-independent a language as is possible, and that will generalize forward when we start looking at more complicated geometries than just geometry of special relativity.

So I want to spend the next roughly 10 or 15 minutes talking about volumes and volume integrals. And my goal here is to try to-- I'm going to start by just doing stuff that comes from the journal of duh. It's stuff you have seen over and over and over again, but I want to re-express it using mathematical formulation that maybe-- you have seen all the symbols, but perhaps not used in quite this way. And then it'll carry forward in a framework that generalizes in a very useful way for us.

So let's begin with just simple 3D space. So I'm going to begin in 3D, and I'm going to consider a parallelepiped-- parallelepiped. Ha-ha. Got it right-- whose sides are a set of vectors.

So there are three vectors, A, B, and C. OK. So here is a vector A. This one going into the board is vector B. And this one going up here is vector C. OK. So those are my three vectors.

And if I go draw the ghost legs associated with these things-- OK. That's a little bit better. So these three vectors define a particular volume. And you guys have probably all seen-- you know you have three vectors. You can define a volume associated with this.

A really easy way to get that volume given those three vectors is to take A and dot it into the cross product of B and C. This is a quantity which is cyclic, so if you prefer, you can write it as B dotted into C cross A, or C dotted into A cross B.

That can be expressed as a determinant, or-- you guys can look at my notes to see the determinant written out if you like, but it's not that interesting, so I'm not going to use it very much. An equivalent way of writing all that is to use the Levi-Civita symbol.

So that 3-volume is given by $\epsilon_{ijk} A_i B_j C_k$, where-- I'll remind you-- ϵ_{ijk} equals plus 1 if i equals 1, j equals 2, k equals 3, and even permutations.

Even permutations means I swap two pairs of indices. So 123, 231, 312-- those all give me plus 1. It gives me minus 1 for any odd permutations of those. So 132, 213, et cetera. Those will all give me minus 1. And it's 0 if any index is repeated.

So you probably have all seen things like this before. This is fairly basic vector geometry. We are going to regard the Levi-Civita symbol as the components of a 0, 3 tensor.

OK. Bear in mind for just a moment here I'm working only in-- sorry, just ran out of good chalk-- I'm only working in 3-space, so my tensor definition is slightly different. It's not going to be a set of things that maps to Lorentz invariants, but it's going to be invariant with respect to things like rotations and translations in three-dimensional space.

So I'm going to regard these as the components of a 0, 3 tensor that basically takes in vectors and spits out the volume associated with the element whose edges are bounded by those vectors.

OK. So I could say, in this abstract form I wrote down earlier, imagine a boldfaced epsilon which is my volume tensor. I put these slots into it, and voila. I get the volume out of it.

Now, with that in mind, remember some of the games that we played with tensors in the previous lecture. So when I was talking about spacetime tensors, if I filled up all of their slots, I got out a Lorentz invariant number. In this case, I'm in 3-space, so I fill up all its slots, I get an invariant number in this 3-space.

Suppose I only put in two vectors. So suppose I do something like, I plug in-- let's leave the first slot blank-- and I put in vectors B and C. OK, well, writing this out in component form, I know this is epsilon ijk B_j C_k.

That's just B cross C, right? That's the area spanned by the side that is B cross C. And you guys have learned in other classes that you have an extra index left over, so it's a vector that has a direction associated to it. So it's sort of an oriented surface.

We put the index in the downstairs position, so we're actually going to think about this as a 1-form corresponding to the side whose edges are B and C. So let's call this

side 1-form σ . This is a 1-form whose magnitude is the area of the side spanned by the vectors B and C .

And although I can still tell it hasn't quite gelled yet, it's useful to think of 1-forms as being associated with surfaces. Guess what? This is the side of a parallelepiped. That's a surface. So it actually holds together.

All right. So using all of this, if I wanted to write down how to do something like Gauss's theorem in this kind of geometric language-- and again, we emphasize this is very much in the spirit right now of mosquito with a sledgehammer.

We don't need all this sort of stuff, but we're about to step up to something a little bit more complicated. So what you would do is say, OK, well, I know Gauss's theorem. I pick a particular 3-volume. I say the divergence of some vector field integrated over that volume is given by integrating the flux of that vector over the surface of this thing.

So what you might want to do then at this point is say, oh, OK, well, what I'm going to do, then, is say that my volume-- oops, pardon me a second. First thing I'll do is define a differential triple. I'll define some x_1 that points along one direction I care about, and x_2 , and an x_3 .

And then I will say dV equals $\epsilon_{ijk} dx^i dx^j dx^k$. I can likewise define a 1-form associated with my area element, as I have done over here. I'm not going to actually write this out. The key thing which I want to say is you have all the pieces-- you put all these things together, and you can define this thing. It's now very easy for you to prove Gauss's theorem using this kind of ingredients.

What I want to move onto-- there's a few more details in my notes. It's not super difficult or interesting to go through this. What I want to now start doing is generalize all of these ideas to the way we're going to approach them in spacetime.

Basically, we're going to do exactly the same kind of operations that I just did in space-- three-dimensional space-- but I'm going to put an extra index on things, and I'm going to do all of my quantities in spacetime.

OK. So imagine a parallelepiped with sides A , B , C , and D . Four dimensions, so it's

going to point along four different-- these can be mutually orthogonal. I'm going to define the invariant 4-volume associated with these things like so.

Where now my four-index Levi-Civita is defined such that ϵ_{0123} equals plus 1. If I do an odd permutation of those-- I exchange one pair of them-- ϵ_{1023} -- this equals minus 1.

If I repeat any index, I get 0. And likewise, all even permutations of this give me plus 1. All odd permutations of this give me minus 1. Or likewise, just do even permutations of this one.

So that is how I'm going to generalize my Levi-Civita symbol. As they say on *The Simpsons*, it's a perfectly cromulent object. I'm going to need to talk about the area associated with the faces of each of these things. So what is the area of the face of a 4-volume? A 3-volume.

So you can define a 1-form that tells me about the-- you can either call it the 3-volume or the 4-area. Knock yourselves out as to how you want to call it. And the obvious generalization-- let's say I leave off edge A. I want to get something like this.

I do my similar exercise of defining a-- in this case, it'll be a differential quartet associated with directions 0, 1, 2, and 3. And so by going through a procedure very similar to this, you get a generalization of Gauss's theorem. It says that if I integrate the spacetime divergence of some 4-vector over a four-dimensional volume, it looks like what I get when I sum up the flux of that guy over all of the little faces.

So I'm not going to step through the proof of that. It's fairly elementary, and basically it's just like proofs of Gauss's law that you have seen elsewhere, but there's an extra dimension attached to it. Really nothing new that's going on here. The thing which is new is, this is now being done in an additional dimension, and where this tends to be useful is when there is some kind of a conservation law that tells you something about this left-hand side here.

So the whole starting point of this discussion was I wrote down, on intuitive grounds, that the rate of change of the total amount of-- so what I did was I had an integral of dust density over a 3-volume, and I said, d by dt of that was balanced by the flux through the surfaces on the edge of that 3-volume.

As I argued, that's an observer-dependent statement, because you have chosen a particular time. You have chosen a particular space to make those volumes. This, on the other hand-- so let's switch my general vector field V to be N that we started this discussion off with.

So this guy-- oops. Let me write it the way I wrote it over there. d^4x over some 4-volume. This must be 0, because we are going to require that this thing be-- so if my number density is conserved, this must be 0.

So this tells me that when I do this flux integral, it's going to have to be 0. Let me now break this integral up, and actually write this into a form that is a little closer to the way that you may have seen something like this before.

So what I want to do is actually zoom in and think about the four-dimensional volume that I'm doing this integral over. So I'm just going to do a two-dimensional cut of it here on the blackboard. Let's let the time axis go up, and let's define the edges of my volumes to be t_1 and t_2 . x -axis going across here. Boundaries are x_1 and x_2 .

OK. So here is my V_4 . And every face here is an example of my boundary of d^4 . OK. Of course there's one over here as well, and over here. So if I were to do the top line here, I know that's got to get me 0, so I'm going to take advantage of this and say, let's just look at what happens when I do the integral of the flux of this thing across the many different faces.

So in a four-dimensional parallelepiped-- an n -dimensional parallelepiped has $2n$ phases, so there will be eight integrals we need to do. They're pretty obvious though, so I'm just going to write down a handful of them.

OK. So over all those faces-- $N_\alpha \sigma_\alpha$. So what I'm going to do now is say, OK, let's evaluate this on the face that is at t equals t_2 . So when I do this, I'm going to get $N_0 dx dy dz$.

Let's do next the contribution at moment-- actually, let me move this over, because I want a little bit more room. I also do one on the slice t equals t_1 . When I do this, though, in the same way that when you guys do fluxes in three-dimensional space, you get a sign associated with the orientation of these things, because the Levi-

Civita symbol here has sort of a right-hand rule built into it.

And so when I do it for the side that is on the future side of the box, I'm going to get a plus sign. Do all the analysis carefully on the side that's on the past side the box, you get a minus sign. Then I'm going to do it-- I'll pick out the component N_1 , and I'm going to do this along the face that's at $x = x_1$. Sorry. This is going to be x_2 .

OK. You can write down four other integrals according to when you do the y side and the z side that I've not drawn. OK. What I'm going to do now is imagine that this box gets very small in the timelike direction.

So let's let t_2 go to $t_1 + dt$. And I'm going to do this and then rearrange things a little bit. Again, apologies. I'm eating chalk here.

So I'm going to rearrange this integral so that it's of the form-- this side's being evaluated at t_1 plus t_2 . So I'm going to rearrange stuff so that I can then write-- just one moment.

Oh. Apologies. This whole thing, of course, equals 0. I'm just looking over my notes and went-- there was a magic sign flip. I was just looking at it and going, where the hell did that come from?

OK. So I left off my equals 0 there. I'm basically moving a bunch of terms to the other side. So now I'm integrating this over the face at x_2 . And I do an integral over the face at x_1 . And then I'm not drawing in a bunch of faces along y_2 and y_1 , z_1 and z_2 .

Divide both sides by dt . Take the limit of dt going to 0. What is this? This is the derivative of the volume integral of N_0 . So let's just write that out explicitly.

So if I've got integral $t_1 + dt$ $N_0 dx dy dz$ minus the integral at t_1 -- all this divided by dt . Take the limit as dt goes to 0. That becomes this. That i is a bad erasure. My apologies.

What do I get for the other term? Well, when I divide out that dt , I am just left with the flux of the spatial component of the number vector N through all six sides at moment t , or t_1 .

And at last-- so I will wrap things up in just a moment here. What this finally leads us to is that this fully covariant form that I have at the top there-- this becomes-- so in the language that you probably learned about in an earlier class, you can think of this as the area element of each side. These are completely equivalent to one another after you have chosen a particular Lorentz frame and invoked this four-dimensional analog of Gauss's law.

So that was a lot to take, but I wanted to do it for this one particularly-- to be blunt, this was a particularly simple example, because-- so right now, the only matter that I've introduced beyond particle kinematics is dust. And dust is actually surprisingly important.

When we actually get to cosmology, there are essentially two forms of matter that we consider when we study cosmology, and one of them is dust. In cosmological situations, a dust particle is basically a galaxy, or even a cluster of galaxies. We're thinking big.

So it's not trivial to do this, but we are soon going to need to introduce mathematical tools that are a little better for describing things like fluids, and stuff like that. And when we do that, we are at last going to have things like conservation of energy and conservation of momentum. It will be really easy to write down a differential conservation law that describes conservation of both energy and momentum using that mathematical object.

Via a process very similar to what I just did here, we can then turn this into integrals that describe how energy is conserved in a particular volume where things may be flowing into or flowing out of it, and how momentum is conserved per unit volume as things flow into it and flow out of it.

I'm not going to go through that in detail on the board, but having done this-- I have some notes that I'm going to post to the website that describe this, and that will be the way we communicate this. And we're going to take advantage-- to be blunt, we will mostly just use this differential form. We'll use the fact there is a particular mathematical object whose divergence-like derivative there is equal to 0.

OK. We have a few minutes left and we've ended early a few times, so I'd like to take advantage of this to switch gears a little bit and talk about another important 4-

vector that plays some role in physics, and allows me to introduce a few other very useful tricks that we will often take advantage of at various points in this class.

I forgot there's a straw. When you tip it, the straw doesn't work. Anyway, so the next example of stuff that we will occasionally talk about is an electric current. So switching gears very, very much now.

So we will describe this as a 4-vector whose timelike component is the charge density as seen by some observer, and whose spatial component is the current density as seen by that observer. Bear in mind, you might look at this and twitch a little bit because the units look wrong. Don't forget c equals 1. So if you use this in other systems of units, sometimes we call this component charge density times c .

So you guys have all-- you know what, let me just go ahead and write it out. So a couple properties about this are important and interesting. One is that the current and the charge density obey a continuity equation. We can think of it as a conservation of charge.

It's expressed-- so in elementary E&M, you guys all presumably learned that the rate of change of charge density is related to the divergence of the current density. Well, this is exactly the same as saying that the spacetime divergence of the current 4-vector is equal to 0.

This is really useful for us. And indeed, bearing this in mind, we find that if we want to express Maxwell's equations in a covariant way, we can do so such that this is built in automatically. So skipping over a few sets of things in my notes, we're going to find-- we're not going to find. We have found that electric fields and magnetic fields are inconvenient objects to describe using geometric objects that are appropriate for spacetime.

If I want a geometric object that is appropriate for spacetime, the first thing you think of is a 4-vector. A 4-vector has four components. E-fields and B-fields have a total of six components among them. So what you going to do, have two 4-vectors and just ignore two of the components? That seems sketchy.

So you think, eh, you know what? Why don't we make a tensor. Ah crap, a tensor has 16 components. That doesn't seem right. Then you go, ooh, I can make it

symmetric.

If you have a symmetric 4-by-4 tensor, well, that basically means that the number of independent numbers that go into this thing-- you have four down the diagonal. And then you count the number that are off the diagonal-- you have six. There's 16 in total, but the ones that are off the diagonal are equal to one another, so you have four on diagonal, six off. Too many.

So you say, well, what about antisymmetric? If I have an antisymmetric object, that means that component $F_{\alpha\beta}$ is the negative of component $F_{\beta\alpha}$. When you do that, that forces you the conclusion that there are, in fact, actually only six independent numbers in that thing.

The diagonal has to be 0, because set β equal to α here-- $F_{\alpha\alpha}$ equals minus $F_{\alpha\alpha}$. That only works if that component is equal to 0. So the diagonal becomes zero, and only the six off-diagonals survive.

And then you go, holy crap. Six. That's exactly what I need to have a geometric object that cleanly holds the three independent electric field components and the three independent magnetic field components.

So many of you have seen all this. If this isn't familiar to you, take a look at a book like Griffiths or something like that. It goes through this. The punchline is that what you find is that the electric and magnetic field is very nicely represented by this antisymmetric 2-index object, whose components, in the units that we are using, are filled with the E and the B like so.

OK. Last semester I had these memorized, but I have totally forgotten them. OK. So this is a geometric object that-- whoops. Anybody know why I just said whoops and had to put a dot on that there?

The point is that electric and magnetic fields look different to different observers. This is its representation according to one particular Lorentz observer. So it's important to get that right.

So in terms of this-- you're all familiar with the four Maxwell's equations. They turn out to be equivalent to-- so if you take a divergence of this F , it is-- depending on your units, so you might want to put μ_0 's in there, and things like that. Basically, I'm

setting everything that I can never remember to 1.

The divergence of that thing-- by the way, actually divergence on the second index-- becomes the current density. This will actually only give you half of the Maxwell's equations. The other half-- what you do is you lower these two indices, and there's this cyclic permutation of derivatives that give you these.

So you put these two things together, you apply it to this form I've written out here, and you will reproduce Maxwell's equations as they are presented in textbooks like Purcell and Griffiths. The thing which I want to emphasize here is, this form that we've got is written in such a way that the conservation of source is built into it. This geometric language requires that J_μ have no divergence.

And let me just show you can do this. All you need to know is that F is an antisymmetric tensor. You don't need to know anything about the properties of the E - and the B -field. So let's just try it.

Let's look at the divergence of the current. So I'm going to do $4\pi d_\mu J_\mu$. So I'm going to take a μ derivative of this. Now bear in mind, μ and ν -- I'm using Einstein's summation convention-- they are dummy indices, so I can relabel them.

I can change them. I can change μ to an α , ν to a β , or I can just change μ to ν , ν to μ . So this-- as long as I do it consistently on all objects that have those things-- this is the exact same thing.

But I also know that this tensor is antisymmetric. So if I switch these guys back, I get a minus sign. OK? Antisymmetry.

What happens if I switch the order of the derivatives? Does it matter whether I take the x derivative first and then the y derivative, or the y derivative and then the x derivative? Partial derivatives commute with each other, right? They are perfectly symmetric.

So this-- so what I've got is $d_\mu d_\nu F_{\mu\nu}$ is minus $d_\mu d_\nu F_{\nu\mu}$, and that only works if the whole thing is 0. What I just did is, I actually just invoked a trick that we are going to use many times.

And the one reason why-- this is a bit of a tangent. Or more than a tangent-- 150-

degree turn from what I've been discussing before-- but I wanted to make sure you saw this little trick. Whenever I have an object that is antisymmetric in its indices and I contract it with an object that is symmetric in its indices, you just get 0.

You can go through the little exercise if you don't feel fluent in this yet. You can do the exercise I just did up there again over. Take advantage of dummy indices, swap one. Work in the antisymmetry, swap the other. Work in the symmetry-- boom. You will necessarily prove that you've got 0.

So I lay this out here because there's going to be several times later in the course, where I'm going to get to a particular calculation and there's going to be some godawful mess. We're going to look at it and go, oh, this is horrible, and then go, wait, symmetry-antisymmetry, boom. We just killed 13 terms. Tricks are fun.

So I am going to stop there for today. There is another symmetry-antisymmetry thing which allows you to-- if you apply it to the equation of motion of a charge in an electromagnetic field, it just shows you that that equation of motion builds in the fact that in spacetime, the acceleration is always orthogonal to the 4-velocity.

I won't do it in class. It is on page 6 of the notes that I'm about to post up. We will pick up next time. We'll begin by talking about the stress energy tensor.